Cisco UCS X-Series for VMware Horizon 8 on VMware vSAN 8 for up to 600 Seats

Deployment Guide for Virtual Desktop Infrastructure built on Cisco UCS 210C M7 X-Series with 4th Generation Intel Xeon processors, Cisco Intersight, VMware Horizon 8 2212, VMware vSAN 8 and VMware vSphere 8.0 Hypervisor

Published: December 2023

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.

This document details the design of the VMware HCI Virtual Desktop Infrastructure for VMware Horizon 8 and VMware vSphere 8.0 developed by Cisco.

The solution covers the deployment of a predesigned, best-practice data center architecture with:

- VMware Horizon and VMware vSphere
- VMware vSAN 8
- Cisco Unified Computing System (Cisco UCS) incorporating the Cisco X-Series modular platform
- Cisco Nexus 9000 family of switches

Additionally, the Cisco Intersight cloud platform delivers monitoring, orchestration, workload optimization, and lifecycle management capabilities for the solution.

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

Solution Overview

This chapter contains the following:

- Audience
- Purpose of this Document
- What's New in this Release?

The joint solution from Cisco and VMware is designed to provide a highly efficient, scalable, and high-performing infrastructure for hosting VMware Horizon 8 virtual desktops. The solution utilizes VMware Virtual SAN 8 as a hyper-converged storage solution, which allows for the seamless integration of storage and compute resources, resulting in simplified management and lower costs.

This document provides guidance on the deployment best practices for this pre-validated architecture, built on the latest technologies from Cisco and VMware. The solution has been rigorously tested and validated, ensuring it can be deployed quickly and efficiently.

The solution is designed to be scalable, allowing organizations to easily add additional resources as needed to meet changing workload demands. Additionally, the solution is highly performant, with optimized resource utilization and efficient data access.

Overall, this Cisco Validated Design provides an efficient and effective architecture for hosting virtual desktop workloads, allowing organizations to quickly and easily deploy a high-performance infrastructure that can scale as needed to meet changing business demands.

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:

- VMware Horizon 8 VDI
- VMware Virtual SAN 8
- Cisco UCS X210c M7 Blade Servers with fourth-generation Intel Xeon processors running VMware vSphere 8
- Cisco Nexus 9000 Series Ethernet Switches

What's New in this Release?

This version of the VDI Design introduces the Cisco UCS X-Series modular platform.

Highlights for this design include:

- Support for Cisco UCS X9508 chassis with Cisco UCS X210c M7 with fourth-generation Intel Xeon processors
- VMware VSAN 8.0 (OSA/ESA)

- VMware Horizon 8 2212 (ESB)
- Support for VMware vSphere 8.0
- Support for VMware vCenter 8.0 to set up and manage the virtual infrastructure as well as integration of the virtual environment with Cisco Intersight software
- Support for Cisco Intersight platform to deploy, maintain, and support solution UCS components
- Support for Cisco Intersight Assist virtual appliance to help connect the VMware vCenter with the Cisco Intersight platform

The use cases include:

- Enterprise Data Center
- Service Provider Data Center
- Large Commercial Data Center

Technology Overview

This chapter contains the following:

- 4th Gen Intel Xeon processors
- Cisco Unified Computing System
- Cisco UCS Fabric Interconnect
- Cisco Unified Computing System X-Series
- Cisco UCS Virtual Interface Cards (VICs)
- Cisco Switching
- VMware Horizon
- VMware vSphere 8.0 Update 1
- VMware vSAN
- Cisco Intersight Assist Device Connector for VMware vCenter

4th Gen Intel Xeon Processors

4th Gen Intel® Xeon® processors, designed to accelerate performance across the fastest-growing workloads that businesses depend on today.

Built-in Intel® Accelerator Engines improve performance across AI, data analytics, networking, storage, and HPC. By making the best use of CPU core resources, built-in accelerators can result in more efficient utilization and power efficiency advantages, helping businesses achieve their sustainability goals.

4th Gen Intel Xeon processors have advanced, hardware-enabled security technologies to help protect data while unlocking new opportunities for business collaboration and insights. No matter the deployment path, these processors enable solutions that help businesses scale infrastructure and achieve value, fast.

Highlights

- Redefine performance with 4th Gen Intel® Xeon® processors—featuring built-in accelerators to improve performance across the fastest-growing workloads in AI, data analytics, networking, storage, and HPC.
- The latest built-in Intel® Accelerator Engines and software optimizations help improve power efficiency—you can achieve a 3x average performance per watt efficiency improvement for targeted workloads utilizing built-in accelerators compared to the previous generation.[1]
- AI gets even better with all-new Intel® Advanced Matrix Extensions (Intel® AMX), delivering exceptional AI training and inference performance through accelerated matrix multiply operations.
- Other built-in accelerators speed up data movement, encryption, and compression for faster networking and storage, boost query throughput for more responsive analytics, and offload scheduling and queue management to dynamically balance loads across multiple cores.
- By making the best use of CPU core resources, built-in accelerators can result in more efficient utilization and power efficiency advantages, helping businesses achieve their sustainability goals.
- Intel® Software Guard Extensions (Intel® SGX) and other hardware-enabled security features help bring a zero-trust security strategy to life while unlocking new opportunities for business collaboration and insights—even with sensitive or regulated data.

- Solutions built on Intel® Xeon® processors offer the most choice and flexibility, no matter the deployment path—on-prem, hybrid cloud, network, or edge.

[1] See E1 at intel.com/processorclaims: 4th Gen Intel® Xeon® processors. Results may vary.

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 blade and rack form factor, managed by Cisco UCS Manager.
- 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 lowers 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 system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers 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 UCS Manager software. Cisco UCS Manager 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 re-sources 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 UCS Manager:

- 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 blade 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 blade 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: Once a compute resource is discovered by Cisco UCS Manager, 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 showcases the policy-based resource classification of Cisco UCS Manager.
- Combined Rack and Blade Server Management: Cisco UCS Manager can manage Cisco UCS B-series blade 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 UCS Manager 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 UCS Manager with other management systems.
- Policies, Pools, Templates: The management approach in Cisco UCS Manager 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 UCS Manager, a service 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. This provides different subject matter experts to work independently from each other. This provides great flexibility where different experts from different domains, such as network, storage, security, server, and virtualization work together to accomplish a complex task.
- Policy Resolution: In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the re-al-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.
- Service Profiles and Stateless Computing: A service 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 UCS Manager 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 UCS Manager 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 1. Cisco Intersight

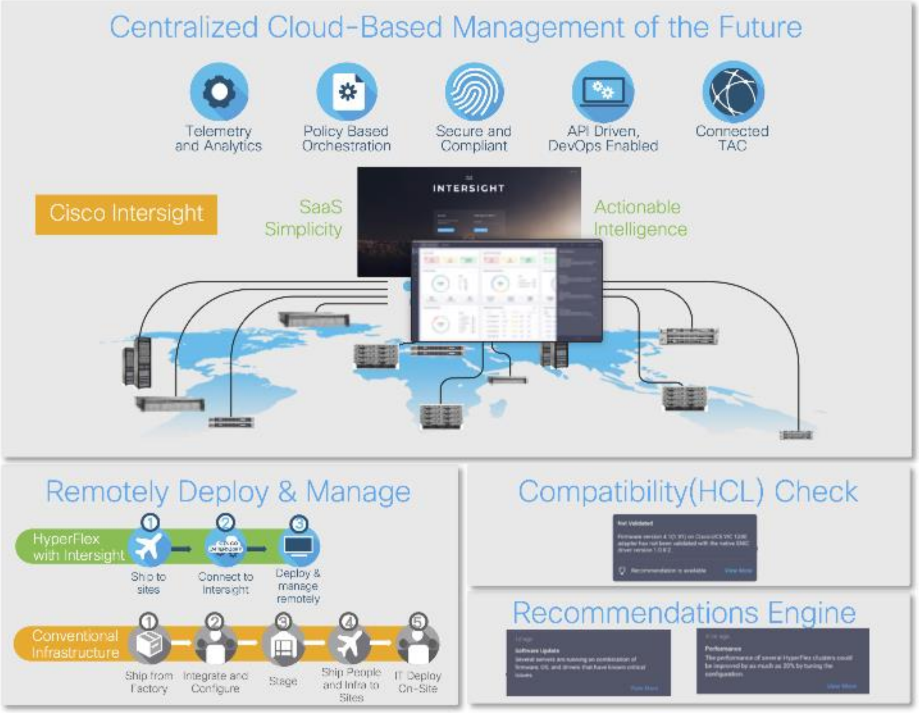

- Automate your infrastructure

  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 and 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 of 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.

- 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 Inter-sight Mobile App.

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

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, S-Series, and HX-Series Rack-Mount Servers, Cisco UCS B-Series Blade Servers, and Cisco UCS 5100 Series Blade Server Chassis. All servers and chassis, and therefore all blades, 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 6454 Series uses a cut-through architecture, supporting deterministic, low latency, line rate 10/25/40/100 Gigabit Ethernet ports, 3.82 Tbps of switching capacity, and 320 Gbps bandwidth per Cisco 5108 blade chassis when connected through the IOM 2208 model. The product family supports Cisco low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet net-works. The Fabric Interconnect supports multiple traffic classes over the Ethernet fabric from the servers to the uplinks. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6454 Fabric Interconnect

The Cisco UCS 6454 Fabric Interconnect is a one-rack-unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has eight (8) 10/25-Gbps fixed Ethernet ports, which optionally can be configured as 8/16/32-Gbps FC ports (ports 1 to 8), thirty-six (36) 10/25-Gbps fixed Ethernet ports (ports 9 to 44), four (4) 1/10/25-Gbps Ethernet ports (ports 45 to 48), and finally six (6) 40/100-Gbps Ethernet uplink ports (ports 49 to 54). For more information, refer to the Cisco UCS 6454 Fabric Interconnect spec sheet: https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/6400-specsheet.pdf

Figure 2. **Cisco UCS 6454 Fabric Interconnect**

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 customer 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 integrating with third-party devices, including VMware vCenter, to provide visibility, optimization, and orchestration from a single platform, thereby driving agility and deployment consistency.

Why Cisco UCS X-Series for VMware vSAN?

Companies have traditionally chosen rack form-factor servers for VMware vSAN deployments because of the requirements for fast networks, storage capacity, or GPUs. The Cisco UCS X-Series easily meets these requirements. The Cisco UCS X-Series provides the functionalities of both blade and rack servers by offering compute density, storage capacity, and expandability in a single system, embracing a wide range of workloads in your data center.

With six drives, up to 15.3TB each, customers have plenty of storage capacity, up to 91TB per server node. You can quickly and seamlessly add GPUs to any server node with Cisco UCS X-Series X-Fabric Technology. And only Cisco offers a 100G unified fabric. Unified fabric allows customers to easily wire the Cisco UCS X9508 Chassis for needed bandwidth regardless of type – data, storage, or management.

Cisco UCS X-Series Modular System is managed from the cloud using Cisco Intersight. It is designed to meet the needs of modern applications and improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design.

Cisco UCS X-Series can reduce risk by using pre-validated and tested hardware configurations of vSAN ReadyNodes. This simplifies deployment, making it easy and quick to deploy the solution for your hyper-converged infrastructure.

Figure 3. Cisco UCS X9508 Chassis

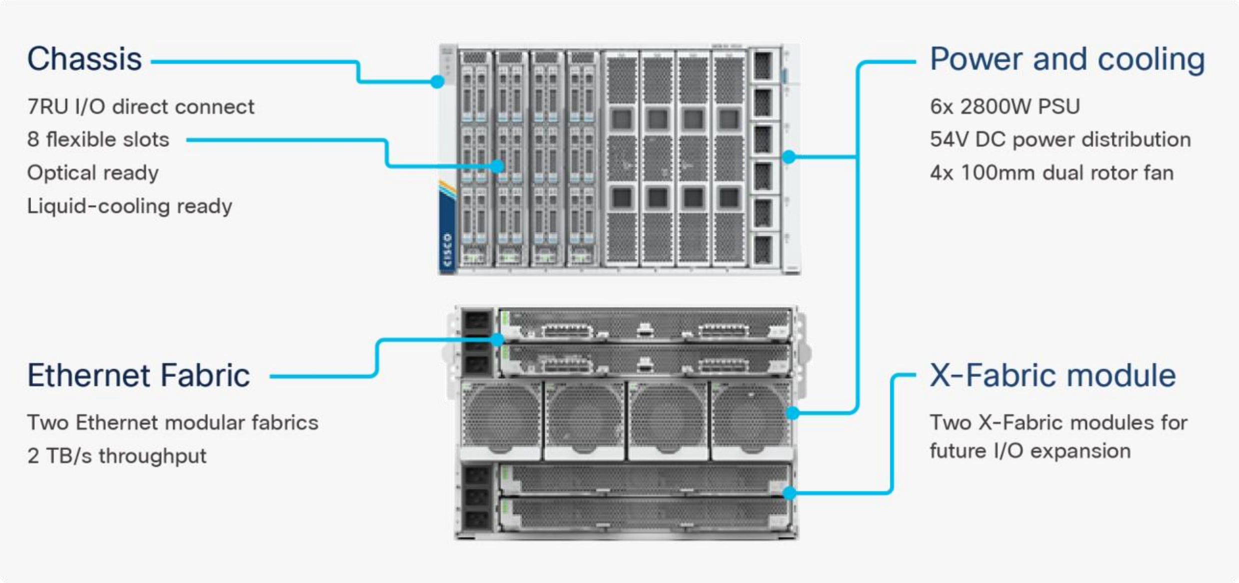

The various components of the Cisco UCS X-Series are described in the following sections.

Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As shown in Figure 4, 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 4. Cisco UCS X9508 Chassis – Midplane Free Design

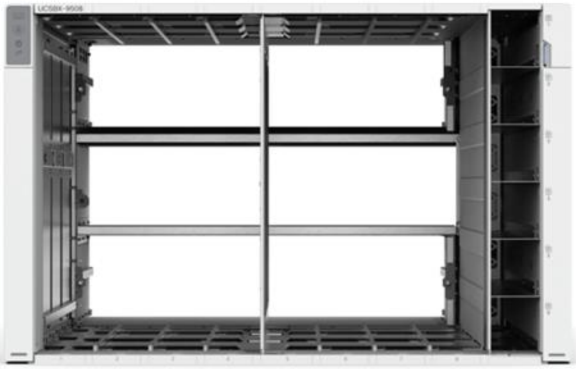

The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes and a pool of future I/O resources that may include GPU accelerators, disk storage, and nonvolatile memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) that connect the chassis to upstream Cisco UCS 6400 Series Fabric Interconnects. At the bottom rear of the chassis are slots ready to house future 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 the customer’s environment.

Cisco UCSX 9108-25G Intelligent Fabric Modules

For the Cisco UCS X9508 Chassis, the network connectivity is provided by a pair of Cisco UCSX 9108-25G 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 6400 Series 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 5. Cisco UCSX 9108-25G Intelligent Fabric Module

Each IFM supports eight 25Gb uplink ports for connecting the Cisco UCS X9508 Chassis to the FIs and 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 UCS FIs, providing up to 400Gbps connectivity across the two IFMs. The unified fabric carries management, VM, and Fibre Channel over Ethernet (FCoE) traffic to the FIs, where management traffic is routed to the Cisco Intersight cloud operations platform, FCoE traffic is forwarded to the native Fibre Channel interfaces through unified ports on the FI (to Cisco MDS switches), and data Ethernet traffic is forwarded upstream to the data center network (via Cisco Nexus switches).

Cisco UCS X210c M7 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to eight Cisco UCS X210c M7 Compute Nodes. The hardware details of the Cisco UCS X210c M7 Compute Nodes are shown in Figure 6.

Figure 6. Cisco UCS X210c M7 Compute Node

The Cisco UCS X210c M7 features:

- CPU: Up to 2x 4th Gen Intel Xeon processors with up to 60 cores per processor and up to 2.625 MB Level 3 cache per core and up to 112.5 MB per CPU.
- Memory: Up to 8TB of main memory with 32x 256 GB DDR5-4800 DIMMs.
- Disk storage: Up to six hot-pluggable, solid-state drives (SSDs), or non-volatile memory express (NVMe) 2.5-inch drives with a choice of enterprise-class redundant array of independent disks (RAIDs) or passthrough controllers, up to two M.2 SATA drives with optional hardware RAID.
- Optional front mezzanine GPU module: The Cisco UCS front mezzanine GPU module is a passive PCIe Gen 4.0 front mezzanine option with support for up to two U.2 NVMe drives and two HHHL GPUs.
- mLOM virtual interface cards:

  - Cisco UCS Virtual Interface Card (VIC) 15420 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 50 Gbps of unified fabric connectivity to each of the chassis intelligent fabric modules (IFMs) for 100 Gbps connectivity per server.
  - Cisco UCS Virtual Interface Card (VIC) 15231 occupies the server's modular LAN on motherboard (mLOM) slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis intelligent fabric modules (IFMs) for 100 Gbps connectivity per server.
- Optional mezzanine card:
  - Cisco UCS 5th Gen Virtual Interface Card (VIC) 15422 can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric technology. An included bridge card extends this VIC's 2x 50 Gbps of network connections through IFM connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per server).
  - Cisco UCS PCI Mezz card for X-Fabric can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric modules and enable connectivity to the Cisco UCS X440p PCIe Node.
  - All VIC mezzanine cards also provide I/O connections from the X210c M7 compute node to the X440p PCIe 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)

The Cisco UCS VIC 15000 series is designed for Cisco UCS X-Series M6/M7 Blade Servers, Cisco UCS B-Series M6 Blade Servers, and Cisco UCS C-Series M6/M7 Rack Servers. The adapters are capable of supporting 10/25/40/50/100/200-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). They incorporate Cisco’s next-generation Converged Network Adapter (CNA) technology and offer a comprehensive feature set, providing investment protection for future feature software releases.

Cisco UCS VIC 15231

The Cisco UCS VIC 15231 (Figure 7) is a 2x100-Gbps Ethernet/FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the Cisco UCS X210 Compute Node. The Cisco UCS VIC 15231 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.

Figure 7. **Cisco UCS VIC 15231**

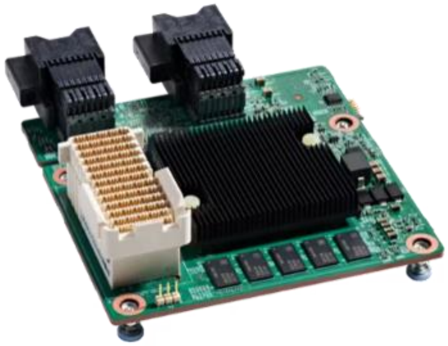

Figure 8. Cisco UCS VIC 15231 Infrastructure

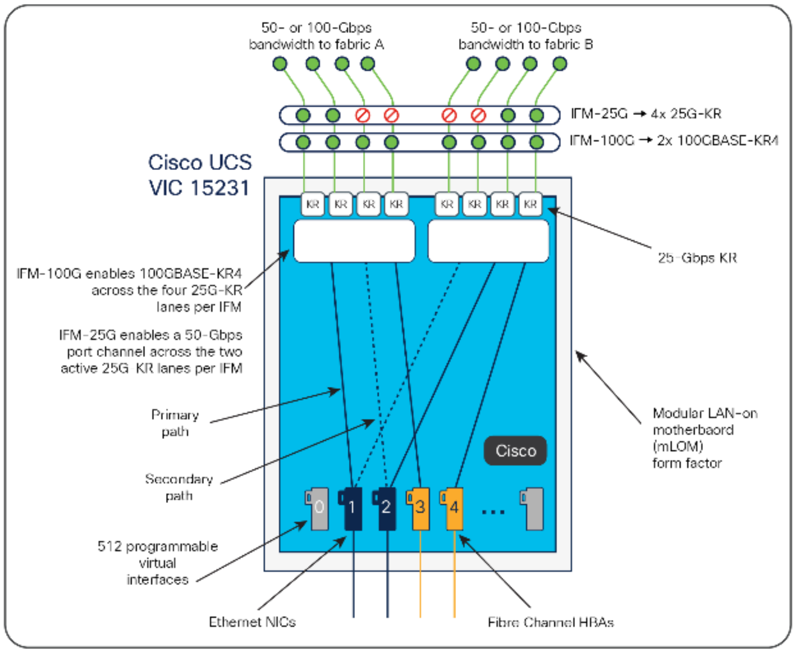
50- or 100-Gbps bandwidth to fabric A
50- or 100-Gbps bandwidth to fabric B
IFM-25G → 4x 25G-KR
IFM-100G → 2x 100GBASE-KR4
Cisco UCS VIC 15231
KR
25-Gbps KR
IFM-100G enables 100GBASE-KR4 across the four 25G-KR lanes per IFM
IFM-25G enables a 50-Gbps port channel across the two active 25G KR lanes per IFM
Primary path
Secondary path
Cisco
Modular LAN-on motherbaord (mLOM) form factor
0
1
2
3
4
...
512 programmable virtual interfaces
Ethernet NICs
Fibre Channel HBAs

Cisco Switching

Cisco Nexus 93180YC-FX Switches

The Cisco Nexus 93180YC-EX 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 infra-structure
    - Python Scripting for programmatic access to the switch command-line interface (CLI)
    - Hot and cold patching, and online diagnostics
- Investment Protection

A Cisco 40 Gbe bidirectional transceiver allows reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Giga-bit Ethernet Support for 1 Gbe and 10 Gbe access connectivity for data centers migrating access switching infrastructure to faster speed. The following is supported:

- 1.8 Tbps of bandwidth in a 1 RU form factor
- 48 fixed 1/10/25-Gbe SFP+ ports
- 6 fixed 40/100-Gbe QSFP+ for uplink connectivity
- Latency of less than 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies
- Hot swappable 3+1 redundant fan trays

Figure 9. **Cisco Nexus 93180YC-EX Switch**

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 2022 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:
        - Install or Upgrade to the latest released version of Connection Server.
        - Navigate to https://<CS-IP//FQDN>rest/swagger-ui.html from any browser.
        - 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.

VMware vSphere 8 Update 1

VMware vSphere is an enterprise workload 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.

Note: VMware vSphere 8 became generally available in November of 2022.

The vSphere 8 Update 1 release delivered enhanced value in operational efficiency for admins, supercharged performance for higher-end AI/ML workloads, and elevated security across the environment. vSphere 8 Update 1 has now achieved general availability.

For more information about VMware vSphere 8 Update 1 three key areas of enhancements, see VMware blog.

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 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.

VMware vSAN

VMware vSAN is a software-defined storage solution that provides hyper-converged infrastructure (HCI) for virtual machines (VMs). vSAN 8 is the latest version of the vSAN platform, released in 2022. It offers new features and enhancements that enable organizations to simplify storage management and accelerate application performance.

One of the most significant features of vSAN 8 is its ability to support modern hardware technologies, such as NVMe and persistent memory. This enables vSAN 8 to deliver high-performance storage capabilities, with faster data access and reduced latency.

As an optional architecture to the Original Storage Architecture (OSA) found in previous versions of vSAN, including vSAN 8, which uses disk groups with caching devices, vSAN 8 introduces a new architecture called the vSAN Express Storage Architecture (ESA). The ESA uses disk pools of high-performing NVMe-based storage devices and delivers all-new capabilities not possible with the OSA.

Here are some of the key features of vSAN 8.0:

- vSAN Express Storage Architecture. vSAN ESA is an alternative architecture that provides the potential for huge boosts in performance with more predictable I/O latencies and optimized space efficiency.
- Increased write buffer. vSAN Original Storage Architecture can support more intensive workloads. You can configure vSAN hosts to increase the write buffer from 600 GB to 1.6 TB.
- Native snapshots with minimal performance impact. vSAN ESA file system has snapshots built in. These native snapshots cause minimal impact to VM performance, even if the snapshot chain gets deep. The snapshots are fully compatible with existing backup applications using VMware VADP.

vSAN 8 also includes enhanced support for VMware Cloud Foundation, which allows organizations to easily manage their HCI infrastructure across multiple clouds.

Overall, vSAN 8.0 is a significant upgrade to the vSAN platform, providing new features and enhancements that enable organizations to simplify storage management and accelerate application performance. With the ability to run stateful containerized applications, support for native file services, and enhanced security features, vSAN 8.0 is an attractive option for organizations looking to build a modern, hyper-converged infrastructure.

Cisco Intersight Assist Device Connector for VMware vCenter

Cisco Intersight has the capability to integrate with VMware vCenter. This integration is made possible through the use of the device connector, which runs within the Cisco Intersight Assist virtual appliance. The device connector is responsible for facilitating communication between Cisco Intersight and VMware vCenter, enabling administrators to manage both their Cisco infrastructure and VMware vSphere environments from a single management console. This integration provides a comprehensive management experience, allowing administrators to monitor and manage virtualization and hardware resources, provision and manage virtual machines, and automate routine management tasks. Overall, the integration between Cisco Intersight, VMware vCenter helps organizations to improve their operational efficiency, reduce management complexity, and optimize their resource utilization.

Solution Design

This chapter contains the following:

Design Considerations for Desktop Virtualization

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 located.
- 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 as 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: physical device with a locally installed operating system.
- Remoted Desktop Server Hosted Sessions: 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 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 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 type 1 hypervisor and is synced with the data center when the device is connected to the network.

Note: For the purposes of the validation represented in this document, both Single-session OS and Multi-session OS VDAs were validated.

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 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

The following key project and solution sizing questions should be considered:

- 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 Single-session OS version?
- 32 bit or 64 bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, 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 Multi-session OS version?
- What is a method be used for virtual desktop deployment?
- 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 there a 3rd party graphics component?
- Is anti-virus a part of the image?
- What is the SQL server version for the database?
- Is user profile management (for example, non-roaming profile-based) part of the solution?
- What is the fault tolerance, failover, and disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Hypervisor Selection

VMware vSphere comprises the management infrastructure or virtual center server software and the hypervisor software that virtualizes the server's hardware resources. It offers features like Distributed Resource Scheduler, vMotion, high availability, Storage vMotion, VMFS, and a multi-pathing storage layer.

VMware vSphere supports VMware vSAN, a hyper-converged infrastructure (HCI) solution. vSAN uses local storage devices to establish a virtual storage area network (SAN) and produce a single shared data store. This data store is automatically provisioned to the hosts in the cluster.

VMware vSphere 8.0, with the latest enhancement in all areas of vSphere, has been selected as the hypervisor for this VMware vSAN with VMware Horizon Virtual Desktops and Remote Server Desktop Hosted (RDSH) Sessions deployment.

More information on VMware vSphere can be obtained at the VMware web site.

Storage Considerations

When considering storage for vSphere ESXi and VMware Horizon deployments using vSAN, there are several important considerations to keep in mind:

- vSAN Requirements: Ensure that your environment meets the requirements for vSAN deployment, such as having a minimum number of hosts in the cluster, supported hardware, and compatible firmware versions.
- Capacity Planning: Estimate your storage capacity needs based on the number of virtual machines (VMs), their sizes, and anticipated growth. Consider factors such as VM templates, linked clones, and user profiles that impact storage requirements.
- Performance Requirements: Determine the expected workload and performance requirements of your VMs and desktops, such as IOPS, throughput, latency, and the number of concurrent users.

- vSAN Configuration: Properly configure vSAN to align with your storage requirements. This includes setting up disk groups, selecting the appropriate RAID level, configuring cache and capacity tiers, and enabling features like deduplication and compression.
- Network Considerations: Ensure that your network infrastructure can support the traffic generated by vSAN. Consider factors such as bandwidth, latency, network redundancy, and quality of service (QoS) settings.
- Disaster Recovery: Plan for data protection and disaster recovery by configuring vSAN stretched clusters or using vSAN replication to replicate data to another site.
- Monitoring and Management: Implement monitoring and management tools to monitor the health, performance, and capacity of your vSAN storage.
- Compatibility: Ensure that your hardware and software components, including server hardware, storage devices, network switches, and ESXi and Horizon versions, are compatible with vSAN. Check VMware's compatibility guide for detailed information.
- Scalability: Consider future scalability requirements and plan for expansion by following best practices for adding hosts, disks, or expanding the vSAN cluster.
- Backup and Recovery: Implement a backup and recovery solution for your VMs and desktops to protect against data loss or corruption. This can involve using VMware Data Protection or third-party backup solutions.

Note: For specific guidance, consult VMware's documentation, best practice guides, or seek assistance from VMware support or Cisco.

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 10. **VMware Horizon Design Overview**

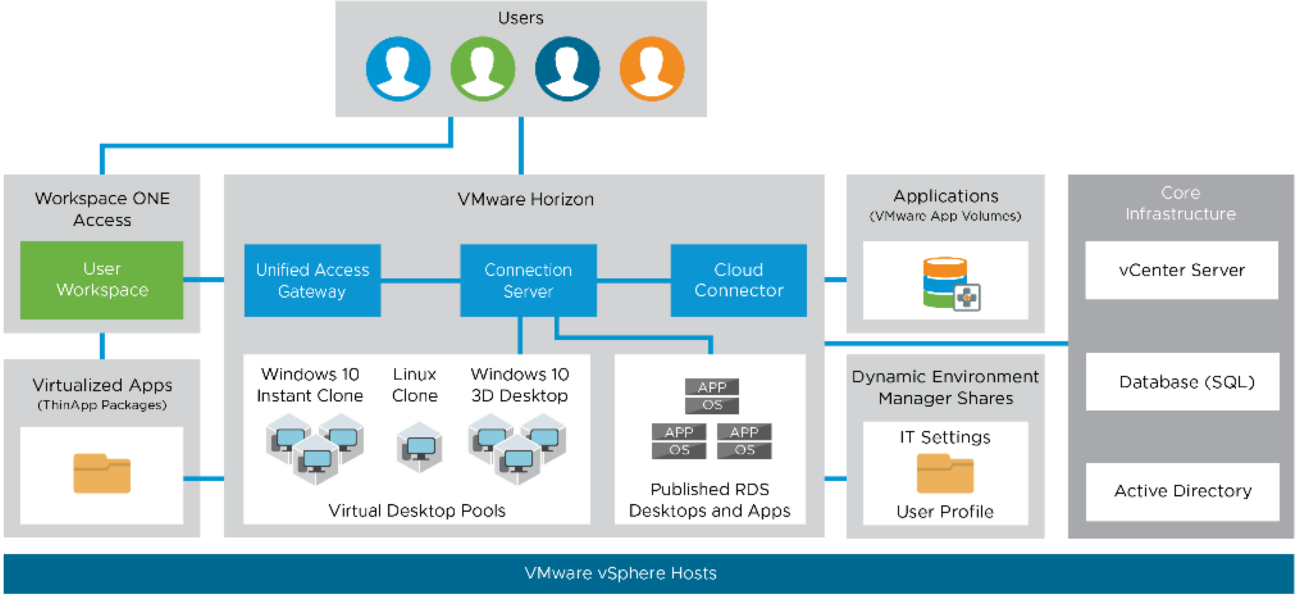

Horizon VDI Pool and RDSH Servers Pool

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

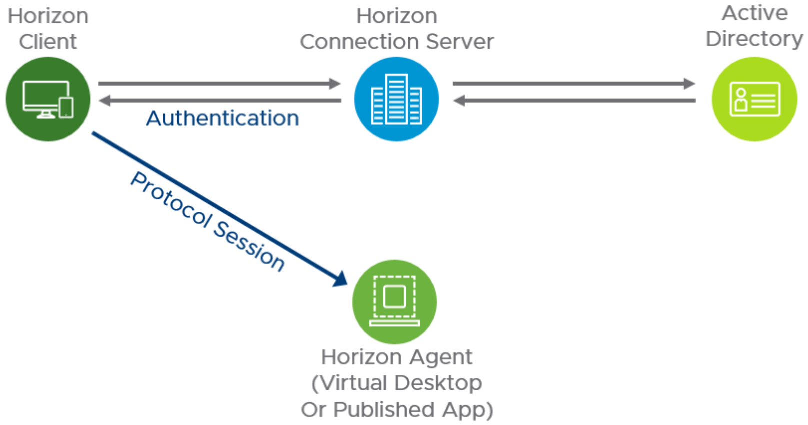
Horizon Client
Horizon Connection Server
Active Directory
Authentication
Protocol Session
Horizon Agent (Virtual Desktop Or Published App)

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 12 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 12. Multisite Configuration Overview

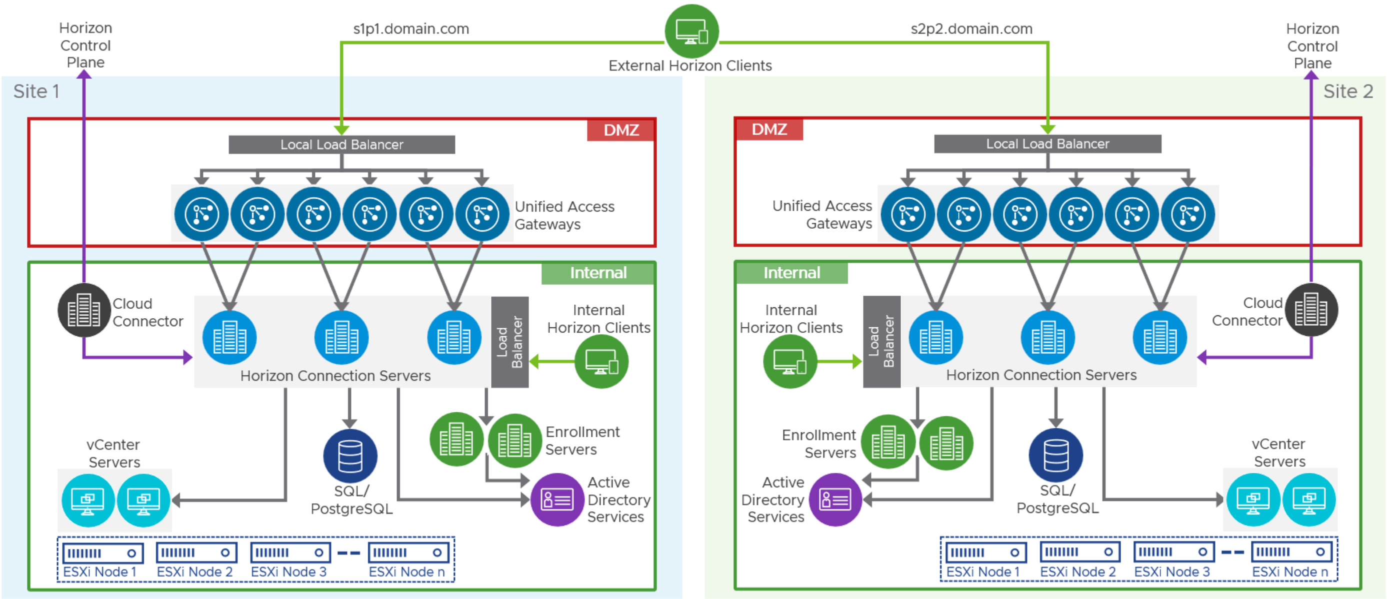

Based on the requirement and the number of data centers or remote locations, we can choose any available load-balancing software to increase security and optimize the user experience.

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 1. Desktop type and user experience

Desktop Type	User Experience
Server OS machines	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. 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. Application types: Any application.
Desktop OS machines	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. Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted

Desktop Type	User Experience
	applications. 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. 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.
Remote PC Access	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. 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. Host: The same as Desktop OS machines. Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.

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 of Hardware and Software

This chapter contains the following:

- Architecture
- Products Deployed
- Physical Topology
- Logical Architecture
- Configuration Guidelines

Architecture

This architecture delivers a Virtual Desktop Infrastructure that is redundant and uses the best practices of Cisco and VMware.

It includes:

- VMware vSphere 8.0U1 hypervisor installed on the Cisco UCS x210C M7 compute nodes.
- vSAN 8.0 provides storage for the nodes in the cluster.
- VDI workload delivered by VMware Horizon 8 2212.
- Cisco Intersight provides UCS infrastructure management with lifecycle management capabilities.

The architecture deployed is highly modular. While each customer's 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).

Products Deployed

This CVD details the deployment of up to 1440 multi-session OS, 585 Single-session OS VDI users featuring the following software:

- VMware vSphere ESXi 8.0U1 hypervisor
- VMware vCenter 8.0 to set up and manage the virtual infrastructure as well as integration of the virtual environment with Cisco Intersight software
- Microsoft SQL Server 2019
- Microsoft Windows Server 2019 and Windows 11 64-bit virtual machine Operating Systems
- 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

Physical Topology

The VMware vSAN on Cisco UCS X-Series Modular System physical connectivity details are explained below.

The system is composed of a pair of Cisco UCS 6454 Fabric Interconnects connected to Cisco UCS X9508 chassis, allowing the use of Cisco UCSX 210C M7 blades. Upstream network connections, also referred to as "northbound" network connections, are made from the Fabric Interconnects to the customer datacenter network at the time of installation.

For this study, we uplinked the Cisco UCS 6454 Fabric Interconnects to Cisco Nexus 93180YC-FX switches.

Figure 13. **VMware vSAN on Cisco UCS X-Series Modular System**

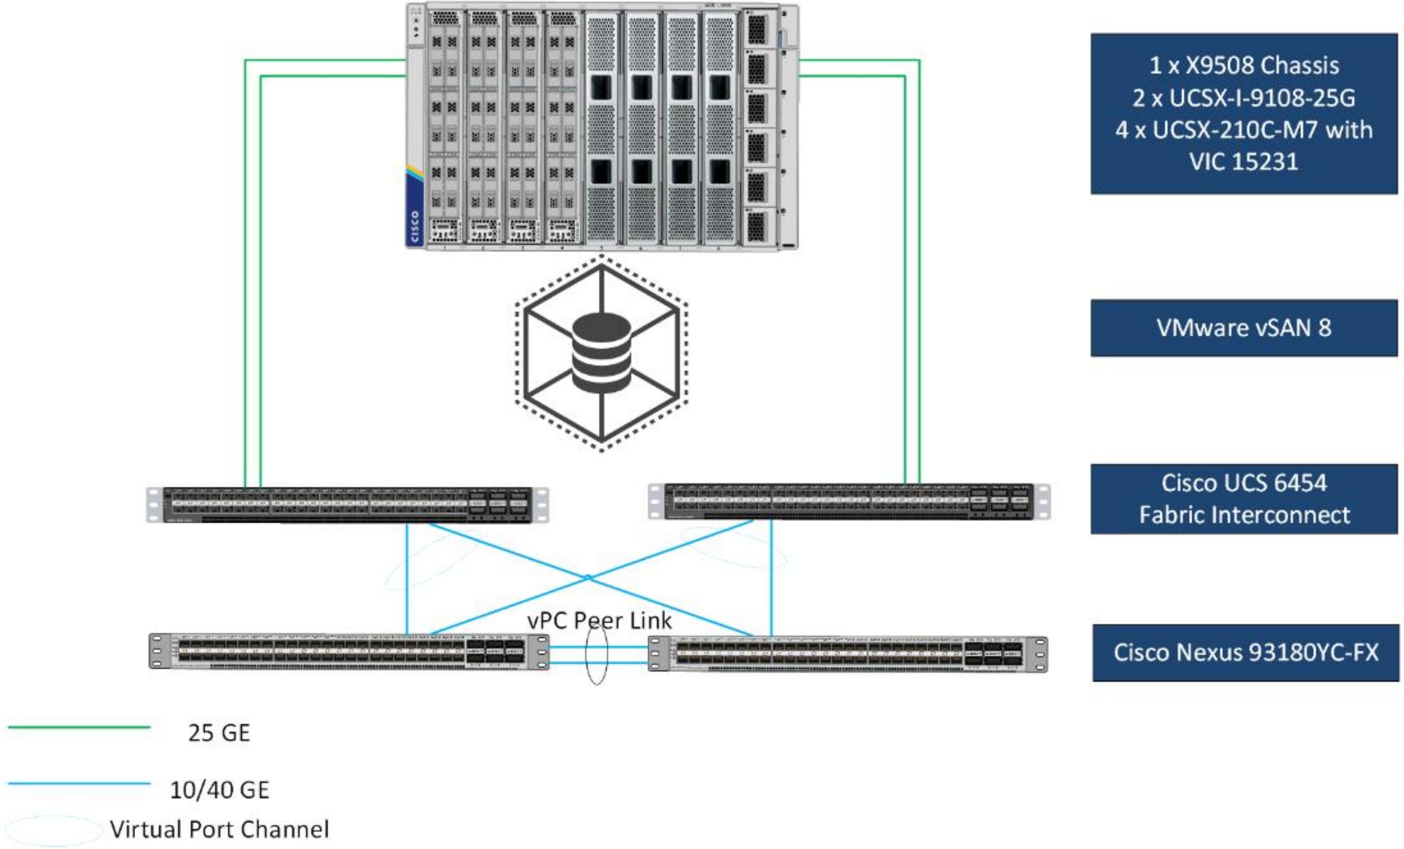

Figure 13 details the physical hardware and cabling deployed to enable this solution:

- Two Cisco Nexus 93180YC-FX Switches in NX-OS Mode.
- One Cisco UCS X9508 Chassis with two Cisco UCSX 9108 25G IF Modules.
- OSA vSAN cluster: Four Cisco UCS X210c M7 Compute Nodes with Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 1TB 4800 MHz RAM, and one Cisco VIC15231 mezzanine card, one 1.6TB High Performance High Endurance, and five 1.92TB High Performance Medium Endurance NVMe drives, providing N+1 server fault tolerance.
- ESA vSAN cluster: Three Cisco UCS X210c M7 Compute Nodes with Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 1TB 4800 MHz RAM, and one Cisco VIC15231 mezzanine card, five 3.2TB High Performance High Endurance NVMe drives, providing N+1 server fault tolerance.

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.

Table 2 lists the software versions of the primary products installed in the environment.

Table 2. Software and Firmware Versions

Vendor	Product/Component	Version/Build/Code
Cisco	UCS Component Firmware	4.2(3e)
Cisco	UCS x210c Compute Node	5.1(1.230052)
Cisco	VIC 15231	5.3(1.230046)
Cisco	Cisco Nexus 93180YC-FX	9.3(3)
VMware	vCenter Server Appliance	8.0.0-21216066
VMware	vSphere 8.0 Update 1a	8.0.1, 21813344
VMware	vSAN 8	8.0.1, 21383123
VMware	VMware Horizon 8 2212 Connection server	8.8.0-21073894
VMware	VMware Horizon 8 2212 Agent	8.8.0-21067308
Cisco	Intersight Assist	1.0.11-759
Microsoft	FSLogix 2210 hotfix 1	2.9.8440.42104
VMware	Tools	12.2.0.21223074

Logical Architecture

The logical architecture of the validated solution which is designed to run desktop and RDSH server VMs supporting up to 1400 users on a single chassis containing four blades configured with vSAN OSA or up to 960 users on a single chassis containing three blades configured with vSAN ESA, with physical redundancy for the blade servers for each workload type and have a separate vSphere cluster to host core services and management components, is illustrated in Figure 14.

Note: Separating management components and desktops is a best practice for the large environments.

Figure 14. **Logical Architecture Overview**

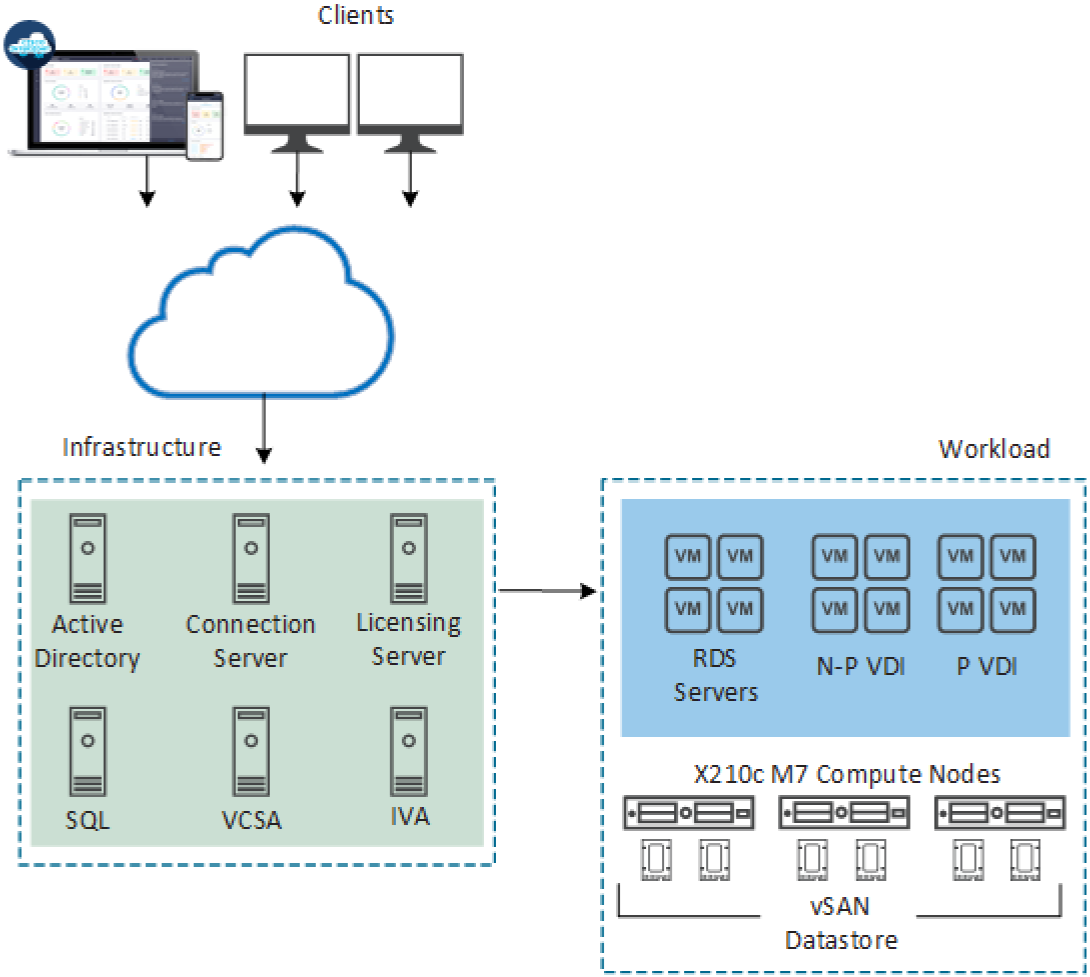

Configuration Guidelines

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, Cisco Nexus A and Cisco Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are configured similarly.

Note: This document is intended to allow the reader to configure the VMware Horizon customer environment as a stand-alone solution.

VLANs

The VLAN configuration recommended for the environment includes a total of six VLANs as outlined in Table 3.

Table 3. VLANs Configured in this study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN

VLAN Name	VLAN ID	VLAN Purpose
InBand-Mgmt_70	70	In-Band management interfaces
Infra-Mgmt_71	71	Infrastructure Virtual Machines
VDI_72	72	RDSH, VDI Persistent and Non-Persistent
vMotion_73	73	VMware vMotion
vSAN_74	74	vSAN Storage
OOB-Mgmt	164	Out-of-Band management interfaces

Solution Configuration

This chapter contains the following:

- Solution Cabling

Solution Cabling

The following sections detail the physical connectivity configuration of the HCI VMware Horizon VDI environment on vSAN.

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.

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.

Note: Be sure to follow the cabling directions in this section. Failure to do so will result in problems with your deployment.

Figure 15 details the cable connections used in the validation lab for vSAN topology based on the Cisco UCS 6454 fabric interconnect. 25Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the VMware HCI infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

Figure 15. vSAN cabling diagram

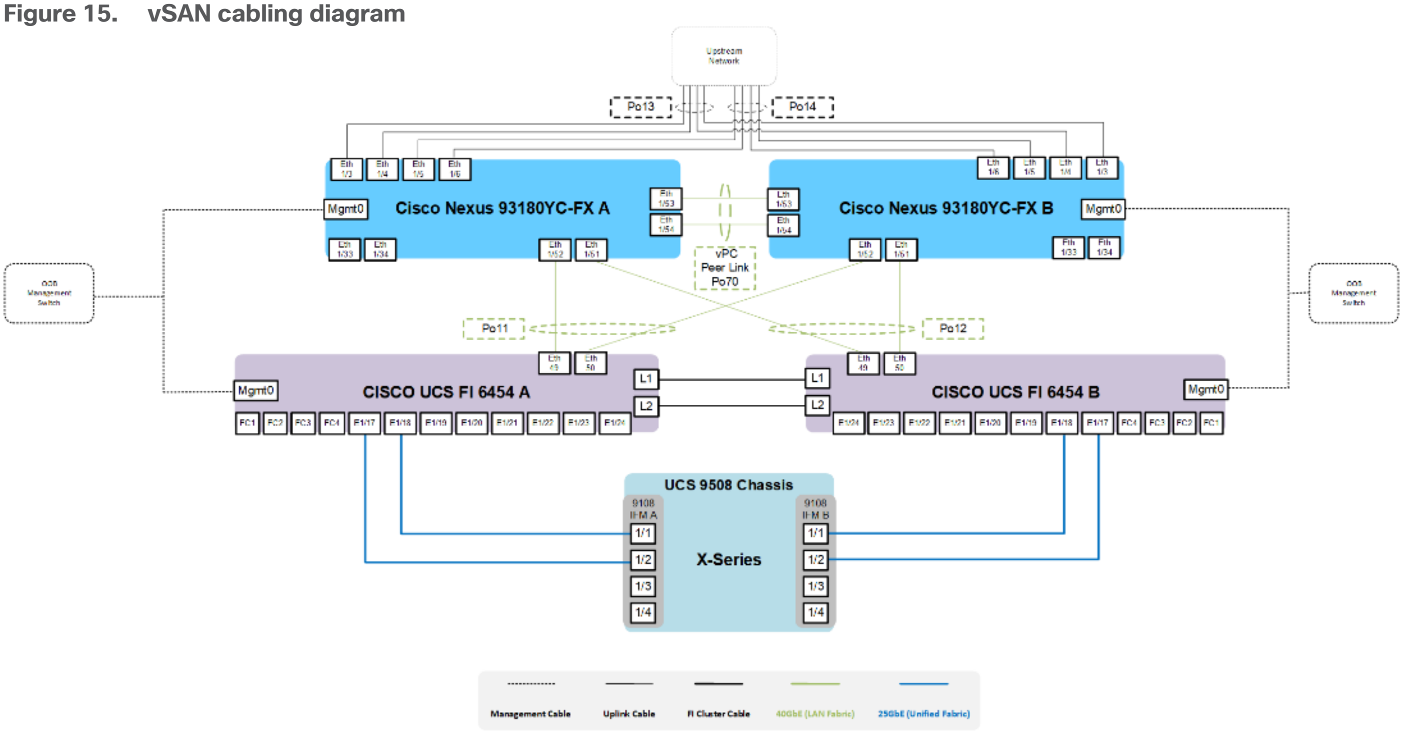

Configuration and Installation

This chapter contains the following:

- Cisco UCS X-Series Configuration – Intersight Managed Mode (IMM)
- Install and Configure VMware ESXi 8.0
- Cisco Intersight Orchestration
- Create vSAN Cluster

Note: Set of Ansible scripts developed for setting up Cisco converged infrastructure could be adopted to automate deployment of this environment. The scripts can be found here but automated deployment is not in scope of this CVD.

Cisco Intersight Managed Mode standardizes policy and operation management for Cisco UCS X-Series. The compute nodes in Cisco UCS X-Series are configured using server profiles defined in Cisco Intersight. These server profiles derive all the server characteristics from various policies and templates. At a high level, configuring Cisco UCS using Cisco Intersight Managed Mode consists of the steps shown in Figure 16.

Figure 16. Configuration Steps for Cisco Intersight Managed Mode

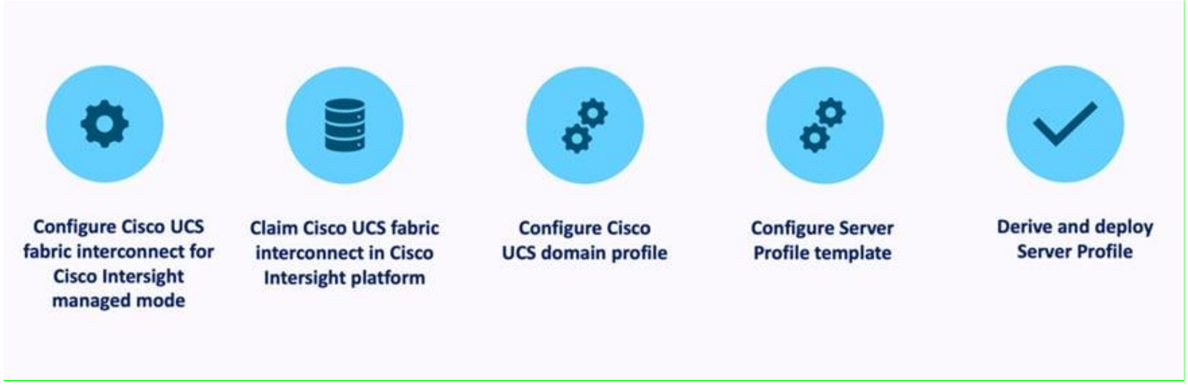
Configure Cisco UCS fabric interconnect for Cisco Intersight managed mode
Claim Cisco UCS fabric interconnect in Cisco Intersight platform
Configure Cisco UCS domain profile
Configure Server Profile template
Derive and deploy Server Profile

Cisco UCS X-Series Configuration – Intersight Managed Mode (IMM)

Procedure 1. Configure Cisco UCS Fabric Interconnects for IMM

Step 1. Verify the following physical connections on the fabric interconnect:

- The management Ethernet port (mgmt0) is connected to an external hub, switch, or router
- The L1 ports on both fabric interconnects are directly connected to each other
- The L2 ports on both fabric interconnects are directly connected to each other

Step 2. Connect to the console port on the first Fabric Interconnect.

Step 3. Configure Fabric Interconnect A (FI-A). On the Basic System Configuration Dialog screen, set the management mode to Intersight. All the remaining settings are similar to those for the Cisco UCS Manager managed mode (UCSM-Managed).

Cisco UCS Fabric Interconnects

Procedure 1. Configure the Cisco UCS for use in Intersight Managed Mode

Step 1. 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 2. 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 3. 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 B

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-6454)

  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. Claim a Cisco UCS Fabric Interconnect in the Cisco Intersight Platform

If you do not already have a Cisco Intersight account, you need to set up a new account in which to claim your Cisco UCS deployment. Start by connecting to https://intersight.com.

All information about Cisco Intersight features, configurations can be accessed in the Cisco Intersight Help Center.

Step 1. Click Create an account.

Step 2. Sign in with your Cisco ID.

Step 3. Read, scroll through, and accept the end-user license agreement. Click Next.

Step 4. Enter an account name and click Create.

If you have an existing Cisco Intersight account, connect to https://intersight.com and sign in with your Cisco ID, select the appropriate account.

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

Step 5. Log into the Cisco Intersight portal as a user with account administrator role.

Step 6. From the Service Selector drop-down list, select System.

Step 7. Navigate to Settings > General > Resource Groups.

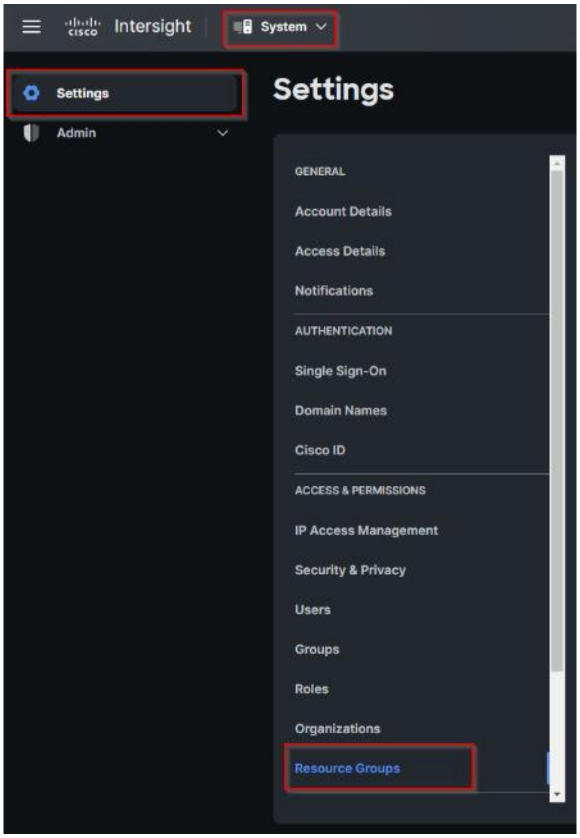

Step 8. On Resource Groups panel click + Create Resource Group in the top-right corner.

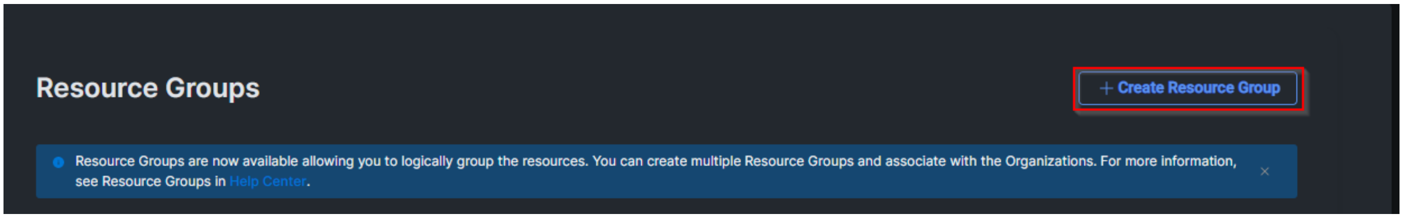

Resource Groups are now available allowing you to logically group the resources. You can create multiple Resource Groups and associate with the Organizations. For more information, see Resource Groups in Help Center.

Step 9. Provide a name for the Resource Group (for example, L151-DMZ). You can select either ALL or Custom.

Step 10. Click Create.

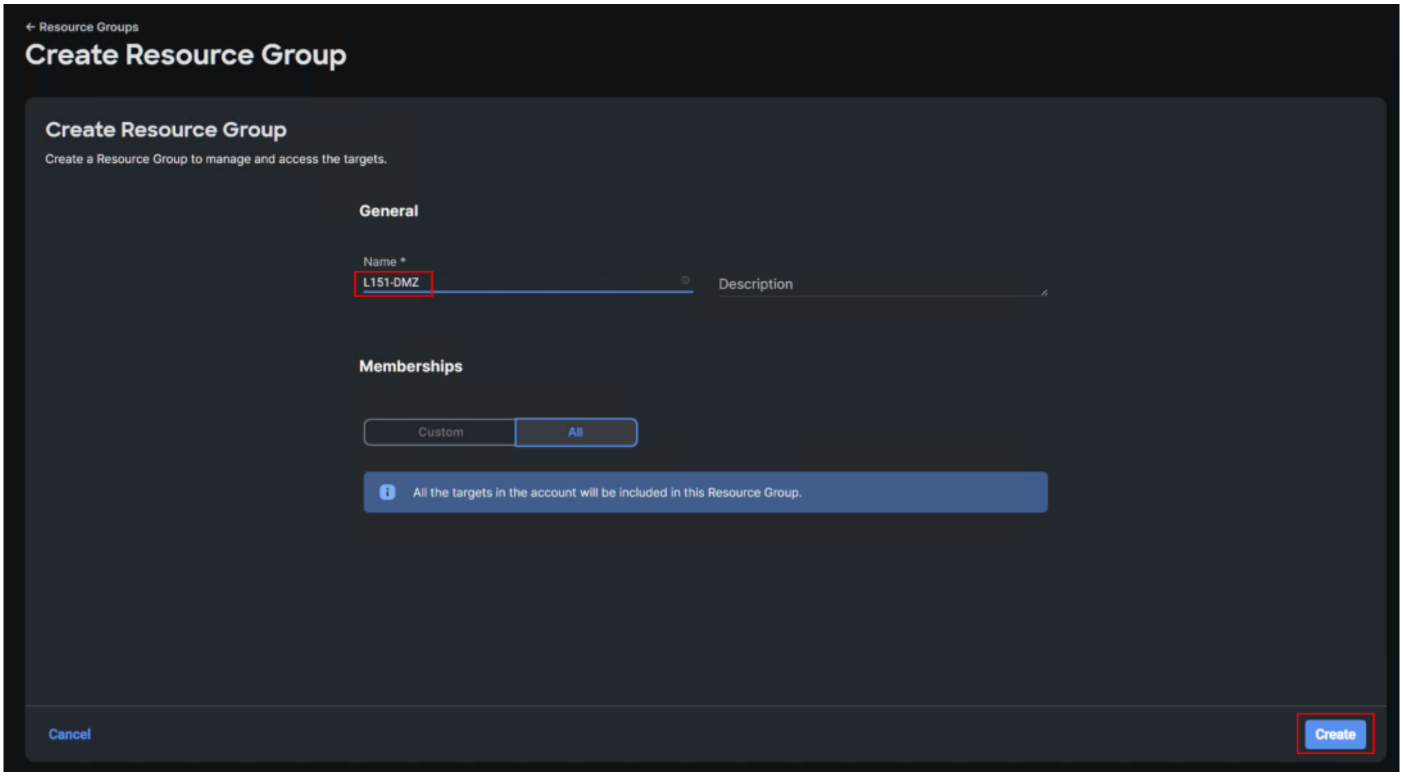

← Resource Groups

Create Resource Group

Create Resource Group

Create a Resource Group to manage and access the targets.

General

Name * L151-DMZ

Description

Memberships

Custom | All

All the targets in the account will be included in this Resource Group.

Step 11. Navigate to Settings > General > Organizations.

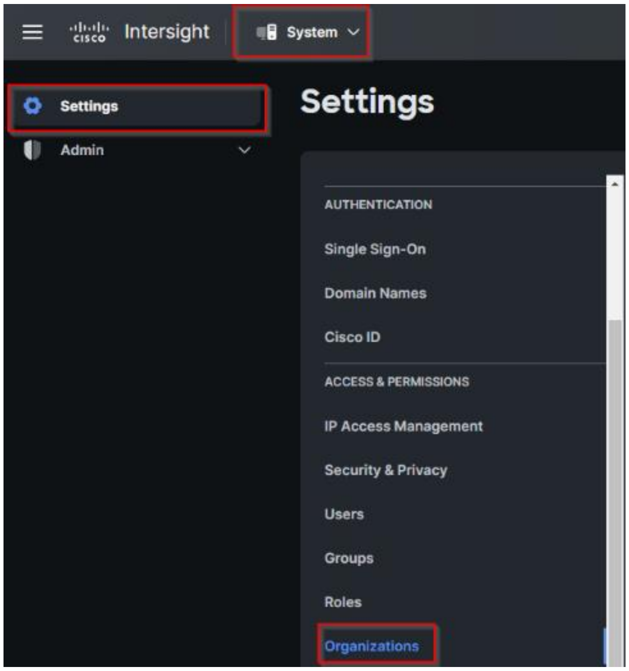

Step 12. On Organizations panel click + Create Organization in the top-right corner.

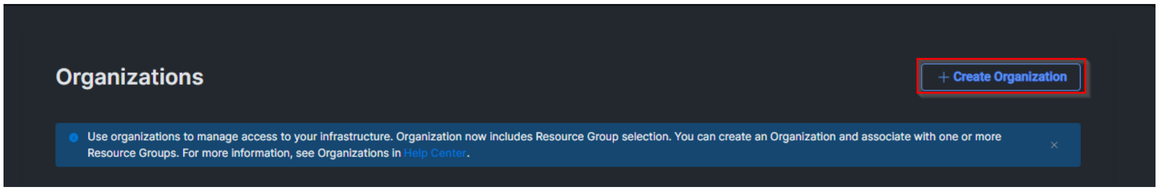
Organizations

+ Create Organization

Use organizations to manage access to your infrastructure. Organization now includes Resource Group selection. You can create an Organization and associate with one or more Resource Groups. For more information, see Organizations in Help Center.

Step 13. Provide a name for the organization in your environment (SJC2-L151).

Step 14. Select the Resource Group created in the last step (for example, SJC2-L151).

Step 15. Click Create.

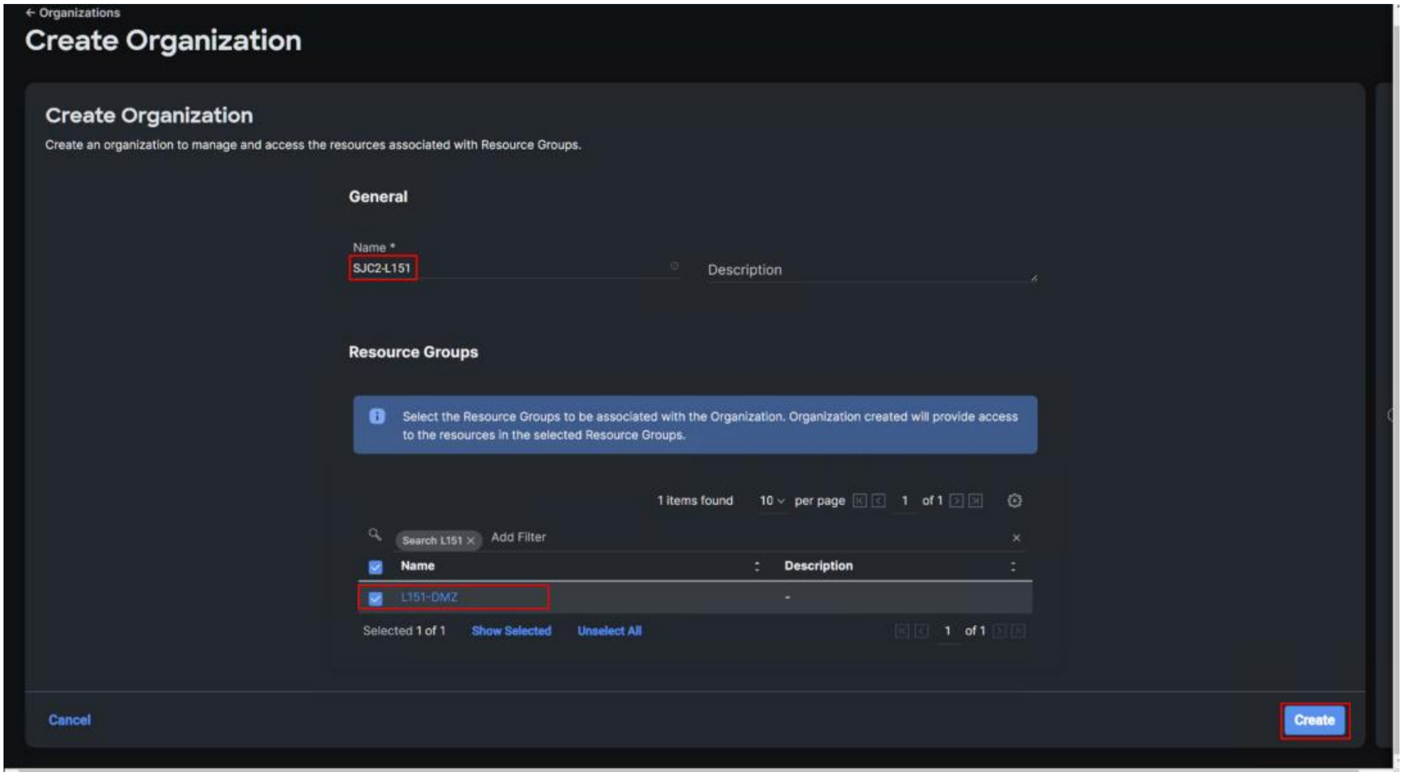
← Organizations

Create Organization

Create Organization

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

General

Name * SJC2-L151

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.

1 items found 10 per page 1 of 1

Search L151 × Add Filter

☑ Name	Description
☑ L151-DMZ	-

Selected 1 of 1 Show Selected Unselect All 1 of 1

Cancel Create

Step 16. 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 17. 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.

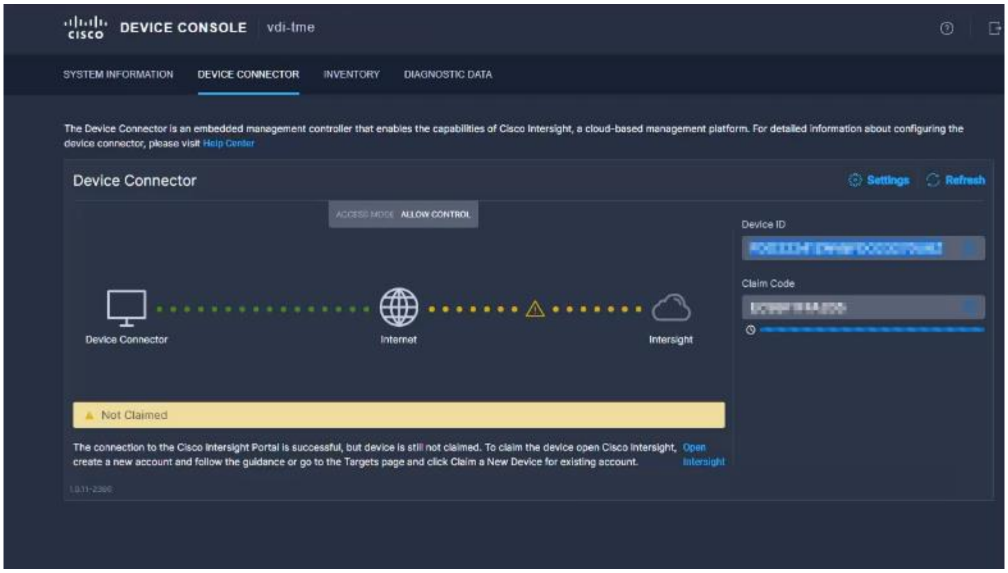

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

Device Connector

Device ID [redacted]

Claim Code [redacted]

Not Claimed

The connection to the Cisco Intersight Portal is successful, but device is still not claimed. To claim the device open Cisco Intersight, create a new account and follow the guidance or go to the Targets page and click Claim a New Device for existing account. Open Intersight

Step 18. Navigate to Admin > General > Targets.

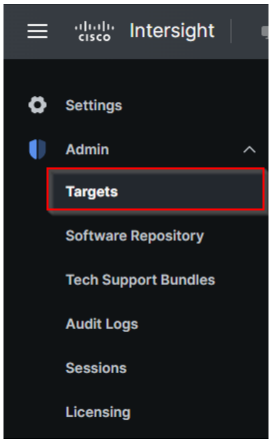

Step 19. On Targets panel click Claim a New Target in the top-right corner.

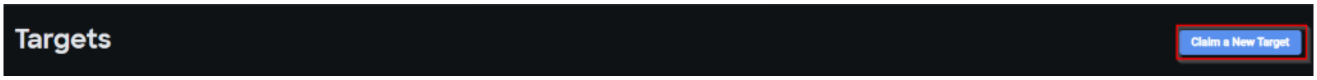

Step 20. Select Cisco UCS Domain (Intersight Managed) and click Start.

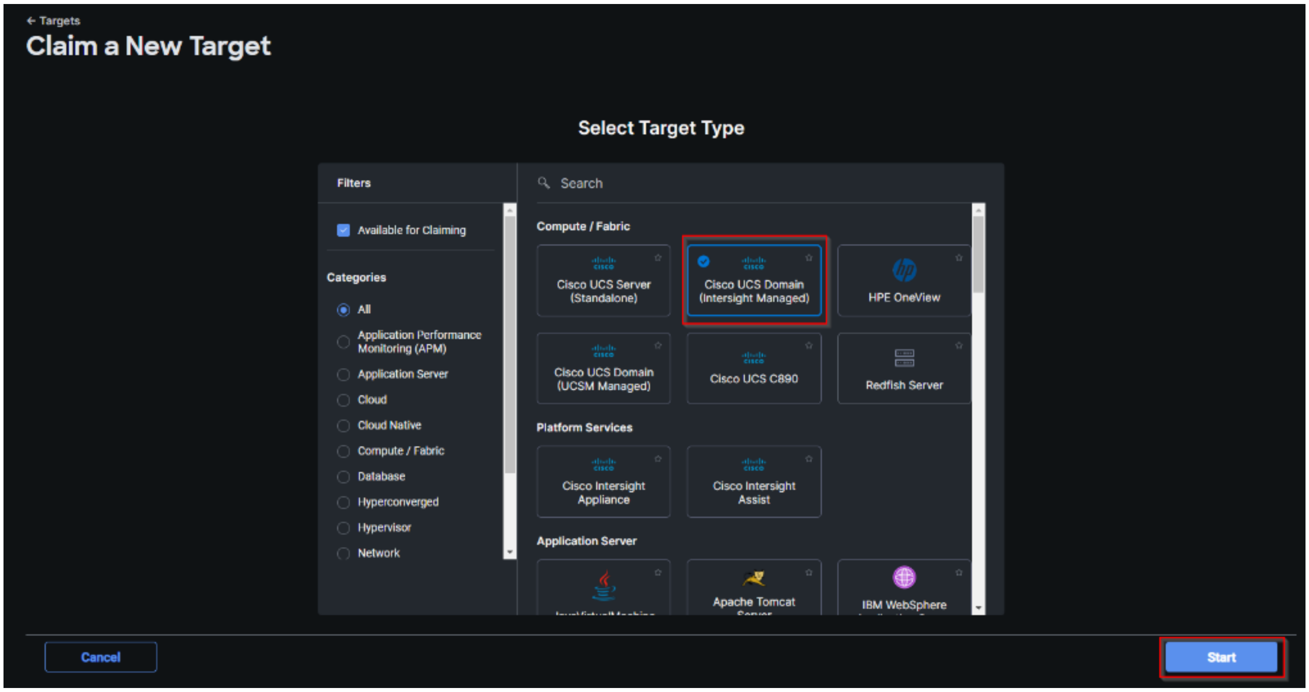
← Targets

Claim a New Target

Select Target Type

Step 21. Enter the Device ID and Claim Code captured from the Cisco UCS FI.

Step 22. Select the previously created Resource Group and click Claim.

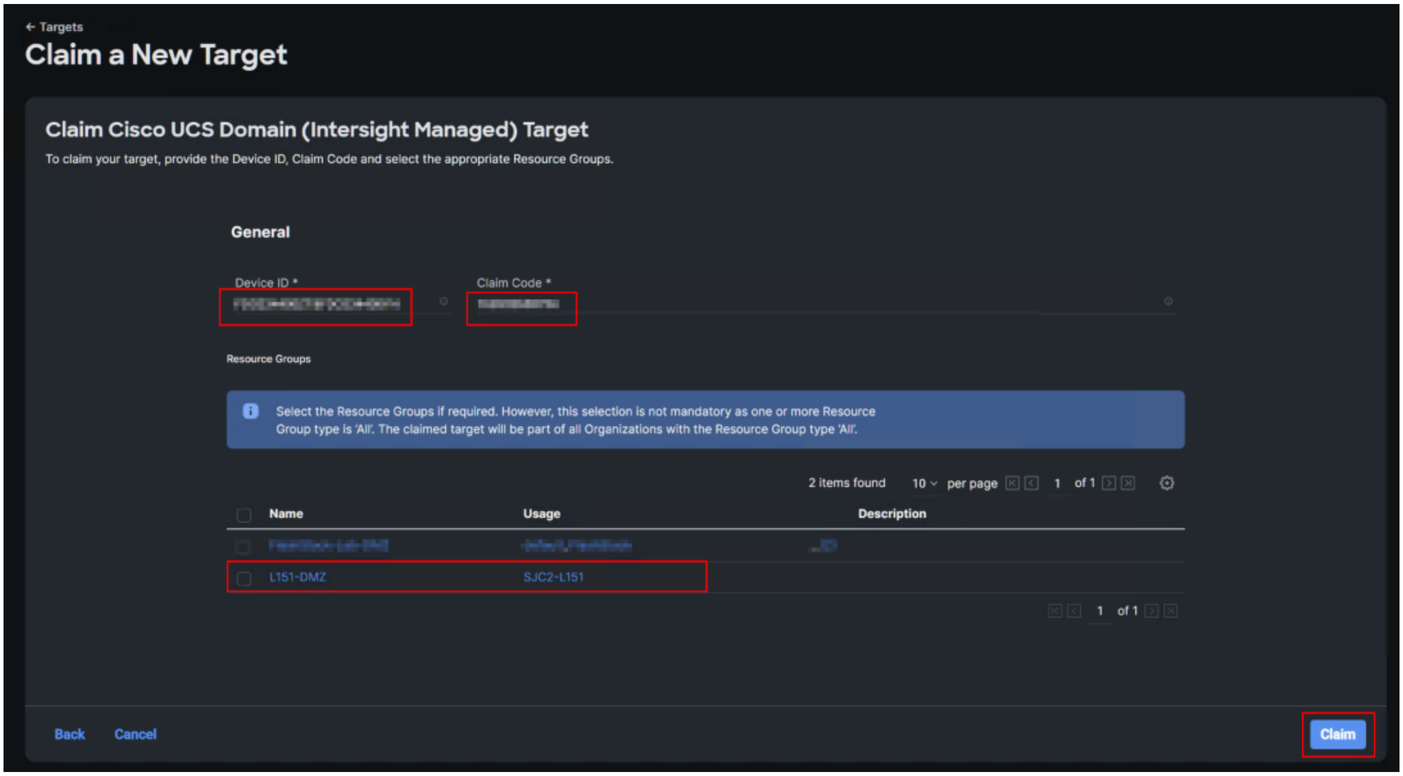
← Targets

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 * [redacted]

Claim Code * [redacted]

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'.

2 items found 10 per page 1 of 1

	Name	Usage	Description
□	[redacted]	[redacted]	[redacted]
□	L151-DMZ	SJC2-L151	

Step 23. On successfully device claim, Cisco UCS FI should appear as a target in Cisco Intersight.

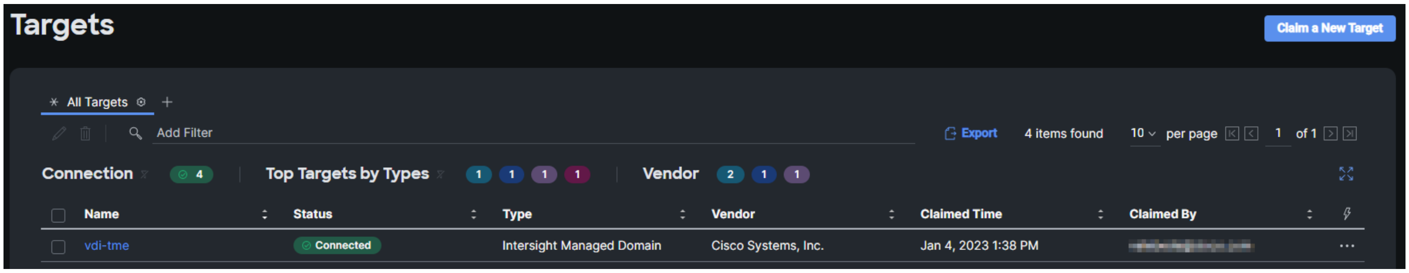

Targets

Name	Status	Type	Vendor	Claimed Time	Claimed By
vdi-tme	Connected	Intersight Managed Domain	Cisco Systems, Inc.	Jan 4, 2023 1:38 PM	[redacted]

Configure a Cisco UCS Domain Profile

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 and VSANs 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.

After the Cisco UCS domain profile has been successfully created and deployed, the policies including the port policies are pushed to Cisco UCS Fabric Interconnects. Cisco UCS domain profile can easily be cloned to install additional Cisco UCS systems. When cloning the UCS domain profile, the new UCS domains utilize the existing policies for consistent deployment of additional Cisco UCS systems at scale.

Note: A number of Domain policies required for the UCS Domain Profile. These policies can be created beforehand or created and assigned at UCS Domain Profile creation.

Procedure 1. Create a Domain Profile

Step 1. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Configure > Profiles, to launch the Profiles Table view.

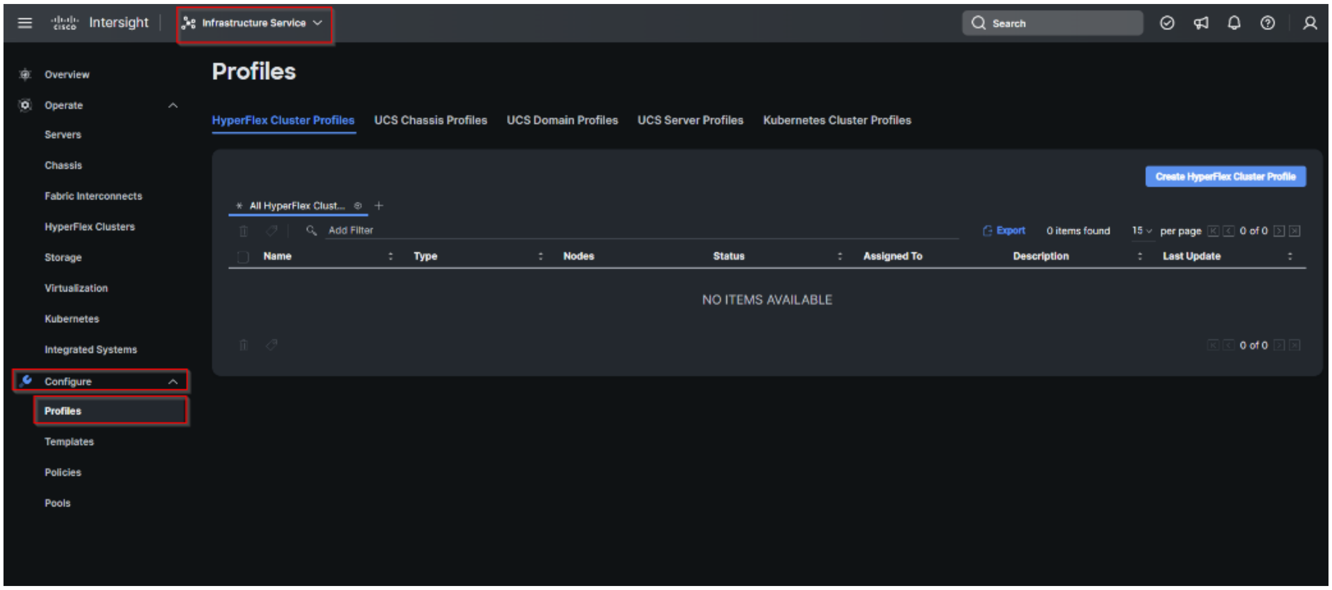

Step 2. Navigate UCS Domain Profiles tab and click Create UCS Domain Profile.

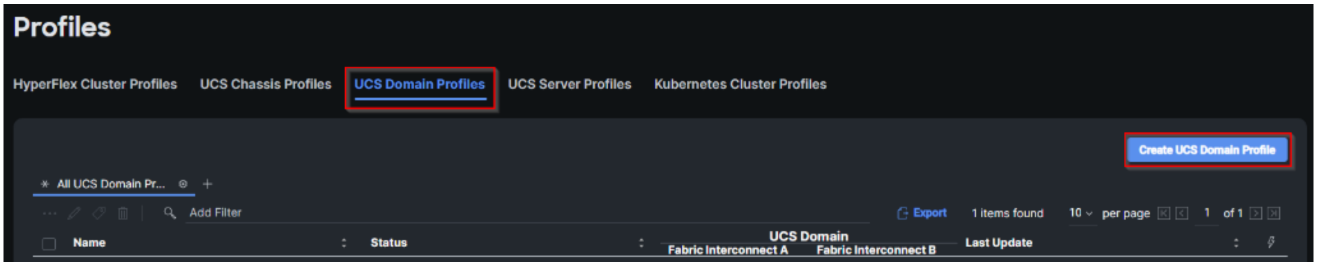

Step 3. On the Create UCS Domain Profile screen, click Start.

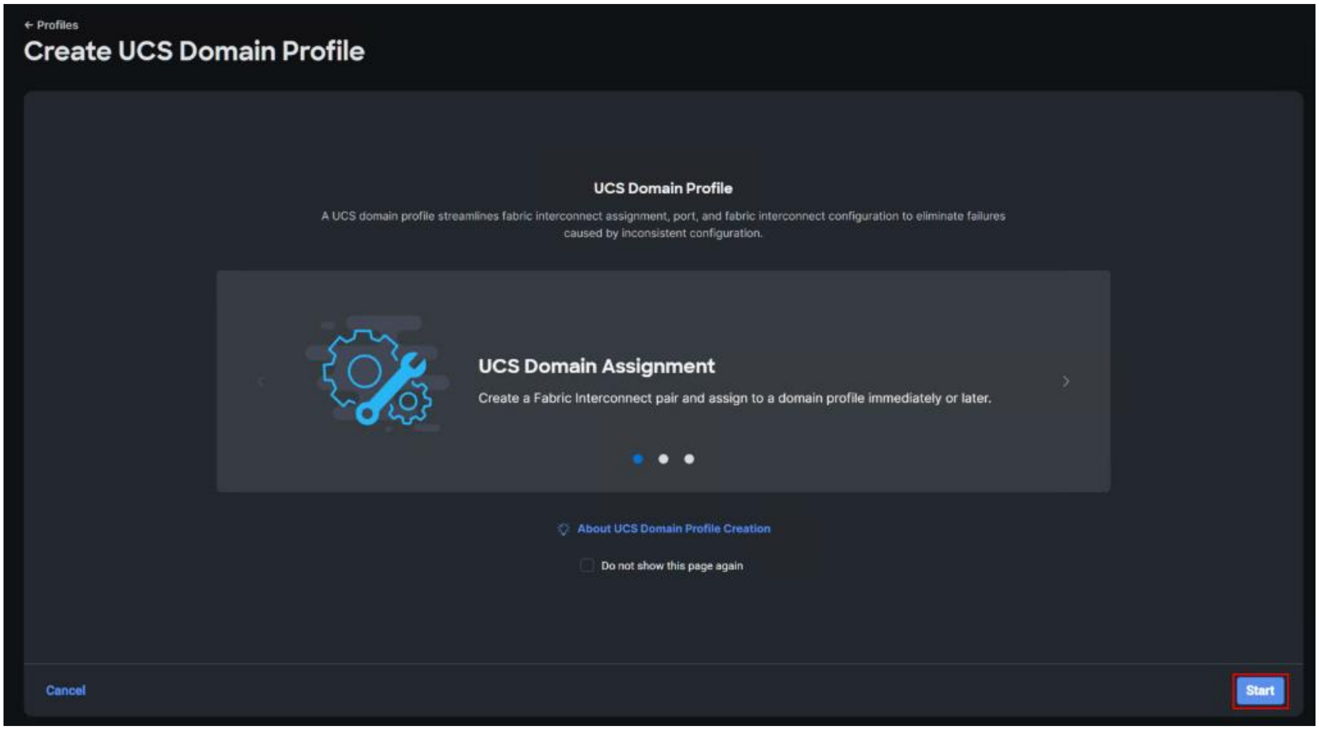
← Profiles
Create UCS Domain Profile
UCS Domain Profile
UCS Domain Assignment
About UCS Domain Profile Creation
Do not show this page again
Cancel
Start

A UCS domain profile streamlines fabric interconnect assignment, port, and fabric interconnect configuration to eliminate failures caused by inconsistent configuration.

Create a Fabric Interconnect pair and assign to a domain profile immediately or later.

Step 4. On the General page, select the organization created before and enter a name for your profile (for example, L152-J5). Optionally, include a short description and tag information to help identify the profile. Tags must be in the key:value format. For example, Org: IT or Site: APJ. Click Next.

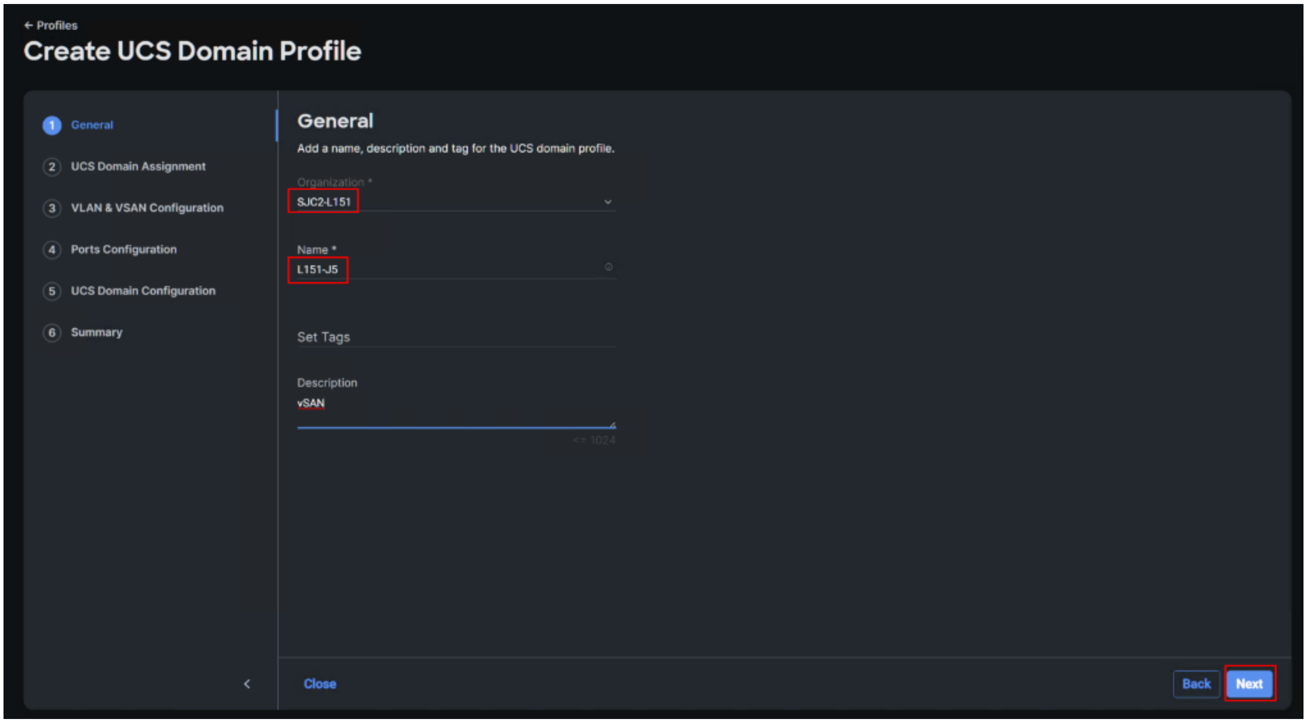

Create UCS Domain Profile

General

Add a name, description and tag for the UCS domain profile.

Organization * SJC2-L151

Name * L151-J5

Set Tags

Description vSAN

<= 1024

Step 5. On the Domain Assignment page, assign a switch pair to the Domain profile. Click Next.

Note: You can also click Assign Later and assign a switch pair to the Domain profile at a later time.

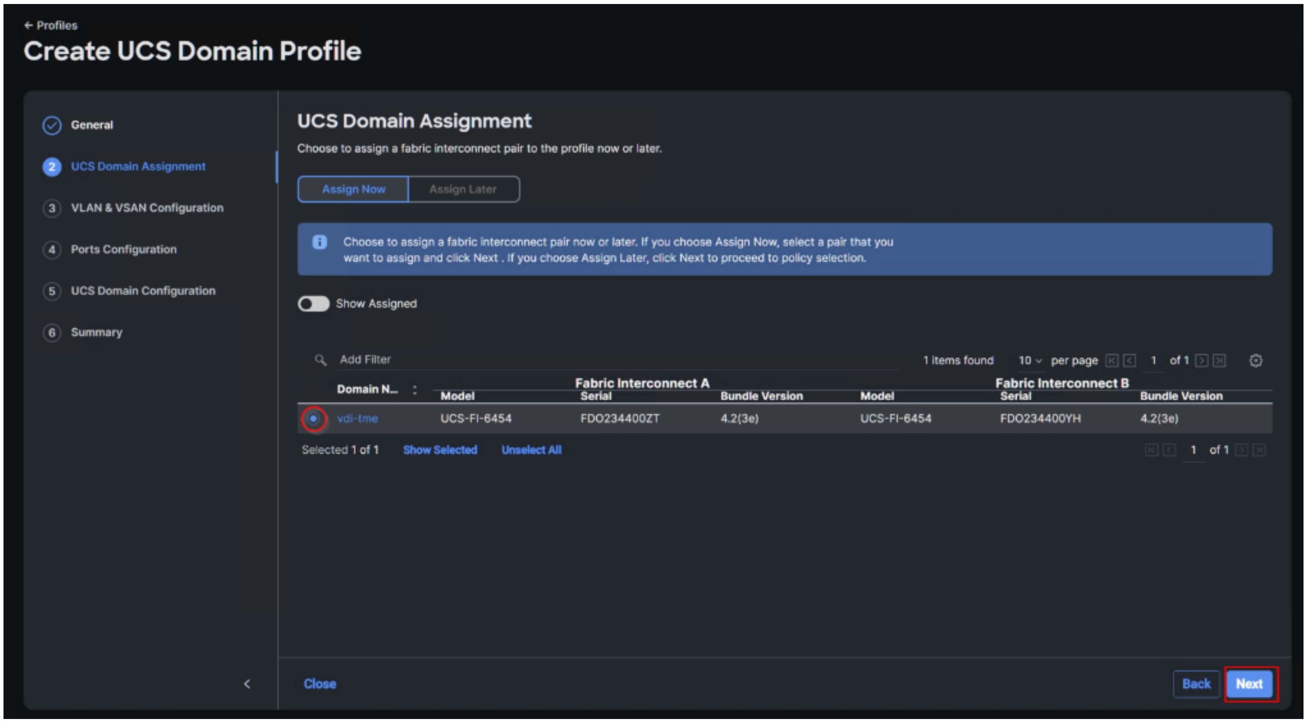

Create UCS Domain Profile

UCS Domain Assignment

Choose to assign a fabric interconnect pair to the profile now or later.

Choose to assign a fabric interconnect pair now or later. If you choose Assign Now, select a pair that you want to assign and click Next . If you choose Assign Later, click Next to proceed to policy selection.

1 items found

		Fabric Interconnect A			Fabric Interconnect B		
	Domain N...	Model	Serial	Bundle Version	Model	Serial	Bundle Version
●	vdi-tme	UCS-FI-6454	FDO234400ZT	4.2(3e)	UCS-FI-6454	FDO234400YH	4.2(3e)

Selected 1 of 1

Step 6. On the VLAN & VSAN Configuration page, attach VLAN and VSAN policies for each switch to the UCS Domain Profile.

Note: A VSAN policy is only needed when configuring Fibre Channel and can be skipped when configuring IP-only storage access.

Note: In this step, a single VLAN policy is created for both fabrics.

Step 7. Click Select Policy next to VLAN Configuration under Fabric Interconnect A.

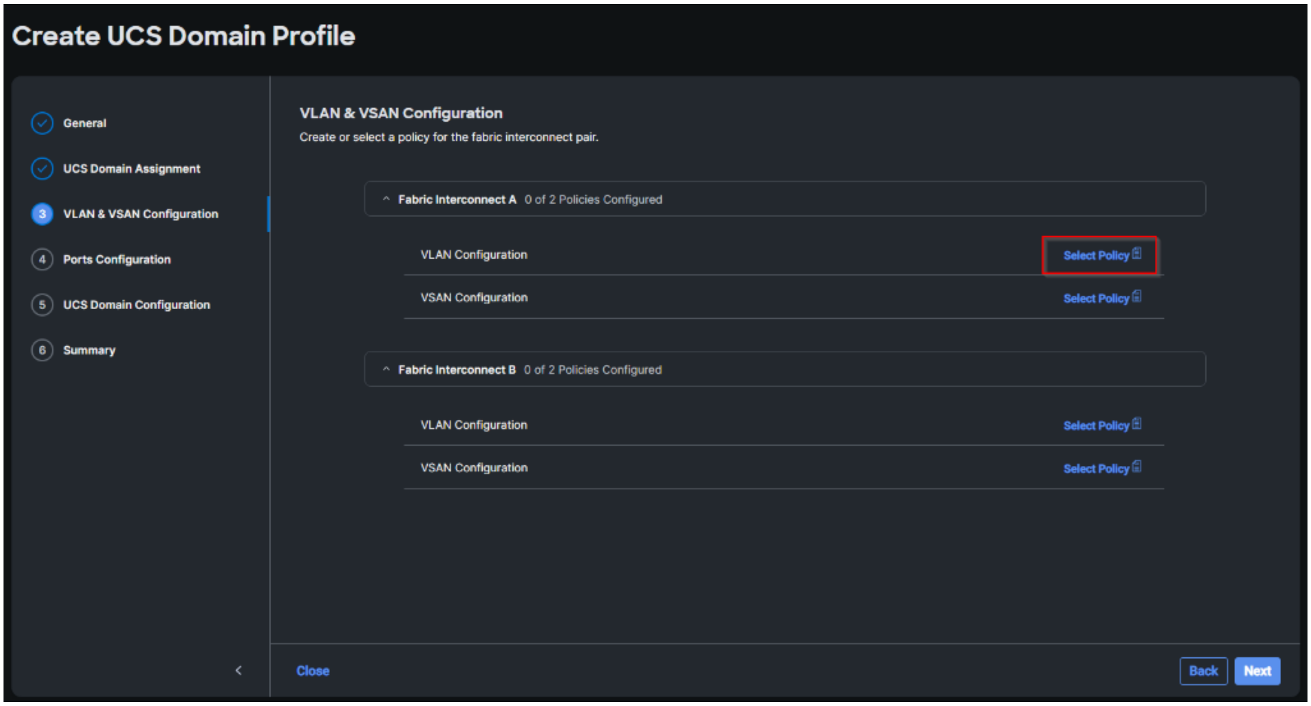

Create UCS Domain Profile

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 8. In the pane on the right, click Create New.

Step 9. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, L152-J5-VLANs). Click Next.

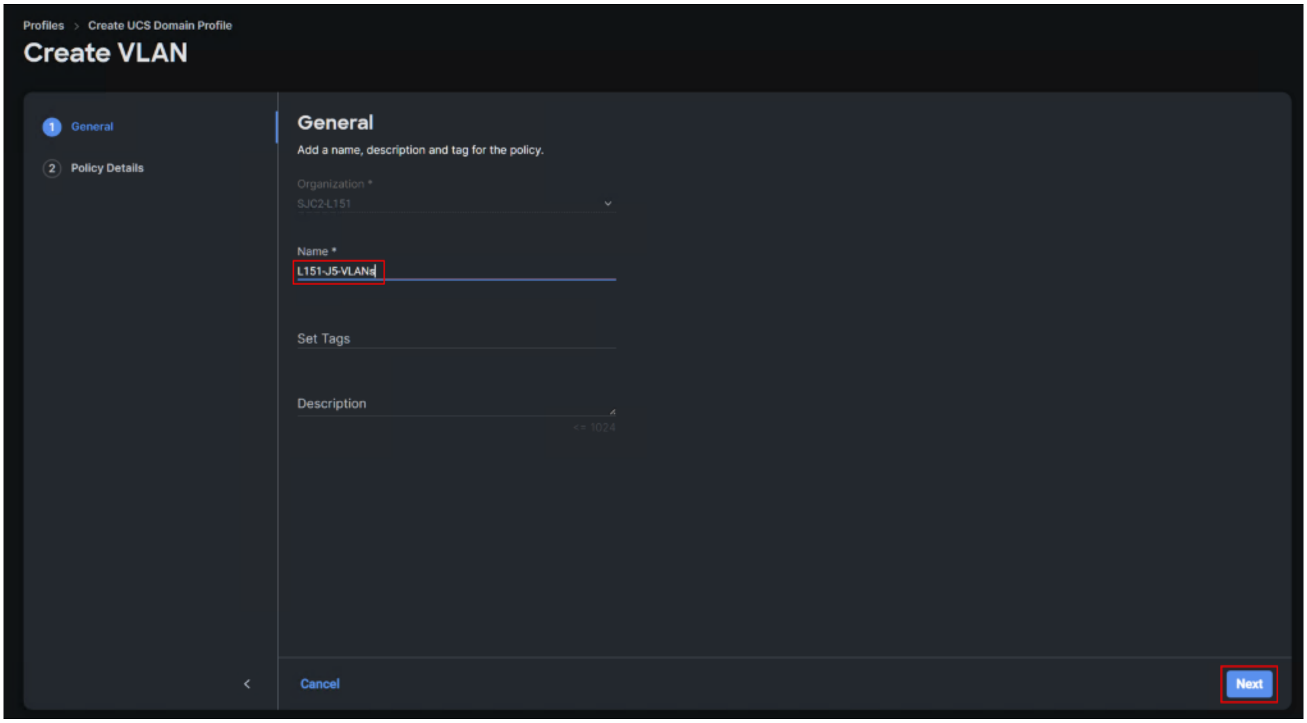

Profiles > Create UCS Domain Profile

Create VLAN

General

Add a name, description and tag for the policy.

Organization *
SJC2-L151

Name *
L151-J5-VLANs

Set Tags

Description
<= 1024

Step 10. Click Add VLANs.

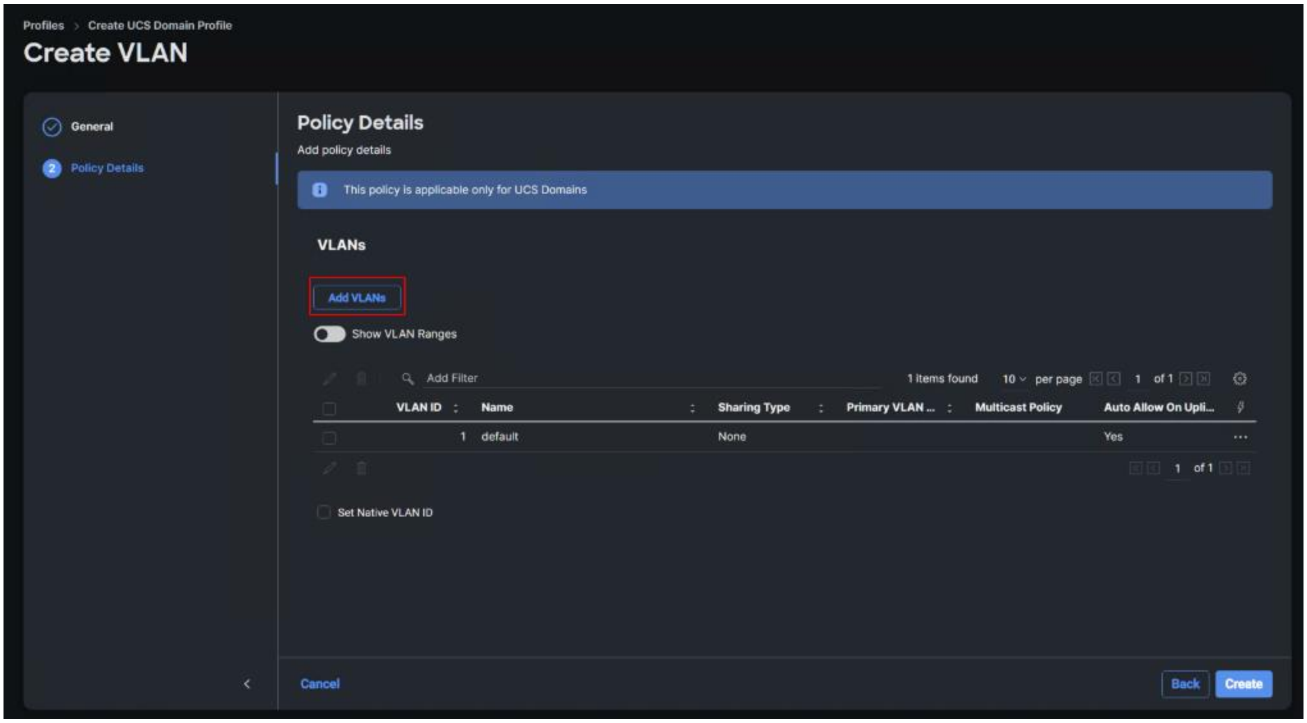

Profiles > Create UCS Domain Profile

Create VLAN

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

Show VLAN Ranges

Add Filter | 1 items found | 10 per page | 1 of 1

VLAN ID	Name	Sharing Type	Primary VLAN ...	Multicast Policy	Auto Allow On Upli...
1	default	None			Yes

1 of 1

□ Set Native VLAN ID

Step 11. Provide a name and VLAN ID for the VLAN from you list (for example, 70, 71, 72, 73, 74). Enable Auto Allow On Uplinks. To create the required Multicast policy, click Select Policy under Multicast*.

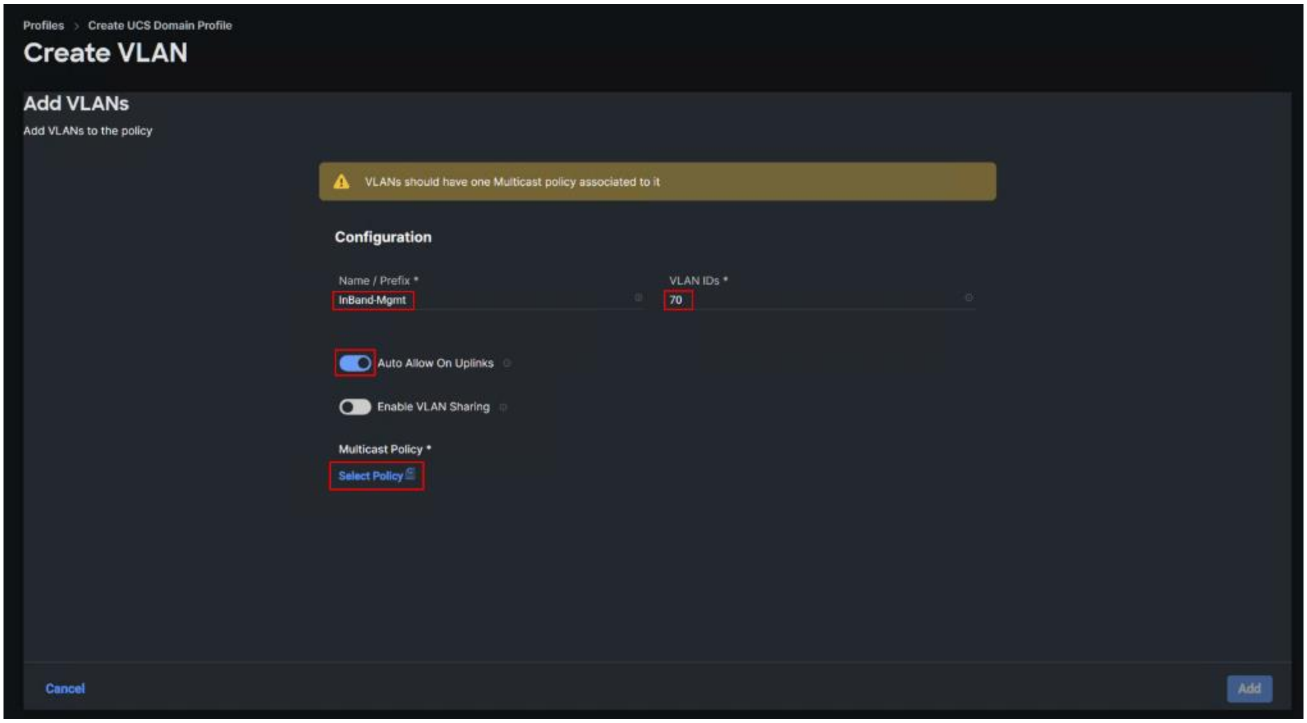
Profiles > Create UCS Domain Profile

Create VLAN

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix *: InBand-Mgmt

VLAN IDs *: 70

Auto Allow On Uplinks: enabled

Enable VLAN Sharing: disabled

Multicast Policy *: Select Policy

Step 12. In the window on the right, click Create New to create a new Multicast Policy.

Step 13. Provide a Name for the Multicast Policy (for example, L152-J5-McastPol). Provide optional Description and click Next.

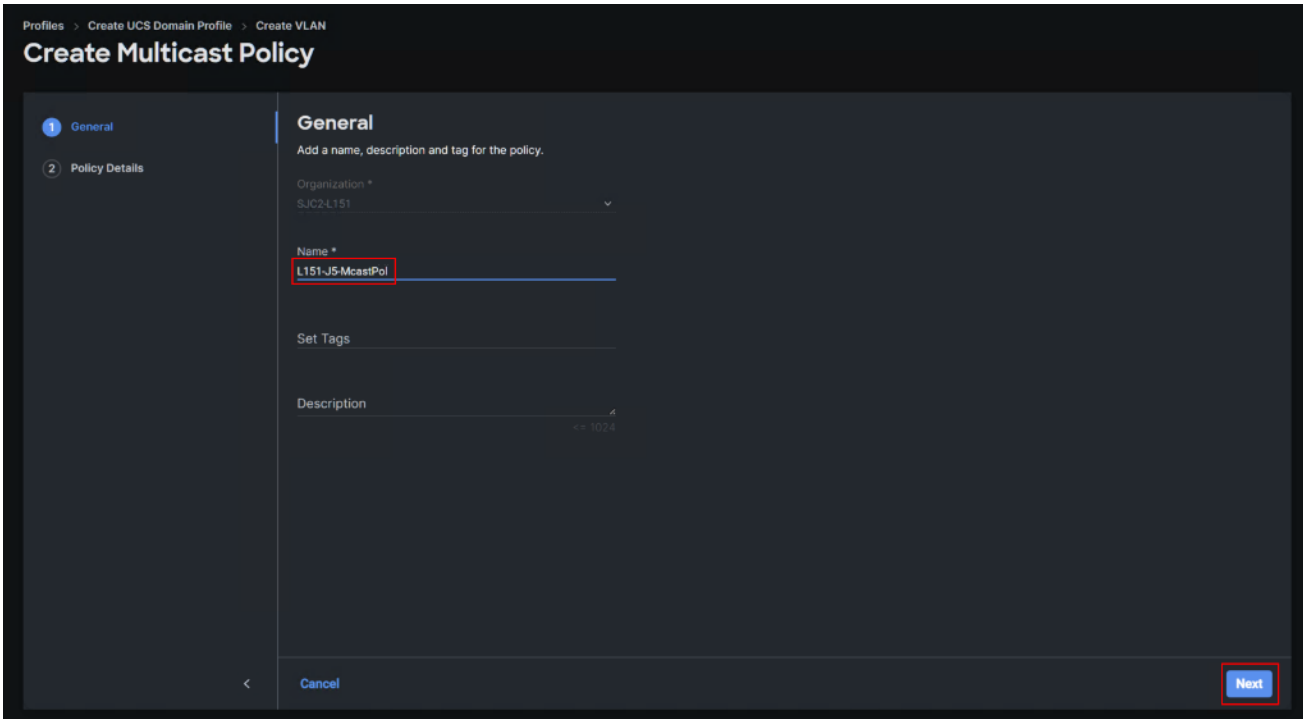
Profiles > Create UCS Domain Profile > Create VLAN

Create Multicast Policy

General

Add a name, description and tag for the policy.

Organization *: SJC2-L151

Name *: L151-J5-McastPol

Set Tags

Description

<= 1024

Step 14. Leave defaults selected and click Create.

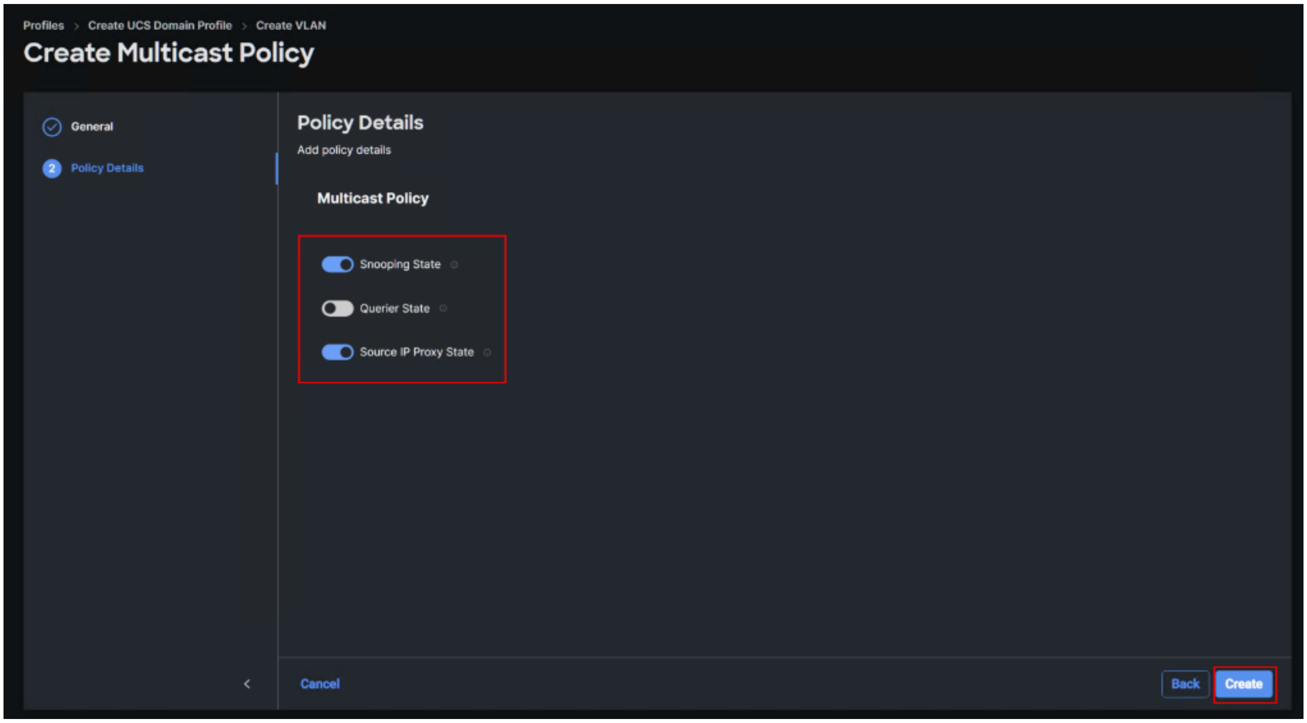

Profiles > Create UCS Domain Profile > Create VLAN

Create Multicast Policy

General

2 Policy Details

Policy Details

Add policy details

Multicast Policy

Snooping State: on

Querier State: off

Source IP Proxy State: on

Step 15. Click Add to add the VLAN.

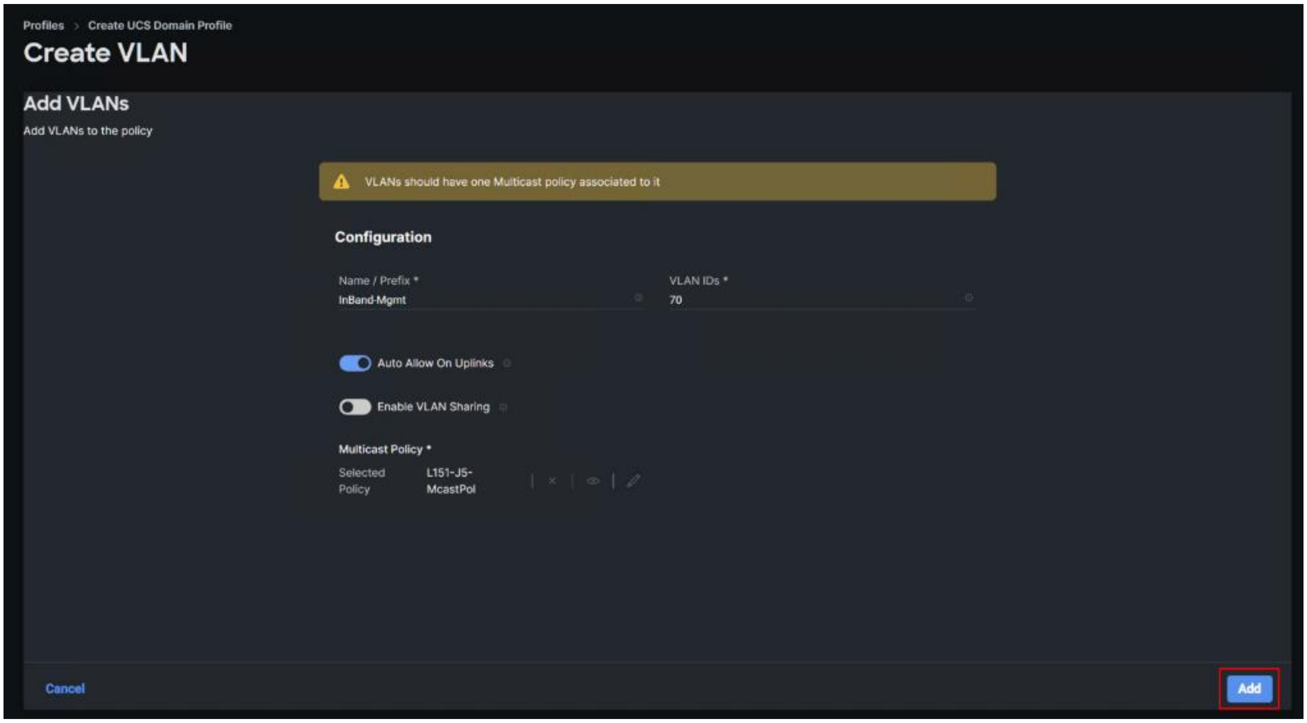

Profiles > Create UCS Domain Profile

Create VLAN

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix * InBand-Mgmt

VLAN IDs * 70

Auto Allow On Uplinks: on

Enable VLAN Sharing: off

Multicast Policy *

Selected Policy L151-J5-McastPol

Step 16. Add the remaining VLANs from you list 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:

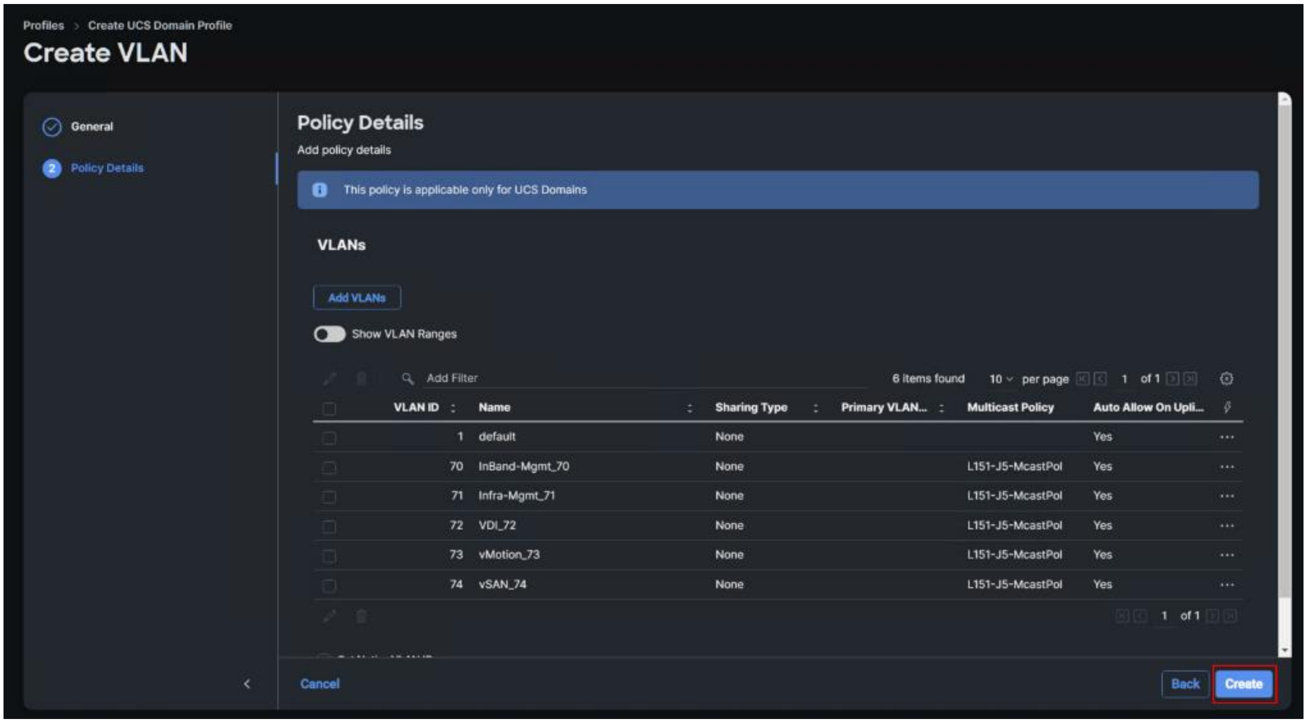

Profiles > Create UCS Domain Profile

Create VLAN

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

6 items found 10 per page 1 of 1

VLAN ID	Name	Sharing Type	Primary VLAN...	Multicast Policy	Auto Allow On Upli...
1	default	None			Yes
70	InBand-Mgmt_70	None		L151-J5-McastPol	Yes
71	Infra-Mgmt_71	None		L151-J5-McastPol	Yes
72	VDI_72	None		L151-J5-McastPol	Yes
73	vMotion_73	None		L151-J5-McastPol	Yes
74	vSAN_74	None		L151-J5-McastPol	Yes

Step 17. Repeat steps 7 and 8 for fabric interconnect B assigning the VLAN policy created previously.

Step 18. Verify that a common VLAN policy is associated with the two fabric interconnects. Click Next.

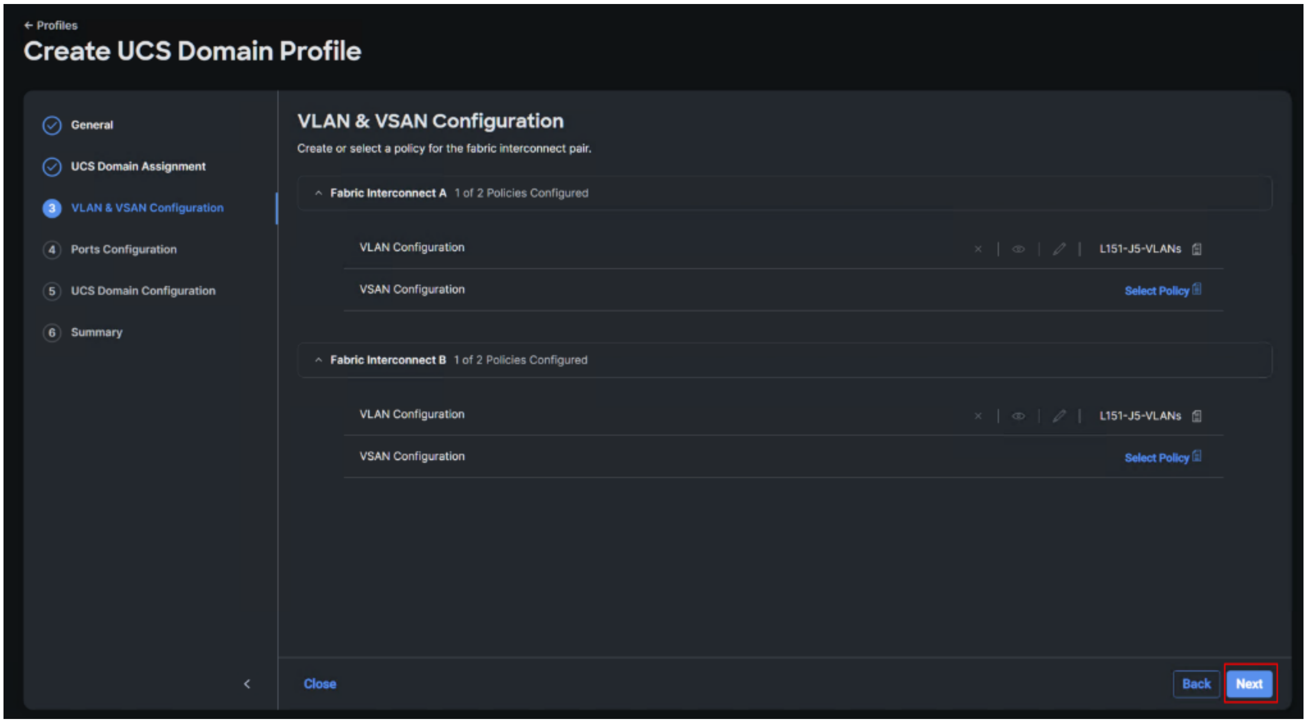

← Profiles

Create UCS Domain Profile

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 1 of 2 Policies Configured

VLAN Configuration L151-J5-VLANs

VSAN Configuration Select Policy

Fabric Interconnect B 1 of 2 Policies Configured

VLAN Configuration L151-J5-VLANs

VSAN Configuration Select Policy

Step 19. On the Ports Configuration page, attach port policies for each switch to the UCS Domain Profile.

Note: Use two separate port policies for the fabric interconnects. Using separate policies provides flexibility when port configuration (port numbers or speed) differs between the two FIs.

Step 20. Click Select Policy for Fabric Interconnect A.

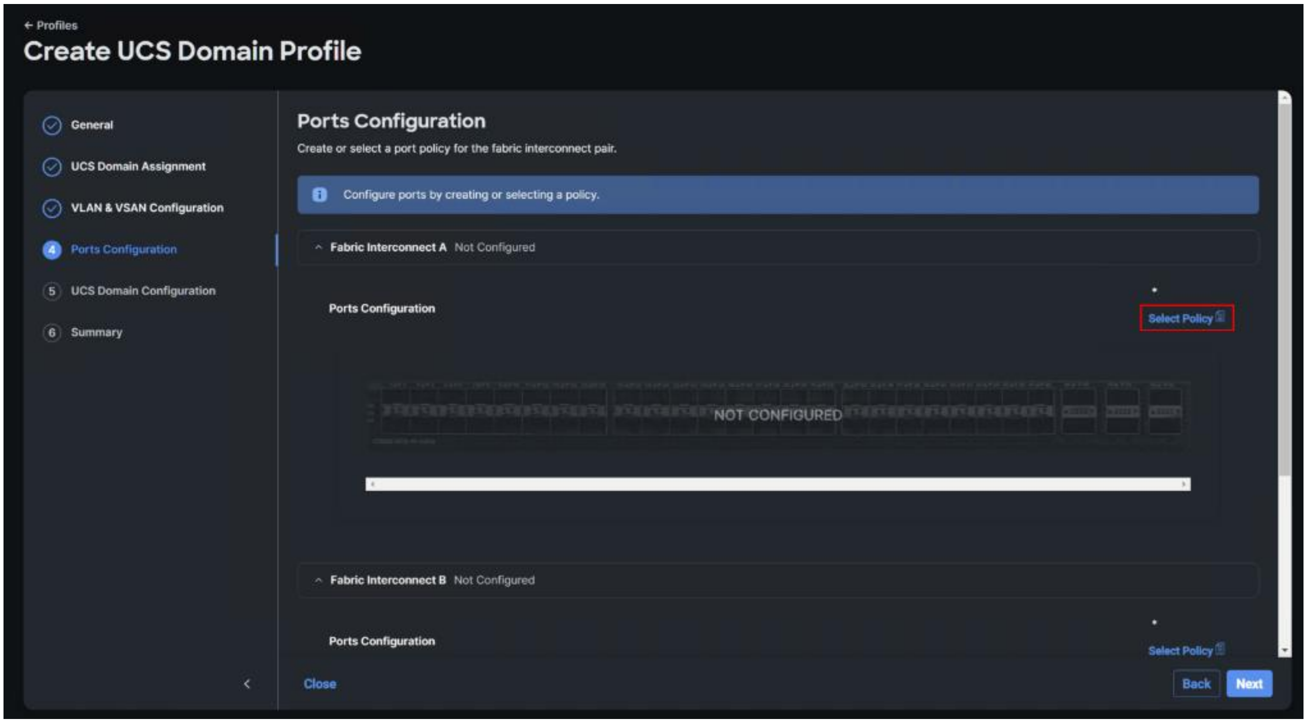

← Profiles

Create UCS Domain Profile

General

UCS Domain Assignment

VLAN & VSAN Configuration

4 Ports Configuration

5 UCS Domain Configuration

6 Summary

Ports Configuration

Create or select a port policy for the fabric interconnect pair.

Configure ports by creating or selecting a policy.

Fabric Interconnect A Not Configured

Ports Configuration

Select Policy

NOT CONFIGURED

Fabric Interconnect B Not Configured

Ports Configuration

Select Policy

Close

Back

Next

Step 21. Click Create New.

Step 22. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L152-J5-FI-A). Click Next.

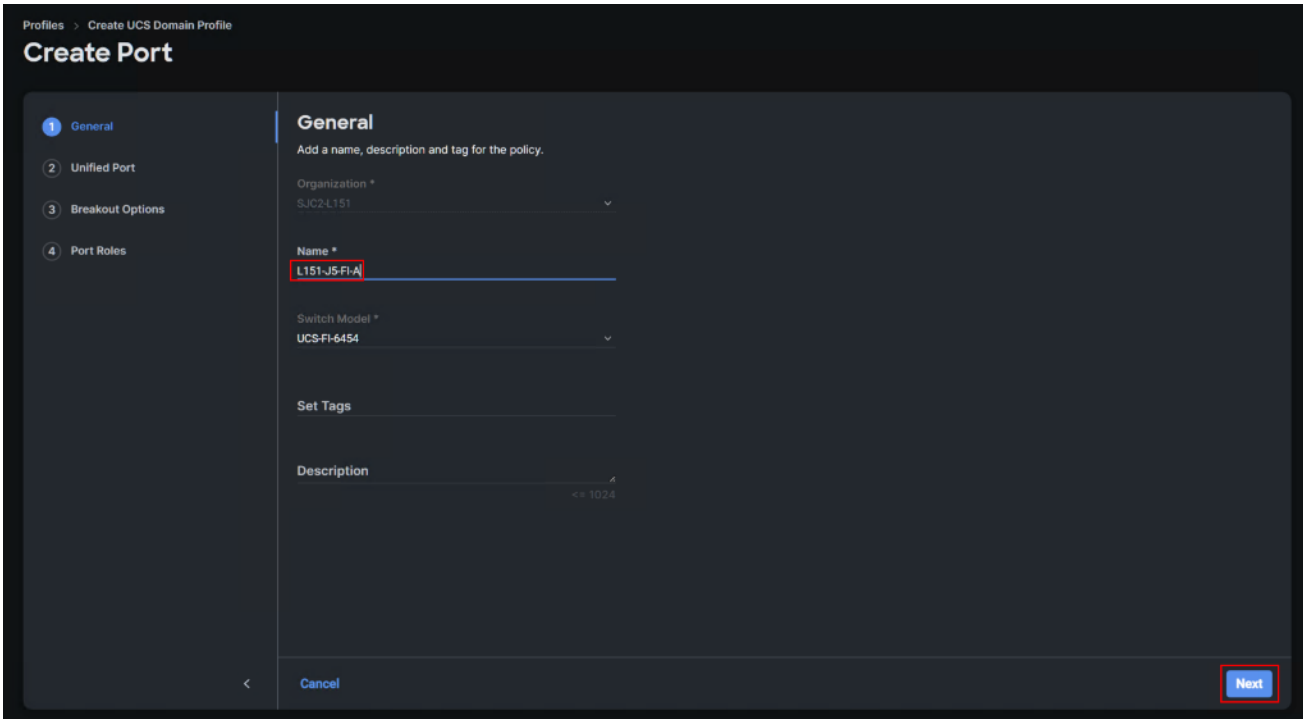

Profiles > Create UCS Domain Profile

Create Port

1 General

2 Unified Port

3 Breakout Options

4 Port Roles

General

Add a name, description and tag for the policy.

Organization *
SJC2-L151

Name *
L151-J5-FI-A

Switch Model *
UCS-FI-6454

Set Tags

Description

<= 1024

Cancel

Next

Step 23. On the unified ports page click Next.

Note: No Fibre Channel ports were used in this validation.

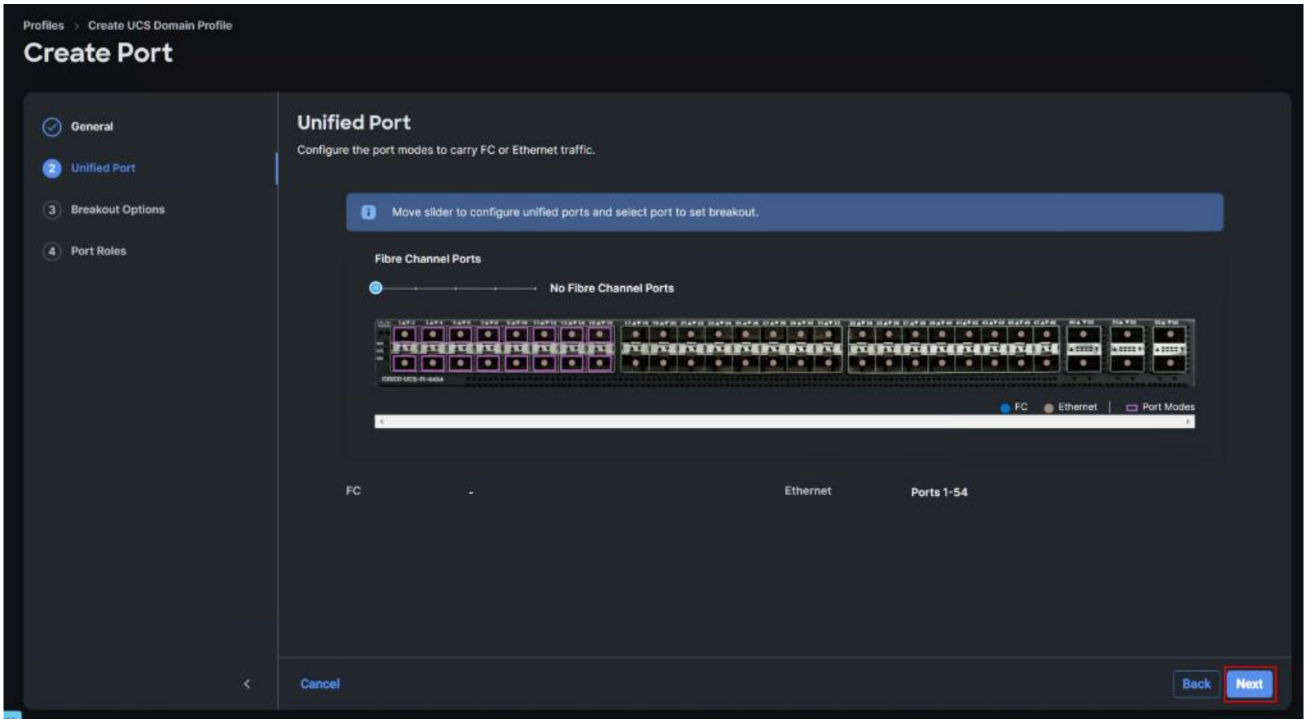

Profiles > Create UCS Domain Profile

Create Port

- General
- 2 Unified Port
- 3 Breakout Options
- 4 Port Roles

Unified Port

Configure the port modes to carry FC or Ethernet traffic.

Move slider to configure unified ports and select port to set breakout.

Fibre Channel Ports

No Fibre Channel Ports

● FC ● Ethernet | ▭ Port Modes

FC -

Ethernet Ports 1-54

Cancel

Back Next

Step 24. On the Breakout Options page click Next.

Note: No Ethernet/Fibre Channel breakouts were used in this validation.

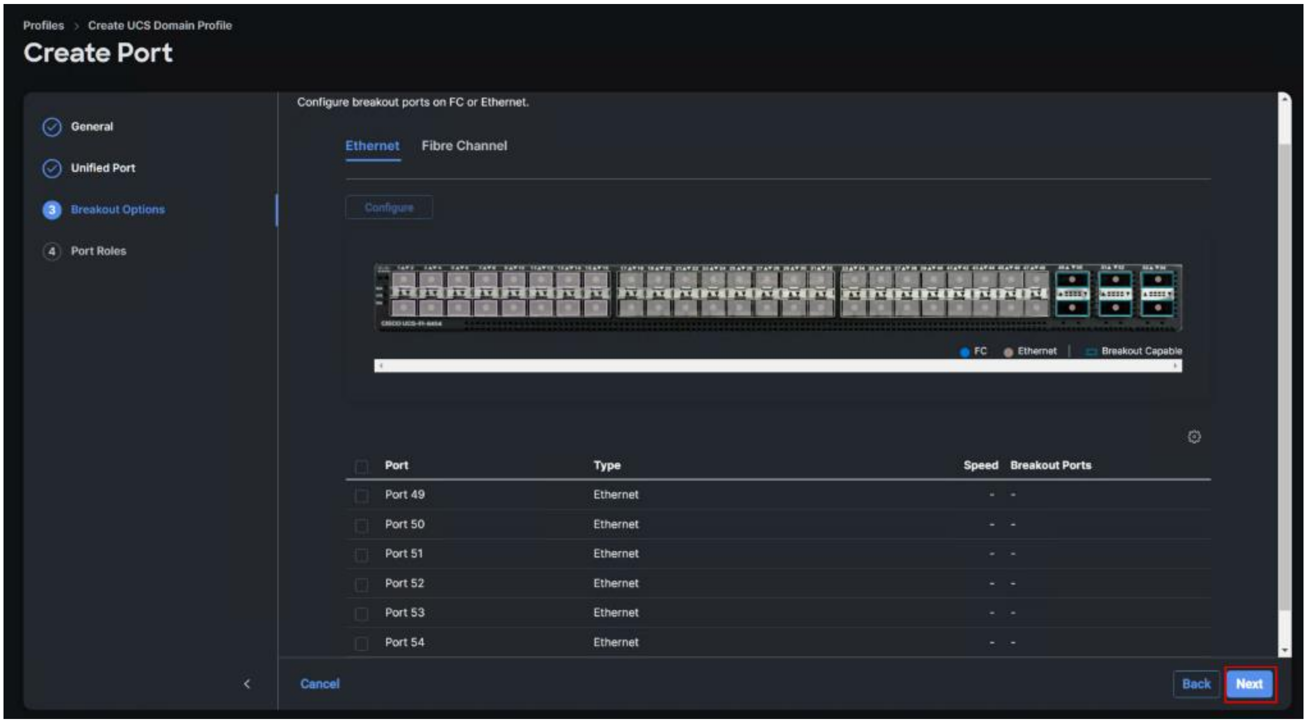

Profiles > Create UCS Domain Profile

Create Port

- General
- Unified Port
- 3 Breakout Options
- 4 Port Roles

Configure breakout ports on FC or Ethernet.

Ethernet Fibre Channel

Configure

● FC ● Ethernet | ▭ Breakout Capable

	Port	Type	Speed	Breakout Ports
□	Port 49	Ethernet	-	-
□	Port 50	Ethernet	-	-
□	Port 51	Ethernet	-	-
□	Port 52	Ethernet	-	-
□	Port 53	Ethernet	-	-
□	Port 54	Ethernet	-	-

Cancel

Back Next

Step 25. 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.

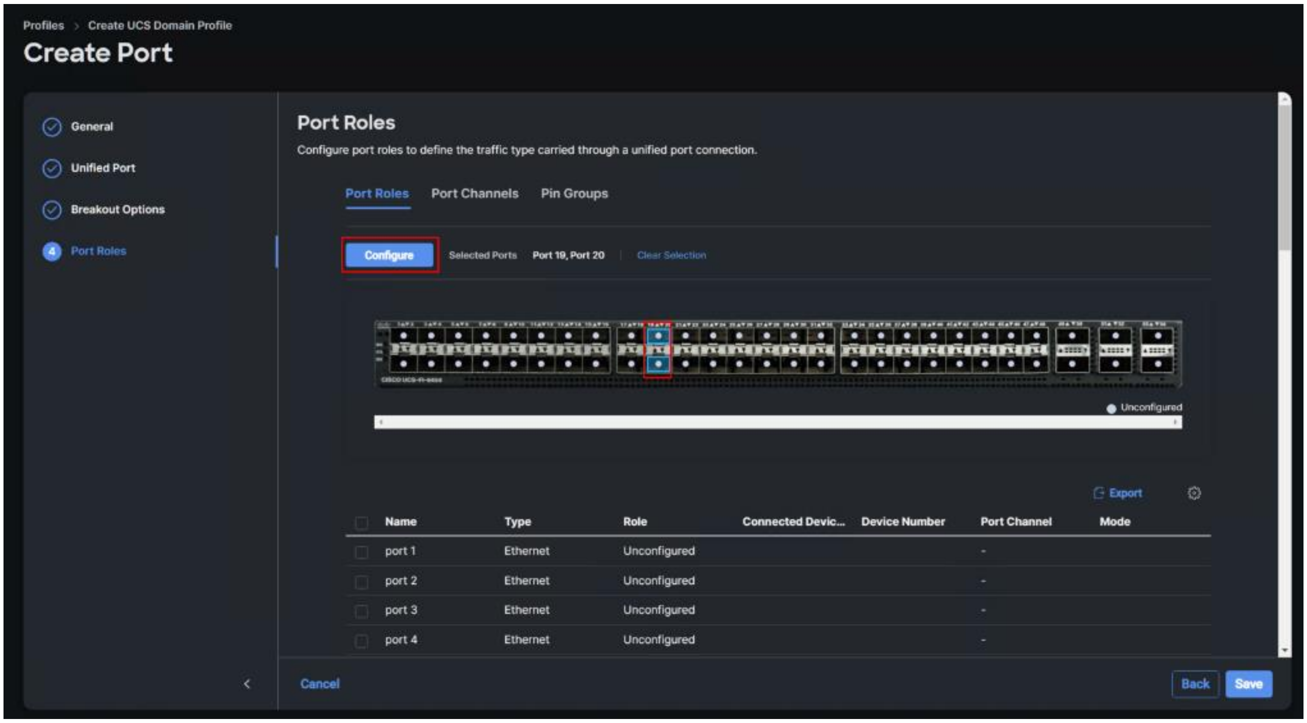

Profiles > Create UCS Domain Profile

Create Port

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

Selected Ports Port 19, Port 20

Name	Type	Role	Connected Devic...	Device Number	Port Channel	Mode
port 1	Ethernet	Unconfigured			-	
port 2	Ethernet	Unconfigured			-	
port 3	Ethernet	Unconfigured			-	
port 4	Ethernet	Unconfigured			-	

Step 26. From the drop-down list, select Server as the role. Click Save.

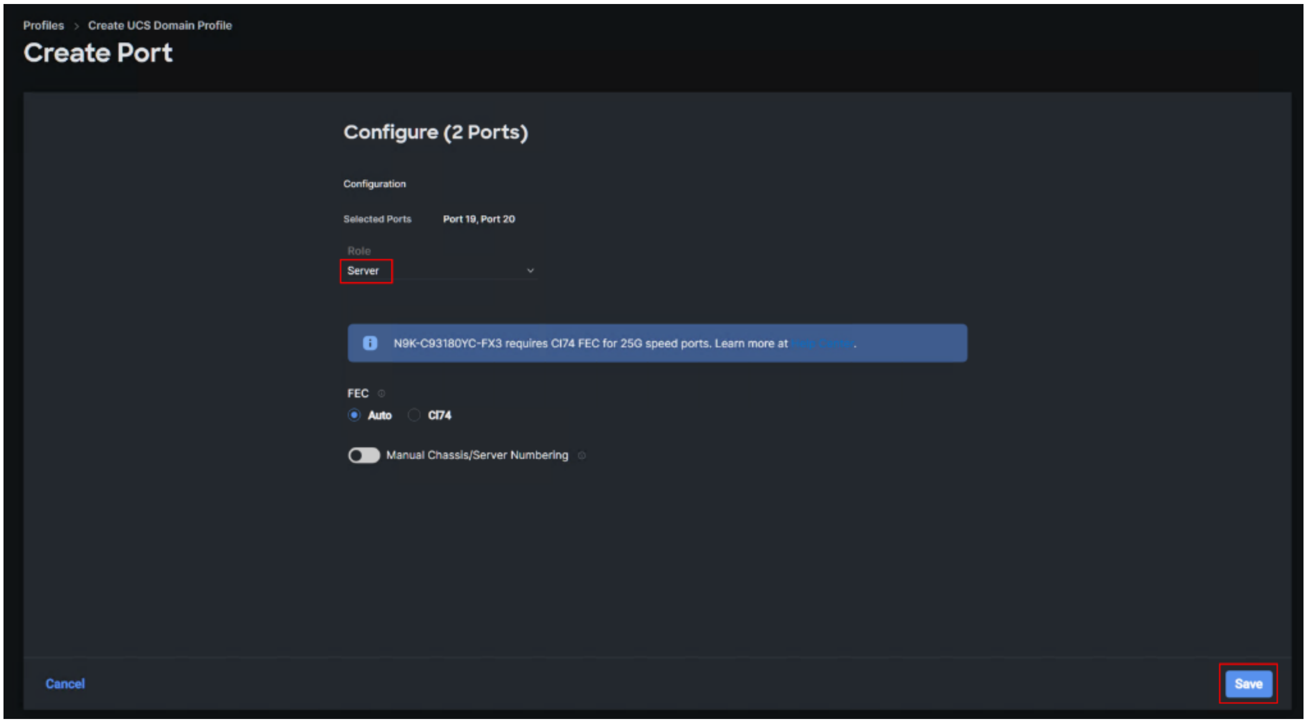

Profiles > Create UCS Domain Profile

Create Port

Configure (2 Ports)

Configuration

Selected Ports Port 19, Port 20

Role: Server

N9K-C93180YC-FX3 requires CI74 FEC for 25G speed ports. Learn more at Help Center.

FEC

● Auto ○ CI74

Manual Chassis/Server Numbering

Step 27. Configure the Ethernet uplink port channel by selecting the Port Channel in the main pane and then click Create Port Channel.

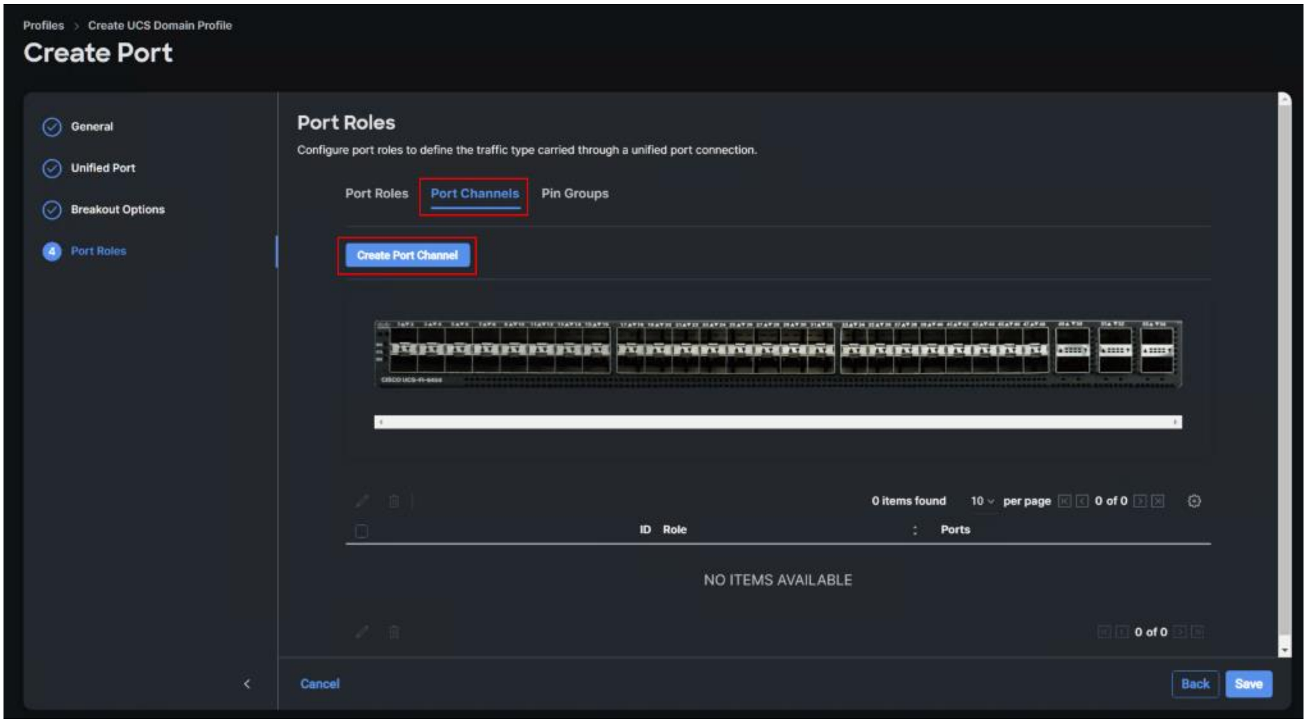

Step 28. Select Ethernet Uplink Port Channel as the role, provide a port-channel ID (for example, 11).

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 29. Scroll down and select uplink ports from the list of available ports (for example, port 49 and 50).

Step 30. Click Save.

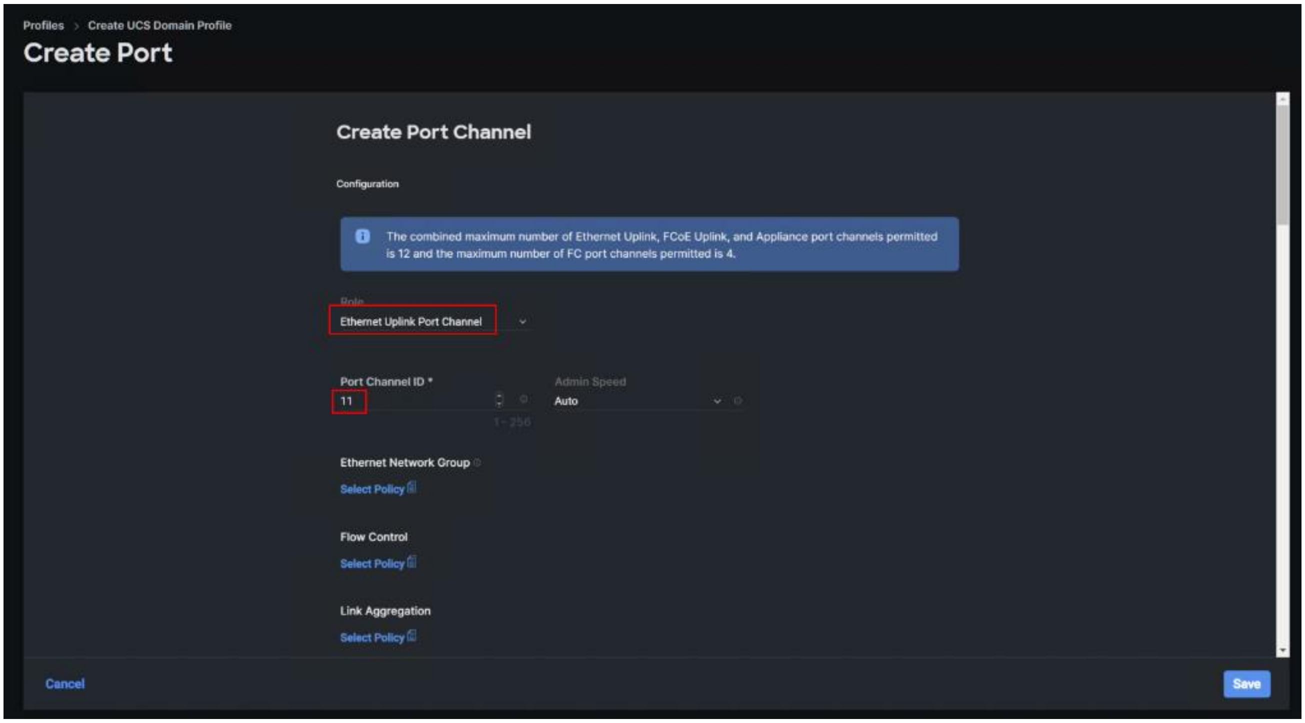

Profiles > Create UCS Domain Profile

Create Port

Create Port Channel

Configuration

The combined maximum number of Ethernet Uplink, FCoE Uplink, and Appliance port channels permitted is 12 and the maximum number of FC port channels permitted is 4.

Role: Ethernet Uplink Port Channel

Port Channel ID *: 11 (1 - 256)

Admin Speed: Auto

Ethernet Network Group: Select Policy

Flow Control: Select Policy

Link Aggregation: Select Policy

Step 31. Click Save.

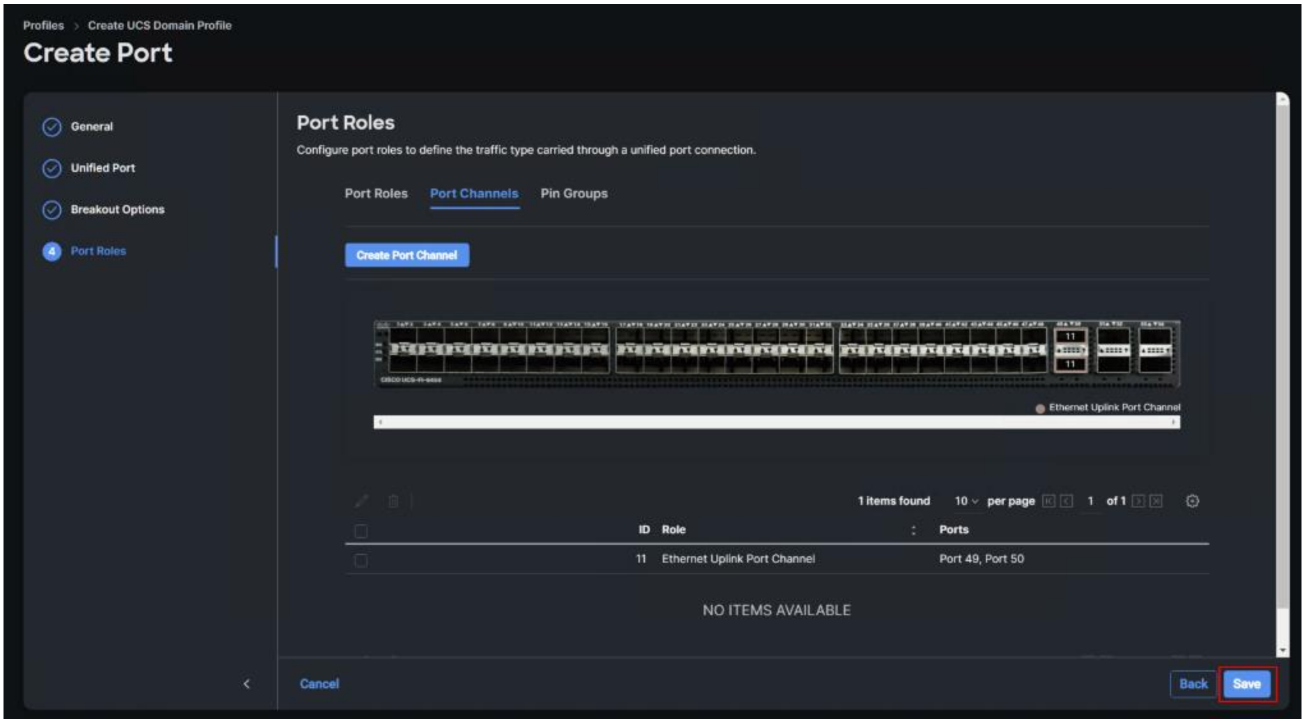

Profiles > Create UCS Domain Profile

Create Port

Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

1 items found 10 per page 1 of 1

ID	Role	Ports
11	Ethernet Uplink Port Channel	Port 49, Port 50

NO ITEMS AVAILABLE

Step 32. Repeat steps 21 - 31 to create the port policy for Fabric Interconnect B. Use unique values for various parameters:

- For example, name of the port policy: L152-J5-FI-B
- For example, Ethernet port-Channel ID: 12

Step 33. When the port configuration for both fabric interconnects is complete and looks good, click Next.

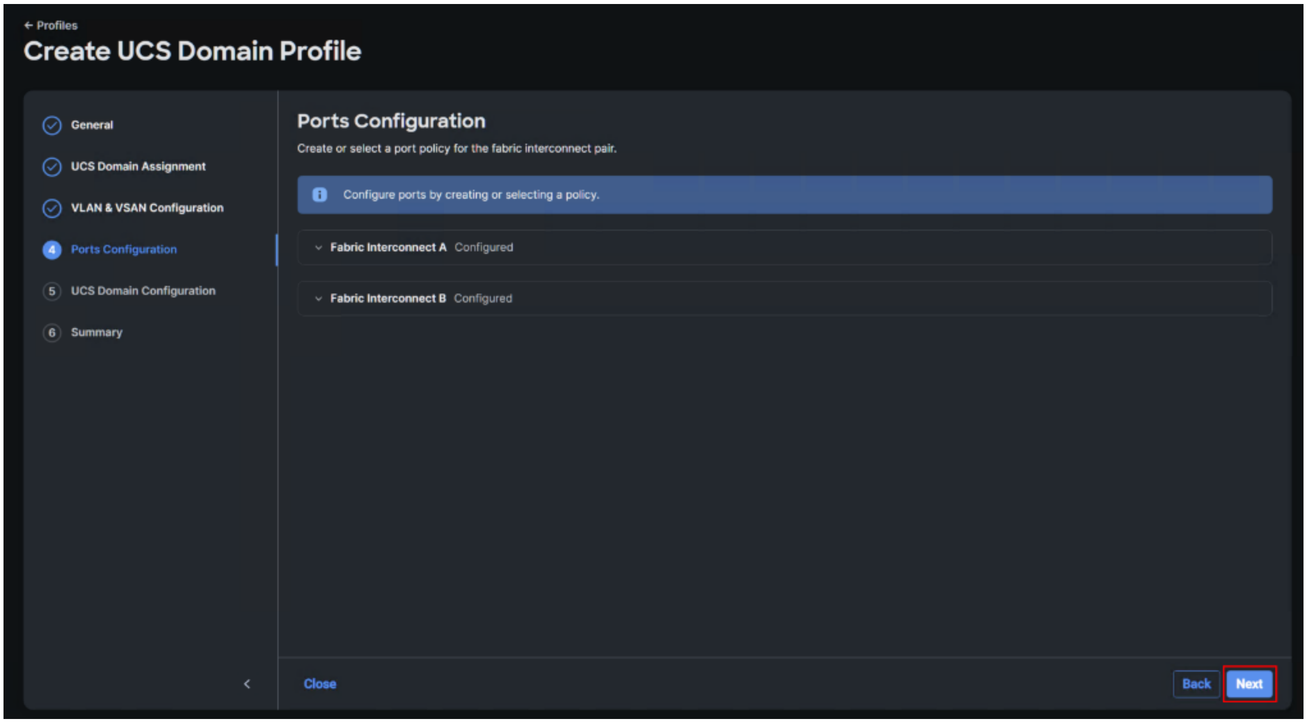

← Profiles

Create UCS Domain Profile

Ports Configuration

Create or select a port policy for the fabric interconnect pair.

Configure ports by creating or selecting a policy.

Fabric Interconnect A Configured

Fabric Interconnect B Configured

Step 34. 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). System QoS is required.

Step 35. Click Select Policy next to System QoS* and click Create New to define the System QOS policy.

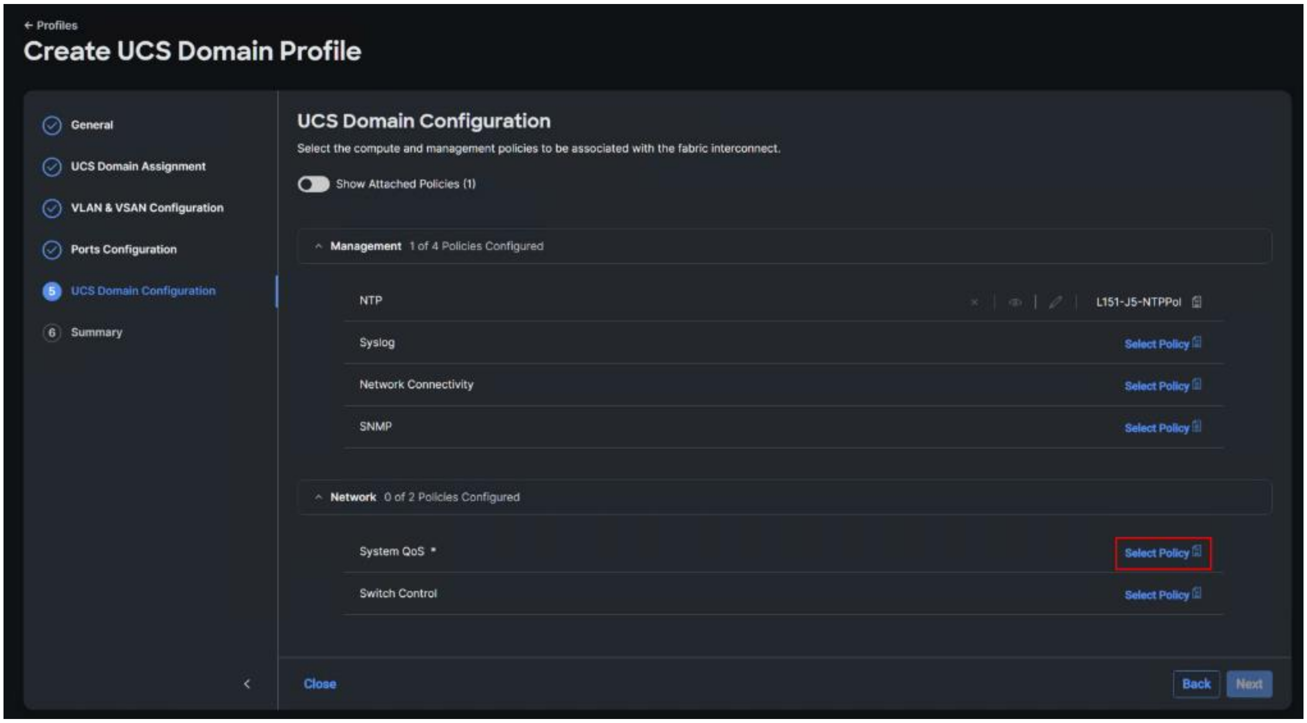

← Profiles

Create UCS Domain Profile

UCS Domain Configuration

Select the compute and management policies to be associated with the fabric interconnect.

Show Attached Policies (1)

Management 1 of 4 Policies Configured

NTP L151-J5-NTPPol

Syslog Select Policy

Network Connectivity Select Policy

SNMP Select Policy

Network 0 of 2 Policies Configured

System QoS * Select Policy

Switch Control Select Policy

Step 36. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, L152-J5-QoSPol). Click Next.

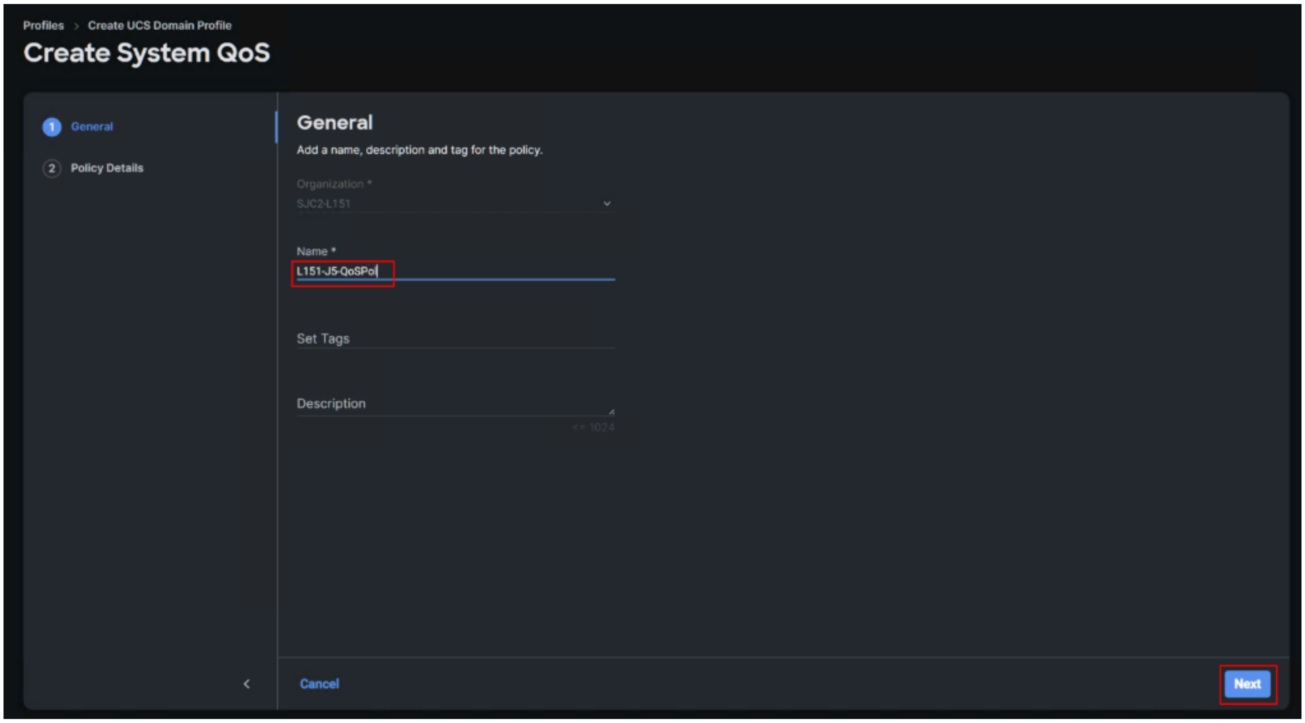

Profiles > Create UCS Domain Profile

Create System QoS

1 General

2 Policy Details

General

Add a name, description and tag for the policy.

Organization *
SJC2-L151

Name *
L151-J5-QoSPol

Set Tags

Description
<= 1024

Cancel

Next

Step 37. Change the MTU for Best Effort class to 9216. Keep the rest default selections. Click Create.

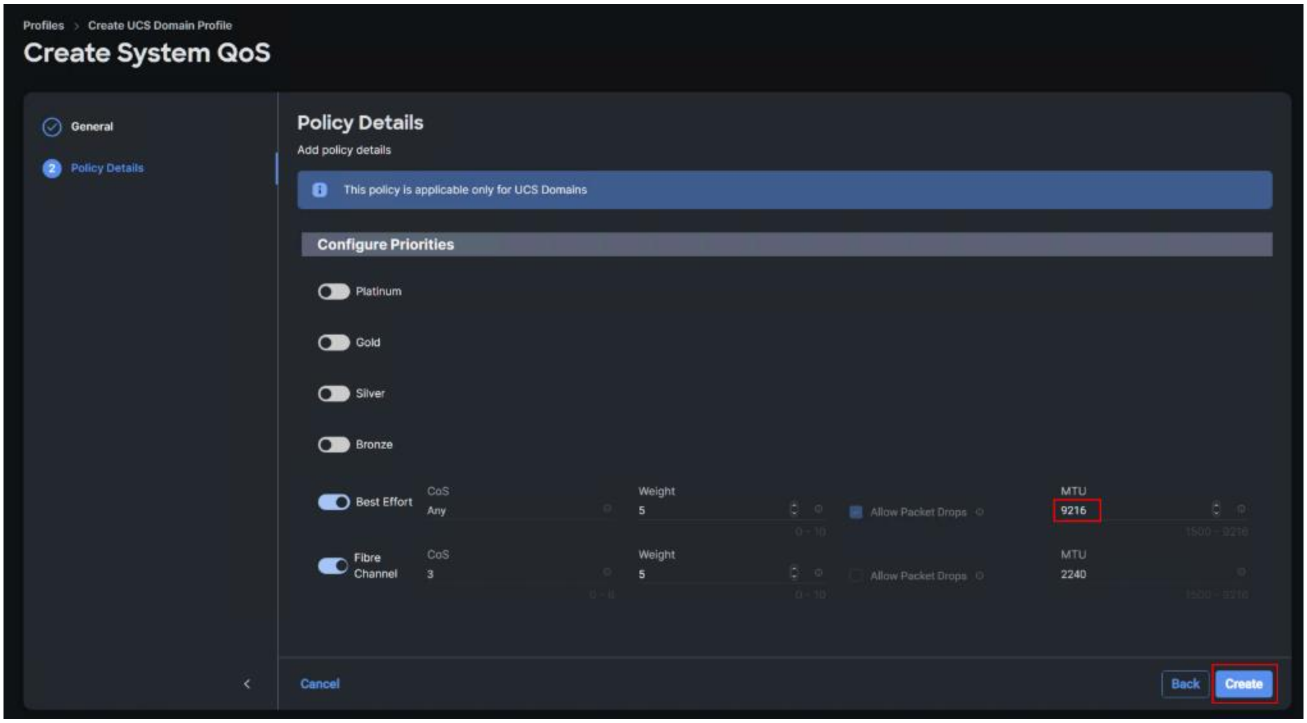

Profiles > Create UCS Domain Profile

Create System QoS

General

2 Policy Details

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Configure Priorities

	CoS	Weight		MTU
Platinum				
Gold				
Silver				
Bronze				
Best Effort	Any	5 (0 - 10)	☑ Allow Packet Drops	9216 (1500 - 9216)
Fibre Channel	3 (0 - 6)	5 (0 - 10)	□ Allow Packet Drops	2240 (1500 - 9216)

Cancel

Back

Create

Step 38. Click Next.

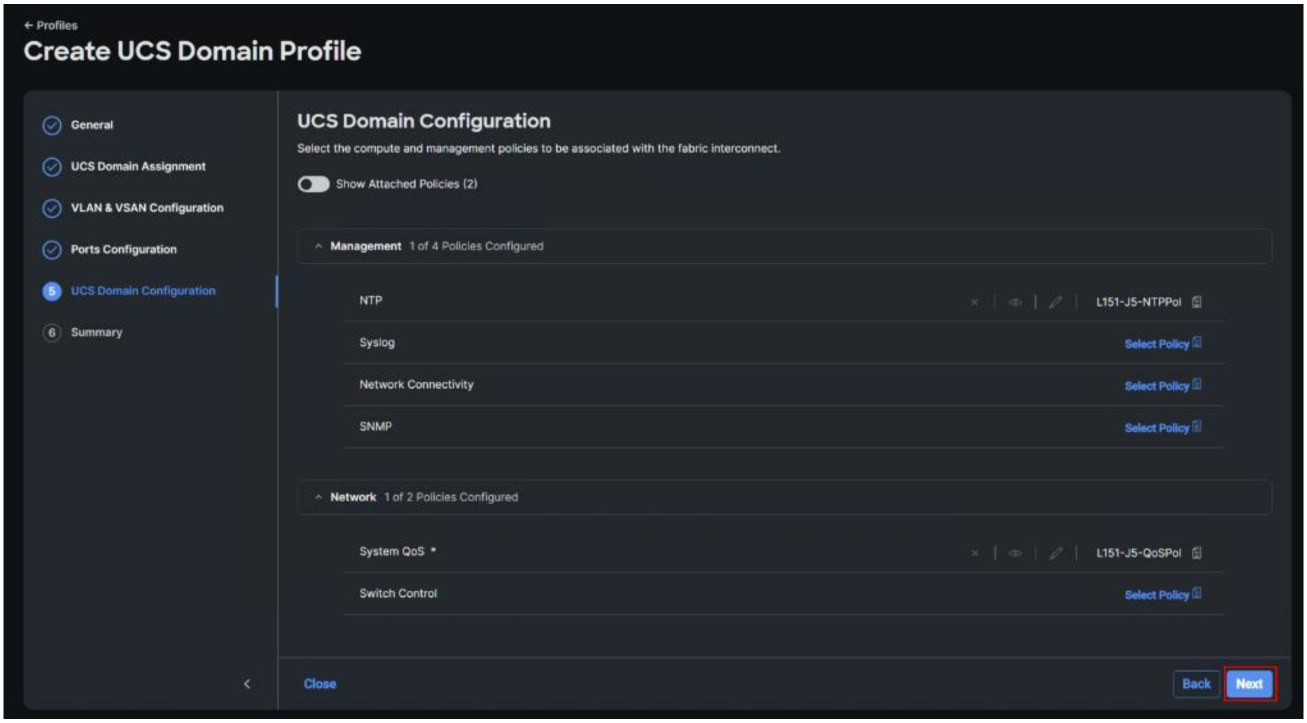

← Profiles

Create UCS Domain Profile

UCS Domain Configuration

Select the compute and management policies to be associated with the fabric interconnect.

Show Attached Policies (2)

Management 1 of 4 Policies Configured

NTP	L151-J5-NTPPol
Syslog	Select Policy
Network Connectivity	Select Policy
SNMP	Select Policy

Network 1 of 2 Policies Configured

System QoS *	L151-J5-QoSPol
Switch Control	Select Policy

Step 39. From the UCS domain profile Summary view, Verify all the settings including the fabric interconnect settings, by expanding the settings and make sure that the configuration is correct. Click Deploy.

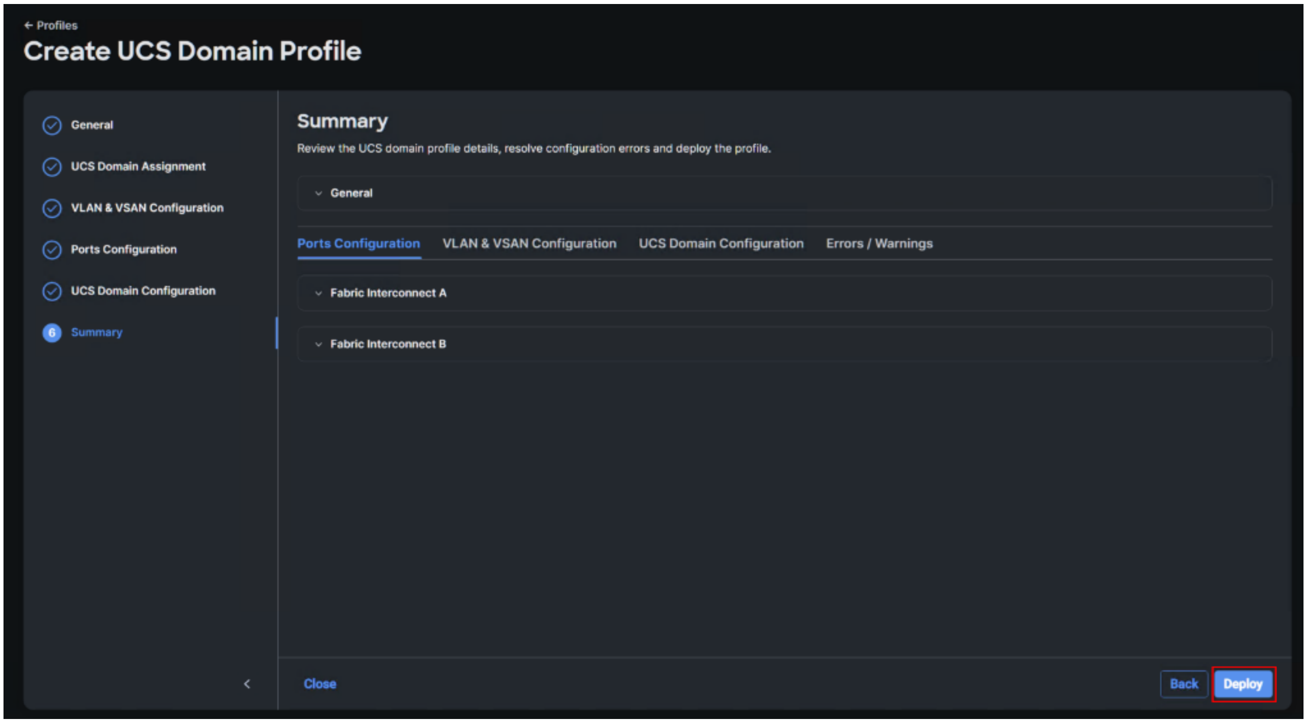

← Profiles

Create UCS Domain Profile

Summary

Review the UCS domain profile details, resolve configuration errors and deploy the profile.

General

Fabric Interconnect A

Fabric Interconnect B

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.

When the Cisco UCS domain profile has been successfully deployed, the Cisco UCS chassis and the blades should be successfully discovered.

It takes a while to discover the blades for the first time. Cisco Intersight provides an ability to view the progress in the Requests page:

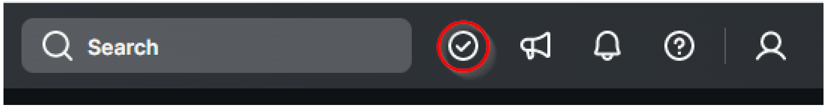

Step 40. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Configure > Profiles, select UCS Domain Profiles, verify that the domain profile has been successfully deployed.

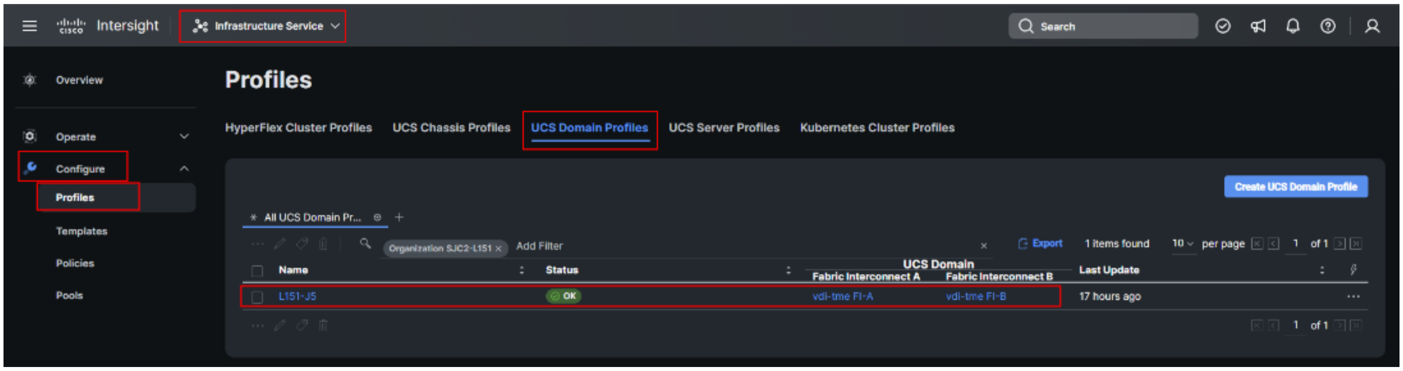

Profiles

		UCS Domain		
Name	Status	Fabric Interconnect A	Fabric Interconnect B	Last Update
L151-J5	OK	vdi-tme FI-A	vdi-tme FI-B	17 hours ago

Step 41. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Chassis, verify that the chassis has been discovered.

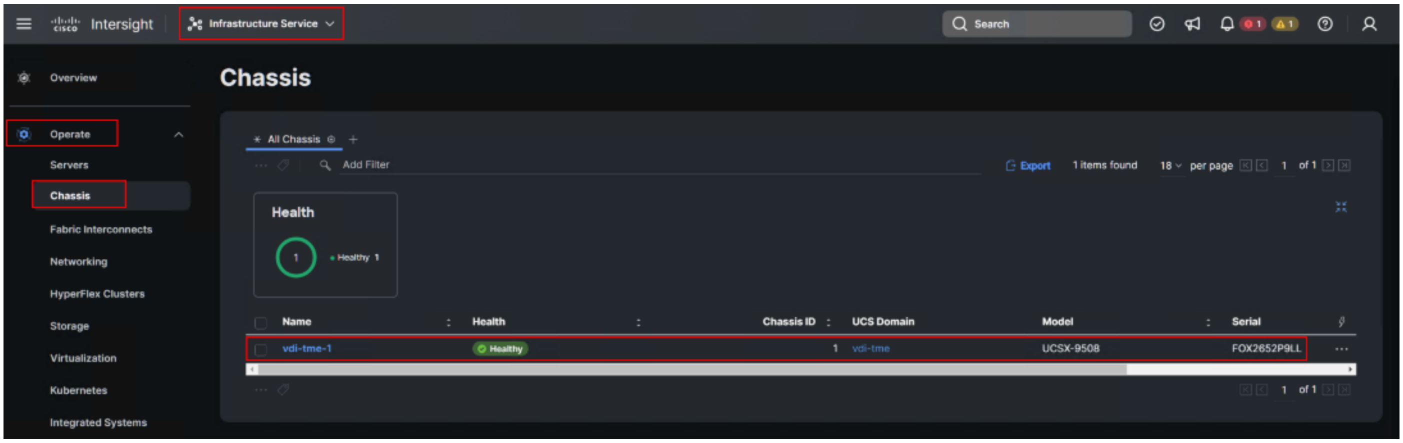

Chassis

Name	Health	Chassis ID	UCS Domain	Model	Serial
vdi-tme-1	Healthy	1	vdi-tme	UCSX-9508	FOX2652P9LL

Step 42. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Servers, verify that the servers have been successfully discovered.

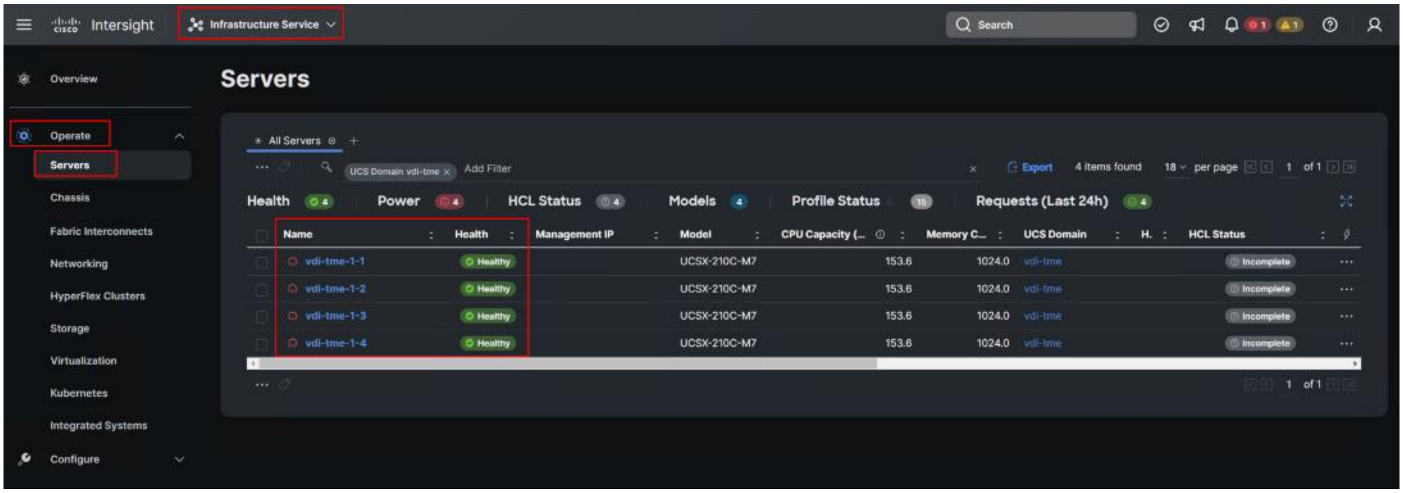

Servers

4 items found 18 per page 1 of 1

Health 4 | Power 4 | HCL Status 4 | Models 4 | Profile Status 15 | Requests (Last 24h) 4

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status
vdi-tme-1-1	Healthy		UCSX-210C-M7	153.6	1024.0	vdi-tme		Incomplete
vdi-tme-1-2	Healthy		UCSX-210C-M7	153.6	1024.0	vdi-tme		Incomplete
vdi-tme-1-3	Healthy		UCSX-210C-M7	153.6	1024.0	vdi-tme		Incomplete
vdi-tme-1-4	Healthy		UCSX-210C-M7	153.6	1024.0	vdi-tme		Incomplete

Configure Cisco UCS Chassis Profile

Cisco UCS Chassis profile in Cisco Intersight allows you to configure various parameters for the chassis, including:

- IMC Access Policy: IP configuration for the in-band chassis connectivity. This setting is independent of Server IP connectivity and only applies to communication to and from chassis.
- SNMP Policy, and SNMP trap settings.
- Power Policy to enable power management and power supply redundancy mode.
- Thermal Policy to control the speed of FANs (only applicable to Cisco UCS 5108).

A chassis policy can be assigned to any number of chassis profiles to provide a configuration baseline for a chassis. In this deployment, chassis profile was created and attached to the chassis with following settings shown in Figure 17.

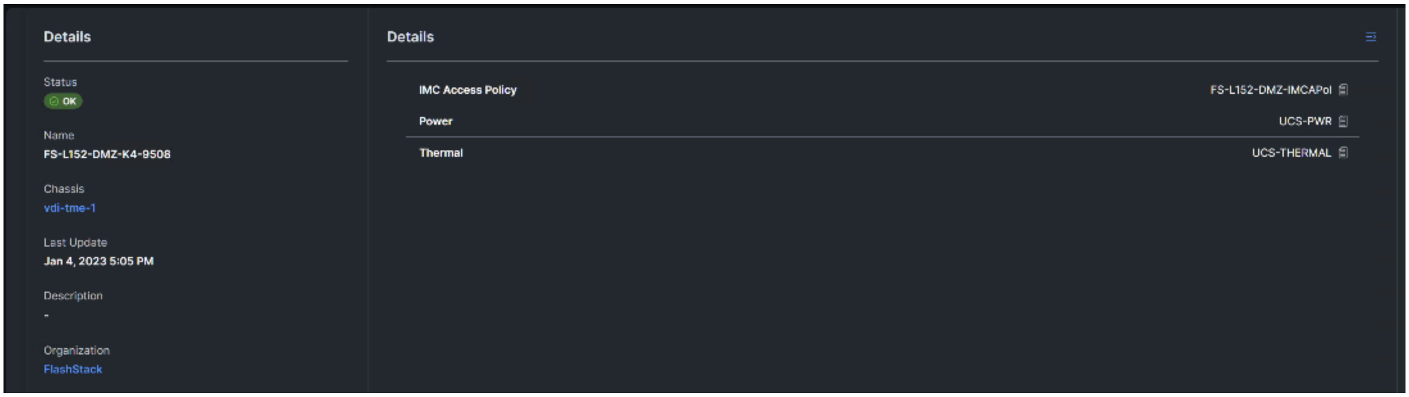

Details

Status
OK

Name
FS-L152-DMZ-K4-9508

Chassis
vdi-tme-1

Last Update
Jan 4, 2023 5:05 PM

Description
-

Organization
FlashStack

Details

IMC Access Policy	FS-L152-DMZ-IMCAPol
Power	UCS-PWR
Thermal	UCS-THERMAL

Figure 17. Chassis policy detail

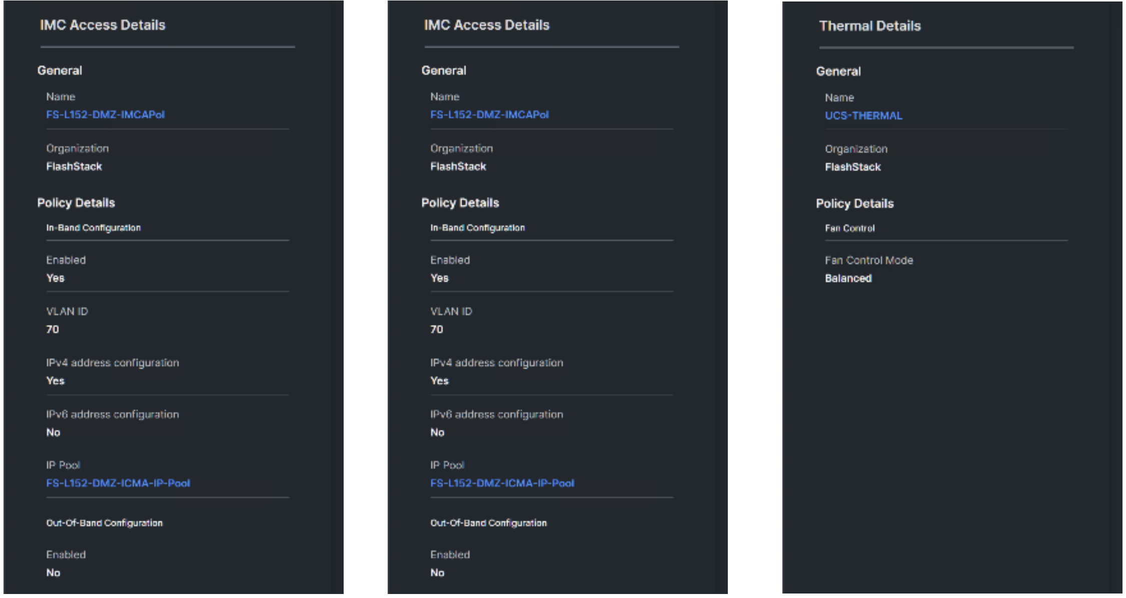

IMC Access Details

General

Name
FS-L152-DMZ-IMCAPol

Organization
FlashStack

Policy Details

In-Band Configuration

Enabled
Yes

VLAN ID
70

IPv4 address configuration
Yes

IPv6 address configuration
No

IP Pool
FS-L152-DMZ-ICMA-IP-Pool

Out-Of-Band Configuration

Enabled
No

IMC Access Details

General

Name
FS-L152-DMZ-IMCAPol

Organization
FlashStack

Policy Details

In-Band Configuration

Enabled
Yes

VLAN ID
70

IPv4 address configuration
Yes

IPv6 address configuration
No

IP Pool
FS-L152-DMZ-ICMA-IP-Pool

Out-Of-Band Configuration

Enabled
No

Thermal Details

General

Name
UCS-THERMAL

Organization
FlashStack

Policy Details

Fan Control

Fan Control Mode
Balanced

Configure Server Profiles

Configure Server Profile Template

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.

Note: The server profile captured in this deployment guide supports both Cisco UCS X-Series blade servers and Cisco UCS X210c M7 compute nodes.

Procedure 1. Create Server Profile Template

In this deployment, four vNICs are configured. These devices are manually placed as listed in Table 4.

Table 4. vNIC placement

vNIC/vHBA Name	Slot	Switch ID	PCI Order
01-vSwitch0-A	MLOM	A	0
02-vSwitch0-B	MLOM	B	1
03-vDS0-A	MLOM	A	2
04-vDS0-B	MLOM	B	3

Step 1. Log into the Cisco Intersight portal as a user with account administrator role.

Step 2. Navigate to Configure > Templates and click Create UCS Server Profile Template.

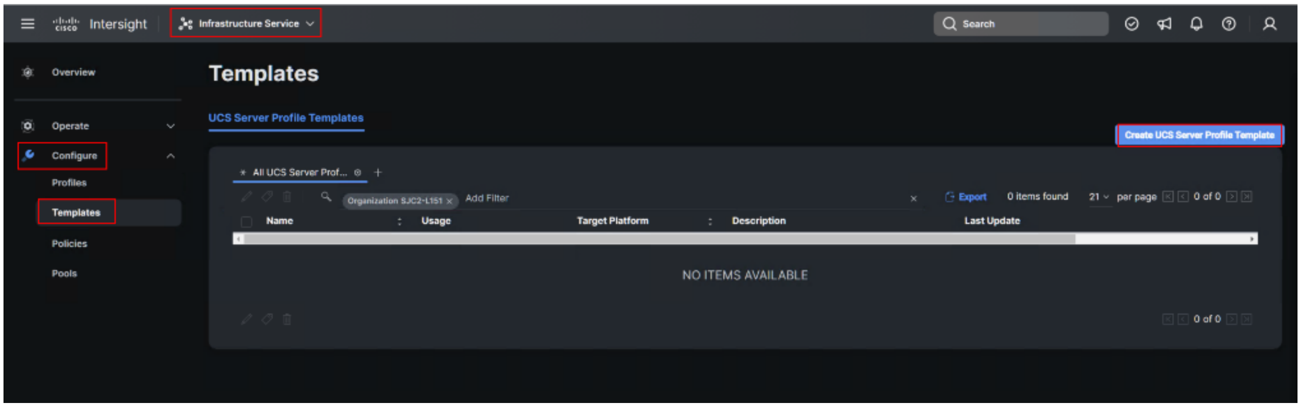

Step 3. Select the organization from the drop-down list. Provide a name for the server profile template (for example, L151-J5-X210CM7-vSANnode) for FI-Attached UCS Server. Click Next.

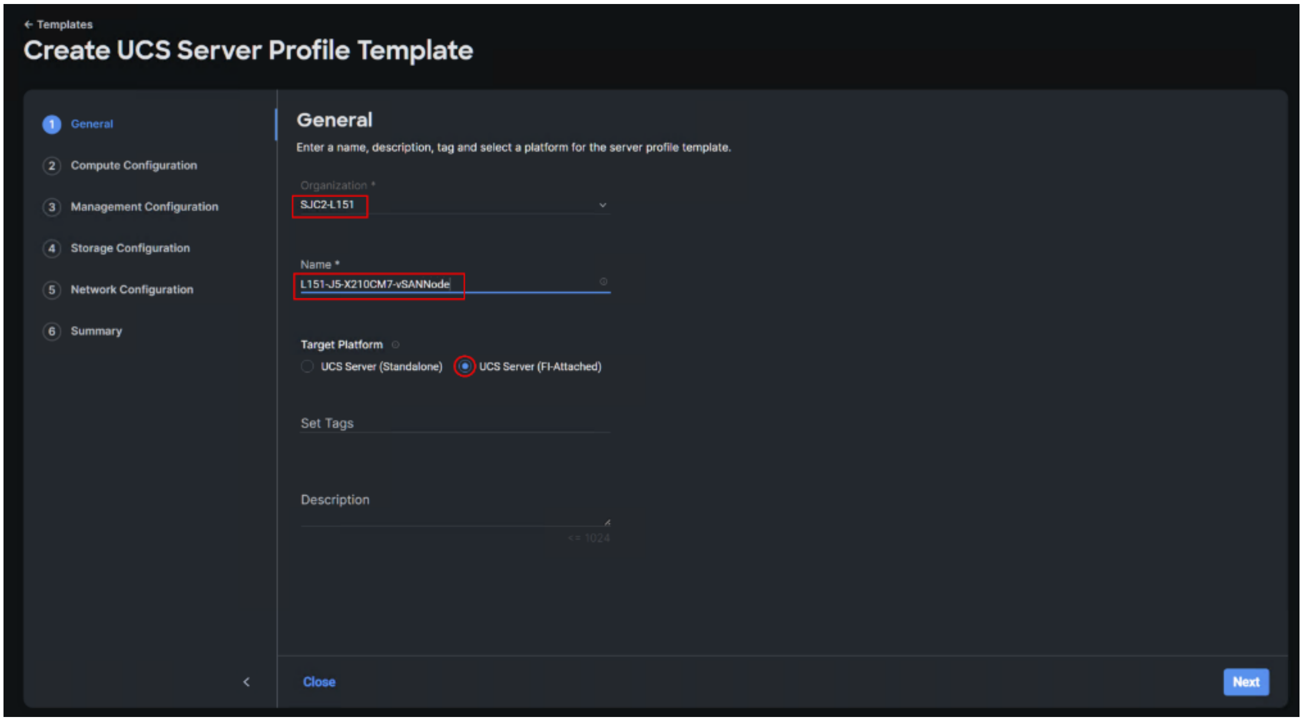

General

Enter a name, description, tag and select a platform for the server profile template.

Organization * SJC2-L151

Name * L151-J5-X210CM7-vSANNode

Target Platform ○ UCS Server (Standalone) ● UCS Server (FI-Attached)

Step 4. Click Select Pool under UUID Pool and then click Create New.

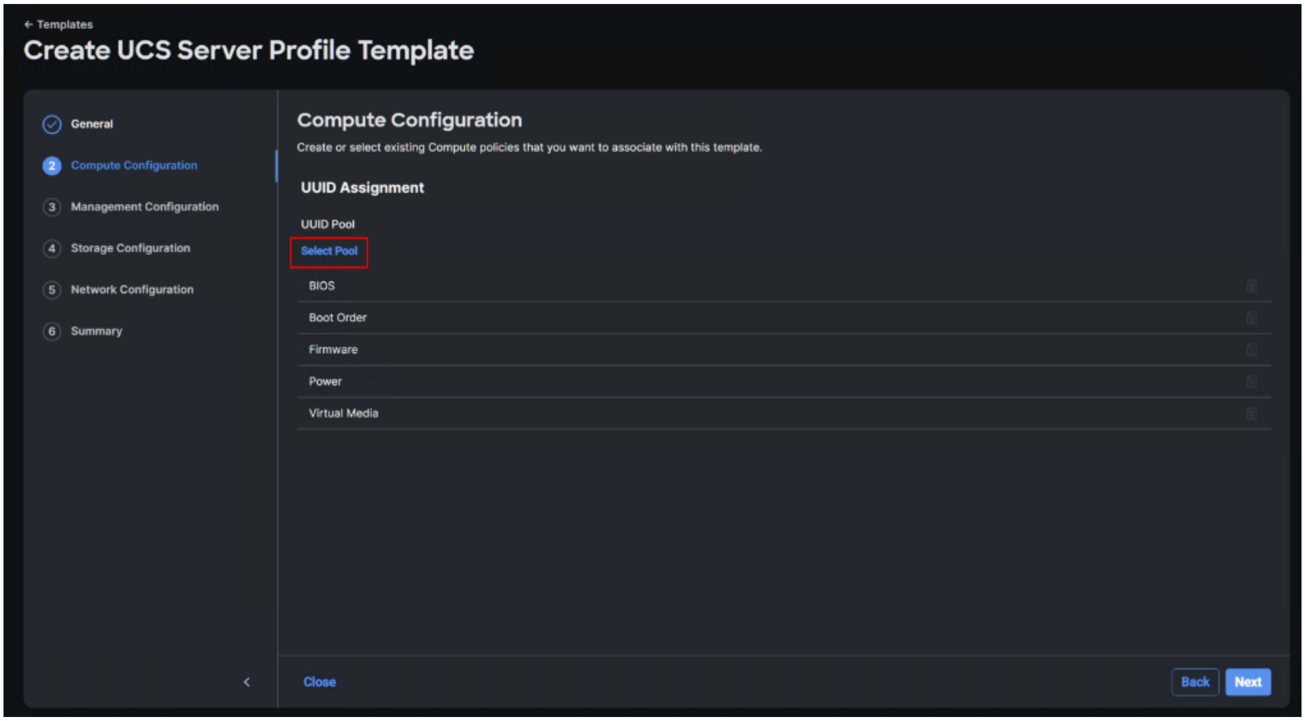

← Templates

Create UCS Server Profile Template

Compute Configuration

Create or select existing Compute policies that you want to associate with this template.

UUID Assignment

UUID Pool

Step 5. Verify correct organization is selected from the drop-down list and provide a name for the UUID Pool (for example, L151-J5-UUID-Pool). Provide an optional Description and click Next.

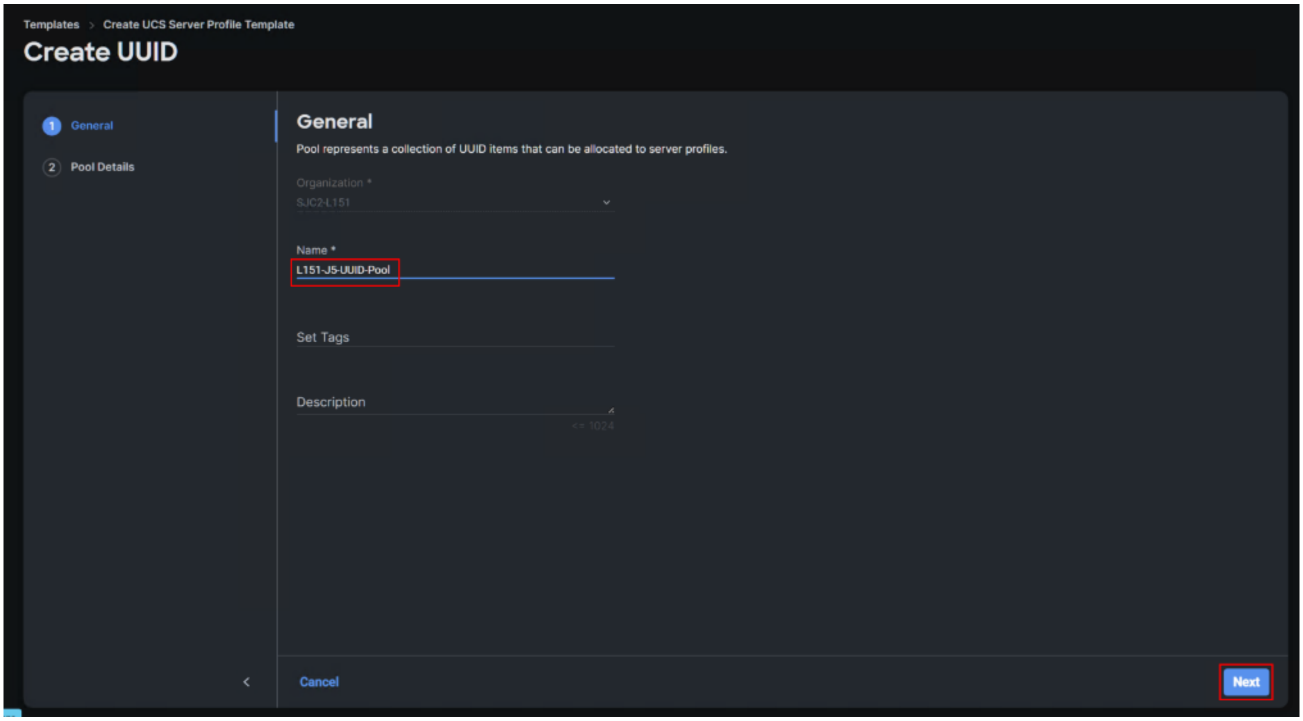

Templates > Create UCS Server Profile Template

Create UUID

General

Pool represents a collection of UUID items that can be allocated to server profiles.

Organization * SJC2-L151

Name * L151-J5-UUID-Pool

Set Tags

Description

<= 1024

Step 6. Provide a UUID Prefix (for example, a random prefix of B2151000-2023-0515 was used). Add a UUID block of appropriate size. Click Create.

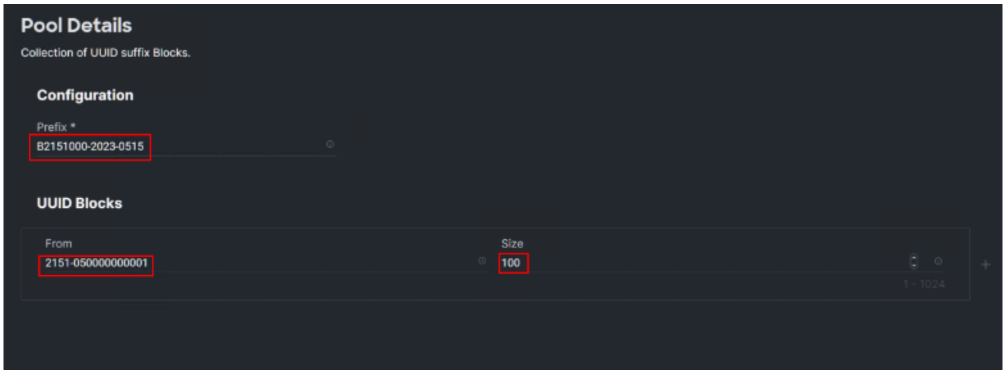

Pool Details

Collection of UUID suffix Blocks.

Configuration

Prefix *: B2151000-2023-0515

UUID Blocks

From: 2151-050000000001

Size: 100

1 - 1024

Step 7. Click Select Policy next to BIOS and in the pane on the right, click Create New.

Step 8. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-M7-vSANnode).

Step 9. Click Next.

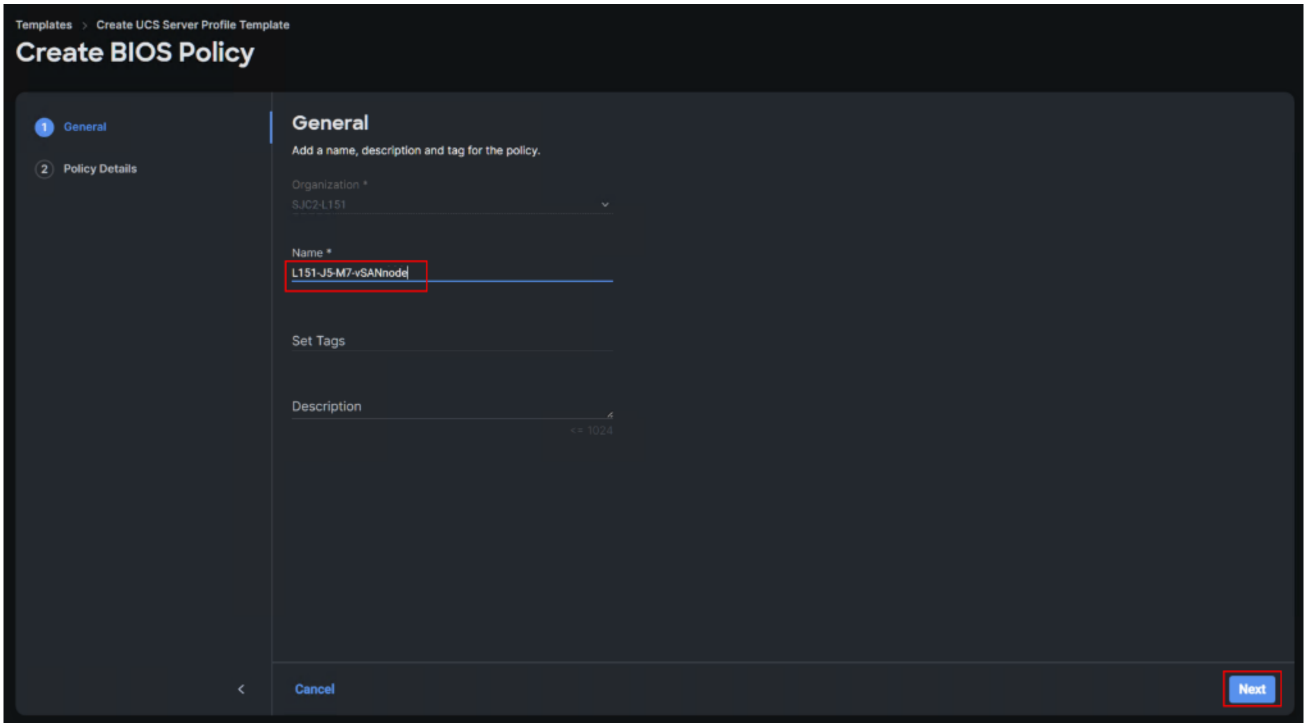

Templates > Create UCS Server Profile Template

Create BIOS Policy

1 General

2 Policy Details

General

Add a name, description and tag for the policy.

Organization *: SJC2-L151

Name *: L151-J5-M7-vSANnode

Set Tags

Description

<= 1024

Step 10. On the Policy Details screen, select appropriate values for the BIOS settings. Click Create.

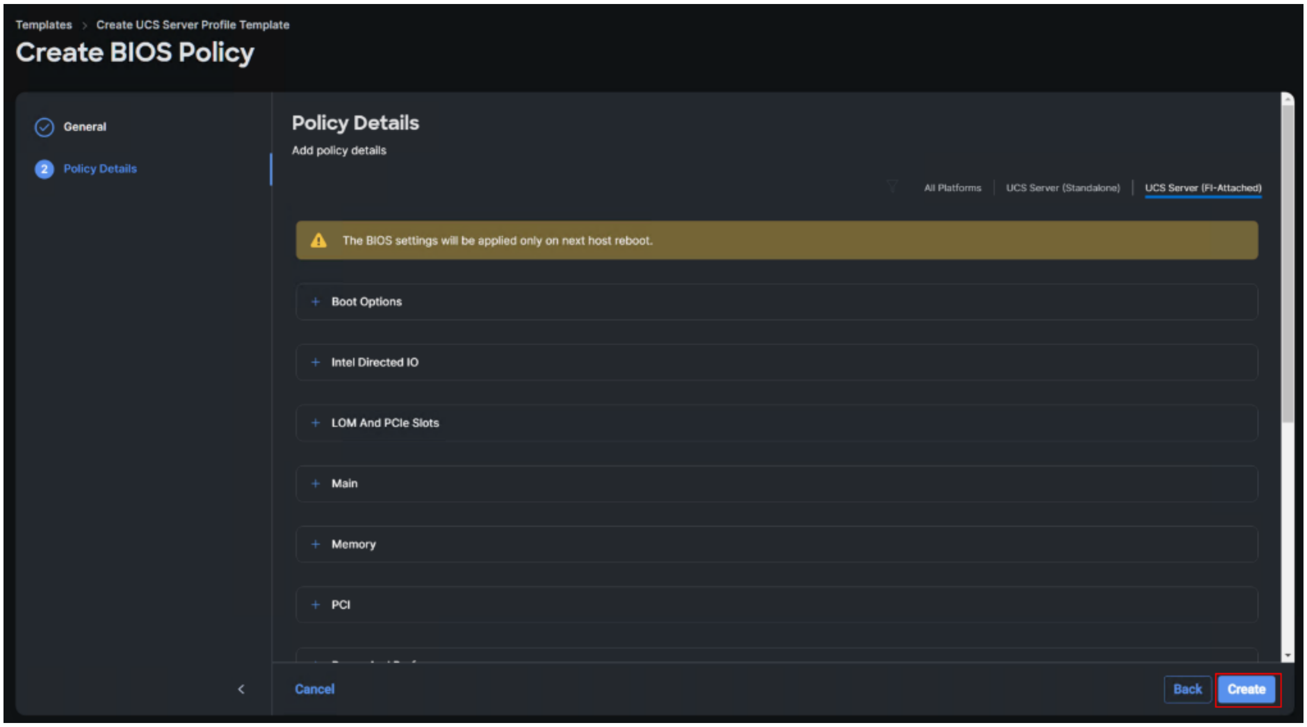

The BIOS settings will be applied only on next host reboot.

Note: 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-blade-servers/performance-tuning-guide-ucs-M7-servers.html and listed in Table 5. Table 6 provides alternative values to support an energy-efficient policy. This policy will lower power consumption when the desktops are not running.

Table 5. L151-J5-M7-vSANnode token values

BIOS Token	Value
Intel Directed IO	
Intel VT for Directed IO	enabled
Memory	
Memory RAS Configuration	maximum-performance
Power And Performance	
Core Performance Boost	Auto
Enhanced CPU Performance	Auto
LLC Dead Line	disabled
UPI Link Enablement	1
UPI Power Management	enabled
Processor	

BIOS Token	Value
Altitude	auto
Boot Performance Mode	Max Performance
Core Multi Processing	all
CPU Performance	enterprise
Power Technology	performance
Direct Cache Access Support	enabled
DRAM Clock Throttling	Performance
Enhanced Intel Speedstep(R) Technology	enabled
Execute Disable Bit	enabled
IMC Interleaving	1-way Interleave
Intel HyperThreading Tech	Enabled
Intel Turbo Boost Tech	enabled
Intel(R) VT	enabled
DCU IP Prefetcher	enabled
Processor C1E	disabled
Processor C3 Report	disabled
Processor C6 Report	disabled
CPU C State	disabled
Sub Numa Clustering	enabled
DCU Streamer Prefetch	enabled

Table 6. L151-J5Z-M7-BIOS-EnergyEfficient token values

BIOS Token	Value
Processor	
Processor C1E	enabled
Processor C6 Report	enabled

Step 11. Click Select Policy next to Boot Order and then click Create New.

Step 12. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-vSANnode-BootOrderPol). Click Next.

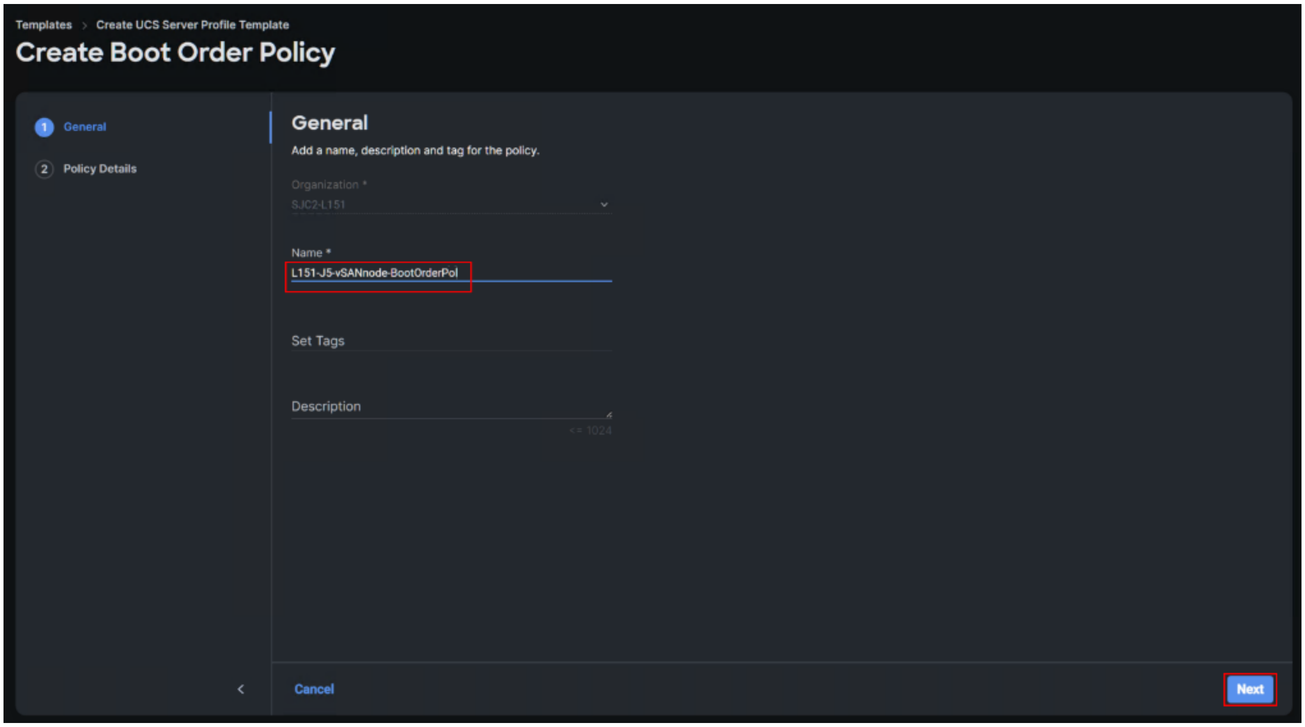

Templates > Create UCS Server Profile Template

Create Boot Order Policy

General

Add a name, description and tag for the policy.

Organization *

SJC2-L151

Name *

L151-J5-vSANnode-BootOrderPol

Set Tags

Description

<= 1024

Step 13. For Configured Boot Mode, select Unified Extensible Firmware Interface (UEFI).

Step 14. Turn on Enable Secure Boot.

Step 15. Click Add Boot Device drop-down list and select Virtual Media.

Step 16. Provide a device name (for example, vKVM-DVD) and then, for the subtype, select KVM Mapped DVD.

Step 17. Click Add Boot Device drop-down list and select Local Disk.

Step 18. Provide a device name (for example, Local-Disk).

Step 19. Verify the order of the boot policies and adjust the boot order as necessary using the arrows next to delete icon. Click Create.

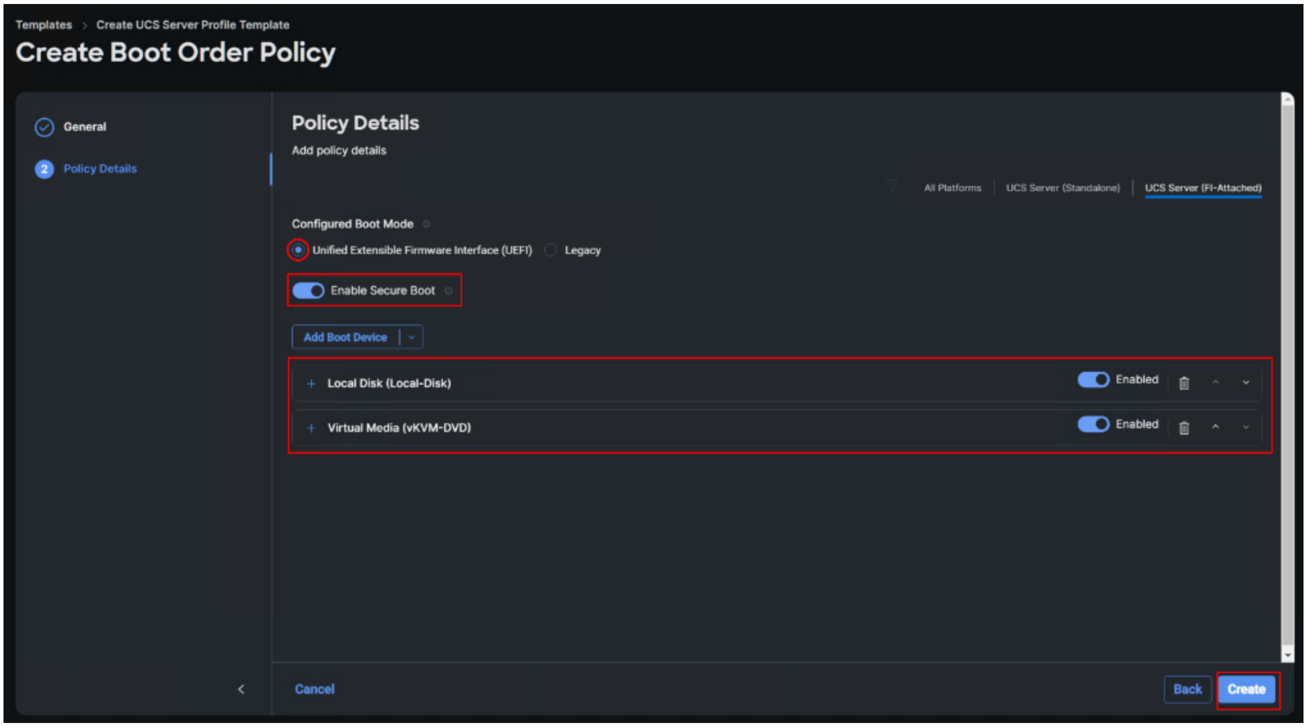

Templates > Create UCS Server Profile Template

Create Boot Order Policy

Policy Details

Add policy details

Configured Boot Mode: ● Unified Extensible Firmware Interface (UEFI) ○ Legacy

Enable Secure Boot: on

Local Disk (Local-Disk): Enabled

Virtual Media (vKVM-DVD): Enabled

Step 20. Click Select Policy next to Power and in the pane on the right, click Create New.

Step 21. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-X210CM7-PwrPol). Click Next.

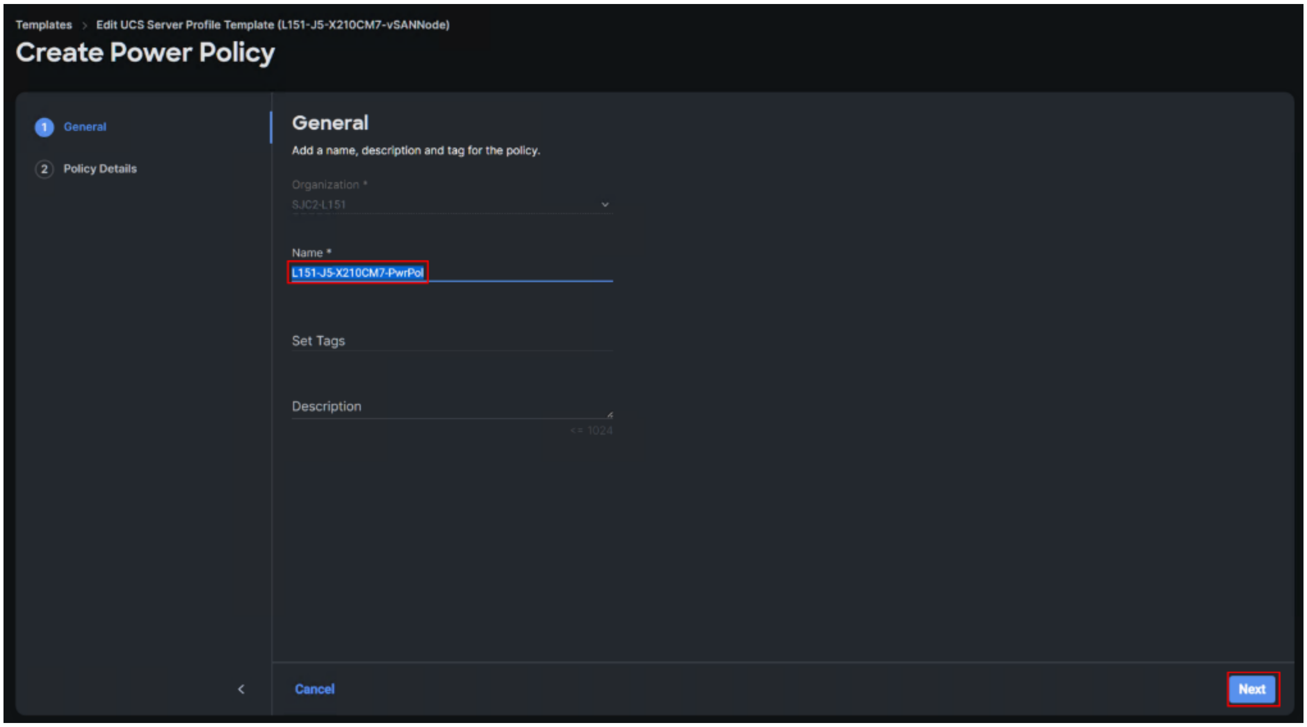

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Power Policy

General

Add a name, description and tag for the policy.

Organization *: SJC2-L151

Name *: L151-J5-X210CM7-PwrPol

Step 22. Enable Power Profiling and select High from the Power Priority drop-down list. Click Create.

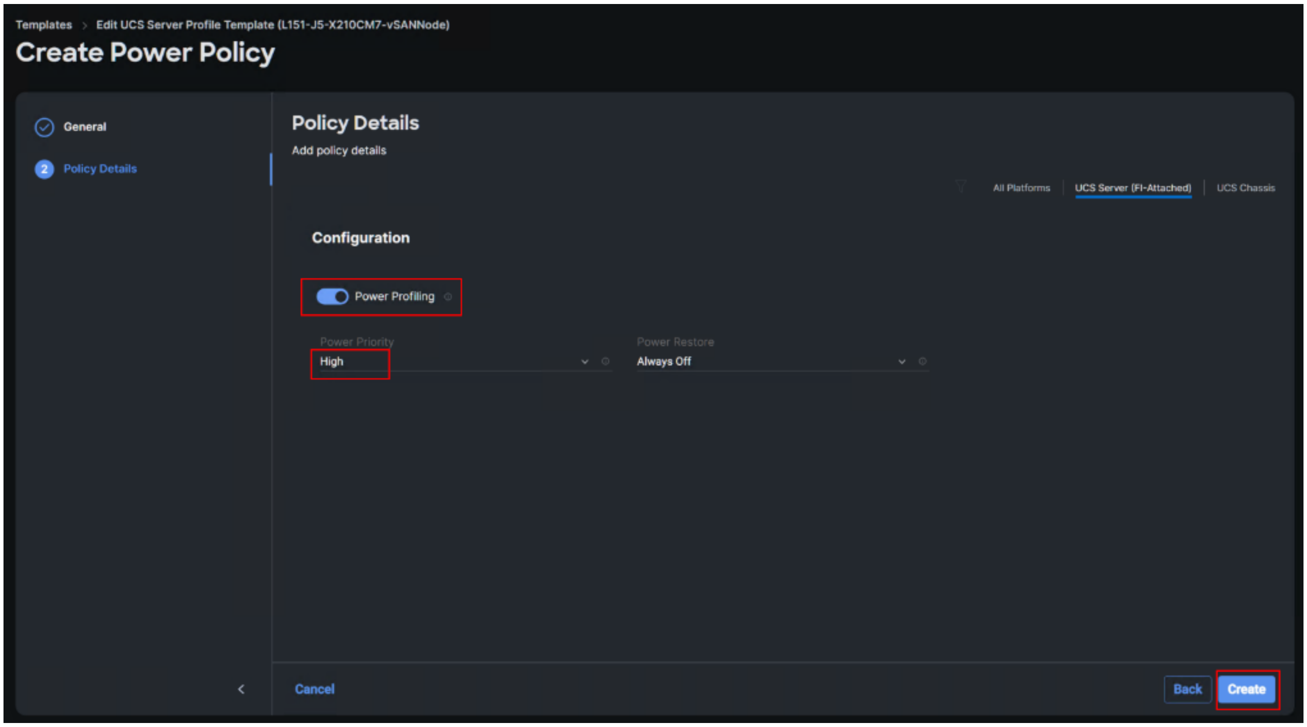
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Power Policy

Policy Details

Add policy details

Configuration

Power Profiling: enabled

Power Priority: High

Power Restore: Always Off

Step 23. Click Select Policy next to Virtual Media and in the pane on the right, click Create New (Optional)

Step 24. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-X210CM7-VMediaPol). Click Next.

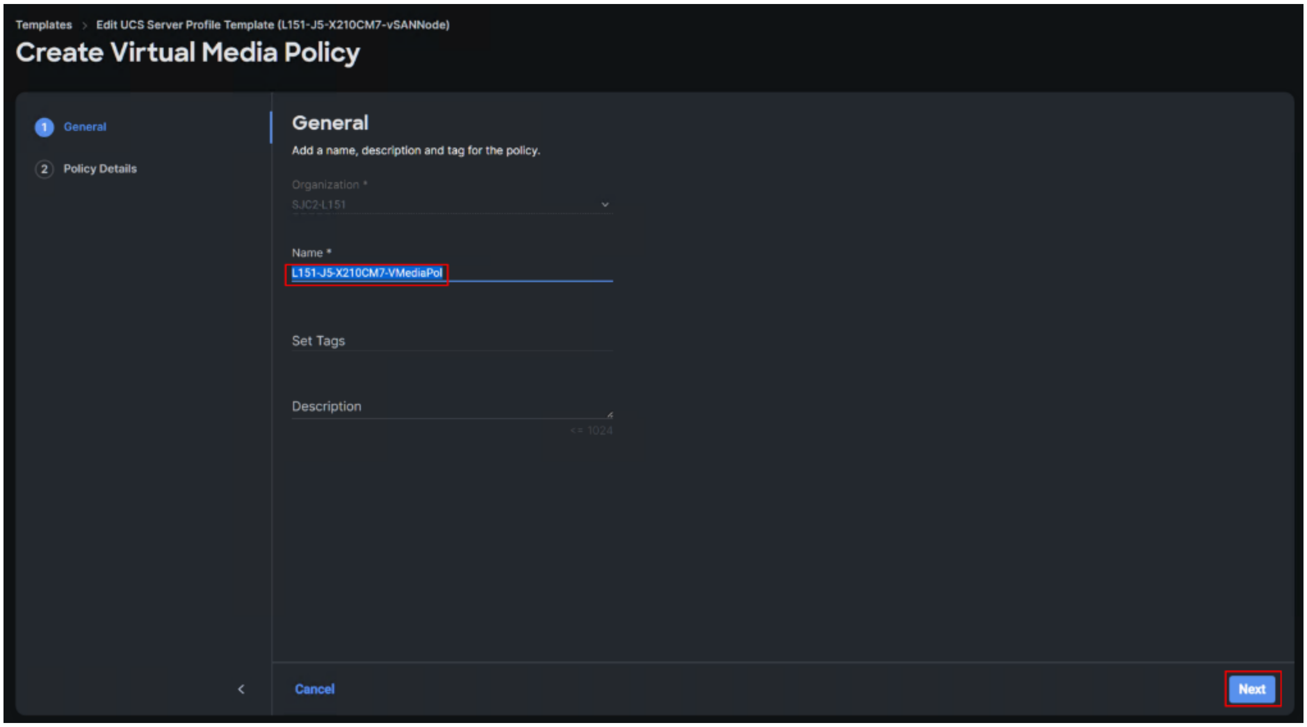
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Virtual Media Policy

General

Add a name, description and tag for the policy.

Organization *: SJC2-L151

Name *: L151-J5-X210CM7-VMediaPol

Set Tags

Description

<= 1024

Step 25. Disable Lower Power USB and click Create.

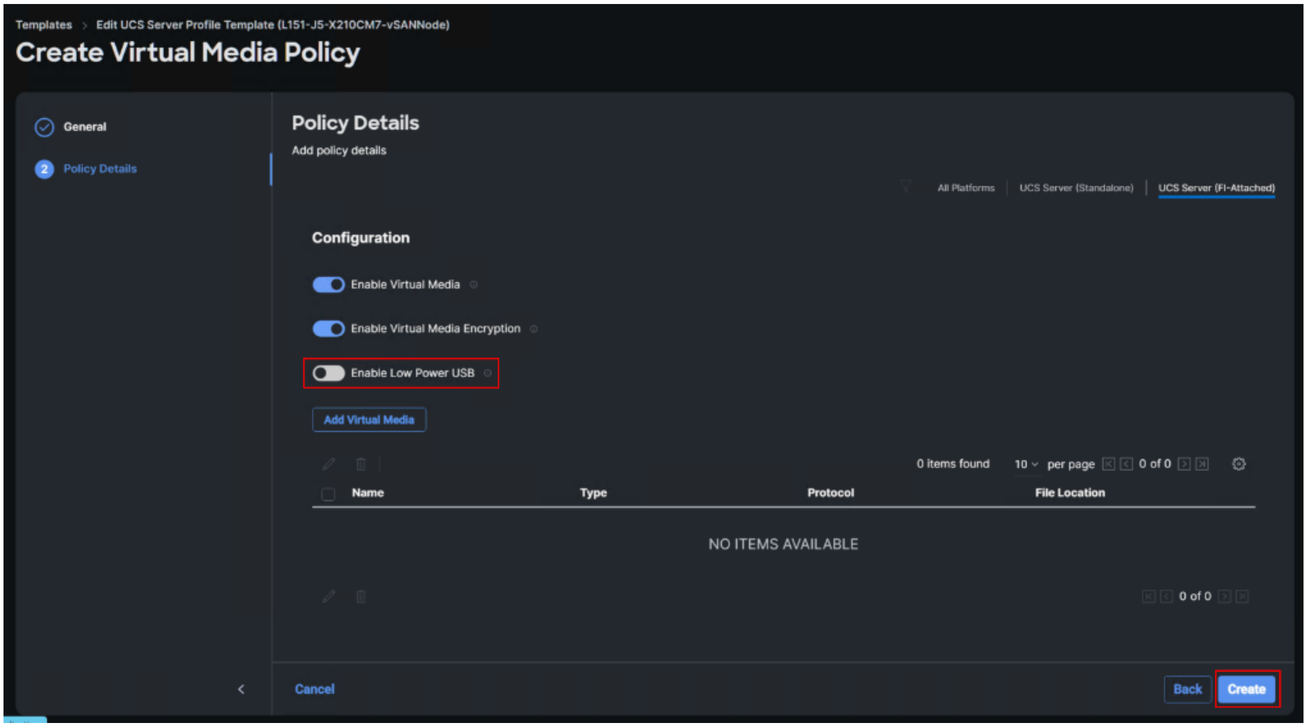

Create Virtual Media Policy

Policy Details

Add policy details

Configuration

Enable Virtual Media: on

Enable Virtual Media Encryption: on

Enable Low Power USB: off

0 items found 10 per page 0 of 0

Name	Type	Protocol	File Location

NO ITEMS AVAILABLE

0 of 0

Step 26. Click Next to go to Management Configuration.

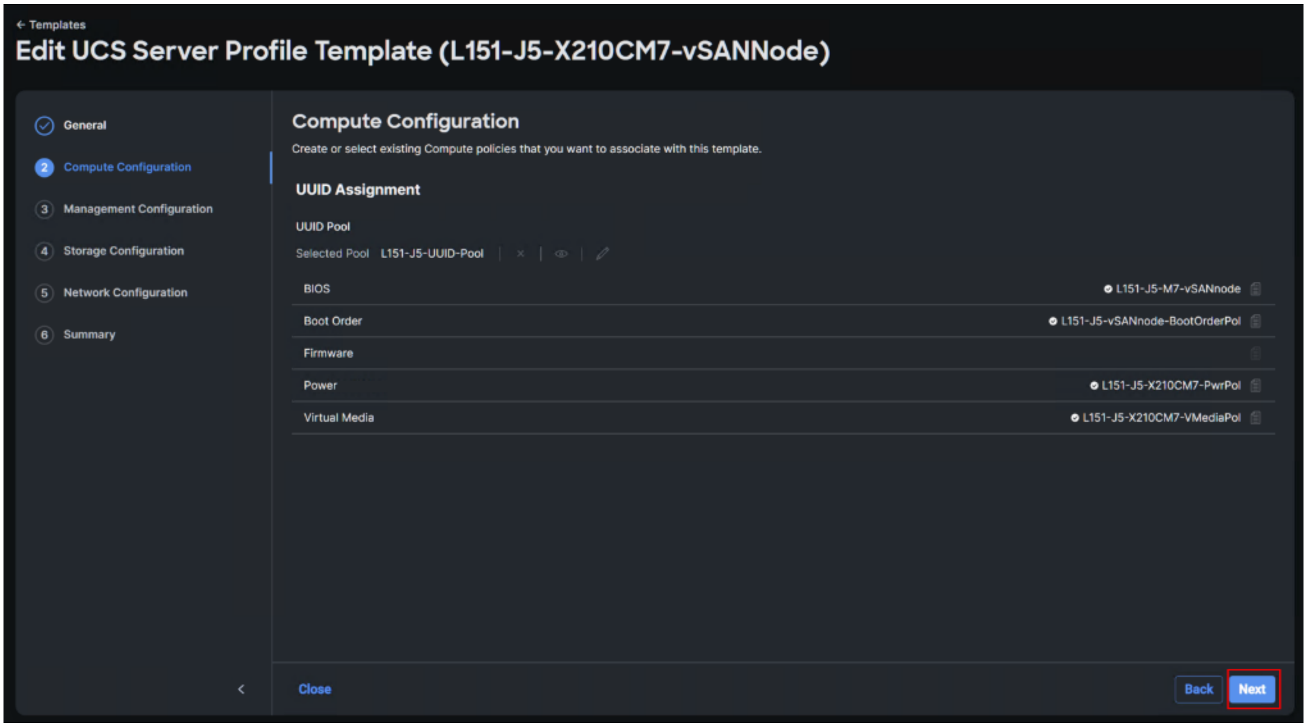

Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Compute Configuration

Create or select existing Compute policies that you want to associate with this template.

UUID Assignment

UUID Pool

Selected Pool L151-J5-UUID-Pool

BIOS	L151-J5-M7-vSANnode
Boot Order	L151-J5-vSANnode-BootOrderPol
Firmware	
Power	L151-J5-X210CM7-PwrPol
Virtual Media	L151-J5-X210CM7-VMediaPol

Step 27. Click Select Policy next to IMC Access and then click Create New.

Step 28. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L152-J5-IMCAccPol). Click Next.

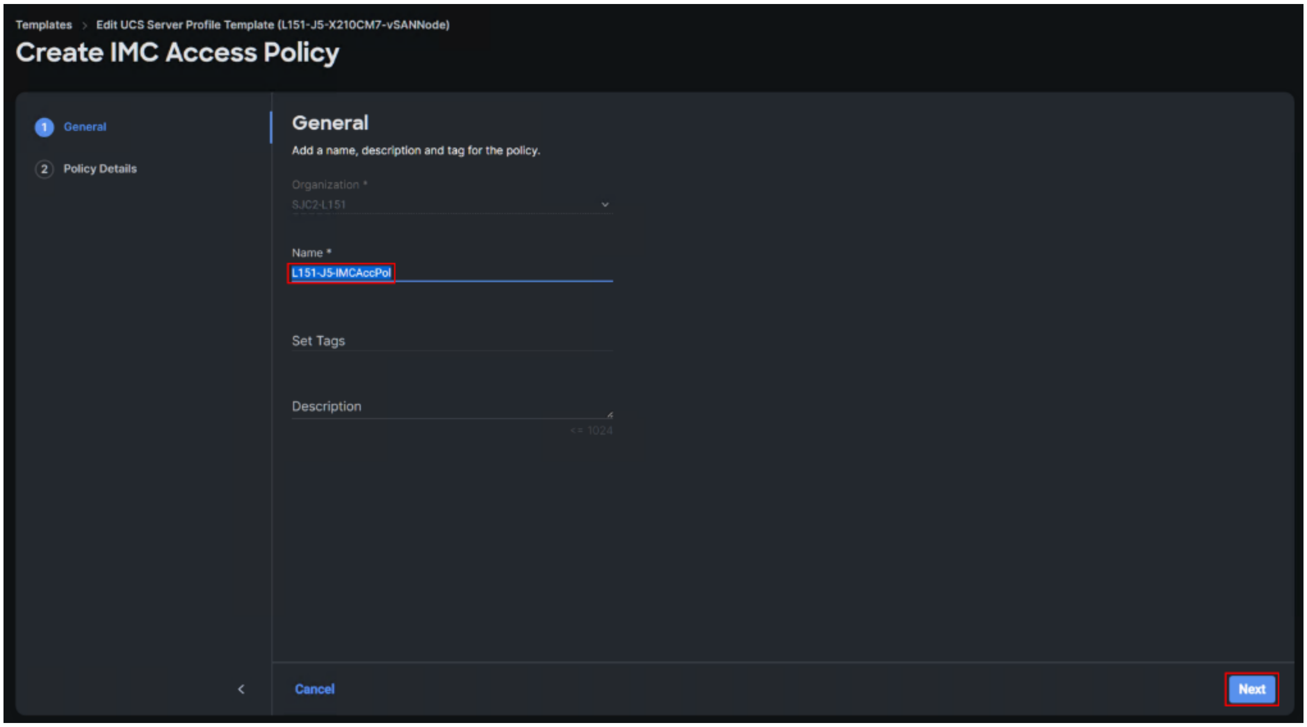
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)
Create IMC Access Policy
1 General
2 Policy Details
General
Cancel
Next

Add a name, description and tag for the policy.

Organization * SJC2-L151

Name * L151-J5-IMCAccPol

Set Tags

Description

<= 1024

Note: You can select in-band management access to the compute node using an in-band management VLAN (for example, VLAN 70) 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.

Step 29. Click UCS Server (FI-Attached). Enable In-Band Configuration and type VLAN Id designated for the In-Band management (for example, 70).

Step 30. Under IP Pool, click Select IP Pool and then click Create New.

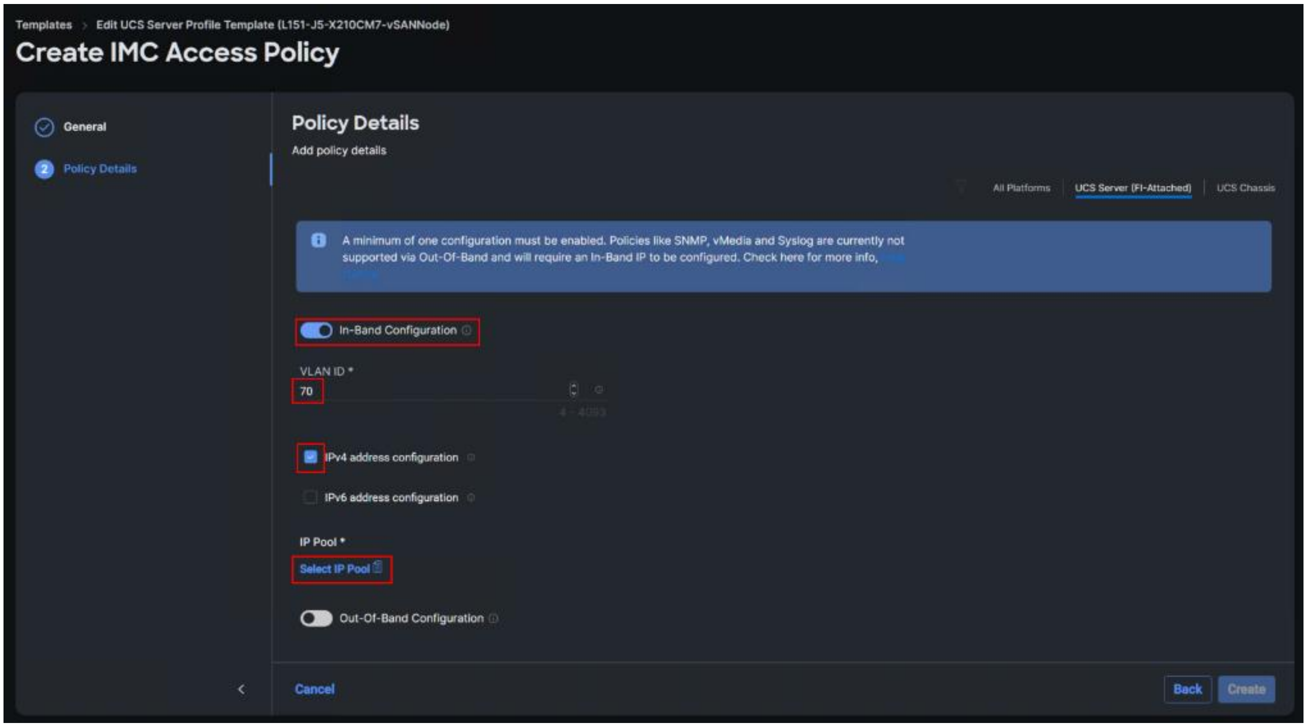

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create IMC Access Policy

General

2 Policy Details

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,

In-Band Configuration

VLAN ID *
70
4 - 4093

☑ IPv4 address configuration

□ IPv6 address configuration

IP Pool *
Select IP Pool

Out-Of-Band Configuration

Cancel | Back | Create

Step 31. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L152-J5-ICMA-IP-Pool). Click Next.

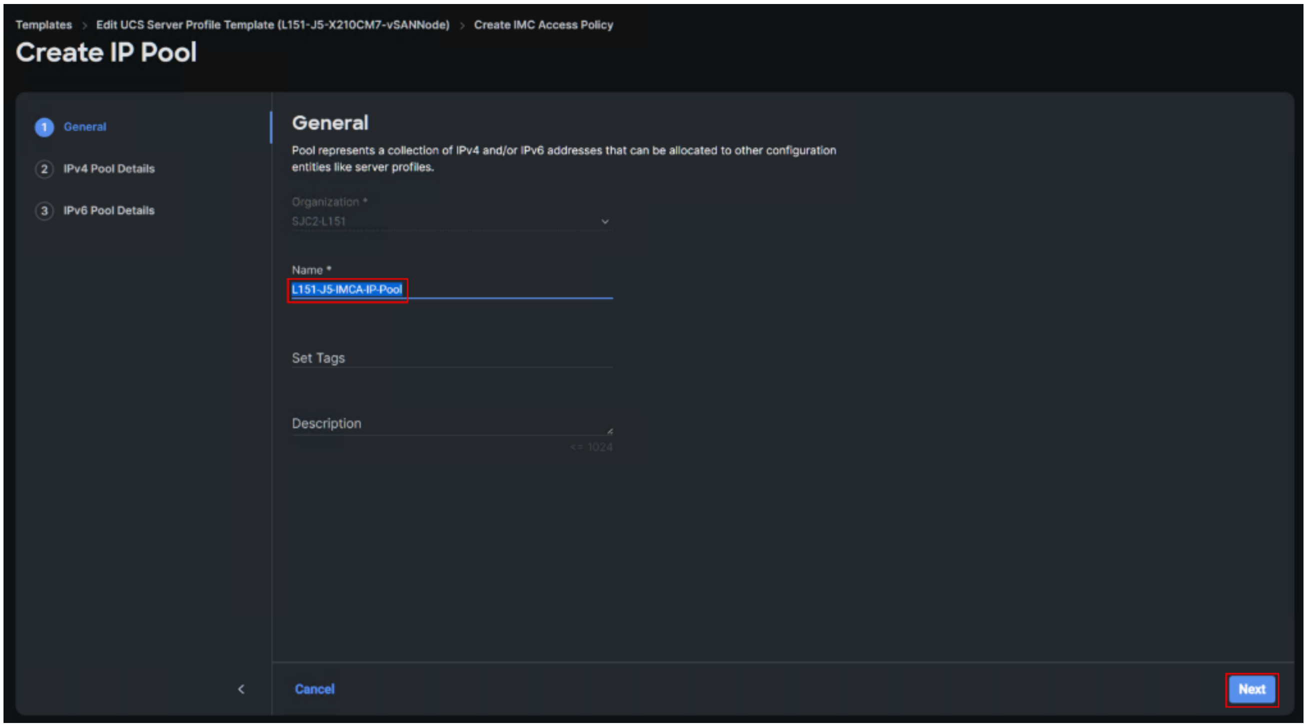

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create IMC Access Policy

Create IP Pool

1 General

2 IPv4 Pool Details

3 IPv6 Pool Details

General

Pool represents a collection of IPv4 and/or IPv6 addresses that can be allocated to other configuration entities like server profiles.

Organization *
SJC2-L151

Name *
L151-J5-IMCA-IP-Pool

Set Tags

Description
<= 1024

Cancel | Next

Step 32. Select Configure IPv4 Pool and provide the information to define a pool for KVM IP address assignment including an IP Block.

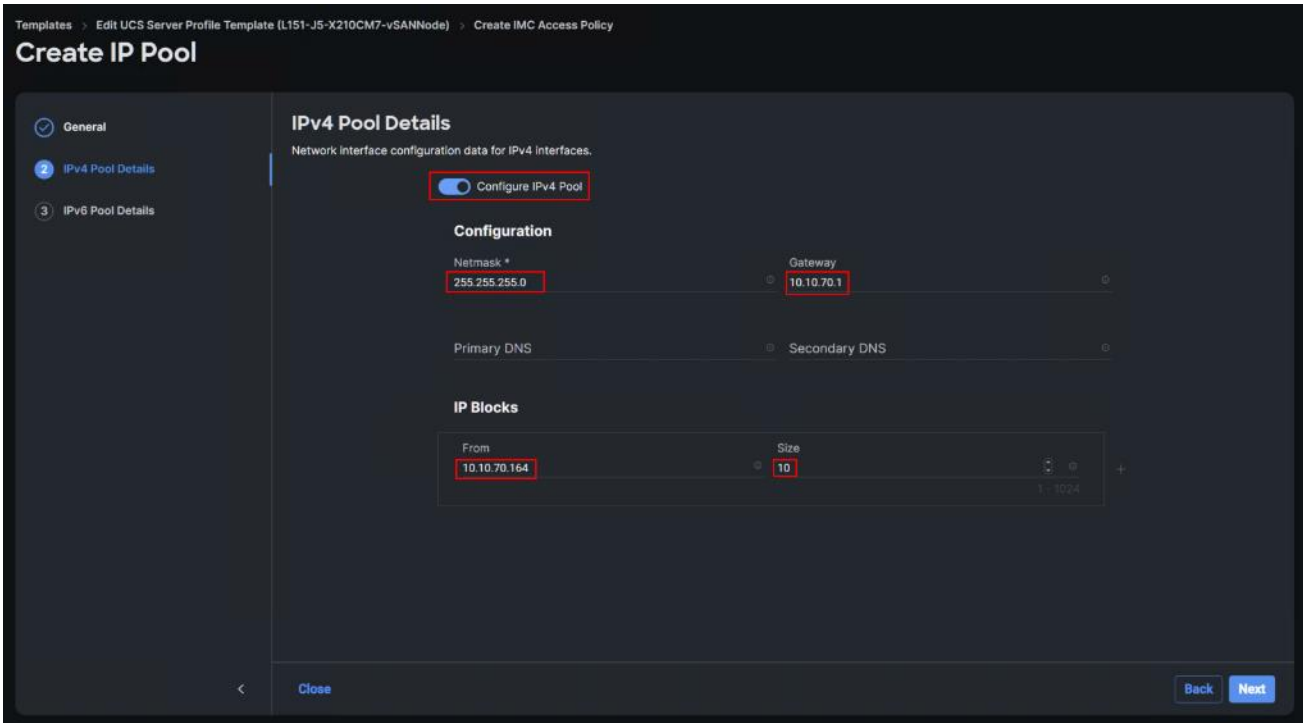

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create IMC Access Policy

Create IP Pool

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Configuration

Netmask *: 255.255.255.0

Gateway: 10.10.70.1

Primary DNS

Secondary DNS

IP Blocks

From: 10.10.70.164

Size: 10

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 10.10.70.0/24 subnet.

Step 33. Leave Configure IPv6 Pool disabled. Click Create.

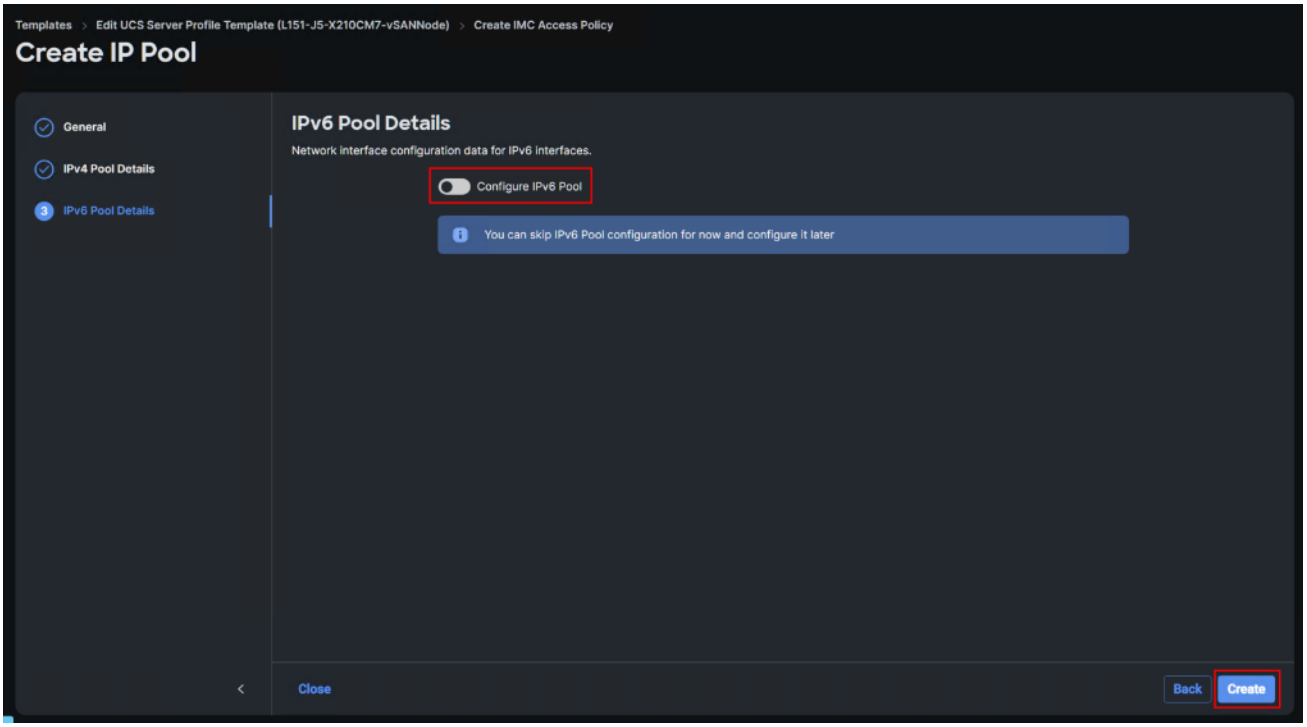

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create IMC Access Policy

Create IP Pool

IPv6 Pool Details

Network interface configuration data for IPv6 interfaces.

Configure IPv6 Pool

You can skip IPv6 Pool configuration for now and configure it later

Step 34. Click Create.

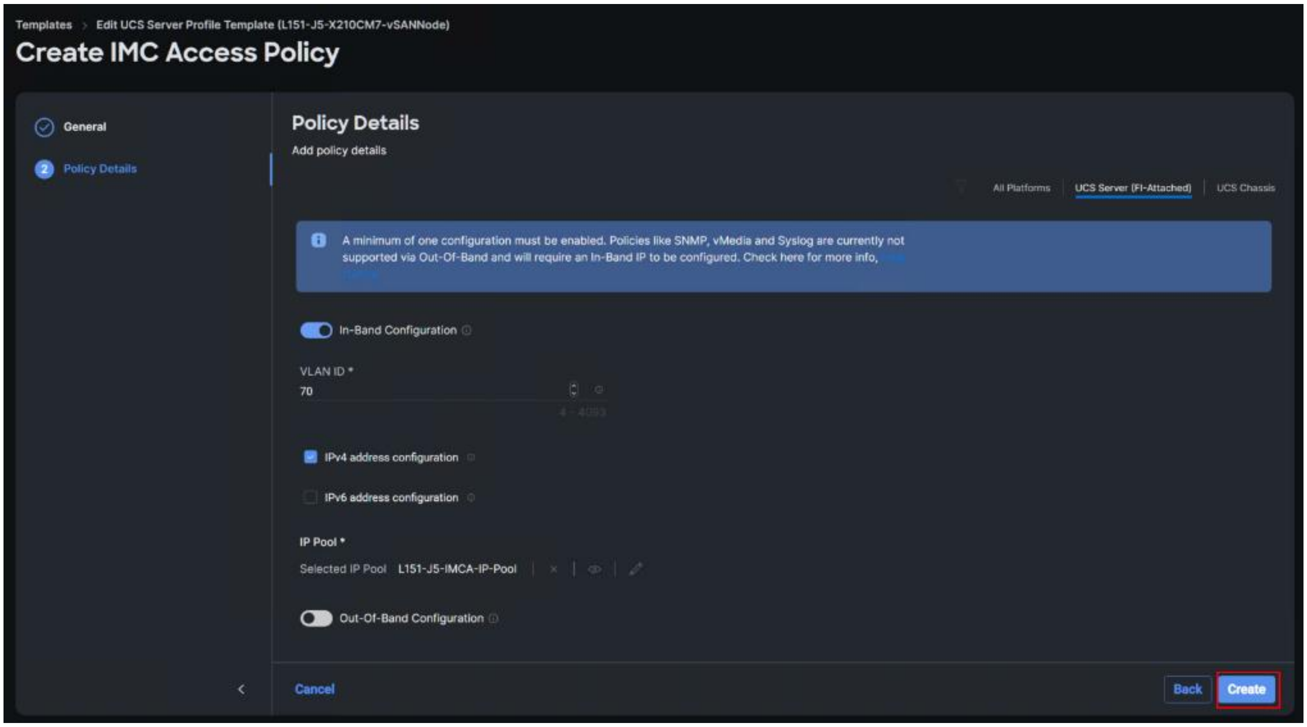

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create IMC Access Policy

Policy Details

Add policy details

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,

In-Band Configuration: on

VLAN ID *: 70 (4 - 4093)

☑ IPv4 address configuration

□ IPv6 address configuration

IP Pool *

Selected IP Pool L151-J5-IMCA-IP-Pool

Out-Of-Band Configuration: off

Step 35. Click Select Policy next to IPMI Over LAN and then click Create New.

Step 36. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-IPMIoLANPol). Click Next.

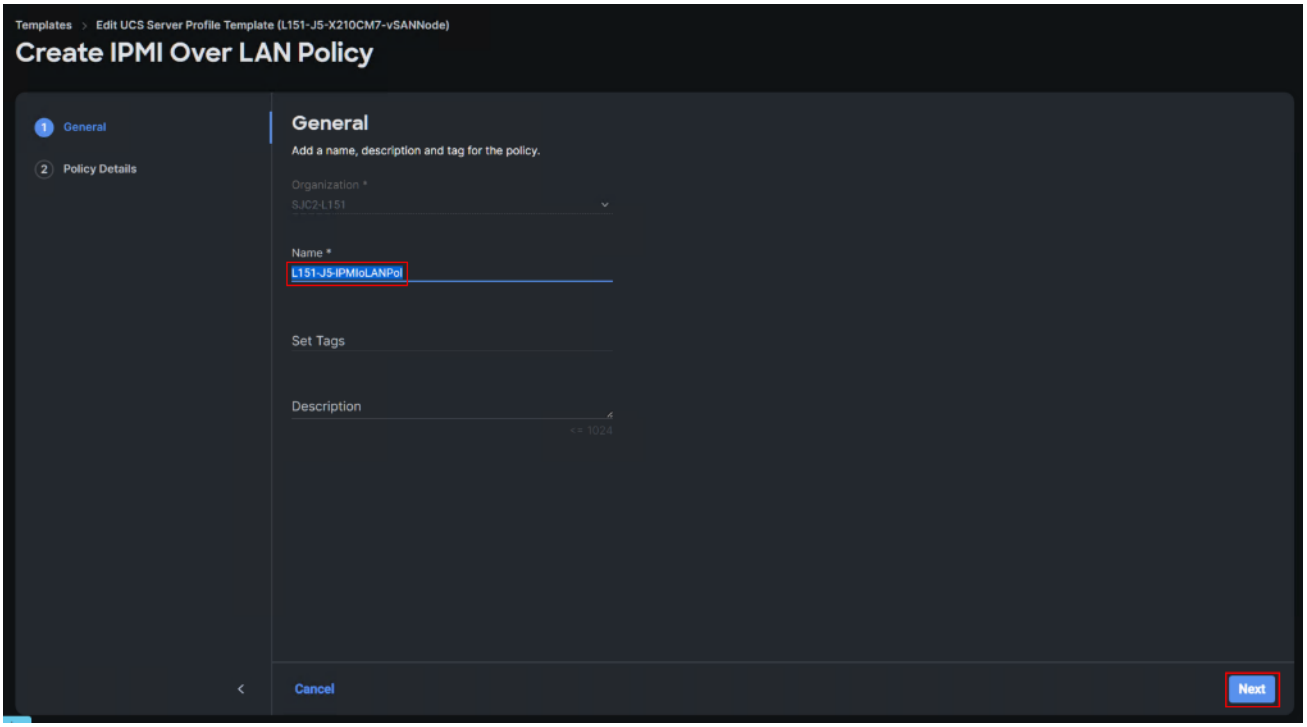

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create IPMI Over LAN Policy

General

Add a name, description and tag for the policy.

Organization *: SJC2-L151

Name *: L151-J5-IPMIoLANPol

Set Tags

Description (<= 1024)

Step 37. Turn on Enable IPMI Over LAN.

Step 38. From the Privilege Level drop-down list, select admin.

Step 39. Click Create.

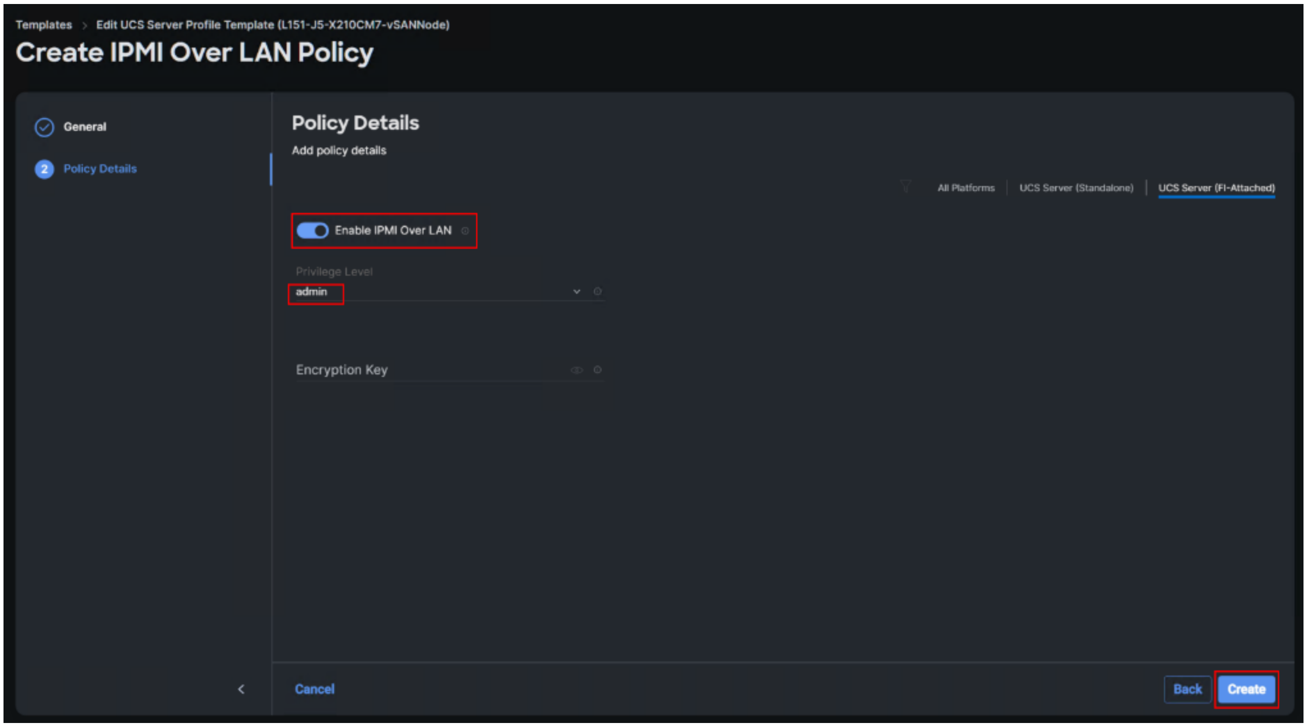

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create IPMI Over LAN Policy

Enable IPMI Over LAN

Privilege Level: admin

Encryption Key

Step 40. Click Select Policy next to Local User and the, in the pane on the right, click Create New.

Step 41. Verify the correct organization is selected from the drop-down list and provide a name for the policy. (for example, L151-J5-LocalUsrPol). Click Next.

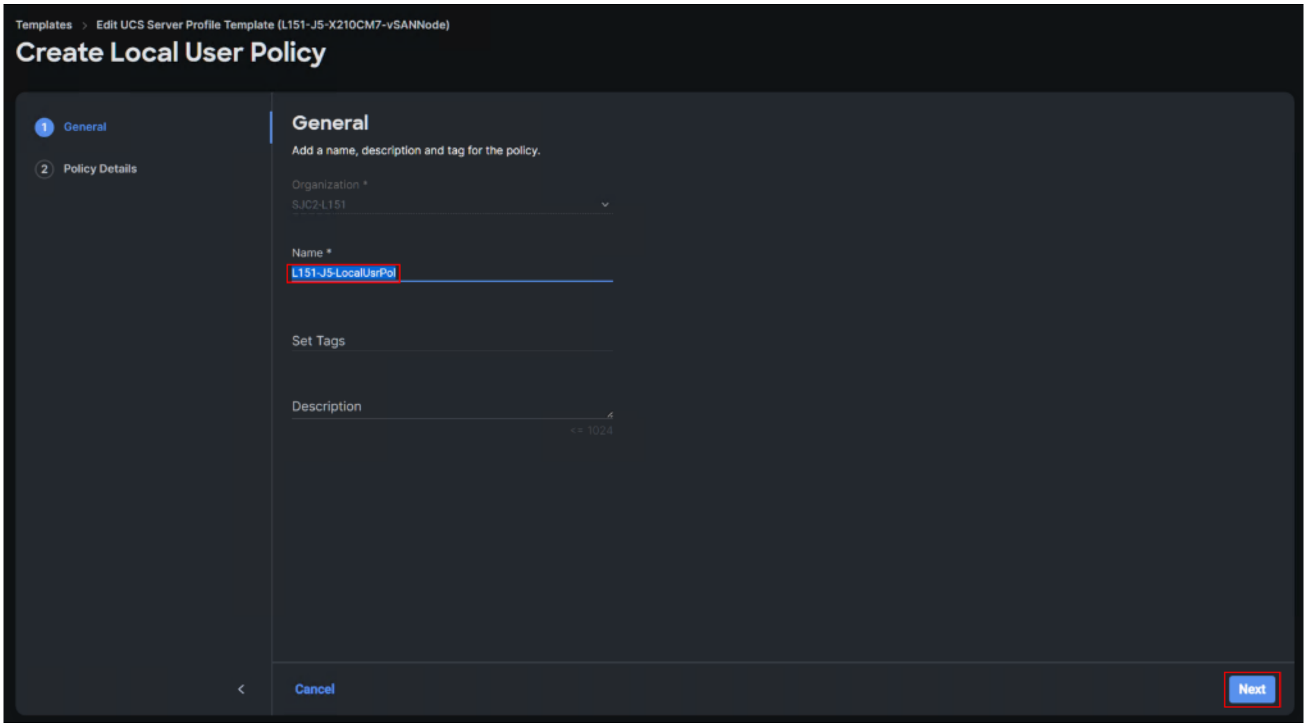

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Local User Policy

Organization *: SJC2-L151

Name *: L151-J5-LocalUsrPol

Set Tags

Description

Step 42. Verify that UCS Server (FI-Attached) is selected.

Step 43. Verify that Enforce Strong Password is selected.

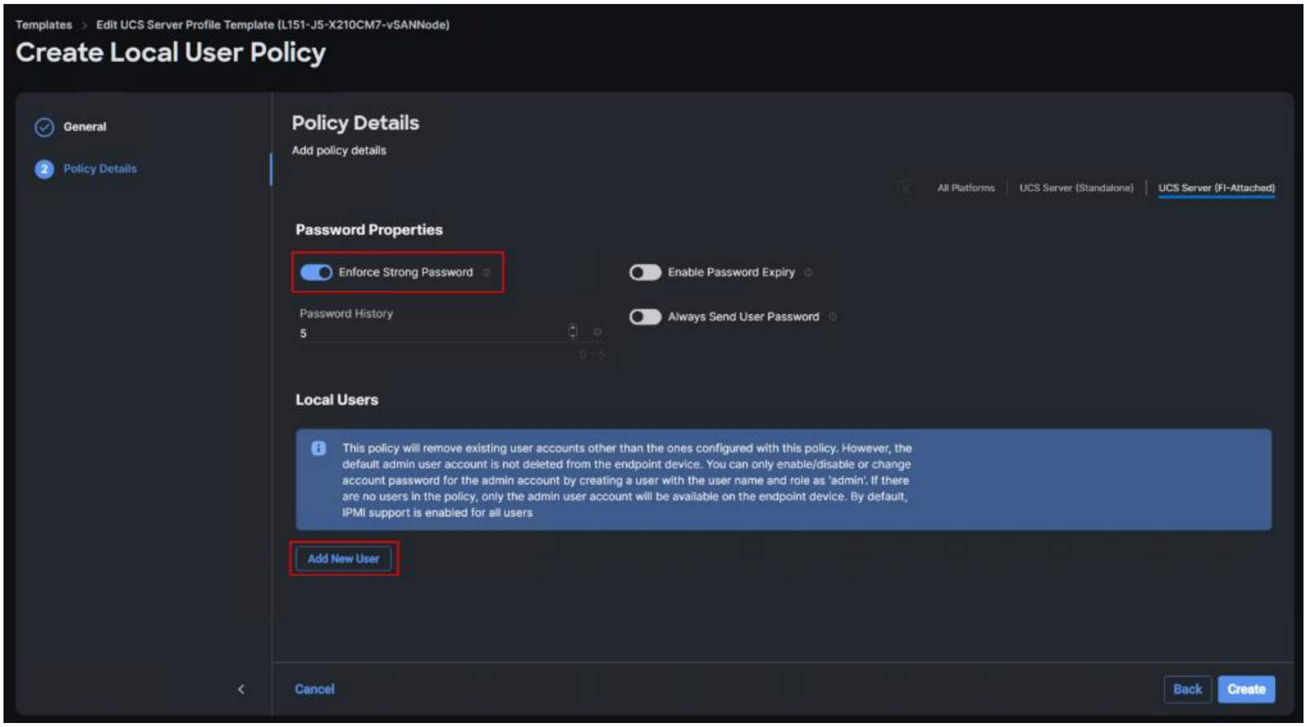

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Local User Policy

Policy Details

Add policy details

Password Properties

Enforce Strong Password: on

Enable Password Expiry: off

Password History: 5

Always Send User Password: off

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

Step 44. Click Add New User and then click + next to the New User.

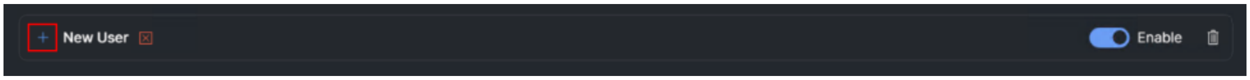

Step 45. Provide the username (for example, L151admin), choose a role for example, admin), and provide a password.

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 46. Click Create to finish configuring the user policy.

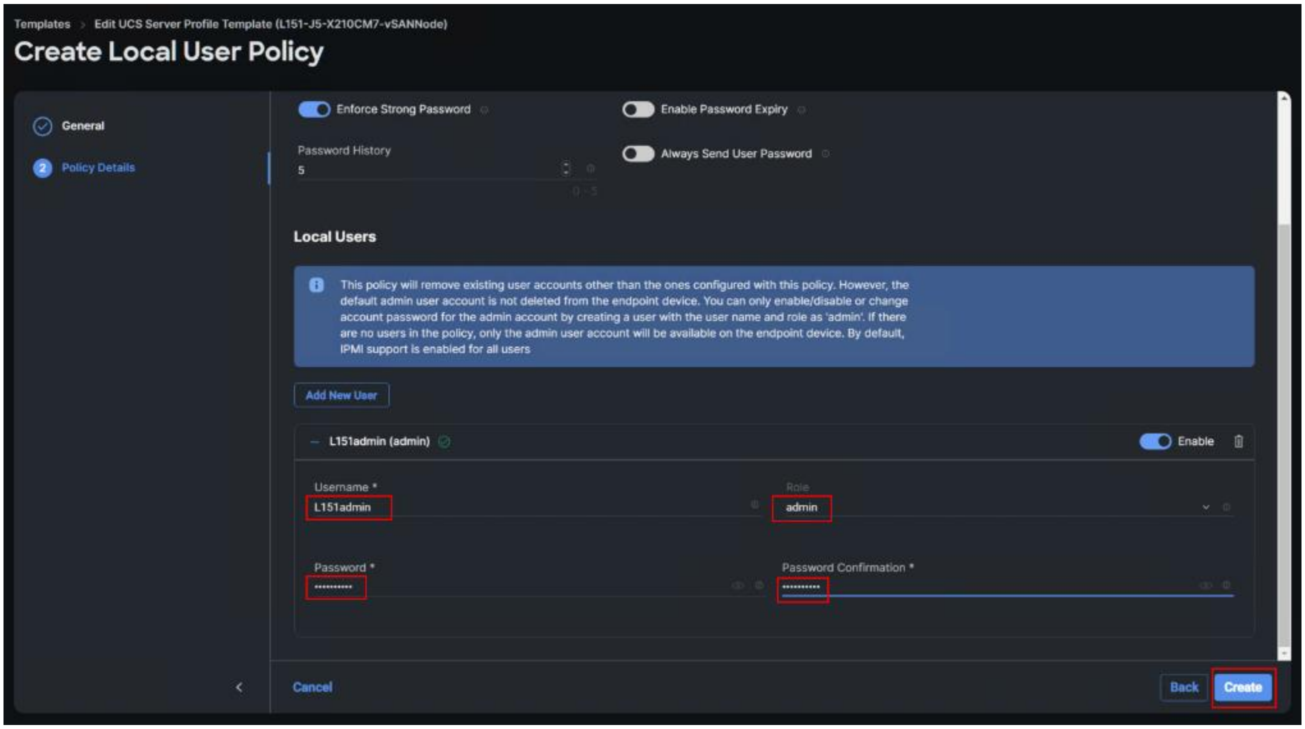

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Local User Policy

- General
- 2 Policy Details

Enforce Strong Password: enabled

Enable Password Expiry: disabled

Password History: 5 (0 - 5)

Always Send User Password: disabled

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

L151admin (admin) — Enable: enabled

Username *: L151admin

Role: admin

Password *: ••••••••••

Password Confirmation *: ••••••••••

Step 47. Click Select Policy next to Virtual KVM and the, in the pane on the right, click Create New.

Step 48. Verify the correct organization is selected from the drop-down list and provide a name for the policy. (for example, L151-J5-vKVMPol). Click Next.

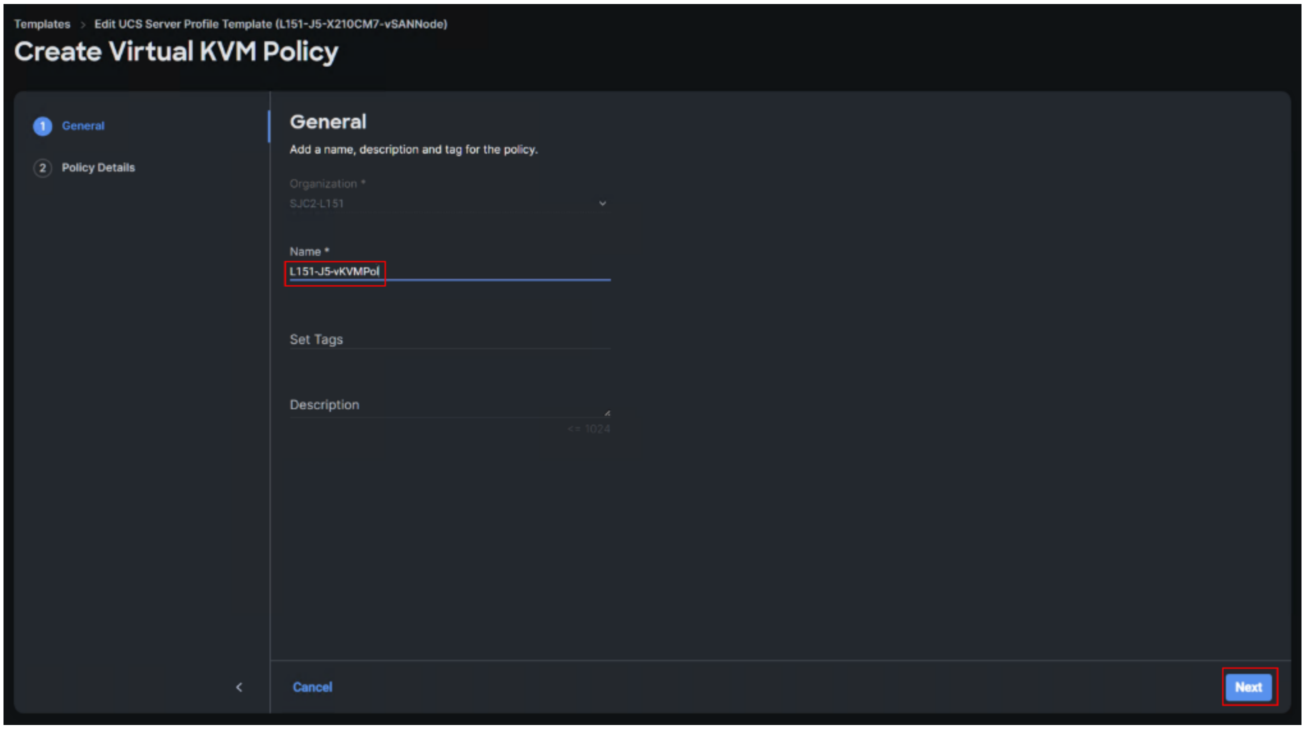

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Virtual KVM Policy

- 1 General
- 2 Policy Details

General

Add a name, description and tag for the policy.

Organization *: SJC2-L151

Name *: L151-J5-vKVMPol

Set Tags

Description (<= 1024)

Step 49. Enable all available options. Click Create.

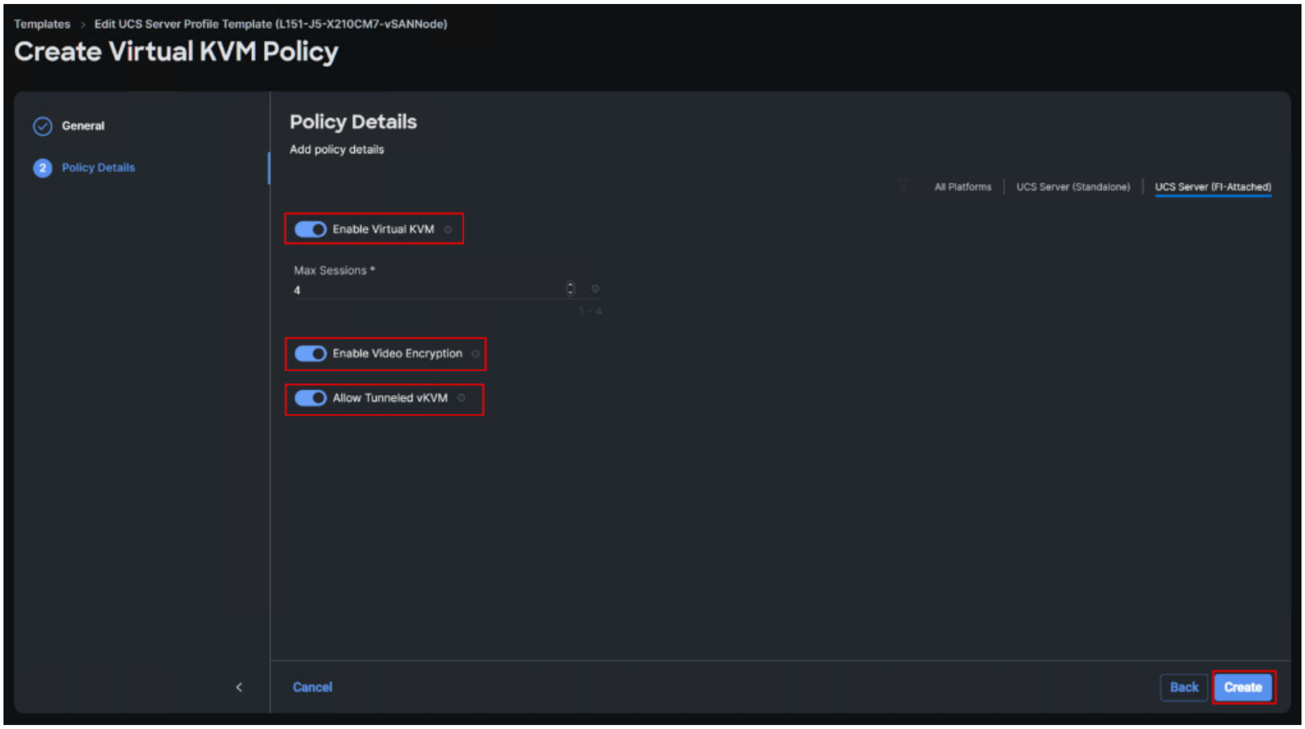

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Virtual KVM Policy

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

Enable Virtual KVM

Max Sessions *

4

1 - 4

Enable Video Encryption

Allow Tunneled vKVM

Step 50. Click Next to move to Storage Configuration.

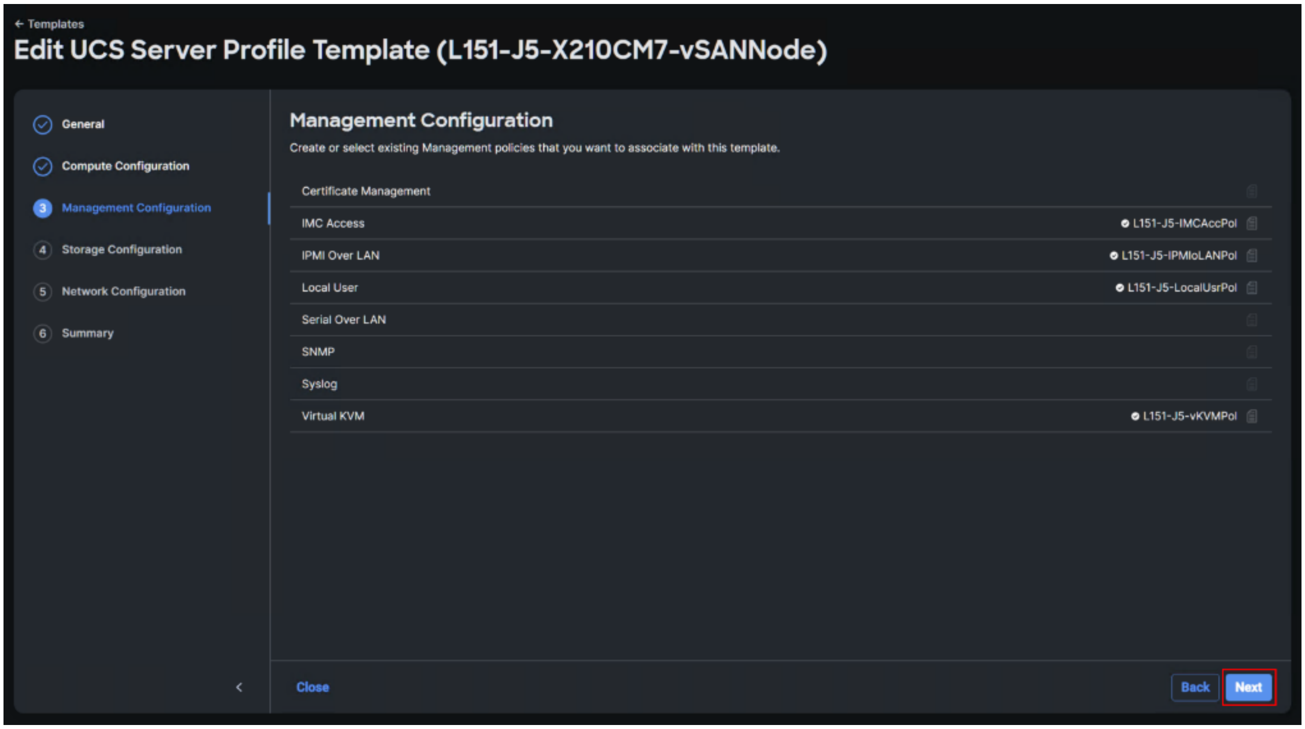

← Templates

Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Management Configuration

Create or select existing Management policies that you want to associate with this template.

Certificate Management	
IMC Access	L151-J5-IMCAccPol
IPMI Over LAN	L151-J5-IPMIoLANPol
Local User	L151-J5-LocalUsrPol
Serial Over LAN	
SNMP	
Syslog	
Virtual KVM	L151-J5-vKVMPol

Step 51. Click Select Policy next to Storage and the, in the pane on the right, click Create New.

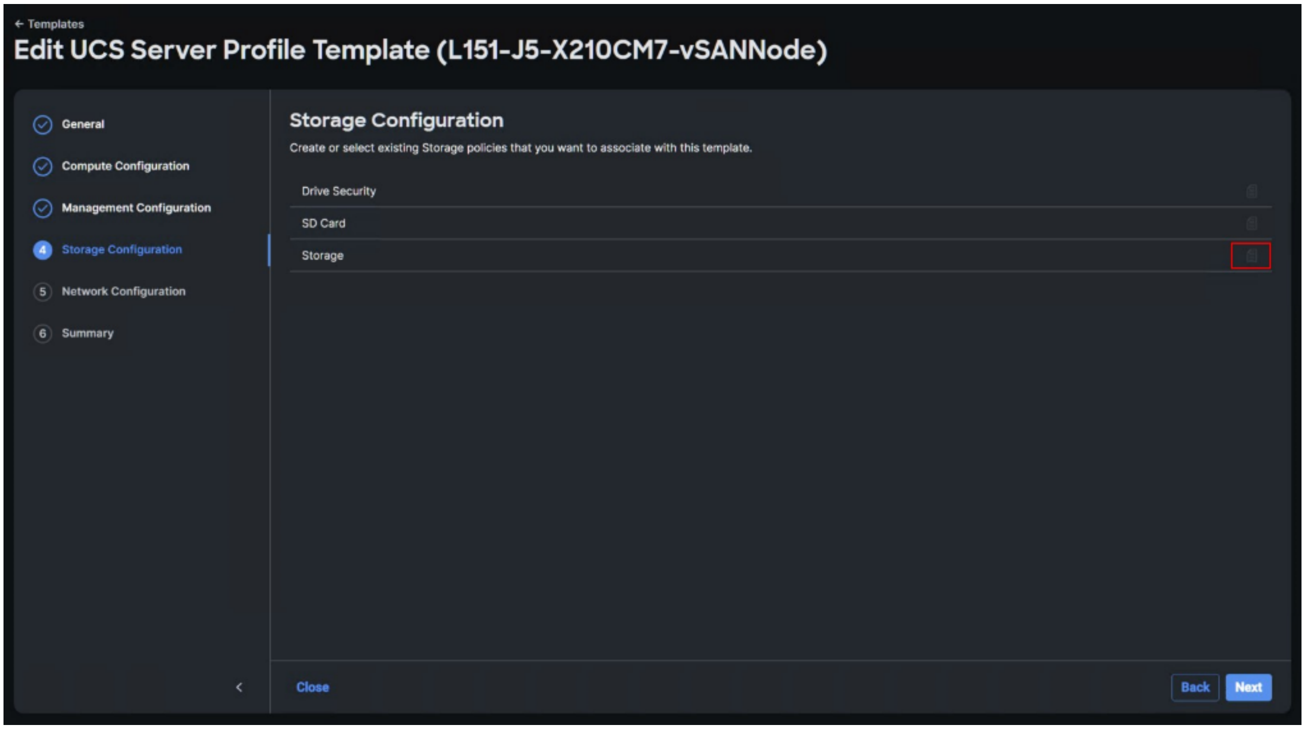

Step 52. Verify the correct organization is selected from the drop-down list and provide a name for the policy. (for example, L151-J5-X210CM7-vSANnode-M2ConfPol). Click Next.

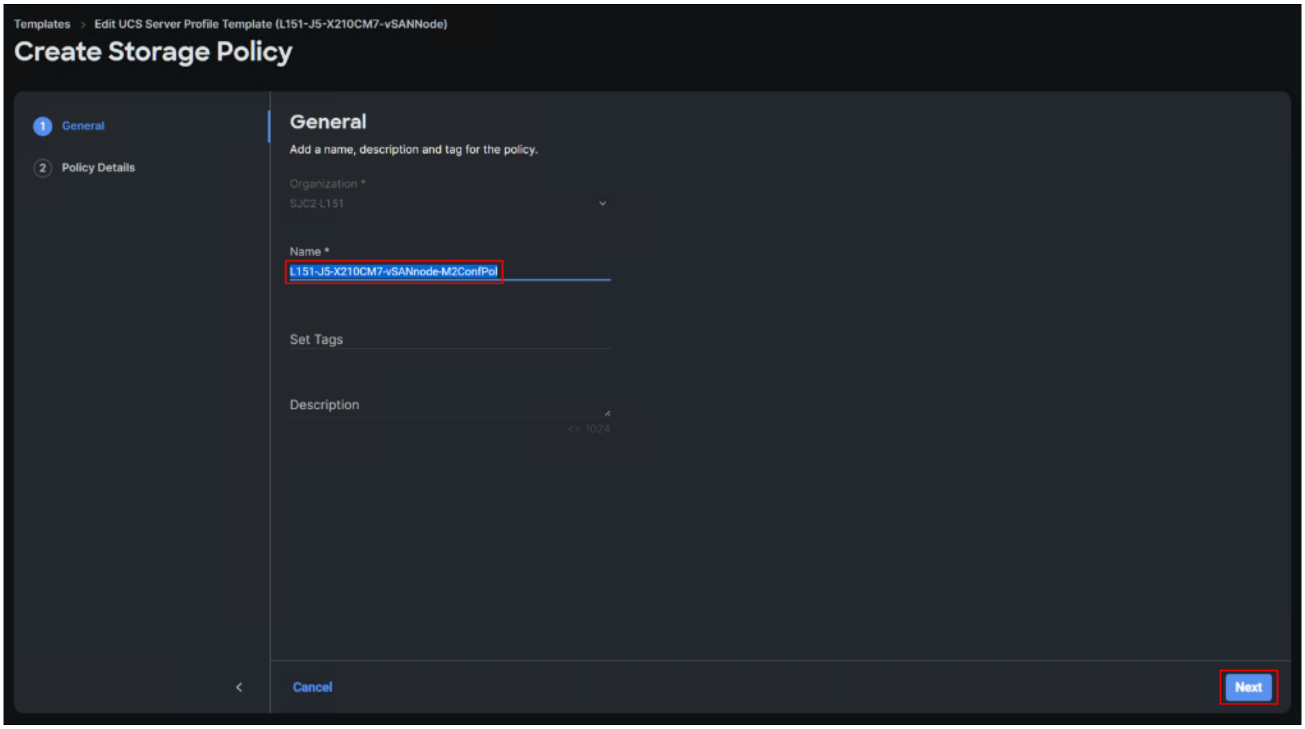
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)
Create Storage Policy
1 General
2 Policy Details
General
Add a name, description and tag for the policy.
Organization *
SJC2-L151
Name *
L151-J5-X210CM7-vSANnode-M2ConfPol
Set Tags
Description
<= 1024
Cancel
Next

Step 53. Enable M.2 RAID Configuration and select MSTOR-RAID-1 from virtual drive configuration drop-down list. Click Next.

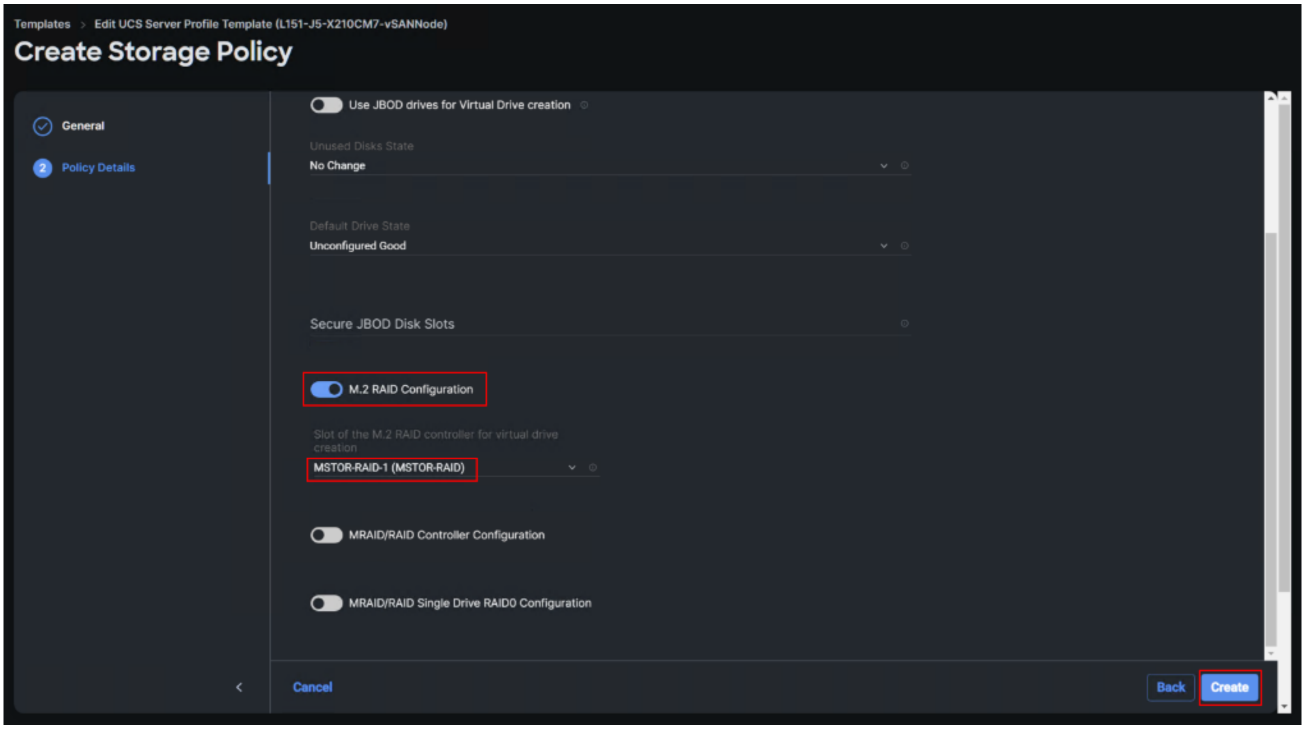

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create Storage Policy

Use JBOD drives for Virtual Drive creation

Unused Disks State: No Change

Default Drive State: Unconfigured Good

Secure JBOD Disk Slots

M.2 RAID Configuration

Slot of the M.2 RAID controller for virtual drive creation: MSTOR-RAID-1 (MSTOR-RAID)

MRAID/RAID Controller Configuration

MRAID/RAID Single Drive RAID0 Configuration

Step 54. Click Next to move to Network Configuration.

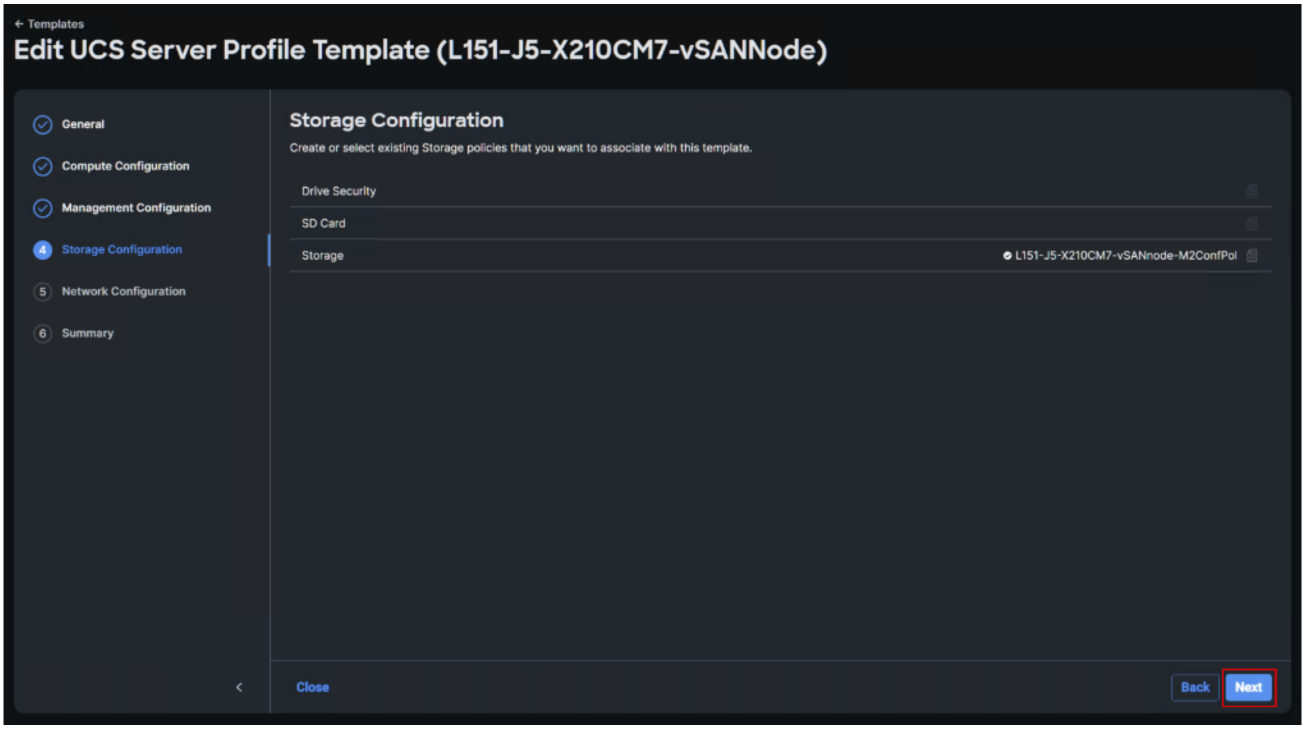

← Templates

Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Storage Configuration

Create or select existing Storage policies that you want to associate with this template.

Drive Security

SD Card

Storage: L151-J5-X210CM7-vSANnode-M2ConfPol

Step 55. Click Select Policy next to LAN Connectivity and then click Create New.

Note: LAN connectivity policy defines the connections and network communication resources between the server and the LAN. This policy uses pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For consistent vNIC placement, manual vNIC placement is utilized.

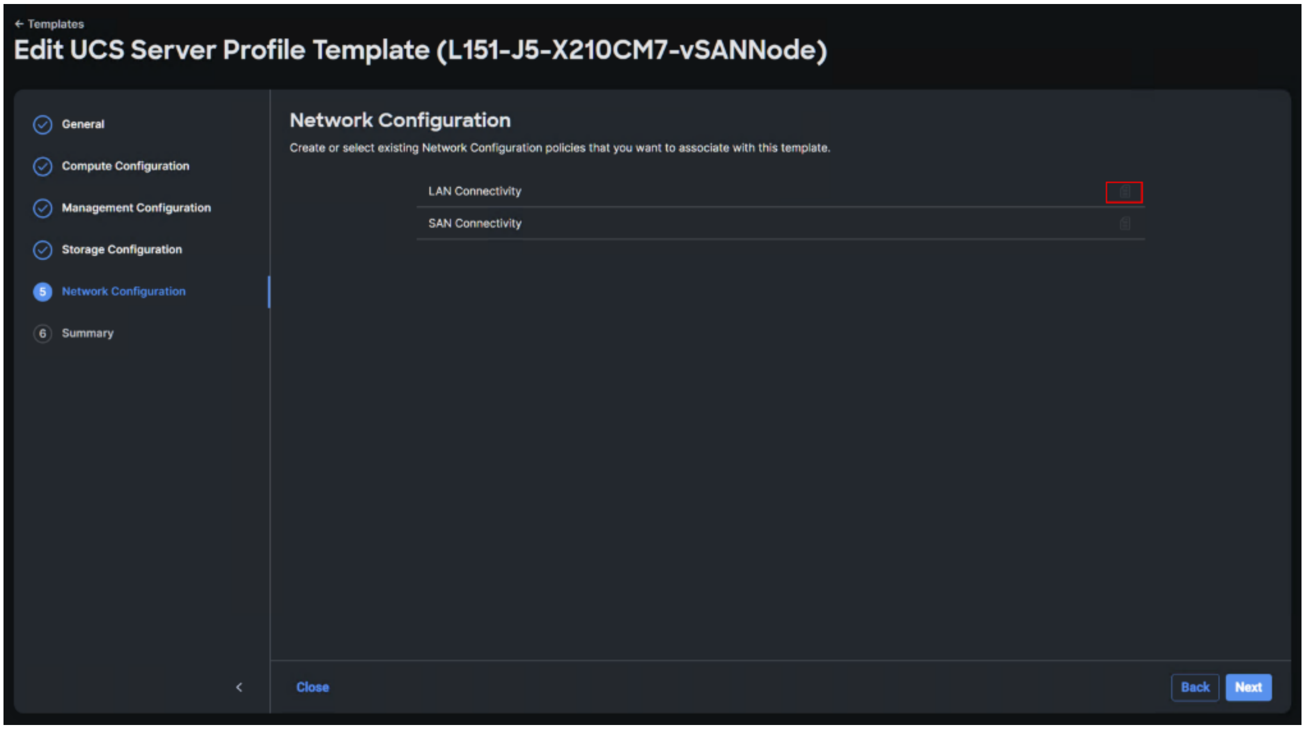

← Templates

Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Network Configuration

Create or select existing Network Configuration policies that you want to associate with this template.

LAN Connectivity

SAN Connectivity

The vSAN node uses 4 vNICs configured as listed in Table 7.

Table 7. vNICs for LAN Connectivity

vNIC	Slot ID	Switch ID	PCI Order	VLANs
vSwitch0-A	MLOM	A	0	InBand-Mgmt_70
vSwitch0-B	MLOM	B	1	InBand-Mgmt_70
VDS0-A	MLOM	A	2	VDI_72, vMotion_73, vSAN_74
VDS0-B	MLOM	B	3	VDI_72, vMotion_73, vSAN_74

Step 56. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-LANConPol). Click Next.

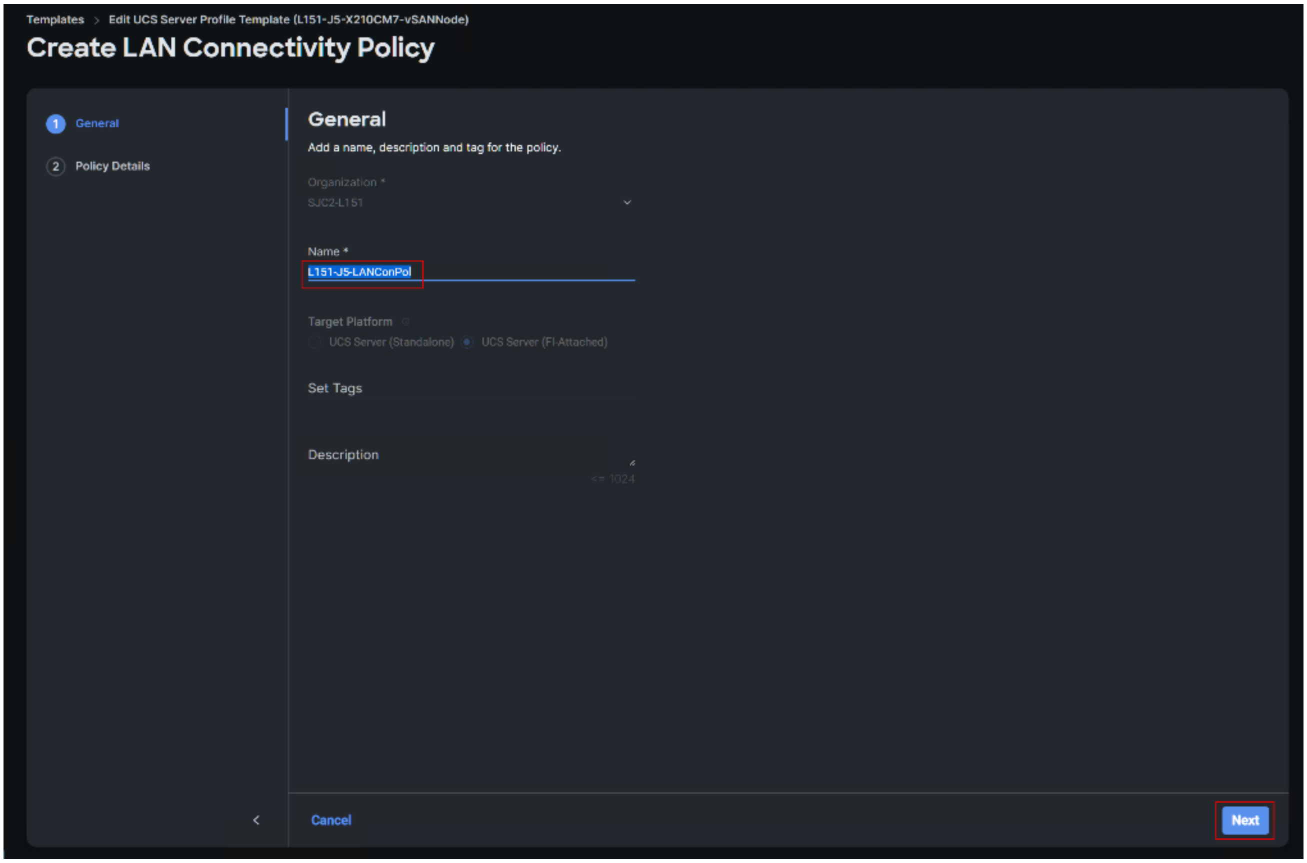
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

General

Add a name, description and tag for the policy.

Organization * SJC2-L151

Name * L151-J5-LANConPol

Target Platform ○ UCS Server (Standalone) ● UCS Server (FI-Attached)

Set Tags

Description <= 1024

Step 57. Under vNIC Configuration, select Manual vNICs Placement.

Step 58. Click Add vNIC.

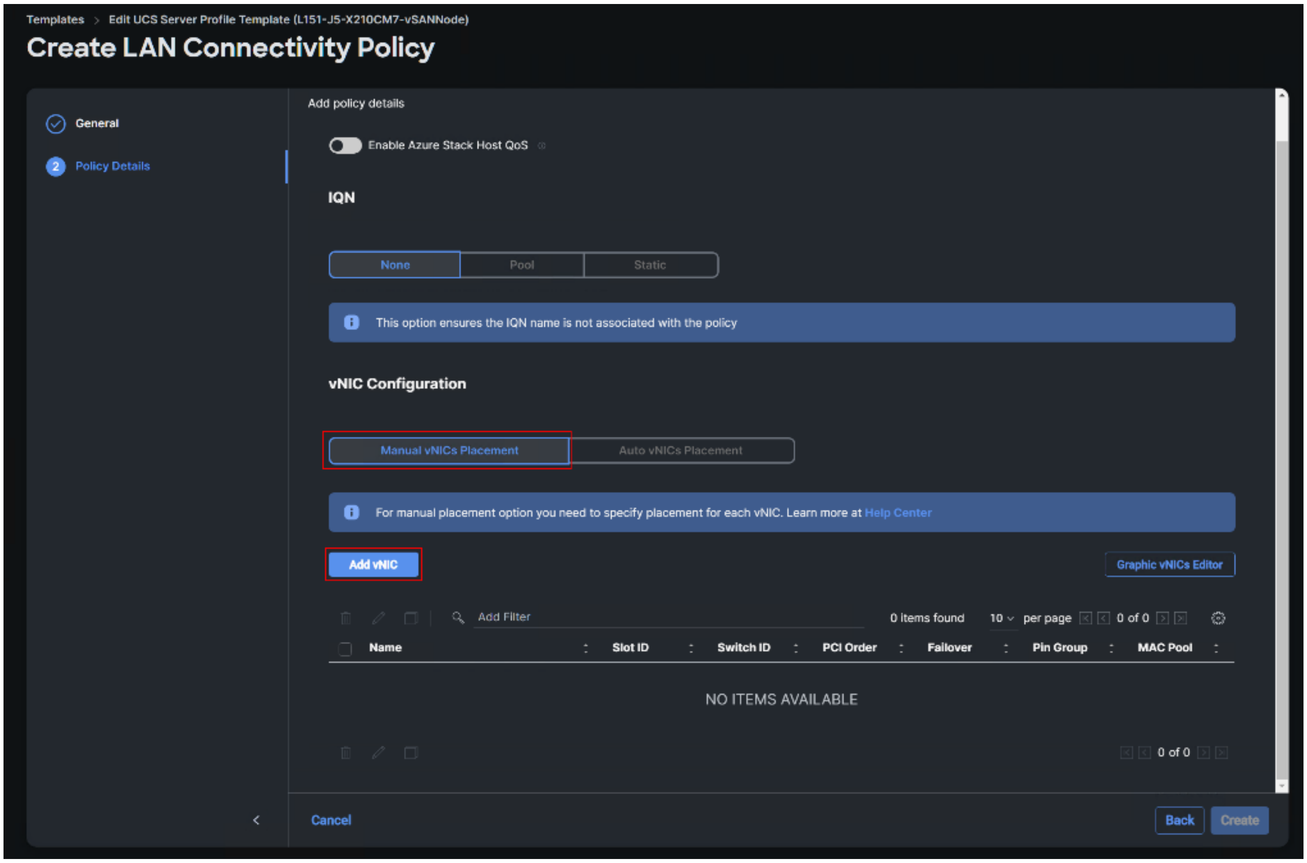

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Add policy details

Enable Azure Stack Host QoS

IQN

[✓ None] Pool Static

This option ensures the IQN name is not associated with the policy

vNIC Configuration

[✓ Manual vNICs Placement] Auto vNICs Placement

For manual placement option you need to specify placement for each vNIC. Learn more at Help Center

0 items found 10 per page 0 of 0

Name	Slot ID	Switch ID	PCI Order	Failover	Pin Group	MAC Pool

NO ITEMS AVAILABLE

0 of 0

Step 59. Click Select Pool under MAC Pool and then click Create New.

Note: 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. MAC-Pool-A will be reused for all Fabric-A vNICs, and MAC-Pool-B will be reused for all Fabric-B vNICs.

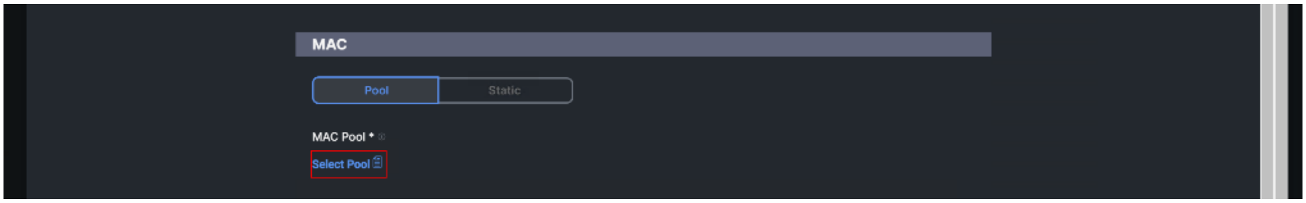

MAC

[✓ Pool] Static

MAC Pool *

Table 8. MAC Address Pools

Pool Name	Starting MAC Address	Size	vNICs
L151-J5-MAC-Pool-A	00:25:B5:04:A0:00	32*	vSwitch0-A, VDS0-A
L151-J5-MAC-Pool-B	00:25:B5:04:B0:00	32*	vSwitch0-B, VDS0-B

Step 60. Verify the correct organization is selected from the drop-down list and provide a name for the pool from Table 8 depending on the vNIC being created (for example, L151-J5-MAC-Pool-A for Fabric A).

Step 61. Click Next.

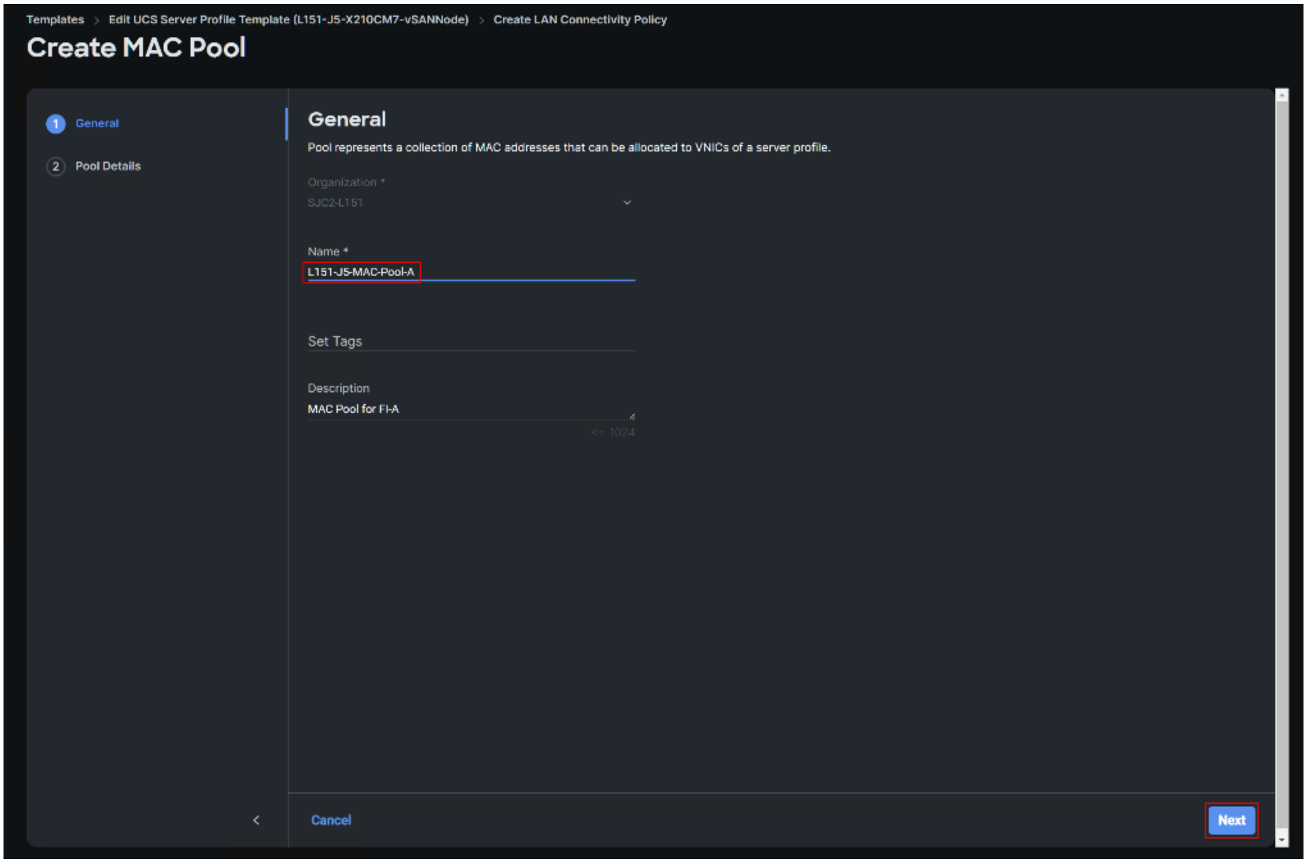

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create MAC Pool

General

Pool represents a collection of MAC addresses that can be allocated to VNICs of a server profile.

Organization *
SJC2-L151

Name *
L151-J5-MAC-Pool-A

Set Tags

Description
MAC Pool for FI-A
<= 1024

Step 62. Provide the starting MAC address from Table 8 (for example, 00:25:B5:05:A0:00) and the size of the MAC address pool (for example, 32). Click Create to finish creating the MAC address pool.

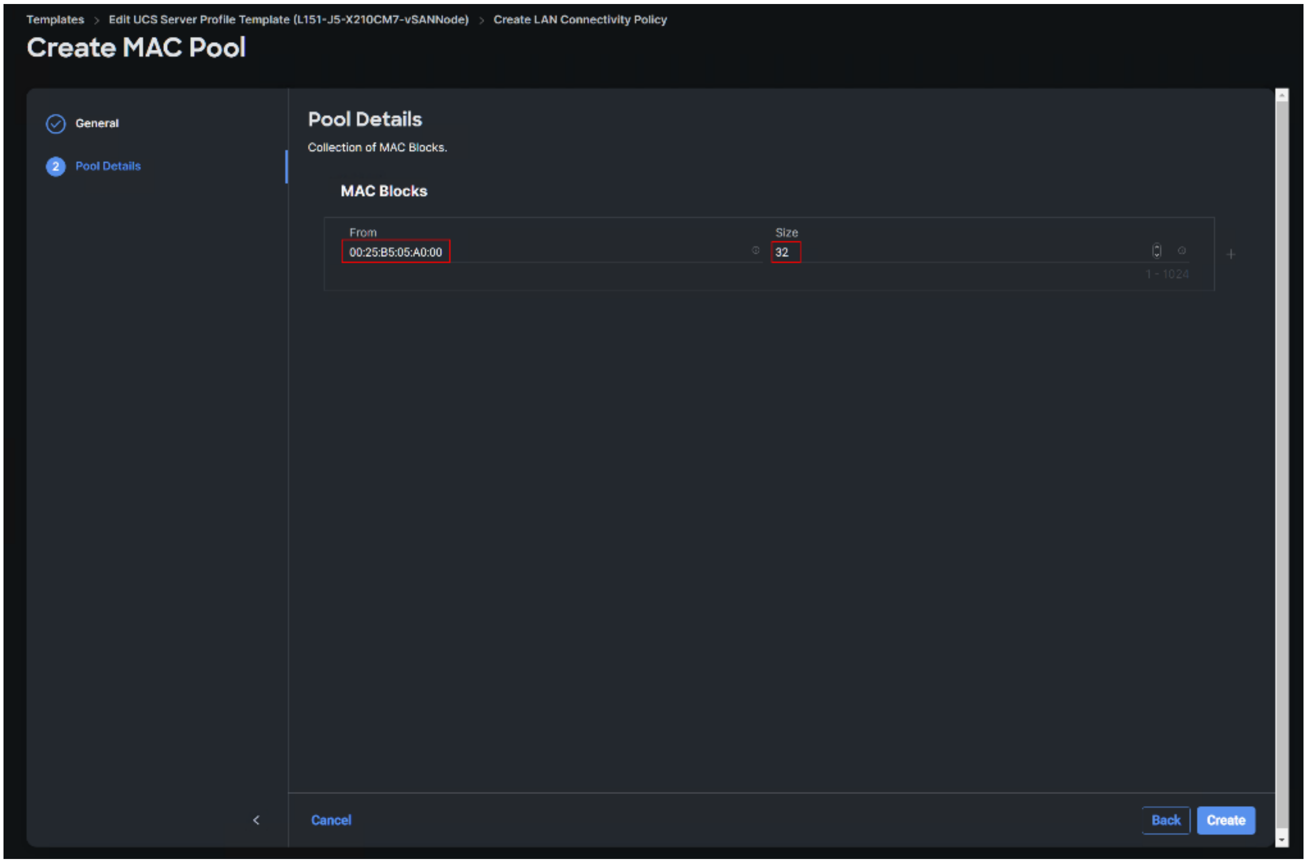
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy
Create MAC Pool
General
2 Pool Details
Pool Details
Collection of MAC Blocks.
MAC Blocks
1 - 1024
Cancel
Back
Create

From: 00:25:B5:05:A0:00

Size: 32

Step 63. From the Add vNIC window, provide vNIC Name, Slot ID, Switch ID, and PCI Order information under Advanced placement section from Table 8.

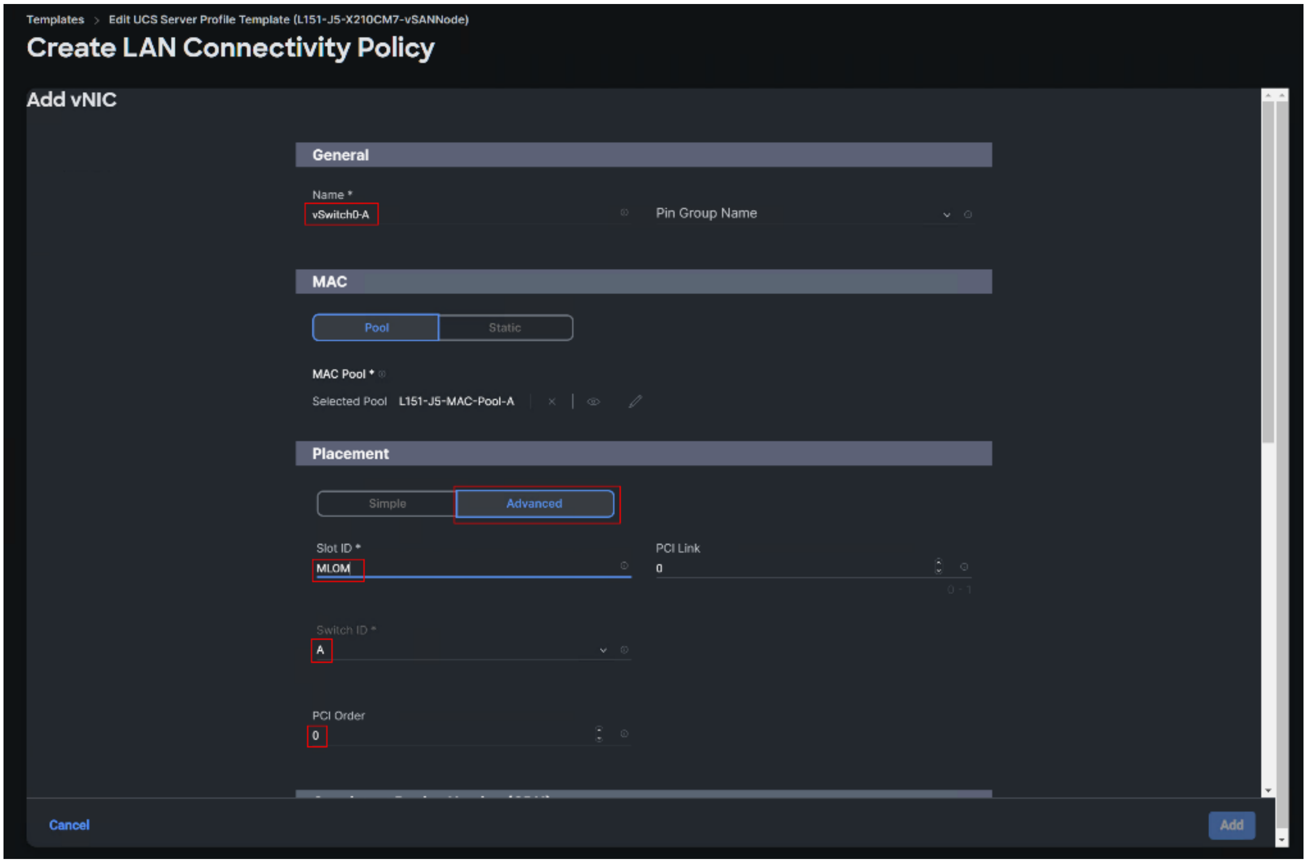
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Add vNIC

General

Name * vSwitch0-A

Pin Group Name

MAC

MAC Pool *

Selected Pool L151-J5-MAC-Pool-A

Placement

Slot ID * MLOM

PCI Link 0

Switch ID * A

PCI Order 0

Step 64. For Consistent Device Naming (CDN), from the drop-down list, select vNIC Name.

Step 65. Verify that Failover is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and VDS.

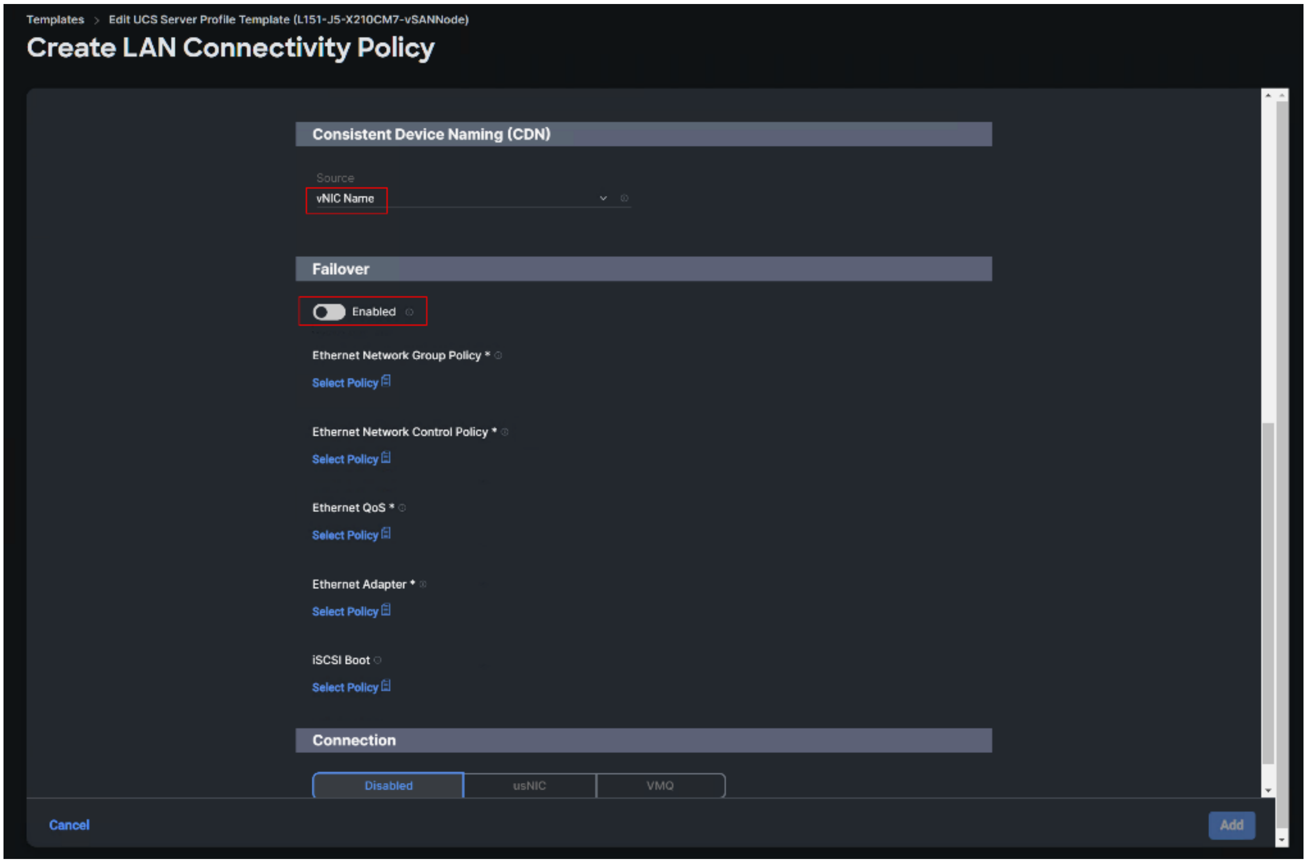

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled

Step 66. Click Select Policy under Ethernet Network Group Policy and then click Create New.

Note: 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 listed in Table 9.

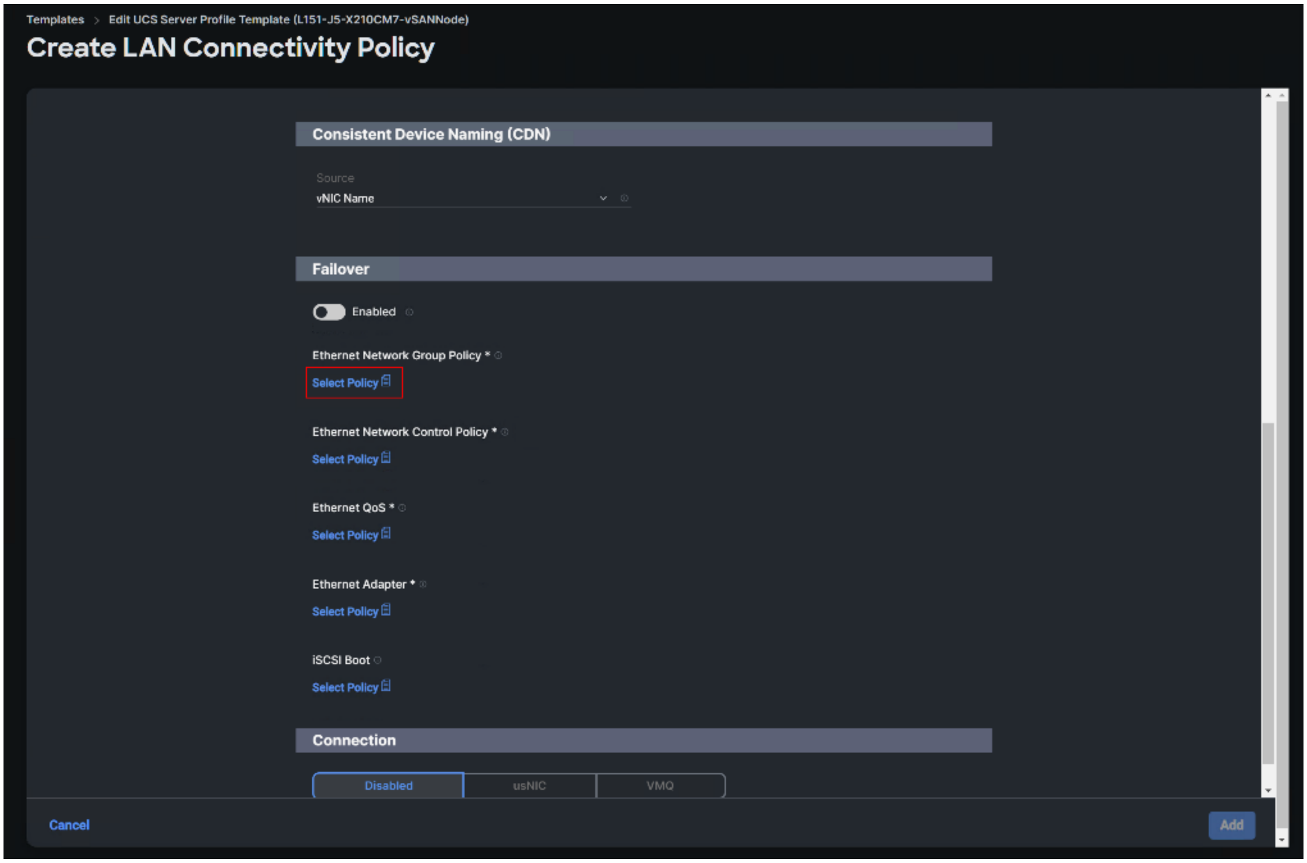
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled

Ethernet Network Group Policy *

Ethernet Network Control Policy *

Ethernet QoS *

Ethernet Adapter *

iSCSI Boot

Connection

Table 9. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs	VLANs
L151-J5-vSwitch0-NetGrp-Pol	Native-VLAN (1)	vSwitch0-A, vSwitch0-B	InBand-Mgmt_70
L151-J5-vDS0-NetGrp-Pol	Native-VLAN (1)	VDS0-A, VDS0-B	VDI_72, vMotion_73, vSAN_74

Step 67. Verify the correct organization is selected from the drop-down list and provide a name for the policy from Table 9 (for example, L151-J5-vSwitch0-NetGrp-Pol). Click Next.

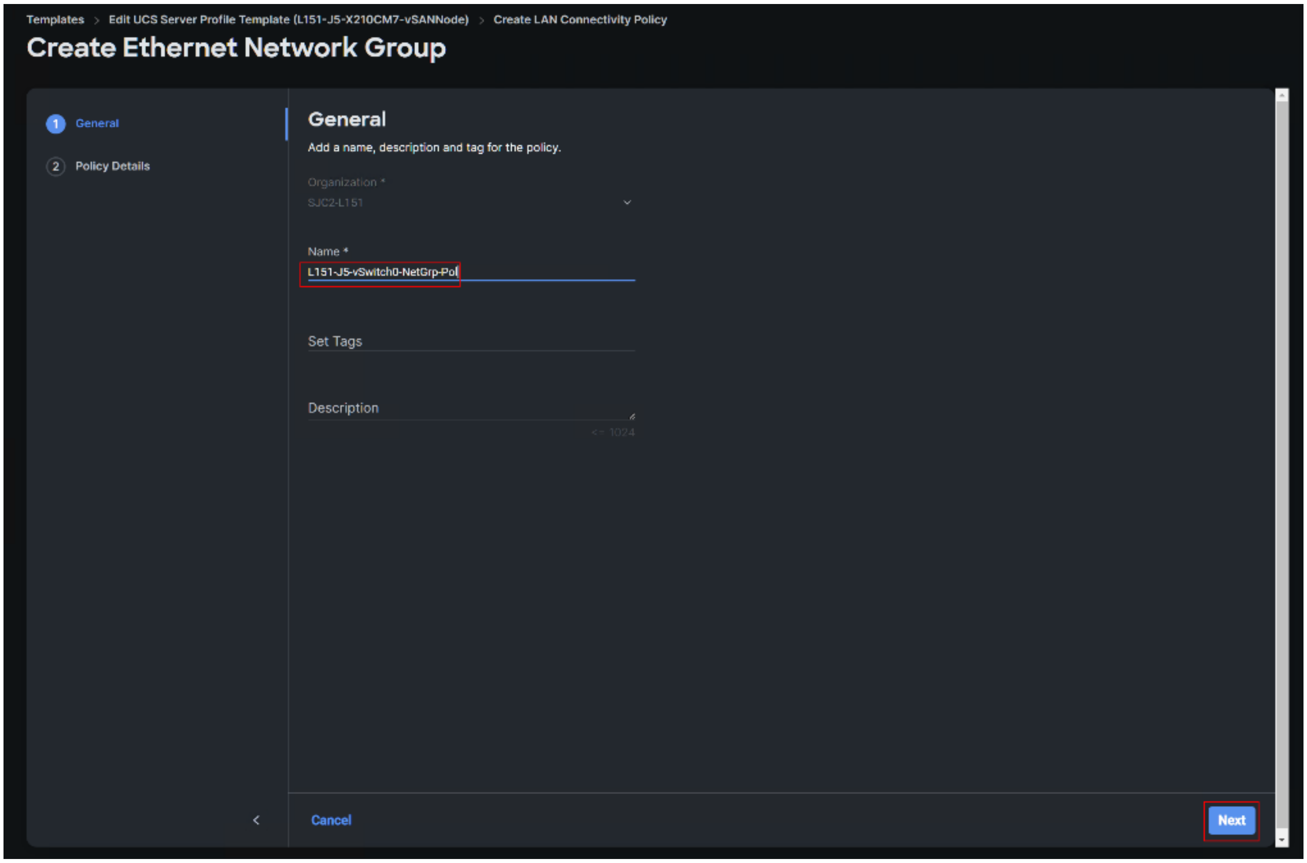
Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet Network Group

1 General
2 Policy Details

General

Add a name, description and tag for the policy.

Organization * SJC2-L151

Name * L151-J5-vSwitch0-NetGrp-Pol

Set Tags

Description

Step 68. Enter the allowed VLANs from Table 8 (for example, 70) and the native VLAN ID from Table 9 (for example, 1). Click Create.

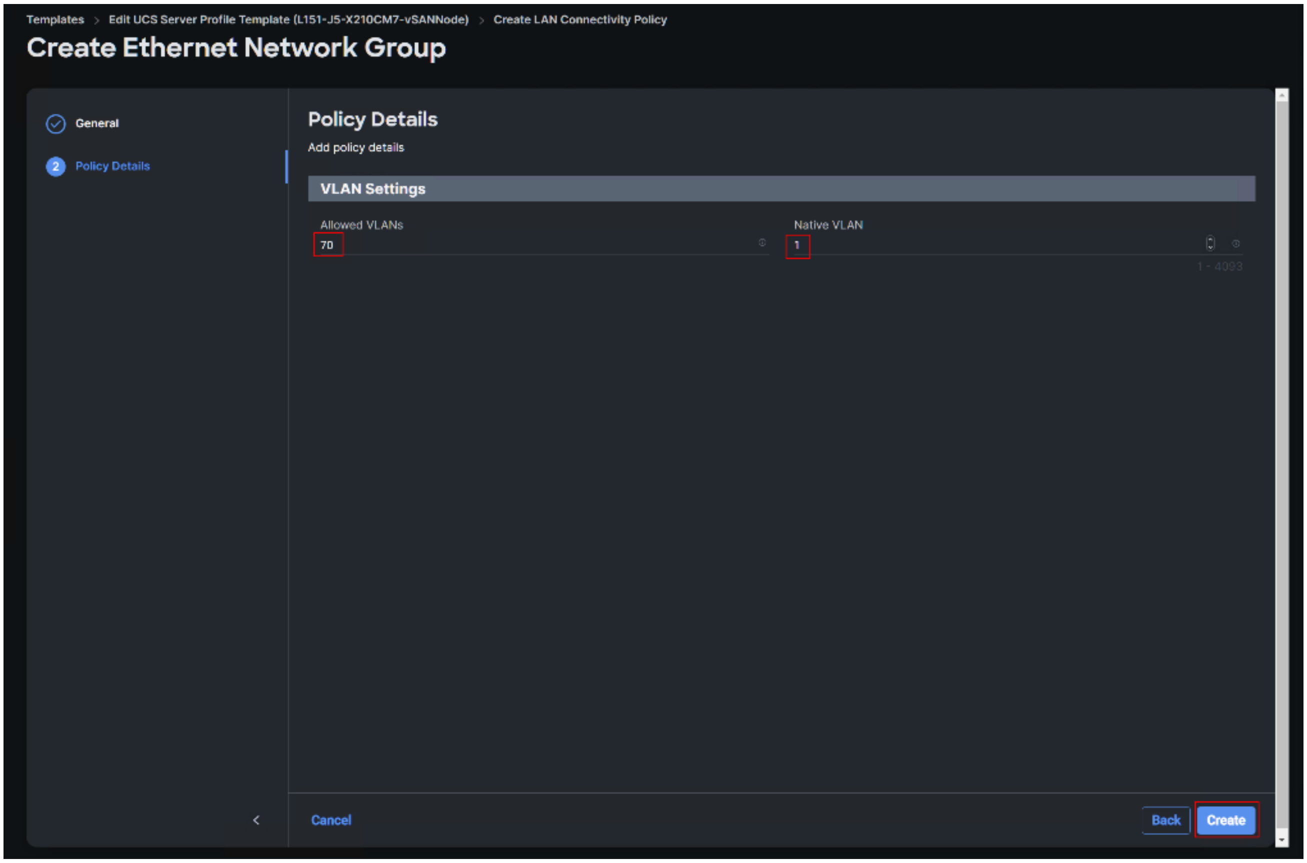

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet Network Group

Policy Details

Add policy details

VLAN Settings

Allowed VLANs: 70

Native VLAN: 1

1 - 4093

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, just click Select Policy and pick the previously defined ethernet group policy from the list on the right.

Step 69. Click Select Policy under Ethernet Network Control Policy and then click Create New.

Note: 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.

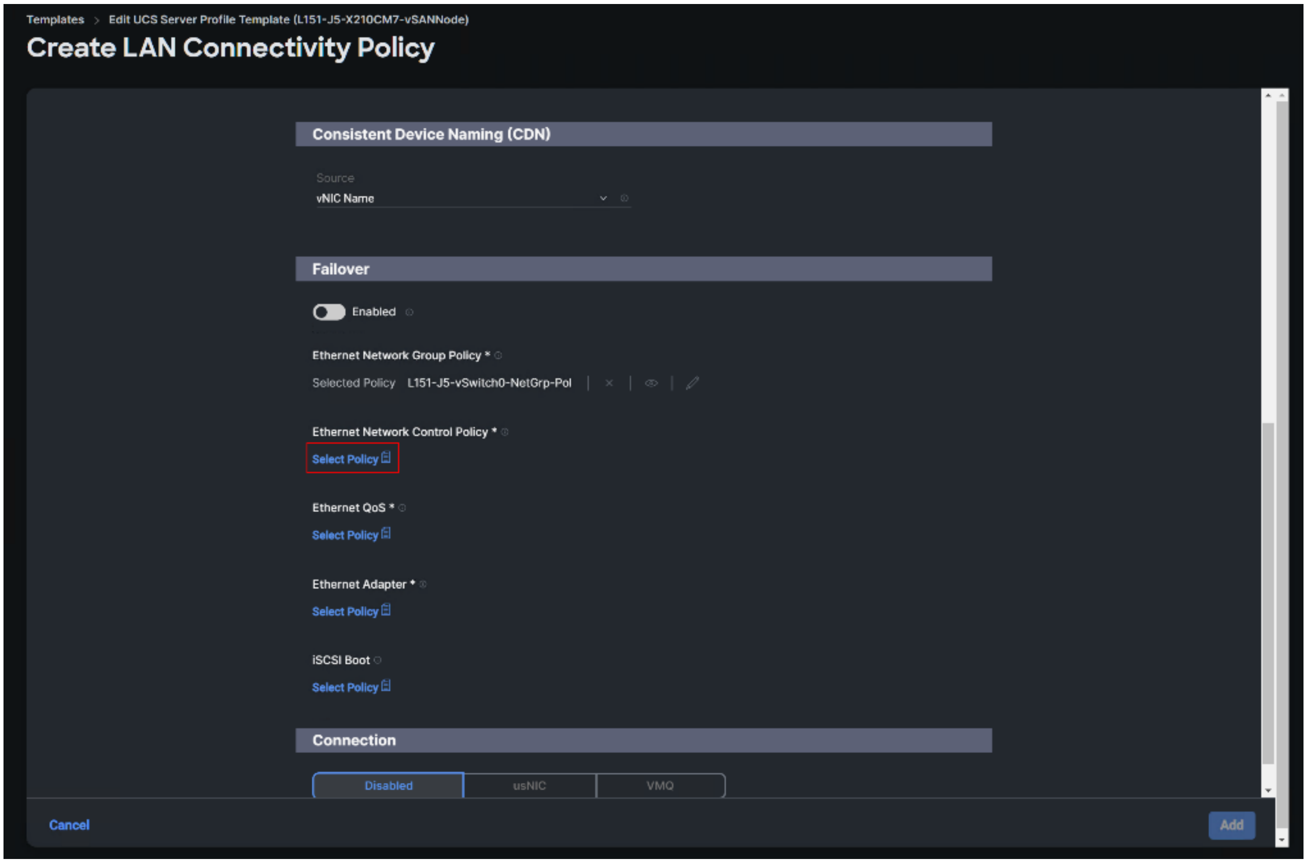

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled

Ethernet Network Group Policy *

Selected Policy L151-J5-vSwitch0-NetGrp-Pol

Ethernet Network Control Policy *

Select Policy

Ethernet QoS *

Select Policy

Ethernet Adapter *

Select Policy

iSCSI Boot

Select Policy

Connection

Step 70. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-NetCtrlPol).

Step 71. Click Next.

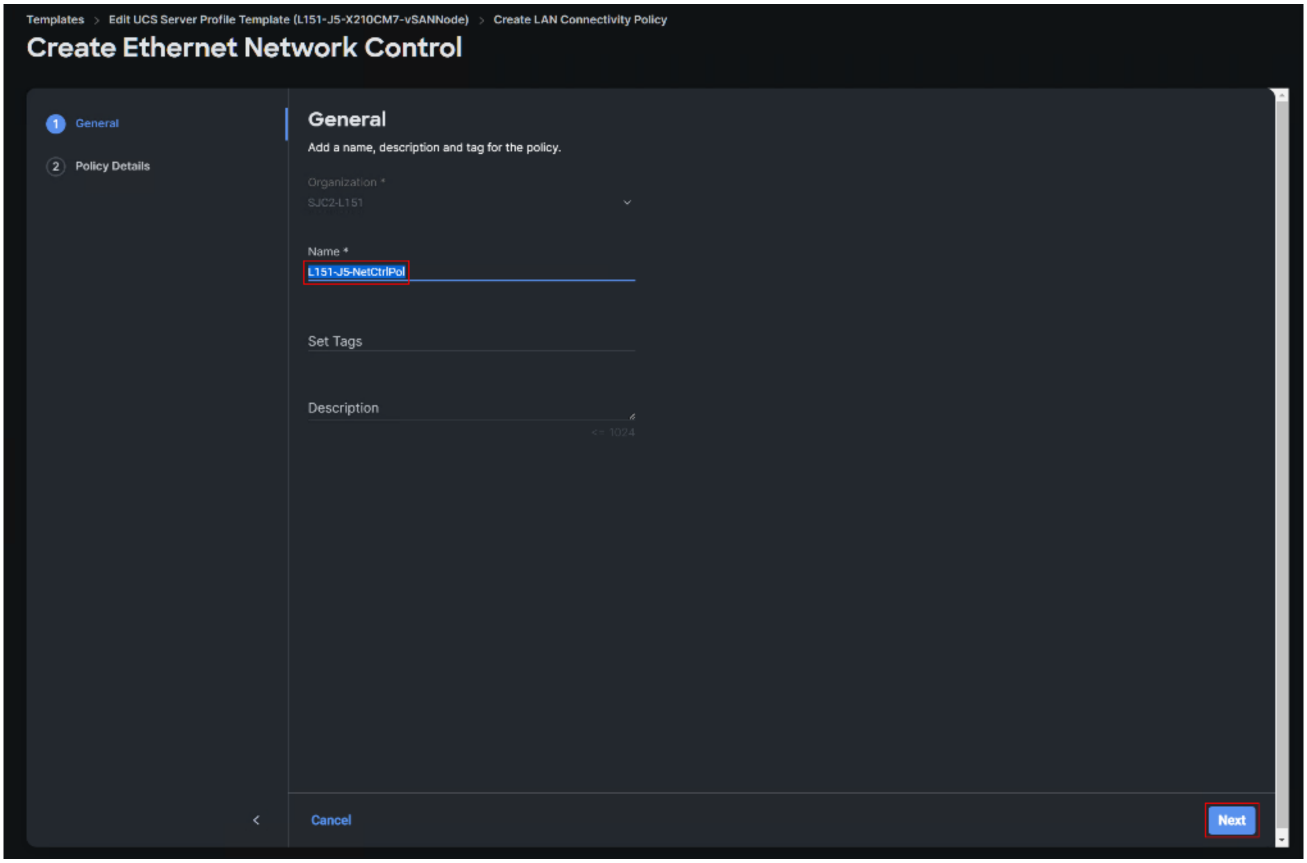

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet Network Control

General

Add a name, description and tag for the policy.

Organization *
SJC2-L151

Name *
L151-J5-NetCtrlPol

Step 72. Enable Cisco Discovery Protocol and both Enable Transmit and Enable Receive under LLDP. Click Create.

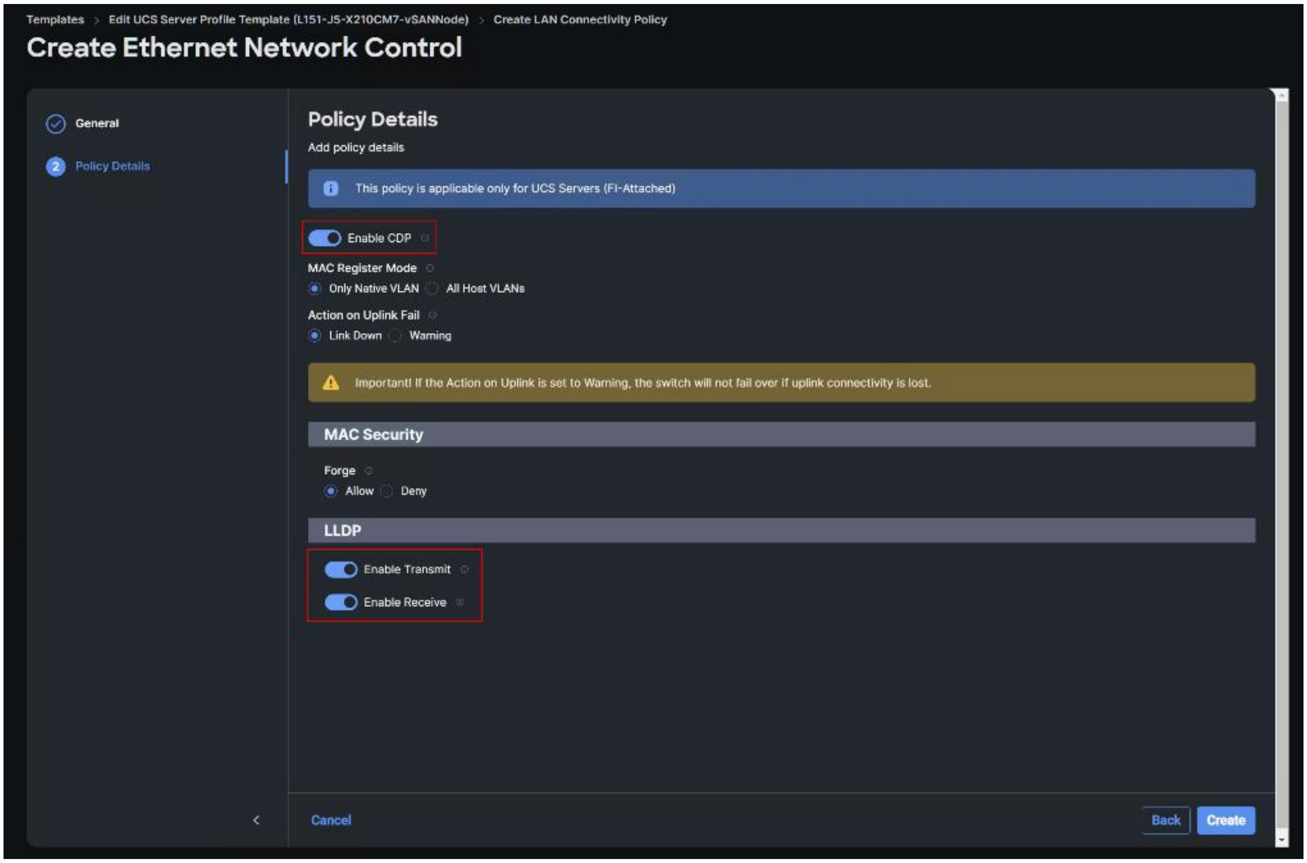

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet Network Control

Policy Details

Add policy details

This policy is applicable only for UCS Servers (FI-Attached)

Enable CDP: on

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: on

Enable Receive: on

Step 73. Click Select Policy under Ethernet QoS and click Create New.

Note: The 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.

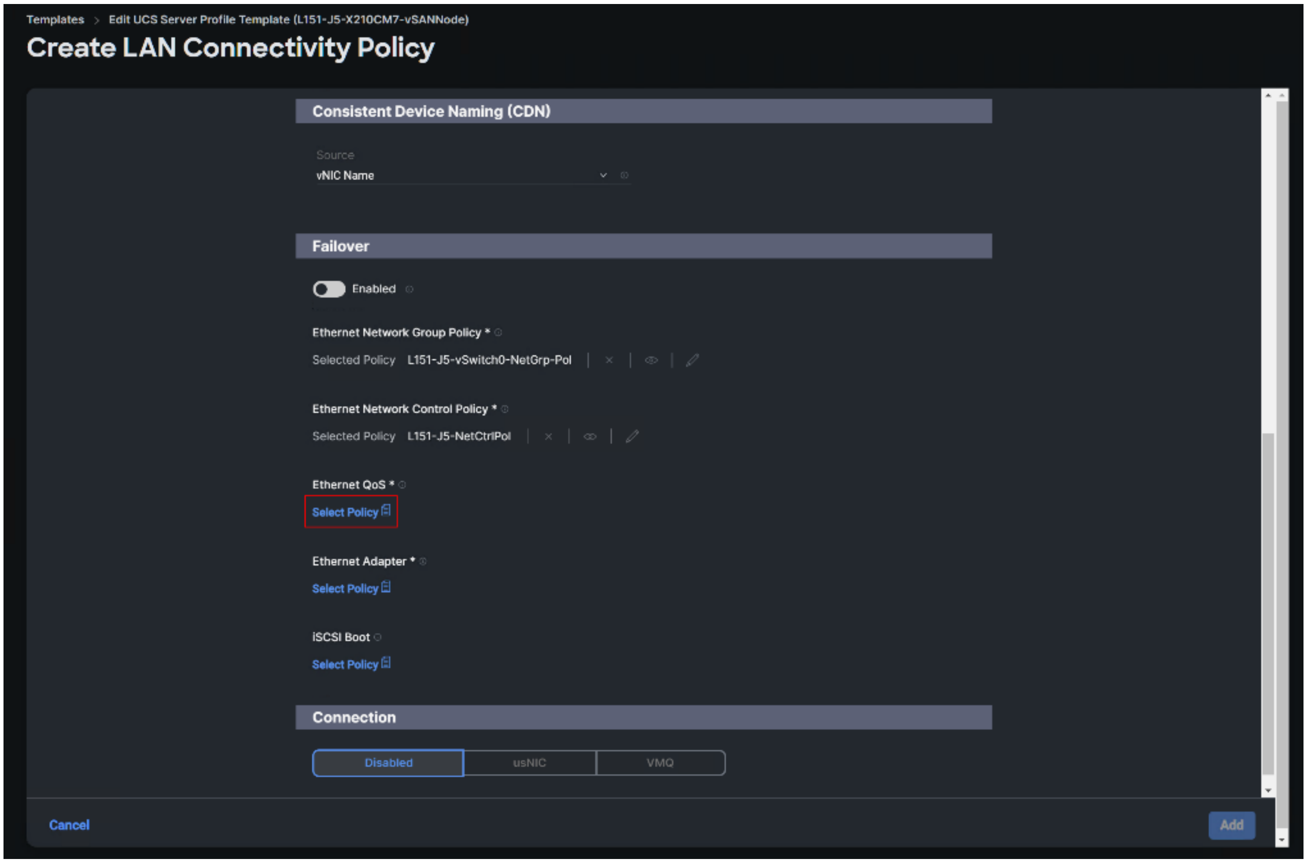

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled

Ethernet Network Group Policy *

Selected Policy L151-J5-vSwitch0-NetGrp-Pol

Ethernet Network Control Policy *

Selected Policy L151-J5-NetCtrlPol

Ethernet QoS *

Select Policy

Ethernet Adapter *

Select Policy

iSCSI Boot

Select Policy

Connection

Step 74. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-QOSPol).

Step 75. Click Next.

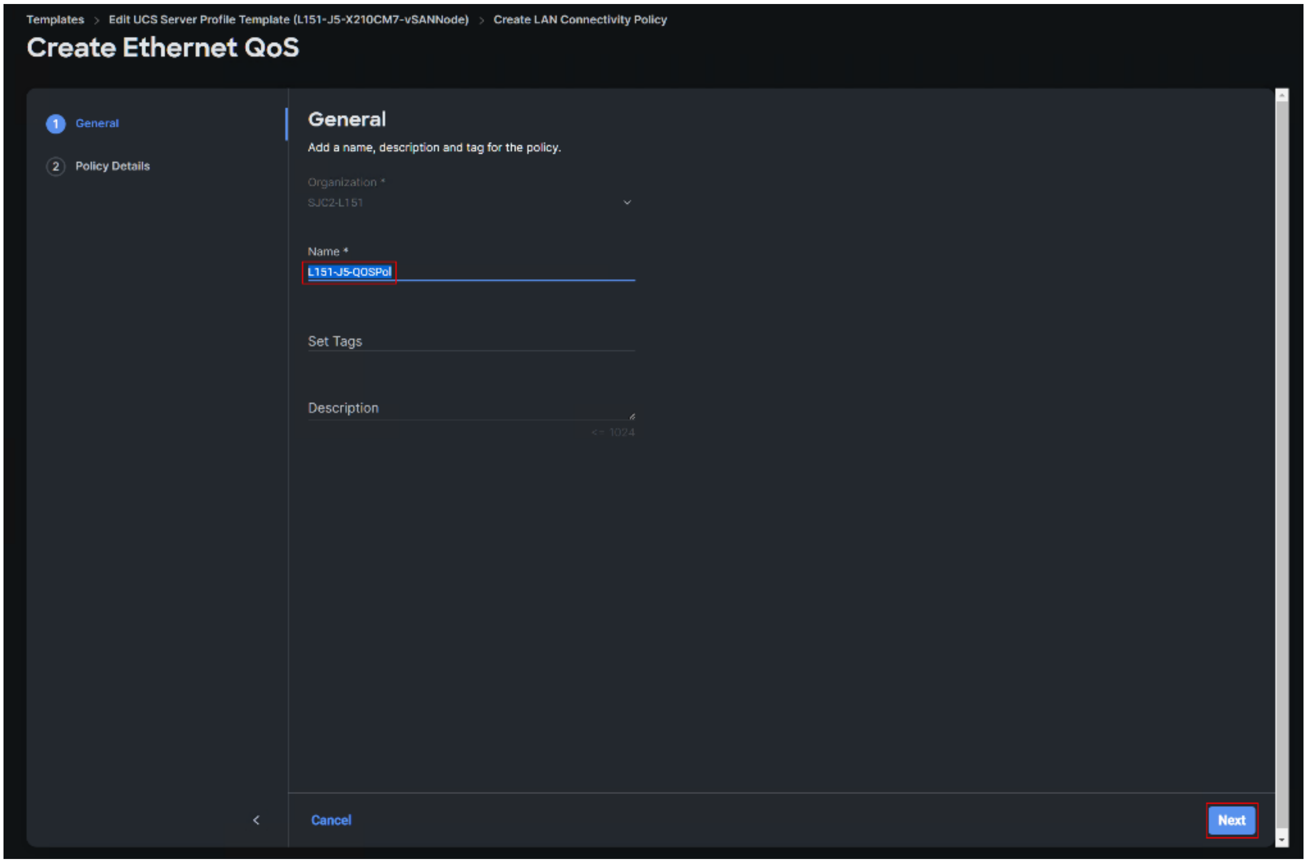

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy
Create Ethernet QoS
1 General
2 Policy Details
General
Add a name, description and tag for the policy.
Organization *
SJC2-L151
Name *
L151-J5-QOSPol
Set Tags
Description
<= 1024
Cancel
Next

Step 76. Change the MTU Bytes value to 9000. Click Create.

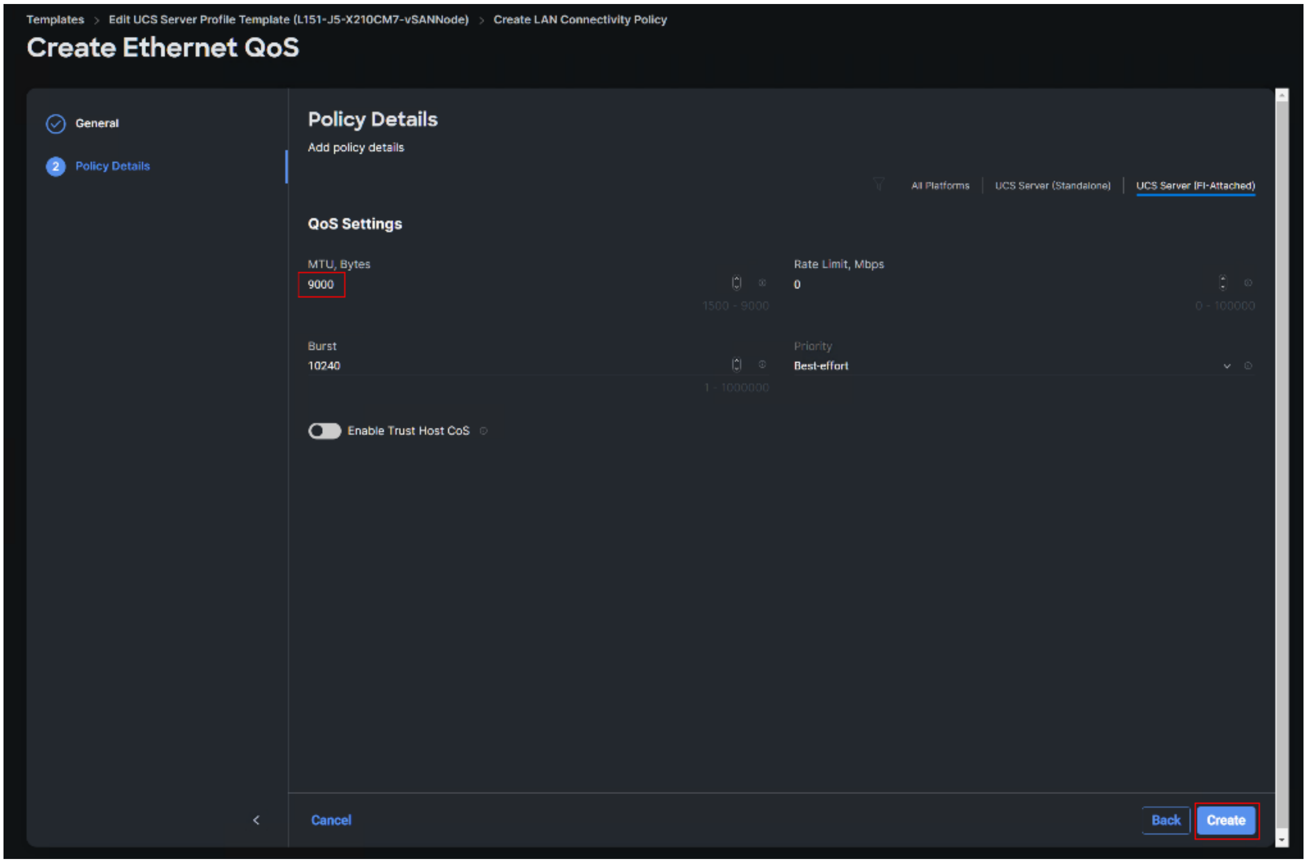

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet QoS

Policy Details

Add policy details

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 77. Click Select Policy under Ethernet Adapter and then click Create New.

Note: The 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, L151-J5-EthAdapt-VMware-HiTraffic, is created and attached to the VDS0-A and VDS0-B interfaces which handle vMotion.

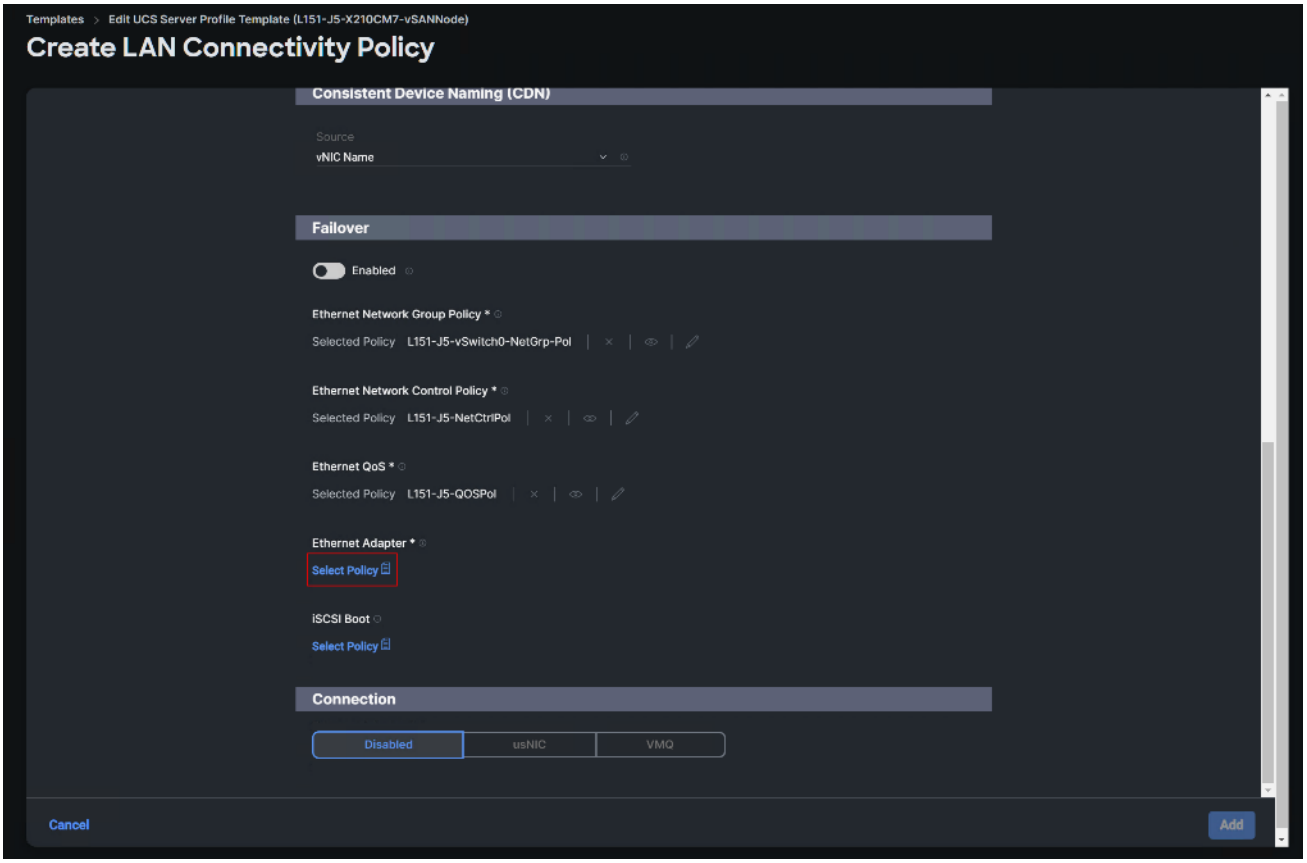

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled

Ethernet Network Group Policy *
Selected Policy L151-J5-vSwitch0-NetGrp-Pol

Ethernet Network Control Policy *
Selected Policy L151-J5-NetCtrlPol

Ethernet QoS *
Selected Policy L151-J5-QOSPol

Ethernet Adapter *
Select Policy

iSCSI Boot
Select Policy

Connection

Table 10. Ethernet Adapter Policy association to vNICs

Policy Name	vNICS
L151-J5-EthAdapt-VMwarePol	vSwitch0-A, vSwitch0-B
L151-J5-EthAdapt-VMware-HiTrafficPol	VDS0-A, VDS0-B,

Step 78. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, L151-J5-EthAdapt-VMware).

Step 79. Click Select Default Configuration under Ethernet Adapter Default Configuration.

Step 80. From the list, select VMware. Click Next.

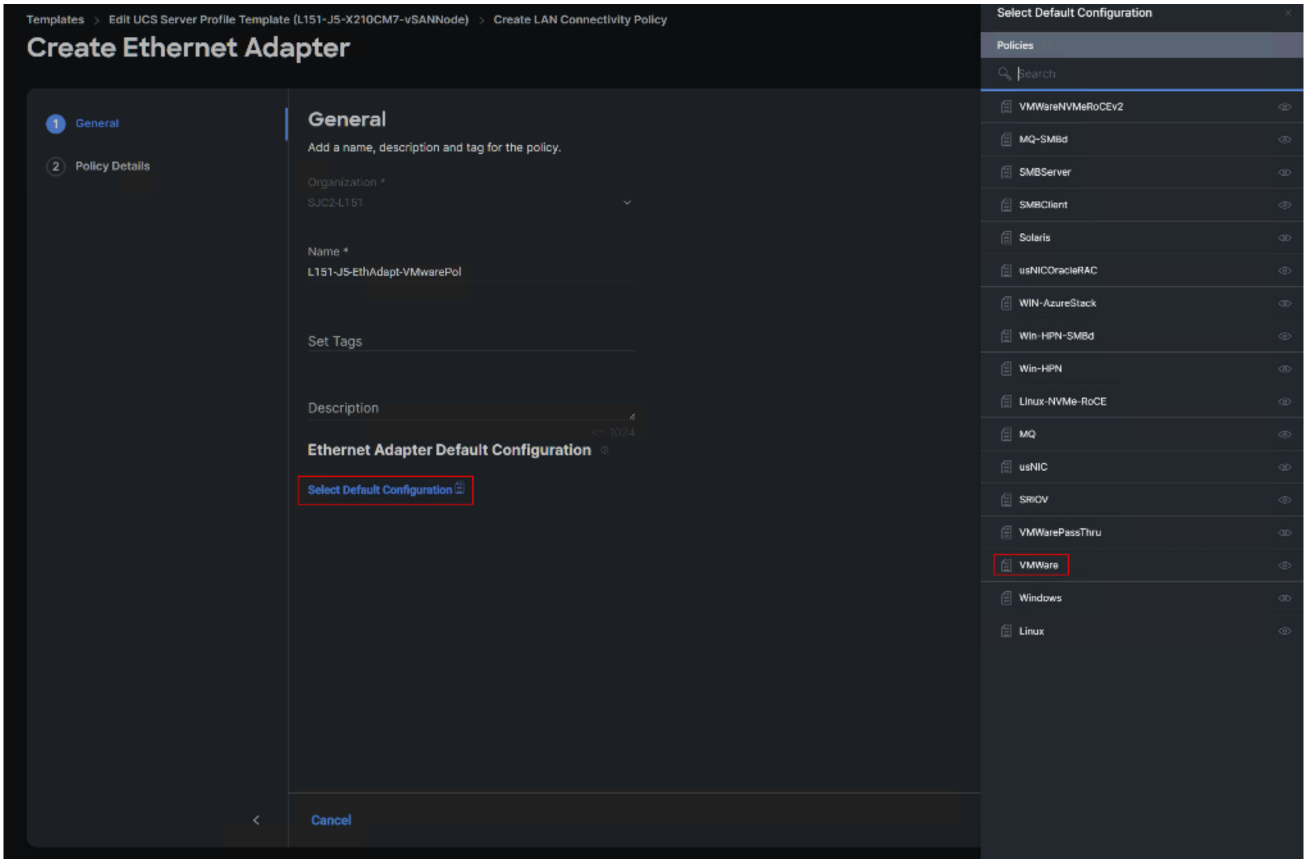

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet Adapter

1 General

2 Policy Details

General

Add a name, description and tag for the policy.

Organization *
SJC2-L151

Name *
L151-J5-EthAdapt-VMwarePol

Set Tags

Description

<= 1024

Ethernet Adapter Default Configuration

Select Default Configuration

Cancel

Select Default Configuration

Policies

Search

- VMWareNVMeRoCEv2
- MQ-SMBd
- SMBServer
- SMBClient
- Solaris
- usNICOracleRAC
- WIN-AzureStack
- Win-HPN-SMBd
- Win-HPN
- Linux-NVMe-RoCE
- MQ
- usNIC
- SRIOV
- VMWarePassThru
- VMWare
- Windows
- Linux

Step 81. For the L151-J5-EthAdapt-VMware policy, click Create and skip the rest of the steps in this section.

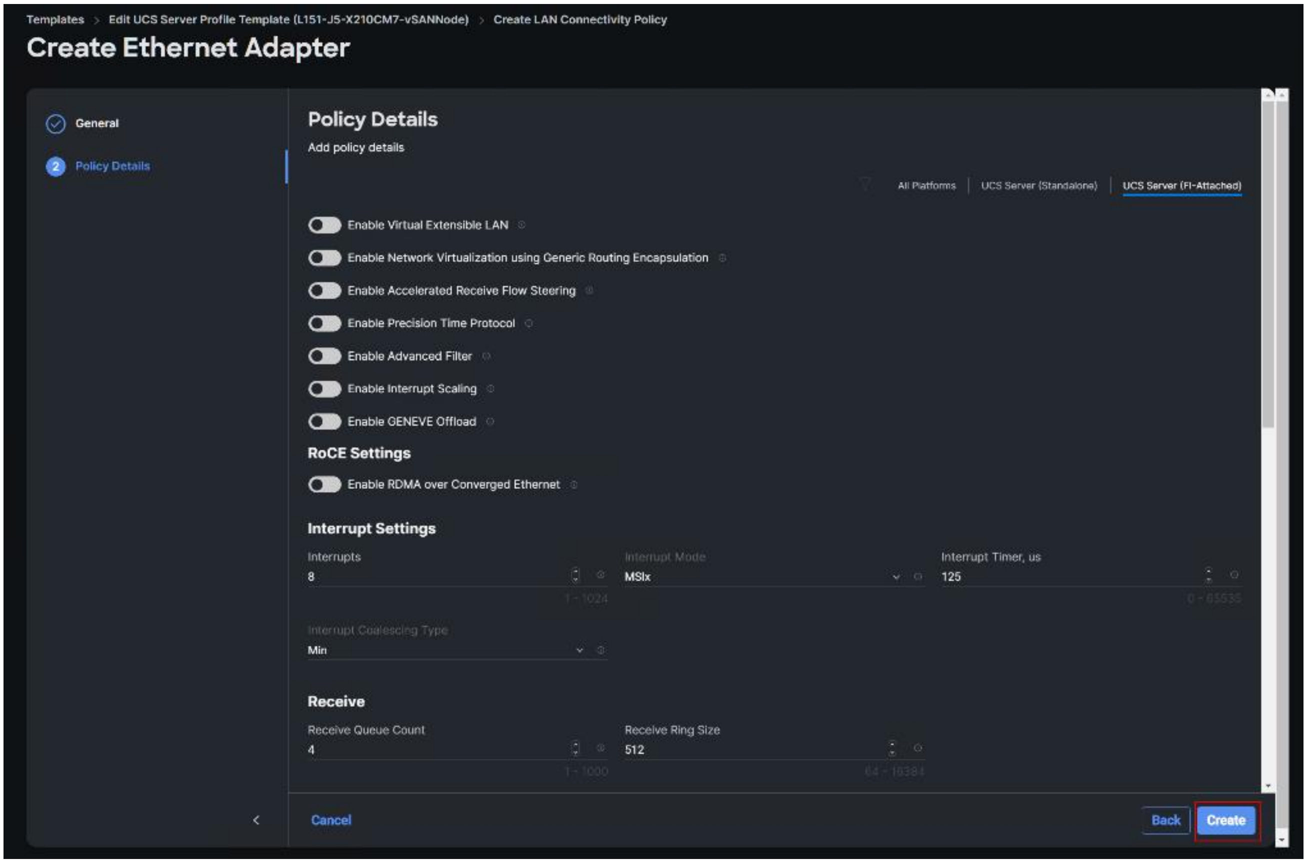

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode) > Create LAN Connectivity Policy

Create Ethernet Adapter

Policy Details

Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

- Enable Virtual Extensible LAN
- Enable Network Virtualization using Generic Routing Encapsulation
- Enable Accelerated Receive Flow Steering
- Enable Precision Time Protocol
- Enable Advanced Filter
- Enable Interrupt Scaling
- Enable GENEVE Offload

RoCE Settings

- Enable RDMA over Converged Ethernet

Interrupt Settings

Interrupts: 8 (1 - 1024)

Interrupt Mode: MSIx

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type: Min

Receive

Receive Queue Count: 4 (1 - 1000)

Receive Ring Size: 512 (64 - 16384)

Step 82. For the optional L151-J5-EthAdapt-VMware-HiTraffic policy used 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

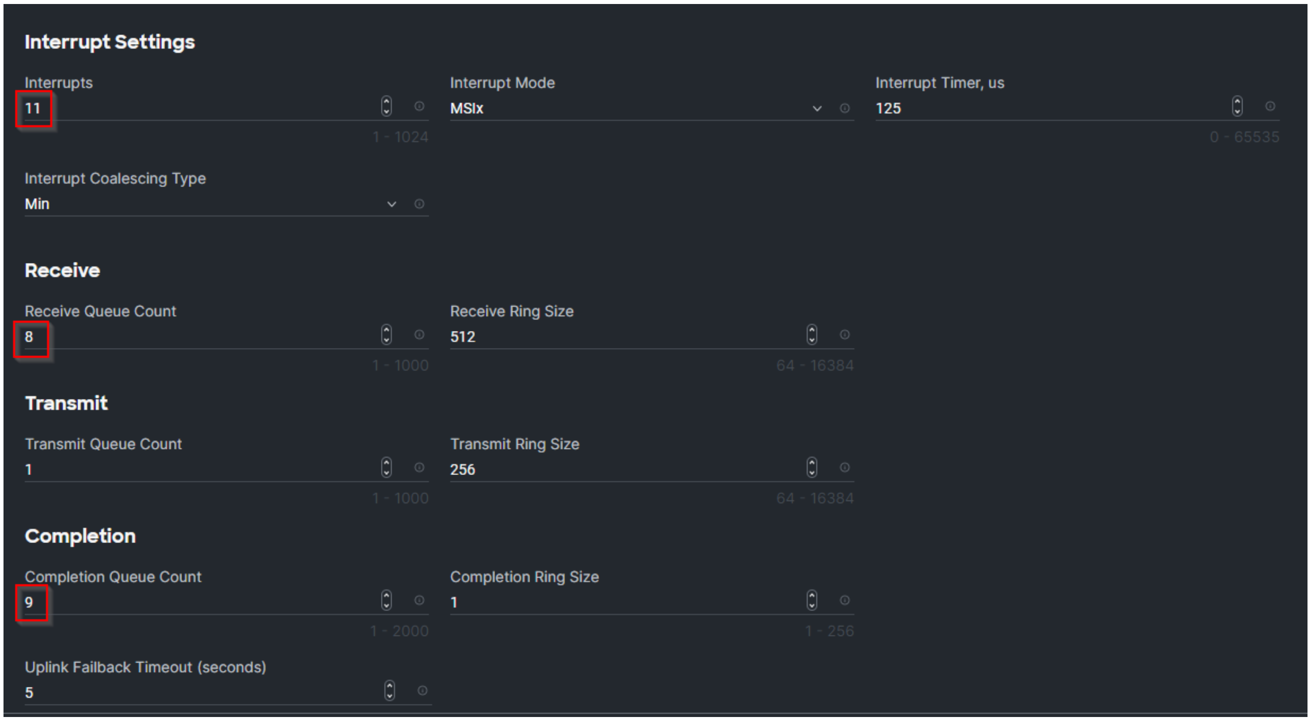

Interrupt Settings

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: 512 (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

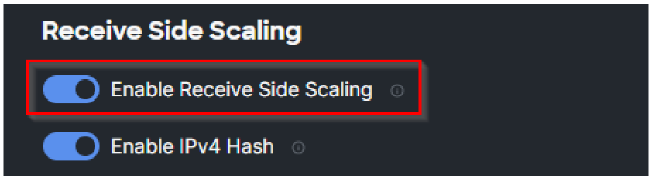

Receive Side Scaling

Enable Receive Side Scaling

Enable IPv4 Hash

Step 83. Click Create.

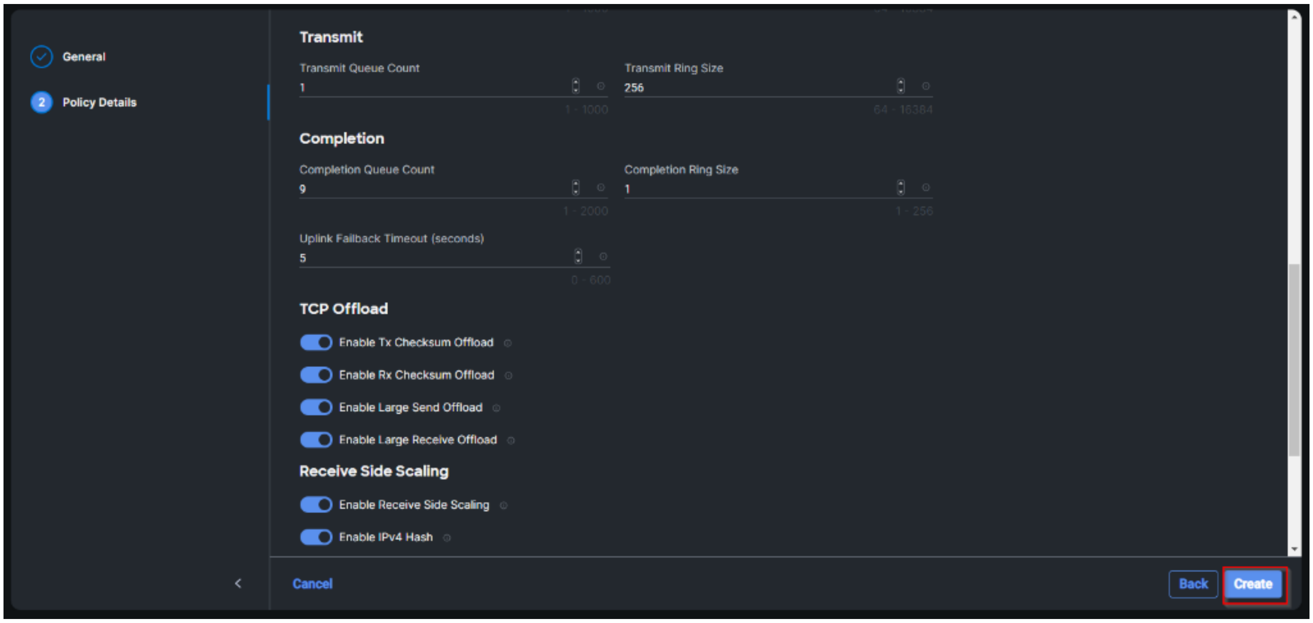

Transmit

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

Step 84. Click Add to finish creating the vNIC.

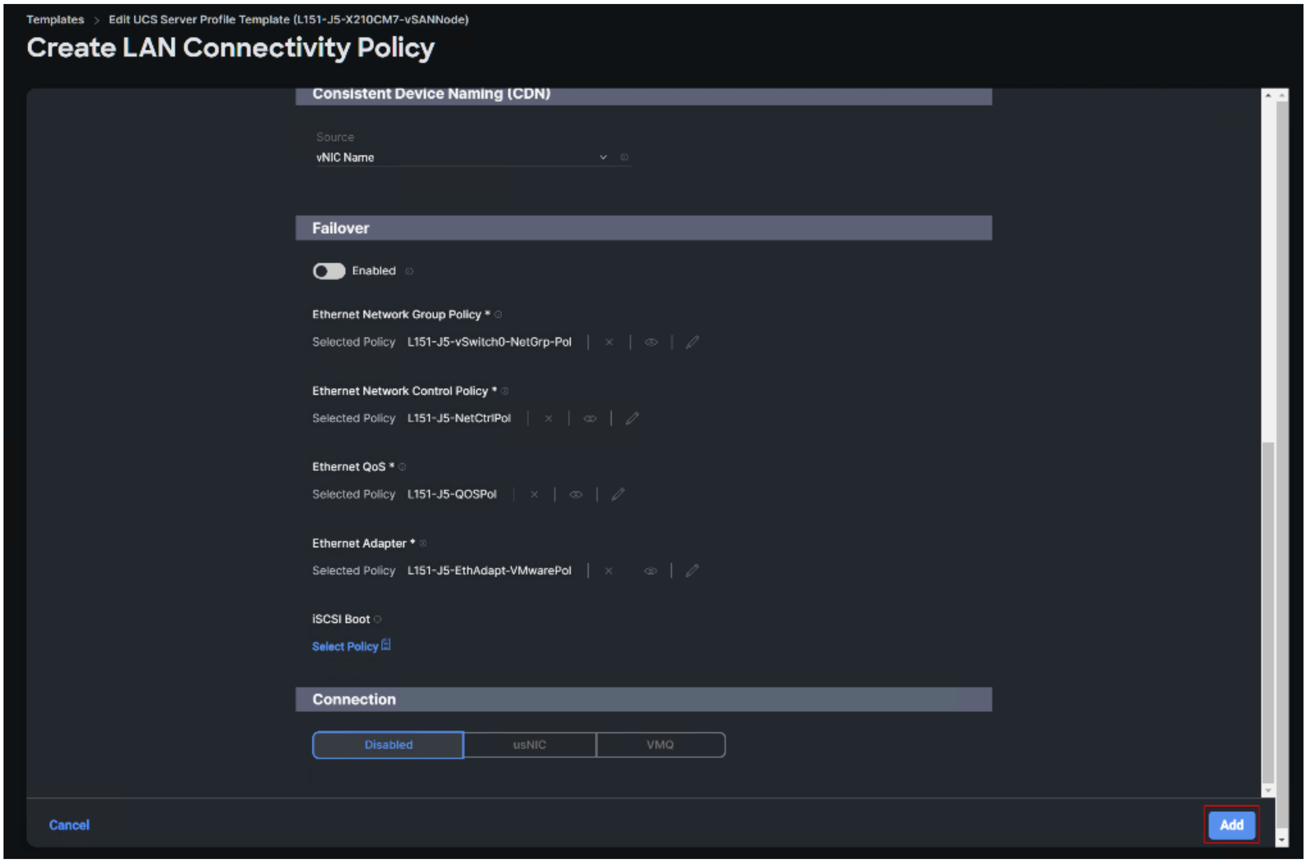

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Consistent Device Naming (CDN)

Source: vNIC Name

Failover

Enabled

Ethernet Network Group Policy *

Selected Policy L151-J5-vSwitch0-NetGrp-Pol

Ethernet Network Control Policy *

Selected Policy L151-J5-NetCtrlPol

Ethernet QoS *

Selected Policy L151-J5-QOSPol

Ethernet Adapter *

Selected Policy L151-J5-EthAdapt-VMwarePol

iSCSI Boot

Select Policy

Connection

Step 85. Repeat the vNIC creation steps for the rest of vNICs. Verify all four vNICs were successfully created. Click Create.

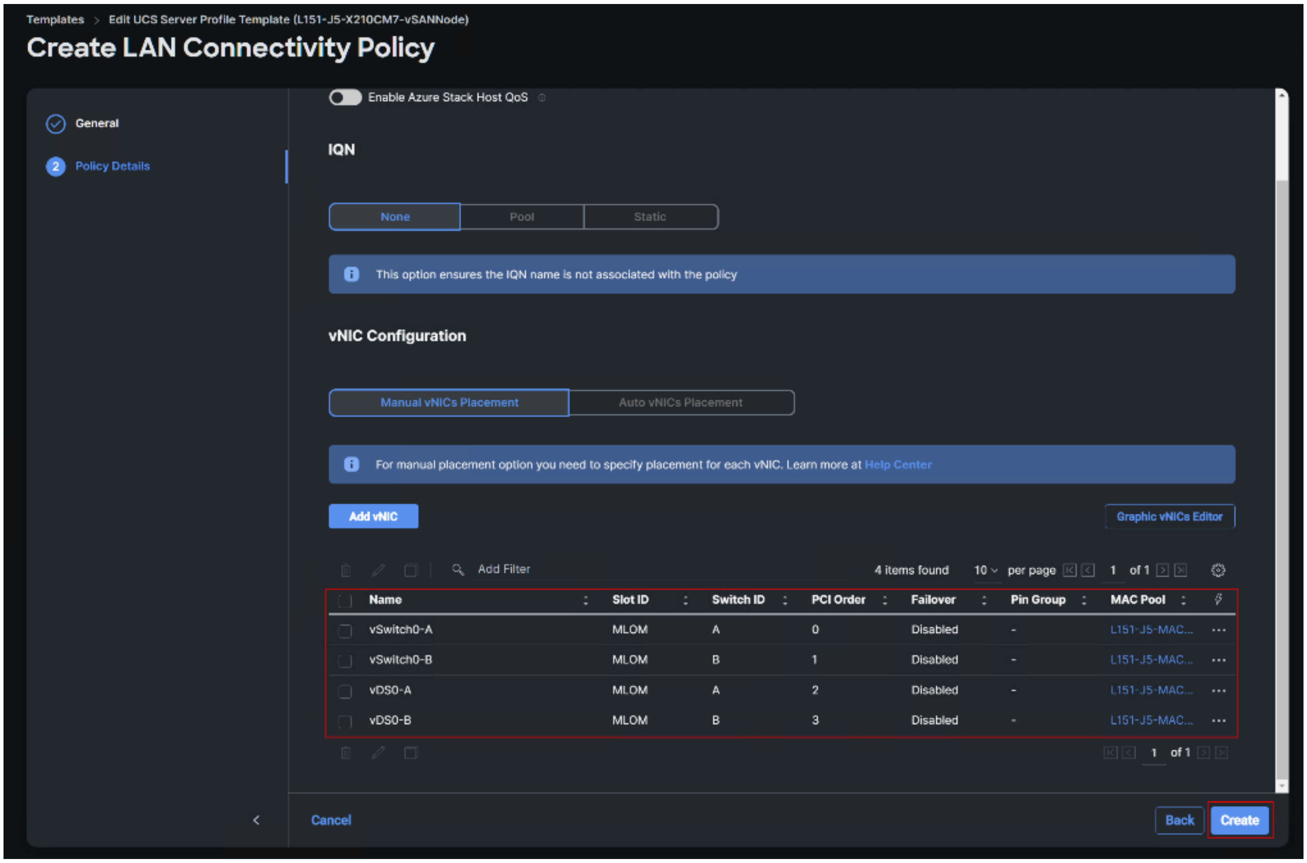

Templates > Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Create LAN Connectivity Policy

Enable Azure Stack Host QoS

IQN

This option ensures the IQN name is not associated with the policy

vNIC Configuration

For manual placement option you need to specify placement for each vNIC. Learn more at Help Center

4 items found 10 per page 1 of 1

Name	Slot ID	Switch ID	PCI Order	Failover	Pin Group	MAC Pool
vSwitch0-A	MLOM	A	0	Disabled	-	L151-J5-MAC...
vSwitch0-B	MLOM	B	1	Disabled	-	L151-J5-MAC...
vDS0-A	MLOM	A	2	Disabled	-	L151-J5-MAC...
vDS0-B	MLOM	B	3	Disabled	-	L151-J5-MAC...

Step 86. When the LAN connectivity policy and SAN connectivity policy are created and assigned, click Next to move to the Summary screen.

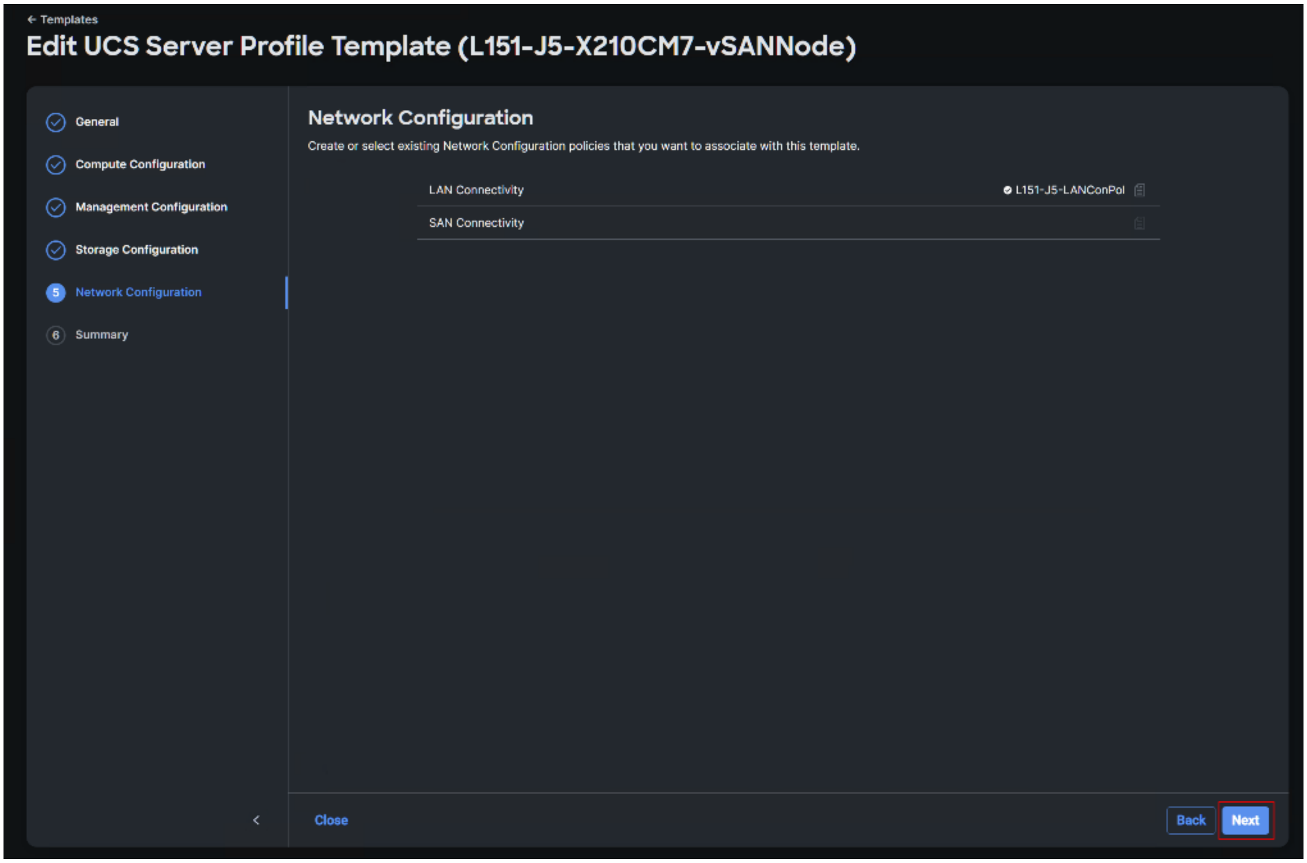
← Templates

Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Network Configuration

Create or select existing Network Configuration policies that you want to associate with this template.

LAN Connectivity: L151-J5-LANConPol

SAN Connectivity

Step 87. 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.

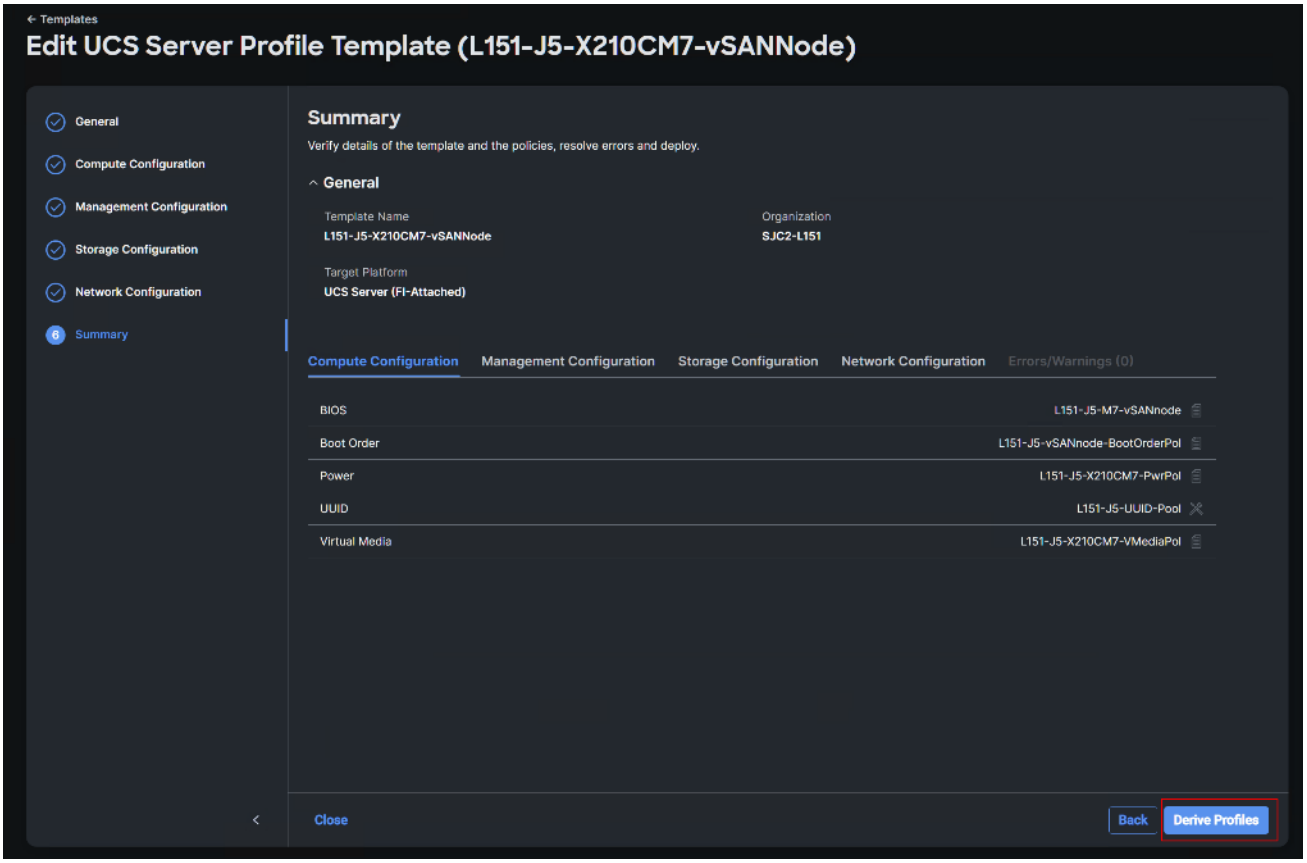

← Templates

Edit UCS Server Profile Template (L151-J5-X210CM7-vSANNode)

Summary

Verify details of the template and the policies, resolve errors and deploy.

General

Template Name: L151-J5-X210CM7-vSANNode

Organization: SJC2-L151

Target Platform: UCS Server (FI-Attached)

BIOS	L151-J5-M7-vSANnode
Boot Order	L151-J5-vSANnode-BootOrderPol
Power	L151-J5-X210CM7-PwrPol
UUID	L151-J5-UUID-Pool
Virtual Media	L151-J5-X210CM7-VMediaPol

Step 88. Under the Server Assignment, select Assign Now and select Cisco UCS X210c M7 Nodes. You can select one or more servers depending on the number of profiles to be deployed. Click Next.

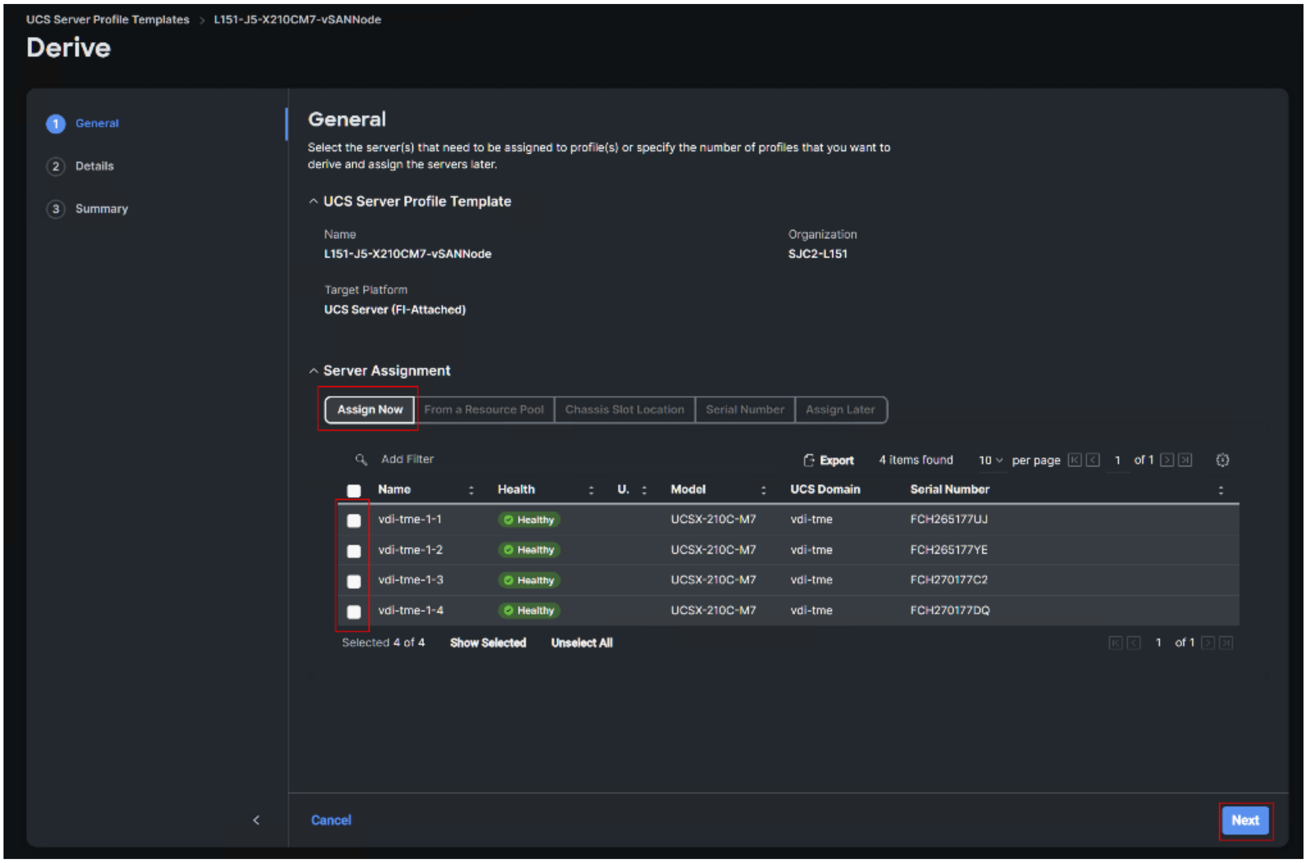

UCS Server Profile Templates > L151-J5-X210CM7-vSANNode

Derive

General

Select the server(s) that need to be assigned to profile(s) or specify the number of profiles that you want to derive and assign the servers later.

UCS Server Profile Template

Name: L151-J5-X210CM7-vSANNode

Organization: SJC2-L151

Target Platform: UCS Server (FI-Attached)

Server Assignment

☐	Name	Health	U.	Model	UCS Domain	Serial Number
☐	vdi-tme-1-1	Healthy		UCSX-210C-M7	vdi-tme	FCH265177UJ
☐	vdi-tme-1-2	Healthy		UCSX-210C-M7	vdi-tme	FCH265177YE
☐	vdi-tme-1-3	Healthy		UCSX-210C-M7	vdi-tme	FCH270177C2
☐	vdi-tme-1-4	Healthy		UCSX-210C-M7	vdi-tme	FCH270177DQ

Selected 4 of 4

Cisco Intersight will fill the default information for the number of servers selected.

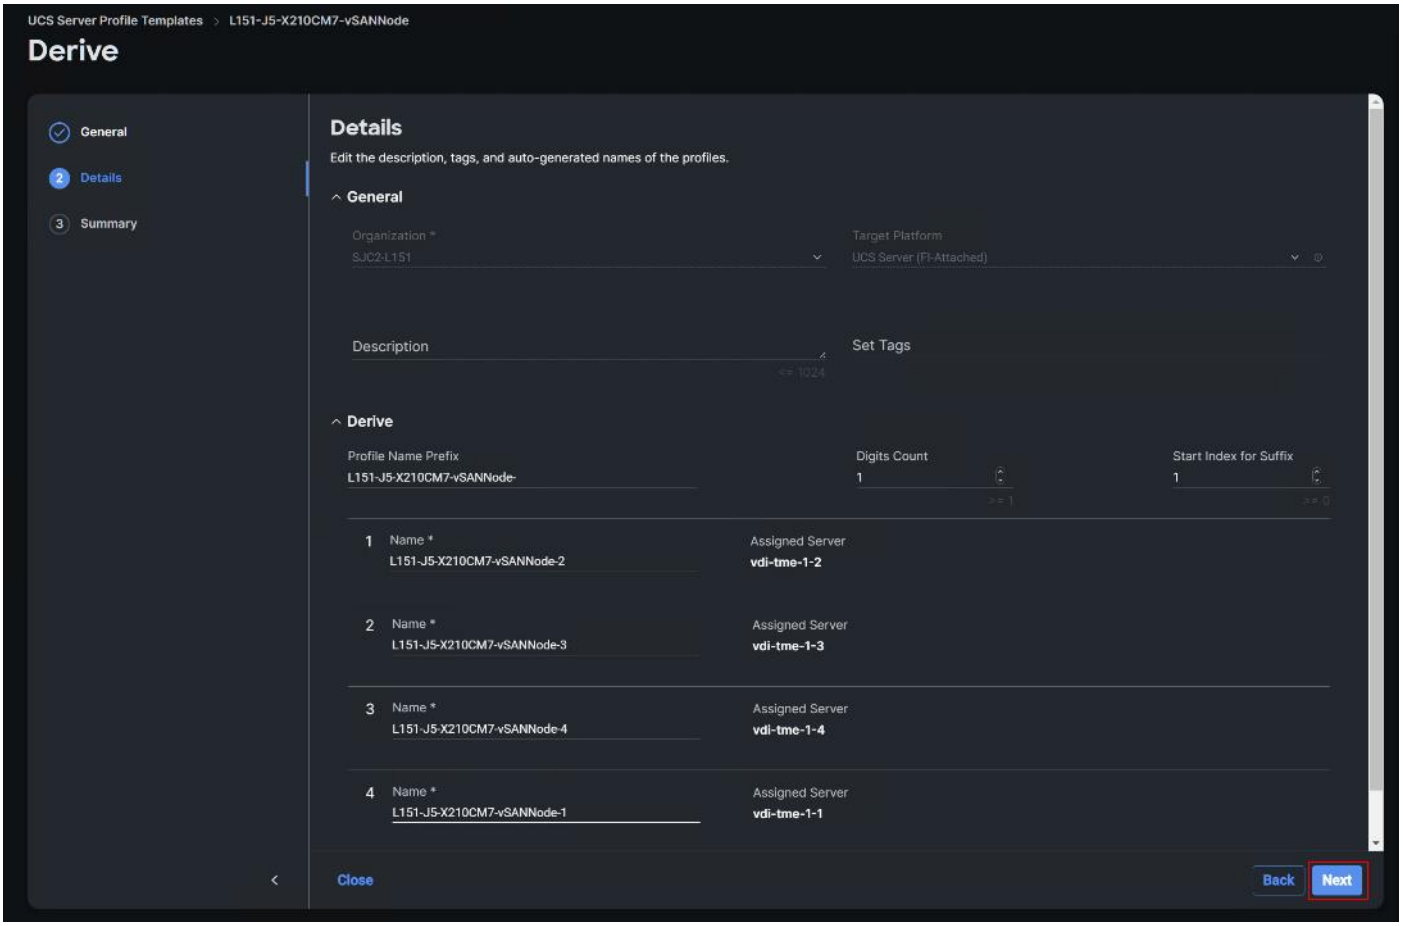

Details

Edit the description, tags, and auto-generated names of the profiles.

General

Organization *: SJC2-L151

Target Platform: UCS Server (FI-Attached)

Description

<= 1024

Set Tags

Derive

Profile Name Prefix: L151-J5-X210CM7-vSANNode-

Digits Count: 1 (>= 1)

Start Index for Suffix: 1 (>= 0)

	Name *	Assigned Server
1	L151-J5-X210CM7-vSANNode-2	vdi-tme-1-2
2	L151-J5-X210CM7-vSANNode-3	vdi-tme-1-3
3	L151-J5-X210CM7-vSANNode-4	vdi-tme-1-4
4	L151-J5-X210CM7-vSANNode-1	vdi-tme-1-1

Step 89. Adjust the Prefix and number as needed. Click Next.

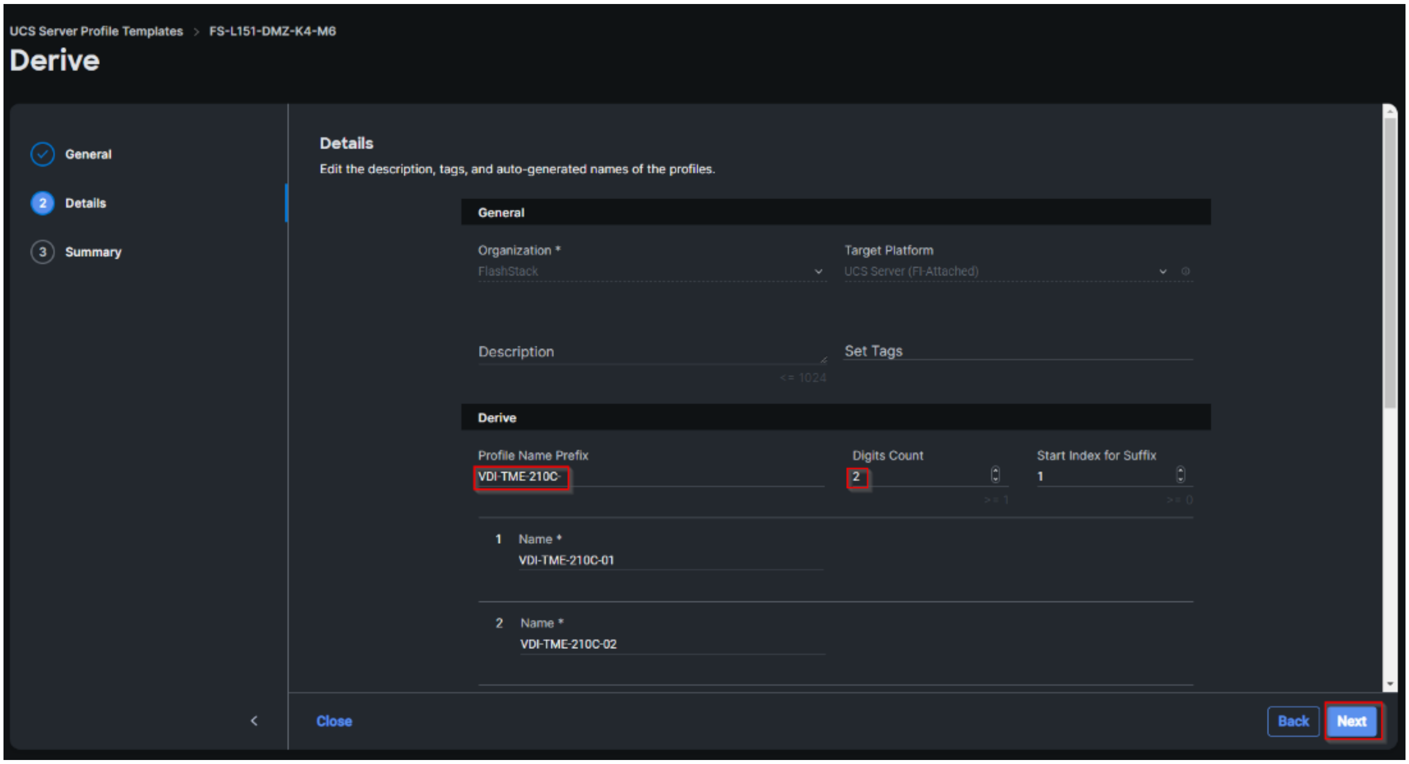

UCS Server Profile Templates > FS-L151-DMZ-K4-M6

Derive

General

2 Details

3 Summary

Details

Edit the description, tags, and auto-generated names of the profiles.

General

Organization * FlashStack

Target Platform UCS Server (FI-Attached)

Description <= 1024

Set Tags

Derive

Profile Name Prefix VDI-TME-210C-

Digits Count 2

Start Index for Suffix 1

1 Name * VDI-TME-210C-01

2 Name * VDI-TME-210C-02

Step 90. Verify the information and click Derive to create the Server Profiles.

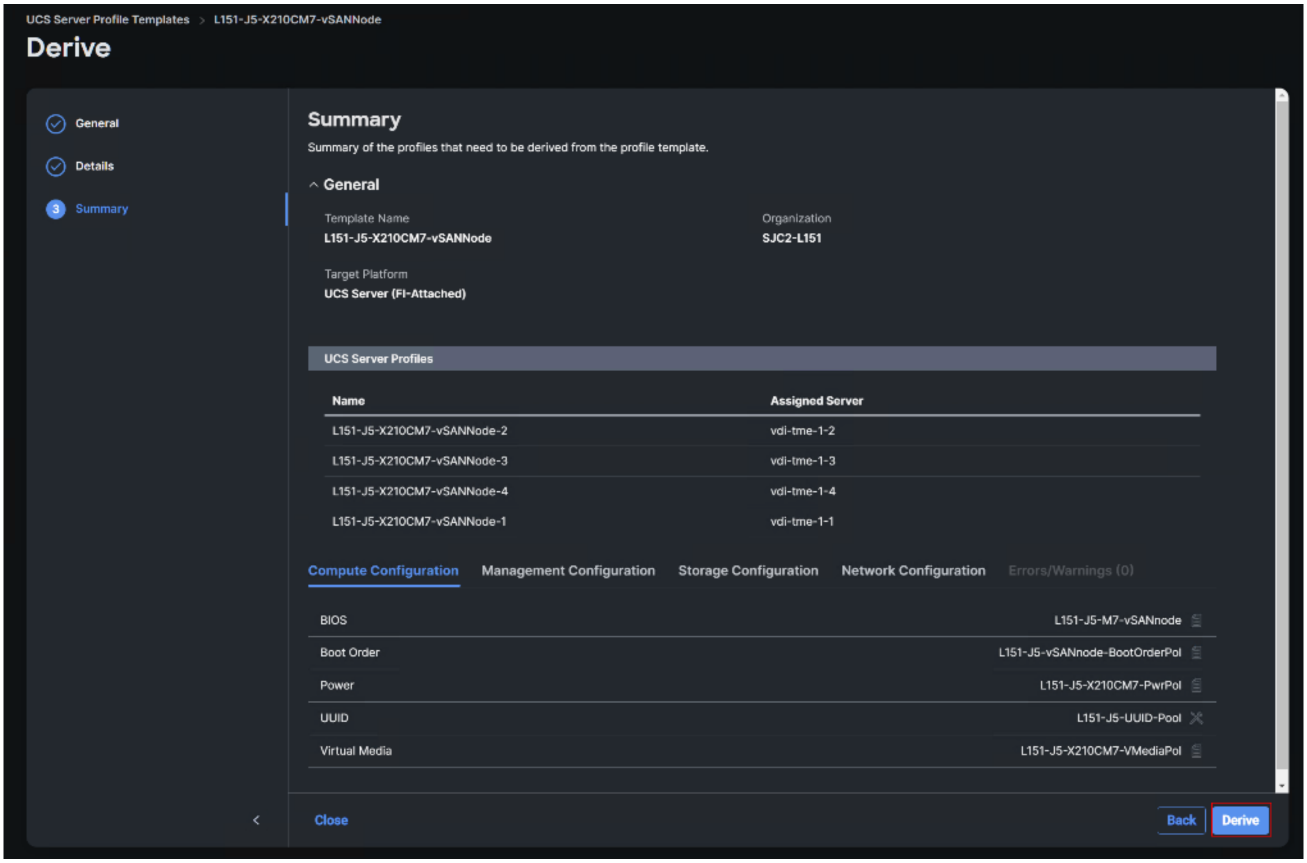

UCS Server Profile Templates > L151-J5-X210CM7-vSANNode

Derive

General

Details

3 Summary

Summary

Summary of the profiles that need to be derived from the profile template.

General

Template Name L151-J5-X210CM7-vSANNode

Organization SJC2-L151

Target Platform UCS Server (FI-Attached)

UCS Server Profiles

Name	Assigned Server
L151-J5-X210CM7-vSANNode-2	vdi-tme-1-2
L151-J5-X210CM7-vSANNode-3	vdi-tme-1-3
L151-J5-X210CM7-vSANNode-4	vdi-tme-1-4
L151-J5-X210CM7-vSANNode-1	vdi-tme-1-1

Compute Configuration | Management Configuration | Storage Configuration | Network Configuration | Errors/Warnings (0)

BIOS	L151-J5-M7-vSANnode
Boot Order	L151-J5-vSANnode-BootOrderPol
Power	L151-J5-X210CM7-PwrPol
UUID	L151-J5-UUID-Pool
Virtual Media	L151-J5-X210CM7-VMediaPol

Deploy Server Profile

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

Procedure 1. Deploy Server Profile

Step 1. Log into the Cisco Intersight portal as a user with account administrator role.

Step 2. Navigate to Configure > Templates and click Create UCS Server Profile Template

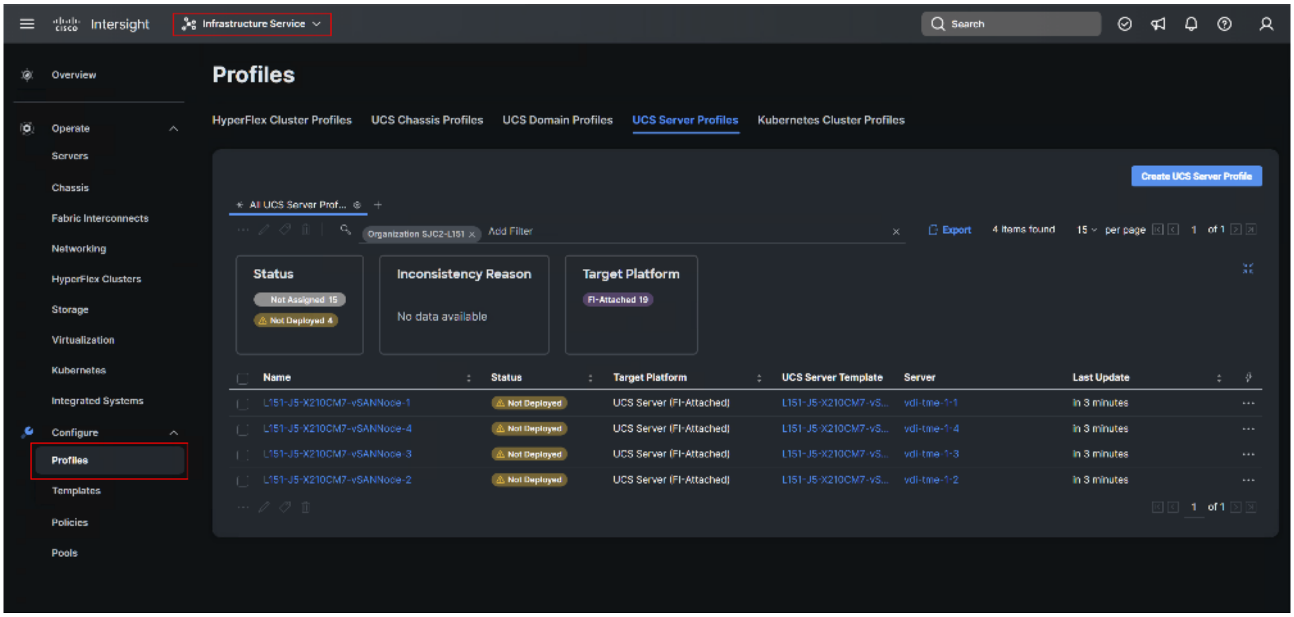

Profiles

Name	Status	Target Platform	UCS Server Template	Server	Last Update
L151-J5-X210CM7-vSANNode-1	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-1	in 3 minutes
L151-J5-X210CM7-vSANNode-4	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-4	in 3 minutes
L151-J5-X210CM7-vSANNode-3	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-3	in 3 minutes
L151-J5-X210CM7-vSANNode-2	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-2	in 3 minutes

Step 3. Click on Action section next to profile to be deployed. Click Deploy.

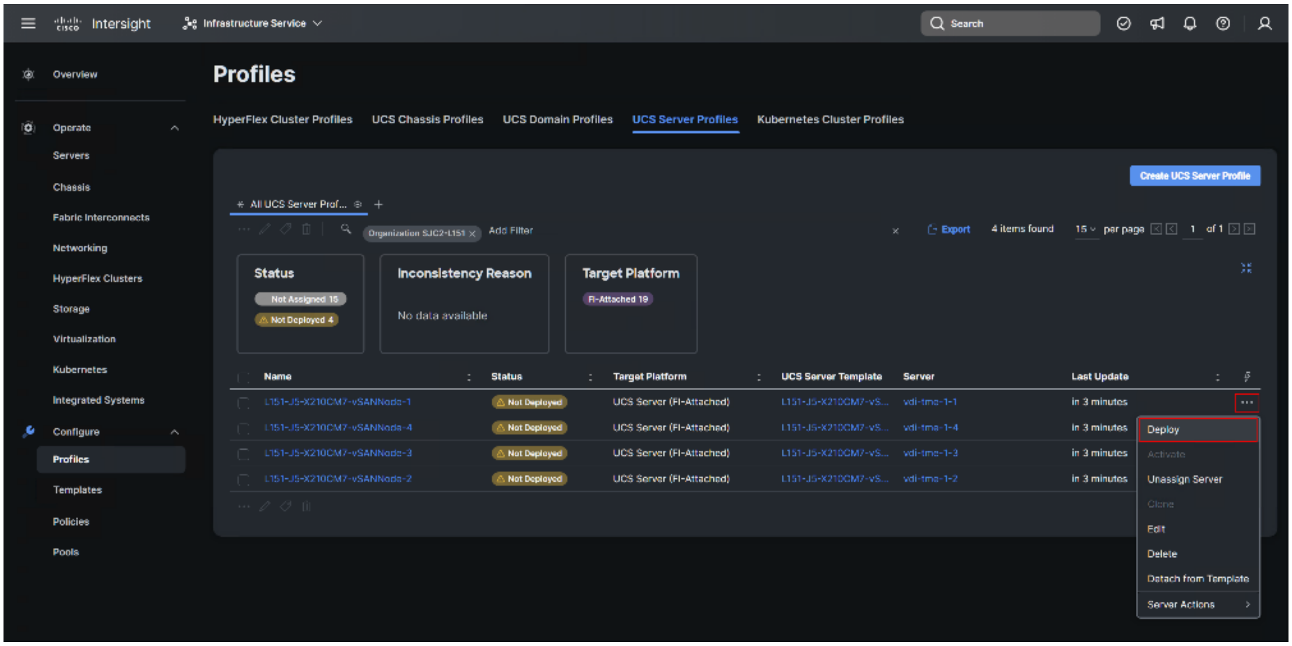

Profiles

Name	Status	Target Platform	UCS Server Template	Server	Last Update
L151-J5-X210CM7-vSANNode-1	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-1	in 3 minutes
L151-J5-X210CM7-vSANNode-4	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-4	in 3 minutes
L151-J5-X210CM7-vSANNode-3	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-3	in 3 minutes
L151-J5-X210CM7-vSANNode-2	⚠ Not Deployed	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-2	in 3 minutes

Step 4. On Acknowledgement screen select Reboot Immediately to Activate. Click Deploy.

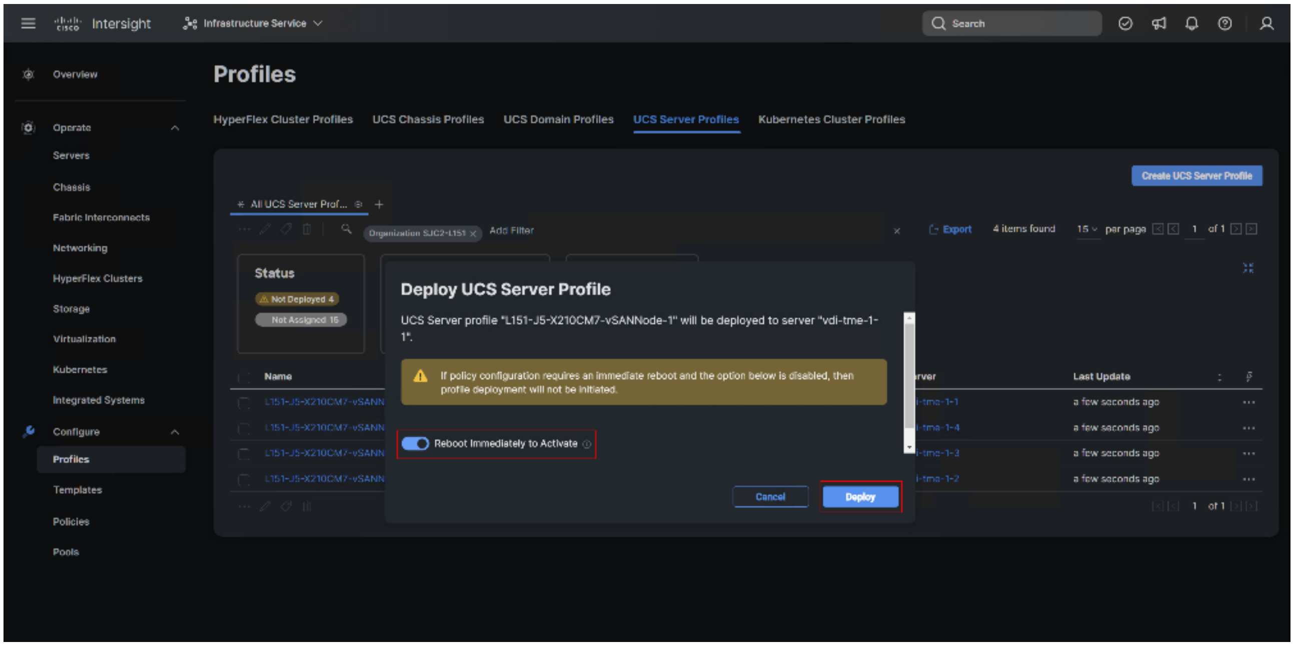

Deploy UCS Server Profile

UCS Server profile "L151-J5-X210CM7-vSANNode-1" will be deployed to server "vdi-tme-1-1".

If policy configuration requires an immediate reboot and the option below is disabled, then profile deployment will not be initiated.

Reboot Immediately to Activate

Name	Server	Last Update
L151-J5-X210CM7-vSANN	vdi-tme-1-1	a few seconds ago
L151-J5-X210CM7-vSANN	vdi-tme-1-4	a few seconds ago
L151-J5-X210CM7-vSANN	vdi-tme-1-3	a few seconds ago
L151-J5-X210CM7-vSANN	vdi-tme-1-2	a few seconds ago

Step 5. Verify deployment status.

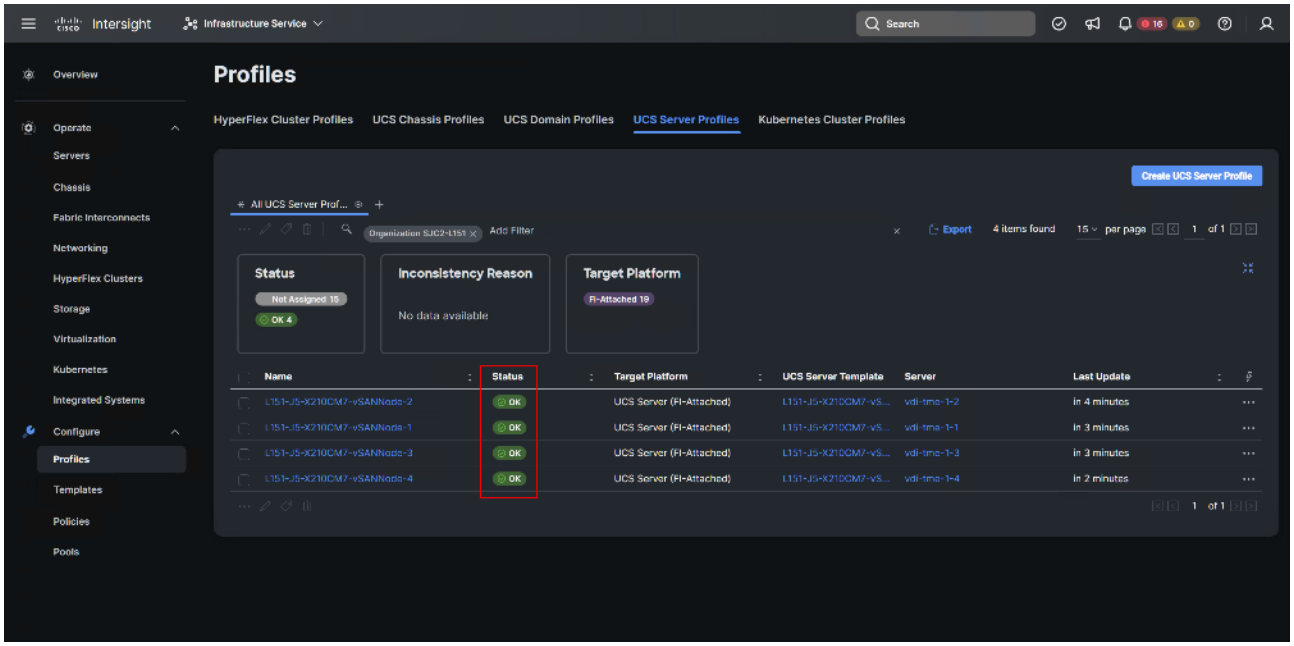

Name	Status	Target Platform	UCS Server Template	Server	Last Update
L151-J5-X210CM7-vSANNode-2	OK	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-2	in 4 minutes
L151-J5-X210CM7-vSANNode-1	OK	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-1	in 3 minutes
L151-J5-X210CM7-vSANNode-3	OK	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-3	in 3 minutes
L151-J5-X210CM7-vSANNode-4	OK	UCS Server (FI-Attached)	L151-J5-X210CM7-vS...	vdi-tme-1-4	in 2 minutes

Configure Cisco Nexus 93180YC-FX Switches

This section details the steps for the Cisco Nexus 93180YC-FX switch configuration.

Procedure 1. Configure Global Settings for Cisco Nexus A and Cisco Nexus B

Step 1. Log in as admin user into the Cisco Nexus Switch A and run the following commands to set global configurations and jumbo frames in QoS:

```
conf terminal

policy-map type network-qos jumbo
```

```
class type network-qos class-default
mtu 9216
exit
class type network-qos class-fcoe
pause no-drop
mtu 2158
exit
exit
system qos
service-policy type network-qos jumbo
exit
copy running-config startup-config
```

Step 2. Log in as admin user into the Cisco Nexus Switch B and run the same commands (above) to set global configurations and jumbo frames in QoS.

Procedure 2. Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

Note: For this solution, we created VLAN 70, 71, 72, 73, 74 and 76.

Step 1. Log in as admin user into the Cisco Nexus Switch A.

Step 2. Create VLAN 70:

```
config terminal
VLAN 70
name InBand-Mgmt
no shutdown
exit
copy running-config startup-config
```

Step 3. Log in as admin user into the Nexus Switch B and create VLANs.

Virtual Port Channel (vPC) Summary for Data and Storage Network

In the Cisco Nexus 93180YC-FX switch topology, a single vPC feature is enabled to provide HA, faster convergence in the event of a failure, and greater throughput. Cisco Nexus 93180YC-FX vPC configurations with the vPC domains and corresponding vPC names and IDs for vSAN ESXi Servers is listed in Table 11.

Table 11. vPC Summary

vPC Domain	vPC Name	vPC ID
50	Peer-Link	1
50	vPC Port-Channel to FI-A	11
50	vPC Port-Channel to FI-B	12

As listed in Table 11, a single vPC domain with Domain ID 50 is created across two Cisco Nexus 93180YC-FX member switches to define vPC members to carry specific VLAN network traffic. In this topology, a total number of three vPCs were defined:

- vPC ID 1 is defined as Peer link communication between two Nexus switches in Fabric A and B.
- vPC IDs 11 and 12 are defined for traffic from Cisco UCS fabric interconnects.

Cisco Nexus 93180YC-FX Switch Cabling Details

The following tables list the cabling information.

Table 12. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch A	Eth1/51	40Gbe	Cisco UCS fabric interconnect B	Eth1/49
	Eth1/52	40Gbe	Cisco UCS fabric interconnect A	Eth1/49
	Eth1/1	10Gbe	Cisco Nexus 93180YC-FX B	Eth1/1
	Eth1/2	10Gbe	Cisco Nexus 93180YC-FX B	Eth1/2
	Eth1/3	10Gbe	Cisco Nexus 93180YC-FX B	Eth1/3
	Eth1/4	10Gbe	Cisco Nexus 93180YC-FX B	Eth1/4
	MGMT0	1Gbe	Gbe management switch	Any

Table 13. Cisco Nexus 93180YC-FX-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180YC-FX Switch B	Eth1/51	40Gbe	Cisco UCS fabric interconnect B	Eth1/50
	Eth1/52	40Gbe	Cisco UCS fabric interconnect A	Eth1/50
	Eth1/1	10Gbe	Cisco Nexus 93180YC-FX A	Eth1/1
	Eth1/2	10Gbe	Cisco Nexus 93180YC-FX A	Eth1/2
	Eth1/3	10Gbe	Cisco Nexus 93180YC-FX A	Eth1/3
	Eth1/4	10Gbe	Cisco Nexus 93180YC-FX A	Eth1/4
	MGMT0	1Gbe	Gbe management switch	Any

Cisco UCS Fabric Interconnect 6454 Cabling

The following tables list the Cisco UCS FI 6454 cabling information.

Table 14. Cisco UCS Fabric Interconnect (FI) A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-A	Eth1/17-18	25Gbe	UCS X9508 Chassis IFM-A Chassis 1	Intelligent Fabric Module 1 Port1-2
	Eth1/49	40Gbe	Cisco Nexus 93180YC-FX Switch A	Eth1/52
	Eth1/50	40Gbe	Cisco Nexus 93180YC-FX Switch B	Eth1/52
	Mgmt 0	1Gbe	Management Switch	Any
	L1	1Gbe	Cisco UCS FI - A	L1
	L2	1Gbe	Cisco UCS FI - B	L2

Table 15. Cisco UCS Fabric Interconnect (FI) B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS FI-6454-B	Eth1/17-18	25Gbe	UCS X9508 Chassis IFM-B Chassis 1	Intelligent Fabric Module 1 Port1-2
	Eth1/49	40Gbe	Cisco Nexus 93180YC-FX Switch A	Eth1/51
	Eth1/50	40Gbe	Cisco Nexus 93180YC-FX Switch B	Eth1/51
	Mgmt 0	1Gbe	Management Switch	Any
	L1	1Gbe	Cisco UCS FI - A	L1
	L2	1Gbe	Cisco UCS FI - B	L2

Procedure 1. Create vPC Peer-Link Between the Two Cisco Nexus Switches

Step 1. Log in as “admin” user into the Cisco Nexus Switch A.

Note: For vPC 1 as Peer-link, we used interfaces 53-54 for Peer-Link. You may choose the appropriate number of ports for your needs.

Step 2. Create the necessary port channels between devices by running these commands on both Cisco Nexus switches:

```
config terminal
feature vpc
feature lacp
vpc domain 50
peer-keepalive destination 173.37.52.104 source 173.37.52.103
exit
interface port-channel 10
description VPC peer-link
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type network
vpc peer-link
interface Ethernet1/1
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,70-76,132
channel-group 10 mode active
no shutdown
exit

interface Ethernet1/2
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,70-76,132
channel-group 10 mode active
no shutdown
exit

interface Ethernet1/3
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,70-76,132
channel-group 10 mode active
no shutdown
exit

interface Ethernet1/4
description VPC to K23-N9K-A
switchport mode trunk
switchport trunk allowed vlan 1,70-76,132
```

```
channel-group 10 mode active
no shutdown
exit
copy running-config startup-config
```

Step 3. Log in as admin user into the Nexus Switch B and repeat the above steps to configure second Cisco Nexus switch.

Step 4. Make sure to change the peer-keepalive destination and source IP address appropriately for Cisco Nexus Switch B.

Procedure 2. Create vPC Configuration Between Cisco Nexus 93180YC-FX and Cisco Fabric Interconnects

Create and configure vPC 11 and 12 for the data network between the Cisco Nexus switches and fabric interconnects.

Note: Create the necessary port channels between devices, by running the following commands on both Cisco Nexus switches.

Step 1. Log in as admin user into Cisco Nexus Switch A and enter the following:

```
config terminal
interface port-channel11
description FI-A-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 11
no shutdown
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
```

```
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

Step 2. Log in as admin user into the Nexus Switch B and complete the following for the second switch configuration:

```
config Terminal
interface port-channel11
description FI-A-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 11
no shutdown
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
no shutdown
exit
```

```
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

Verify all vPC Status is up on both Cisco Nexus Switches

Figure 18 shows the verification of the vPC status on both Cisco Nexus Switches.

Figure 18. vPC Description for Cisco Nexus Switch A and B

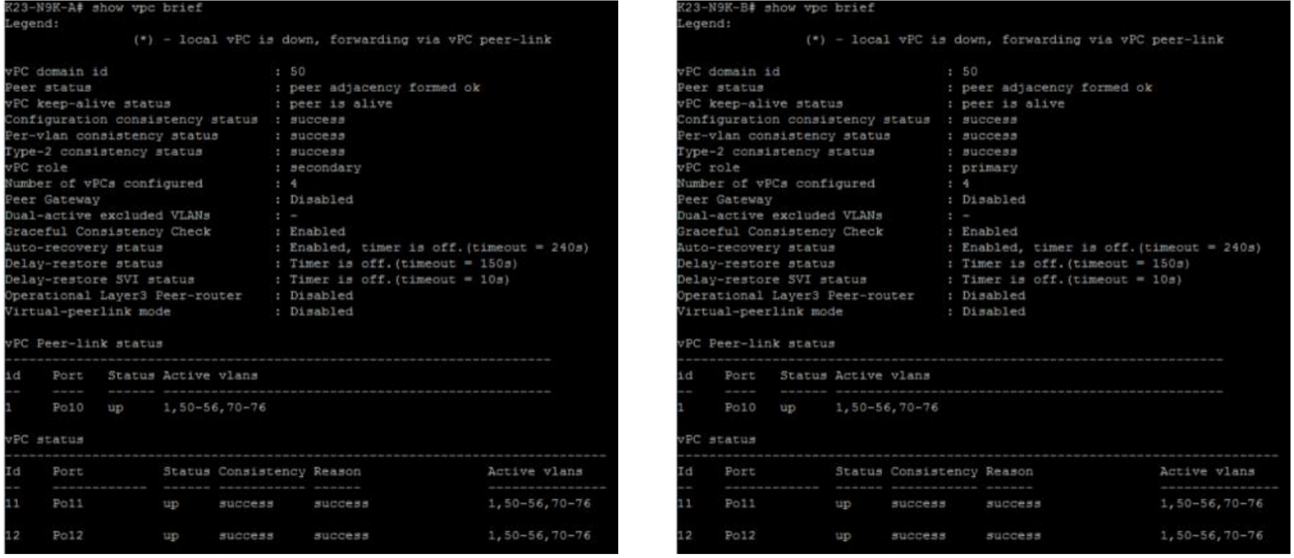

```
K23-N9K-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 50
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 150s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
Virtual-peerlink mode             : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ -------------------------------------------------
1     Po10   up     1,50-56,70-76

vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
11    Po11          up     success     success               1,50-56,70-76

12    Po12          up     success     success               1,50-56,70-76
```

```
K23-N9K-B# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 50
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status  : success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured         : 4
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled
Auto-recovery status              : Enabled, timer is off.(timeout = 240s)
Delay-restore status              : Timer is off.(timeout = 150s)
Delay-restore SVI status          : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router    : Disabled
Virtual-peerlink mode             : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id    Port   Status Active vlans
--    ----   ------ -------------------------------------------------
1     Po10   up     1,50-56,70-76

vPC status
----------------------------------------------------------------------------
Id    Port          Status Consistency Reason                Active vlans
--    ------------  ------ ----------- ------                ---------------
11    Po11          up     success     success               1,50-56,70-76

12    Po12          up     success     success               1,50-56,70-76
```

Install and Configure VMware ESXi 8.0

This section explains how to install VMware ESXi 8.0 Update 1a in an environment.

There are several methods to install 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 UCS Manager to map remote installation media to individual servers and install ESXi on local disk.

Download Cisco Custom Image for VMware vSphere ESXi 8.0

To download the Cisco Custom Image for VMware ESXi 8.0 Update 1a, from the VMware vSphere Hypervisor 8.0 page click the Custom ISOs tab.

Procedure 1. Install VMware vSphere ESXi 8.08.0

Step 1. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Servers.

Step 2. Right-click on ... icon for the server being access and select Launch vKVM.

Step 3. Click Boot Device and then select vKVM Mapped vDVD.

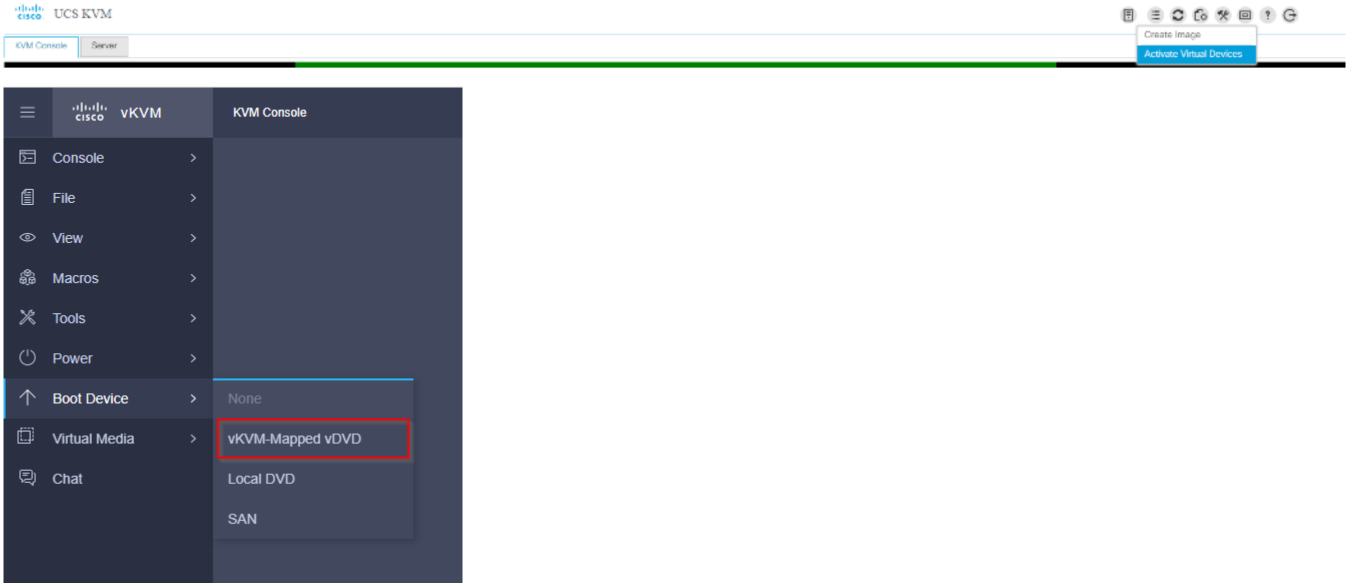

Step 4. Browse to the ESXi iso image file. Click Map Drive to mount the ESXi ISO image.

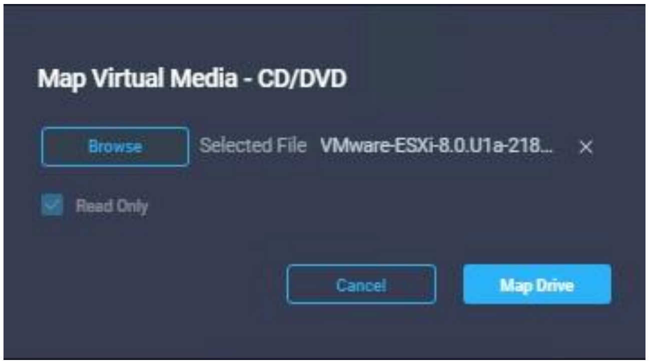

Selected File VMware-ESXi-8.0.U1a-218...

Step 5. Boot into ESXi installer and follow the prompts to complete installing VMware vSphere ESXi hypervisor.

Step 6. When selecting a storage device to install ESXi, select local disk.

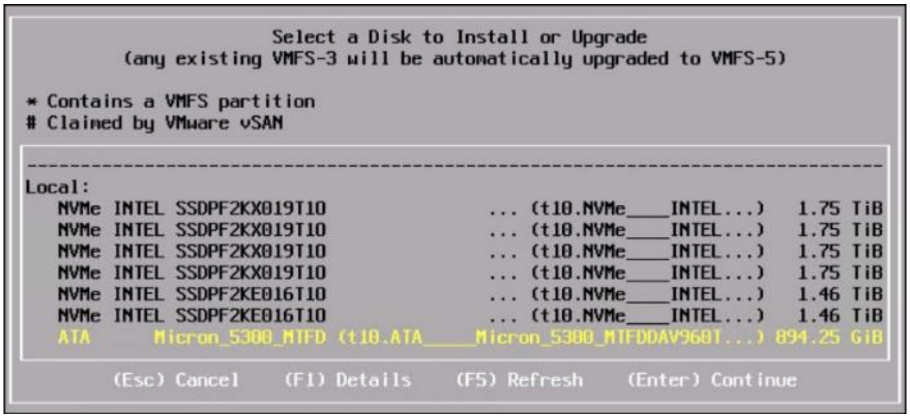

```
                   Select a Disk to Install or Upgrade
        (any existing VMFS-3 will be automatically upgraded to VMFS-5)

 * Contains a VMFS partition
 # Claimed by VMware vSAN

 -------------------------------------------------------------------------------
 Local:
    NVMe  INTEL  SSDPF2KX019T10             ... (t10.NVMe____INTEL...)   1.75 TiB
    NVMe  INTEL  SSDPF2KX019T10             ... (t10.NVMe____INTEL...)   1.75 TiB
    NVMe  INTEL  SSDPF2KX019T10             ... (t10.NVMe____INTEL...)   1.75 TiB
    NVMe  INTEL  SSDPF2KX019T10             ... (t10.NVMe____INTEL...)   1.75 TiB
    NVMe  INTEL  SSDPF2KE016T10             ... (t10.NVMe____INTEL...)   1.46 TiB
    NVMe  INTEL  SSDPF2KE016T10             ... (t10.NVMe____INTEL...)   1.46 TiB
    ATA       Micron_5300_MTFD (t10.ATA_____Micron_5300_MTFDDAV960T...) 894.25 GiB

        (Esc) Cancel     (F1) Details     (F5) Refresh     (Enter) Continue
```

Procedure 2. Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host and connection to vCenter Server. Select the IP address that can communicate with existing or new vCenter Server.

Step 1. After the server has finished rebooting, press F2 to enter in to configuration wizard for ESXi Hypervisor.

Step 2. Log in as root and enter the corresponding password.

Step 3. Select the Configure the Management Network option and press Enter.

Step 4. Select the VLAN (Optional) option and press Enter. Enter the VLAN In-Band management ID and press Enter.

Step 5. From the Configure Management Network menu, select IP Configuration and press Enter.

Step 6. Select the Set Static IP Address and Network Configuration option by using the space bar. Enter the IP address to manage the first ESXi host. Enter the subnet mask for the first ESXi host. Enter the default gateway for the first ESXi host. Press Enter to accept the changes to the IP configuration.

Note: IPv6 Configuration is set to automatic.

Step 7. Select the DNS Configuration option and press Enter.

Step 8. Enter the IP address of the primary and secondary DNS server. Enter Hostname

Step 9. Enter DNS Suffixes.

Note: Since the IP address is assigned manually, the DNS information must also be entered manually.

Note: The steps provided vary based on the configuration. Please make the necessary changes according to your configuration.

Figure 19. Sample ESXi Configure Management Network

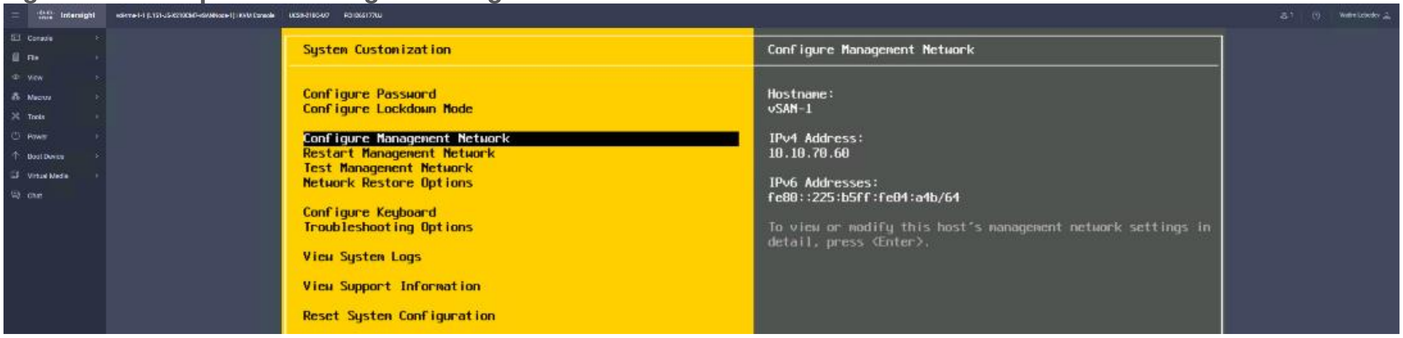

```
System Customization

Configure Password
Configure Lockdown Mode

Configure Management Network
Restart Management Network
Test Management Network
Network Restore Options

Configure Keyboard
Troubleshooting Options

View System Logs

View Support Information

Reset System Configuration
```

```
Configure Management Network

Hostname:
vSAN-1

IPv4 Address:
10.10.70.60

IPv6 Addresses:
fe80::225:b5ff:fe04:a4b/64

To view or modify this host's management network settings in
detail, press <Enter>.
```

Update Cisco VIC Drivers for ESXi

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.

Additionally, Cisco Intersight incorporates an HCL check.

Figure 20. **Servers HCL Status in Cisco Intersight Infrastructure Services**

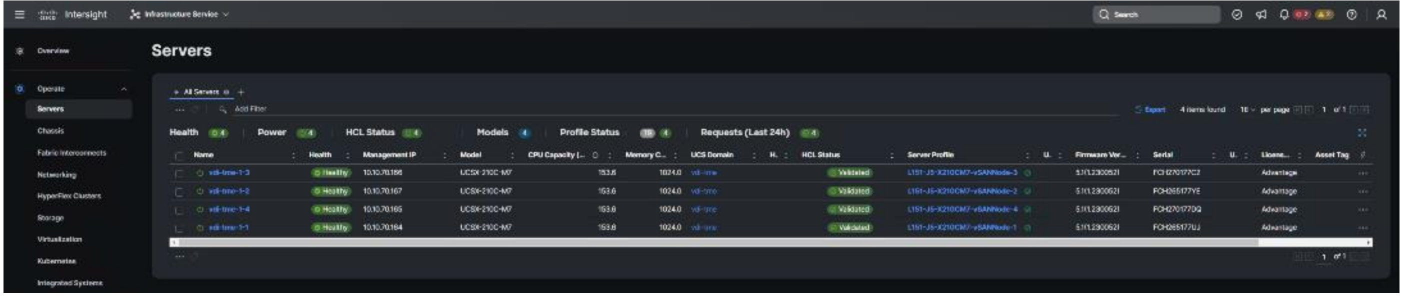

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status	Server Profile	U.	Firmware Ver...	Serial	U.	License...	Asset Tag
vdi-tme-1-3	Healthy	10.10.70.166	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-3		5.1(1.230052)	FCH270177C2		Advantage	
vdi-tme-1-2	Healthy	10.10.70.167	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-2		5.1(1.230052)	FCH265177YE		Advantage	
vdi-tme-1-4	Healthy	10.10.70.165	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-4		5.1(1.230052)	FCH270177DQ		Advantage	
vdi-tme-1-1	Healthy	10.10.70.164	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-1		5.1(1.230052)	FCH265177UJ		Advantage	

In this Cisco Validated Design the following drivers were used:

- Cisco-nenic_2.0.11.0-1OEM.800.1.0.20143090_22197096

Note: For additional information on how to update Cisco VIC drivers on ESXi refer to the Cisco UCS Virtual Interface Card Drivers for ESX Installation Guide.

Cisco Intersight Orchestration

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. This environment includes 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, and helps you add devices into Cisco Intersight.

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. You can install the appliance on an ESXi server. For more information, see the Cisco Intersight Virtual Appliance Getting Started Guide.

After claiming Cisco Intersight Assist into Cisco Intersight, you can claim endpoint devices using the Claim Through Intersight Assist option.

Procedure 1. Configure Cisco Intersight Assist Virtual Appliance

Step 1. To install Cisco Intersight Assist from an Open Virtual Appliance (OVA) in your VMware Management Cluster, first download the latest release of the OVA from: https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-342.

Step 2. To set up the DNS entries for the Cisco Intersight Assist hostname as specified under Before you Begin, go to: https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/cisco-intersight-assist-getting-started-guide/m-installing-cisco-intersight-assist.html.

Step 3. From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the Management cluster and click Deploy OVF Template.

Step 4. Specify a URL or browse to the intersight-appliance-installer-vsphere-1.0.9-342.ova file. Click NEXT.

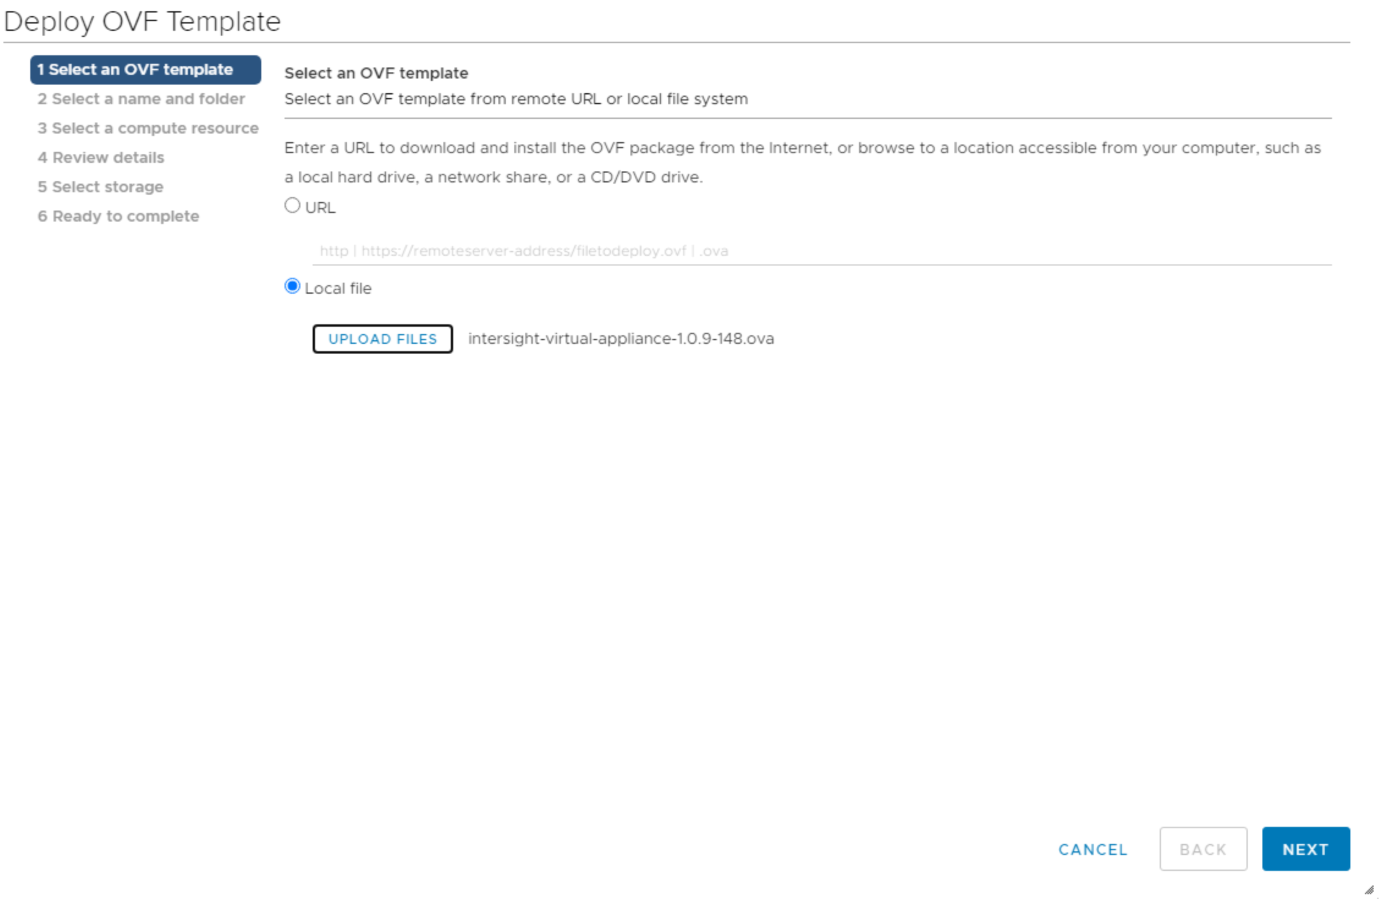

Deploy OVF Template

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

http | https://remoteserver-address/filetodeploy.ovf | .ova

● Local file

UPLOAD FILES intersight-virtual-appliance-1.0.9-148.ova

Step 5. Name the Cisco Intersight Assist VM and choose the location. Click NEXT.

Step 6. Select the Management cluster and click NEXT.

Step 7. Review details and click NEXT.

Step 8. Select a deployment configuration (Tiny recommended) and click NEXT.

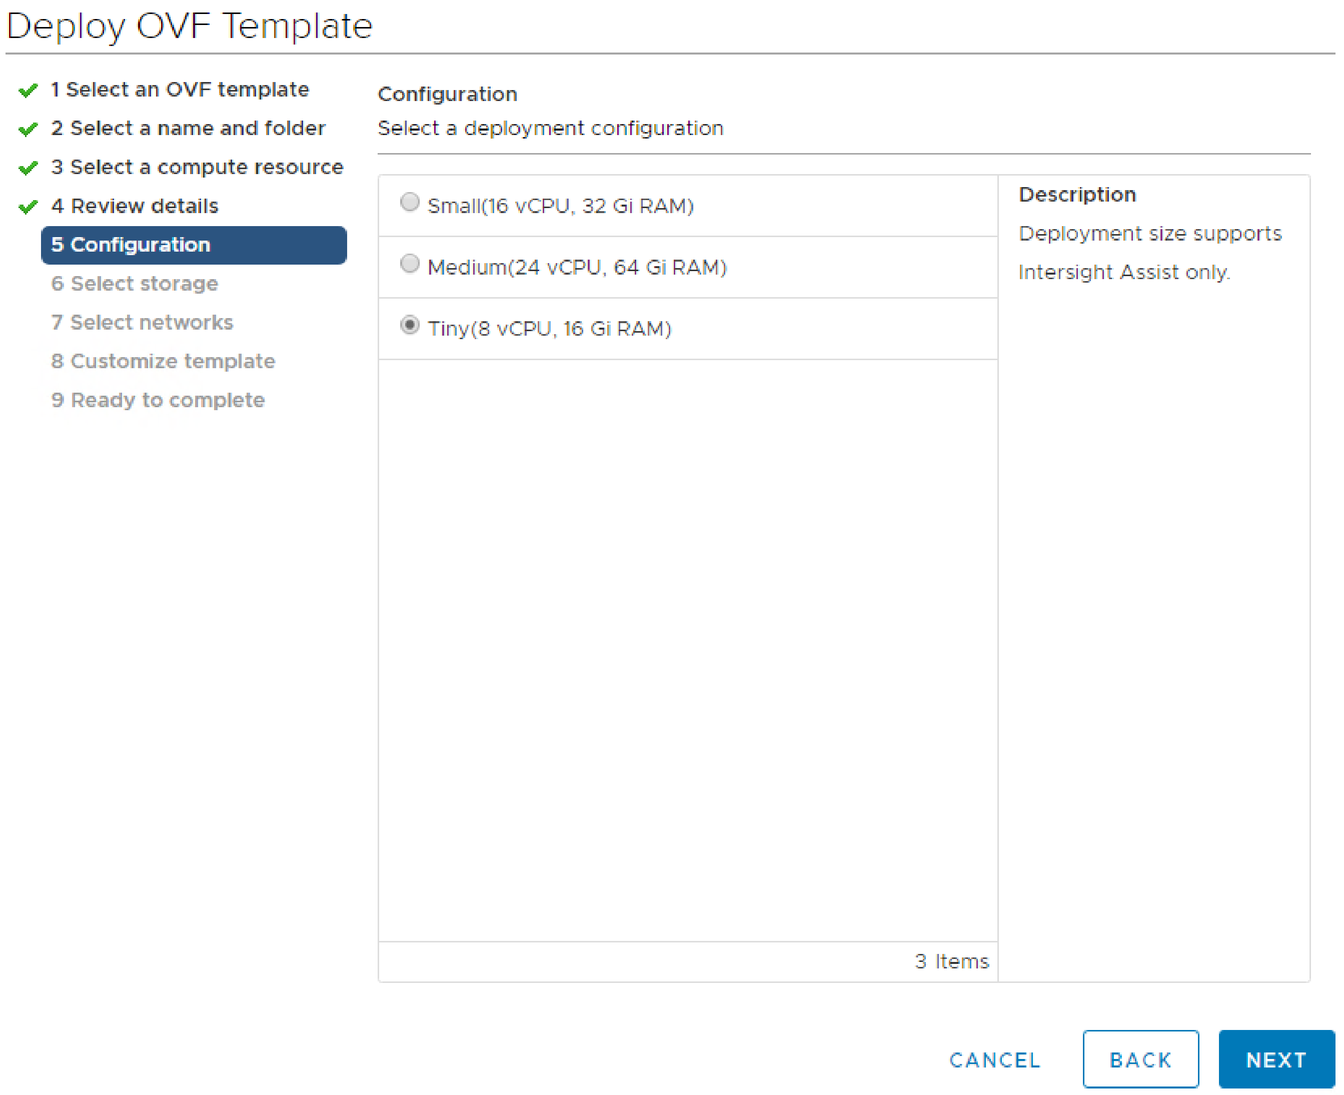

Deploy OVF Template

Configuration

Select a deployment configuration

○ Small(16 vCPU, 32 Gi RAM)

○ Medium(24 vCPU, 64 Gi RAM)

● Tiny(8 vCPU, 16 Gi RAM)

3 Items

Description

Deployment size supports Intersight Assist only.

Step 9. Select the appropriate datastore for storage and select the Thin Provision virtual disk format. Click NEXT.

Step 10. Select IB-MGMT Network for the VM Network. Click NEXT.

Step 11. Fill in all values to customize the template. Click NEXT.

Step 12. Review the deployment information and click FINISH to deploy the appliance.

Step 13. Once the OVA deployment is complete, right-click the Cisco Intersight Assist VM and click Edit Settings.

Step 14. Expand CPU and adjust the Cores per Socket so that 2 Sockets are shown. Click OK.

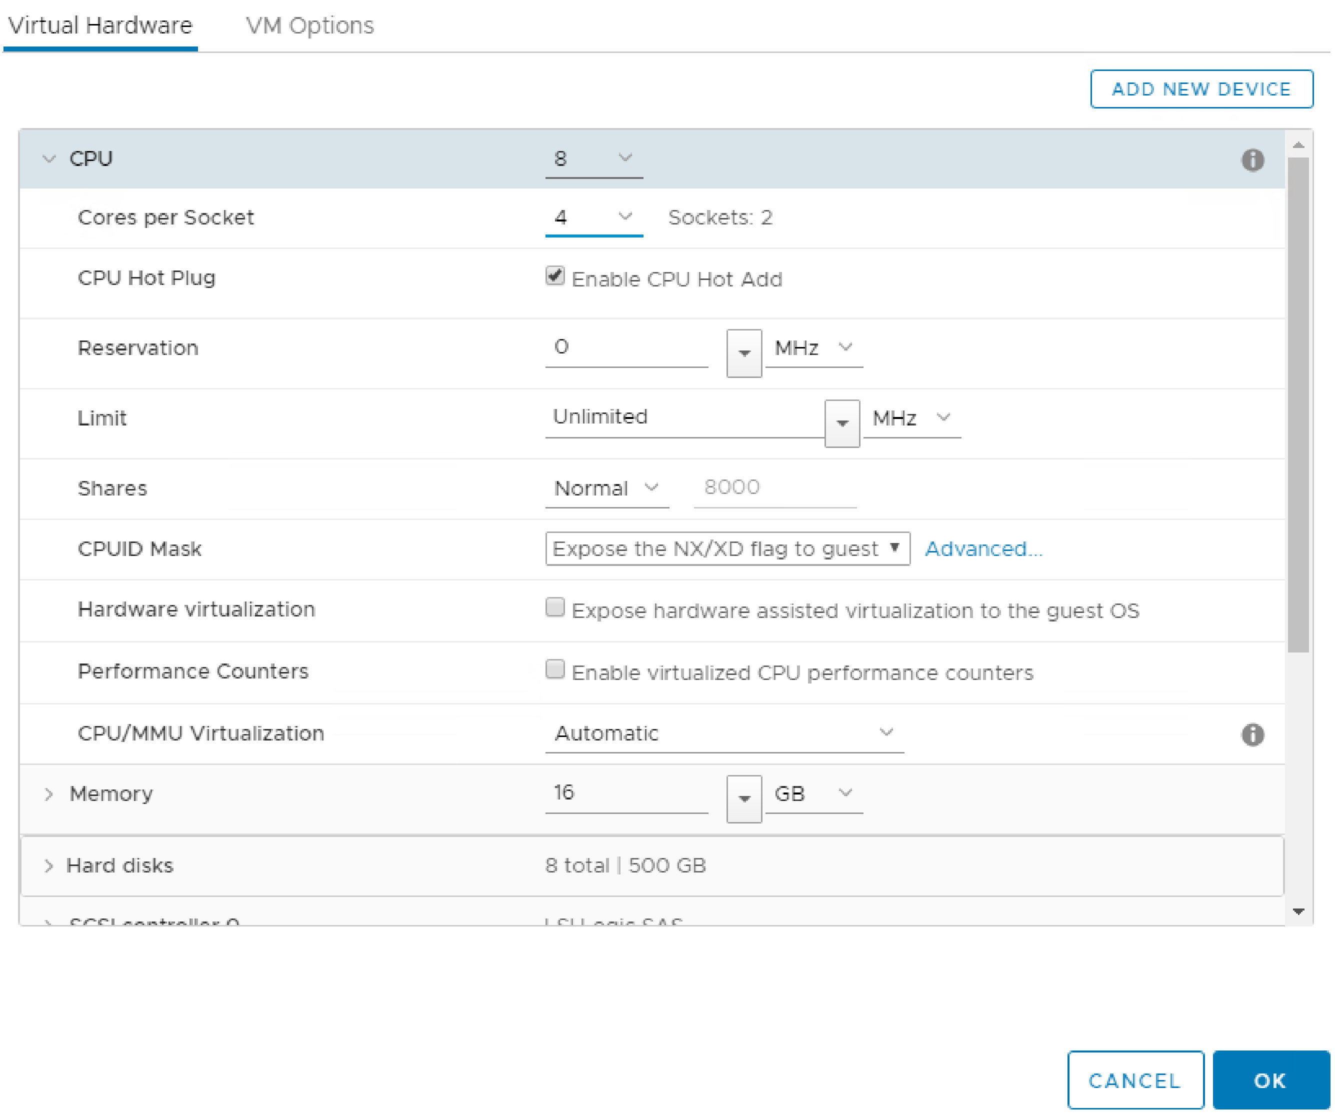

Virtual Hardware | VM Options

CPU	8
Cores per Socket	4 Sockets: 2
CPU Hot Plug	☑ Enable CPU Hot Add
Reservation	0 MHz
Limit	Unlimited MHz
Shares	Normal 8000
CPUID Mask	Expose the NX/XD flag to guest Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS
Performance Counters	□ Enable virtualized CPU performance counters
CPU/MMU Virtualization	Automatic
Memory	16 GB
Hard disks	8 total | 500 GB
SCSI controller 0	LSI Logic SAS

Step 15. Right-click the Cisco Intersight Assist VM and select Open Remote Console.

Step 16. Power on the VM.

Step 17. When you see the login prompt, close the Remote Console, and connect to https://intersight-assist-fqdn.

Note: It may take a few minutes for https://intersight-assist-fqdn to respond.

Step 18. Navigate the security prompts and select Intersight Assist. Click Proceed.

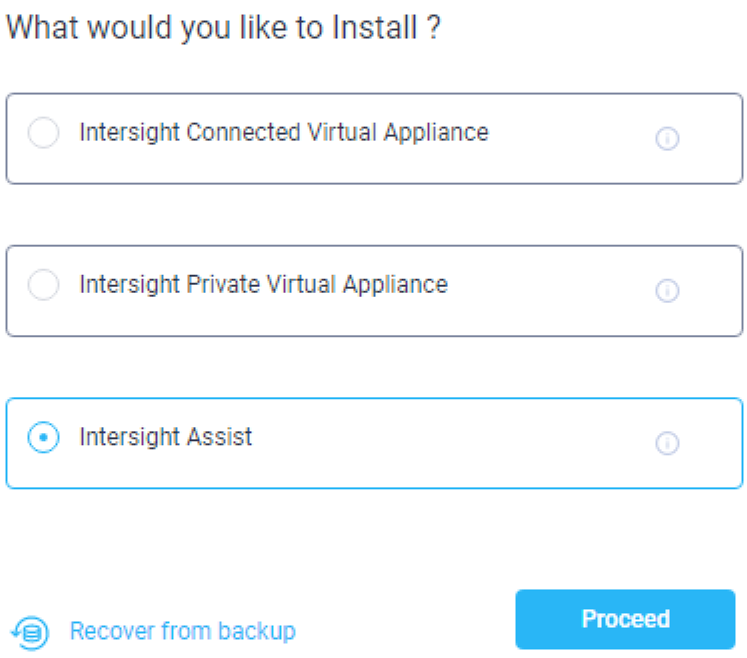

What would you like to Install ?

○ Intersight Connected Virtual Appliance

○ Intersight Private Virtual Appliance

● Intersight Assist

Step 19. From Cisco Intersight, click ADMIN > Devices. Click Claim a New Device. Copy and paste the Device ID and Claim Code shown in the Cisco Intersight Assist web interface to the Cisco Intersight Device Claim Direct Claim window. In Cisco Intersight, click Claim.

Step 20. In the Cisco Intersight Assist web interface, click Continue.

Note: The Cisco Intersight Assist software will now be downloaded and installed into the Cisco Intersight Assist VM. This can take up to an hour to complete.

Note: The Cisco Intersight Assist VM will reboot during the software download process. It will be necessary to refresh the Web Browser after the reboot is complete to follow the status of the download process.

Step 21. When the software download is complete, navigate the security prompts and a Cisco Intersight Assist login screen will appear. Log into Cisco Intersight Assist with the admin user and the password supplied in the OVA installation. Check the Cisco Intersight Assist status and log out of Intersight Assist.

Procedure 2. Claim Intersight Assist into Cisco Intersight

Step 1. To claim the Intersight assist appliance, from the Service Selector drop-down list, select System.

Step 2. From Cisco Intersight, click ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select Cisco Intersight Assist under Platform Services and click Start.

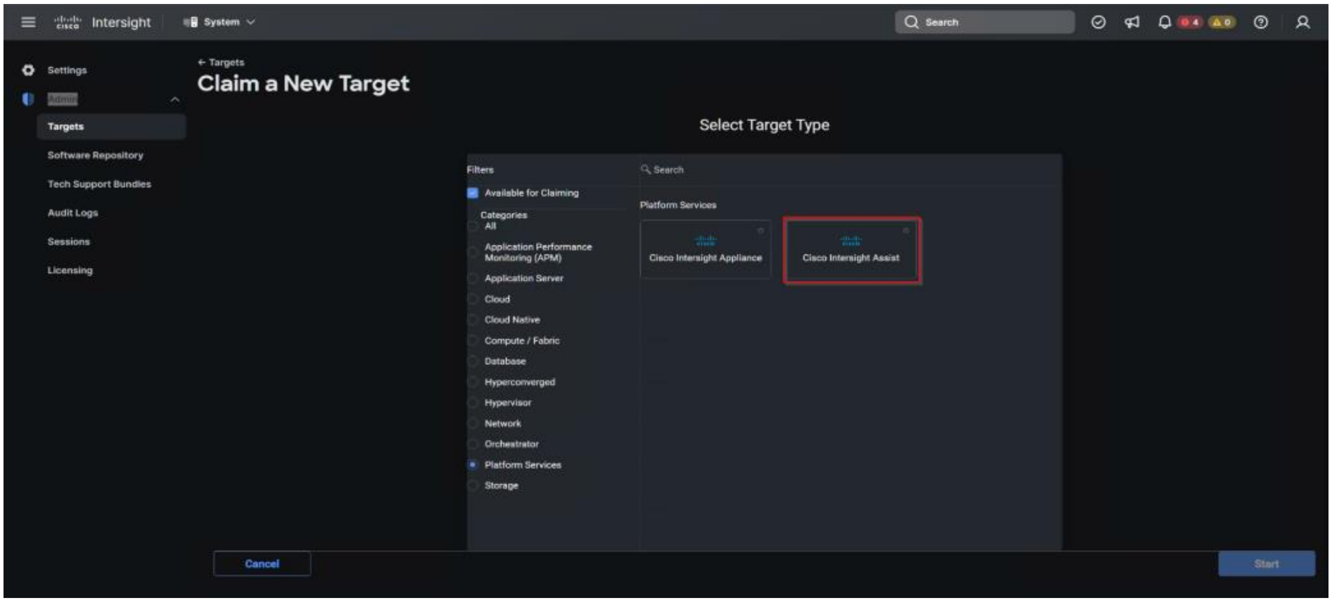

Step 3. Fill in the Intersight Assist information and click Claim.

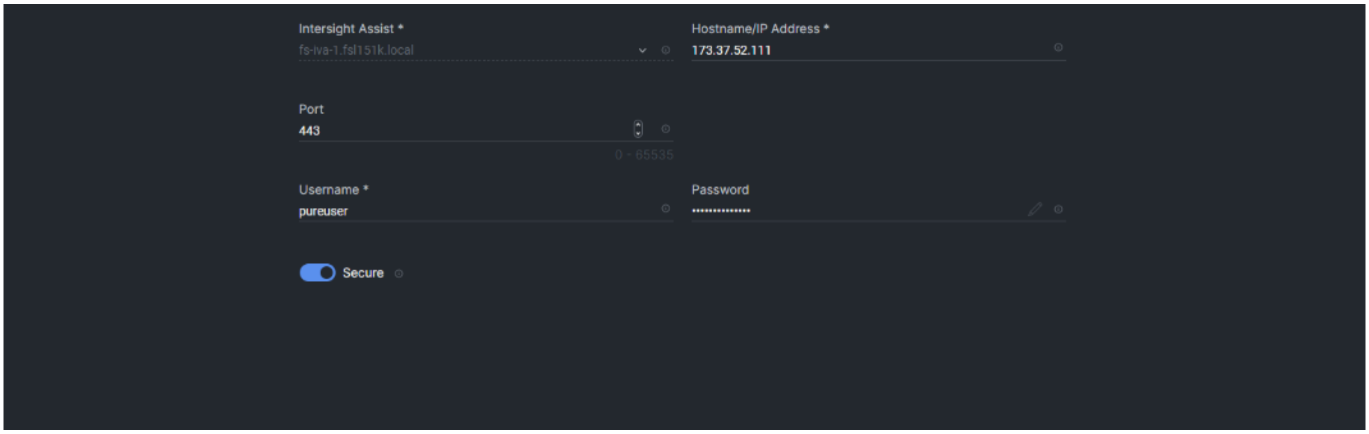

Intersight Assist *
fs-iva-1.fsl151k.local

Hostname/IP Address *
173.37.52.111

Port
443
0 - 65535

Username *
pureuser

Password
•••••••••••••••

Secure

After a few minutes, Cisco Intersight Assist will appear in the Targets list.

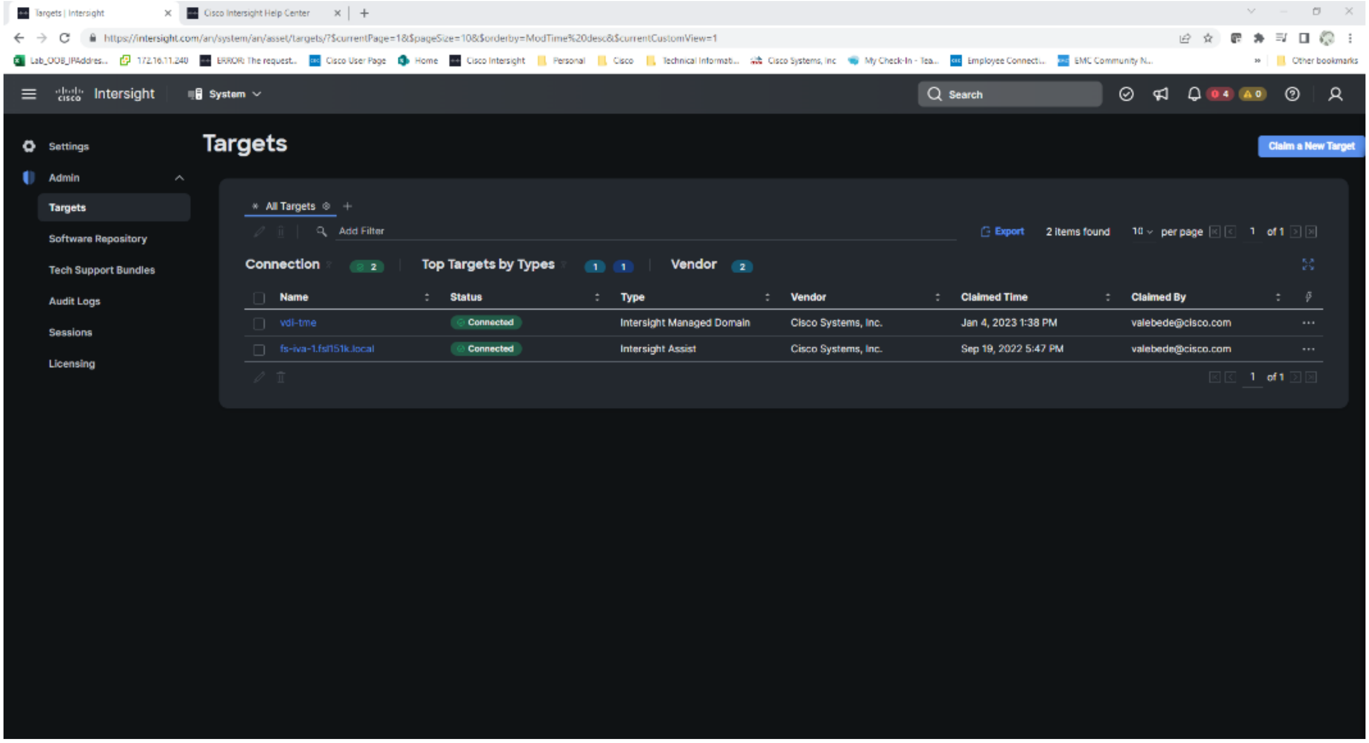

Targets

Name	Status	Type	Vendor	Claimed Time	Claimed By
vdi-tme	Connected	Intersight Managed Domain	Cisco Systems, Inc.	Jan 4, 2023 1:38 PM	valebede@cisco.com
fs-iva-1.fsl151k.local	Connected	Intersight Assist	Cisco Systems, Inc.	Sep 19, 2022 5:47 PM	valebede@cisco.com

Procedure 3. Claim vCenter in Cisco Intersight

Step 1. To claim the vCenter, from Cisco Intersight, click ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select VMware vCenter under Hypervisor and click Start.

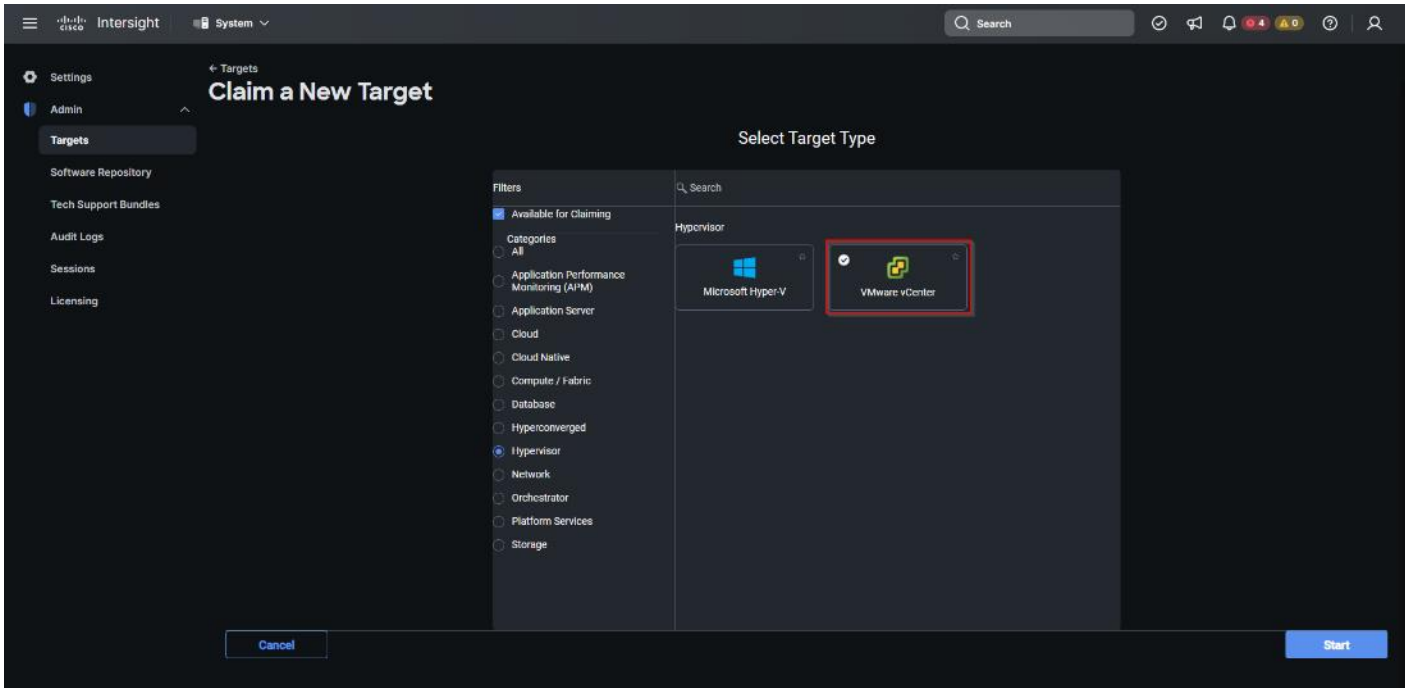

Step 2. In the VMware vCenter window, make sure the Intersight Assist is correctly selected, fill in the vCenter information, and click Claim.

Note: Enabling HSM will give escalated privileges to the vCenter target to perform firmware operations on UCS servers claimed in Cisco Intersight.

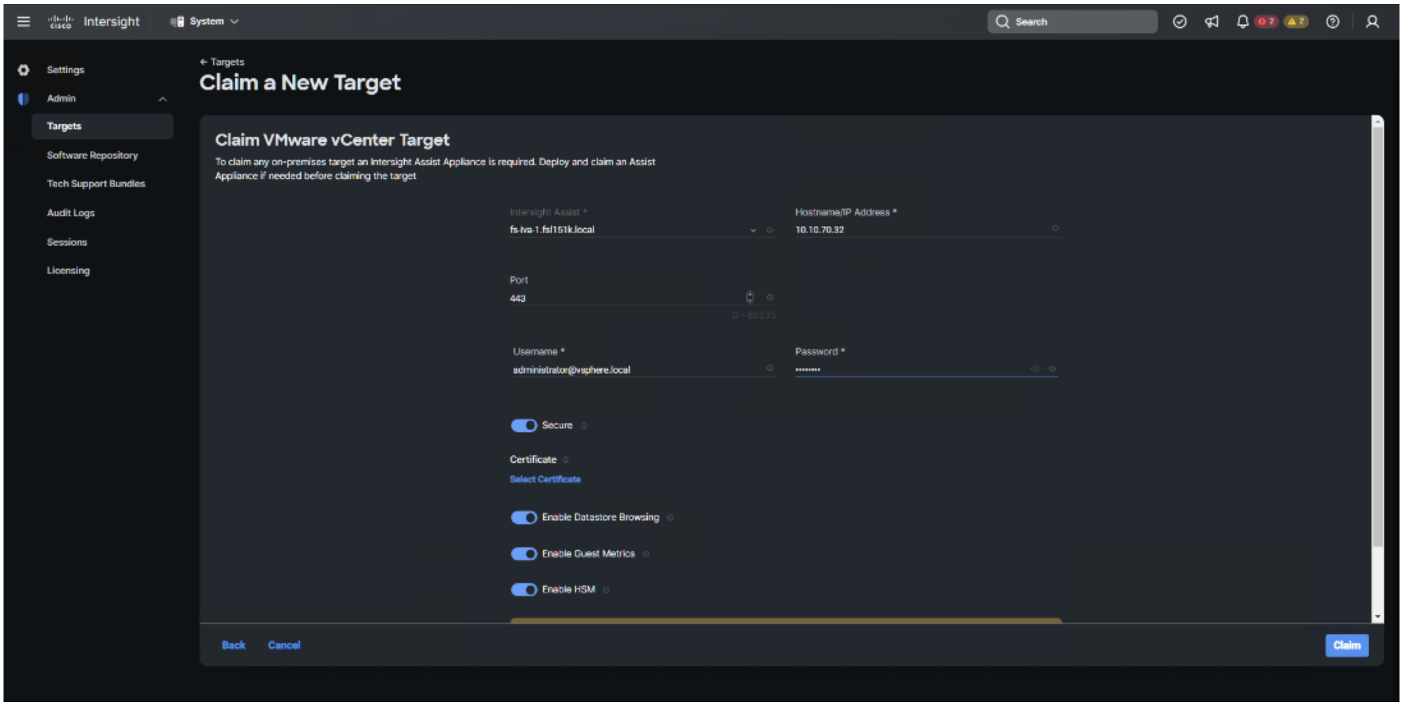

Claim a New Target

Claim VMware vCenter Target

To claim any on-premises target an Intersight Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target

Intersight Assist *: fs-iva-1.fsl151k.local

Hostname/IP Address *: 10.10.70.32

Port: 443 (0 - 65535)

Username *: administrator@vsphere.local

Password *: ••••••••

Secure: on

Certificate: Select Certificate

Enable Datastore Browsing: on

Enable Guest Metrics: on

Enable HSM: on

Step 3. After a few minutes, the VMware vCenter will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

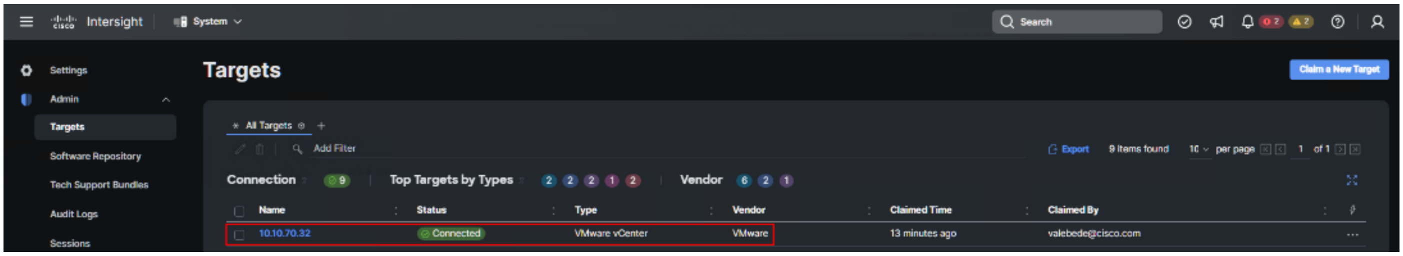

Targets

Name	Status	Type	Vendor	Claimed Time	Claimed By
10.10.70.32	Connected	VMware vCenter	VMware	13 minutes ago	valebede@cisco.com

Step 4. Detailed information obtained from the vCenter can now be viewed by clicking Virtualization from the Infrastructure service > Operate menu.

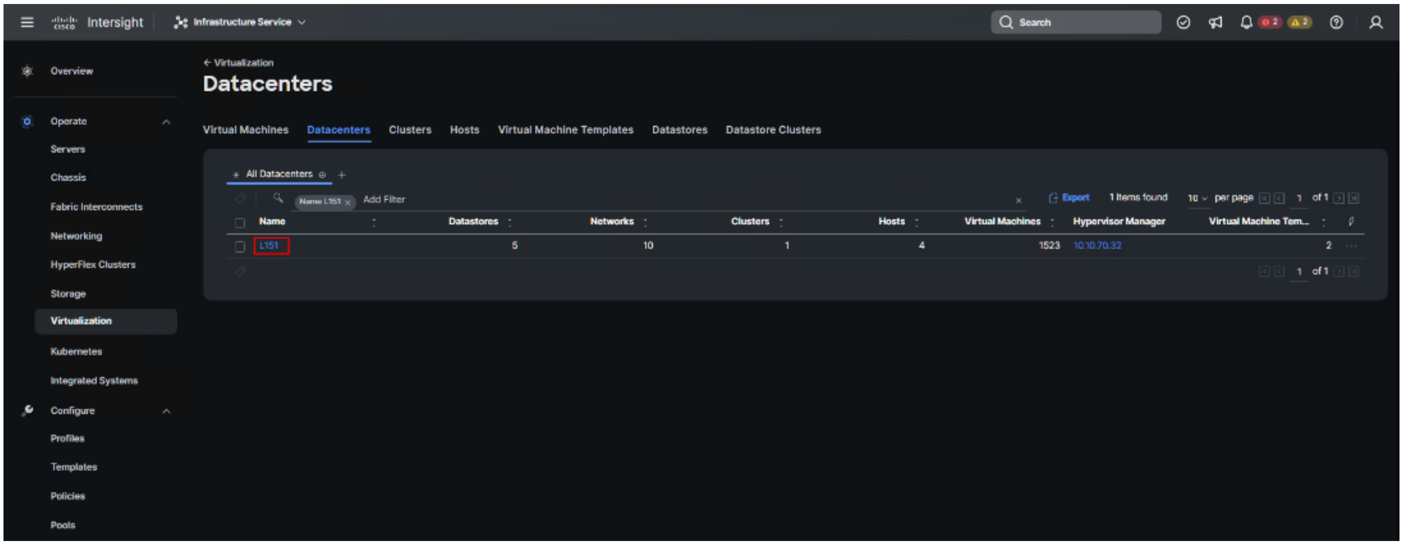

← Virtualization

Datacenters

Name	Datastores	Networks	Clusters	Hosts	Virtual Machines	Hypervisor Manager	Virtual Machine Tem...
L151	5	10	1	4	1523	10.10.70.32	2

Procedure 4. Verify Cisco UCS Server HCL Status using Cisco Intersight

Step 1. From the Infrastructure Service click Operate >Servers, HCL Status field will provide the status overview.

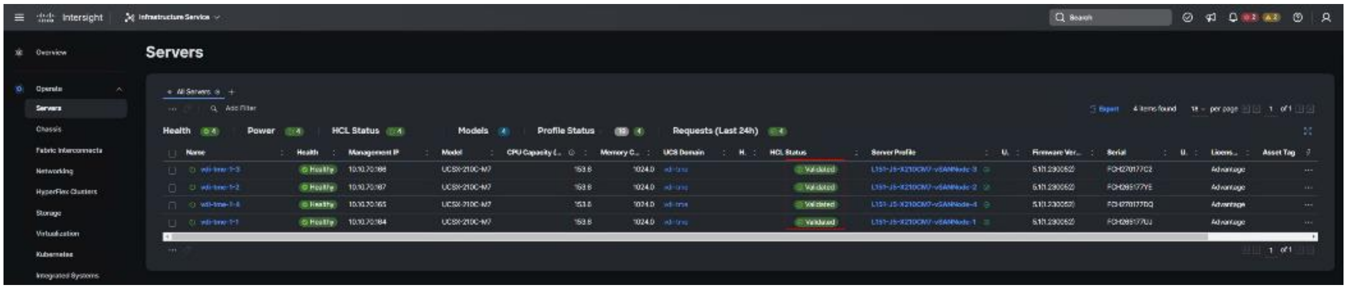

Servers

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status	Server Profile	U.	Firmware Ver...	Serial	U.	Licens...	Asset Tag
vdi-tme-1-3	Healthy	10.10.70.188	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-3		5.1(1.230052)	FCH270177C2		Advantage	
vdi-tme-1-2	Healthy	10.10.70.187	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-2		5.1(1.230052)	FCH285177YE		Advantage	
vdi-tme-1-4	Healthy	10.10.70.165	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-4		5.1(1.230052)	FCH270177DQ		Advantage	
vdi-tme-1-1	Healthy	10.10.70.184	UCSX-210C-M7	153.6	1024.0	vdi-tme		Validated	L151-J5-X210CM7-vSANNode-1		5.1(1.230052)	FCH285177UJ		Advantage	

Step 2. Select a server and click the HCL tab to view validation details.

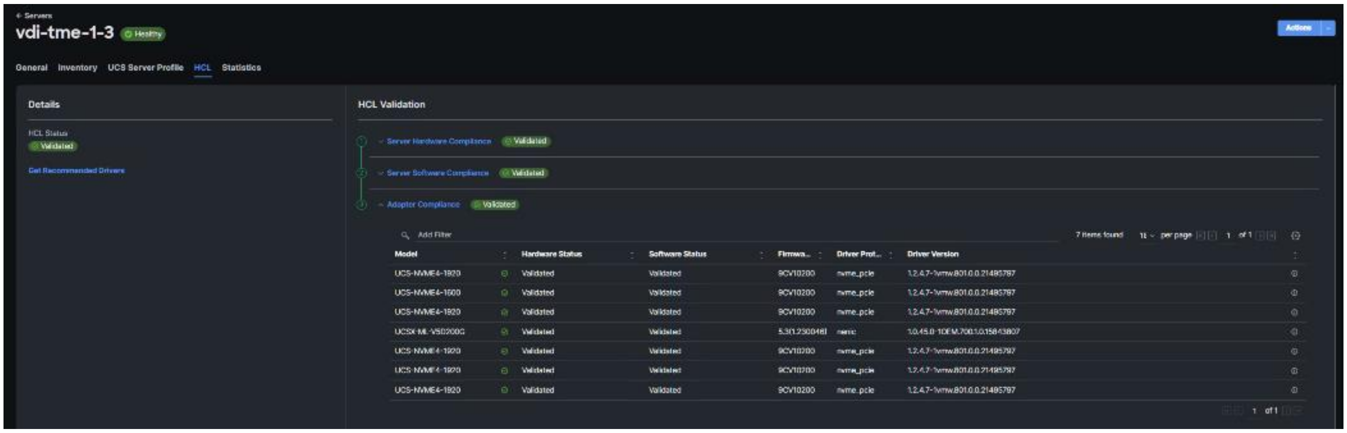

← Servers

vdi-tme-1-3 Healthy

Details

HCL Status: Validated

Get Recommended Drivers

HCL Validation

- Server Hardware Compliance: Validated
- Server Software Compliance: Validated
- Adapter Compliance: Validated

Model	Hardware Status	Software Status	Firmwa...	Driver Prot...	Driver Version
UCS-NVME4-1920	Validated	Validated	9CV10200	nvme_pcie	1.2.4.7-1vmw.801.0.0.21495797
UCS-NVME4-1600	Validated	Validated	9CV10200	nvme_pcie	1.2.4.7-1vmw.801.0.0.21495797
UCS-NVME4-1920	Validated	Validated	9CV10200	nvme_pcie	1.2.4.7-1vmw.801.0.0.21495797
UCSX-ML-V5D200G	Validated	Validated	5.3(1.230048)	nenic	1.0.45.0-1OEM.700.1.0.15843807
UCS-NVME4-1920	Validated	Validated	9CV10200	nvme_pcie	1.2.4.7-1vmw.801.0.0.21495797
UCS-NVME4-1920	Validated	Validated	9CV10200	nvme_pcie	1.2.4.7-1vmw.801.0.0.21495797
UCS-NVME4-1920	Validated	Validated	9CV10200	nvme_pcie	1.2.4.7-1vmw.801.0.0.21495797

Create vSAN Cluster

Note: There are many considerations when provisioning a VMware vSAN software-defined storage solution. This section describes a manual process to create the vSAN cluster. Refer to official VMware documentation if you want set up the vSAN cluster using Quickstart workflow.

Procedure 1. Create and Prepare vSphere Cluster for vSAN

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.

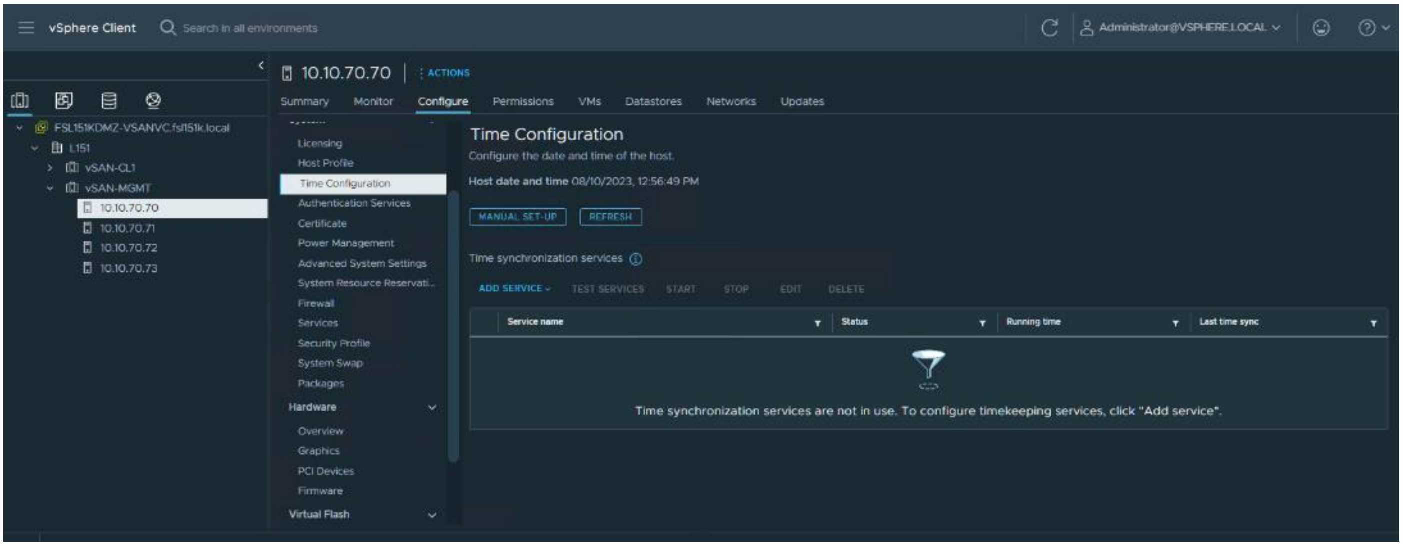

Time Configuration

Configure the date and time of the host.

Host date and time 08/10/2023, 12:56:49 PM

Time synchronization services

Time synchronization services are not in use. To configure timekeeping services, click "Add service".

Step 2. From Add Service drop down list select Network Time protocol option.

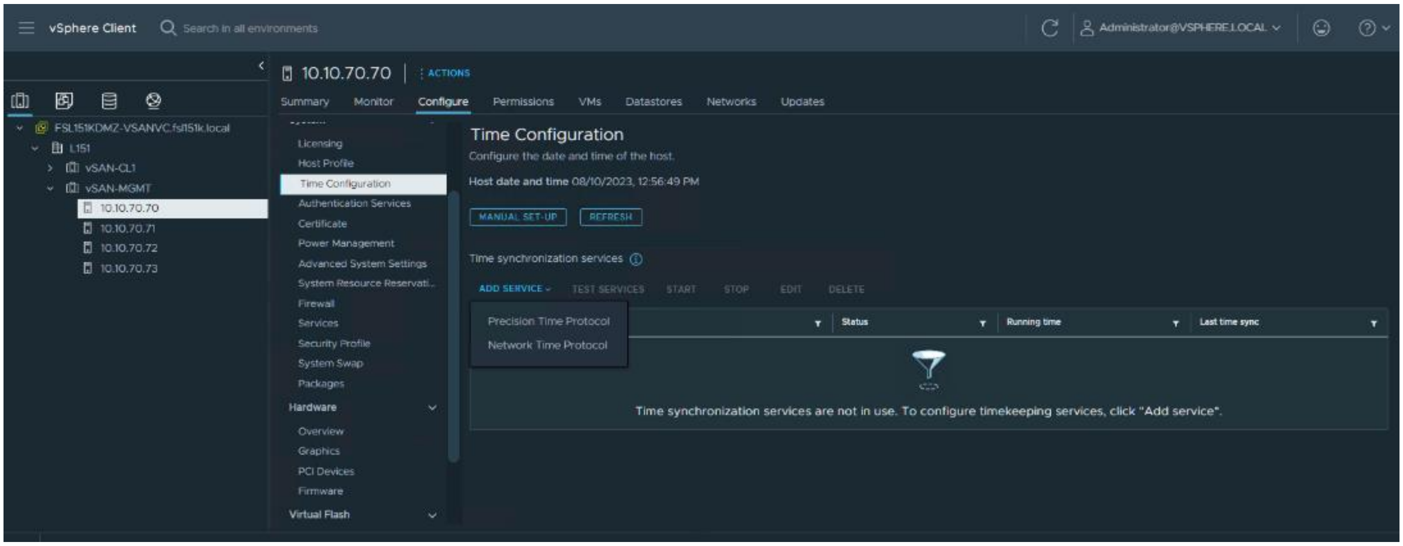

Time Configuration

Configure the date and time of the host.

Host date and time 08/10/2023, 12:56:49 PM

Time synchronization services

Time synchronization services are not in use. To configure timekeeping services, click "Add service".

Step 3. Provide the IP addresses for the NTP servers in your environment. Click OK.

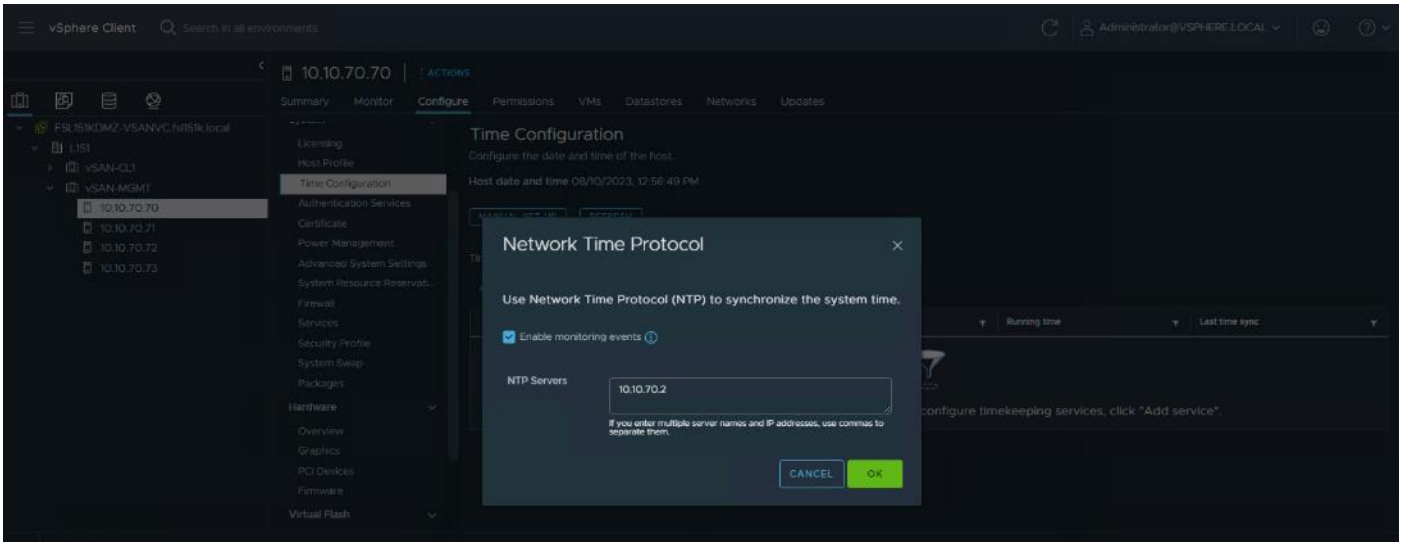

Time Configuration

Configure the date and time of the host.

Host date and time 08/10/2023, 12:56:49 PM

Network Time Protocol

Use Network Time Protocol (NTP) to synchronize the system time.

☑ Enable monitoring events

NTP Servers: 10.10.70.2

If you enter multiple server names and IP addresses, use commas to separate them.

CANCEL OK

...configure timekeeping services, click "Add service".

Step 4. Test the service configuration.

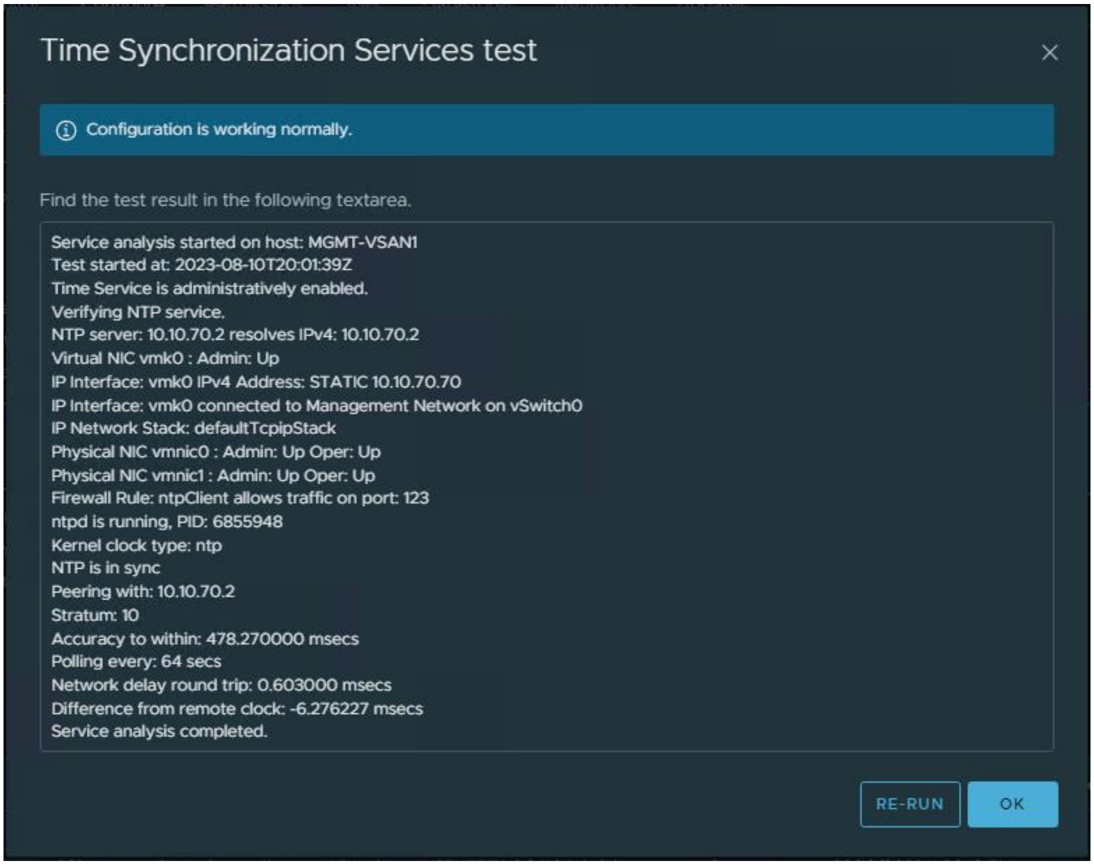

Time Synchronization Services test

Configuration is working normally.

Find the test result in the following textarea.

```
Service analysis started on host: MGMT-VSAN1
Test started at: 2023-08-10T20:01:39Z
Time Service is administratively enabled.
Verifying NTP service.
NTP server: 10.10.70.2 resolves IPv4: 10.10.70.2
Virtual NIC vmk0 : Admin: Up
IP Interface: vmk0 IPv4 Address: STATIC 10.10.70.70
IP Interface: vmk0 connected to Management Network on vSwitch0
IP Network Stack: defaultTcpipStack
Physical NIC vmnic0 : Admin: Up Oper: Up
Physical NIC vmnic1 : Admin: Up Oper: Up
Firewall Rule: ntpClient allows traffic on port: 123
ntpd is running, PID: 6855948
Kernel clock type: ntp
NTP is in sync
Peering with: 10.10.70.2
Stratum: 10
Accuracy to within: 478.270000 msecs
Polling every: 64 secs
Network delay round trip: 0.603000 msecs
Difference from remote clock: -6.276227 msecs
Service analysis completed.
```

Step 5. Review vCenter configuration.

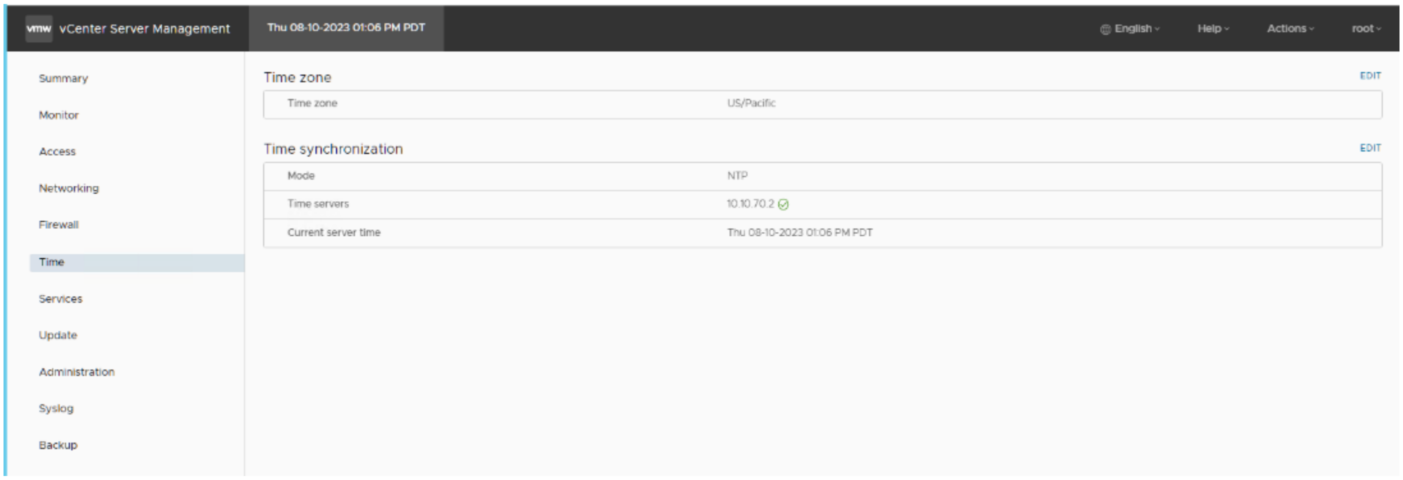

Time zone

Time zone	US/Pacific

Time synchronization

Mode	NTP
Time servers	10.10.70.2 ✓
Current server time	Thu 08-10-2023 01:06 PM PDT

Note: SSH was also enabled to support data collection.

Step 6. Create Distributed Switch for vSAN Traffic.

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.

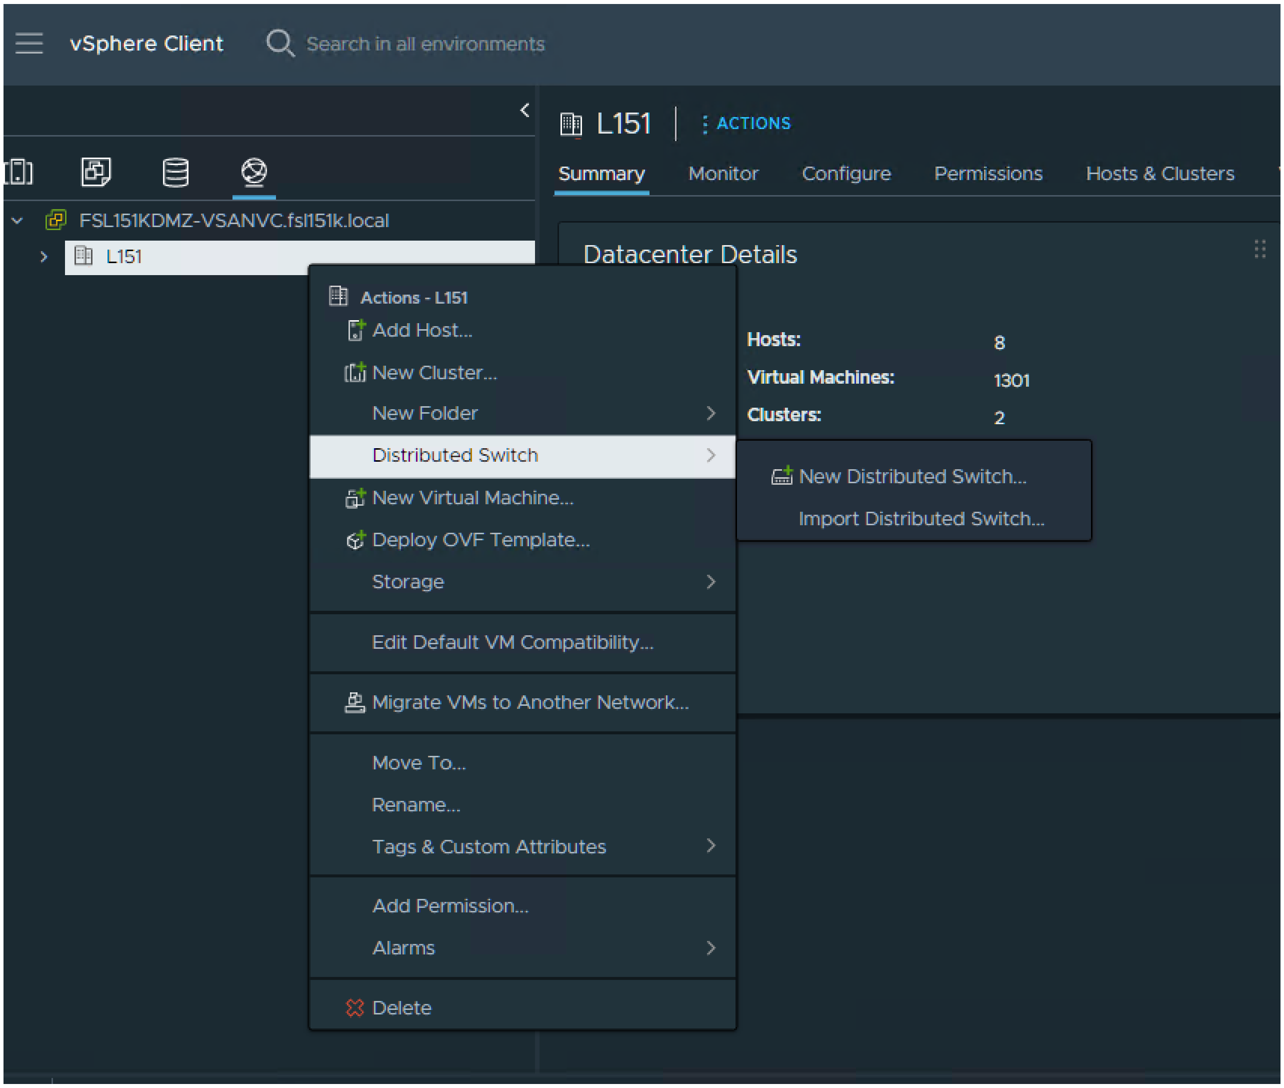

Step 2. Provide a name for the new distributed switch and select the location within the vCenter inventory where you would like to store the new vDS (a data center object or a folder). Click NEXT.

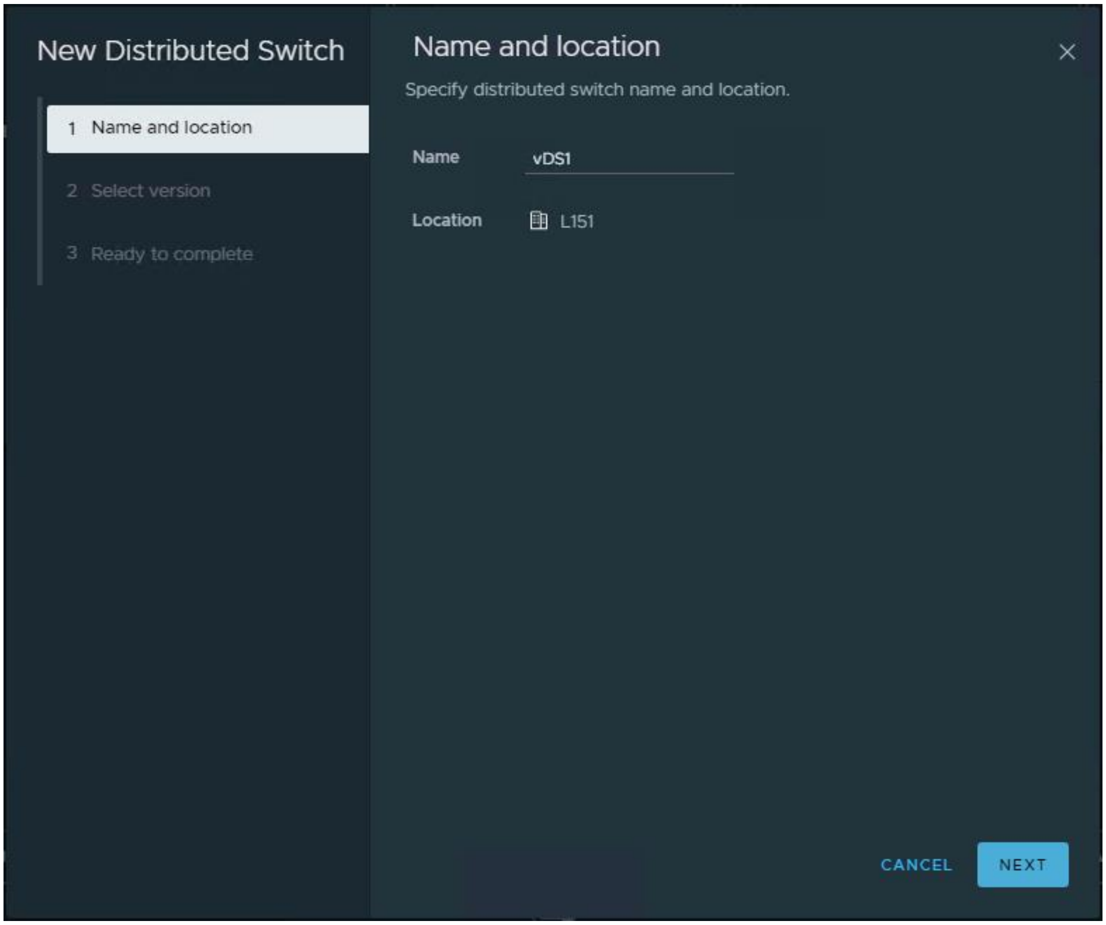

New Distributed Switch

Name and location

Specify distributed switch name and location.

Name: vDS1

Location: L151

Step 3. Select the version of the vDS to create. Click NEXT.

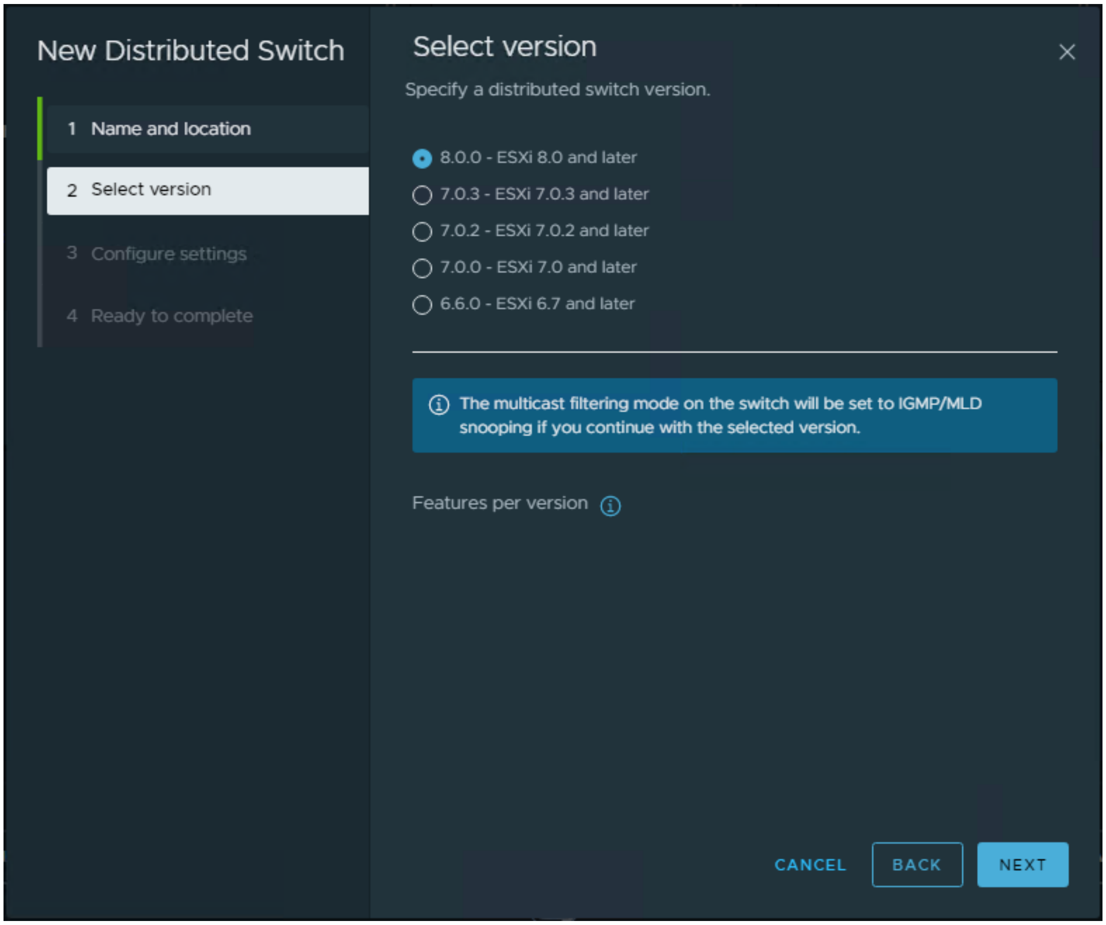

New Distributed Switch

Select version

Specify a distributed switch version.

● 8.0.0 - ESXi 8.0 and later
○ 7.0.3 - ESXi 7.0.3 and later
○ 7.0.2 - ESXi 7.0.2 and later
○ 7.0.0 - ESXi 7.0 and later
○ 6.6.0 - ESXi 6.7 and later

ⓘ The multicast filtering mode on the switch will be set to IGMP/MLD snooping if you continue with the selected version.

Features per version ⓘ

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.

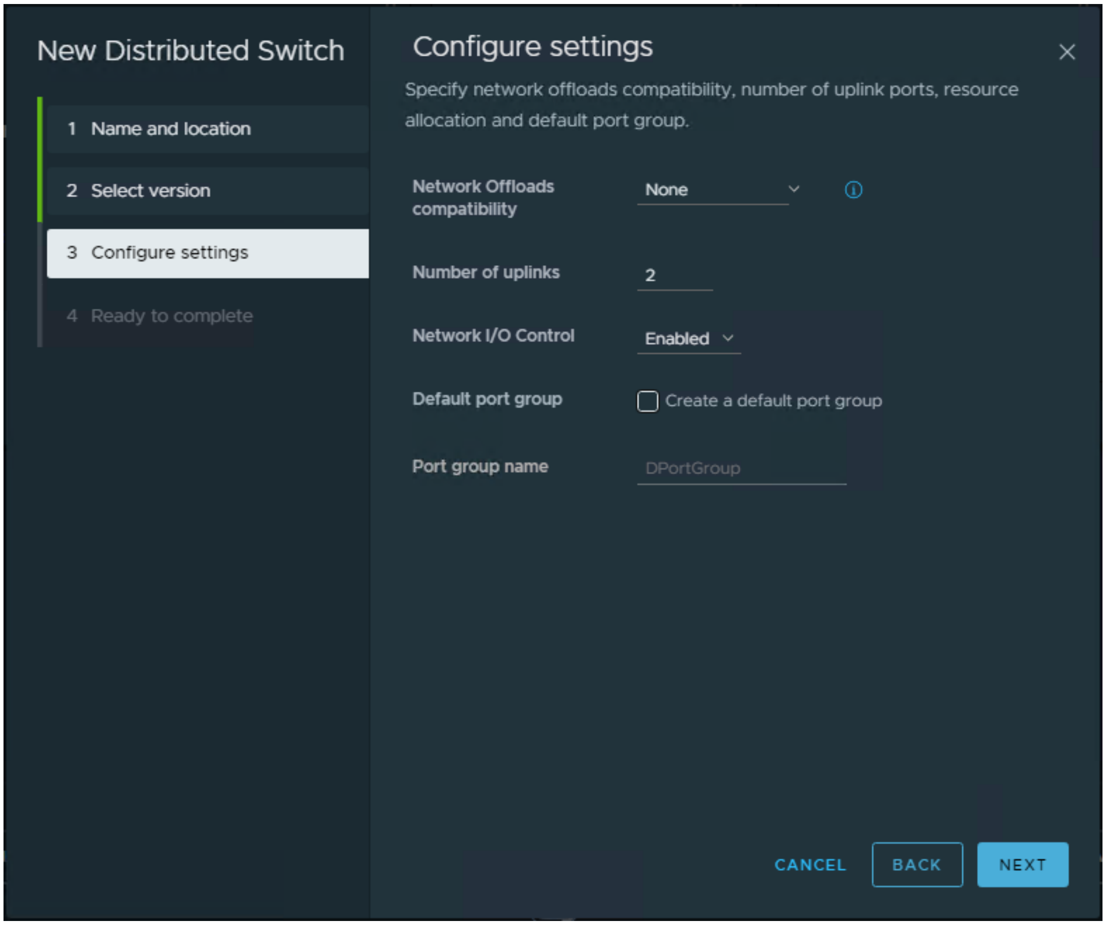

Configure settings

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

Step 5. Click Finish.

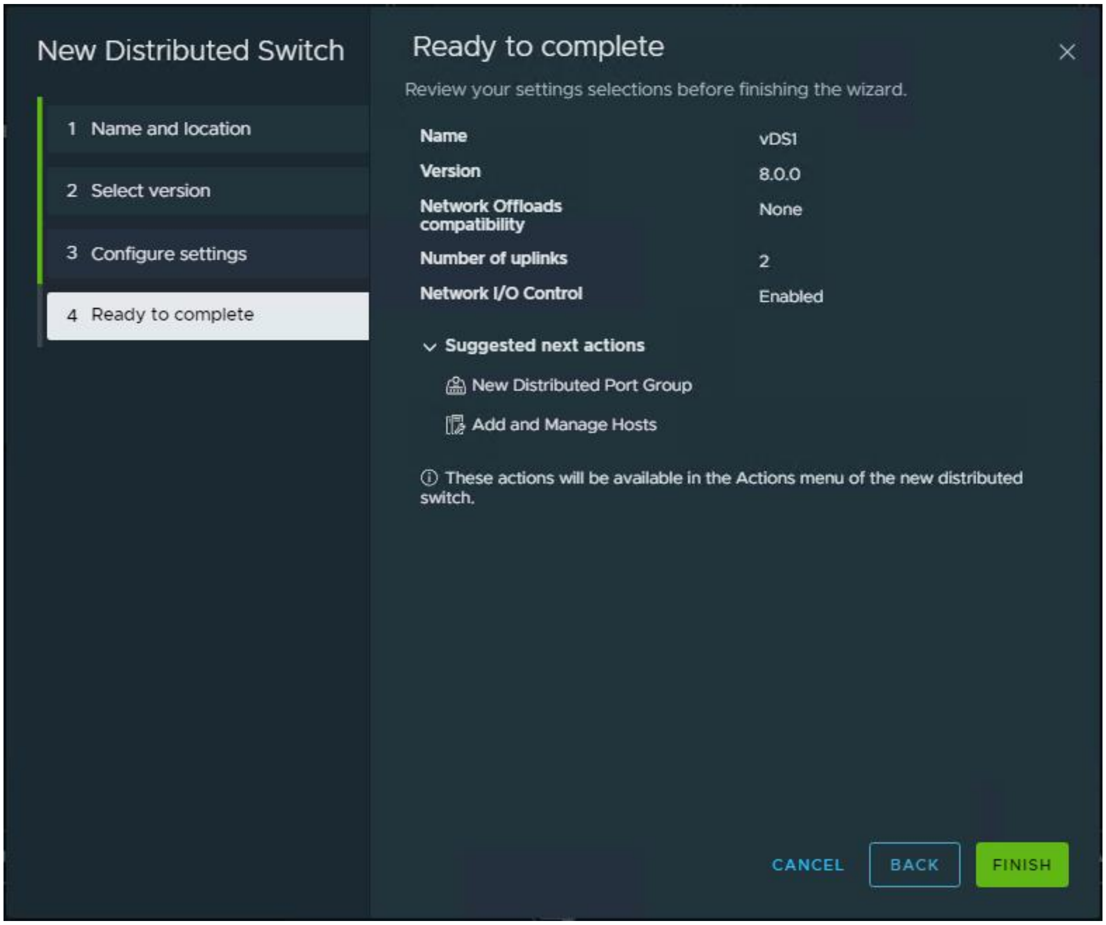

Ready to complete

Review your settings selections before finishing the wizard.

Name	vDS1
Version	8.0.0
Network Offloads compatibility	None
Number of uplinks	2
Network I/O Control	Enabled

Suggested next actions

- New Distributed Port Group
- Add and Manage Hosts

ⓘ These actions will be available in the Actions menu of the new distributed switch.

Step 6. Right-click the new distributed switch in the list of objects and select Settings > Edit Settings....

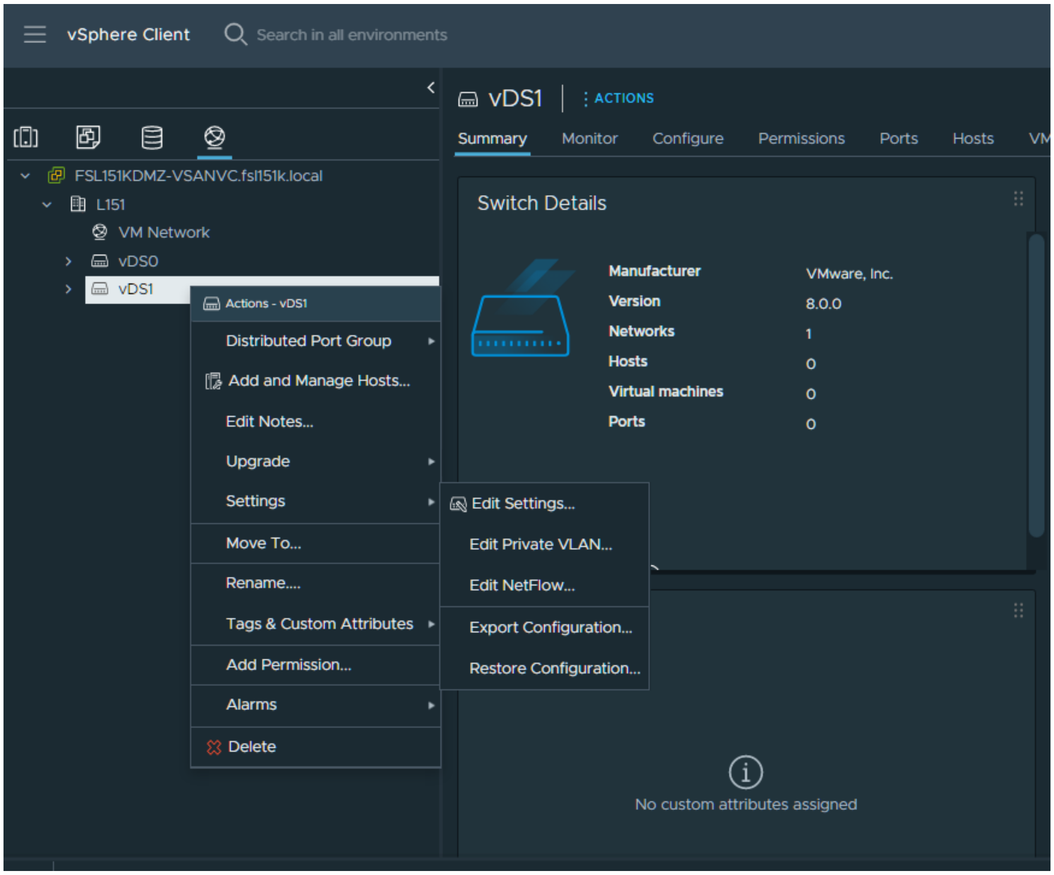

Step 7. In the Distributed Switch-Edit Settings dialog box Advanced tab, set the MTU to 9000, Discovery protocol to Link Layer Discovery Protocol and Operation to Both. Click OK.

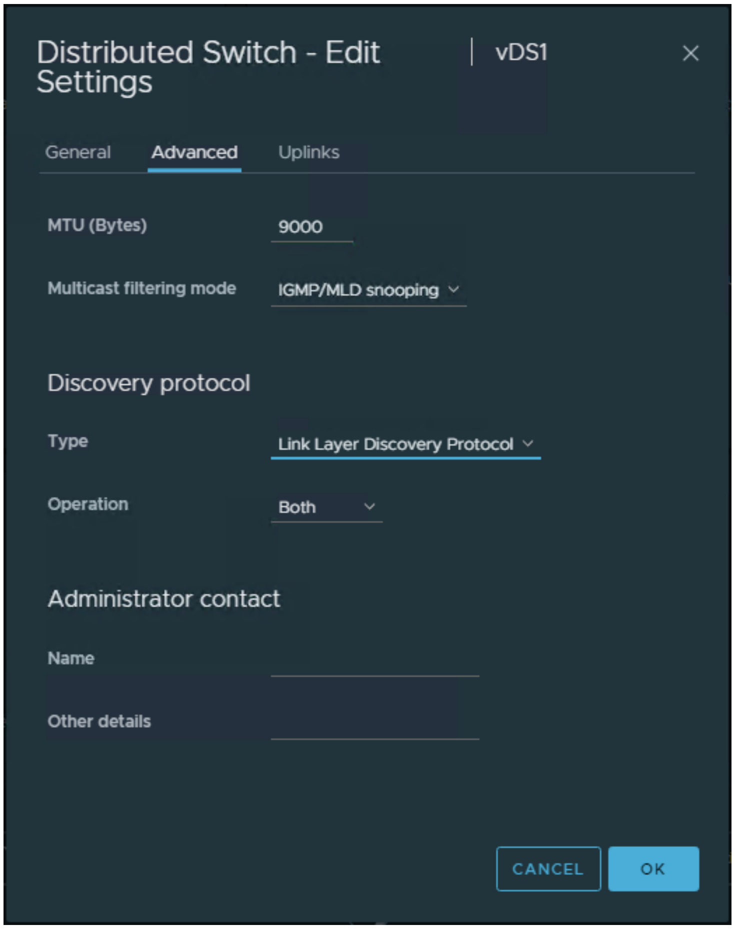

Distributed Switch - Edit Settings | vDS1

MTU (Bytes): 9000

Multicast filtering mode: IGMP/MLD snooping

Discovery protocol

Type: Link Layer Discovery Protocol

Operation: Both

Administrator contact

Name:

Other details:

Note: In server profiles vmnic2 and vmnic3 are created for the use as vDS uplinks.

Step 8. Right-click the new distributed switch in the list of objects and select Add and Manage Hosts from the Actions menu.

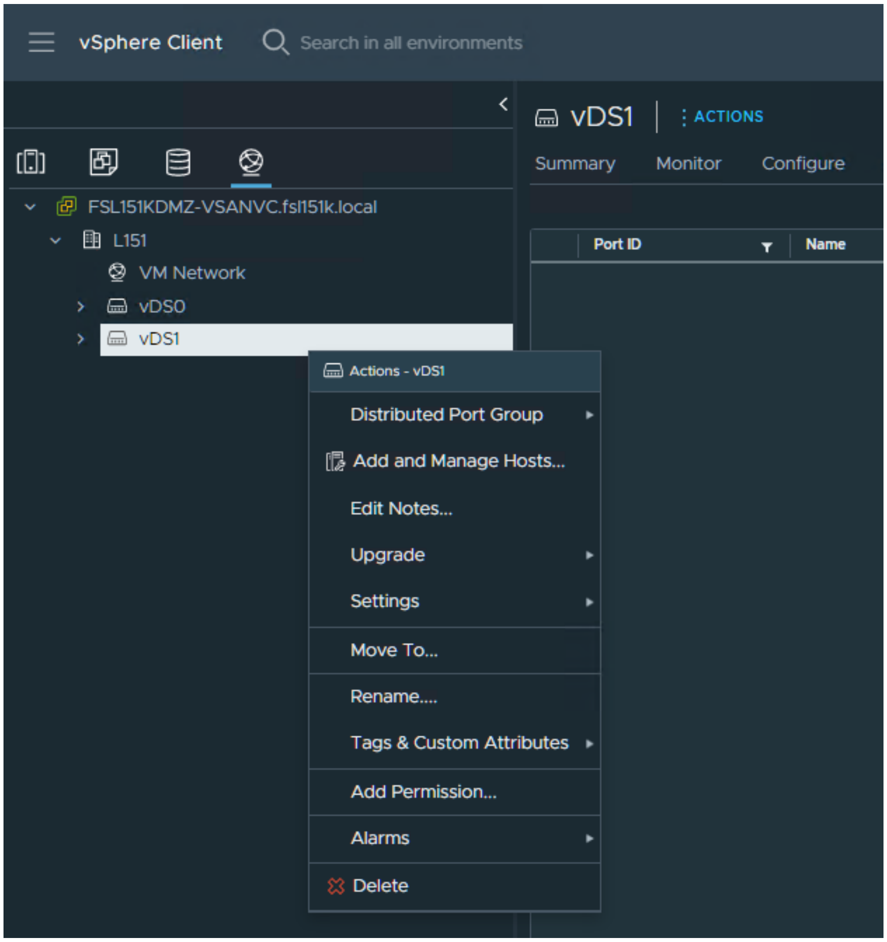

Step 9. Select the Add hosts button and click NEXT.

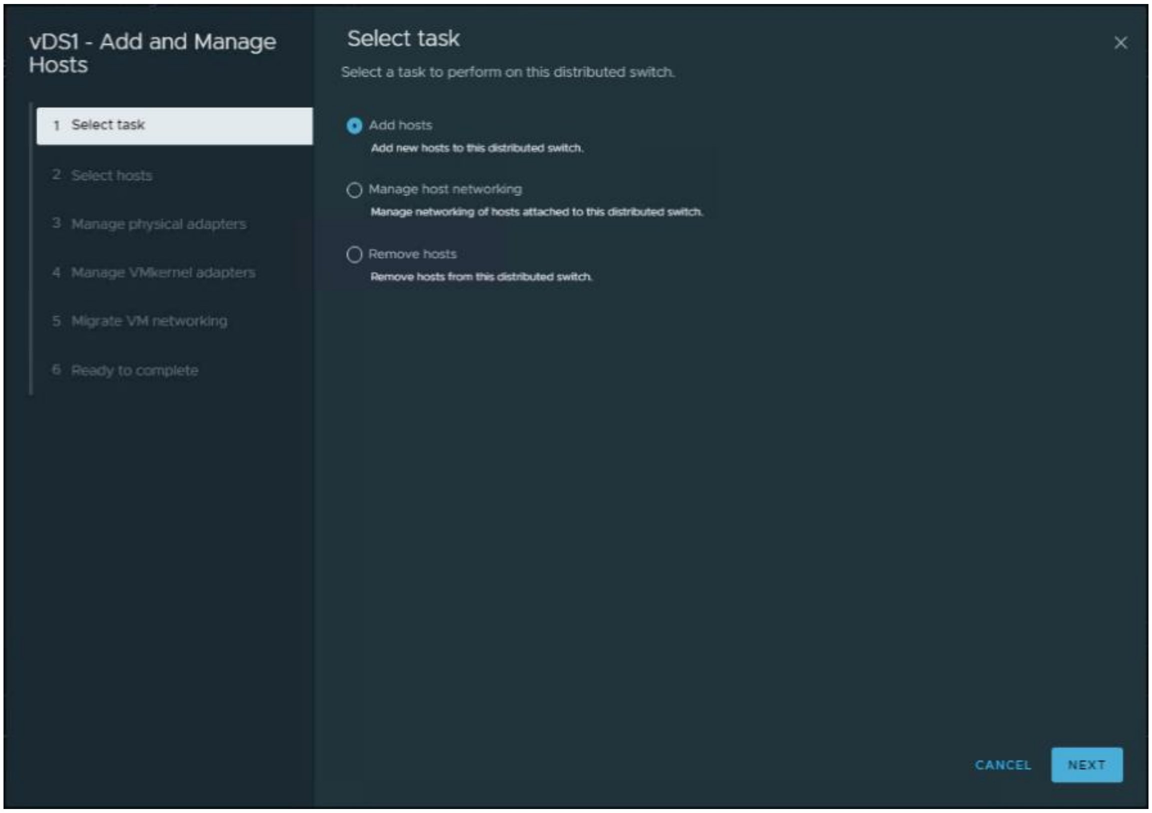
vDS1 - Add and Manage Hosts
1 Select task
2 Select hosts
3 Manage physical adapters
4 Manage VMkernel adapters
5 Migrate VM networking
6 Ready to complete
CANCEL
NEXT

Select task

Select a task to perform on this distributed switch.

● Add hosts
Add new hosts to this distributed switch.

○ Manage host networking
Manage networking of hosts attached to this distributed switch.

○ Remove hosts
Remove hosts from this distributed switch.

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.

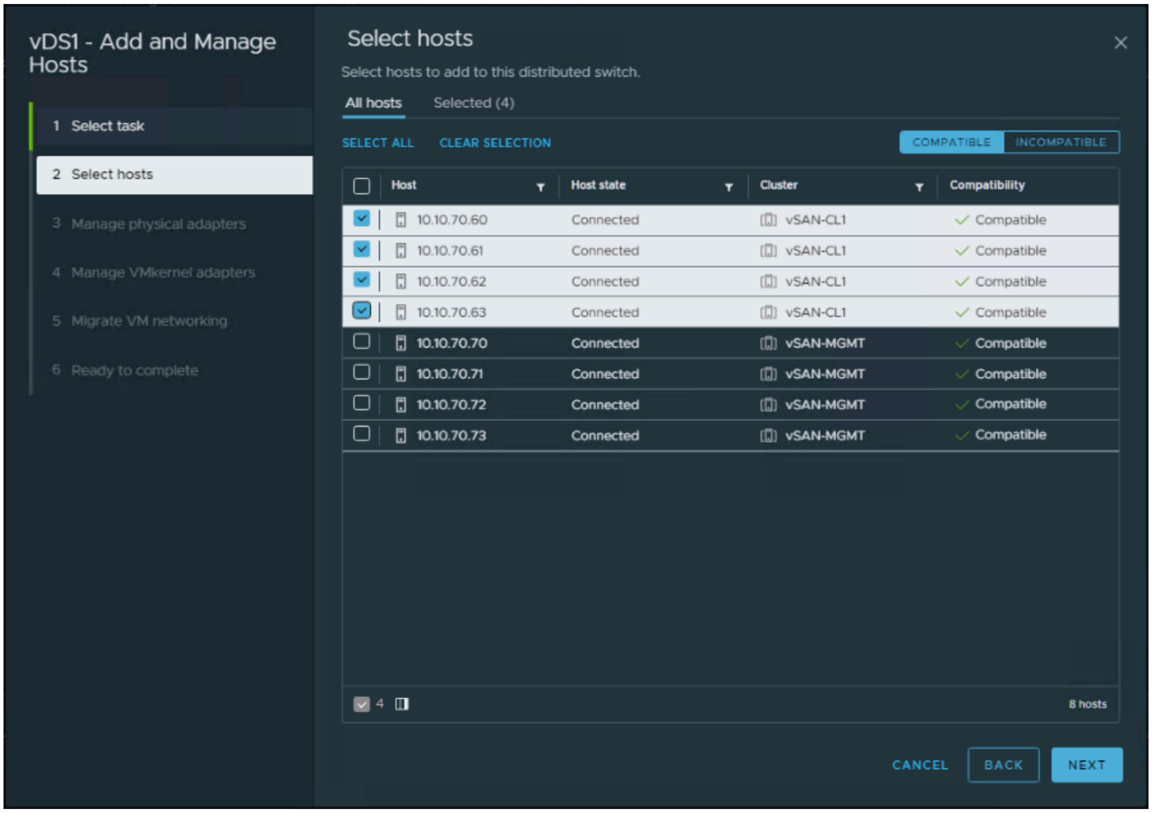

Select hosts

Select hosts to add to this distributed switch.

□	Host	Host state	Cluster	Compatibility
☑	10.10.70.60	Connected	vSAN-CL1	✓ Compatible
☑	10.10.70.61	Connected	vSAN-CL1	✓ Compatible
☑	10.10.70.62	Connected	vSAN-CL1	✓ Compatible
☑	10.10.70.63	Connected	vSAN-CL1	✓ Compatible
□	10.10.70.70	Connected	vSAN-MGMT	✓ Compatible
□	10.10.70.71	Connected	vSAN-MGMT	✓ Compatible
□	10.10.70.72	Connected	vSAN-MGMT	✓ Compatible
□	10.10.70.73	Connected	vSAN-MGMT	✓ Compatible

☑ 4 | 8 hosts

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.

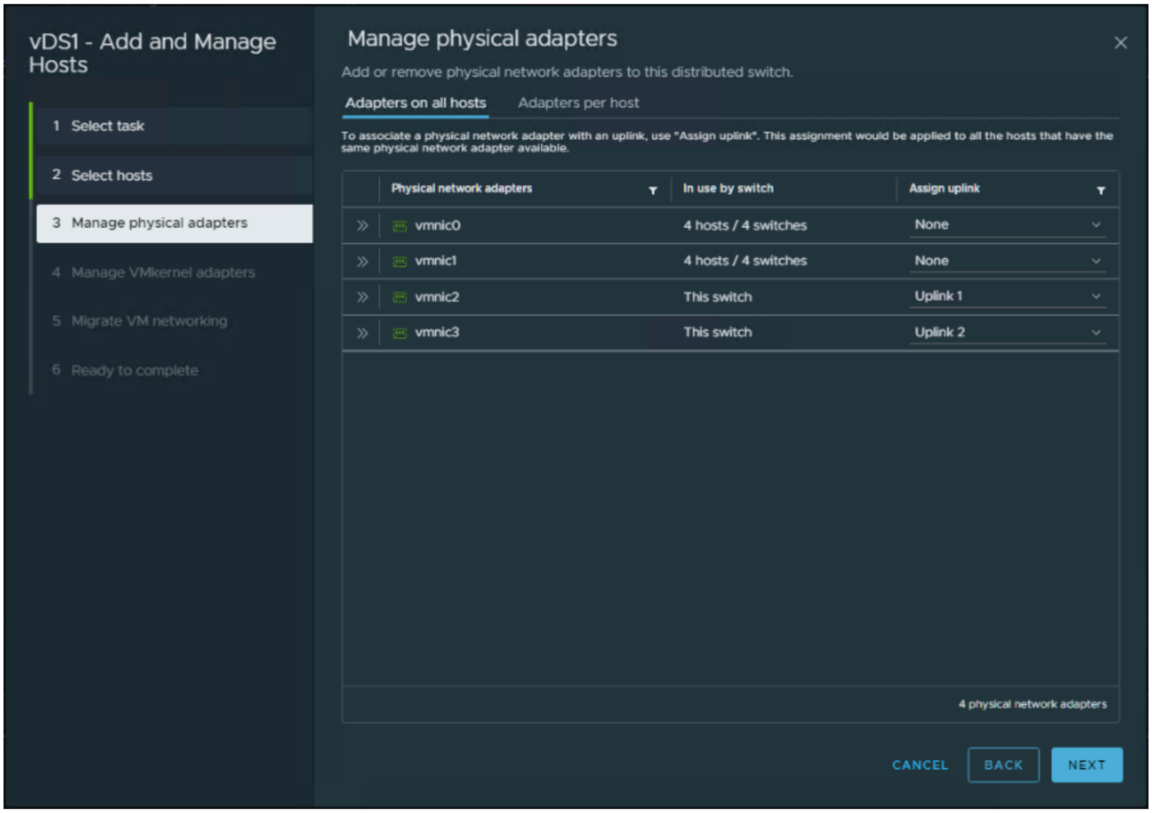

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

To associate a physical network adapter with an uplink, use "Assign uplink". This assignment would be applied to all the hosts that have the same physical network adapter available.

Physical network adapters	In use by switch	Assign uplink
vmnic0	4 hosts / 4 switches	None
vmnic1	4 hosts / 4 switches	None
vmnic2	This switch	Uplink 1
vmnic3	This switch	Uplink 2

4 physical network adapters

Step 12. In the next Manage VMkernel adapters and Migrate VM networking menus, click NEXT to continue.

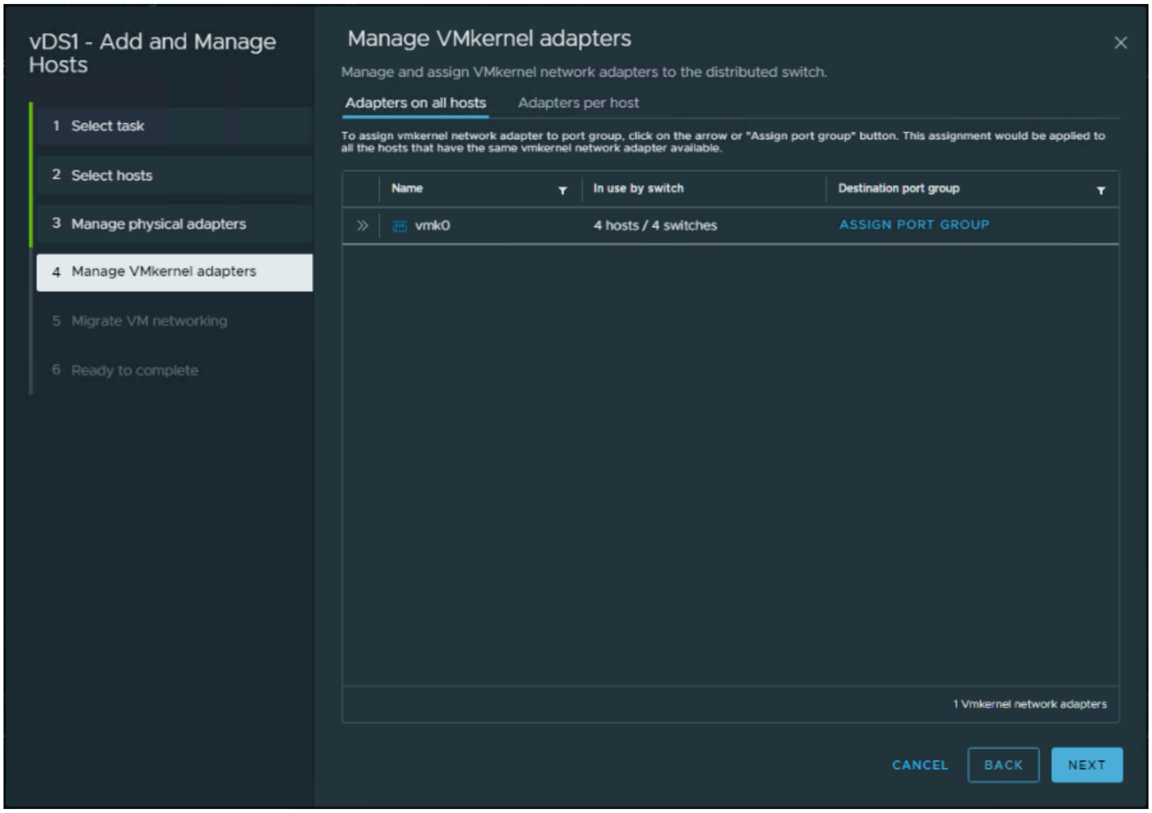

Manage VMkernel adapters

Manage and assign VMkernel network adapters to the distributed switch.

To assign vmkernel network adapter to port group, click on the arrow or "Assign port group" button. This assignment would be applied to all the hosts that have the same vmkernel network adapter available.

Name	In use by switch	Destination port group
vmk0	4 hosts / 4 switches	ASSIGN PORT GROUP

1 Vmkernel network adapters

Step 13. In the next Manage VMkernel adapters and Migrate VM networking menus, click NEXT to continue.

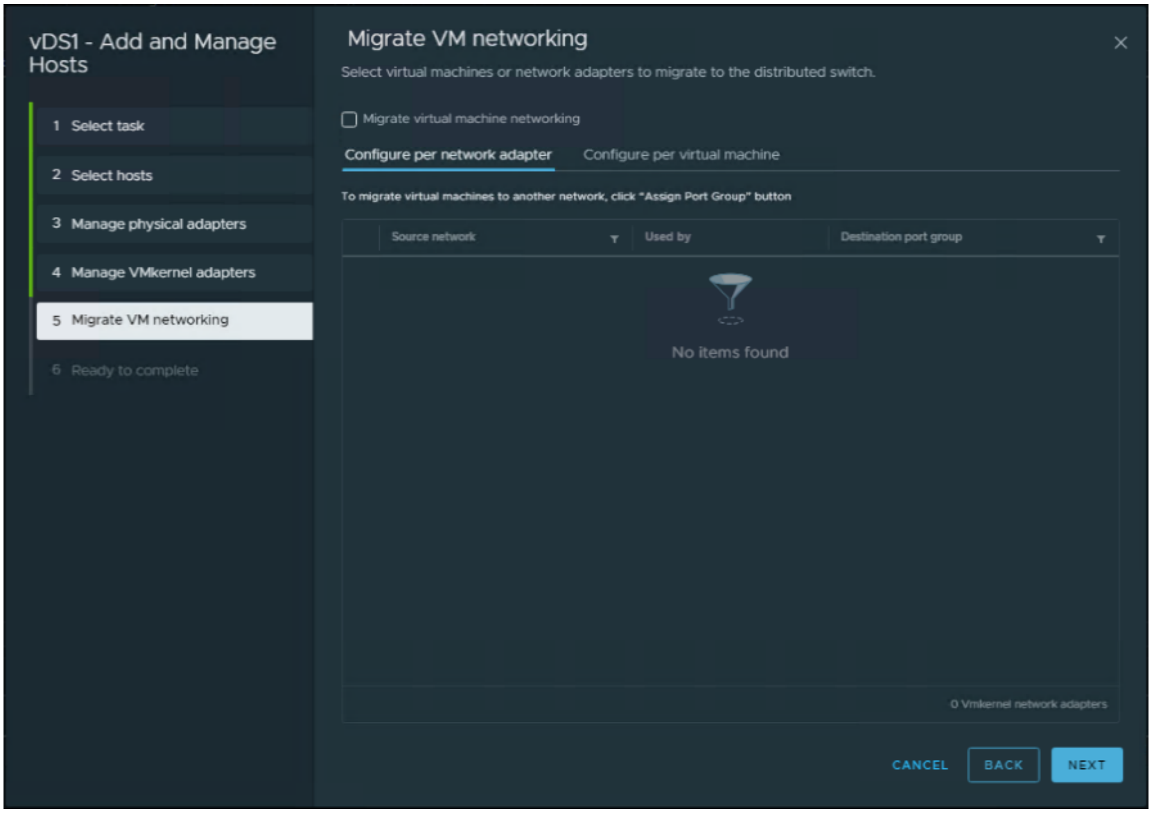
vDS1 - Add and Manage Hosts
1 Select task
2 Select hosts
3 Manage physical adapters
4 Manage VMkernel adapters
5 Migrate VM networking
6 Ready to complete
Configure per network adapter
Configure per virtual machine
CANCEL
BACK
NEXT

Migrate VM networking

Select virtual machines or network adapters to migrate to the distributed switch.

☐ Migrate virtual machine networking

To migrate virtual machines to another network, click "Assign Port Group" button

Source network	Used by	Destination port group

No items found

0 Vmkernel network adapters

Step 14. Click FINISH.

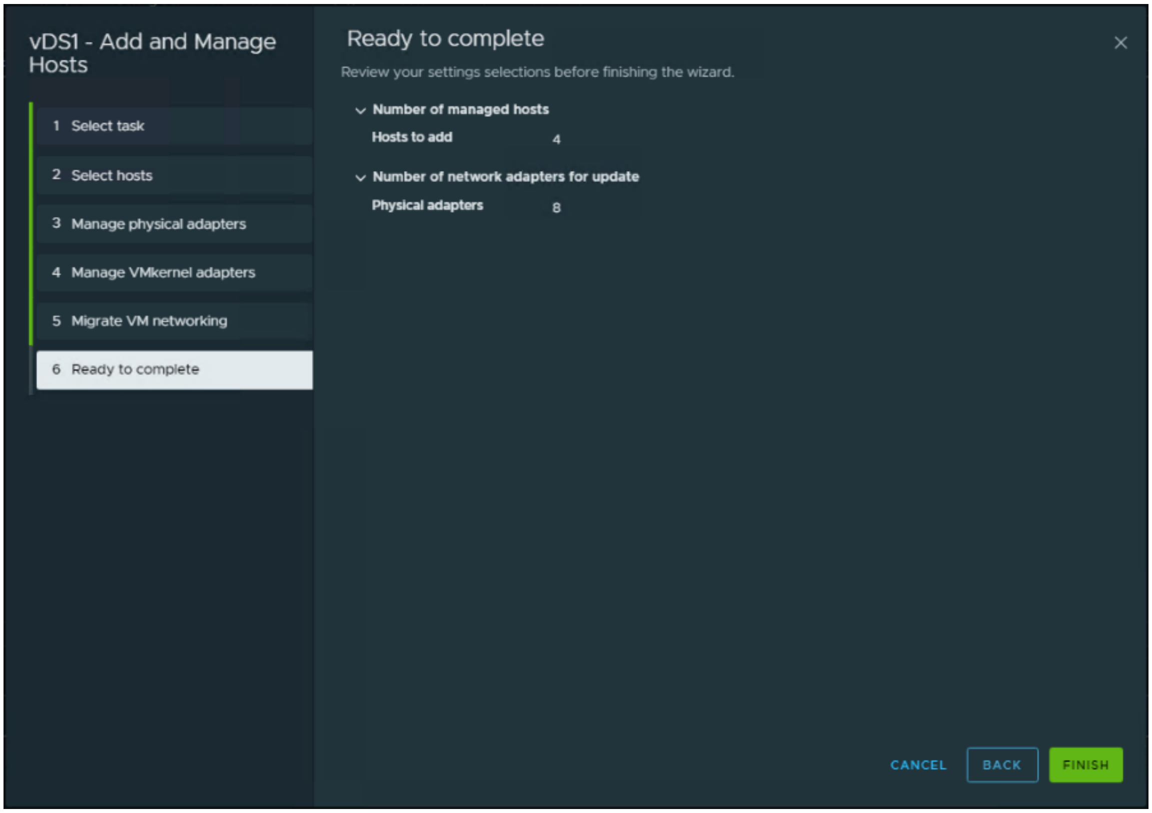

vDS1 - Add and Manage Hosts

Ready to complete

Review your settings selections before finishing the wizard.

Number of managed hosts

Hosts to add 4

Number of network adapters for update

Physical adapters 8

Procedure 4. Creating a Distributed Port Group for vMotion traffic

Step 1. Right-click Distributed switch and select Distributed Port Group click New Distributed Port Group.

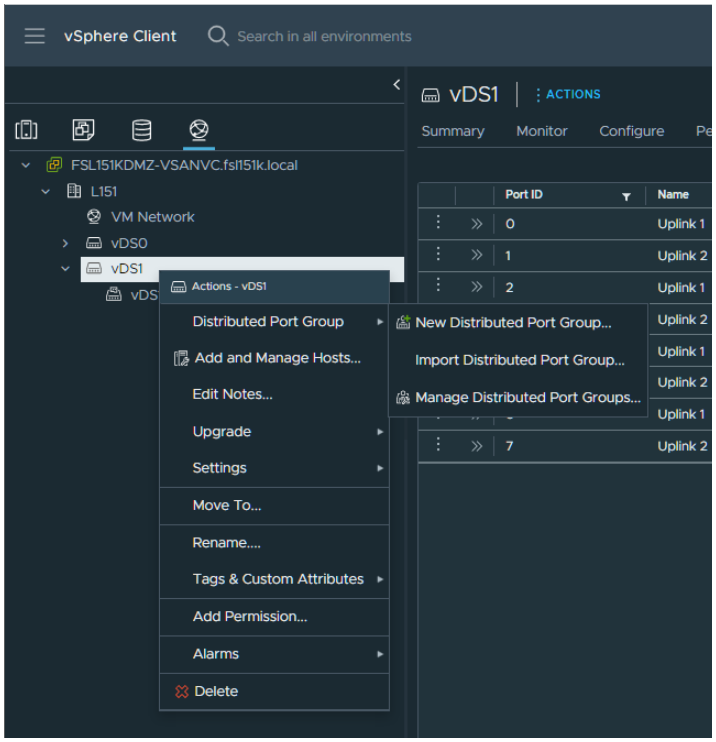

Port ID	Name
0	Uplink 1
1	Uplink 2
2	Uplink 1
	Uplink 2
	Uplink 1
	Uplink 2
	Uplink 1
7	Uplink 2

Step 2. On the New Distributed Port Group dialog box, enter a Name (for example vMotion), and click NEXT.

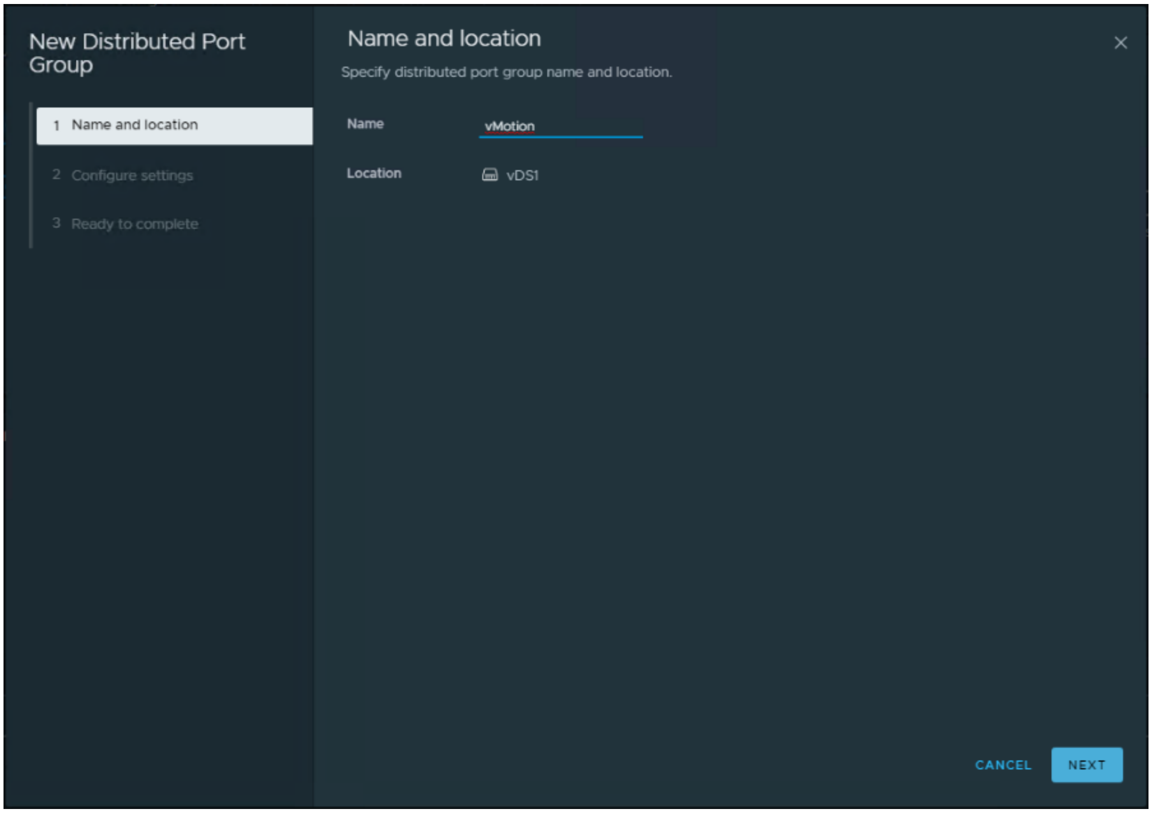

New Distributed Port Group

Name and location

Specify distributed port group name and location.

Name: vMotion

Location: vDS1

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.

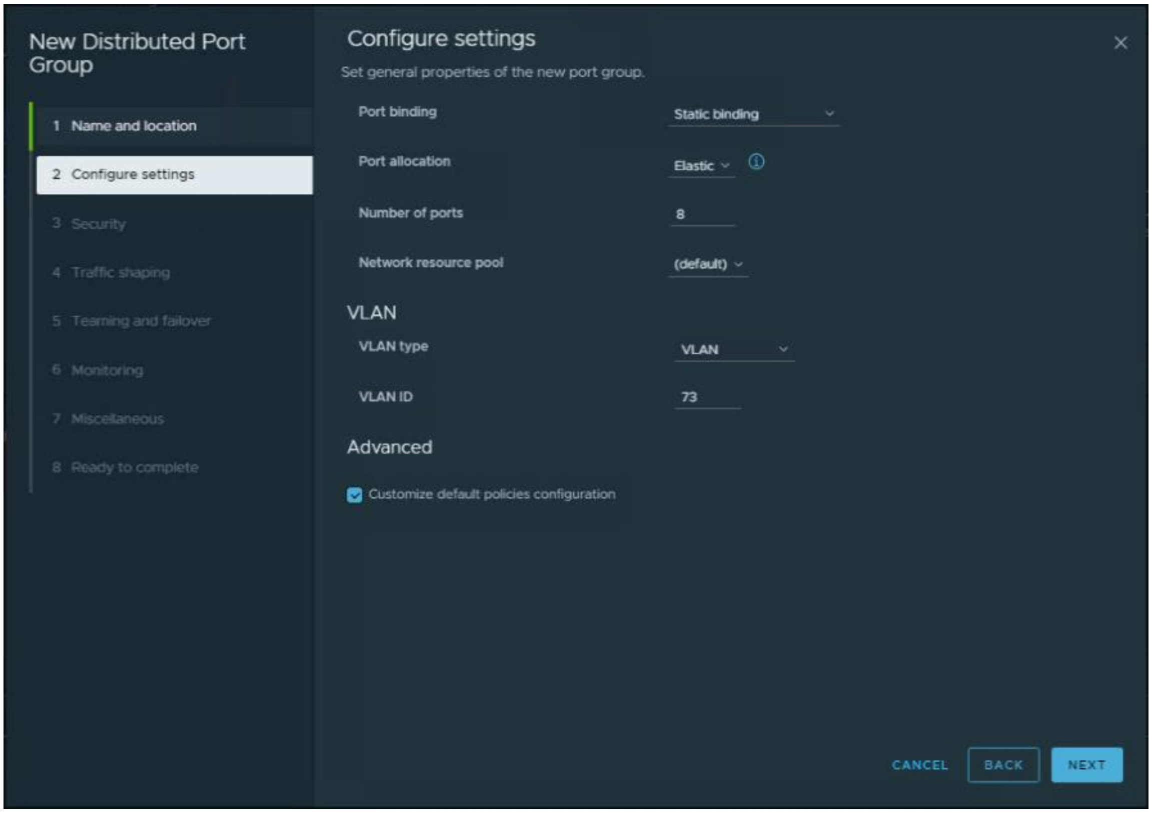

New Distributed Port Group

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Configure settings

Set general properties of the new port group.

Port binding: Static binding

Port allocation: Elastic

Number of ports: 8

Network resource pool: (default)

VLAN

VLAN type: VLAN

VLAN ID: 73

Advanced

☑ Customize default policies configuration

Step 4. On the Security dialog box, click NEXT.

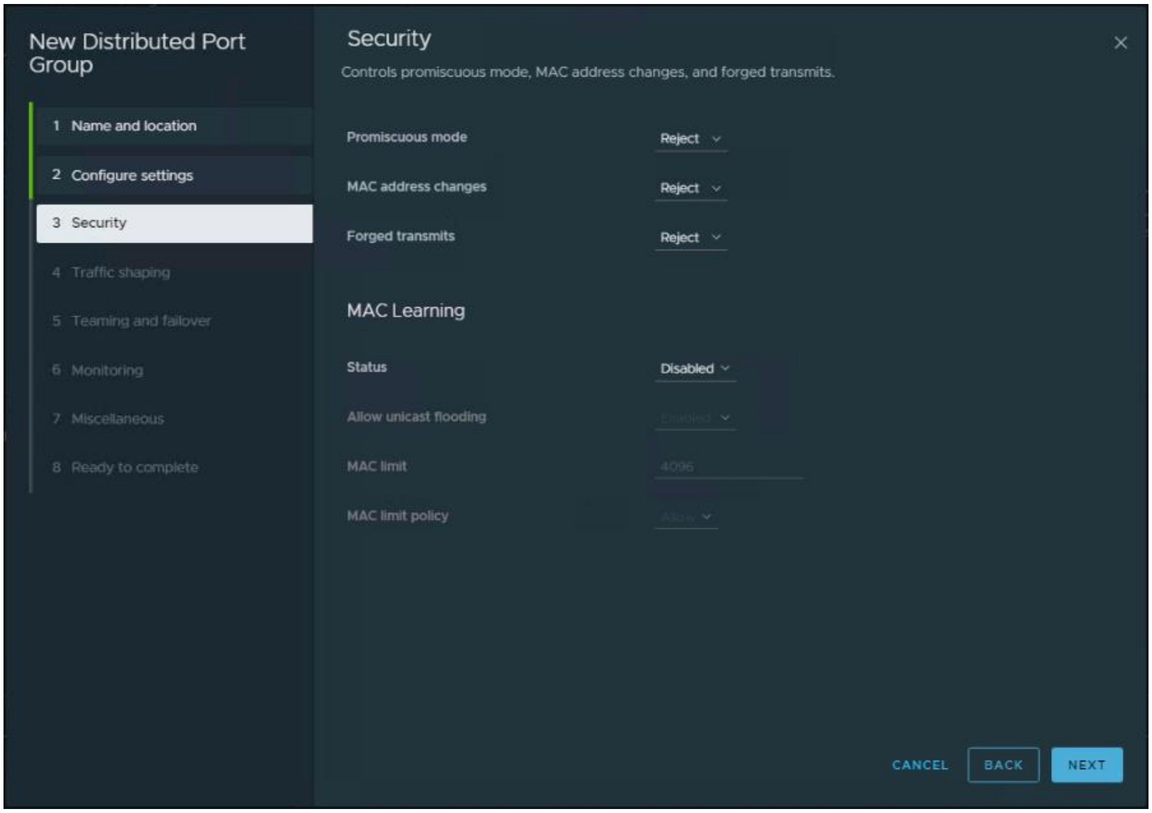

New Distributed Port Group

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Security

Controls promiscuous mode, MAC address changes, and forged transmits.

Promiscuous mode: Reject

MAC address changes: Reject

Forged transmits: Reject

MAC Learning

Status: Disabled

Allow unicast flooding: Enabled

MAC limit: 4096

MAC limit policy: Allow

Step 5. On the Traffic shaping dialog box, click NEXT.

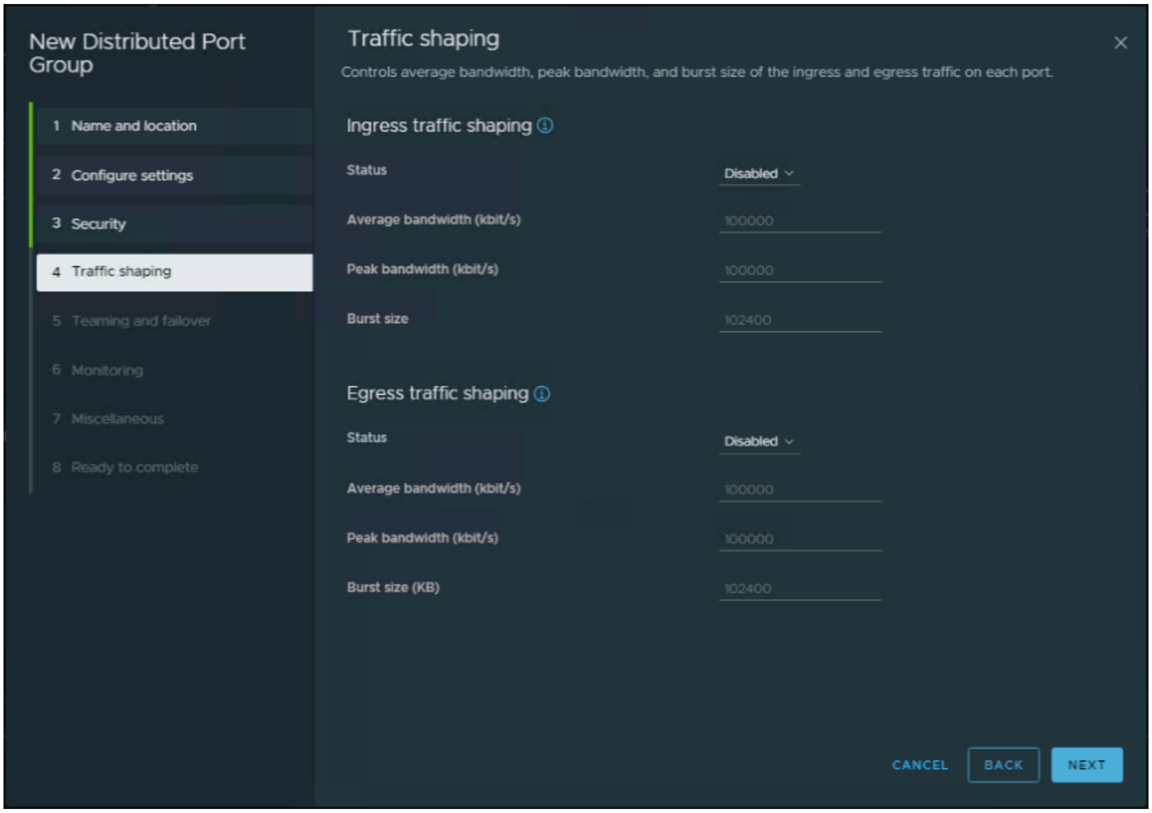

New Distributed Port Group

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Traffic shaping

Controls average bandwidth, peak bandwidth, and burst size of the ingress and egress traffic on each port.

Ingress traffic shaping

Status: Disabled
Average bandwidth (kbit/s): 100000
Peak bandwidth (kbit/s): 100000
Burst size: 102400

Egress traffic shaping

Status: Disabled
Average bandwidth (kbit/s): 100000
Peak bandwidth (kbit/s): 100000
Burst size (KB): 102400

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.

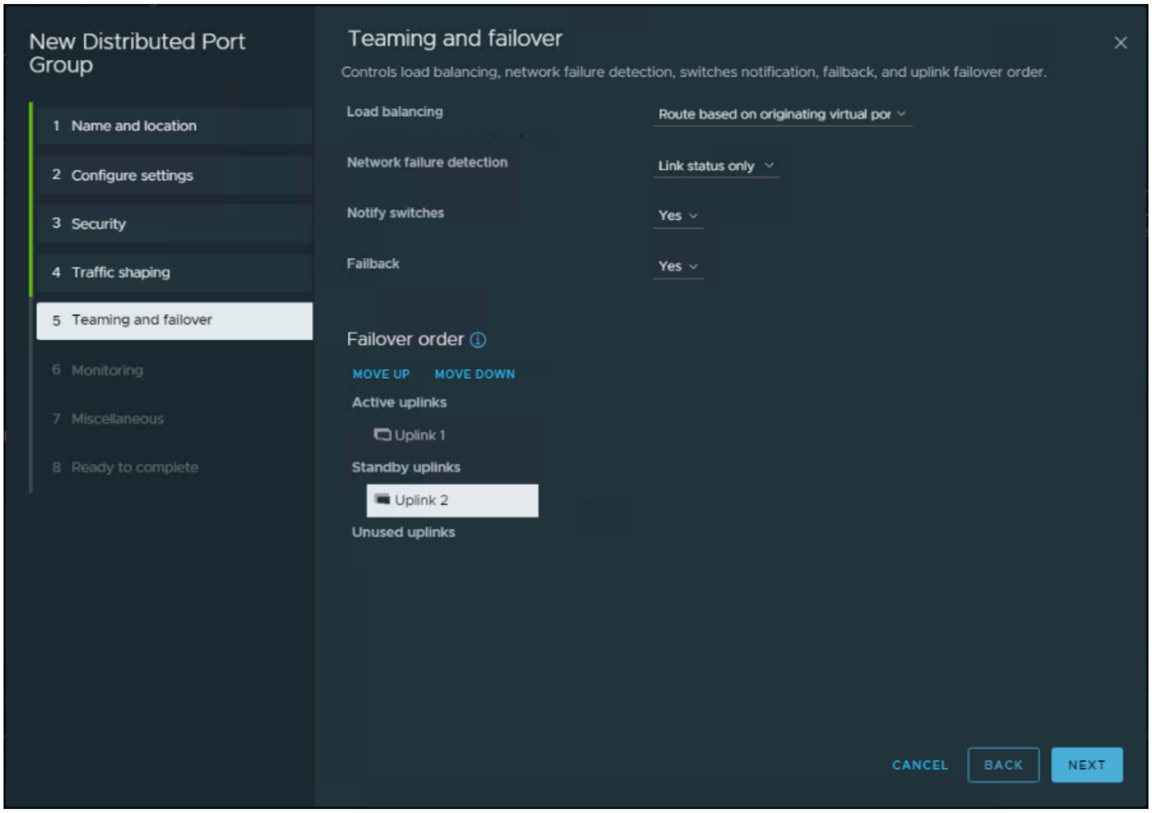

Teaming and failover

Controls load balancing, network failure detection, switches notification, failback, and uplink failover order.

Load balancing: Route based on originating virtual por

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active uplinks
- Uplink 1

Standby uplinks
- Uplink 2

Unused uplinks

Step 7. In the Monitoring dialog box, set NetFlow to Disabled, and click NEXT.

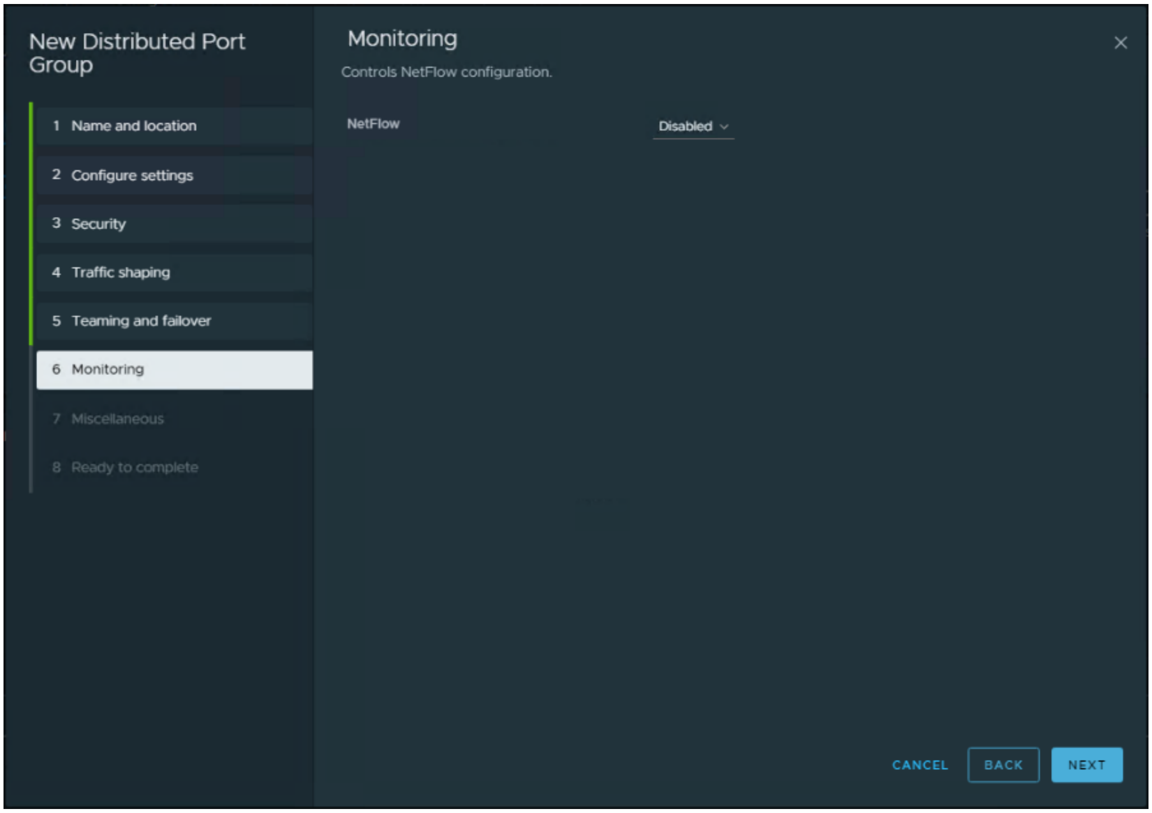

Monitoring

Controls NetFlow configuration.

NetFlow: Disabled

Step 8. In the Miscellaneous dialog box, set Block All Ports to No, and click NEXT.

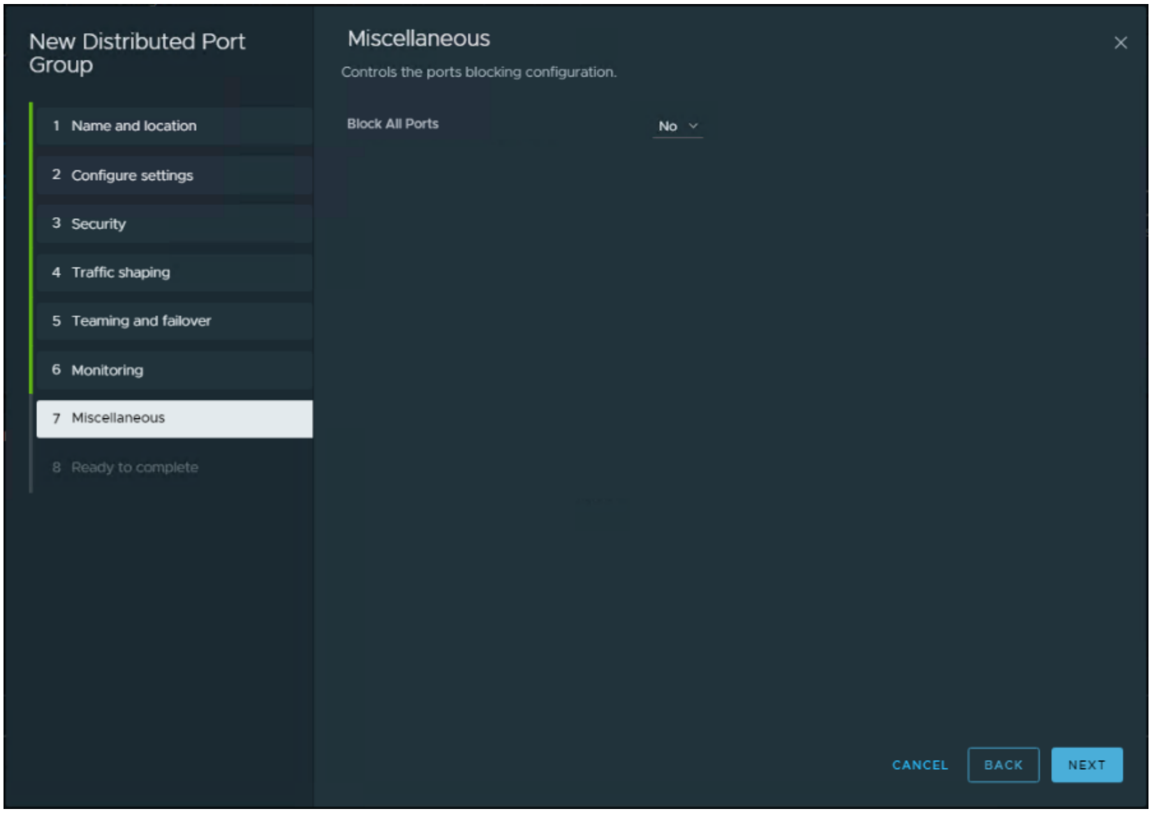

New Distributed Port Group

Miscellaneous

Controls the ports blocking configuration.

Block All Ports: No

Step 9. In the Ready to complete dialog box, review all the changes, and click FINISH.

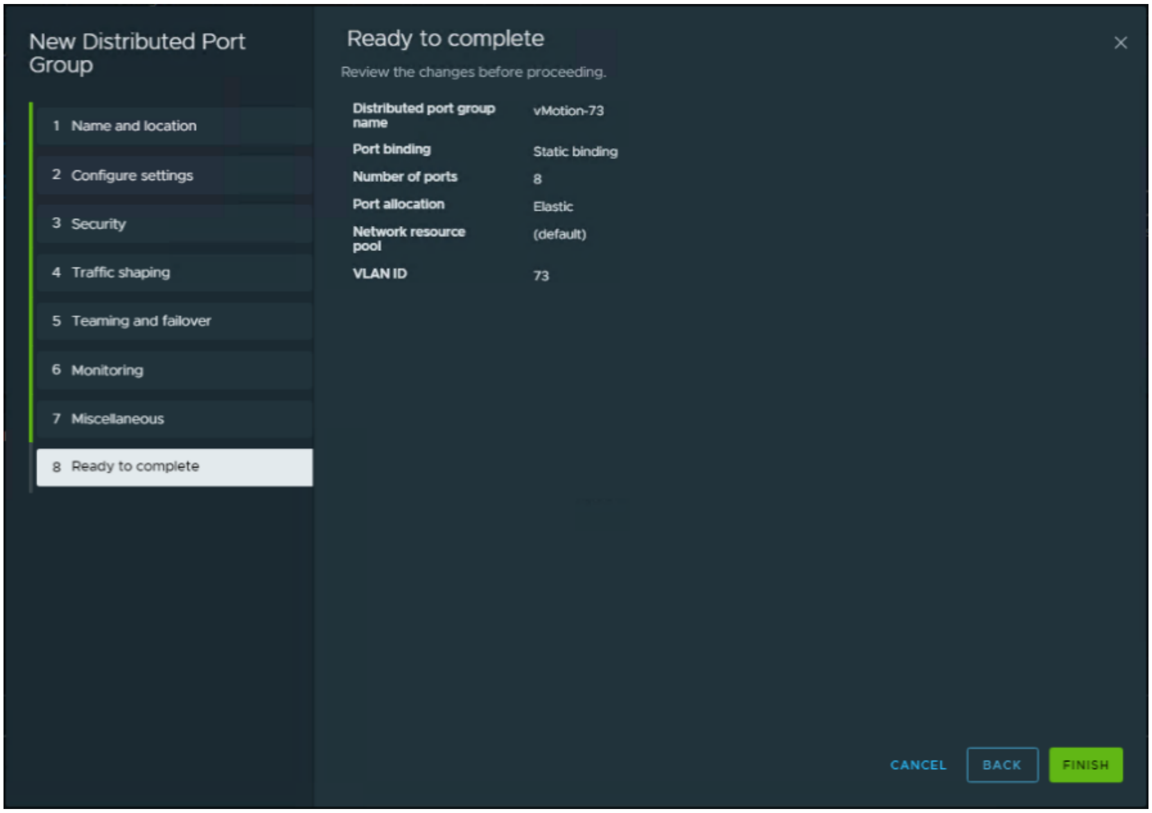
New Distributed Port Group

Ready to complete

Review the changes before proceeding.

Distributed port group name	vMotion-73
Port binding	Static binding
Number of ports	8
Port allocation	Elastic
Network resource pool	(default)
VLAN ID	73

Procedure 5. Creating a Distributed Port Group for Storage traffic

Step 1. Right-click Distributed switch and select Distributed Port Group click New Distributed Port Group.

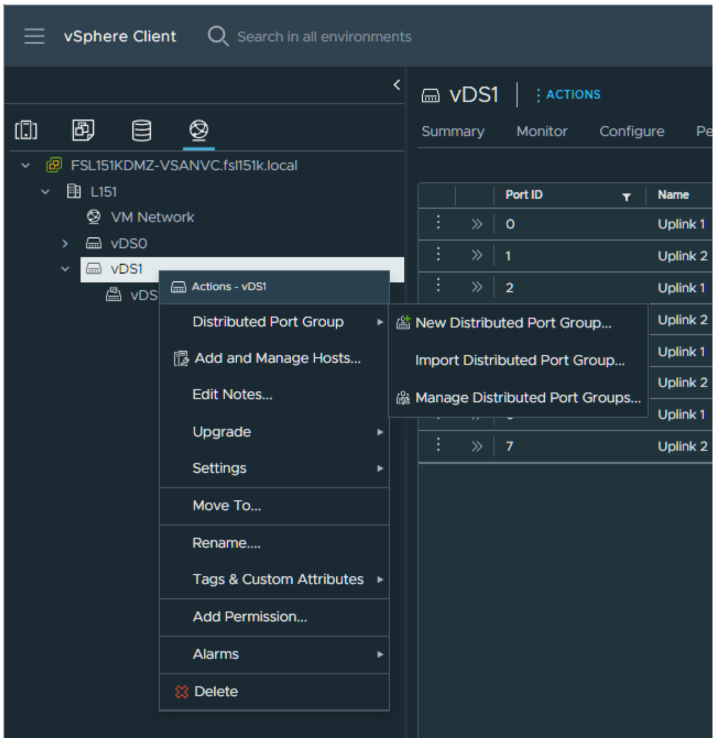

Step 2. On the New Distributed Port Group dialog box, enter a Name (for example vSAN), and click NEXT.

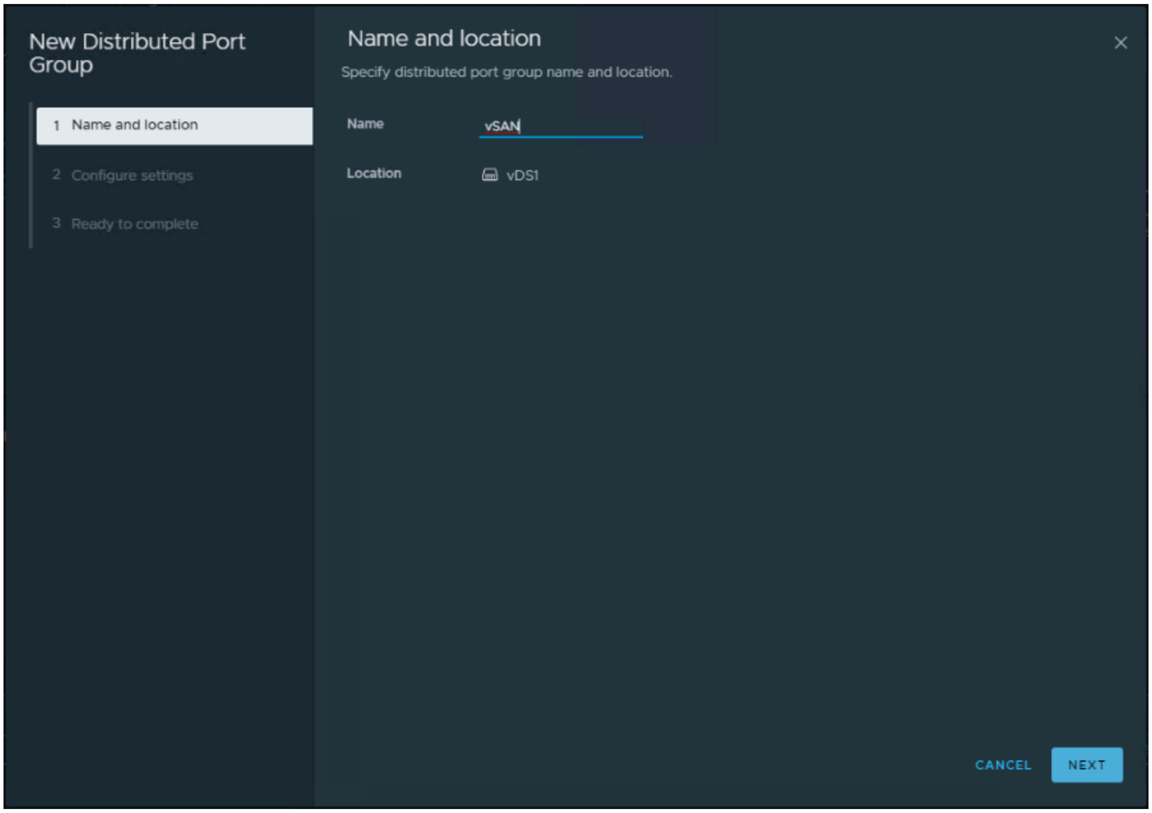

New Distributed Port Group

Name and location

Specify distributed port group name and location.

Name: vSAN

Location: vDS1

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.

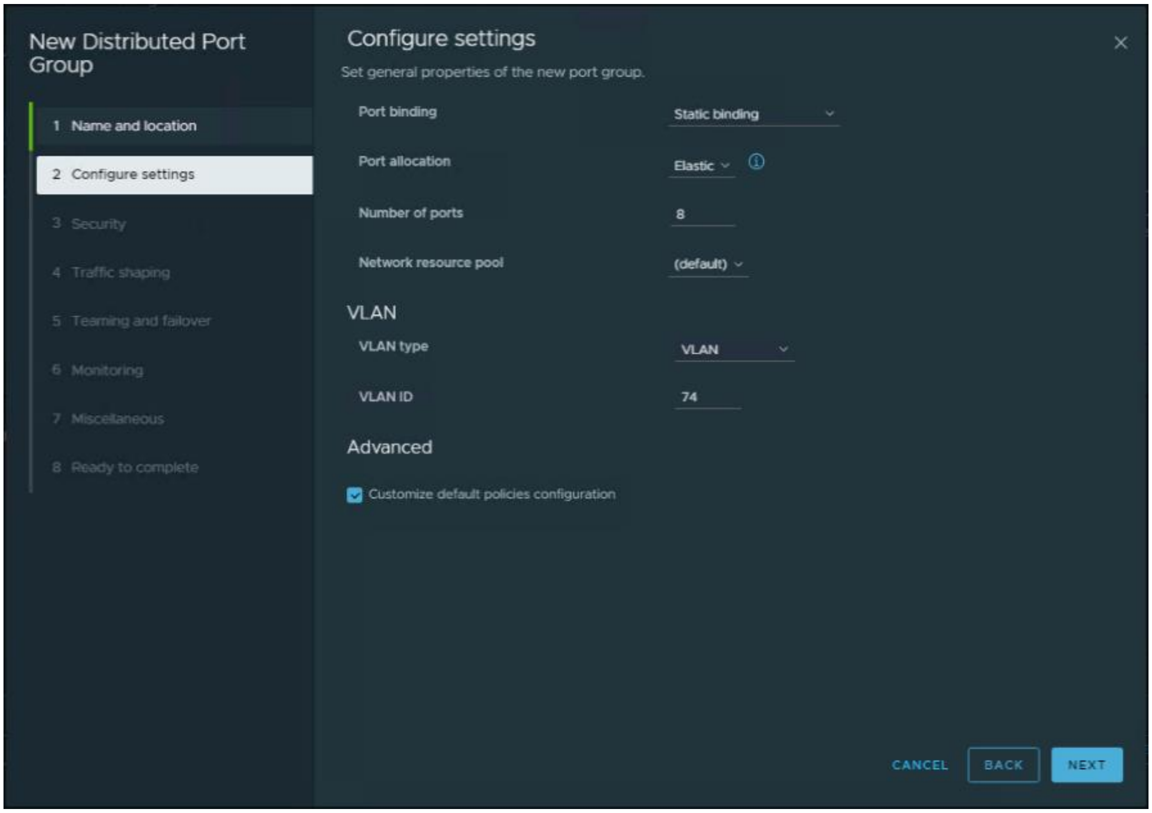

New Distributed Port Group

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Configure settings

Set general properties of the new port group.

Port binding: Static binding

Port allocation: Elastic

Number of ports: 8

Network resource pool: (default)

VLAN

VLAN type: VLAN

VLAN ID: 74

Advanced

☑ Customize default policies configuration

Step 4. On the Security dialog box, click NEXT.

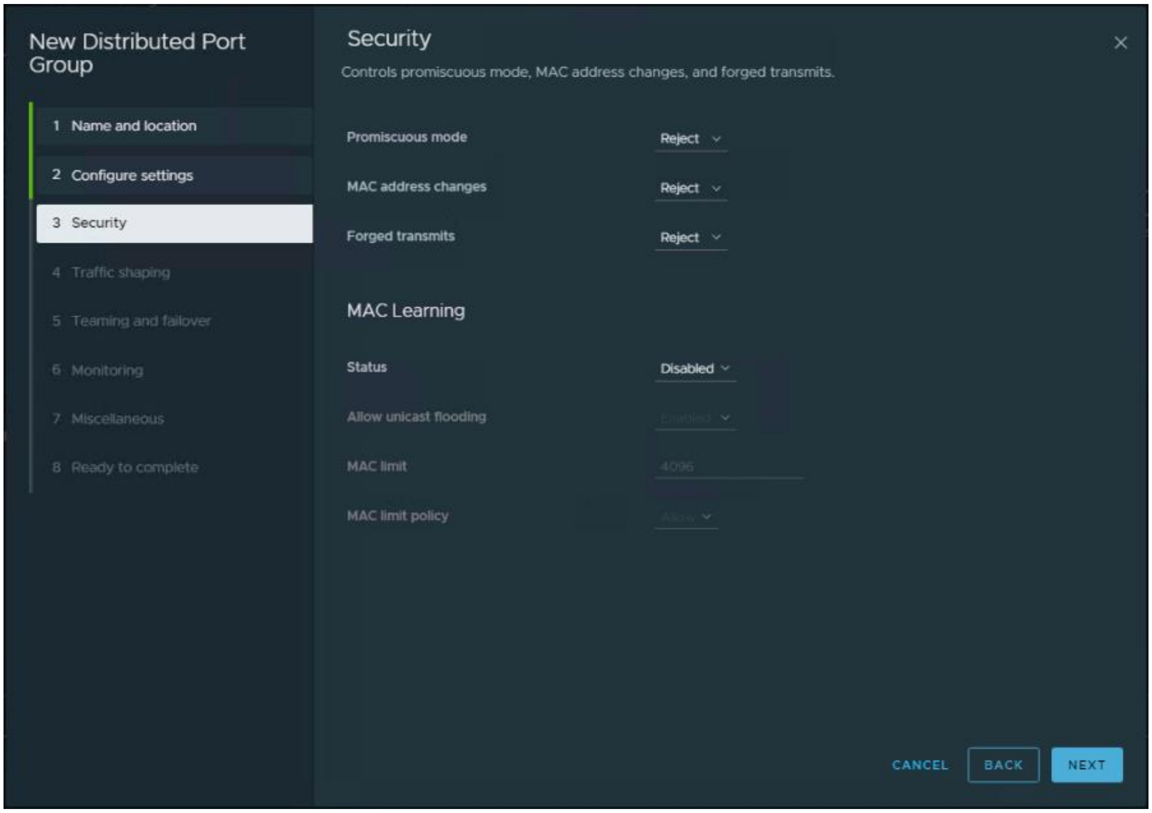

New Distributed Port Group

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Security

Controls promiscuous mode, MAC address changes, and forged transmits.

Promiscuous mode: Reject

MAC address changes: Reject

Forged transmits: Reject

MAC Learning

Status: Disabled

Allow unicast flooding: Enabled

MAC limit: 4096

MAC limit policy: Allow

Step 5. On the Traffic shaping dialog box, click NEXT.

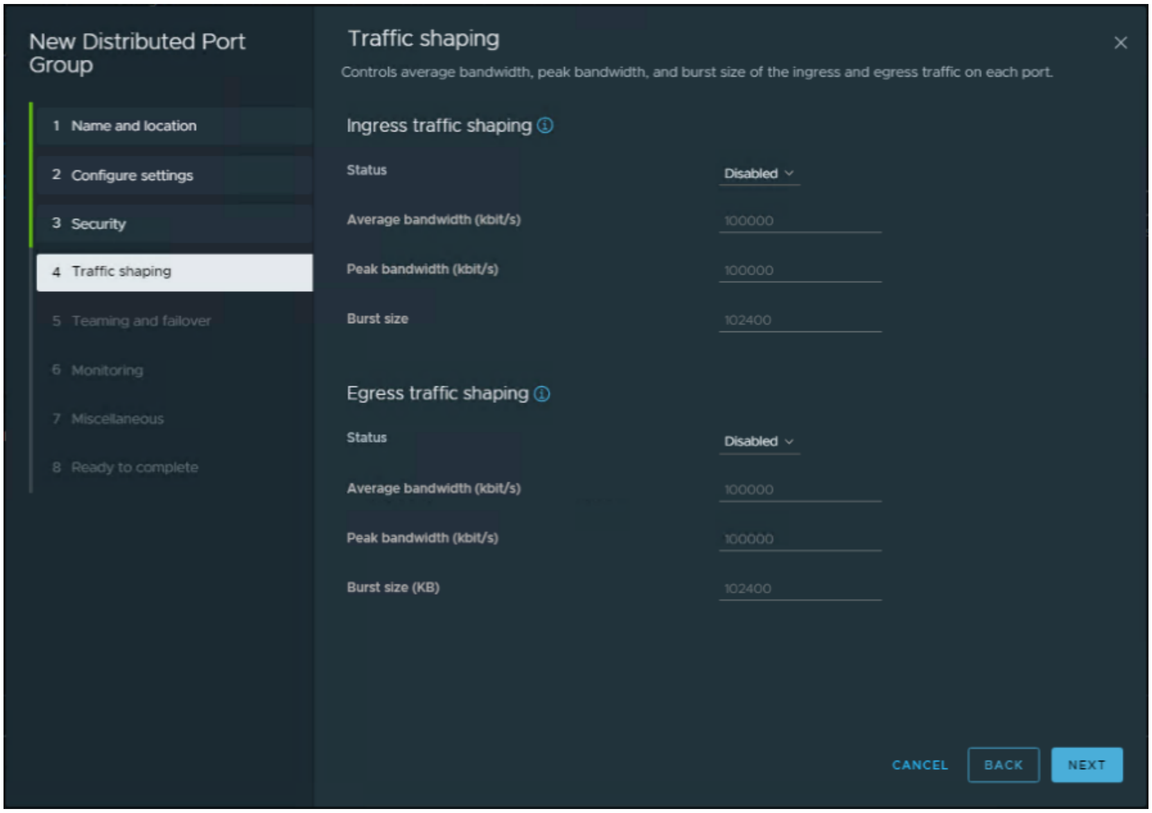

New Distributed Port Group

1 Name and location
2 Configure settings
3 Security
4 Traffic shaping
5 Teaming and failover
6 Monitoring
7 Miscellaneous
8 Ready to complete

Traffic shaping

Controls average bandwidth, peak bandwidth, and burst size of the ingress and egress traffic on each port.

Ingress traffic shaping

Status: Disabled
Average bandwidth (kbit/s): 100000
Peak bandwidth (kbit/s): 100000
Burst size: 102400

Egress traffic shaping

Status: Disabled
Average bandwidth (kbit/s): 100000
Peak bandwidth (kbit/s): 100000
Burst size (KB): 102400

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.

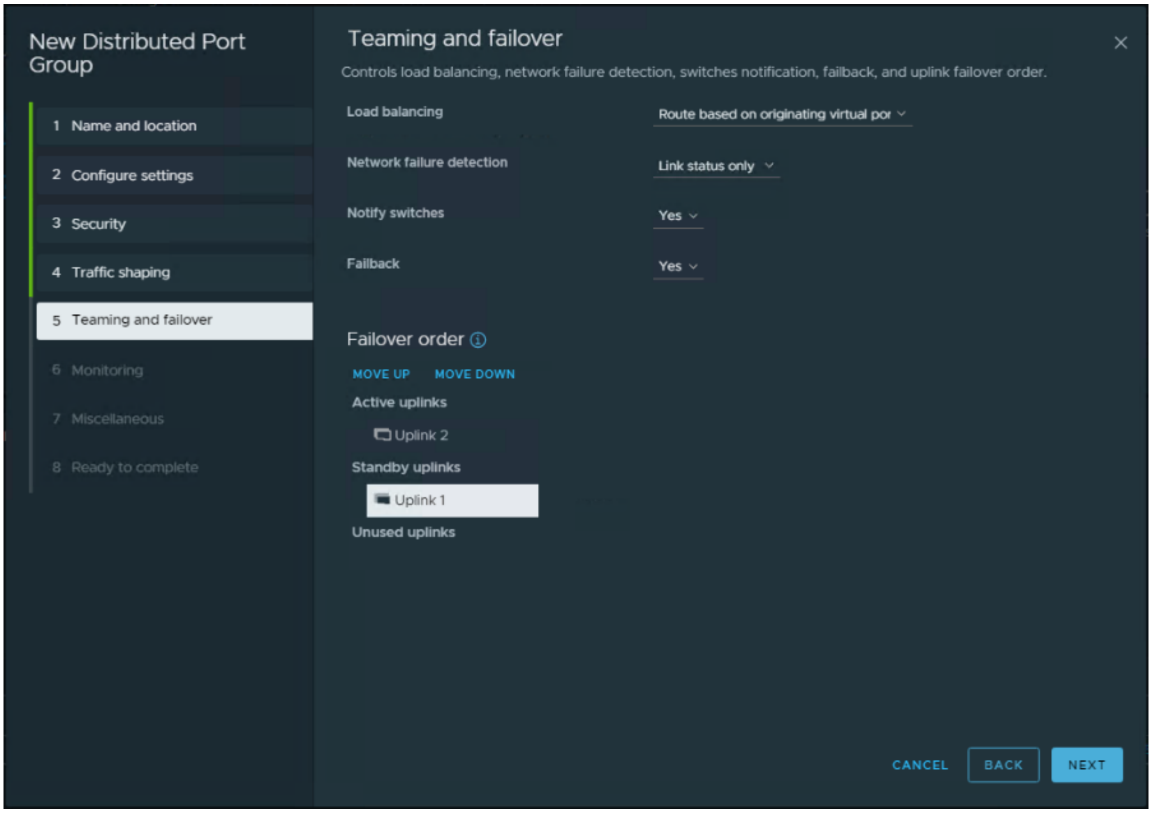

Teaming and failover

Controls load balancing, network failure detection, switches notification, failback, and uplink failover order.

Load balancing: Route based on originating virtual por

Network failure detection: Link status only

Notify switches: Yes

Failback: Yes

Failover order

Active uplinks
- Uplink 2

Standby uplinks
- Uplink 1

Unused uplinks

Step 7. In the Monitoring dialog box, set NetFlow to Disabled, and click NEXT.

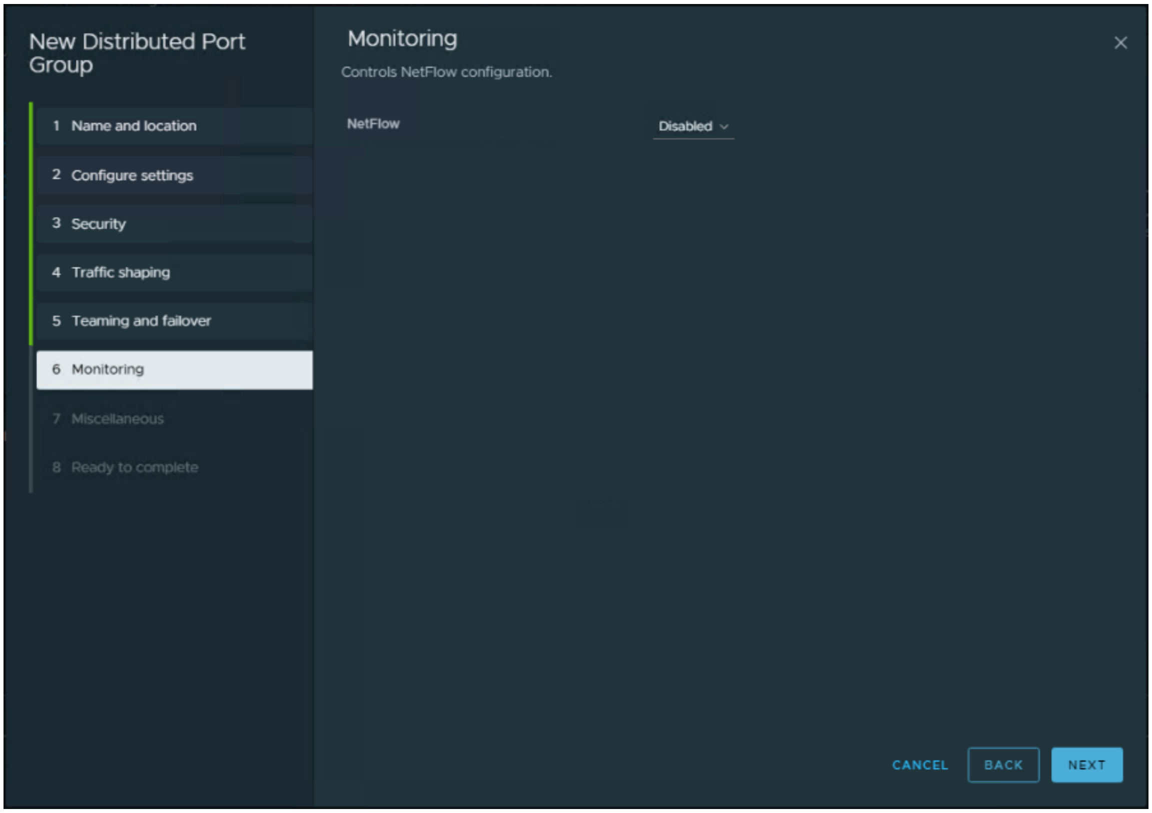

Monitoring

Controls NetFlow configuration.

NetFlow: Disabled

Step 8. In the Miscellaneous dialog box, set Block All Ports to No, and click NEXT.

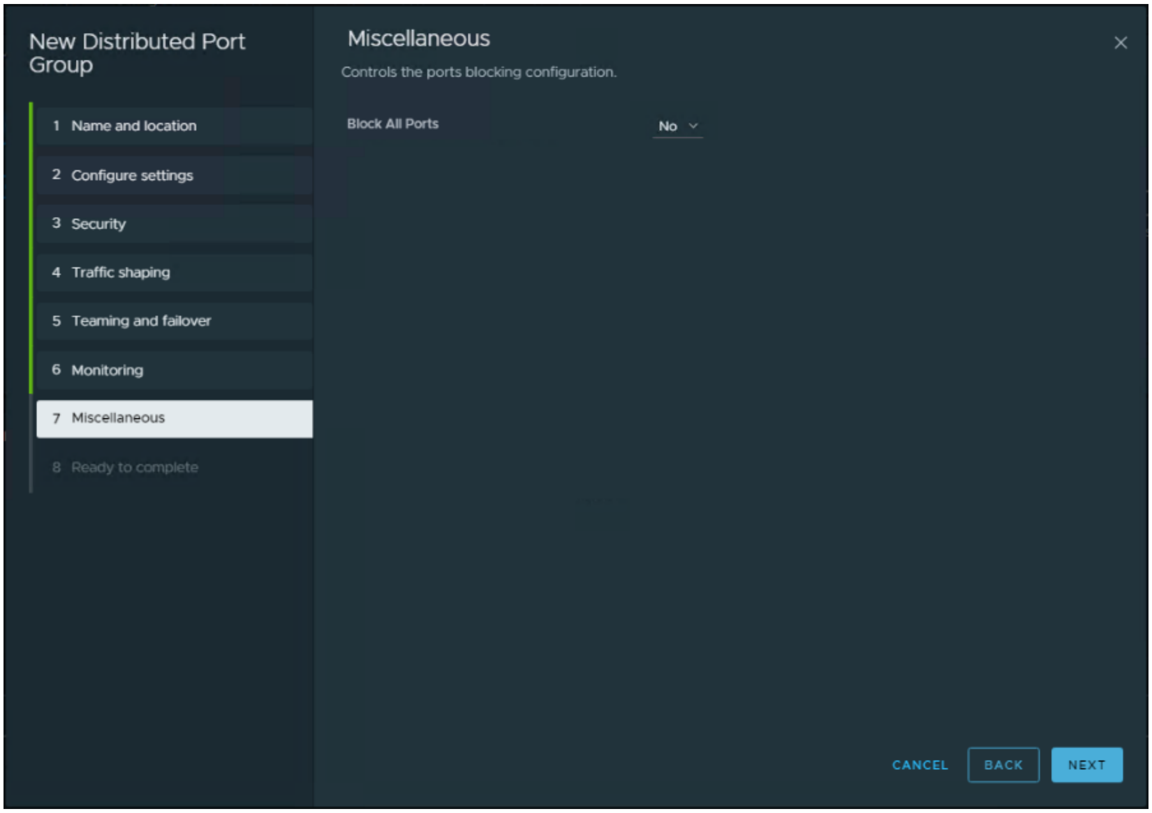

Miscellaneous

Controls the ports blocking configuration.

Block All Ports: No

Step 9. In the Ready to complete dialog box, review all the changes, and click FINISH.

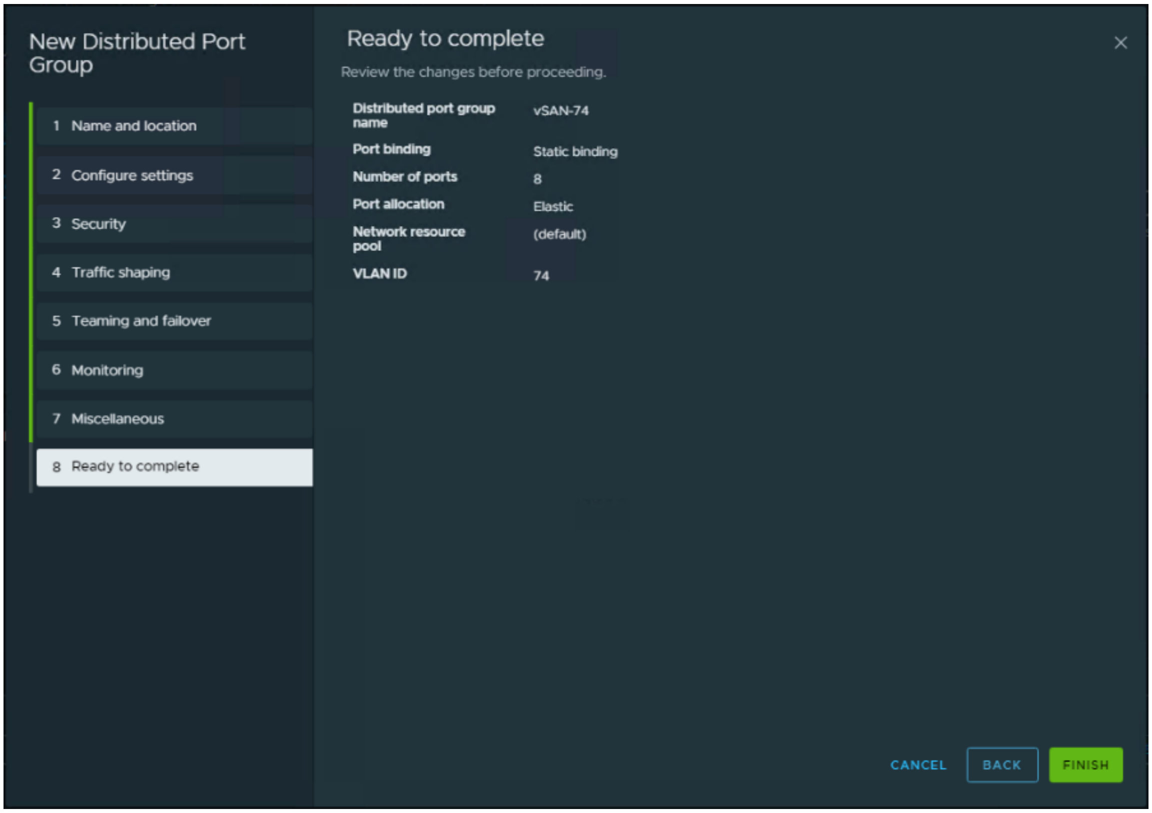

New Distributed Port Group

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	vSAN-74
Port binding	Static binding
Number of ports	8
Port allocation	Elastic
Network resource pool	(default)
VLAN ID	74

Procedure 6. Adding a VMkernel Adapters to distributed port groups.

Step 1. Right-click the distributed port group and select Add VMkernel Adapters....

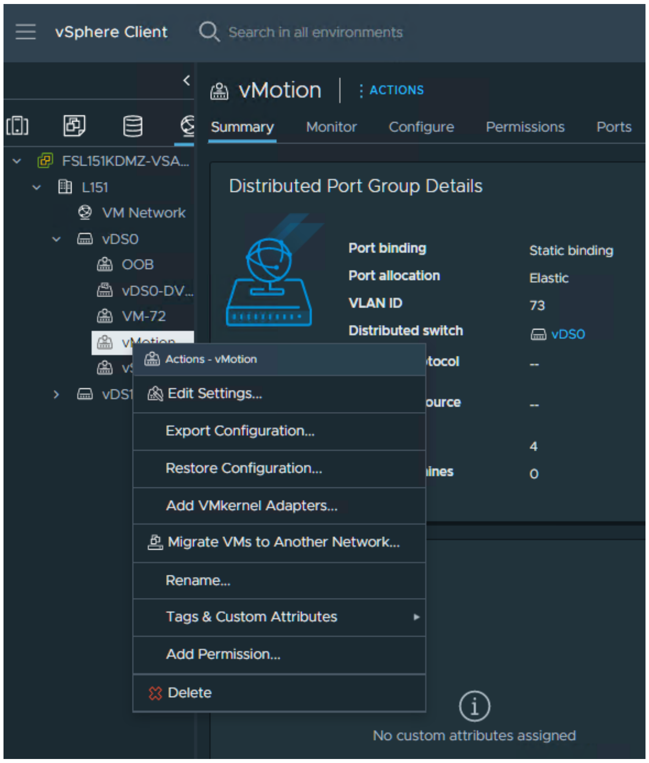

Distributed Port Group Details

Port binding	Static binding
Port allocation	Elastic
VLAN ID	73
Distributed switch	vDS0

Step 2. Select Attached Hosts and click NEXT.

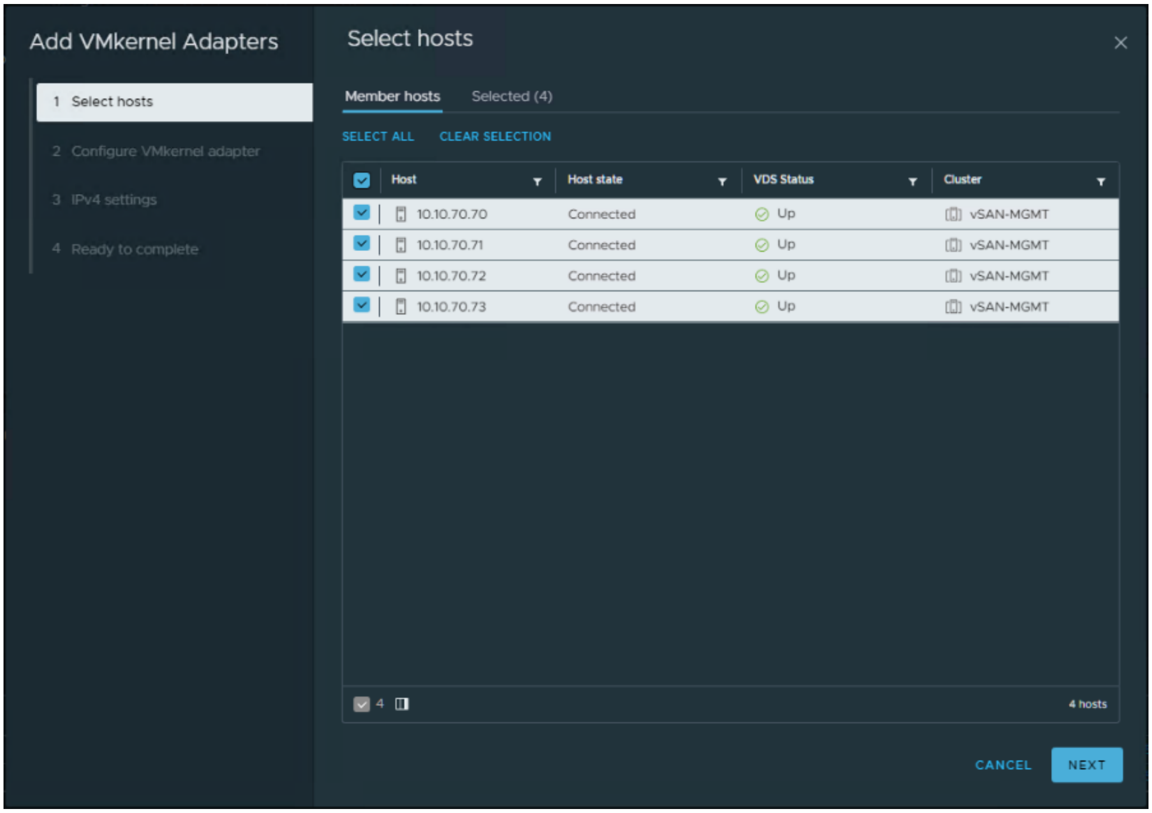

Add VMkernel Adapters

1 Select hosts

2 Configure VMkernel adapter

3 IPv4 settings

4 Ready to complete

Select hosts

☑	Host	Host state	VDS Status	Cluster
☑	10.10.70.70	Connected	✓ Up	vSAN-MGMT
☑	10.10.70.71	Connected	✓ Up	vSAN-MGMT
☑	10.10.70.72	Connected	✓ Up	vSAN-MGMT
☑	10.10.70.73	Connected	✓ Up	vSAN-MGMT

☑ 4

4 hosts

Step 3. Select service (for example vMotion) in Available services and click NEXT.

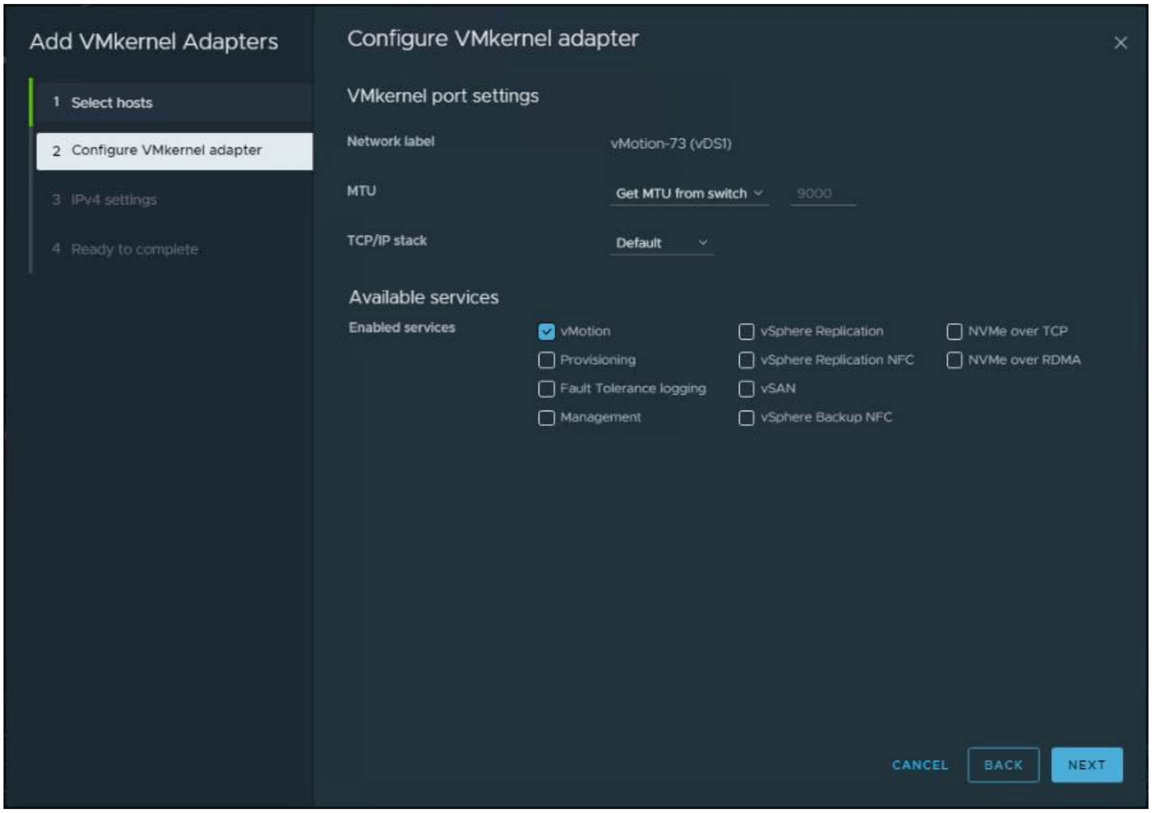

Configure VMkernel adapter

VMkernel port settings

Network label: vMotion-73 (vDS1)

MTU: Get MTU from switch 9000

TCP/IP stack: Default

Available services

Enabled services

☑ vMotion
□ Provisioning
□ Fault Tolerance logging
□ Management
□ vSphere Replication
□ vSphere Replication NFC
□ vSAN
□ vSphere Backup NFC
□ NVMe over TCP
□ NVMe over RDMA

Step 4. Enter the Network Settings and Gateway details and click NEXT.

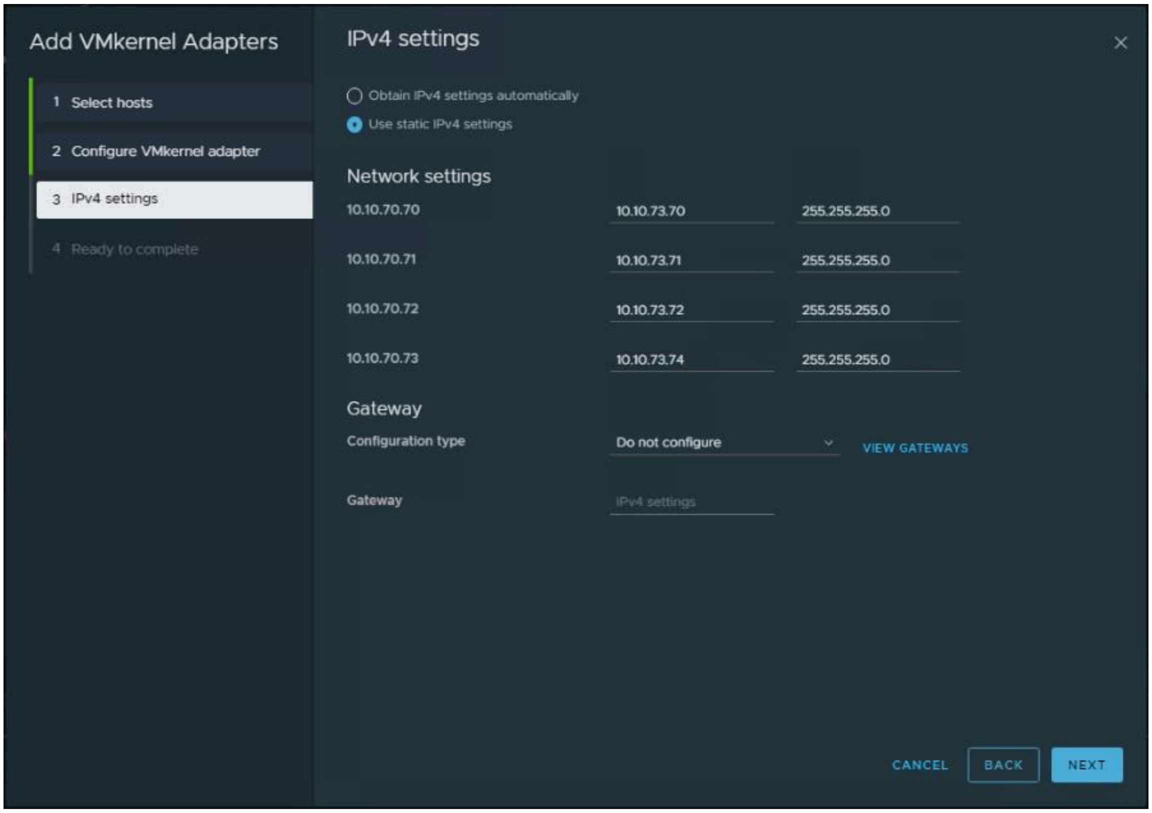

IPv4 settings

○ Obtain IPv4 settings automatically
● Use static IPv4 settings

Network settings

10.10.70.70	10.10.73.70	255.255.255.0
10.10.70.71	10.10.73.71	255.255.255.0
10.10.70.72	10.10.73.72	255.255.255.0
10.10.70.73	10.10.73.74	255.255.255.0

Gateway

Configuration type: Do not configure

Gateway: IPv4 settings

Step 5. Click FINISH.

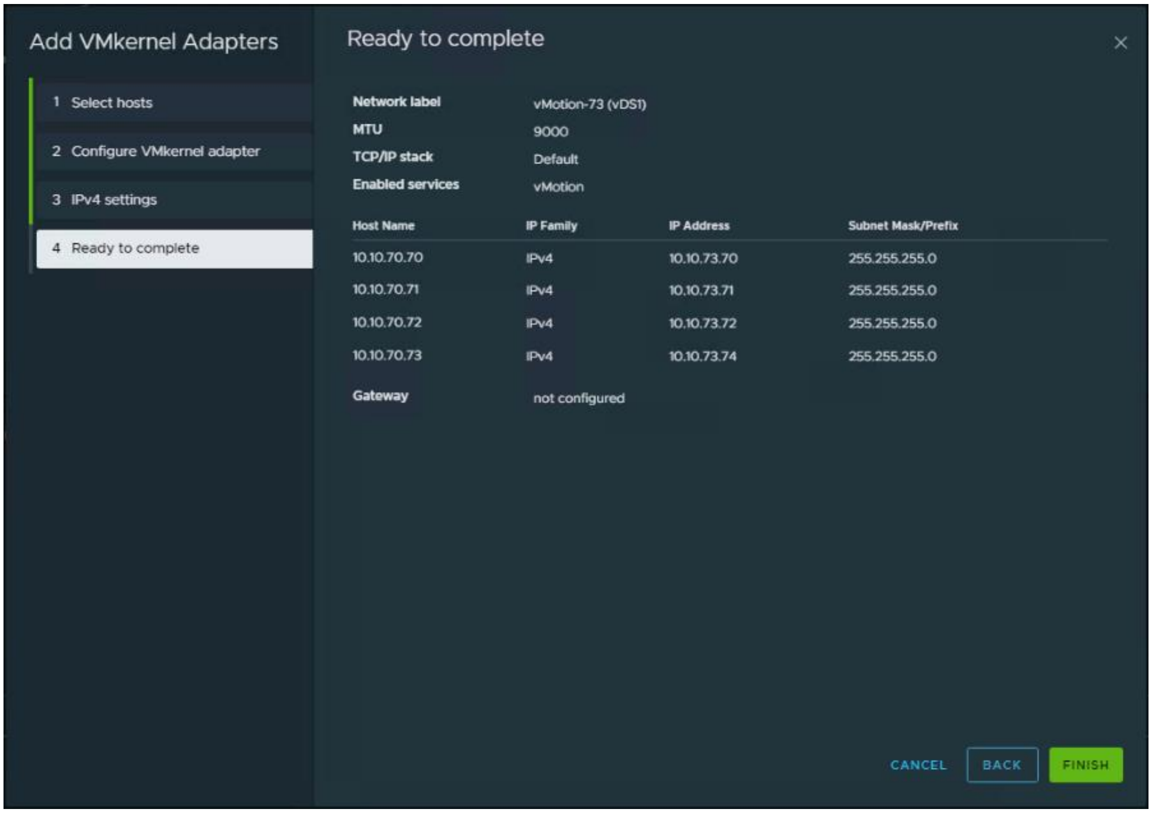

Ready to complete

Network label	vMotion-73 (vDS1)
MTU	9000
TCP/IP stack	Default
Enabled services	vMotion

Host Name	IP Family	IP Address	Subnet Mask/Prefix
10.10.70.70	IPv4	10.10.73.70	255.255.255.0
10.10.70.71	IPv4	10.10.73.71	255.255.255.0
10.10.70.72	IPv4	10.10.73.72	255.255.255.0
10.10.70.73	IPv4	10.10.73.74	255.255.255.0

Gateway not configured

Step 6. Repeat the Add VMkernel Adapters... steps for the vSAN traffic, and any other required port groups.

Note: Completed configuration can be verified under the Distributed Switch Topology tab**.**

Enable and Configure vSAN on the ESX Cluster

Procedure 1. Manually enable and configure vSAN OSA on the ESX Cluster

Note: Follow KB (https://kb.vmware.com/s/article/89485) to support vSAN higher cache tier capacities, up to 1.6TB.

Step 1. Navigate to an existing host cluster.

Step 2. On Configure tab select vSAN and then Services. Select Standard vSAN cluster radio button and click Configure.

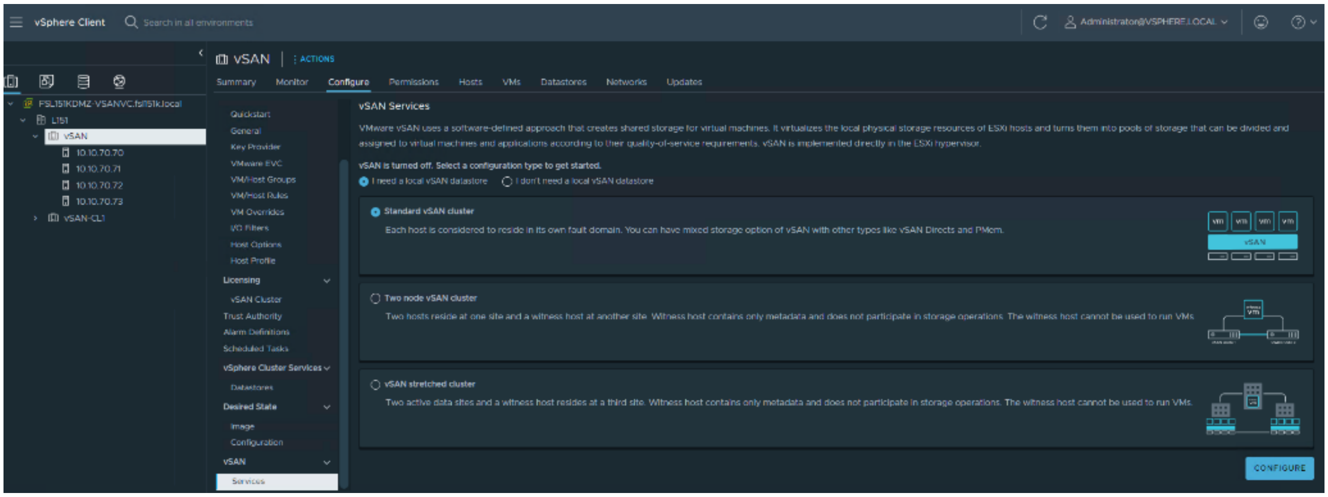

vSAN Services

VMware vSAN uses a software-defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi hosts and turns them into pools of storage that can be divided and assigned to virtual machines and applications according to their quality-of-service requirements. vSAN is implemented directly in the ESXi hypervisor.

vSAN is turned off. Select a configuration type to get started.

● I need a local vSAN datastore ○ I don't need a local vSAN datastore

● Standard vSAN cluster

Each host is considered to reside in its own fault domain. You can have mixed storage option of vSAN with other types like vSAN Directs and PMem.

○ Two node vSAN cluster

Two hosts reside at one site and a witness host at another site. Witness host contains only metadata and does not participate in storage operations. The witness host cannot be used to run VMs.

○ vSAN stretched cluster

Two active data sites and a witness host resides at a third site. Witness host contains only metadata and does not participate in storage operations. The witness host cannot be used to run VMs.

Step 3. Click Next.

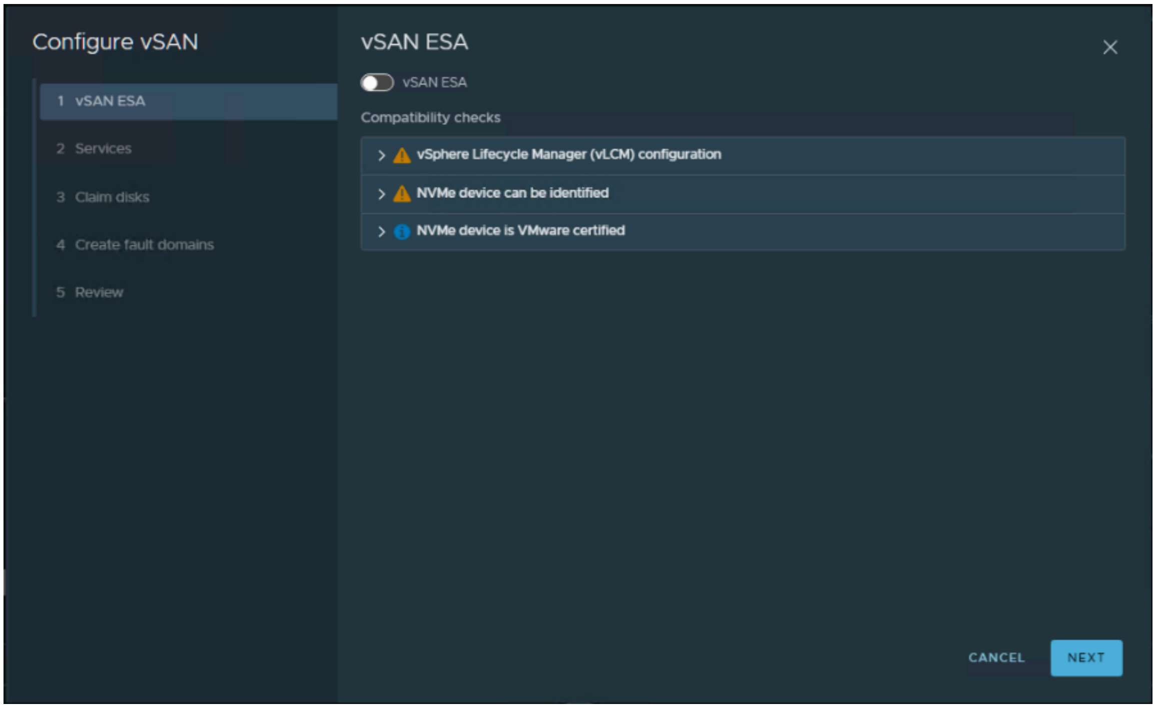

vSAN ESA

vSAN ESA

Compatibility checks

- vSphere Lifecycle Manager (vLCM) configuration
- NVMe device can be identified
- NVMe device is VMware certified

Step 4. Configure the data management features, including Data-At-Rest encryption, Data-In-Transit encryption and RDMA support (for example Deduplication and Compression). Click NEXT.

Note: Deduplication and compression was enabled to support the full clone desktop deployment at full scale.

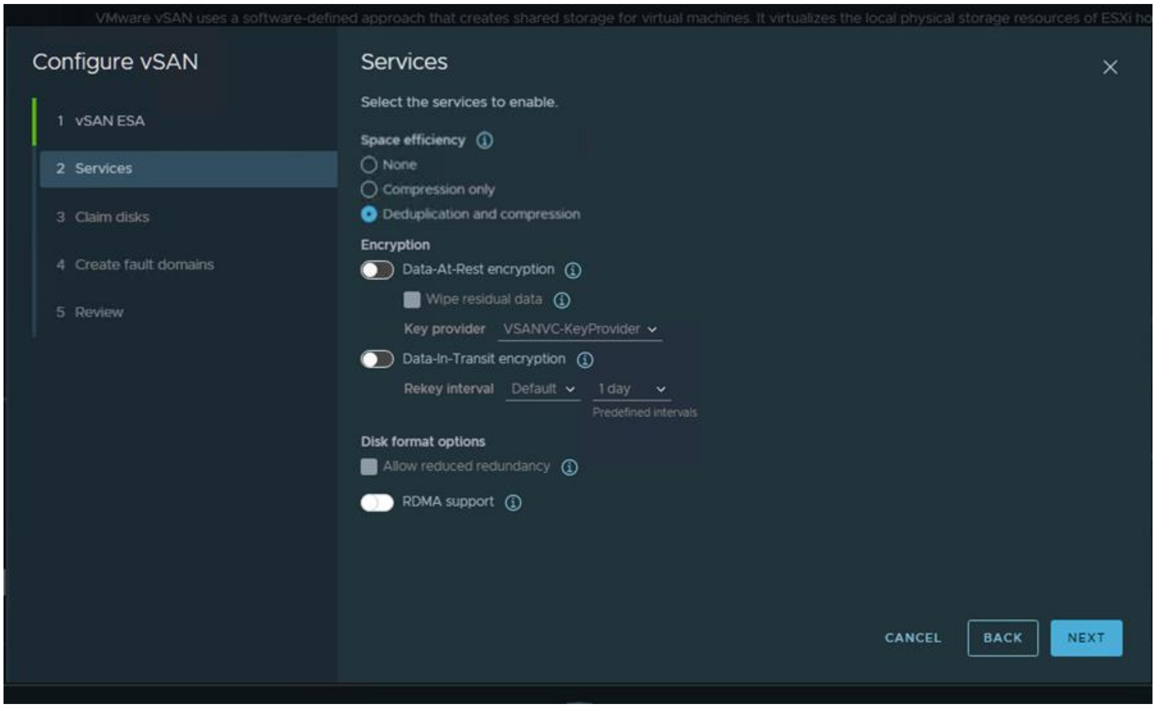

VMware vSAN uses a software-defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi ho

Configure vSAN

1 vSAN ESA
2 Services
3 Claim disks
4 Create fault domains
5 Review

Services

Select the services to enable.

Space efficiency

○ None
○ Compression only
● Deduplication and compression

Encryption

Data-At-Rest encryption

□ Wipe residual data

Key provider VSANVC-KeyProvider

Data-In-Transit encryption

Rekey interval Default 1 day

Predefined intervals

Disk format options

□ Allow reduced redundancy

RDMA support

Step 5. Claim disks for the vSAN cluster and click NEXT. Each host requires at least one flash device in the cache tier, and one or more devices in the capacity tier.

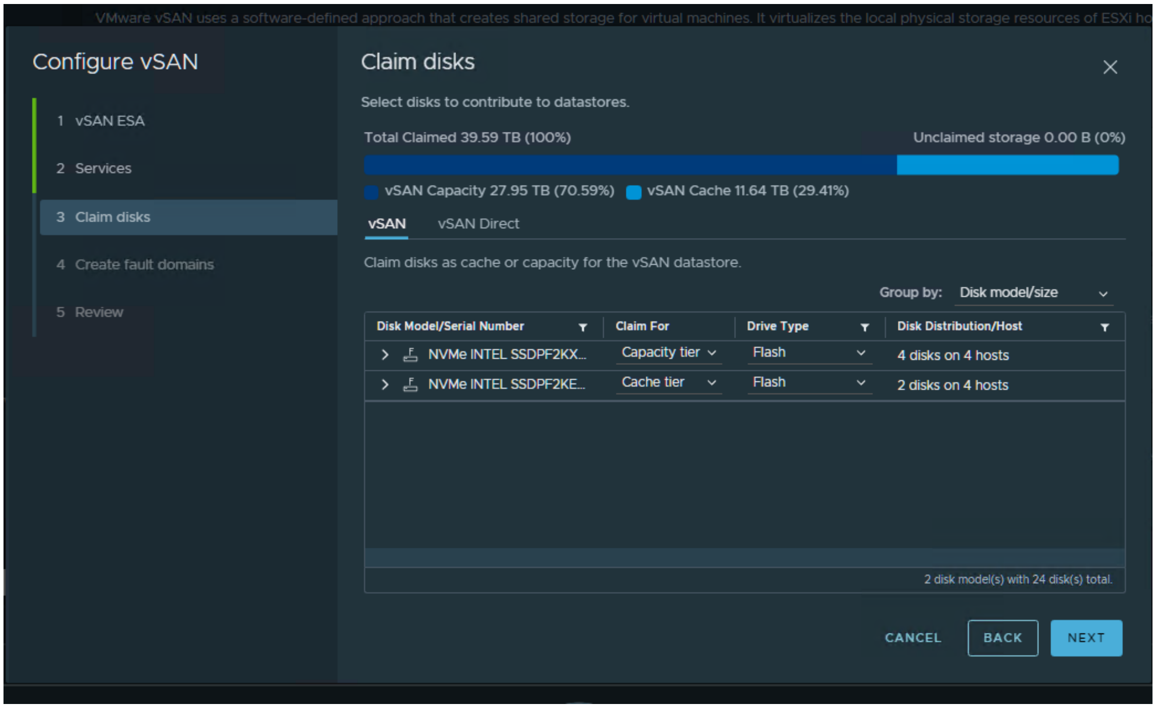

VMware vSAN uses a software-defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi ho

Configure vSAN

1 vSAN ESA
2 Services
3 Claim disks
4 Create fault domains
5 Review

Claim disks

Select disks to contribute to datastores.

Total Claimed 39.59 TB (100%)

Unclaimed storage 0.00 B (0%)

vSAN Capacity 27.95 TB (70.59%) vSAN Cache 11.64 TB (29.41%)

Claim disks as cache or capacity for the vSAN datastore.

Group by: Disk model/size

Disk Model/Serial Number	Claim For	Drive Type	Disk Distribution/Host
NVMe INTEL SSDPF2KX...	Capacity tier	Flash	4 disks on 4 hosts
NVMe INTEL SSDPF2KE...	Cache tier	Flash	2 disks on 4 hosts

2 disk model(s) with 24 disk(s) total.

Step 6. Create a fault domain, in case you need it, and click NEXT.

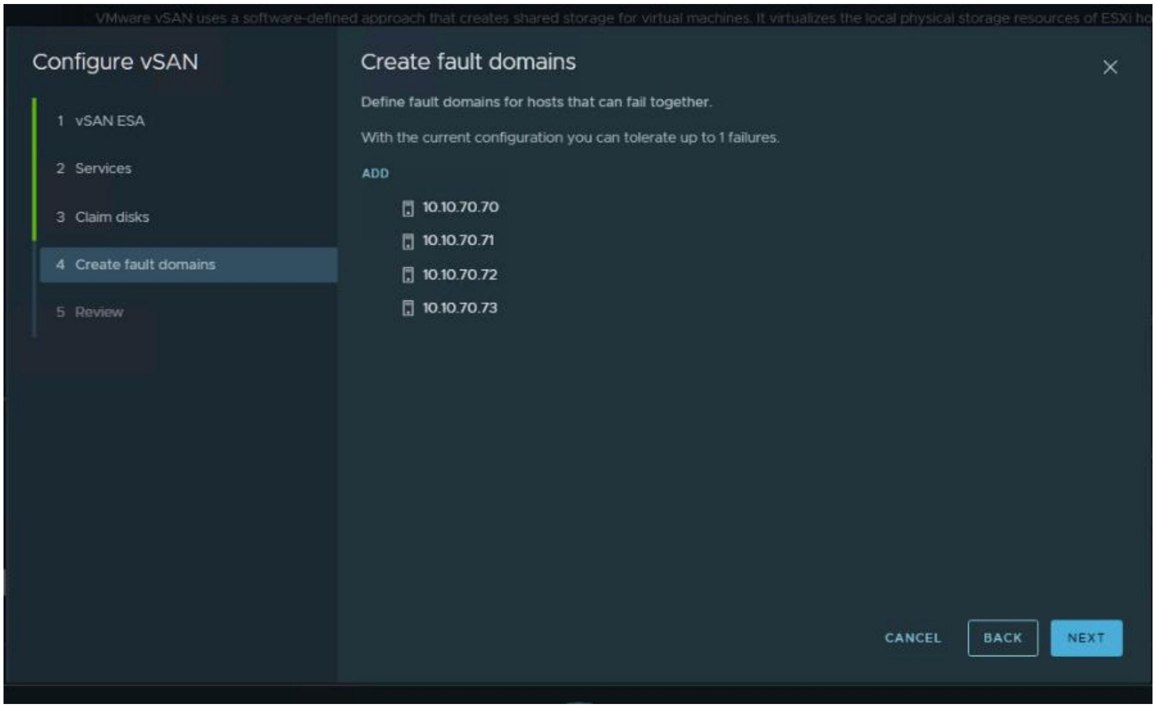
VMware vSAN uses a software-defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi ho
Configure vSAN
1 vSAN ESA
2 Services
3 Claim disks
4 Create fault domains
5 Review
CANCEL
BACK
NEXT

Create fault domains

Define fault domains for hosts that can fail together.

With the current configuration you can tolerate up to 1 failures.

ADD

- 10.10.70.70
- 10.10.70.71
- 10.10.70.72
- 10.10.70.73

Step 7. Review the configuration and click FINISH.

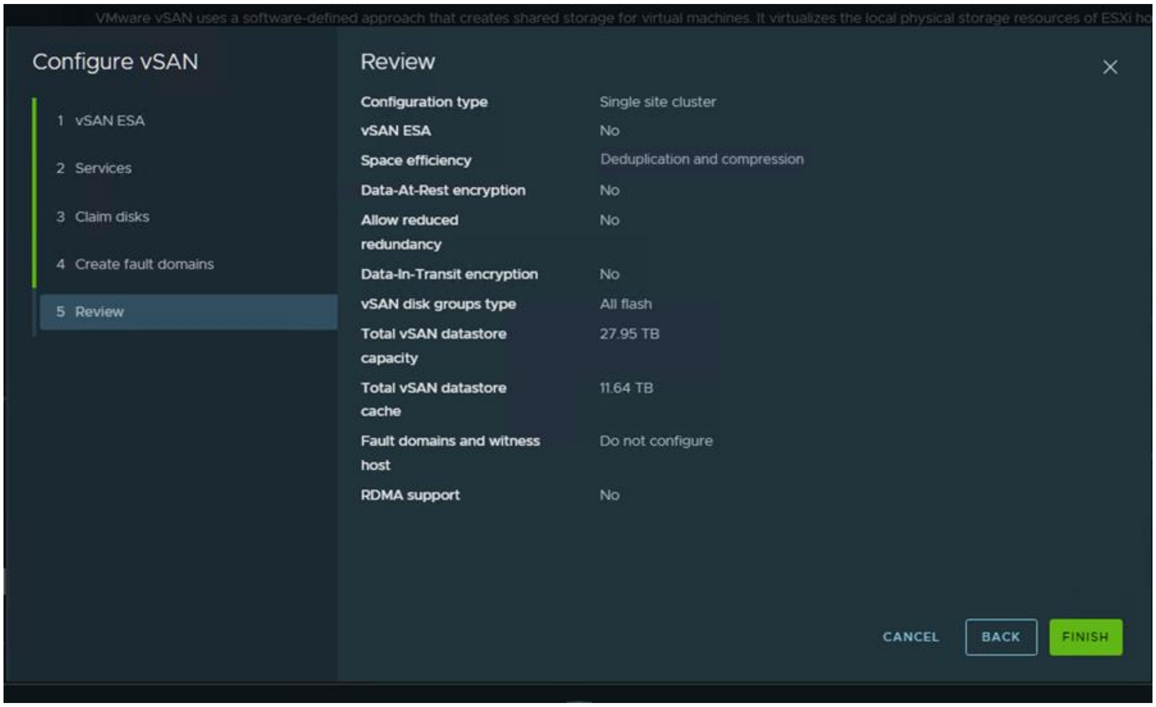
VMware vSAN uses a software-defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi ho
Configure vSAN
1 vSAN ESA
2 Services
3 Claim disks
4 Create fault domains
5 Review
CANCEL
BACK
FINISH

Review

Configuration type	Single site cluster
vSAN ESA	No
Space efficiency	Deduplication and compression
Data-At-Rest encryption	No
Allow reduced redundancy	No
Data-In-Transit encryption	No
vSAN disk groups type	All flash
Total vSAN datastore capacity	27.95 TB
Total vSAN datastore cache	11.64 TB
Fault domains and witness host	Do not configure
RDMA support	No

Procedure 2. Manually enable and configure vSAN ESA on the ESX Cluster

Note: The 3 node cluster was created. For more information about vSAN ESA, refer to the vSAN Design Guide

Step 1. Navigate to an existing host cluster.

Step 2. On Configure tab select vSAN and then Services. Select Standard vSAN cluster radio button and click Configure.

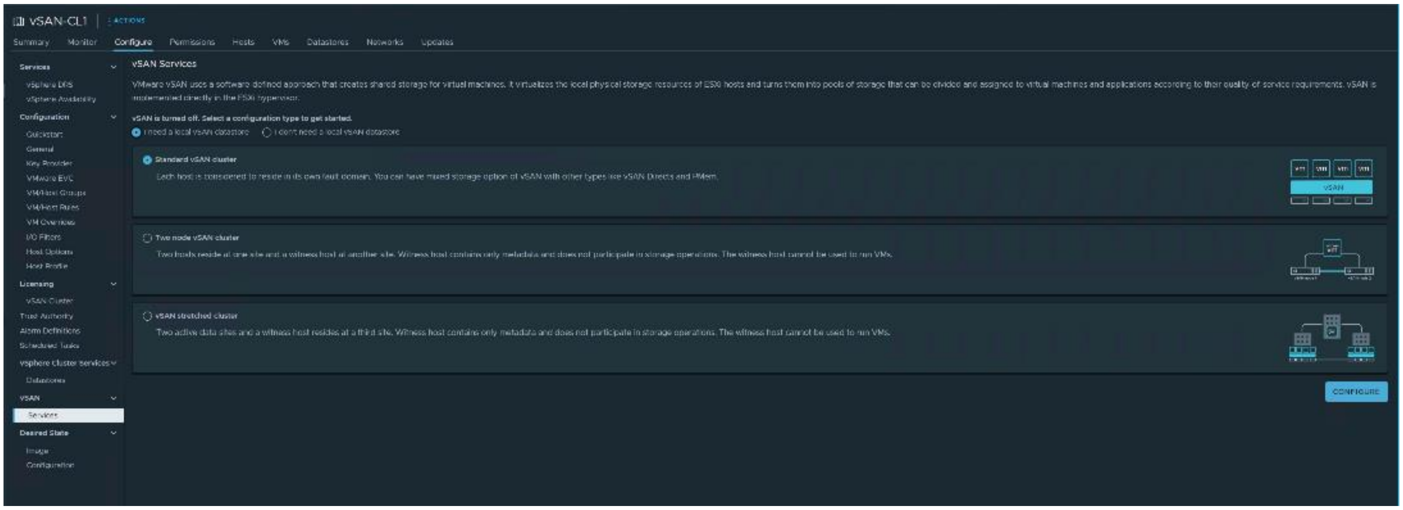

vSAN Services

VMware vSAN uses a software defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi hosts and turns them into pools of storage that can be divided and assigned to virtual machines and applications according to their quality of service requirements. vSAN is implemented directly in the ESXi hypervisor.

vSAN is turned off. Select a configuration type to get started.

● I need a local vSAN datastore ○ I don't need a local vSAN datastore

● Standard vSAN cluster

Each host is considered to reside in its own fault domain. You can have mixed storage option of vSAN with other types like vSAN Directs and PMem.

○ Two node vSAN cluster

Two hosts reside at one site and a witness host at another site. Witness host contains only metadata and does not participate in storage operations. The witness host cannot be used to run VMs.

○ vSAN stretched cluster

Two active data sites and a witness host resides at a third site. Witness host contains only metadata and does not participate in storage operations. The witness host cannot be used to run VMs.

Step 3. Click Next.

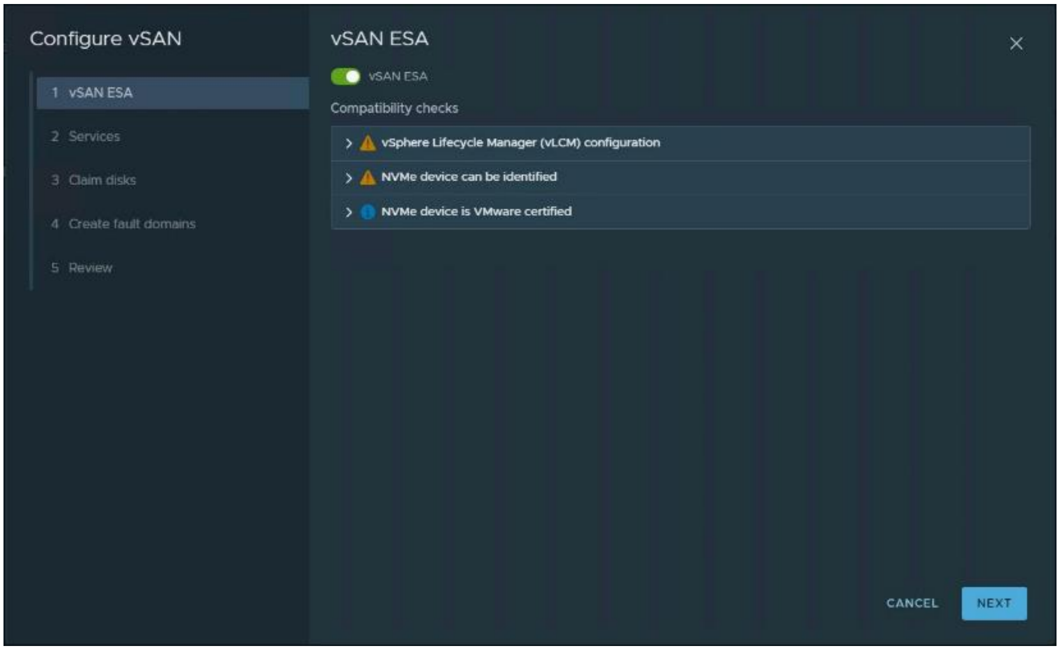

vSAN ESA

vSAN ESA

Compatibility checks

- vSphere Lifecycle Manager (vLCM) configuration
- NVMe device can be identified
- NVMe device is VMware certified

Step 4. Configure the data management features, including Data-At-Rest encryption, Data-In-Transit encryption and RDMA support (for example Storage policy managed compression). Click NEXT.

Note: Encryption and RDMA were not used in this deployment.

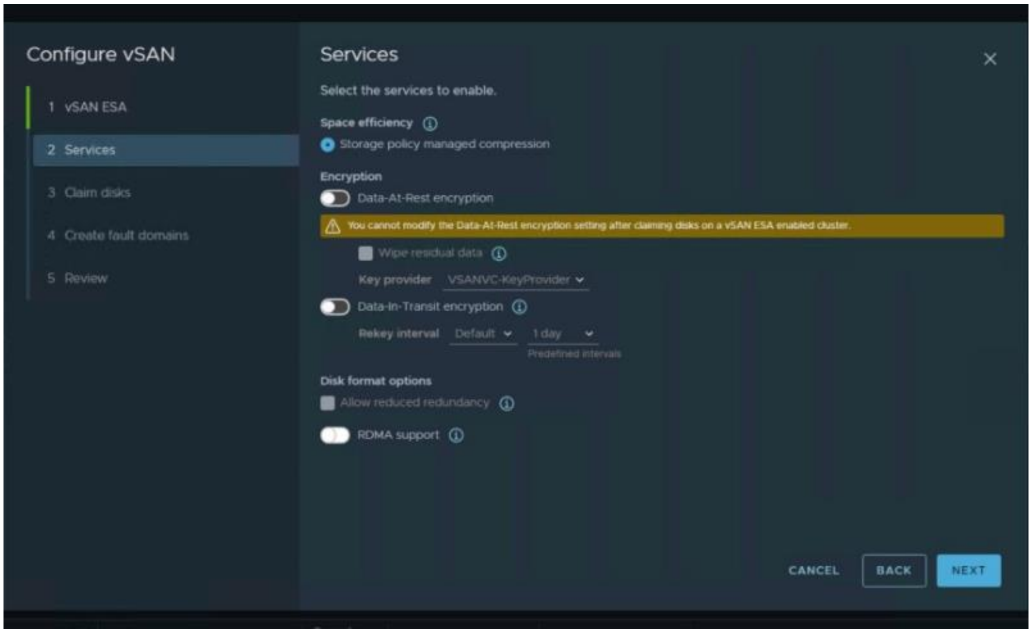

Configure vSAN

1 vSAN ESA
2 Services
3 Claim disks
4 Create fault domains
5 Review

Services

Select the services to enable.

Space efficiency
● Storage policy managed compression

Encryption
Data-At-Rest encryption

⚠ You cannot modify the Data-At-Rest encryption setting after claiming disks on a vSAN ESA enabled cluster.

□ Wipe residual data
Key provider: VSANVC-KeyProvider

Data-In-Transit encryption
Rekey interval: Default, 1 day (Predefined intervals)

Disk format options
□ Allow reduced redundancy

RDMA support

Step 5. Claim disks for the vSAN cluster and click NEXT. Each host requires at least one flash device in the cache tier, and one or more devices in the capacity tier.

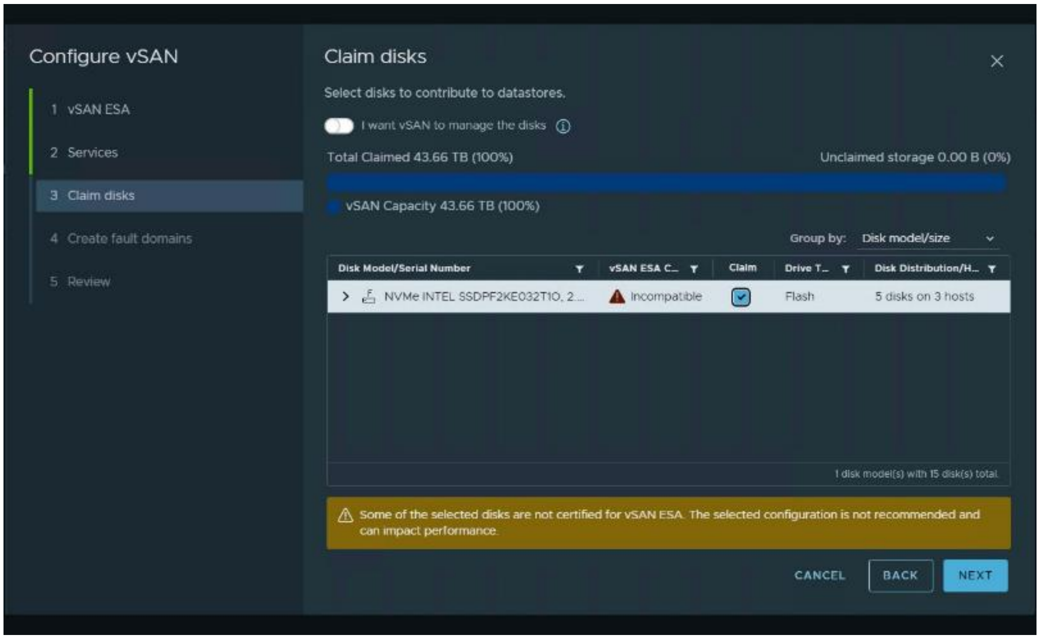

Configure vSAN

1 vSAN ESA
2 Services
3 Claim disks
4 Create fault domains
5 Review

Claim disks

Select disks to contribute to datastores.

I want vSAN to manage the disks

Total Claimed 43.66 TB (100%)

Unclaimed storage 0.00 B (0%)

vSAN Capacity 43.66 TB (100%)

Group by: Disk model/size

Disk Model/Serial Number	vSAN ESA C...	Claim	Drive T...	Disk Distribution/H...
NVMe INTEL SSDPF2KE032T1O, 2...	Incompatible	☑	Flash	5 disks on 3 hosts

1 disk model(s) with 15 disk(s) total

⚠ Some of the selected disks are not certified for vSAN ESA. The selected configuration is not recommended and can impact performance.

Step 6. Create a fault domain, in case you need it, and click NEXT.

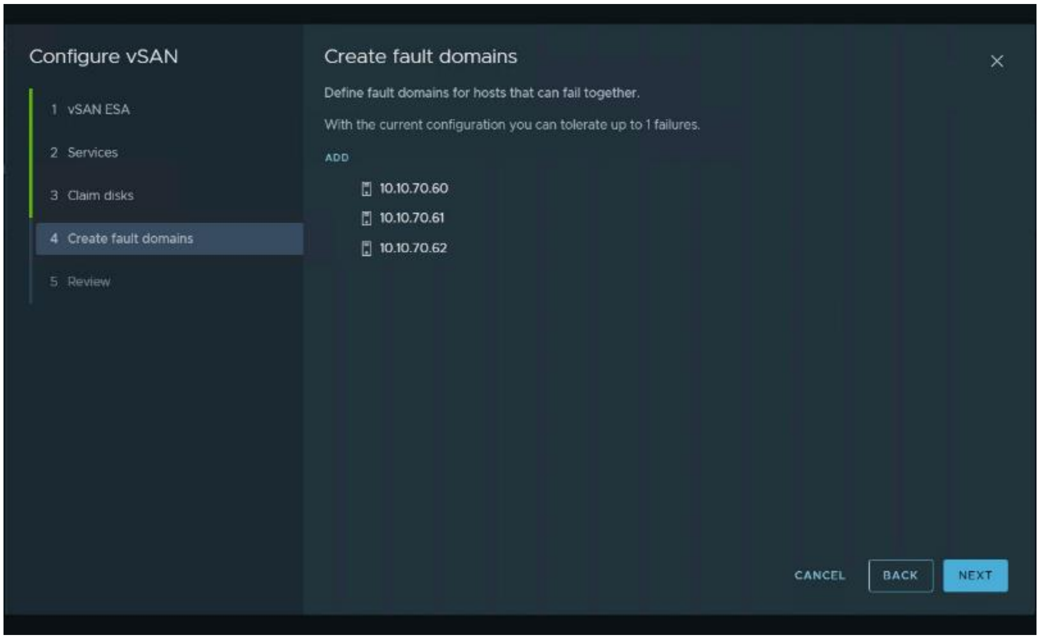

Create fault domains

Define fault domains for hosts that can fail together.

With the current configuration you can tolerate up to 1 failures.

ADD

- 10.10.70.60
- 10.10.70.61
- 10.10.70.62

Step 7. Review the configuration and click FINISH.

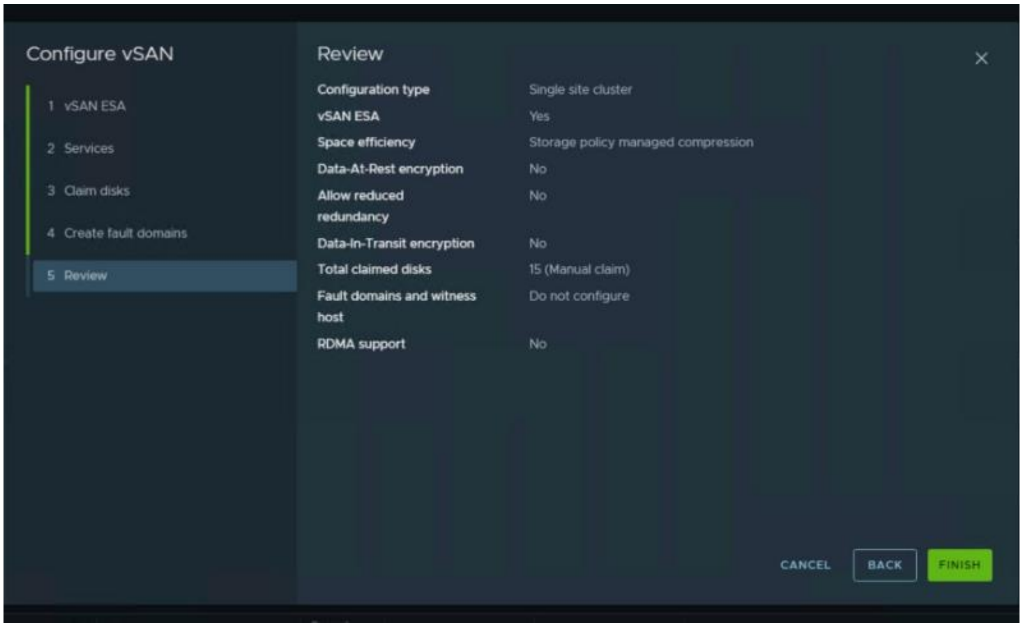

Review

Configuration type	Single site cluster
vSAN ESA	Yes
Space efficiency	Storage policy managed compression
Data-At-Rest encryption	No
Allow reduced redundancy	No
Data-In-Transit encryption	No
Total claimed disks	15 (Manual claim)
Fault domains and witness host	Do not configure
RDMA support	No

Procedure 3. Configure License Settings for a vSAN Cluster

Note: You must assign a license to a vSAN cluster before its evaluation period, or its currently assigned license expire.

Step 1. Navigate to your vSAN cluster.

Step 2. Right-click the Cluster and select Assign vSAN Cluster License... from the Actions menu.

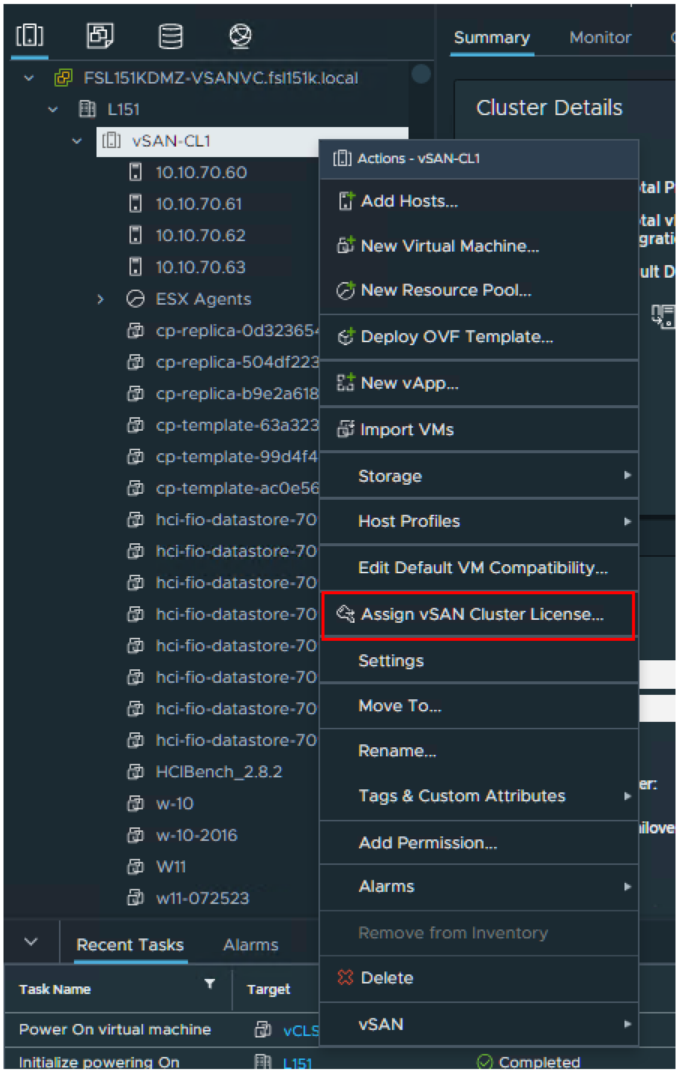

Task Name	Target	
Power On virtual machine	vCLS	
Initialize powering On	L151	Completed

Step 3. Select an existing license and click OK.

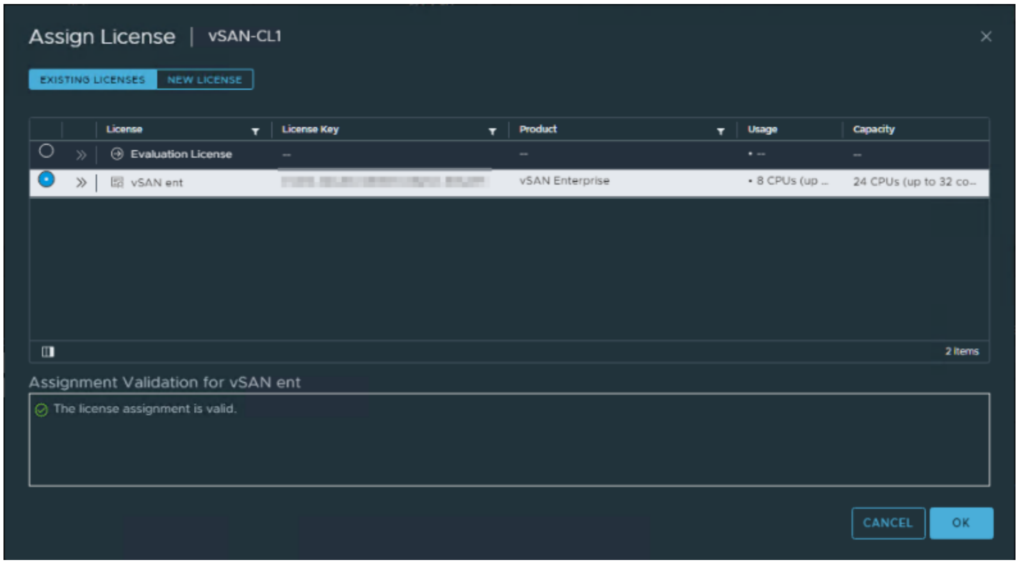

		License	License Key	Product	Usage	Capacity
○	»	Evaluation License	--	--	• --	--
●	»	vSAN ent	[redacted]	vSAN Enterprise	• 8 CPUs (up ...	24 CPUs (up to 32 co...

2 items

Assignment Validation for vSAN ent

✓ The license assignment is valid.

Step 4. Validate the Assigned License.

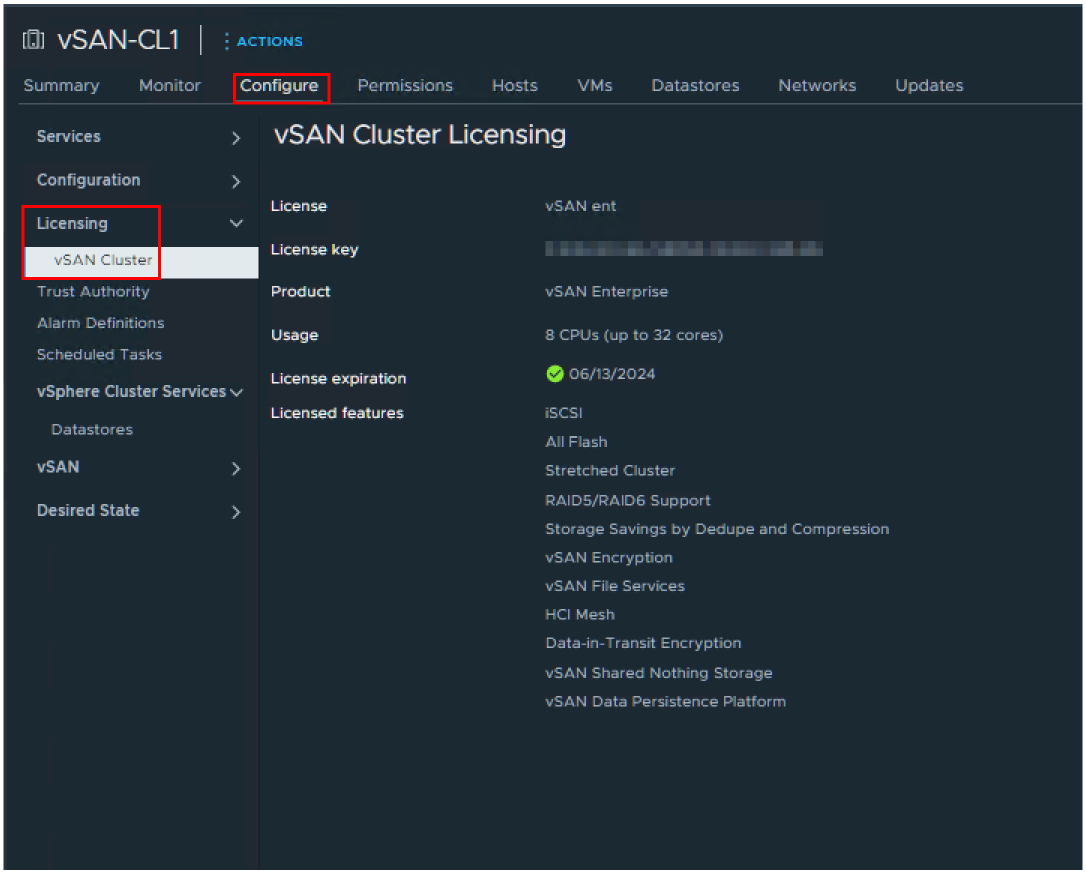

vSAN Cluster Licensing

License	vSAN ent
License key	[redacted]
Product	vSAN Enterprise
Usage	8 CPUs (up to 32 cores)
License expiration	✓ 06/13/2024
Licensed features	iSCSI All Flash Stretched Cluster RAID5/RAID6 Support Storage Savings by Dedupe and Compression vSAN Encryption vSAN File Services HCI Mesh Data-in-Transit Encryption vSAN Shared Nothing Storage vSAN Data Persistence Platform

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 the necessary DHCP scopes for the environment and set the Scope Options.

Figure 21. Example of the DHCP Scopes used in this CVD

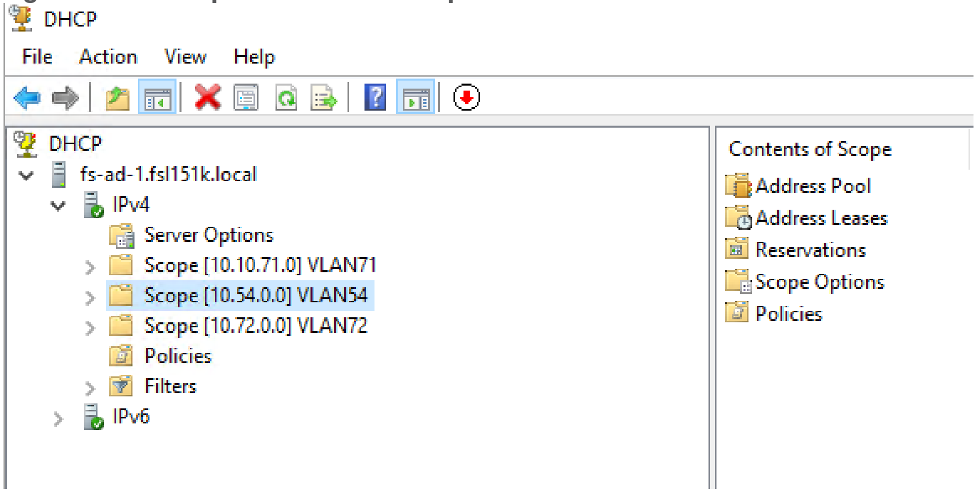

Software Infrastructure Configuration

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 16.

Table 16. Test Infrastructure Virtual Machine Configuration

Configuration	Microsoft Active Directory DCs Virtual Machine	vCenter Server Appliance Virtual Machine
Operating system	Microsoft Windows Server 2019	VCSA – SUSE Linux
Virtual CPU amount	4	16
Memory amount	8 GB	32 GB
Network	VMXNET3 Infra-Mgmt_71	VMXNET3 Infra-Mgmt_71
Disk-1 (OS) size	40 GB	698.84 GB (across 13 VMDKs)
Disk-2 size		

Configuration	Microsoft SQL Server Virtual Machine	VMware Horizon Connection Server Virtual Machine
Operating system	Microsoft Windows Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	6	4
Memory amount	24GB	8 GB
Network	VMXNET3 Infra-Mgmt_71	VMXNET3 Infra-Mgmt_71
Disk-1 (OS) size	40 GB	40
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 2022 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 17.

Table 17. 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
Network	VMXNET3 VDI_72	VMXNET3 VDI_72
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	Configure DHCP

Configuration	Single-session OS Virtual Machine	Mutli-session OS Virtual Machine
	Add to domain Install VMWare tools Install .Net 3.5 Activate Office Install Horizon Agent Install FSLogix 2210 hotfix 1	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.

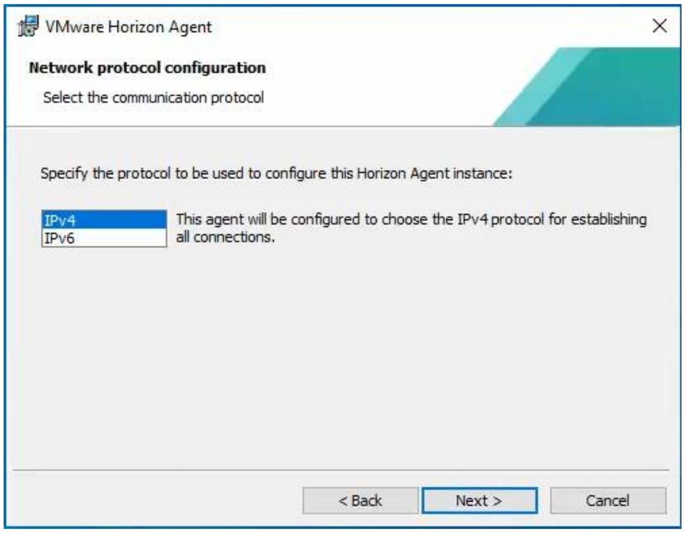

VMware Horizon Agent

Network protocol configuration

Select the communication protocol

Specify the protocol to be used to configure this Horizon Agent instance:

IPv4 (selected)
IPv6

This agent will be configured to choose the IPv4 protocol for establishing all connections.

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.

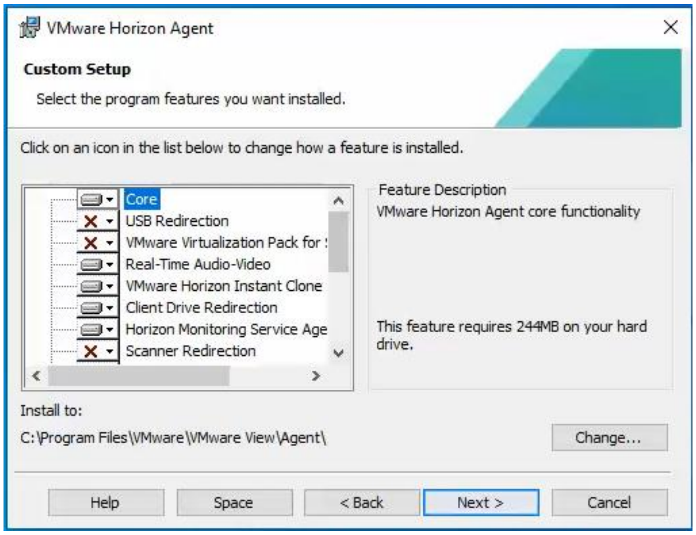

VMware Horizon Agent

Custom Setup

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

- ▭ Core
- ✕ USB Redirection
- ✕ VMware Virtualization Pack for :
- ▭ Real-Time Audio-Video
- ▭ VMware Horizon Instant Clone
- ▭ Client Drive Redirection
- ▭ Horizon Monitoring Service Age
- ✕ Scanner Redirection

Feature Description

VMware Horizon Agent core functionality

This feature requires 244MB on your hard drive.

Install to:

C:\Program Files\VMware\VMware View\Agent\

Step 3. Enable the remote Desktop Protocol.

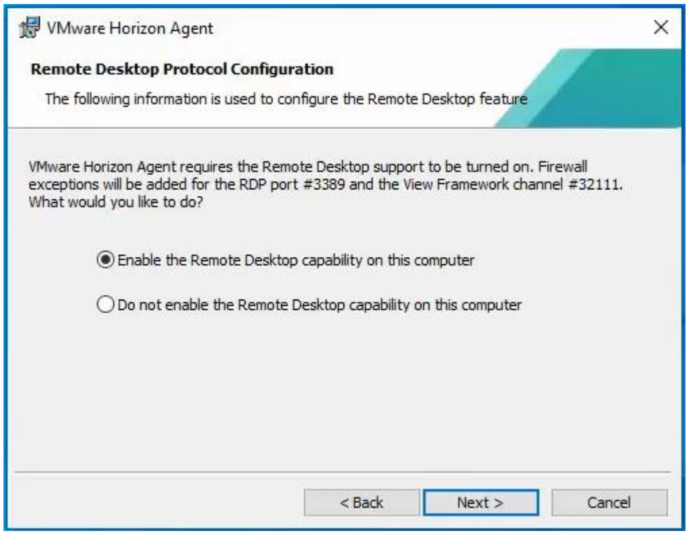

VMware Horizon Agent

Remote Desktop Protocol Configuration

The following information is used to configure the Remote Desktop feature

VMware Horizon Agent requires the Remote Desktop support to be turned on. Firewall exceptions will be added for the RDP port #3389 and the View Framework channel #32111. What would you like to do?

● Enable the Remote Desktop capability on this computer

○ Do not enable the Remote Desktop capability on this computer

Step 4. During the VMware Horizon Agent installation on the Windows server select RDS Mode otherwise select Desktop Mode.

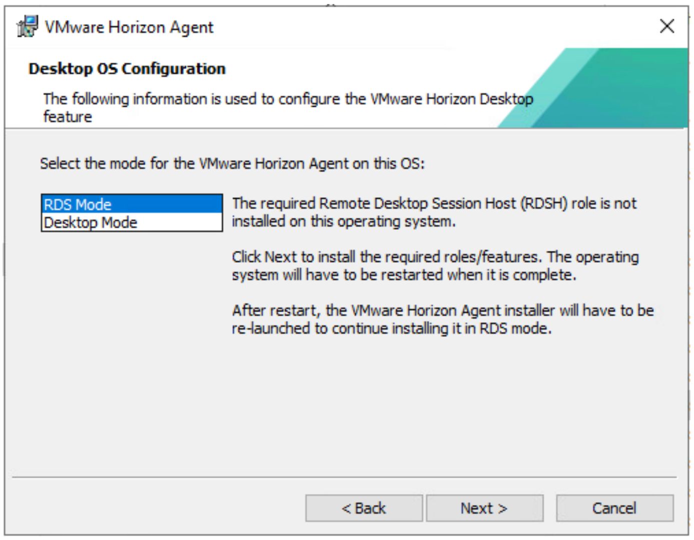

VMware Horizon Agent

Desktop OS Configuration

The following information is used to configure the VMware Horizon Desktop feature

Select the mode for the VMware Horizon Agent on this OS:

- RDS Mode
- Desktop Mode

The required Remote Desktop Session Host (RDSH) role is not installed on this operating system.

Click Next to install the required roles/features. The operating system will have to be restarted when it is complete.

After restart, the VMware Horizon Agent installer will have to be re-launched to continue installing it in RDS 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.

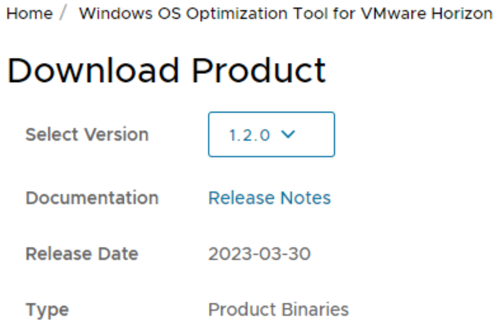

Home / Windows OS Optimization Tool for VMware Horizon

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 2022.

To successfully run the Login VSI knowledge worker workload the ‘Disable animation in web pages – Machine Policy’ option in Default Template under Programs -> Internal Explorer was disabled.

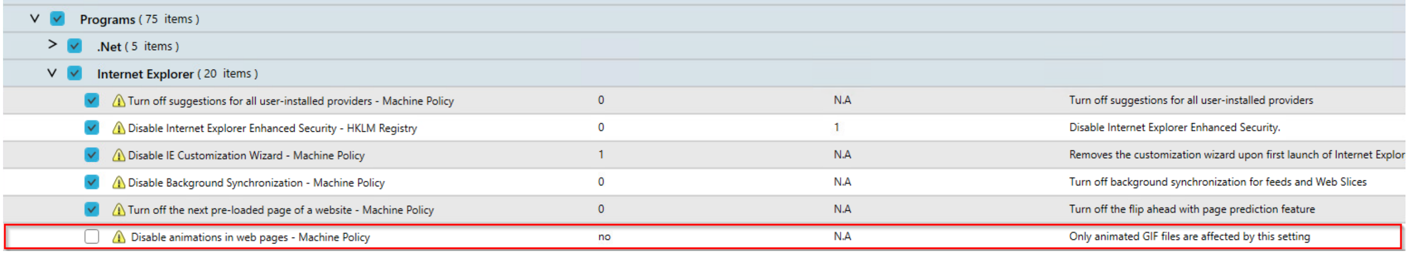

V ☑ **Programs** (75 items)			
> ☑ **.Net** (5 items)			
V ☑ **Internet Explorer** (20 items)			
☑ ⚠ Turn off suggestions for all user-installed providers - Machine Policy	0	N.A	Turn off suggestions for all user-installed providers
☑ ⚠ Disable Internet Explorer Enhanced Security - HKLM Registry	0	1	Disable Internet Explorer Enhanced Security.
☑ ⚠ Disable IE Customization Wizard - Machine Policy	1	N.A	Removes the customization wizard upon first launch of Internet Explor
☑ ⚠ Disable Background Synchronization - Machine Policy	0	N.A	Turn off background synchronization for feeds and Web Slices
☑ ⚠ Turn off the next pre-loaded page of a website - Machine Policy	0	N.A	Turn off the flip ahead with page prediction feature
□ ⚠ Disable animations in web pages - Machine Policy	no	N.A	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.

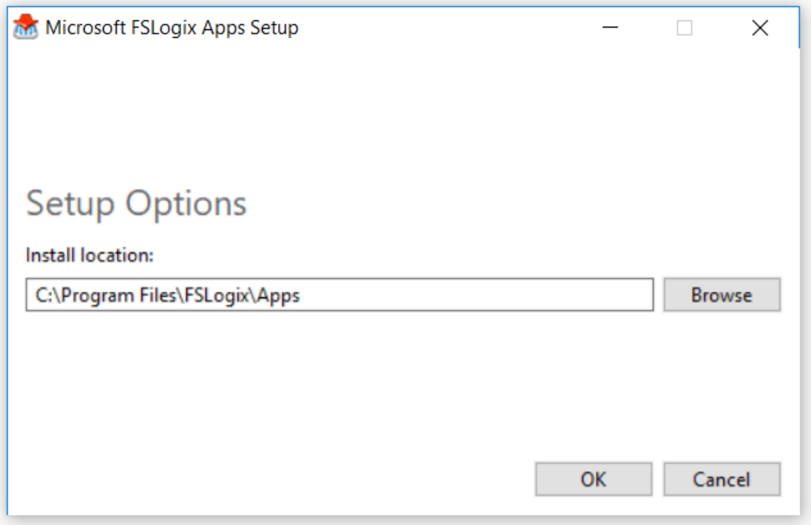

Install location: C:\Program Files\FSLogix\Apps

Step 4. Review and accept the license agreement.

Step 5. Click Install.

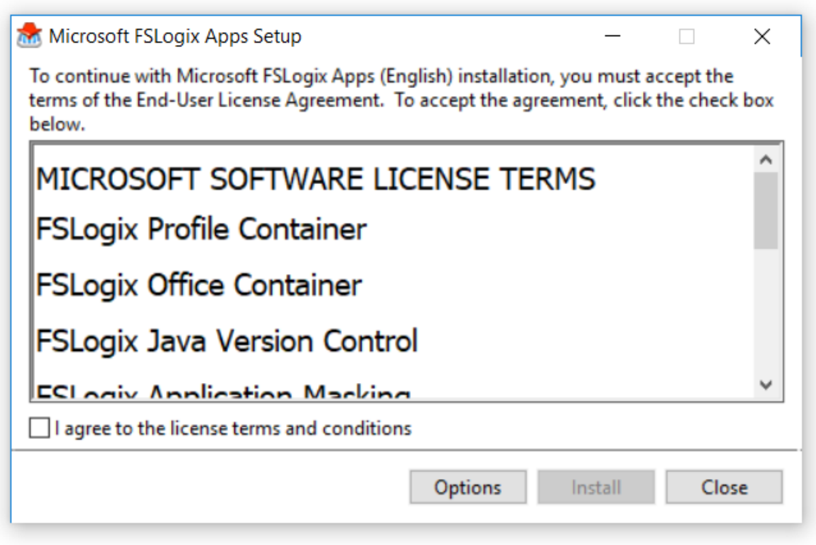

To continue with Microsoft FSLogix Apps (English) installation, you must accept the terms of the End-User License Agreement. To accept the agreement, click the check box below.

MICROSOFT SOFTWARE LICENSE TERMS

FSLogix Profile Container

FSLogix Office Container

FSLogix Java Version Control

FSLogix Application Masking

☐ I agree to the license terms and conditions

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.
- 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 22. Example of FSLogix Policy

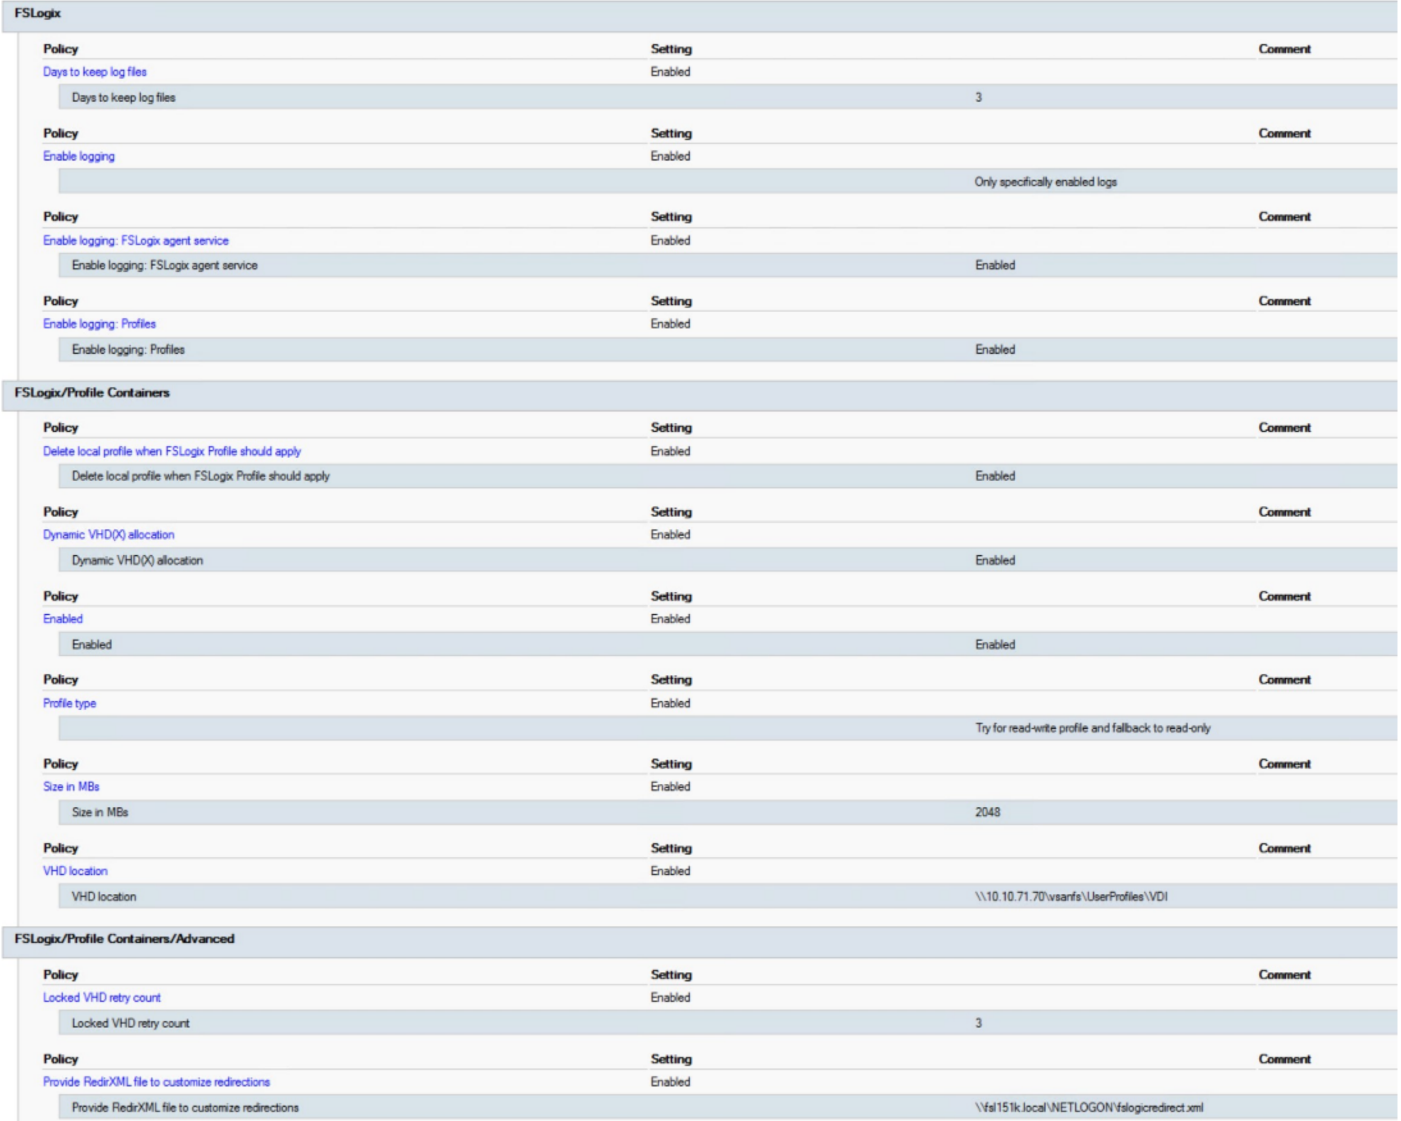

FSLogix

Policy	Setting	Comment
Days to keep log files	Enabled	
Days to keep log files	3	

Policy	Setting	Comment
Enable logging	Enabled	
	Only specifically enabled logs	

Policy	Setting	Comment
Enable logging: FSLogix agent service	Enabled	
Enable logging: FSLogix agent service	Enabled	

Policy	Setting	Comment
Enable logging: Profiles	Enabled	
Enable logging: Profiles	Enabled	

FSLogix/Profile Containers

Policy	Setting	Comment
Delete local profile when FSLogix Profile should apply	Enabled	
Delete local profile when FSLogix Profile should apply	Enabled	

Policy	Setting	Comment
Dynamic VHD(X) allocation	Enabled	
Dynamic VHD(X) allocation	Enabled	

Policy	Setting	Comment
Enabled	Enabled	
Enabled	Enabled	

Policy	Setting	Comment
Profile type	Enabled	
	Try for read-write profile and fallback to read-only	

Policy	Setting	Comment
Size in MBs	Enabled	
Size in MBs	2048	

Policy	Setting	Comment
VHD location	Enabled	
VHD location	\\10.10.71.70\vsanfs\UserProfiles\VDI	

FSLogix/Profile Containers/Advanced

Policy	Setting	Comment
Locked VHD retry count	Enabled	
Locked VHD retry count	3	

Policy	Setting	Comment
Provide RedirXML file to customize redirections	Enabled	
Provide RedirXML file to customize redirections	\\fsl151k.local\NETLOGON\fslogicredirect.xml	

Figure 23. **FSLogix policy exclusions list**

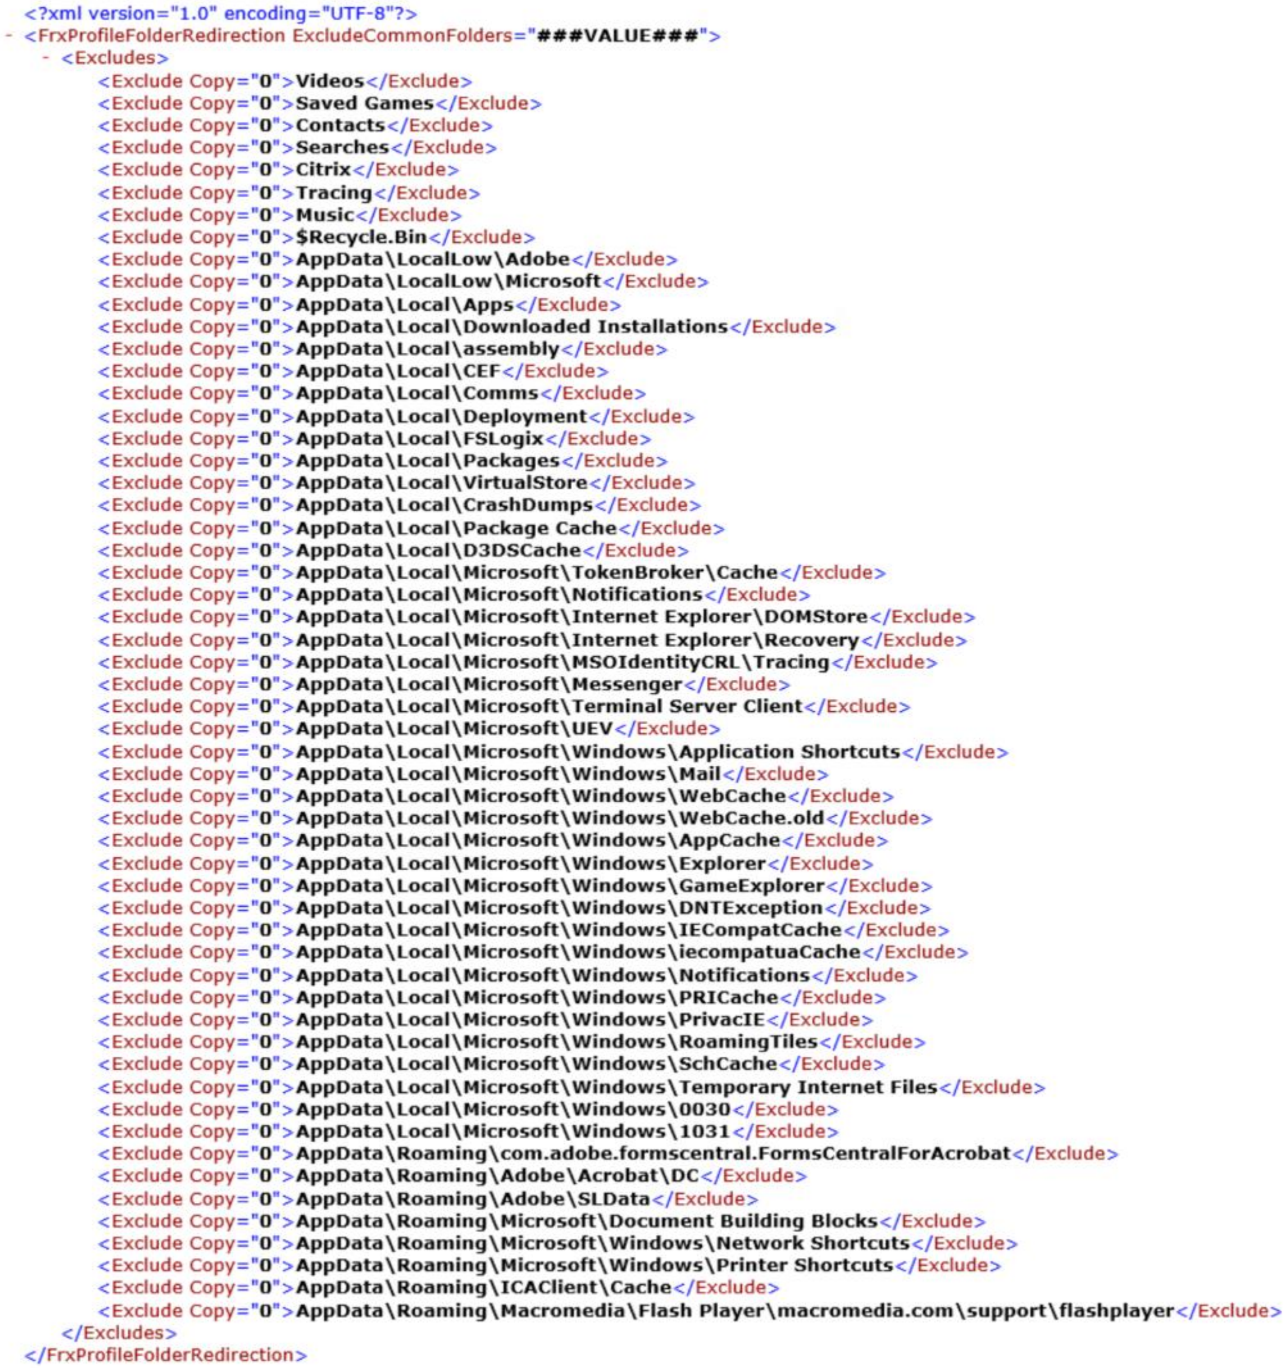

```
  <?xml version="1.0" encoding="UTF-8"?>
- <FrxProfileFolderRedirection ExcludeCommonFolders="###VALUE###">
  - <Excludes>
      <Exclude Copy="0">Videos</Exclude>
      <Exclude Copy="0">Saved Games</Exclude>
      <Exclude Copy="0">Contacts</Exclude>
      <Exclude Copy="0">Searches</Exclude>
      <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>
  </FrxProfileFolderRedirection>
```

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.

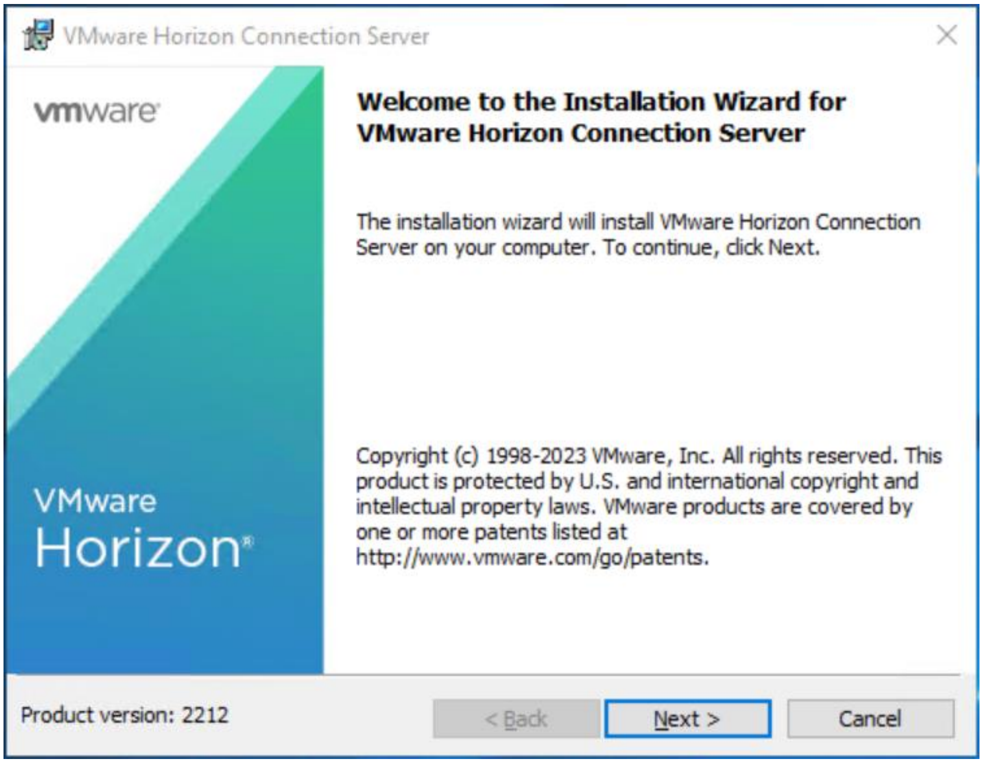

VMware Horizon Connection Server

Welcome to the Installation Wizard for VMware Horizon Connection Server

The installation wizard will install VMware Horizon Connection Server on your computer. To continue, click Next.

Copyright (c) 1998-2023 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 2212

Step 3. Read and accept the End User License Agreement and click Next.

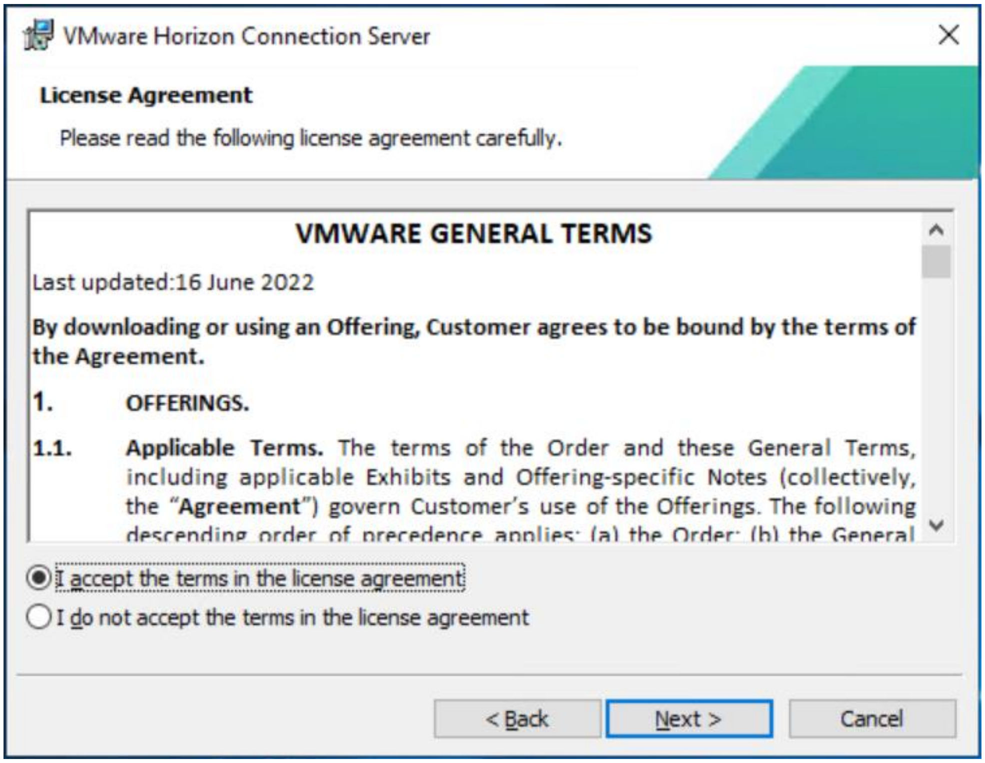

VMware Horizon Connection Server

License Agreement

Please read the following license agreement carefully.

VMWARE GENERAL TERMS

Last updated:16 June 2022

By downloading or using an Offering, Customer agrees to be bound by the terms of the Agreement.

1. OFFERINGS.

1.1. Applicable Terms. The terms of the Order and these General Terms, including applicable Exhibits and Offering-specific Notes (collectively, the "**Agreement**") govern Customer's use of the Offerings. The following descending order of precedence applies: (a) the Order; (b) the General

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 4. Select the destination folder where you want to install the application and click Next.

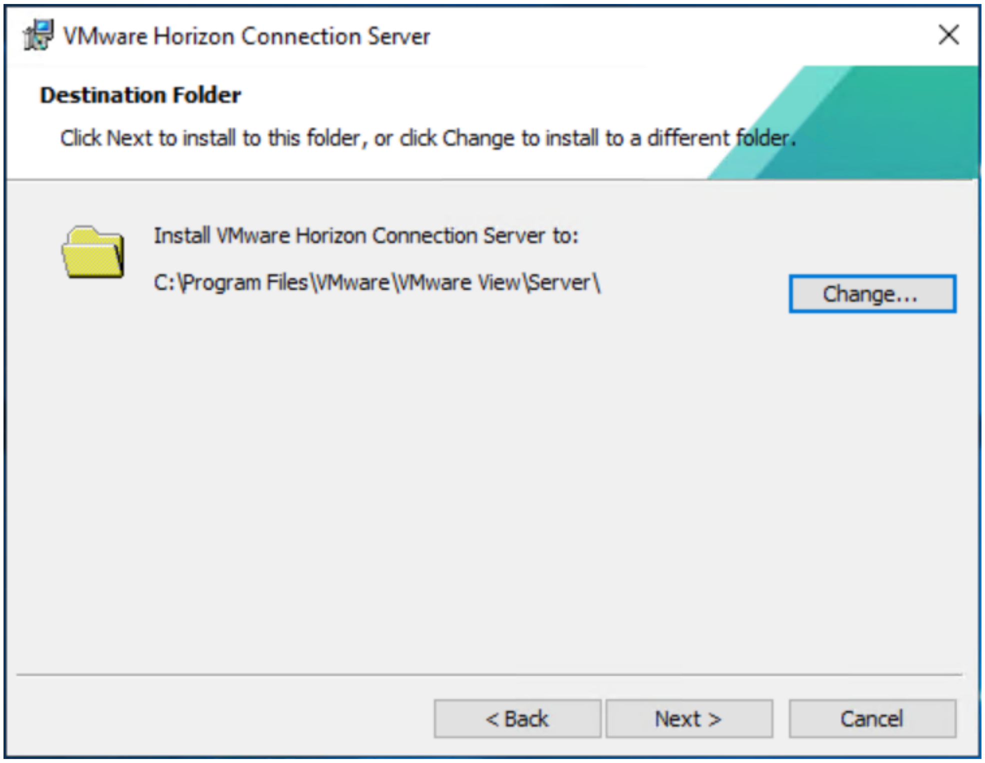

VMware Horizon Connection Server

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware Horizon Connection Server to:

C:\Program Files\VMware\VMware View\Server\

Step 5. Select the Standard Server and IPv4 for the IP protocol version.

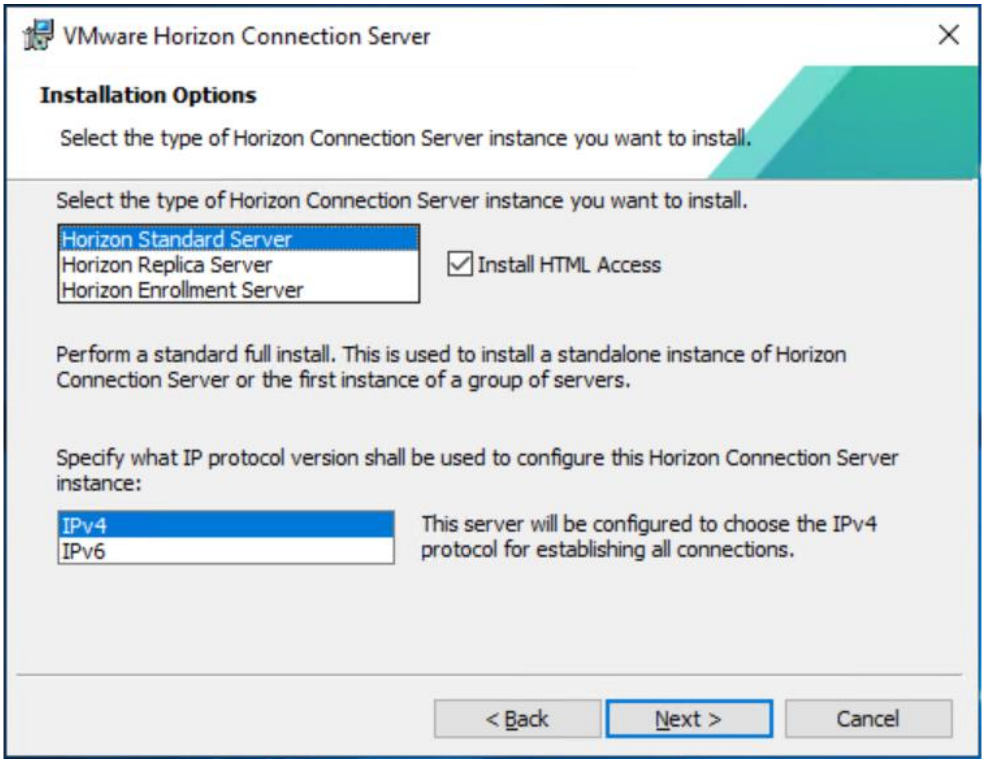

VMware Horizon Connection Server

Installation Options

Select the type of Horizon Connection Server instance you want to install.

Select the type of Horizon Connection Server instance you want to install.

- Horizon Standard Server (selected)
- Horizon Replica Server
- Horizon Enrollment Server

☑ Install HTML Access

Perform a standard full install. This is used to install a standalone instance of Horizon Connection Server or the first instance of a group of servers.

Specify what IP protocol version shall be used to configure this Horizon Connection Server instance:

- IPv4 (selected)
- IPv6

This server will be configured to choose the IPv4 protocol for establishing all connections.

Step 6. Provide data recovery details.

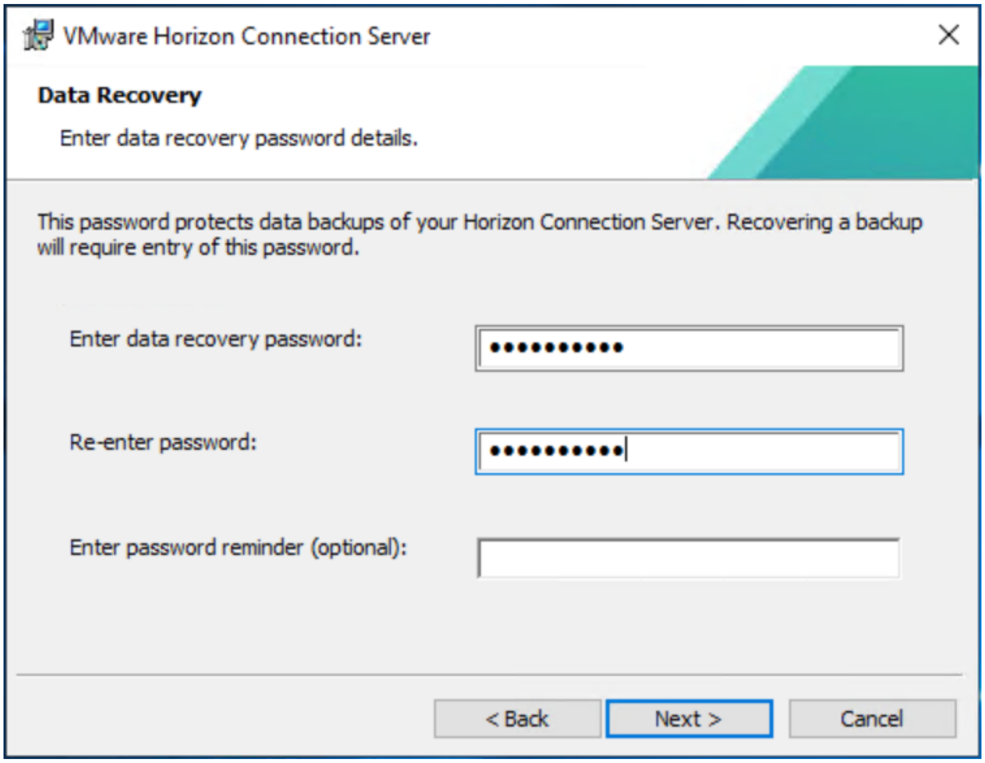

VMware Horizon Connection Server

Data Recovery

Enter data recovery password details.

This password protects data backups of your Horizon Connection Server. Recovering a backup will require entry of this password.

Enter data recovery password: ●●●●●●●●●●●

Re-enter password: ●●●●●●●●●●●

Enter password reminder (optional):

Step 7. Select Configure Windows Firewall automatically. Click Next.

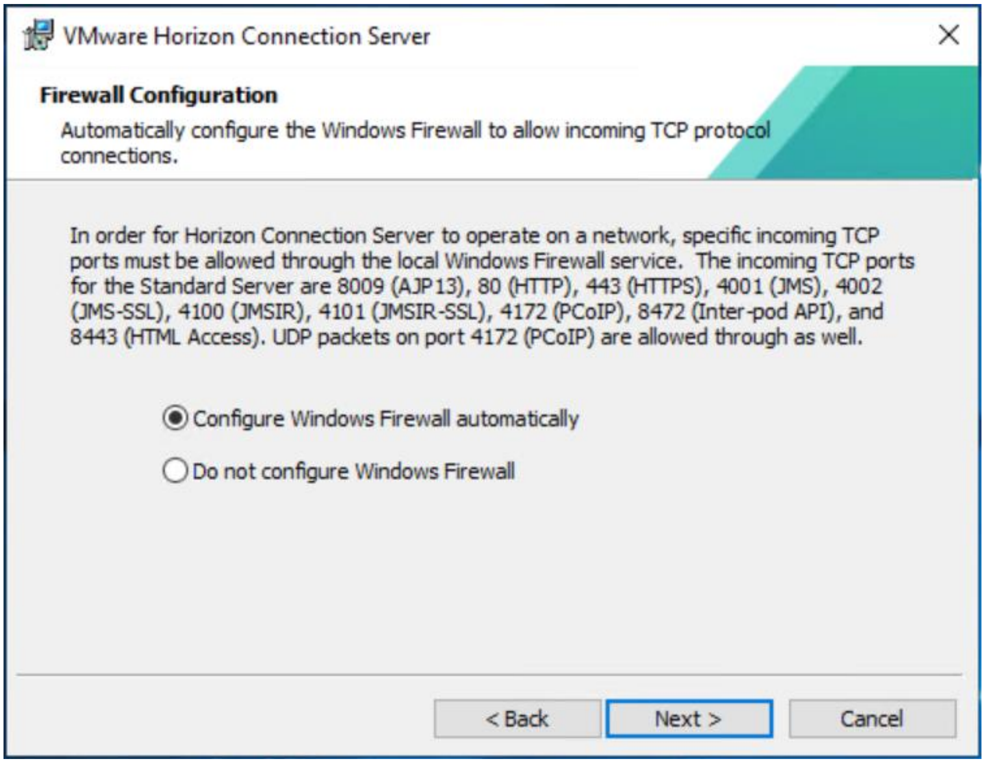

VMware Horizon Connection Server

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for Horizon Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Standard Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), 4002 (JMS-SSL), 4100 (JMSIR), 4101 (JMSIR-SSL), 4172 (PCoIP), 8472 (Inter-pod API), and 8443 (HTML Access). UDP packets on port 4172 (PCoIP) are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

Step 8. Authorize Domain Admins to be VMware Horizon administrators.

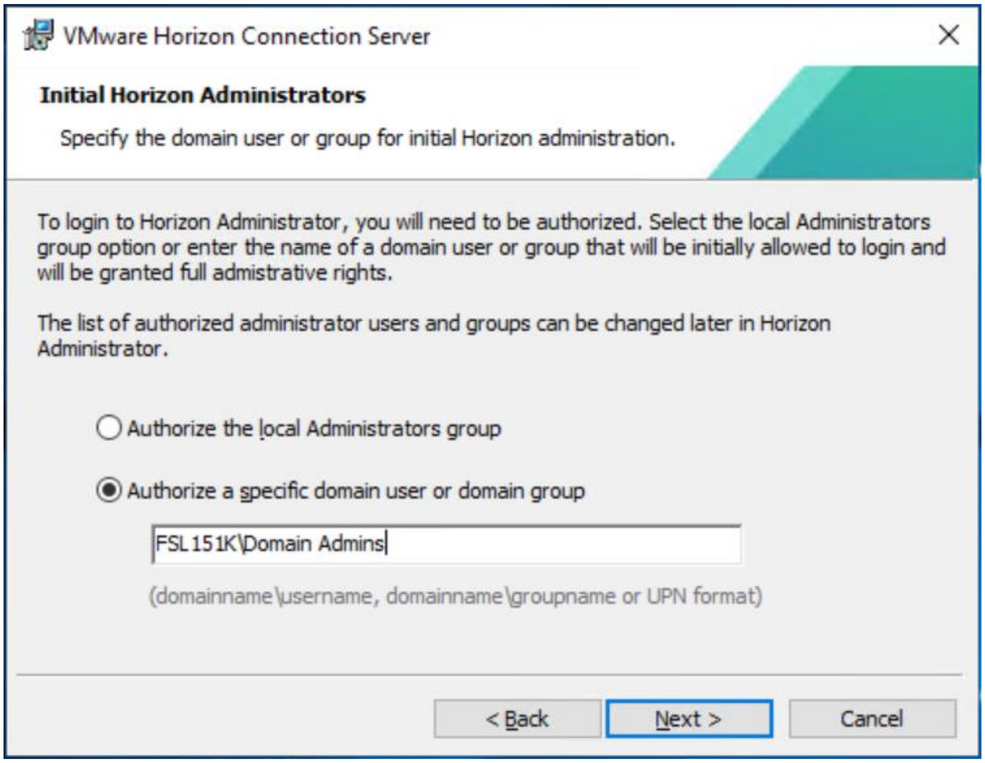

VMware Horizon Connection Server

Initial Horizon Administrators

Specify the domain user or group for initial Horizon administration.

To login to Horizon Administrator, you will need to be authorized. Select the local Administrators group option or enter the name of a domain user or group that will be initially allowed to login and will be granted full admistrative rights.

The list of authorized administrator users and groups can be changed later in Horizon Administrator.

○ Authorize the local Administrators group

● Authorize a specific domain user or domain group

FSL151K\Domain Admins

(domainname\username, domainname\groupname or UPN format)

Step 9. (Optional) Join Customer Experience Program.

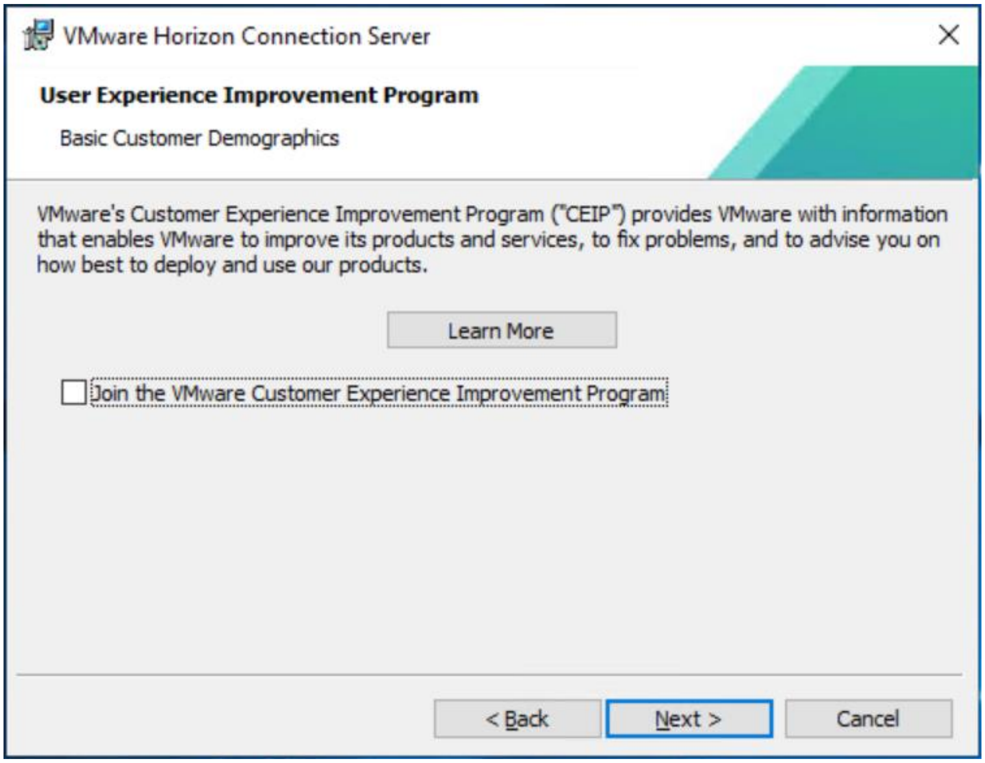

VMware Horizon Connection Server

User Experience Improvement Program

Basic Customer Demographics

VMware's Customer Experience Improvement Program ("CEIP") provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products.

□ Join the VMware Customer Experience Improvement Program

Step 10. Click Next.

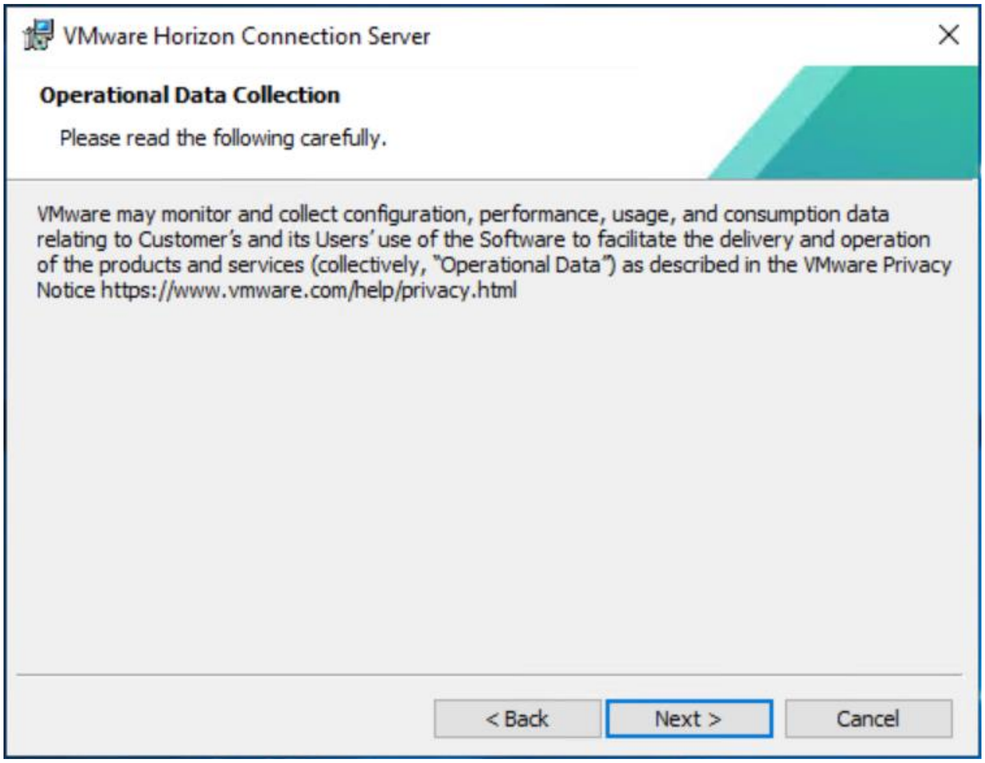

Operational Data Collection

Please read the following carefully.

VMware may monitor and collect configuration, performance, usage, and consumption data relating to Customer's and its Users' use of the Software to facilitate the delivery and operation of the products and services (collectively, "Operational Data") as described in the VMware Privacy Notice https://www.vmware.com/help/privacy.html

Step 11. Select General for the type of the type of installation. Click Install.

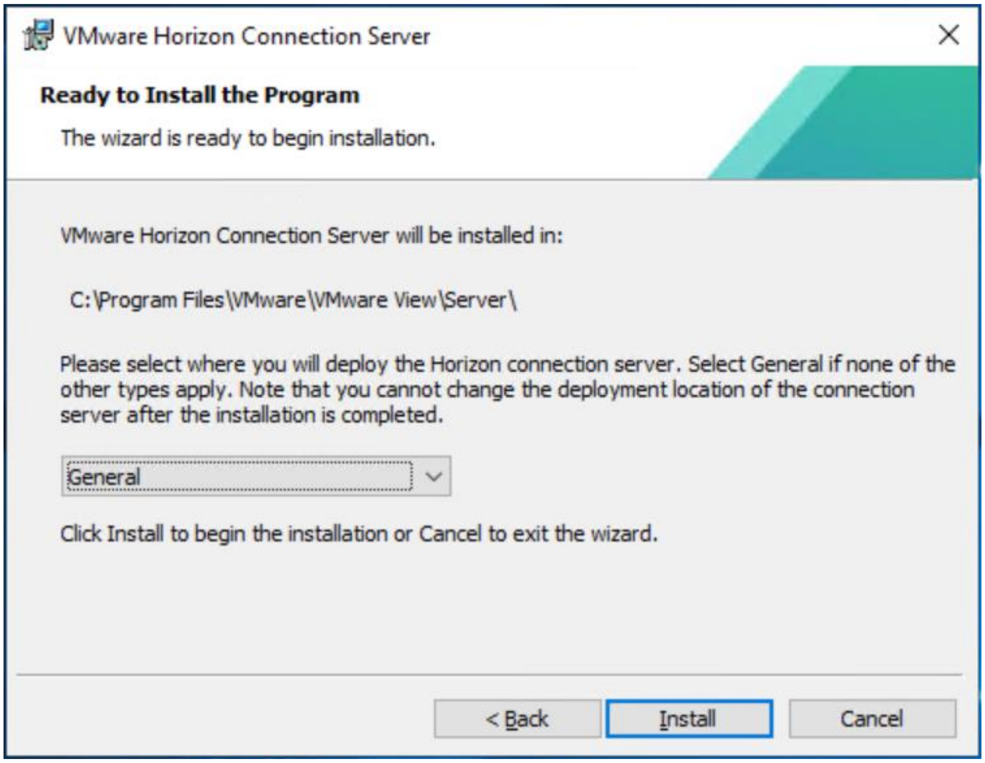

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon Connection Server will be installed in:

C:\Program Files\VMware\VMware View\Server\

Please select where you will deploy the Horizon connection server. Select General if none of the other types apply. Note that you cannot change the deployment location of the connection server after the installation is completed.

General

Click Install to begin the installation or Cancel to exit the wizard.

Step 12. After Horizon Connection Server installation is complete, click Finish.

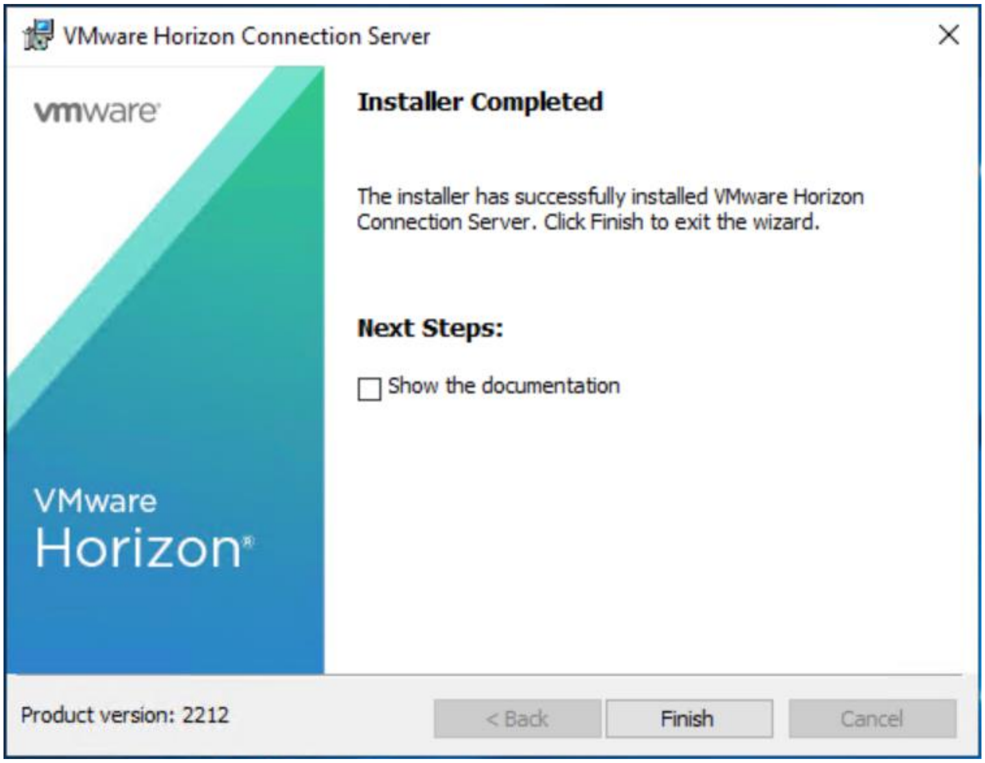

VMware Horizon Connection Server

Installer Completed

The installer has successfully installed VMware Horizon Connection Server. Click Finish to exit the wizard.

Next Steps:

□ Show the documentation

Product version: 2212

Procedure 2. Install VMware Horizon Replica Server

Step 1. Click the Connection Server installer based on your Operating System.

Step 2. Click Next.

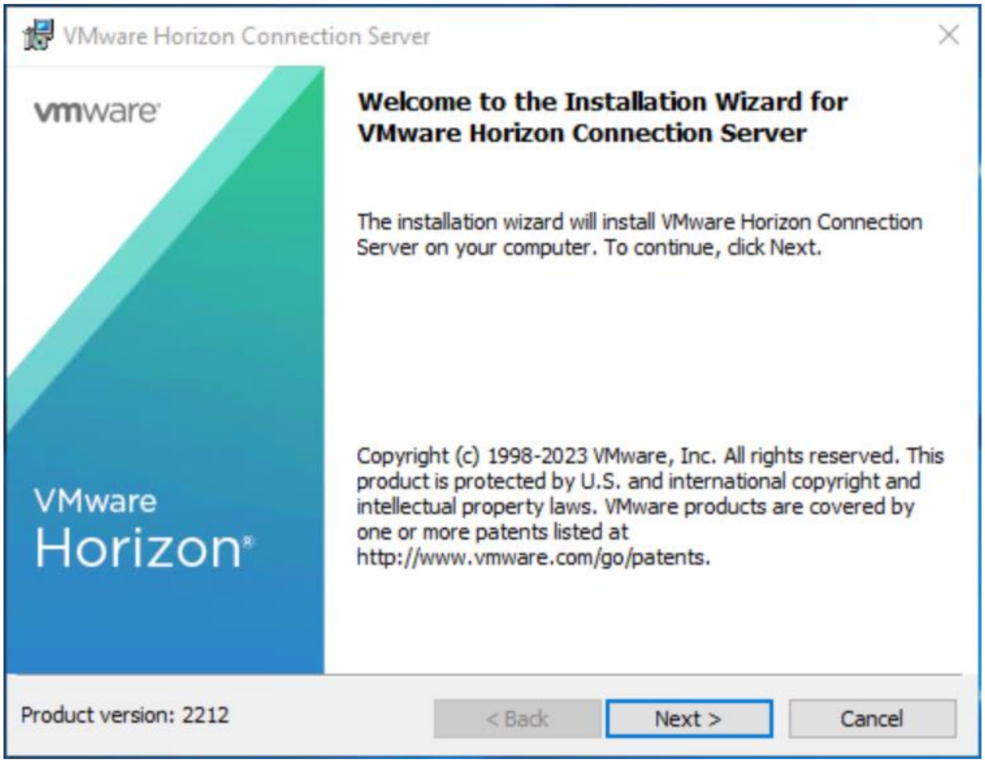

VMware Horizon Connection Server

Welcome to the Installation Wizard for VMware Horizon Connection Server

The installation wizard will install VMware Horizon Connection Server on your computer. To continue, click Next.

Copyright (c) 1998-2023 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents.

Product version: 2212

Step 3. Read and accept the End User License Agreement and click Next.

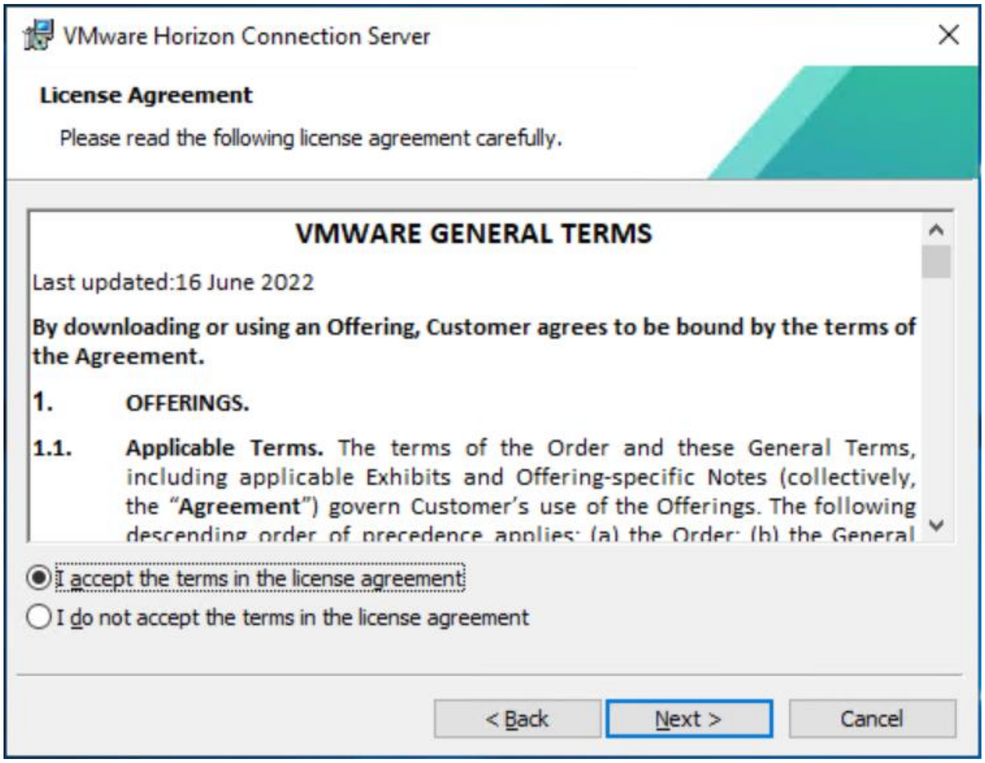

VMware Horizon Connection Server

License Agreement

Please read the following license agreement carefully.

VMWARE GENERAL TERMS

Last updated:16 June 2022

By downloading or using an Offering, Customer agrees to be bound by the terms of the Agreement.

1. OFFERINGS.

1.1. Applicable Terms. The terms of the Order and these General Terms, including applicable Exhibits and Offering-specific Notes (collectively, the "**Agreement**") govern Customer's use of the Offerings. The following descending order of precedence applies: (a) the Order; (b) the General

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 4. Select the destination folder where you want to install the application and click Next.

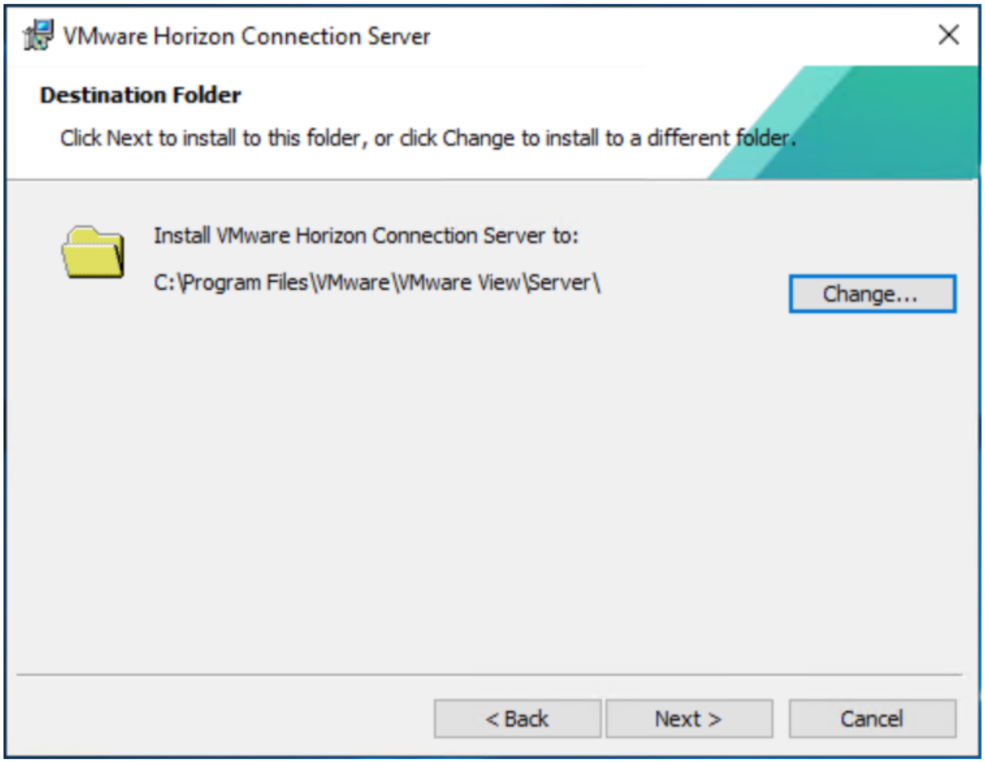

VMware Horizon Connection Server

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install VMware Horizon Connection Server to:

C:\Program Files\VMware\VMware View\Server\

Step 5. Select the Replica Server and IPv4 for the IP protocol version.

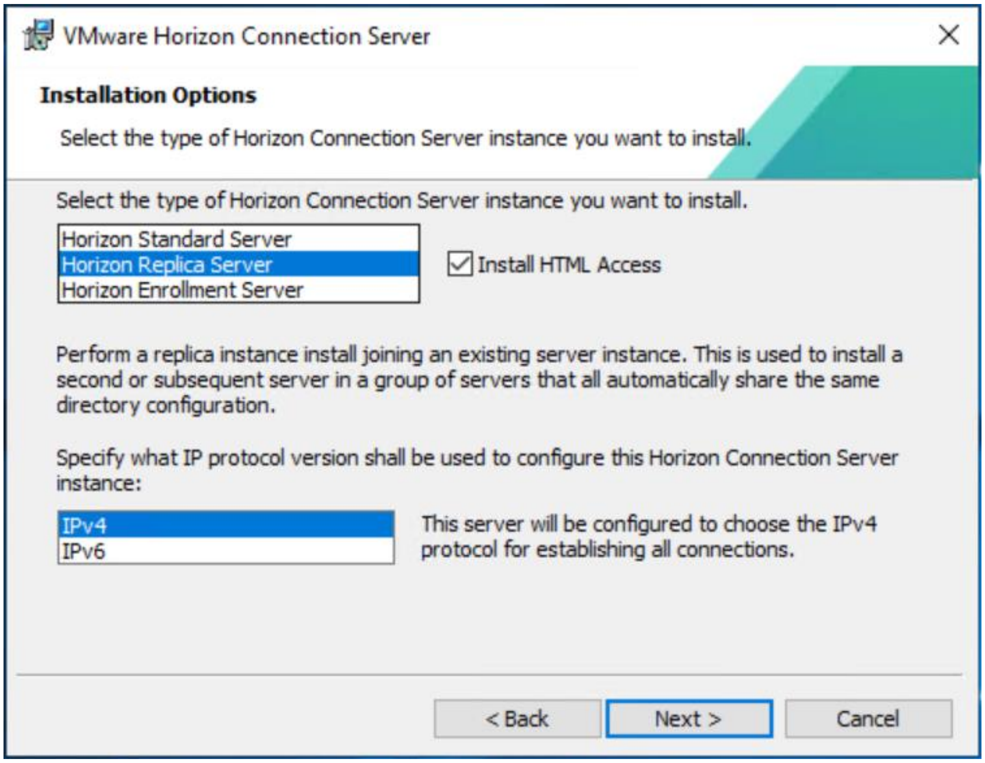

VMware Horizon Connection Server

Installation Options

Select the type of Horizon Connection Server instance you want to install.

Select the type of Horizon Connection Server instance you want to install.

- Horizon Standard Server
- Horizon Replica Server
- Horizon Enrollment Server

☑ Install HTML Access

Perform a replica instance install joining an existing server instance. This is used to install a second or subsequent server in a group of servers that all automatically share the same directory configuration.

Specify what IP protocol version shall be used to configure this Horizon Connection Server instance:

- IPv4
- IPv6

This server will be configured to choose the IPv4 protocol for establishing all connections.

Step 6. Provide the existing Standard View Connection Server's FQDN or IP address and click Next.

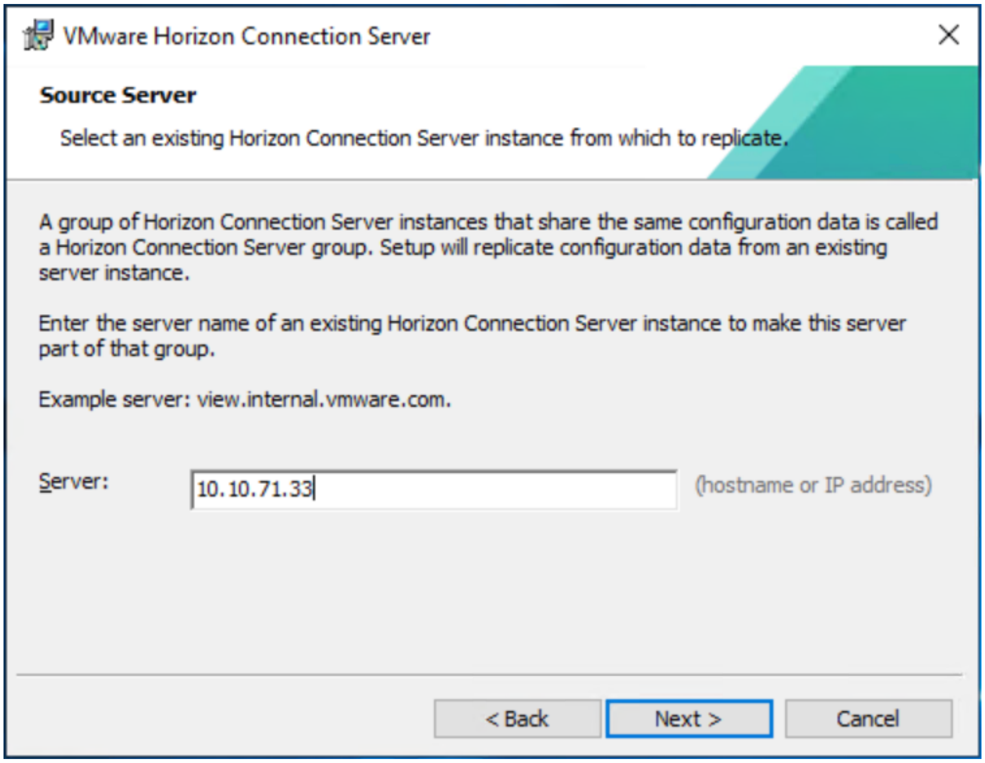

VMware Horizon Connection Server

Source Server

Select an existing Horizon Connection Server instance from which to replicate.

A group of Horizon Connection Server instances that share the same configuration data is called a Horizon Connection Server group. Setup will replicate configuration data from an existing server instance.

Enter the server name of an existing Horizon Connection Server instance to make this server part of that group.

Example server: view.internal.vmware.com.

Server: 10.10.71.33 (hostname or IP address)

Step 7. Select Configure the Windows Firewall automatically.

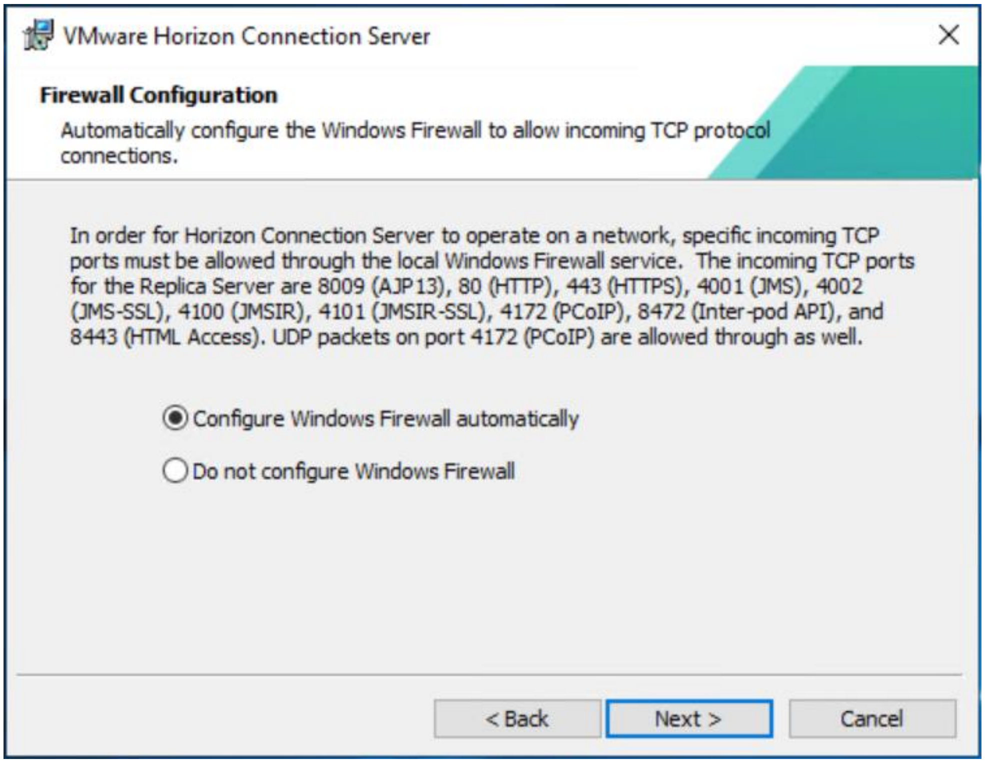

Firewall Configuration

Automatically configure the Windows Firewall to allow incoming TCP protocol connections.

In order for Horizon Connection Server to operate on a network, specific incoming TCP ports must be allowed through the local Windows Firewall service. The incoming TCP ports for the Replica Server are 8009 (AJP13), 80 (HTTP), 443 (HTTPS), 4001 (JMS), 4002 (JMS-SSL), 4100 (JMSIR), 4101 (JMSIR-SSL), 4172 (PCoIP), 8472 (Inter-pod API), and 8443 (HTML Access). UDP packets on port 4172 (PCoIP) are allowed through as well.

● Configure Windows Firewall automatically

○ Do not configure Windows Firewall

Step 8. Click Install to begin the installation process.

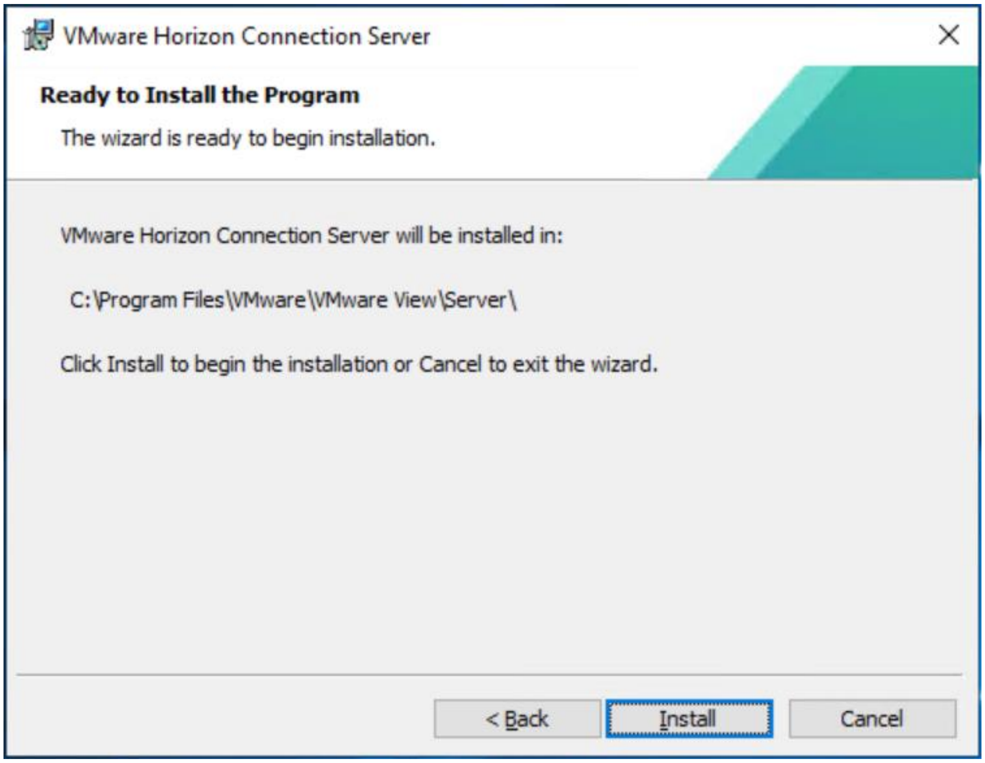

Ready to Install the Program

The wizard is ready to begin installation.

VMware Horizon Connection Server will be installed in:

C:\Program Files\VMware\VMware View\Server\

Click Install to begin the installation or Cancel to exit the wizard.

Step 9. After installation is complete, click Finish.

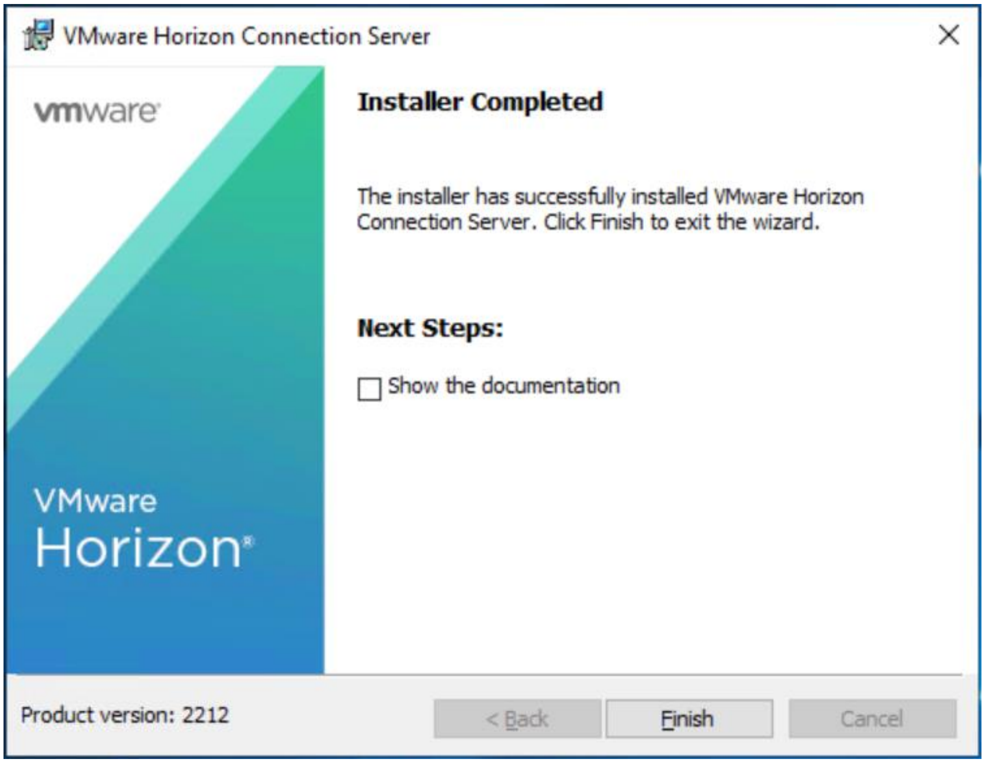

Installer Completed

The installer has successfully installed VMware Horizon Connection Server. Click Finish to exit the wizard.

Next Steps:

□ Show the documentation

Product version: 2212

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.

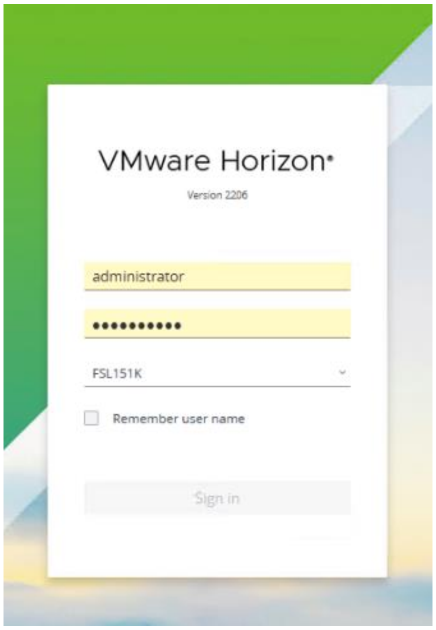

administrator

••••••••••

FSL151K

Step 2. In Horizon Console, expand Settings and click Servers.

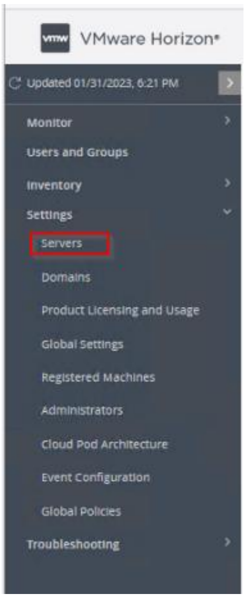

Step 3. Select the vCenter Settings tab and click Add.

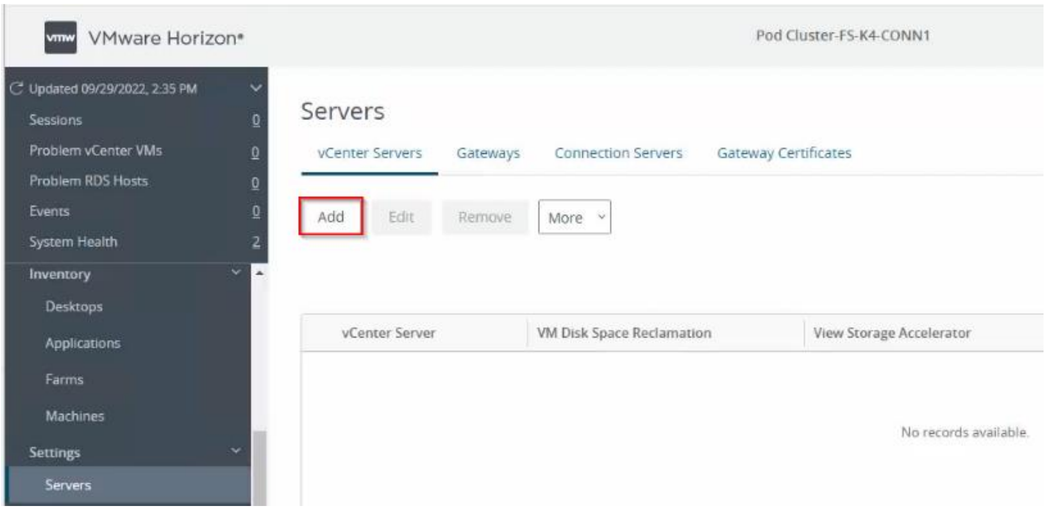

Step 4. Provide Server Address (IP or FQDN) and credentials that Horizon will use to login to vCenter, then click Next.

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.

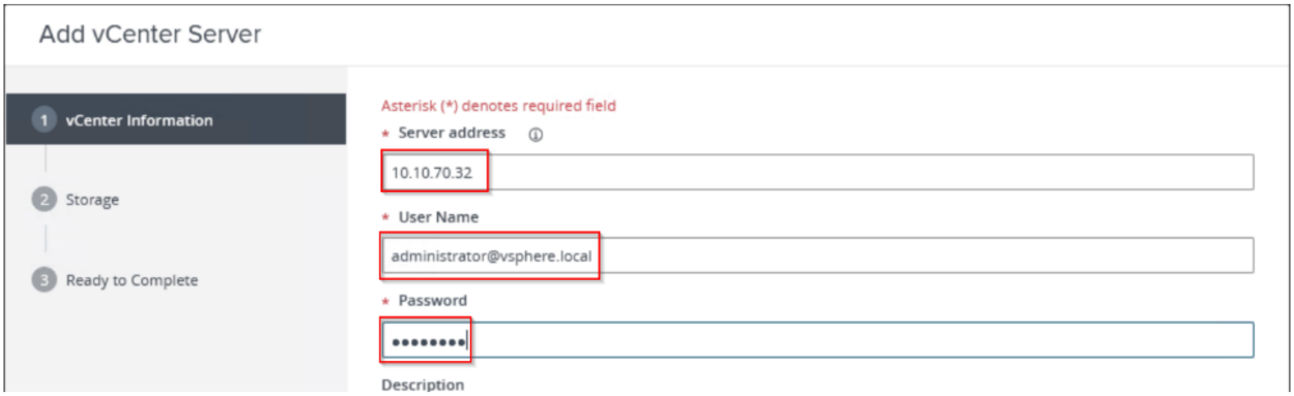

Add vCenter Server

Asterisk (*) denotes required field

* Server address: 10.10.70.32

* User Name: administrator@vsphere.local

* Password: ••••••••

Description

Step 5. If you receive a message stating an invalid certificate, click View Certificate.

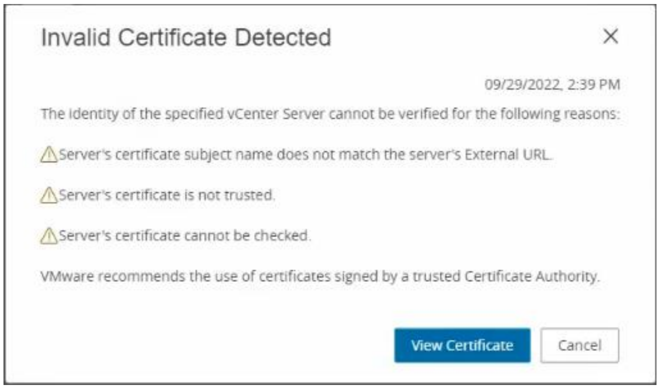

Invalid Certificate Detected

09/29/2022, 2:39 PM

The identity of the specified vCenter Server cannot be verified for the following reasons:

⚠Server's certificate subject name does not match the server's External URL.

⚠Server's certificate is not trusted.

⚠Server's certificate cannot be checked.

VMware recommends the use of certificates signed by a trusted Certificate Authority.

Step 6. Click Accept.

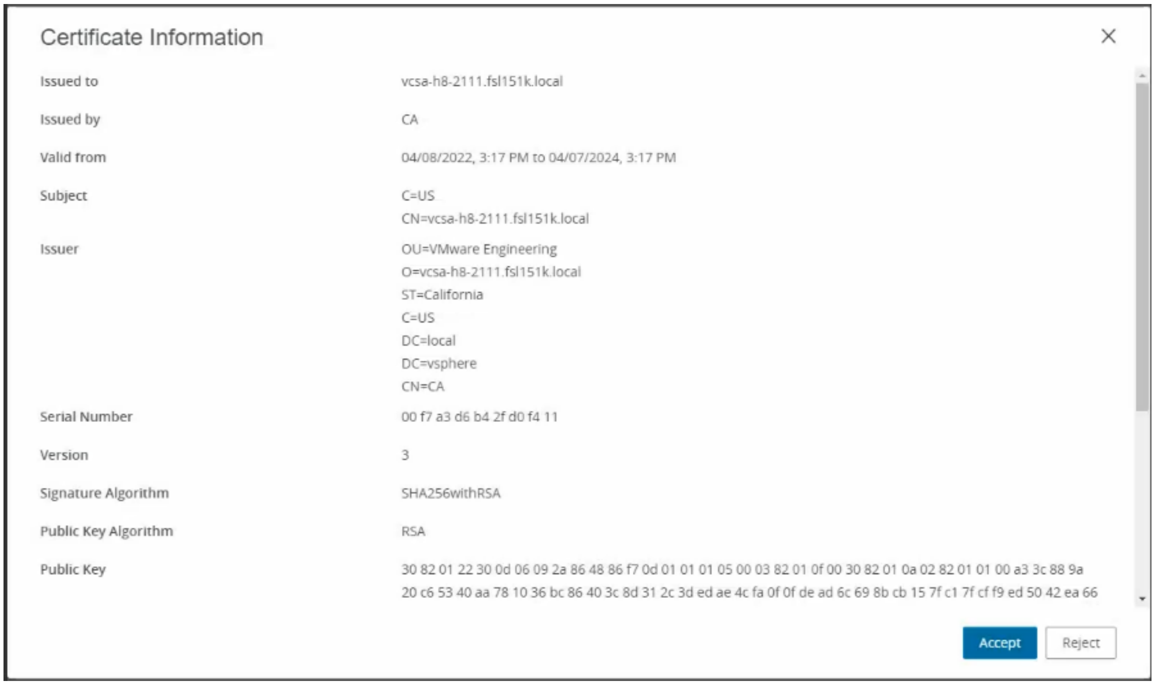

Certificate Information

Issued to	vcsa-h8-2111.fsl151k.local
Issued by	CA
Valid from	04/08/2022, 3:17 PM to 04/07/2024, 3:17 PM

Subject	C=US
	CN=vcsa-h8-2111.fsl151k.local

Issuer	OU=VMware Engineering
	O=vcsa-h8-2111.fsl151k.local
	ST=California
	C=US
	DC=local
	DC=vsphere
	CN=CA

Serial Number	00 f7 a3 d6 b4 2f d0 f4 11
Version	3
Signature Algorithm	SHA256withRSA
Public Key Algorithm	RSA
Public Key	30 82 01 22 30 0d 06 09 2a 86 48 86 f7 0d 01 01 01 05 00 03 82 01 0f 00 30 82 01 0a 02 82 01 01 00 a3 3c 88 9a 20 c6 53 40 aa 78 10 36 bc 86 40 3c 8d 31 2c 3d ed ae 4c fa 0f 0f de ad 6c 69 8b cb 15 7f c1 7f cf f9 ed 50 42 ea 66

Step 7. Keep the defaults, select Reclaim VM disk space and Enable Horizon Storage Accelerator with cache size of 1024MB. Click Next.

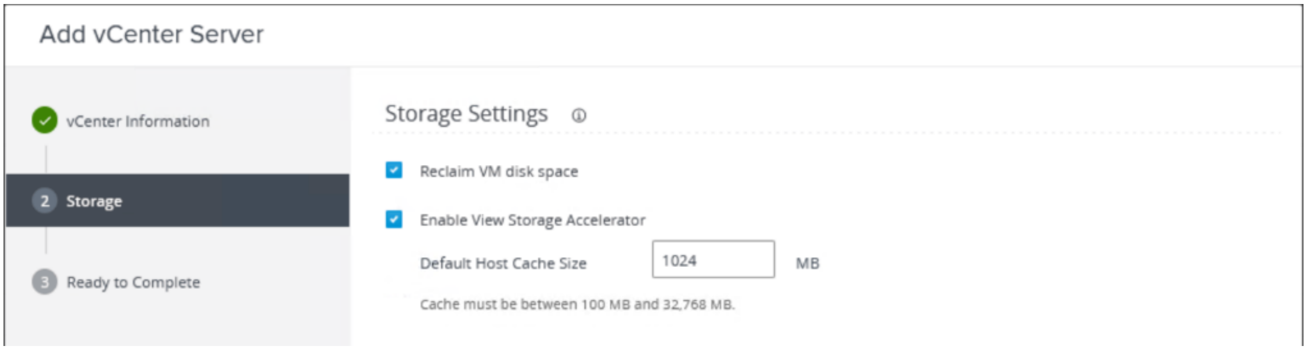

Add vCenter Server

Storage Settings

☑ Reclaim VM disk space

☑ Enable View Storage Accelerator

Default Host Cache Size: 1024 MB

Cache must be between 100 MB and 32,768 MB.

Step 8. Review the information you provided and click Submit.

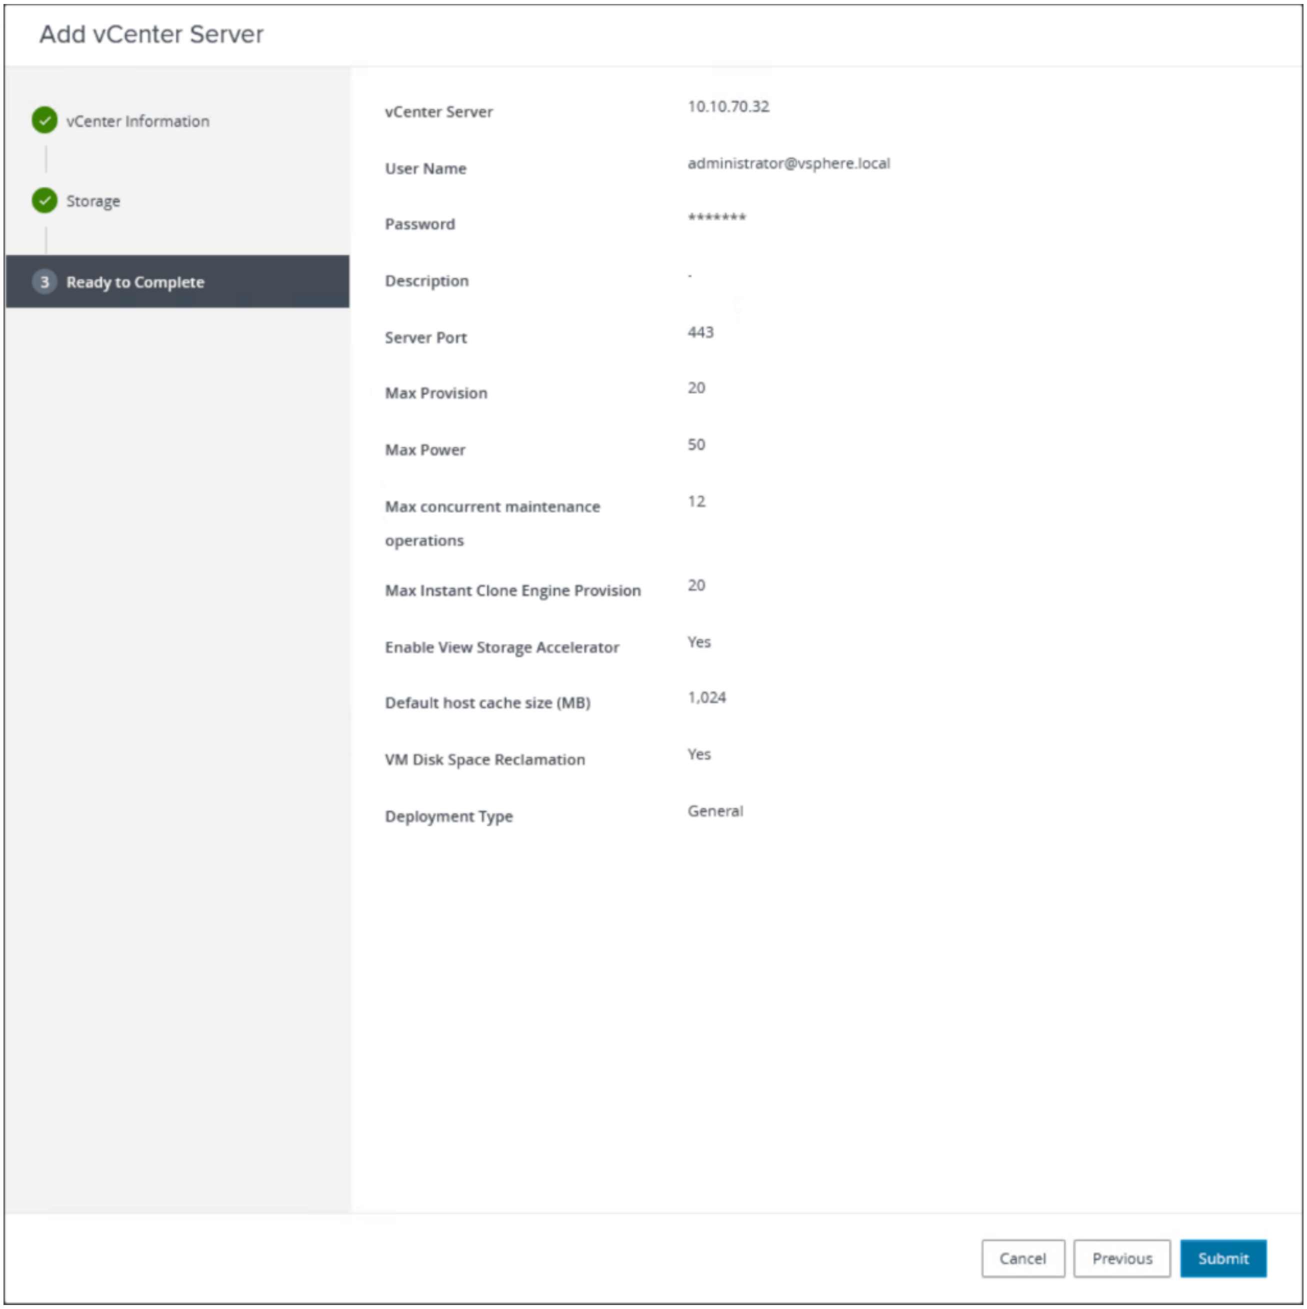

Add vCenter Server

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.

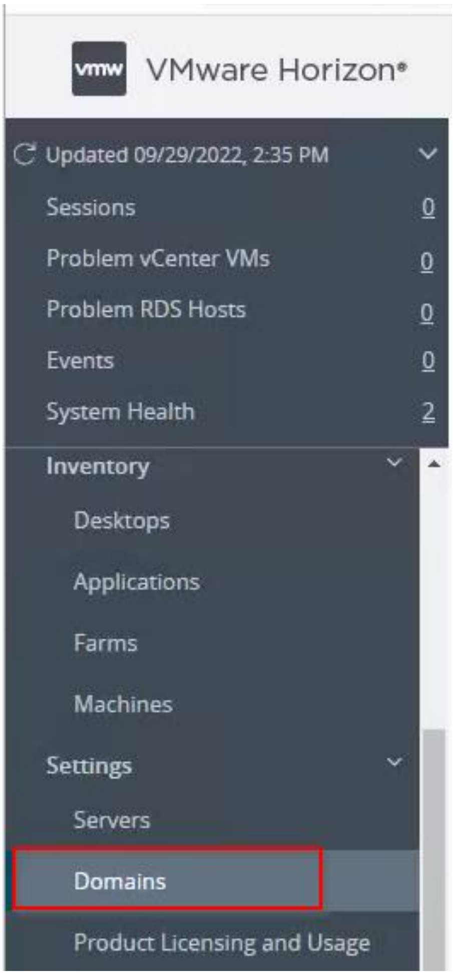

Step 10. Select the Instant Clone Engine Domain Accounts tab and click Add.

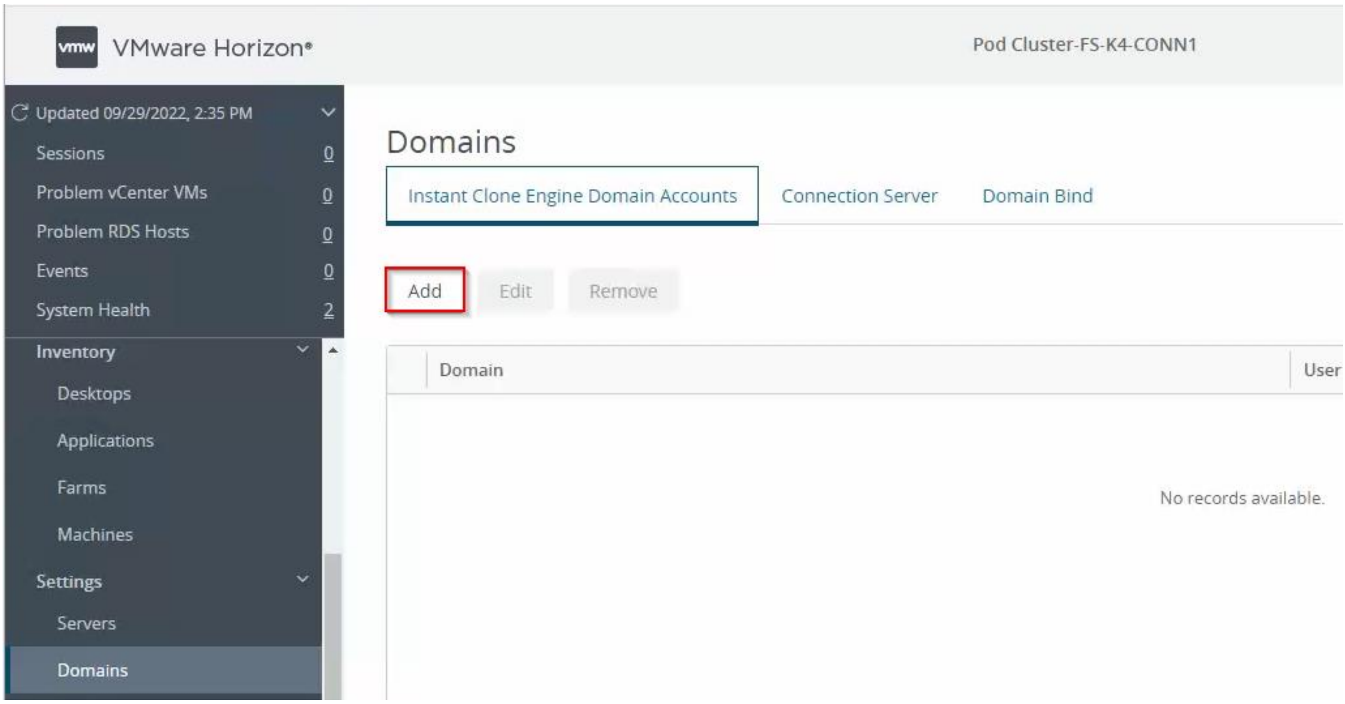

Step 11. Provide a domain name and credentials that Horizon will use to login to AD during Instant Clone management tasks, then click OK.

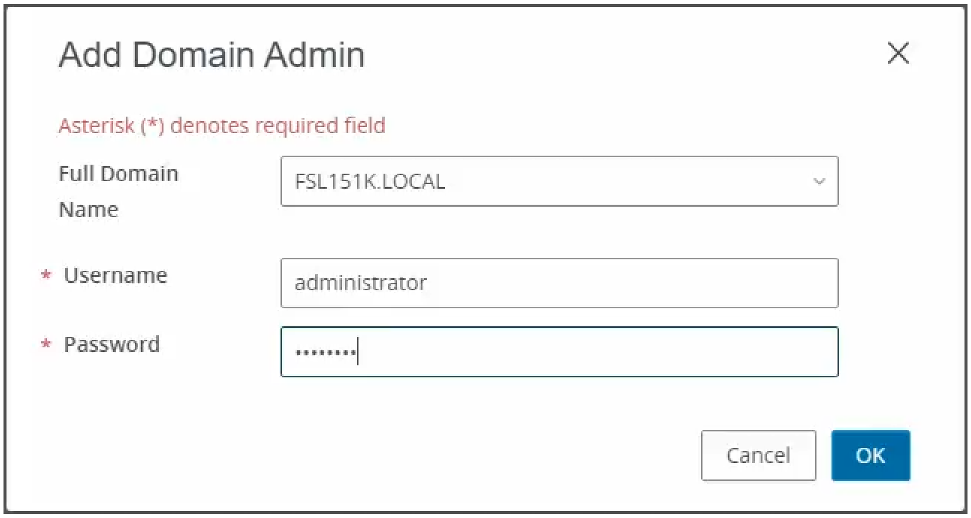

Full Domain Name: FSL151K.LOCAL

* Username: administrator

* Password: ••••••••

Procedure 2. Create VDI Instant Clone Desktop Pool

Step 1. In Horizon Console on the left plane, expand Inventory, select Desktops. Click Add.

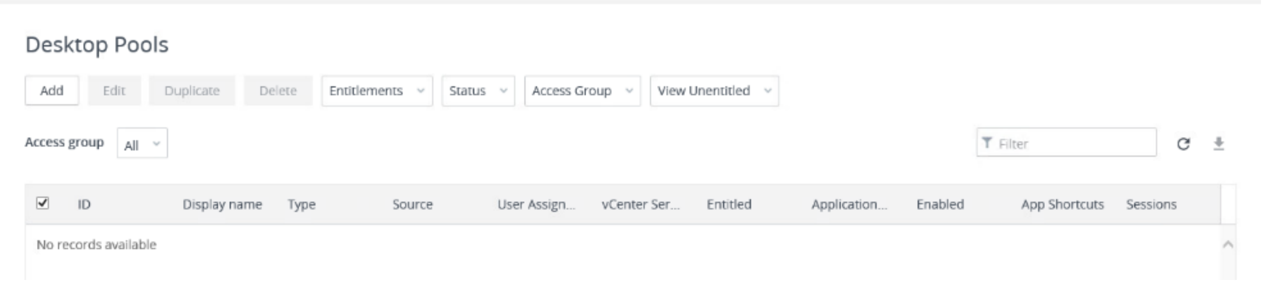

☑	ID	Display name	Type	Source	User Assign...	vCenter Ser...	Entitled	Application...	Enabled	App Shortcuts	Sessions
No records available											

Step 2. Select Type of Desktop pool to be created. Click Next.

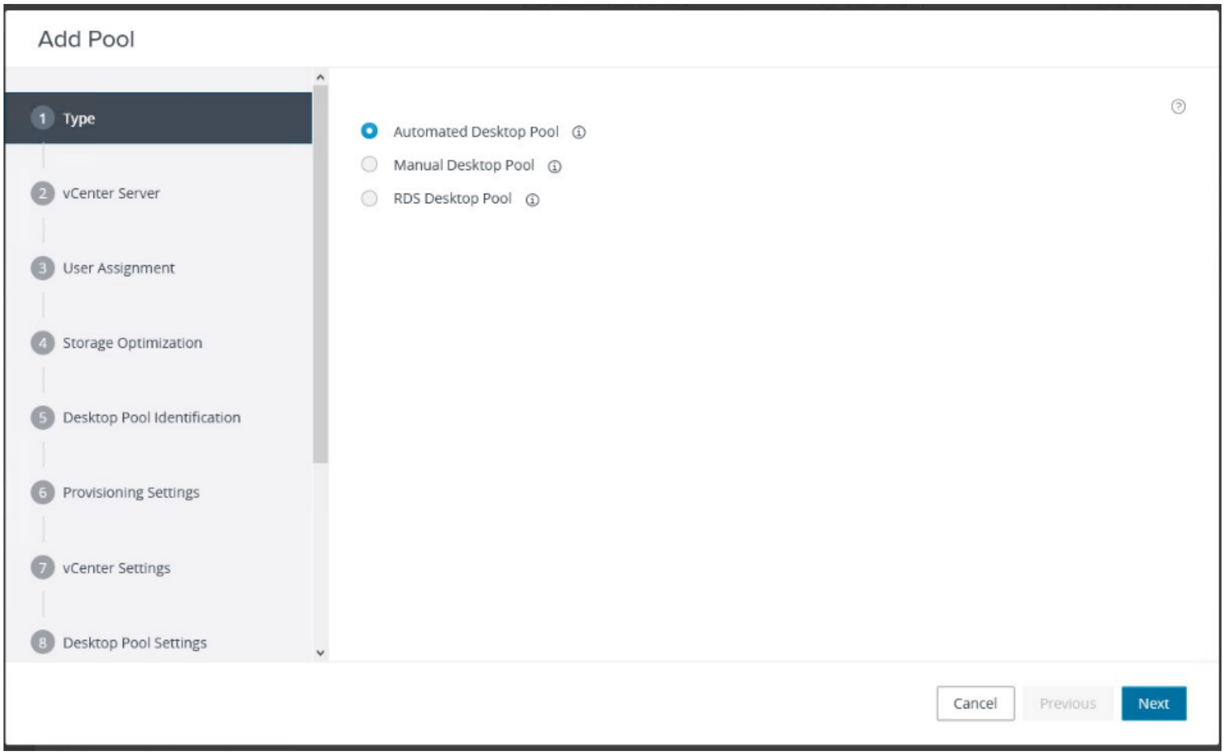

Add Pool

○ Automated Desktop Pool (selected)
○ Manual Desktop Pool
○ RDS Desktop Pool

Step 3. Select the provisioning type as Instant Clones for the desktops in the pool. Click Next.

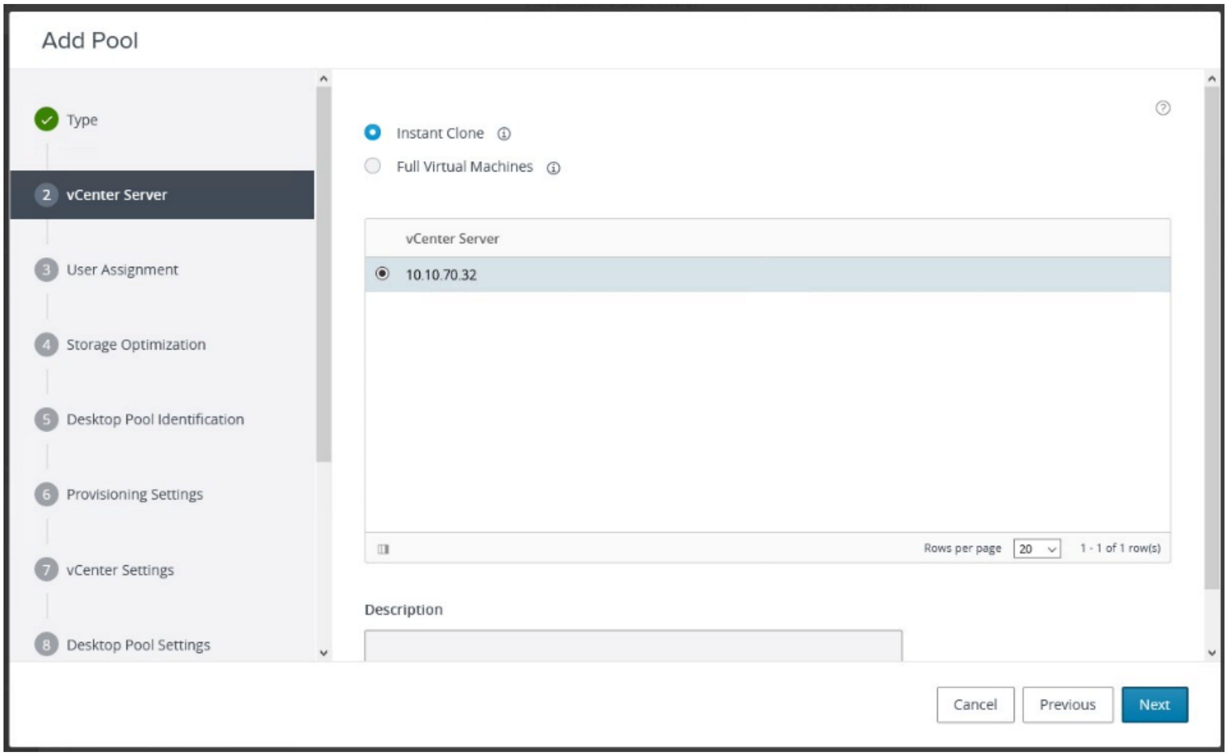

Add Pool

● Instant Clone
○ Full Virtual Machines

	vCenter Server
●	10.10.70.32

Rows per page 20 1 - 1 of 1 row(s)

Description

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.

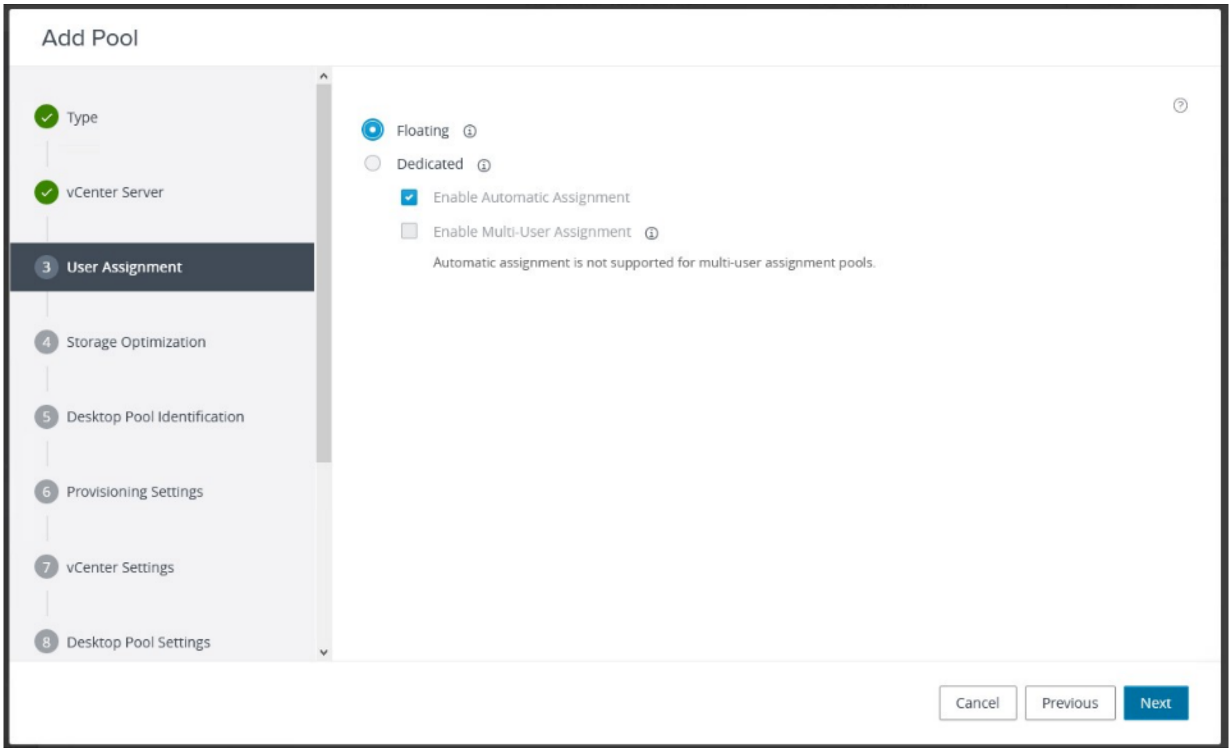

Add Pool

● Floating

○ Dedicated

☑ Enable Automatic Assignment

□ Enable Multi-User Assignment

Automatic assignment is not supported for multi-user assignment pools.

Step 5. On Storage Optimization screen, pick Use VMware Virtual SAN as Storage Policy Management, and click Next.

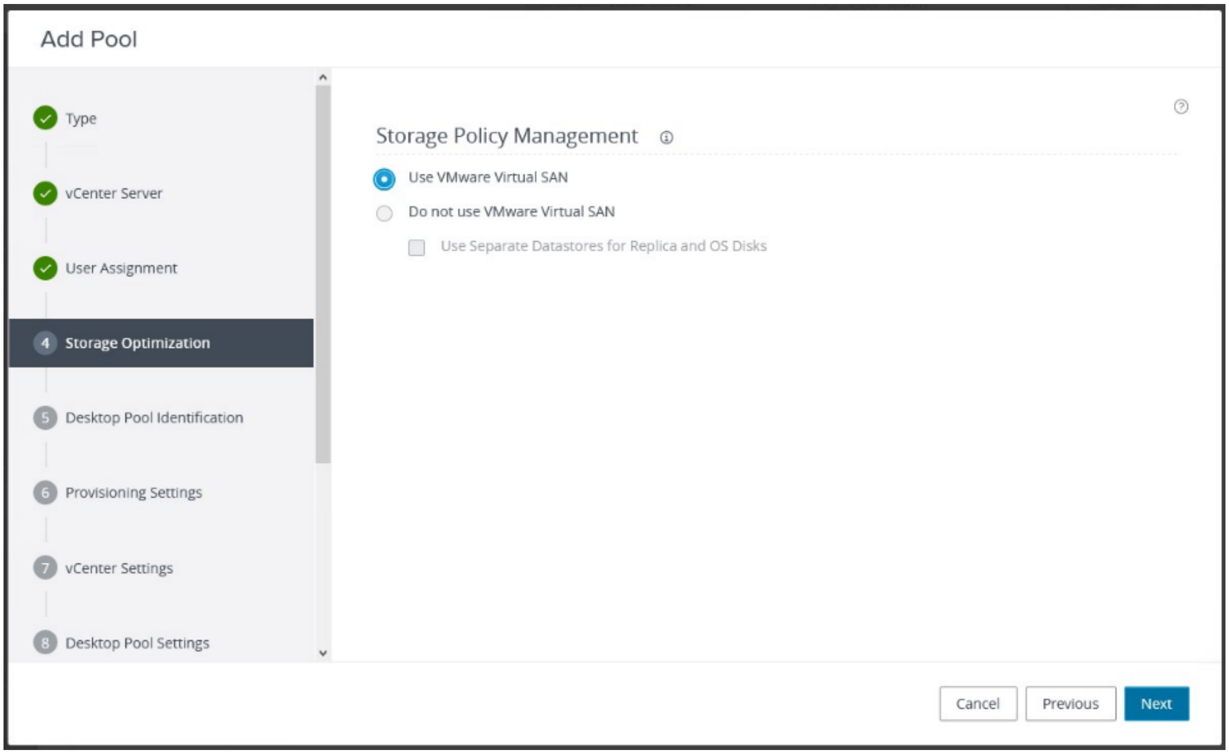

Add Pool

Storage Policy Management

● Use VMware Virtual SAN

○ Do not use VMware Virtual SAN

□ Use Separate Datastores for Replica and OS Disks

Step 6. Provide Desktop Pool ID and virtual display name. Click Next.

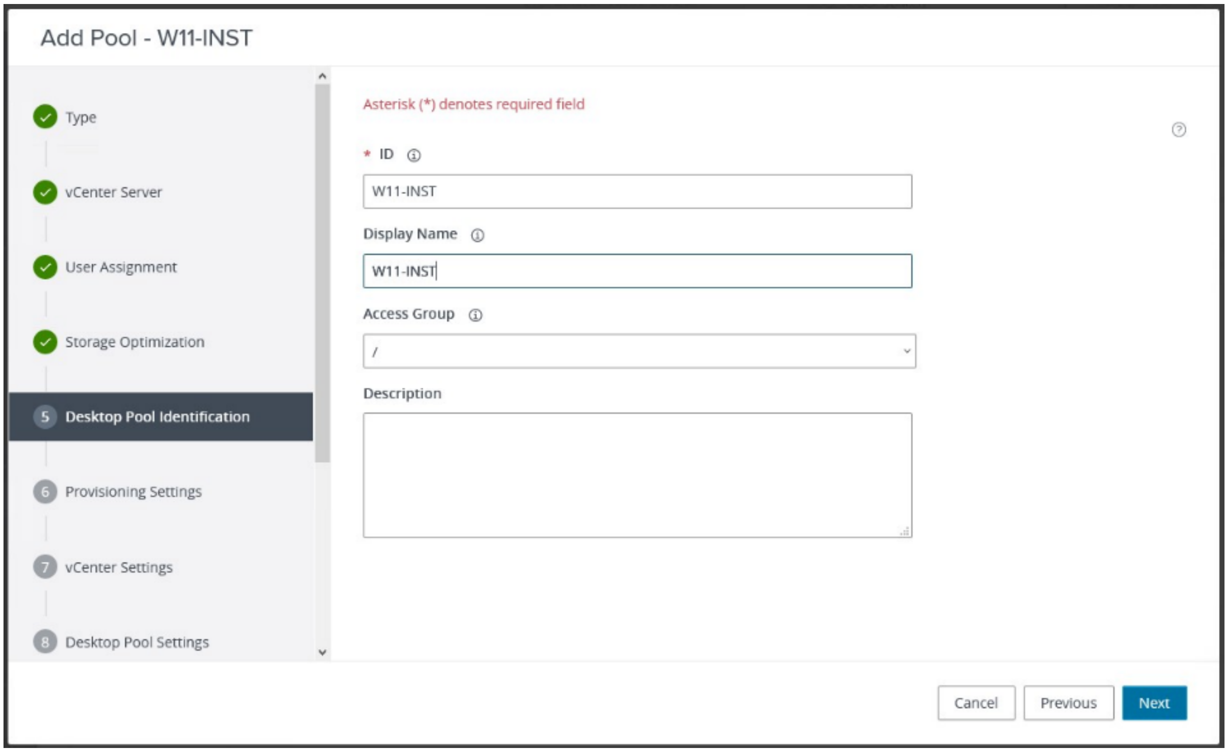

Add Pool - W11-INST

Asterisk (*) denotes required field

* ID: W11-INST

Display Name: W11-INST

Access Group: /

Description:

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 pool – 195
Cluster pool – 585

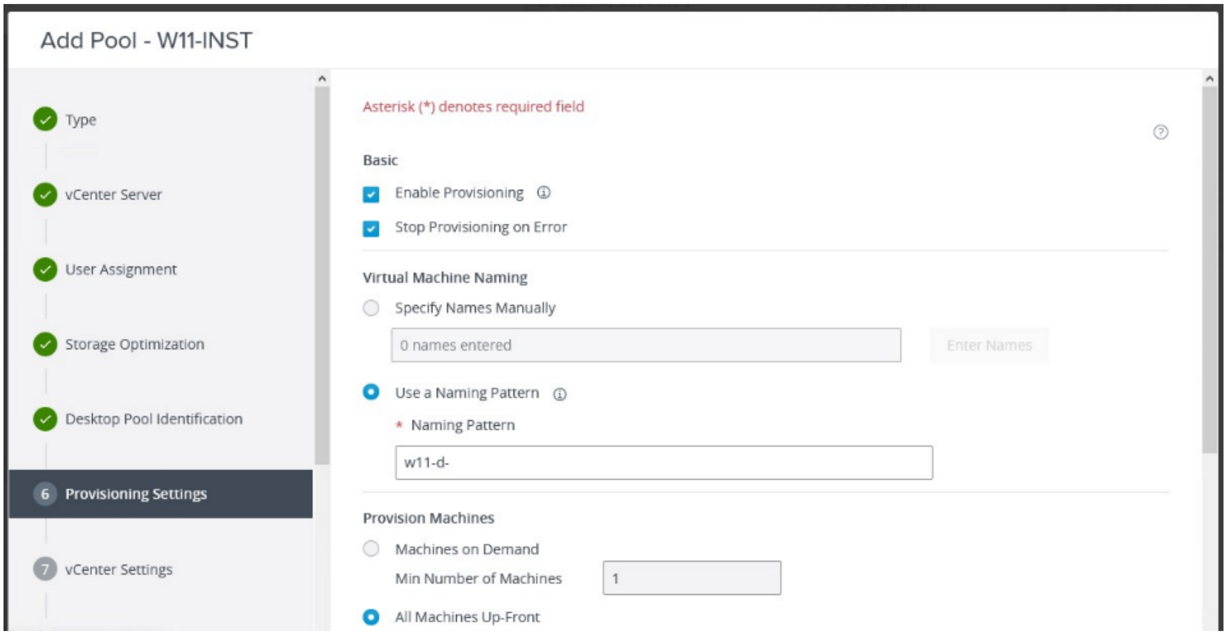

Add Pool - W11-INST

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-d-

Provision Machines

○ Machines on Demand

Min Number of Machines: 1

● All Machines Up-Front

Step 8. Provide the parent VM, snapshot and host/cluster info, and data store information for the virtual machines to create. Click Next.

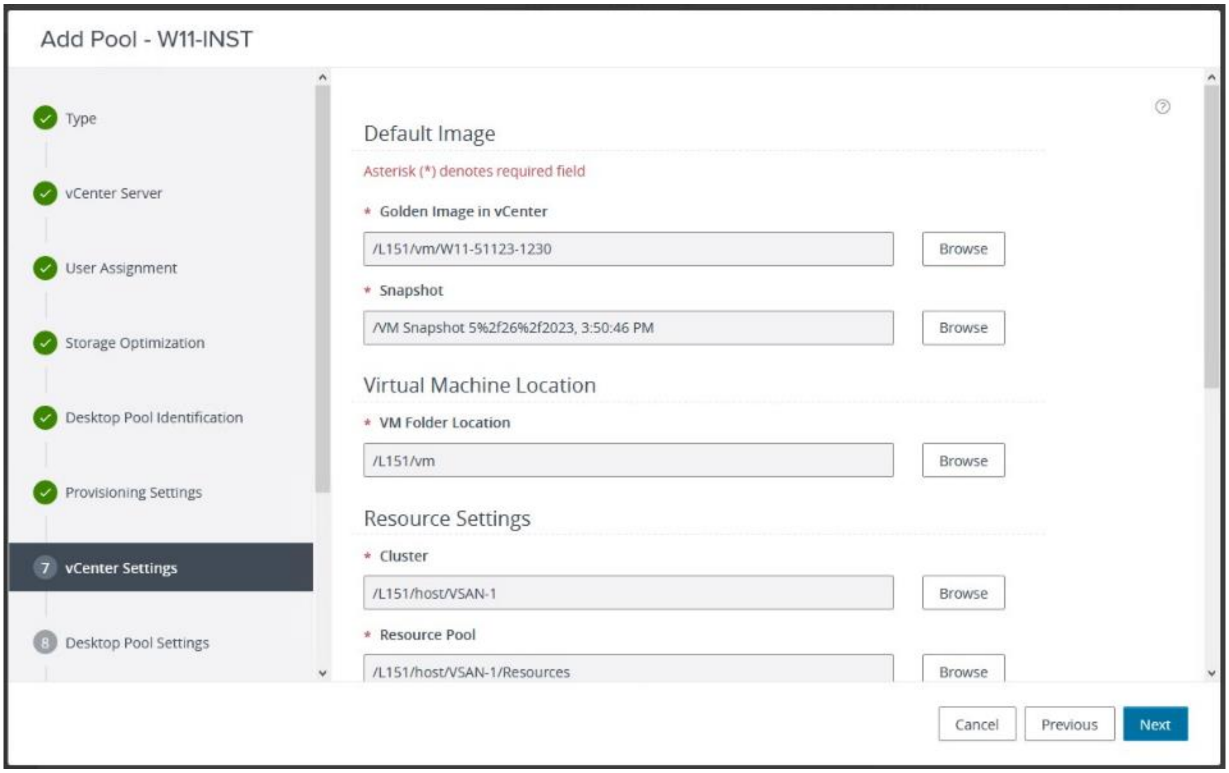

Default Image

Asterisk (*) denotes required field

* Golden Image in vCenter: /L151/vm/W11-51123-1230

* Snapshot: /VM Snapshot 5%2f26%2f2023, 3:50:46 PM

Virtual Machine Location

* VM Folder Location: /L151/vm

Resource Settings

* Cluster: /L151/host/VSAN-1

* Resource Pool: /L151/host/VSAN-1/Resources

Step 9. Configure the State and Session Type for Desktop Pool Settings. Click Next.

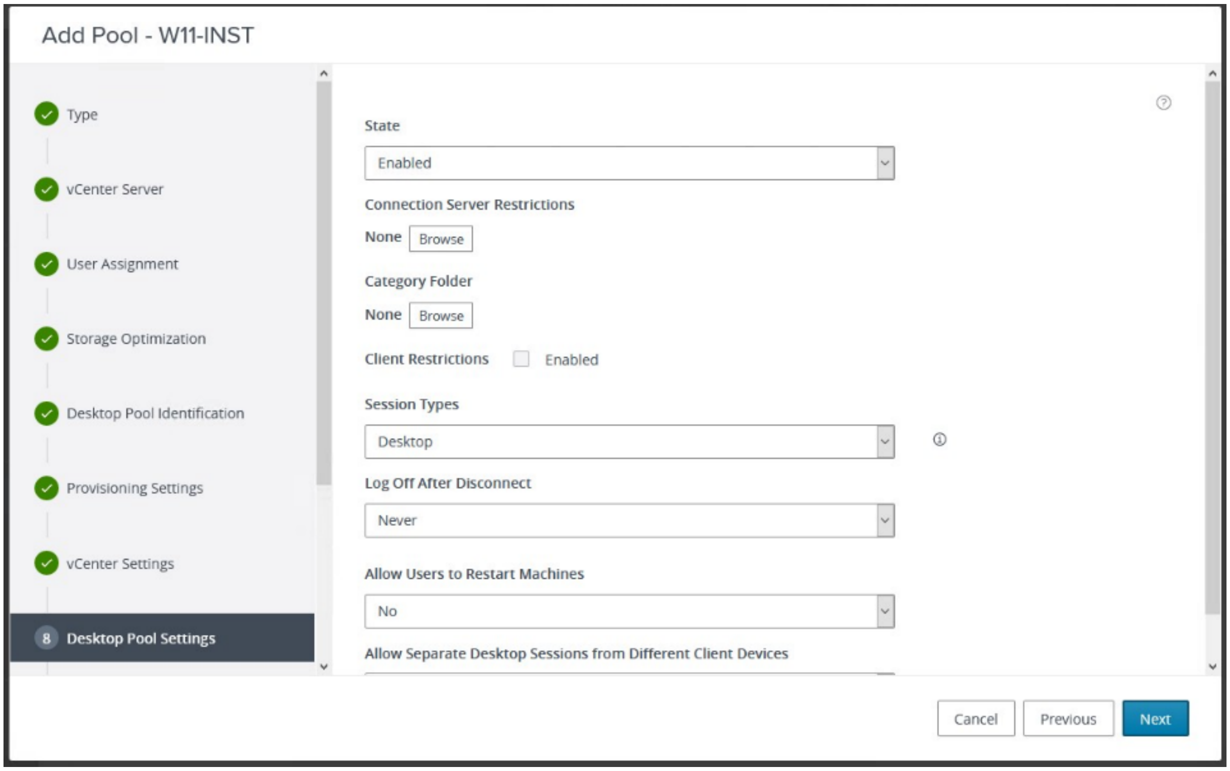

State: Enabled

Connection Server Restrictions: None

Category Folder: None

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

Step 10. Configure the Remote Display Protocol. Click Next.

Note: We used the VMware Blast for the Instant Clone pool.

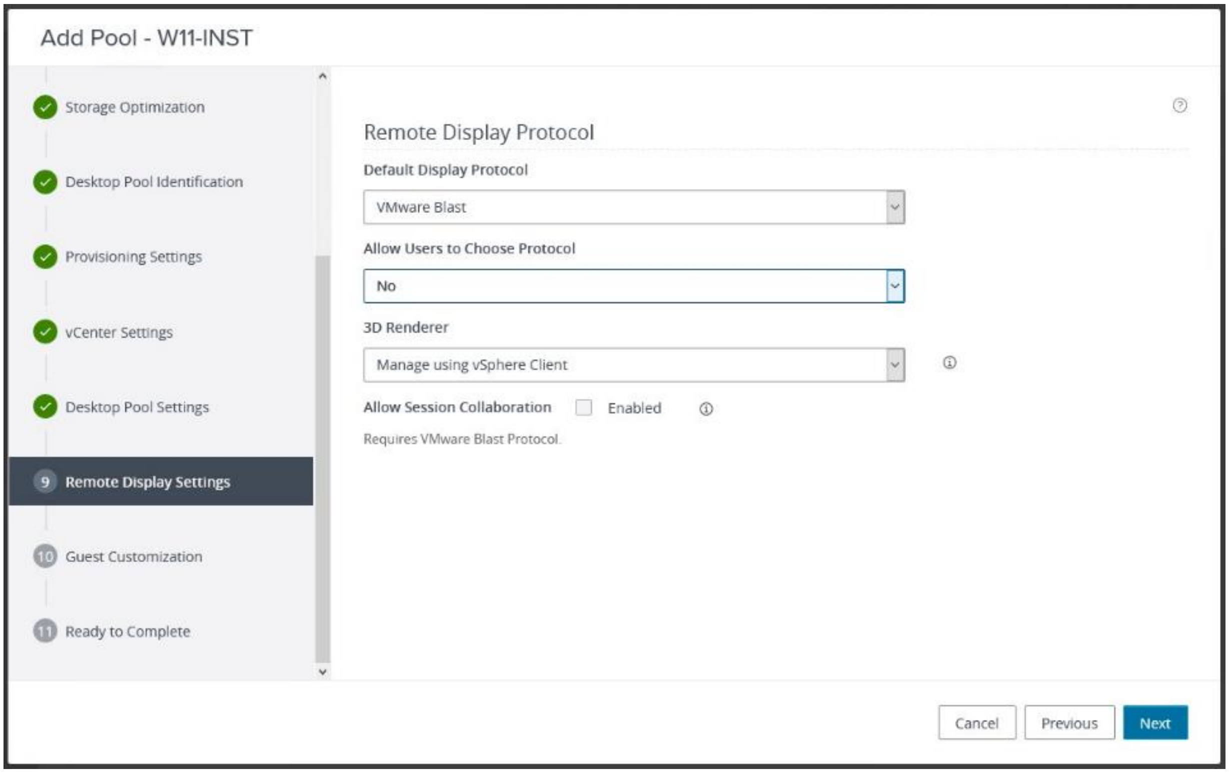

Add Pool - W11-INST

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.

Step 11. Select the AD Container for desktops to place in a Domain Controller computer location.

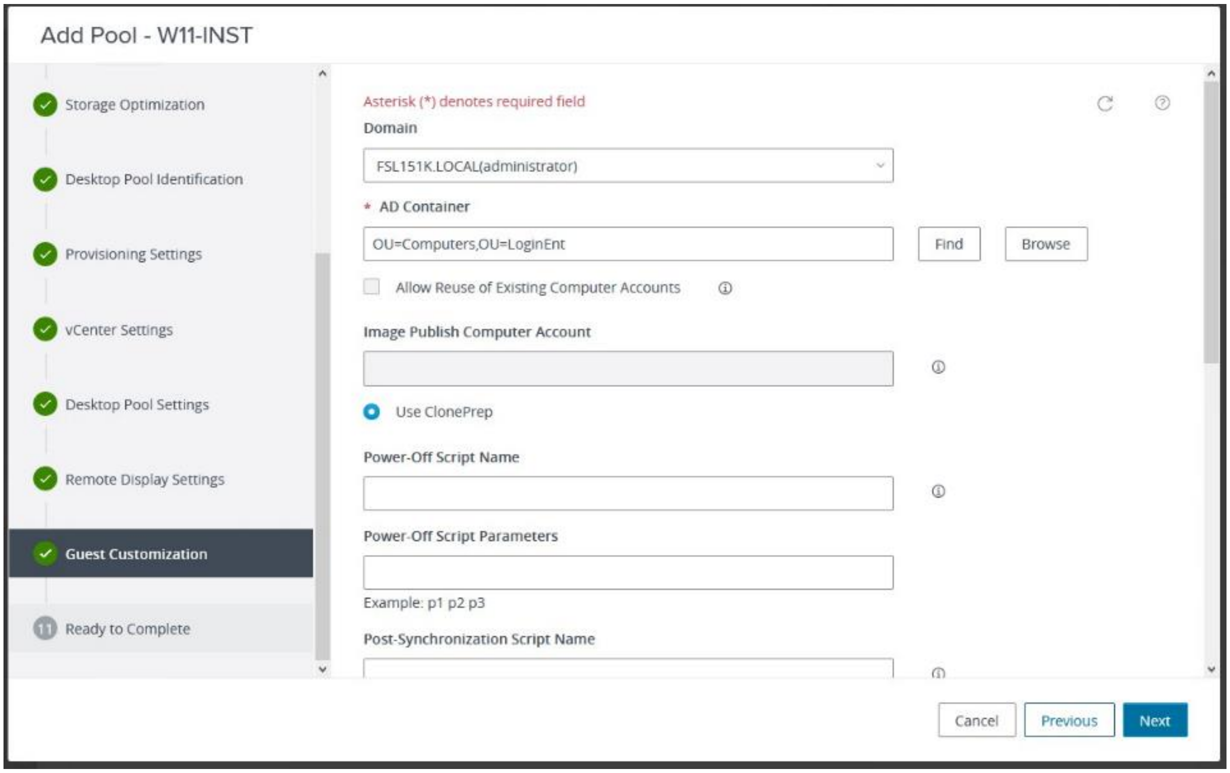

Add Pool - W11-INST

Asterisk (*) denotes required field

Domain: FSL151K.LOCAL(administrator)

* AD Container: OU=Computers,OU=LoginEnt

□ 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

Step 12. Review the deployment specifications and click Submit to complete the deployment.

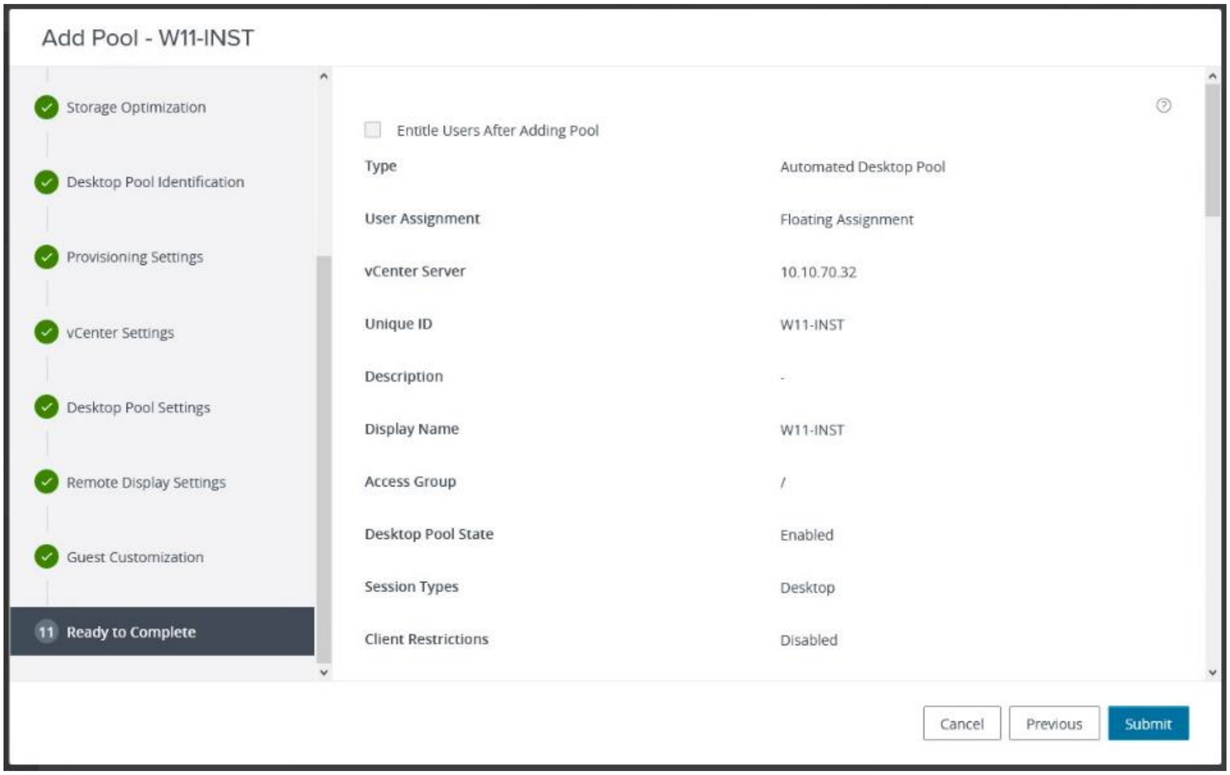

Add Pool - W11-INST

□ Entitle Users After Adding Pool

Type	Automated Desktop Pool
User Assignment	Floating Assignment
vCenter Server	10.10.70.32
Unique ID	W11-INST
Description	-
Display Name	W11-INST
Access Group	/
Desktop Pool State	Enabled
Session Types	Desktop
Client Restrictions	Disabled

Step 13. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

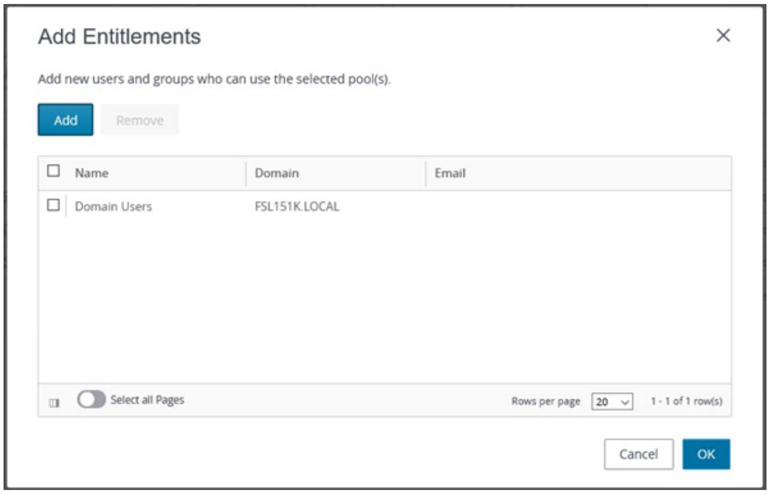

Add Entitlements

Add new users and groups who can use the selected pool(s).

□ Name	Domain	Email
□ Domain Users	FSL151K.LOCAL	

Procedure 3. Create VDI Full Clone Desktop Pool

Step 1. Select Type of Desktop pool to be created. Click Next.

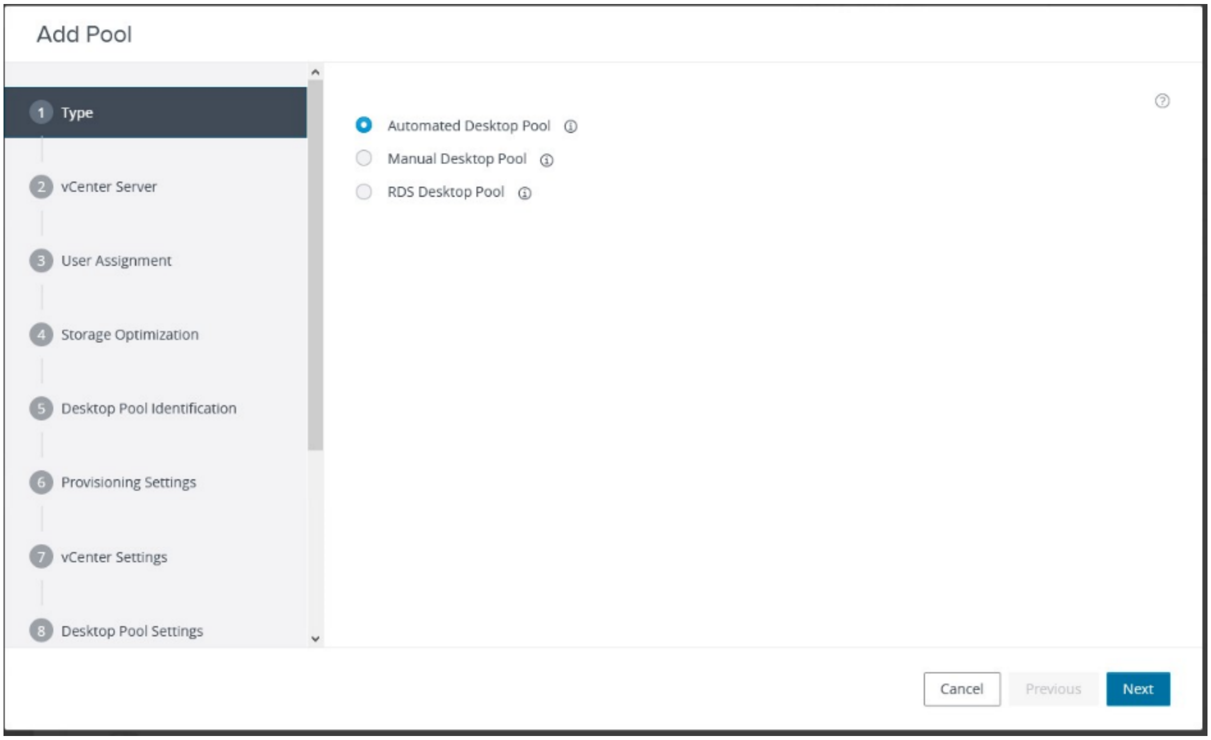
Add Pool

1 Type
2 vCenter Server
3 User Assignment
4 Storage Optimization
5 Desktop Pool Identification
6 Provisioning Settings
7 vCenter Settings
8 Desktop Pool Settings

● Automated Desktop Pool
○ Manual Desktop Pool
○ RDS Desktop Pool

Step 2. Select the provisioning type as Full Virtual Machines for the desktops in the pool. Click Next.

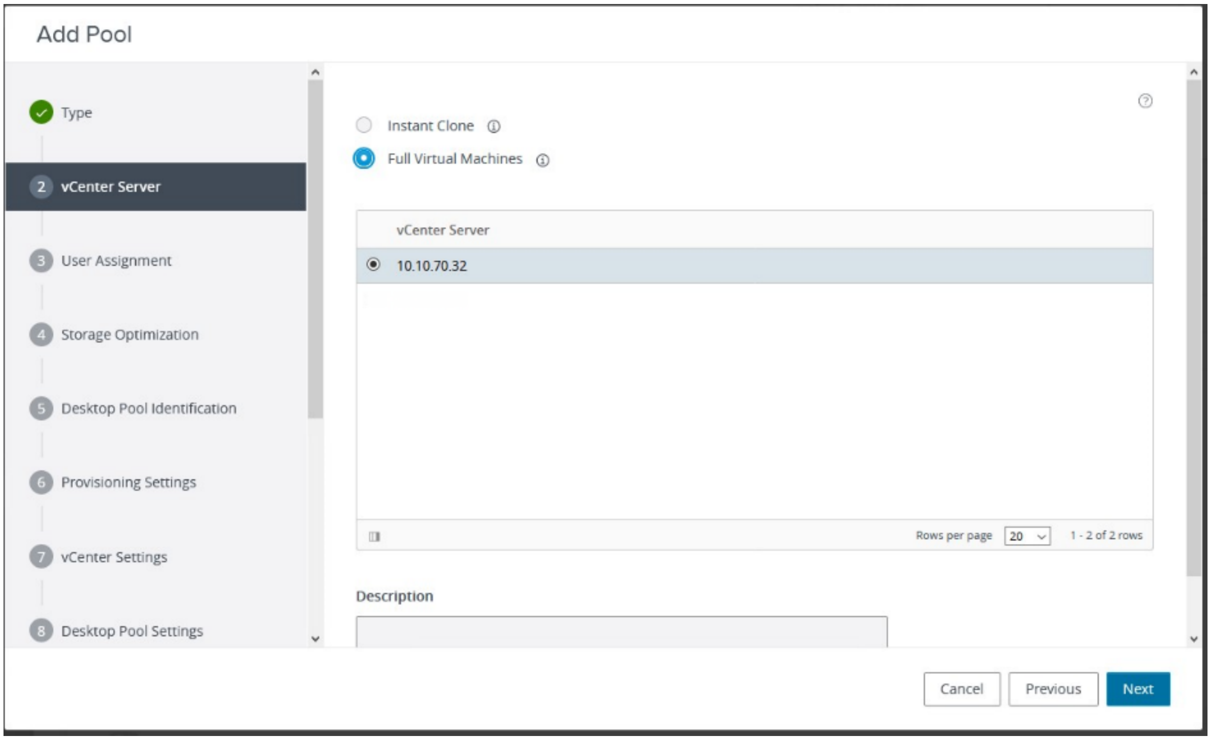
Add Pool

Type
2 vCenter Server
3 User Assignment
4 Storage Optimization
5 Desktop Pool Identification
6 Provisioning Settings
7 vCenter Settings
8 Desktop Pool Settings

○ Instant Clone
● Full Virtual Machines

vCenter Server
● 10.10.70.32

Rows per page 20 1 - 2 of 2 rows

Description

Step 3. Select the User assignment to be used by the desktop pool, we used Dedicated assignment for Full Cone pool. Click Next.

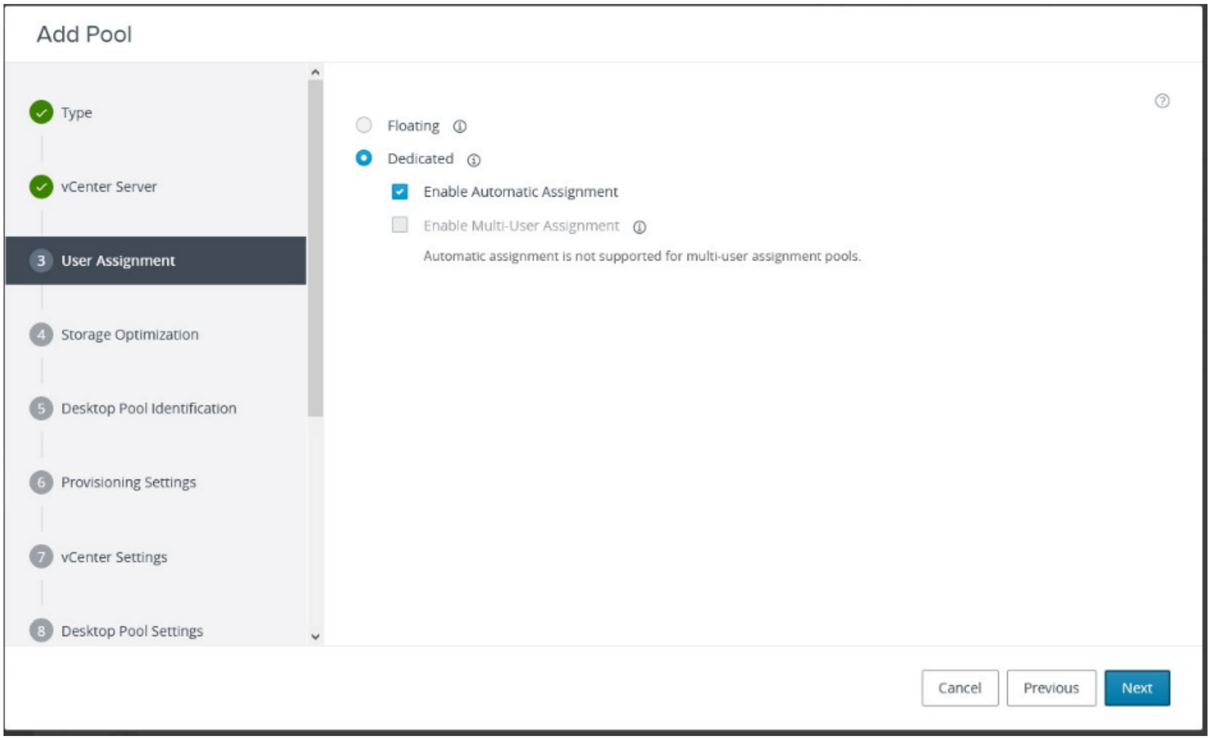

○ Floating

● Dedicated

☑ Enable Automatic Assignment

□ Enable Multi-User Assignment

Automatic assignment is not supported for multi-user assignment pools.

Step 4. On Storage Optimization screen, select Use VMware Virtual SAN for the Storage Policy Management and click Next.

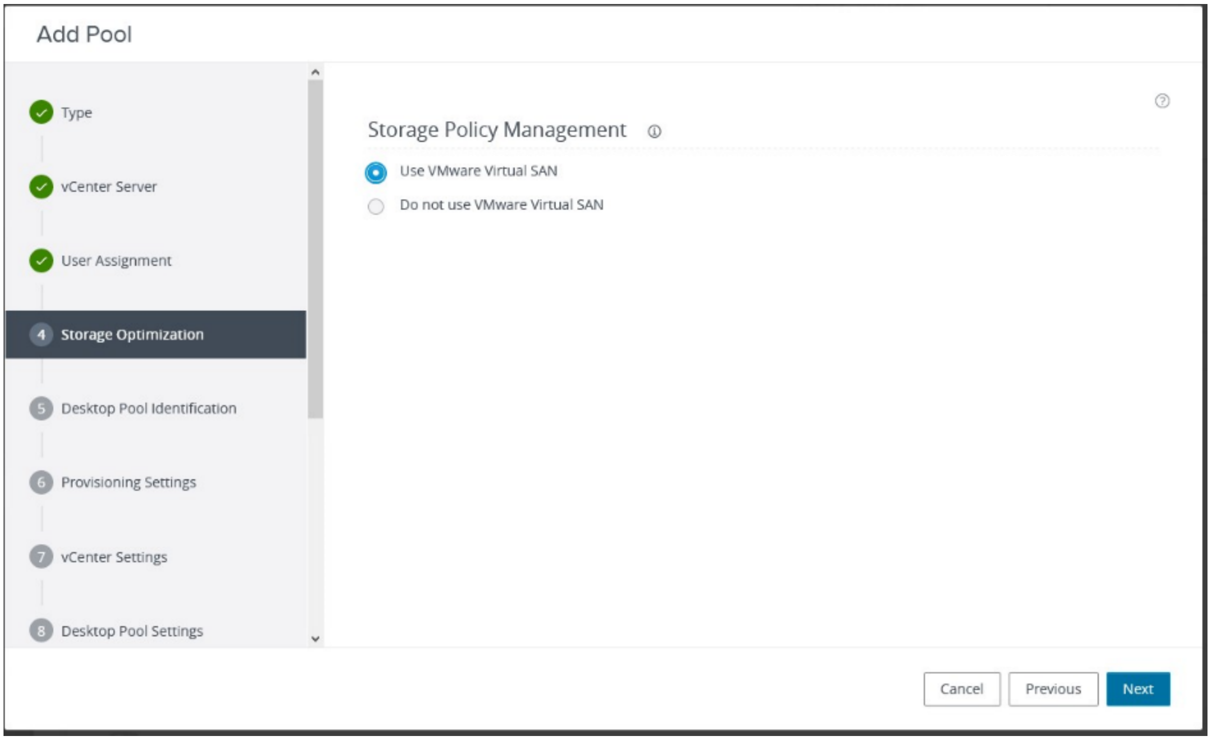

Storage Policy Management

● Use VMware Virtual SAN

○ Do not use VMware Virtual SAN

Step 5. Provide the Desktop Pool ID and Display Name. Click Next.

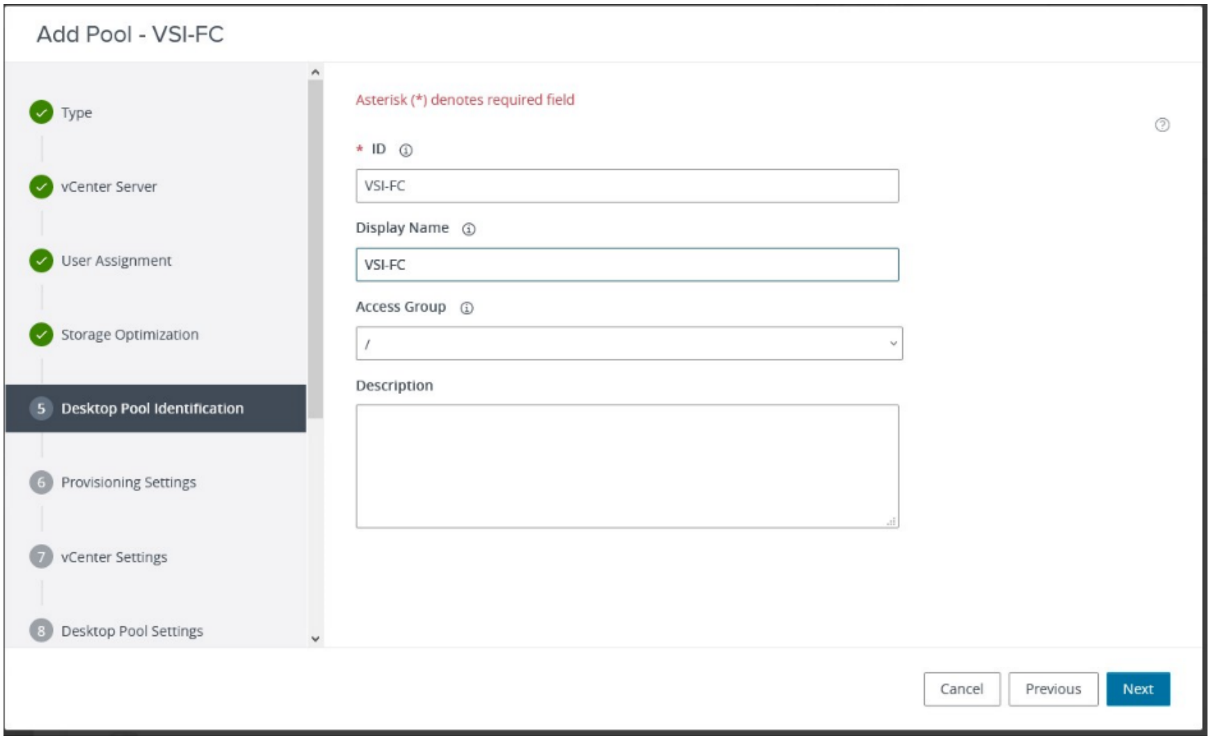

Add Pool - VSI-FC

Asterisk (*) denotes required field

* ID: VSI-FC

Display Name: VSI-FC

Access Group: /

Description:

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 pool – 195
Cluster pool – 585

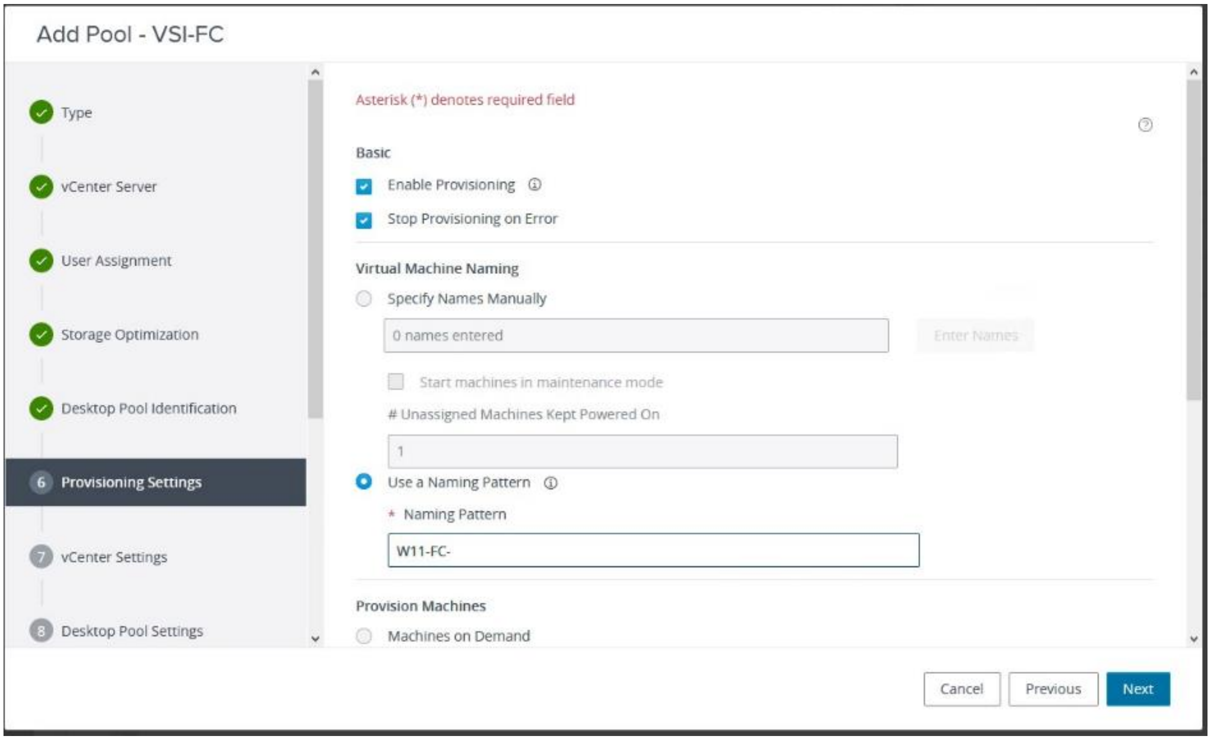

Add Pool - VSI-FC

Asterisk (*) denotes required field

Basic

☑ Enable Provisioning

☑ Stop Provisioning on Error

Virtual Machine Naming

○ Specify Names Manually

0 names entered

□ Start machines in maintenance mode

# Unassigned Machines Kept Powered On: 1

● Use a Naming Pattern

* Naming Pattern: W11-FC-

Provision Machines

○ Machines on Demand

Step 7. Provide the parent VM, snapshot and host/cluster info, data store information for the virtual machines to create.

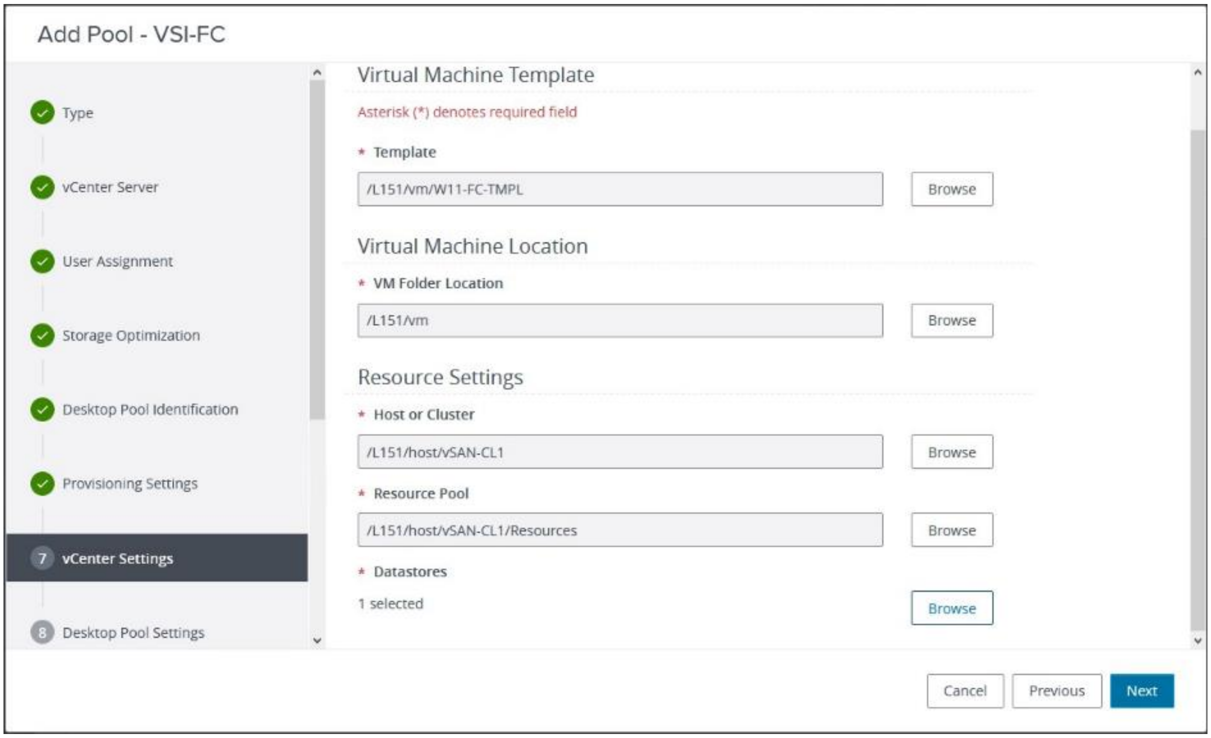

Add Pool - VSI-FC

Virtual Machine Template

Asterisk (*) denotes required field

* Template: /L151/vm/W11-FC-TMPL

Virtual Machine Location

* VM Folder Location: /L151/vm

Resource Settings

* Host or Cluster: /L151/host/vSAN-CL1

* Resource Pool: /L151/host/vSAN-CL1/Resources

* Datastores: 1 selected

Step 8. Configure Desktop Pool settings.

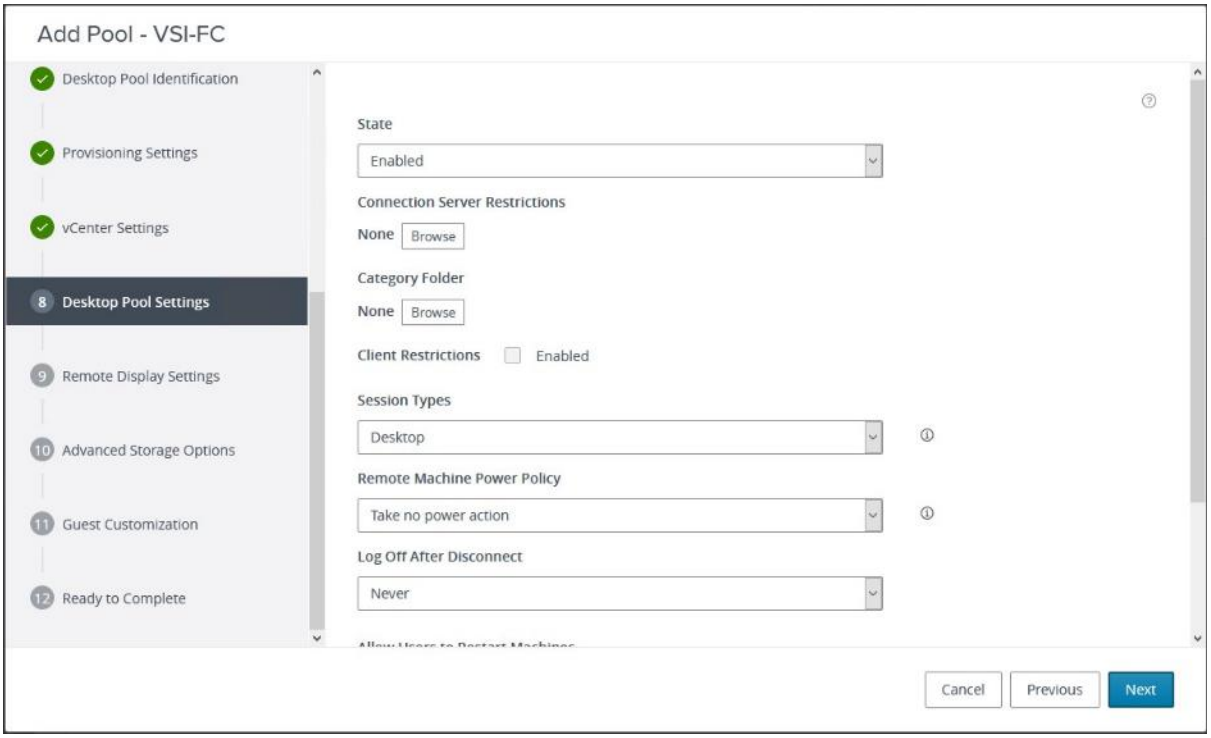

Add Pool - VSI-FC

State: Enabled

Connection Server Restrictions: None

Category Folder: None

Client Restrictions ☐ Enabled

Session Types: Desktop

Remote Machine Power Policy: Take no power action

Log Off After Disconnect: Never

Allow Users to Restart Machines

Step 9. Provide the customizations to remote display protocol to be used by the desktops in the pool.

Note: We used the defaults in this deployment.

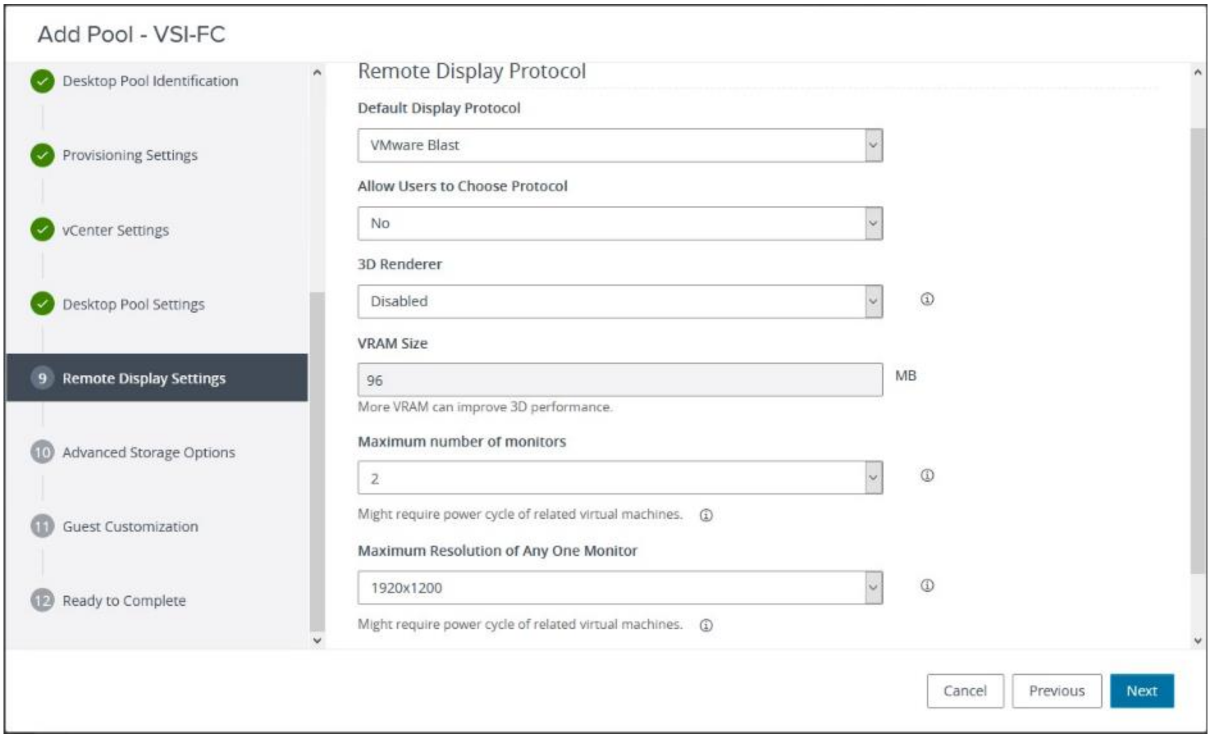

Add Pool - VSI-FC

- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- 9 Remote Display Settings
- 10 Advanced Storage Options
- 11 Guest Customization
- 12 Ready to Complete

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.

Note: For Advanced Storage Options, we used defaults in this deployment.

Step 10. Click Next.

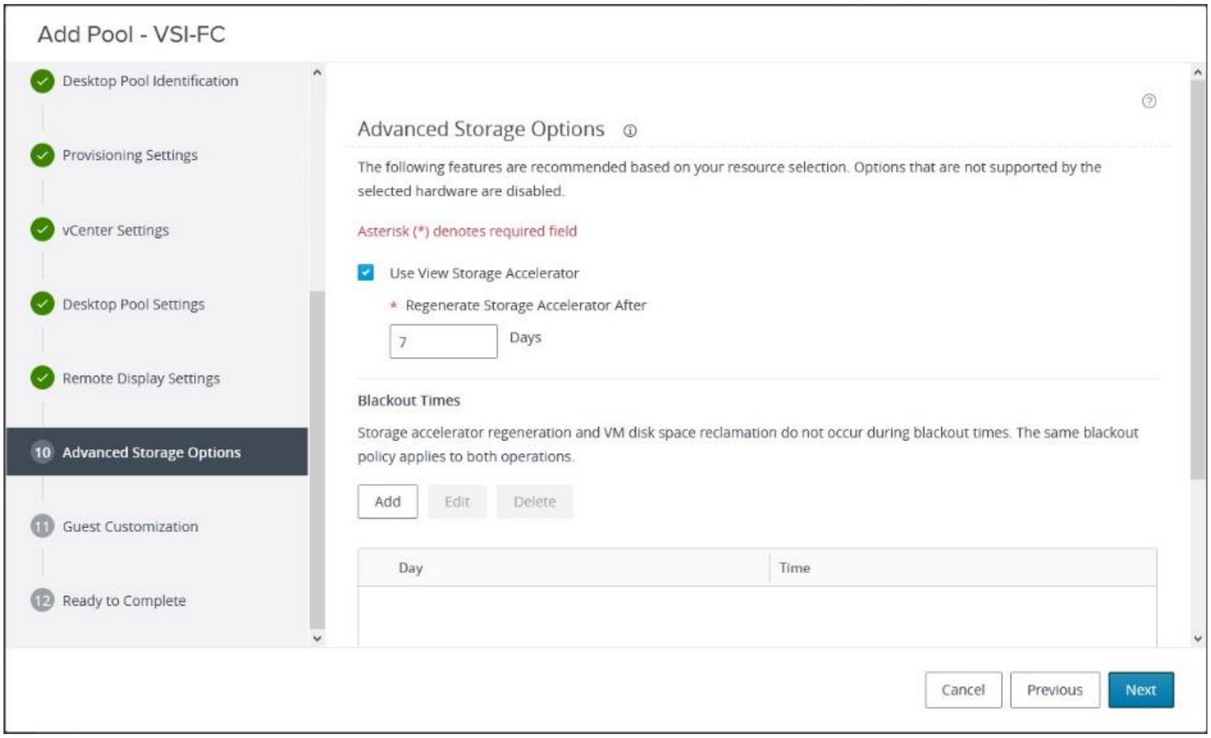

Add Pool - VSI-FC

- ✓ Desktop Pool Identification
- ✓ Provisioning Settings
- ✓ vCenter Settings
- ✓ Desktop Pool Settings
- ✓ Remote Display Settings
- 10 Advanced Storage Options
- 11 Guest Customization
- 12 Ready to Complete

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

7 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

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.

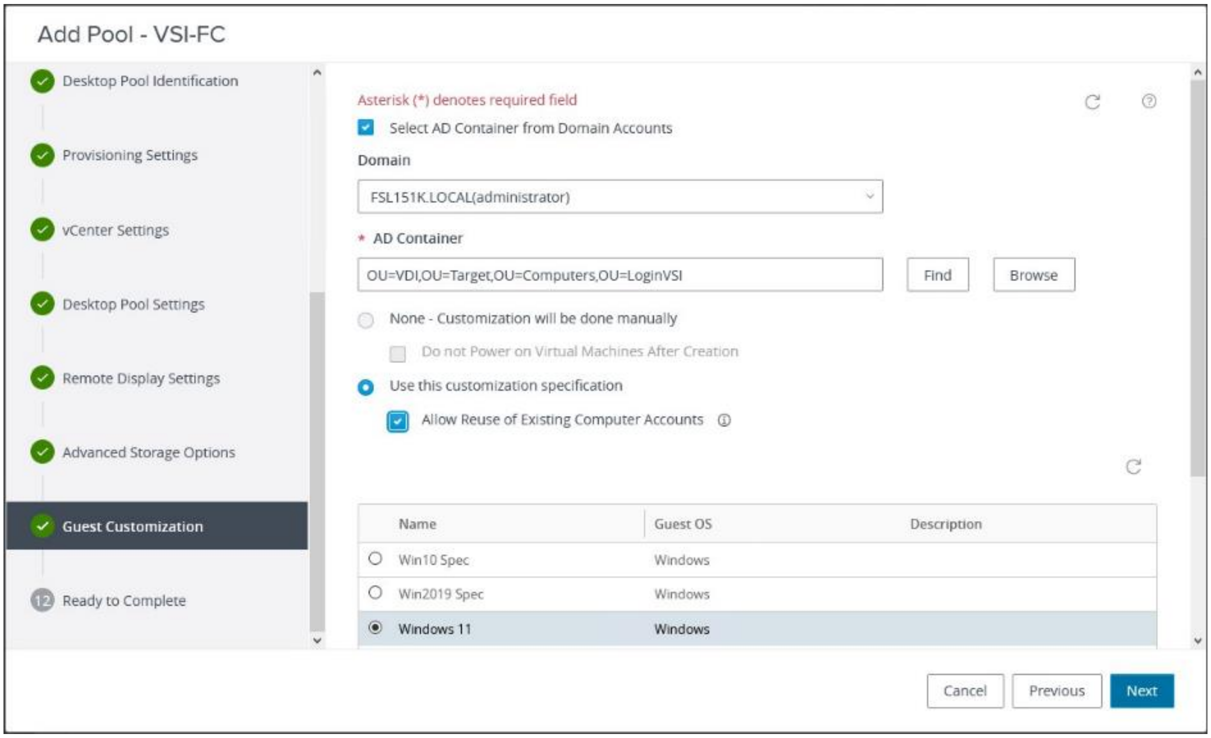

Add Pool - VSI-FC

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
○	Win10 Spec	Windows	
○	Win2019 Spec	Windows	
●	Windows 11	Windows	

Note: The following VM Customization Specifications were used:

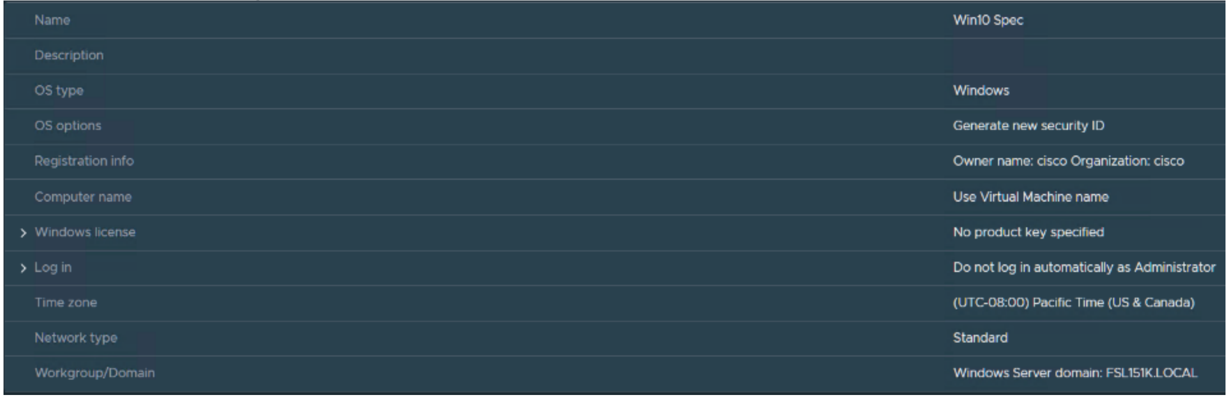

Name	Win10 Spec
Description	
OS type	Windows
OS options	Generate new security ID
Registration info	Owner name: cisco Organization: cisco
Computer name	Use Virtual Machine name
› Windows license	No product key specified
› Log in	Do not log in automatically as Administrator
Time zone	(UTC-08:00) Pacific 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.

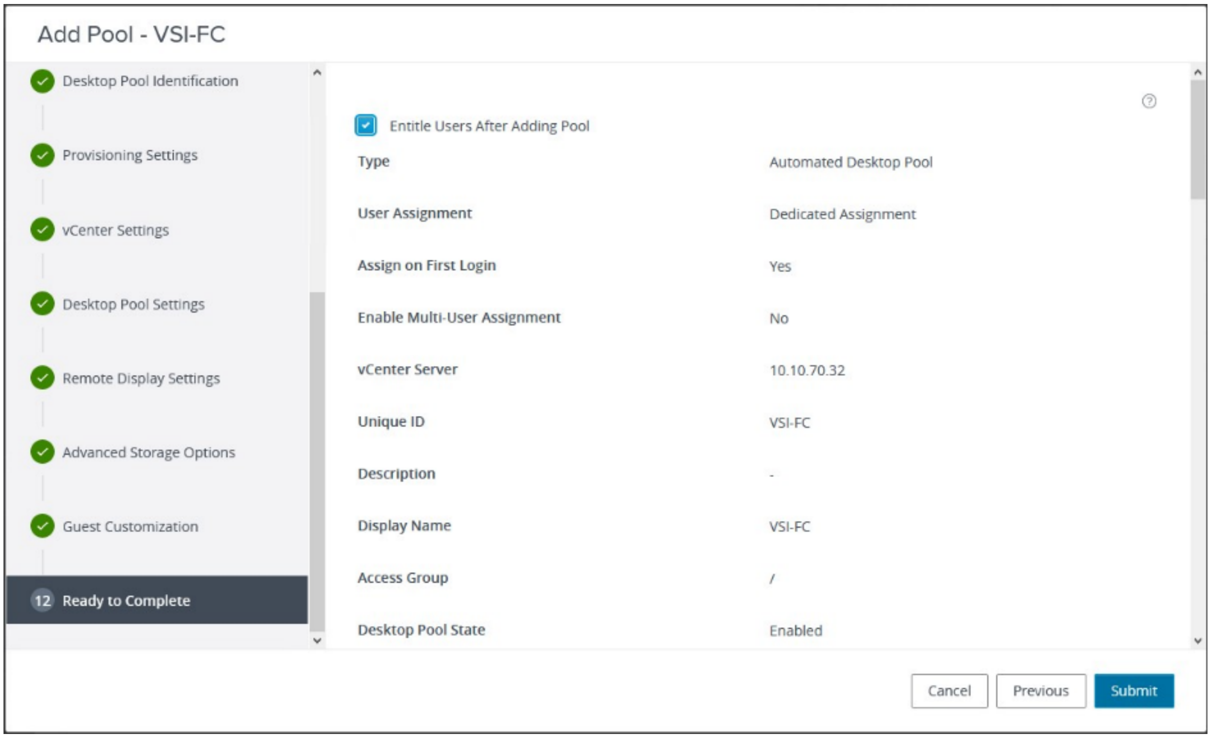

Add Pool - VSI-FC

- Desktop Pool Identification
- Provisioning Settings
- vCenter Settings
- Desktop Pool Settings
- Remote Display Settings
- Advanced Storage Options
- Guest Customization
- 12 Ready to Complete

☑ 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	VSI-FC
Description	-
Display Name	VSI-FC
Access Group	/
Desktop Pool State	Enabled

Step 13. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

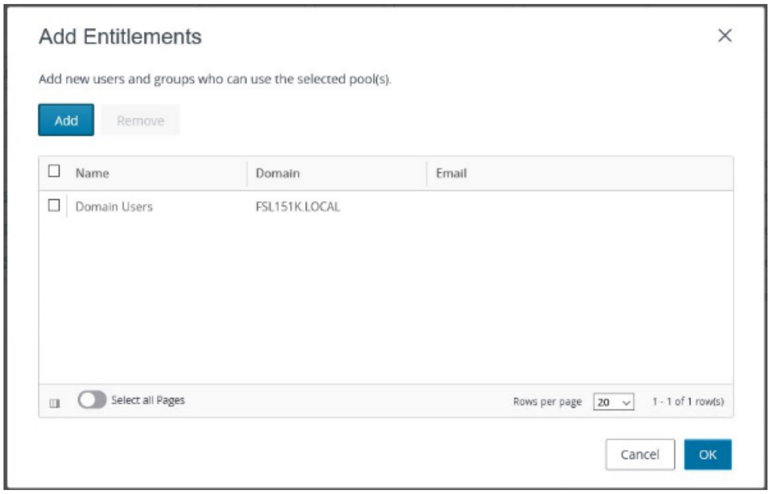

Add Entitlements

Add new users and groups who can use the selected pool(s).

☐	Name	Domain	Email
☐	Domain Users	FSL151K.LOCAL	

Procedure 4. Create RDSH Farm and Pool

Step 1. Select the 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 the Farm. We used Automated Farm the RDS desktops in this design. Click Next.

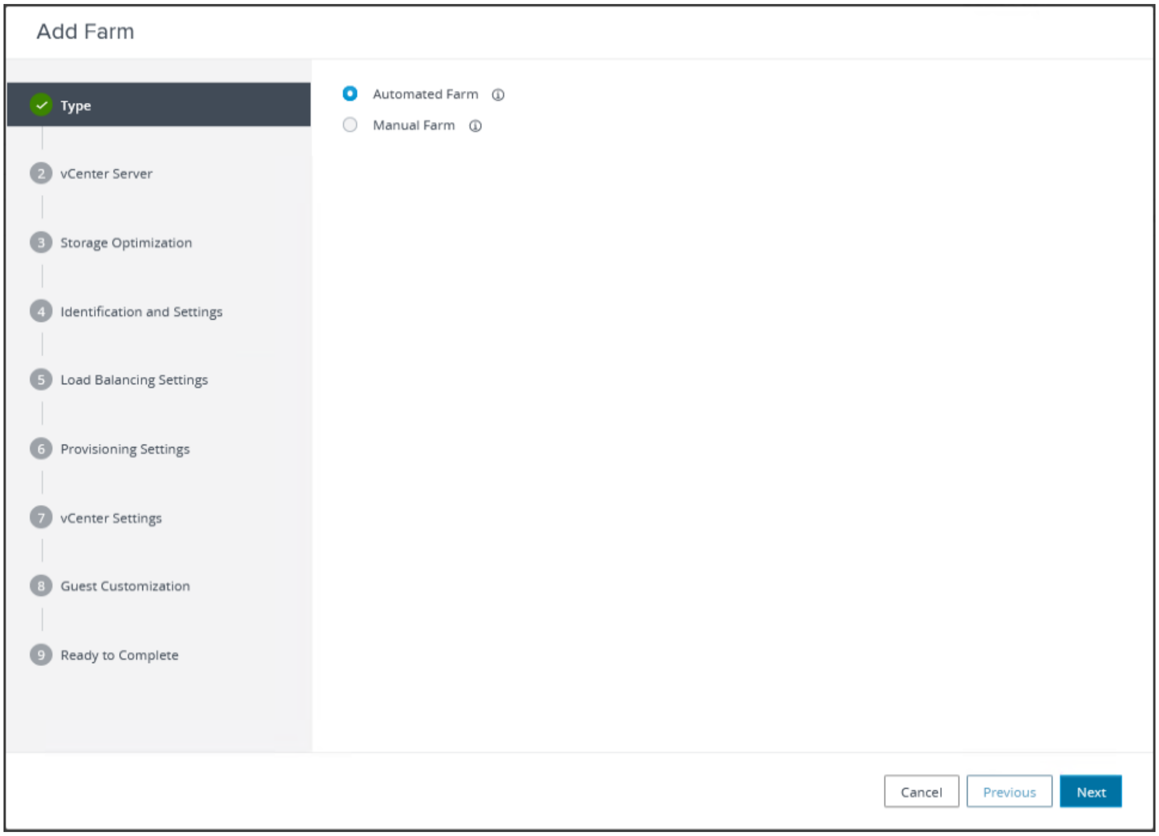
Add Farm

● Automated Farm ⓘ

○ Manual Farm ⓘ

Step 3. Select the provisioning type and vCenter Server for the desktops in the pool. Click Next.

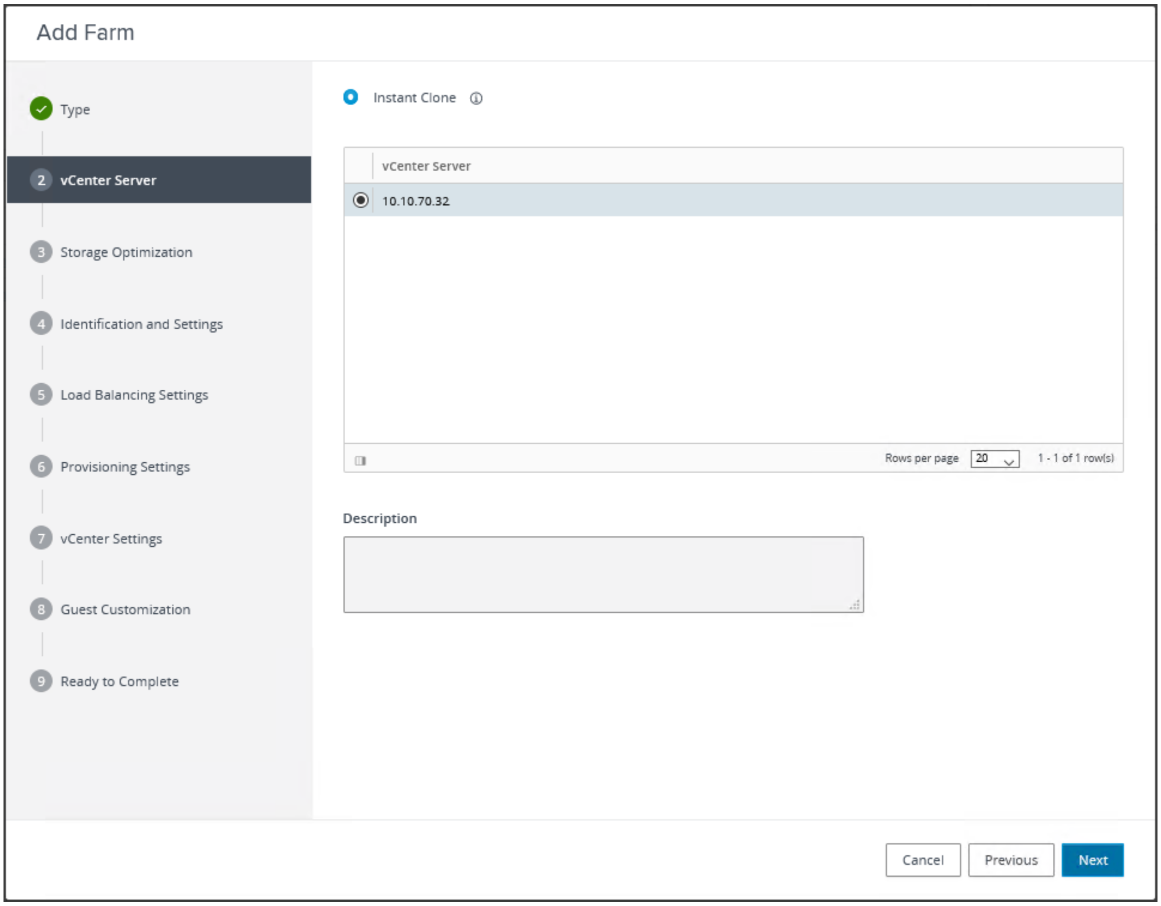
Add Farm

● Instant Clone

	vCenter Server
●	10.10.70.32

Rows per page 20 1 - 1 of 1 row(s)

Description

Step 4. Select Use VMware Virtual SAN on the Storage Optimization screen click Next.

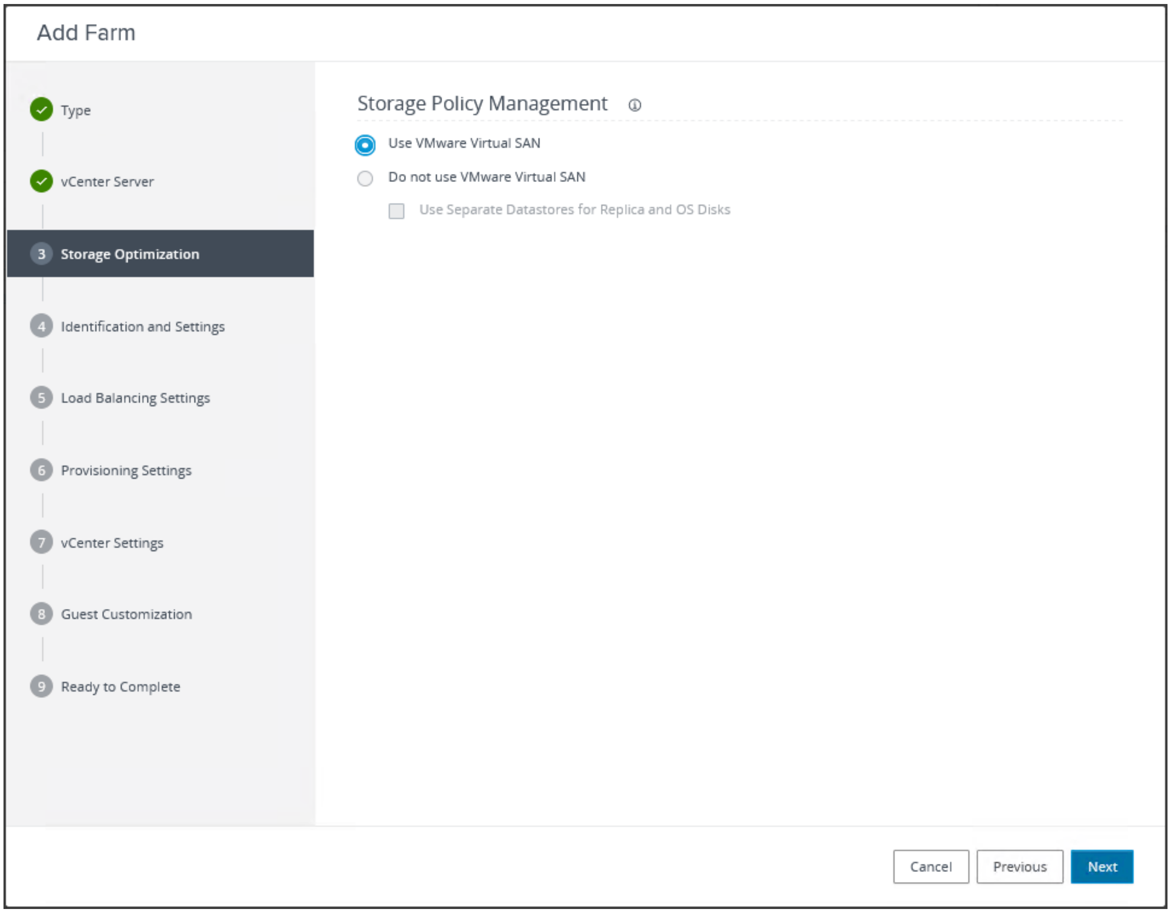

Add Farm

Storage Policy Management

● Use VMware Virtual SAN

○ Do not use VMware Virtual SAN

□ Use Separate Datastores for Replica and OS Disks

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.

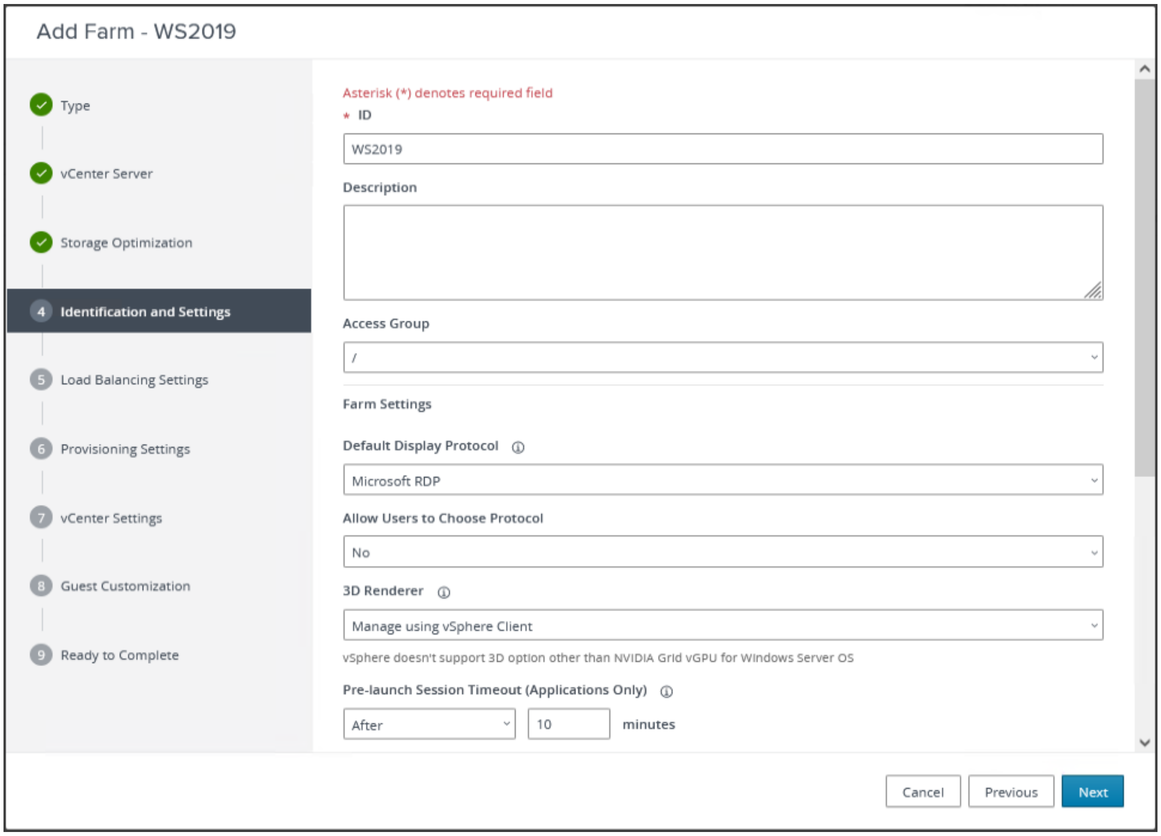

Add Farm - WS2019

Asterisk (*) denotes required field

* ID: WS2019

Description:

Access Group: /

Farm Settings

Default Display Protocol: Microsoft RDP

Allow Users to Choose Protocol: No

3D Renderer: Manage using vSphere Client

vSphere doesn't support 3D option other than NVIDIA Grid vGPU for Windows Server OS

Pre-launch Session Timeout (Applications Only): After 10 minutes

Step 6. Select Load Balancing Settings. Click Next.

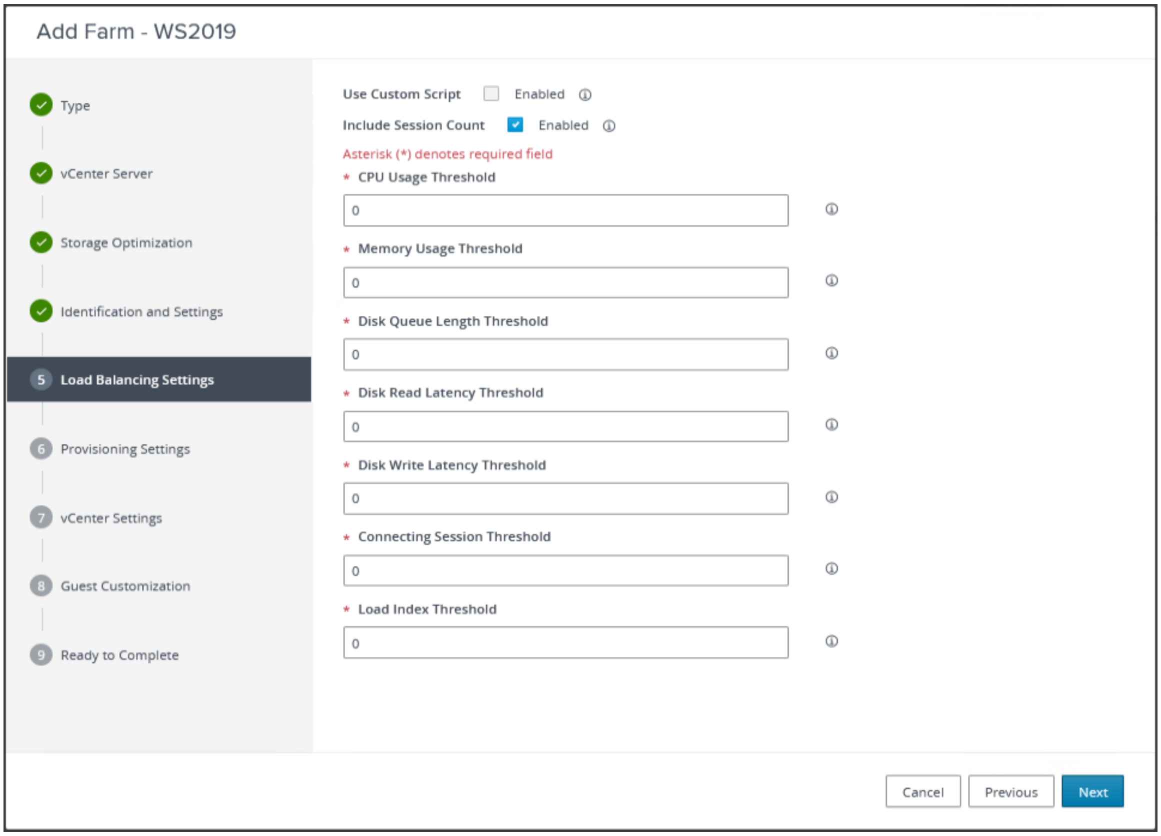

Add Farm - WS2019

Use Custom Script: □ Enabled

Include Session Count: ☑ Enabled

Asterisk (*) denotes required field

* CPU Usage Threshold: 0

* Memory Usage Threshold: 0

* Disk Queue Length Threshold: 0

* Disk Read Latency Threshold: 0

* Disk Write Latency Threshold: 0

* Connecting Session Threshold: 0

* Load Index Threshold: 0

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 – 96

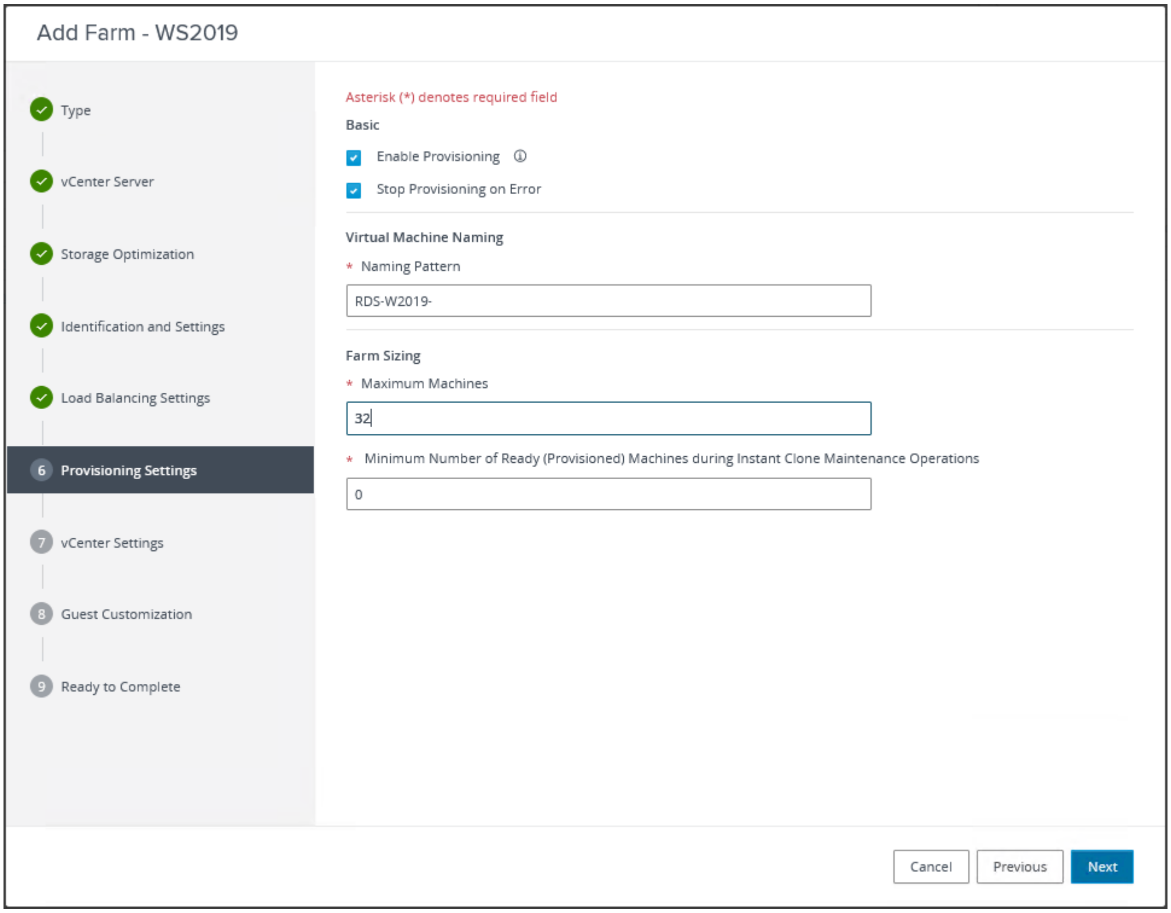

Add Farm - WS2019

Asterisk (*) denotes required field

Basic

☑ Enable Provisioning

☑ Stop Provisioning on Error

Virtual Machine Naming

* Naming Pattern: RDS-W2019-

Farm Sizing

* Maximum Machines: 32

* Minimum Number of Ready (Provisioned) Machines during Instant Clone Maintenance Operations: 0

Step 8. Select the previously created golden image to be used as RDS host. Select datastore where RDS hosts will be deployed. Click Next.

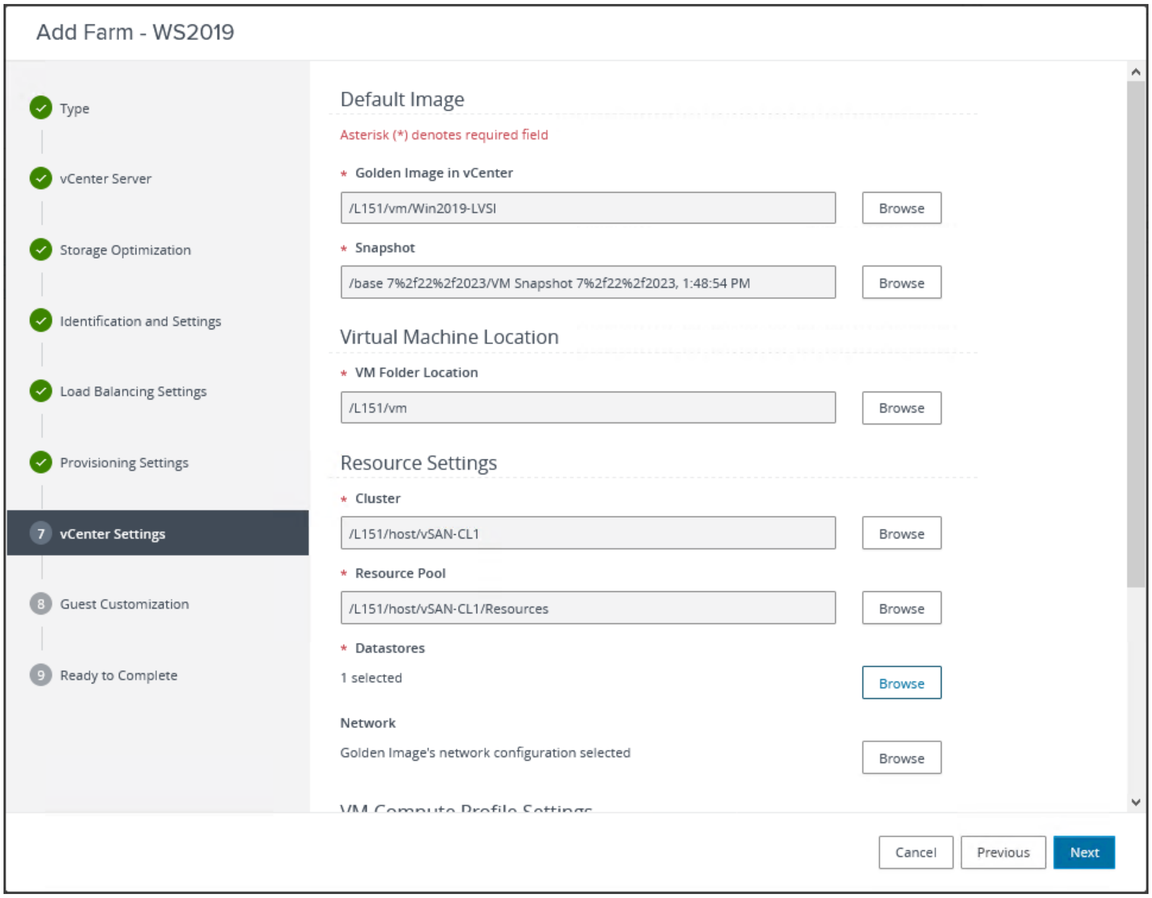

Add Farm - WS2019

Default Image

Asterisk (*) denotes required field

* Golden Image in vCenter: /L151/vm/Win2019-LVSI

* Snapshot: /base 7%2f22%2f2023/VM Snapshot 7%2f22%2f2023, 1:48:54 PM

Virtual Machine Location

* VM Folder Location: /L151/vm

Resource Settings

* Cluster: /L151/host/vSAN-CL1

* Resource Pool: /L151/host/vSAN-CL1/Resources

* Datastores: 1 selected

Network: Golden Image's network configuration selected

VM Compute Profile Settings

Step 9. Select the AD Container for desktops to place in a Domain Controller computer location.

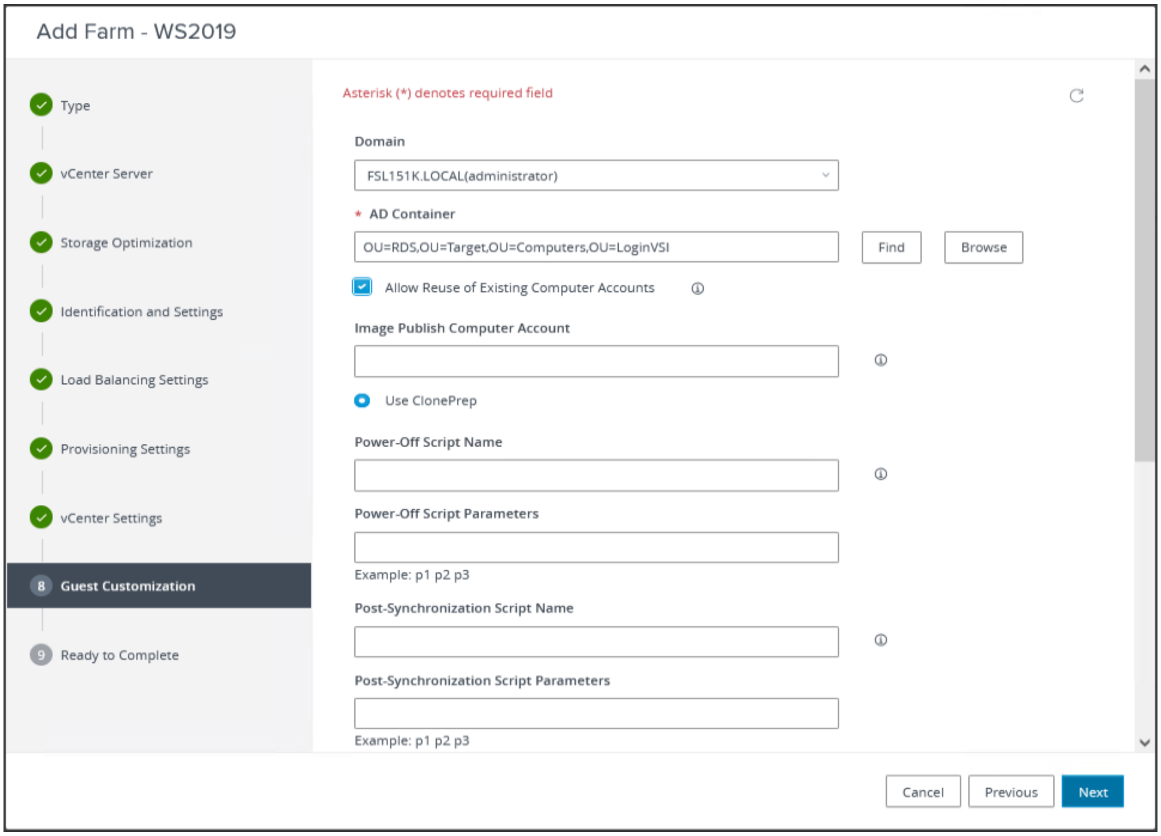

Add Farm - WS2019

Asterisk (*) denotes required field

Domain

FSL151K.LOCAL(administrator)

* AD Container

OU=RDS,OU=Target,OU=Computers,OU=LoginVSI

☑ 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

Step 10. Review the Farm information and click Submit to complete the RDS Farm creation.

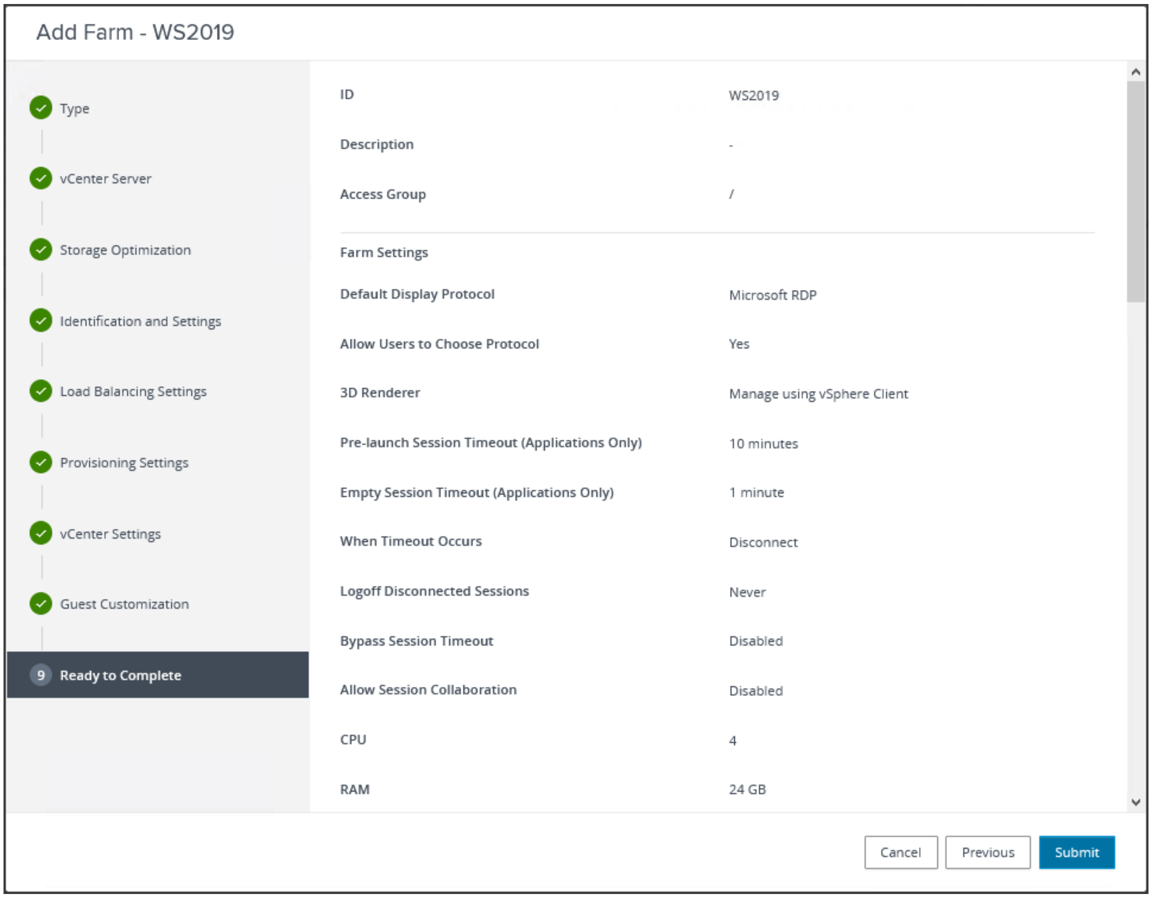

Add Farm - WS2019

ID	WS2019
Description	-
Access Group	/
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

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.

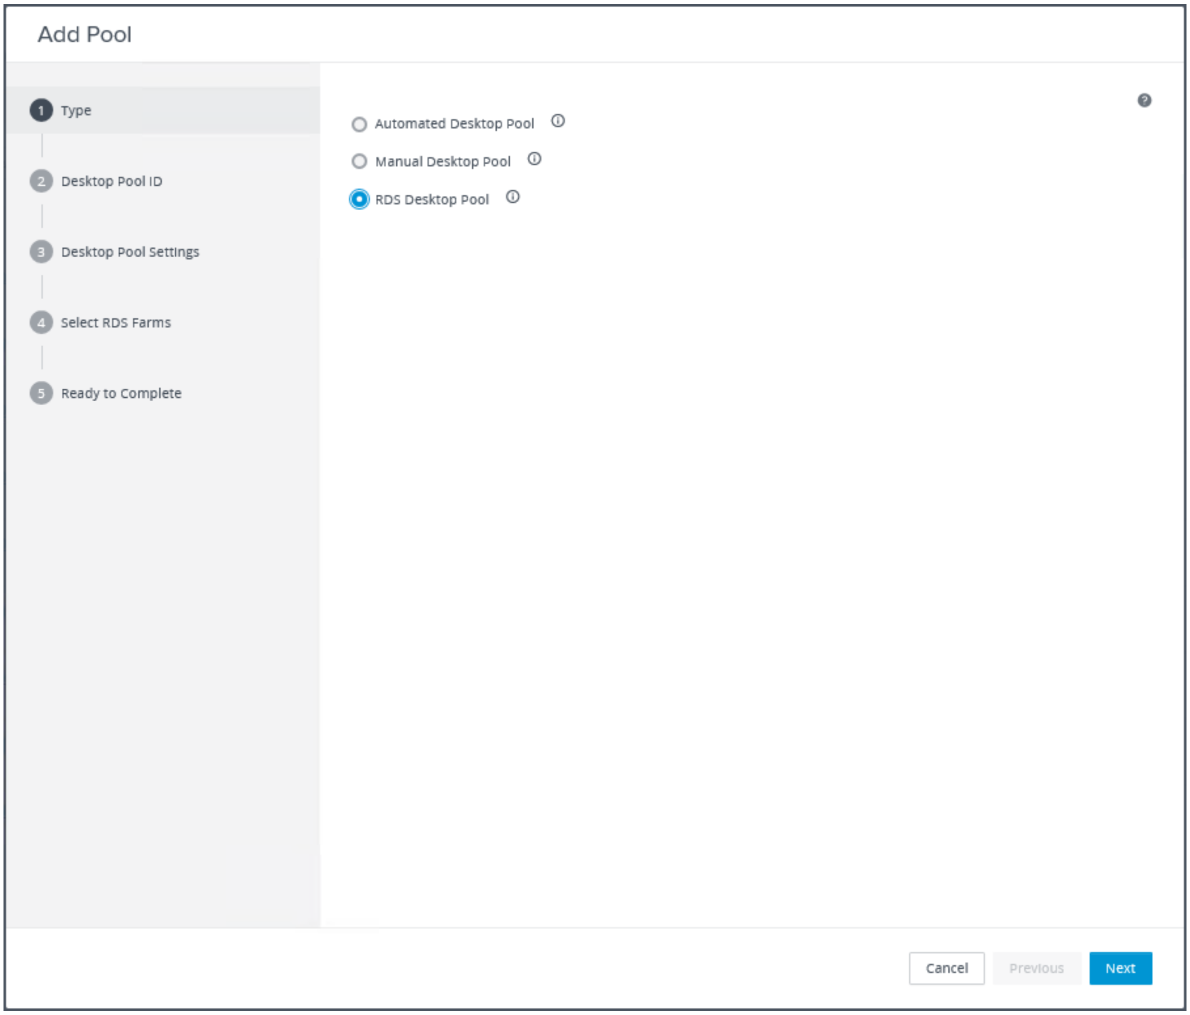

Add Pool

○ Automated Desktop Pool

○ Manual Desktop Pool

● RDS Desktop Pool

Step 2. Provide an ID and Display Name for the Pool. Click Next.

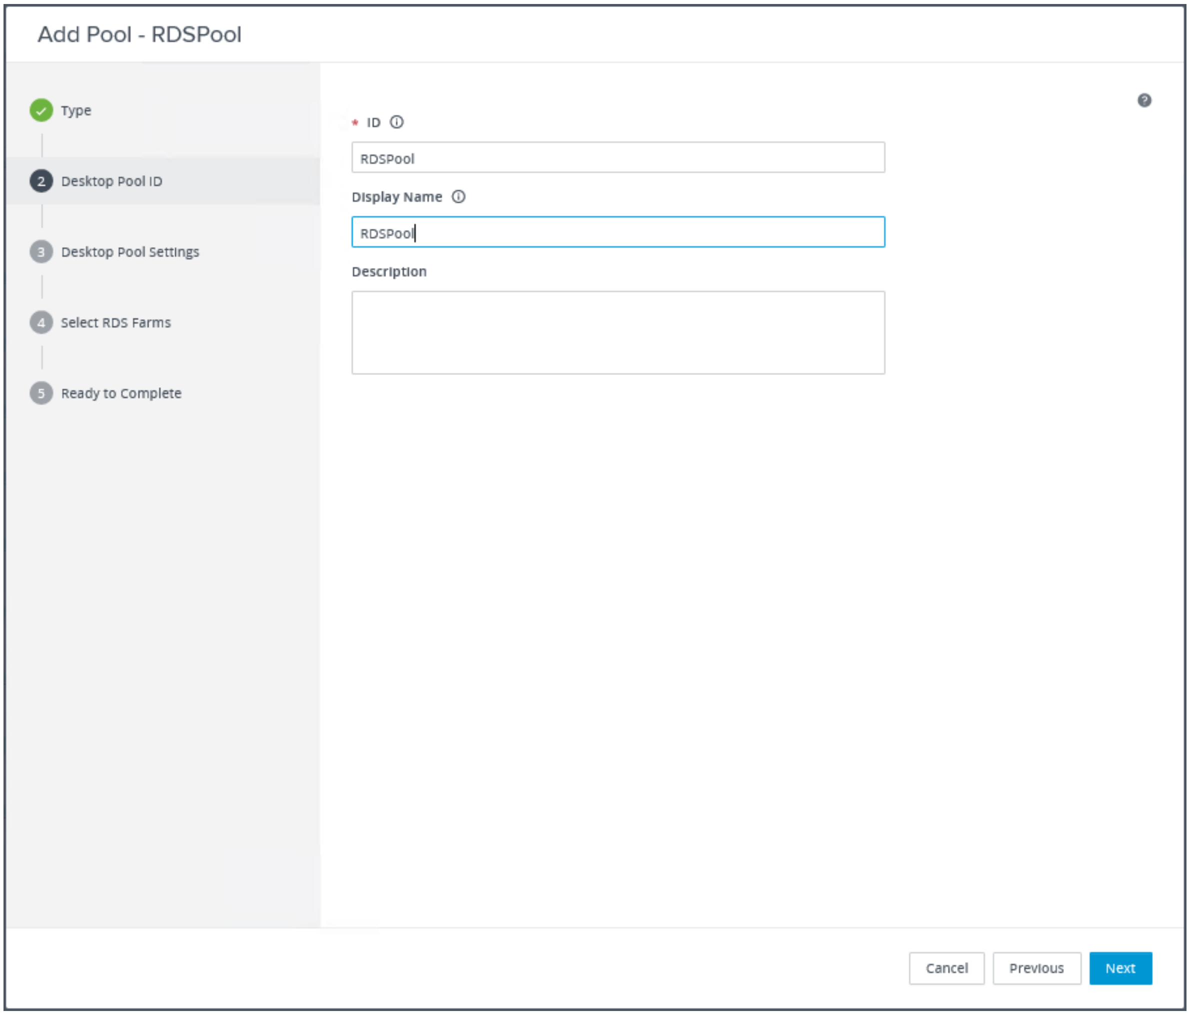
Add Pool - RDSPool

* ID: RDSPool

Display Name: RDSPool

Description:

Step 3. Leave the default settings for the Desktop Pool Settings. Click Next.

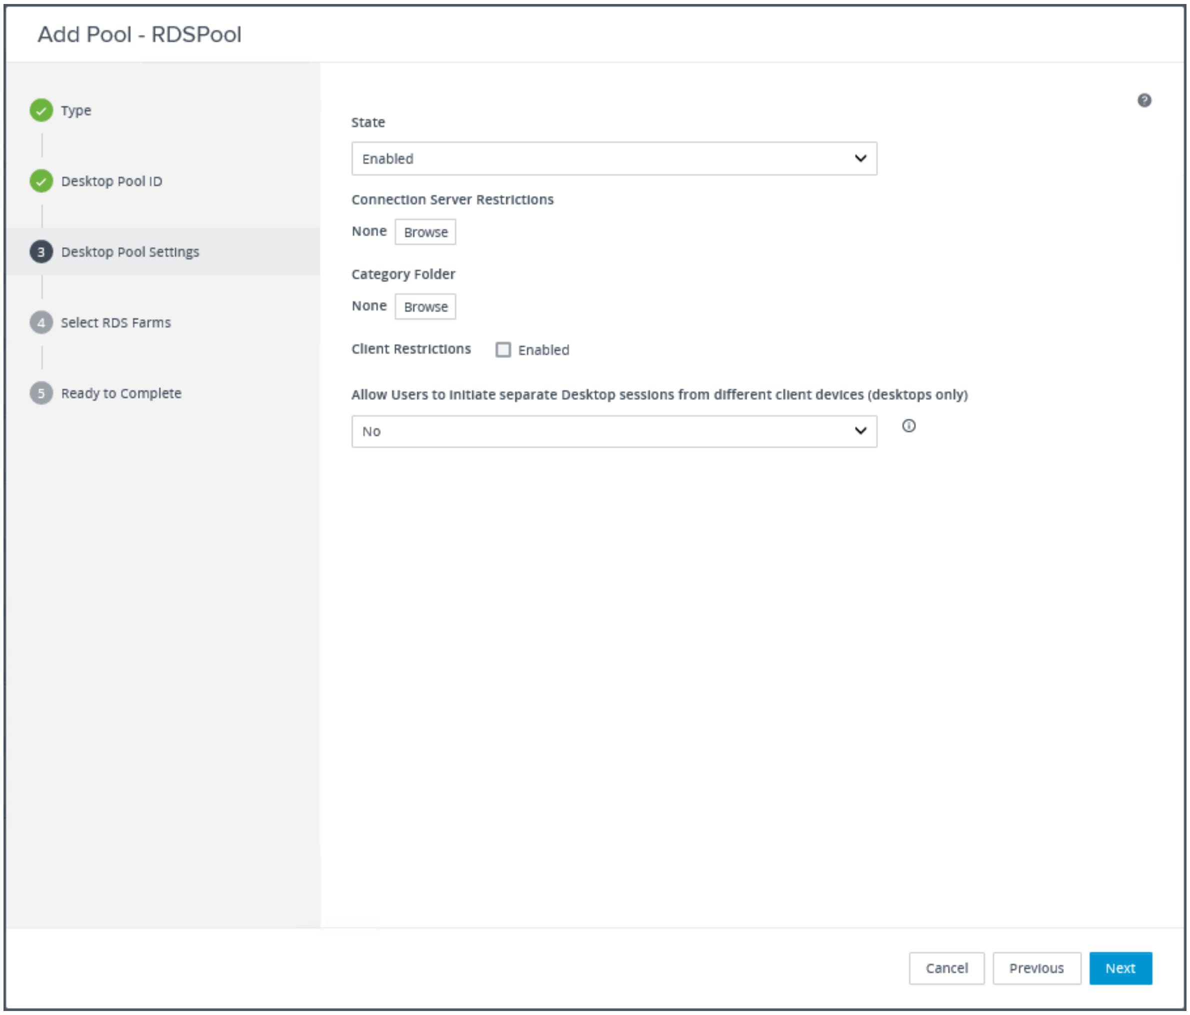

Add Pool - RDSPool

State

Enabled

Connection Server Restrictions

None Browse

Category Folder

None Browse

Client Restrictions □ Enabled

Allow Users to initiate separate Desktop sessions from different client devices (desktops only)

No

Step 4. Select the RDS Farm. Select the farm which was already created for this desktop pool. Click Next.

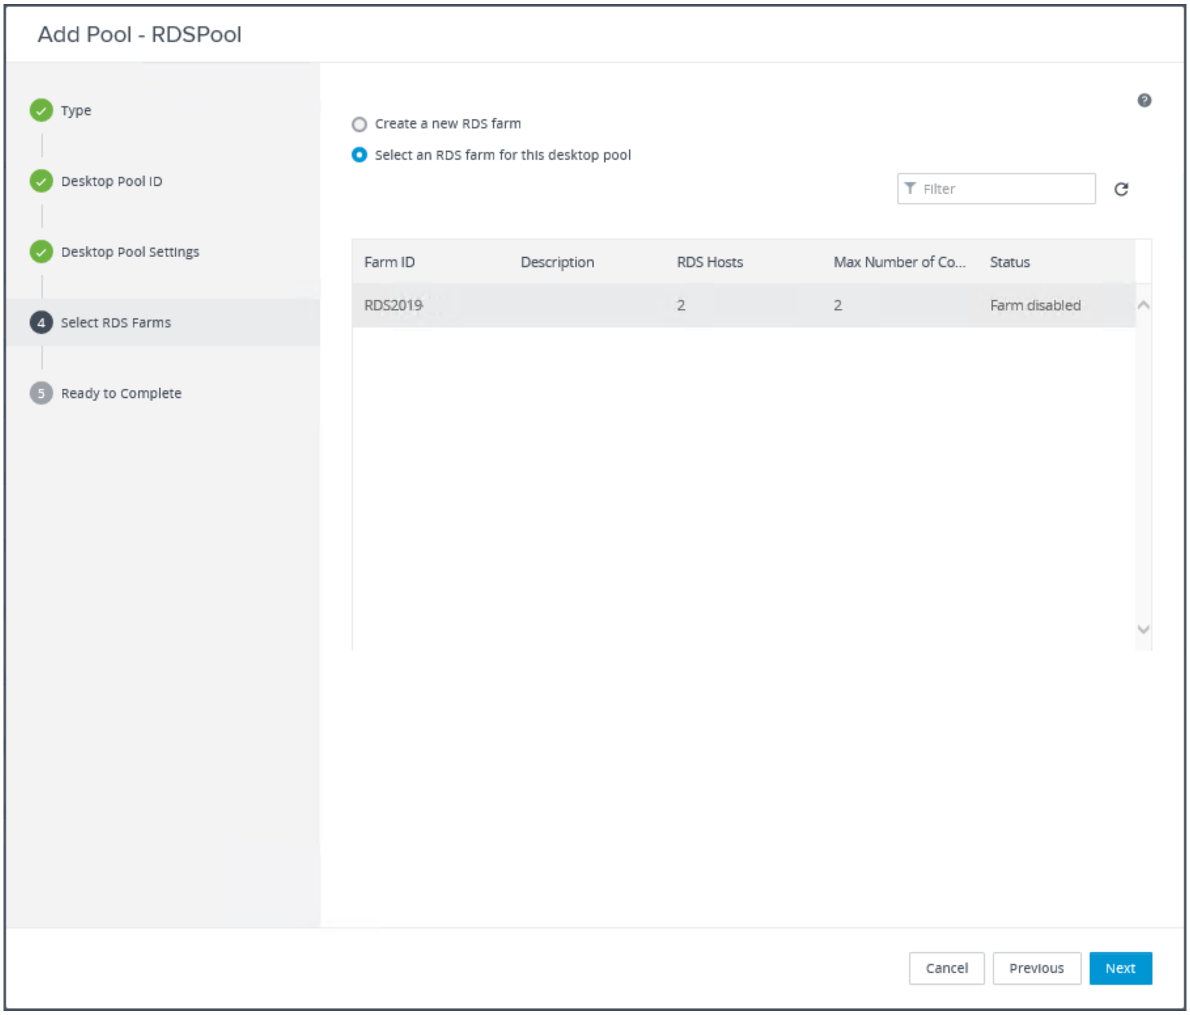
Add Pool - RDSPool

○ Create a new RDS farm

● Select an RDS farm for this desktop pool

Farm ID	Description	RDS Hosts	Max Number of Co...	Status
RDS2019		2	2	Farm disabled

Step 5. Review the RDS Pool deployment specifications and click Next to complete the RDS pool deployment.

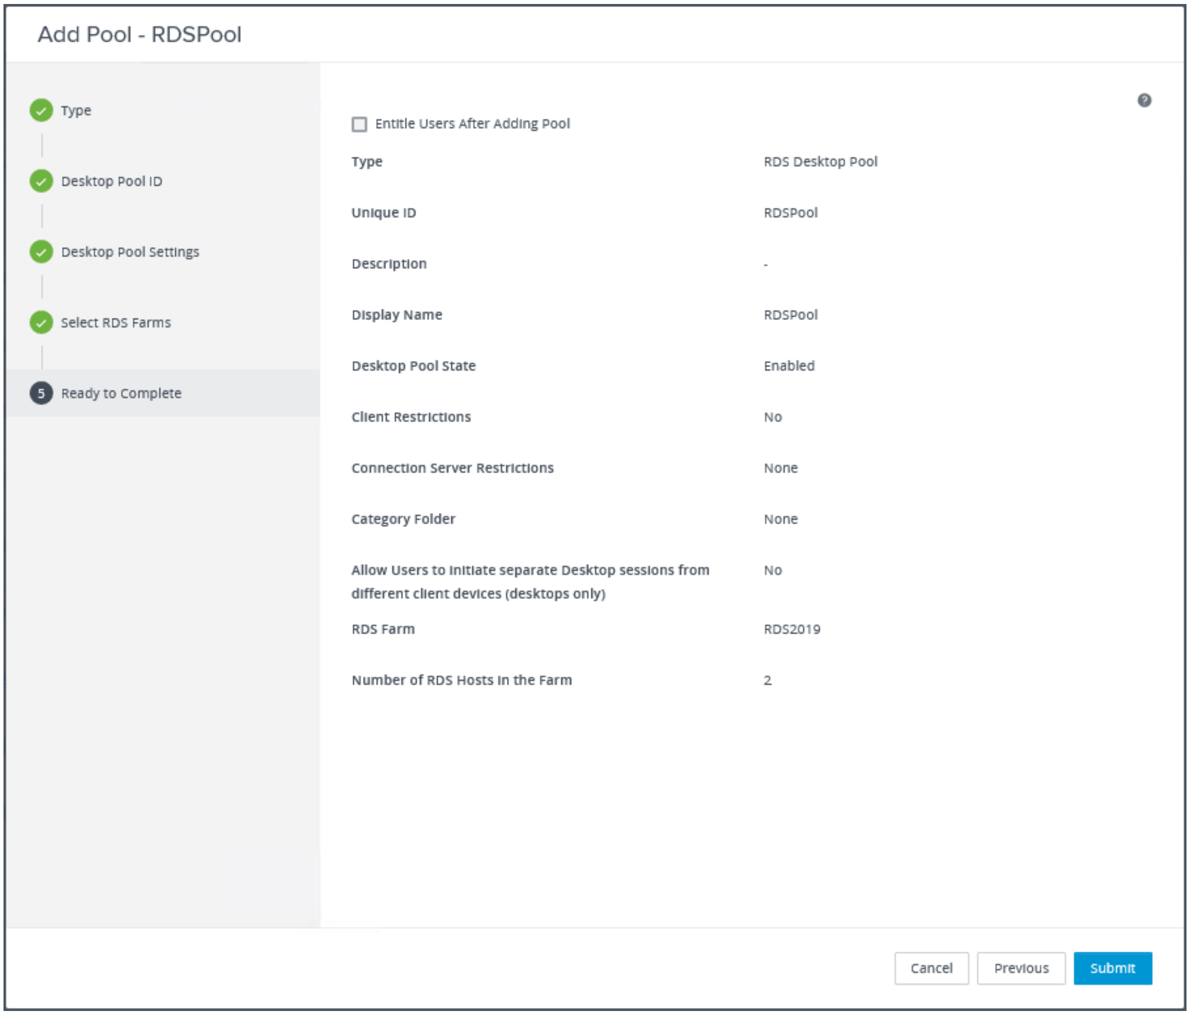

Add Pool - RDSPool

□ Entitle Users After Adding Pool

Type	RDS Desktop Pool
Unique ID	RDSPool
Description	-
Display Name	RDSPool
Desktop Pool State	Enabled
Client Restrictions	No
Connection Server Restrictions	None
Category Folder	None
Allow Users to initiate separate Desktop sessions from different client devices (desktops only)	No
RDS Farm	RDS2019
Number of RDS Hosts in the Farm	2

Step 6. Select Entitle users after this wizard finishes, to enable desktop user group/users to access this pool.

Test Setup, Configuration, and Load Recommendation

This chapter contains the following:

- Cisco UCS Test Configuration for Single Blade Scalability
- Cisco UCS Test Configuration for Full Scale Testing
- Test Methodology and Success Criteria

We tested a single Cisco UCS X210c M7 Compute Node to validate against the performance of one, three (ESA vSAN cluster) and four (OSA vSAN cluster) Cisco UCS X210c M7 Compute Nodes on a single chassis to illustrate linear scalability for each workload use case studied.

Cisco UCS Test Configuration for Single Blade Scalability

This test case validates the Recommended Maximum Workload per host server using VMware Horizon 8 2212 with 480 Multi-session OS sessions and 195 Single-session OS sessions.

Figure 24. Test Configuration for Single Server Scalability VMware Horizon 8 2212 Non-persistent (NP) Single-session OS machine VDAs

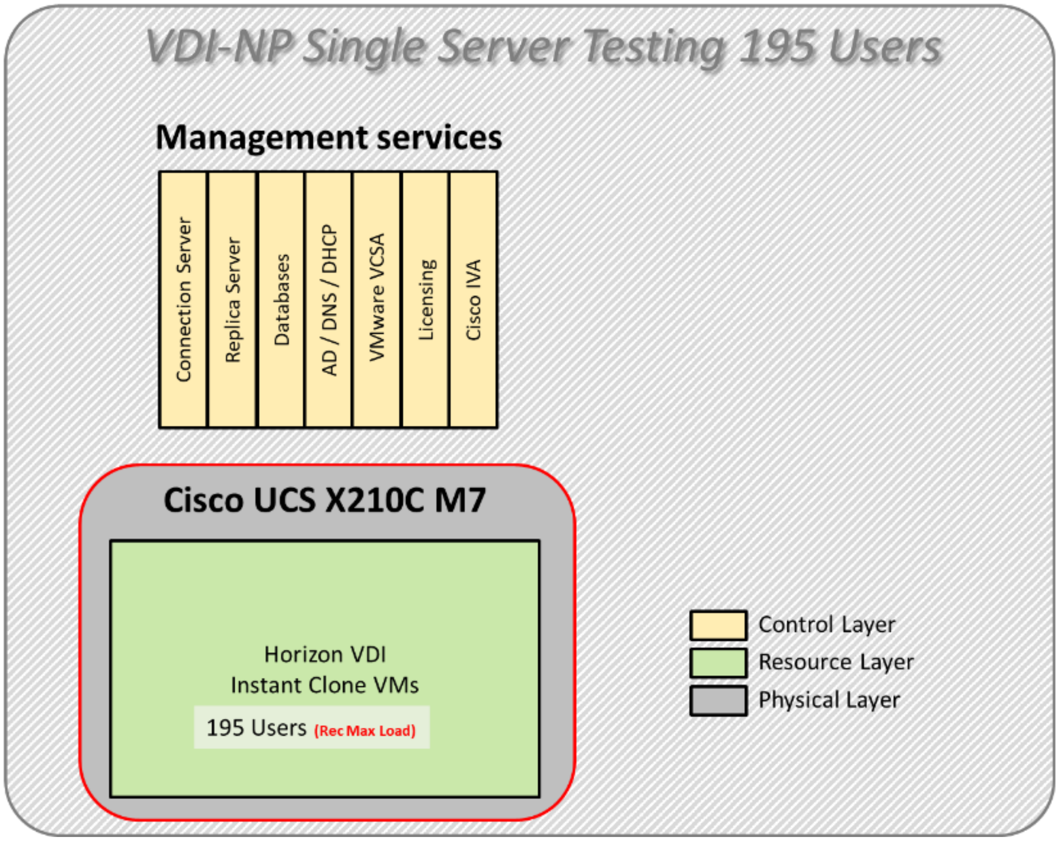

Figure 25. Test configuration for Single Server Scalability VMware Horizon 8 2212 Persistent (P) Single-session OS machine VDAs

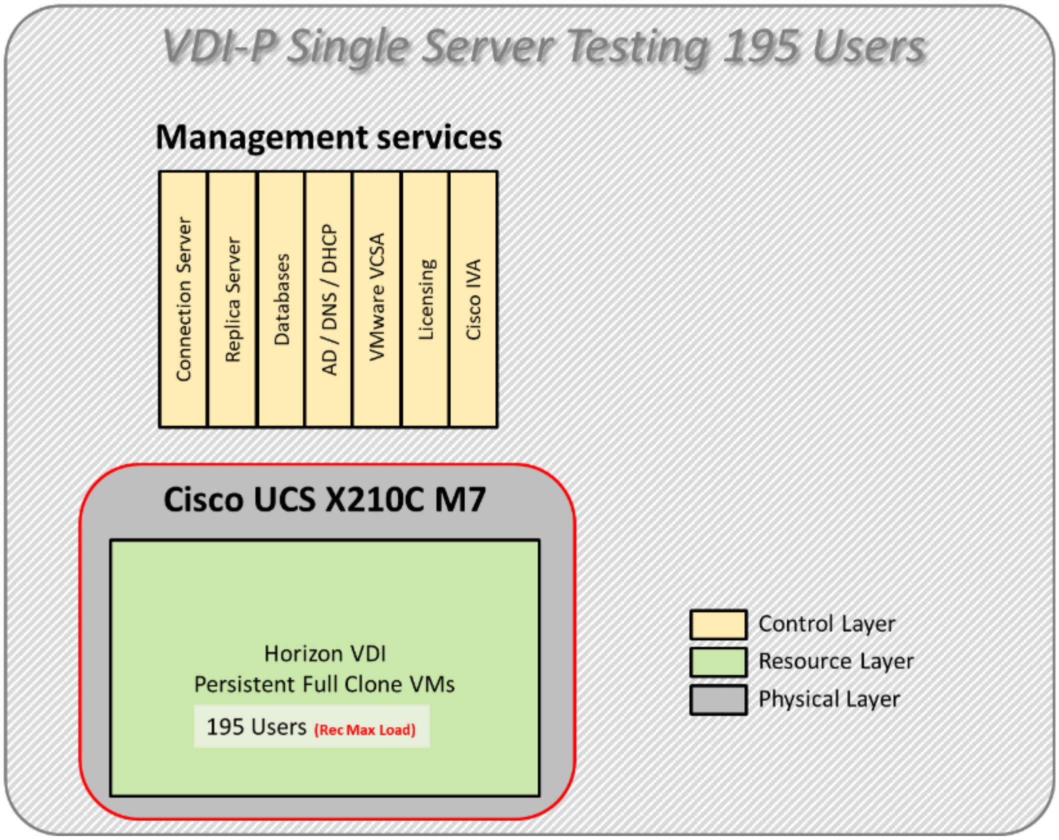

Figure 26. **Test configuration for Single Server Scalability VMware Horizon 8 2212 Instant-clone Multi-session OS machine VDAs**

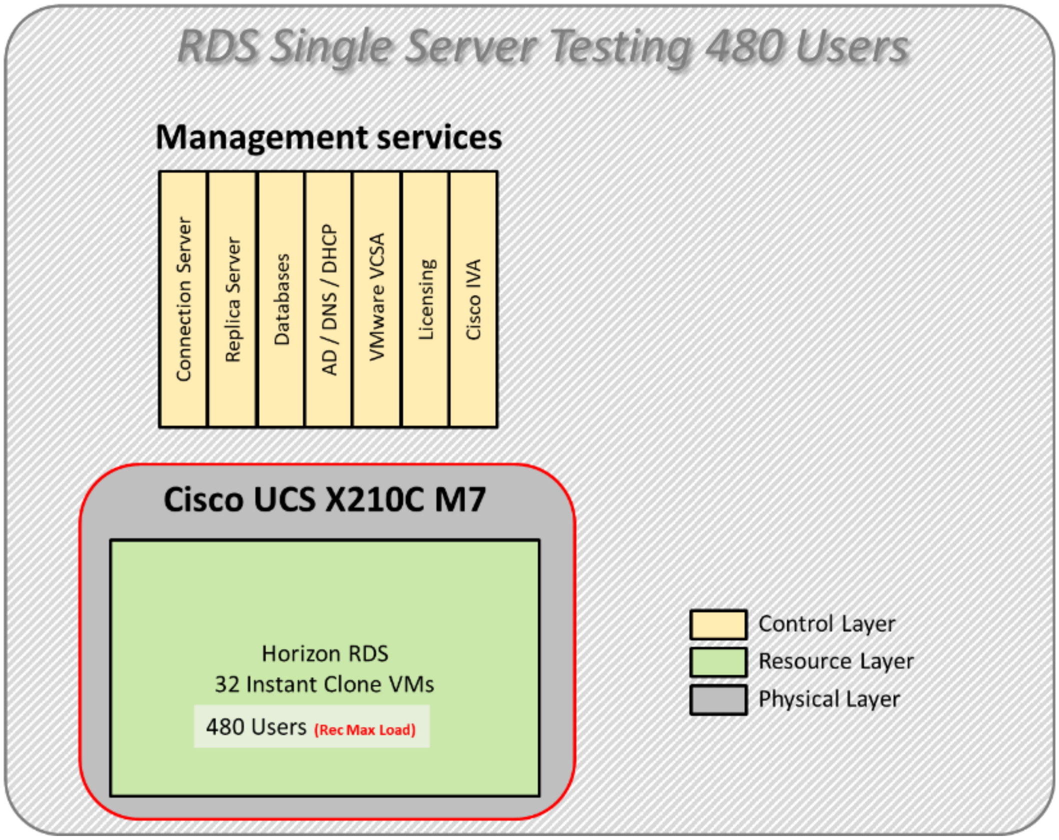

Hardware components:

- 1 Cisco UCS X9508 Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 1 Cisco UCS X210c M7 Compute Node Servers with Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 1TB 4800 MHz RAM for all host blades
- Cisco UCS VIC 15231 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 1 1.6TB High Performance High Endurance and 5 1.92TB High Performance Medium Endurance NVMe drives (per blade, OSA vSAN cluster)
- 5 3.2TB High Perf High Endurance NVMe drives (per blade, ESA vSAN cluster)

Software components:

- Cisco UCS firmware 5.1(1.230052)
- VMware vSAN 8 (OSA/ESA)
- VMware ESXi 8.0 Update 1a 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

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:

- 585 VDI-NP (instant clones) Single-session OS sessions (OSA vSAN cluster)
- 585 VDI-P (full clones) Single-session OS sessions (OSA vSAN cluster)
- 1440 Instant-clone RDS sessions (OSA vSAN cluster)
- 390 VDI-NP (instant clones) Single-session OS sessions (ESA vSAN cluster)
- 390 VDI-P (full clones) Single-session OS sessions (ESA vSAN cluster)
- 960 Instant-clone RDS sessions (ESA vSAN cluster)

Note: Server N+1 fault tolerance is factored into this solution for each cluster/workload.

Figure 27. Test Configuration for Full Scale VMware Horizon 8 2212 non-persistent Single-session OS machine VDAs (OSA vSAN cluster)

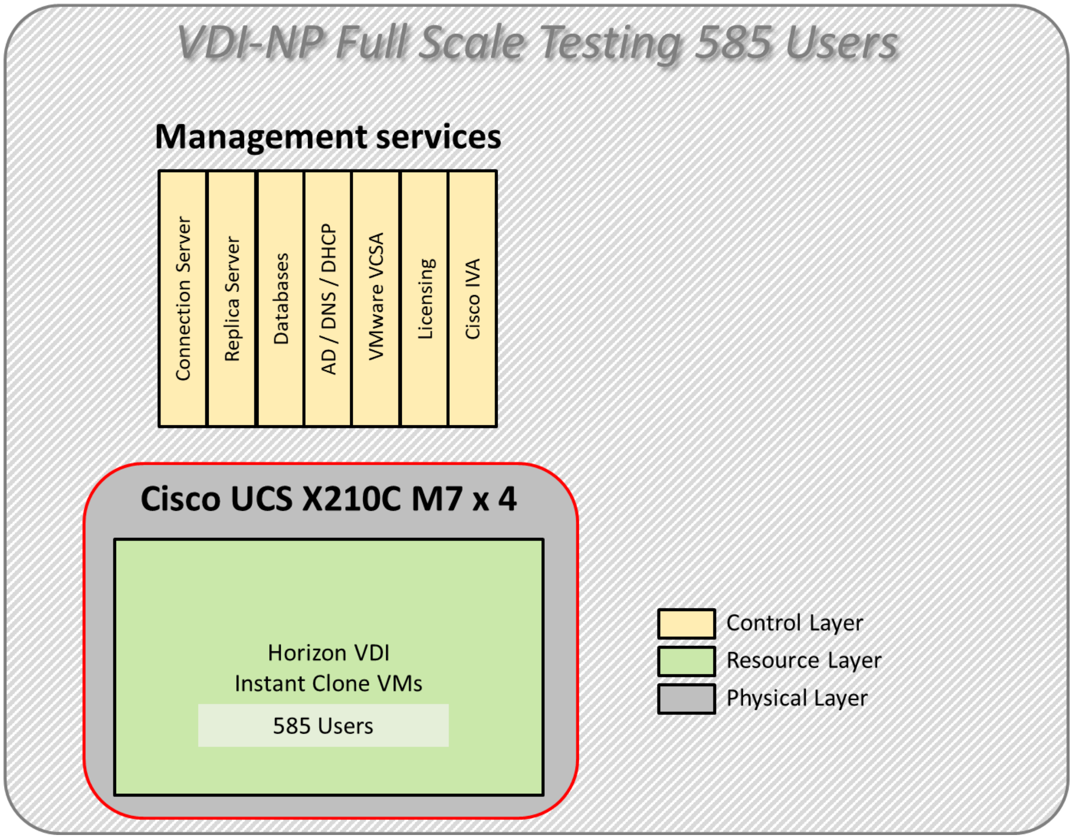

Figure 28. Test Configuration for Full Scale VMware Horizon 8 2212 persistent Single-session OS machine VDAs (OSA vSAN cluster)

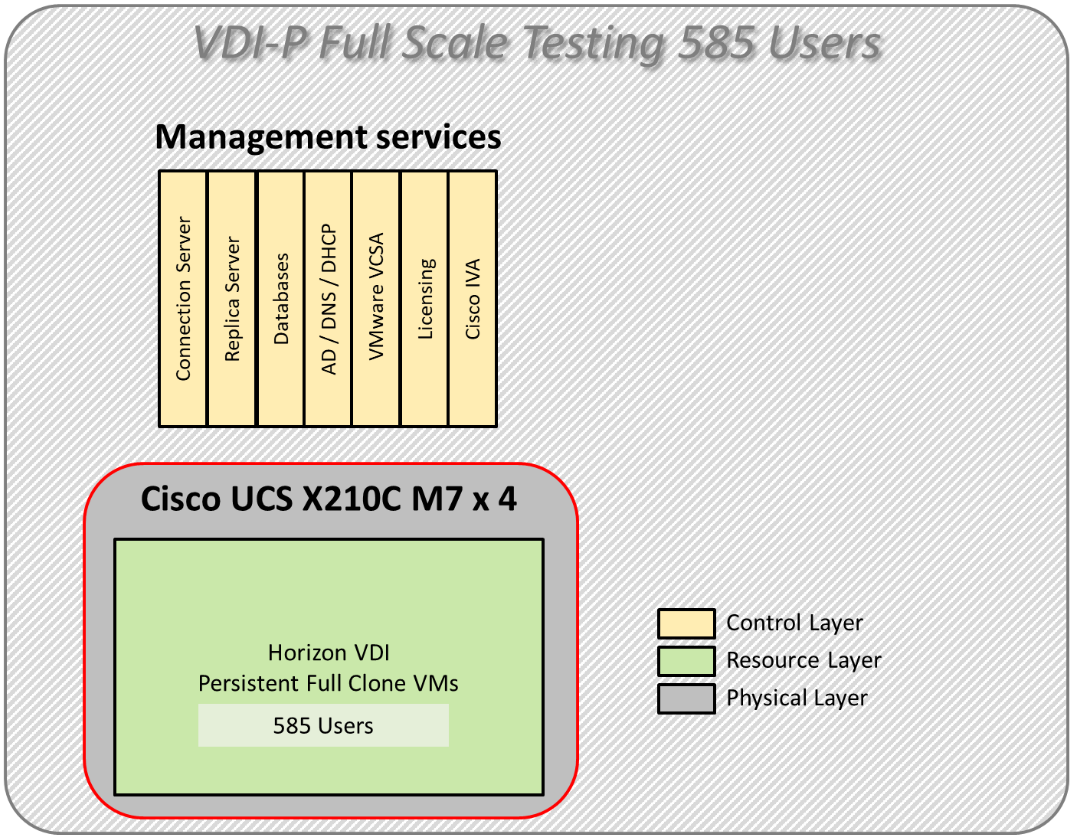

Figure 29. **Test Configuration for Full Scale VMware Horizon 8 2212 instant-clones Multi-session OS machine VDAs (OSA vSAN cluster)**

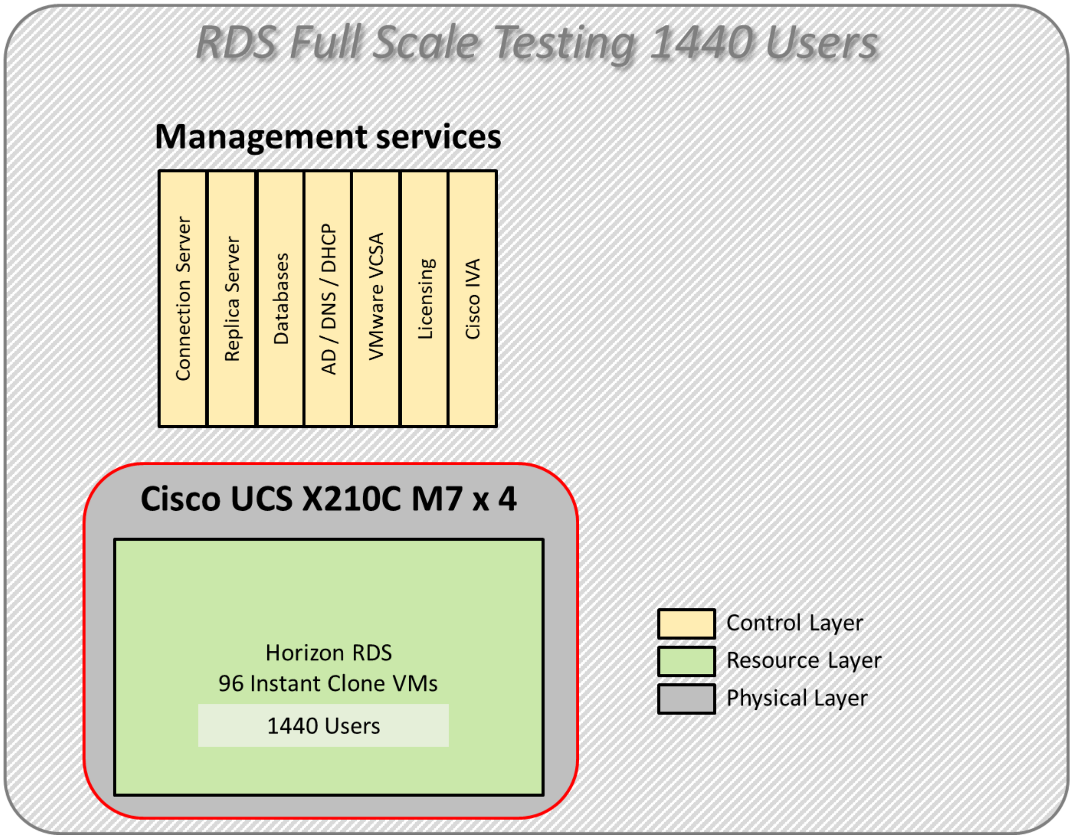

Figure 30. Test Configuration for Full Scale VMware Horizon 8 2212 non-persistent Single-session OS machine VDAs (ESA vSAN cluster)

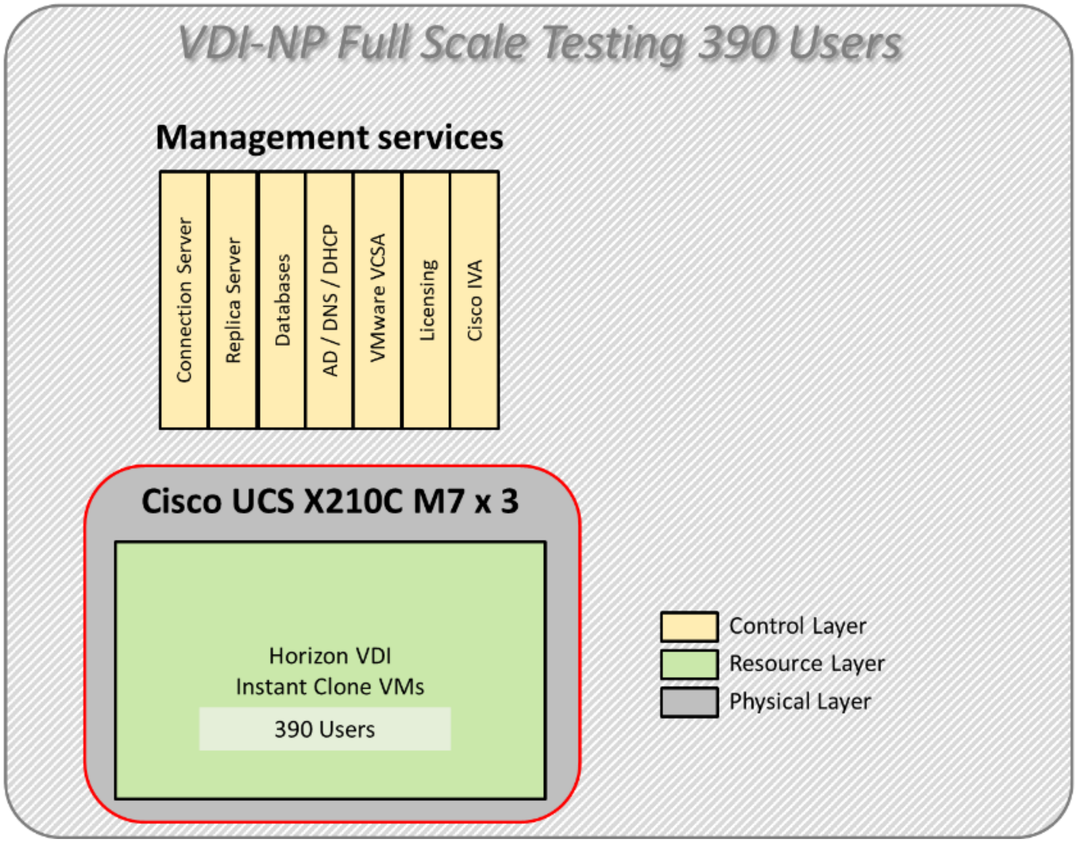

Figure 31. **Test Configuration for Full Scale VMware Horizon 8 2212 persistent Single-session OS machine VDAs (ESA vSAN cluster)**

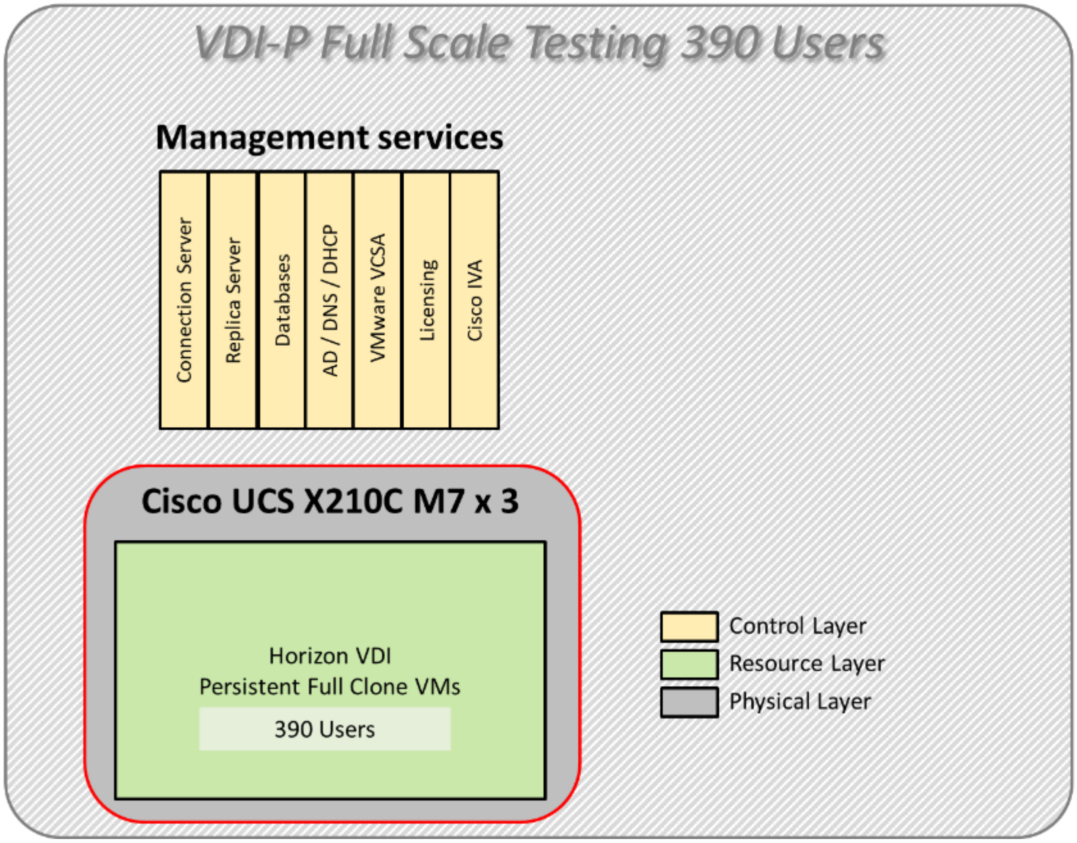

Figure 32. **Test Configuration for Full Scale VMware Horizon 8 2212 instant-clones Multi-session OS machine VDAs (ESA vSAN cluster)**

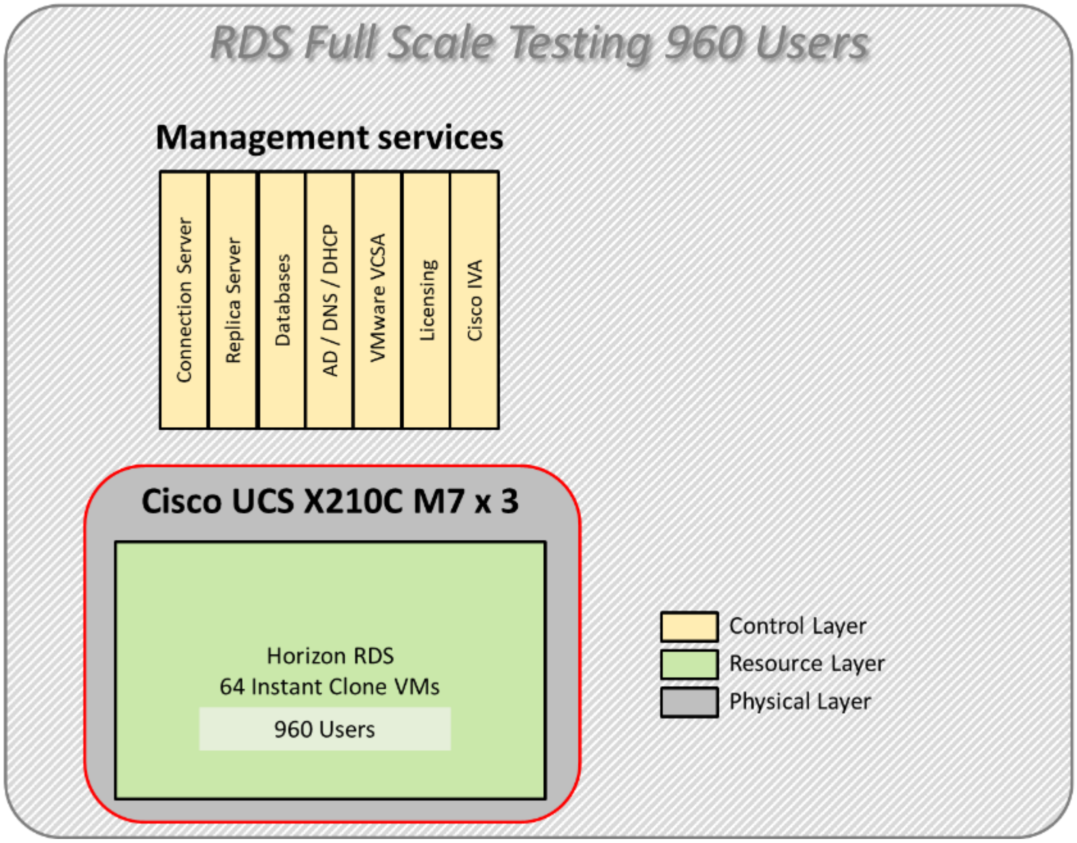

Hardware components:

- 1 Cisco UCS X9508 Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 4 Cisco UCS X210c M7 Compute Node Servers with Intel(R) Xeon(R) Gold 6448H CPU 2.40GHz 32-core processors, 1TB 4800 MHz RAM for all host blades
- Cisco UCS VIC 15231 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 1 1.6TB High Performance High Endurance and 5 1.92TB High Performance Medium Endurance NVMe drives (per blade, OSA vSAN cluster)
- 5 3.2TB High Perf High Endurance NVMe drives (per blade, ESA vSAN cluster)

Software components:

- Cisco UCS firmware 5.1(1.230052)
- VMware vSAN 8 (OSA/ESA)
- VMware ESXi 8.0 Update 1a 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

Test 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 RDSH/VDI Session under test.

Test metrics were gathered from the 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

Test Procedure

This chapter contains the following:

- Pre-Test Setup for Single and Multi-Blade Testing
- Test Run Protocol
- Success Criteria
- VSImax 4.1.x Description
- Login Enterprise Performance Testing

The following protocol was used for each test cycle in this study to ensure consistent results.

Pre-Test Setup for Single and Multi-Blade 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 blades to be tested were restarted prior to each test cycle.

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. For testing where the user session count exceeds 1000 users, we will now deem the test run successful with up to 1% session failure rate.

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. To do so, follow these steps:

1. Time 0:00:00 Start PerfMon/Esxtop Logging on the following system:
   - Infrastructure and VDI Host Blades used in the test run
2. vCenter used in the test run.
3. All Infrastructure virtual machines used in test run (AD, SQL, brokers, image mgmt., and so on)
4. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
5. Time 0:05: Boot Virtual Desktops/RDS Virtual Machines using View Connection server.
6. The boot rate should be around 10-12 virtual machines per minute per server.
7. Time 0:06 First machines boot.
8. Time 0:30 Single Server or Scale target number of desktop virtual machines booted on 1 or more blades.
9. No more than 30 minutes for boot up of all virtual desktops is allowed.
10. Time 0:35 Single Server or Scale target number of desktop virtual machines desktops available on View Connection Server.
11. Virtual machine settling time.

12. No more than 60 Minutes of rest time is allowed after the last desktop is registered on the XD Studio or available in the View Connection Server dashboard. Typically, a 30-45-minute rest period is sufficient.

13. Time 1:35 Start Login VSI 4.1.x Office Worker Benchmark Mode Test, setting auto-logoff time at 15 minutes, with Single Server or Scale target number of desktop virtual machines utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).

14. Time 2:23 Single Server or Scale target number of desktop virtual machines desktops launched (48 minute benchmark launch rate).

15. Time 2:25 All launched sessions must become active. id test run within this window.

16. Time 2:40 Login VSI Test Ends (based on Auto Logoff 15 minutes period designated above).

17. Time 2:55 All active sessions logged off.

18. Time 2:57 All logging terminated; Test complete.

19. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows machines.

20. Time 3:30 Reboot all hypervisor hosts.

21. Time 3:45 Ready for the new test sequence.

Success Criteria

Our pass criteria for this testing is as follows:

- Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 4.1.x 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 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 state
- Within 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 in our testing. VMware HCI with Cisco UCS and VMware Horizon 8 2212 on VMware ESXi 8.0 Update 1a Test Results.

The purpose of this testing is to provide the data needed to validate VMware Horizon Remote Desktop Sessions (RDS) and VMware Horizon Virtual Desktop (VDI) instant-clones and VMware Horizon Virtual Desktop (VDI) full-clones models using ESXi and vCenter to virtualize Microsoft Windows 11 desktops and Microsoft Windows Server 2019 sessions on Cisco UCS X210c M7 Compute Node Servers using the VMware vSAN cluster.

The information contained in this section provides data points that a customer may reference in designing their 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 products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish single blade performance and multi-blade, 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 HSD 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 number 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 shows 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 its true maximum user capacity is.

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.

Calculating 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 33. Sample of a VSI Max Response Time Graph, Representing a Normal Test

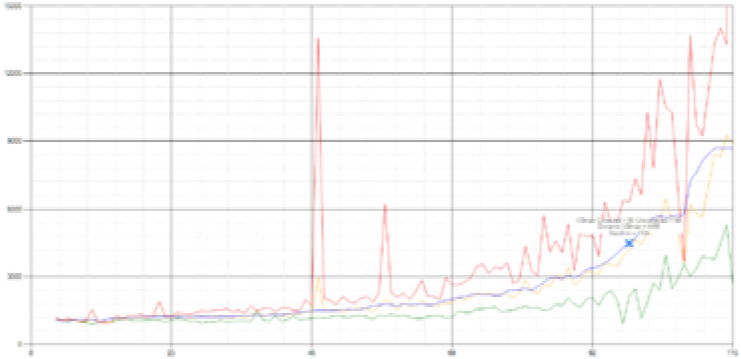

Figure 34. Sample of a VSI Test Response Time Graph with a Performance Issue

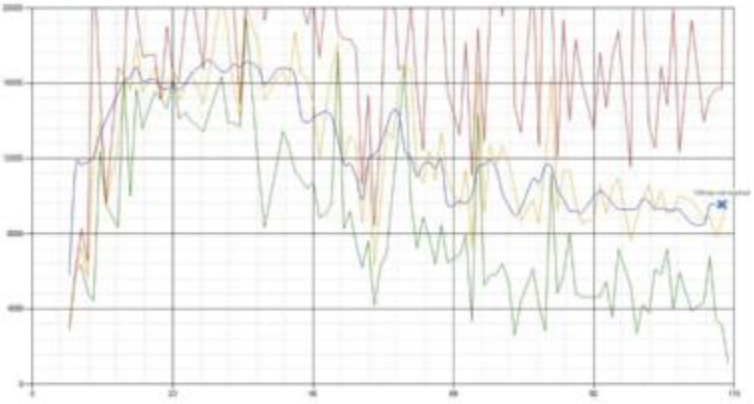

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached, and the number of sessions ran successfully.

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 times 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 VSImax models, this weighting much better represents system performance. All actions have very similar weight in the VSImax total. The following weighting of the response times is applied.

The following actions are part of the VSImax v4.1.x 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.x, 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 the 15 lowest VSI response time samples are taken from the entire test; the lowest two samples are removed. and the 13 remaining samples are averaged. The result is the Baseline.

To summarize:

- 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 number 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 are 60, then latest 5 + 24 (=40 percent of 60) = 31 response time measurement is 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 the 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 VSIbase + a 1000 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 as the maximum performance degradation to be considered acceptable.

In VSImax v4.1.x this latency threshold is fixed to 1000ms, 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 1000ms (weighted).

The threshold for the total response time is average weighted baseline response time + 1000ms.

When the system has a weighted baseline response time average of 1500ms, the maximum average response time may not be greater than 2500ms (1500+1000). If the average baseline is 3000 the maximum average response time may not be greater than 4000ms (3000+1000).

When the threshold is not exceeded by the average VSI response time during the test, VSImax is not hit, and the number 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.x 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 and 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 VSImax v4.1.x, and the higher VSImax is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSImax method is introduced: VSImax v4.1.x. This methodology gives much better insight into system performance and scales to extremely large systems.

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 blade 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 ensure 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 18. 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 registered on the XD Studio. (typically, a 30-45 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 the 15-minute duration)
Logoff	Sessions finish executing the Login VSI workload and logoff

Login Enterprise Performance Testing

The Login Enterprise Platform is the next step in the evolution of Login VSI in testing VDI environments. It focuses on monitoring and optimizing the End-User Experience (EUX) in virtual desktop infrastructure. EUX refers to the overall experience that end-users have while interacting with their virtual desktops, applications, and other IT resources.

Login Enterprise workload, like Login VSI workload before, uses a standard collection of desktop application software on each VDI desktop testing instance. A configurable launcher system connects a specified number of simulated users to available desktops within the environment. When the simulated user is connected, a login script configures the user environment and starts a defined workload. These simulations help evaluate response times, latency, and overall system performance. Each launcher system can launch connections to several VDI desktops (target machines). A centralized management console configures and manages the launchers and the Login Enterprise environment.

The Knowledge Worker the Login Enterprise workloads used in testing:

- Microsoft Outlook—Browse messages
- Microsoft Edge—Browse websites and open a YouTube style video (480p movie trailer) three times in every loop
- Microsoft Word—Start one instance to measure response time and another to review and edit a document
- Microsoft Excel—Open a large, randomized sheet

- Microsoft PowerPoint—Review and edit a presentation

Test Results

This chapter contains the following:

- Single-Server Recommended Maximum Workload Testing
- Full Scale Workload Testing

Single-Server Recommended Maximum Workload Testing

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:

- 195 VDI Non-Persistent sessions (Instant Clone machines with floating assignment)
- 195 VDI Persistent sessions (Full Virtual machines with dedicated assignments)
- 480 instant clones Multi-session OS RDS sessions (floating assignment)

OSA vSAN: Single-Server Recommended Maximum Workload for Non-persistent Single-session OS Random Sessions with 195 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, 1TB 4800 MHz RAM is 195 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 35. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | VSI Score

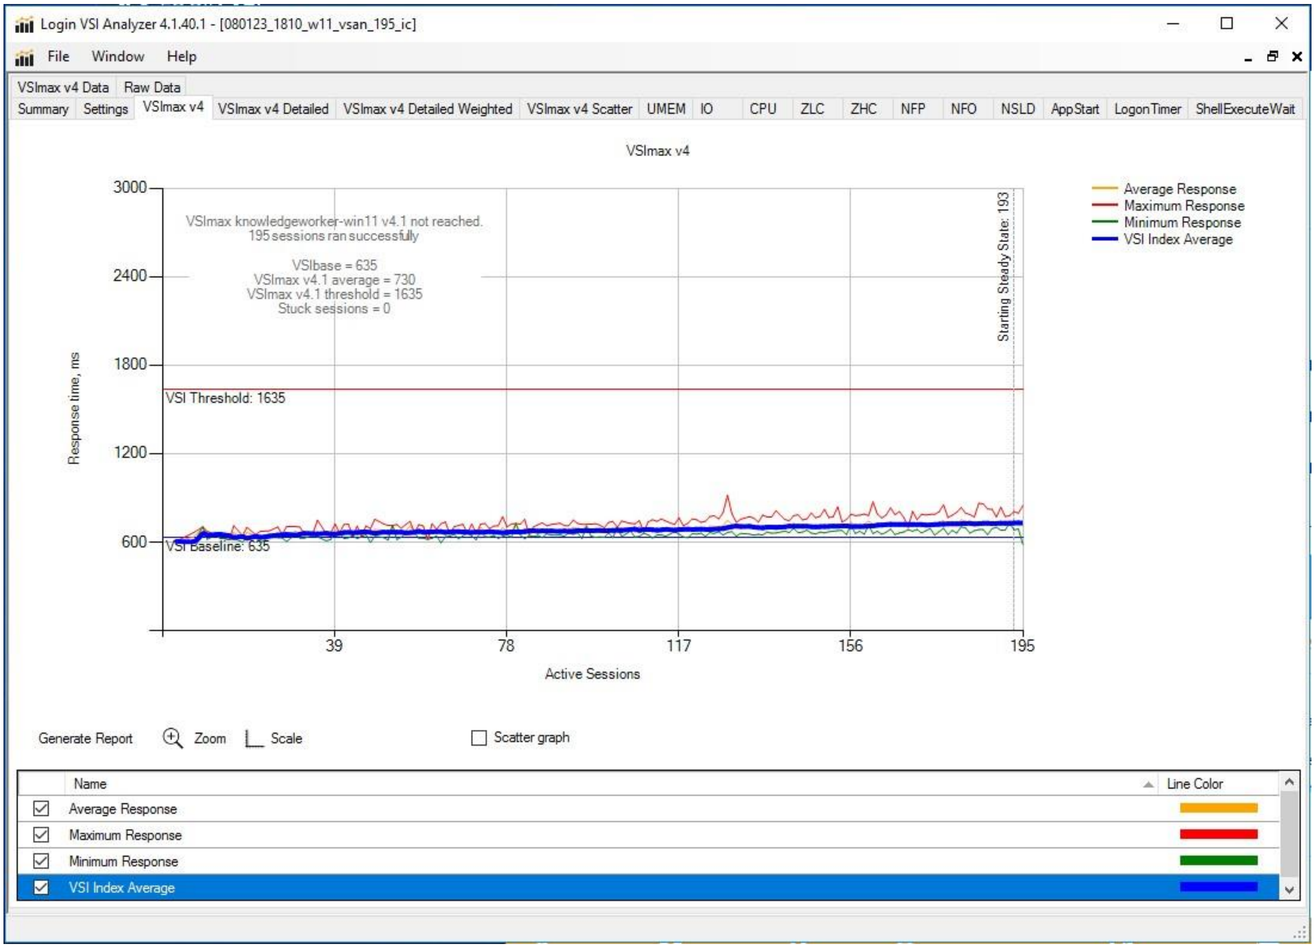

```
VSImax knowledgeworker-win11 v4.1 not reached.
195 sessions ran successfully

VSIbase = 635
VSImax v4.1 average = 730
VSImax v4.1 threshold = 1635
Stuck sessions = 0
```

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Performance data for the server running the workload is shown below:

Figure 36. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host CPU Utilization

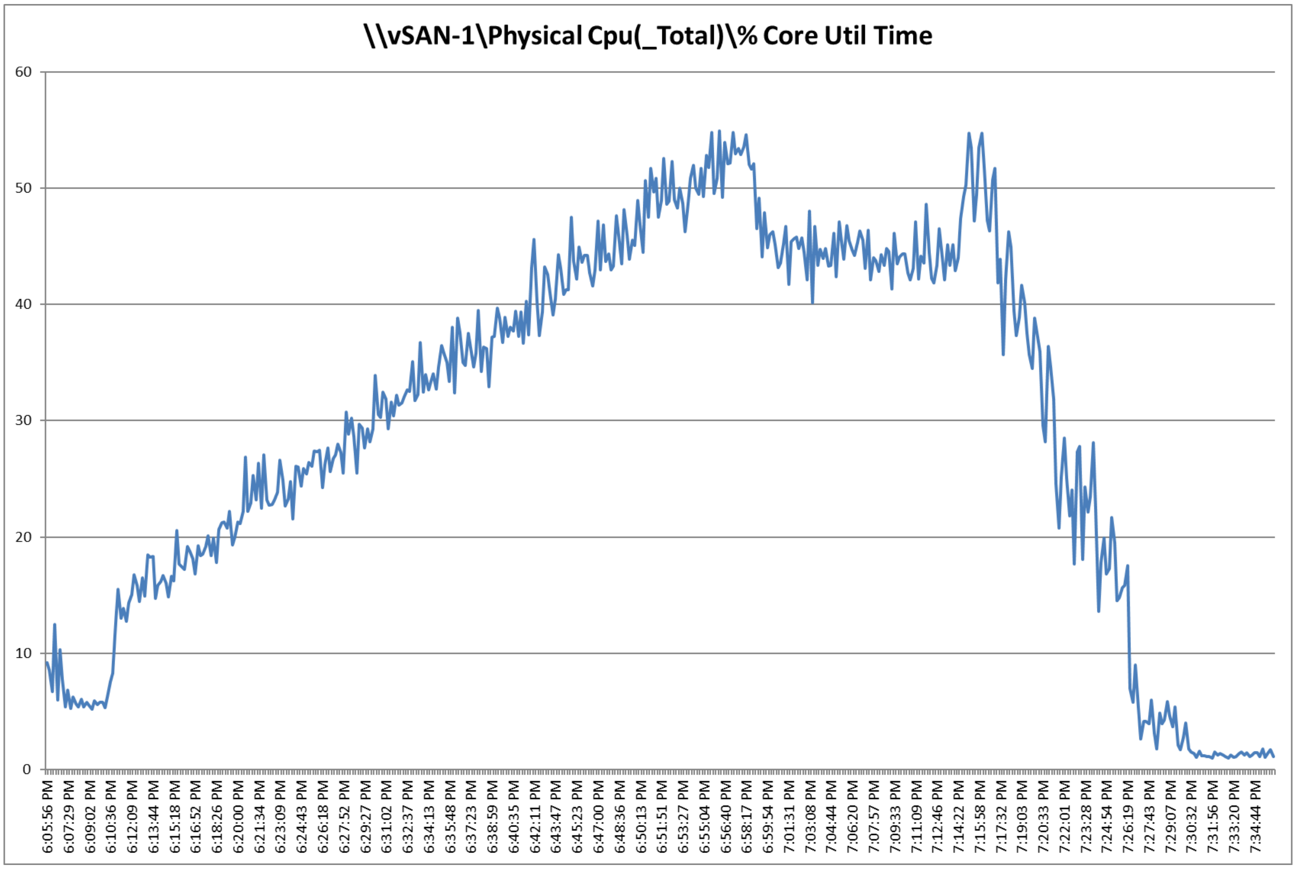

Figure 37. **Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Memory Utilization**

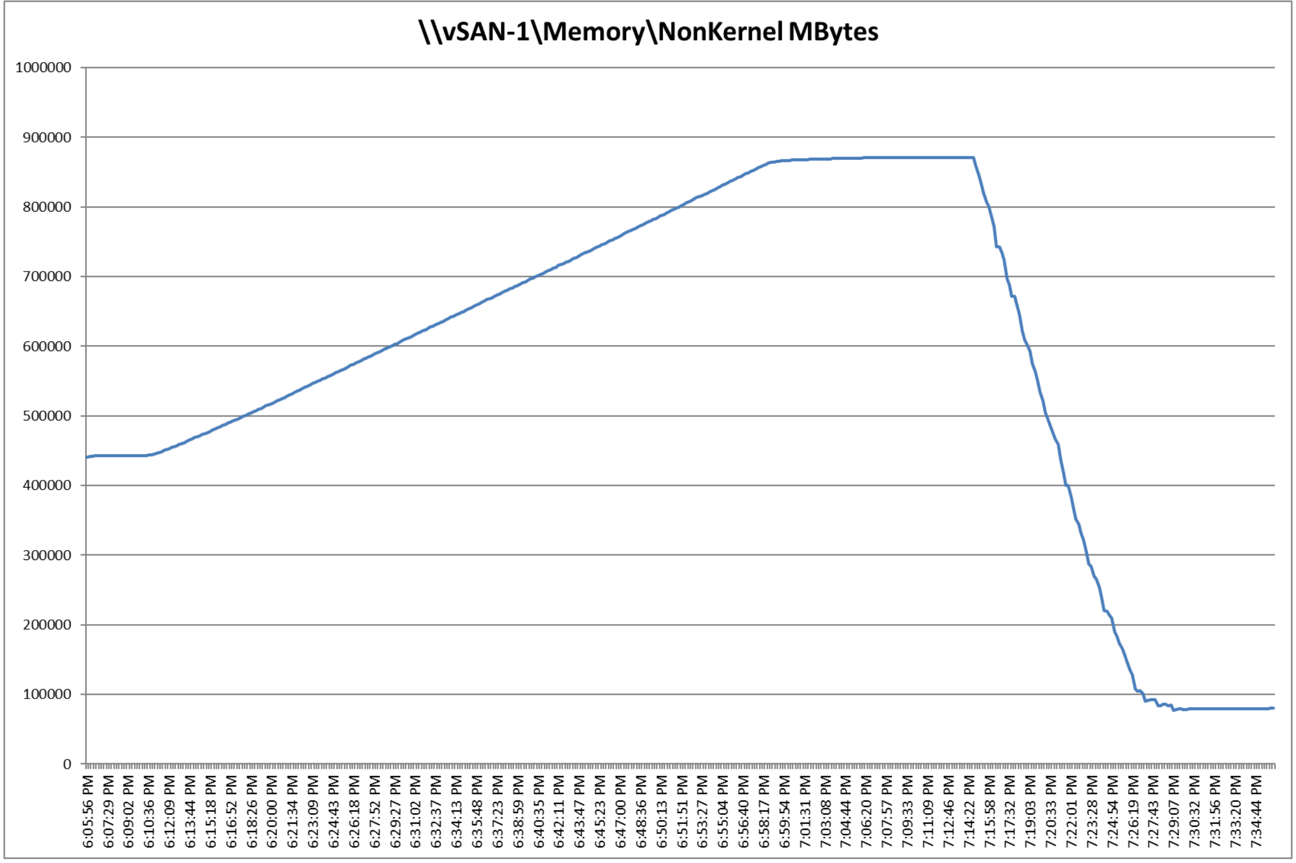

Figure 38. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Network Utilization

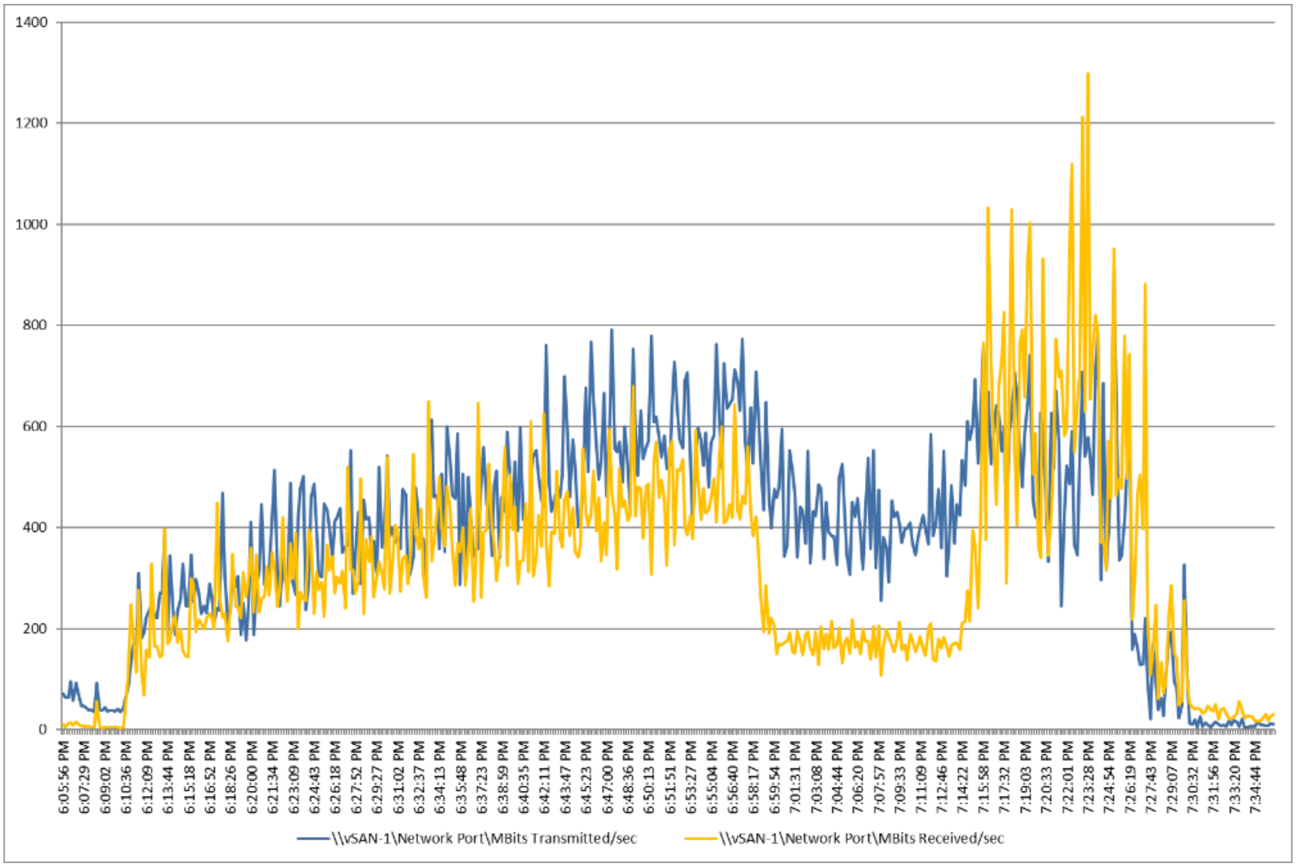

Login Enterprise performance data is shown below:

Figure 39. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | EUX Score

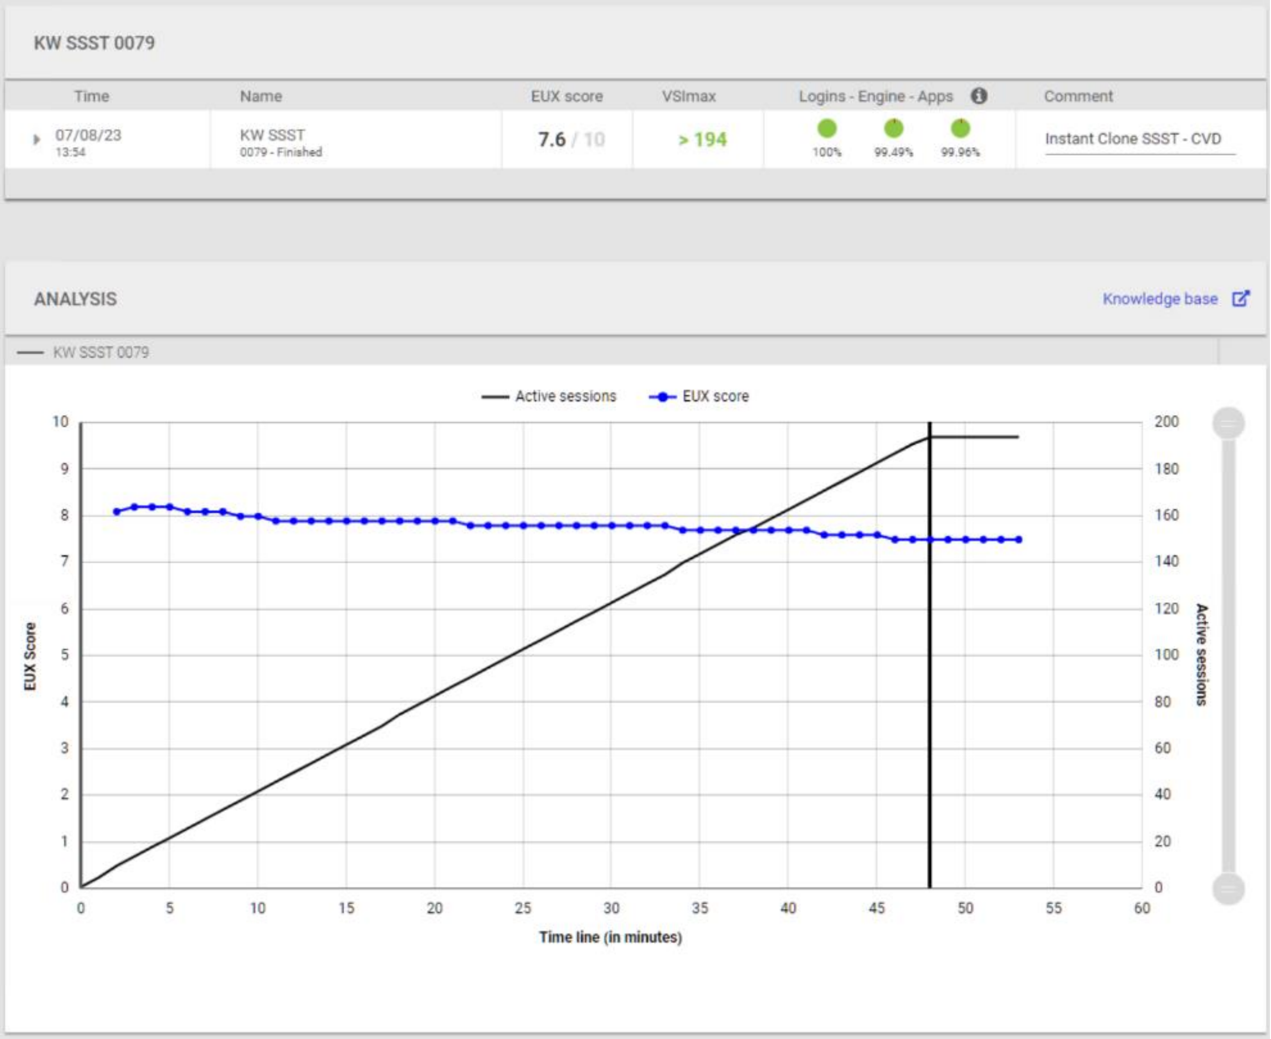

KW SSST 0079

Time	Name	EUX score	VSImax	Logins - Engine - Apps	Comment
07/08/23 13:54	KW SSST 0079 - Finished	7.6 / 10	> 194	100% 99.49% 99.96%	Instant Clone SSST - CVD

ANALYSIS

Knowledge base

OSA vSAN: Single-Server Recommended Maximum Workload for Persistent Single-session OS dedicated sessions with 195 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, 1TB 4800 MHz RAM is 195 Windows 11 64-bit VDI Persistent virtual machines with 2 vCPU and 4GB RAM.

Login VSI performance data is as shown below:

Figure 40. Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Score

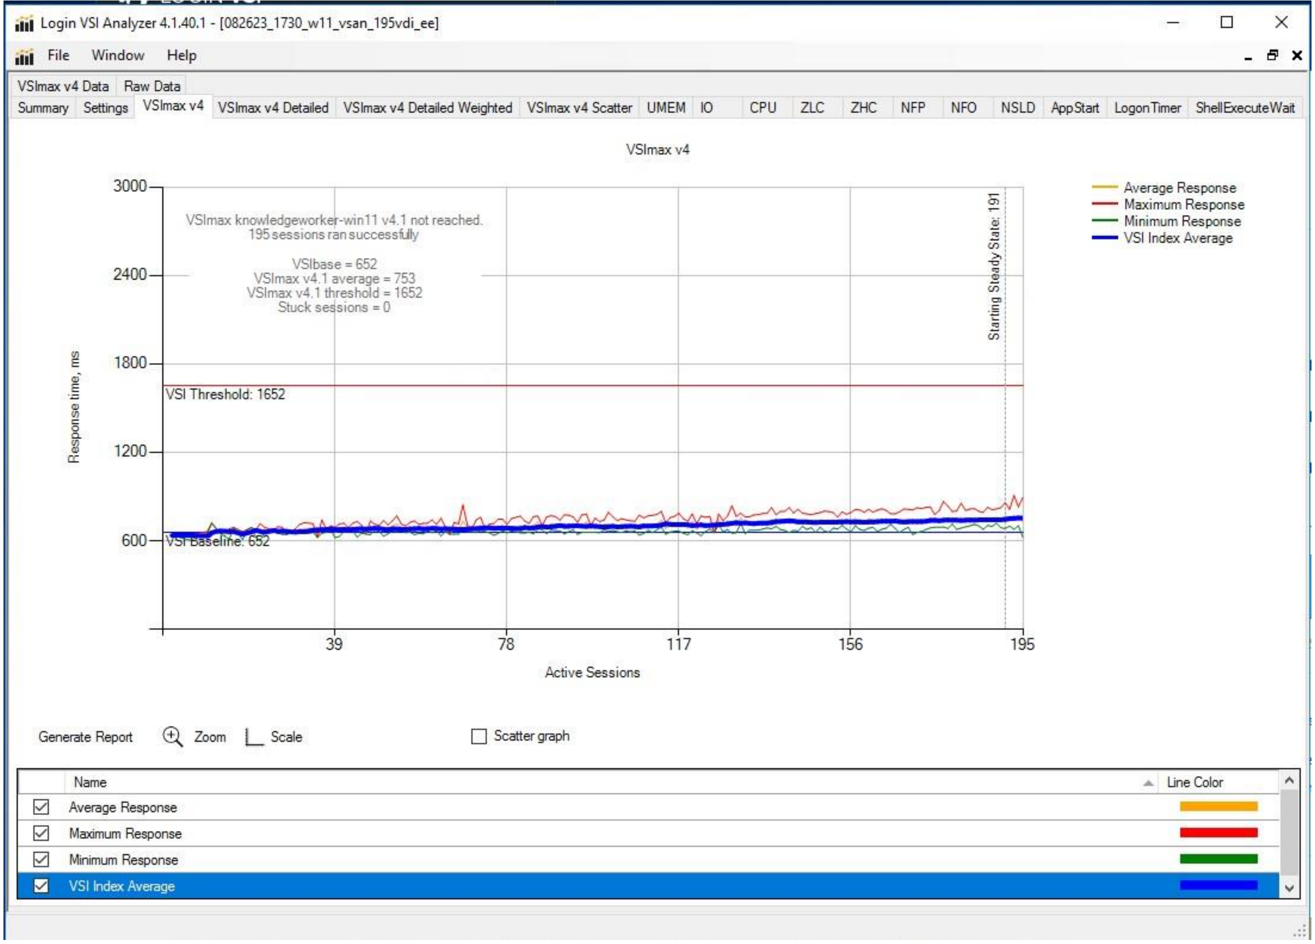

VSImax knowledgeworker-win11 v4.1 not reached.
195 sessions ran successfully

VSIbase = 652
VSImax v4.1 average = 753
VSImax v4.1 threshold = 1652
Stuck sessions = 0

☐ Scatter graph

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Performance data for the server running the workload is shown below:

Figure 41. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host CPU Utilization**

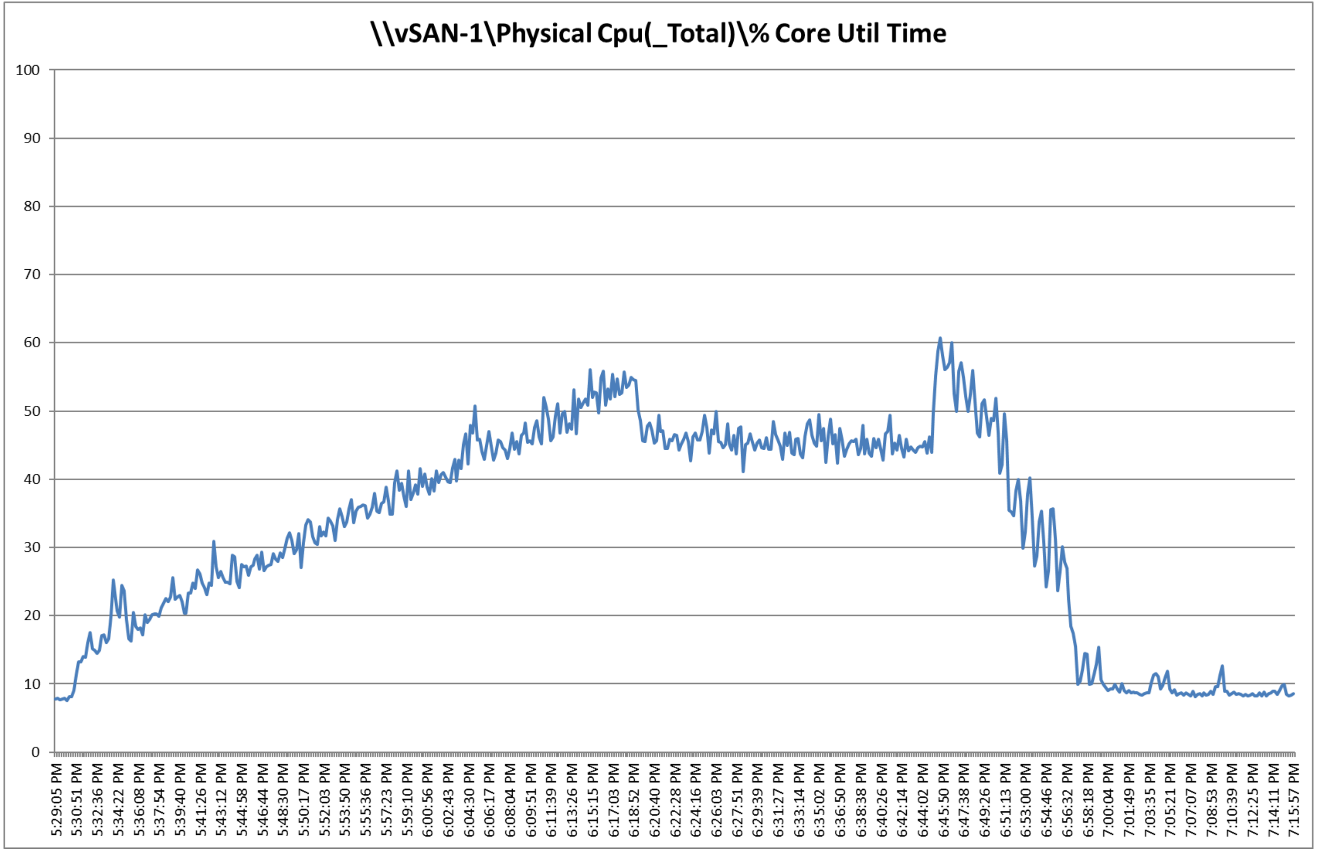

Figure 42. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Memory Utilization**

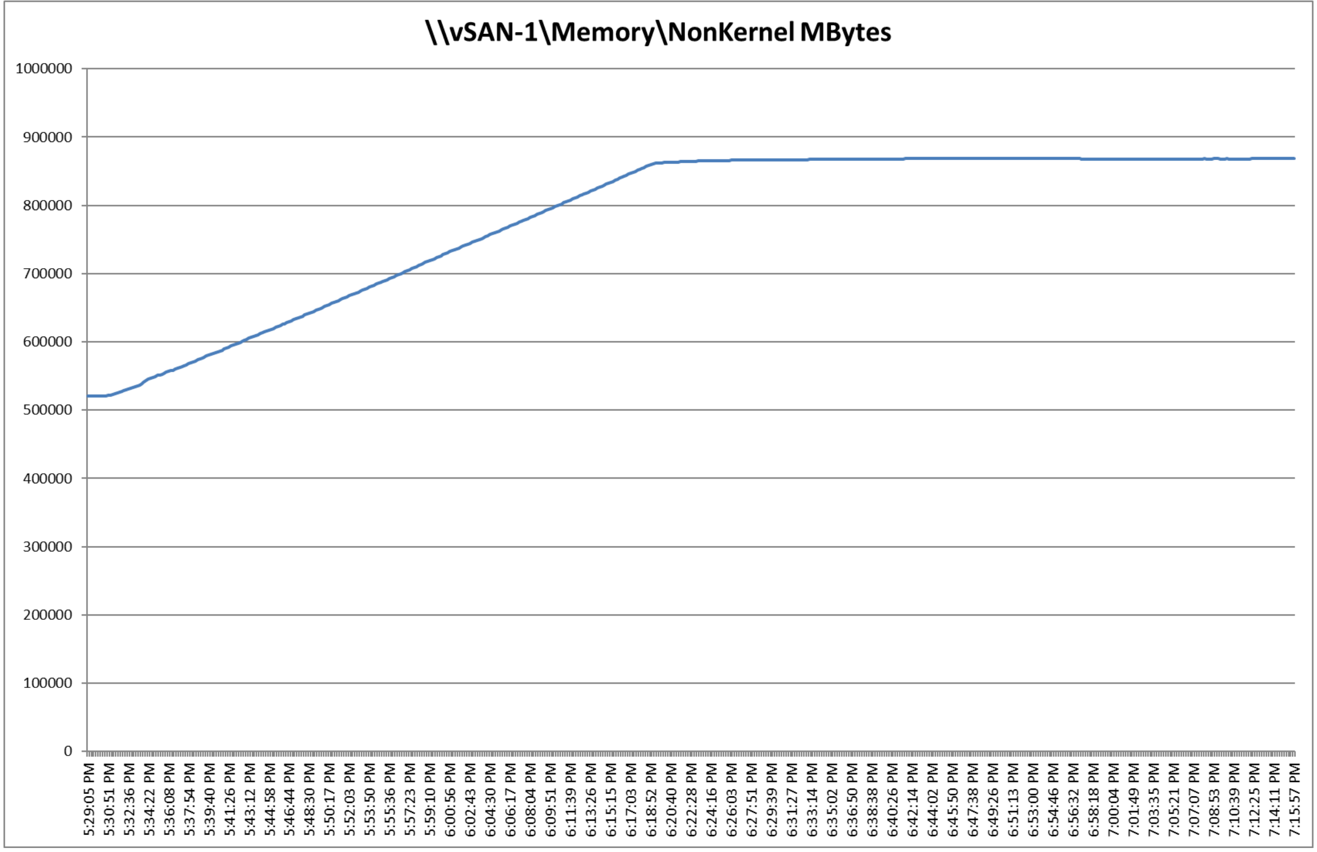

Figure 43. **Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Network Utilization**

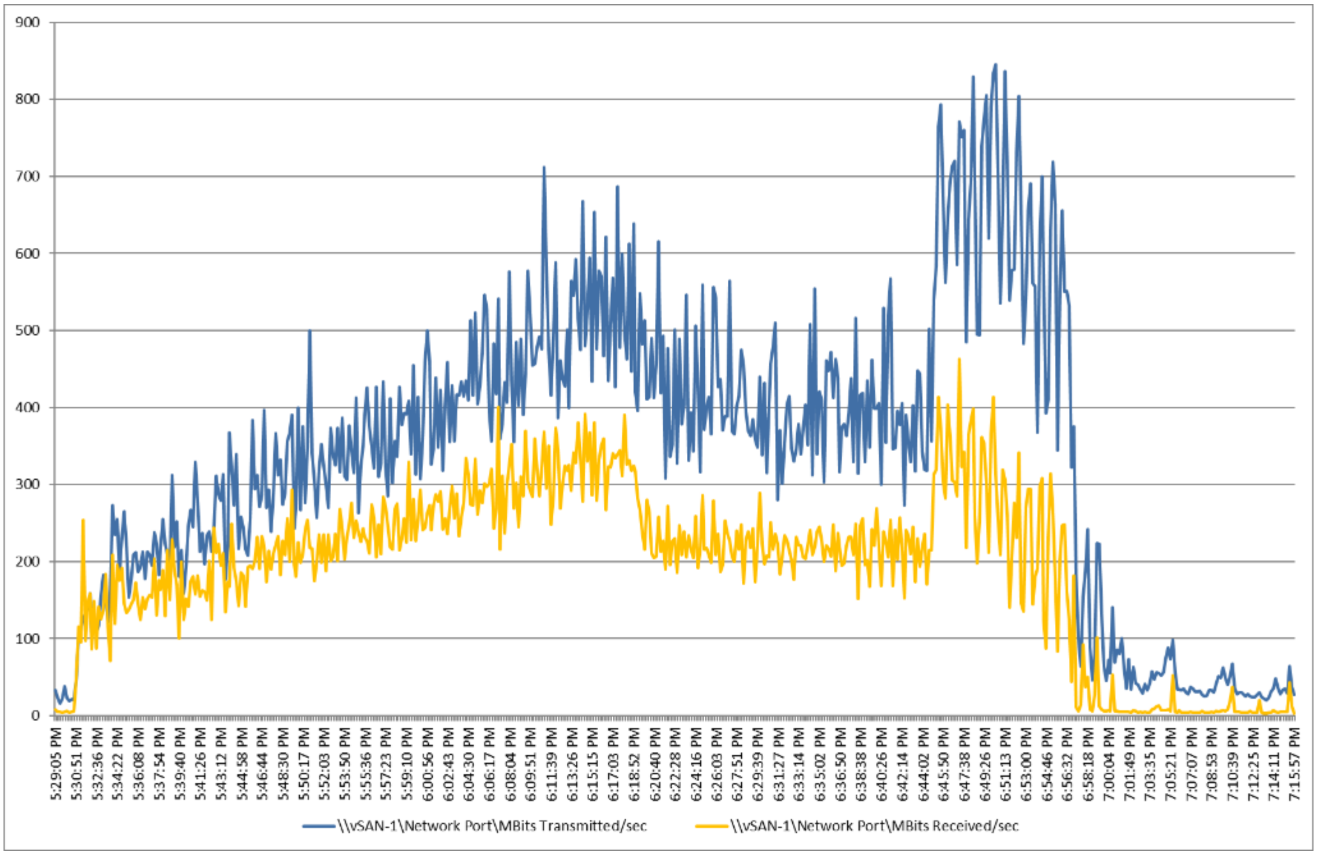

Login Enterprise performance data is shown below:

Figure 44. Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | EUX Score

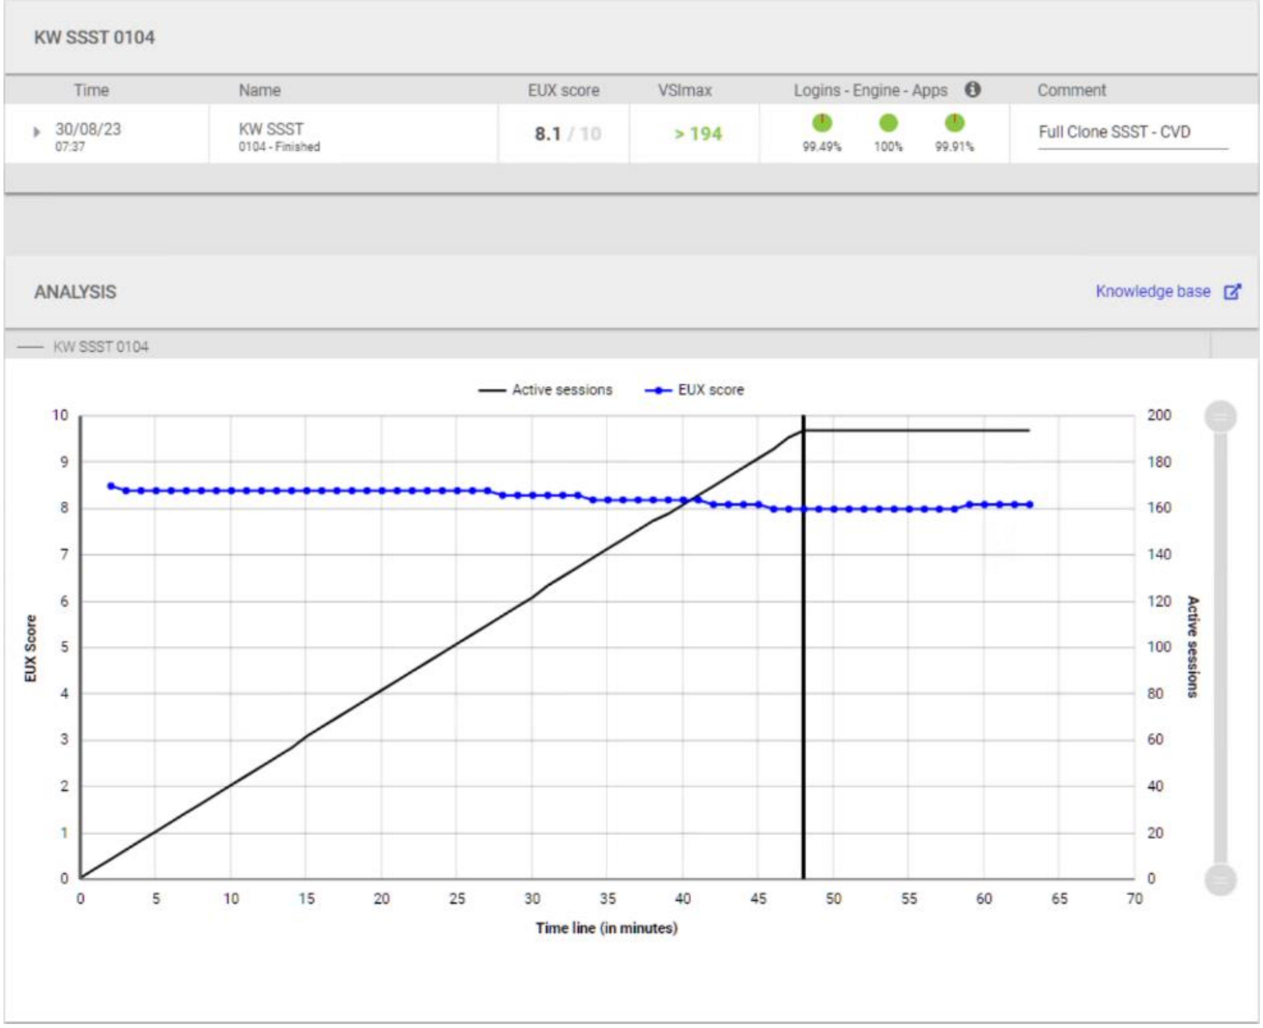

KW SSST 0104

Time	Name	EUX score	VSImax	Logins - Engine - Apps	Comment
30/08/23 07:37	KW SSST 0104 - Finished	8.1 / 10	> 194	99.49% 100% 99.91%	Full Clone SSST - CVD

ANALYSIS

Knowledge base

OSA vSAN: 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, 1TB 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 45. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Score**

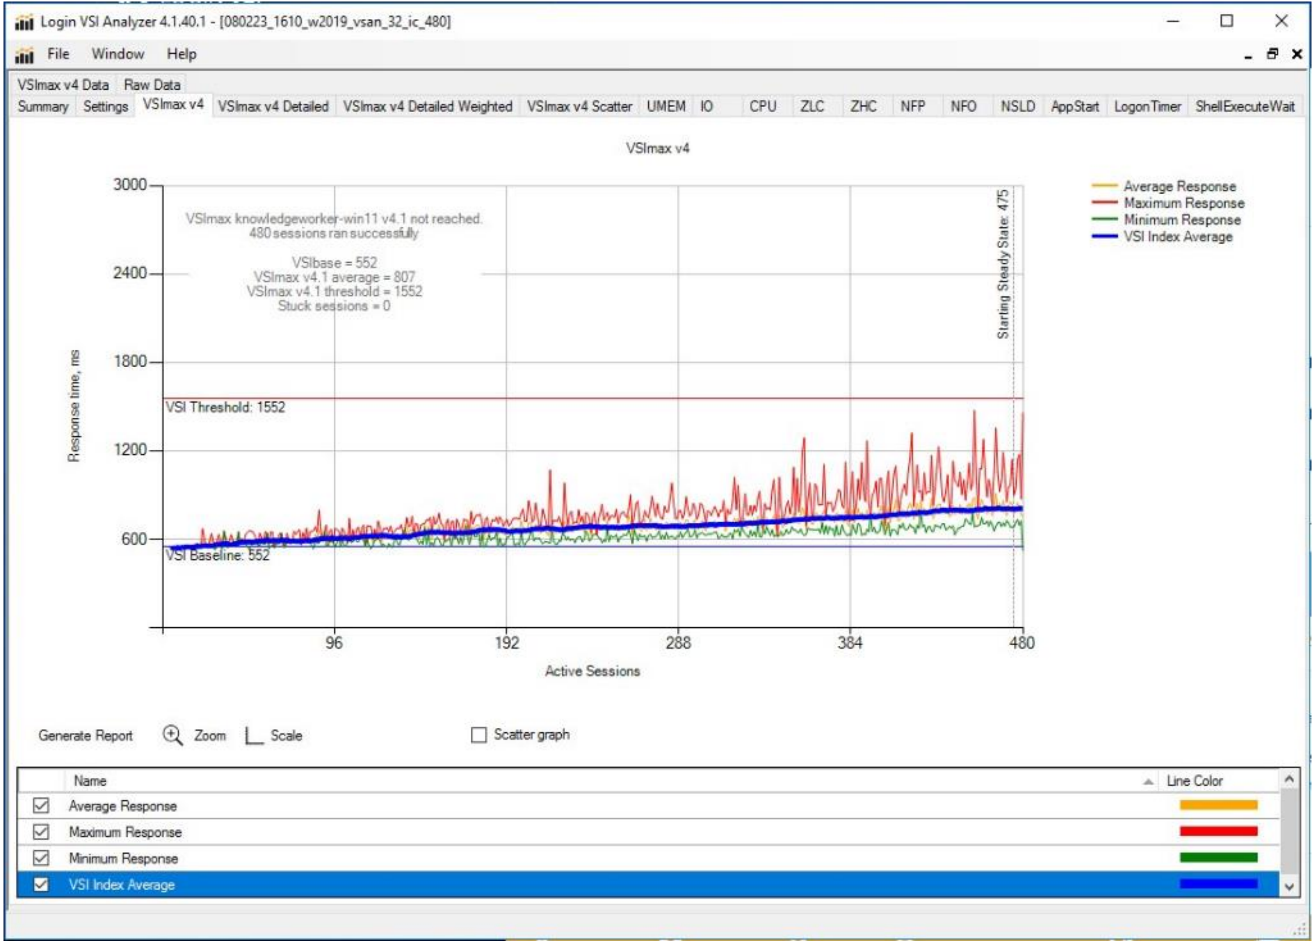

VSImax knowledgeworker-win11 v4.1 not reached.
480 sessions ran successfully

VSIbase = 552
VSImax v4.1 average = 807
VSImax v4.1 threshold = 1552
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Performance data for the server running the workload is shown below:

Figure 46. Single Server Recommended Maximum Workload VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host CPU Utilization

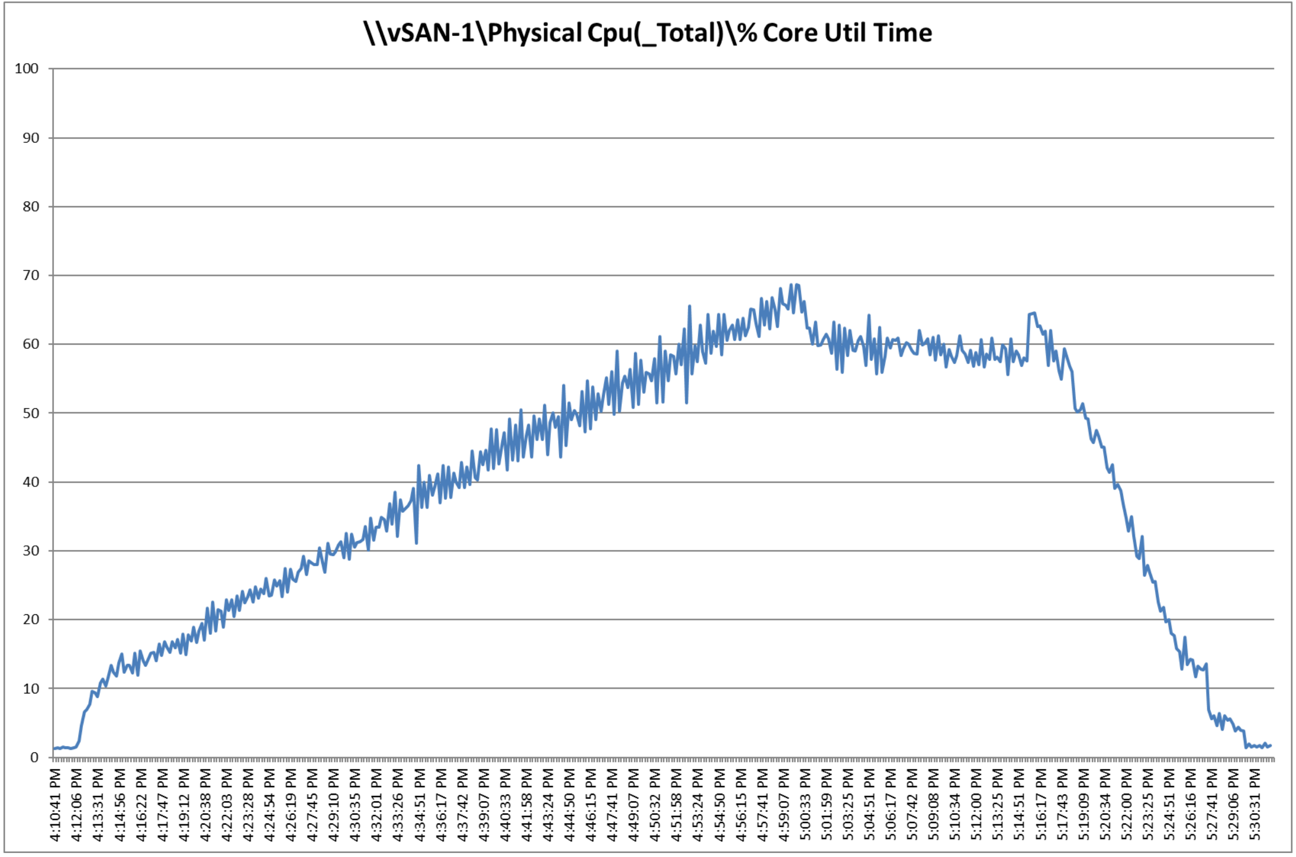

Figure 47. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Memory Utilization**

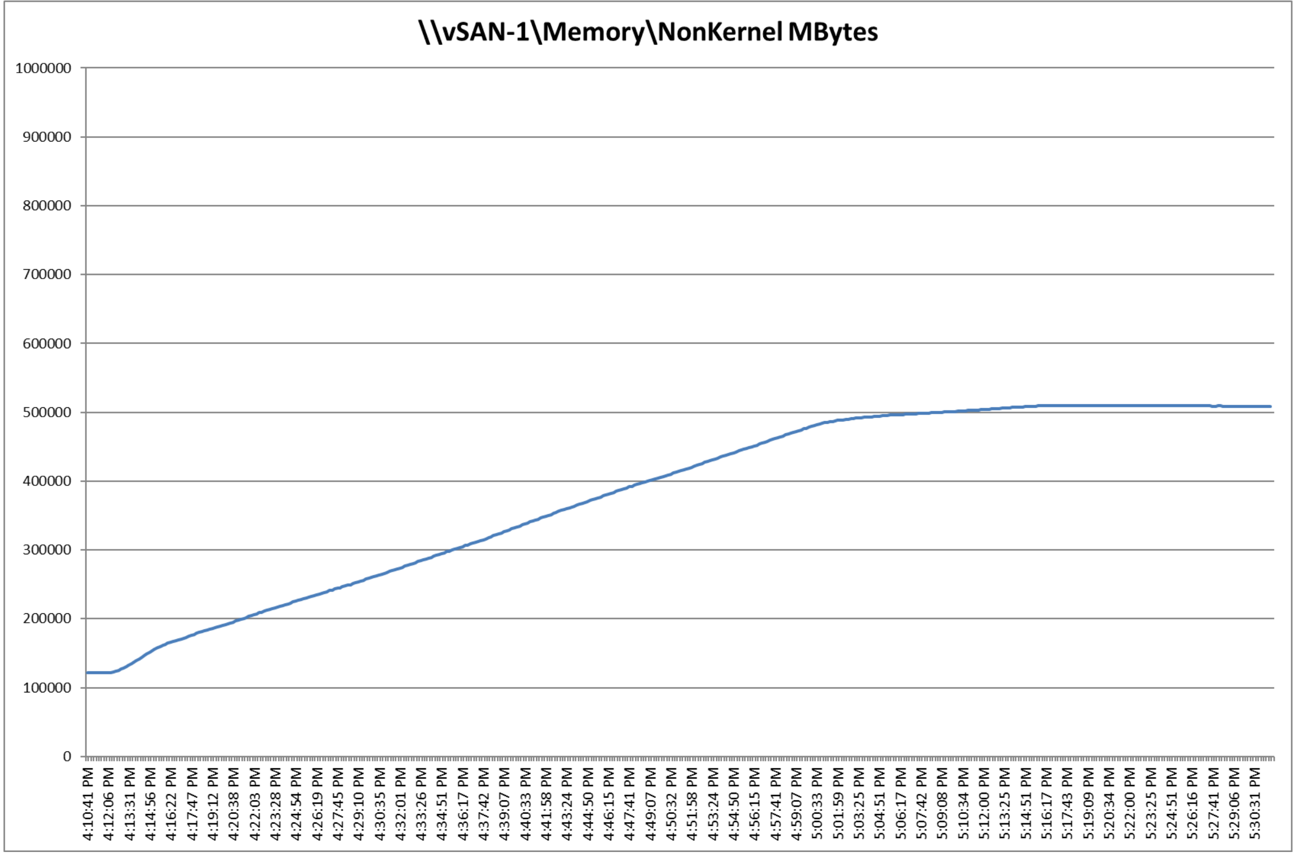

Figure 48. **Single Server | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Network Utilization**

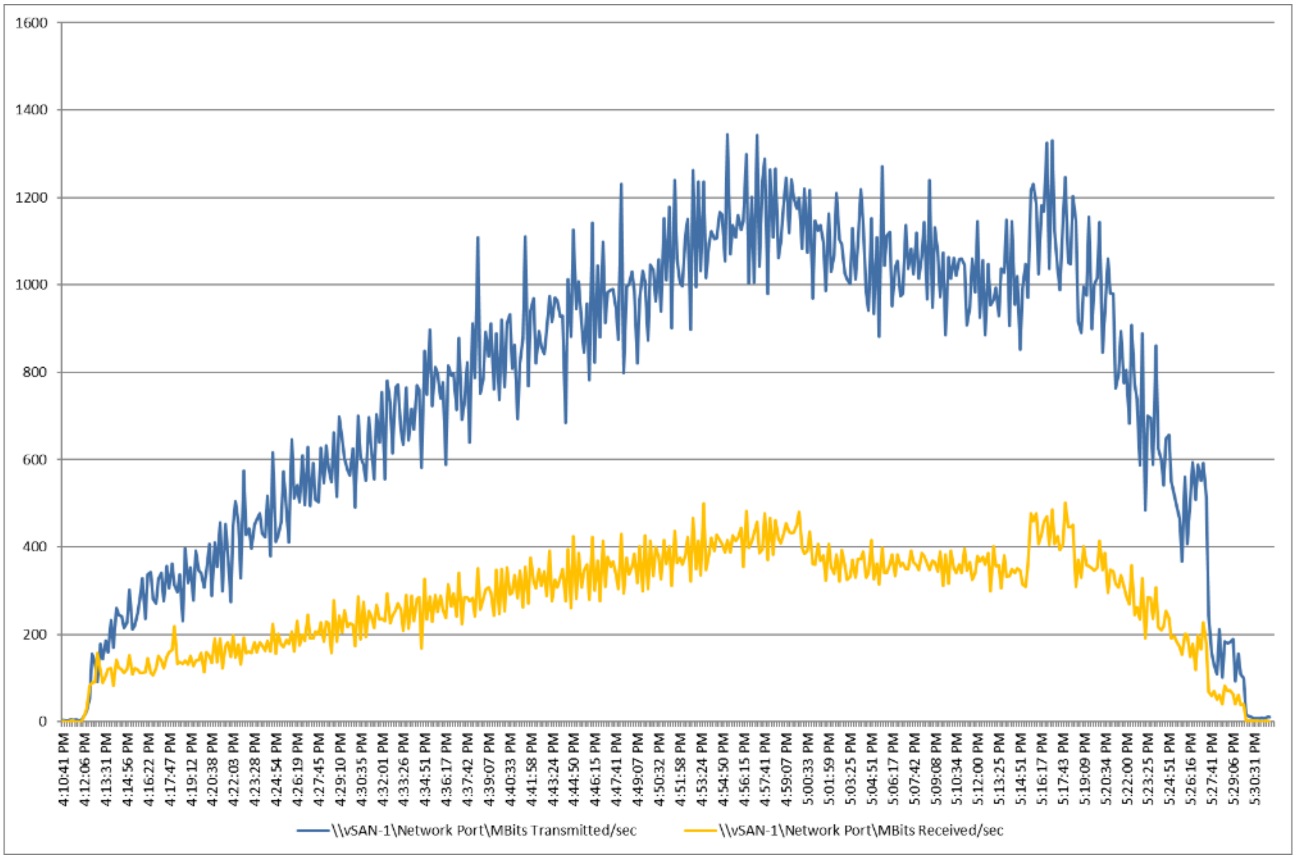

Performance data for the RDS Virtual Machine running the workload is shown below:

Figure 49. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Virtual Machine CPU Utilization**

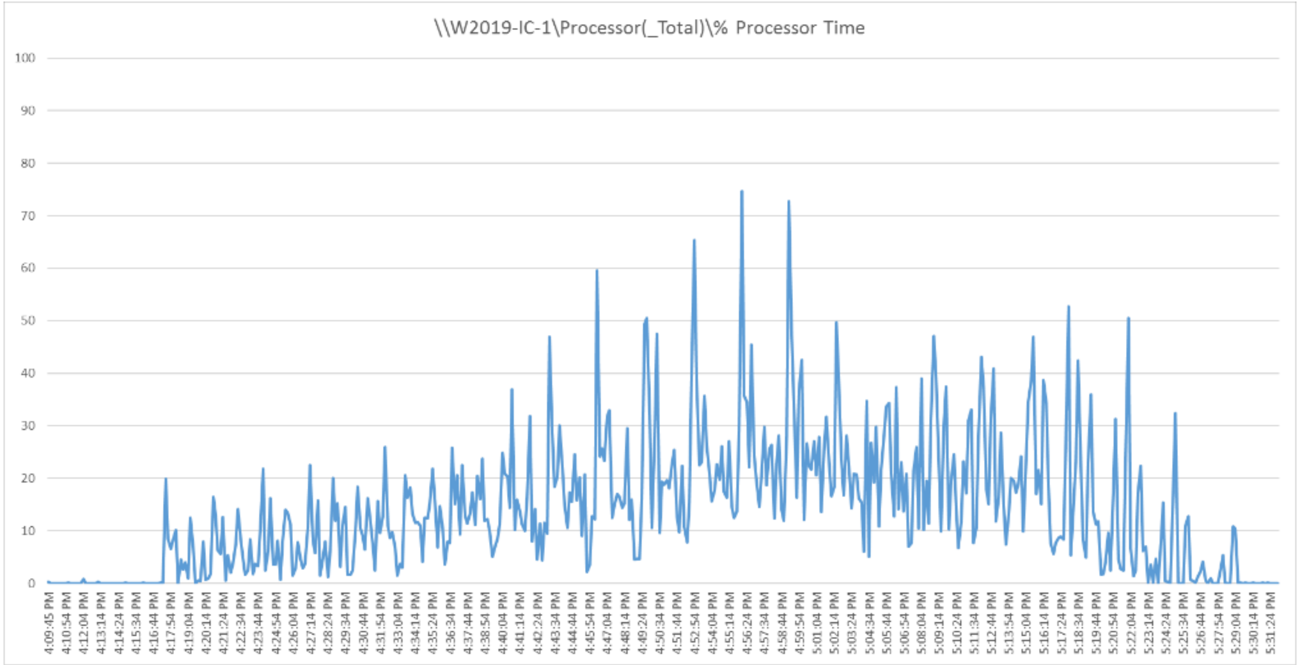

Figure 50. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Virtual Machine Memory Utilization

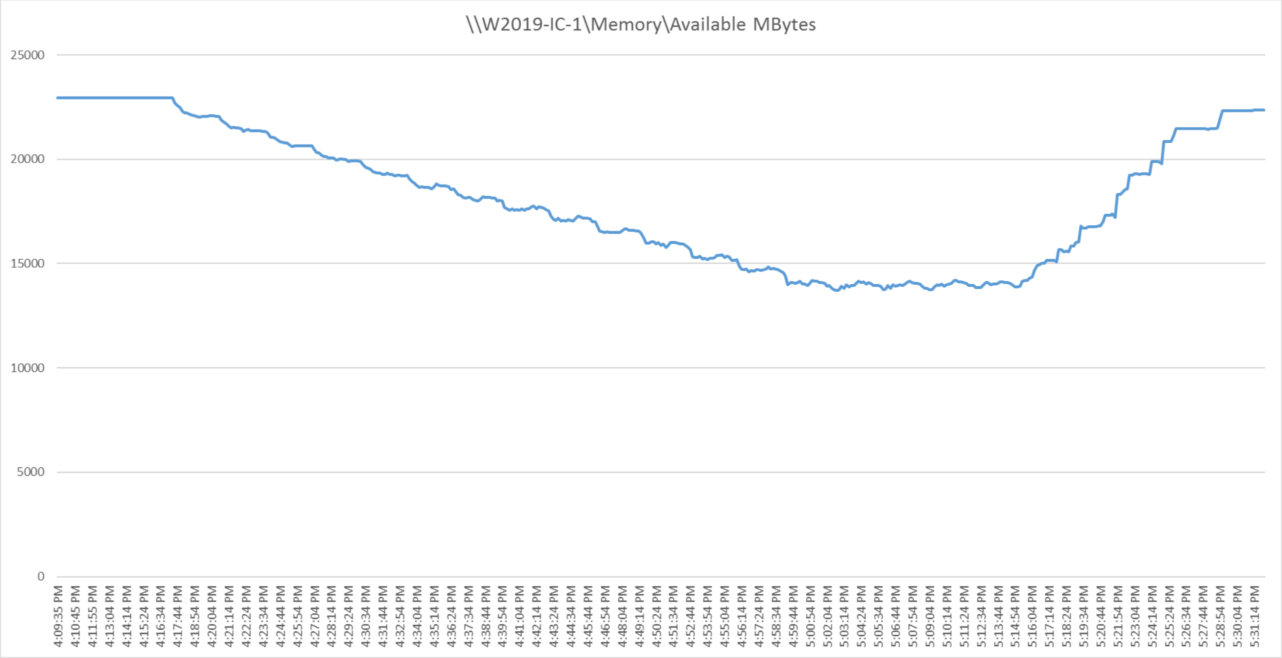

Figure 51. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Network Utilization

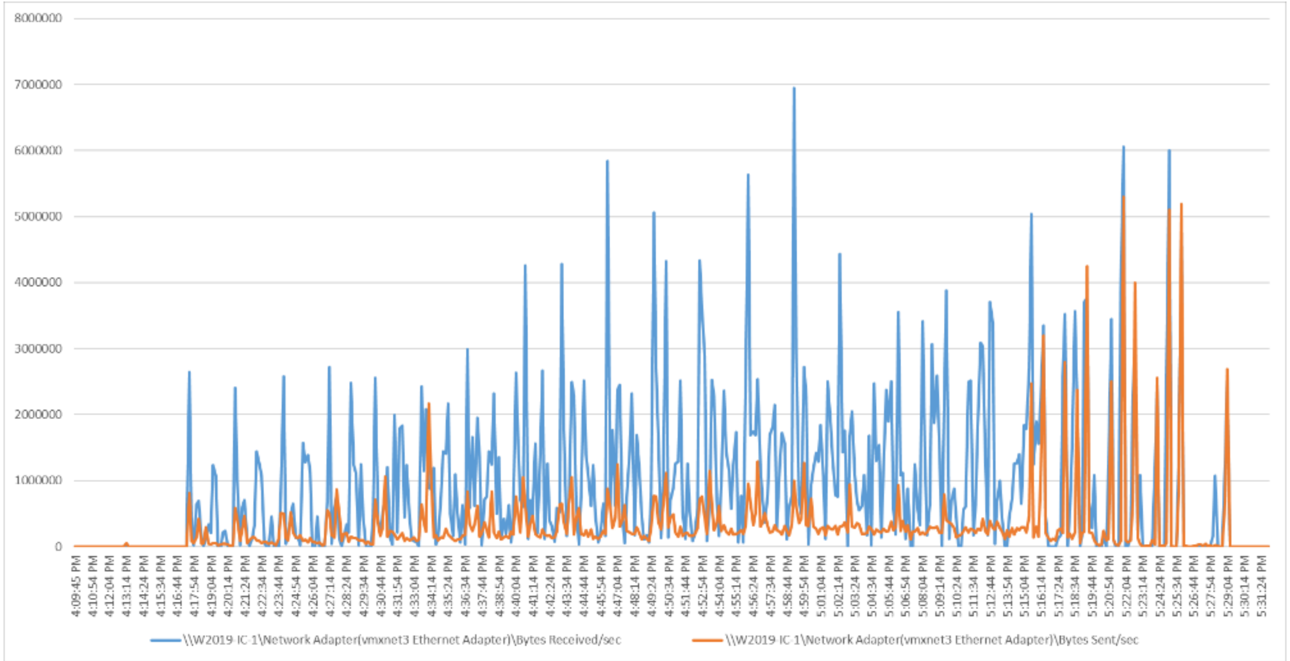

Login Enterprise performance data is shown below:

Figure 52. Single Server | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | EUX Score

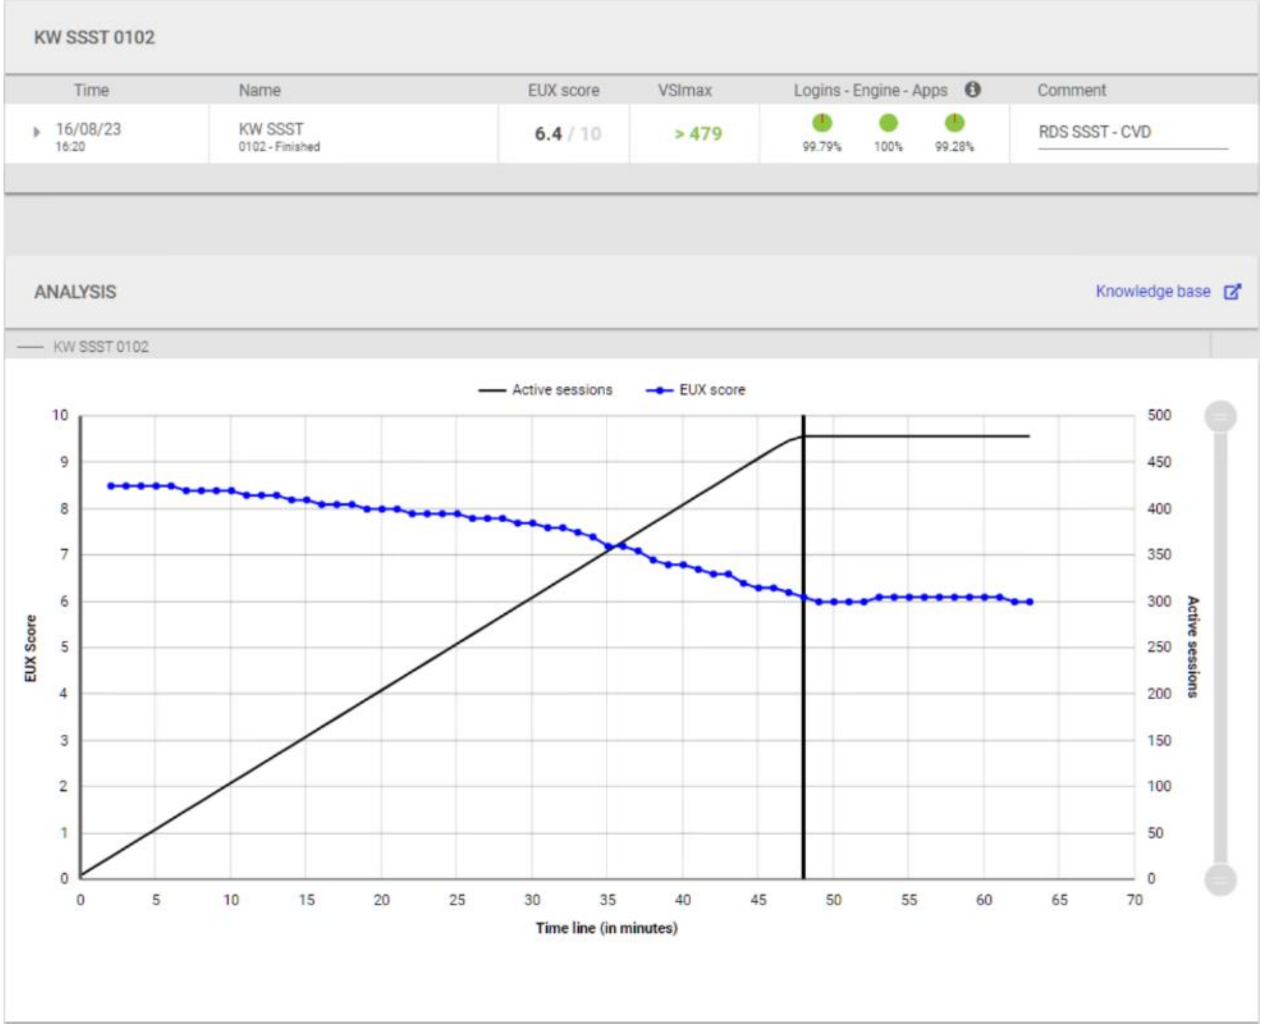

KW SSST 0102

Time	Name	EUX score	VSImax	Logins - Engine - Apps	Comment
16/08/23 16:20	KW SSST 0102 - Finished	6.4 / 10	> 479	99.79% 100% 99.28%	RDS SSST - CVD

ANALYSIS

Knowledge base

ESA vSAN: Single-Server Recommended Maximum Workload for Non-persistent Single-session OS Random Sessions with 195 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, 1TB 4800 MHz RAM is 195 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 53. Single Server | VMware Horizon 8 2212 non-persistent Single-session OS machine VDAs | VSI Score

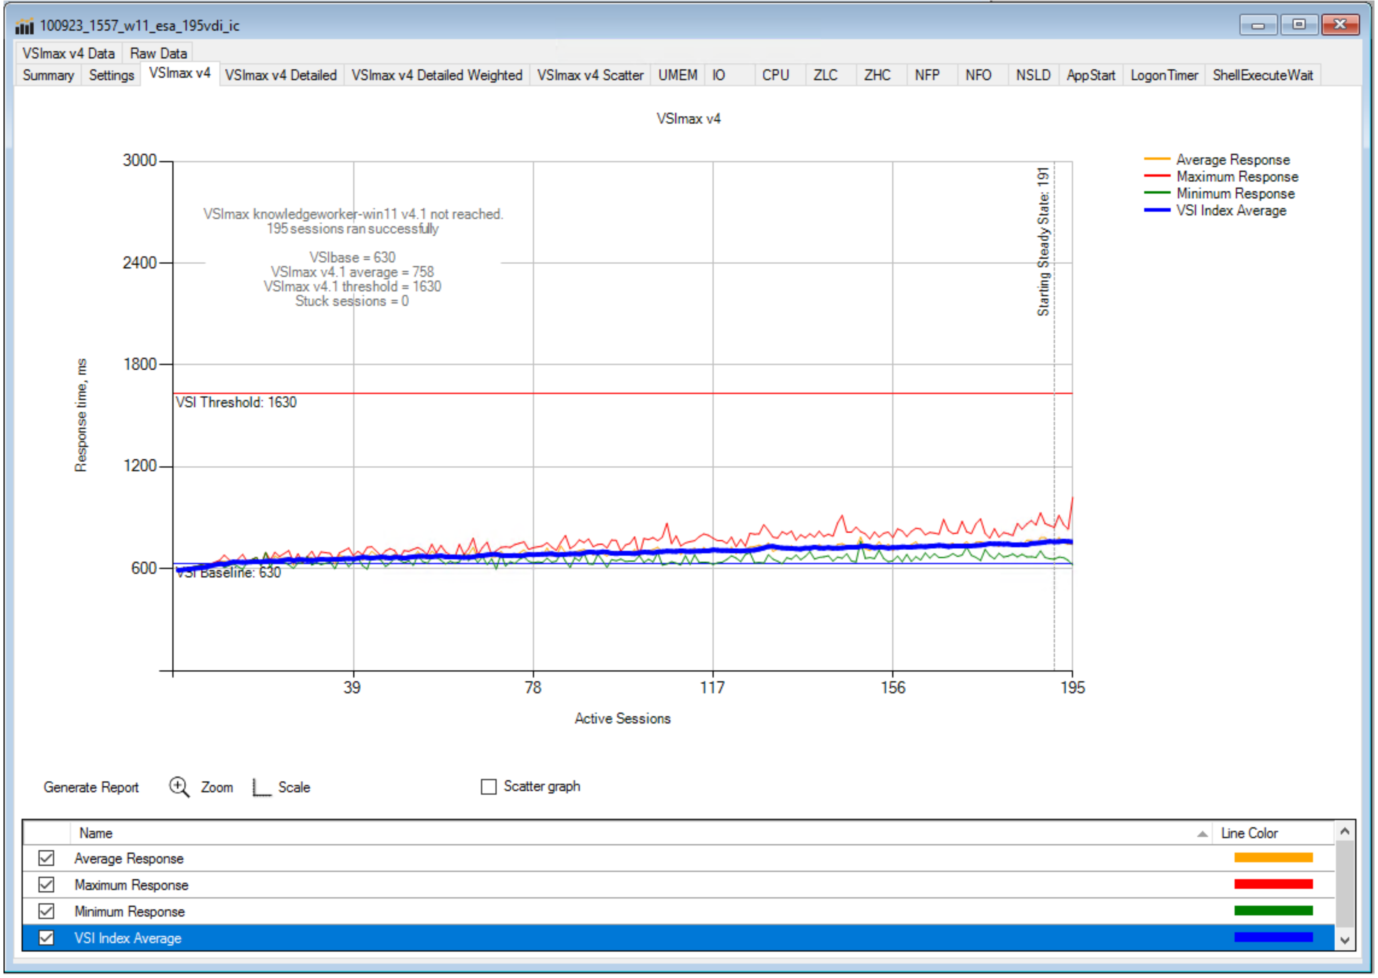

VSImax knowledgeworker-win11 v4.1 not reached.
195 sessions ran successfully

VSIbase = 630
VSImax v4.1 average = 758
VSImax v4.1 threshold = 1630
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Performance data for the server running the workload is shown below:

Figure 54. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host CPU Utilization

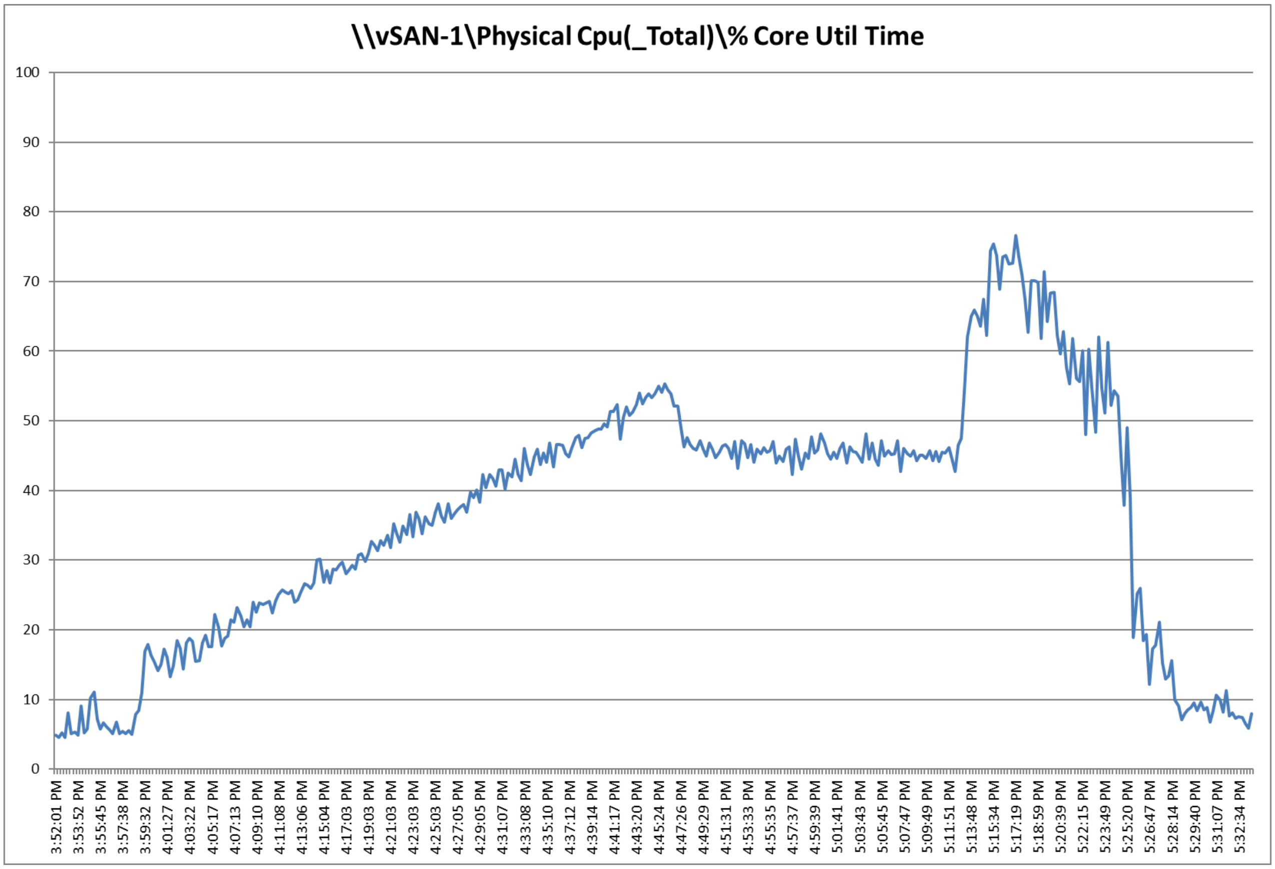

Figure 55. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Memory Utilization

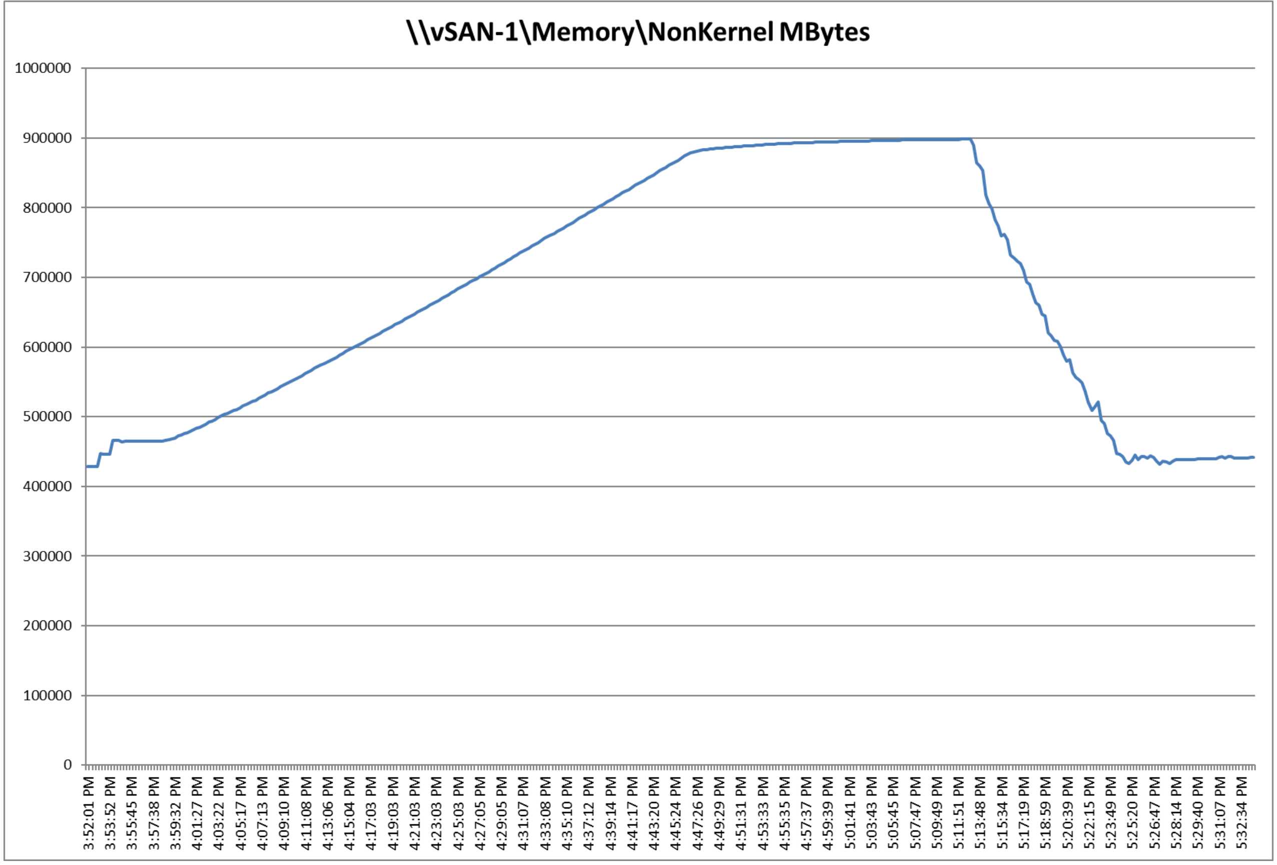

Figure 56. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Network Utilization

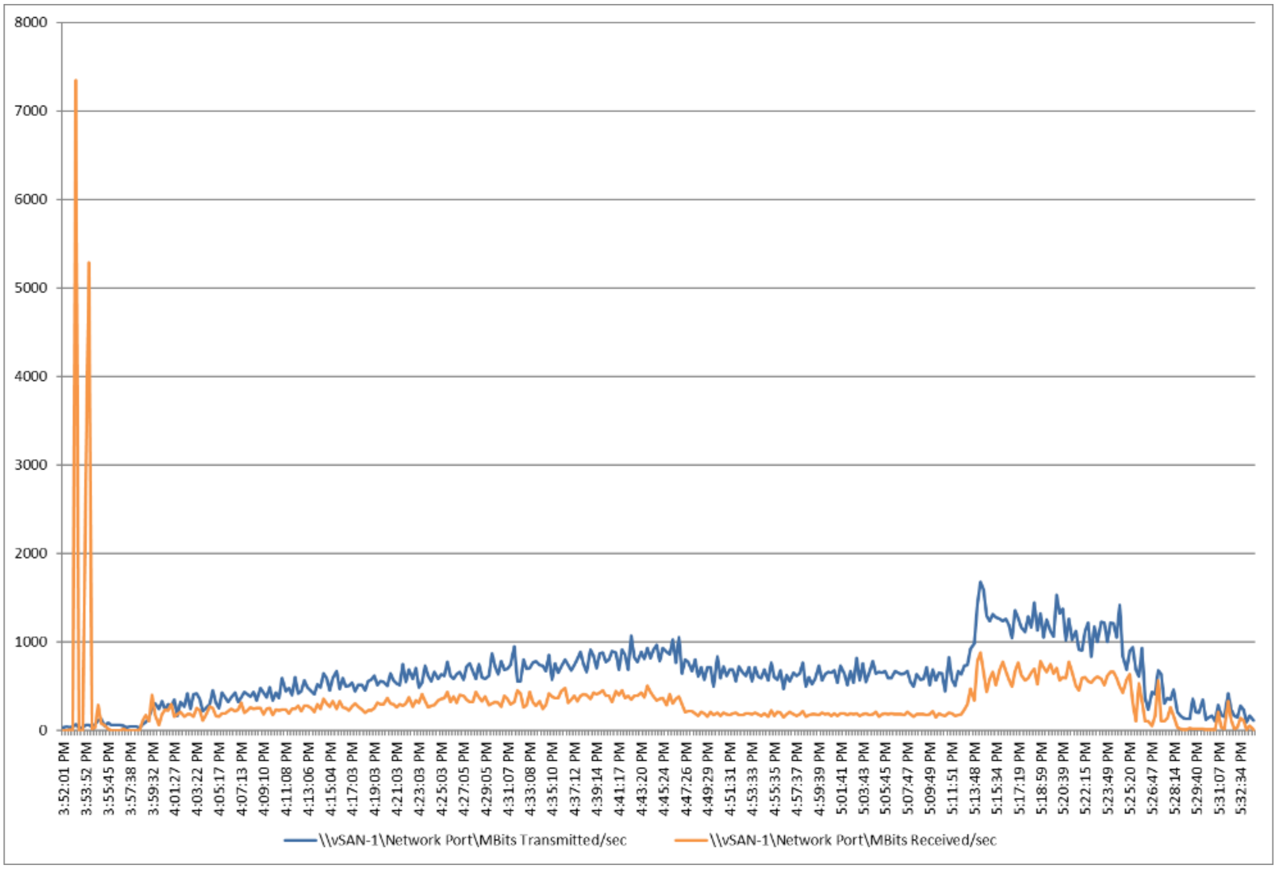

Login Enterprise performance data is shown below:

Figure 57. Single Server | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | EUX Score

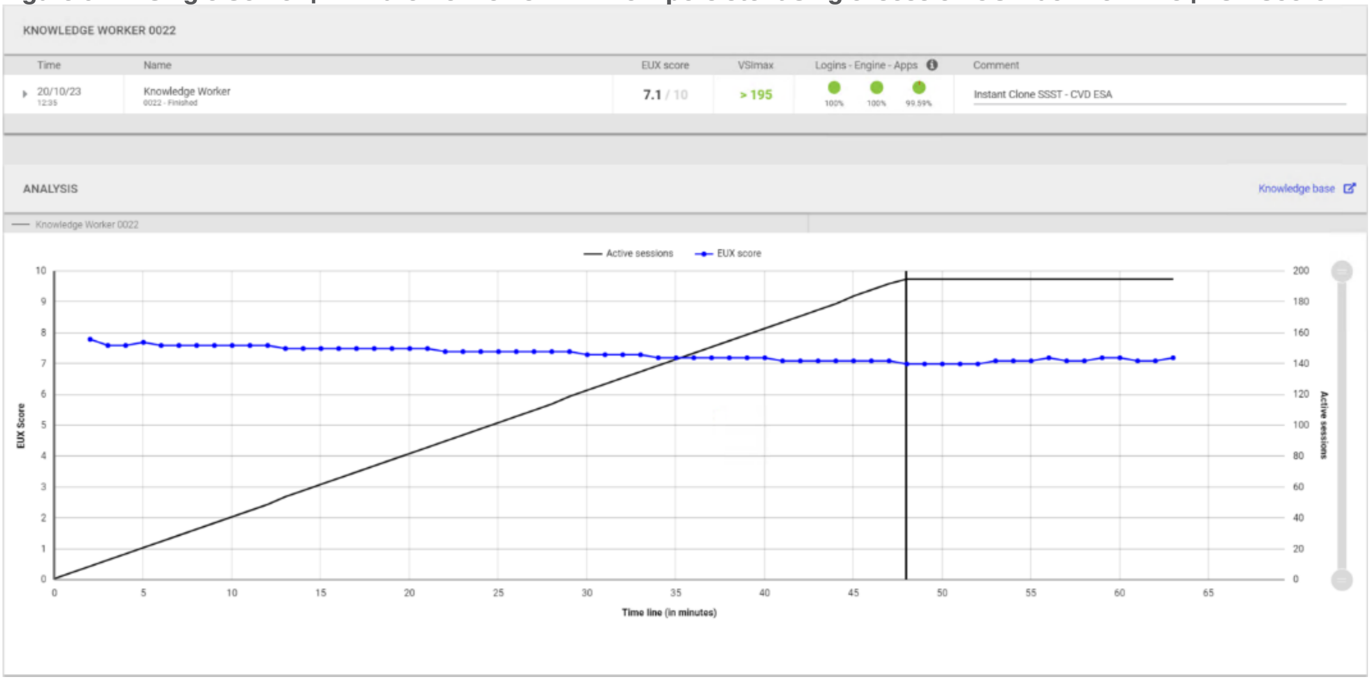

KNOWLEDGE WORKER 0022

Time	Name	EUX score	VSImax	Logins - Engine - Apps	Comment
20/10/23 12:35	Knowledge Worker 0022 - Finished	7.1 / 10	> 195	● 100% ● 100% ● 99.59%	Instant Clone SSST - CVD ESA

ANALYSIS

Knowledge base

Knowledge Worker 0022

ESA vSAN: Single-Server Recommended Maximum Workload for Persistent Single-session OS Dedicated Sessions with 195 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, 1TB 4800 MHz RAM is 195 Windows 11 64-bit VDI Persistent virtual machines with 2 vCPU and 4GB RAM.

Login VSI performance data is as shown below:

Figure 58. **Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Score**

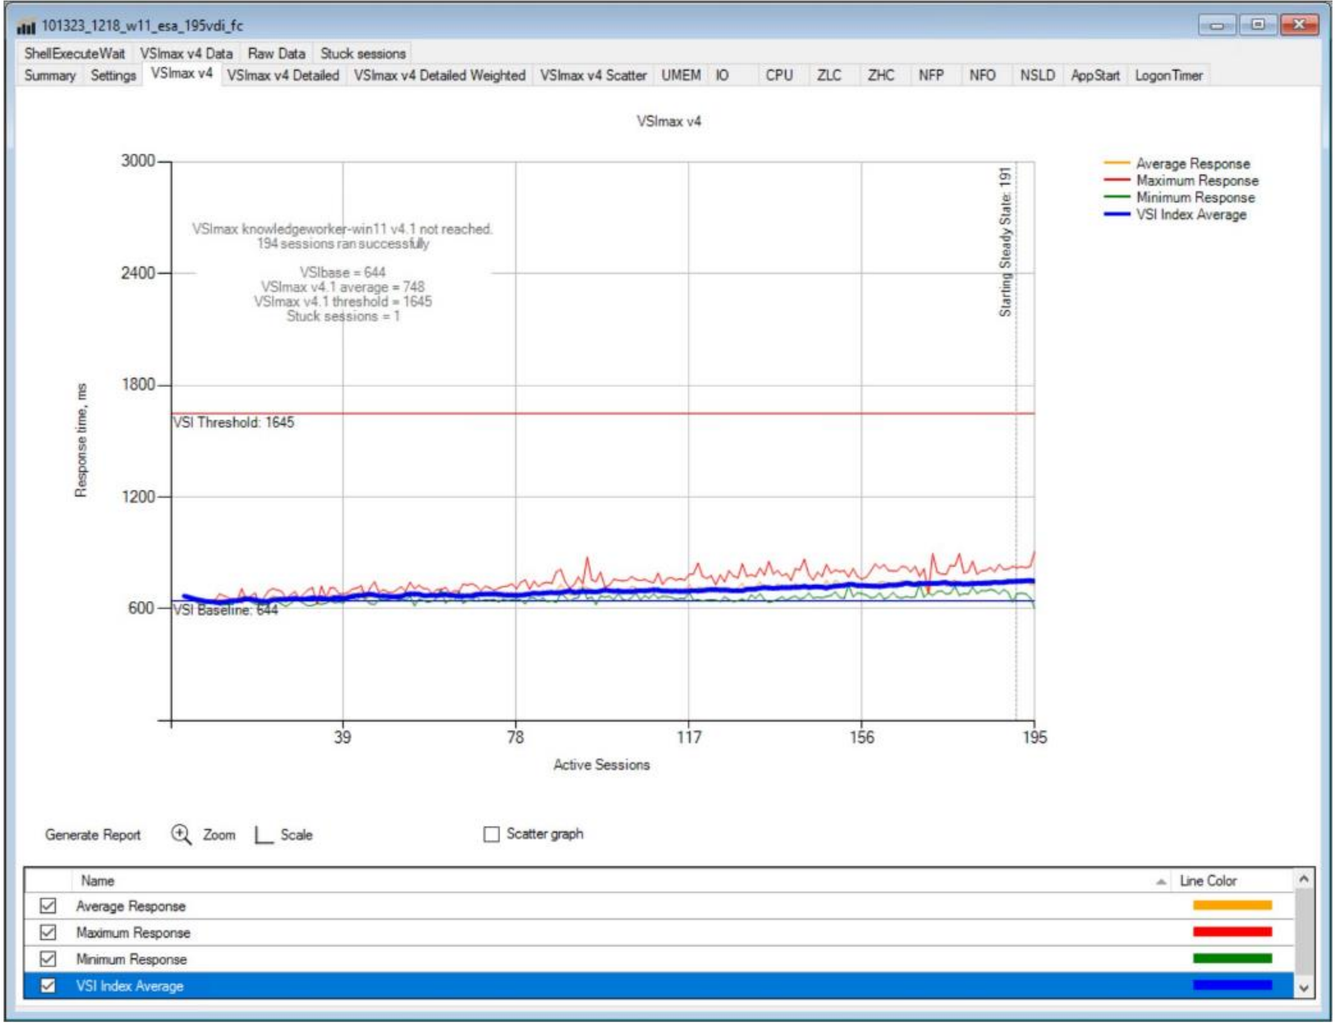

VSImax knowledgeworker-win11 v4.1 not reached.
194 sessions ran successfully

VSIbase = 644
VSImax v4.1 average = 748
VSImax v4.1 threshold = 1645
Stuck sessions = 1

☐ Scatter graph

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Performance data for the server running the workload is shown below:

Figure 59. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host CPU Utilization

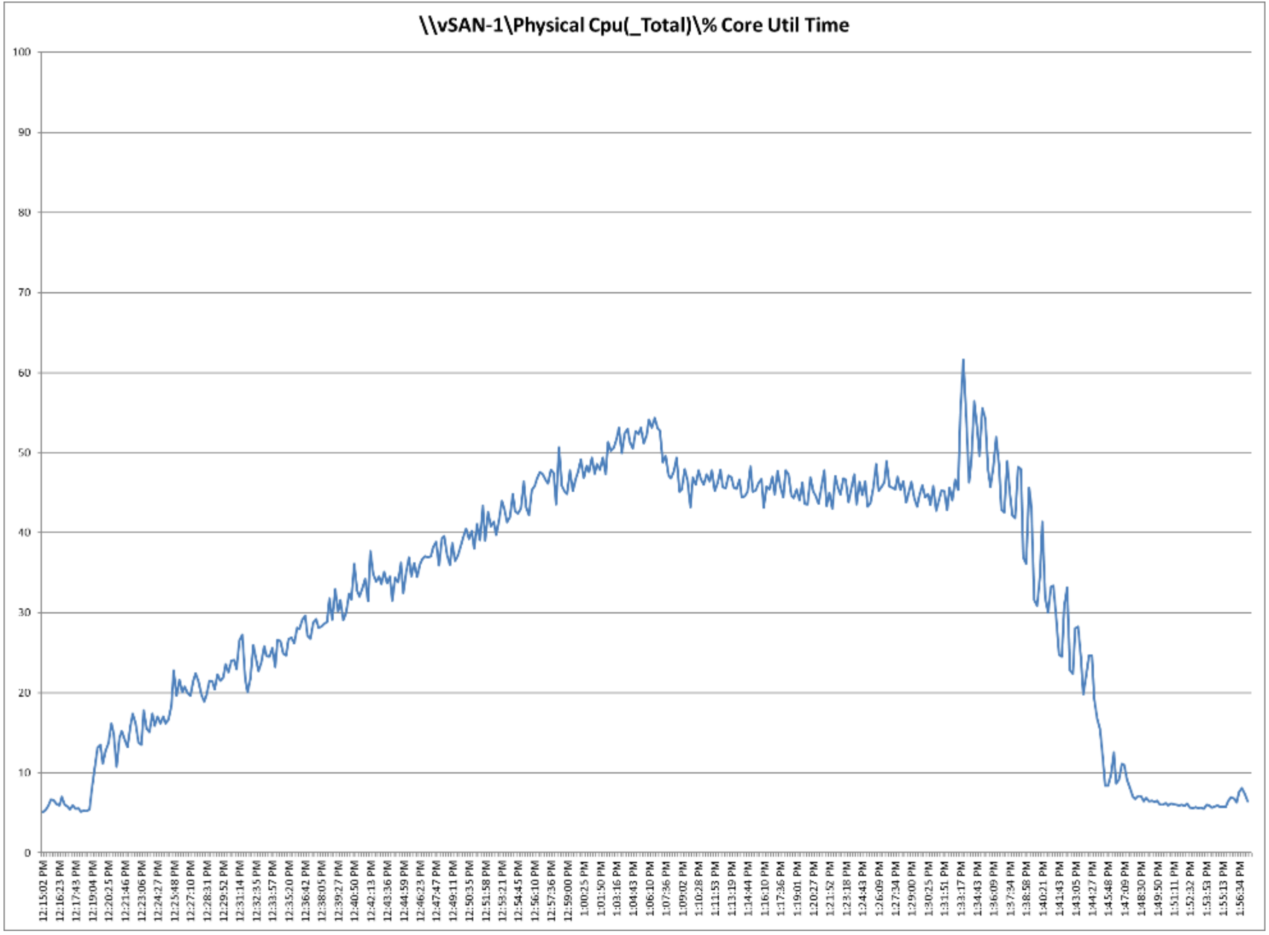

Figure 60. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Memory Utilization**

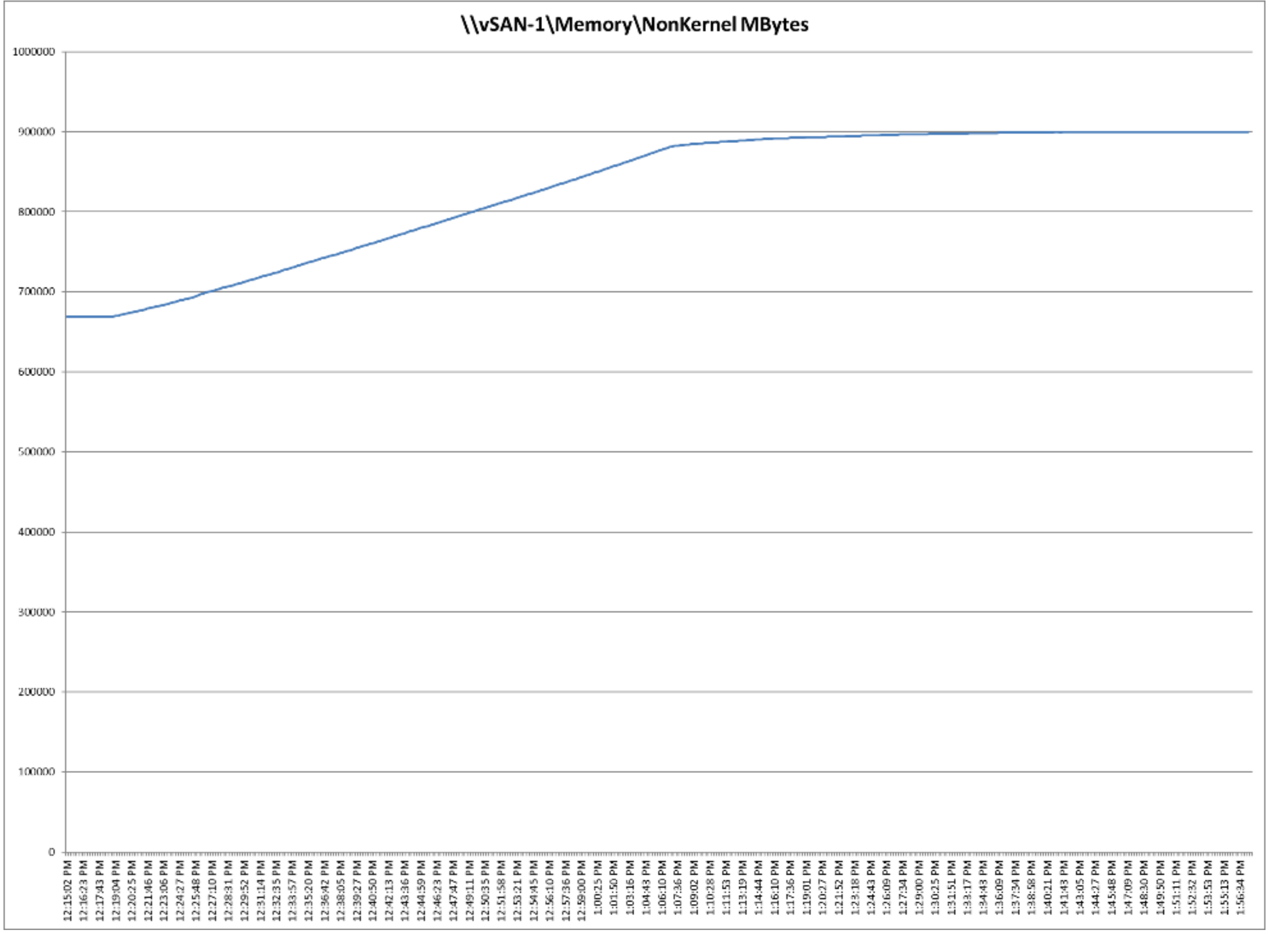

Figure 61. **Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Network Utilization**

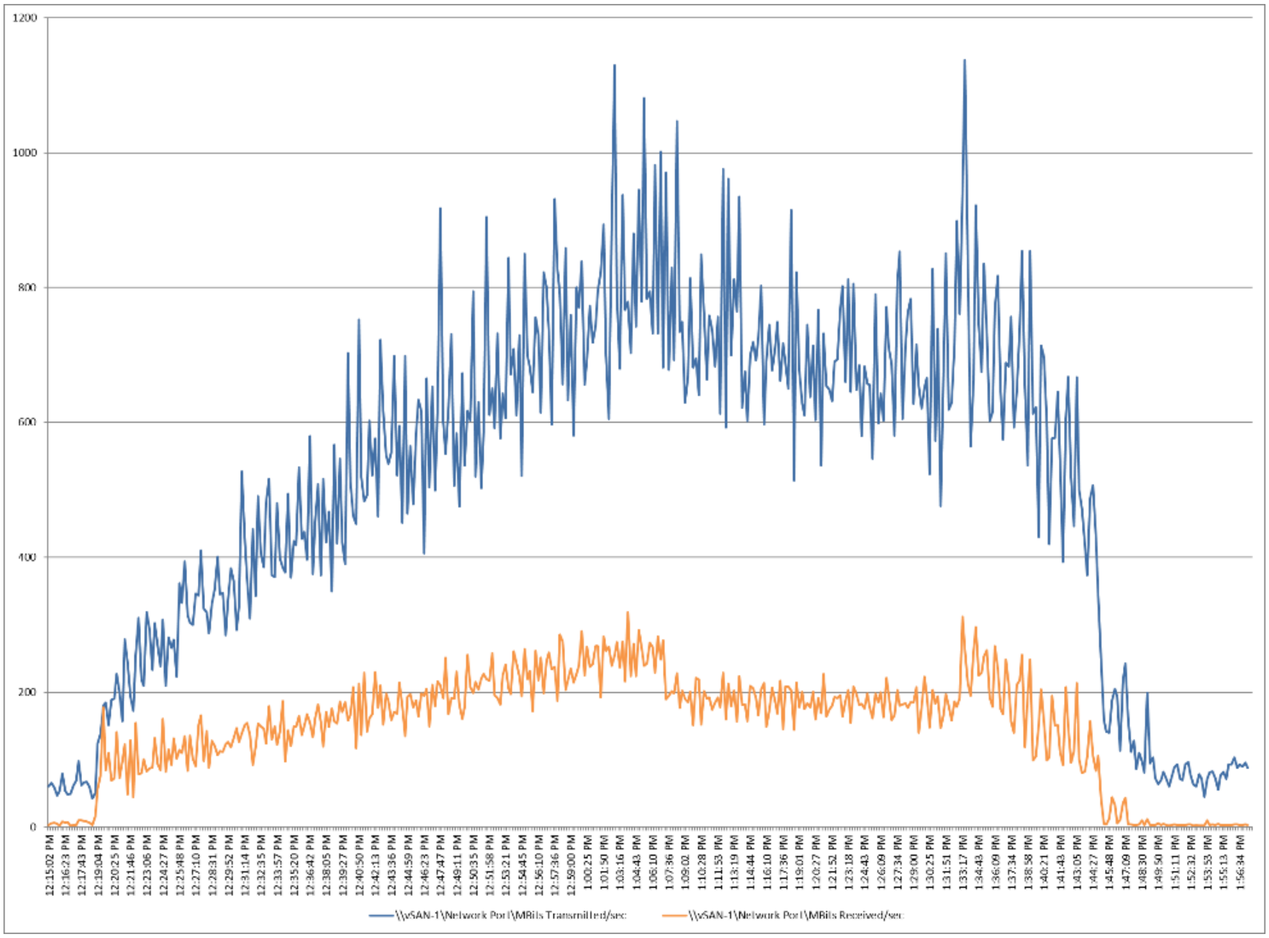

Login Enterprise performance data is shown below:

Figure 62. Single Server | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | EUX Score

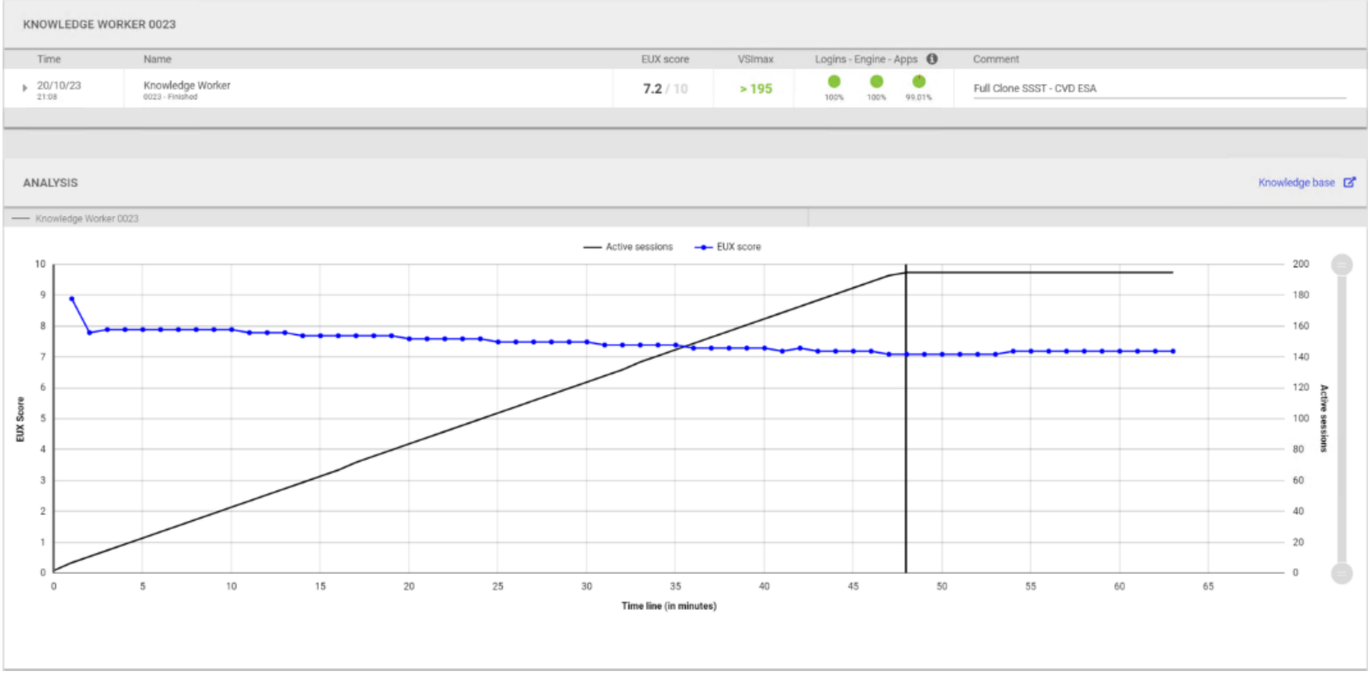

KNOWLEDGE WORKER 0023

Time	Name	EUX score	VSImax	Logins - Engine - Apps	Comment
20/10/23 21:08	Knowledge Worker 0023 - Finished	7.2 / 10	> 195	● 100% ● 100% ● 99.01%	Full Clone SSST - CVD ESA

ANALYSIS

ESA vSAN: 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, 1TB 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 63. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Score

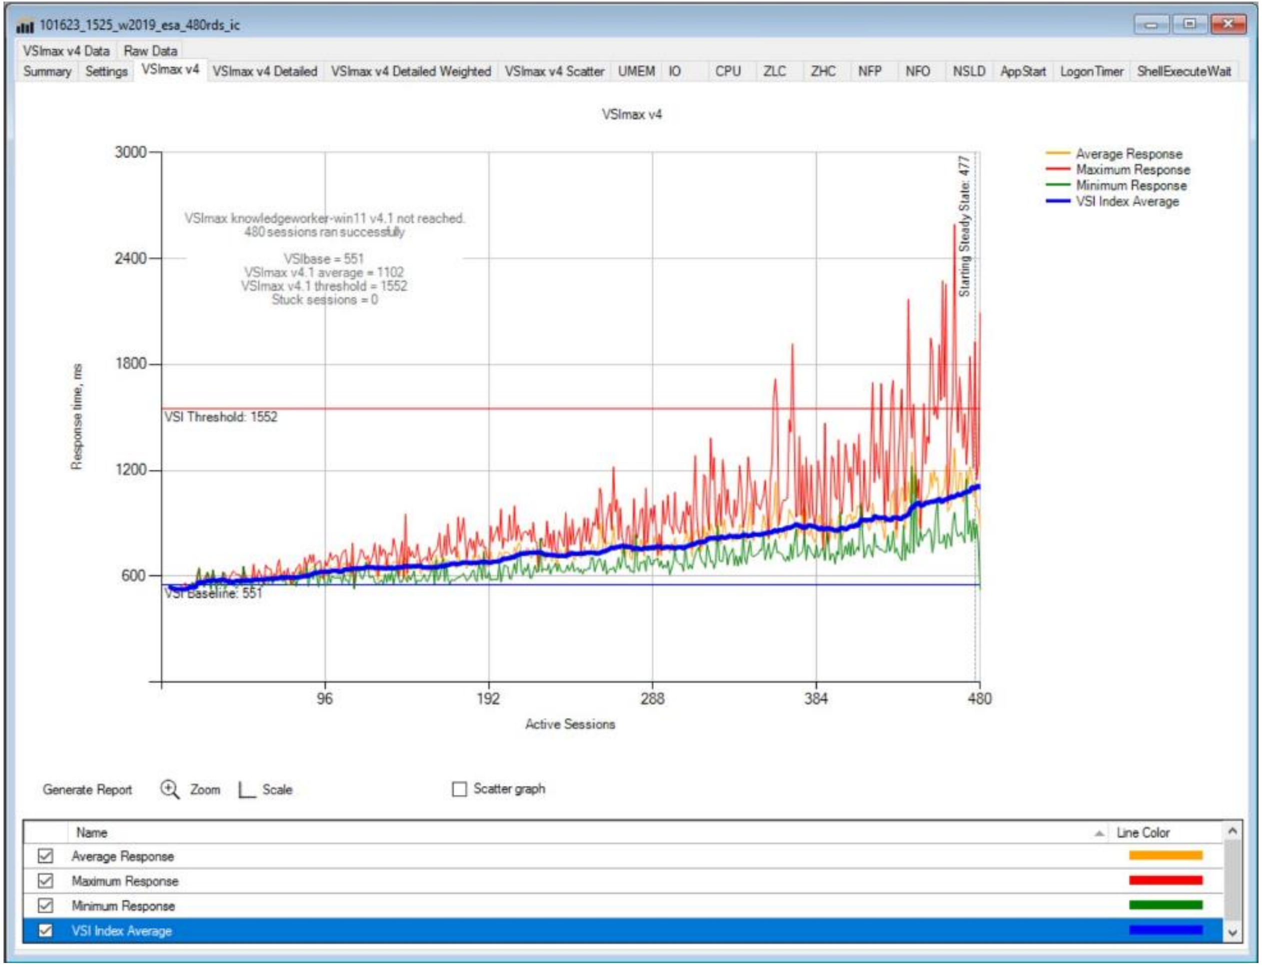

VSImax knowledgeworker-win11 v4.1 not reached.
480 sessions ran successfully

VSIbase = 551
VSImax v4.1 average = 1102
VSImax v4.1 threshold = 1552
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Performance data for the server running the workload is shown below:

Figure 64. **Single Server Recommended Maximum Workload VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host CPU Utilization**

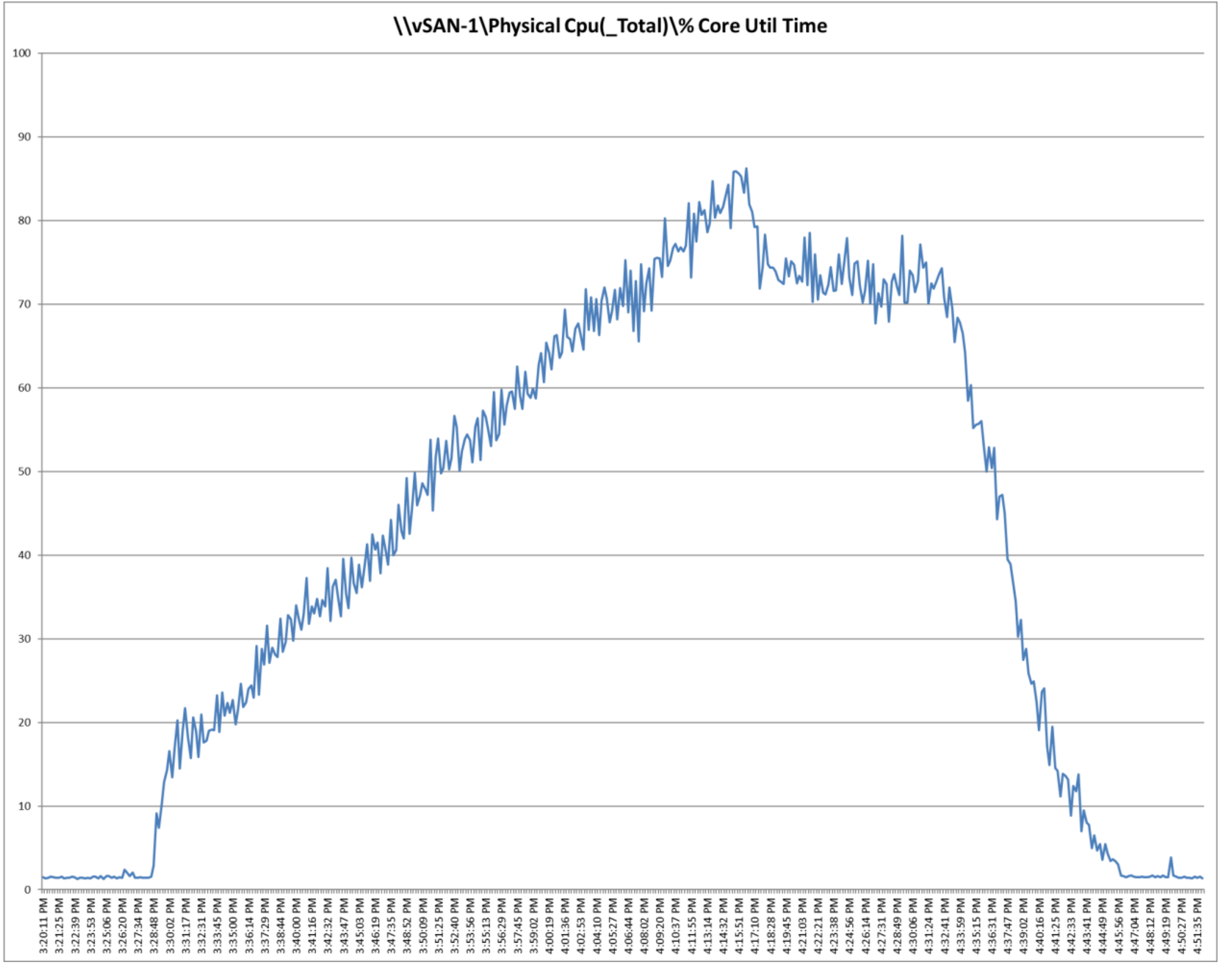

Figure 65. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Memory Utilization

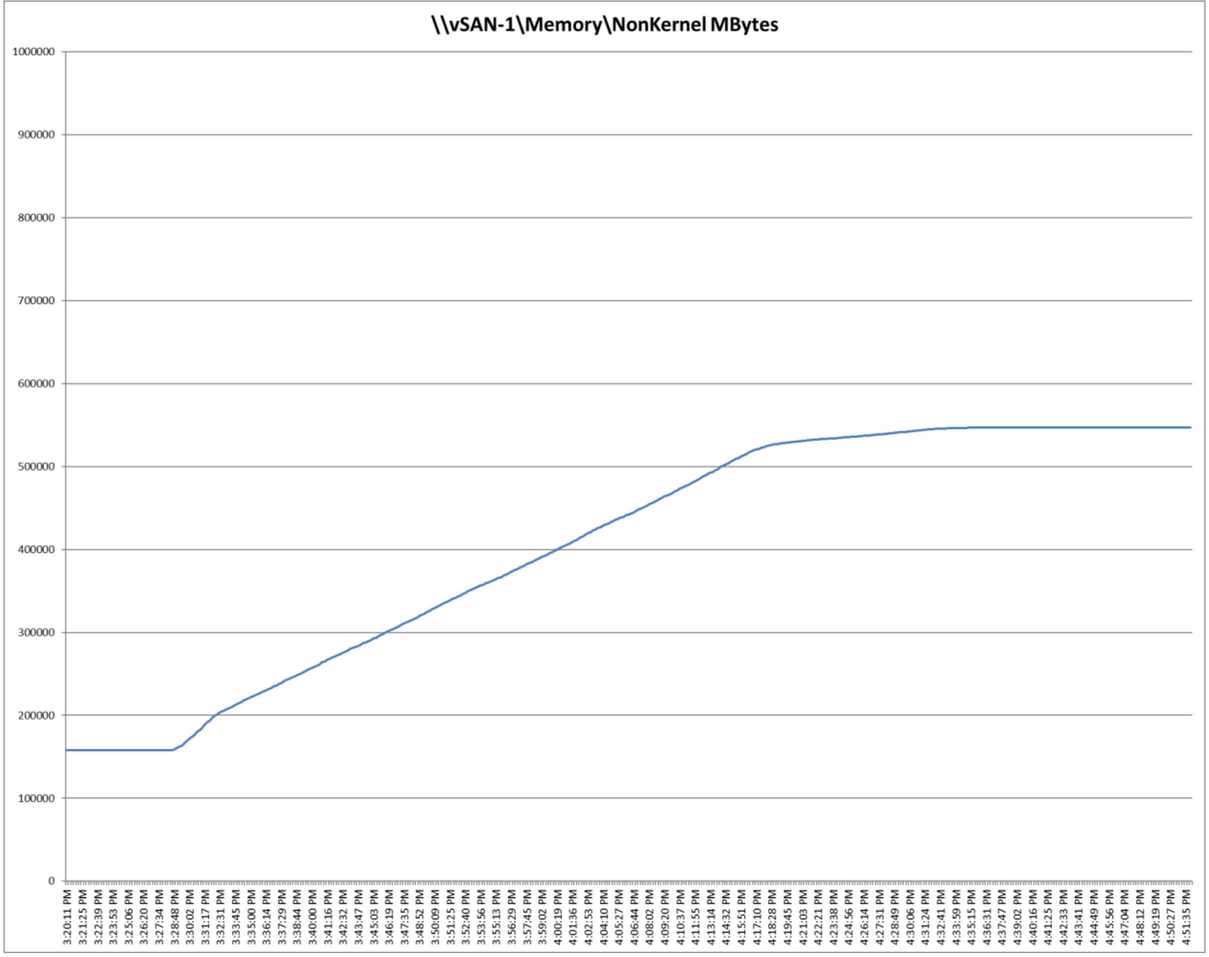

Figure 66. Single Server | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Network Utilization

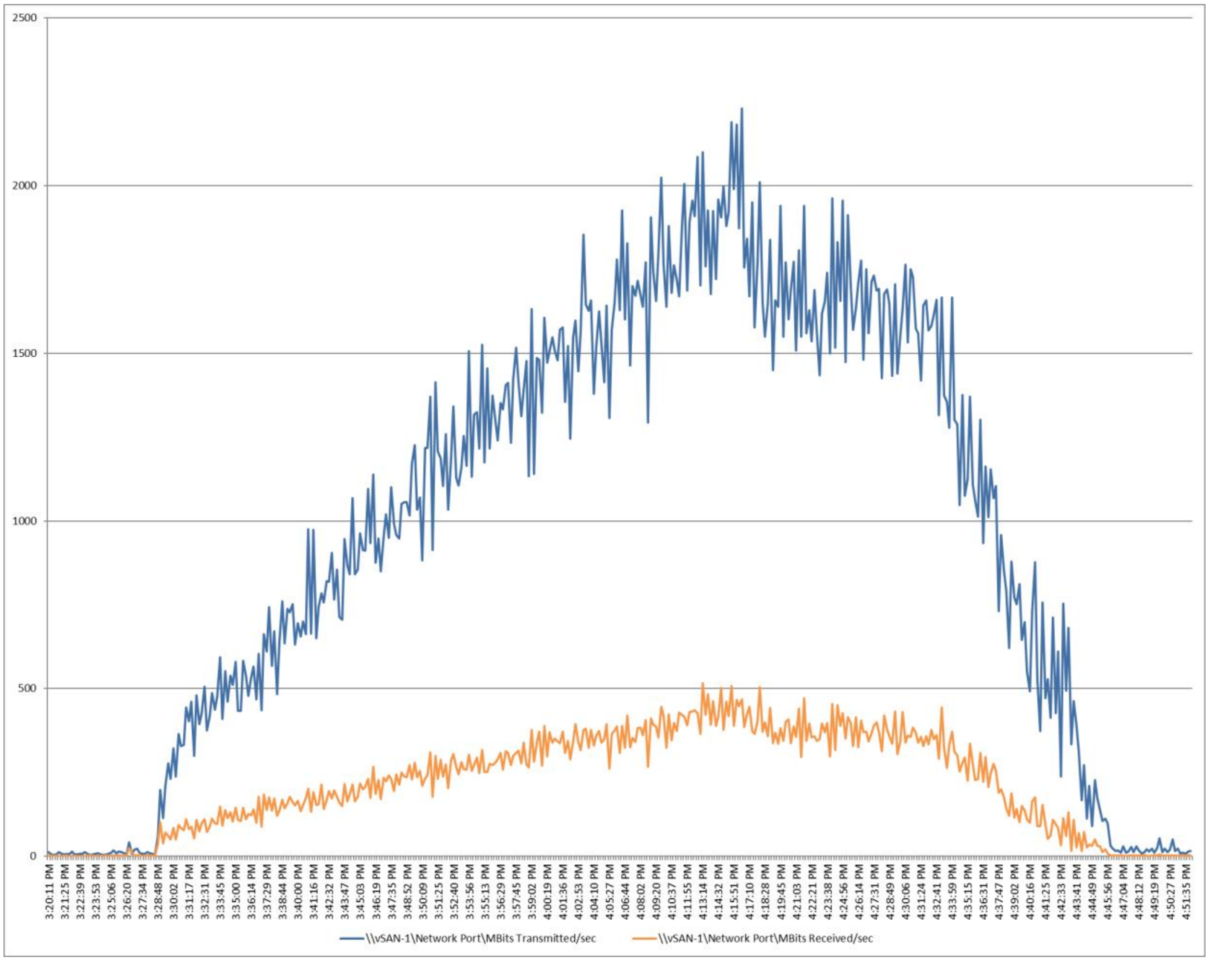

Performance data for the RDS Virtual Machine running the workload is shown below:

Figure 67. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Virtual Machine CPU Utilization**

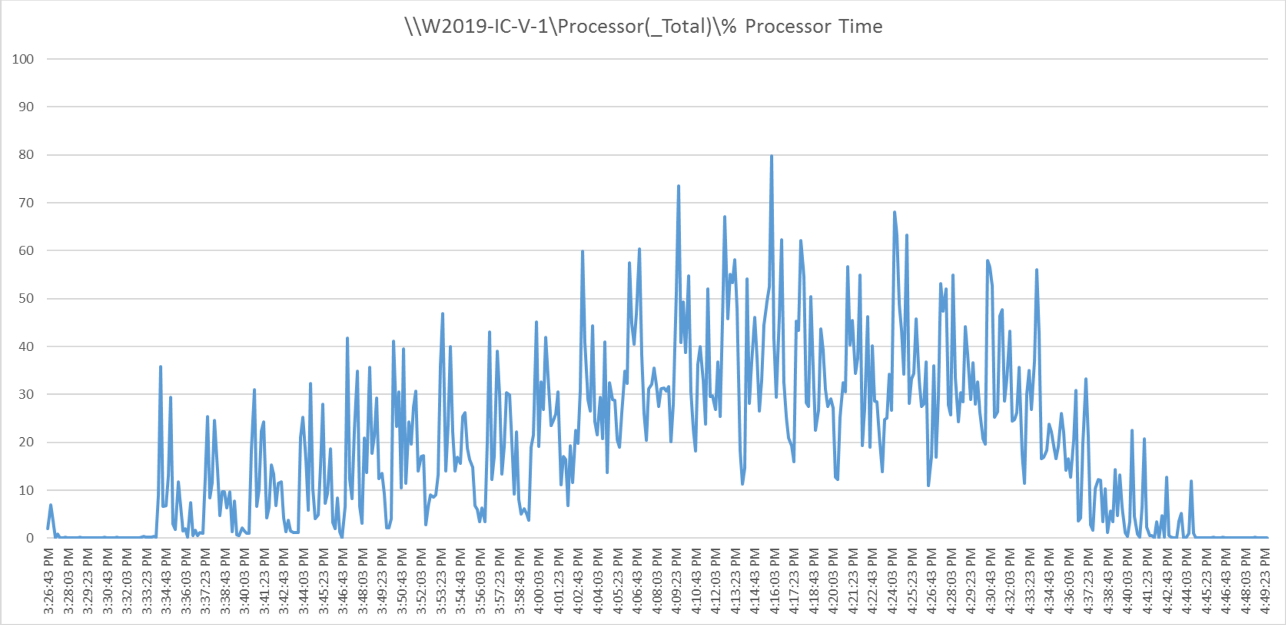

Figure 68. **Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Virtual Machine Memory Utilization**

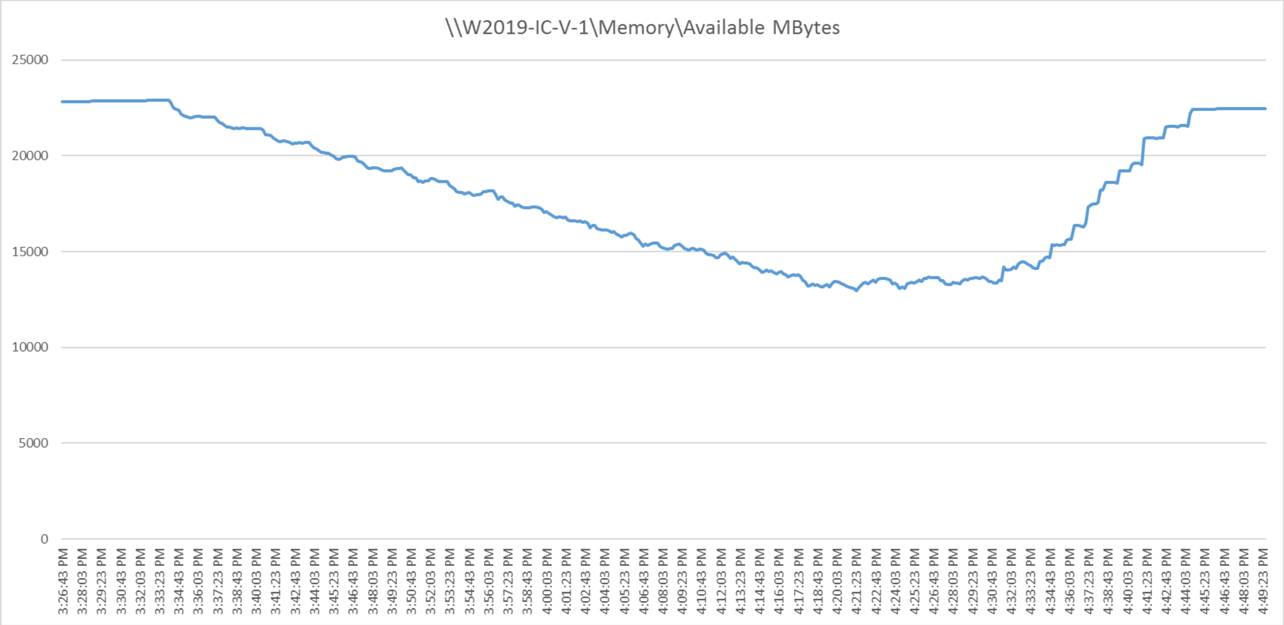

Figure 69. Single Server Recommended Maximum Workload | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Network Utilization

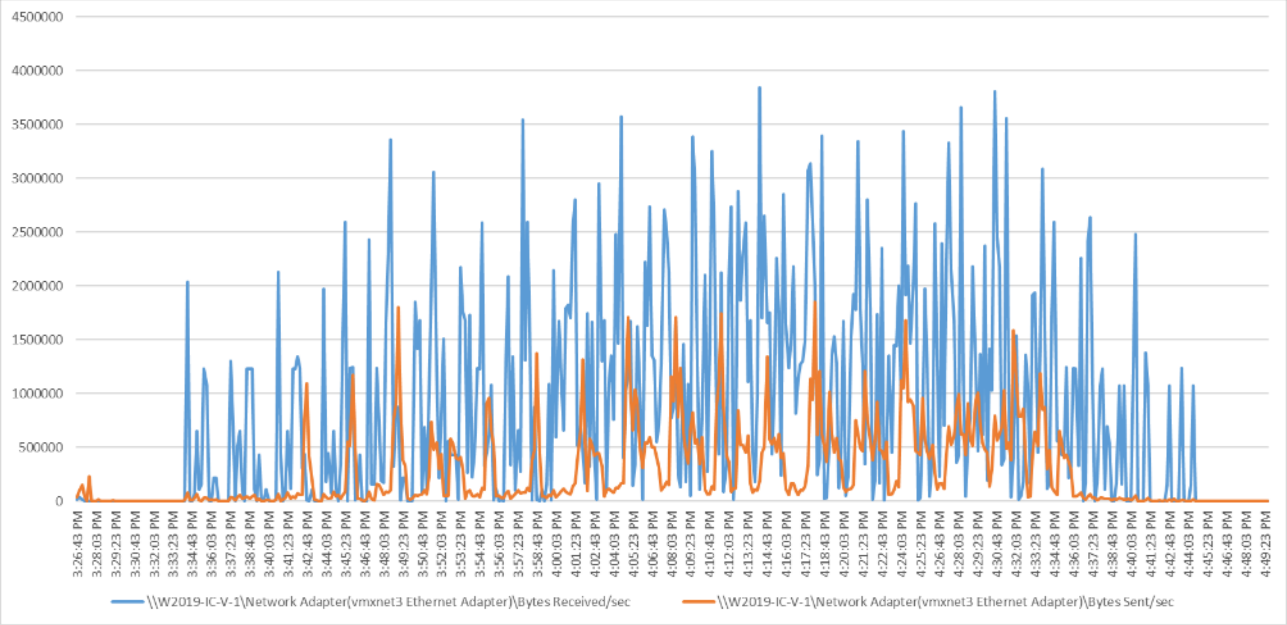

Login Enterprise performance data is shown below:

Figure 70. Single Server | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | EUX Score

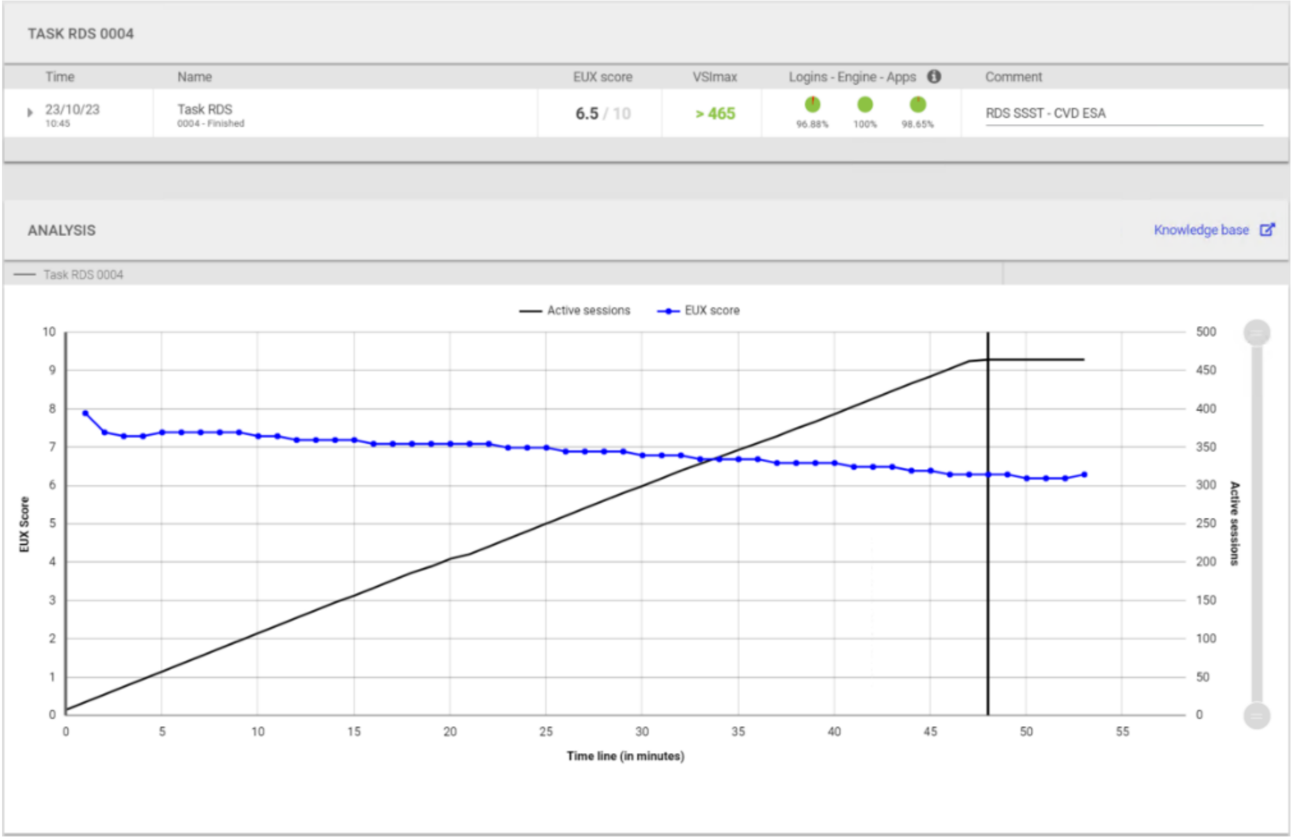

TASK RDS 0004

Time	Name	EUX score	VSImax	Logins - Engine - Apps	Comment
23/10/23 10:45	Task RDS 0004 - Finished	6.5 / 10	> 465	96.88% 100% 98.65%	RDS SSST - CVD ESA

ANALYSIS

Knowledge base

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 ESXi Host vSAN cluster, and designed to support single Host failure (N+1 Fault tolerance).

Four-server vSAN OSA cluster

- 585 Non-persistent Single-session OS sessions
- 585 Persistent Single-session OS sessions
- 1440 Non-persistent Multi-session OS sessions

Three-server vSAN ESA cluster

- 390 Non-persistent Single-session OS sessions
- 390 Persistent Single-session OS sessions
- 960 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.

OSA vSAN: Full Scale Recommended Maximum Workload Testing for Non-persistent Single-session OS Machine VDAs with 585 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware vSAN cluster during the full scale testing with 585 Non-persistent Single-session OS machines using four blades in a single pool.

The workload for the test is 585 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 71. Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs| VSI Score

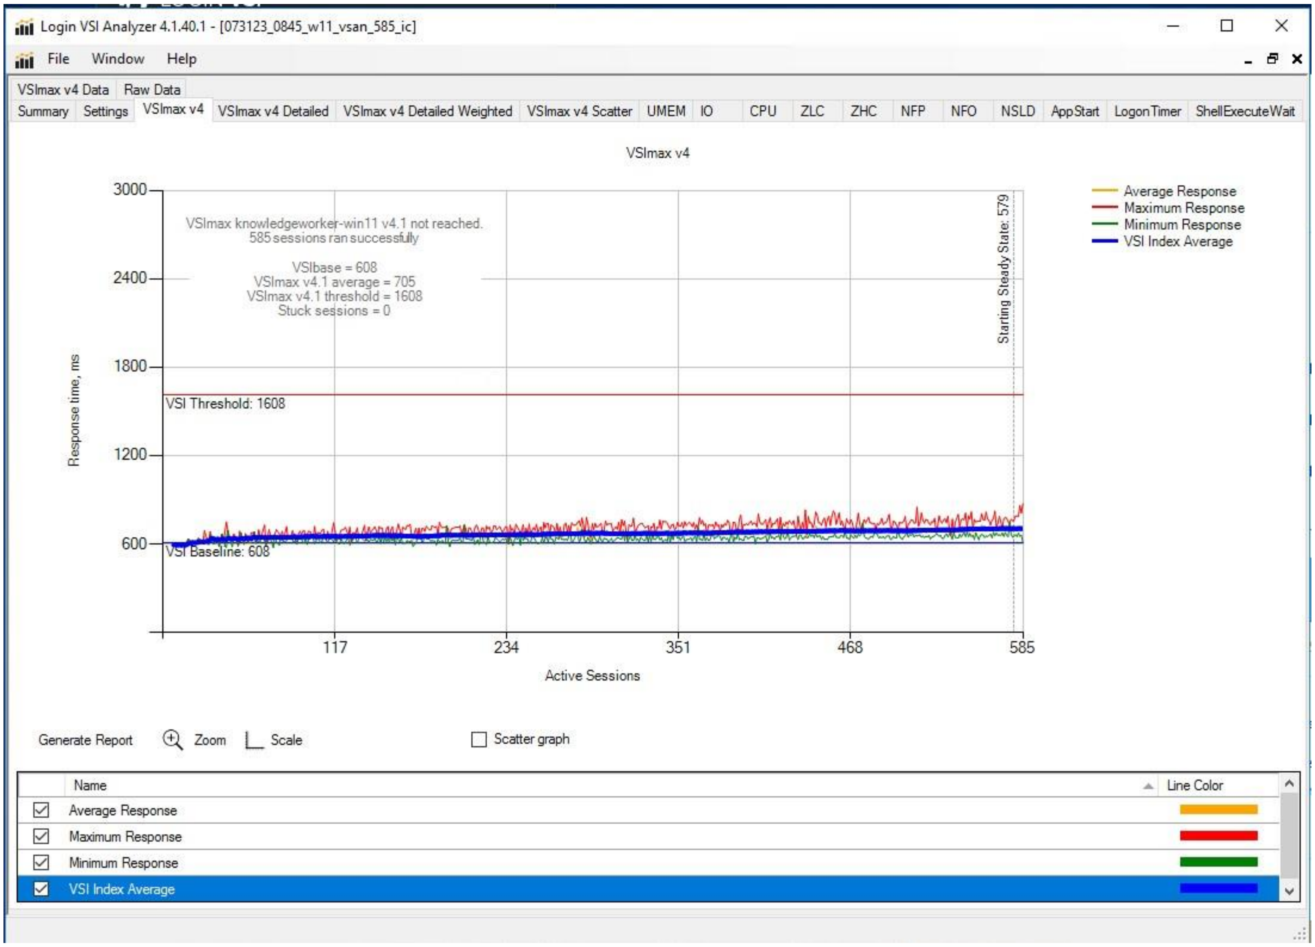

VSImax knowledgeworker-win11 v4.1 not reached.
585 sessions ran successfully

VSIbase = 608
VSImax v4.1 average = 705
VSImax v4.1 threshold = 1608
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 72. **Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host CPU Utilization**

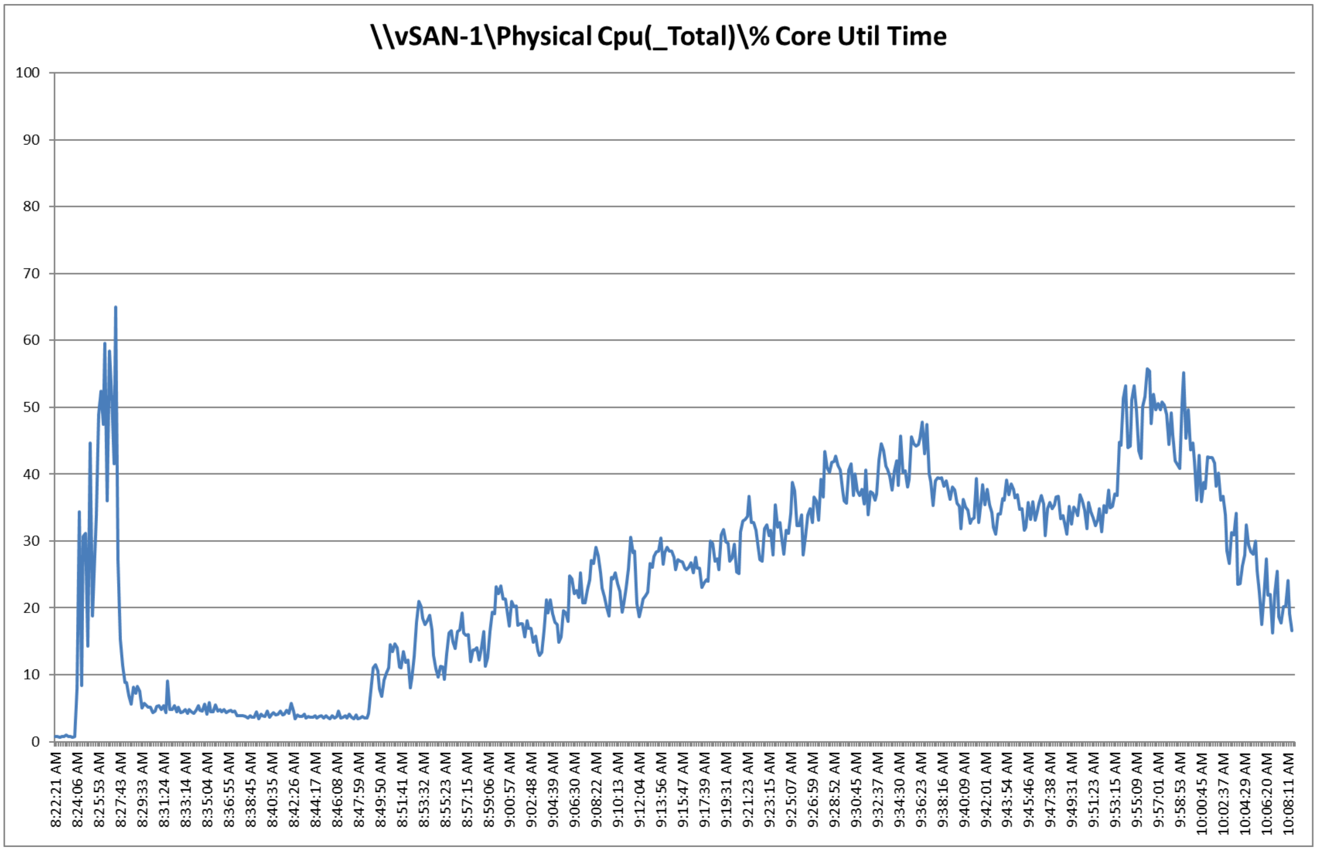

Figure 73. **Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Memory Utilization**

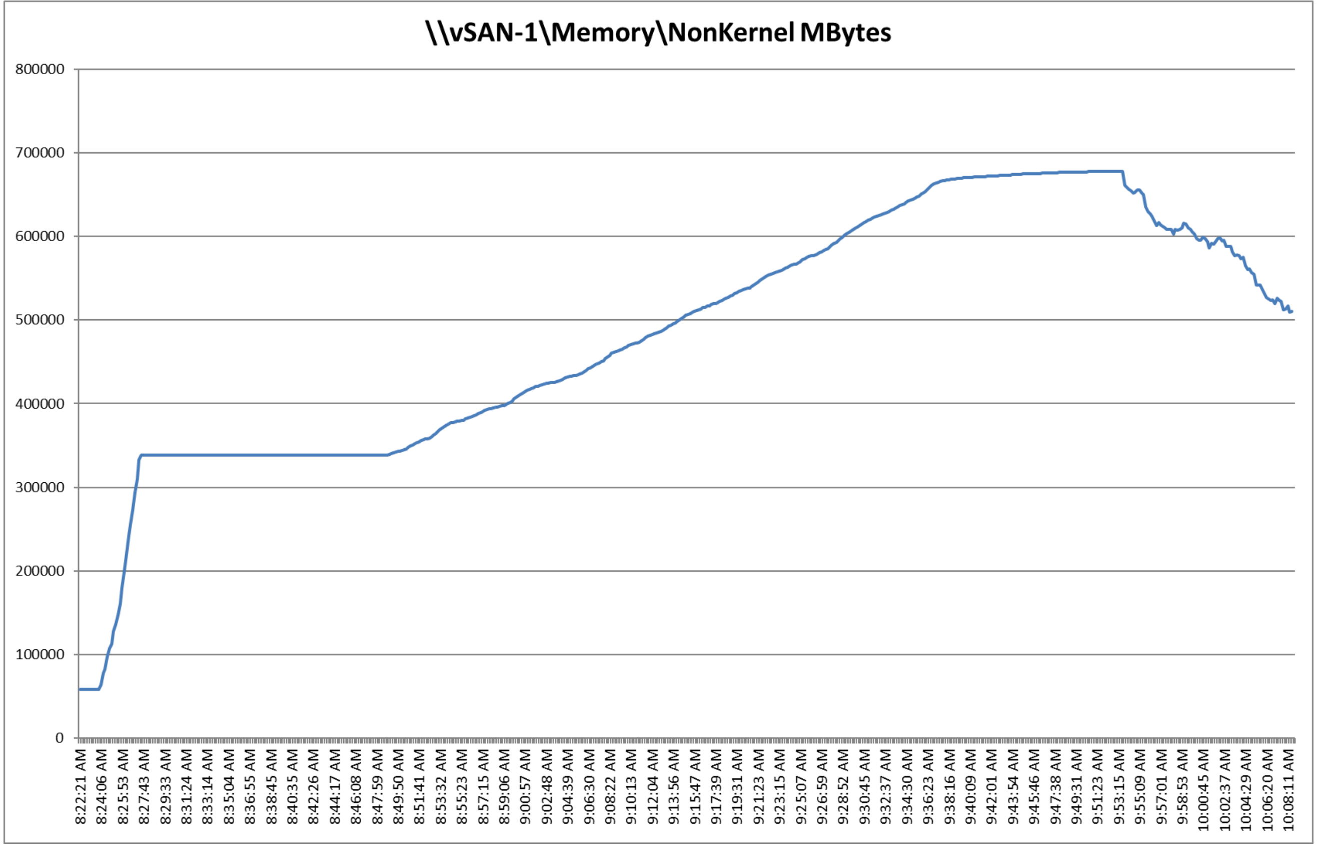

Figure 74. **Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Network Utilization**

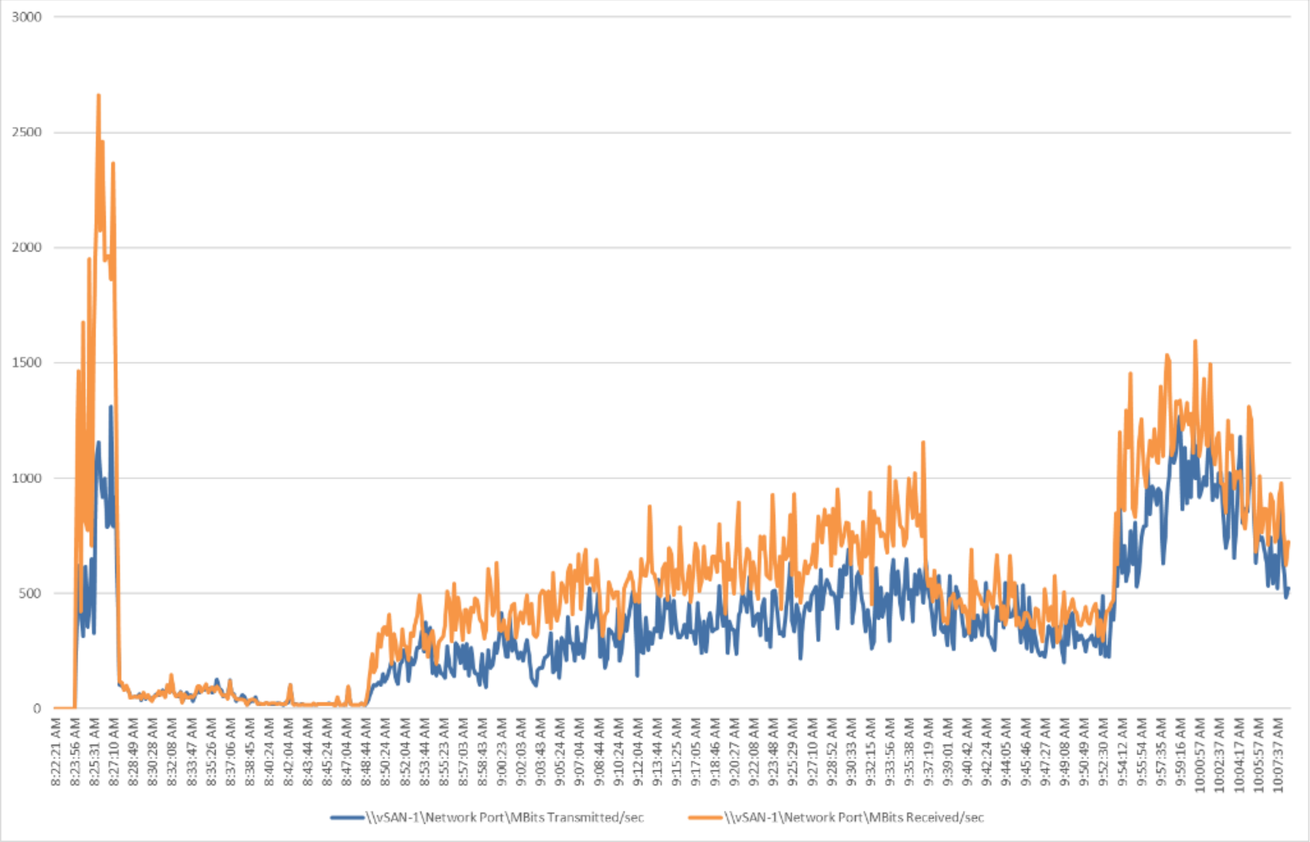

Figure 75. **Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | vSAN Latency Chart**

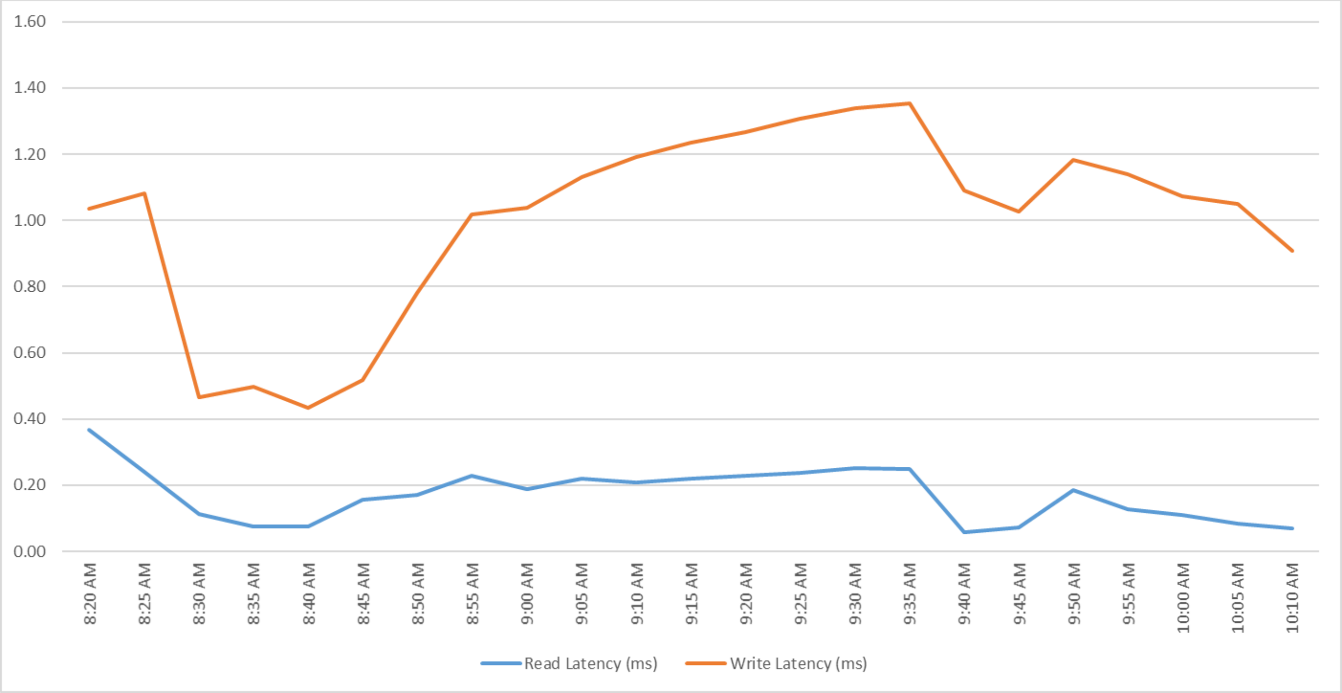

Figure 76. Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | vSAN IOPS Chart

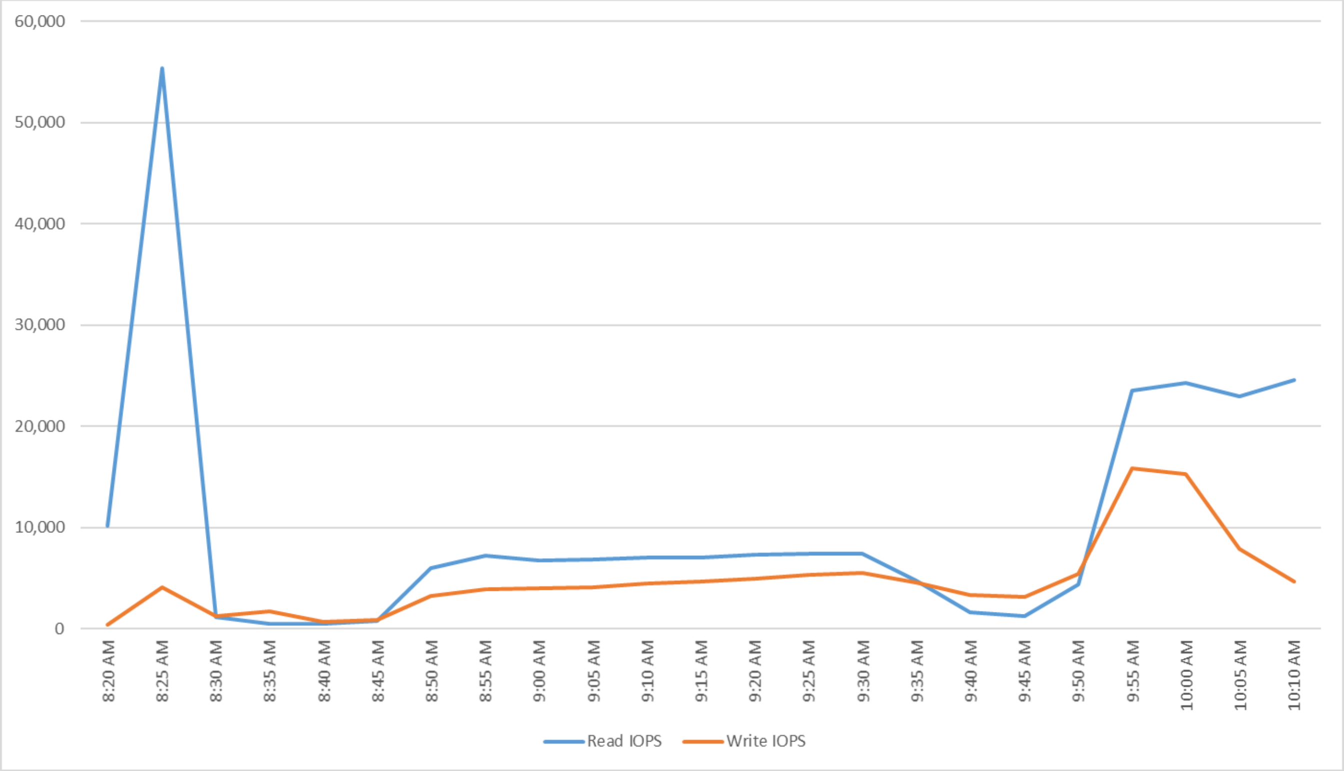

Figure 77. Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | vSAN Throughput Chart

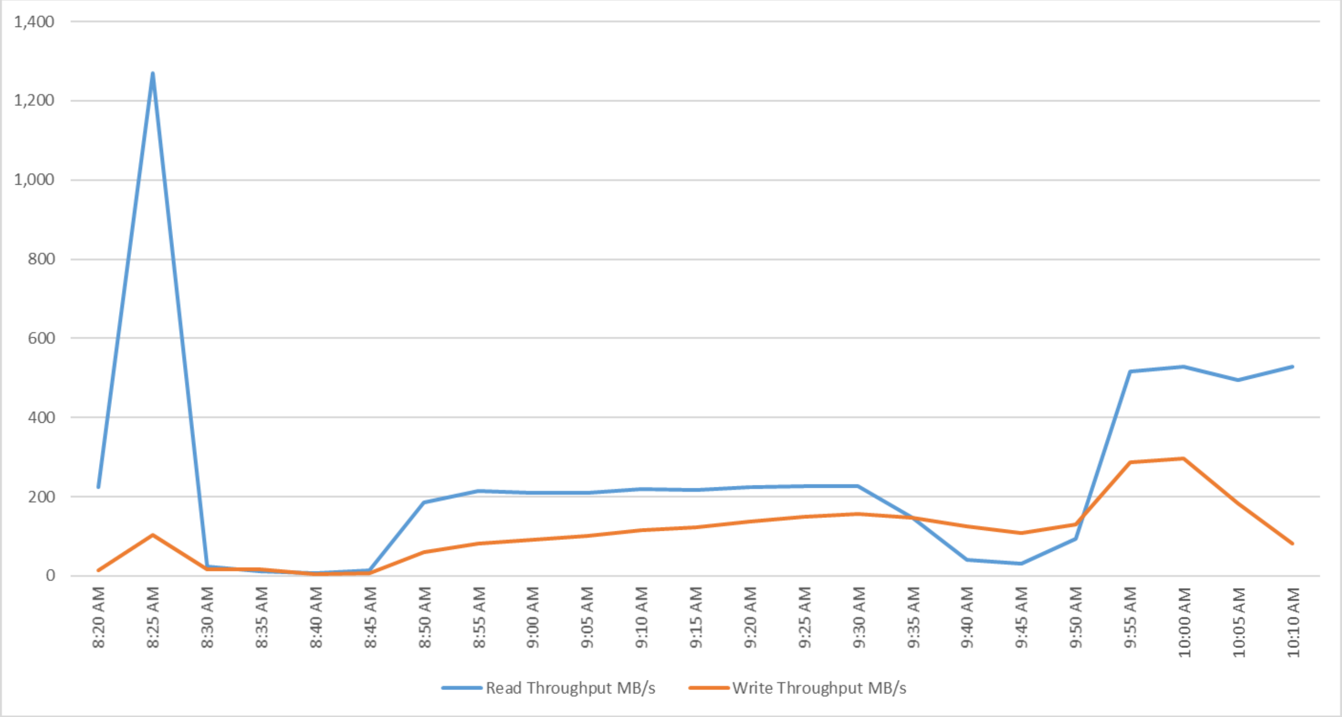

OSA vSAN: Full Scale Recommended Maximum Workload Testing for Persistent Single-session OS Machine VDAs with 585 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware vSAN cluster during the persistent desktop full-scale testing with 585 Persistent Single-session OS machines using 4 blades in a single cluster.

The workload for the test is 585 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 78. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Score

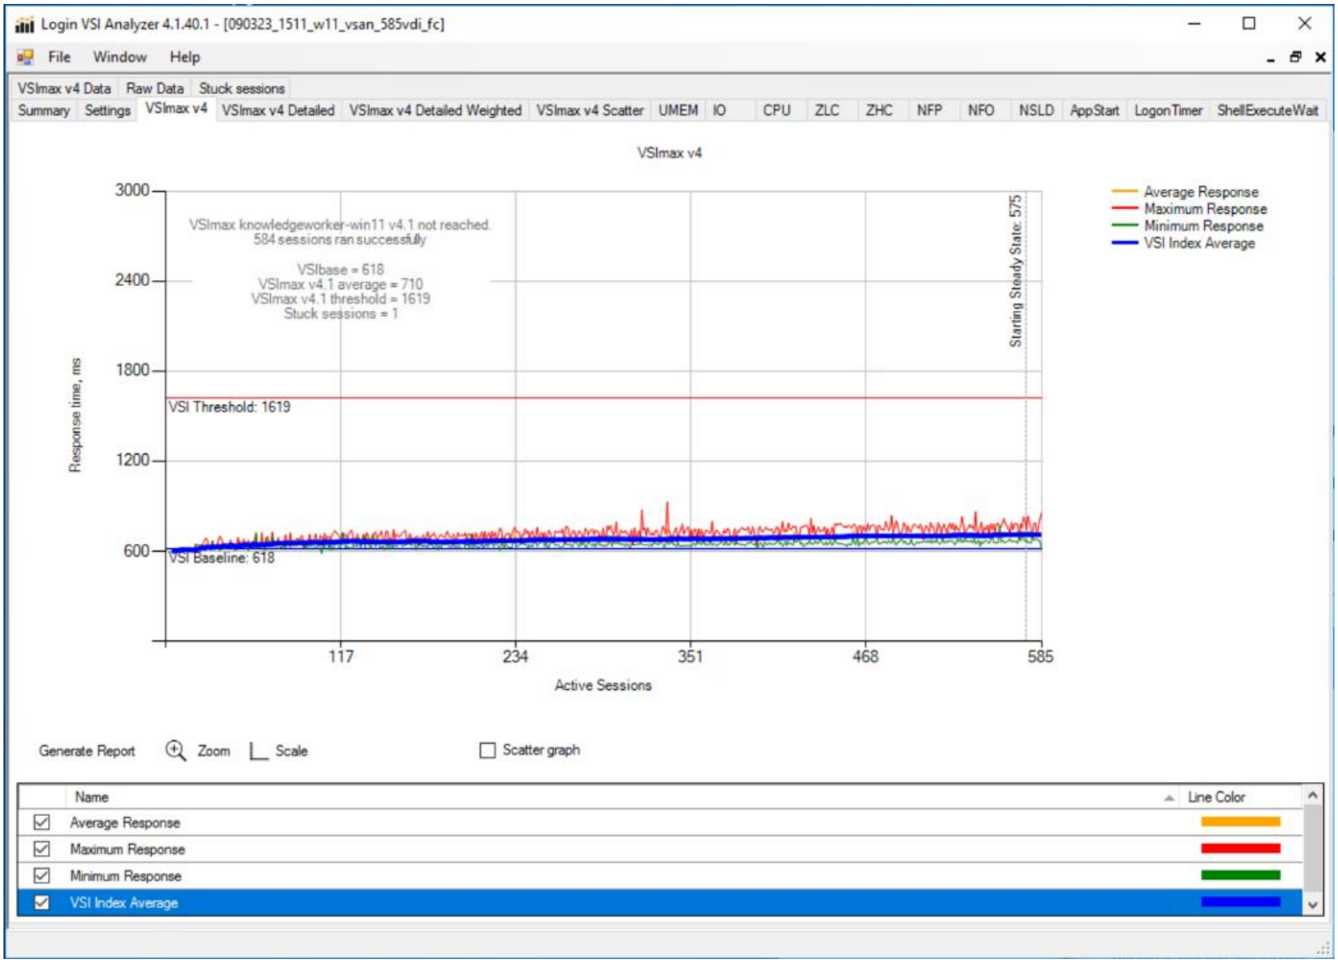

Login VSI Analyzer 4.1.40.1 - [090323_1511_w11_vsan_585vdi_fc]

VSImax knowledgeworker-win11 v4.1 not reached.
584 sessions ran successfully

VSIbase = 618
VSImax v4.1 average = 710
VSImax v4.1 threshold = 1619
Stuck sessions = 1

□ Scatter graph

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 79. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host CPU Utilization

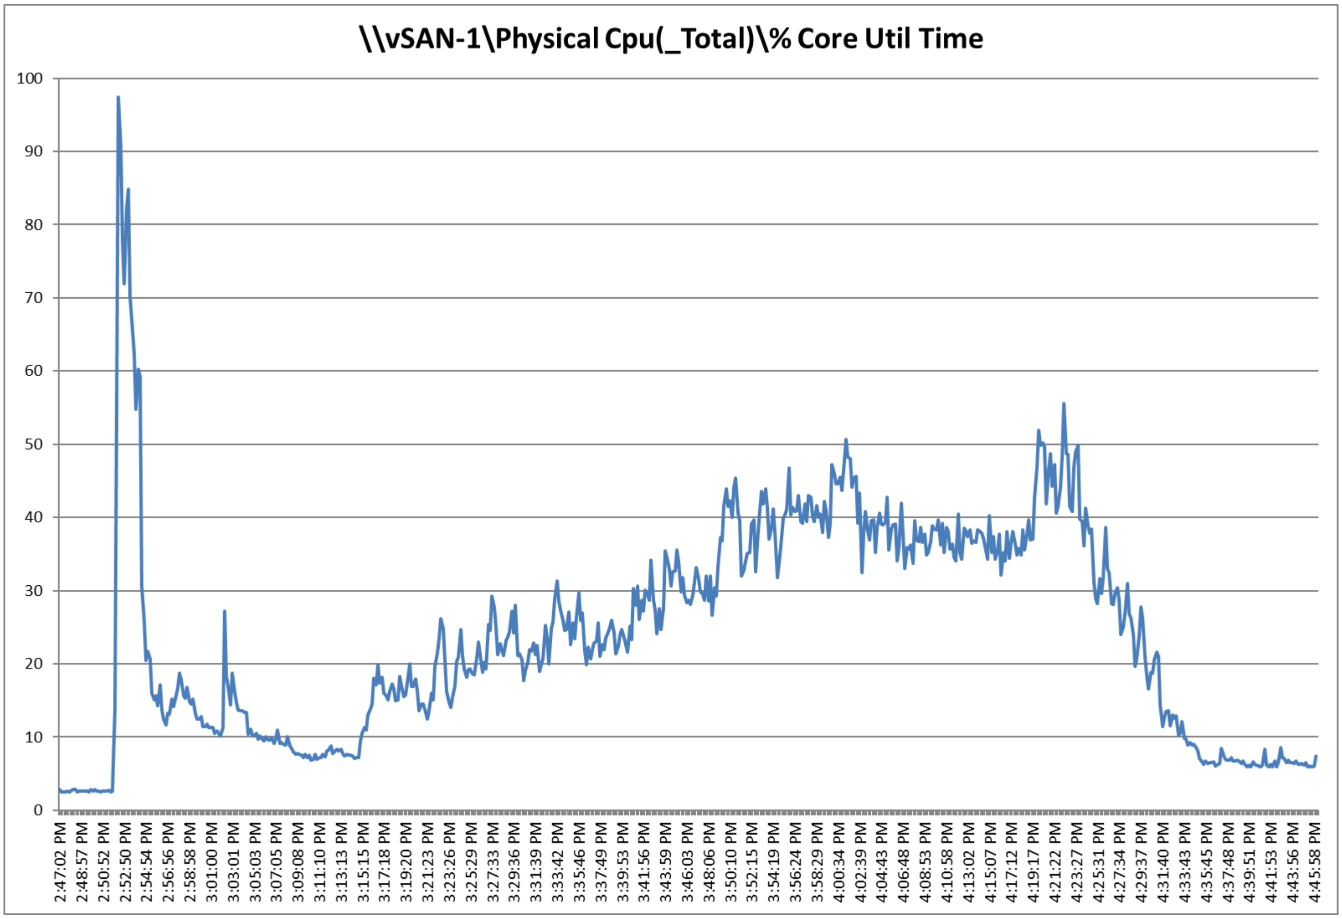

Figure 80. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Memory Utilization

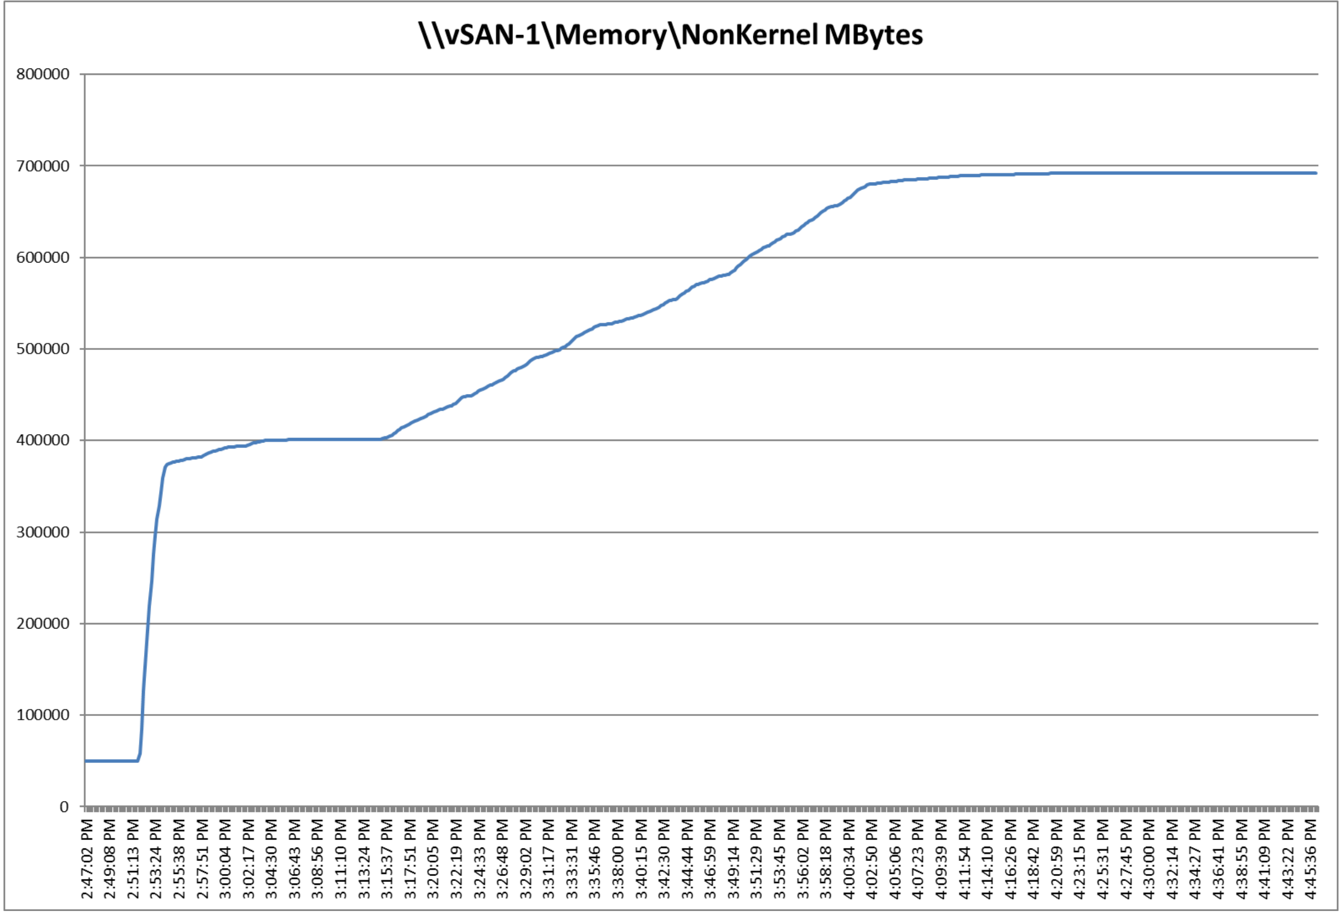

Figure 81. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Network Utilization

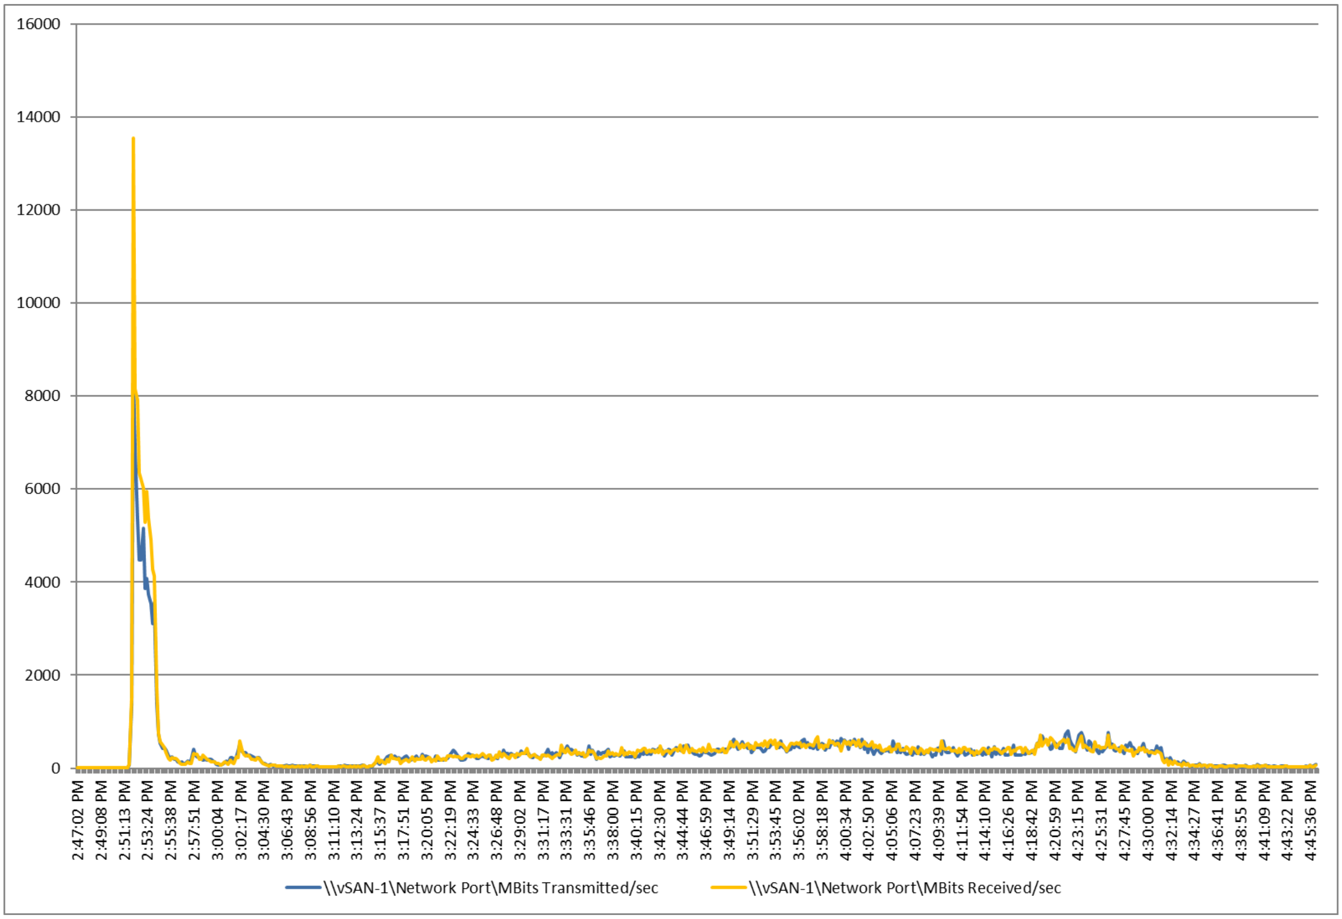

Figure 82. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VMware vSAN cluster Latency Chart

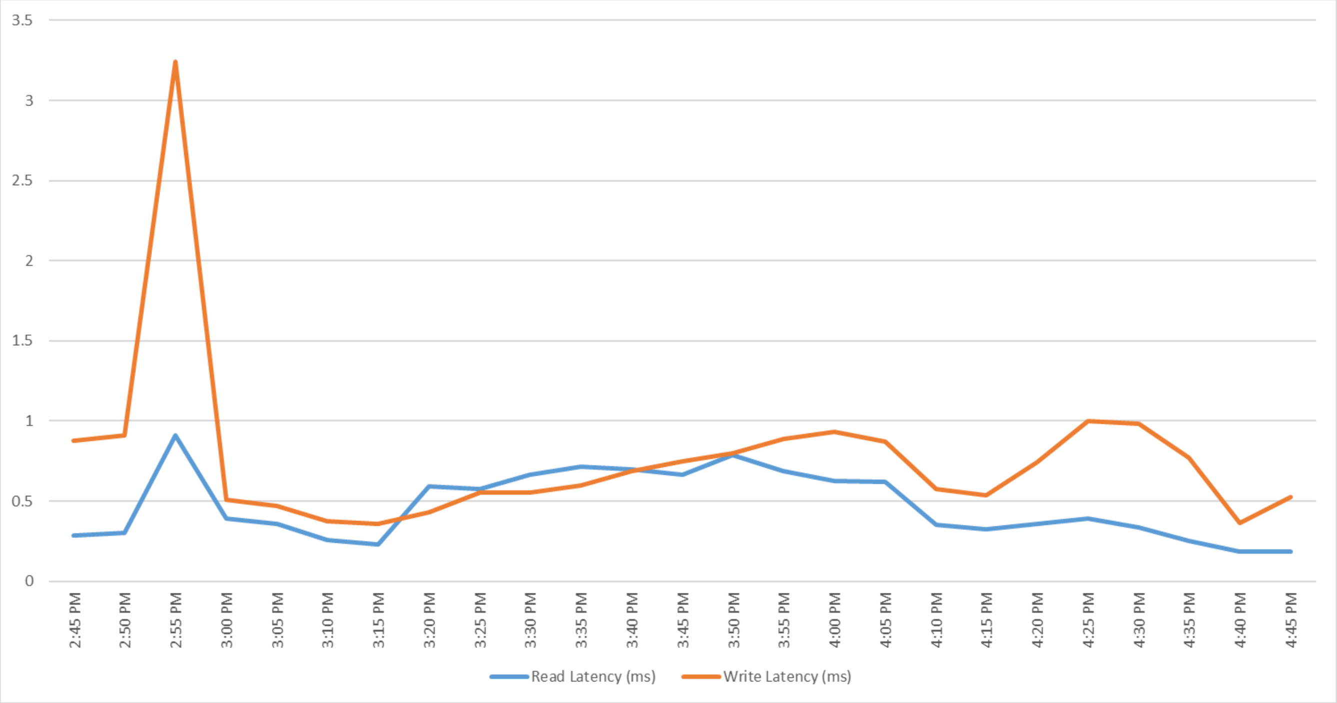

Figure 83. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VMware vSAN cluster IOPS Chart

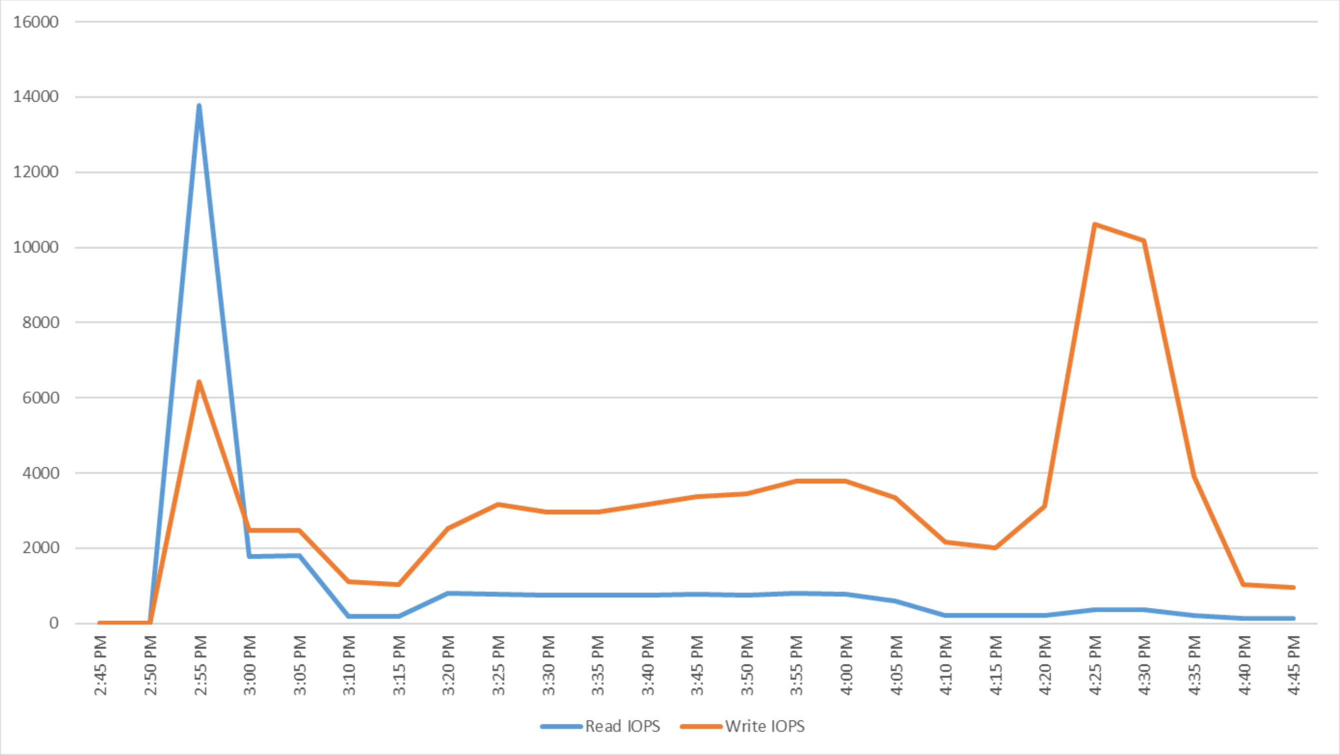

Figure 84. Full Scale | 585 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VMware vSAN cluster Throughput Chart

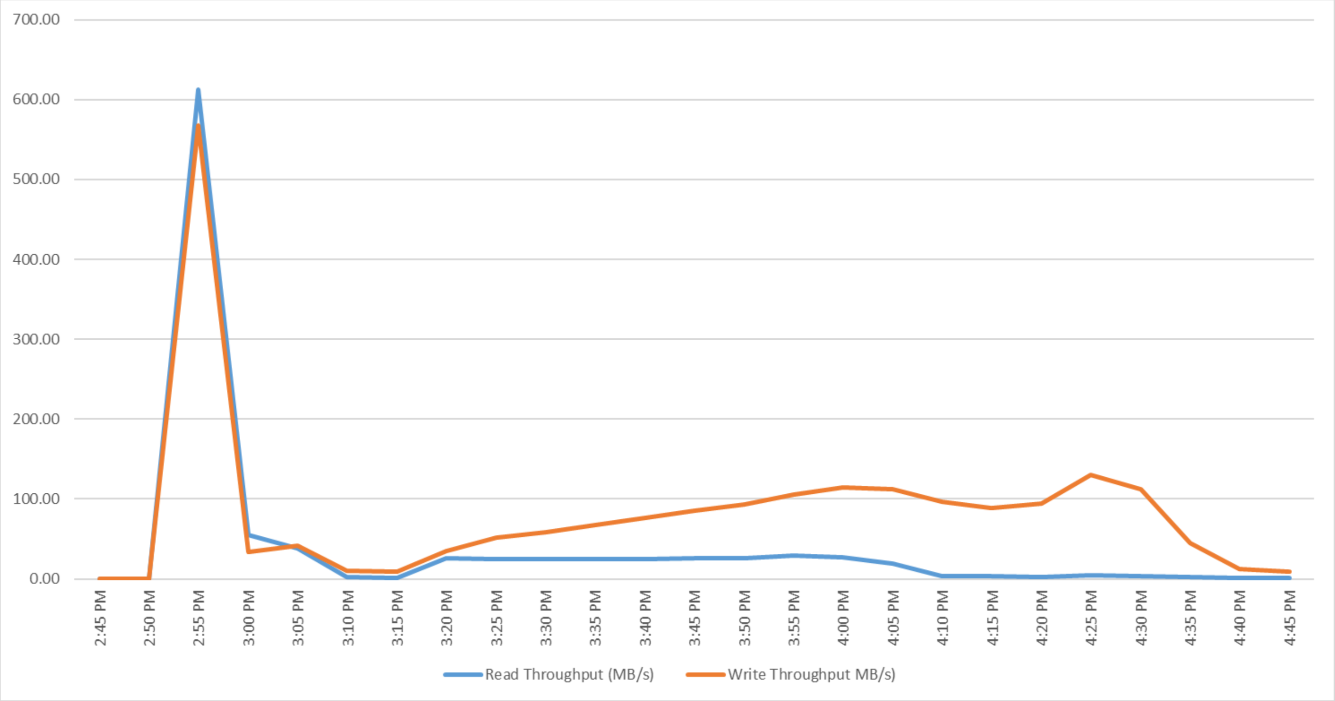

OSA vSAN: Full Scale Recommended Maximum Workload for Non-persistent Multi-session OS Random Sessions with 1440 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware vSAN cluster, during the Non-persistent Multi-session OS full-scale testing with 1440 Desktop Sessions using 8 blades configured in single Host Pool.

The Multi-session OS workload for the solution is 1440 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 85. Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Score

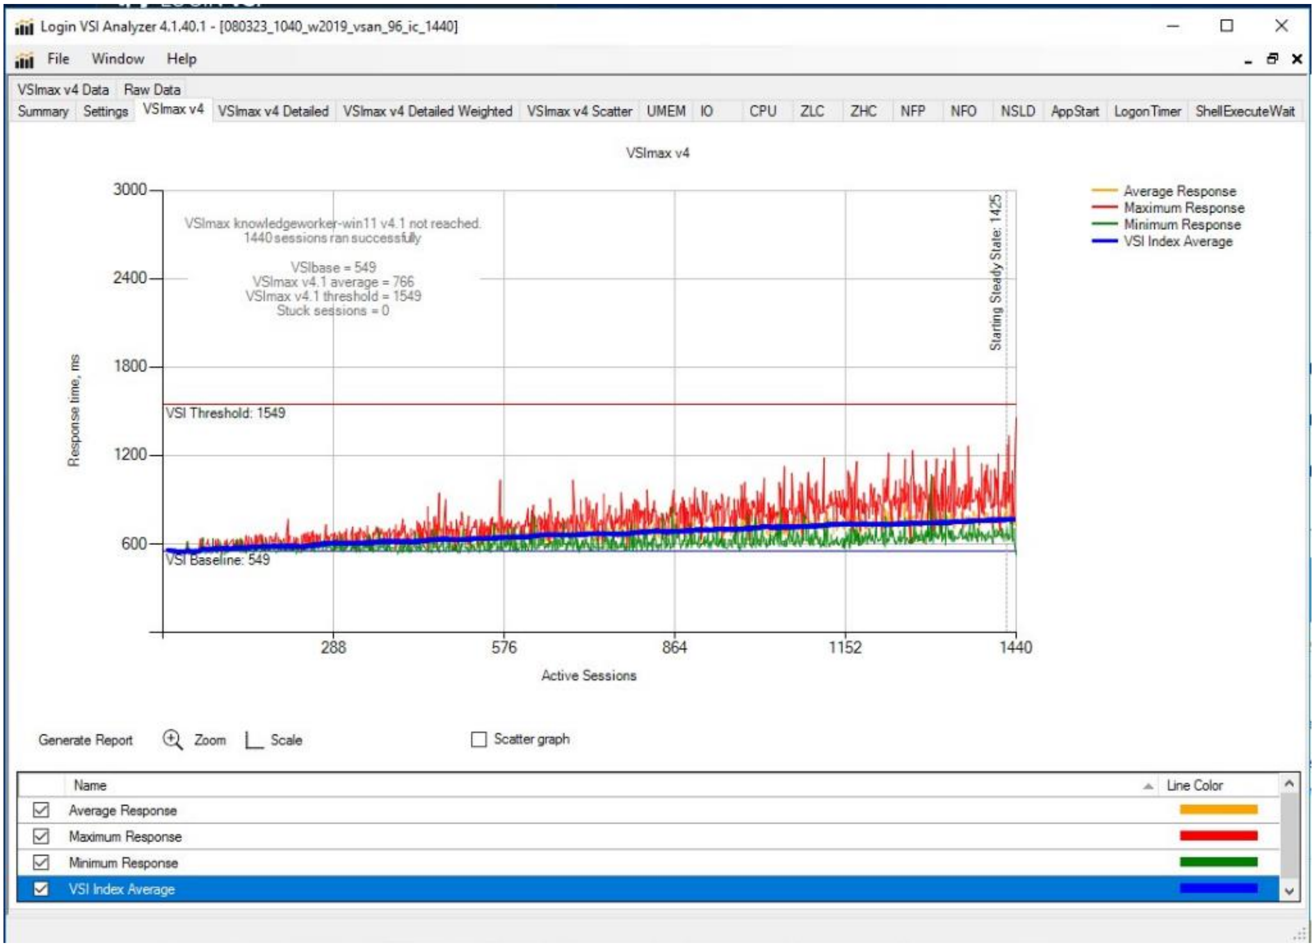

VSImax knowledgeworker-win11 v4.1 not reached.
1440 sessions ran successfully

VSIbase = 549
VSImax v4.1 average = 766
VSImax v4.1 threshold = 1549
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 86. Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host CPU Utilization

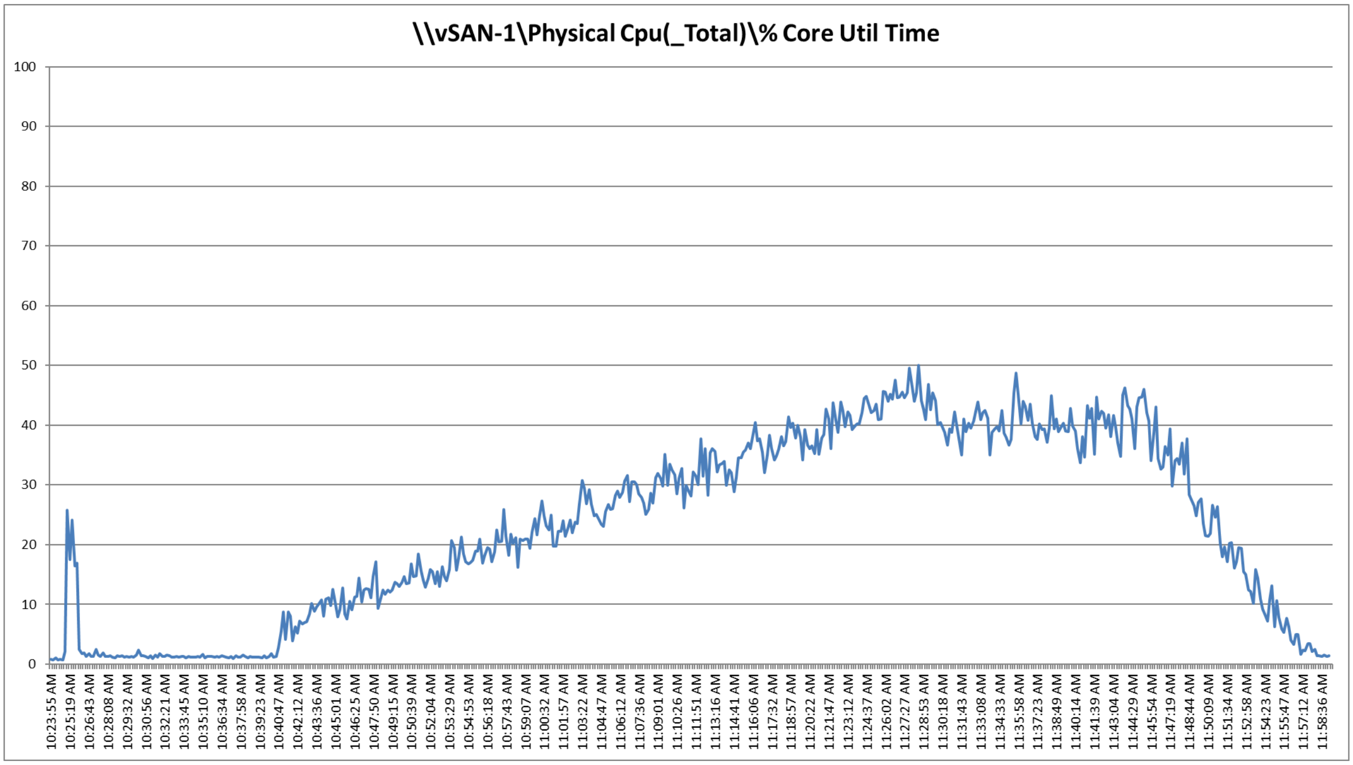

Figure 87. Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Memory Utilization

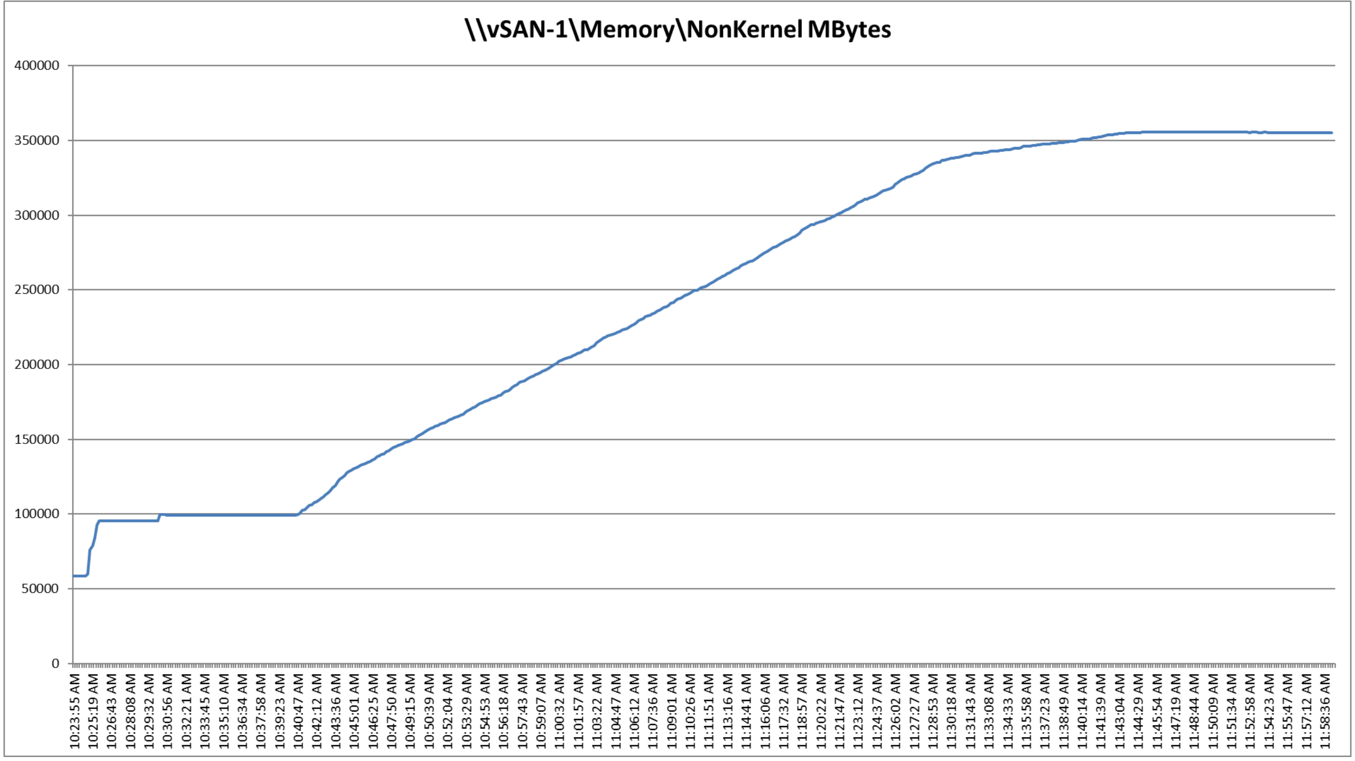

Figure 88. **Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Network Utilization**

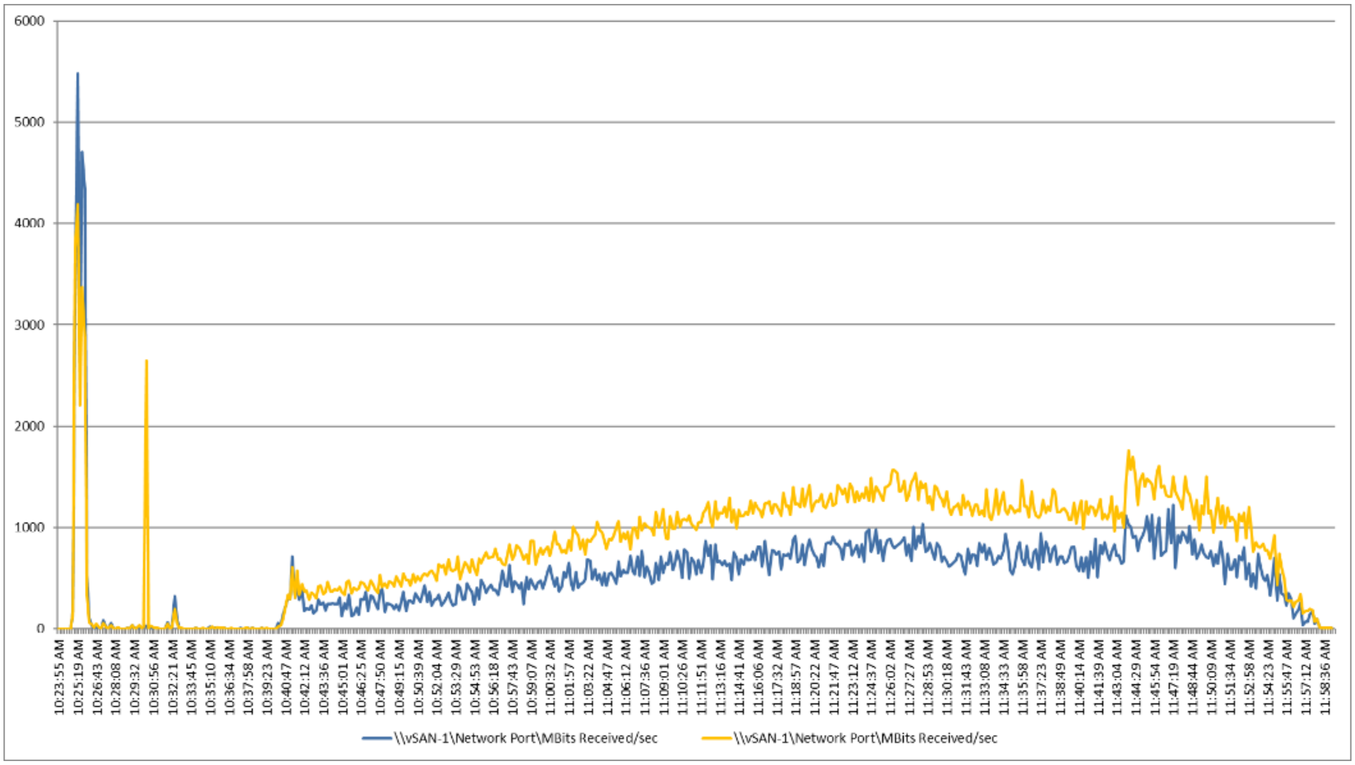

Figure 89. **Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VMware vSAN cluster Latency Chart**

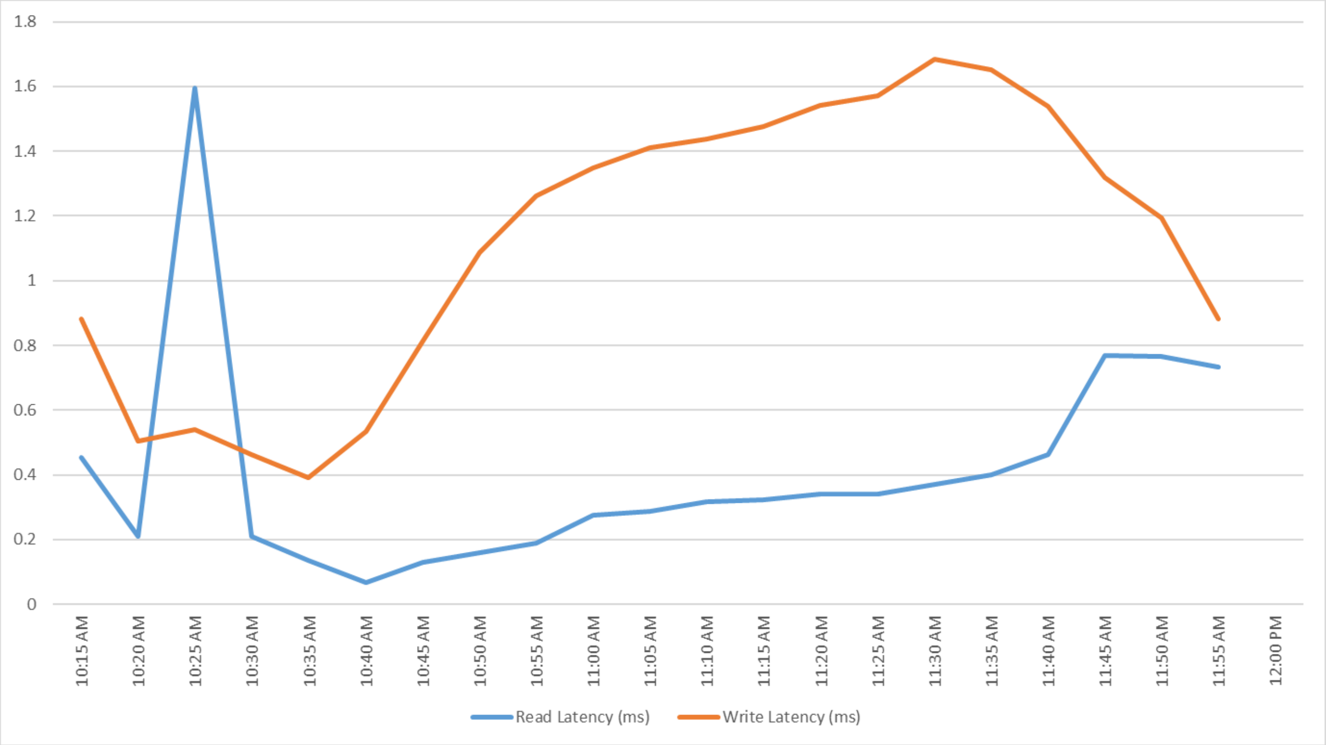

Figure 90. Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VMware vSAN cluster IOPS Chart

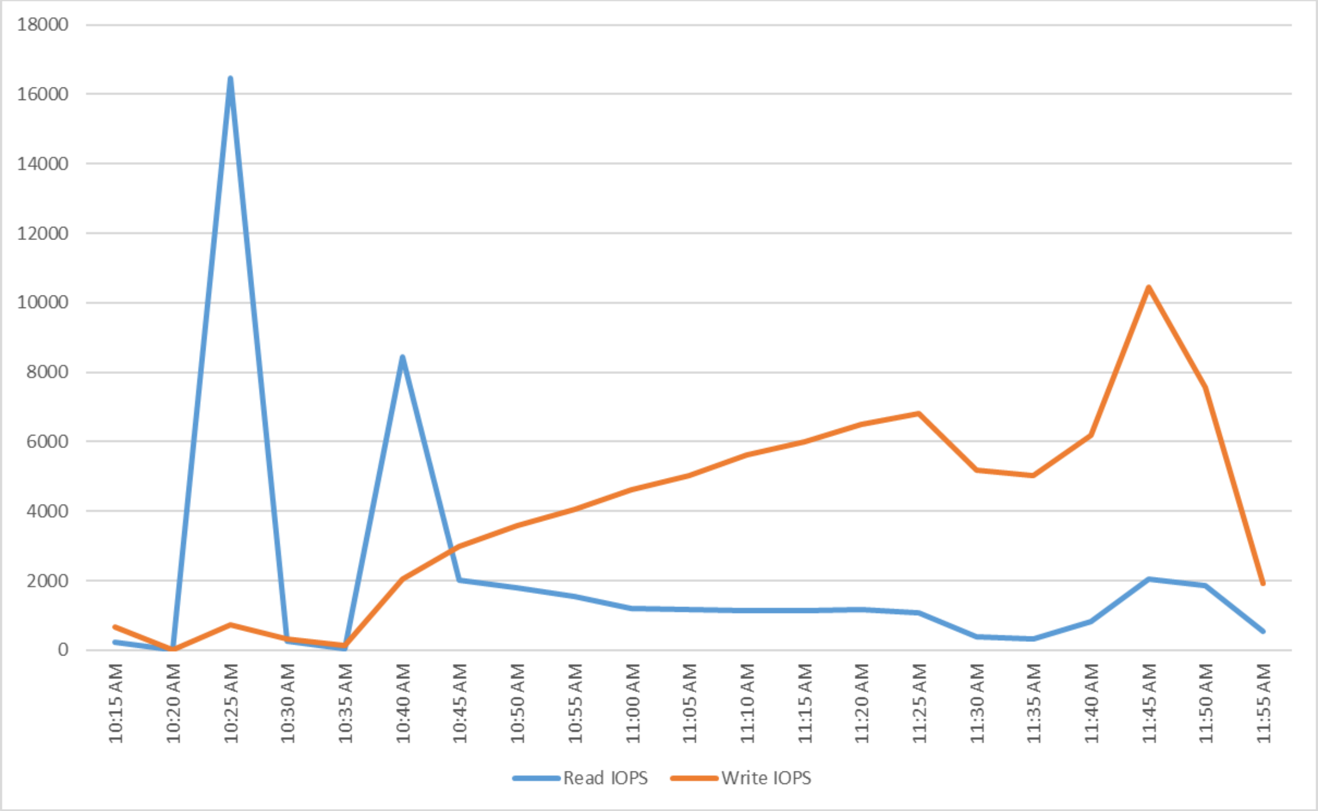

Figure 91. Full Scale | 1440 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VMware vSAN cluster Throughput Chart

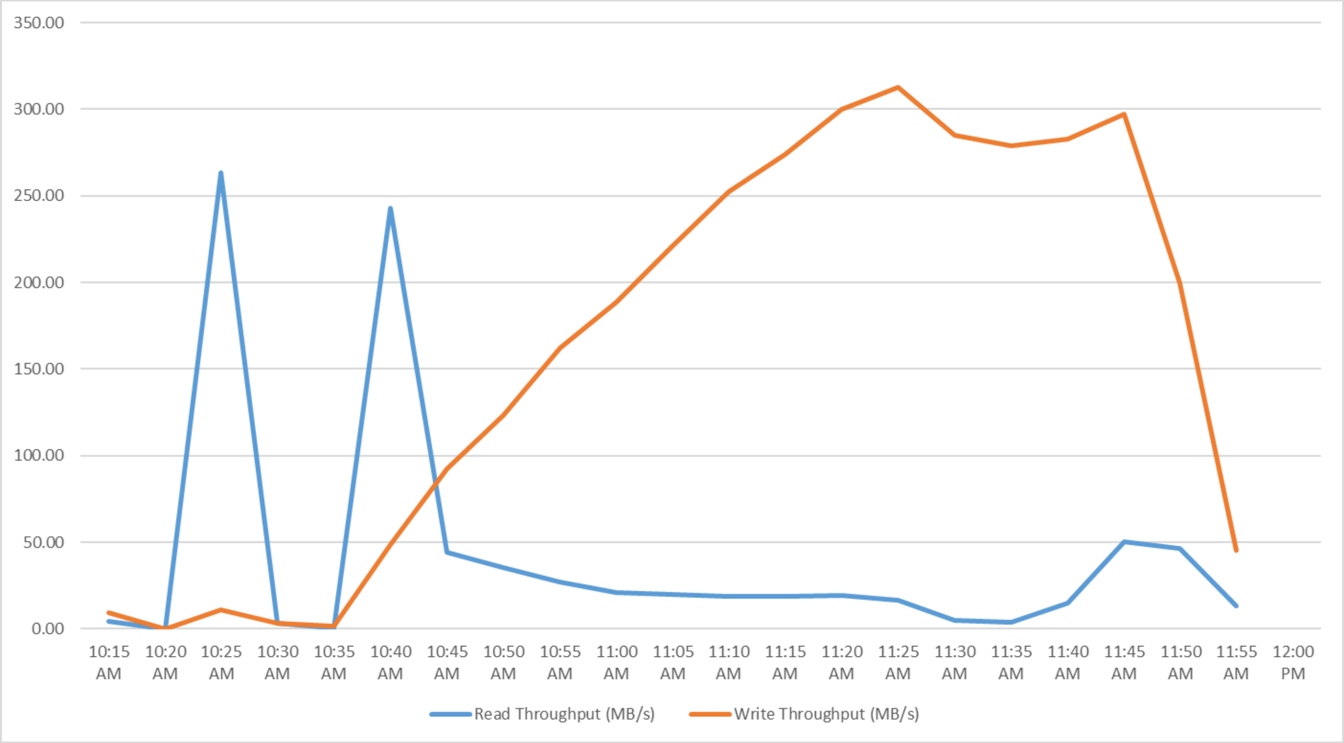

ESA vSAN: Full Scale Recommended Maximum Workload Testing for Non-persistent Single-session OS Machine VDAs with 390 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware vSAN cluster during the full-scale testing with 390 Non-persistent Single-session OS machines using four blades in a single pool.

The workload for the test is 390 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 92. Full Scale | 585 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs| VSI Score

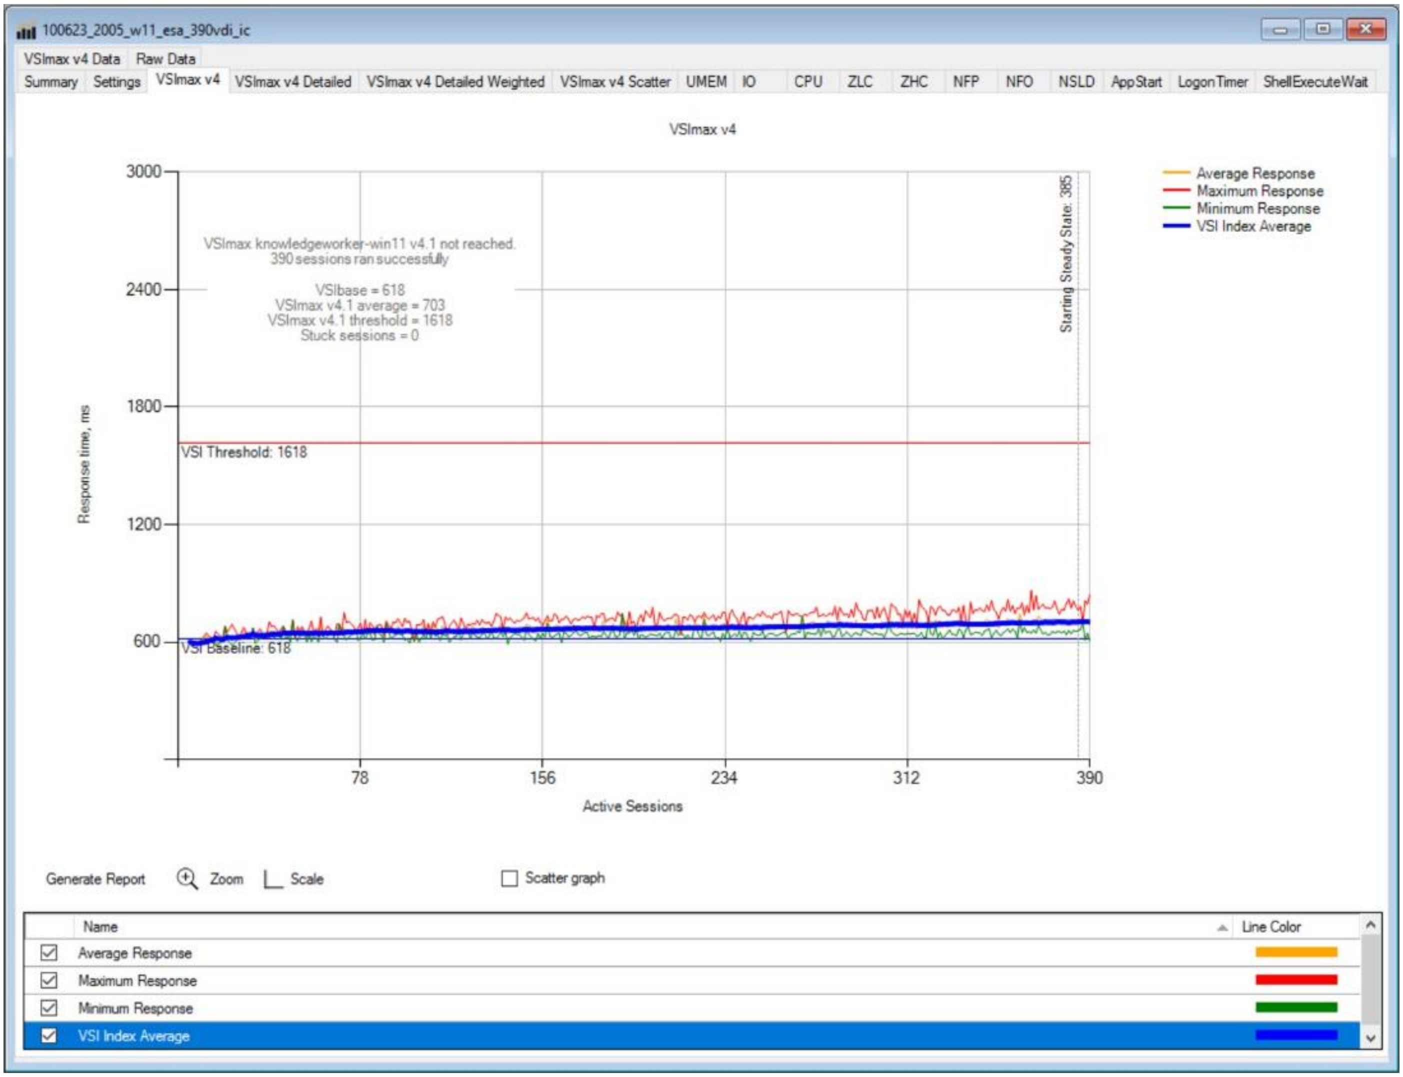

VSImax knowledgeworker-win11 v4.1 not reached.
390 sessions ran successfully

VSIbase = 618
VSImax v4.1 average = 703
VSImax v4.1 threshold = 1618
Stuck sessions = 0

□ Scatter graph

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 93. Full Scale | 390 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host CPU Utilization

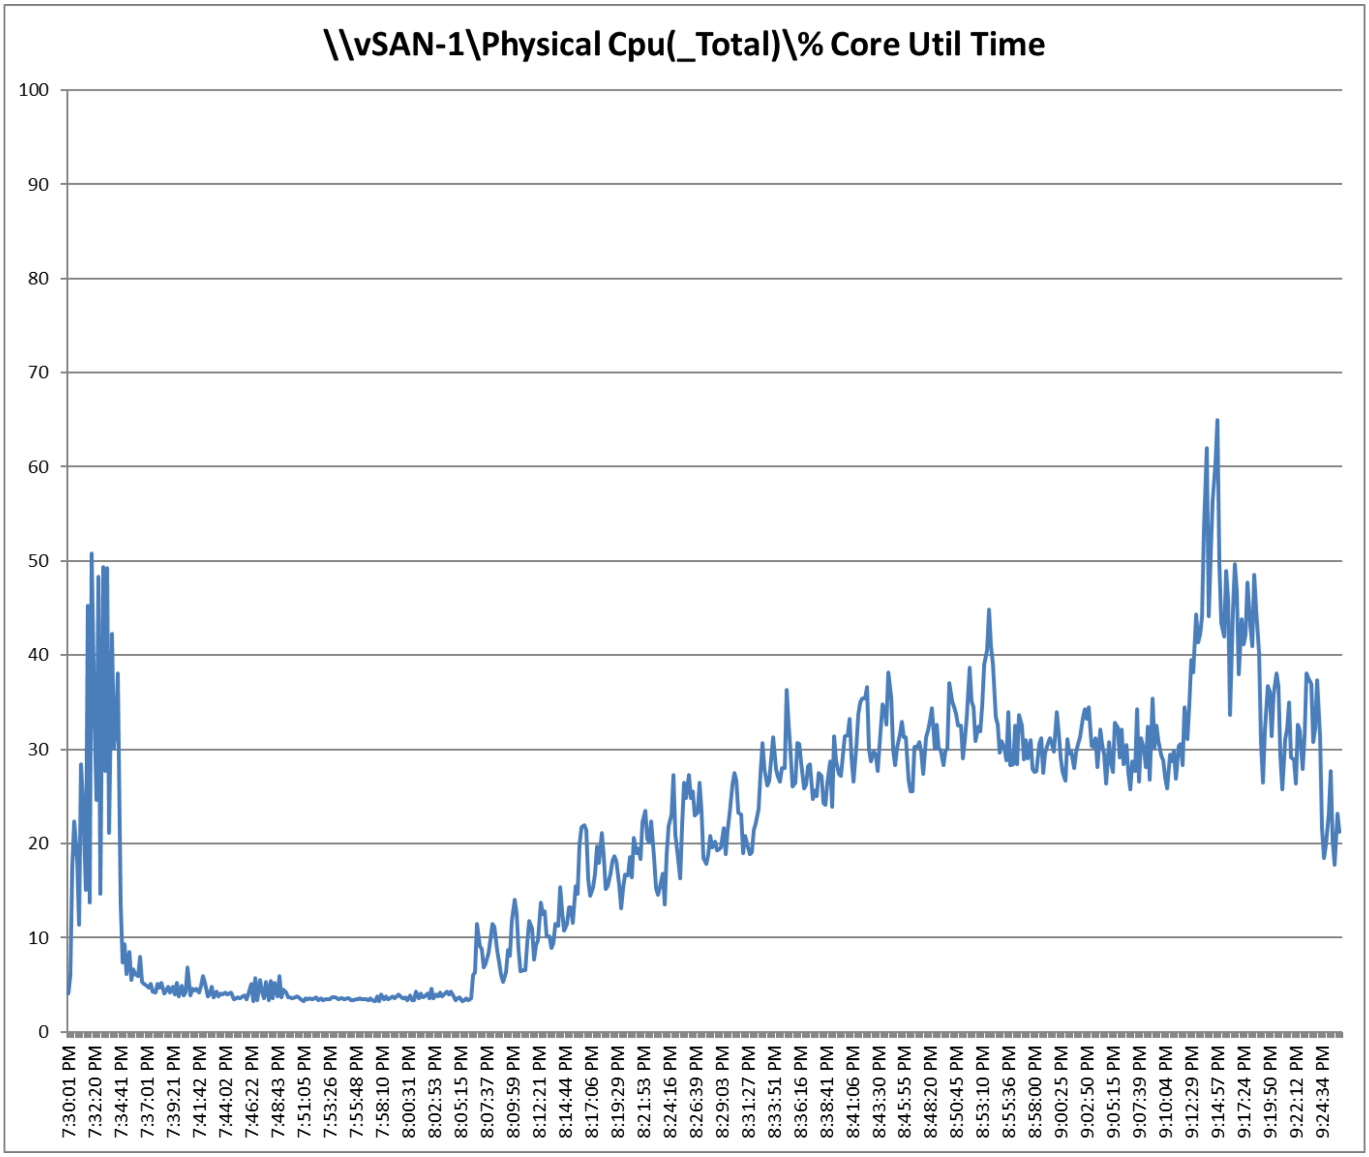

Figure 94. Full Scale | 390 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Memory Utilization

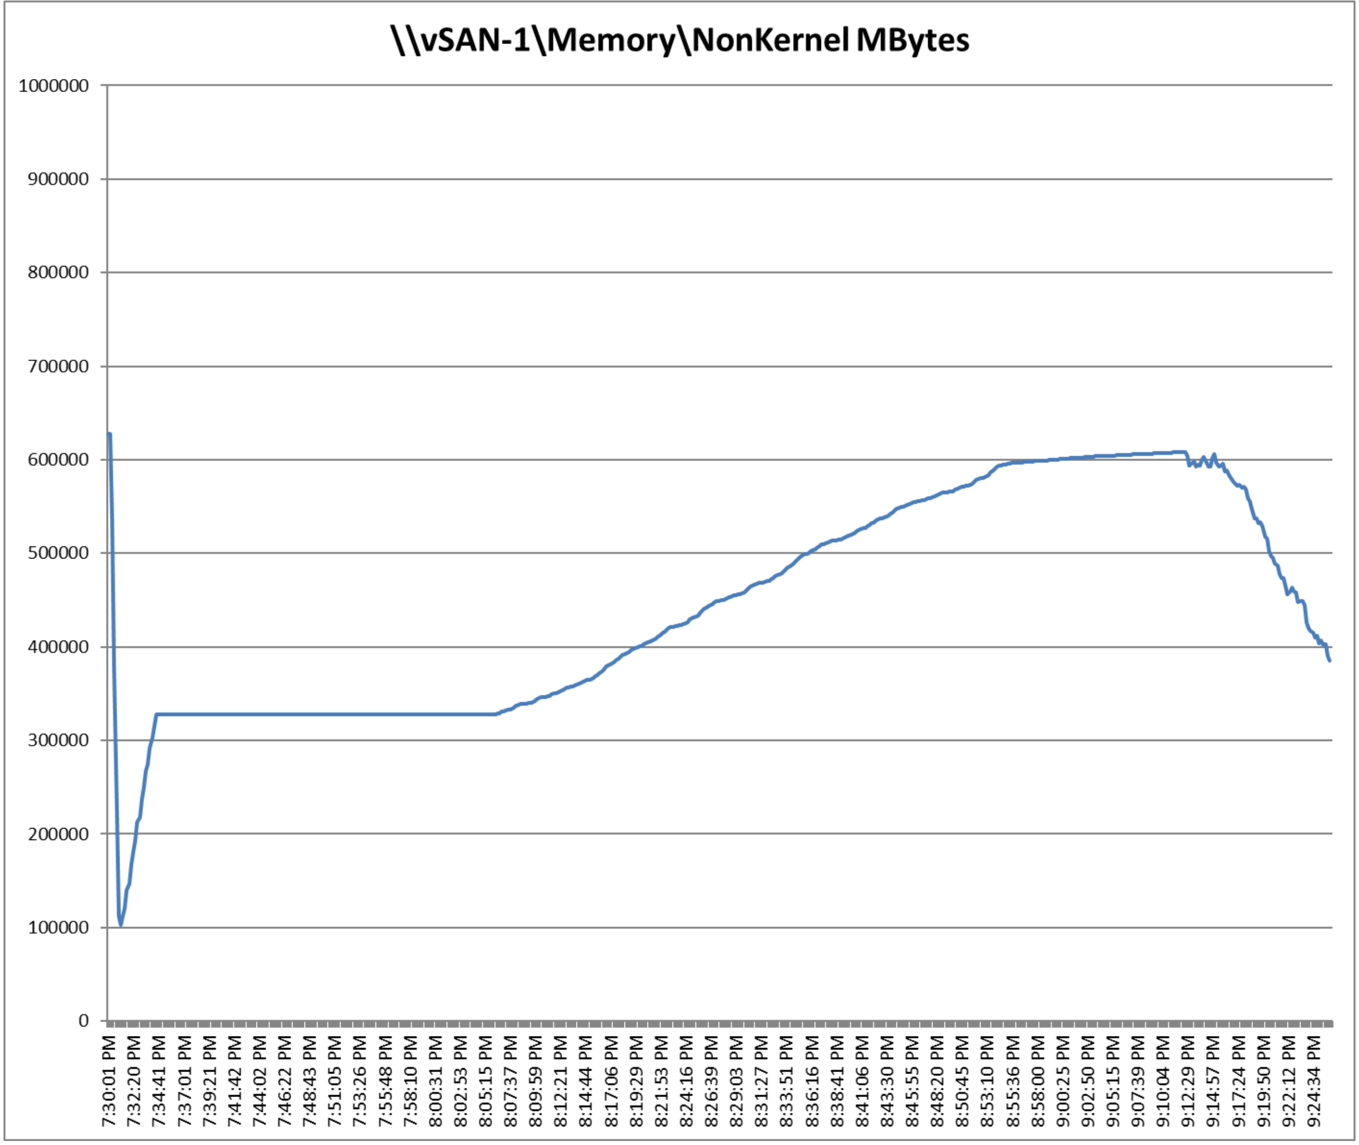

Figure 95. Full Scale | 390 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | Host Network Utilization

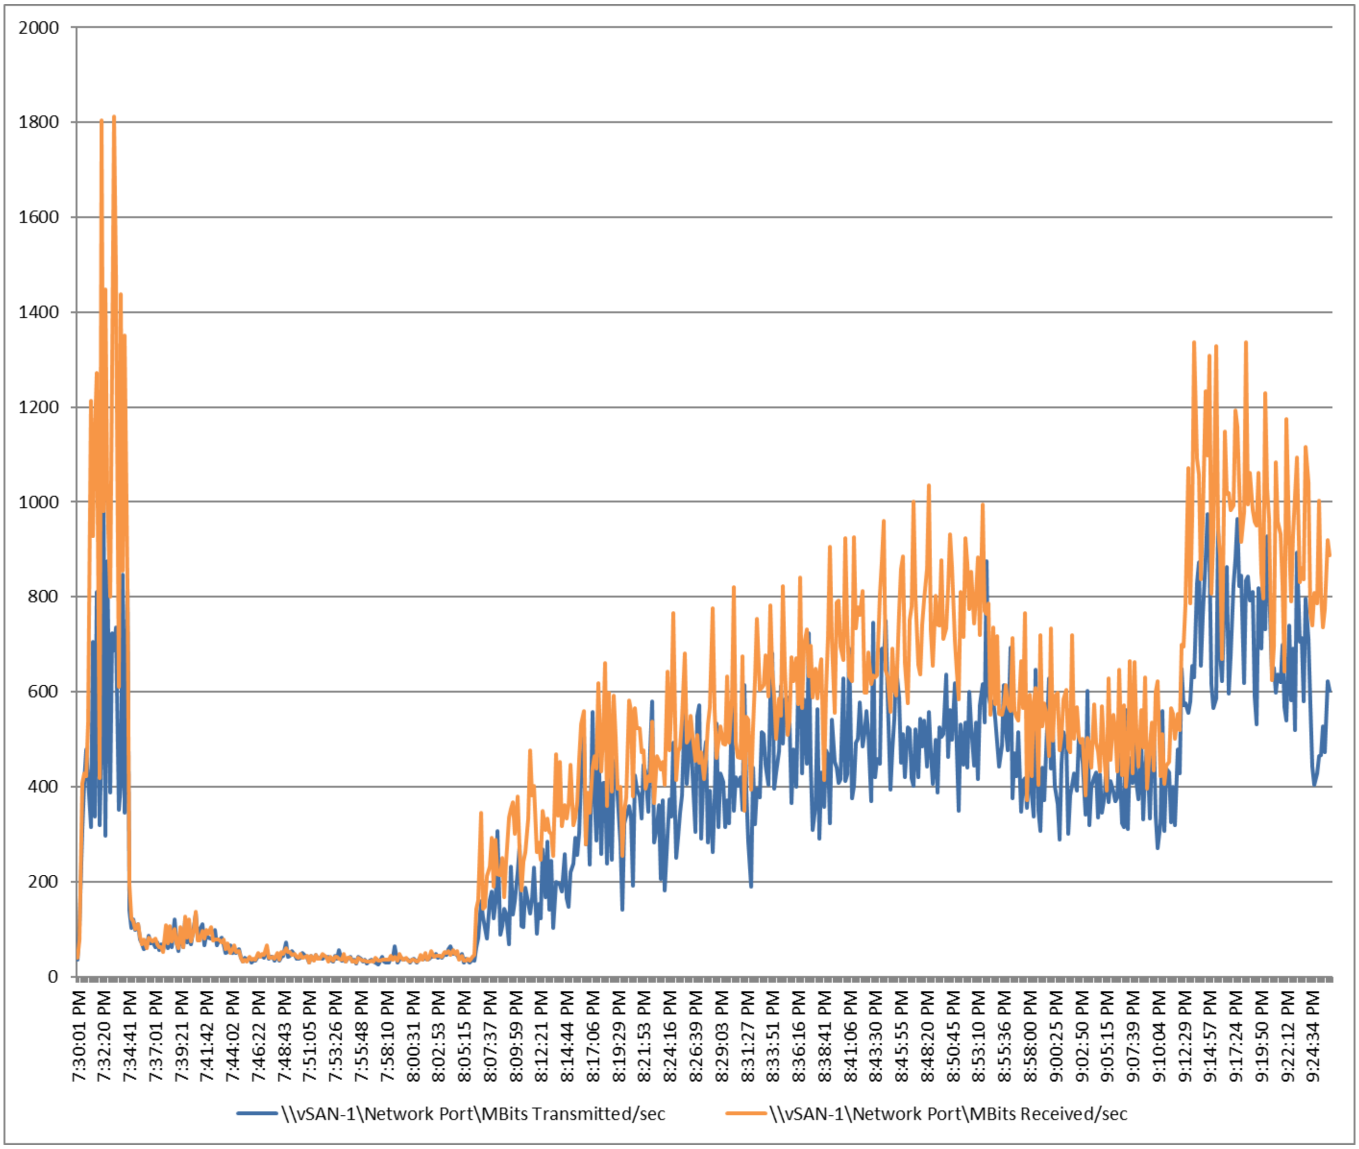

Figure 96. Full Scale | 390 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | vSAN Latency Chart

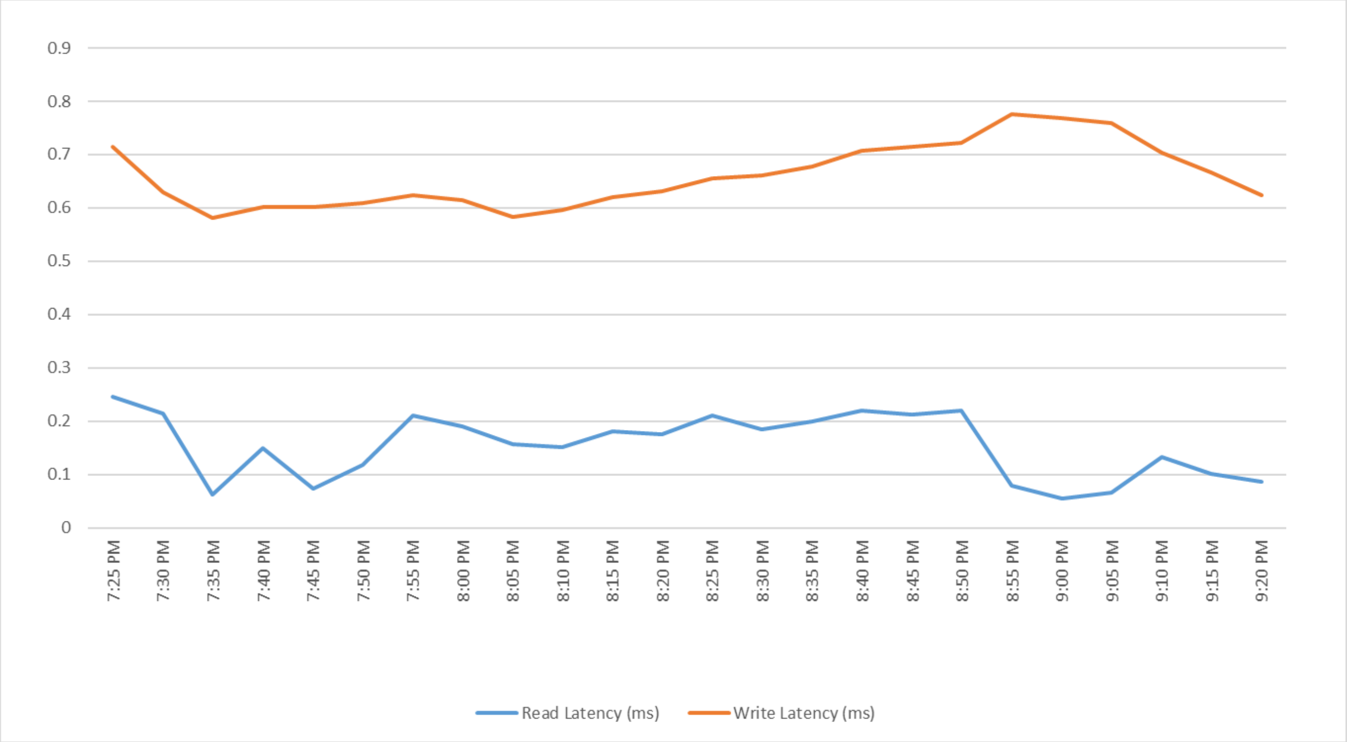

Figure 97. Full Scale | 390 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | vSAN IOPS Chart

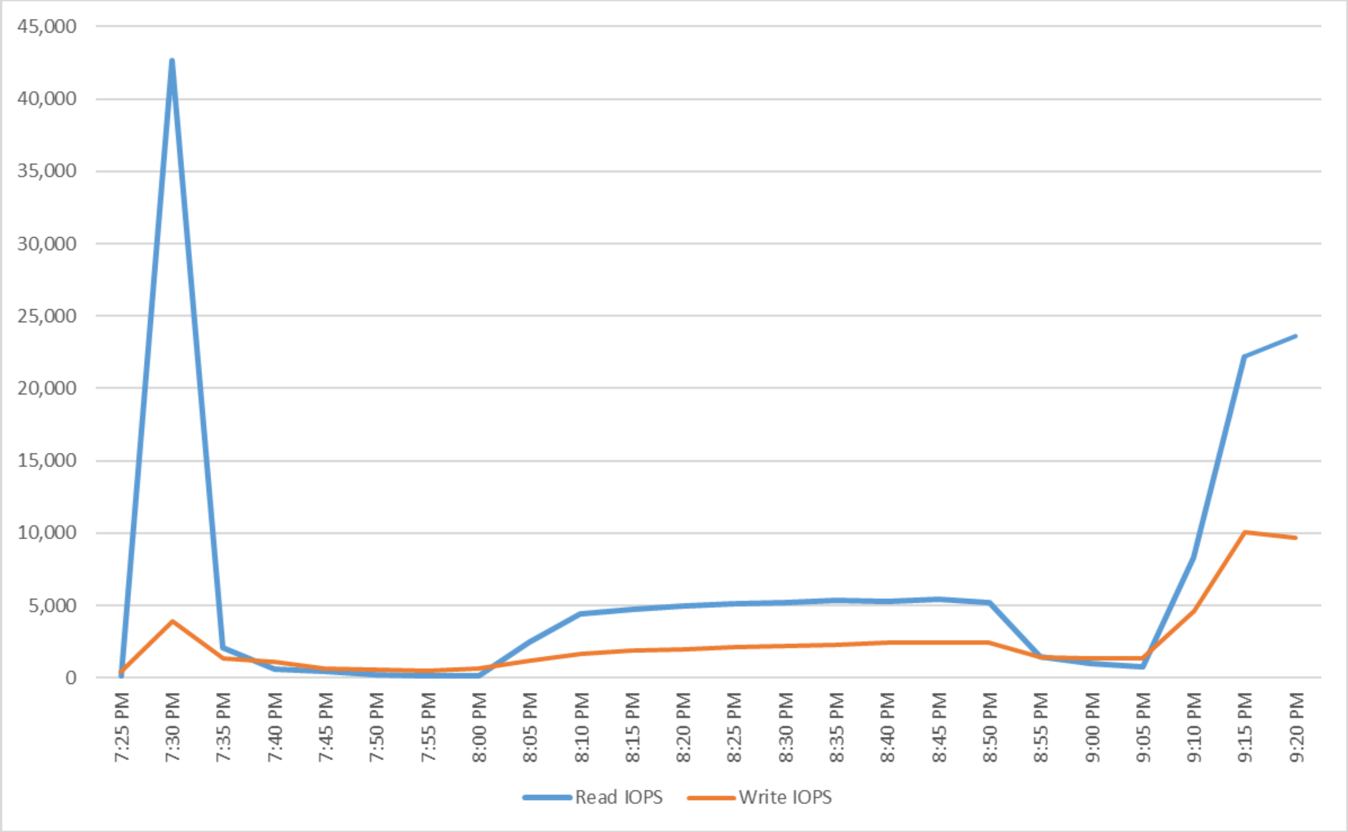

Figure 98. Full Scale | 390 Users | VMware Horizon 8 2212 Non-persistent Single-session OS machine VDAs | vSAN Throughput Chart

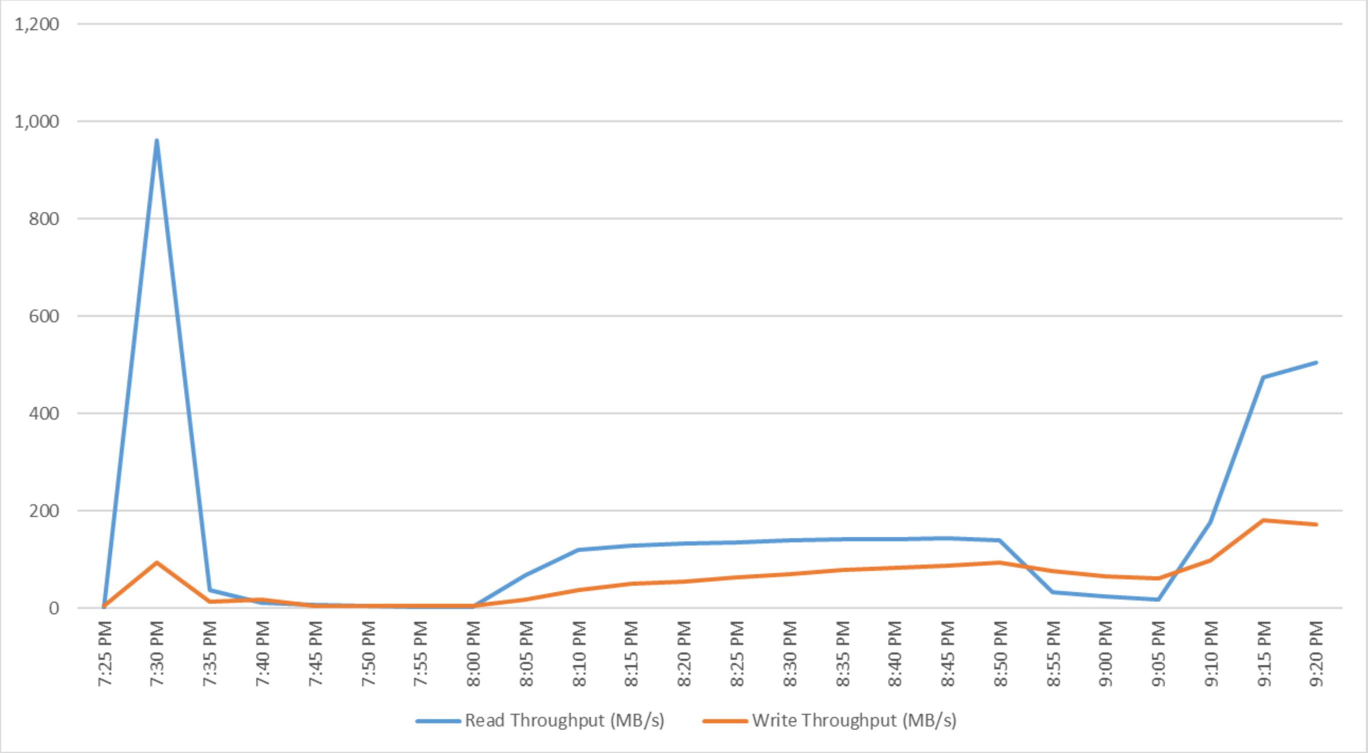

ESA vSAN: Full Scale Recommended Maximum Workload Testing for Persistent Single-session OS Machine VDAs with 390 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware vSAN cluster during the persistent desktop full-scale testing with 390 Persistent Single-session OS machines using 4 blades in a single cluster.

The workload for the test is 390 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 99. **Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VSI Score**

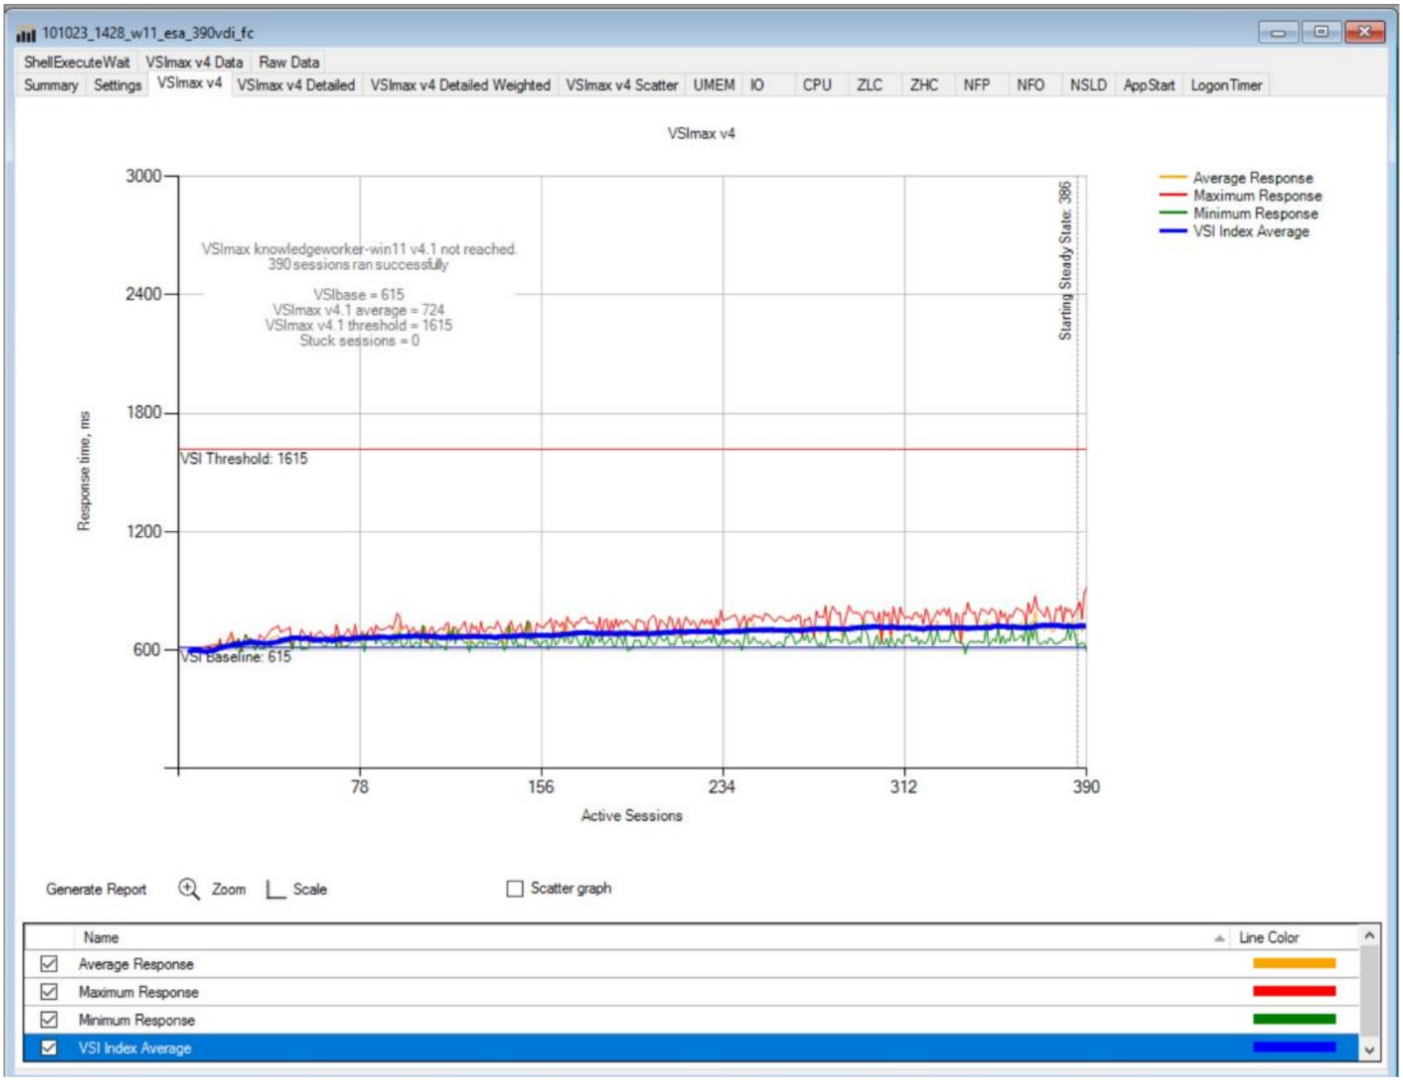

VSImax knowledgeworker-win11 v4.1 not reached.
390 sessions ran successfully

VSIbase = 615
VSImax v4.1 average = 724
VSImax v4.1 threshold = 1615
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 100. **Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host CPU Utilization**

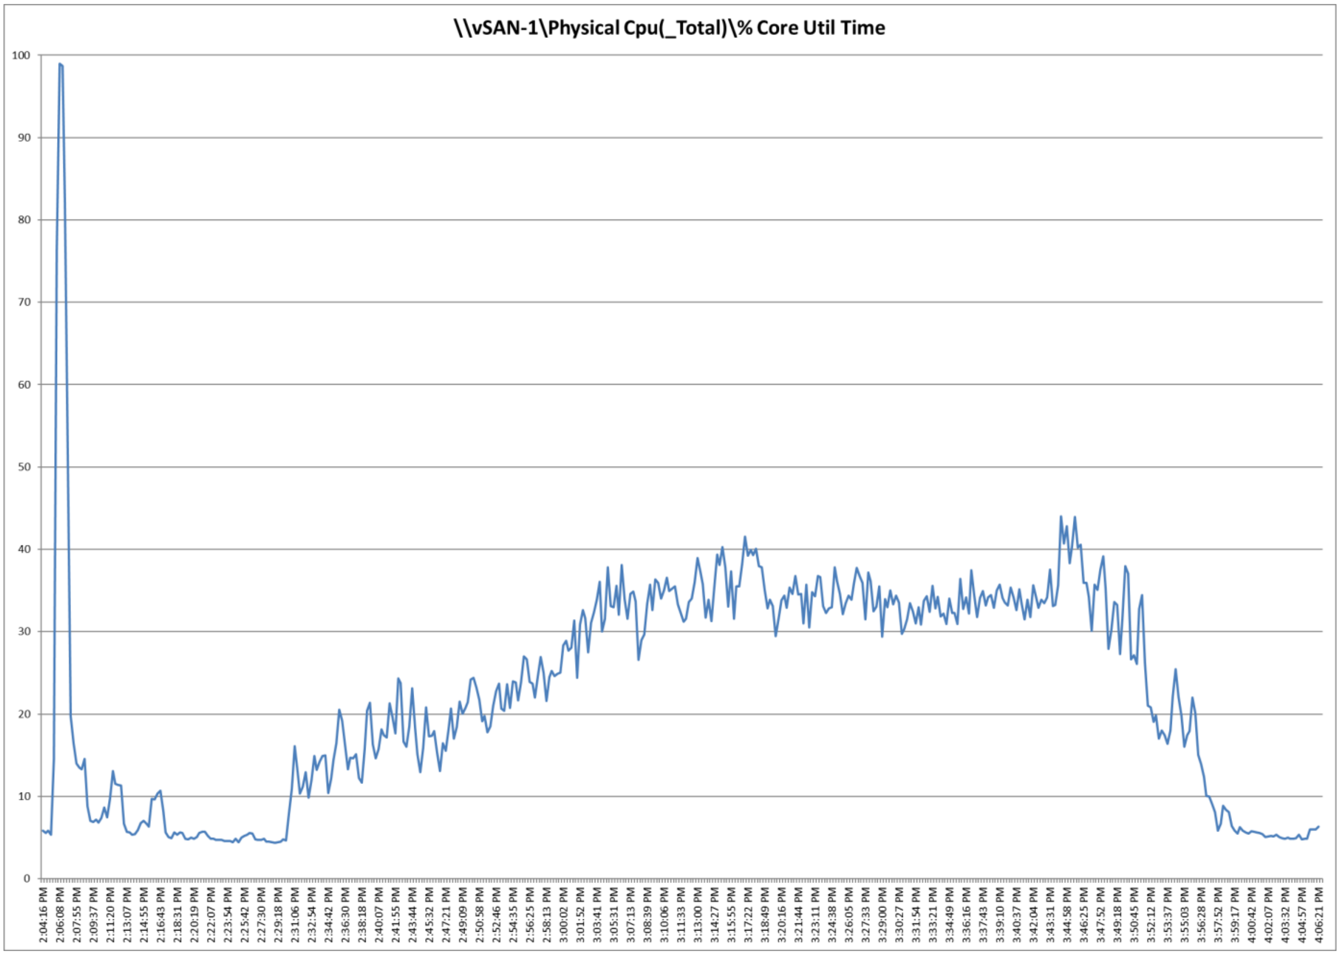

Figure 101. **Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Memory Utilization**

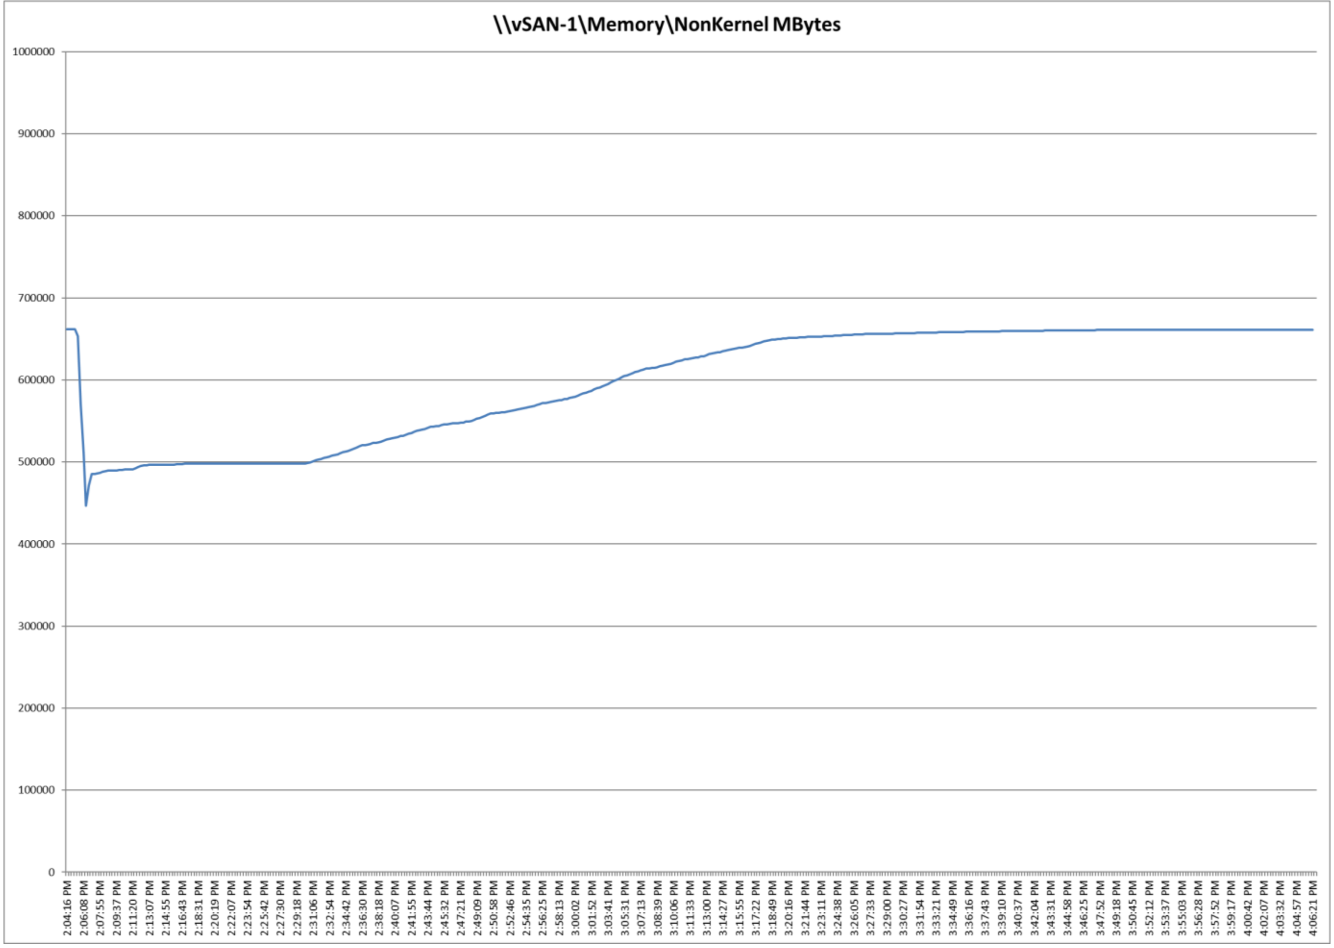

Figure 102. Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | Host Network Utilization

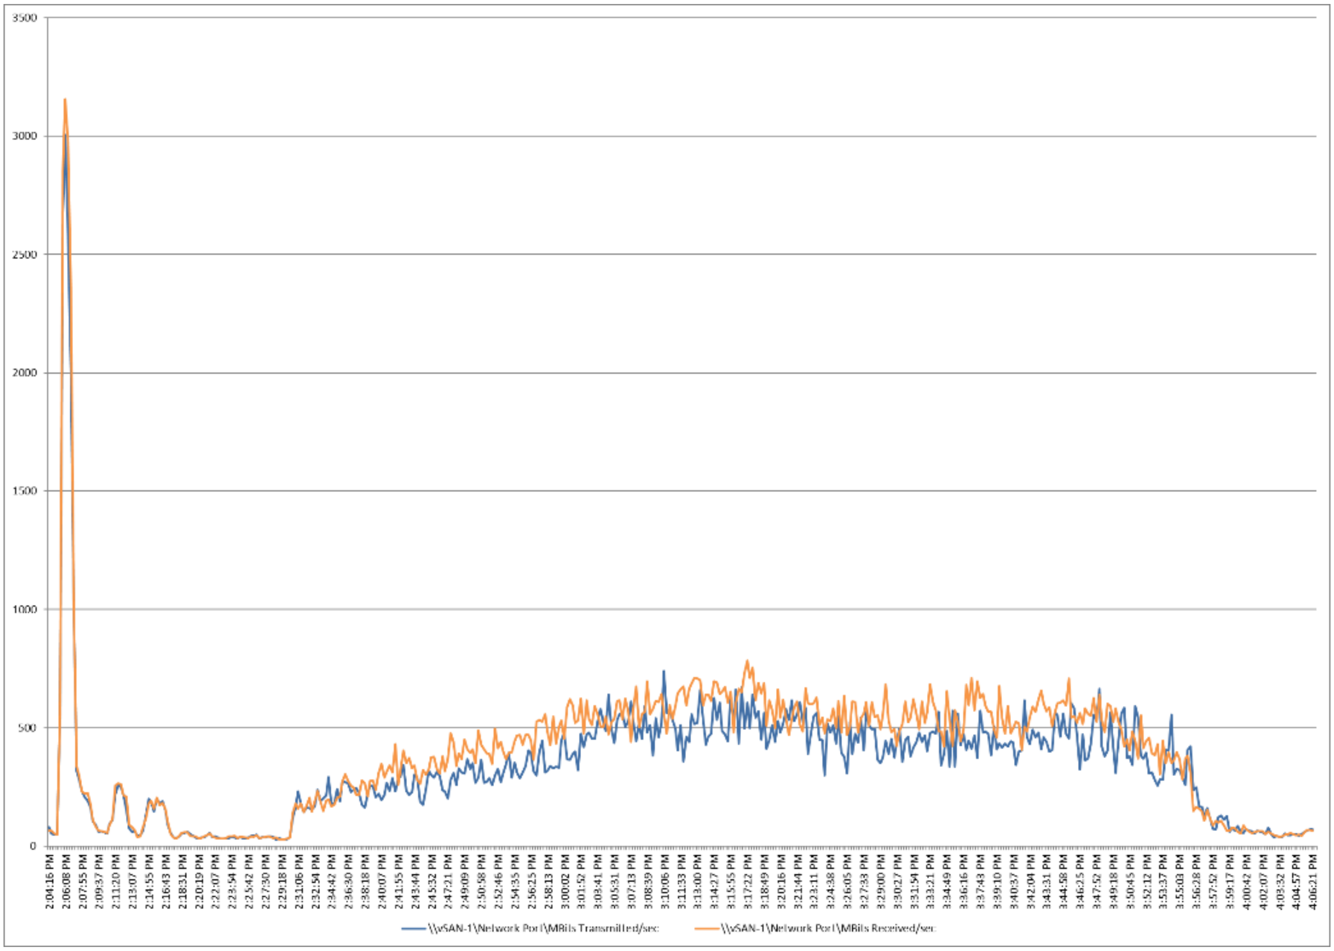

Figure 103. Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VMware vSAN cluster Latency Chart

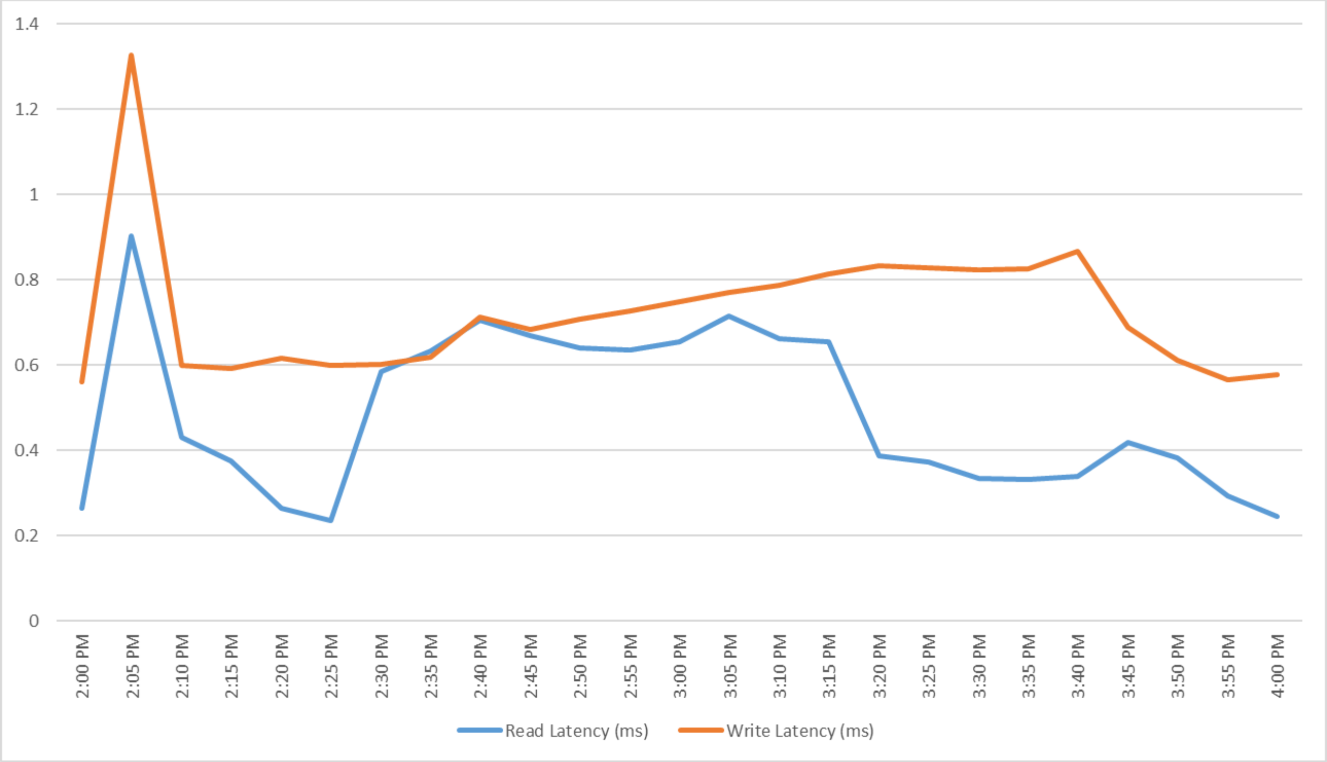

Figure 104. Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VMware vSAN cluster IOPS Chart

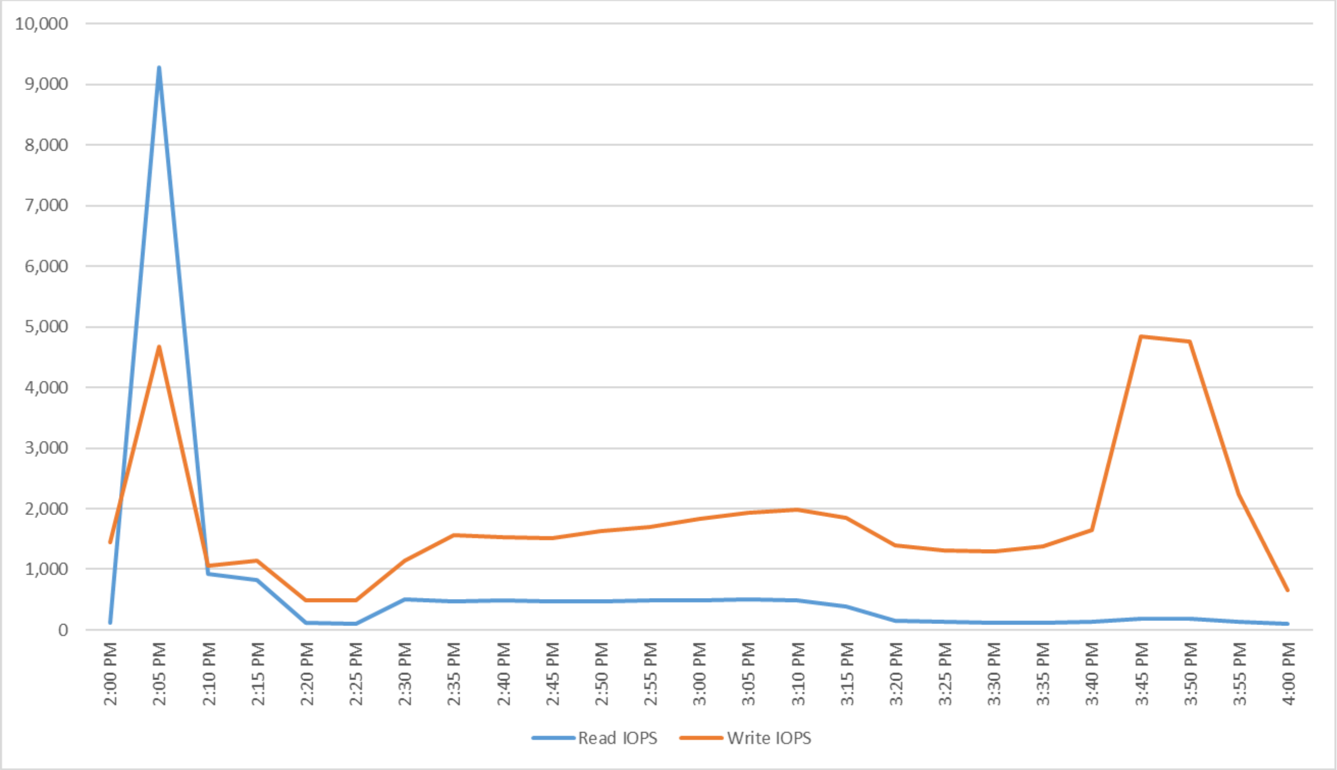

Figure 105. Full Scale | 390 Users | VMware Horizon 8 2212 Persistent Single-session OS machine VDAs | VMware vSAN cluster Throughput Chart

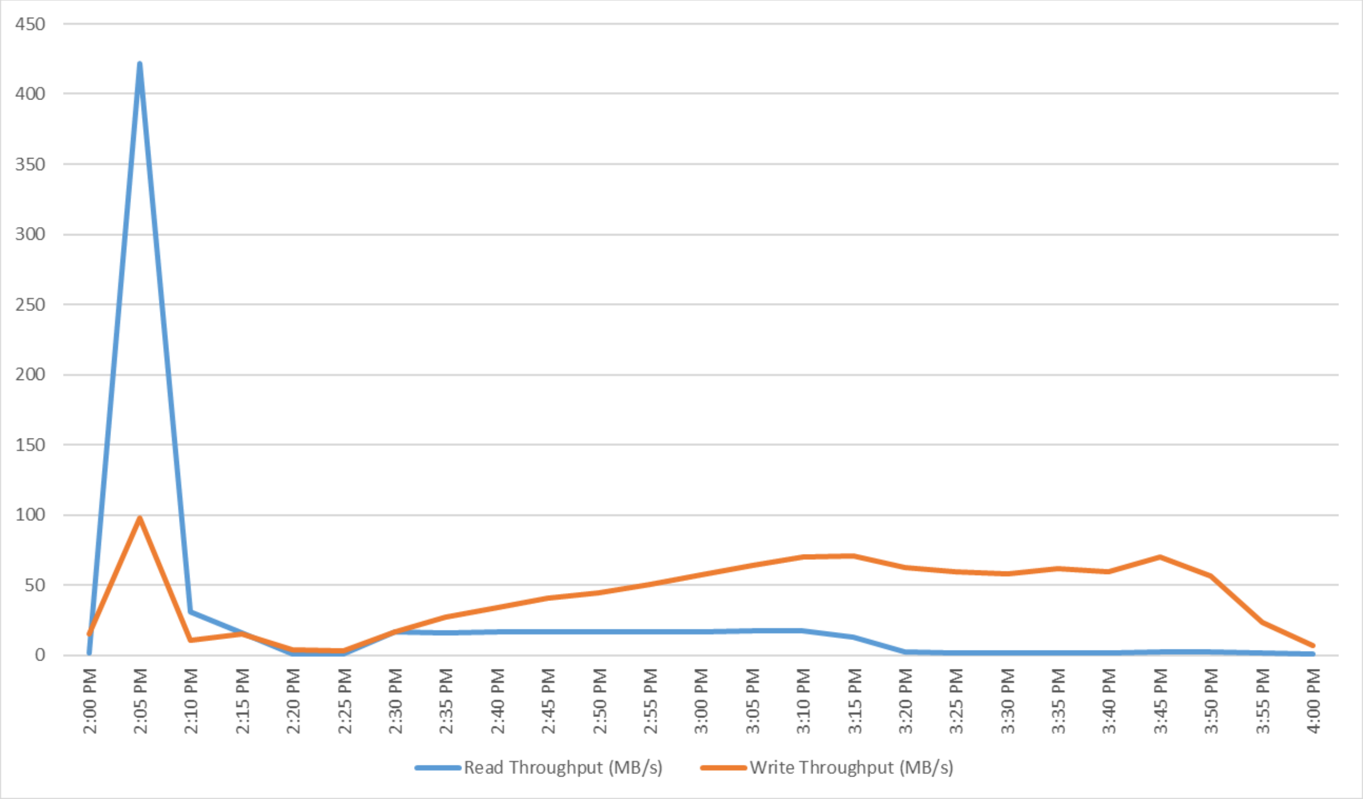

ESA vSAN: Full Scale Recommended Maximum Workload for Non-persistent Multi-session OS Random Sessions with 960 Users

This section describes the key performance metrics that were captured on the Cisco UCS and VMware vSAN cluster, during the Non-persistent Multi-session OS full-scale testing with 960 Desktop Sessions using 8 blades configured in single Host Pool.

The Multi-session OS workload for the solution is 960 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 106. **Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VSI Score**

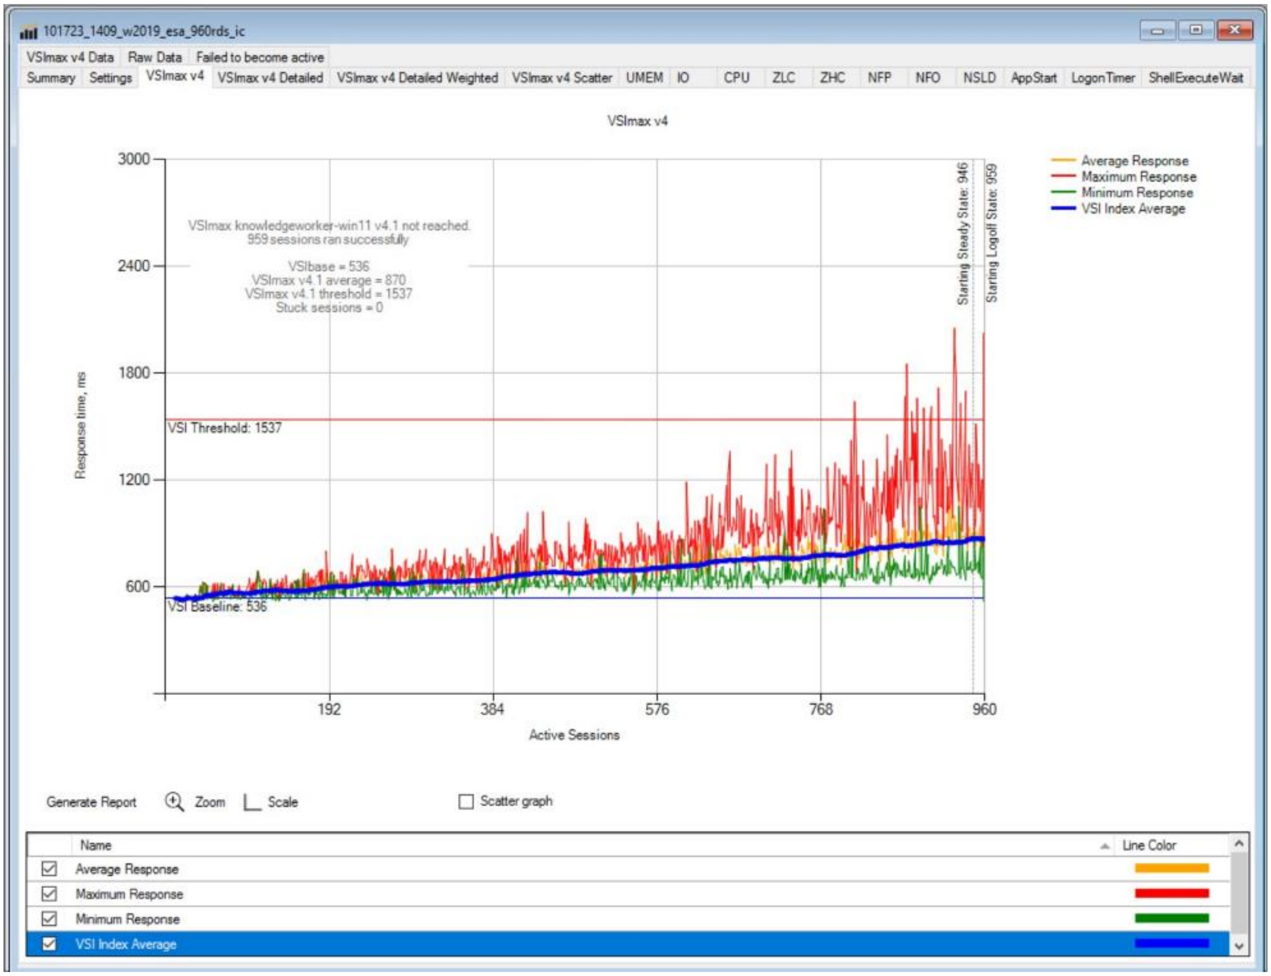

VSImax knowledgeworker-win11 v4.1 not reached.
959 sessions ran successfully

VSIbase = 536
VSImax v4.1 average = 870
VSImax v4.1 threshold = 1537
Stuck sessions = 0

	Name	Line Color
☑	Average Response	🟧
☑	Maximum Response	🟥
☑	Minimum Response	🟩
☑	VSI Index Average	🟦

Figure 107. Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host CPU Utilization

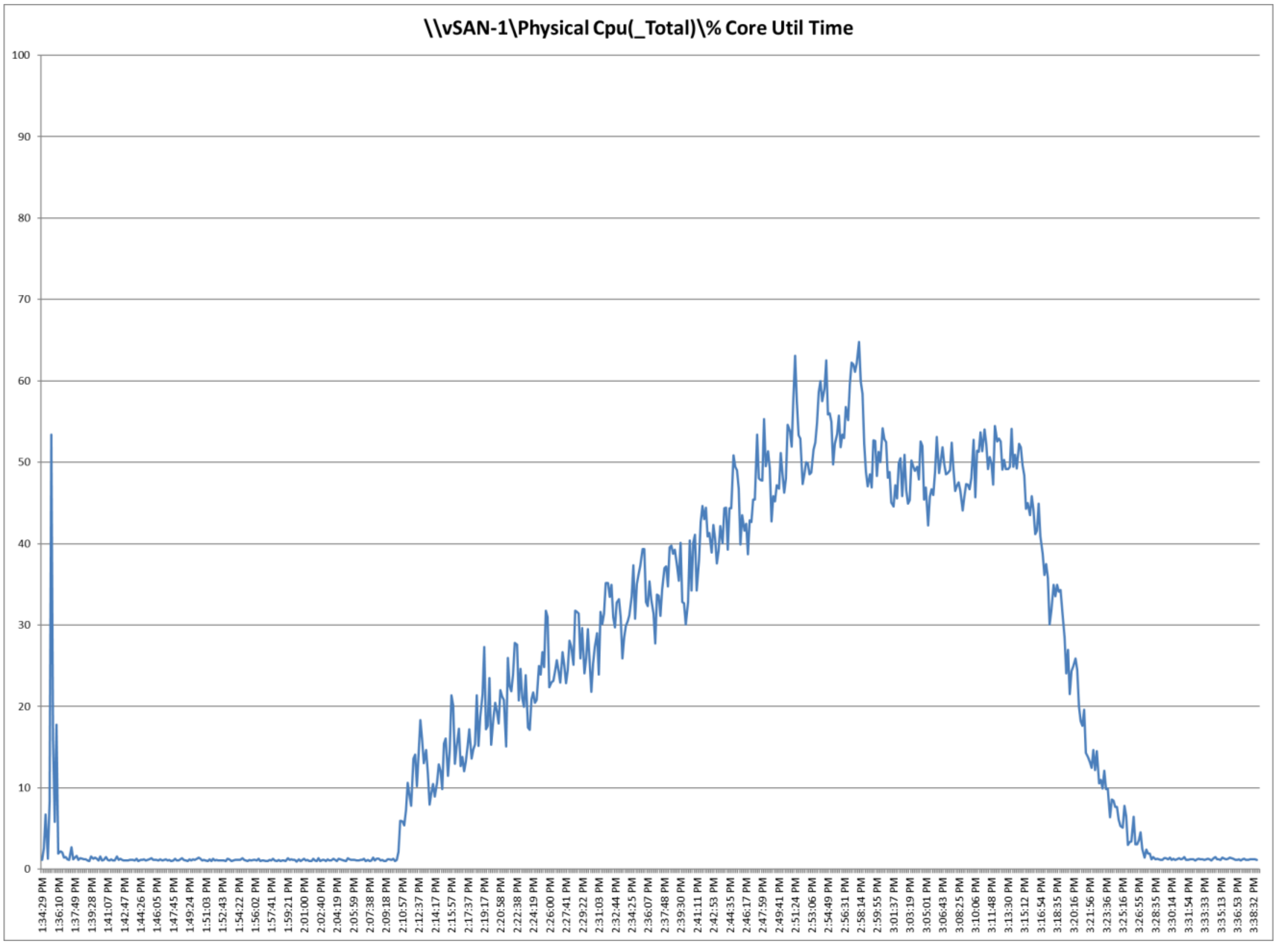

Figure 108. Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Memory Utilization

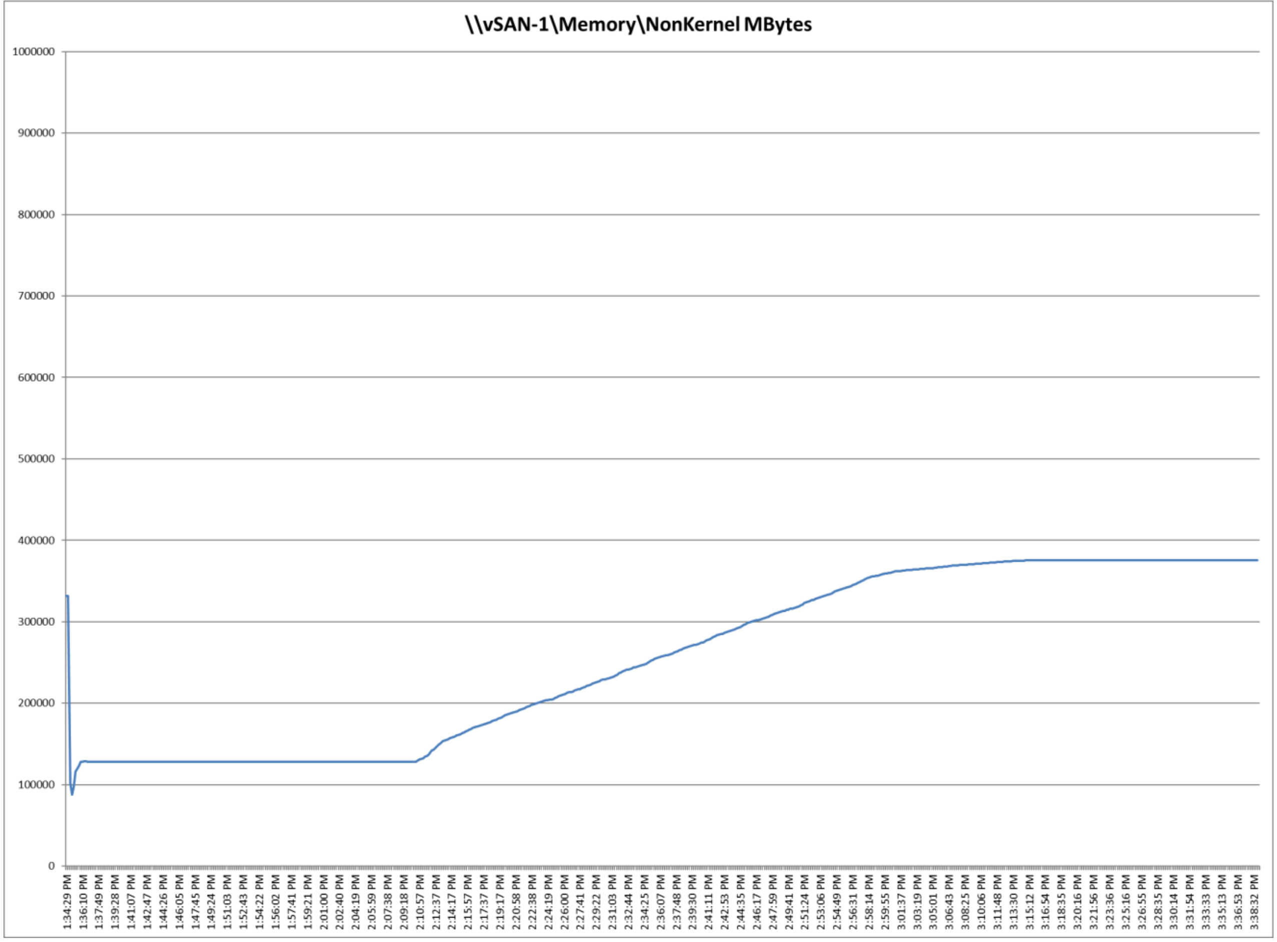

Figure 109. Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | Host Network Utilization

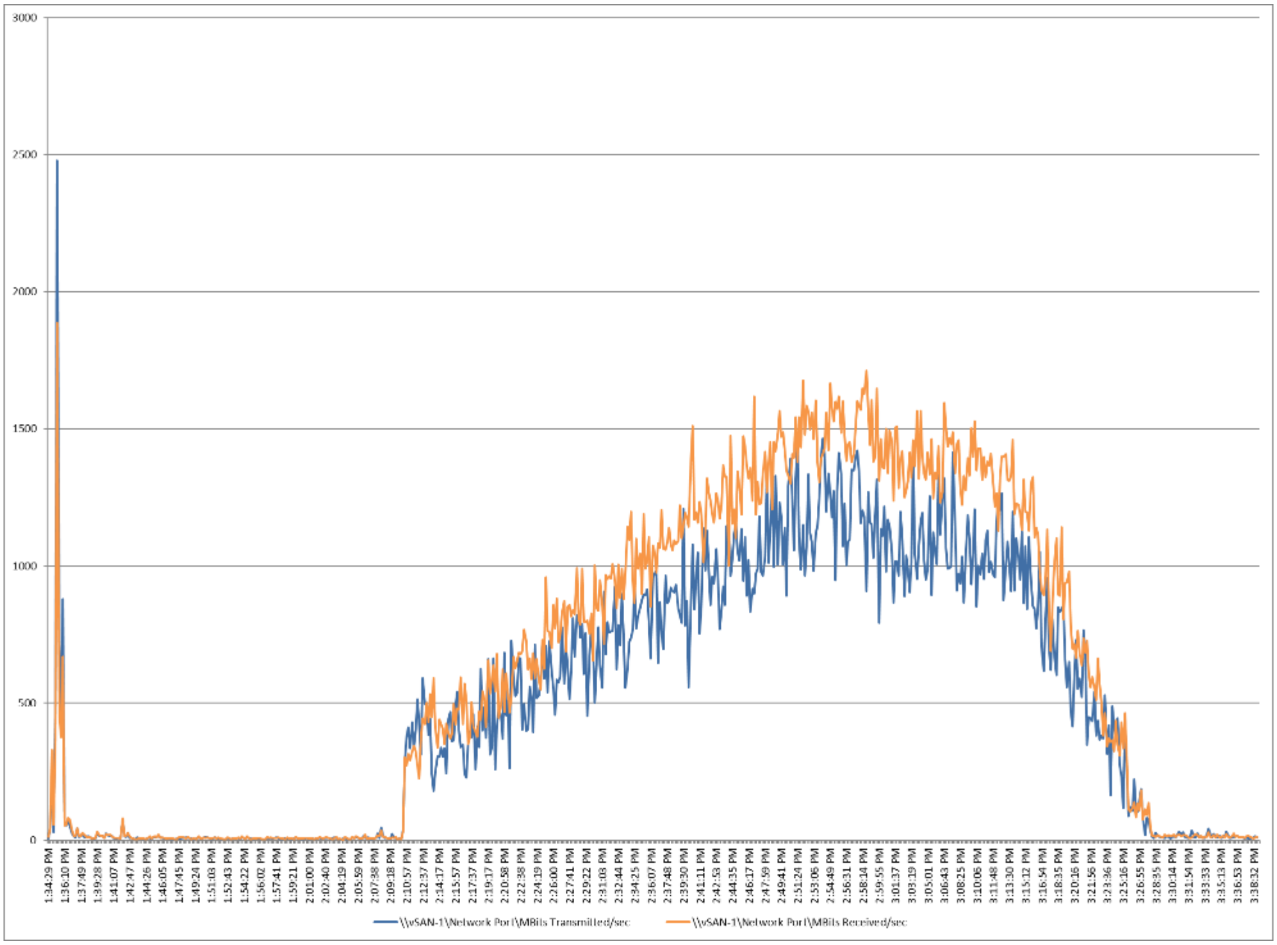

Figure 110. Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VMware vSAN cluster Latency Chart

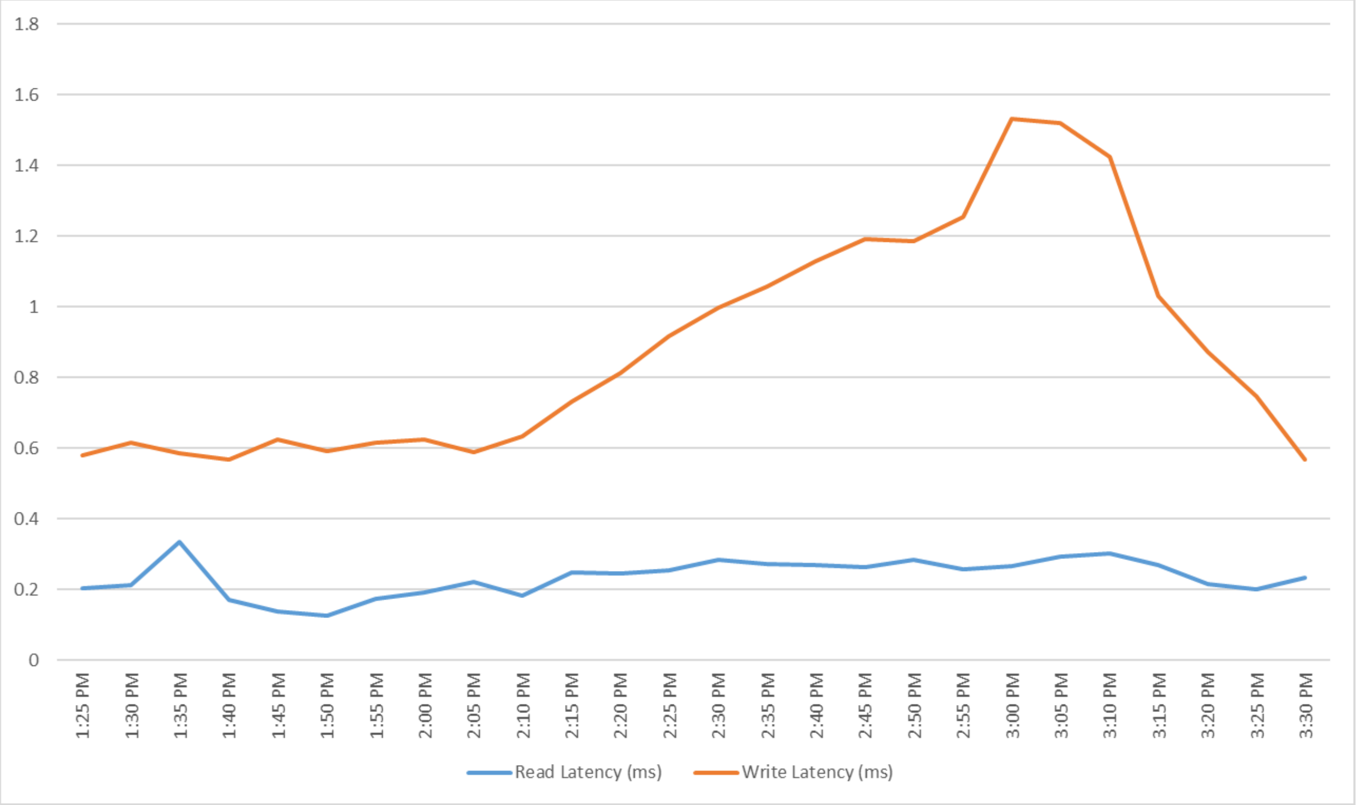

Figure 111. Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VMware vSAN cluster IOPS Chart

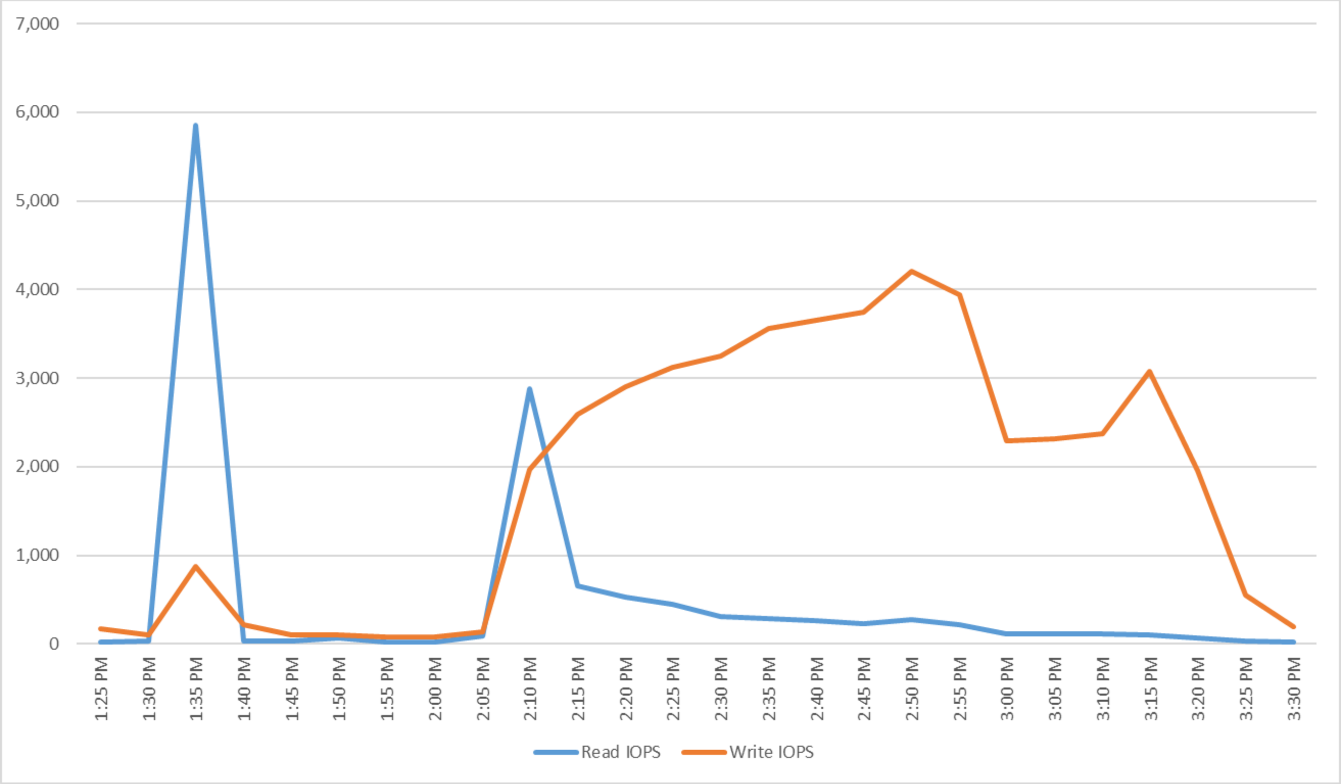

Figure 112. Full Scale | 960 Users | VMware Horizon 8 2212 Non-persistent Multi-session OS machine VDAs | VMware vSAN cluster Throughput Chart

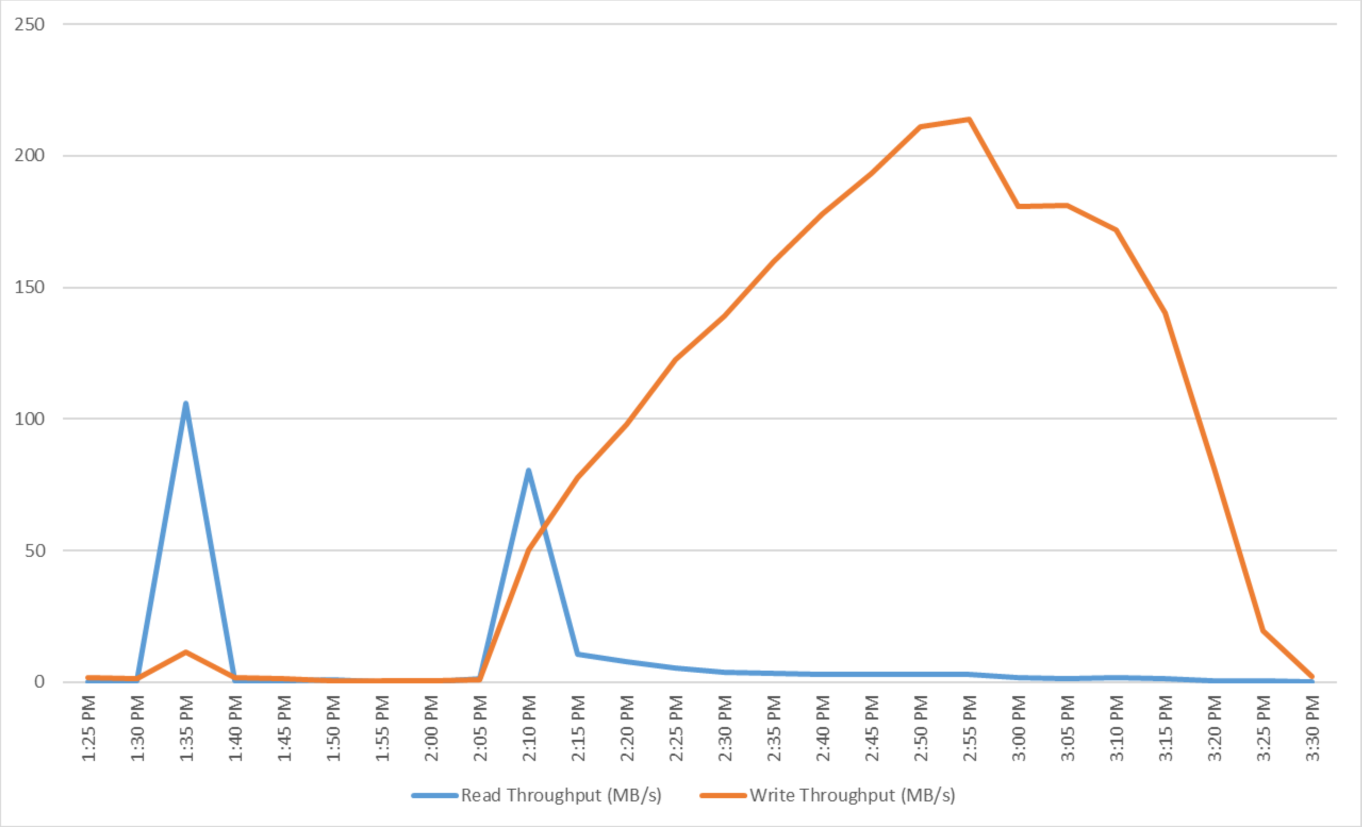

VMware vSAN Capacity

Figure 113. OSA VMware vSAN Capacity and optimization

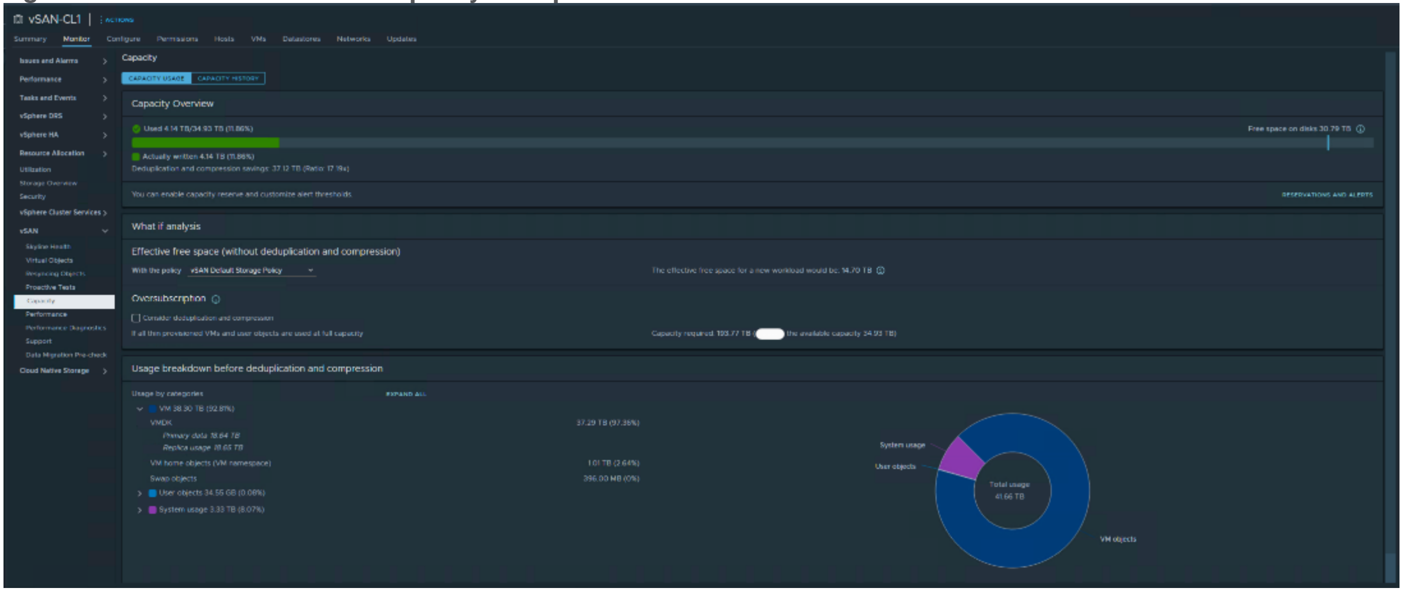

Capacity

Capacity Overview

Used 4.14 TB/34.93 TB (11.86%)

Free space on disks 30.79 TB

Actually written 4.14 TB (11.86%)

Deduplication and compression savings: 37.12 TB (Ratio: 17.19x)

You can enable capacity reserve and customize alert thresholds.

What if analysis

Effective free space (without deduplication and compression)

With the policy vSAN Default Storage Policy

The effective free space for a new workload would be: 14.70 TB

Oversubscription

□ Consider deduplication and compression

If all thin provisioned VMs and user objects are used at full capacity

Capacity required: 193.77 TB (the available capacity 34.93 TB)

Usage breakdown before deduplication and compression

Usage by categories

VM 38.30 TB (92.81%)

VMDK	37.29 TB (97.36%)
Primary data 18.64 TB	
Replica usage 18.65 TB	
VM home objects (VM namespace)	1.01 TB (2.64%)
Swap objects	396.00 MB (0%)

User objects 34.55 GB (0.08%)

System usage 3.33 TB (8.07%)

Figure 114. ESA VMware vSAN Capacity and optimization

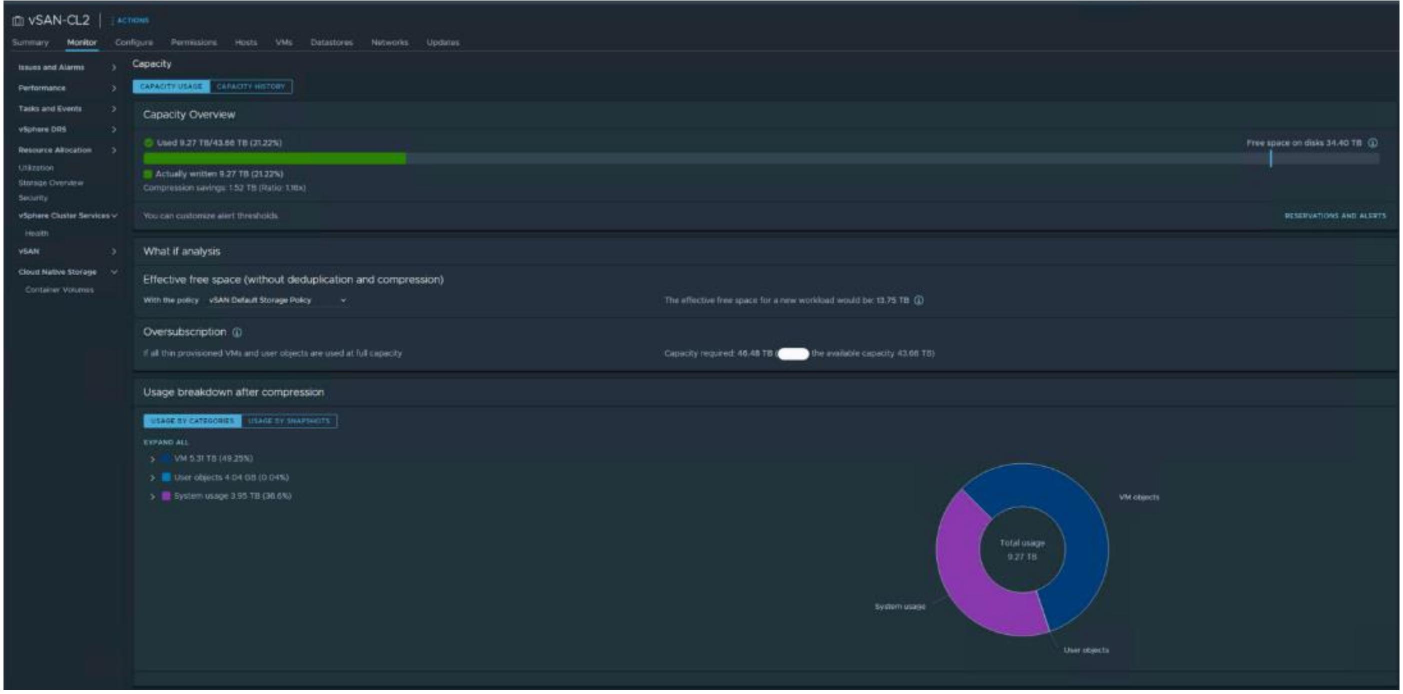

Capacity

Capacity Overview

Used 9.27 TB/43.66 TB (21.22%)

Free space on disks 34.40 TB

Actually written 9.27 TB (21.22%)

Compression savings: 1.52 TB (Ratio: 1.16x)

You can customize alert thresholds.

What if analysis

Effective free space (without deduplication and compression)

With the policy vSAN Default Storage Policy

The effective free space for a new workload would be: 13.75 TB

Oversubscription

If all thin provisioned VMs and user objects are used at full capacity

Capacity required: 46.48 TB ([redacted] the available capacity 43.66 TB)

Usage breakdown after compression

- VM 5.31 TB (49.25%)
- User objects 4.04 GB (0.04%)
- System usage 3.95 TB (36.6%)

Summary

The VMware vSAN platform, combined with the Cisco UCS X-Series featuring 4th Gen Xeon processors, offers a reliable and robust solution for enterprise end-user computing deployments. This hyper-converged infrastructure provides compute density, storage capacity, and expandability in a single system to support a wide range of workloads in your data center.

Deploying the Cisco UCS X-Series with VMware vSAN is made easy with the Cisco Intersight. These technologies simplify the deployment process, reducing project risk and IT costs, as well as enhancing visibility and orchestration across the entire data center, allowing you to modernize your infrastructure and operations effectively.

VMware vSAN on Cisco UCS X-Series with 4th Gen Xeon processors has undergone validation using industry-standard benchmarks, ensuring it meets the highest performance, management, scalability, and resilience standards. It makes an ideal choice for customers seeking enterprise-class hyper-converged infrastructure for virtual desktop infrastructure (VDI) deployments and allows you to focus on your core business objectives.

Get More Business Value with Services

Whether you are planning your next-generation environment, need specialized know-how for a major deployment, or want to get the most from your current storage, Cisco Advanced Services, and our certified partners can help. We collaborate with you to enhance your IT capabilities through a complete portfolio of services for your IT lifecycle with:

- Strategy services to align IT with your business goals
- Design services to architect your best storage environment
- Deploy and transition services to implement validated architectures and prepare your storage environment
- Operations services to deliver continuous operations while driving operational excellence and efficiency

Additionally, Cisco Advanced Services provide in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Author

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.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the following for their contribution and expertise that resulted in developing this document:

- Nagesh Dn, Intel Software Optimization Architect, Intel Inc.

Appendix

This appendix contains the following:

- Appendix A – Full Scale Server Performance Chart
- Appendix B – Switch Configuration
- Appendix C – References Used in Guide
- Appendix D – Parts list

Appendix A – Full Scale Server Performance Chart

This section provides a detailed performance chart for ESXi 8.0 Update 1a installed on Cisco UCS X210c M7 Server as part of the workload test with VMware Horizon 8 2212 deployed on VMware vSAN LoginVSI v4.1.40 based knowledge worker workload part of the VMware HCI reference architecture defined here.

The charts below are defined in the set of four hosts in the single performance chart.

OSA vSAN Performance Monitor Data for One Sample Test: 585 Users Non-persistent Single-session OS machine VDAs Scale Testing.

Figure 115. Full Scale | 585 User| Non-persistent Single-session OS machine VDAs | Host CPU Utilization

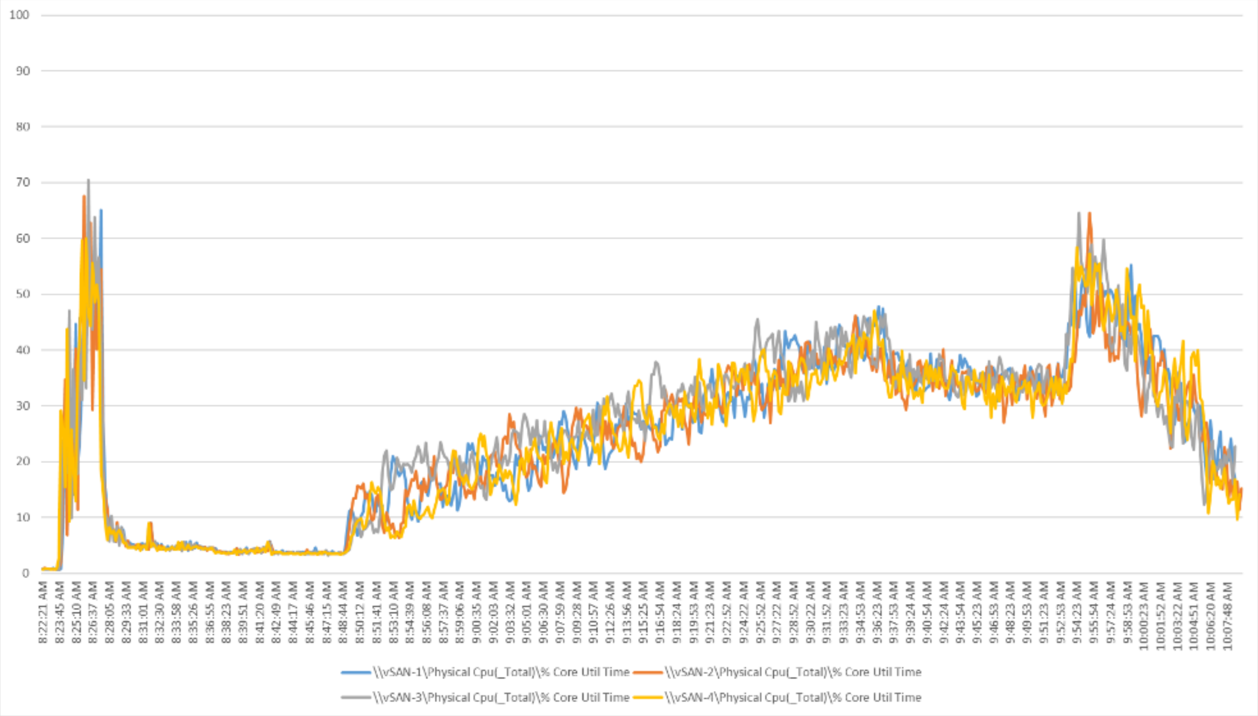

Figure 116. Full Scale | 585 User| Non-persistent Single-session OS machine VDAs | Host Memory Utilization

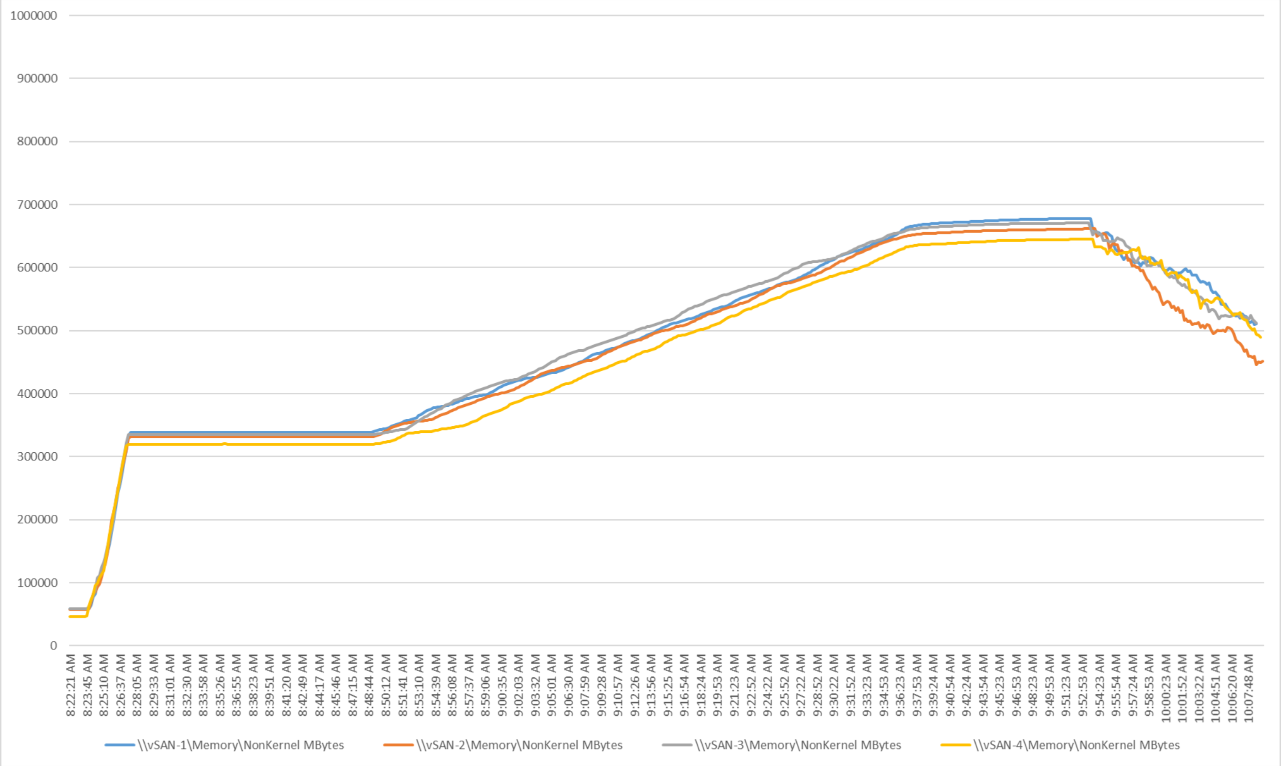

Figure 117. Full Scale | 585 User| Non-persistent Single-session OS machine VDAs | Host Network Utilization

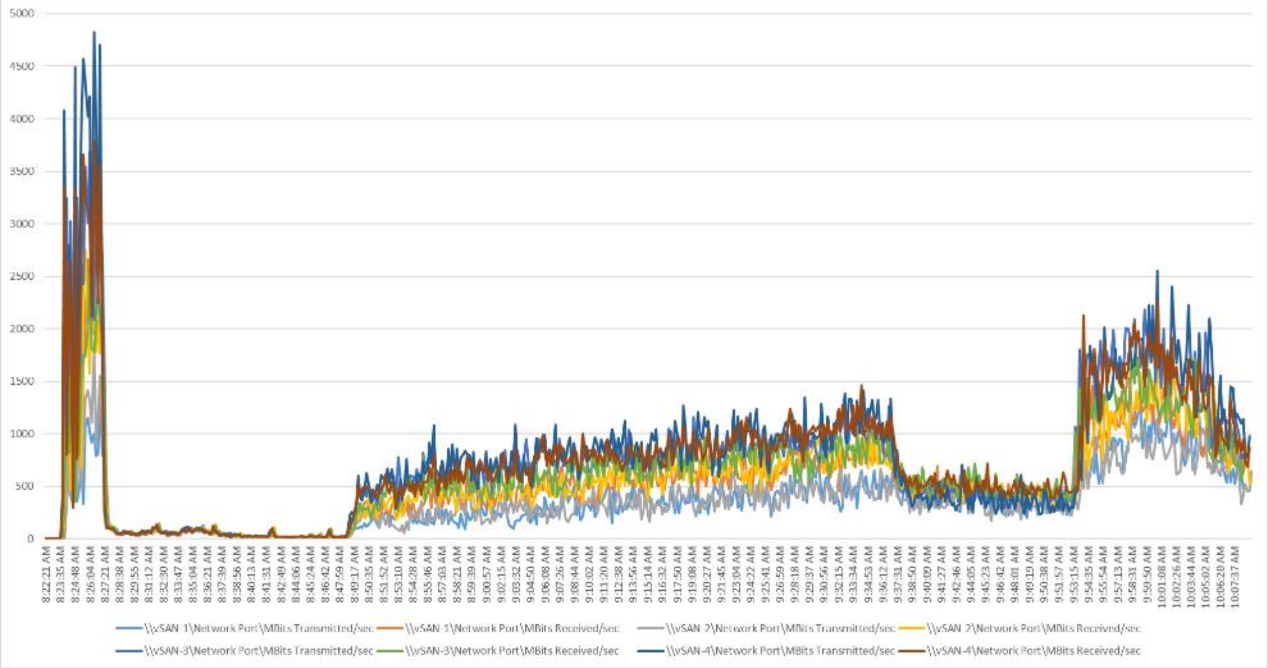

OSA vSAN Performance Monitor Data for One Sample Test: 585 Users Persistent Single-session OS machine VDAs Scale Testing.

Figure 118. Full Scale | 585 User| Persistent Single-session OS machine VDAs | Host CPU Utilization

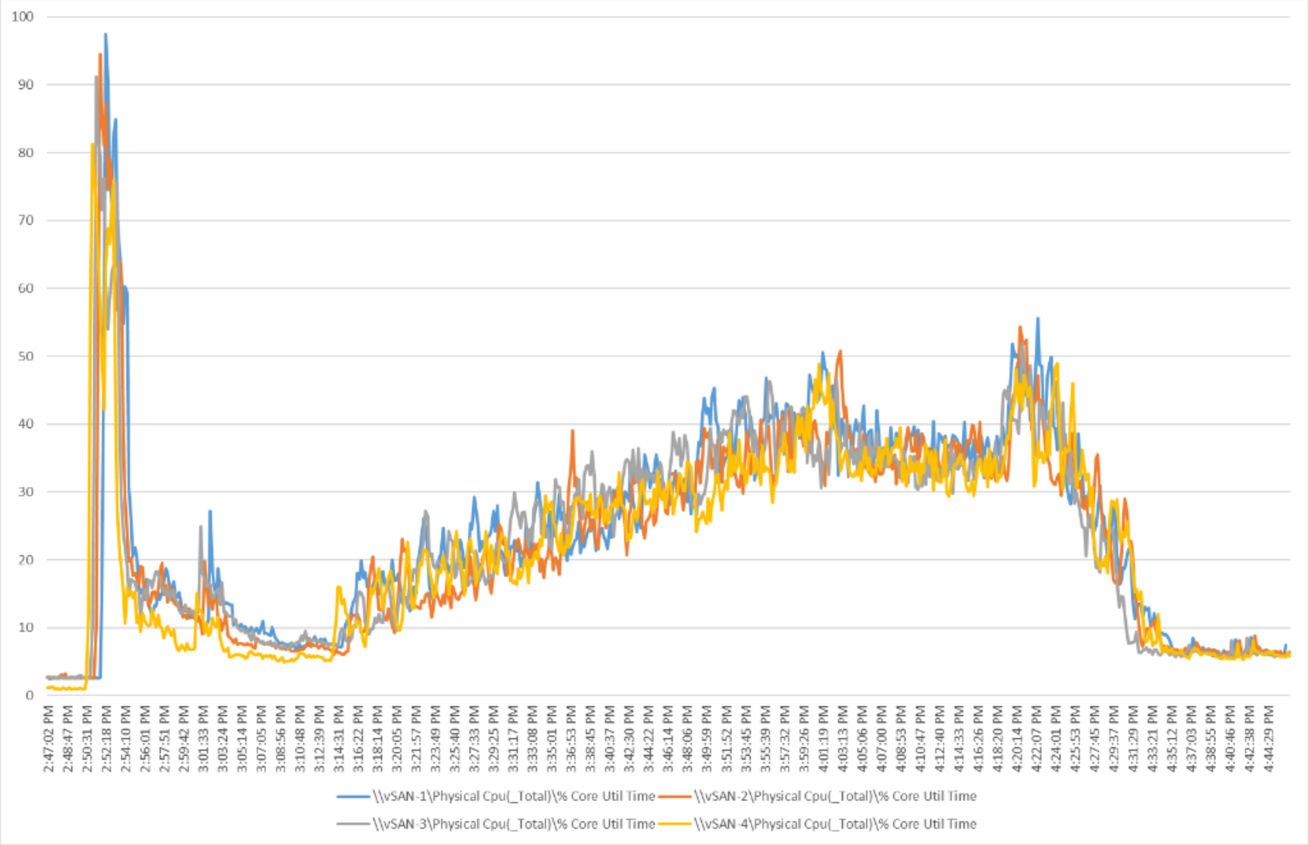

Figure 119. Full Scale | 585 User| Persistent Single-session OS machine VDAs | Host Memory Utilization

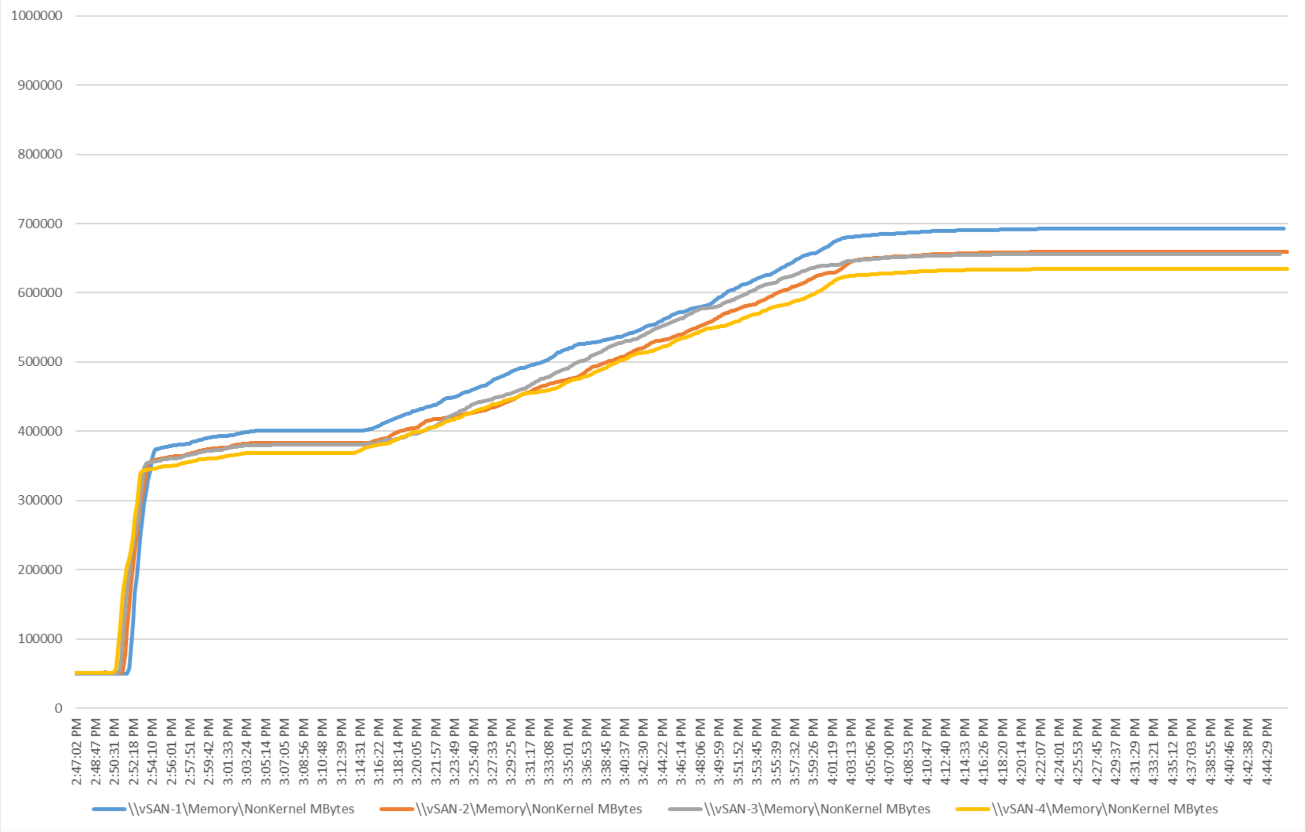

Figure 120. Full Scale | 585 User| Persistent Single-session OS machine VDAs | Host Network Utilization

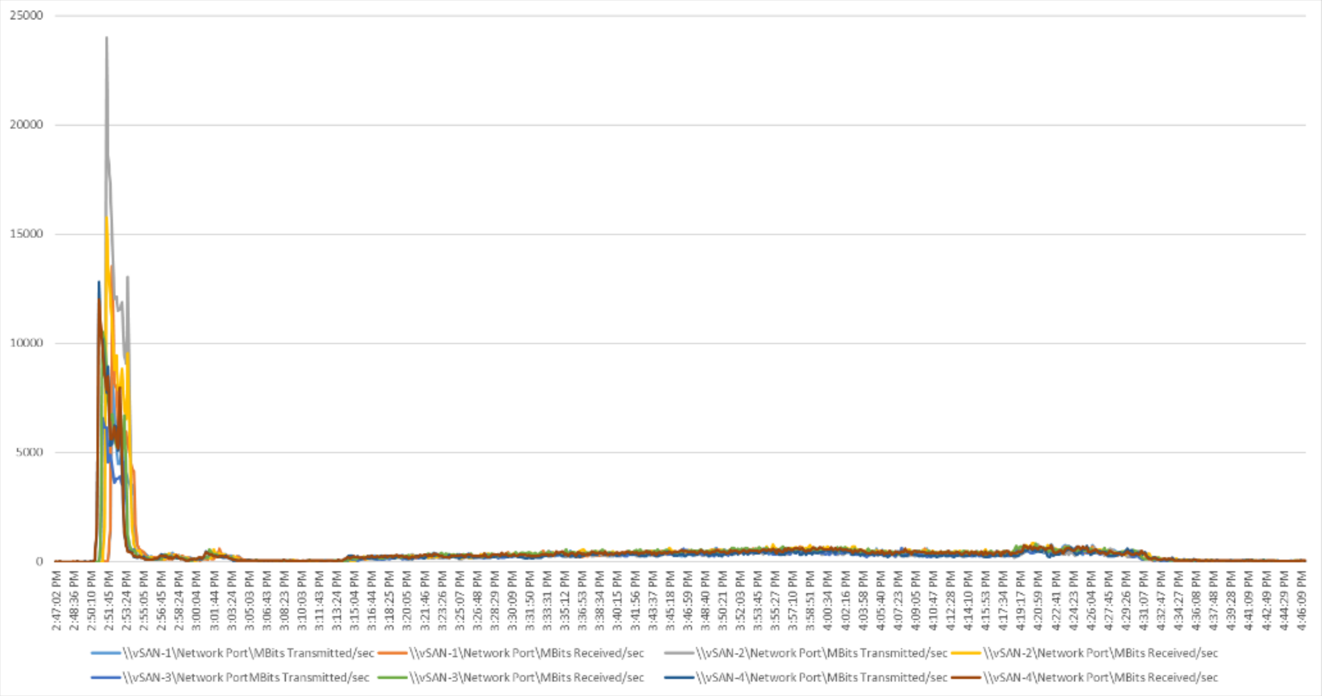

OSA vSAN Performance Monitor Data for One Sample Test: 2688 Users Non-persistent Multi-session OS machine VDAs Scale Testing.

Figure 121. Full Scale | 1440 Users| Multi-session OS machine VDAs | Host CPU Utilization

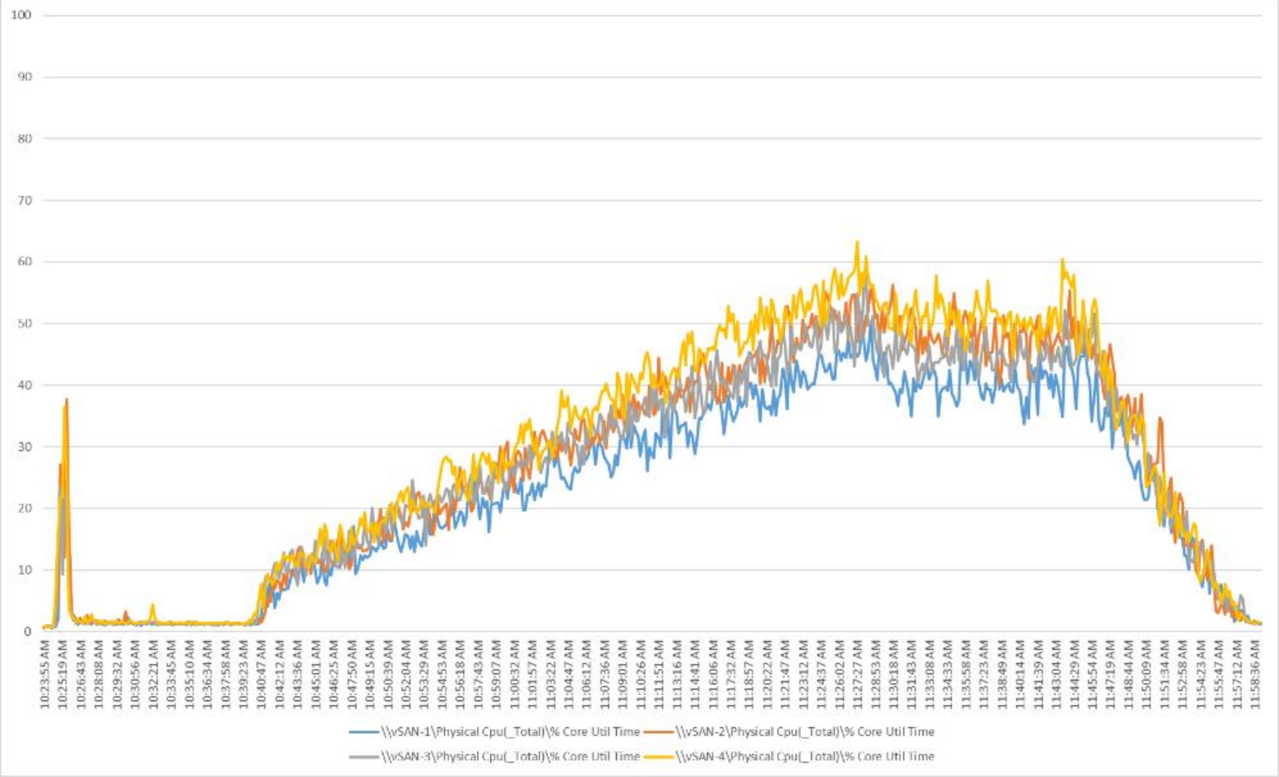

Figure 122. Full Scale | 1440 Users| Multi-session OS machine VDAs | Host Memory Utilization

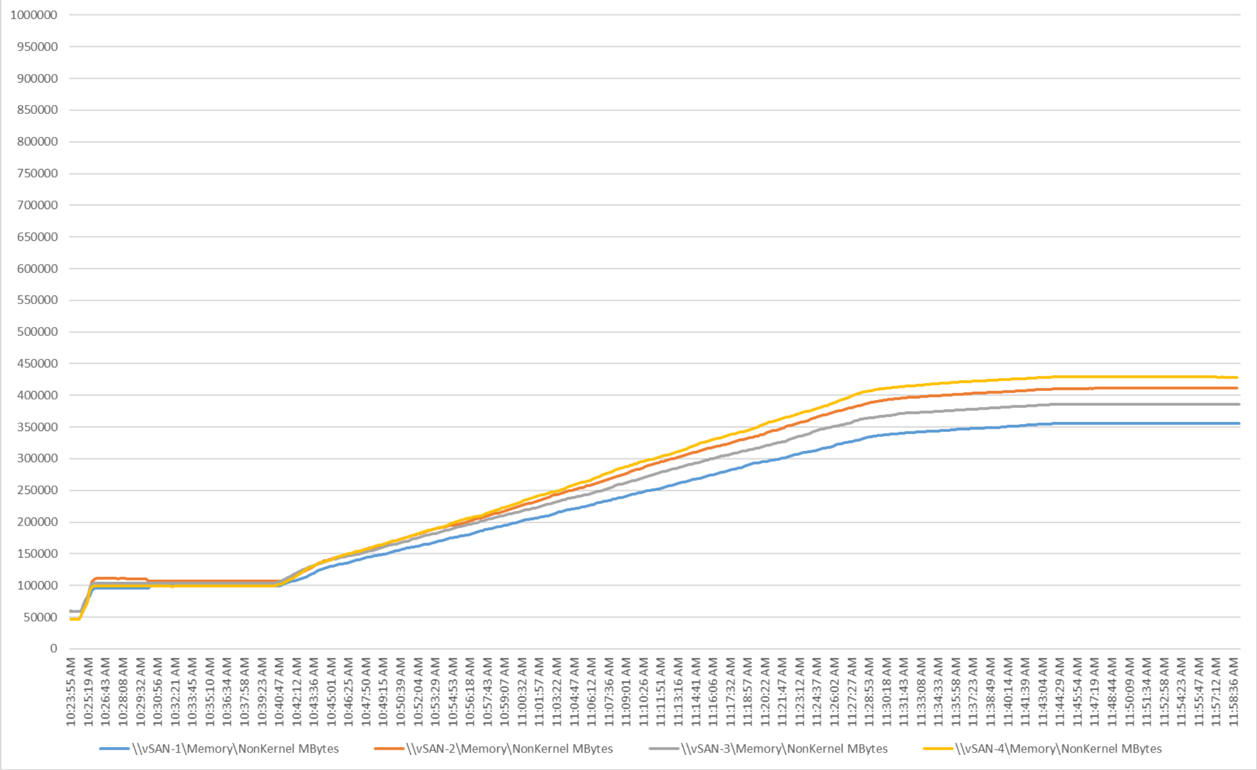

Figure 123. Full Scale | 1440 Users| Multi-session OS machine VDAs | Host Network Utilization

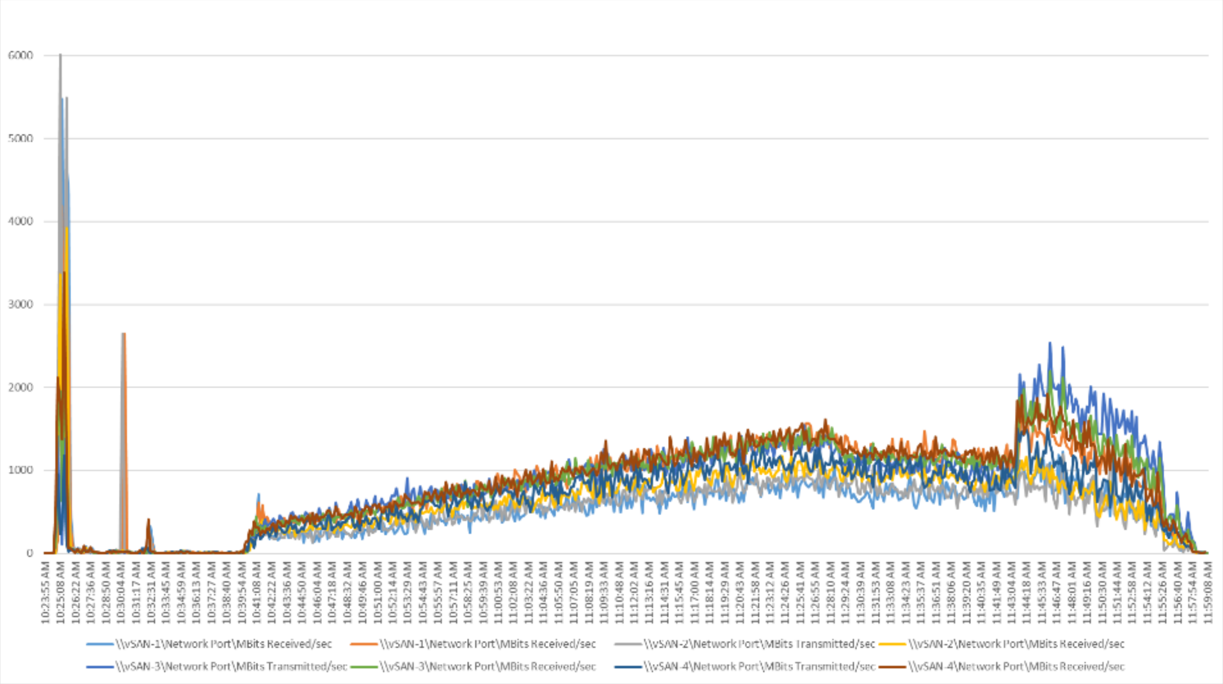

ESA vSAN Performance Monitor Data for One Sample Test: 390 Users Non-persistent Single-session OS machine VDAs Scale Testing.

Figure 124. Full Scale | 390 User| Non-persistent Single-session OS machine VDAs | Host CPU Utilization

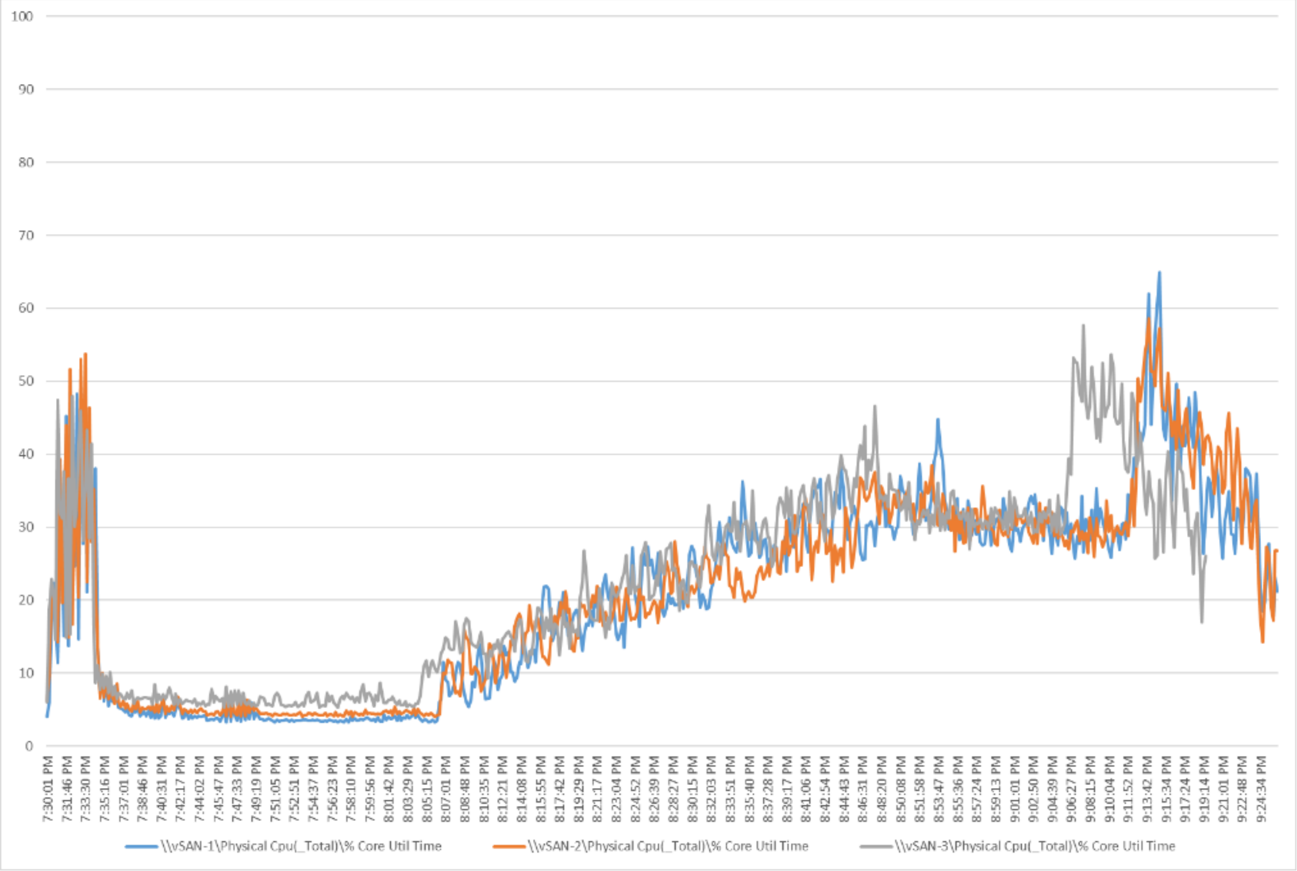

Figure 125. Full Scale | 390 User| Non-persistent Single-session OS machine VDAs | Host Memory Utilization

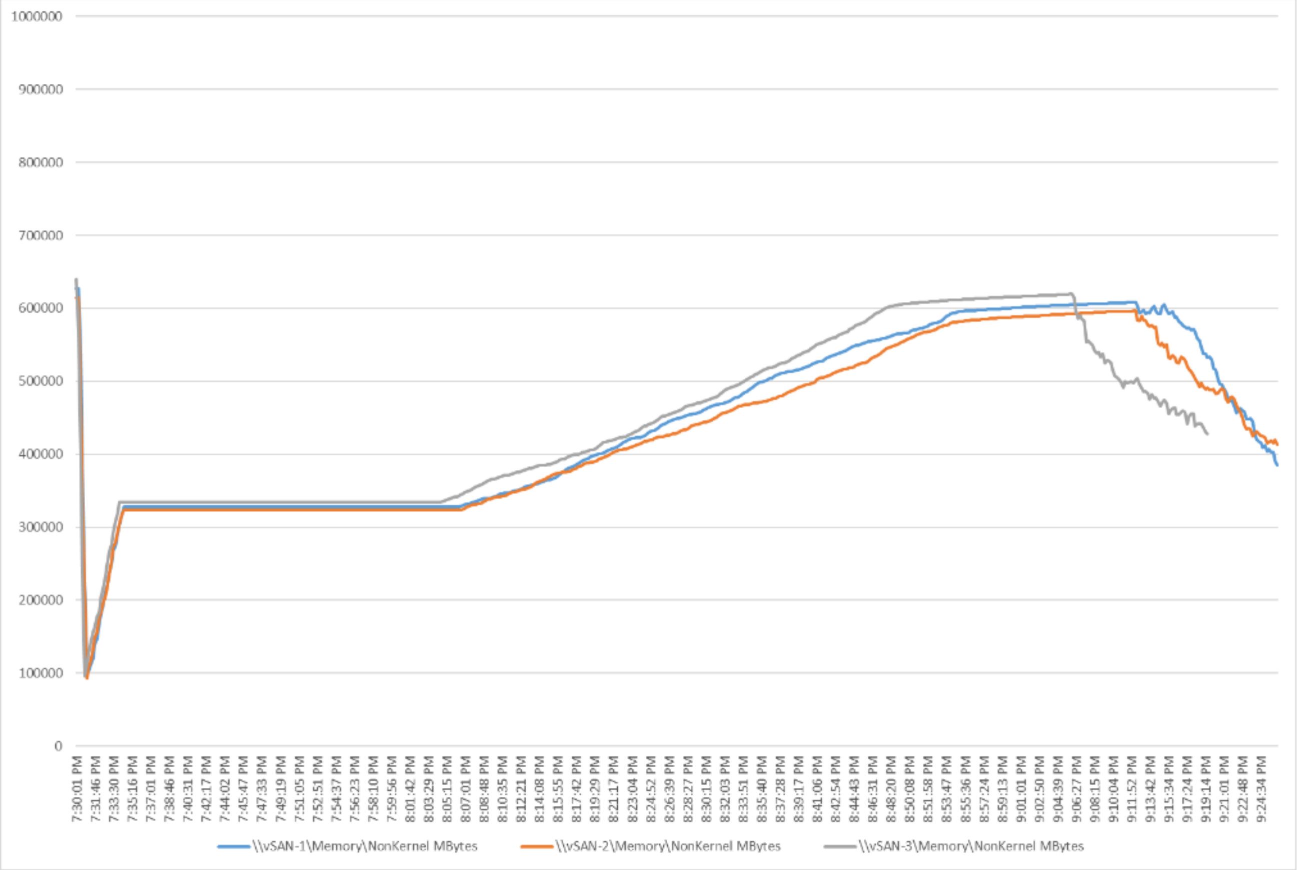

Figure 126. Full Scale | 390 User| Non-persistent Single-session OS machine VDAs | Host Network Utilization

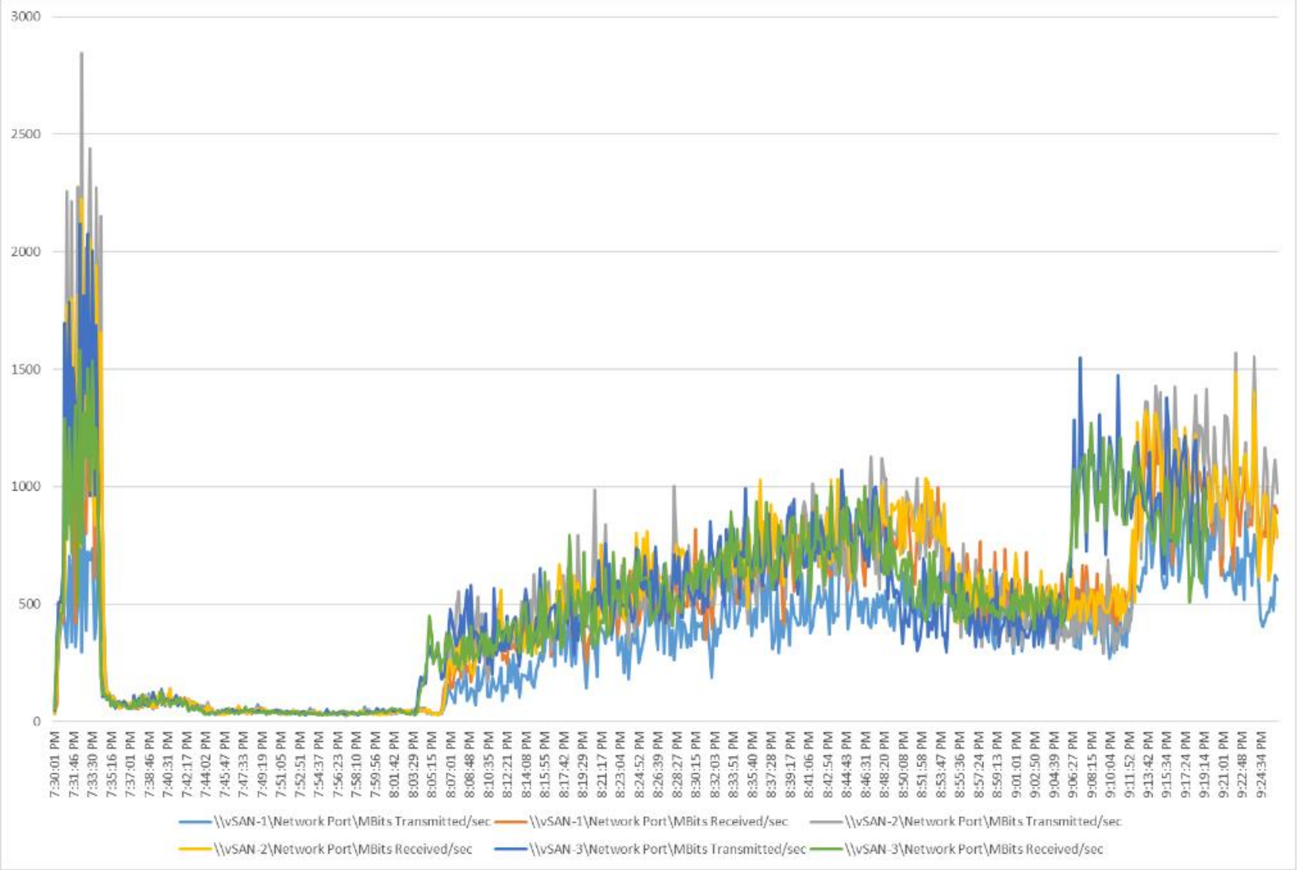

ESA vSAN Performance Monitor Data for One Sample Test: 390 Users Persistent Single-session OS machine VDAs Scale Testing.

Figure 127. Full Scale | 390 User| Persistent Single-session OS machine VDAs | Host CPU Utilization

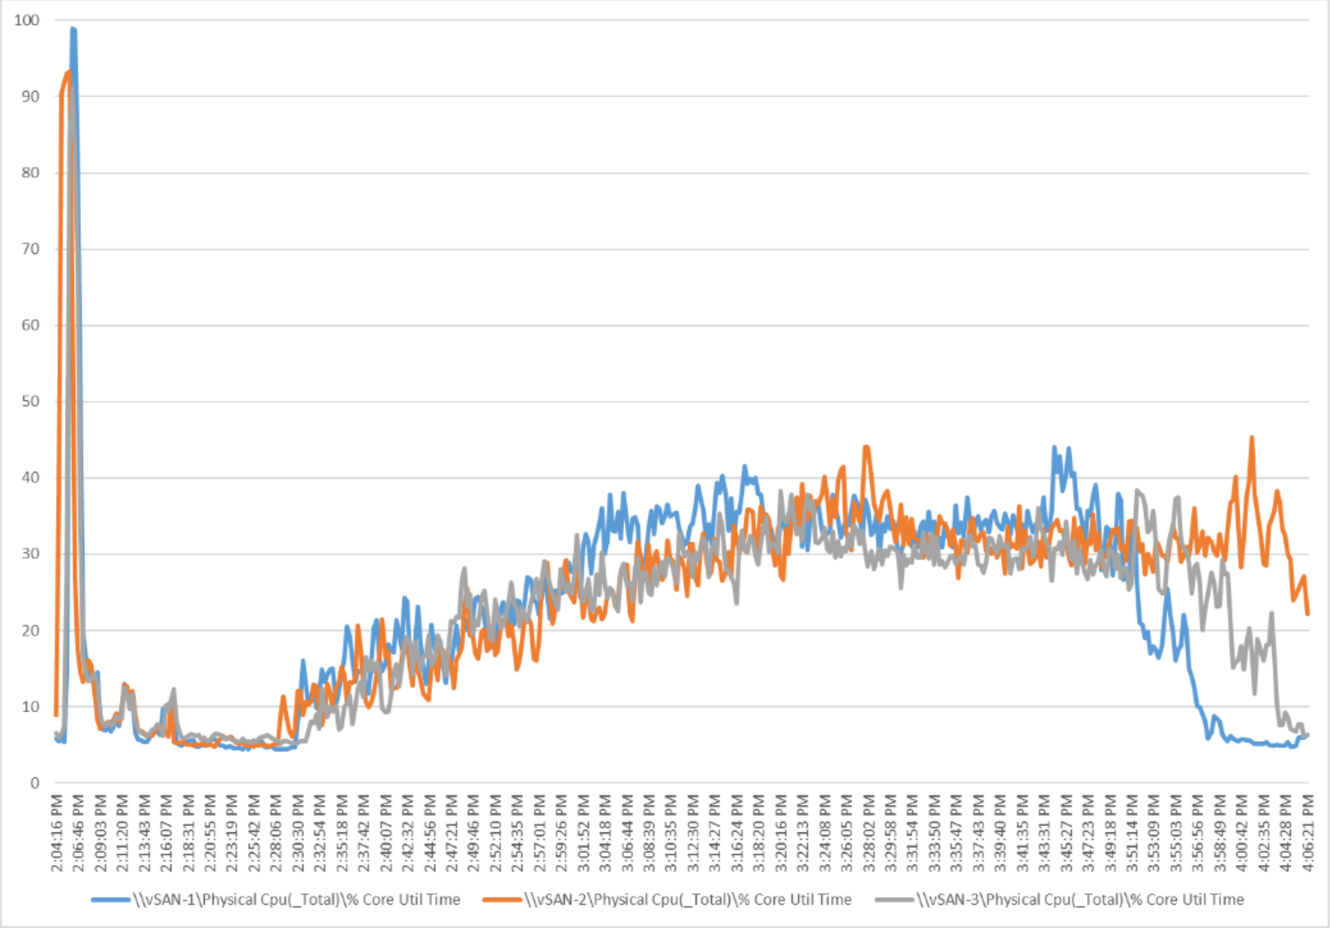

Figure 128. Full Scale | 390 User| Persistent Single-session OS machine VDAs | Host Memory Utilization

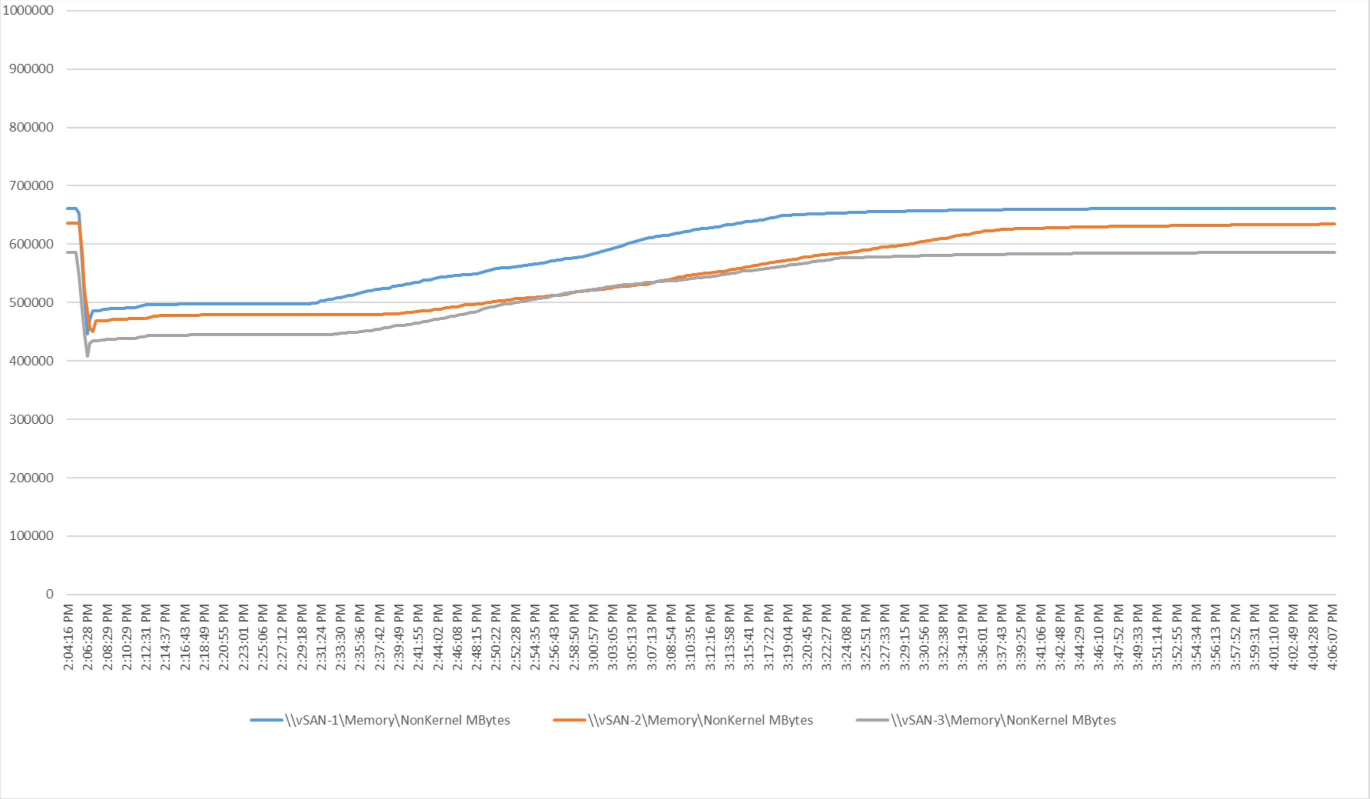

Figure 129. Full Scale | 390 User| Persistent Single-session OS machine VDAs | Host Network Utilization

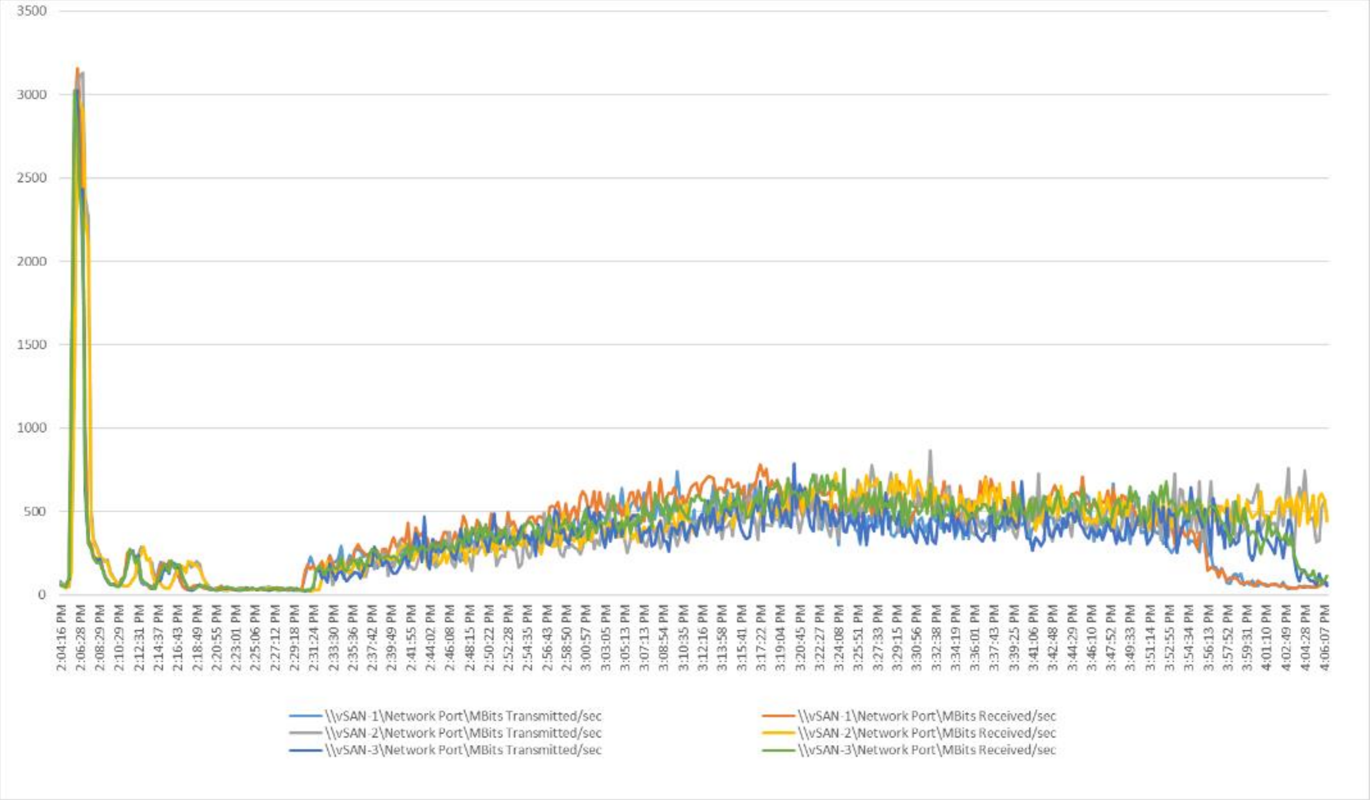

ESA vSAN Performance Monitor Data for One Sample Test: 960 Users Non-persistent Multi-session OS machine VDAs Scale Testing.

Figure 130. Full Scale | 960 Users| Multi-session OS machine VDAs | Host CPU Utilization

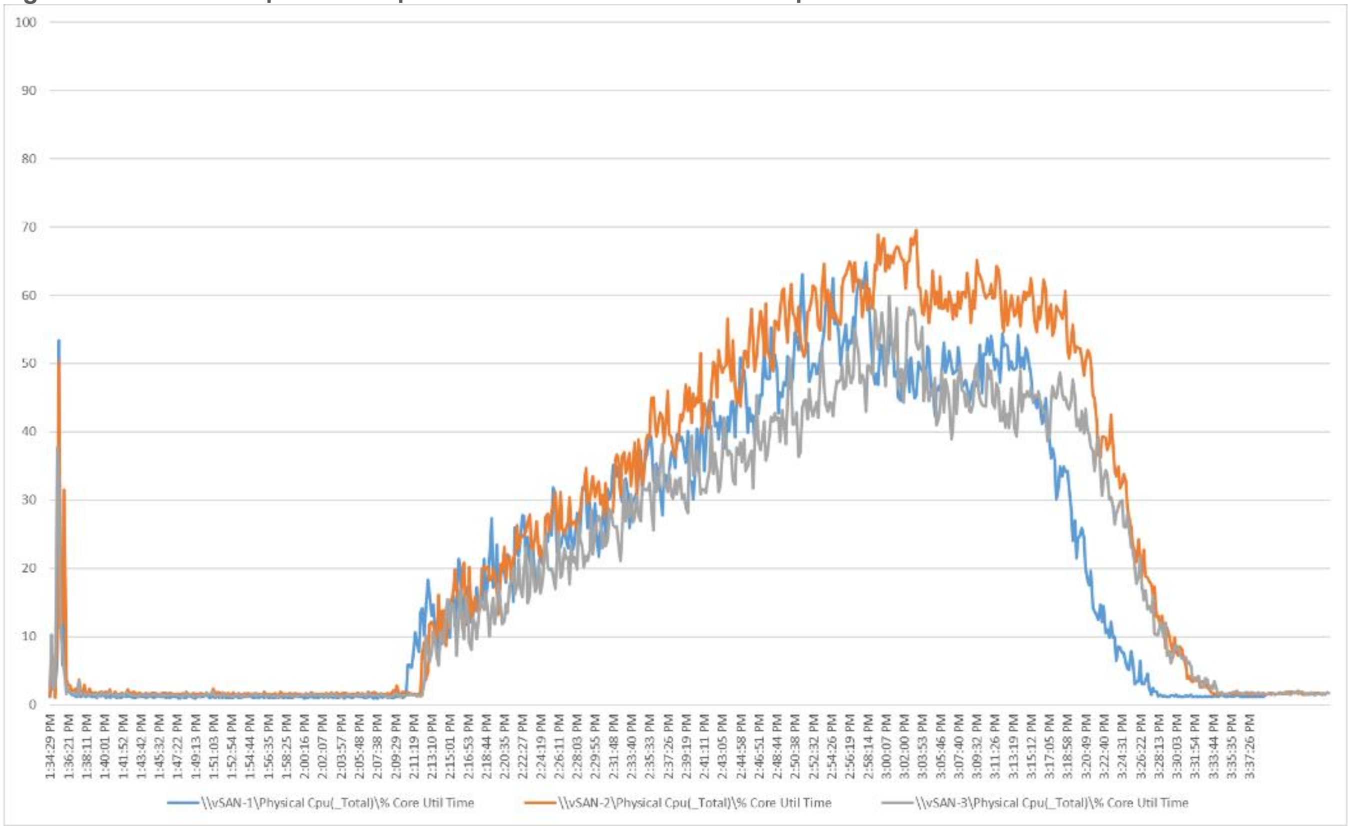

Figure 131. Full Scale | 960 Users| Multi-session OS machine VDAs | Host Memory Utilization

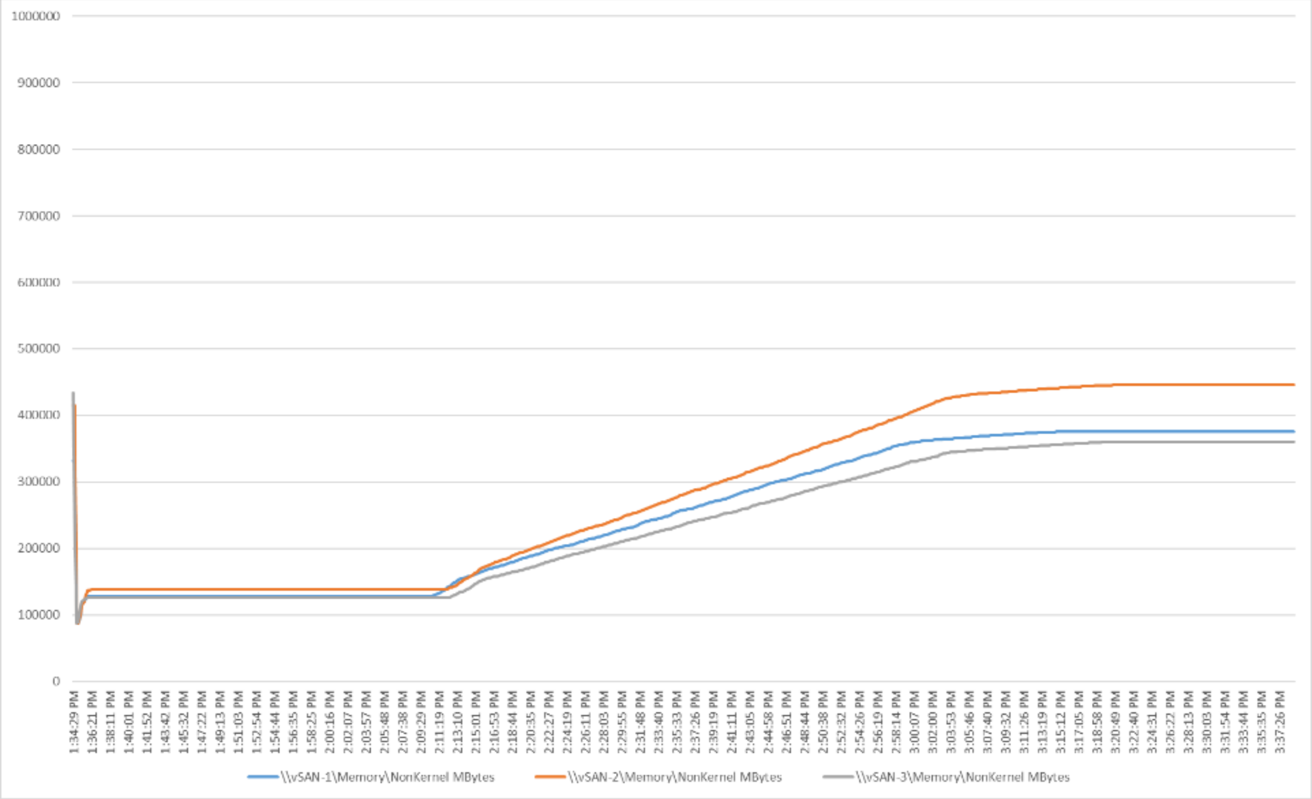

Figure 132. Full Scale | 960 Users| Multi-session OS machine VDAs | Host Network Utilization

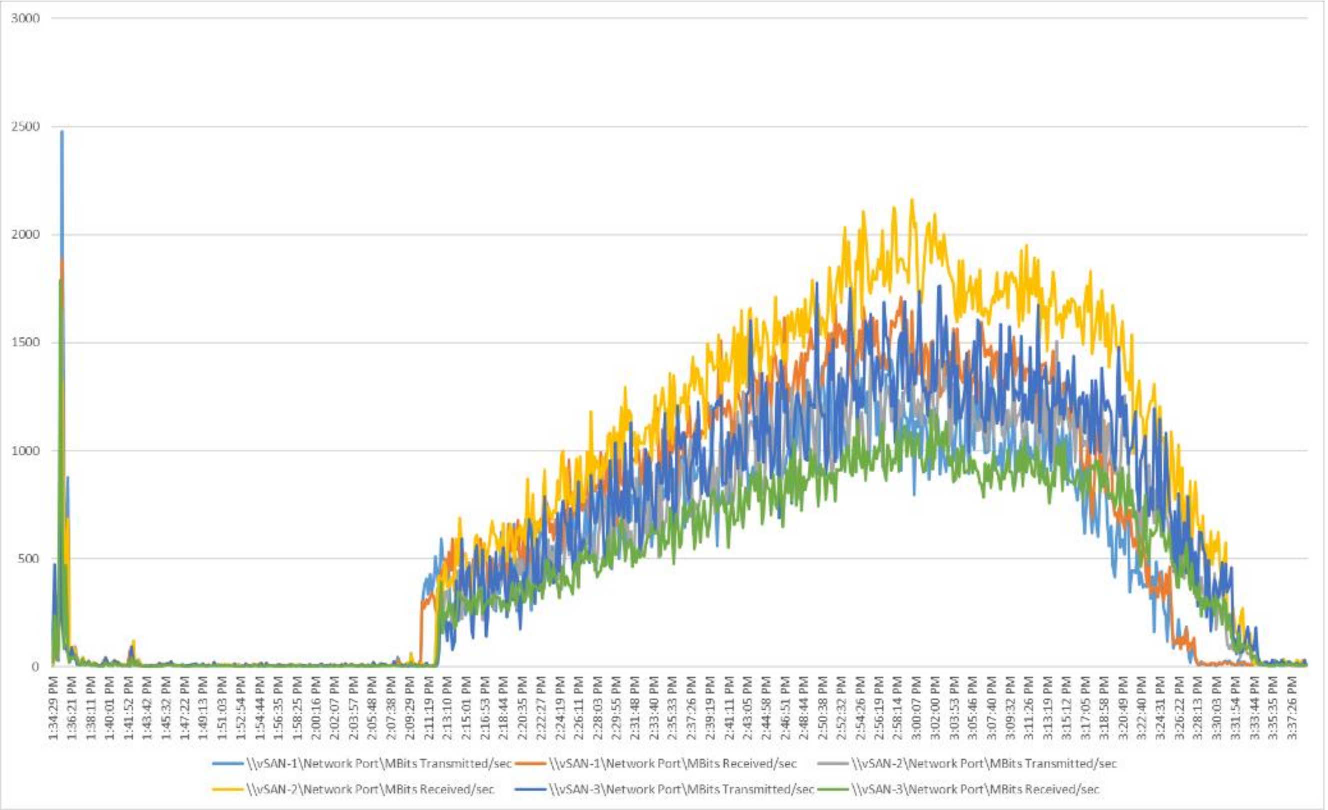

Appendix B – Switch Configurations

Cisco Nexus 93180YC-A Configuration

```
version 9.3(3) Bios:version 05.39
switchname K23-N9K-A
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
vdc K23-N9K-A id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource M7route-mem minimum 8 maximum 8

feature telnet
feature nxapi
feature bash-shell
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature telemetry

no password strength-check
username admin password 5 $5$0BAB7aa4$v07pyr7xw1f5WpD2wZc3qmG3Flb04Wa62aNgxg82hUA  role
network-admin
ip domain-lookup
system default switchport
ip access-list acl1
  10 permit ip 10.10.71.0/24 any
ip access-list acl_oob
  10 permit ip 10.10.71.0/24 any
system qos
  service-policy type network-qos jumbo
copp profile lenient
snmp-server user admin network-admin auth md5 0x83fa863523d7d94fe06388d7669f62f5 priv
0x83fa863523d7d94fe06388d7669f62f5 localizedkey
snmp-server host 173.37.52.102 traps version 2c public udp-port 1163
snmp-server host 192.168.24.30 traps version 2c public udp-port 1163
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 10.10.50.252 use-vrf default
ntp peer 10.10.50.253 use-vrf default
ntp server 171.68.38.65 use-vrf default
ntp logging
ntp master 8

vlan 1,50-56,70-76
vlan 50
  name Inband-Mgmt-C1
vlan 51
  name Infra-Mgmt-C1
vlan 52
  name StorageIP-C1
vlan 53
  name vMotion-C1
vlan 54
  name VM-Data-C1
vlan 55
  name Launcher-C1
vlan 56
  name Launcher-Mgmt-C1
vlan 70
  name InBand-Mgmt-SP
vlan 71
  name Infra-Mgmt-SP
vlan 72
  name VM-Network-SP
vlan 73
  name vMotion-SP
vlan 74
  name Storage_A-SP
vlan 75
  name Storage_B-SP
vlan 76
  name Launcher-SP

service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
```

```
  ip route 0.0.0.0/0 173.37.52.1
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region nat 256
vpc domain 50
  role priority 10
  peer-keepalive destination 173.37.52.104 source 173.37.52.103
  delay restore 150
  auto-recovery


interface Vlan1
  no shutdown

interface Vlan50
  no shutdown
  ip address 10.10.50.252/24
  hsrp version 2
  hsrp 50
    preempt
    priority 110
    ip 10.10.50.1

interface Vlan51
  no shutdown
  ip address 10.10.51.252/24
  hsrp version 2
  hsrp 51
    preempt
    priority 110
    ip 10.10.51.1

interface Vlan52
  no shutdown
  ip address 10.10.52.2/24
  hsrp version 2
  hsrp 52
    preempt
    priority 110
    ip 10.10.52.1

interface Vlan53
  no shutdown
  ip address 10.10.53.2/24
  hsrp version 2
  hsrp 53
```

```
    preempt
    priority 110
    ip 10.10.53.1

interface Vlan54
  no shutdown
  ip address 10.54.0.2/19
  hsrp version 2
  hsrp 54
    preempt
    priority 110
    ip 10.54.0.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface Vlan55
  no shutdown
  ip address 10.10.55.2/23
  hsrp version 2
  hsrp 55
    preempt
    priority 110
    ip 10.10.55.1
  ip dhcp relay address 10.10.51.11
  ip dhcp relay address 10.10.51.12

interface Vlan56
  no shutdown
  ip address 10.10.56.2/24
  hsrp version 2
  hsrp 56
    preempt
    ip 10.10.56.1
  ip dhcp relay address 10.10.51.11
  ip dhcp relay address 10.10.51.12

interface Vlan70
  no shutdown
  ip address 10.10.70.2/24
  hsrp version 2
  hsrp 70
    preempt
    priority 110
    ip 10.10.70.1
```

```
interface Vlan71
  no shutdown
  ip address 10.10.71.2/24
  hsrp version 2
  hsrp 71
    preempt
    priority 110
    ip 10.10.71.1

interface Vlan72
  no shutdown
  ip address 10.72.0.2/19
  hsrp version 2
  hsrp 72
    preempt
    priority 110
    ip 10.72.0.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface Vlan73
  no shutdown
  ip address 10.10.73.2/24
  hsrp version 2
  hsrp 73
    preempt
    priority 110
    ip 10.10.73.1

interface Vlan74
  no shutdown
  ip address 10.10.74.2/24
  hsrp version 2
  hsrp 74
    preempt
    priority 110
    ip 10.10.74.1

interface Vlan75
  no shutdown
  ip address 10.10.75.2/24
  hsrp version 2
  hsrp 75
    preempt
    priority 110
    ip 10.10.75.1

interface Vlan76
  no shutdown
  ip address 10.10.76.2/23
  hsrp version 2
  hsrp 76
    preempt
    priority 110
    ip 10.10.76.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface port-channel10
  description VPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FI-Uplink-K22-B
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel12
  description FI-Uplink-K22-B
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

interface port-channel49
  description FI-Uplink-K23
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 49

interface port-channel50
```

```
  description FI-Uplink-K23
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 50

interface Ethernet1/1
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/2
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/3
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/4
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11
interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33
  switchport access vlan 71
  spanning-tree port type edge
```

```
interface Ethernet1/34
  switchport access vlan 71
  spanning-tree port type edge

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45
  description VLAN 30 access JH
  switchport access vlan 30
  switchport trunk allowed vlan 1,30-36,60-68,132
  speed 1000

interface Ethernet1/46

interface Ethernet1/47
  switchport access vlan 50
  spanning-tree port type edge

interface Ethernet1/48

interface Ethernet1/49
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  mtu 9216
  channel-group 49 mode active
```

```
interface Ethernet1/50
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  mtu 9216
  channel-group 50 mode active

interface Ethernet1/51
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/52
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1,30-36,50-56,60-68,70-76,132

interface Ethernet1/54

interface mgmt0
  vrf member management
  ip address 173.37.52.103/23
line console
line vty
boot nxos bootflash:/nxos.9.3.3.bin
no system default switchport shutdown


telemetry
  certificate /bootflash/home/admin/telemetry-cert.pem localhost
  destination-profile
    use-nodeid timba-640142e96f72612d3459249f
  destination-group timba-640142e96f72612d3459249f-0
    ip address 10.10.71.20 port 443 protocol HTTP encoding JSON
  sensor-group timba-640142e96f72612d3459249f-0
    data-source NX-API
    path "show system resources all-modules" depth 0
  sensor-group timba-640142e96f72612d3459249f-1
```

```
    data-source NX-API
    path "show module" depth 0
  sensor-group timba-640142e96f72612d3459249f-2
    data-source NX-API
    path "show environment power" depth 0
  sensor-group timba-640142e96f72612d3459249f-3
    data-source NX-API
    path "show interface fc regex *" depth 0
  sensor-group timba-640142e96f72612d3459249f-4
    data-source DME
    path sys/ch depth 1 query-condition query-target=subtree&target-subtree-
class=eqptSensor
  sensor-group timba-640142e96f72612d3459249f-5
    data-source DME
    path sys/ch query-condition query-target=subtree&target-subtree-class=eqptSupC
  sensor-group timba-640142e96f72612d3459249f-6
    data-source DME
    path sys/ch query-condition query-target=subtree&target-subtree-class=eqptFt
  sensor-group timba-640142e96f72612d3459249f-7
    data-source DME
    path sys/intf depth 0 query-condition query-target=subtree&target-subtree-
class=ethpmPhysIf filter-condition updated(ethpmPhysIf.operSt)
  subscription 2643
    dst-grp timba-640142e96f72612d3459249f-0
    snsr-grp timba-640142e96f72612d3459249f-0 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-1 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-2 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-3 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-4 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-5 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-6 sample-interval 300000
    snsr-grp timba-640142e96f72612d3459249f-7 sample-interval 0
```

Cisco Nexus 93180YC -B Configuration

```
version 9.3(3) Bios:version 05.39
switchname K23-N9K-B
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
vdc K23-N9K-B id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource M7route-mem minimum 8 maximum 8

feature telnet
feature nxapi
feature bash-shell
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature telemetry

no password strength-check
username admin password 5 $5$5TxyL6Rl$7U4nS.UfzkPgXl5mVqiuHoPLHyAZgnNAiKyz7aEVK05  role
network-admin
ip domain-lookup
system default switchport
system qos
  service-policy type network-qos jumbo
copp profile lenient
snmp-server user admin network-admin auth md5 0x57cdc0fb04a0dd922046cb694508c9b7 priv
0x57cdc0fb04a0dd922046cb694508c9b7 localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 171.68.38.65 use-vrf default

vlan 1,50-56,70-76,132
vlan 50
  name Inband-Mgmt-C1
vlan 51
  name Infra-Mgmt-C1
vlan 52
  name StorageIP-C1
vlan 53
  name vMotion-C1
vlan 54
  name VM-Data-C1
vlan 55
  name Launcher-C1
vlan 56
  name Launcher-Mgmt-C1
```

```
vlan 70
  name InBand-Mgmt-SP
vlan 71
  name Infra-Mgmt-SP
vlan 72
  name VM-Network-SP
vlan 73
  name vMotion-SP
vlan 74
  name Storage_A-SP
vlan 75
  name Storage_B-SP
vlan 76
  name Launcher-SP
vlan 132
  name OOB-Mgmt

service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 173.37.52.1
vpc domain 50
  role priority 10
  peer-keepalive destination 173.37.52.103 source 173.37.52.104
  delay restore 150
  auto-recovery


interface Vlan1
  no shutdown

interface Vlan50
  no shutdown
  ip address 10.10.50.253/24
  hsrp version 2
  hsrp 50
    preempt
    priority 110
    ip 10.10.50.1

interface Vlan51
  no shutdown
  ip address 10.10.51.253/24
  hsrp version 2
  hsrp 51
    preempt
    priority 110
    ip 10.10.51.1

interface Vlan52
  no shutdown
  ip address 10.10.52.3/24
  hsrp version 2
  hsrp 52
    preempt
    priority 110
    ip 10.10.52.1

interface Vlan53
  no shutdown
  ip address 10.10.53.3/24
  hsrp version 2
  hsrp 53
    preempt
    priority 110
    ip 10.10.53.1

interface Vlan54
  no shutdown
  ip address 10.54.0.3/19
  hsrp version 2
  hsrp 54
    preempt
    priority 110
    ip 10.54.0.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface Vlan55
  no shutdown
  ip address 10.10.55.3/23
  hsrp version 2
  hsrp 55
    preempt
    priority 110
    ip 10.10.55.1
  ip dhcp relay address 10.10.51.11
  ip dhcp relay address 10.10.51.12
```

```
interface Vlan56
  no shutdown
  ip address 10.10.56.3/24
  hsrp version 2
  hsrp 56
    preempt
    ip 10.10.56.1
  ip dhcp relay address 10.10.51.11
  ip dhcp relay address 10.10.51.12

interface Vlan70
  no shutdown
  ip address 10.10.70.3/24
  hsrp version 2
  hsrp 70
    preempt
    priority 110
    ip 10.10.70.1

interface Vlan71
  no shutdown
  ip address 10.10.71.3/24
  hsrp version 2
  hsrp 71
    preempt
    priority 110
    ip 10.10.71.1

interface Vlan72
  no shutdown
  ip address 10.72.0.3/19
  hsrp version 2
  hsrp 72
    preempt
    priority 110
    ip 10.72.0.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface Vlan73
  no shutdown
  ip address 10.10.73.3/24
  hsrp version 2
  hsrp 73
```

```
    preempt
    priority 110
    ip 10.10.73.1

interface Vlan74
  no shutdown
  ip address 10.10.74.3/24
  hsrp version 2
  hsrp 74
    preempt
    priority 110
    ip 10.10.74.1

interface Vlan75
  no shutdown
  ip address 10.10.75.3/24
  hsrp version 2
  hsrp 75
    preempt
    priority 110
    ip 10.10.75.1

interface Vlan76
  no shutdown
  ip address 10.10.76.3/23
  hsrp version 2
  hsrp 76
    preempt
    priority 110
    ip 10.10.76.1
  ip dhcp relay address 10.10.71.11
  ip dhcp relay address 10.10.71.12

interface port-channel10
  description VPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type network
  vpc peer-link

interface port-channel11
  description FI-Uplink-K22-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 11

interface port-channel12
  description FI-Uplink-K22-B
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 12

interface port-channel49
  description FI-Uplink-K23-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 49

interface port-channel50
  description FI-Uplink-K23-B
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  vpc 50

interface Ethernet1/1
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/2
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/3
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active
```

```
interface Ethernet1/4
  description VPC to K23-N9K-A
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  channel-group 10 mode active

interface Ethernet1/5

interface Ethernet1/6

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24
interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33
  switchport access vlan 71
  spanning-tree port type edge

interface Ethernet1/34
  switchport access vlan 71
  spanning-tree port type edge

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44
interface Ethernet1/45

interface Ethernet1/46
  description K23-HXVDIJH
  switchport access vlan 70
  spanning-tree port type edge

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  mtu 9216
  channel-group 49 mode active

interface Ethernet1/50
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  mtu 9216
  channel-group 50 mode active

interface Ethernet1/51
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/52
  switchport mode trunk
  switchport trunk allowed vlan 1,50-56,70-76,132
  spanning-tree port type edge trunk
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1,30-36,50-56,60-68,70-76,132

interface Ethernet1/54

interface mgmt0
  vrf member management
  ip address 173.37.52.104/23
line console
line vty
boot nxos bootflash:/nxos.9.3.3.bin
no system default switchport shutdown


telemetry
  certificate /bootflash/home/admin/telemetry-cert.pem localhost
  destination-profile
    use-nodeid timba-640143f86f72612d345931c3
  destination-group timba-640143f86f72612d345931c3-0
    ip address 10.10.71.20 port 443 protocol HTTP encoding JSON
  sensor-group timba-640143f86f72612d345931c3-0
    data-source NX-API
    path "show system resources all-modules" depth 0
  sensor-group timba-640143f86f72612d345931c3-1
    data-source NX-API
    path "show module" depth 0
  sensor-group timba-640143f86f72612d345931c3-2
    data-source NX-API
    path "show environment power" depth 0
  sensor-group timba-640143f86f72612d345931c3-3
    data-source NX-API
    path "show interface fc regex *" depth 0
  sensor-group timba-640143f86f72612d345931c3-4
    data-source DME
    path sys/ch depth 1 query-condition query-target=subtree&target-subtree-
class=eqptSensor
  sensor-group timba-640143f86f72612d345931c3-5
    data-source DME
    path sys/ch query-condition query-target=subtree&target-subtree-class=eqptSupC
  sensor-group timba-640143f86f72612d345931c3-6
    data-source DME
    path sys/ch query-condition query-target=subtree&target-subtree-class=eqptFt
  sensor-group timba-640143f86f72612d345931c3-7
    data-source DME
    path sys/intf depth 0 query-condition query-target=subtree&target-subtree-
class=ethpmPhysIf filter-condition updated(ethpmPhysIf.operSt)
  subscription 1565
    dst-grp timba-640143f86f72612d345931c3-0
    snsr-grp timba-640143f86f72612d345931c3-0 sample-interval 300000
    snsr-grp timba-640143f86f72612d345931c3-1 sample-interval 300000
    snsr-grp timba-640143f86f72612d345931c3-2 sample-interval 300000
    snsr-grp timba-640143f86f72612d345931c3-3 sample-interval 300000
    snsr-grp timba-640143f86f72612d345931c3-4 sample-interval 300000
```

```
        snsr-grp timba-640143f86f72612d345931c3-5 sample-interval 300000
        snsr-grp timba-640143f86f72612d345931c3-6 sample-interval 300000
        snsr-grp timba-640143f86f72612d345931c3-7 sample-interval 0
```

Appendix C – References Used in Guide

This section provides links to additional information for each partner's solution component of this document.

- Cisco UCS X-Series Modular System

  https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-x-series-modular-system/series.html

  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/solution-overview-c22-2432175.html

  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/cisco-ucs-x9508-chassis-aag.html

  https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x210c-M7-compute-node-aag.html

- Cisco UCS Manager Configuration Guides

  http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html

  https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-release-notes-list.html

- Cisco UCS Virtual Interface Cards

  https://www.cisco.com/c/en/us/products/interfaces-modules/unified-computing-system-adapters/index.html

- Cisco Nexus Switching References

  http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html

  https://www.cisco.com/c/en/us/products/switches/nexus-93180yc-fx-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

  https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9100-series-multilayer-fabric-switches/datasheet-c78-739613.html

- Cisco Intersight References

  https://www.cisco.com/c/en/us/products/cloud-systems-management/intersight/index.html

  https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/intersight-ds.html

- FlashStack Cisco Design Guides

  https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/data-center-design-guides-all.html#FlashStack

- Microsoft References

  https://docs.microsoft.com/en-us/fslogix/
- VMware References

  https://docs.vmware.com/en/VMware-vSphere/index.html
- Login VSI Documentation

  https://www.loginvsi.com/resources/

Appendix D – Parts list

Part Number	Description	Quanity
UCSX-FI-6454-D-U	UCS Fabric Interconnect 6454	2
CON-OSP-UCSXOOFI	SNTC-24X7X4OS UCS Fabric Interconnect 6454	2
UCSX-M2-HWRAID=	Cisco Boot optimized M.2 Raid controller	1
UCSX-M2-240G=	240GB SATA M.2	8
UCSX-NVMEI4-I1600=	1.6TB 2.5in U.2 Intel P5600 NVMe High Perf High Endurance	4
UCSX-NVMEI4-I1920=	1.9TB 2.5in U.2 Intel P5500 NVMe High Perf Medium Endurance	20
UCSX-MR-X32G1RW=	32GB RDIMM SRx4 4800 (16Gb)	128
UCSX-9508-U	UCS 9508 Chassis Configured	1
CON-OSP-UCSX95U8	SNTC-24X7X4OS UCS 9508 Chassis Configured	1
N9K-C93180YC-EX-RF	N9300 with 48p 10/25G SFP+ and 6p 100G QSFP28 REMANUFACTURED	2
CON-SNC-93180YCX	SNTC-NCD Nexus 9K,48p 10/25G6p 100G QSFP28,Spare(No Acc kit,	2
DC-MGT-SAAS	Cisco Intersight SaaS	1
UCSX-FI-6454-D-U	UCS Fabric Interconnect 6454	2
CON-OSP-UCSXOOFI	SNTC-24X7X4OS UCS Fabric Interconnect 6454	2
UCSX-M2-HWRAID=	Cisco Boot optimized M.2 Raid controller	1
UCSX-M2-240G=	240GB SATA M.2	8
UCSX-NVMEI4-I1600=	1.6TB 2.5in U.2 Intel P5600 NVMe High Perf High Endurance	4
UCSX-NVMEI4-I1920=	1.9TB 2.5in U.2 Intel P5500 NVMe High Perf Medium Endurance	20
UCSX-MR-X32G1RW=	32GB RDIMM SRx4 4800 (16Gb)	128
UCSX-9508-U	UCS 9508 Chassis Configured	1
CON-OSP-UCSX95U8	SNTC-24X7X4OS UCS 9508 Chassis Configured	1

Part Number	Description	Quanity
N9K-C93180YC-EX-RF	N9300 with 48p 10/25G SFP+ and 6p 100G QSFP28 REMANUFACTURED	2
CON-SNC-93180YCX	SNTC-NCD Nexus 9K,48p 10/25G6p 100G QSFP28,Spare(No Acc kit,	2
DC-MGT-SAAS	Cisco Intersight SaaS	1

Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community 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 IMPLE-MENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

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**.