FlashStack VDI Cisco UCS X-Series M6 with Citrix on VMware vSphere 8 up to 2600 Seats

Deployment Guide

Published: May 2023

In partnership with:

About the Cisco Validated Design Program

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

Executive Summary

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

This document details the design of the FlashStack Virtual Desktop Infrastructure for Citrix Virtual Apps and Desktops VMware vSphere 8.0 Design Guide, which describes a validated Converged Infrastructure (CI) jointly developed by Cisco and Pure Storage.

This solution explains the deployment of a predesigned, best-practice data center architecture with:

- VMware vSphere
- Citrix Virtual Apps
- Cisco Unified Computing System (Cisco UCS) incorporating the Cisco X-Series modular platform
- Cisco Nexus 9000 family of switches
- Cisco MDS 9000 family of Fibre Channel switches
- Pure Storage FlashArray//X R3 all flash array supporting Fibre Channel storage access

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

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

Customers interested in understanding the FlashStack design and deployment details, including the configuration of various elements of design and associated best practices, should refer to Cisco Validated Designs for FlashStack, here: Data Center Design Guides - FlashStack Platforms

Solution Overview

This chapter contains the following:

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

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

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 the following:

- Large-scale Citrix Virtual Apps and Desktops VDI
- Pure Storage FlashArray//X Storage Array
- Cisco UCS X210c M6 Blade Servers running VMware vSphere 8.0
- Cisco Nexus 9000 Series Ethernet Switches
- Cisco MDS 9100 Series Multilayer Fibre Channel Switches

What’s New in this Release?

This version of the FlashStack VDI Design is based on the latest Cisco FlashStack Virtual Server Infrastructure and introduces the Cisco UCS X-Series modular platform.

Highlights for this design include:

- Support for Cisco UCS X9508 chassis with Cisco UCS X210c M6 compute nodes
- Support for Pure Storage FlashArray//X70 R3 with Purity version 6.3.3
- Citrix Virtual Apps and Desktops 2203
- 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 the FlashStack components
- Support for Cisco Intersight Assist virtual appliance to help connect the Pure Storage FlashArray and VMware vCenter with the Cisco Intersight platform

- Fully automated solution deployment describing the FlashStack infrastructure and vSphere virtualization

These factors have led to the need for a predesigned computing, networking and storage building blocks optimized to lower the initial design cost, simplify management, and enable horizontal scalability and high levels of utilization.

The use cases include:

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

Technology Overview

This chapter contains the following:

- FlashStack
- Cisco Unified Computing System
- Cisco UCS Fabric Interconnect
- Cisco Unified Computing System X-Series
- Cisco UCS Virtual Interface Cards (VICs)
- Cisco Switching
- Citrix Virtual Apps and Desktops 2203
- Citrix Cloud
- VMware vSphere 8.0
- Red Hat Ansible
- Cisco Intersight Assist Device Connector for VMware vCenter and Pure Storage FlashArray
- Pure Storage for VDI
- Purity for FlashArray
- Pure1

Cisco and Pure Storage have partnered to deliver over 50 Cisco Validated Designs, which use best-in-class storage, server, and network components to serve as the foundation for virtualized workloads such as Virtual Desktop Infrastructure (VDI), enabling efficient architectural designs that you can deploy quickly and confidently.

FlashStack

The FlashStack architecture was jointly developed by Cisco and Pure Storage. All FlashStack components are integrated, allowing customers can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlashStack is its ability to maintain consistency at scale. Each of the component families shown in Figure 1 (Cisco UCS, Cisco Nexus, Cisco MDS, and Pure Storage FlashArray systems) offers platform and resource options to scale up or scale out the infrastructure while supporting the same features and functions.

Figure 1. FlashStack components

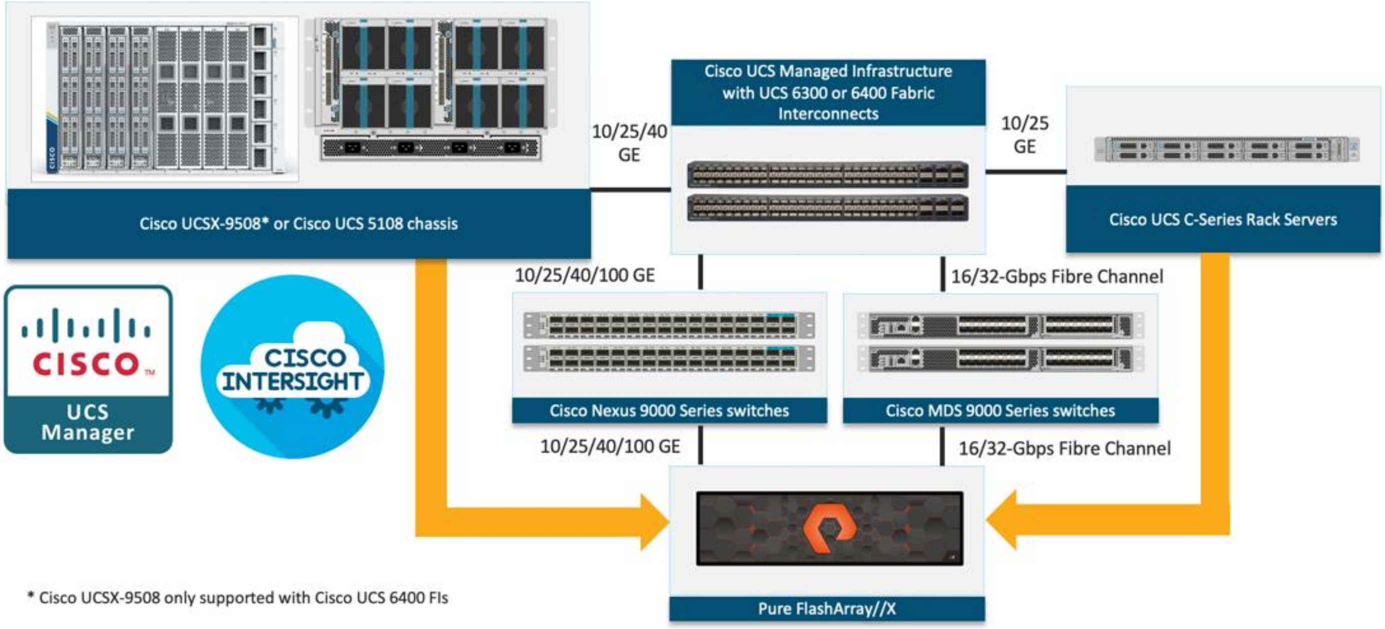
Cisco UCSX-9508* or Cisco UCS 5108 chassis
10/25/40 GE
Cisco UCS Managed Infrastructure with UCS 6300 or 6400 Fabric Interconnects
10/25 GE
Cisco UCS C-Series Rack Servers
10/25/40/100 GE
16/32-Gbps Fibre Channel
CISCO UCS Manager
CISCO INTERSIGHT
Cisco Nexus 9000 Series switches
Cisco MDS 9000 Series switches
10/25/40/100 GE
16/32-Gbps Fibre Channel
* Cisco UCSX-9508 only supported with Cisco UCS 6400 FIs
Pure FlashArray//X

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 Third-Generation Intel Xeon Scalable 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 Cisco UCS Manager discovers a compute resource, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD 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.
- Extended Memory: The enterprise-class Cisco UCS Blade server extends the capabilities of the Cisco Unified Computing System portfolio in a half-width blade form factor. It harnesses the power of the latest Intel Xeon Scalable Series processor family CPUs and Intel Optane DC Persistent Memory (DCPMM) with up to 18TB of RAM (using 256GB DDR4 DIMMs and 512GB DCPMM).
- 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 2. Cisco Intersight

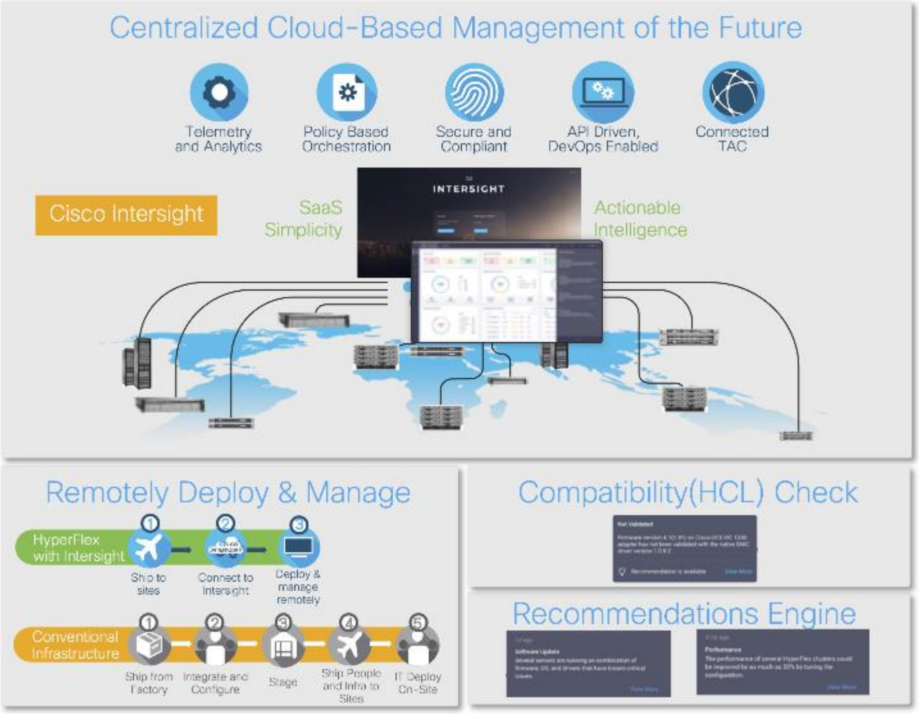
Centralized Cloud-Based Management of the Future
Telemetry and Analytics
Policy Based Orchestration
Secure and Compliant
API Driven, DevOps Enabled
Connected TAC
Cisco Intersight
SaaS Simplicity
INTERSIGHT
Actionable Intelligence
Remotely Deploy & Manage
HyperFlex with Intersight
1
2
3
Ship to sites
Connect to Intersight
Deploy & manage remotely
Conventional Infrastructure
1
2
3
4
5
Ship from Factory
Integrate and Configure
Stage
Ship People and Infra to Sites
IT Deploy On-Site
Compatibility(HCL) Check
Not Validated
Recommendation is available
View More
Recommendations Engine
Software Update
Several servers are running on combination of firmware, OS, and drivers that have known critical issues
View More
Performance
The performance of several HyperFlex clusters could be improved by as much as 20% by tuning the configuration.
View More

- 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 3. Cisco UCS 6454 Fabric Interconnect

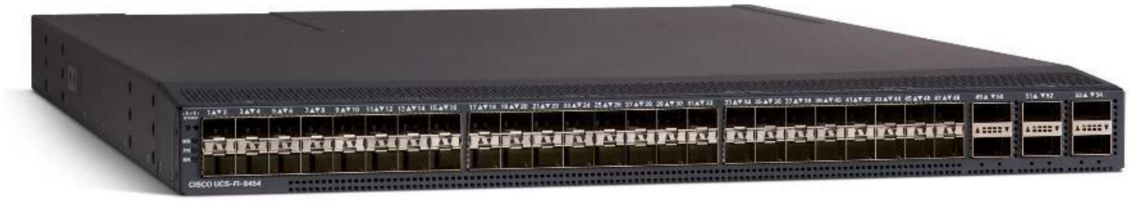

Cisco Unified Computing System X-Series

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 and Pure storage, to provide visibility, optimization, and orchestration from a single platform, thereby driving agility and deployment consistency.

Figure 4.   Cisco UCS X9508 Chassis

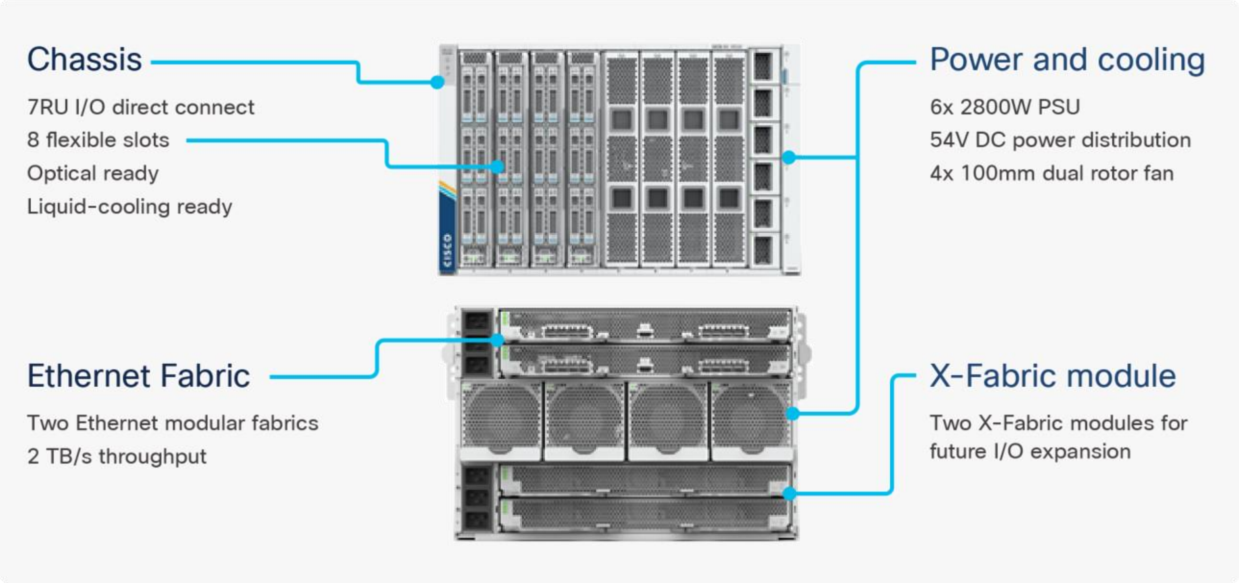

The various components of the Cisco UCS X-Series are described in the following sections.

Cisco UCS X9508 Chassis

The Cisco UCS X-Series chassis is engineered to be adaptable and flexible. As shown in Figure 5, 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 5.   Cisco UCS X9508 Chassis – Midplane Free Design

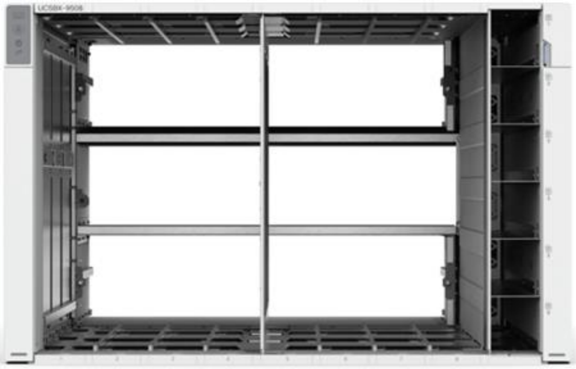

The Cisco UCS X9508 7-Rack-Unit (7RU) chassis has eight flexible slots. These slots can house a combination of compute nodes and a pool of future I/O resources that may include GPU accelerators, disk storage, and nonvolatile memory. At the top rear of the chassis are two Intelligent Fabric Modules (IFMs) 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 6. 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 M6 Compute Node

The Cisco UCS X9508 Chassis is designed to host up to 8 Cisco UCS X210c M6 Compute Nodes. The hardware details of the Cisco UCS X210c M6 Compute Nodes are shown in Figure 7:

Figure 7. Cisco UCS X210c M6 Compute Node

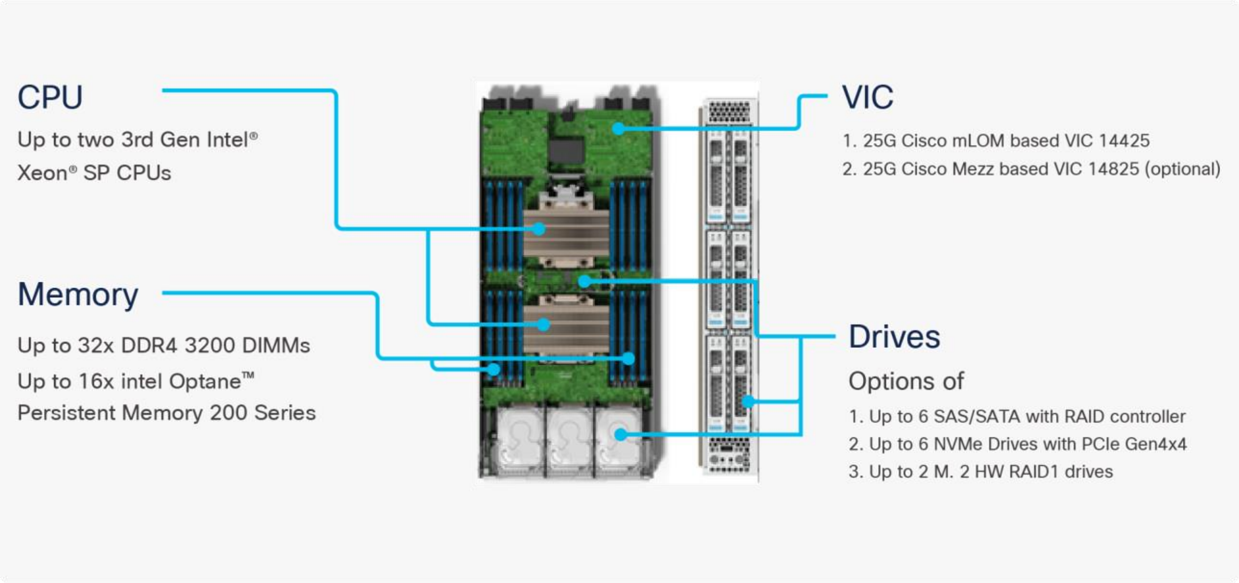

CPU
Up to two 3rd Gen Intel® Xeon® SP CPUs
Memory
Up to 32x DDR4 3200 DIMMs
Up to 16x intel Optane™ Persistent Memory 200 Series
VIC
1. 25G Cisco mLOM based VIC 14425
2. 25G Cisco Mezz based VIC 14825 (optional)
Drives
Options of
1. Up to 6 SAS/SATA with RAID controller
2. Up to 6 NVMe Drives with PCIe Gen4x4
3. Up to 2 M. 2 HW RAID1 drives

The Cisco UCS X210c M6 features:

- CPU: Up to 2x 3rd Gen Intel Xeon Scalable Processors with up to 40 cores per processor and 1.5 MB Level 3 cache per core
- Memory: Up to 32 x 256 GB DDR4-3200 DIMMs for a maximum of 8 TB of main memory. The Compute Node can also be configured for up to 16 x 512-GB Intel Optane persistent memory DIMMs for a maximum of 12 TB of memory
- Disk storage: Up to 6 SAS or SATA drives can be configured with an internal RAID controller, or customers can configure up to 6 NVMe drives. 2 M.2 memory cards can be added to the Compute Node with RAID 1 mirroring.
- Virtual Interface Card (VIC): Up to 2 VICs including an mLOM Cisco VIC 14425 and a mezzanine Cisco VIC card 14825 can be installed in a Compute Node.
- Security: The server supports an optional Trusted Platform Module (TPM). Additional security features include a secure boot FPGA and ACT2 anticounterfeit provisions.

Cisco UCS Virtual Interface Cards (VICs)

Cisco UCS X210c M6 Compute Nodes support the following two Cisco fourth-generation VIC cards:

Cisco VIC 14425

Cisco VIC 14425 fits the mLOM slot in the Cisco X210c Compute Node and enables up to 50 Gbps of unified fabric connectivity to each of the chassis IFMs for a total of 100 Gbps of connectivity per server. Cisco VIC 14425 connectivity to the IFM and up to the fabric interconnects is delivered through 4x 25-Gbps connections, which are configured automatically as 2x 50-Gbps port channels. Cisco VIC 14425 supports 256 virtual interfaces (both Fibre Channel and Ethernet) along with the latest networking innovations such as NVMeoF over RDMA (ROCEv2), VxLAN/NVGRE offload, and so on.

Figure 8. Single Cisco VIC 14425 in Cisco UCS X210c M6

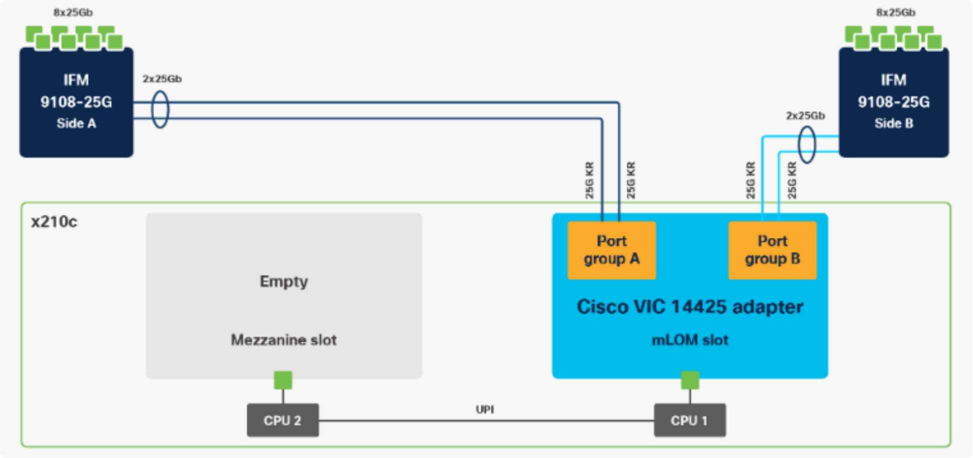

The connections between the 4th generation Cisco VIC (Cisco UCS VIC 14425) in the Cisco UCS B200 blades and the I/O mod-ules in the Cisco UCS 5108 chassis comprise of multiple 10Gbps KR lanes. The same connections between Cisco VIC 14425 and IFMs in Cisco UCS X-Series comprise of multiple 25Gbps KR lanes resulting in 2.5x better connectivity in Cisco UCS X210c M6 Compute Nodes. The network interface speed comparison between VMware ESXi installed on Cisco UCS X210C M6 with Cisco VIC 14425 and Cisco UCS X210c M6 with Cisco VIC 14425 is shown in Figure 9.

Figure 9. Network Interface Speed Comparison

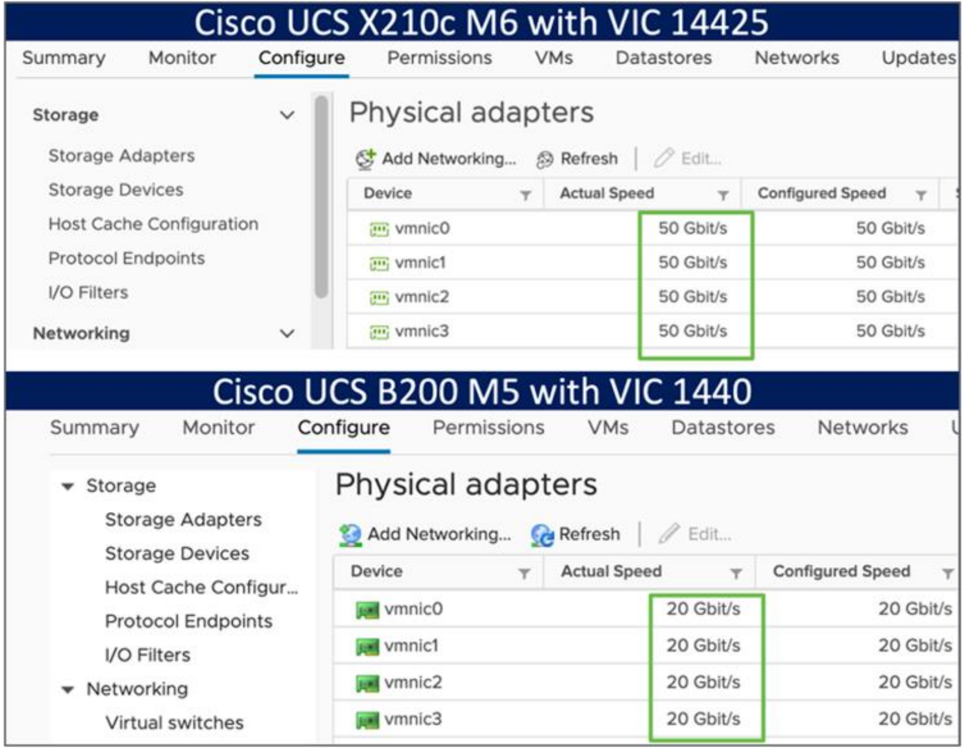

Cisco UCS X210c M6 with VIC 14425

Physical adapters

Device	Actual Speed	Configured Speed
vmnic0	50 Gbit/s	50 Gbit/s
vmnic1	50 Gbit/s	50 Gbit/s
vmnic2	50 Gbit/s	50 Gbit/s
vmnic3	50 Gbit/s	50 Gbit/s

Cisco UCS B200 M5 with VIC 1440

Physical adapters

Device	Actual Speed	Configured Speed
vmnic0	20 Gbit/s	20 Gbit/s
vmnic1	20 Gbit/s	20 Gbit/s
vmnic2	20 Gbit/s	20 Gbit/s
vmnic3	20 Gbit/s	20 Gbit/s

Cisco Switching

Cisco Nexus 93180YC-FX Switches

The 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 10. **Cisco Nexus 93180YC-EX Switch**

Cisco MDS 9132T 32-Gb Fiber Channel Switch

The next-generation Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switch (Figure 11) provides high-speed Fibre Channel connectivity from the server rack to the SAN core. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the dual benefits of greater bandwidth and consolidation.

Small-scale SAN architectures can be built from the foundation using this low-cost, low-power, non-blocking, line-rate, and low-latency, bi-directional airflow capable, fixed standalone SAN switch connecting both storage and host ports.

Medium-size to large-scale SAN architectures built with SAN core directors can expand 32-Gb connectivity to the server rack using these switches either in switch mode or Network Port Virtualization (NPV) mode.

Additionally, investing in this switch for the lower-speed (4- or 8- or 16-Gb) server rack gives you the option to upgrade to 32-Gb server connectivity in the future using the 32-Gb Host Bus Adapter (HBA) that are available today. The Cisco MDS 9132T 32-Gb 32-Port Fibre Channel switch also provides unmatched flexibility through a unique port expansion module (Figure 15) that provides a robust cost-effective, field swappable, port upgrade option.

This switch also offers state-of-the-art SAN analytics and telemetry capabilities that have been built into this next-generation hardware platform. This new state-of-the-art technology couples the next-generation port ASIC with a fully dedicated Network Processing Unit designed to complete analytics calculations in real time. The telemetry data extracted from the inspection of the frame headers are calculated on board (within the switch) and, using an industry-leading open format, can be streamed to any analytics-visualization platform. This switch also includes a dedicated 10/100/1000BASE-T telemetry port to maximize data delivery to any telemetry receiver including Cisco Data Center Network Manager.

Figure 11. Cisco MDS 9132T 32-Gb Fibre Channel Switch

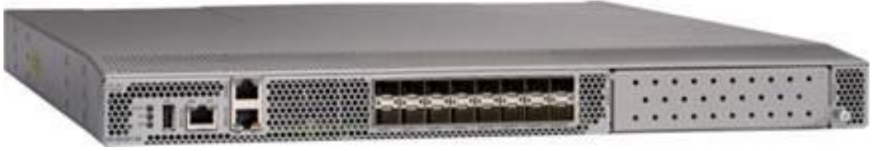

Figure 12. Cisco MDS 9132T 32-Gb 16-Port Fibre Channel Port Expansion Module

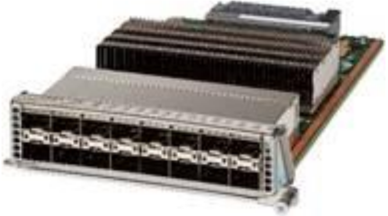

- Features
  - High performance: Cisco MDS 9132T architecture, with chip-integrated nonblocking arbitration, provides consistent 32-Gb low-latency performance across all traffic conditions for every Fibre Channel port on the switch.
  - Capital Expenditure (CapEx) savings: The 32-Gb ports allow users to deploy them on existing 16- or 8-Gb transceivers, reducing initial CapEx with an option to upgrade to 32-Gb transceivers and adapters in the future.
  - High availability: Cisco MDS 9132T switches continue to provide the same outstanding availability and reliability as the previous-generation Cisco MDS 9000 Family switches by providing optional redundancy on all major components such as the power supply and fan. Dual power supplies also facilitate redundant power grids.
  - Pay-as-you-grow: The Cisco MDS 9132T Fibre Channel switch provides an option to deploy as few as eight 32-Gb Fibre Channel ports in the entry-level variant, which can grow by 8 ports to 16 ports, and thereafter with a port expansion module with sixteen 32-Gb ports, to up to 32 ports. This approach results in lower initial investment and power consumption for entry-level configurations of up to 16 ports compared to a fully loaded switch. Upgrading through an expansion module also reduces the overhead of managing multiple instances of port activation licenses on the switch. This unique combination of port upgrade options allow four possible configurations of 8 ports, 16 ports, 24 ports and 32 ports.
  - Next-generation Application-Specific Integrated Circuit (ASIC): The Cisco MDS 9132T Fibre Channel switch is powered by the same high-performance 32-Gb Cisco ASIC with an integrated network processor that powers the Cisco MDS 9700 48-Port 32-Gb Fibre Channel Switching Module. Among all the advanced features that this ASIC enables, one of the most notable is inspection of Fibre Channel and Small Computer System Interface (SCSI) headers at wire speed on every flow in the smallest form-factor Fibre Channel switch without the need for any external taps or appliances. The recorded flows can be analyzed on the switch and also exported using a dedicated 10/100/1000BASE-T port for telemetry and analytics purposes.

- Intelligent network services: Slow-drain detection and isolation, VSAN technology, Access Control Lists (ACLs) for hardware-based intelligent frame processing, smartzoning and fabric wide Quality of Service (QoS) enable migration from SAN islands to enterprise-wide storage networks. Traffic encryption is optionally available to meet stringent security requirements.
- Sophisticated diagnostics: The Cisco MDS 9132T provides intelligent diagnostics tools such as Inter-Switch Link (ISL) diagnostics, read diagnostic parameters, protocol decoding, network analysis tools, and integrated Cisco Call Home capability for greater reliability, faster problem resolution, and reduced service costs.
- Virtual machine awareness: The Cisco MDS 9132T provides visibility into all virtual machines logged into the fabric. This feature is available through HBAs capable of priority tagging the Virtual Machine Identifier (VMID) on every FC frame. Virtual machine awareness can be extended to intelligent fabric services such as analytics[1] to visualize performance of every flow originating from each virtual machine in the fabric.
- Programmable fabric: The Cisco MDS 9132T provides powerful Representational State Transfer (REST) and Cisco NX-API capabilities to enable flexible and rapid programming of utilities for the SAN as well as polling point-in-time telemetry data from any external tool.
- Single-pane management: The Cisco MDS 9132T can be provisioned, managed, monitored, and troubleshot using Cisco Data Center Network Manager (DCNM), which currently manages the entire suite of Cisco data center products.
- Self-contained advanced anticounterfeiting technology: The Cisco MDS 9132T uses on-board hardware that protects the entire system from malicious attacks by securing access to critical components such as the bootloader, system image loader and Joint Test Action Group (JTAG) interface.

Citrix Virtual Apps and Desktops 2203

Citrix Virtual Apps and Desktops is a modern platform for running and delivering virtual desktops and apps across the hybrid cloud. For administrators, this means simple, automated, and secure desktop and app management. For users, it provides a consistent experience across devices and locations.

The virtual app and desktop solution designed for an exceptional experience.

Today's employees spend more time than ever working remotely, causing companies to rethink how IT services should be delivered. To modernize infrastructure and maximize efficiency, many are turning to desktop as a service (DaaS) to enhance their physical desktop strategy, or they are updating on-premises virtual desktop infrastructure (VDI) deployments. Managed in the cloud, these deployments are high-performance virtual instances of desktops and apps that can be delivered from any datacenter or public cloud provider.

DaaS and VDI capabilities provide corporate data protection as well as an easily accessible hybrid work solution for employees. Because all data is stored securely in the cloud or datacenter, rather than on devices, end-users can work securely from anywhere, on any device, and over any network—all with a fully IT-provided experience. IT also gains the benefit of centralized management, so they can scale their environments quickly and easily. By separating endpoints and corporate data, resources stay protected even if the devices are compromised.

As a leading VDI and DaaS provider, Citrix provides the capabilities organizations need for deploying virtual apps and desktops to reduce downtime, increase security, and alleviate the many challenges associated with traditional desktop management.

Citrix Cloud

Citrix Cloud is a platform that hosts and administers Citrix cloud services. It connects to your resources through connectors on any cloud or infrastructure you choose (on-premises, public cloud, private cloud, or hybrid cloud). It allows you to create, manage, and deploy workspaces with apps and data to your end-users from a single console.

VMware vSphere 8.0

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

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

Some limits have been increased in VMware vSphere 8 compared to VMware vSphere 7 U3:

- **The number of vGPU devices** is increased to 8
- **The number of ESXi hosts that can be managed by Lifecycle Manager** is increased from 400 to 1,000
- **VMs per cluster is increased** from 8,000 to 10,000
- **VM DirectPath I/O devices per host** is increased from 8 to 32

Higher limits in vSphere 8.0 allow you to run more VMs, run more powerful VMs, and perform tasks faster.

VMware vSphere 8 introduces the Distributed Services Engine to work with Data Processing Units that allow you to offload a central processing unit (CPU).

A **Data Processing Unit (DPU)** is a new class of programmable processors built on the ARM architecture, which can be used together with CPUs and GPUs (graphics processing units) for computing operations primarily related to networking and communications.

A DPU is now incorporated into a Smart NIC controller, which is plugged into the motherboard. This new approach allows us to improve network performance in a virtual environment built on vSphere 8 and offload a CPU for performing network operations.

A Smart NIC is much more powerful than the traditional NIC (network interface controller) and is especially useful with VMware NSX. VMware says that up to 20% of CPU workloads can be freed up when using DPUs now. The higher transaction rate, lower latency, and security benefits are the advantages you get with this feature.

NSX Distributed Firewall will now use DPUs, this is optimal for east-west network traffic, which is increasingly more common in modern virtualized data centers. This approach offloads security operations from CPU to DPU.

VMware vSphere on DPUs, known as Project Monterey before the release, was implemented in collaboration with hardware vendors like Intel, AMD, NVIDIA, and OEM system partners.

There are also a couple of new security features in VMware vSphere 8.0:

- **SSH timeout**. You can enable SSH access to an ESXi host for a specified period. After this period (timeout) expires, SSH access is disabled automatically. This feature helps avoid accidental SSH access, for example, when an administrator forgets to disable it after completing tasks via SSH.
- **TPM Provision Policy**. This feature gives you the ability to automatically replace a vTPM (Trusted Platform Module) device when cloning VMs. The aim is to improve security and avoid risks associated with TPM secrets being copied.

When configuring VM cloning, you have two options to choose from: copy or replace. Copying virtual machines with the same TPM may cause security issues because the same secret is used on multiple VMs. This new vSphere 8 feature allows you to automatically avoid this issue by choosing the *replace* option.

Note: **TLS 1.2** is the minimum supported version now. You can use higher versions, but legacy TLS versions are not supported.

For more information about VMware vSphere and its components, see: https://www.vmware.com/products/vsphere.html.

VMware vSphere vCenter

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

Red Hat Ansible

Ansible is simple and powerful, allowing users to easily manage various physical devices within FlashStack including the provisioning of Cisco UCS servers, Cisco Nexus switches, Pure Storage and VMware vSphere. Using Ansible's Playbook-based automation is easy and integrates into your current provisioning infrastructure.

Cisco Intersight Assist Device Connector for VMware vCenter and Pure Storage FlashArray

Cisco Intersight integrates with VMware vCenter and Pure Storage FlashArray as follows:

- Cisco Intersight uses the device connector running within Cisco Intersight Assist virtual appliance to communicate with the VMware vCenter.
- Cisco Intersight uses the device connector running within a Cisco Intersight Assist virtual appliance to integrate with all Pure Storage FlashArray models. The newest version 1.1 of Pure Storage integration to Cisco Intersight introduces support for REST API 2.x for FlashArray products (running Purity//FA 6.0.3 or later), along with User Agent support (for telemetry). Intersight Cloud Orchestrator now has new storage tasks for adding/removing a Pure Storage snapshot and copying a Pure Storage volume from snapshot.

Figure 13. Cisco Intersight and vCenter and Pure Storage Integration

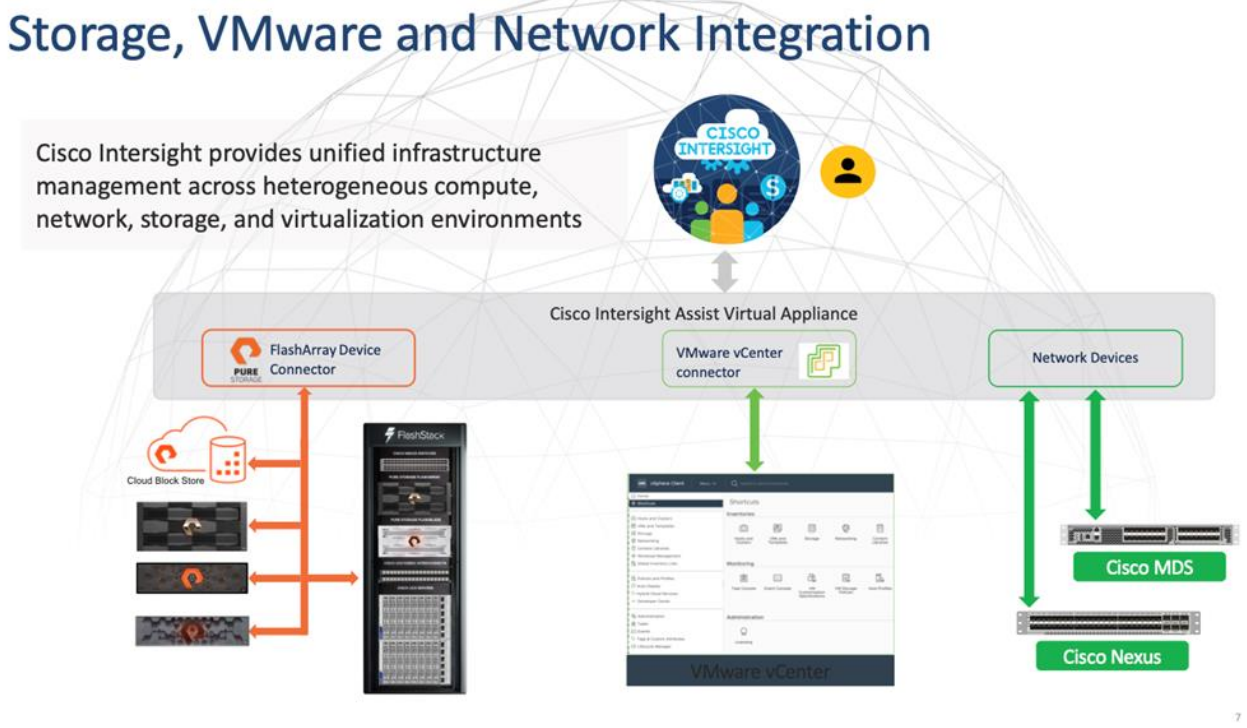
Storage, VMware and Network Integration
Cisco Intersight provides unified infrastructure management across heterogeneous compute, network, storage, and virtualization environments
CISCO INTERSIGHT
Cisco Intersight Assist Virtual Appliance
FlashArray Device Connector
PURE STORAGE
VMware vCenter connector
Network Devices
Cloud Block Store
FlashStack
VMware vCenter
Cisco MDS
Cisco Nexus
7

The device connector provides a safe way for connected targets to send information and receive control instructions from the Cisco Intersight portal using a secure Internet connection. The integration brings the full value and simplicity of Cisco Intersight infrastructure management service to VMware hypervisor and FlashArray storage environments. The integration architecture enables FlashStack customers to use new management capabilities with no compromise in their existing VMware or FlashArray operations. IT users will be able to manage heterogeneous infrastructure from a centralized Cisco Intersight portal. At the same time, the IT staff can continue to use VMware vCenter and the Pure Storage dashboard for comprehensive analysis, diagnostics, and reporting of virtual and storage environments. The next section addresses the functions that this integration provides.

Pure Storage for VDI

Pure Storage helps organizations—of all sizes and across multiple industries—overcome the most common reasons for disappointing results from a VDI. All-flash storage delivers:

- Always-on, always fast and always secure VDI, ensuring a consistently superior end-user experience
- Efficiency with up to 2x better data-reduction rates, lowering capital and operating costs
- Effortless storage management, sharply reducing the demands on IT staff
- Evergreen growth and scalability, incorporating non-disruptive upgrades and clearly defined costs known well in advance.

Whether you're planning a VDI rollout or have already implemented VDI that's delivering sub-par results, this white paper will provide valuable guidance—citing actual end-user deployments—that clearly illustrates how deploying flash storage can optimize your end-user productivity and experience with VDI.

Purity for FlashArray

The essential element of every FlashArray is the Purity Operating Environment software. Purity implements advanced data reduction, storage management, and flash management features, enabling organizations to enjoy Tier 1 data services for all workloads, proven 99.9999% availability over multiple years (inclusive of maintenance and generational upgrades), completely non-disruptive operations, 2X better data reduction versus alternative all-flash solutions, and the power and efficiency of DirectFlash.

Moreover, Purity includes enterprise-grade data security, modern data protection options, and complete business continuity and global disaster recovery through ActiveCluster multi-site stretch cluster and ActiveDR* for continuous replication with near zero RPO. All these features are included with every array.

FlashArray File Services

Pure Storage acquired Compuverde several years ago, and they've been busy integrating this technology into the Purity//FA operating system. They emphasize the "integrating," because they didn't just take the existing product, drop it onto a FlashArray system, and run it on top of Purity. Instead, they incorporated key parts of it into Purity to give you the advantages of a unified black and file storage array.

The SMB and NFS protocols bring consolidated storage to the Purity//FA operating system, complementing its block capabilities, while the file system offers features like directory snapshots, replication and directory-level performance and space monitoring. For the purposes of this reference architecture, we will be focusing on using File Services for User Profile/Home Directory management.

Figure 14. **FlashArray//X Specifications**

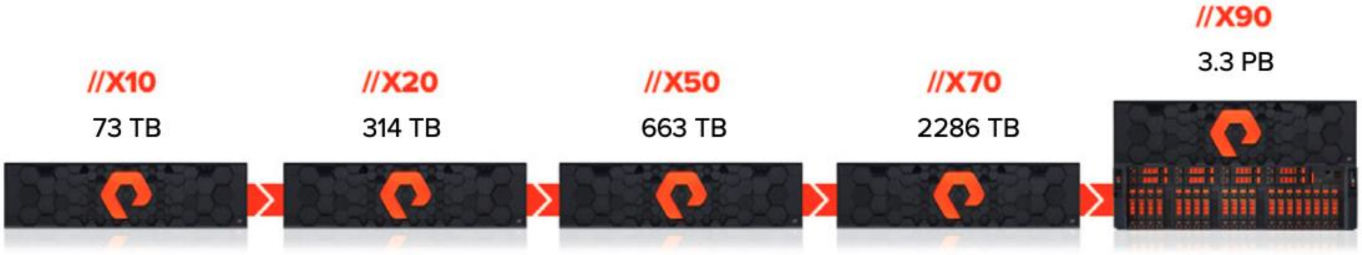

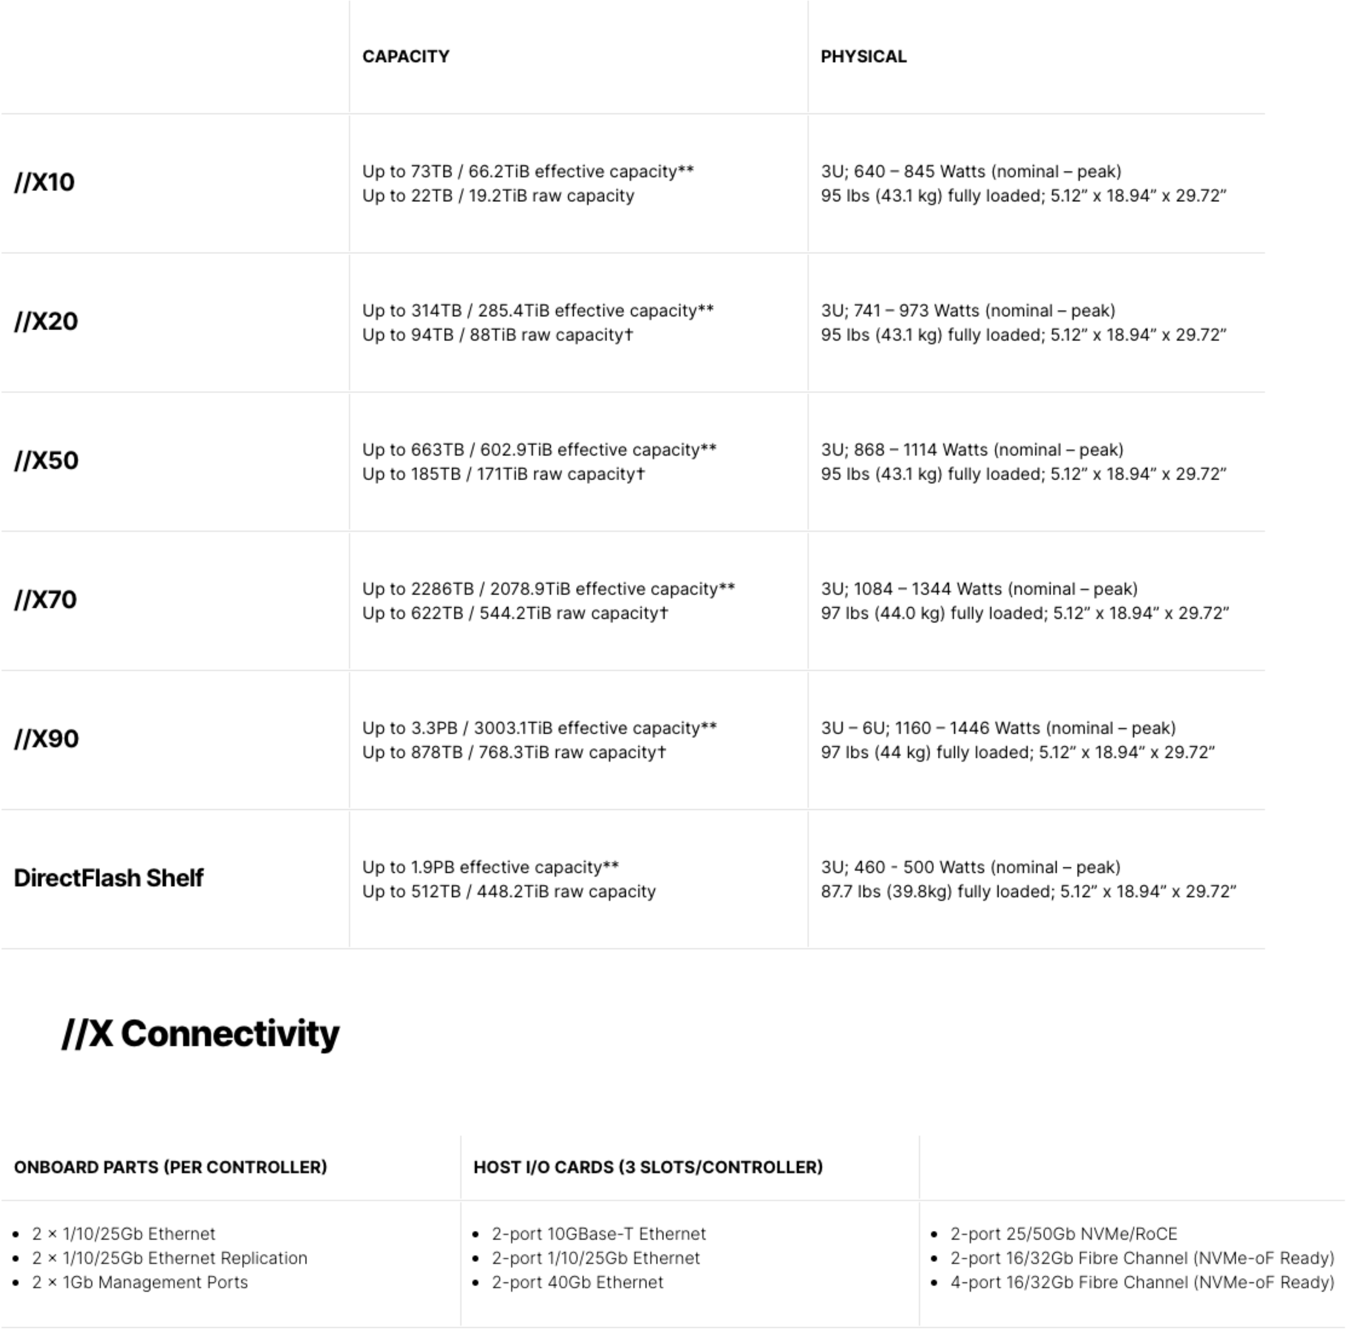

	CAPACITY	PHYSICAL
//X10	Up to 73TB / 66.2TiB effective capacity** Up to 22TB / 19.2TiB raw capacity	3U; 640 – 845 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X20	Up to 314TB / 285.4TiB effective capacity** Up to 94TB / 88TiB raw capacity†	3U; 741 – 973 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X50	Up to 663TB / 602.9TiB effective capacity** Up to 185TB / 171TiB raw capacity†	3U; 868 – 1114 Watts (nominal – peak) 95 lbs (43.1 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X70	Up to 2286TB / 2078.9TiB effective capacity** Up to 622TB / 544.2TiB raw capacity†	3U; 1084 – 1344 Watts (nominal – peak) 97 lbs (44.0 kg) fully loaded; 5.12" x 18.94" x 29.72"
//X90	Up to 3.3PB / 3003.1TiB effective capacity** Up to 878TB / 768.3TiB raw capacity†	3U – 6U; 1160 – 1446 Watts (nominal – peak) 97 lbs (44 kg) fully loaded; 5.12" x 18.94" x 29.72"
DirectFlash Shelf	Up to 1.9PB effective capacity** Up to 512TB / 448.2TiB raw capacity	3U; 460 - 500 Watts (nominal – peak) 87.7 lbs (39.8kg) fully loaded; 5.12" x 18.94" x 29.72"

//X Connectivity

ONBOARD PARTS (PER CONTROLLER)	HOST I/O CARDS (3 SLOTS/CONTROLLER)	
• 2 × 1/10/25Gb Ethernet • 2 × 1/10/25Gb Ethernet Replication • 2 × 1Gb Management Ports	• 2-port 10GBase-T Ethernet • 2-port 1/10/25Gb Ethernet • 2-port 40Gb Ethernet	• 2-port 25/50Gb NVMe/RoCE • 2-port 16/32Gb Fibre Channel (NVMe-oF Ready) • 4-port 16/32Gb Fibre Channel (NVMe-oF Ready)

** Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1 and does not include thin provisioning or snapshots.

† Array accepts Pure Storage DirectFlash Shelf and/or Pure Storage SAS-based expansion shelf.

Evergreen Storage

Customers can deploy storage once and enjoy a subscription to continuous innovation through Pure's Evergreen Storage ownership model: expand and improve performance, capacity, density, and/or features for 10 years or more – all without downtime, performance impact, or data migrations. Pure has disrupted the industry's 3-5-year rip-and-replace cycle by engineering compatibility for future technologies right into its products, notably nondisruptive capability to upgrade from //M to //X with NVMe, DirectMemory, and NVMe-oF capability.

Pure1

Pure1, our cloud-based management, analytics, and support platform, expands the self-managing, plug-n-play design of Pure all-flash arrays with the machine learning predictive analytics and continuous scanning of Pure1 Meta to enable an effortless, worry-free data platform.

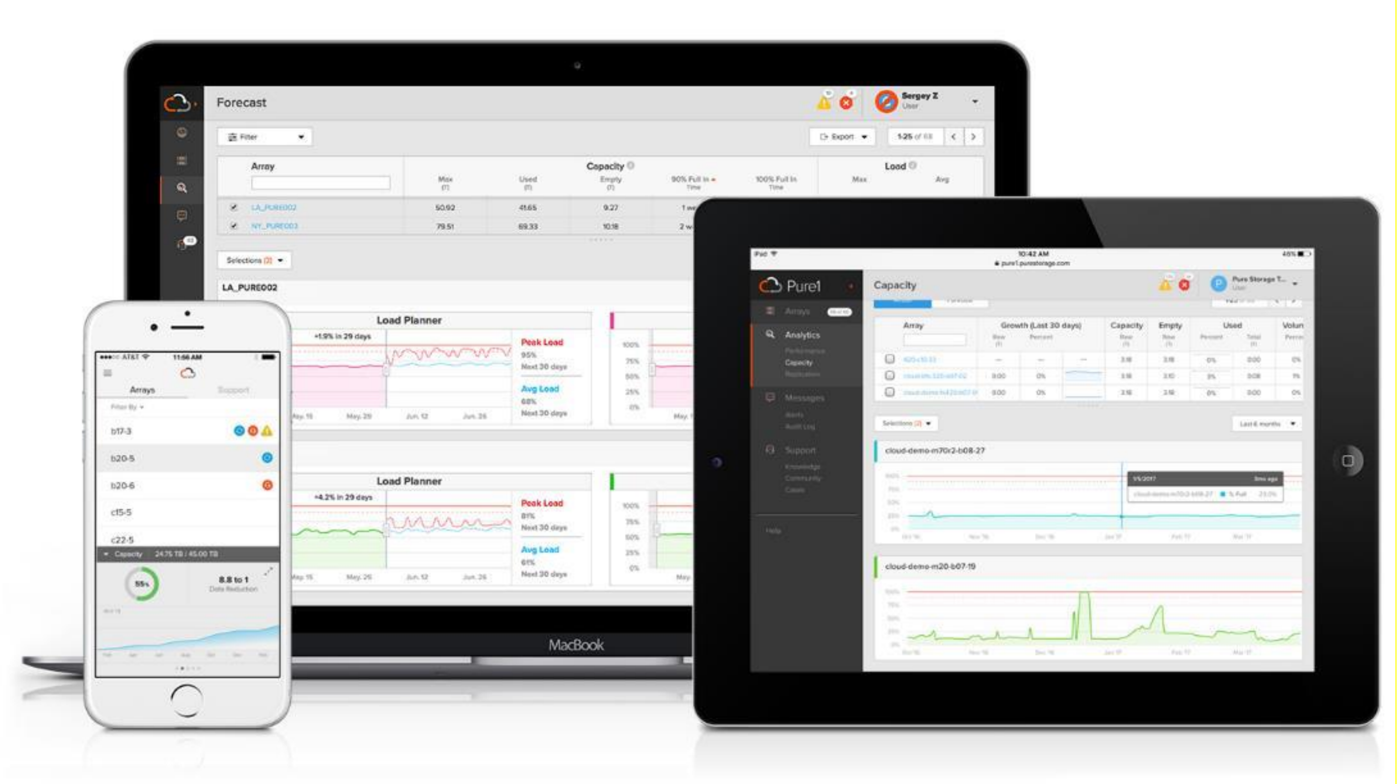

Pure1 Manage

In the Cloud IT operating model, installing, and deploying management software is an oxymoron: you simply login. Pure1 Manage is SaaS-based, allowing you to manage your array from any browser or from the Pure1 Mobile App – with nothing extra to purchase, deploy, or maintain. From a single dashboard you can manage all your arrays, with full visibility on the health and performance of your storage.

Pure1 Analyze

Pure1 Analyze delivers true performance forecasting – giving customers complete visibility into the performance and capacity needs of their arrays – now and in the future. Performance forecasting enables intelligent consolidation and unprecedented workload optimization.

Pure1 Support

Pure combines an ultra-proactive support team with the predictive intelligence of Pure1 Meta to deliver unrivaled support that's a key component in our proven FlashArray 99.9999% availability. Customers are often surprised and delighted when we fix issues they did not even know existed.

Pure1 META

The foundation of Pure1 services, Pure1 Meta is global intelligence built from a massive collection of storage array health and performance data. By continuously scanning call-home telemetry from Pure's installed base, Pure1 Meta uses machine learning predictive analytics to help resolve potential issues and optimize workloads. The result is both a white glove customer support experience and breakthrough capabilities like accurate performance forecasting.

Meta is always expanding and refining what it knows about array performance and health, moving the Data Platform toward a future of self-driving storage.

Pure1 VM Analytics

Pure1 helps you narrow down the troubleshooting steps in your virtualized environment. VM Analytics provides you with a visual representation of the IO path from the VM all the way through to the FlashArray. Other tools and features guide you through identifying where an issue might be occurring in order to help eliminate potential candidates for a problem.

VM Analytics doesn't only help when there's a problem. The visualization allows you to identify which volumes and arrays particular applications are running on. This brings the whole environment into a more manageable domain.

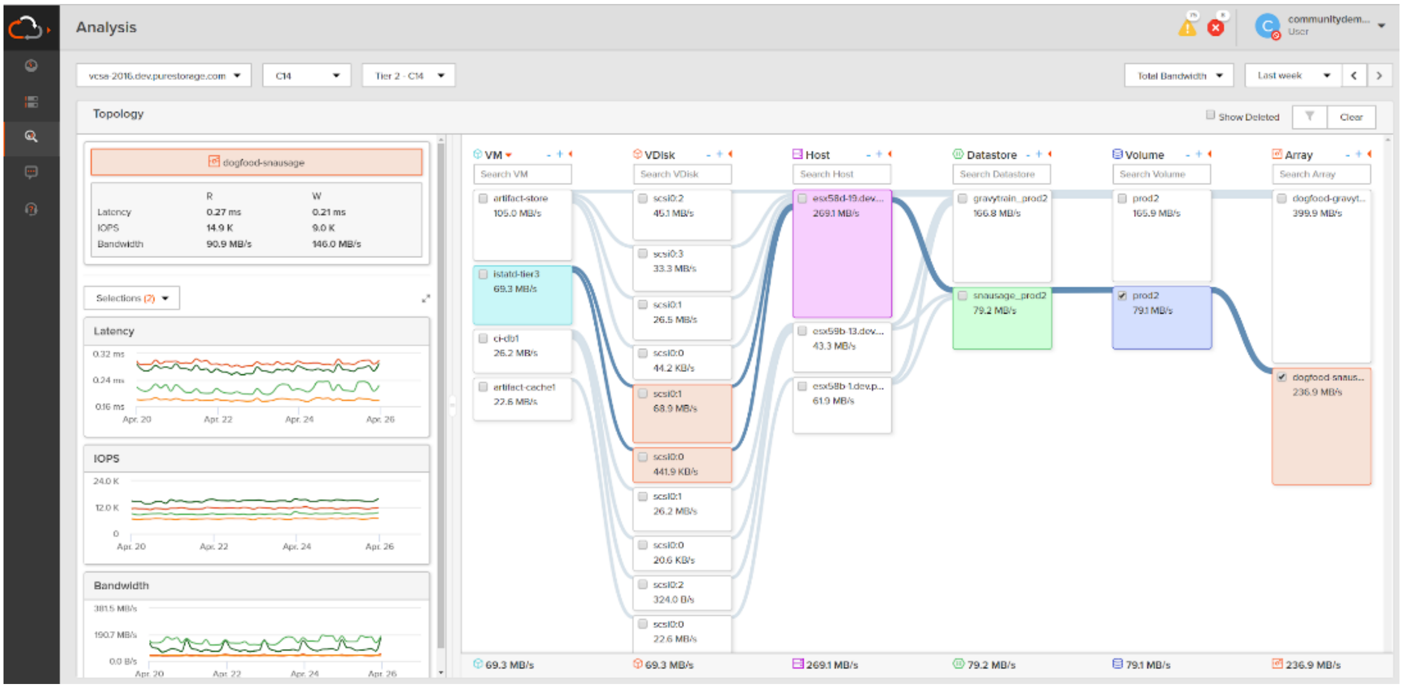

	R	W
Latency	0.27 ms	0.21 ms
IOPS	14.9 K	9.0 K
Bandwidth	90.9 MB/s	146.0 MB/s

CloudSnap

Pure portable snapshots provide simple, built-in, local and cloud protection for Pure FlashArrays. Purity Snapshots enable free movement of space-efficient copies between FlashArrays, to FlashBlade, to 3rd party NFS servers, and to the cloud. Pure's portable snapshot technology encapsulates metadata along with data into the snapshot, making the snapshot portable, so it can be offloaded from a Pure FlashArray to the cloud in a format that is recoverable to any FlashArray.

Benefits

CloudSnap is a self-backup technology built into FlashArray. It does not require the purchase of additional backup software or hardware, nor is there a need to learn and use an additional management interface. CloudSnap is natively managed via Pure FlashArray's GUI, CLI, and REST interfaces and is integrated with the Pure1 Snapshot Catalog. Since FlashArray connects to AWS via https, data is encrypted in transit and stored in an encrypted format in the S3 bucket using server side encryption. Since CloudSnap was built from scratch for FlashArray, it is deeply integrated with the Purity Operating Environment, resulting in highly efficient operation. A few examples of the efficiency of CloudSnap:

- CloudSnap preserves data compression on the wire, and in the S3 bucket, saving network bandwidth and increasing storage space efficiency.
- CloudSnap preserves data reduction across snapshots of a volume. After offloading the initial baseline snapshot of a volume, it only sends delta changes for subsequent snaps of the same volume. The snapshot differencing engine runs within the Purity Operating Environment in FlashArray and uses a local copy of the previous snapshot to compute the delta changes. Therefore, there is no back and forth network traffic between FlashArray and the cloud to compute deltas between snapshots, further reducing network congestion and data access costs in the cloud.
- CloudSnap knows which data blocks already exist on FlashArray, so during restores it only pulls back missing data blocks to rebuild the complete snapshot on FlashArray. In addition, CloudSnap uses dedupe preserving restores, so when data is restored from the offload target to FlashArray, it is deduped to save space on FlashArray.

The highly efficient operation of CloudSnap provides the following benefits:

- Less space is consumed in the S3 bucket
- Network utilization is minimized
- Backup windows are much smaller
- Data retrieval costs from the S3 bucket are lower

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.
- 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.
- Remote 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 2022, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Remote 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 Citrix RDS server virtual machines and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft Office, is shared by multiple users simultaneously. Each user receives an application " session" and works in an isolated memory space.
- Streamed Applications: Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly, but the resources may only be available while they are connected to the network.
- Local Virtual Desktop: A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a 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 of cloud applications, for example, SalesForce.com. This application and data analysis is beyond the scope of this Cisco Validated Design but should not be omitted from the planning process. There are a variety of third-party tools available to assist organizations with this crucial exercise.

Project Planning and Solution Sizing Sample Questions

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 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 database?
- Is user profile management (for example, non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

Hypervisor Selection

VMware vSphere 8.0 has been selected as the hypervisor for this **Citrix Virtual Apps and Desktops** and Remote Server Desktop Hosted (RDSH) Sessions deployment.

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

Storage Considerations

Boot from SAN

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

Pure Storage FlashArray Considerations

Make sure Each FlashArray Controller is connected to BOTH storage fabrics (A/B).

Within Purity, it’s best practice to map Hosts to Host Groups and then Host Groups to Volumes, this ensures the Volume is presented on the same LUN ID to all hosts and allows for simplified management of ESXi Clusters across multiple nodes.

How big should a Volume be? With the Purity Operating Environment, we remove the complexities of aggregates, RAID groups, and so on. When managing storage, you just create a volume based on the size required, availability and performance are taken care of through RAID-HD and DirectFlash Software. As an administrator you can create 1 10TB volume or 10 1TB Volumes and their performance/availability will be the same, so instead of creating volumes for availability or performance you can think about recoverability, manageability, and administrative considerations. For example, what data do I want to present to this application or what data do I want to store together so I can replicate it to another site/system/cloud, and so on.

Port Connectivity

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

16/32Gb Fiber Channel support (N-2 support) – Pure Storage offer up to 32Gb FC support on the latest FlashArray//X series arrays. Always make sure the correct number of HBAs and the speed of SFPs are included in the original FlashArray BOM.

Oversubscription

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

Topology

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

Pure Storage FlashArray Best Practices for VMware vSphere 8.0

The following Pure Storage best practices for VMware vSphere should be followed as part of a design:

- FlashArray Volumes are automatically presented to VMware vSphere using the Round Robin Path Selection Policy (PSP) and appropriate vendor Storage Array Type Plugin (SATP) for vSphere 8.0.
- vSphere 8.0 also uses the Latency SATP that was introduced in vSphere 6.7U1 (This replaces the I/O Operations Limit of 1 SATP, which was the default from vSphere 6.5U1).
- When using iSCSI connected FlashArray volumes, it is recommended to set DelayedAck to false (disabled) and LoginTimeout to 30 seconds. Jumbo Frames are optional when using iSCSI.
- For VMFS-6, keep automatic UNMAP enabled.
- DataMover.HardwareAcceleratedMove, DataMover.HardwareAcceleratedInit, and VMFS3.HardwareAcceleratedLocking should all be enabled.
- Ensure all ESXi hosts are connected to both FlashArray controllers. A minimum of two paths to each. Aim for total redundancy.
- Install VMware tools or Open VM tools whenever possible.
- Queue depths should be left at the default. Changing queue depths on the ESXi host is a tweak and should only be examined if a performance problem (high latency) is observed.
- When mounting snapshots, use the ESXi resignature option and avoid force-mounting.

- Configure Host Groups on the FlashArray identically to clusters in vSphere. For example, if a cluster has four hosts in it, create a corresponding Host Group on the relevant FlashArray with exactly those four hosts—no more, no less.
- When possible, use Paravirtual SCSI adapters for virtual machines.
- Atomic Test and Set (ATS) is required on all Pure Storage volumes. This is a default configuration, and no changes should normally be needed.

For more information about the VMware vSphere Pure Storage FlashArray Best Practices, go to: https://support.purestorage.com/Solutions/VMware_Platform_Guide/001VMwareBestPractices/hhhWeb_Guide%3A_FlashArray_VMware_Best_Practices

Citrix Virtual Apps and Desktops Design Fundamentals

An ever growing and diverse base of user devices, complexity in 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.

Citrix Virtual Apps and Desktops integrates Hosted Shared and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, and manageable solution for delivering Windows applications and desktops as a service.

You can select applications from an easy-to-use "store" that is accessible from tablets, smartphones, PCs, Macs, and thin clients. Virtual Apps and Desktops delivers a native touch-optimized experience with HDX high-definition performance, even over mobile networks.

Machine Catalogs

Collections of identical virtual machines or physical computers are managed as a single entity called a Machine Catalog. In this CVD, virtual machine provisioning relies on Citrix Provisioning Services and Machine Creation Services to make sure that the machines in the catalog are consistent. In this CVD, machines in the Machine Catalog are configured to run either a Multi-session OS VDA (Windows Server OS) or a Single-session OS VDA (Windows Desktop OS).

Delivery Groups

To deliver desktops and applications to users, you create a Machine Catalog and then allocate machines from the catalog to users by creating Delivery Groups. Delivery Groups provide desktops, applications, or a combination of desktops and applications to users. Creating a Delivery Group is a flexible way of allocating machines and applications to users. In a Delivery Group, you can:

- Use machines from multiple catalogs
- Allocate a user to multiple machines
- Allocate multiple users to one machine

As part of the creation process, you specify the following Delivery Group properties:

- Users, groups, and applications allocated to Delivery Groups
- Desktop settings to match users' needs
- Desktop power management options

Figure 15 illustrates how users access desktops and applications through machine catalogs and delivery groups.

Figure 15. Access Desktops and Applications through Machine Catalogs and Delivery Groups

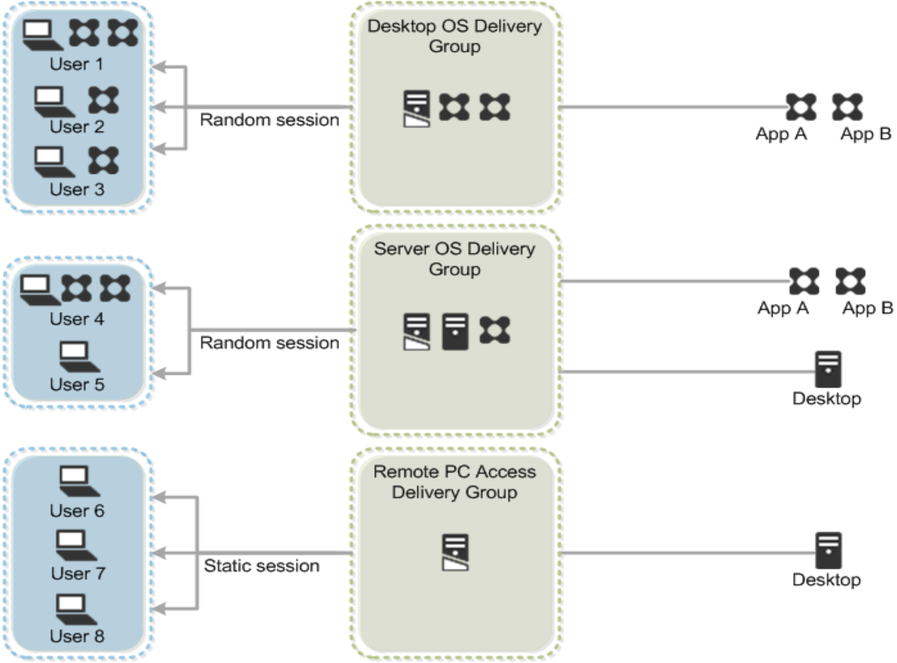

Citrix Provisioning Services

Citrix Virtual Apps and Desktops can be deployed with or without Citrix Provisioning Services (PVS). The advantage of using Citrix PVS is that it allows virtual machines to be provisioned and re-provisioned in real-time from a single shared-disk image. In this way administrators can completely eliminate the need to manage and patch individual systems and reduce the number of disk images that they manage, even as the number of machines continues to grow, simultaneously providing the efficiencies of a centralized management with the benefits of distributed processing.

The Provisioning Services solution's infrastructure is based on software-streaming technology. After installing and configuring Provisioning Services components, a single shared disk image (vDisk) is created from a device's hard drive by taking a snapshot of the OS and application image, and then storing that image as a vDisk file on the network. A device that is used during the vDisk creation process is the Master target device. Devices or virtual machines that use the created vDisks are called target devices.

When a target device is turned on, it is set to boot from the network and to communicate with a Provisioning Server. Unlike thin-client technology, processing takes place on the target device.

Figure 16. **Citrix Provisioning Services Functionality**

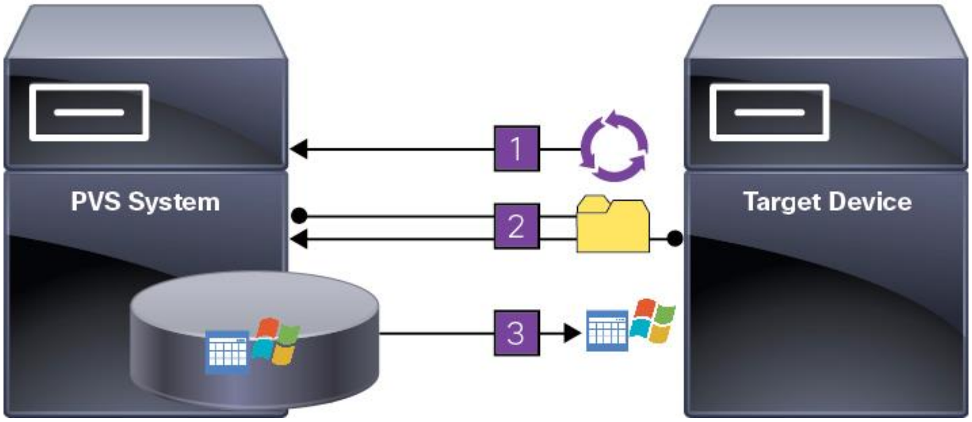

The target device downloads the boot file from a Provisioning Server (Step 2) and boots. Based on the boot configuration settings, the appropriate vDisk is mounted on the Provisioning Server (Step 3). The vDisk software is then streamed to the target device as needed, appearing as a regular hard drive to the system.

Instead of immediately pulling all the vDisk contents down to the target device (as with traditional imaging solutions), the data is brought across the network in real-time as needed. This approach allows a target device to get a completely new operating system and set of software in the time it takes to reboot. This approach dramatically decreases the amount of network bandwidth required and making it possible to support a larger number of target devices on a network without impacting performance

Citrix PVS can create desktops as Pooled or Private:

- Pooled Desktop: A pooled virtual desktop uses Citrix PVS to stream a standard desktop image to multiple desktop instances upon boot.
- Private Desktop: A private desktop is a single desktop assigned to one distinct user.
- The alternative to Citrix Provisioning Services for pooled desktop deployments is Citrix Machine Creation Services (MCS), which is integrated with the Virtual Apps and Desktops Studio console.

Locating the PVS Write Cache

When considering a PVS deployment, there are some design decisions that need to be made regarding the write cache for the target devices that leverage provisioning services. The write cache is a cache of all data that the target device has written. If data is written to the PVS vDisk in a caching mode, the data is not written back to the base vDisk. Instead, it is written to a write cache file in one of the following locations:

- Cache in device RAM. Write cache can exist as a temporary file in the target device's RAM. This provides the fastest method of disk access since memory access is always faster than disk access.
- Cache in device RAM with overflow on hard disk. This method uses VHDX differencing format and is only available for Windows 10 and Server 2008 R2 and later. When RAM is zero, the target device write cache is only written to the local disk. When RAM is not zero, the target device write cache is written to RAM first. When RAM is full, the least recently used block of data is written to the local differencing disk to accommodate newer data on RAM. The amount of RAM specified is the non-paged kernel memory that the target device will consume.
- Cache on a server. Write cache can exist as a temporary file on a Provisioning Server. In this configuration, all writes are handled by the Provisioning Server, which can increase disk I/O and network traffic. For additional security, the Provisioning Server can be configured to encrypt write cache files.

Since the write-cache file persists on the hard drive between reboots, encrypted data provides data protection in the event a hard drive is stolen.

- Cache on server persisted. This cache option allows for the saved changes between reboots. Using this option, a rebooted target device is able to retrieve changes made from previous sessions that differ from the read only vDisk image. If a vDisk is set to this method of caching, each target device that accesses the vDisk automatically has a device-specific, writable disk file created. Any changes made to the vDisk image are written to that file, which is not automatically deleted upon shutdown.

Note: In this CVD, Provisioning Server 2023 was used to manage Pooled/Non-Persistent Single-session OS Machines with “Cache in device RAM with Overflow on Hard Disk” for each virtual machine. This design enables good scalability to many thousands of desktops. Provisioning Server 2023 was used for Active Directory machine account creation and management as well as for streaming the shared disk to the hypervisor hosts.

Example Citrix Virtual Apps and Desktops Deployments

Two examples of typical Virtual Apps and Desktops deployments are as follows:

- A distributed components configuration
- A multiple site configuration

Distributed Components Configuration

You can distribute the components of your deployment among a greater number of servers or provide greater scalability and failover by increasing the number of controllers in your site. You can install management consoles on separate computers to manage the deployment remotely. A distributed deployment is necessary for an infrastructure based on remote access through NetScaler Gateway (formerly called Access Gateway).

Figure 17 shows an example of a distributed components configuration. A simplified version of this configuration is often deployed for an initial proof-of-concept (POC) deployment. The CVD described in this document deploys Citrix Virtual Apps and Desktops in a configuration that resembles this distributed component configuration shown.

Figure 17. Example of a Distributed Components Configuration

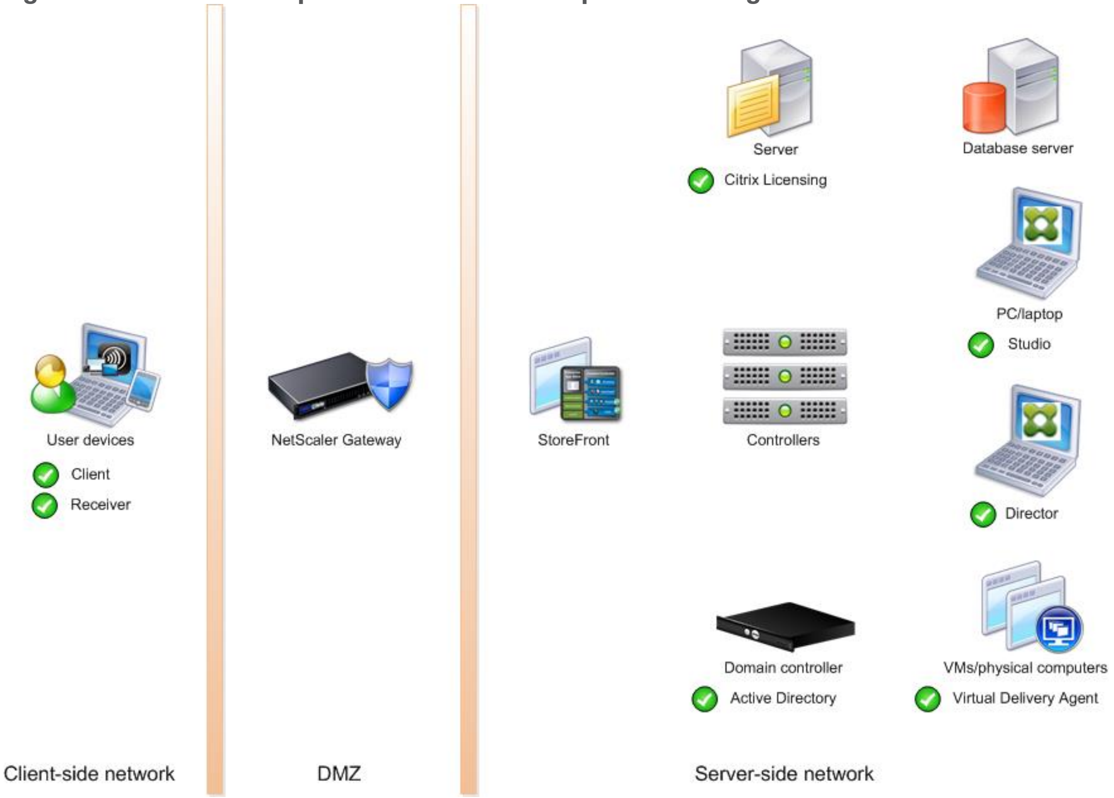

Multiple Site Configuration

If you have multiple regional sites, you can use Citrix NetScaler to direct user connections to the most appropriate site and StoreFront to deliver desktops and applications to users.

Figure 18 depicts multiple sites; a site was created in two data centers. Having two sites globally, rather than just one, minimizes the amount of unnecessary WAN traffic.

Figure 18. **Multiple Sites**

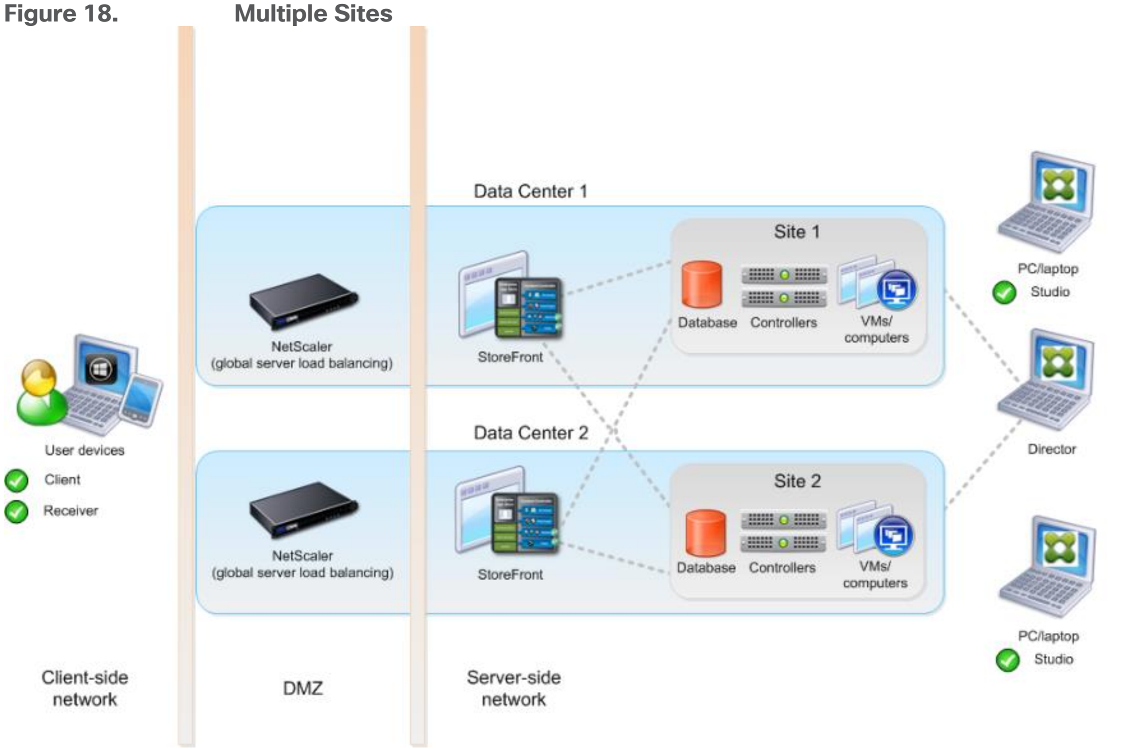

You can use StoreFront to aggregate resources from multiple sites to provide users with a single point of access with NetScaler. A separate Studio console is required to manage each site; sites cannot be managed as a single entity. You can use Director to support users across sites.

Citrix NetScaler accelerates application performance, load balances servers, increases security, and optimizes the user experience. In this example, two NetScalers are used to provide a high availability configuration. The NetScalers are configured for Global Server Load Balancing and positioned in the DMZ to provide a multi-site, fault-tolerant solution.

Note: The CVD was done based on single site and did not use NetScaler for its infrastructure and testing.

Citrix Cloud Services

Easily deliver the Citrix portfolio of products as a service. Citrix Cloud services simplify the delivery and management of Citrix technologies extending existing on-premises software deployments and creating hybrid workspace services.

- Fast: Deploy apps and desktops, or complete secure digital workspaces in hours, not weeks.
- Adaptable: Choose to deploy on any cloud or virtual infrastructure — or a hybrid of both.
- Secure: Keep all proprietary information for your apps, desktops, and data under your control.
- Simple: Implement a fully-integrated Citrix portfolio through a single-management plane to simplify administration

Designing a Virtual Apps and Desktops Environment for Different Workloads

With Citrix Virtual Apps and Desktops, 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.

Desktop Type	
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 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 this Cisco Validated Design, the following designs are included:

- Single-session OS Solution:
  - MCS: 2000 Windows 11 Virtual desktops random pooled were configured and tested
  - PVS: 2000 Windows 11 Virtual desktops random pooled were configured and tested
- Multi-session OS Solution:
  - RDS: 2600 Windows Server 2022 random pooled desktops were configured and tested

Deployment Hardware and Software

This chapter contains the following:

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

Architecture

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

It includes:

- VMware vSphere 8.0 hypervisor installed on the Cisco UCS x210C M6 compute nodes configured for stateless compute design using boot from SAN.
- Pure Storage FlashArray//X70 R3 provides the storage infrastructure required for VMware vSphere hypervisors and the VDI workload delivered by Citrix Virtual Apps and Desktops 2203.
- Cisco Intersight provides Cisco 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 and Pure storage).

Products Deployed

This CVD details the deployment of up to 2600 Multi-session OS, 2000 Single-session OS VDI users featuring the following software:

- VMware vSphere ESXi 8.0 hypervisor
- VMware vCenter 8 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 2022 and Windows 11 64-bit virtual machine Operating Systems
- Microsoft Office 2021
- Citrix Virtual Apps and Desktops 2203
- Citrix Provisioning Services 2203
- FSLogix for User profile management
- Cisco Intersight platform to deploy, maintain, and support the FlashStack components
- Cisco Intersight Assist virtual appliance to help connect the Pure Storage FlashArray and VMware vCenter with the Cisco Intersight platform

Physical Topology

FlashStack VDI with Cisco UCS M6 servers is a Fibre Channel (FC) based storage access design. Pure Storage FlashArray and Cisco UCS are connected through Cisco MDS 9132T switches and storage access utilizes the FC network. For VDI IP based file share storage access Pure Storage FlashArray and Cisco UCS are connected through Cisco Nexus C93180YC-FX switches. The physical connectivity details are explained below.

Figure 19. **FlashStack VDI – Physical Topology for FC**

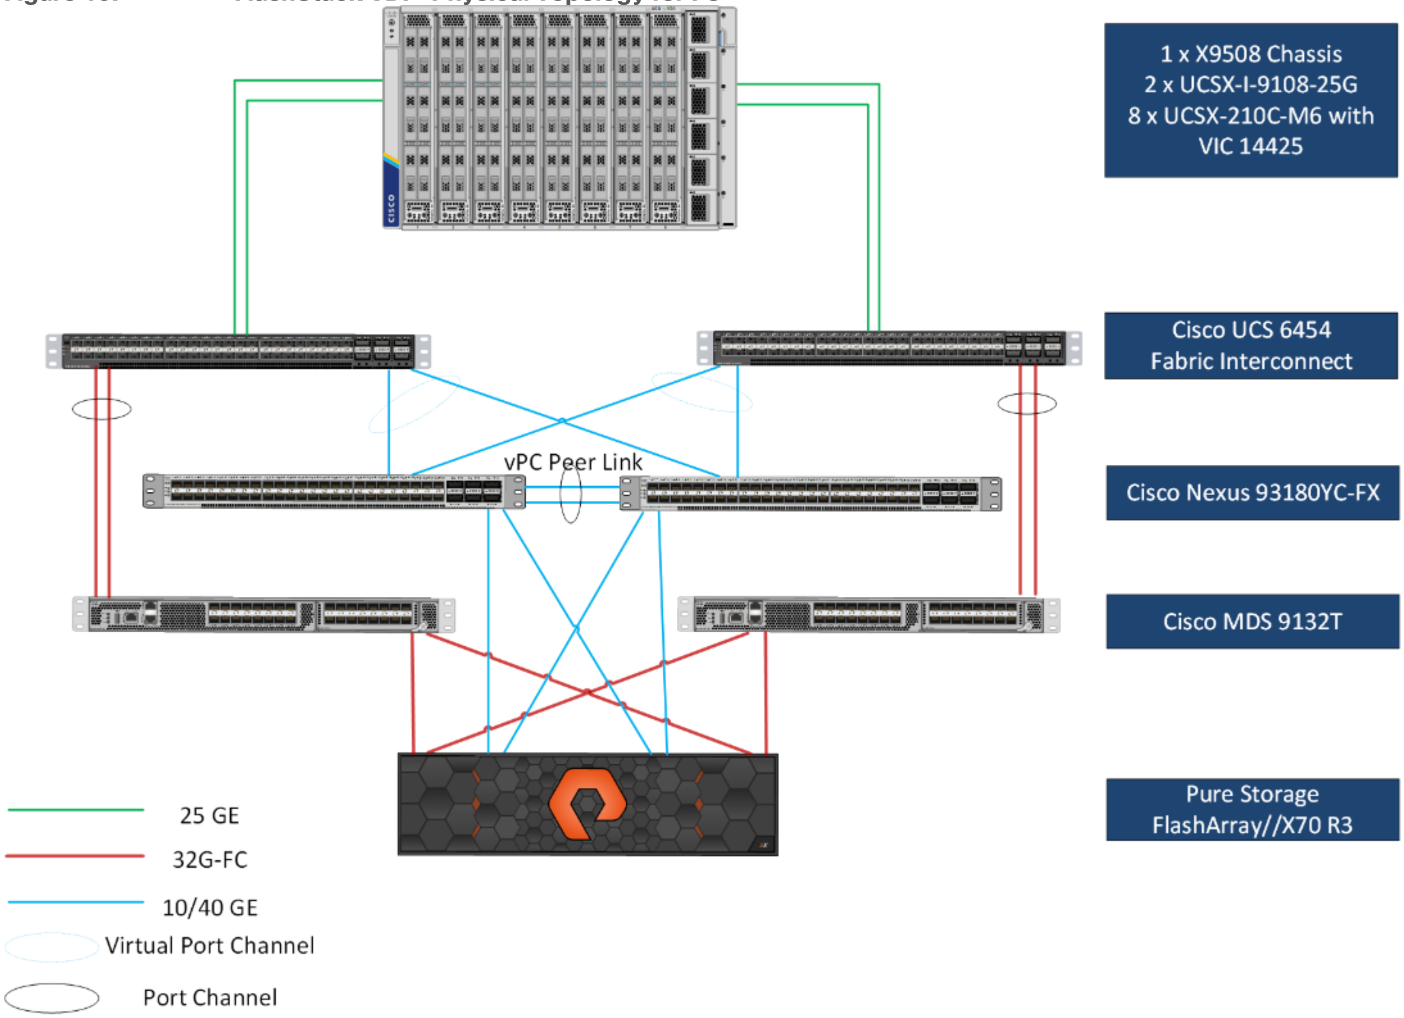

Figure 19 details the physical hardware and cabling deployed to enable this solution:

- Two Cisco Nexus 93180YC-FX Switches in NX-OS Mode.
- Two Cisco MDS 9132T 32-Gb Fibre Channel Switches.
- One Cisco UCS X950808 Chassis with two Cisco UCSX 9108 25G IF Modules.
- Eight Cisco UCS X210c M6 Compute Node s with Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 32-core processors, 1TB 3200MHz RAM, and one Cisco VIC14425 mezzanine card, providing N+1 server fault tolerance.
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with Twenty 1.92TB DirectFlash NVMe drives.

Note: The management components and LoginVSI Test infrastructure are hosted on a separate vSphere cluster and not a part of the physical topology of this solution.

Table 1 lists the software versions of the primary products installed in the environment.

Table 1. Software and Firmware Versions

Vendor	Product/Component	Version/Build/Code
Cisco	UCS Component Firmware	4.2(2d)
Cisco	UCS x210c Compute Node	5.0(2e)
Cisco	VIC 14425	5.2(2e)
Cisco	Cisco Nexus 93180YC-FX	9.3(3)
Cisco	Cisco MDS 9132T	8.4(2d)
Pure Storage	FlashArray//X70 R3	Purity//FA 6.3.3
VMware	vCenter Server Appliance	8
VMware	vSphere 8.0	8
Citrix	Citrix Virtual Apps and Desktops 2203	8.6.0-20099816
Citrix	Citrix Virtual Apps and Desktops 2203 Agent	8.6.0.20088748
Cisco	Intersight Assist	1.0.11-759
Microsoft	FSLogix 2105 HF_01	2.9.7979.62170
VMware	Tools	11.3.5.18557794

Logical Architecture

The logical architecture of the validated solution which is designed to run desktop and RDSH server VMs supporting up to 2600 users on a single chassis containing 8 blades, with physical redundancy for the blade servers for each workload type and have a separate vSphere cluster to host management services, is illustrated in Figure 20.

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

Figure 20. Logical Architecture Overview

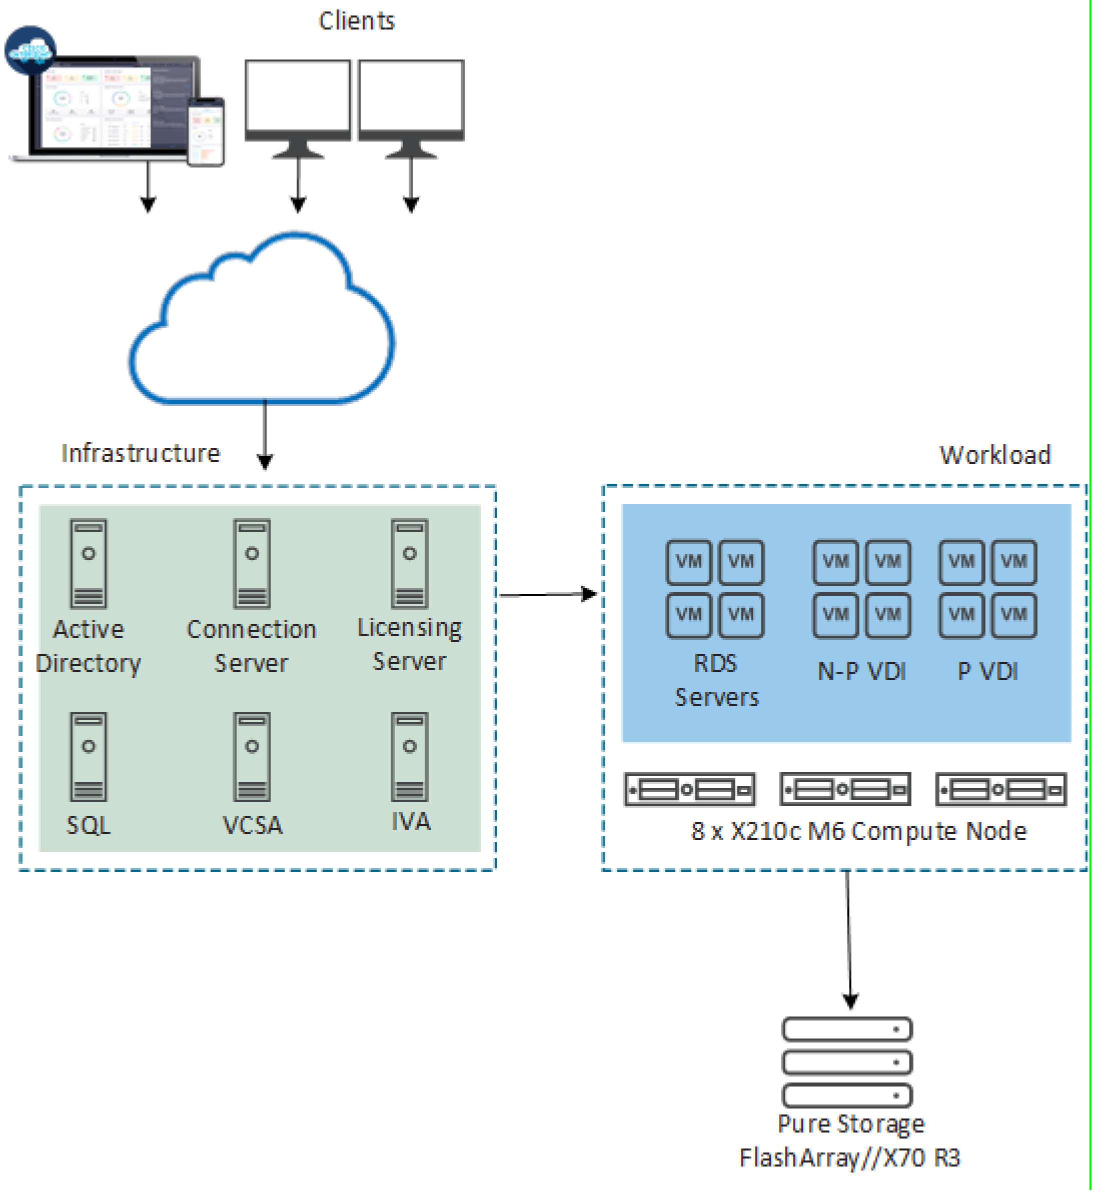

Configuration Guidelines

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

Note: This document is intended to allow the reader to configure the Citrix Virtual Apps and Desktops 2203 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 2.

Table 2. VLANs Configured in this study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt	30	In-Band management interfaces
Infra-Mgmt	31	Infrastructure Virtual Machines
VCC/VM-Network	34	RDSH, VDI Persistent and Non-Persistent
vMotion	33	VMware vMotion
OOB-Mgmt	132	Out of Band management interfaces

VSANs

Table 3 lists the two virtual SANs that were configured for communications and fault tolerance in this design.

Table 3. VSANs Configured in this study

VSAN Name	VSAN ID	VSAN Purpose
VSAN 500	500	VSAN for Primary SAN communication
VSAN 501	501	VSAN for Secondary SAN communication

Solution Configuration

This chapter contains the following:

- Solution Cabling

Solution Cabling

The following sections detail the physical connectivity configuration of the FlashStack VMware & Citrix VDI environment.

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

The tables in this section list the details for the prescribed and supported configuration of the Pure Storage FlashArray//X70 R3 storage array to the Cisco 6454 Fabric Interconnects through Cisco MDS 9132T 32-Gb FC switches.

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

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

Figure 21 details the cable connections used in the validation lab for FlashStack topology based on the Cisco UCS 6454 fabric interconnect. Four 32Gb uplinks connect as port-channels to each Cisco UCS Fabric Interconnect from the MDS switches, and a total of eight 32Gb links connect the MDS switches to the Pure FlashArray//X R3 controllers, four of these have been used for scsi-fc and the other four to support nvme-fc. Also, 25Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the Pure FlashArray//X R3 controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlashStack infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each FlashArray controller has a connection 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 21. FlashStack solution cabling diagram

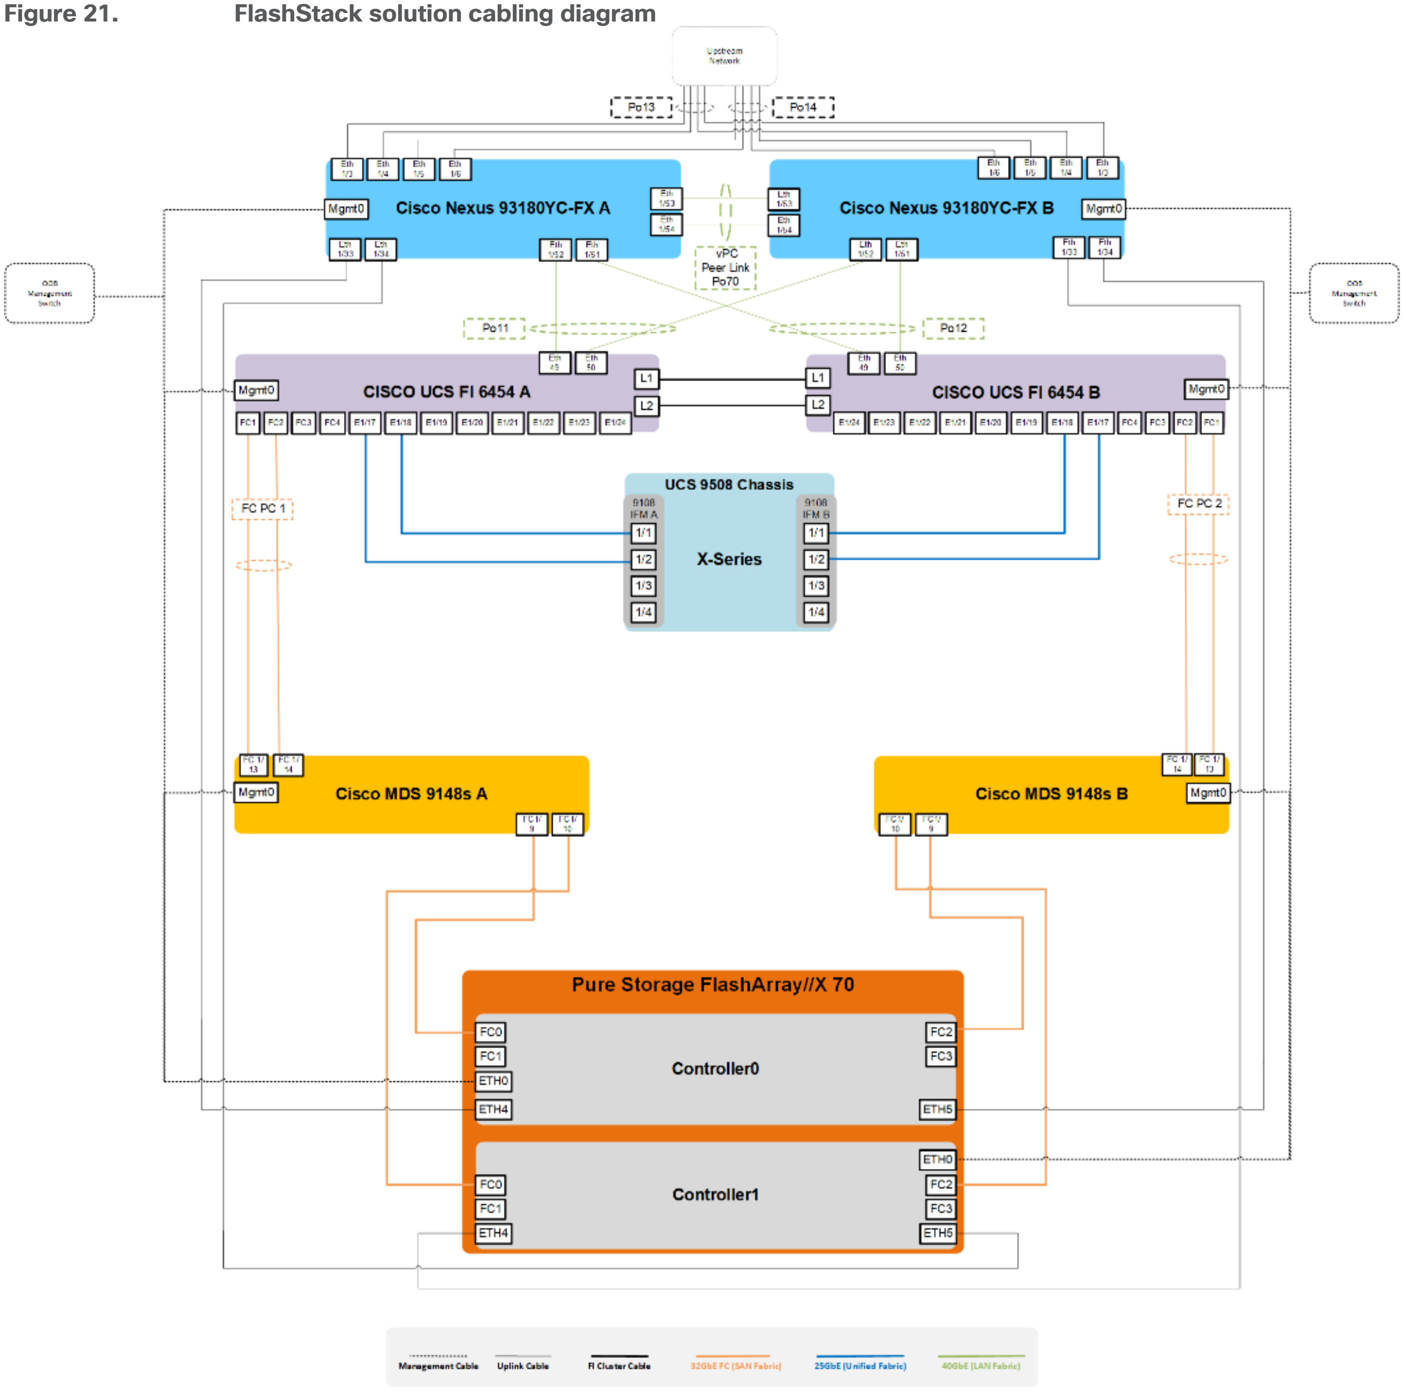

Configuration and Installation

This chapter contains the following:

FlashStack Automated Deployment with Ansible

This solution offers Ansible Playbooks that are made available from a GitHub repository that customers can access to automate the FlashStack deployment.

GitHub repository is available here: https://github.com/ucs-compute-solutions/FlashStack_IMM_Ansible.

This repository contains Ansible playbooks to configure all the components of FlashStack including:

- Cisco UCS in Intersight Managed Mode (IMM)
- Cisco Nexus and MDS Switches
- Pure FlashArray
- VMware ESXi and VMware vCenter.
- FlashStack Manual deployment

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 Intersight Managed Mode consists of the steps shown in Figure 22.

Figure 22. **Configuration Steps for Cisco Intersight Managed Mode**

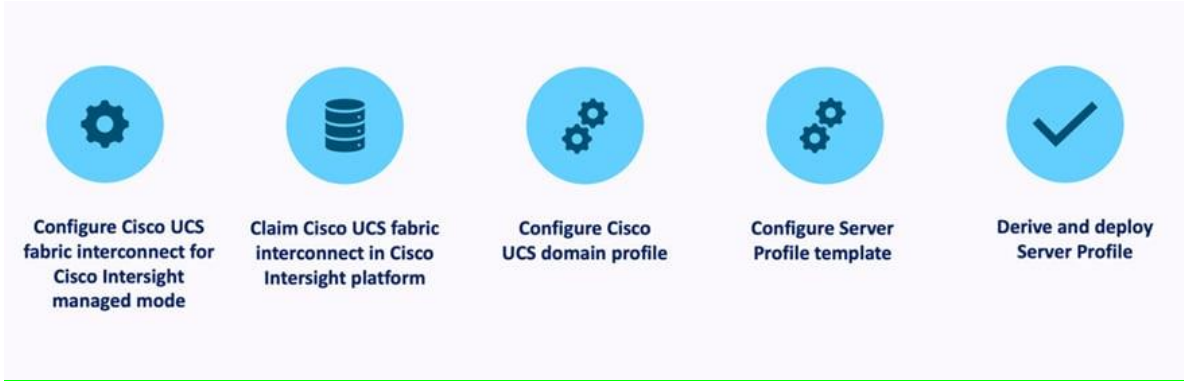

Configure Cisco UCS fabric interconnect for Cisco Intersight managed mode
Claim Cisco UCS fabric interconnect in Cisco Intersight platform
Configure Cisco UCS domain profile
Configure Server Profile template
Derive and deploy Server Profile

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 Interconnect A

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.

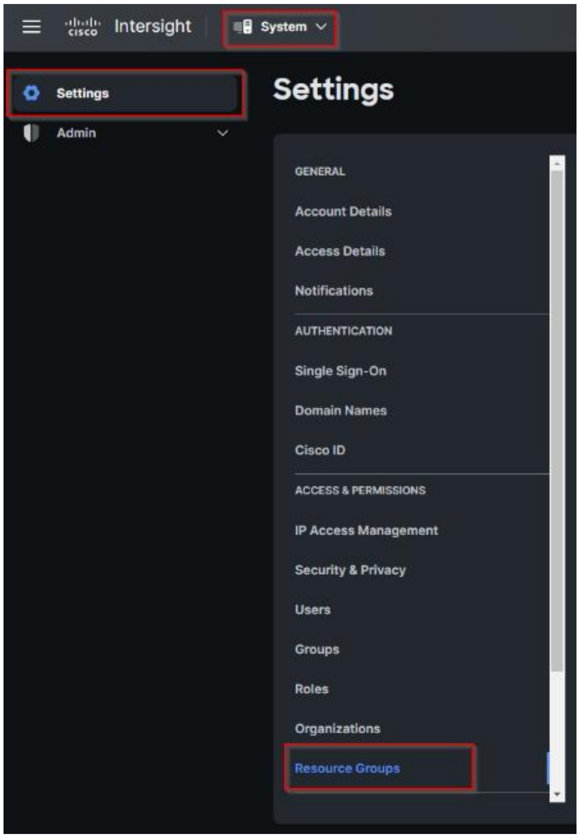

Step 8. On the Resource Groups panel click + Create Resource Group in the top-right corner.

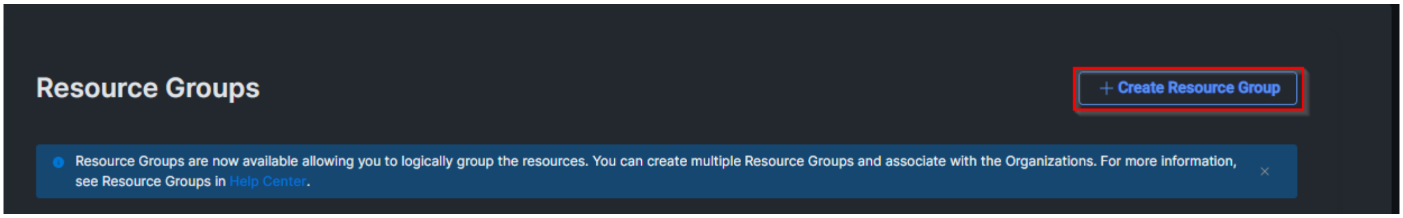

Resource Groups are now available allowing you to logically group the resources. You can create multiple Resource Groups and associate with the Organizations. For more information, see Resource Groups in Help Center.

Step 9. Provide a name for the Resource Group (for example, FlashStack-L151-DMZ).

Step 10. Click Create.

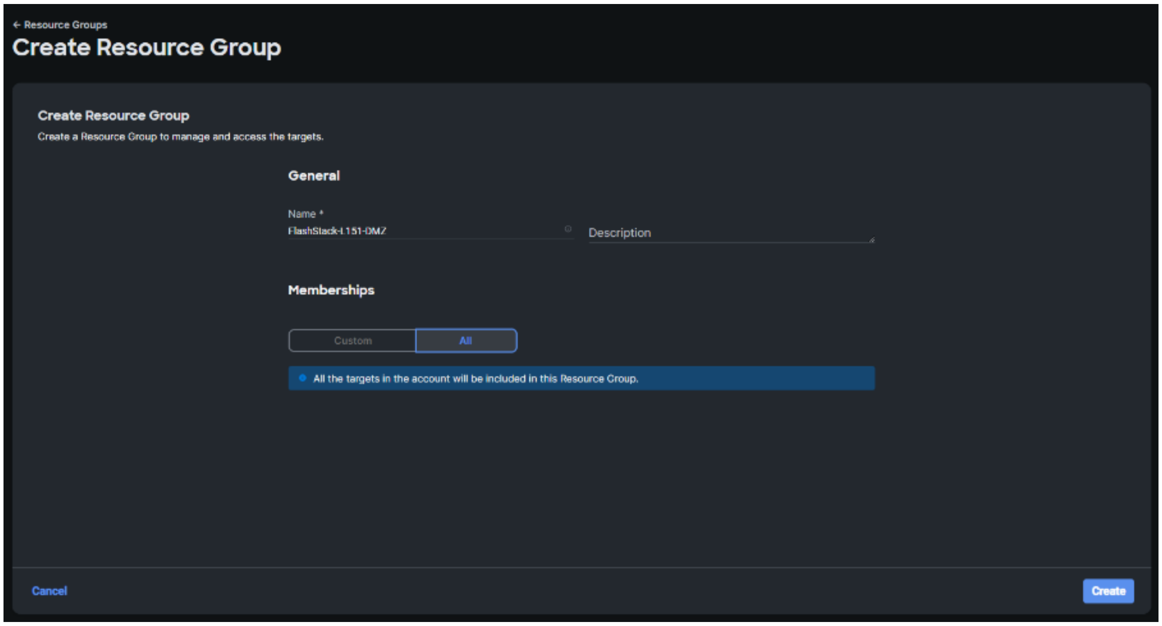

← Resource Groups

Create Resource Group

Create Resource Group

Create a Resource Group to manage and access the targets.

General

Name *: FlashStack-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.

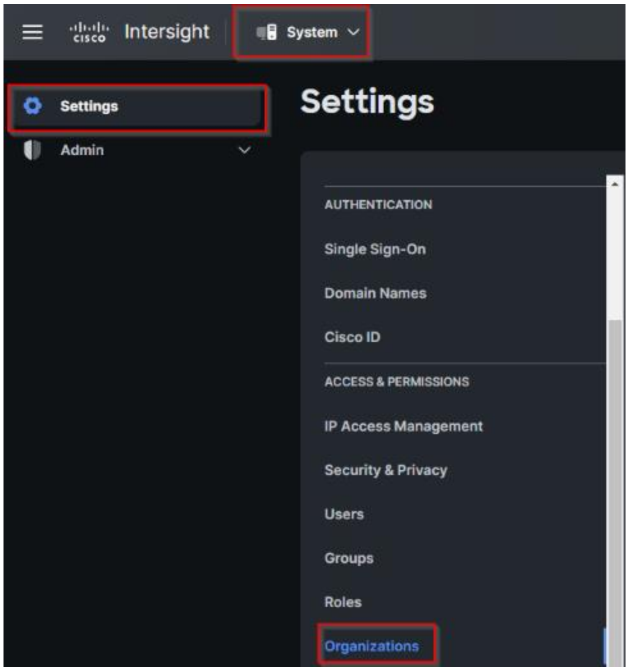

Step 12. On Organizations panel click + Create Organization in the top-right corner.

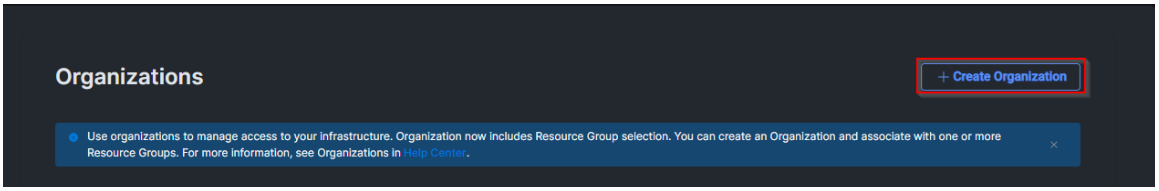

Organizations

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 (FlashStack).

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

Step 15. Click Create.

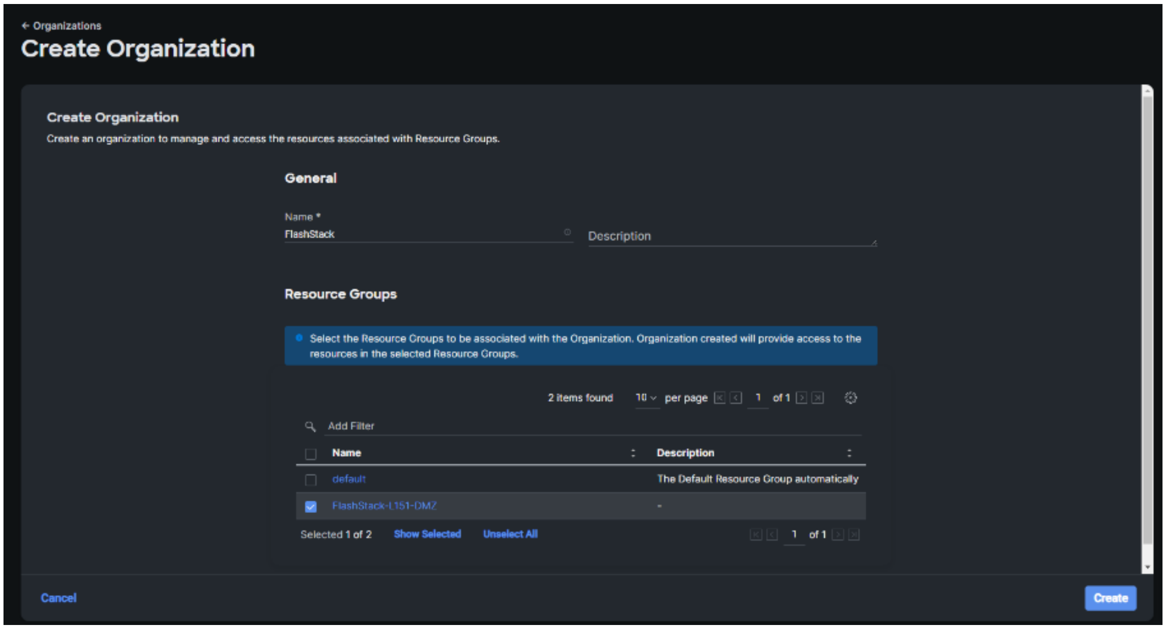

← Organizations

Create Organization

Create Organization

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

General

Name *: FlashStack

Description

Resource Groups

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

2 items found 10 per page 1 of 1

Add Filter

	Name	Description
□	default	The Default Resource Group automatically
☑	FlashStack-L151-DMZ	-

Selected 1 of 2 Show Selected Unselect All 1 of 1

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.

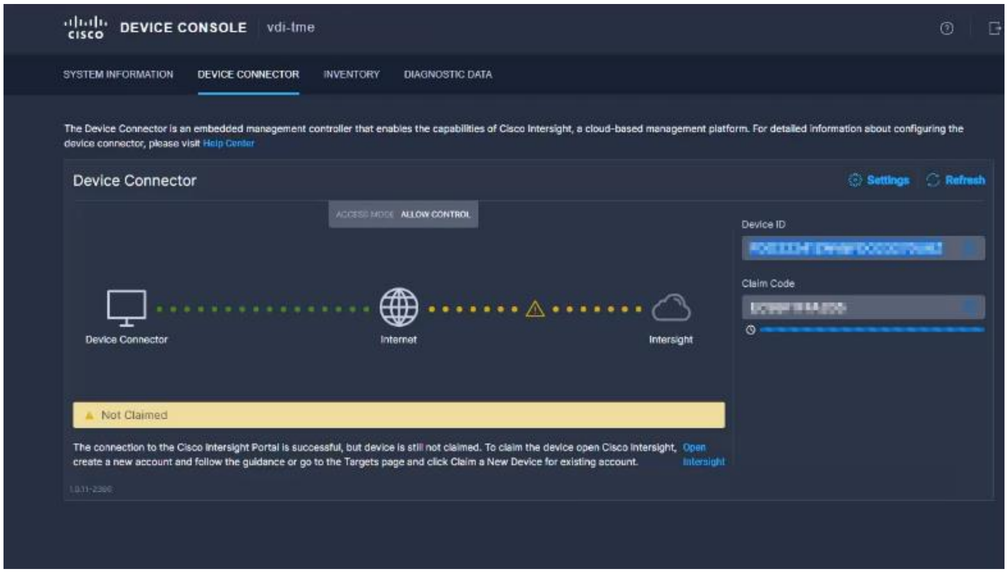

The Device Connector is an embedded management controller that enables the capabilities of Cisco Intersight, a cloud-based management platform. For detailed information about configuring the device connector, please visit Help Center

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

1.0.11-2360

Step 18. Navigate to Admin > General > Targets.

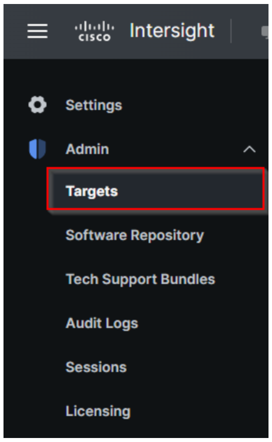

Step 19. On Targets panel click Claim a New Target in the top-right corner.

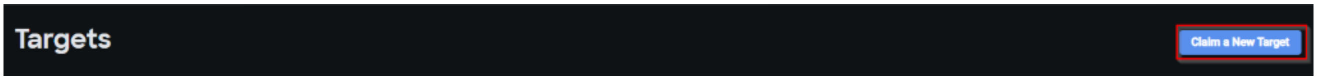

Step 20. Select Cisco UCS Domain (Intersight Managed) and click Start.

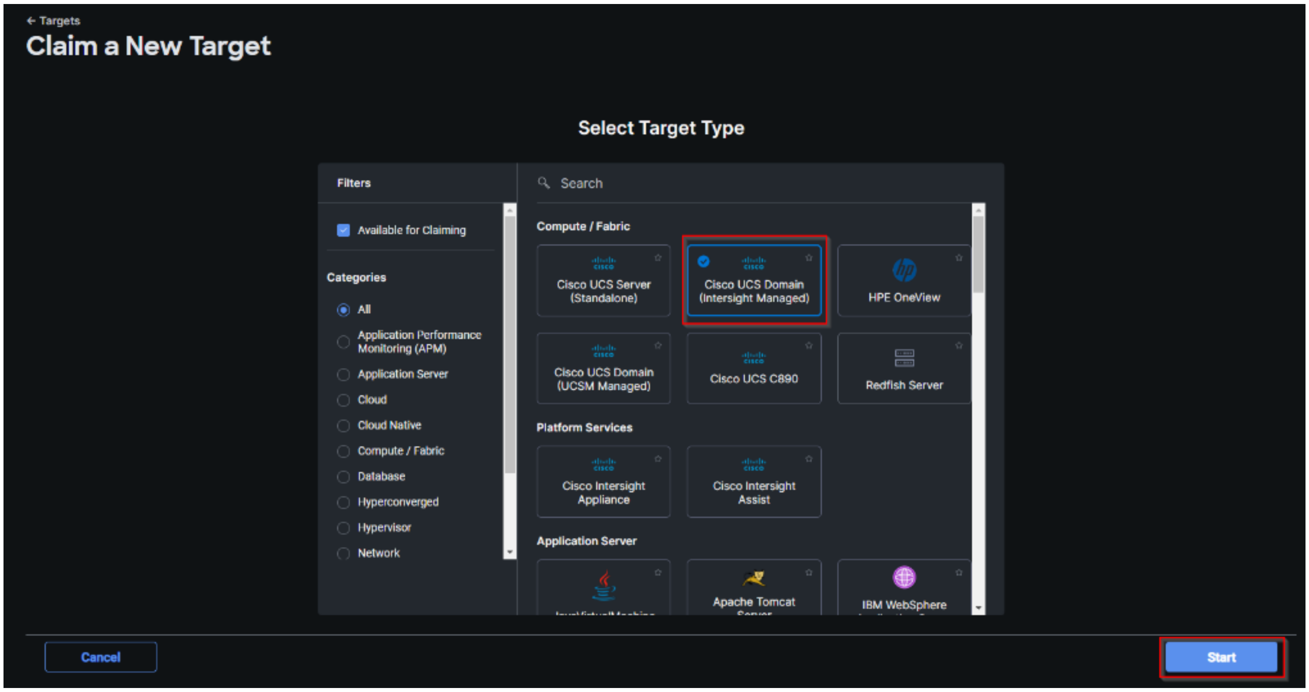

Step 21. Enter the Device ID and Claim Code captured from the Cisco UCS FI.

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

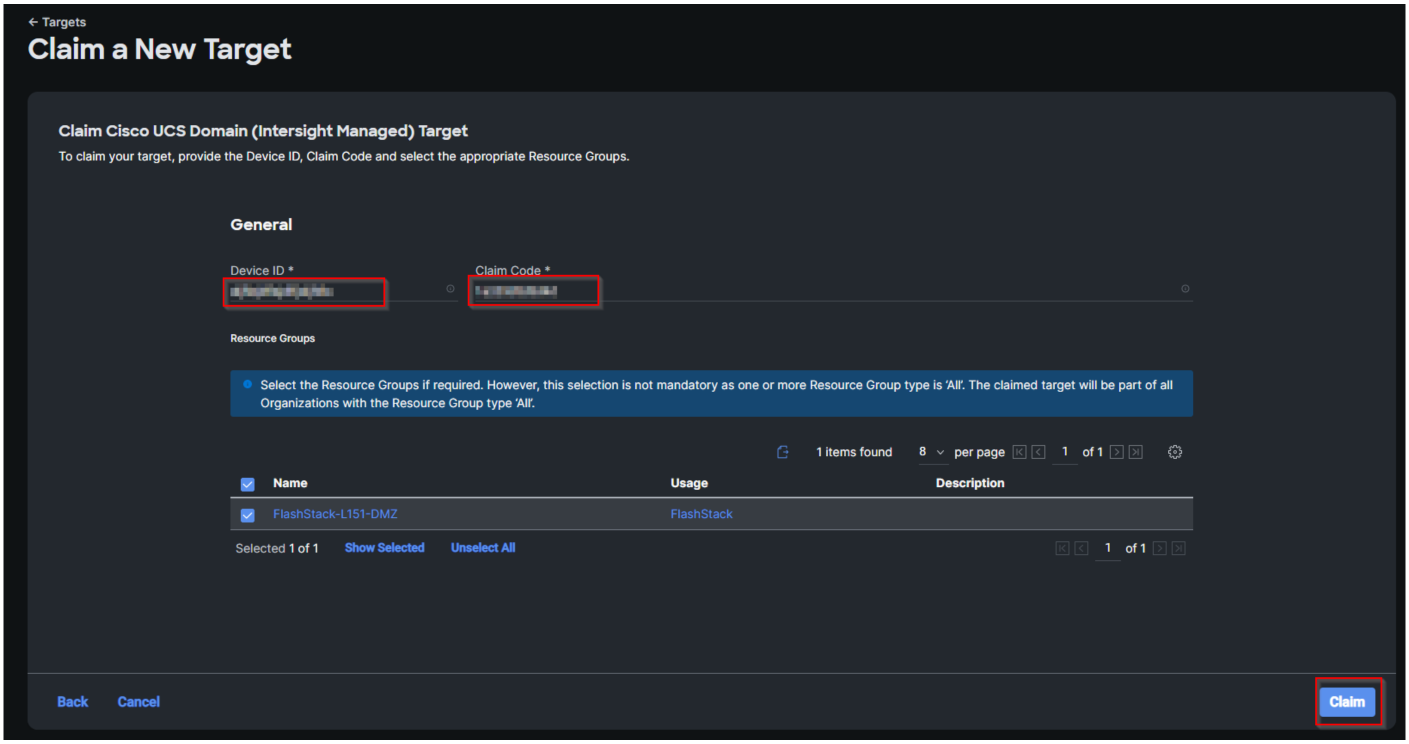

← Targets

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'.

1 items found 8 per page 1 of 1

☑	Name	Usage	Description
☑	FlashStack-L151-DMZ	FlashStack	

Selected 1 of 1 Show Selected Unselect All 1 of 1

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

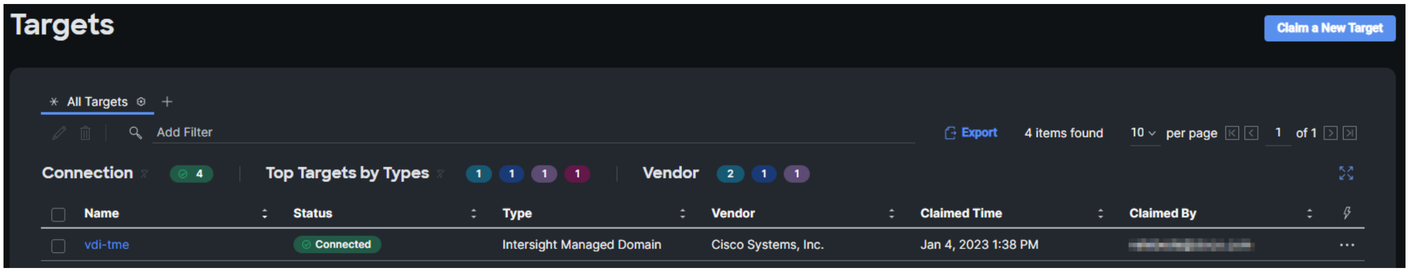

Targets

4 items found 10 per page 1 of 1

Connection 4 Top Targets by Types 1 1 1 1 Vendor 2 1 1

□	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.

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.

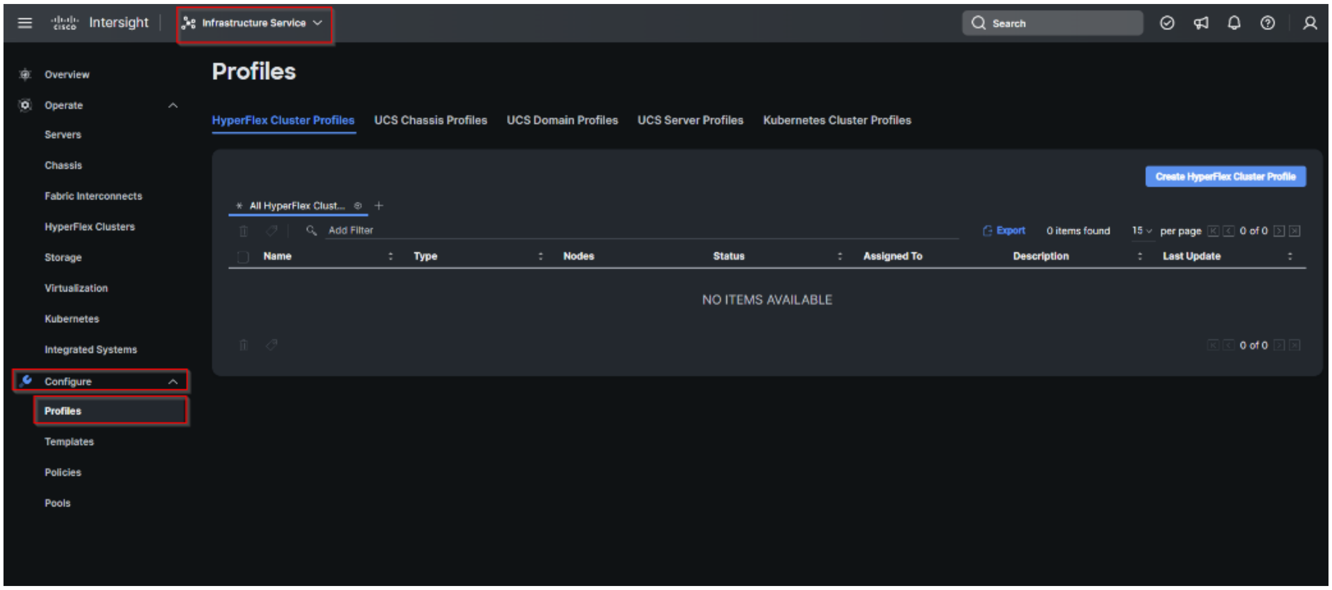

Step 2. Navigate UCS Domain Profiles tab and click Create UCS Domain Profile.

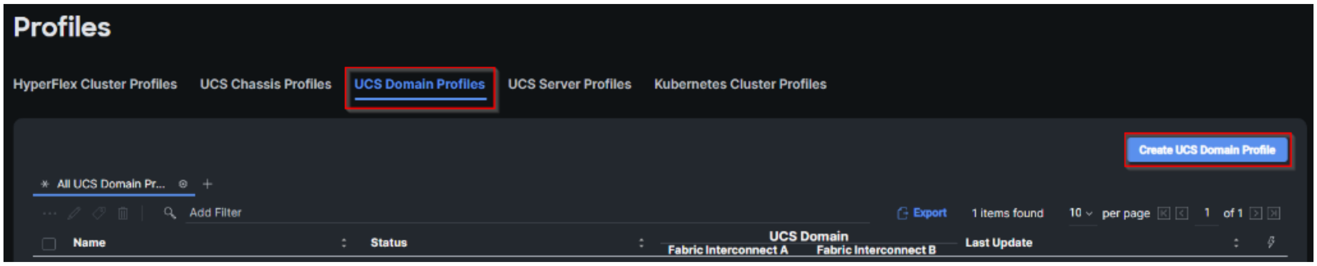

Step 3. On the Create UCS Domain Profile screen, click Start.

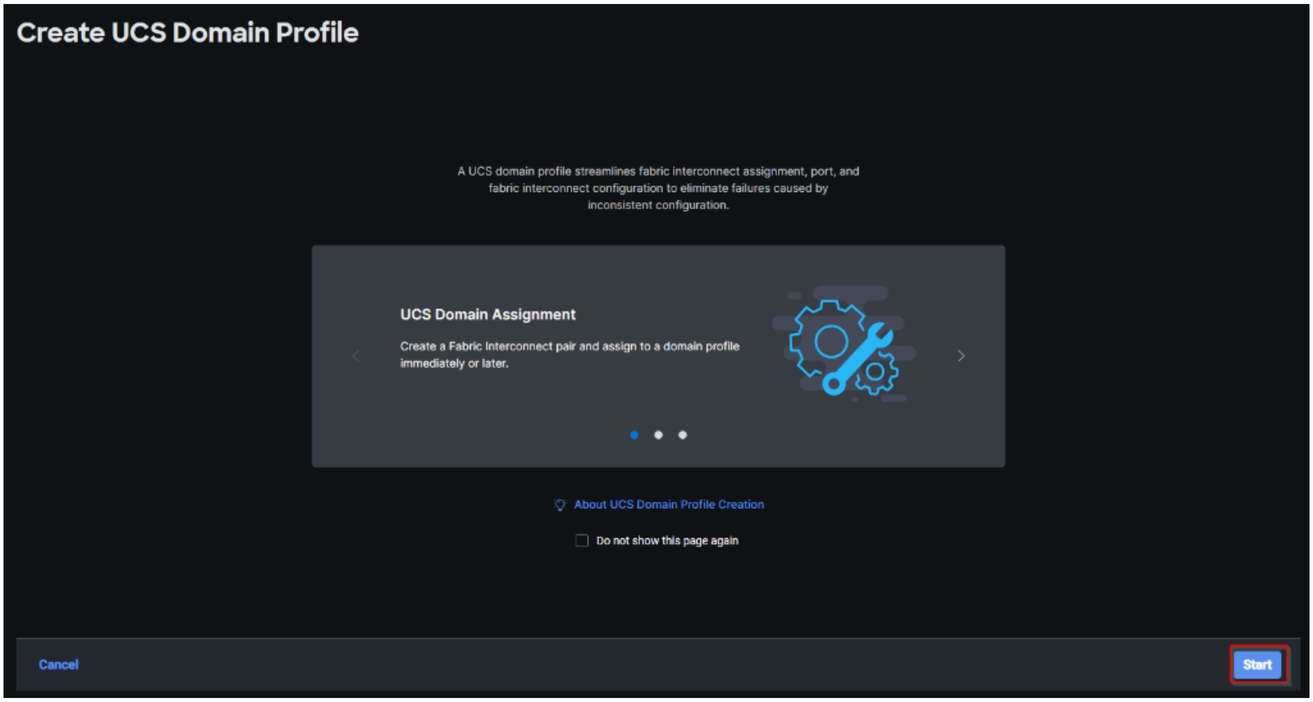
Create UCS Domain Profile

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

UCS Domain Assignment

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

About UCS Domain Profile Creation
□ Do not show this page again
Cancel
Start

Step 4. On the General page, select the organization created before and enter a name for your profile (for example, FS-L152-DMZ-K4). 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.

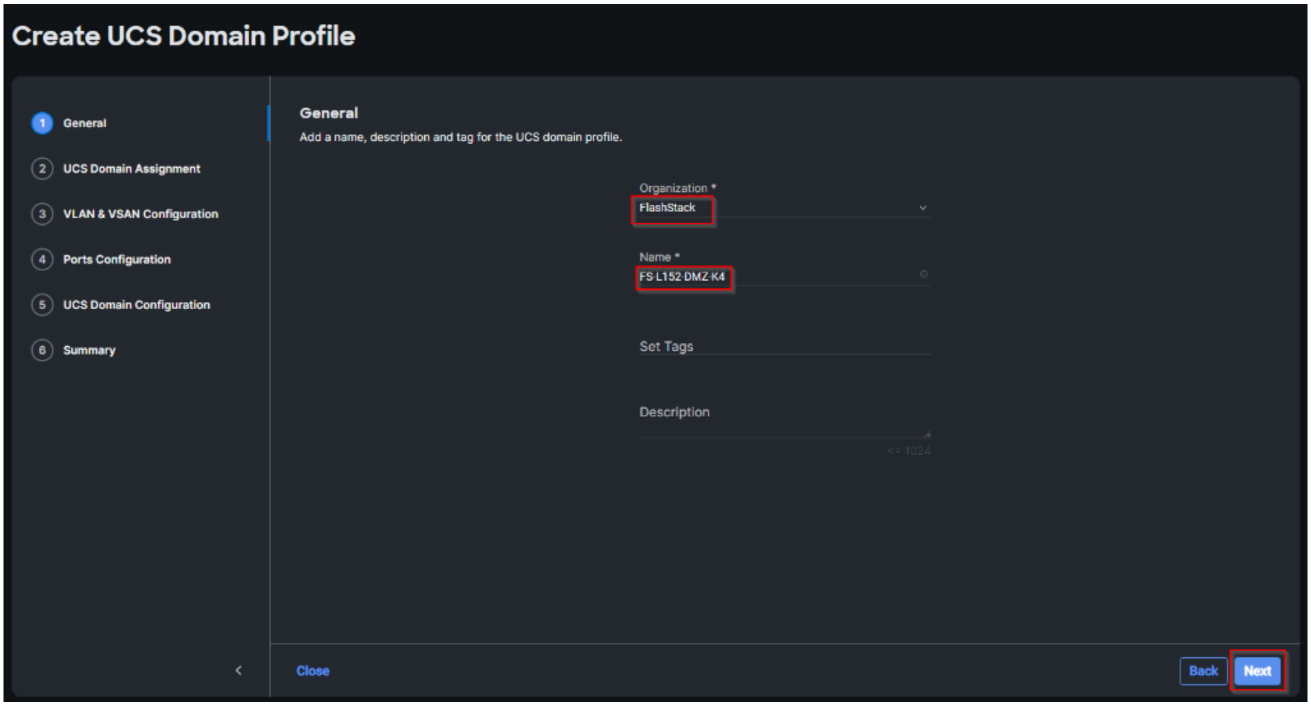

Create UCS Domain Profile

General

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

Organization *: FlashStack

Name *: FS-L152-DMZ-K4

Set Tags

Description

<= 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.

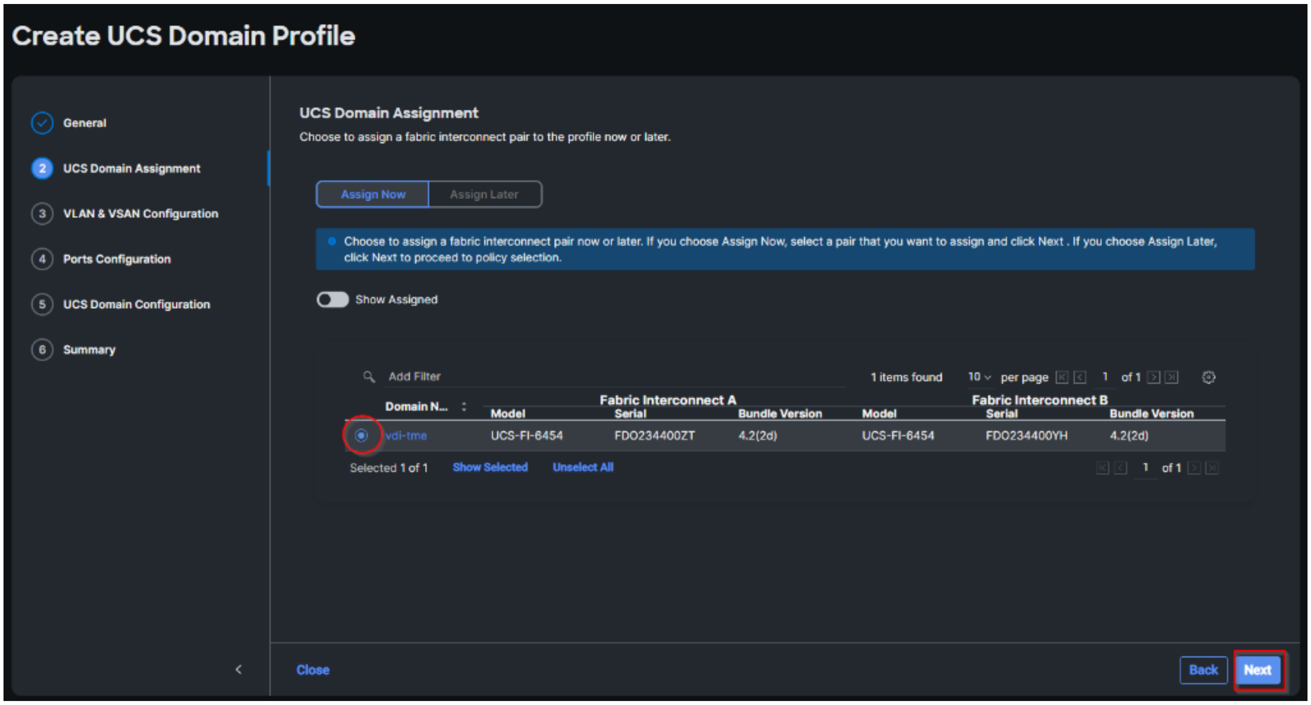

Create UCS Domain Profile

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.

		Fabric Interconnect A			Fabric Interconnect B		
	Domain N...	Model	Serial	Bundle Version	Model	Serial	Bundle Version
●	vdi-tme	UCS-FI-6454	FDO234400ZT	4.2(2d)	UCS-FI-6454	FDO234400YH	4.2(2d)

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: In this step, a single VLAN policy is created for both fabric interconnects and two individual VSAN policies are created because the VSAN IDs are unique for each fabric interconnect.

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

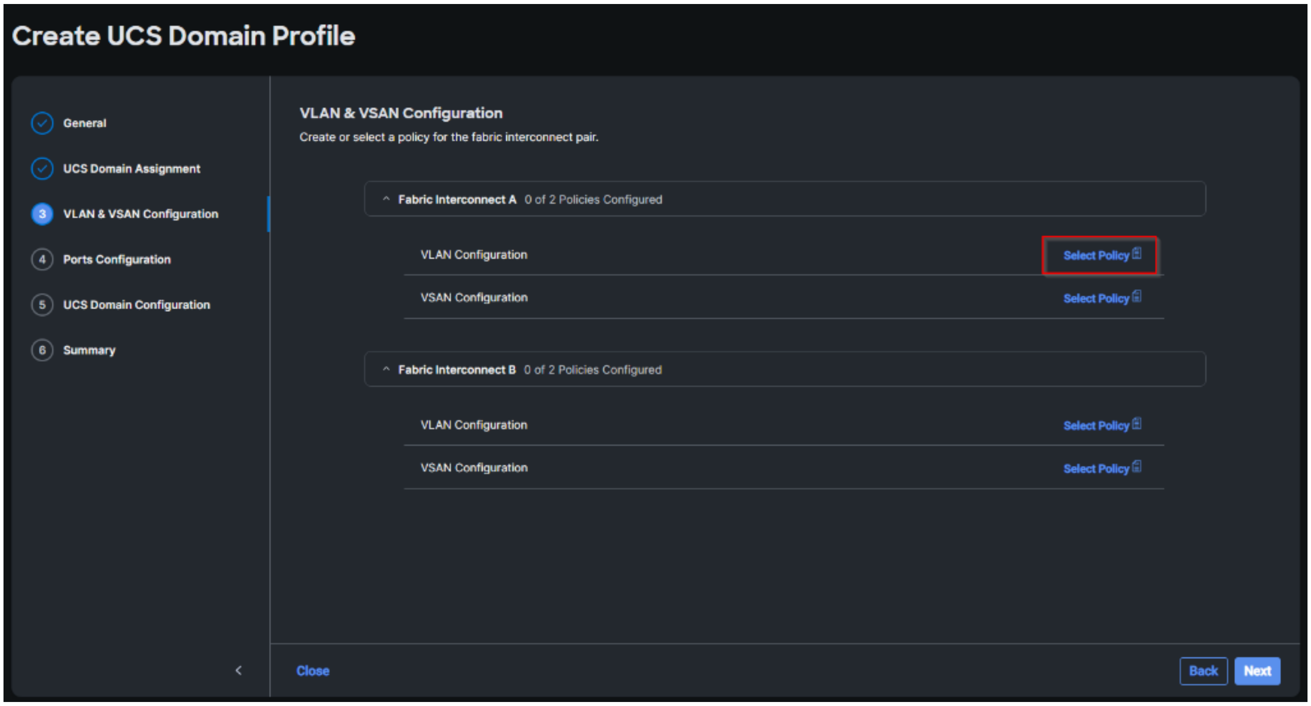

Create UCS Domain Profile

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, FS-L152-DMZ-VLAN). Click Next.

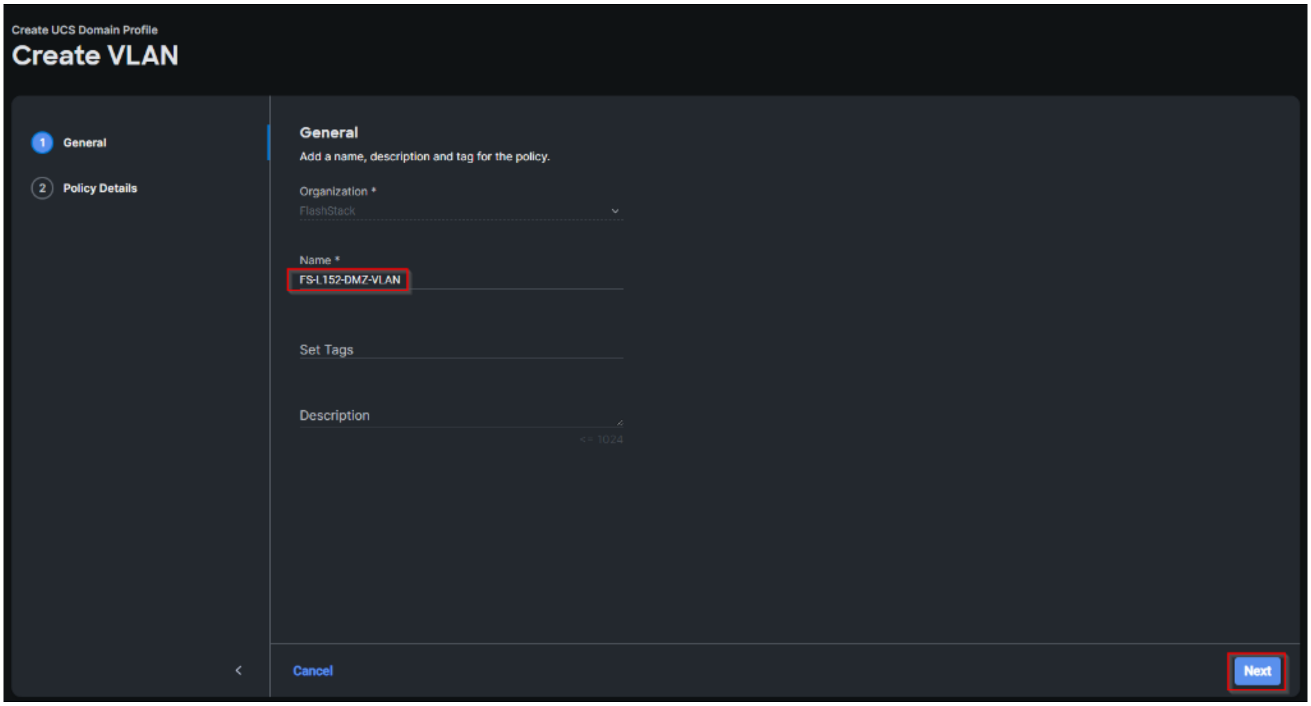

Create UCS Domain Profile

Create VLAN

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-VLAN

Set Tags

Description

<= 1024

Step 10. Click Add VLANs.

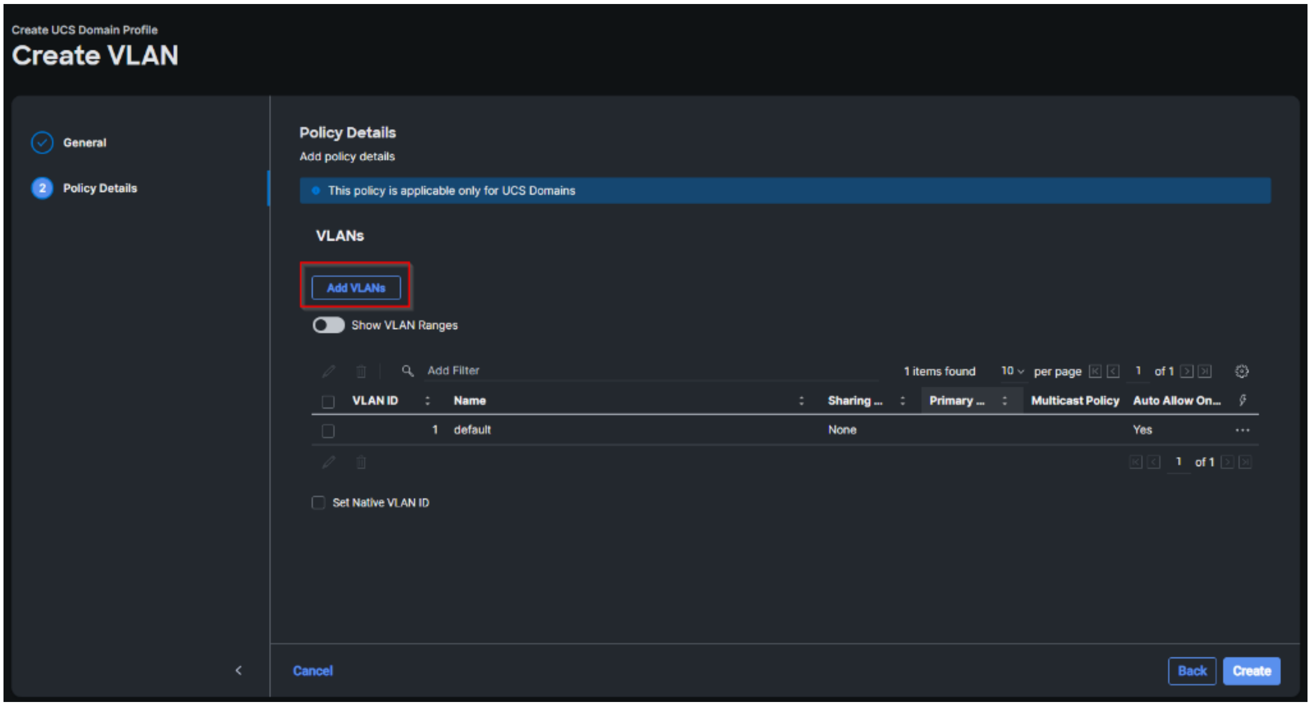

Create UCS Domain Profile

Create VLAN

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

Show VLAN Ranges

1 items found 10 per page 1 of 1

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

☐ Set Native VLAN ID

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

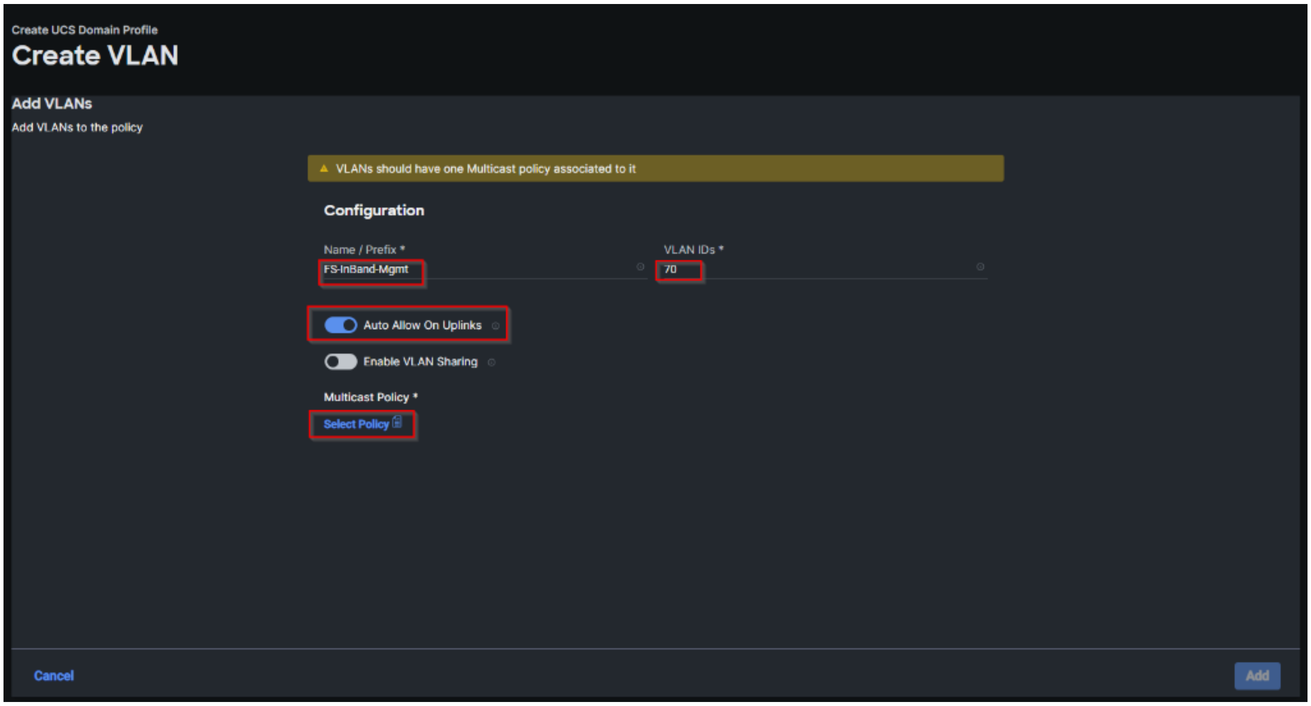

Create UCS Domain Profile

Create VLAN

Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix *: FS-InBand-Mgmt

VLAN IDs *: 70

Auto Allow On Uplinks

Enable VLAN Sharing

Multicast 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, FS-L152-DMZ-McastPol). Provide optional Description and click Next.

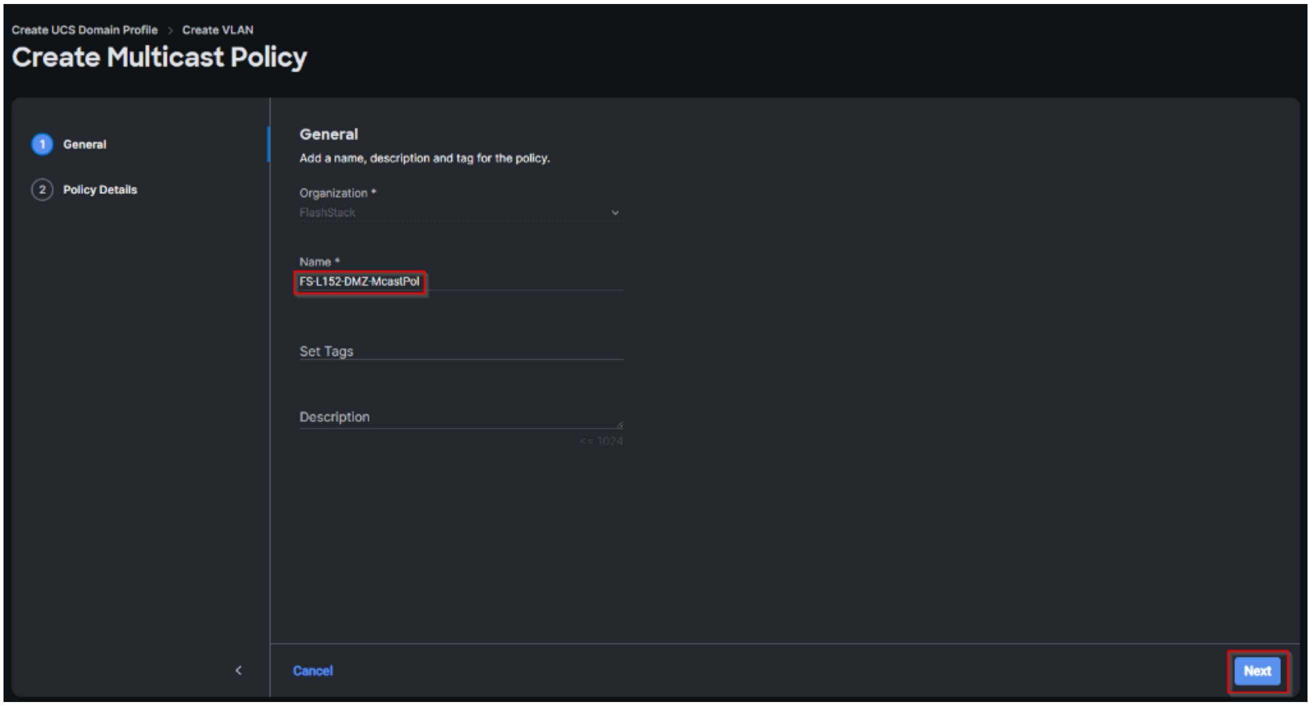

Organization * FlashStack

Name * FS-L152-DMZ-McastPol

Step 14. Leave defaults selected and click Create.

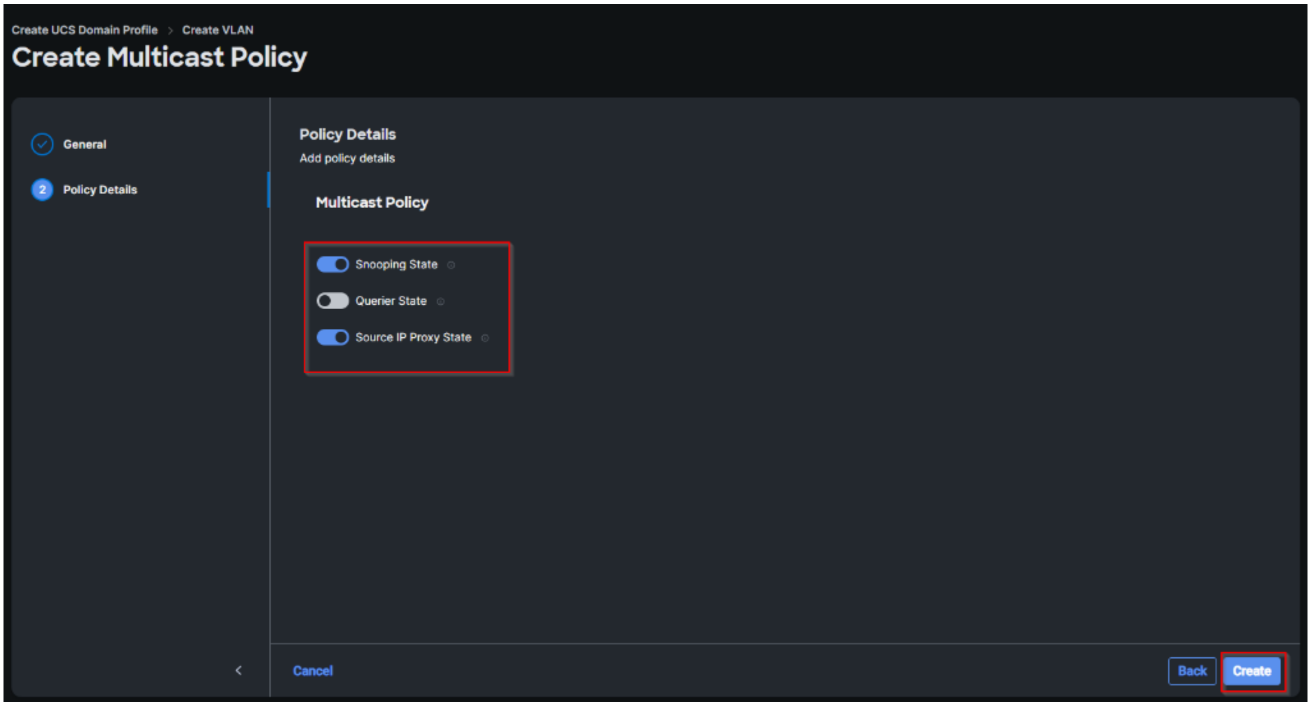

Snooping State: on

Querier State: off

Source IP Proxy State: on

Step 15. Click Add to add the VLAN.

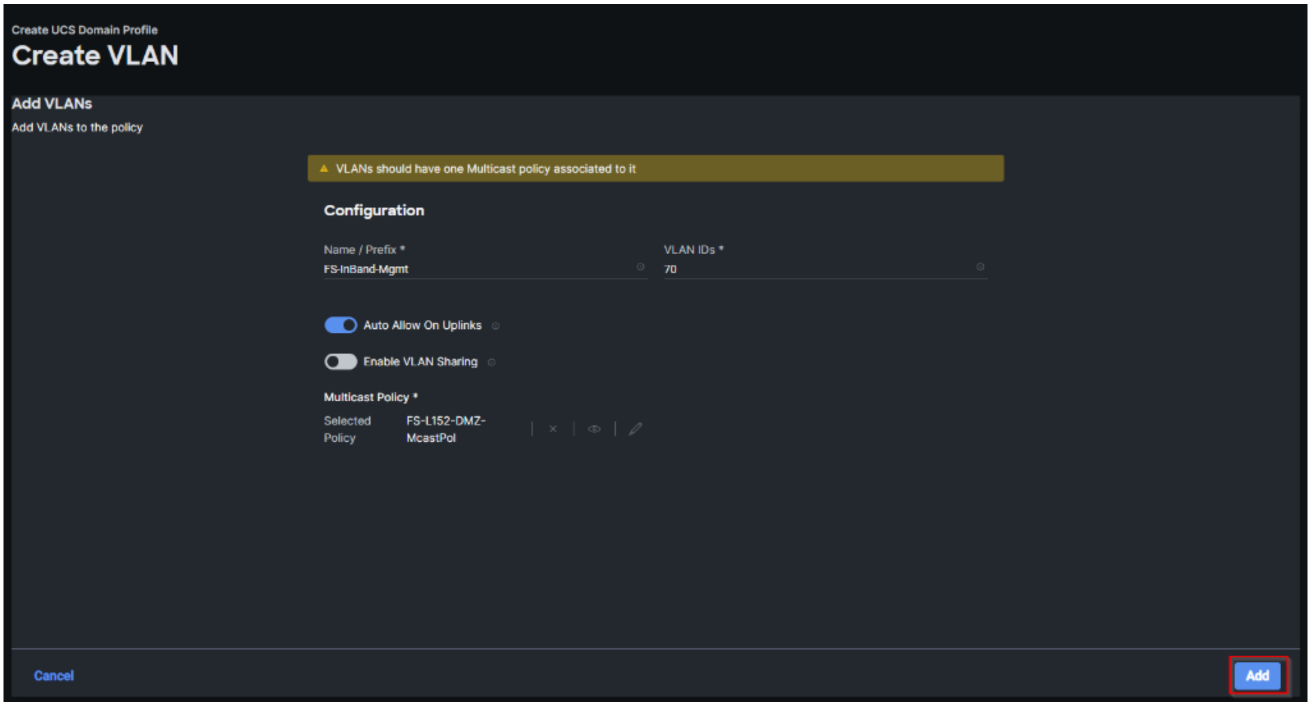

Create UCS Domain Profile

Create VLAN

Add VLANs

Add VLANs to the policy

⚠ VLANs should have one Multicast policy associated to it

Configuration

Name / Prefix * FS-InBand-Mgmt

VLAN IDs * 70

● Auto Allow On Uplinks

○ Enable VLAN Sharing

Multicast Policy *

Selected Policy FS-L152-DMZ-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:

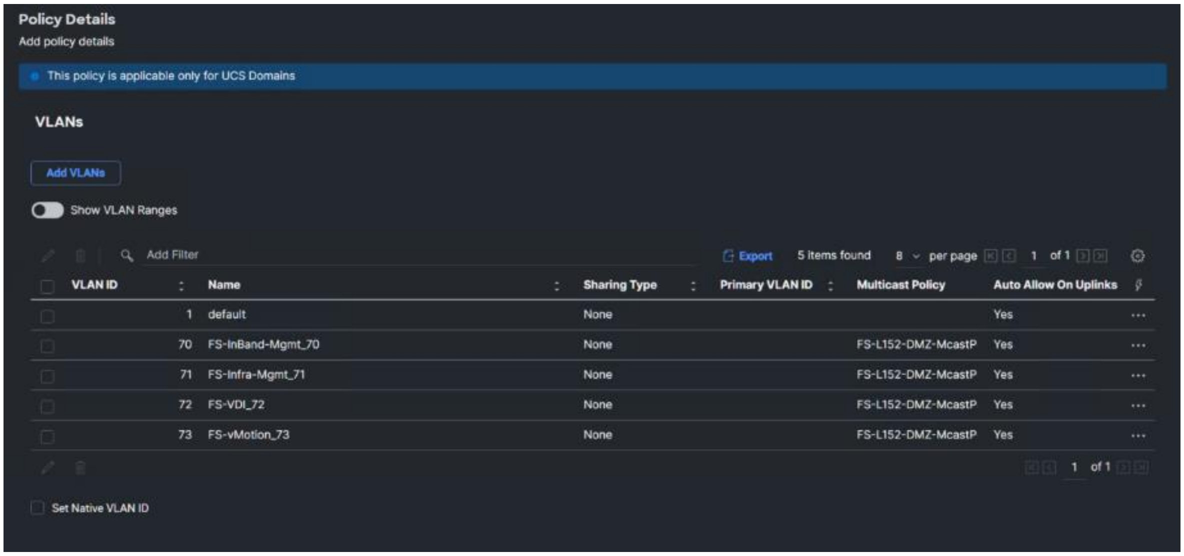

Policy Details

Add policy details

This policy is applicable only for UCS Domains

VLANs

○ Show VLAN Ranges

5 items found 8 per page 1 of 1

VLAN ID	Name	Sharing Type	Primary VLAN ID	Multicast Policy	Auto Allow On Uplinks
1	default	None			Yes
70	FS-InBand-Mgmt_70	None		FS-L152-DMZ-McastP	Yes
71	FS-Infra-Mgmt_71	None		FS-L152-DMZ-McastP	Yes
72	FS-VDI_72	None		FS-L152-DMZ-McastP	Yes
73	FS-vMotion_73	None		FS-L152-DMZ-McastP	Yes

1 of 1

□ Set Native VLAN ID

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

Step 17. Click Select Policy next to VSAN Configuration under Fabric Interconnect A. Click Create New.

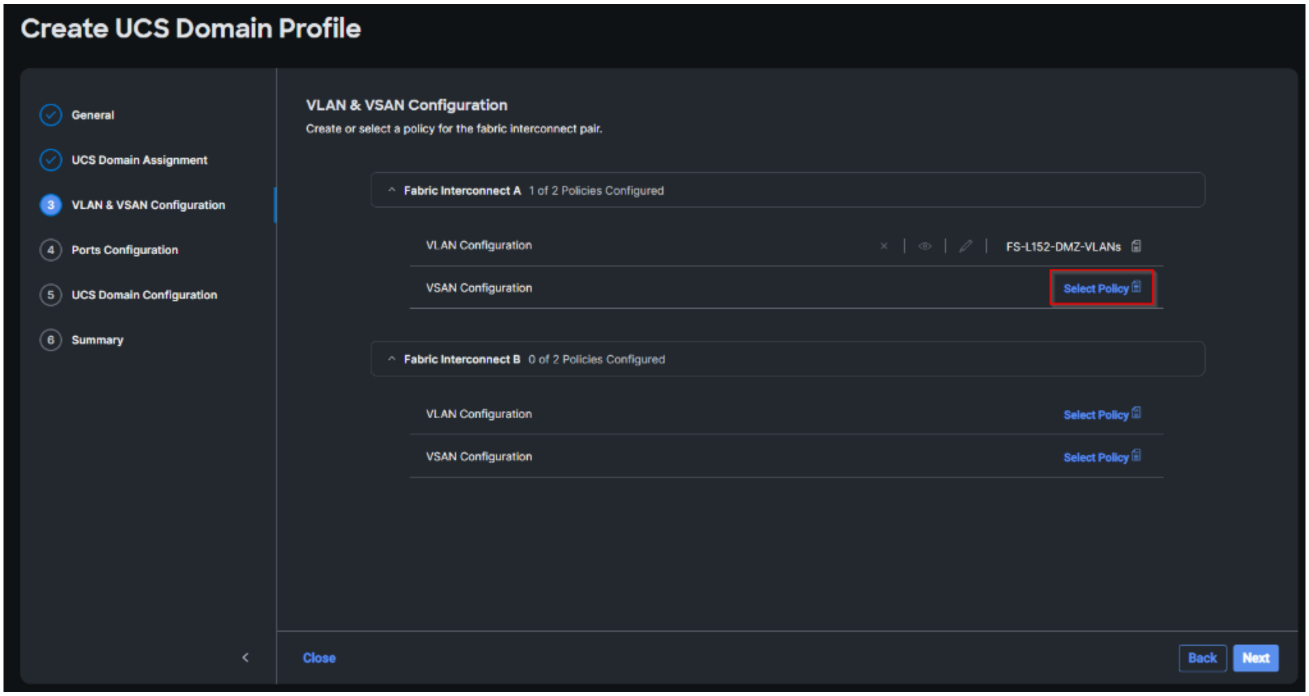

Create UCS Domain Profile

1. General
2. UCS Domain Assignment
3. VLAN & VSAN Configuration
4. Ports Configuration
5. UCS Domain Configuration
6. Summary

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 1 of 2 Policies Configured

VLAN Configuration	FS-L152-DMZ-VLANs
VSAN Configuration	Select Policy

Fabric Interconnect B 0 of 2 Policies Configured

VLAN Configuration	Select Policy
VSAN Configuration	Select Policy

Step 18. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-VSAN-A). Click Next.

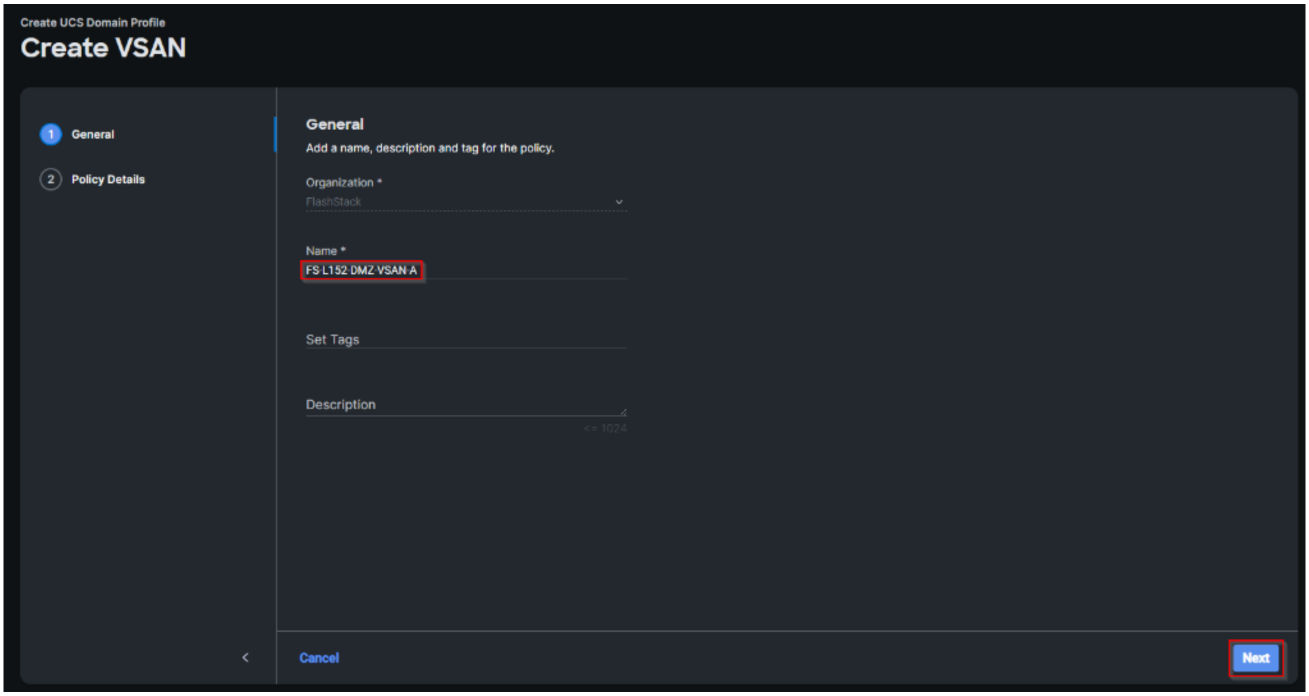

Create UCS Domain Profile

Create VSAN

1. General
2. Policy Details

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-VSAN-A

Set Tags

Description

<= 1024

Step 19. Click Add VSAN.

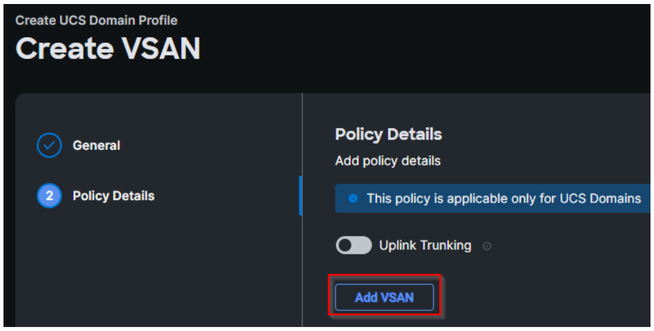
Create UCS Domain Profile

Create VSAN

- General
- 2 Policy Details

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Uplink Trunking

Step 20. Provide a name (for example, VSAN-A), VSAN ID (for example, 100), and associated Fibre Channel over Ethernet (FCoE) VLAN ID (for example, 100) for SAN A.

Step 21. Set VLAN Scope as Uplink.

Step 22. Click Add.

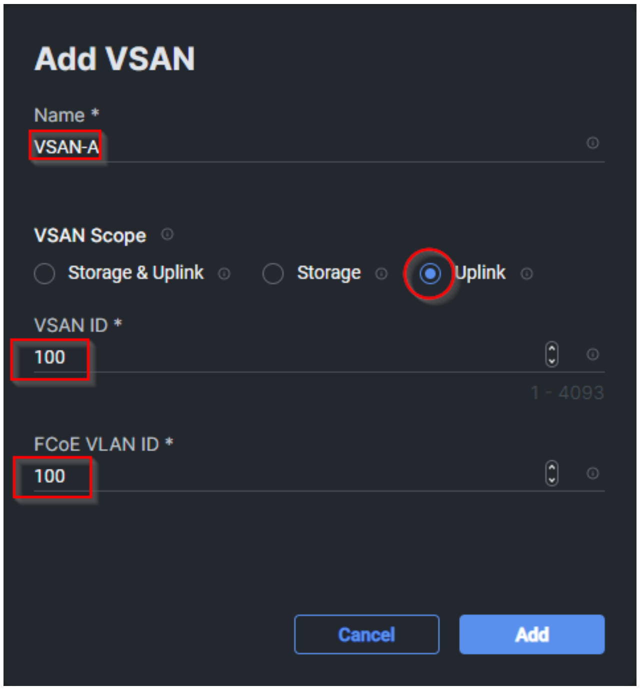
Add VSAN

Name * VSAN-A

VSAN Scope: ○ Storage & Uplink ○ Storage ● Uplink

VSAN ID * 100 (1 - 4093)

FCoE VLAN ID * 100

Step 23. Click Create to finish creating VSAN policy for fabric A.

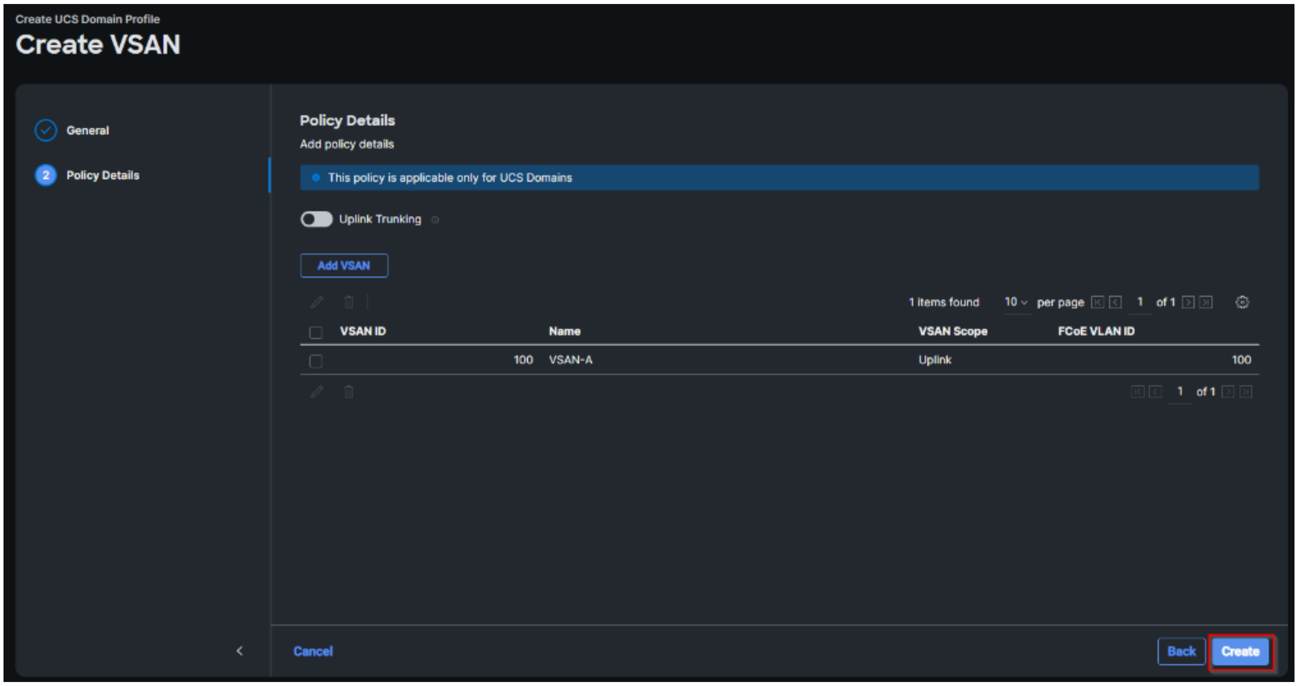

Create UCS Domain Profile

Create VSAN

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Uplink Trunking

VSAN ID	Name	VSAN Scope	FCoE VLAN ID
100	VSAN-A	Uplink	100

Step 24. Repeat steps 7 – 23 for fabric interconnect B assigning the VLAN policy created previously and creating a new VSAN policy for SAN-B. Name the policy to identify the SAN-B configuration (for example, FS-L152-DMZ-VSAN-B) and use appropriate VSAN and FCoE VLAN (for example, 101).

Step 25. Verify that a common VLAN policy and two unique VSAN policies are associated with the two fabric interconnects. Click Next.

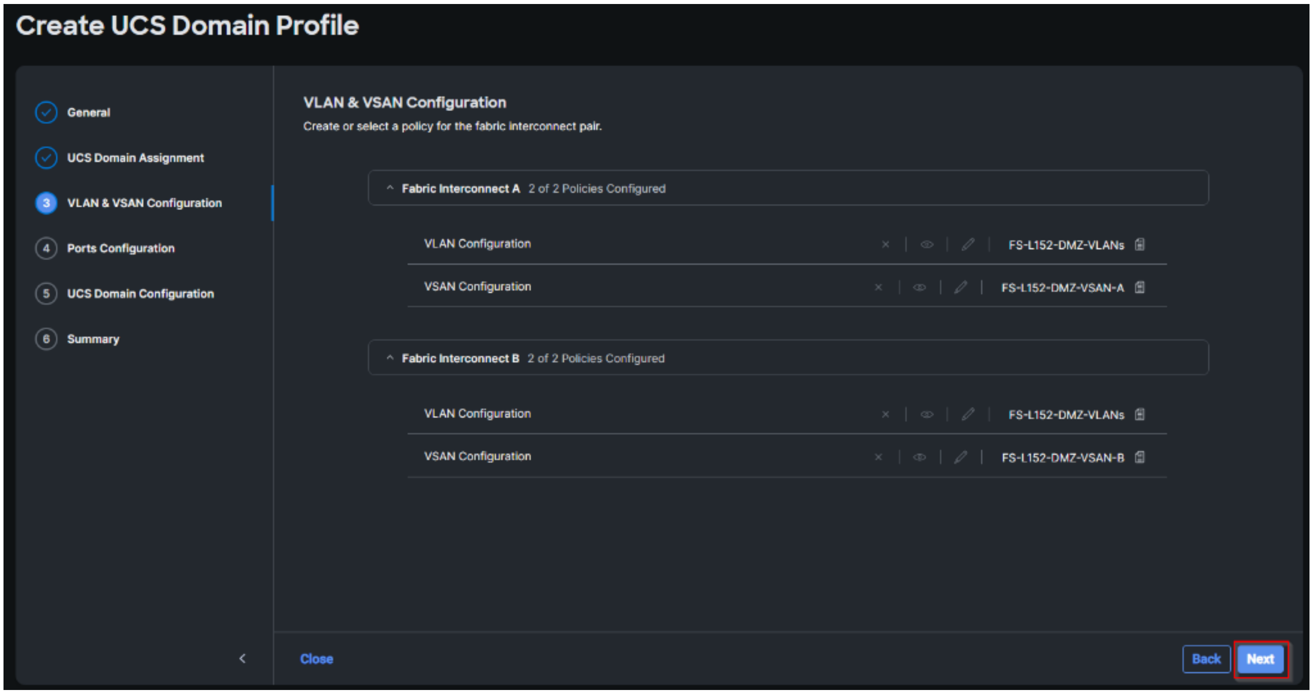

Create UCS Domain Profile

VLAN & VSAN Configuration

Create or select a policy for the fabric interconnect pair.

Fabric Interconnect A 2 of 2 Policies Configured

VLAN Configuration	FS-L152-DMZ-VLANs
VSAN Configuration	FS-L152-DMZ-VSAN-A

Fabric Interconnect B 2 of 2 Policies Configured

VLAN Configuration	FS-L152-DMZ-VLANs
VSAN Configuration	FS-L152-DMZ-VSAN-B

Step 26. 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 provide flexibility when port configuration (port numbers or speed) differs between the two FIs. When configuring

Fibre Channel, two port policies are required because each fabric interconnect uses unique Fibre Channel VSAN ID.

Step 27. Click Select Policy for Fabric Interconnect A.

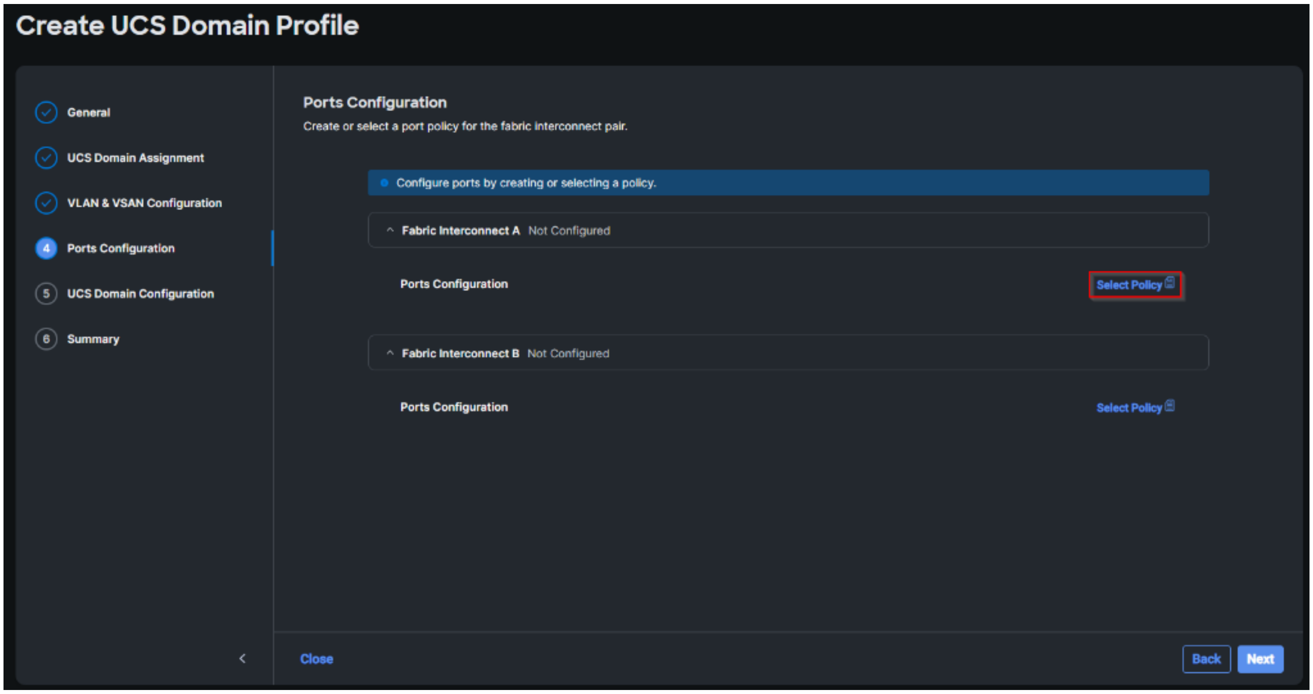

Create UCS Domain Profile

Ports Configuration

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

Configure ports by creating or selecting a policy.

Fabric Interconnect A Not Configured

Ports Configuration — Select Policy

Fabric Interconnect B Not Configured

Ports Configuration — Select Policy

Step 28. Click Create New.

Step 29. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-K4-FI-A). Click Next.

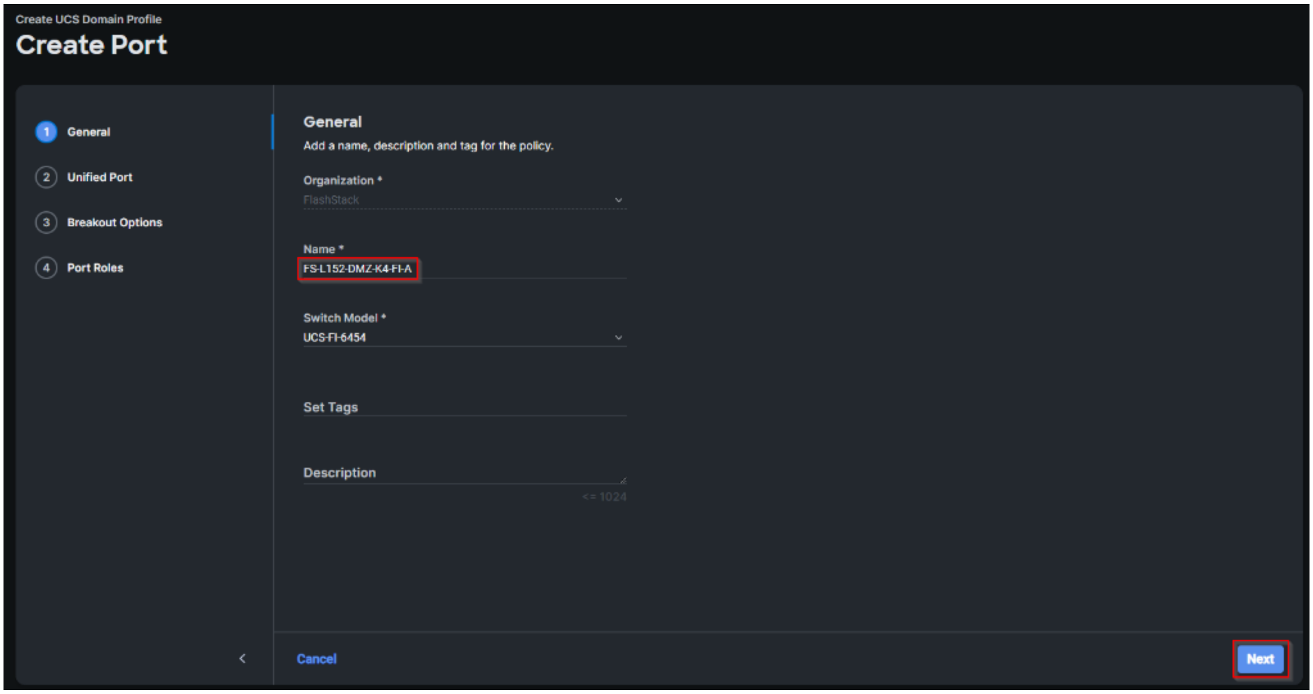

Create UCS Domain Profile

Create Port

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-K4-FI-A

Switch Model * UCS-FI-6454

Set Tags

Description

<= 1024

Step 30. Move the slider to set up unified ports. In this deployment, the first four ports were selected as Fibre Channel ports. Click Next.

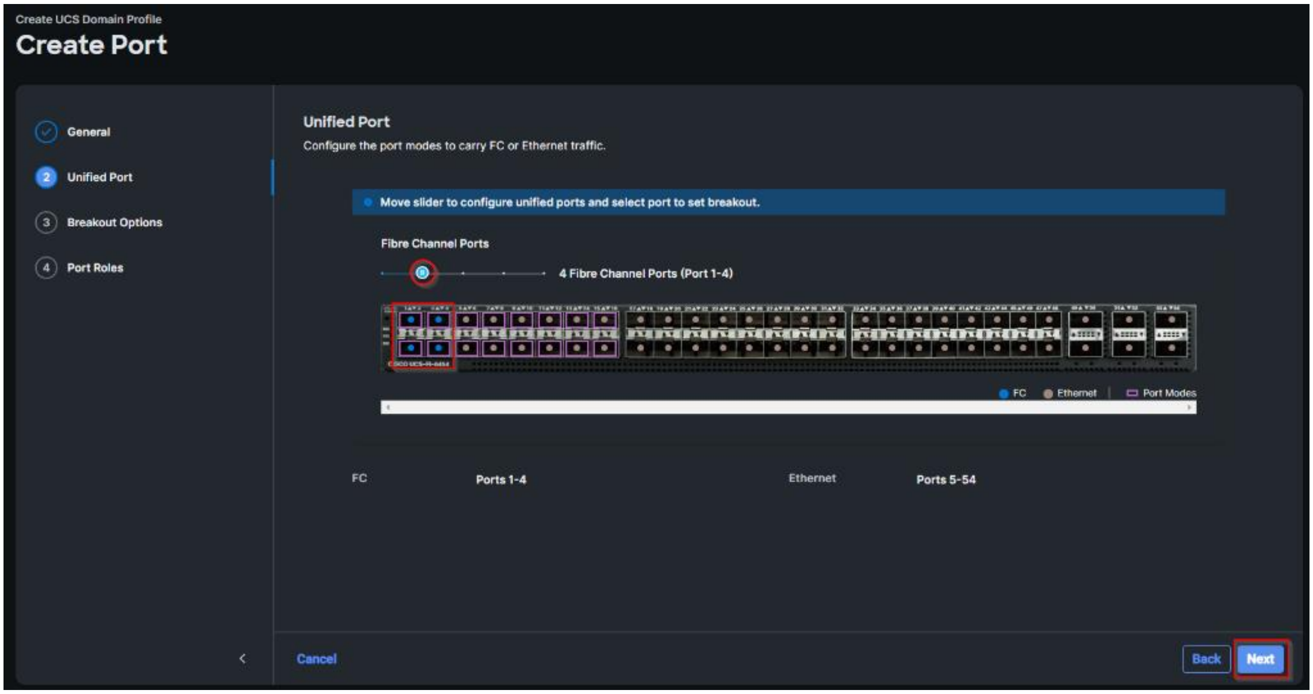

Create UCS Domain Profile

Create Port

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

4 Fibre Channel Ports (Port 1-4)

● FC ● Ethernet | ▭ Port Modes

FC	Ports 1-4	Ethernet	Ports 5-54

Step 31. On the breakout Options page click Next.

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

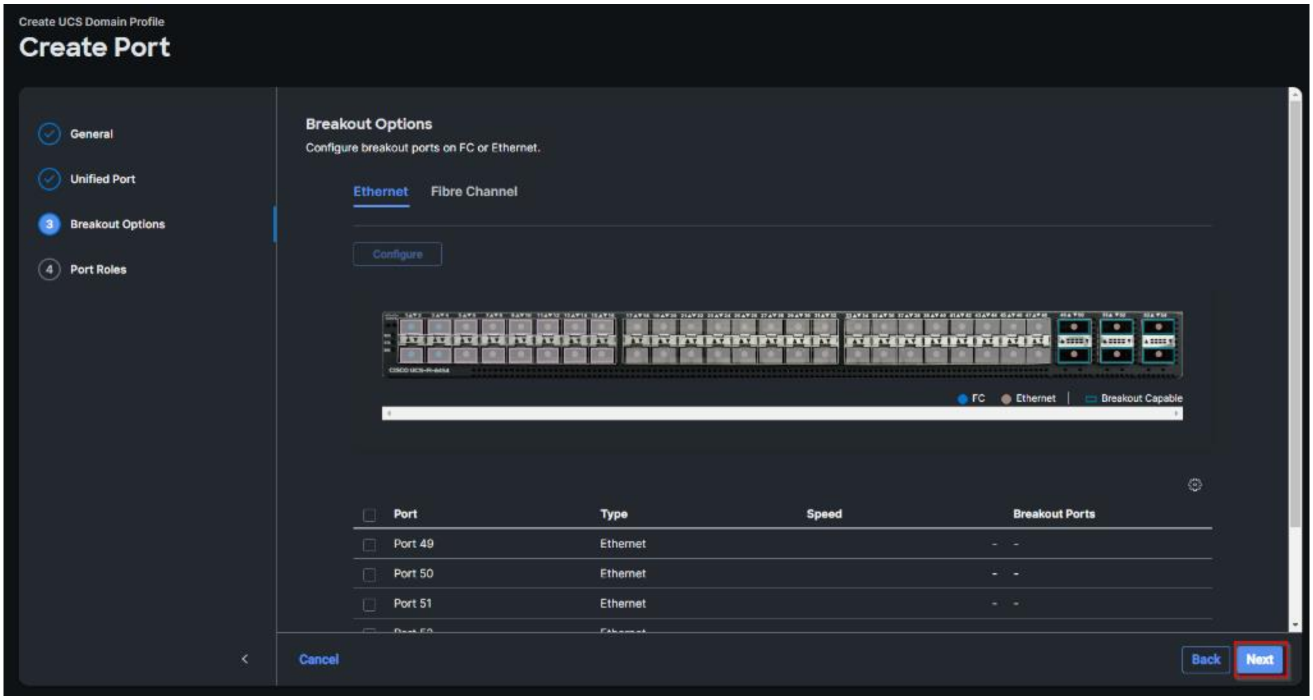

Create UCS Domain Profile

Create Port

Breakout Options

Configure breakout ports on FC or Ethernet.

● FC ● Ethernet | ▭ Breakout Capable

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

Step 32. 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.

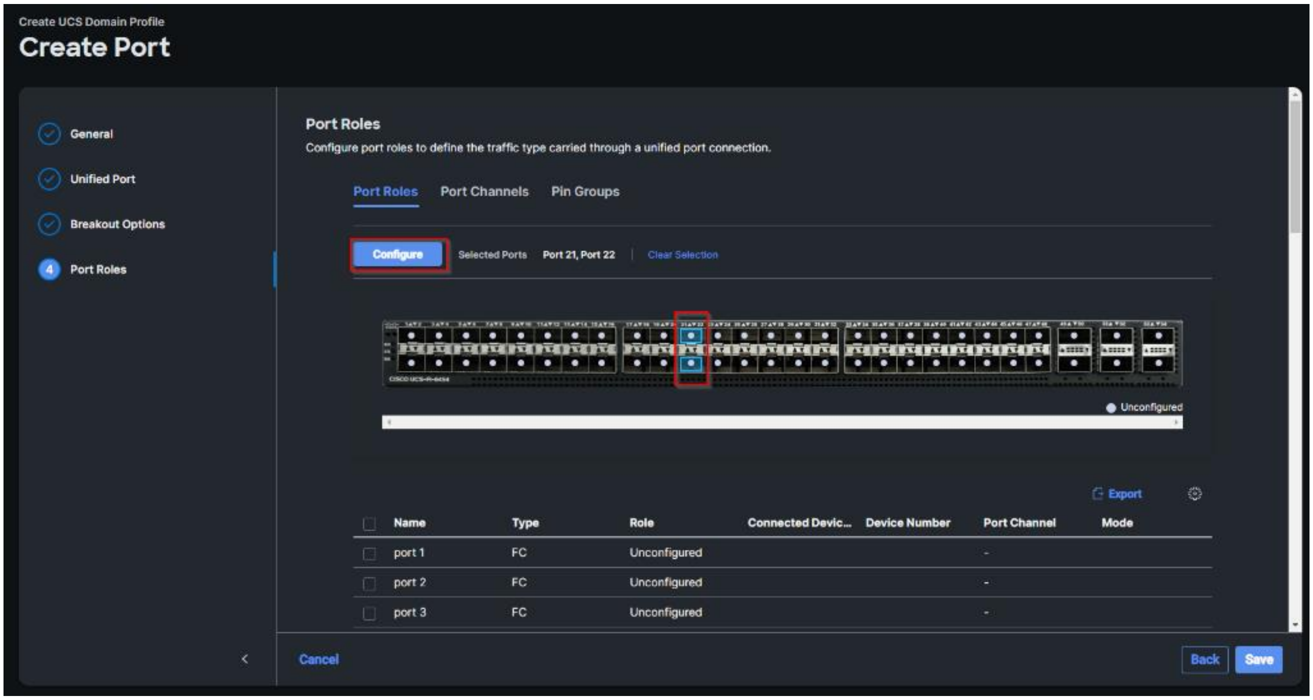

Port Roles

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

Selected Ports Port 21, Port 22

● Unconfigured

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

Step 33. From the drop-down list, select Server as the role. Click Save.

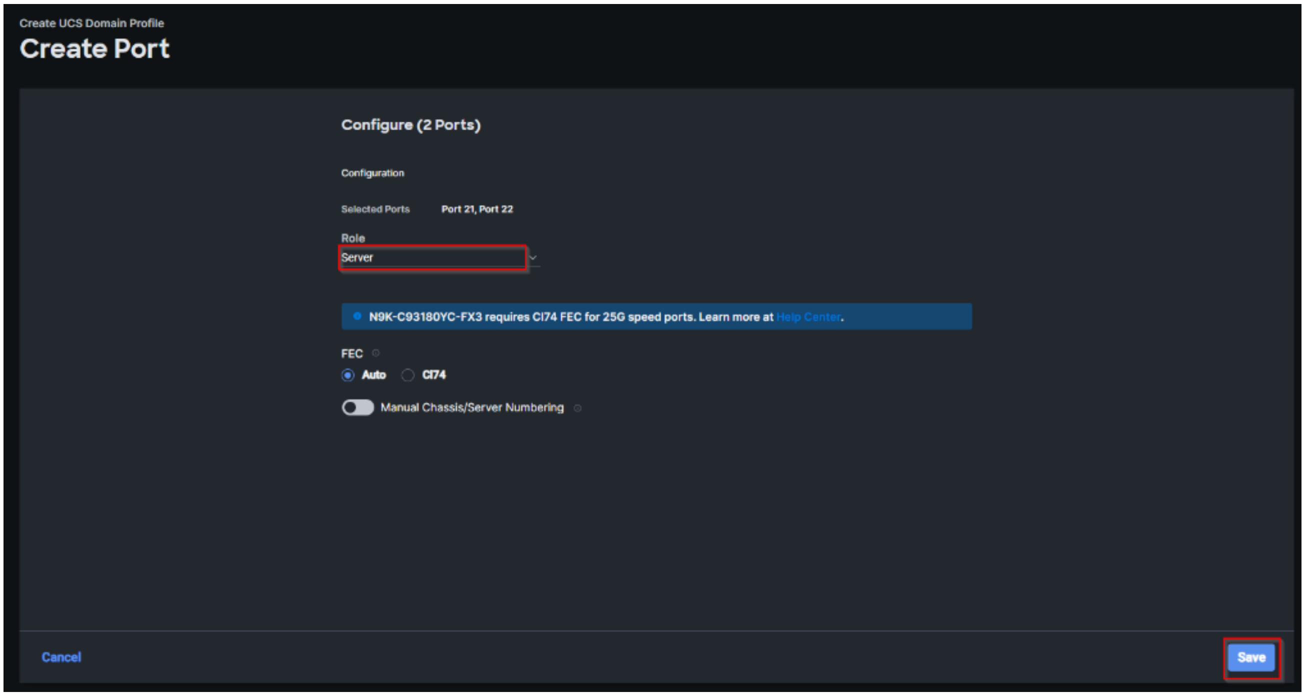

Configure (2 Ports)

Configuration

Selected Ports Port 21, Port 22

Role: Server

N9K-C93180YC-FX3 requires Cl74 FEC for 25G speed ports. Learn more at Help Center.

FEC

● Auto ○ Cl74

Manual Chassis/Server Numbering

Step 34. Configure the Ethernet uplink port channel by selecting the Port Channel in the main pane and then clicking Create Port Channel.

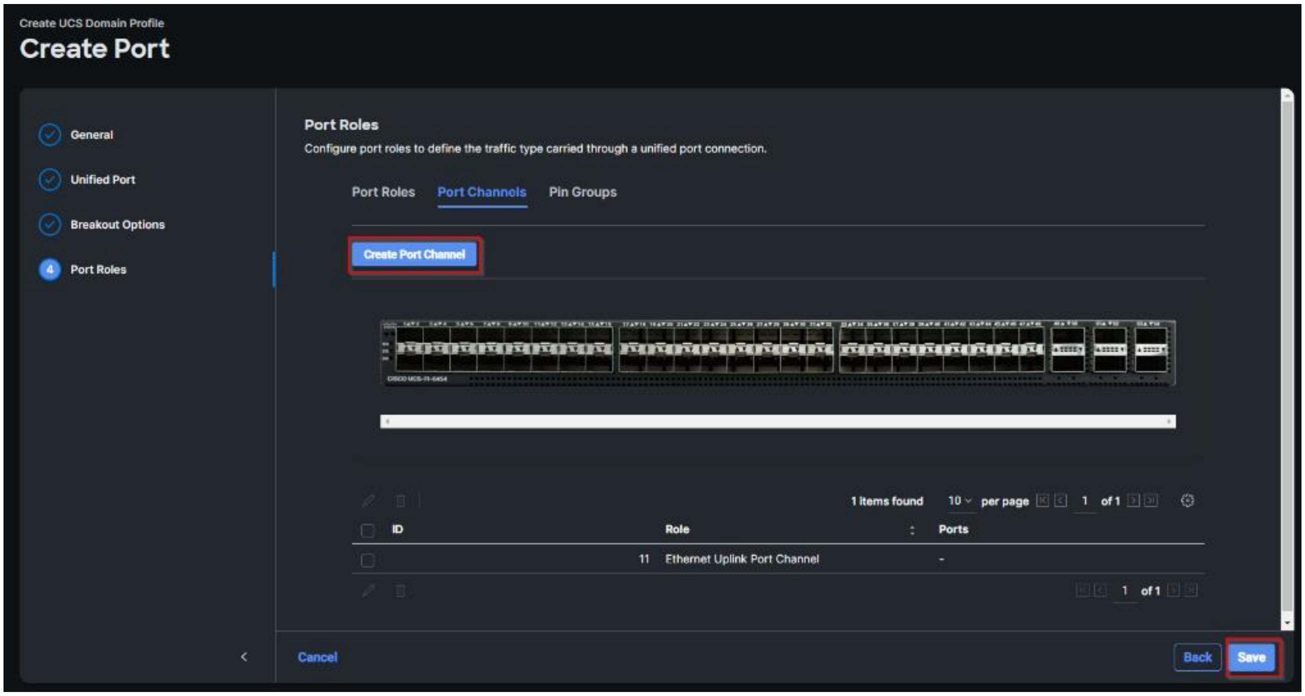

Create UCS Domain Profile

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	-

Step 35. 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 36. Scroll down and select uplink ports from the list of available ports (for example, port 49 and 50).

Step 37. Click Save.

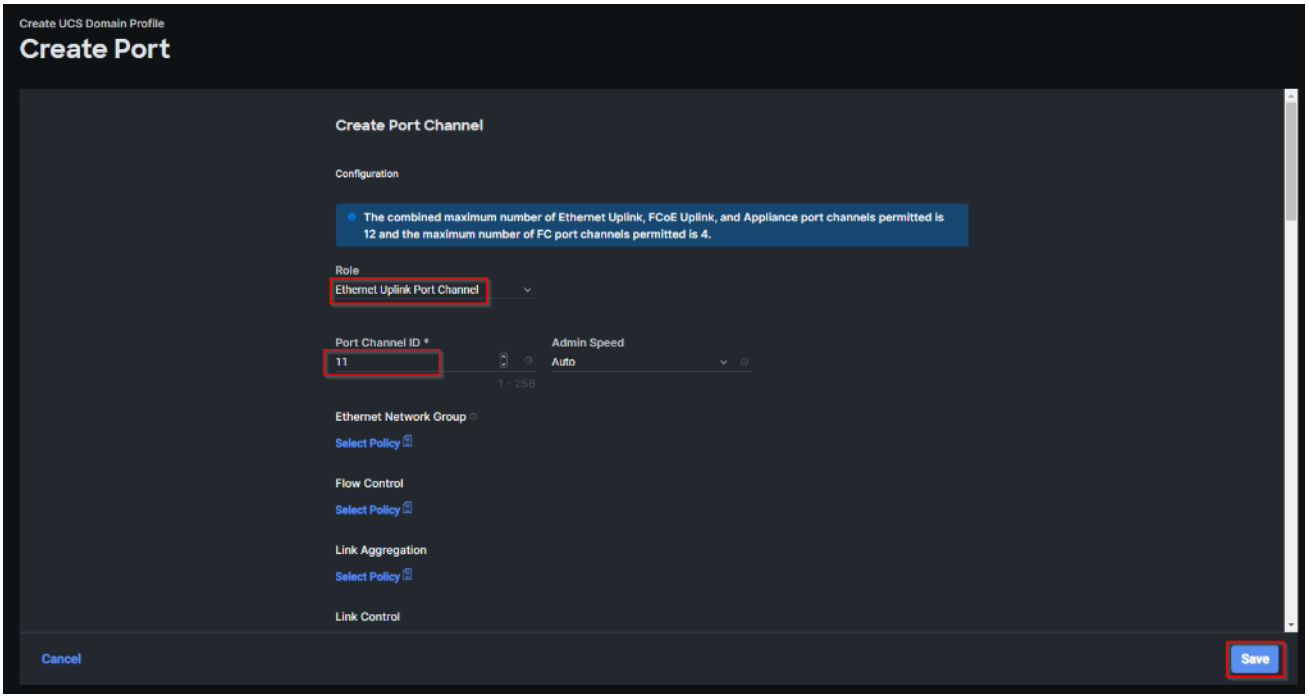

Create UCS Domain Profile

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

Flow Control

Link Aggregation

Link Control

Step 38. Repeat steps 1 - 37 to create the port policy for Fabric Interconnect B. Use the following values for various parameters:

- Name of the port policy: FS-L152-DMZ-K4-FI-B
- Ethernet port-Channel ID: 12

Step 39. When the port configuration for both fabric interconnects is complete and looks good, click Next.

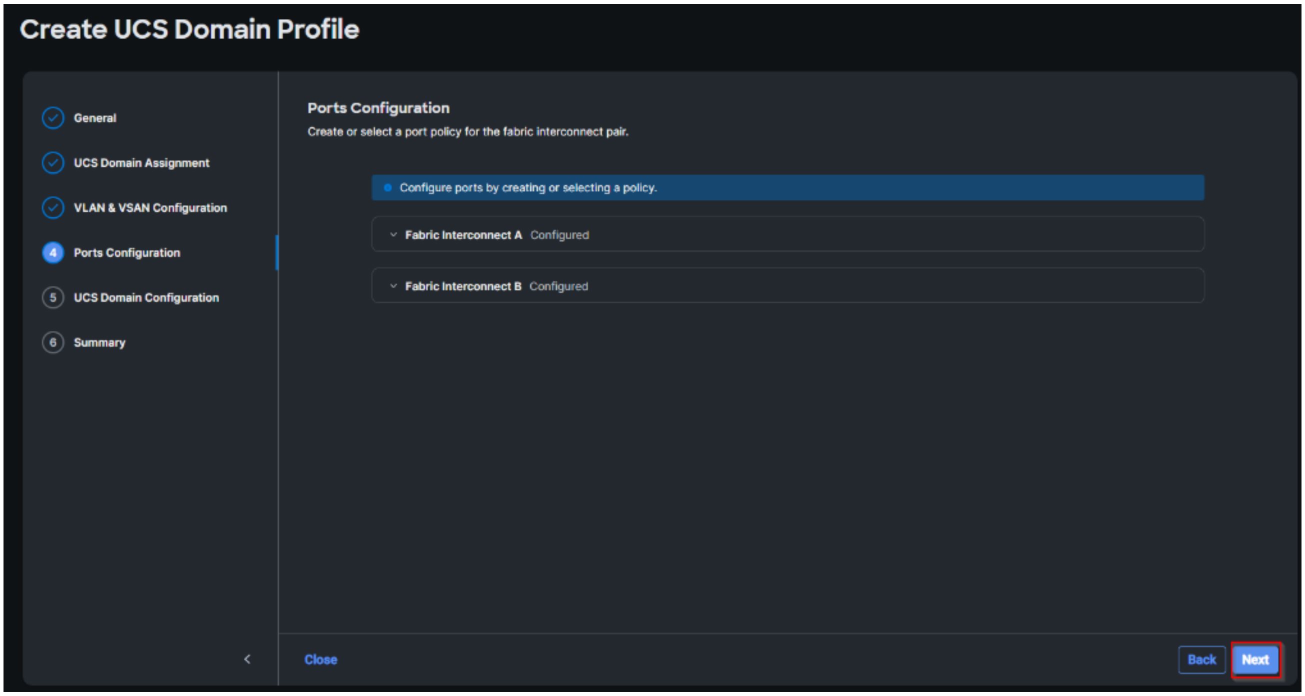

Create UCS Domain Profile

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 40. Under UCS domain configuration, additional policies can be configured to setup NTP, Syslog, DNS settings, SNMP, QoS and UCS operating mode (end host or switch mode). For this deployment, System QoS will be configured.

Step 41. Click Select Policy next to System QoS* and click Create New to define the System QOS policy.

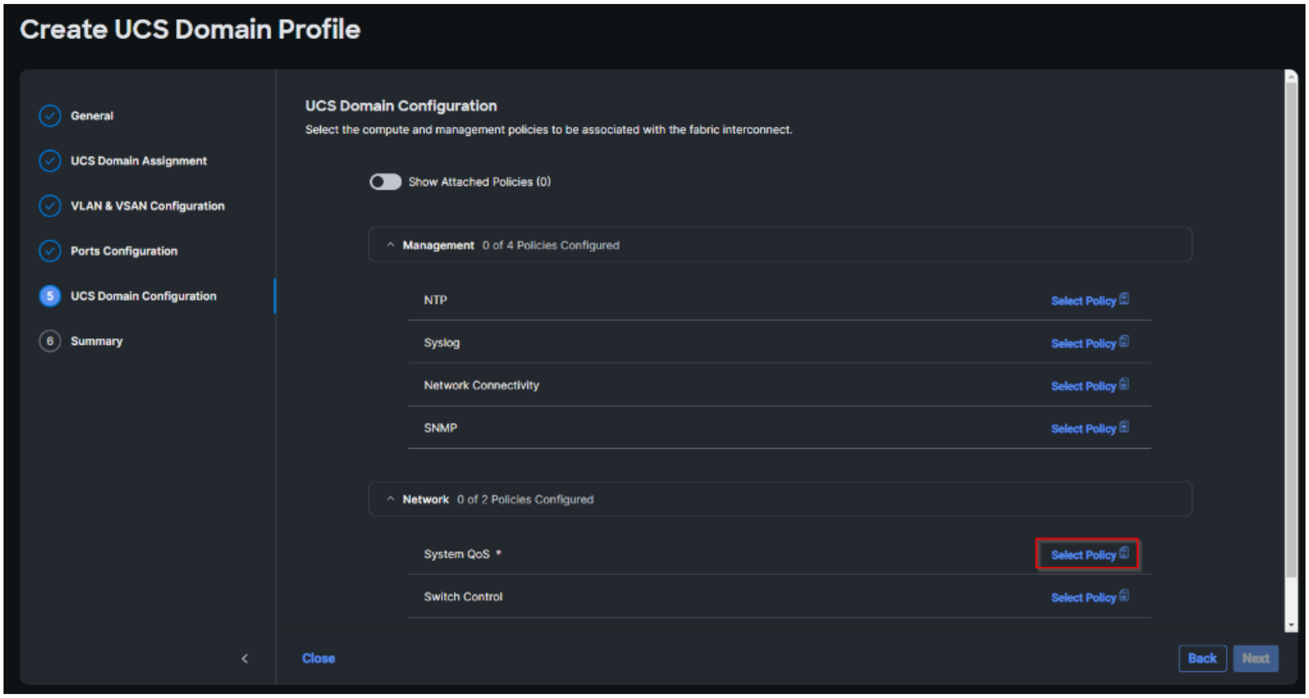

Create UCS Domain Profile

UCS Domain Configuration

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

Show Attached Policies (0)

Management 0 of 4 Policies Configured

NTP	Select Policy
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 42. Verify correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-QosPol). Click Next.

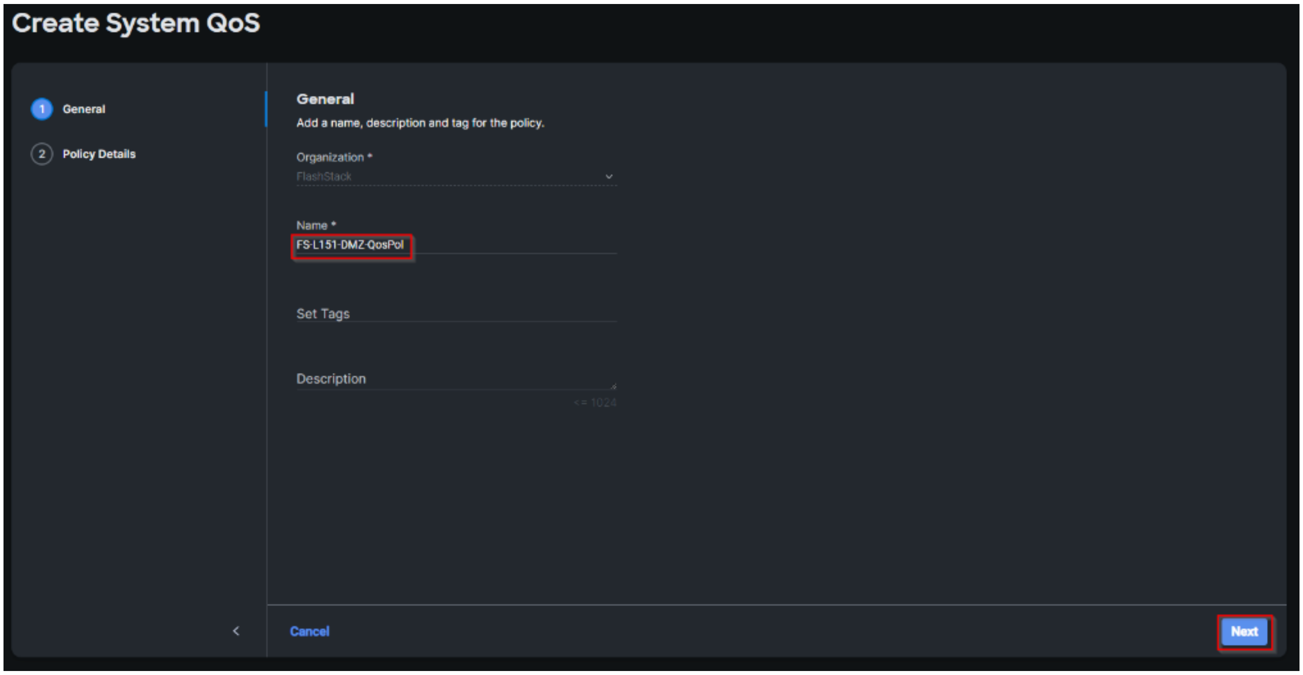

Create System QoS

General

Add a name, description and tag for the policy.

Organization *

FlashStack

Name *

FS-L151-DMZ-QosPol

Set Tags

Description

<= 1024

Step 43. Change the MTU for Best Effort class to 9216. Keep the rest default selections. Click Create.

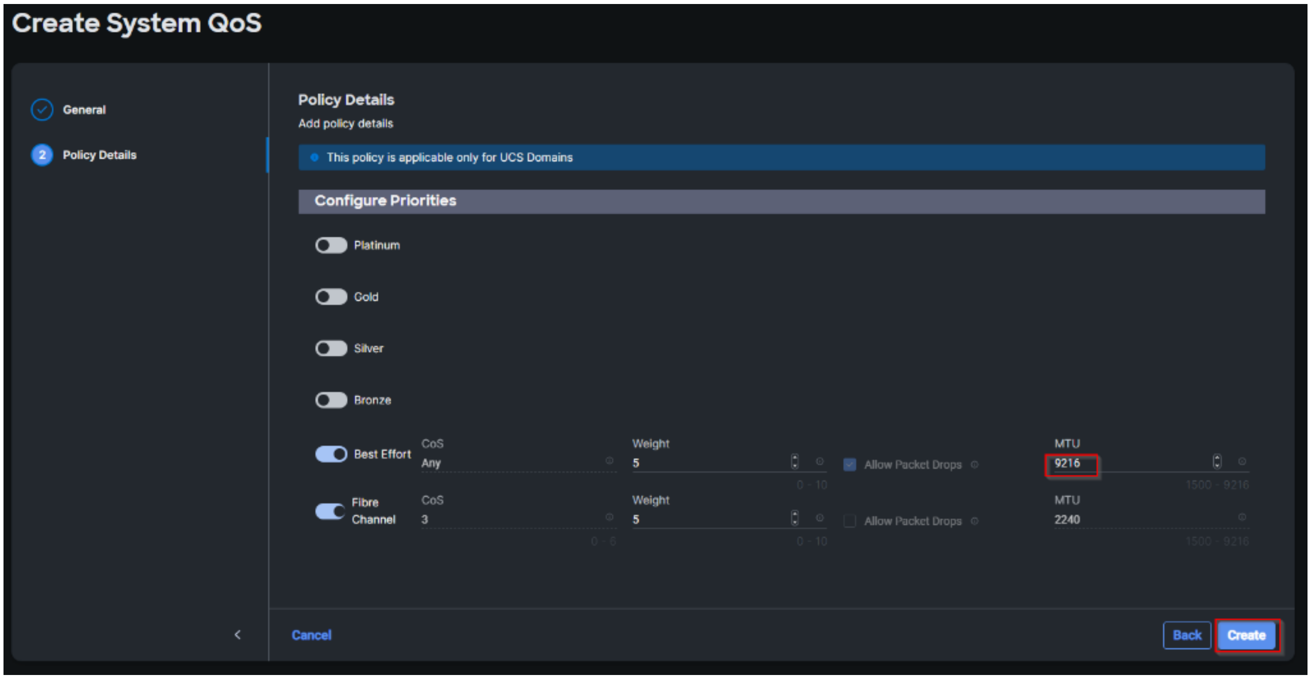

Create System QoS

General

2 Policy Details

Policy Details

Add policy details

This policy is applicable only for UCS Domains

Configure Priorities

Platinum

Gold

Silver

Bronze

	CoS	Weight		MTU
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

Step 44. Click Next.

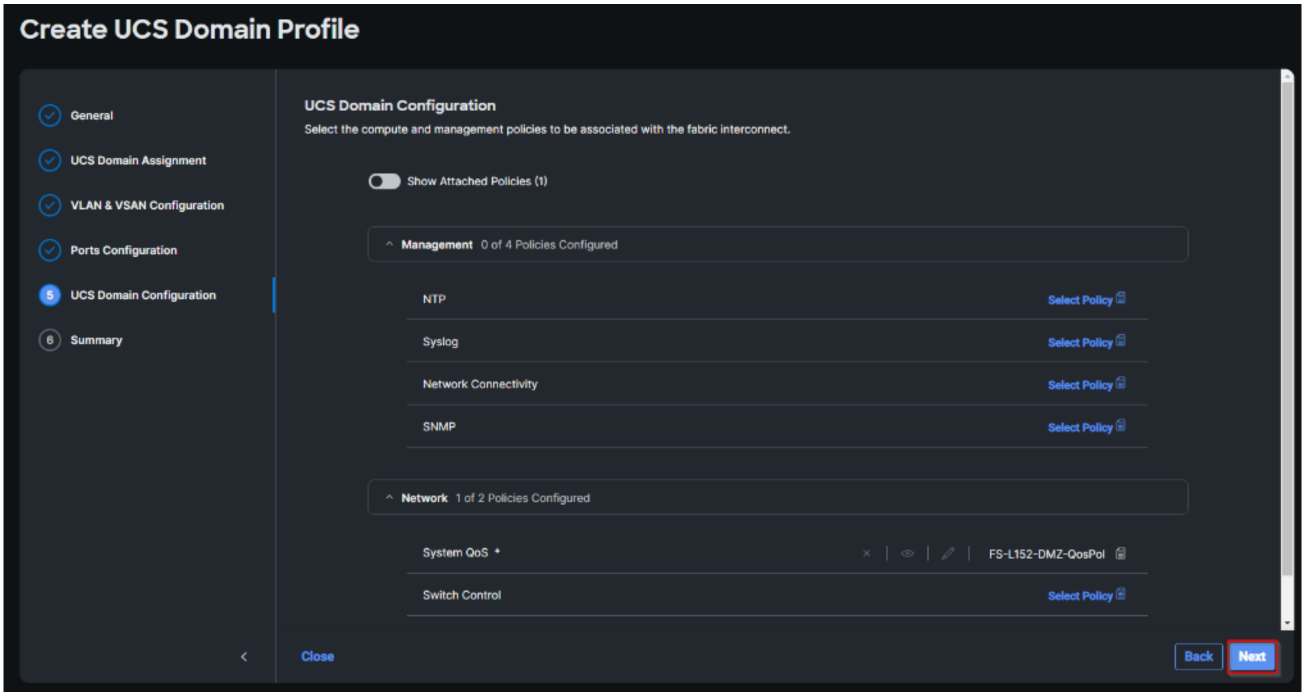

Create UCS Domain Profile

General

UCS Domain Assignment

VLAN & VSAN Configuration

Ports Configuration

5 UCS Domain Configuration

6 Summary

UCS Domain Configuration

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

Show Attached Policies (1)

Management 0 of 4 Policies Configured

NTP	Select Policy
Syslog	Select Policy
Network Connectivity	Select Policy
SNMP	Select Policy

Network 1 of 2 Policies Configured

System QoS *	FS-L152-DMZ-QosPol
Switch Control	Select Policy

Step 45. 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.

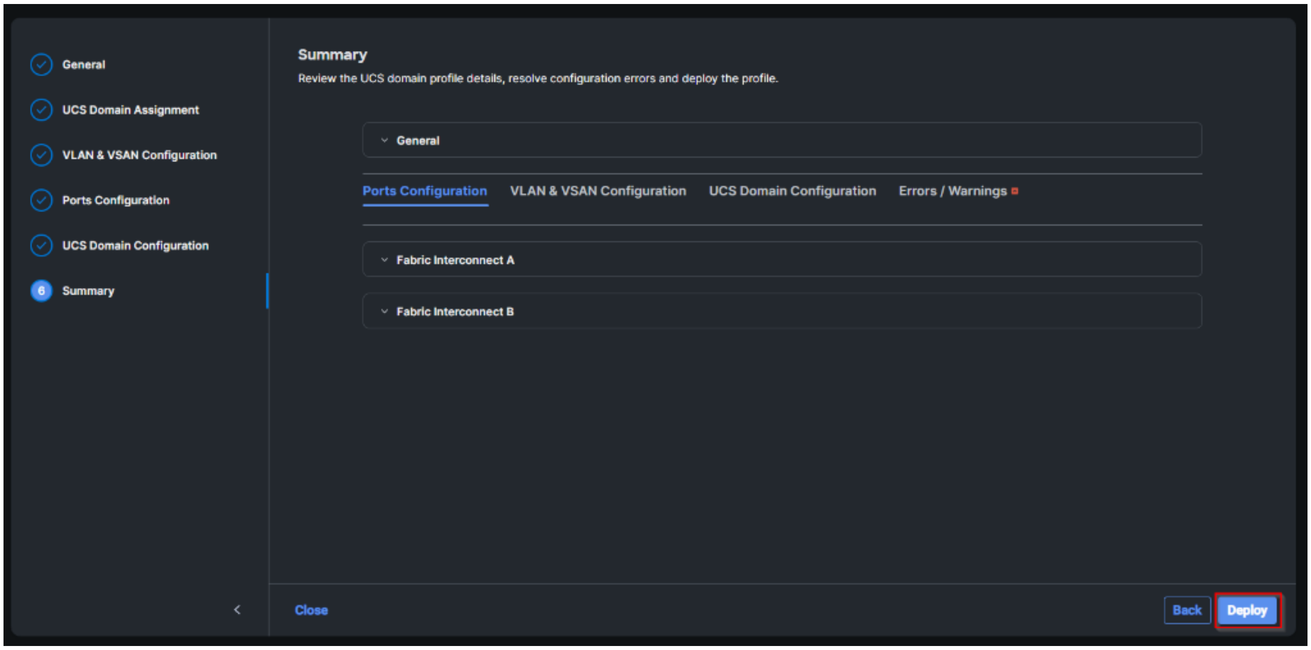

Summary

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

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:

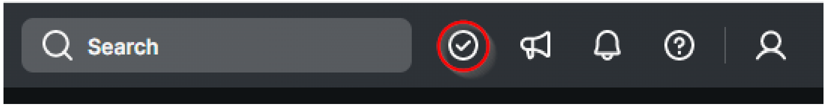

Step 46. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Configure > Profiles, select UCS Domain Profiles, verify that the domain profile has been successfully deployed.

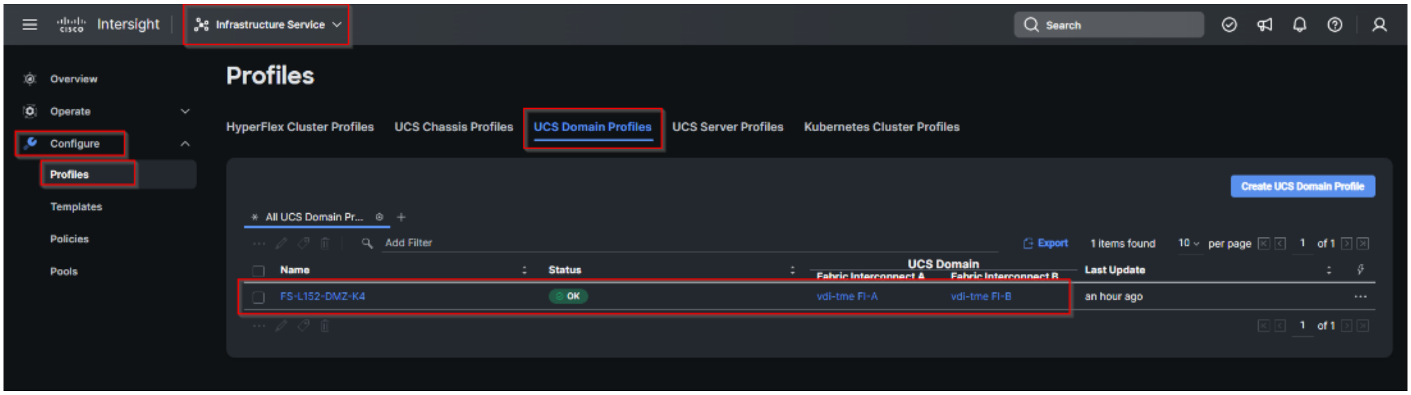

Profiles

Name	Status	UCS Domain Fabric Interconnect A	UCS Domain Fabric Interconnect B	Last Update
FS-L152-DMZ-K4	OK	vdi-tme FI-A	vdi-tme FI-B	an hour ago

Step 47. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Chassis, verify that the chassis has been discovered.

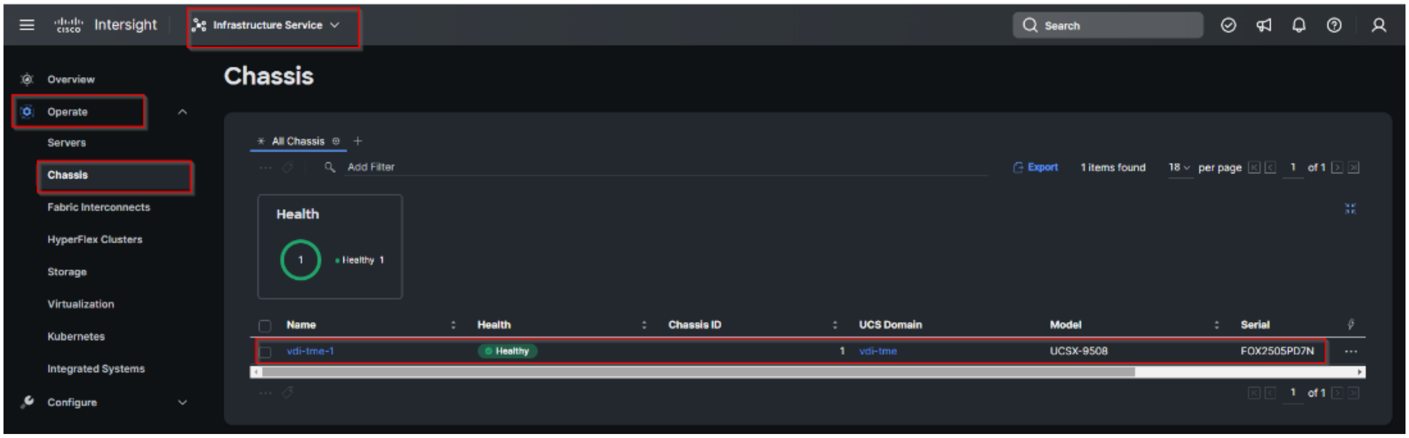

Chassis

Name	Health	Chassis ID	UCS Domain	Model	Serial
vdi-tme-1	Healthy	1	vdi-tme	UCSX-9508	FOX2505PD7N

Step 48. From the Service Selector drop-down list, select Infrastructure Service. Navigate to Operate > Servers, verify that the servers have been successfully discovered.

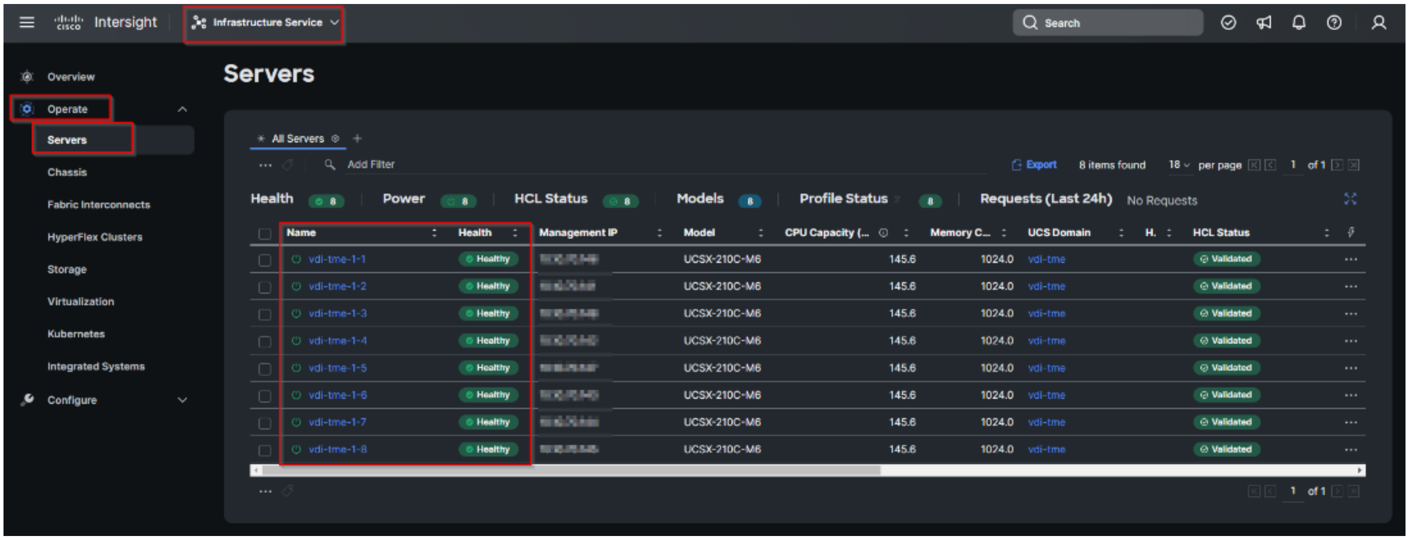

Servers

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status
vdi-tme-1-1	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-2	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-3	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-4	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-5	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-6	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-7	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-8	Healthy	[redacted]	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated

Configure Cisco UCS Chassis Profile

The 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:

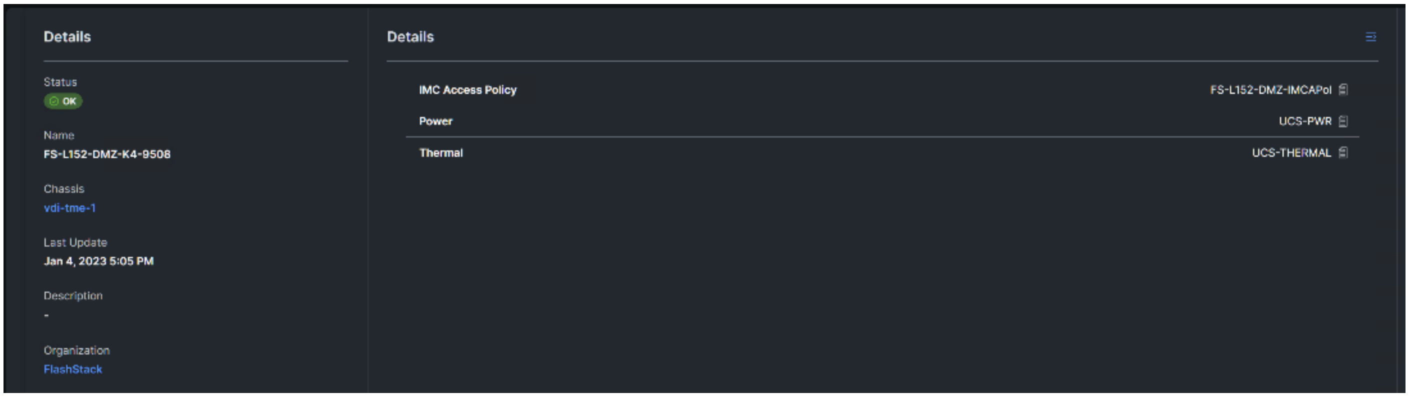

Details

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 23. Chassis policy detail

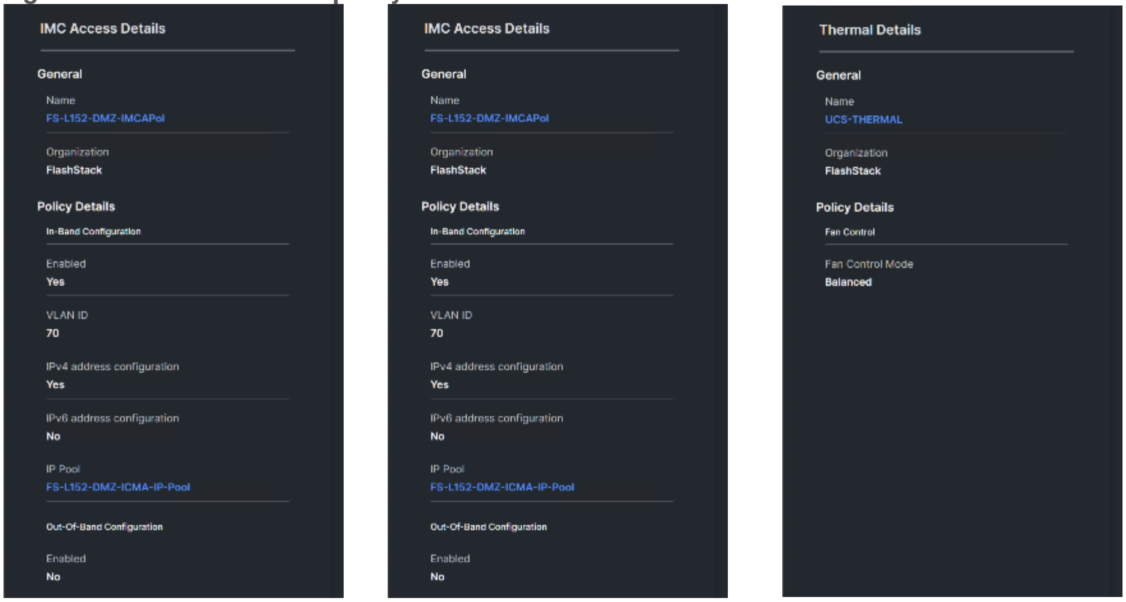

IMC Access Details

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 M6 compute nodes.

Procedure 1. Create vNIC and vHBA Placement for the Server Profile Template

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

Table 4. vHBA and vNIC placement for FC connected storage

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-A	MLOM	A	0
vHBA-B	MLOM	B	1
01-vSwitch0-A	MLOM	A	2
02-vSwitch0-B	MLOM	B	3
03-VDS0-A	MLOM	A	4
04-VDS0-B	MLOM	B	5

Note: Two vHBAs (vHBA-A and vHBA-B) are configured to support FC boot from SAN.

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.

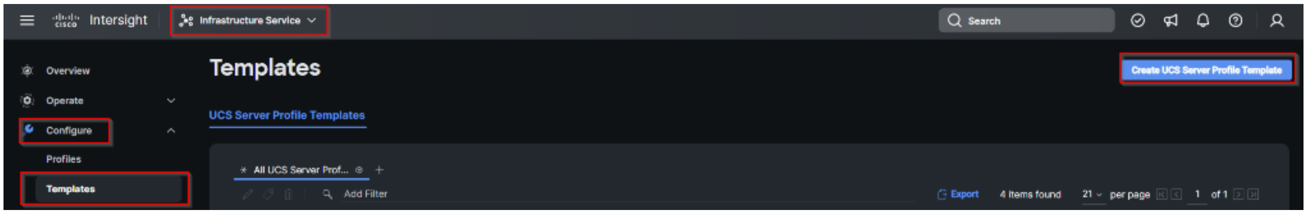

Step 3. Select the organization from the drop-down list. Provide a name for the server profile template (for example, FS-L151-DMZ-K4-X210CM6) for FI-Attached UCS Server. Click Next.

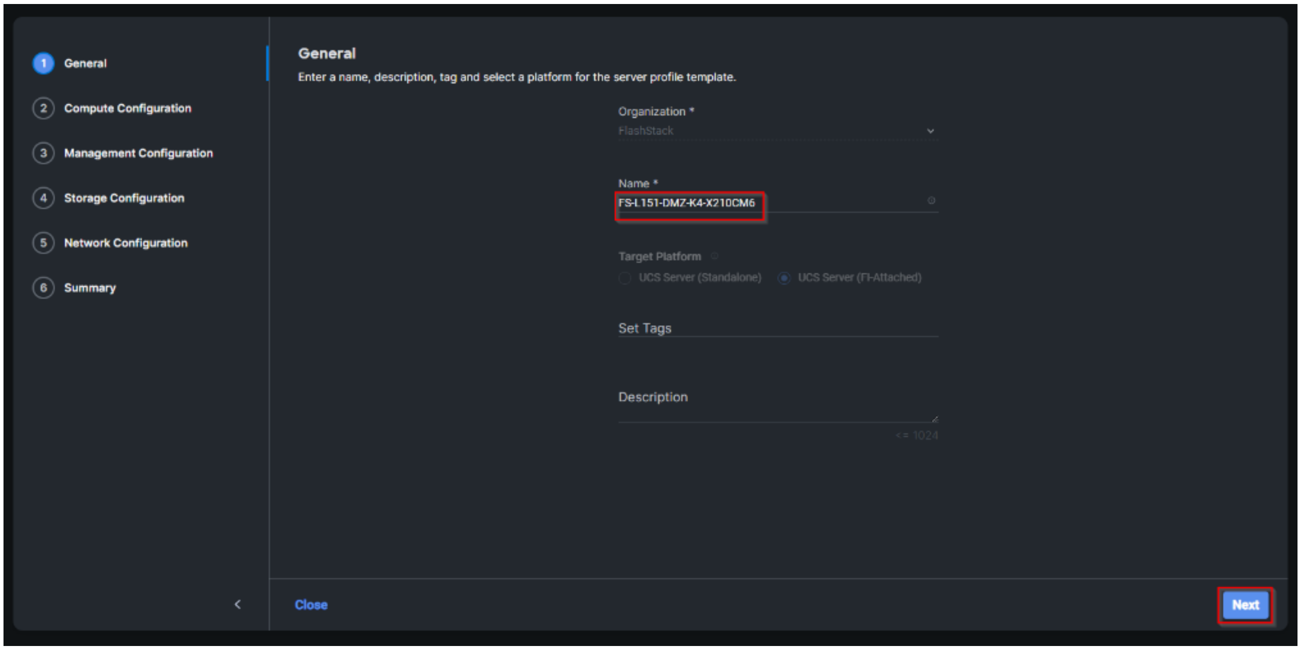

General

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

Organization *: FlashStack

Name *: FS-L151-DMZ-K4-X210CM6

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

Set Tags

Description

<= 1024

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

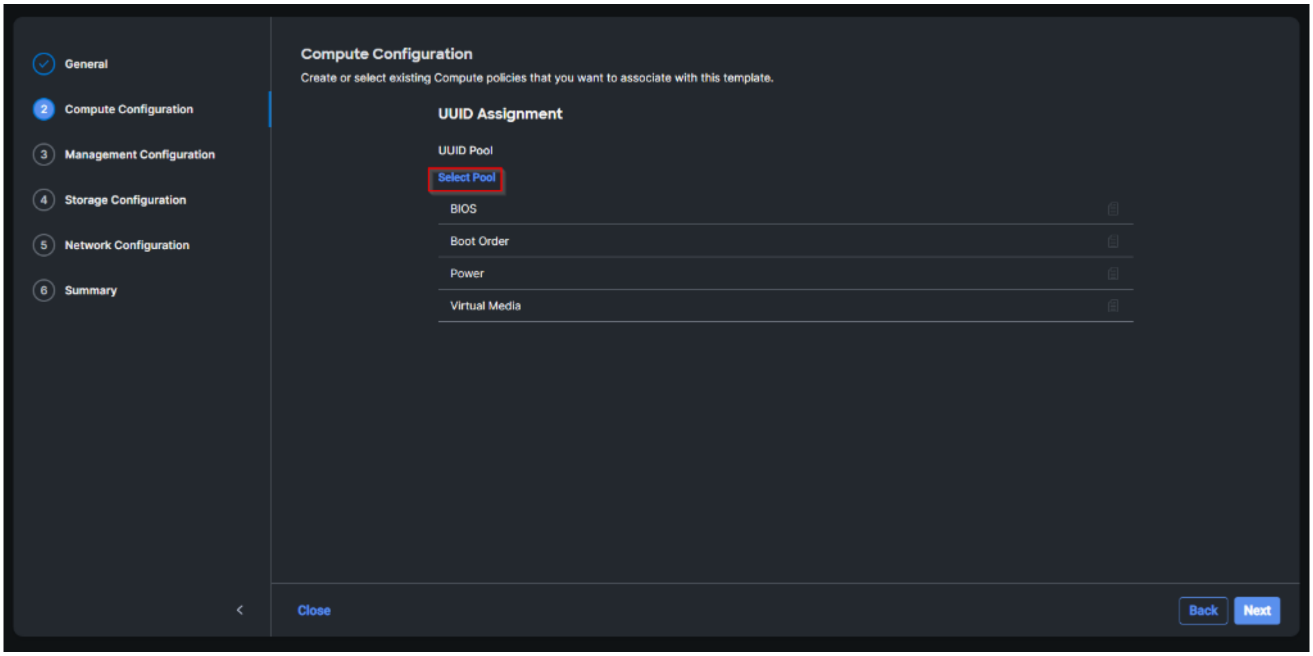

Compute Configuration

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

UUID Assignment

UUID Pool

Select Pool

BIOS

Boot Order

Power

Virtual Media

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

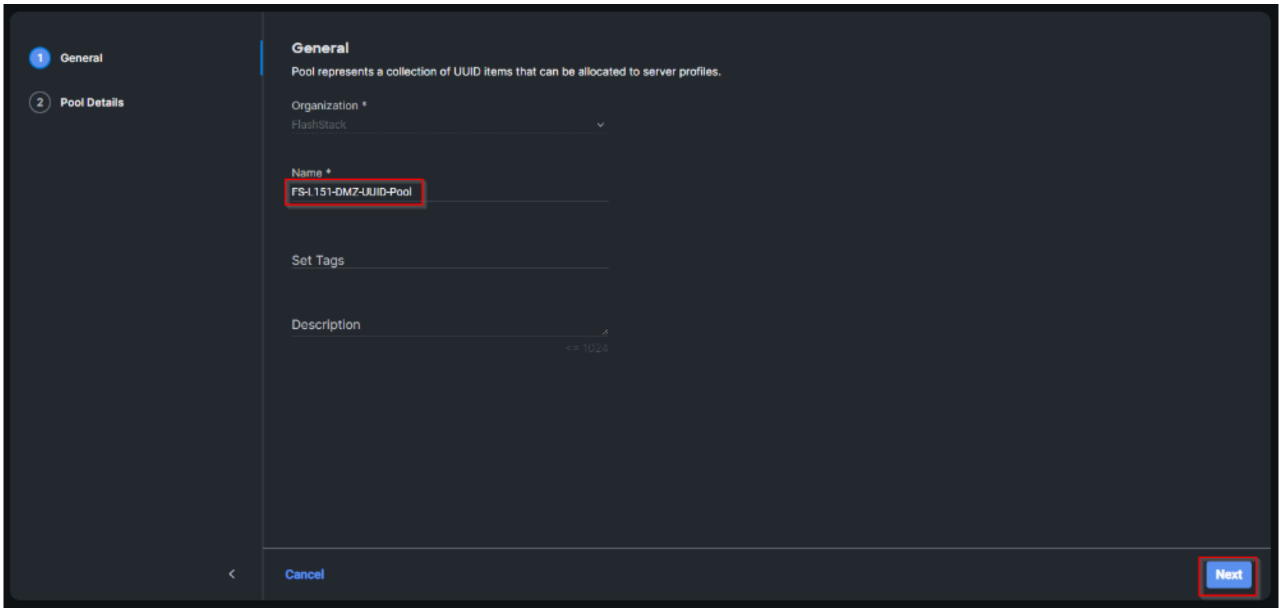

General

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

Organization *
FlashStack

Name *
FS-L151-DMZ-UUID-Pool

Set Tags

Description
<= 1024

Step 6. Provide a UUID Prefix (for example, a random prefix of A11A14B6-B193-49C7 was used). Add a UUID block of appropriate size. Click Create.

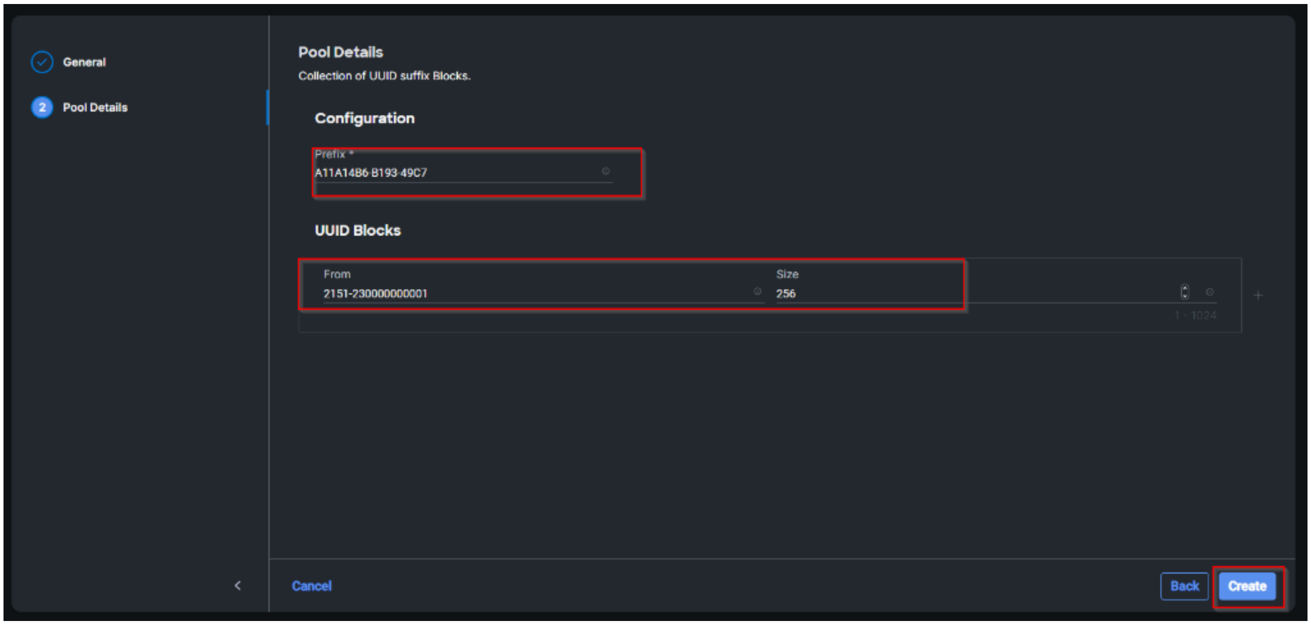

Pool Details

Collection of UUID suffix Blocks.

Configuration

Prefix *
A11A14B6-B193-49C7

UUID Blocks

From
2151-230000000001

Size
256

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, FS-L151-DMZ-M6-BIOS-Perf).

Step 9. Click Next.

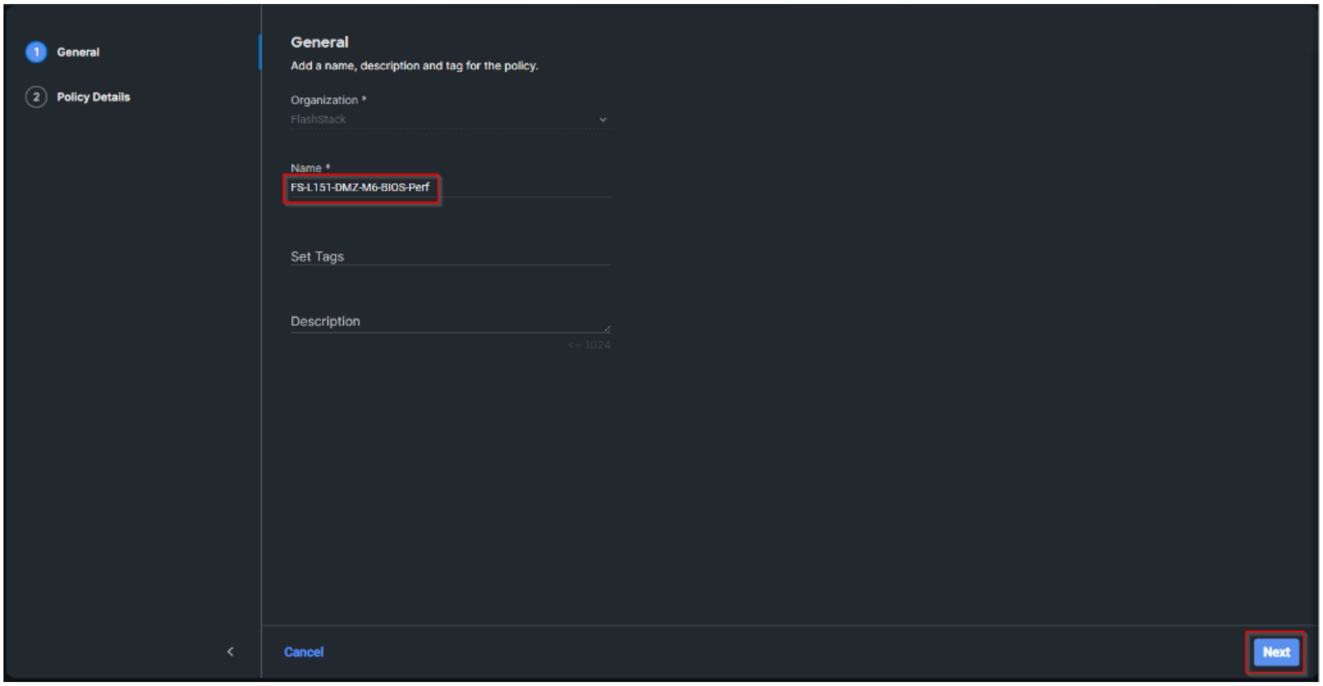

General

Add a name, description and tag for the policy.

Organization *
FlashStack

Name *
FS-L151-DMZ-M6-BIOS-Perf

Set Tags

Description

<= 1024

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

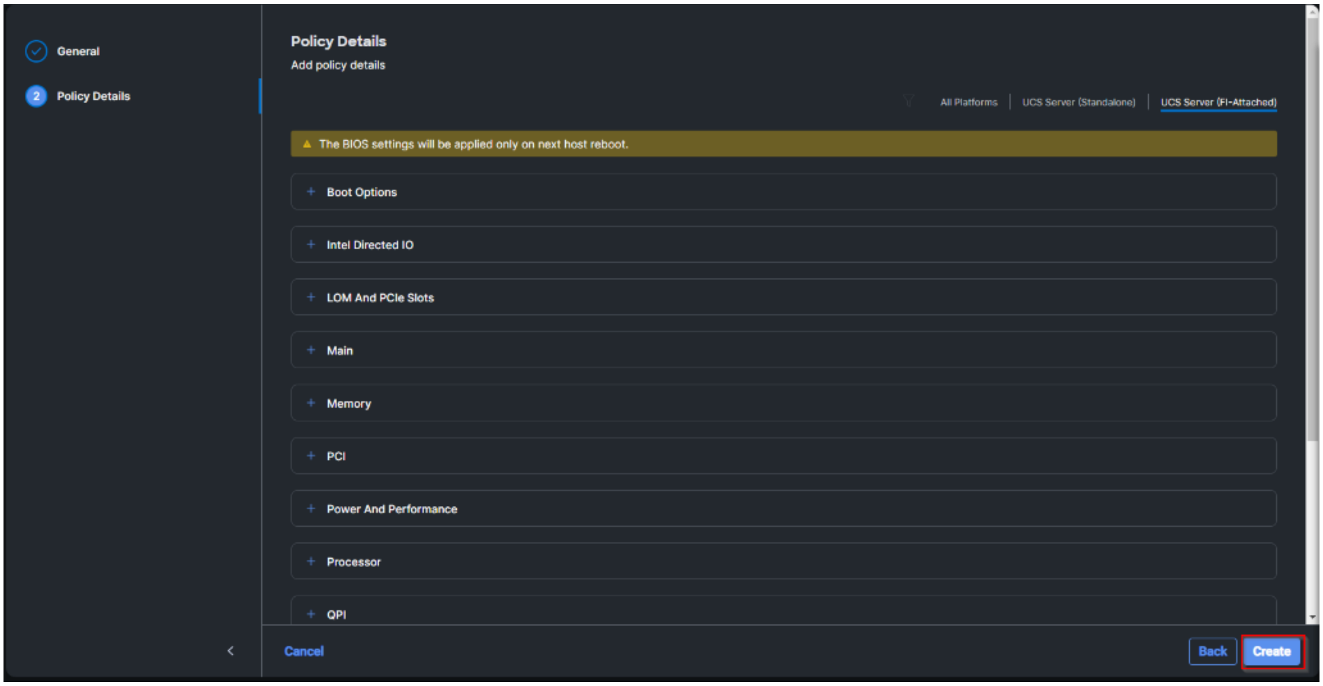

The BIOS settings will be applied only on next host reboot.

Note: In this deployment, the BIOS values were selected based on recommendations in the performance tuning guide for Cisco UCS M6 BIOS: https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/performance-tuning-guide-ucs-m6-servers.html.

Table 5. FS-L151-DMZ-M6-BIOS-Perf 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	
Altitude	auto
Boot Performance Mode	Max Performance

BIOS Token	Value
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

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, FS-L151-DMZ-BootPol). Click Next.

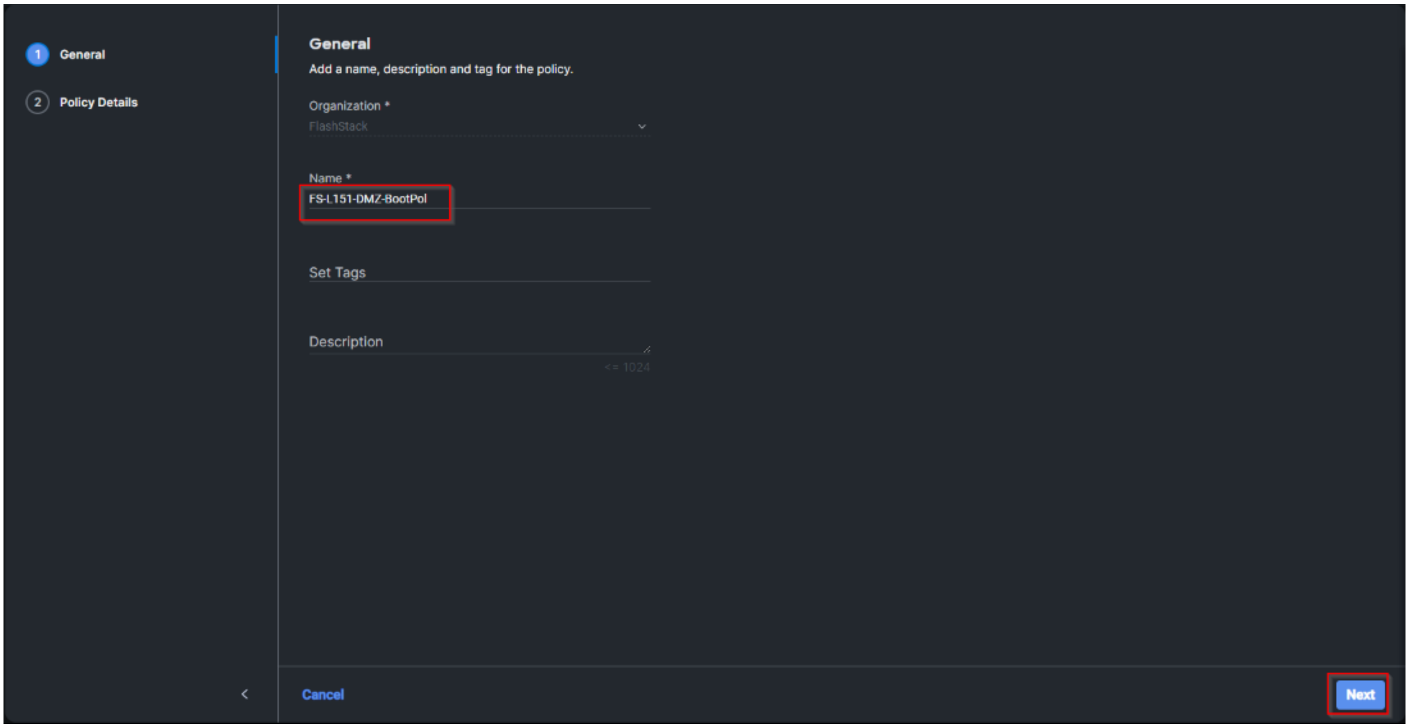

1 General

2 Policy Details

General

Add a name, description and tag for the policy.

Organization *
FlashStack

Name *
FS-L151-DMZ-BootPol

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.

For Fibre Channel SAN boot, four connected FC ports on Pure Storage FlashArray//X70 R3 controllers will be added as boot options. The four FC ports are as follows:

- CT0.FC0, CT1.FC0 are connected to SAN-A
- CT1.FC2, CT0.FC2 are connected to SAN-B.

Figure 24. **Pure Storage FlashArray//X70 R3**

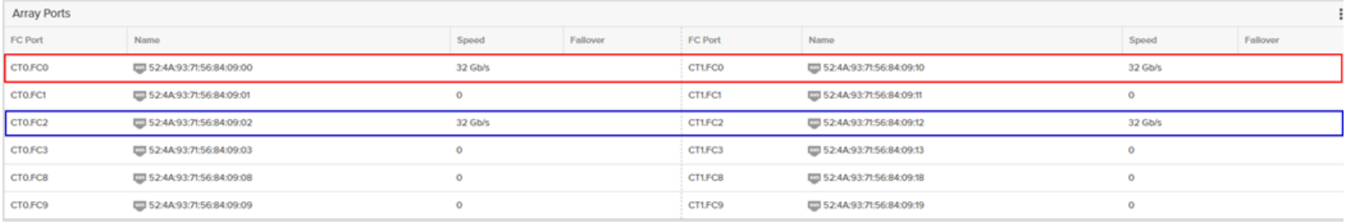

Array Ports

FC Port	Name	Speed	Failover	FC Port	Name	Speed	Failover
CT0.FC0	52:4A:93:71:56:84:09:00	32 Gb/s		CT1.FC0	52:4A:93:71:56:84:09:10	32 Gb/s	
CT0.FC1	52:4A:93:71:56:84:09:01	0		CT1.FC1	52:4A:93:71:56:84:09:11	0	
CT0.FC2	52:4A:93:71:56:84:09:02	32 Gb/s		CT1.FC2	52:4A:93:71:56:84:09:12	32 Gb/s	
CT0.FC3	52:4A:93:71:56:84:09:03	0		CT1.FC3	52:4A:93:71:56:84:09:13	0	
CT0.FC8	52:4A:93:71:56:84:09:08	0		CT1.FC8	52:4A:93:71:56:84:09:18	0	
CT0.FC9	52:4A:93:71:56:84:09:09	0		CT1.FC9	52:4A:93:71:56:84:09:19	0	

Step 17. From the Add Boot Device drop-down list, select SAN Boot (Repeat steps for all 4 FC ports)

Step 18. Provide the Device Name: CT0FC0 and the Logical Unit Number (LUN) value (for example, 1).

Step 19. Provide an interface name vHBA-A. This value is important and should match the vHBA name.

Note: vHBA-A is used to access CT0.FC0, CT1.FC0 and vHBA-B is used to access CT1.FC2, CT0.FC2.

Step 20. Add the appropriate World Wide Port Name (WWPN) as the Target WWPN (for example, 52:4A:93:71:56:84:09:00).

Step 21. Provide bootloader name as BOOTX64.EFI.

Step 22. Provide bootloader name as \EFI\BOOT.

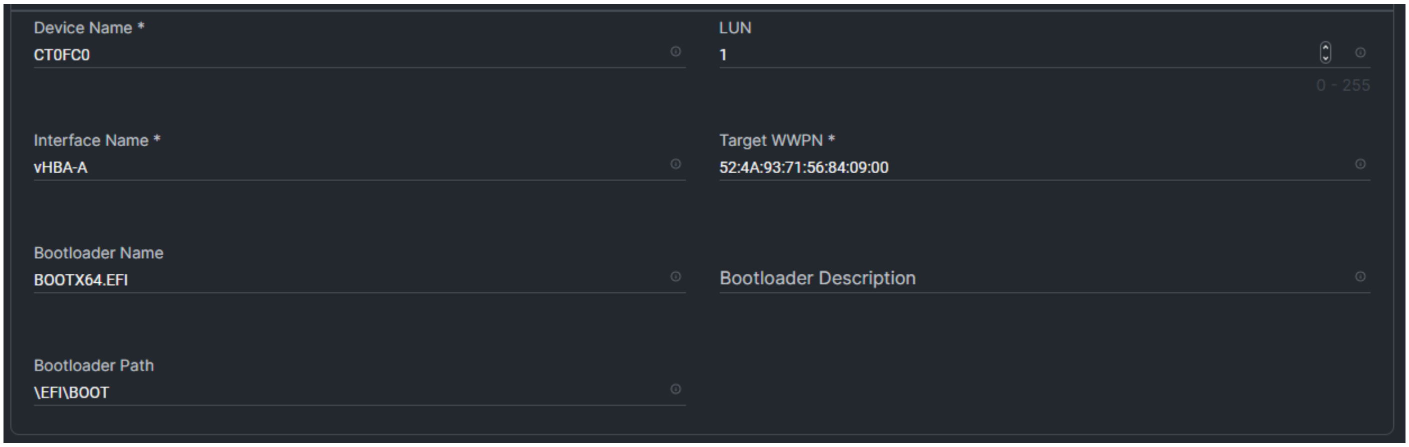

Device Name *
CT0FC0

LUN
1
0 - 255

Interface Name *
vHBA-A

Target WWPN *
52:4A:93:71:56:84:09:00

Bootloader Name
BOOTX64.EFI

Bootloader Description

Bootloader Path
\EFI\BOOT

Step 23. Verify the order of the boot policies and adjust the boot order as necessary using the arrows next to delete icon. Click Create.

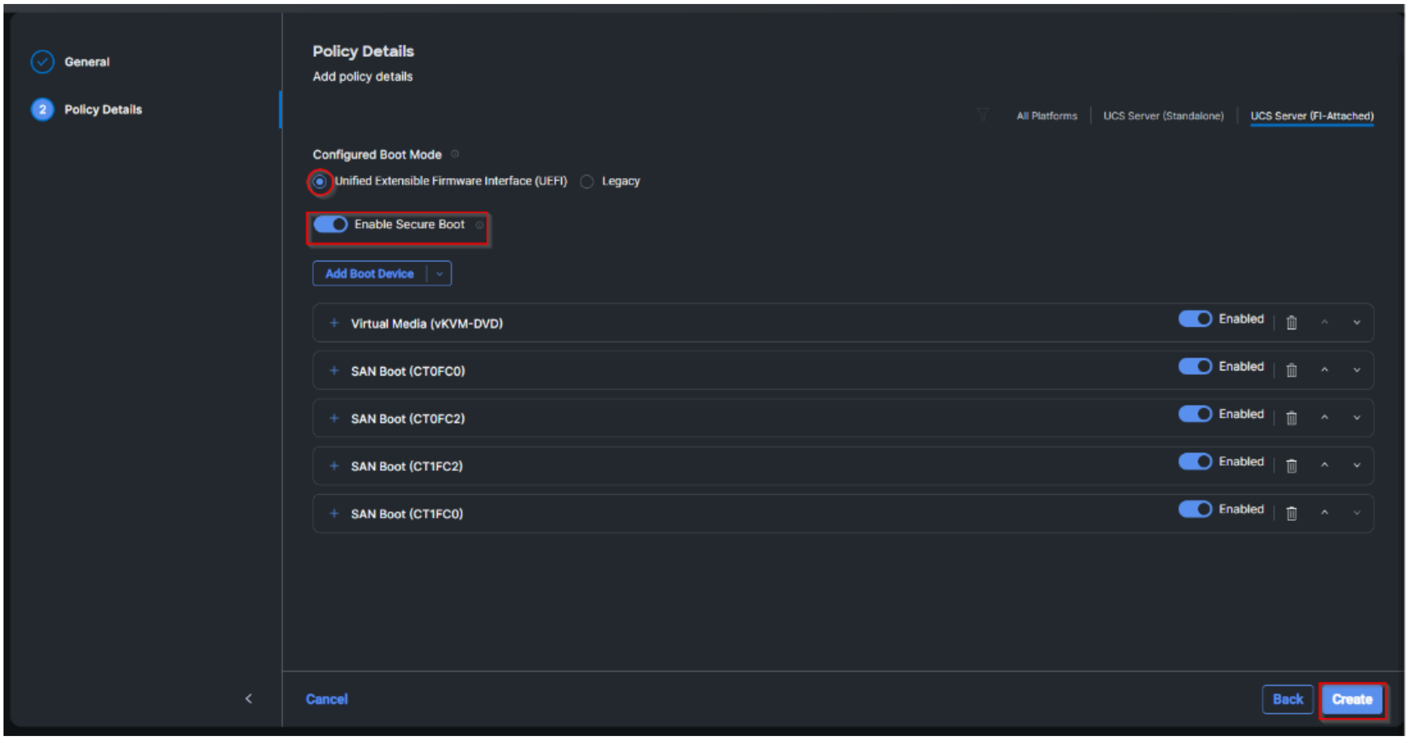

Policy Details
Add policy details

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

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

Enable Secure Boot

+ Virtual Media (vKVM-DVD) Enabled
+ SAN Boot (CT0FC0) Enabled
+ SAN Boot (CT0FC2) Enabled
+ SAN Boot (CT1FC2) Enabled
+ SAN Boot (CT1FC0) Enabled

Step 24. Click Select Policy next to Power and in the pane on the right, click Create New.

Step 25. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, UCS-PWR). Click Next.

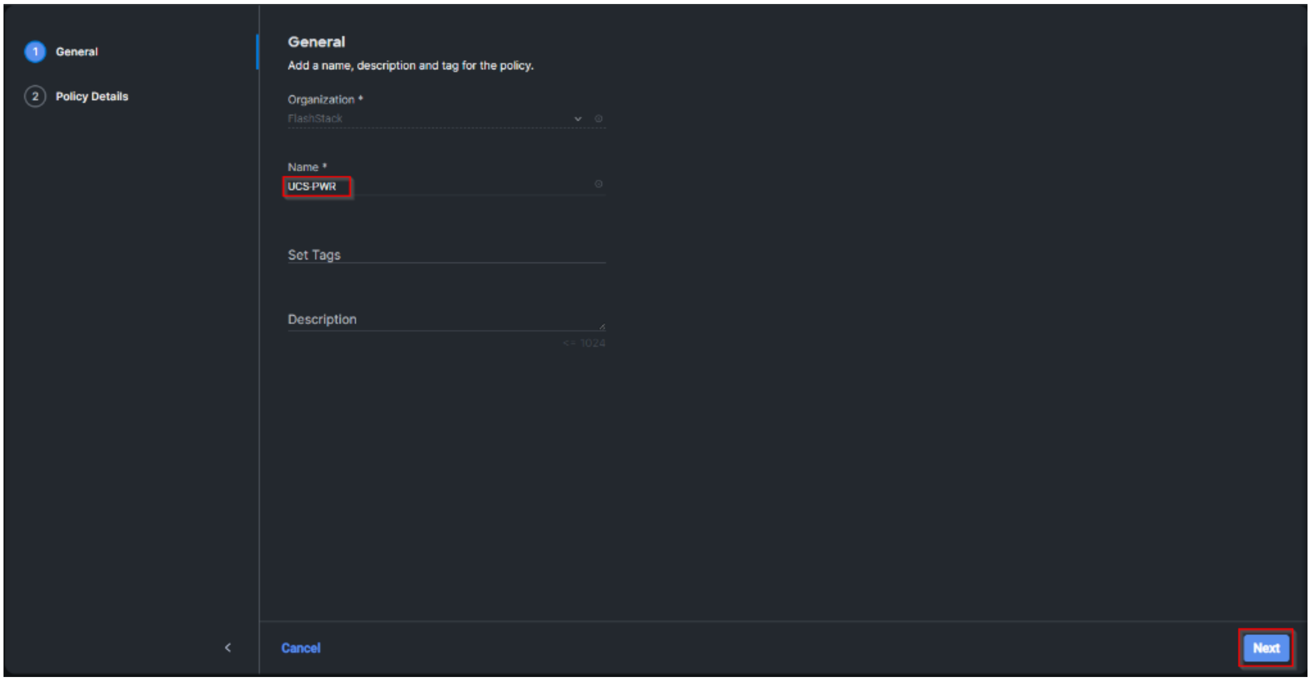

General

Add a name, description and tag for the policy.

Organization *: FlashStack

Name *: UCS-PWR

Set Tags

Description

<= 1024

Step 26. Enable Power Profiling and select High from the Power Priority drop-down list. Click Create.

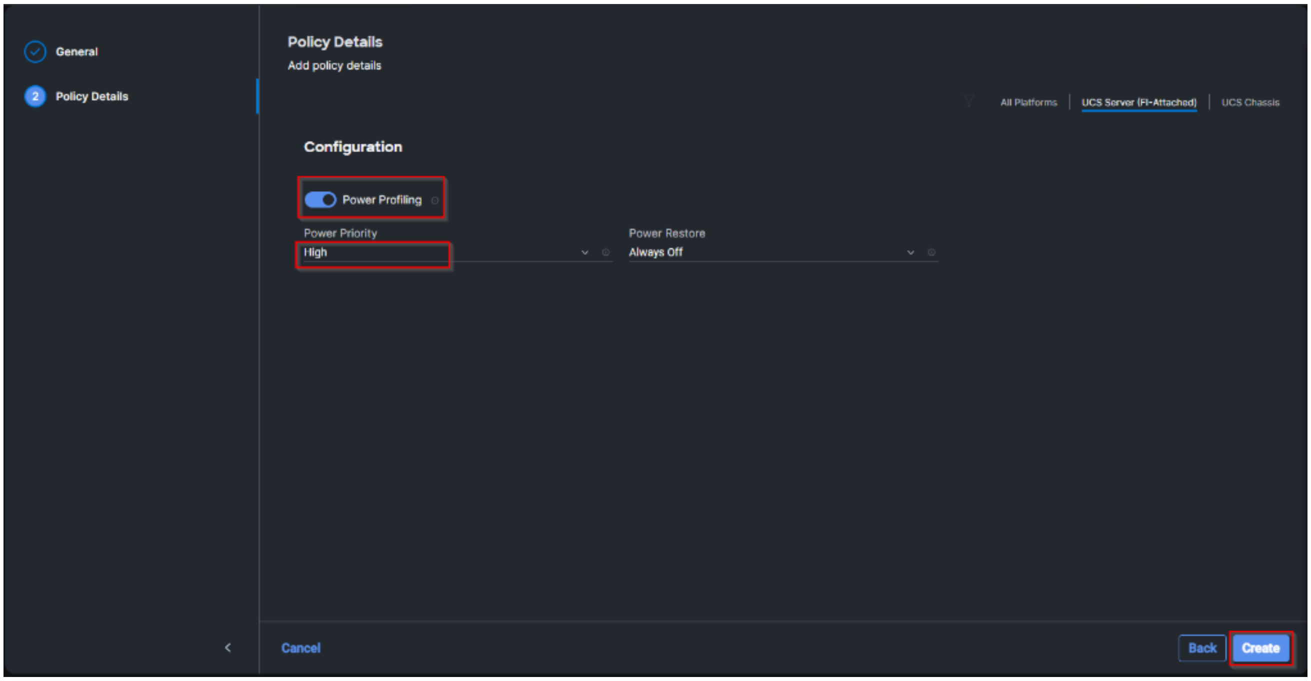

Policy Details

Add policy details

Configuration

Power Profiling: On

Power Priority: High

Power Restore: Always Off

Step 27. Click Select Policy next to Virtual Media and in the pane on the right, click Create New (Optional).

Step 28. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-Vmedia-Pol). Click Next.

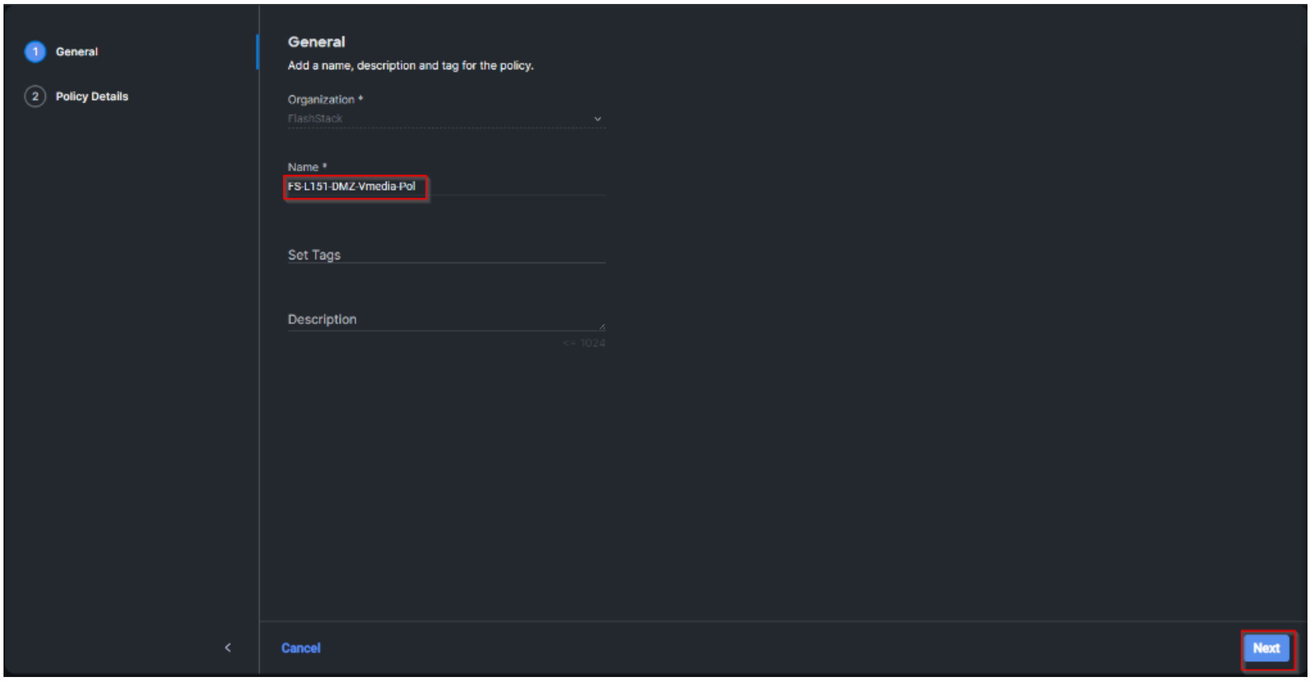

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L151-DMZ-Vmedia-Pol

Set Tags

Description

<= 1024

Step 29. Disable Lower Power USB and click Create.

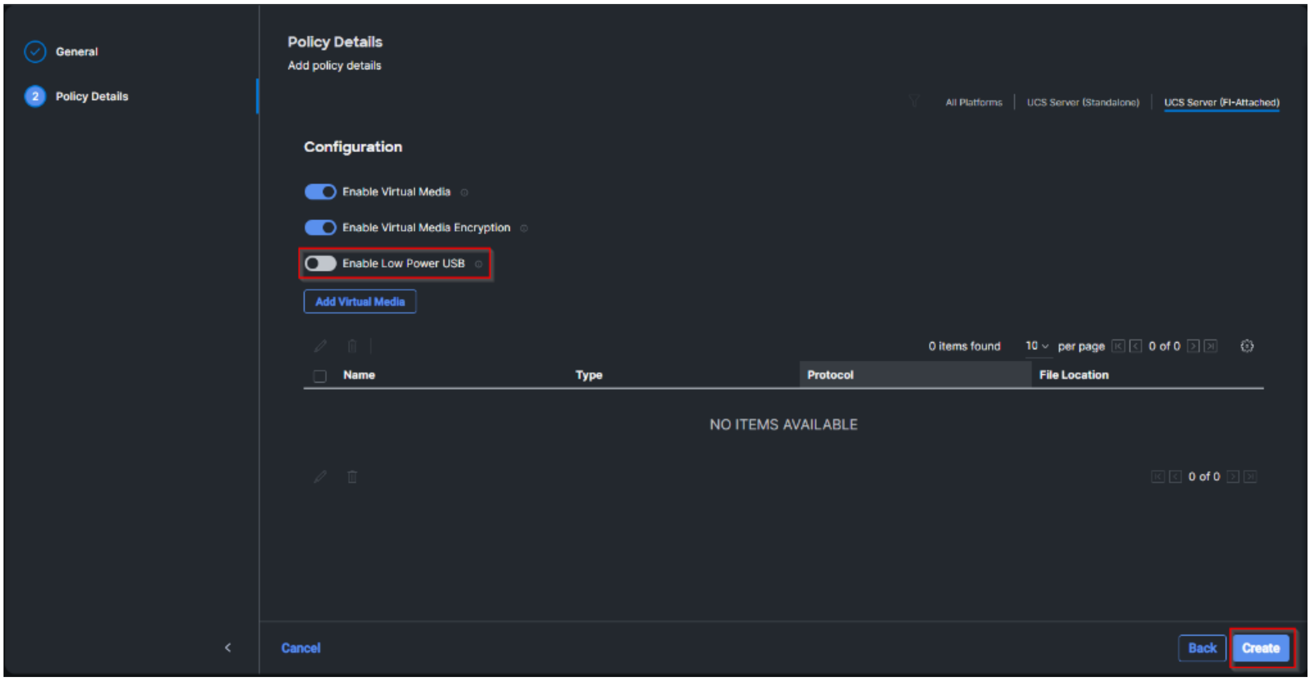

Policy Details

Add policy details

Configuration

● Enable Virtual Media

● Enable Virtual Media Encryption

○ Enable Low Power USB

0 items found 10 per page 0 of 0

□ Name	Type	Protocol	File Location

NO ITEMS AVAILABLE

0 of 0

Step 30. Click Next to go to Management Configuration.

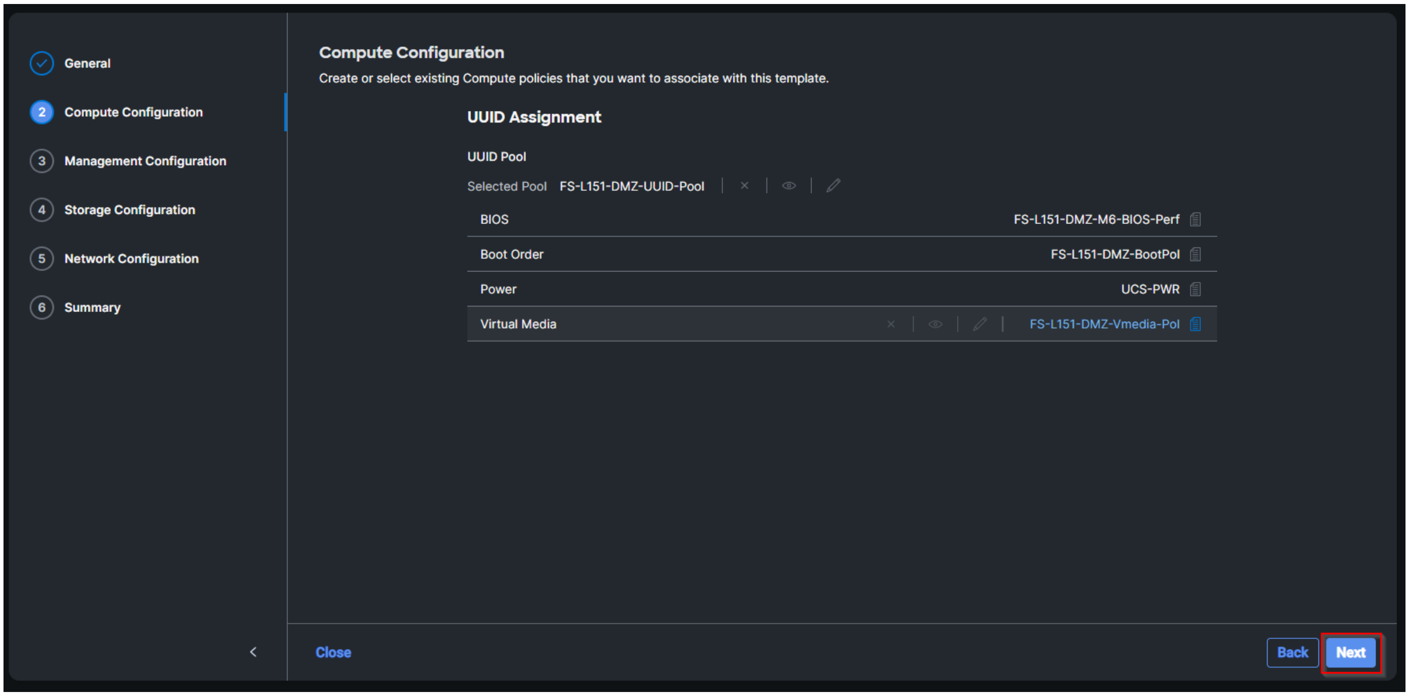

Compute Configuration

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

UUID Assignment

UUID Pool

Selected Pool FS-L151-DMZ-UUID-Pool

BIOS	FS-L151-DMZ-M6-BIOS-Perf
Boot Order	FS-L151-DMZ-BootPol
Power	UCS-PWR
Virtual Media	FS-L151-DMZ-Vmedia-Pol

Step 31. Click Select Policy next to IMC Access and then click Create New.

Step 32. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-IMCAPol). Click Next.

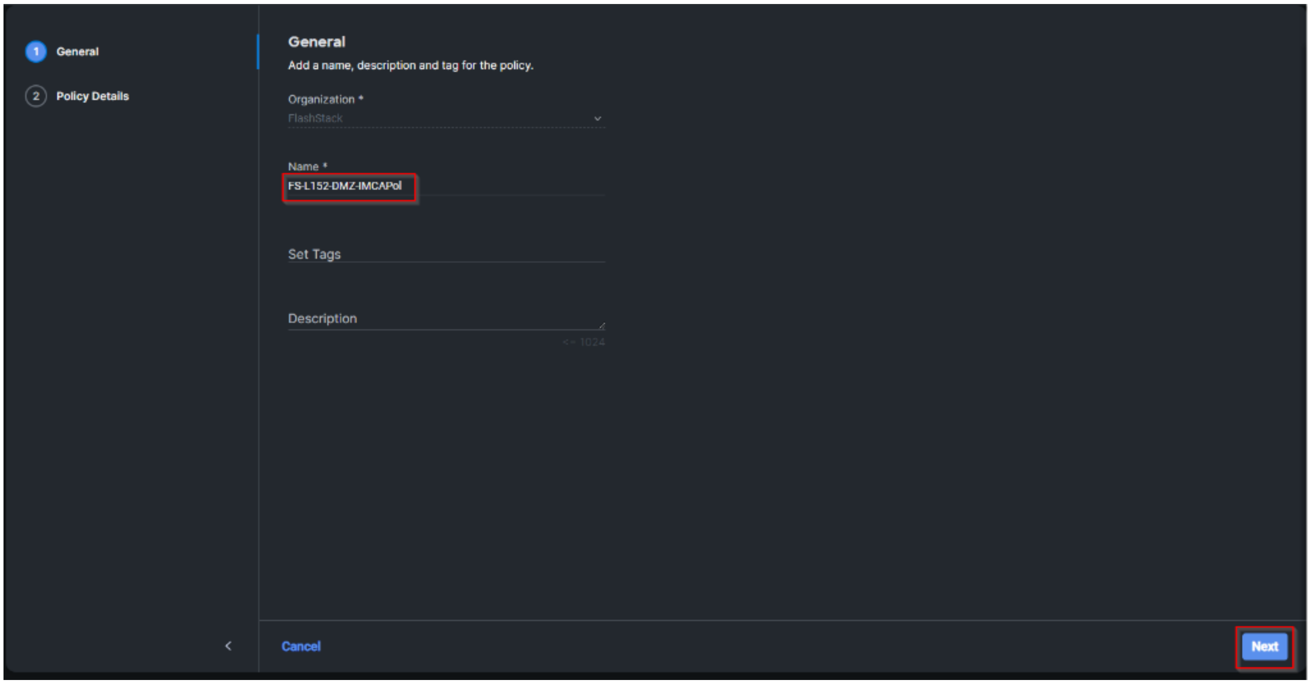

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * FS-L152-DMZ-IMCAPol

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

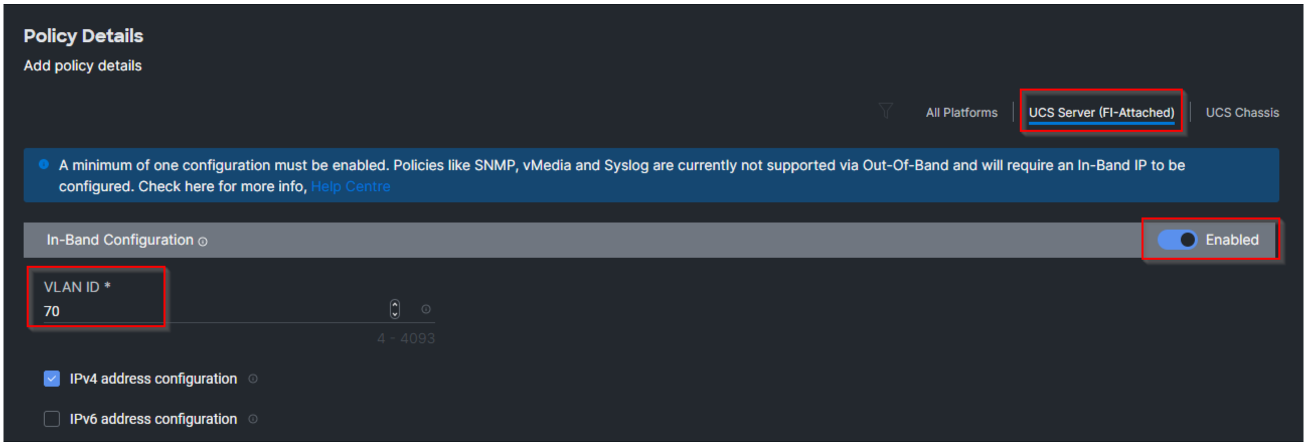

Policy Details

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, Help Centre

In-Band Configuration: Enabled

VLAN ID *: 70

4 - 4093

☑ IPv4 address configuration

□ IPv6 address configuration

Step 34. Under IP Pool, click Select IP Pool and then click Create New.

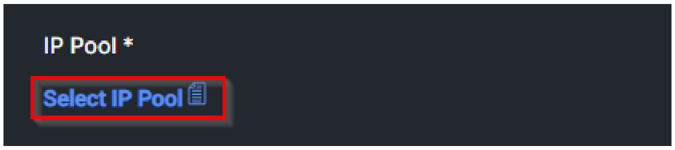

IP Pool *

Step 35. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L152-DMZ-ICMA-IP-Pool). Click Next.

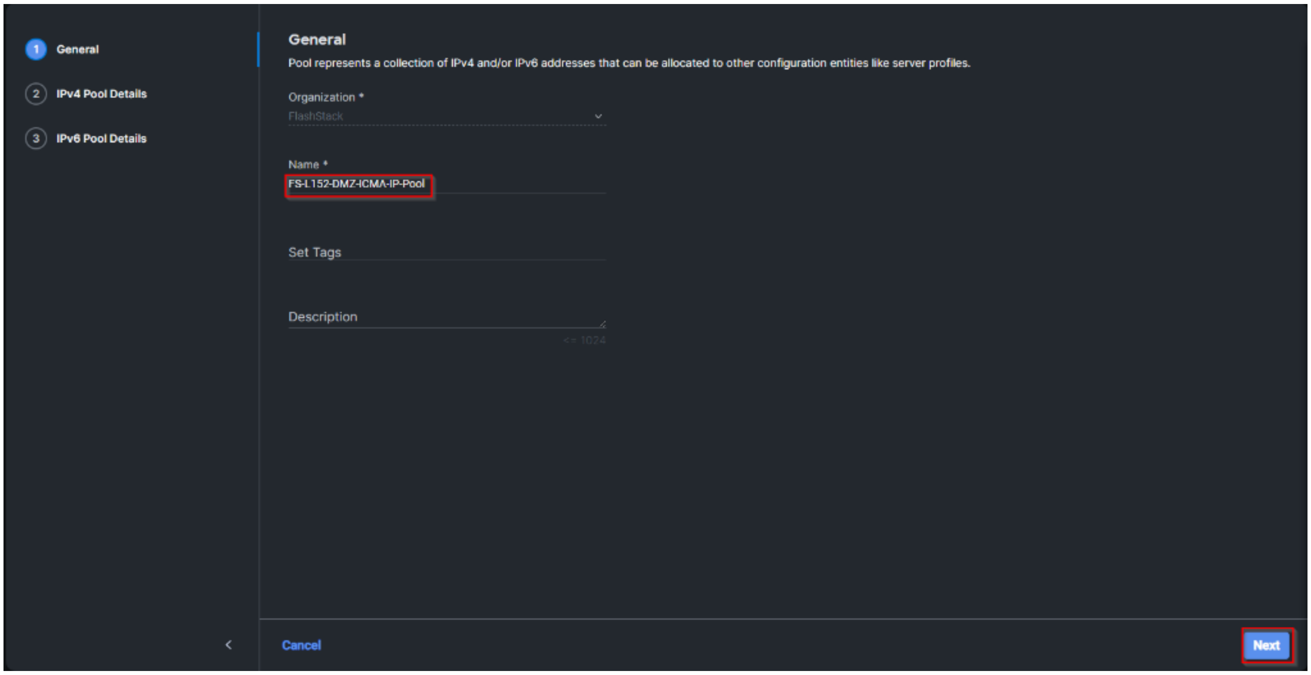

General

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

Organization *: FlashStack

Name *: FS-L152-DMZ-ICMA-IP-Pool

Set Tags

Description

<= 1024

Step 36. Select Configure IPv4 Pool and provide the information to define a pool for KVM IP address assignment including an IP Block.

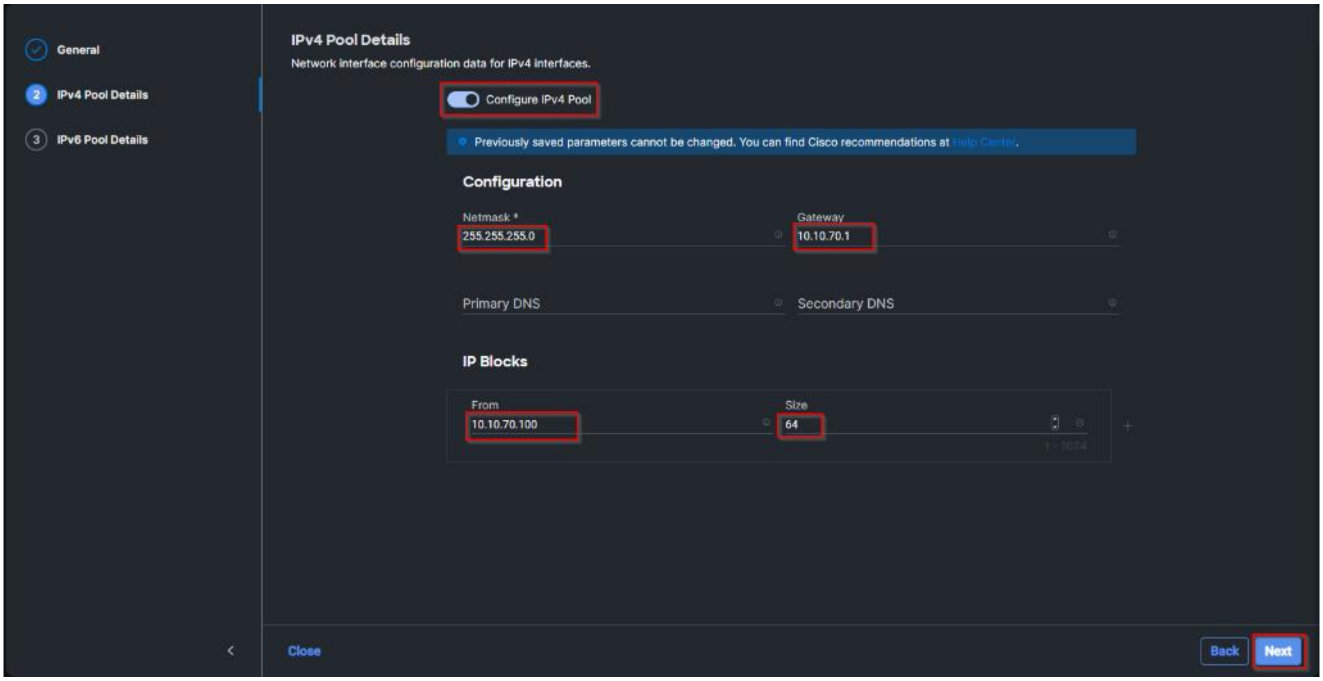

IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

Configure IPv4 Pool

Previously saved parameters cannot be changed. You can find Cisco recommendations at Help Center.

Configuration

Netmask * 255.255.255.0

Gateway 10.10.70.1

Primary DNS

Secondary DNS

IP Blocks

From 10.10.70.100

Size 64

1 - 1024

Note: The management IP pool subnet should be accessible from the host that is trying to open the KVM connection. In the example shown here, the hosts trying to open a KVM connection would need to be able to route to 10.10.70.0/24 subnet.

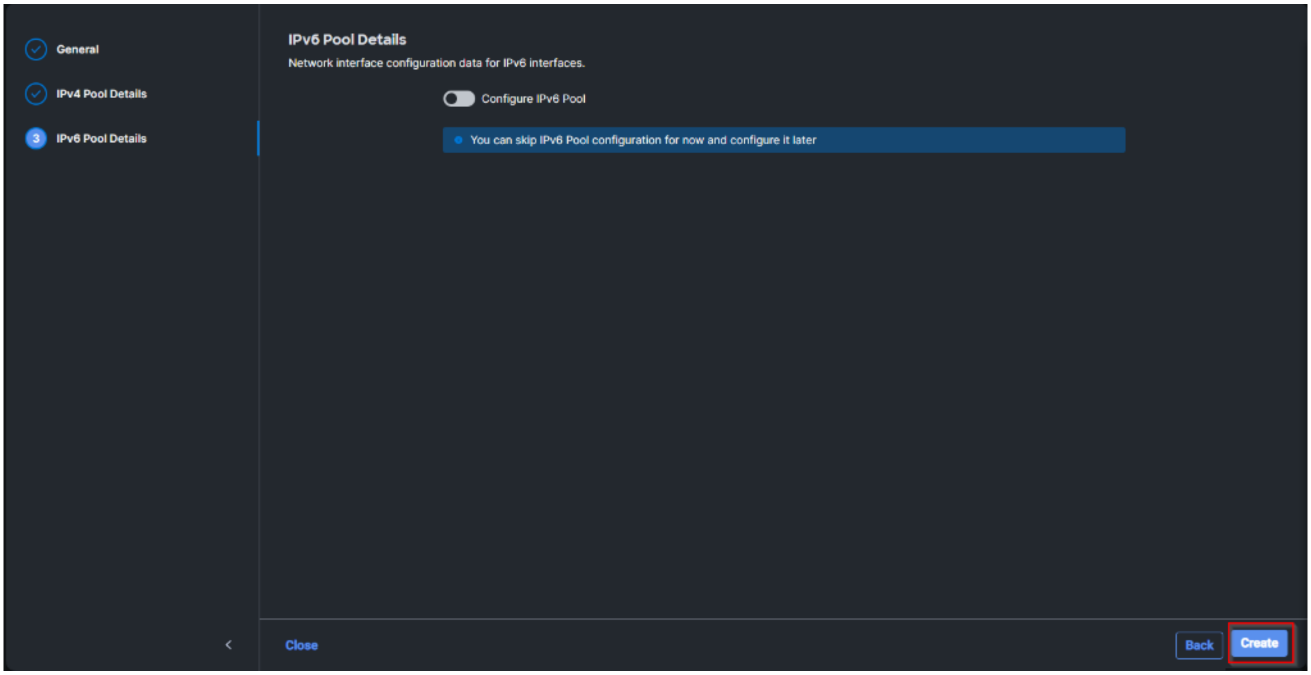

IPv6 Pool Details

Network interface configuration data for IPv6 interfaces.

Configure IPv6 Pool

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

Step 37. Click Select Policy next to IPMI Over LAN and then click Create New.

Step 38. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, Enable-IPMIoLAN). Click Next.

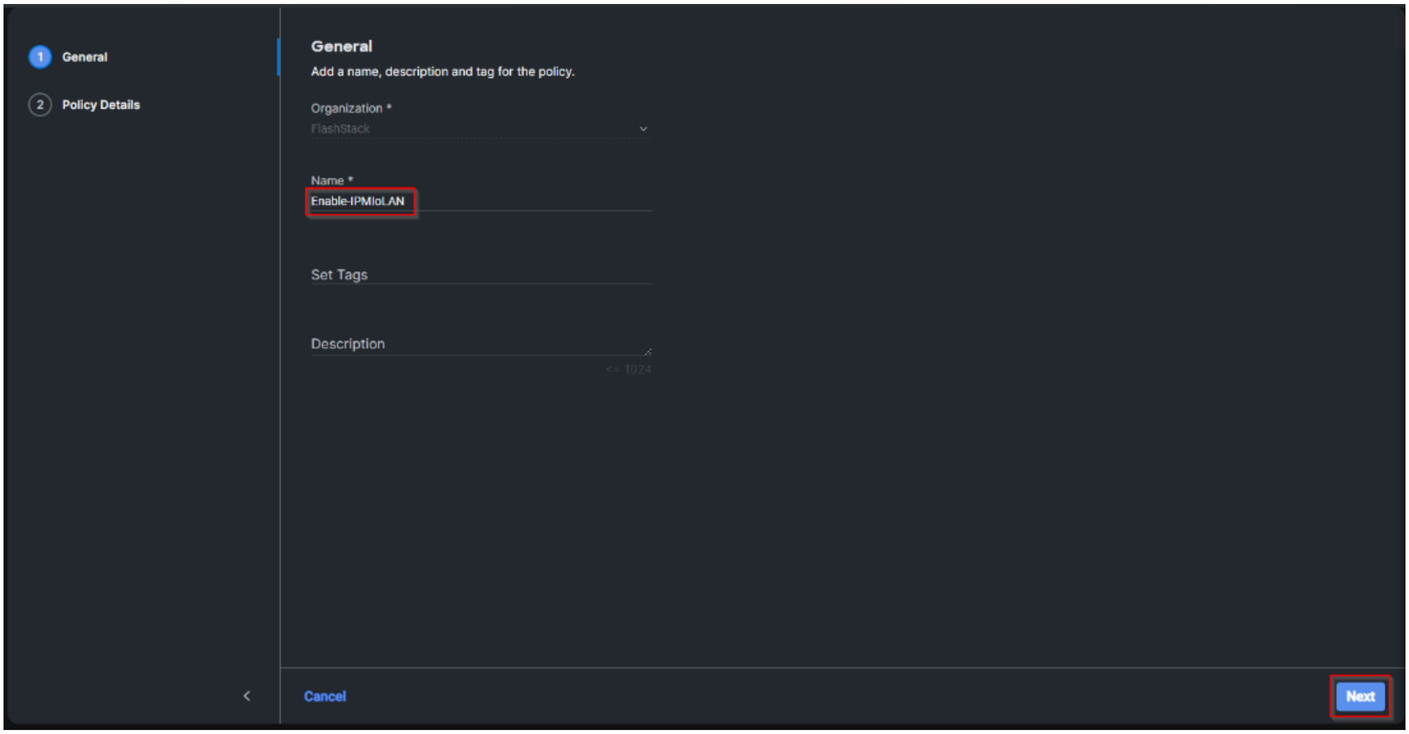

General

Add a name, description and tag for the policy.

Organization * FlashStack

Name * Enable-IPMIoLAN

Set Tags

Description

Step 39. Turn on Enable IPMI Over LAN.

Step 40. From the Privilege Level drop-down list, select admin.

Step 41. Click Create.

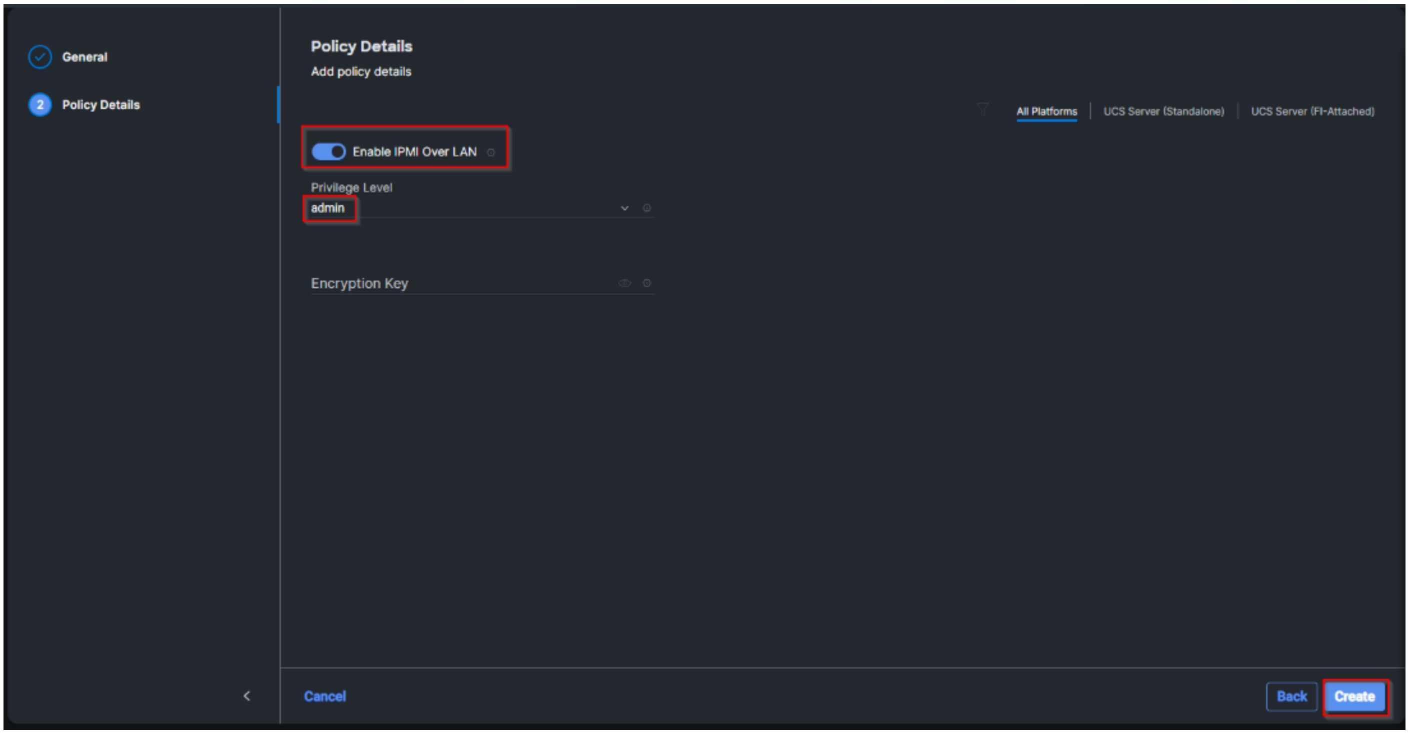

Policy Details

Add policy details

Enable IPMI Over LAN

Privilege Level admin

Encryption Key

Step 42. Click Select Policy next to Local User and the, in the pane on the right, click Create New.

Step 43. Verify the correct organization is selected from the drop-down list and provide a name for the policy.

Step 44. Verify that UCS Server (FI-Attached) is selected.

Step 45. Verify that Enforce Strong Password is selected.

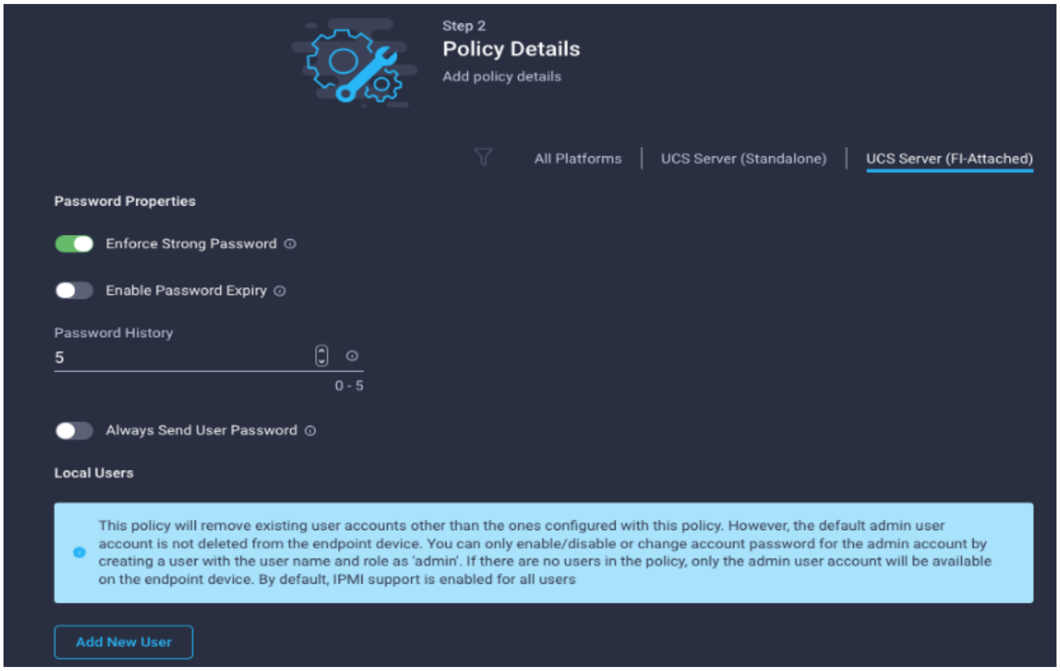

Step 2
Policy Details
Add policy details

Password Properties

Enforce Strong Password: ● (on)

Enable Password Expiry: ○ (off)

Password History: 5 (0 - 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 46. Click Add New User and then click + next to the New User.

Step 47. Provide the username (for example, fpadmin), choose a role for example, admin), and provide a password.

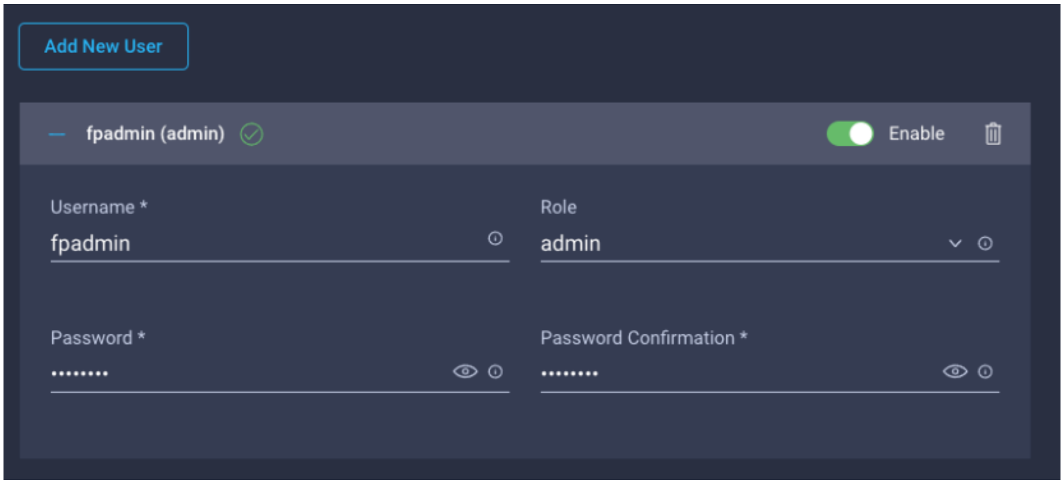

Username *: fpadmin

Role: admin

Password *: ••••••••

Password Confirmation *: ••••••••

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 48. Click Create to finish configuring the user.

Step 49. Click Create to finish configuring local user policy.

Step 50. Click Next to move to Storage Configuration.

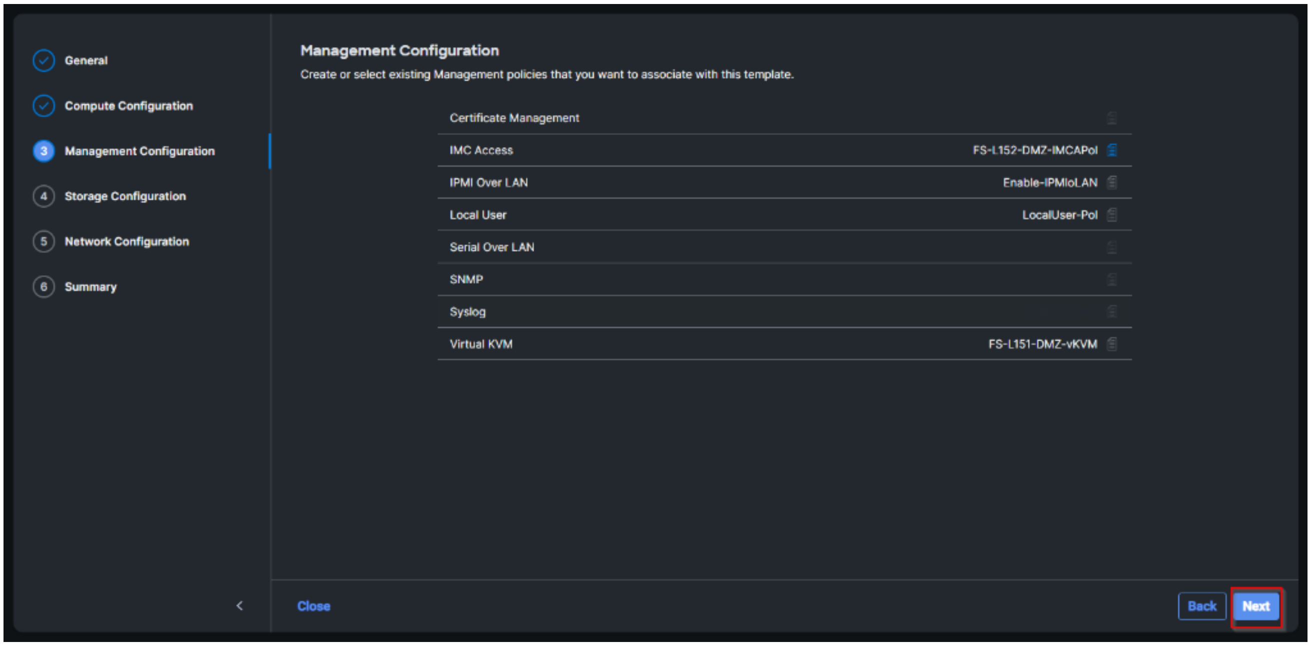

Management Configuration

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

Certificate Management	
IMC Access	FS-L152-DMZ-IMCAPol
IPMI Over LAN	Enable-IPMIoLAN
Local User	LocalUser-Pol
Serial Over LAN	
SNMP	
Syslog	
Virtual KVM	FS-L151-DMZ-vKVM

Step 51. Click Next on the Storage Configuration screen. No configuration is needed in the local storage system.

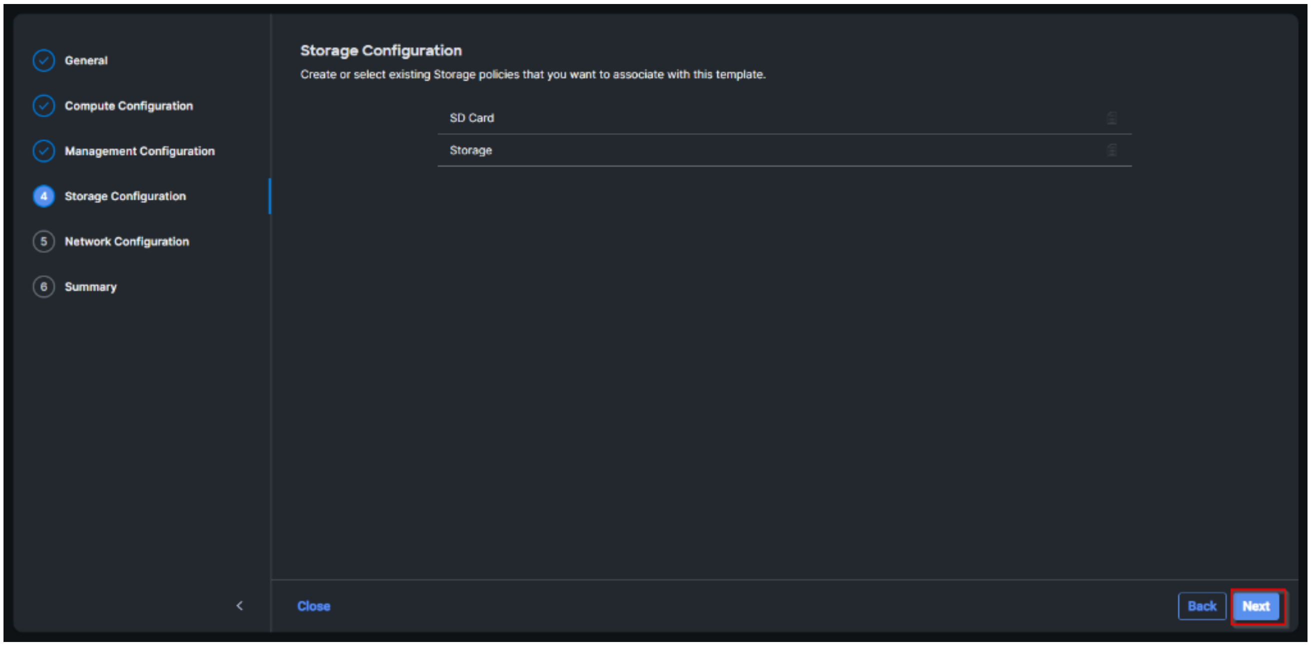

Storage Configuration

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

SD Card
Storage

Step 52. 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.

The FC boot from SAN hosts uses 4 vNICs configured as listed in Table 6:

Table 6. vNICs for FC LAN Connectivity

vNIC	Slot ID	Switch ID	PCI Order	VLANs
vSwitch0-A	MLOM	A	2	IB-MGMT, NFS
vSwitch0-B	MLOM	B	3	IB-MGMT, NFS
VDS0-A	MLOM	A	4	VM Traffic, vMotion
VDS0-B	MLOM	B	5	VM Traffic, vMotion

Note: The PCI order 0 and 1 will be used in the SAN Connectivity policy to create vHBA-A and vHBA-B.

Step 53. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-LAN-Conn-Pol). Click Next.

Step 54. Under vNIC Configuration, select Manual vNICs Placement.

Step 55. Click Add vNIC.

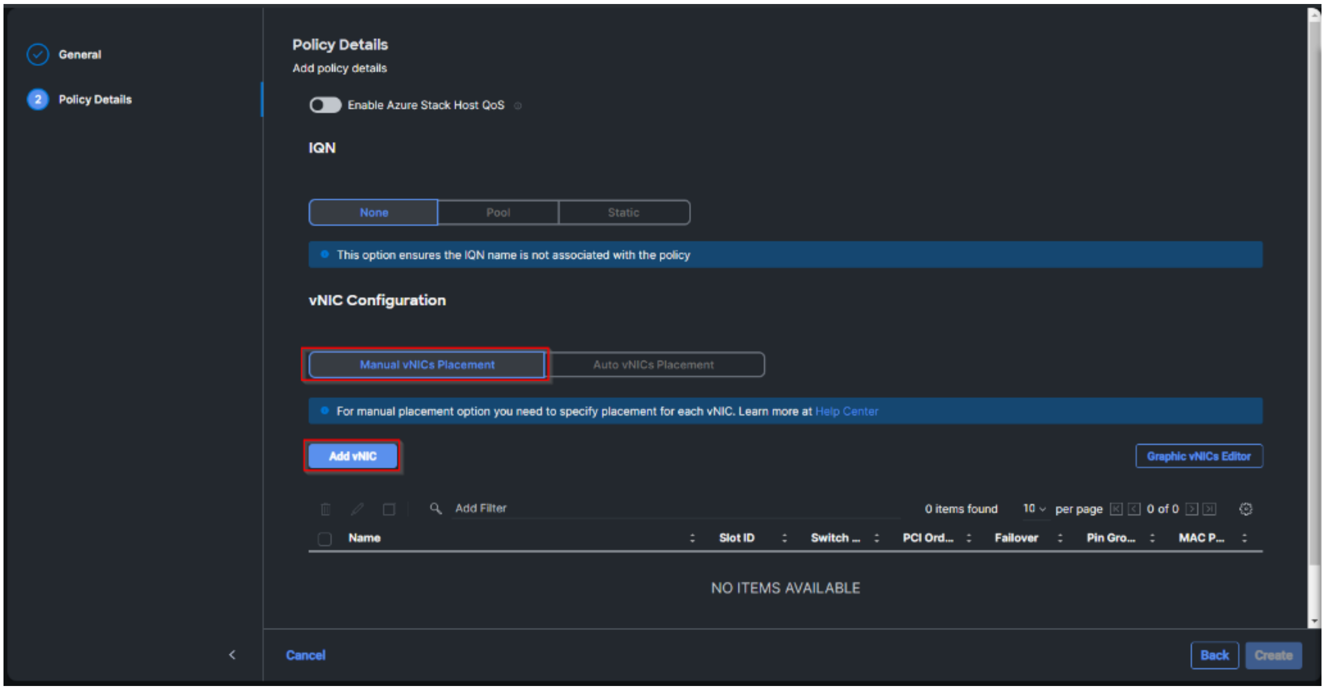

Policy Details

Add policy details

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

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

NO ITEMS AVAILABLE

Step 56. Click Select Pool under MAC Address 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.

Table 7. MAC Address Pools

Pool Name	Starting MAC Address	Size	vNICs
FS-L151-DMZ-MAC-Pool-A	00:25:B5:04:0A:00	256*	vSwitch0-A, VDS0-A
FS-L151-DMZ-MAC-Pool-B	00:25:B5:04:0B:00	256*	vSwitch0-B, VDS0-B

Step 57. Verify the correct organization is selected from the drop-down list and provide a name for the pool from Table 7 depending on the vNIC being created (for example, FS-L151-DMZ-MAC-Pool-A for Fabric A).

Step 58. Click Next.

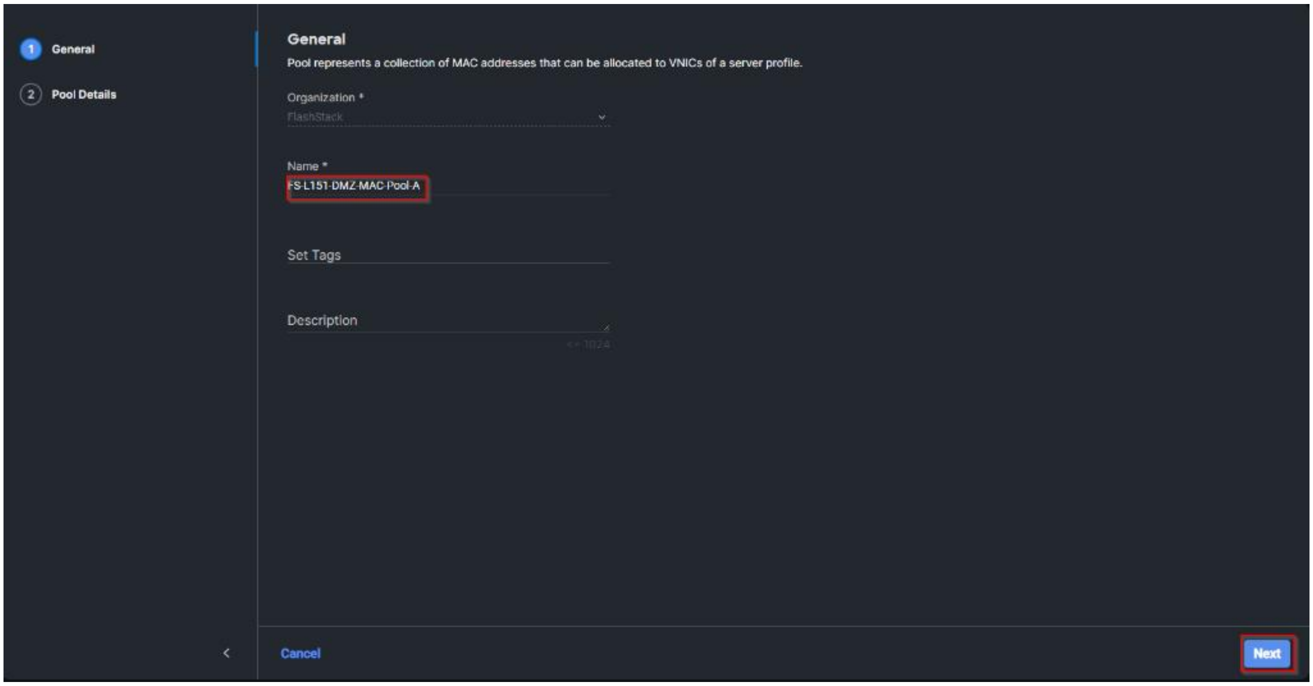

General

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

Organization * FlashStack

Name * FS-L151-DMZ-MAC-Pool-A

Set Tags

Description

<= 1024

Step 59. Provide the starting MAC address from Table 7 (for example, 00:25:B5:04:0A:00) and the size of the MAC address pool (for example, 256). Click Create to finish creating the MAC address pool.

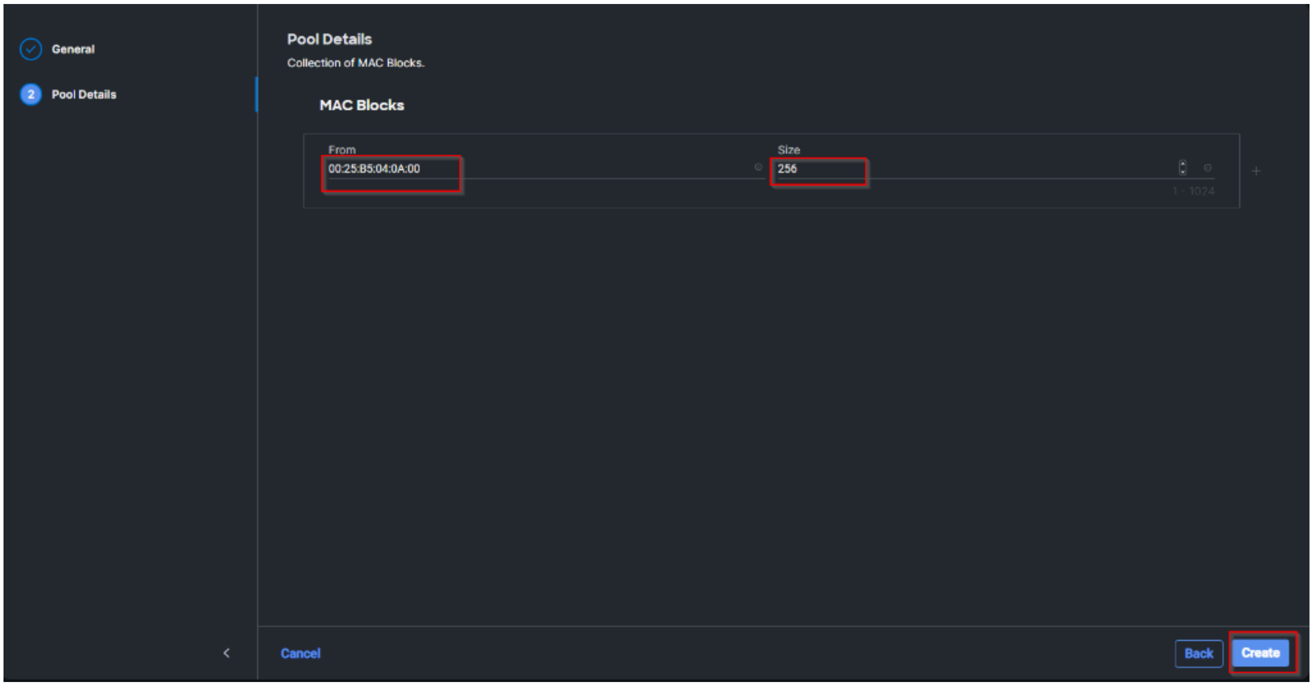

Pool Details

Collection of MAC Blocks.

MAC Blocks

From 00:25:B5:04:0A:00

Size 256

1 - 1024

Step 60. From the Add vNIC window, provide vNIC Name, Slot ID, Switch ID, and PCI Order information from Table 7.

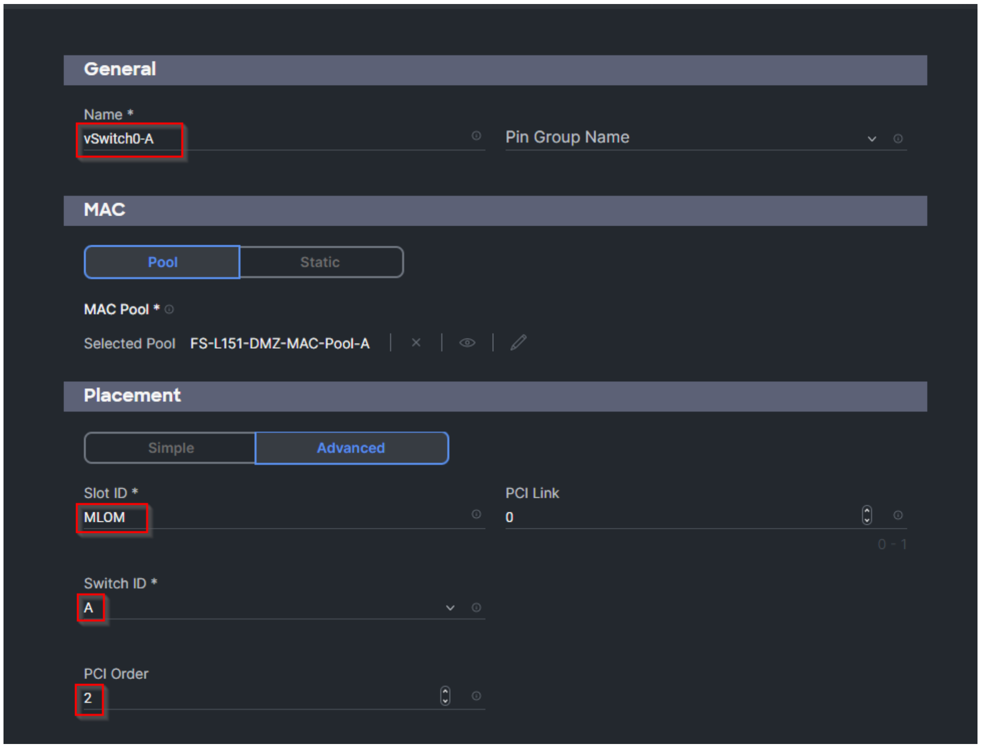

General

Name * vSwitch0-A

Pin Group Name

MAC

MAC Pool *

Selected Pool FS-L151-DMZ-MAC-Pool-A

Placement

Slot ID * MLOM

PCI Link 0 (0 - 1)

Switch ID * A

PCI Order 2

Step 61. For Consistent Device Naming (CDN), from the drop-down list, select vNIC Name.

Step 62. Verify that Failover is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and VDS.

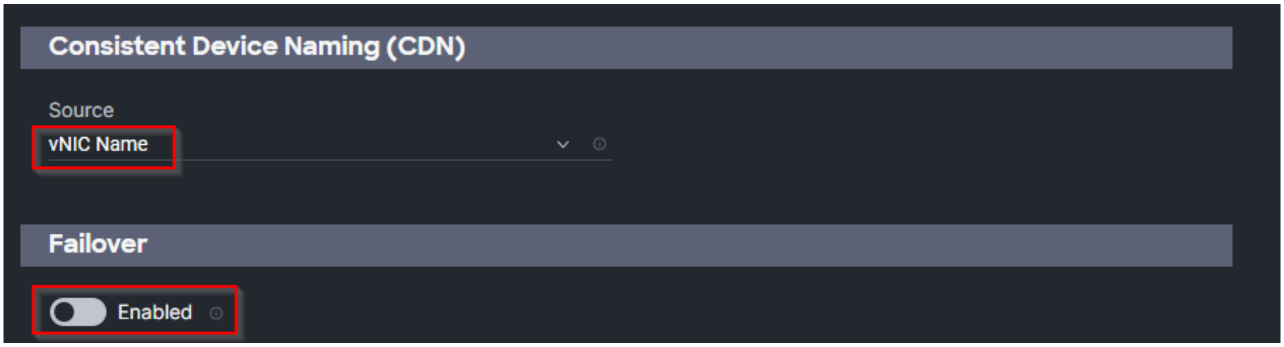

Consistent Device Naming (CDN)

Source vNIC Name

Failover

Enabled

Step 63. 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 follows:

Table 8. Ethernet Group Policy Values

Group Policy Name	Native VLAN	Apply to vNICs	VLANs
FS-L151-DMZ-vSwitch0-NetGrp-Pol	Native-VLAN (1)	vSwitch0-A, vSwitch0-B	IB-MGMT
FS-L151-DMZ-vSwitch1-NetGrp-Pol	Native-VLAN (1)	VDS0-A, VDS0-B	VM Traffic, vMotion

Step 64. Verify the correct organization is selected from the drop-down list and provide a name for the policy from Table 8 (for example, FS-L151-DMZ-vSwitch0-NetGrp-Pol). Click Next.

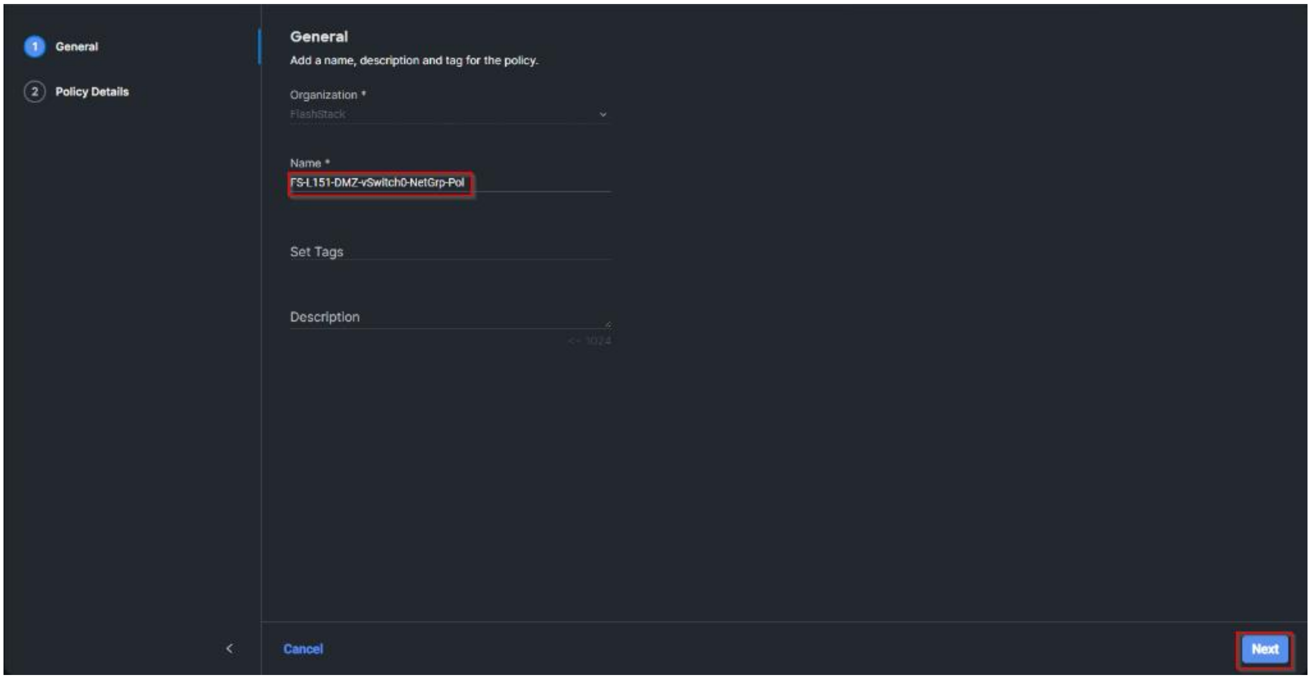

Organization * FlashStack

Name * FS-L151-DMZ-vSwitch0-NetGrp-Pol

Step 65. Enter the allowed VLANs from Table 7 (for example, 70) and the native VLAN ID from Table 8 (for example, 1). Click Create.

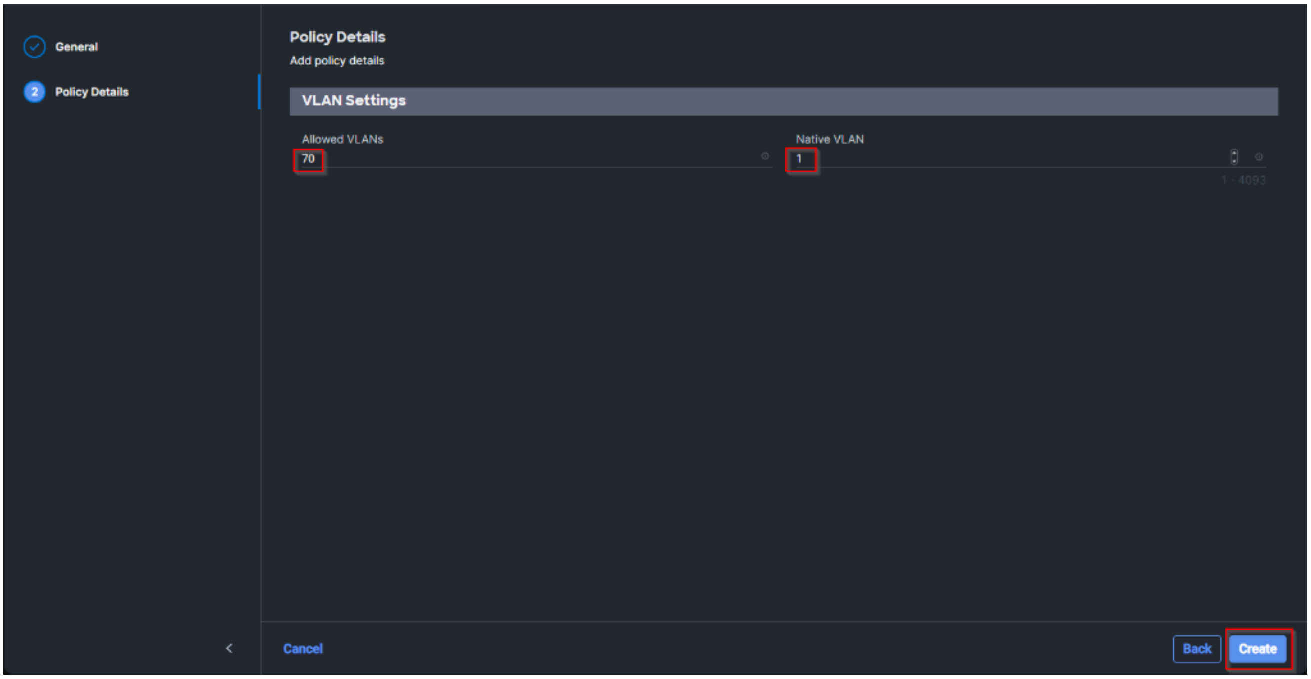

Allowed VLANs 70

Native VLAN 1

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 66. 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.

Step 67. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-NetCtrl-Pol).

Step 68. Click Next.

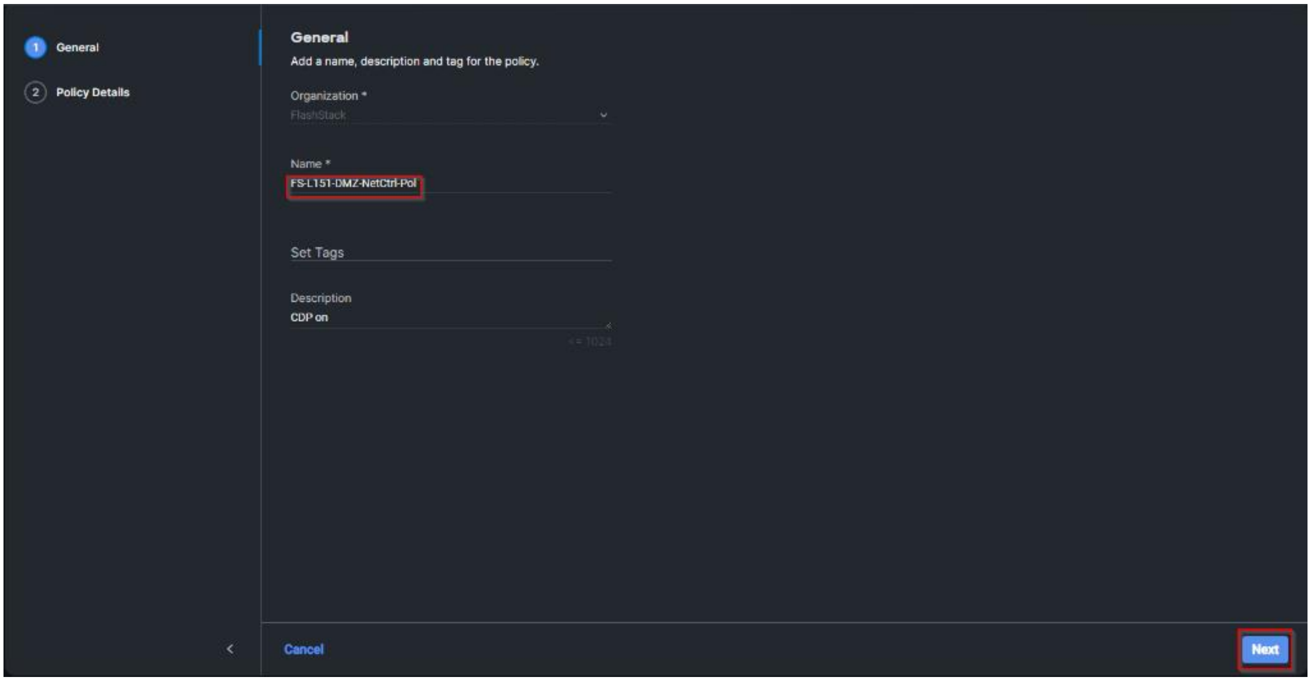

Organization * FlashStack

Name * FS-L151-DMZ-NetCtrl-Pol

Set Tags

Description CDP on

Step 69. Enable Cisco Discovery Protocol and both Enable Transmit and Enable Receive under LLDP. Click Create.

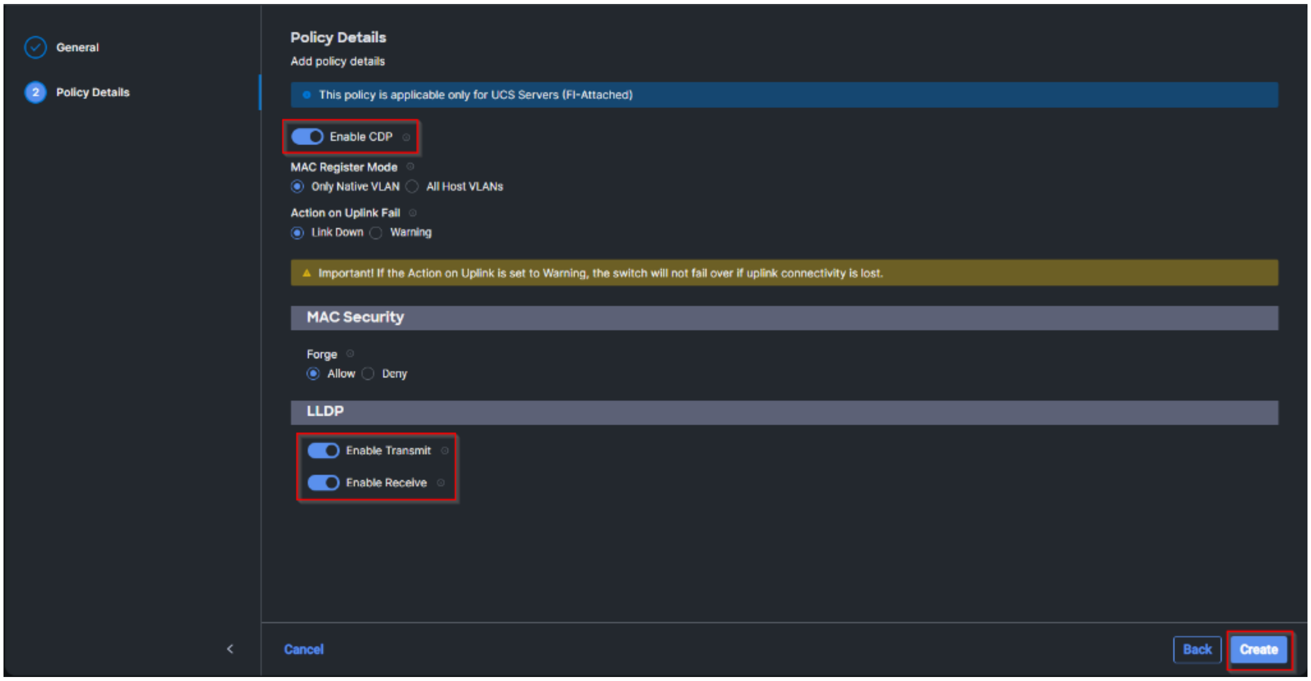

This policy is applicable only for UCS Servers (FI-Attached)

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 70. 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.

Step 71. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-QOS).

Step 72. Click Next.

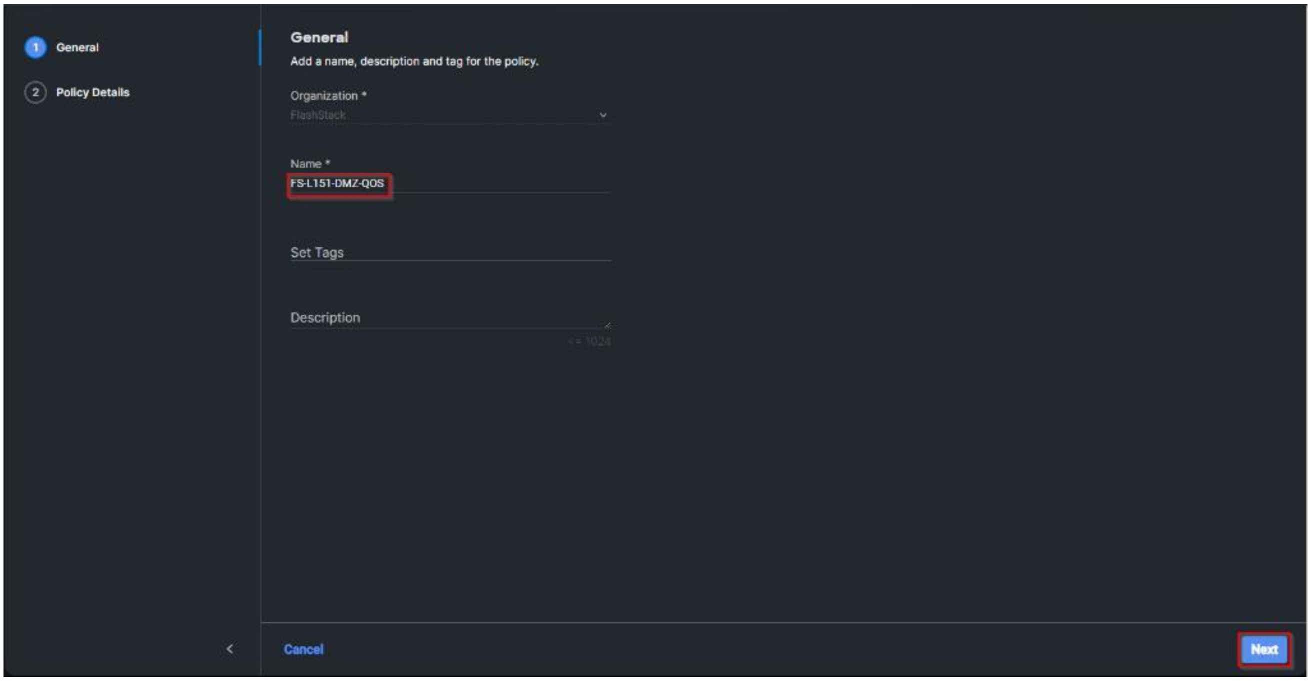

General

Add a name, description and tag for the policy.

Organization *
FlashStack

Name *
FS-L151-DMZ-QOS

Set Tags

Description
<= 1024

Step 73. Change the MTU Bytes value to 9000. Click Create.

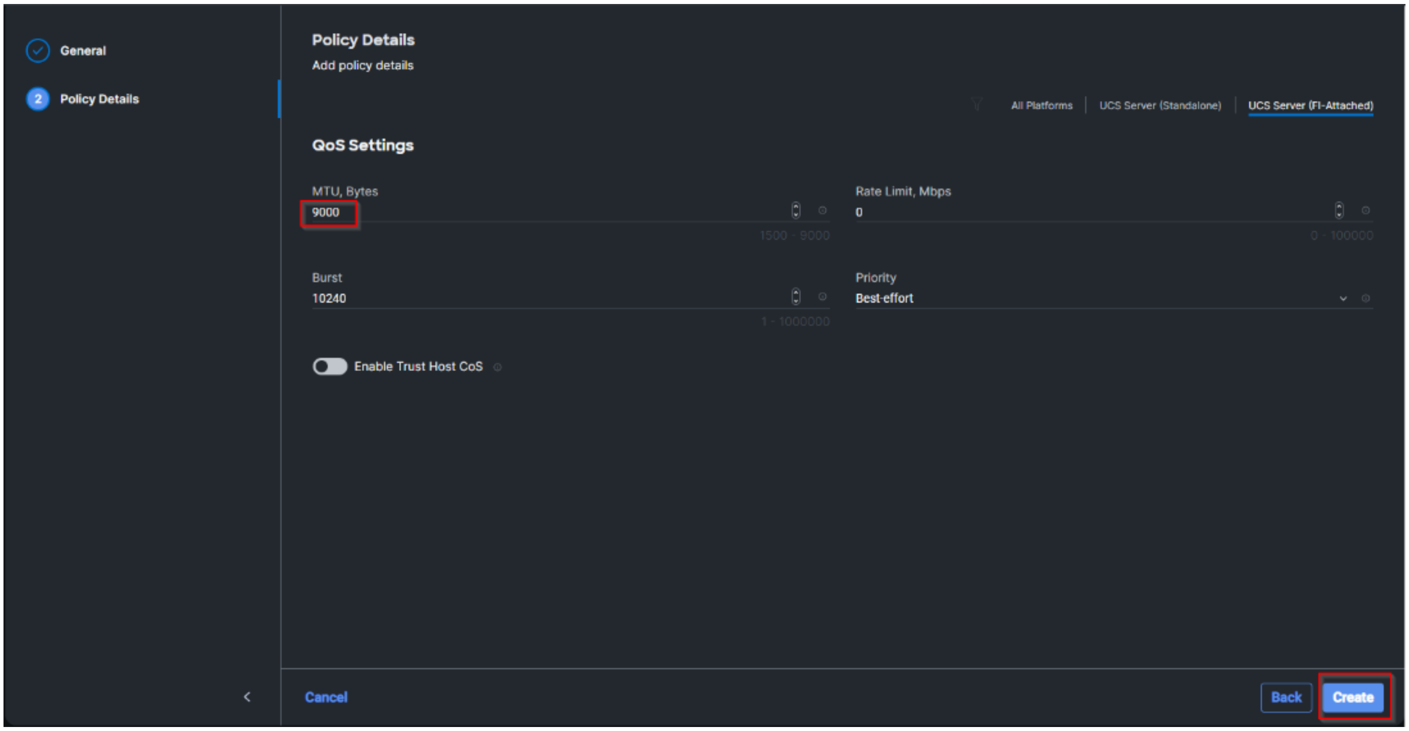

Policy Details

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 74. 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, FS-L151-DMZ-EthAdapt-VMware-HiTraffic, is created and attached to the VDS0-A and VDS0-B interfaces which handle vMotion.

Table 9. Ethernet Adapter Policy association to vNICs

Policy Name	vNICS
FS-L151-DMZ-EthAdapt-VMware	vSwitch0-A, vSwitch0-B
FS-L151-DMZ-EthAdapt-VMware-HiTraffic	VDS0-A, VDS0-B,

Step 75. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-EthAdapt-VMware).

Step 76. Click Select Default Configuration under Ethernet Adapter Default Configuration.

Step 77. From the list, select VMware. Click Next.

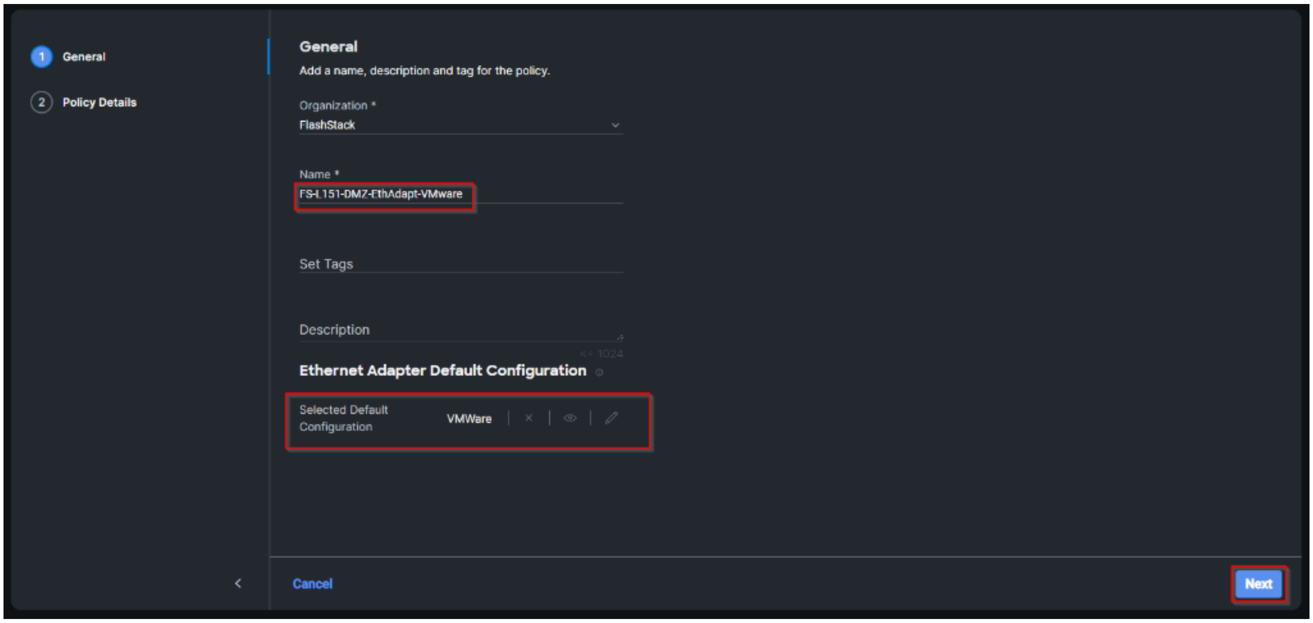

Organization *
FlashStack

Name *
FS-L151-DMZ-EthAdapt-VMware

Ethernet Adapter Default Configuration

Selected Default Configuration VMWare

Step 78. For the FS-L151-DMZ-EthAdapt-VMware policy, click Create and skip the rest of the steps in this section.

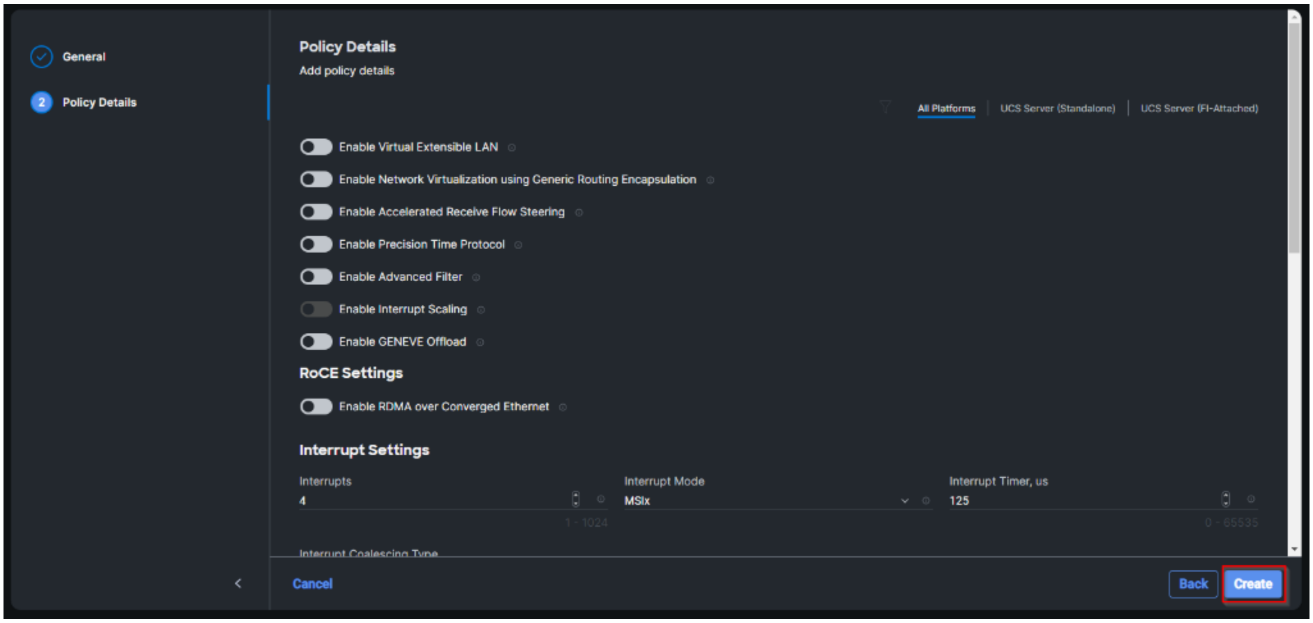

Policy Details

Add policy details

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: 4 (1 - 1024)

Interrupt Mode: MSIx

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type

Step 79. For the optional FS-L151-DMZ-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

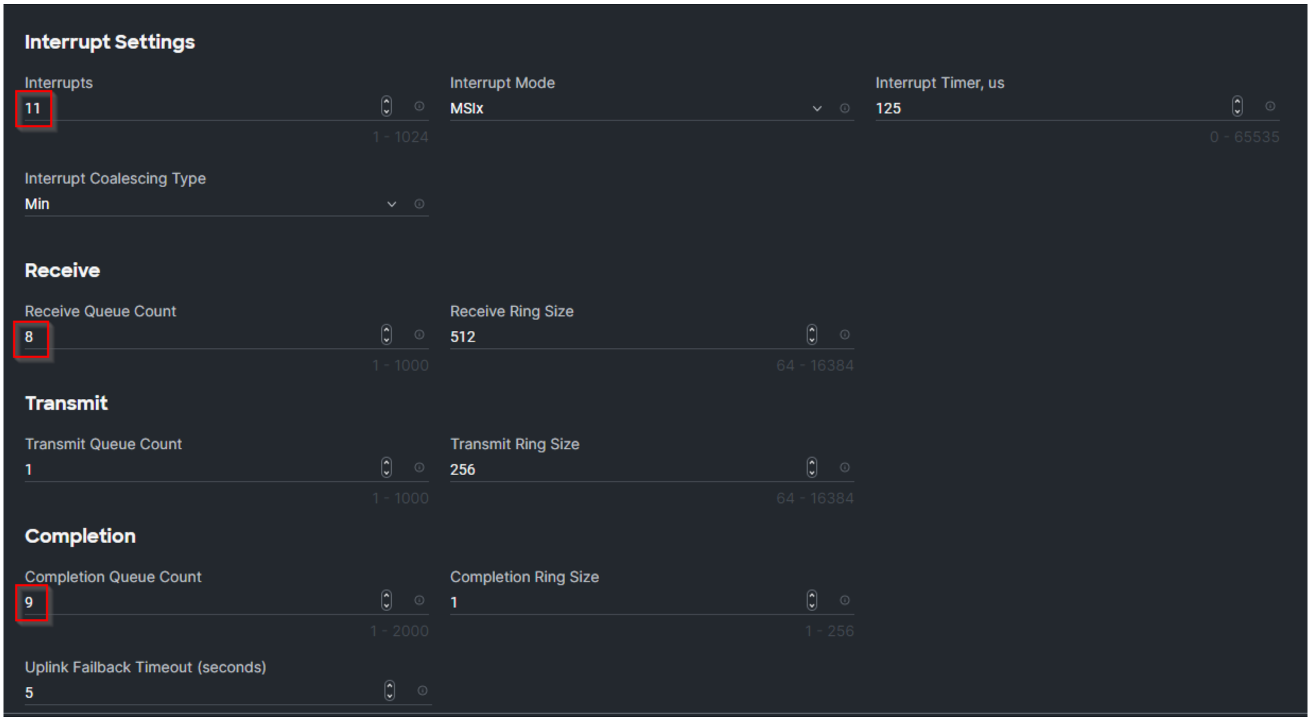

Interrupt Settings

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

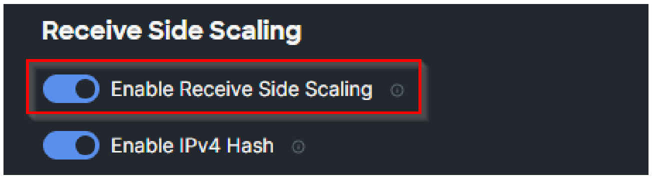

Receive Side Scaling

Enable Receive Side Scaling

Enable IPv4 Hash

Step 80. Click Create.

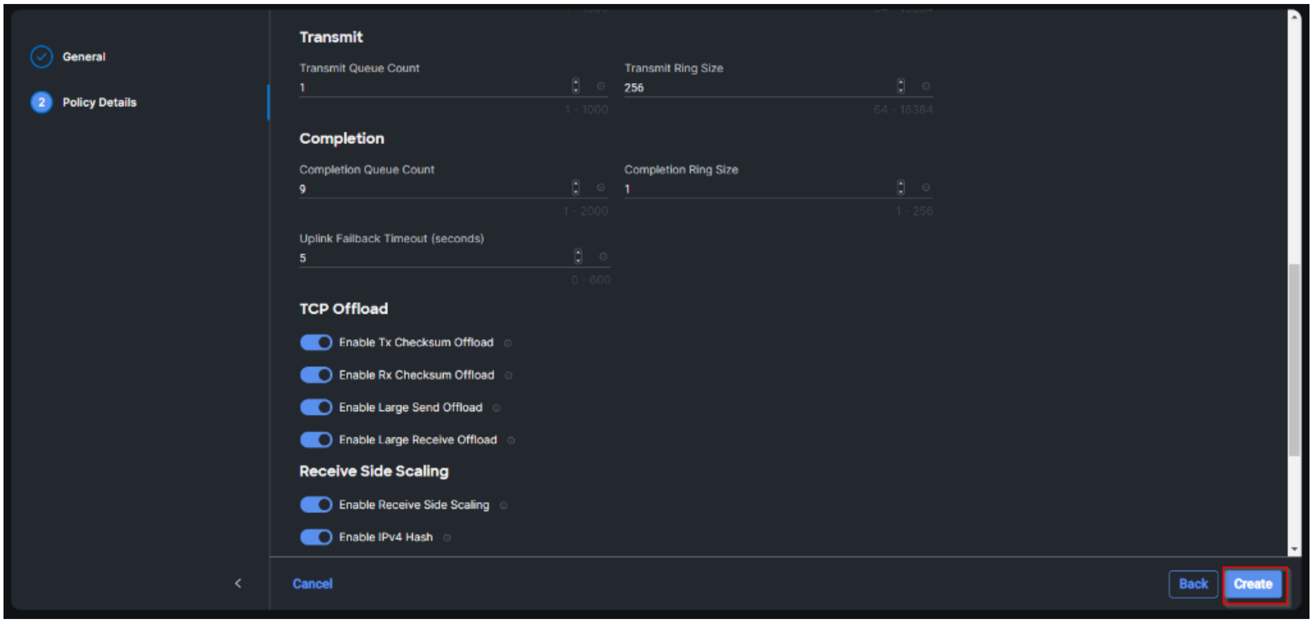

Transmit

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 81. Click Create to finish creating the vNIC.

Step 82. Repeat the vNIC creation steps for the rest of vNICs. Verify all four vNICs were successfully created. Click Create.

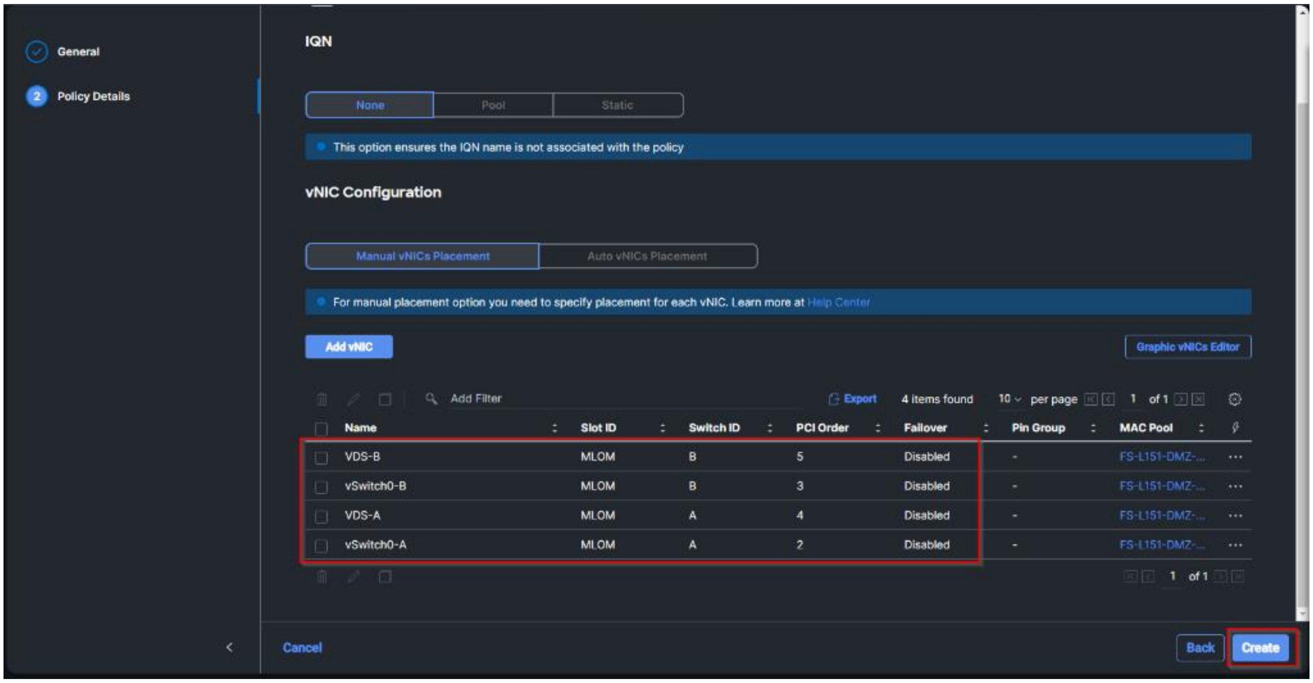

IQN

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
VDS-B	MLOM	B	5	Disabled	-	FS-L151-DMZ-...
vSwitch0-B	MLOM	B	3	Disabled	-	FS-L151-DMZ-...
VDS-A	MLOM	A	4	Disabled	-	FS-L151-DMZ-...
vSwitch0-A	MLOM	A	2	Disabled	-	FS-L151-DMZ-...

Step 83. Click Select Policy next to SAN Connectivity and then click Create New.

Note: A SAN connectivity policy determines the network storage resources and the connections between the server and the storage device on the network. This policy enables customers to configure the vHBAs that the servers use to communicate with the SAN.

Table 10. vHBA for boot from FC SAN

vNIC/vHBA Name	Slot	Switch ID	PCI Order
vHBA-A	MLOM	A	0
vHBA-B	MLOM	B	1

Step 84. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-FC-SAN-Conn-Pol).

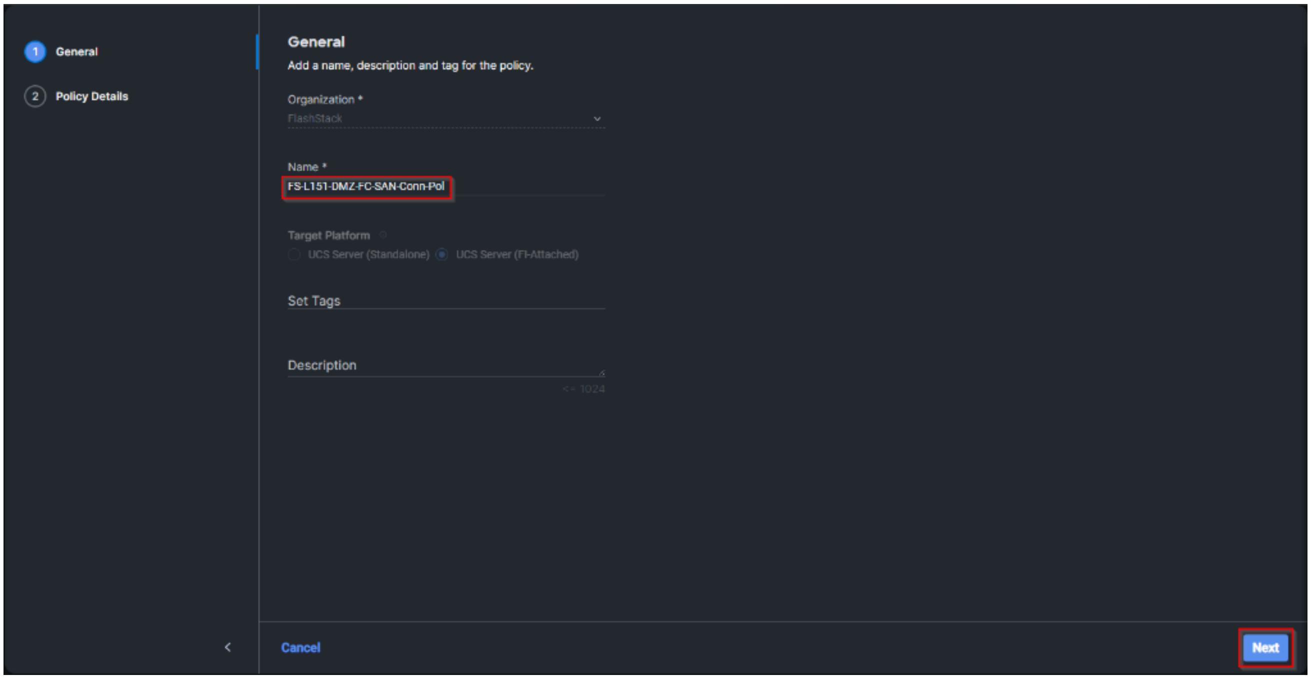

Organization *
FlashStack

Name *
FS-L151-DMZ-FC-SAN-Conn-Pol

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

Step 85. Select Manual vHBAs Placement.

Step 86. Select Pool under WWNN.

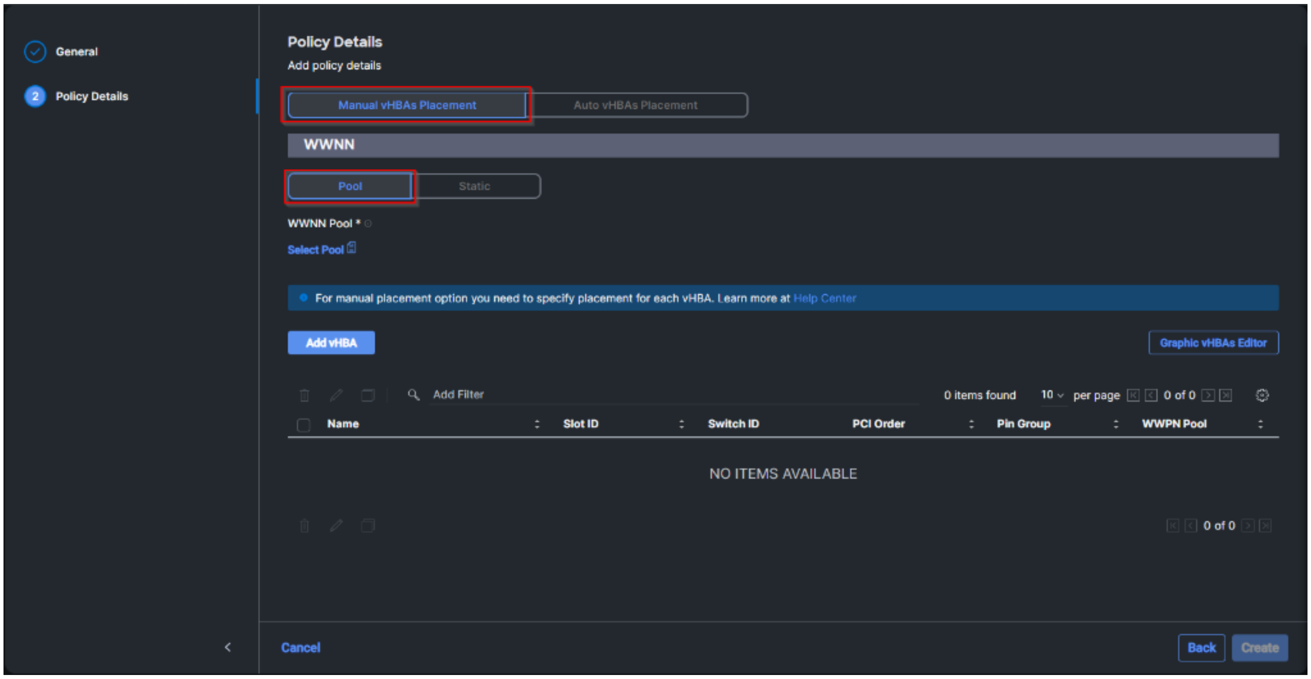

General
2 Policy Details

Policy Details
Add policy details

WWNN Pool *
Select Pool

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

Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN Pool

NO ITEMS AVAILABLE

Note: The WWNN address pools have not been defined yet therefore a new WWNN address pool has to be defined.

Step 87. Click Select Pool under WWNN Pool and then click Create New.

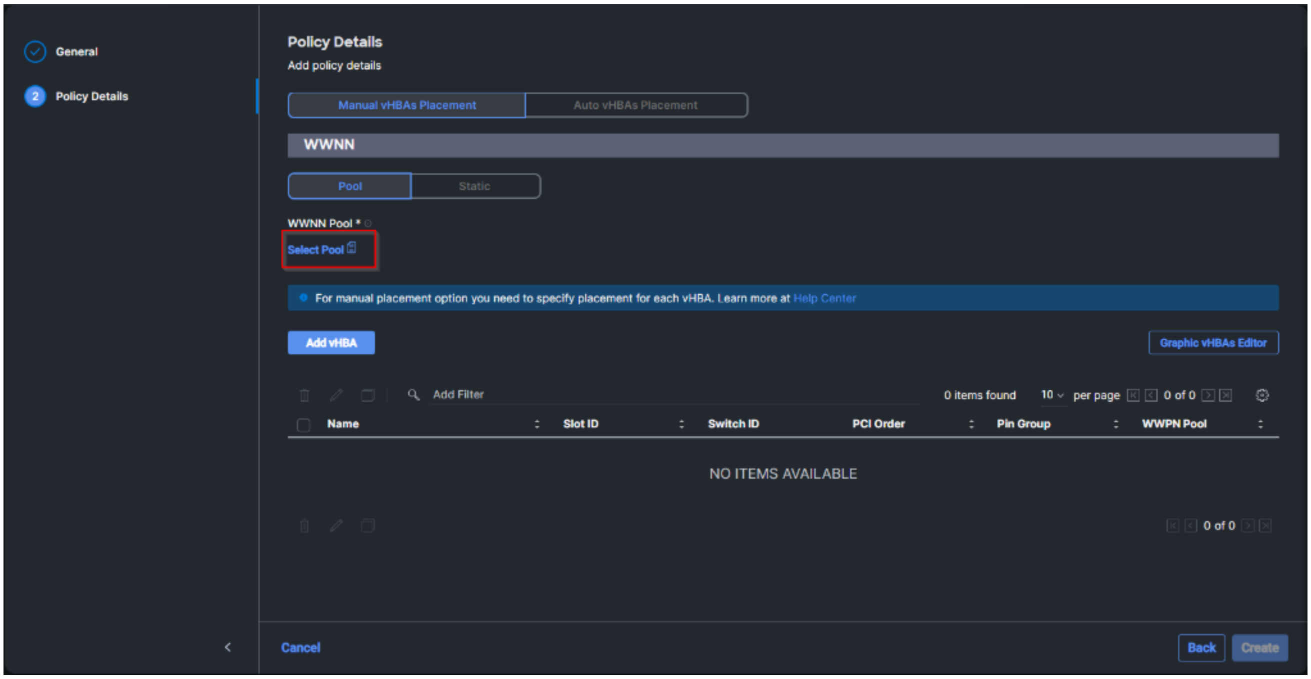

General
2 Policy Details

Policy Details
Add policy details

WWNN Pool *
Select Pool

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

Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN Pool

NO ITEMS AVAILABLE

Step 88. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-WWN-Pool).

Step 89. Click Next.

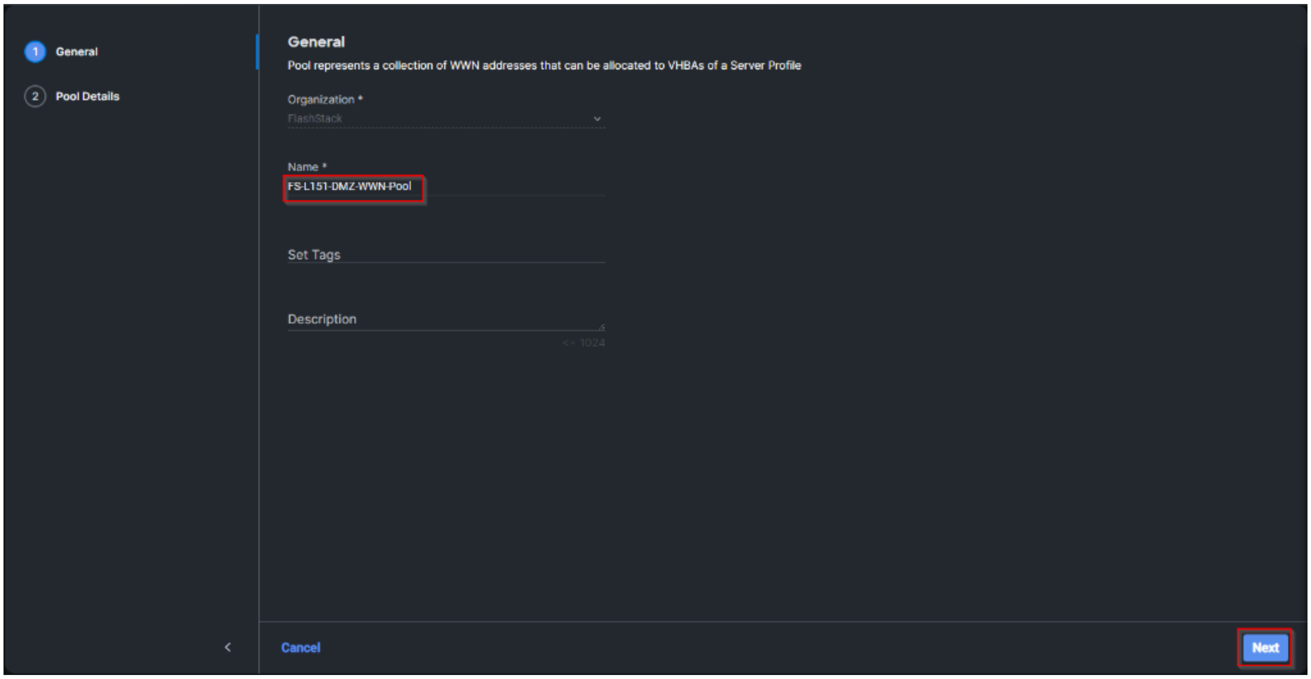

General

Pool represents a collection of WWN addresses that can be allocated to VHBAs of a Server Profile

Organization *: FlashStack

Name *: FS-L151-DMZ-WWN-Pool

Set Tags

Description

Step 90. Provide the starting WWNN block address and the size of the pool. Click Create.

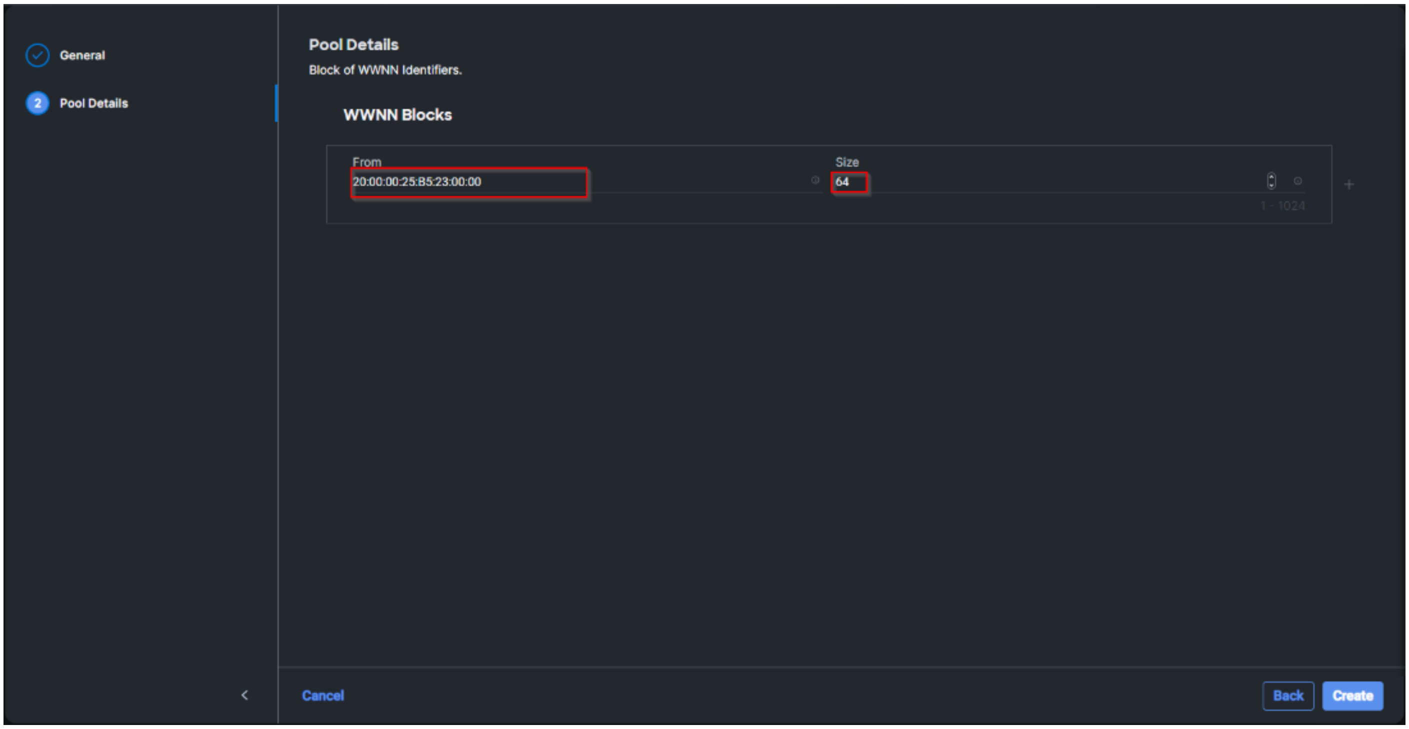

Pool Details

Block of WWNN Identifiers.

WWNN Blocks

From: 20:00:00:25:B5:23:00:00

Size: 64

Note: As a best practice, additional information should always be coded into the WWNN address pool for troubleshooting. For example, in the address 20:00:00:25:B5:23:00:00, 23 is the rack ID.

Step 91. Click Add vHBA.

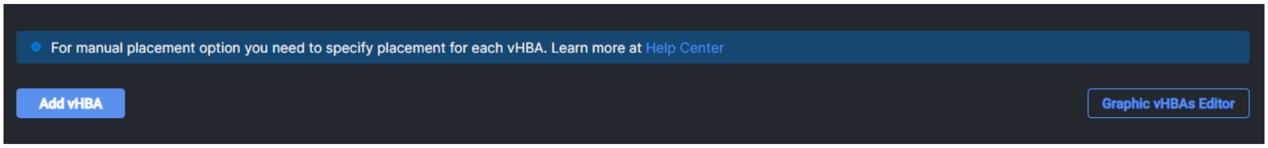

For manual placement option you need to specify placement for each vHBA. Learn more at Help Center

Step 92. Enter vHBA-A for the Name and select fc-initiator from the drop-down list.

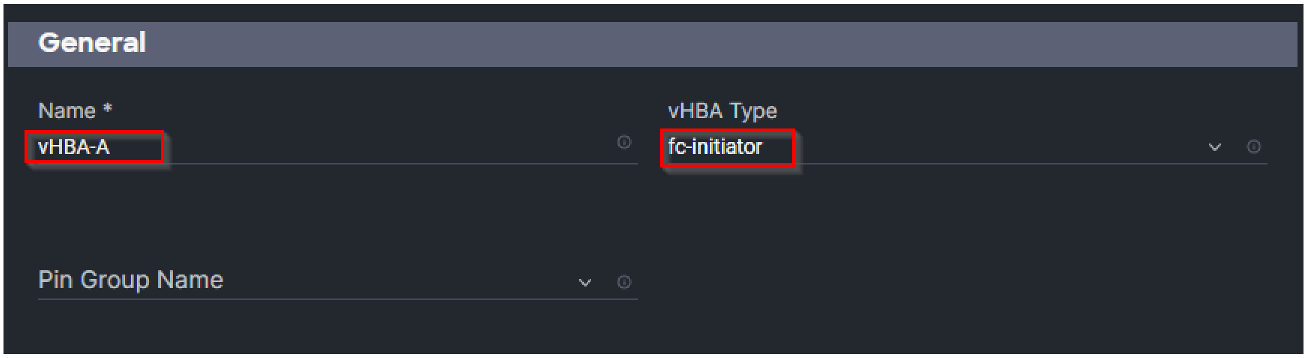

General

Name * vHBA-A

vHBA Type fc-initiator

Note: The WWPN address pool has not been defined yet therefore a WWPN address pool for Fabric A will be defined.

Step 93. Click Select Pool under WWPN Address Pool and then click Create New.

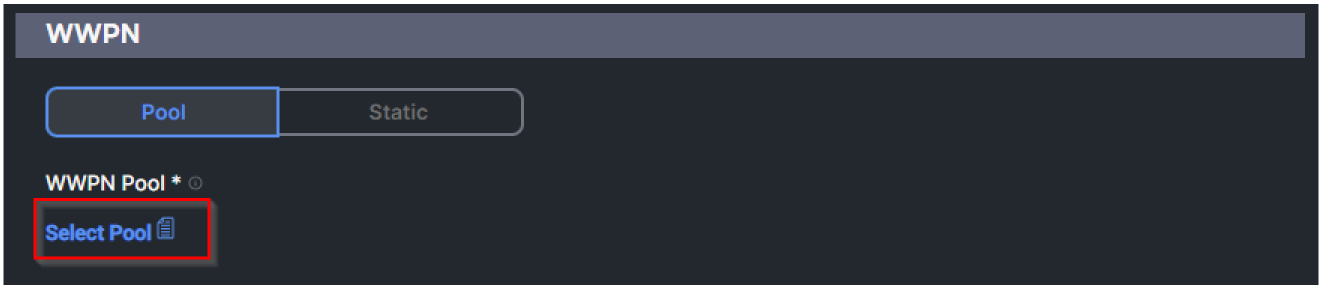

WWPN

Step 94. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-WWPN-Pool-A).

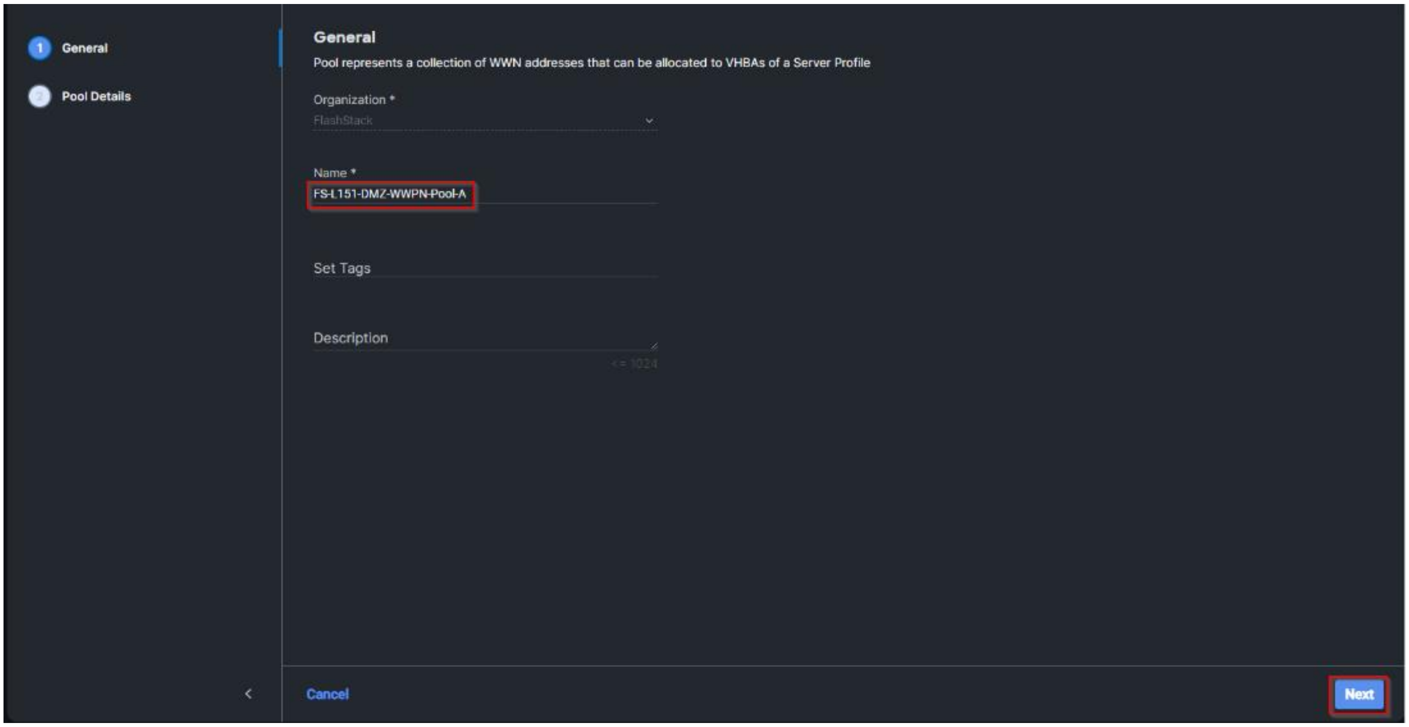

General

Pool represents a collection of WWN addresses that can be allocated to VHBAs of a Server Profile

Organization * FlashStack

Name * FS-L151-DMZ-WWPN-Pool-A

Set Tags

Description

<= 1024

Step 95. Provide the starting WWPN block address for SAN A and the size. Click Create.

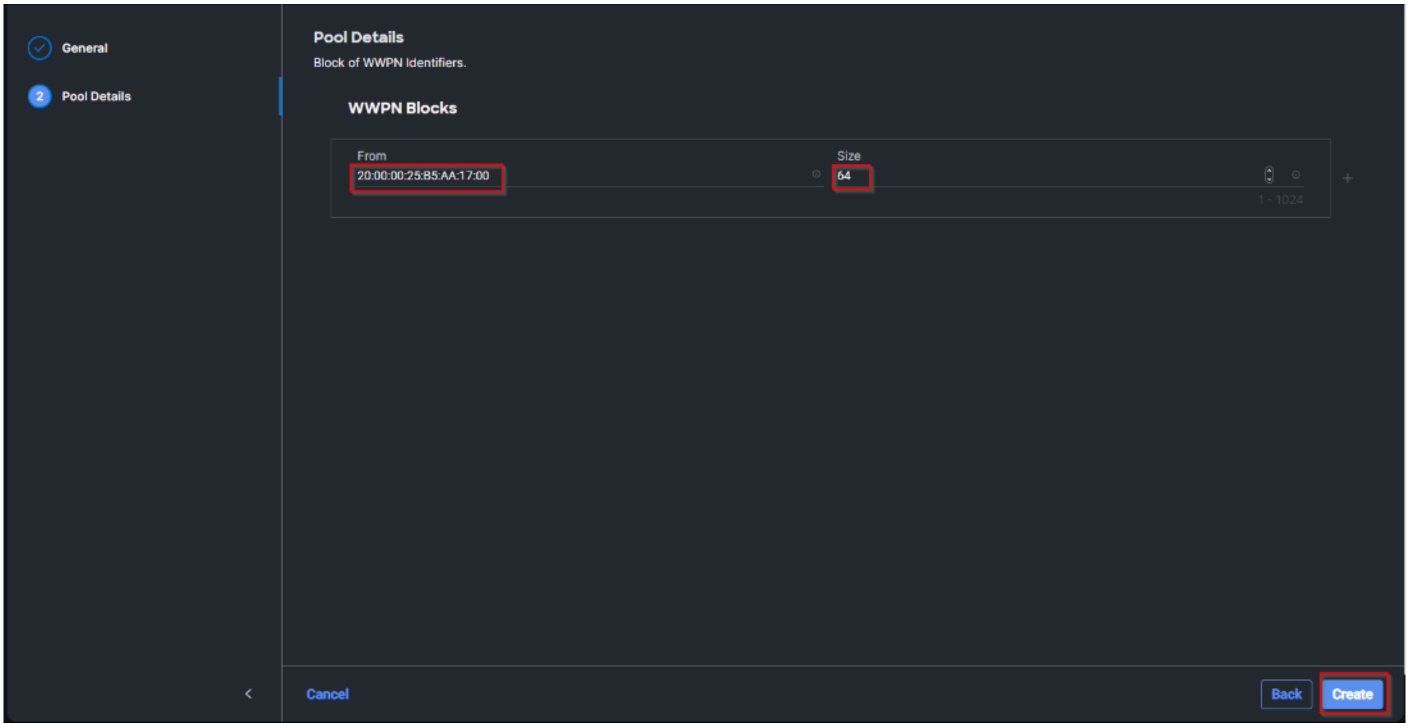

Pool Details

Block of WWPN Identifiers.

WWPN Blocks

From 20:00:00:25:B5:AA:17:00

Size 64

1 - 1024

Step 96. Provide the Switch ID (for example, A) and PCI Order (for example, 0) from Table 9.

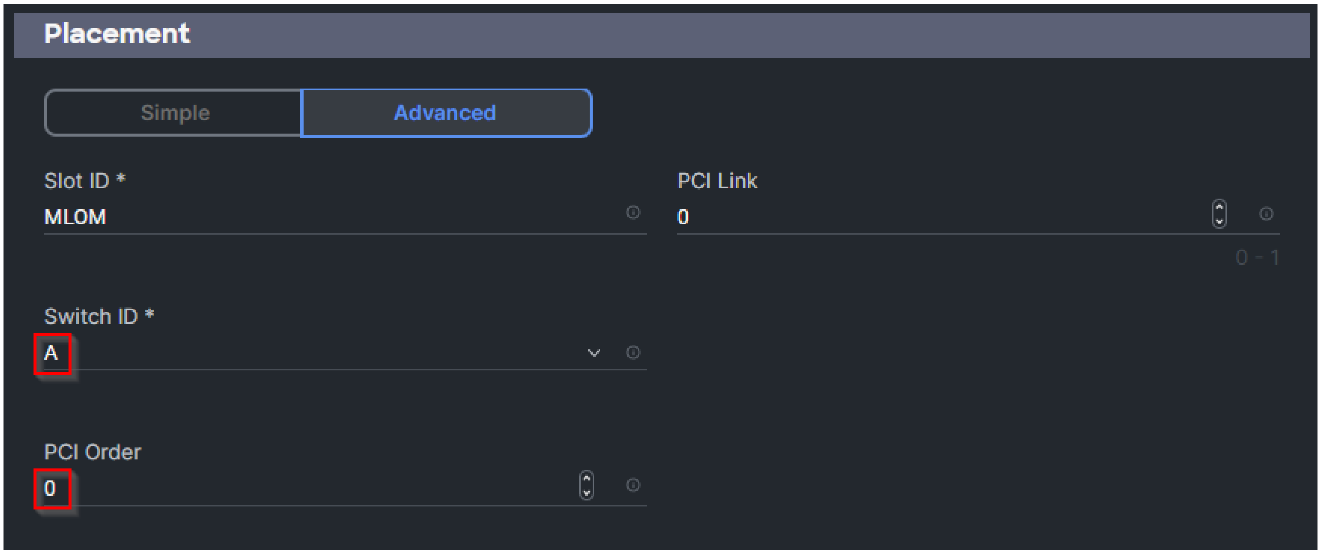

Slot ID *
MLOM

PCI Link
0
0 - 1

Switch ID *
A

PCI Order
0

Step 97. Click Select Policy under Fibre Channel Network and then click Create New.

Note: A Fibre Channel network policy governs the VSAN configuration for the virtual interfaces. In this deployment, VSAN 100 will be used for vHBA-A and VSAN 101 will be used for vHBA-B.

Step 98. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-K4-FCN-A). Click Next.

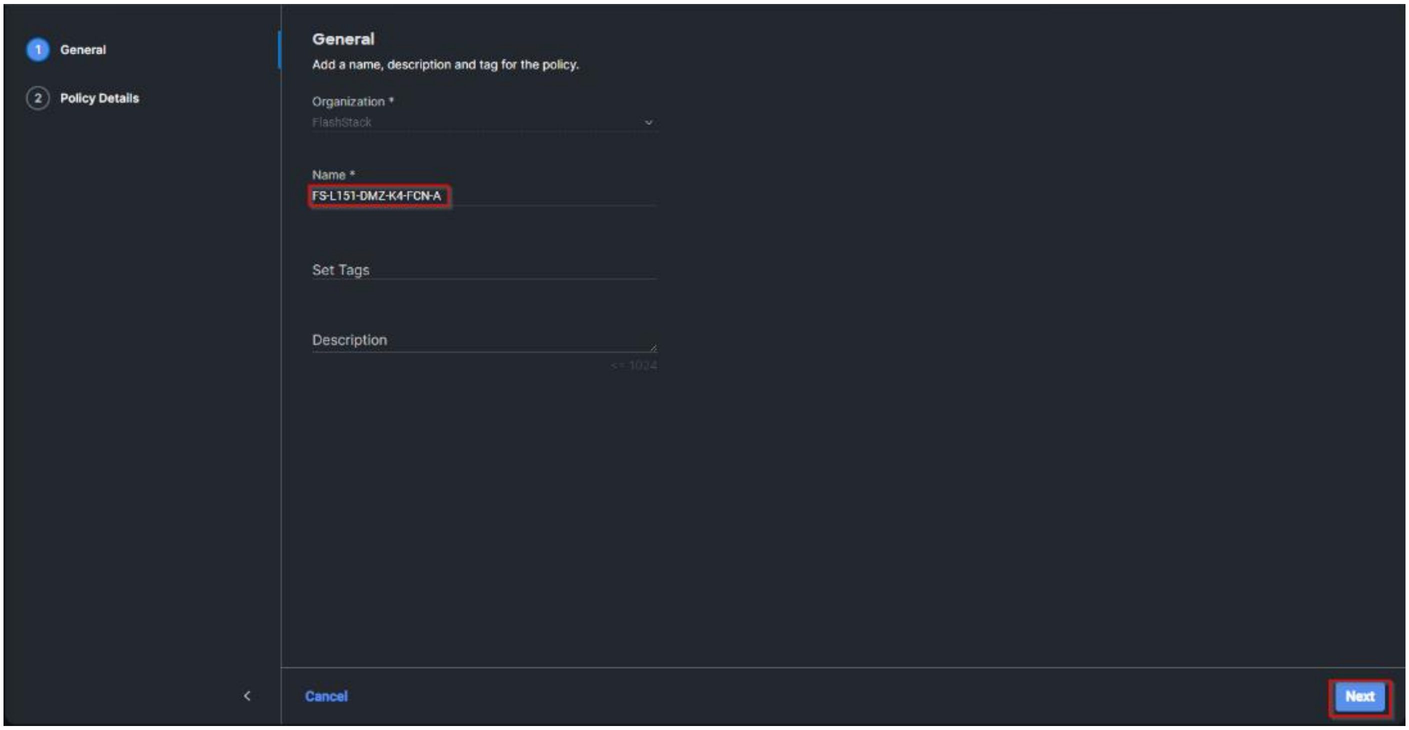

General

Add a name, description and tag for the policy.

Organization *
FlashStack

Name *
FS-L151-DMZ-K4-FCN-A

Set Tags

Description
<= 1024

Step 99. For the scope, select UCS Server (FI-Attached).

Step 100. Under VSAN ID, provide the VSAN information (for example, 100).

Step 101. Click Create.

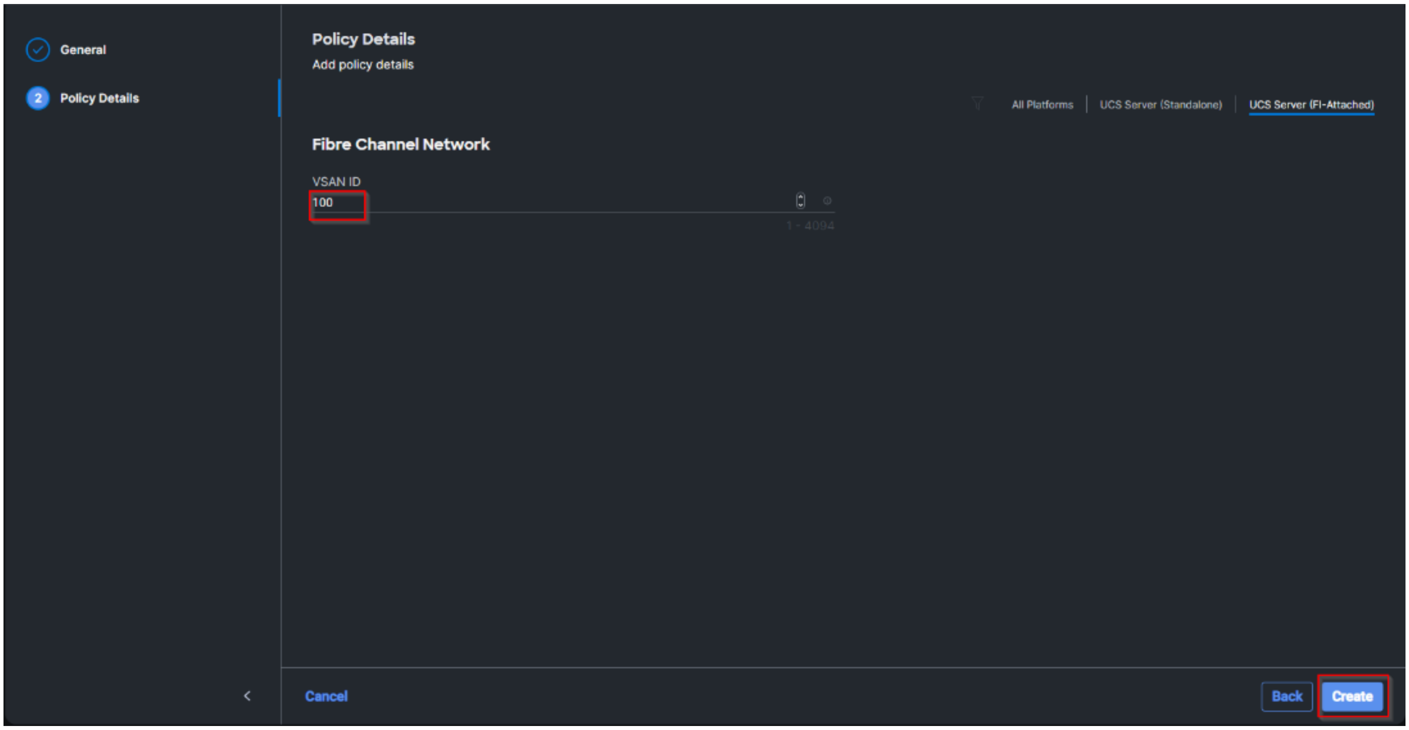

VSAN ID: 100

Step 102. Click Select Policy under Fibre Channel QoS and then click Create New.

Note: The Fibre Channel QoS policy assigns a system class to the outgoing traffic for a vHBA. This system class determines the quality of service for the outgoing traffic. The Fibre Channel QoS policy used in this deployment uses default values and will be shared by all vHBAs.

Step 103. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-FCQOS-Pol). Click Next.

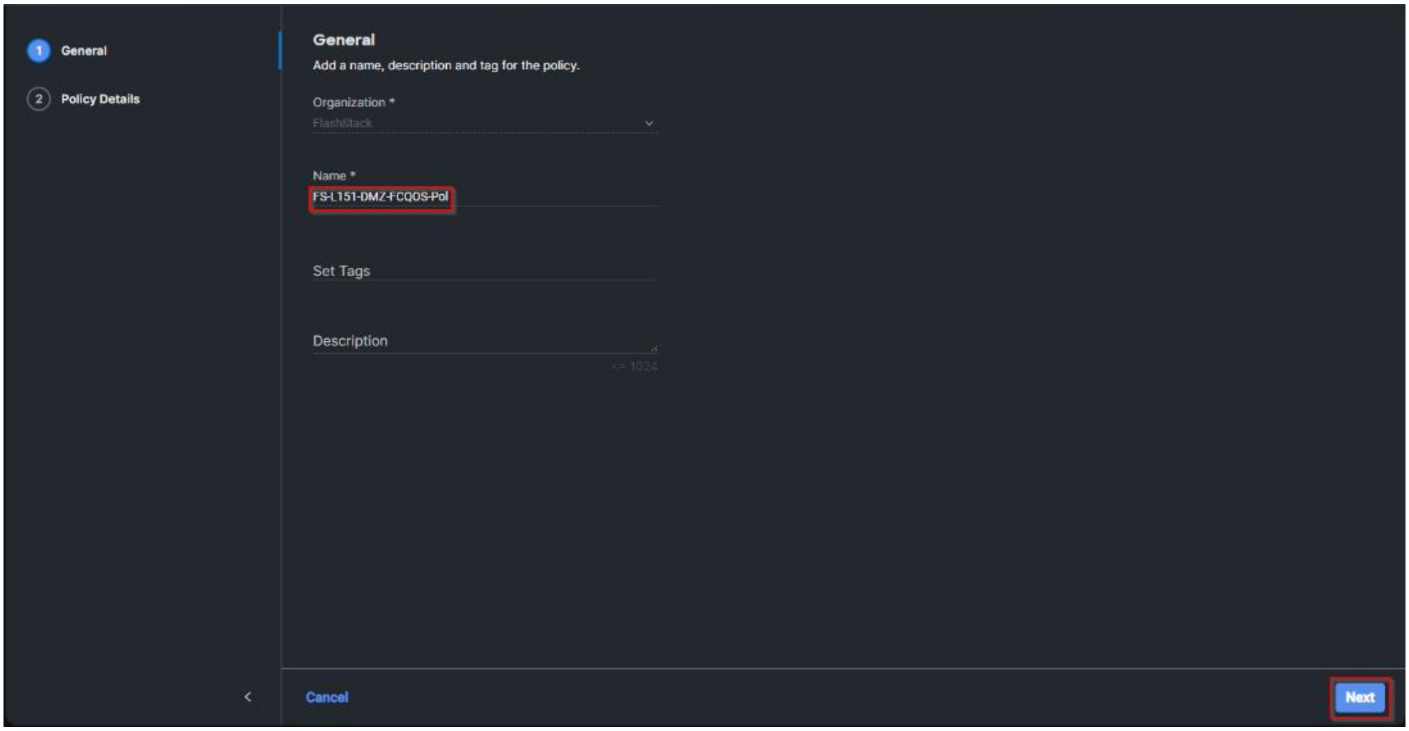

Organization *: FlashStack

Name *: FS-L151-DMZ-FCQOS-Pol

Step 104. For the scope, select UCS Server (FI-Attached).

Note: Do not change the default values on the Policy Details screen.

Step 105. Click Create.

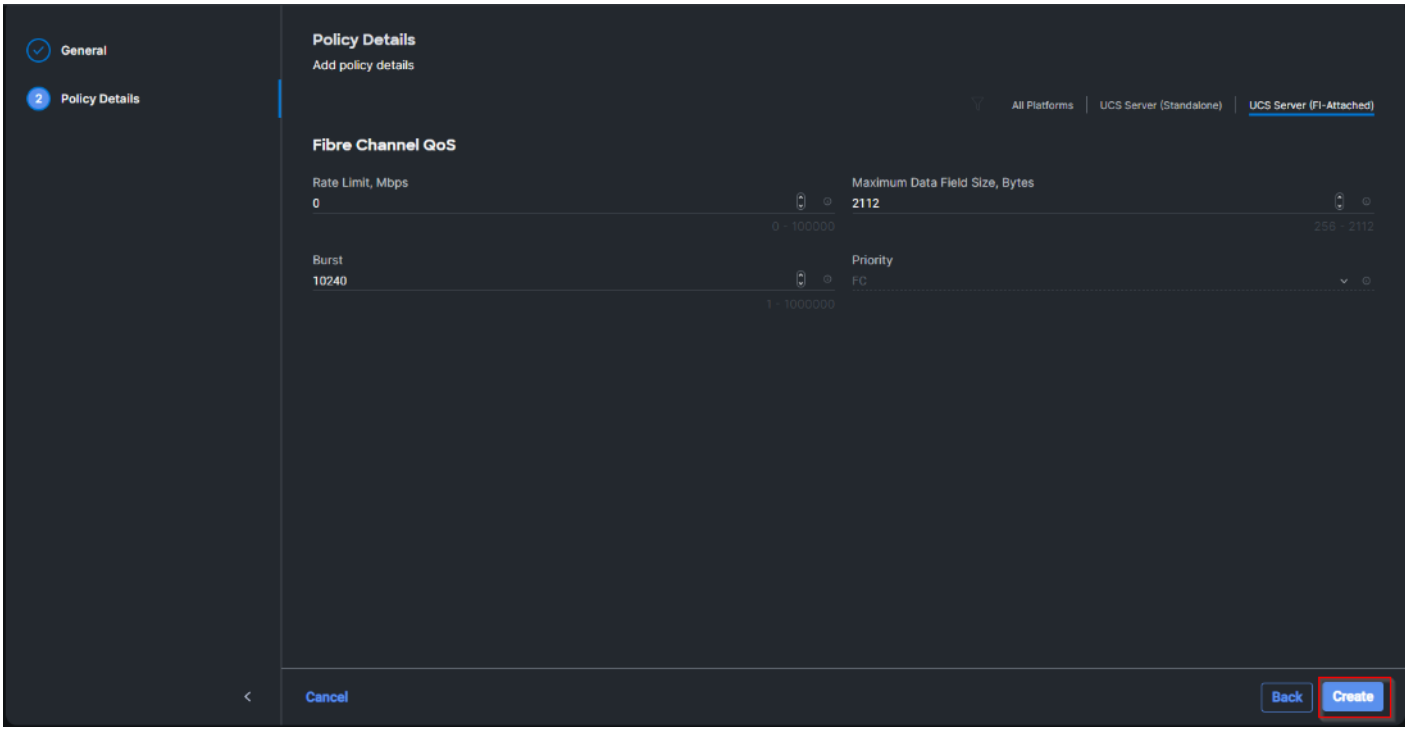

General

2 Policy Details

Policy Details

Add policy details

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

Fibre Channel QoS

Rate Limit, Mbps: 0 (0 - 100000)

Maximum Data Field Size, Bytes: 2112 (256 - 2112)

Burst: 10240 (1 - 1000000)

Priority: FC

Step 106. Click Select Policy under Fibre Channel Adapter and then click Create New.

Note: A Fibre Channel adapter policy governs the host-side behavior of the adapter, including the way that the adapter handles traffic. This validation uses the default values for the adapter policy, and the policy will be shared by all the vHBAs.

Step 107. Verify the correct organization is selected from the drop-down list and provide a name for the policy (for example, FS-L151-DMZ-FC-Adapter-Pol).

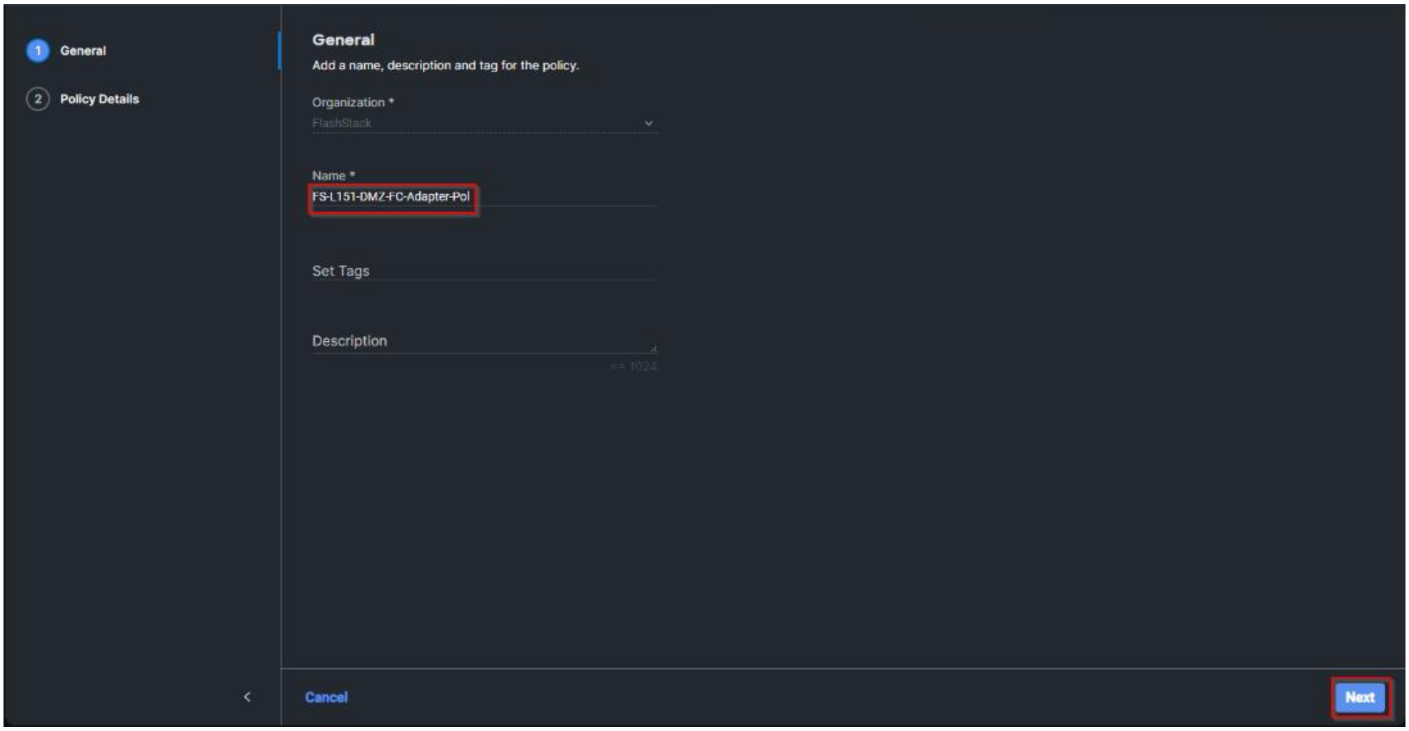

General

Policy Details

General

Add a name, description and tag for the policy.

Organization *
FlashStack

Name *
FS-L151-DMZ-FC-Adapter-Pol

Set Tags

Description
<= 1024

Step 108. For the scope, select UCS Server (FI-Attached).

Note: Do not change the default values on the Policy Details screen.

Step 109. Click Create.

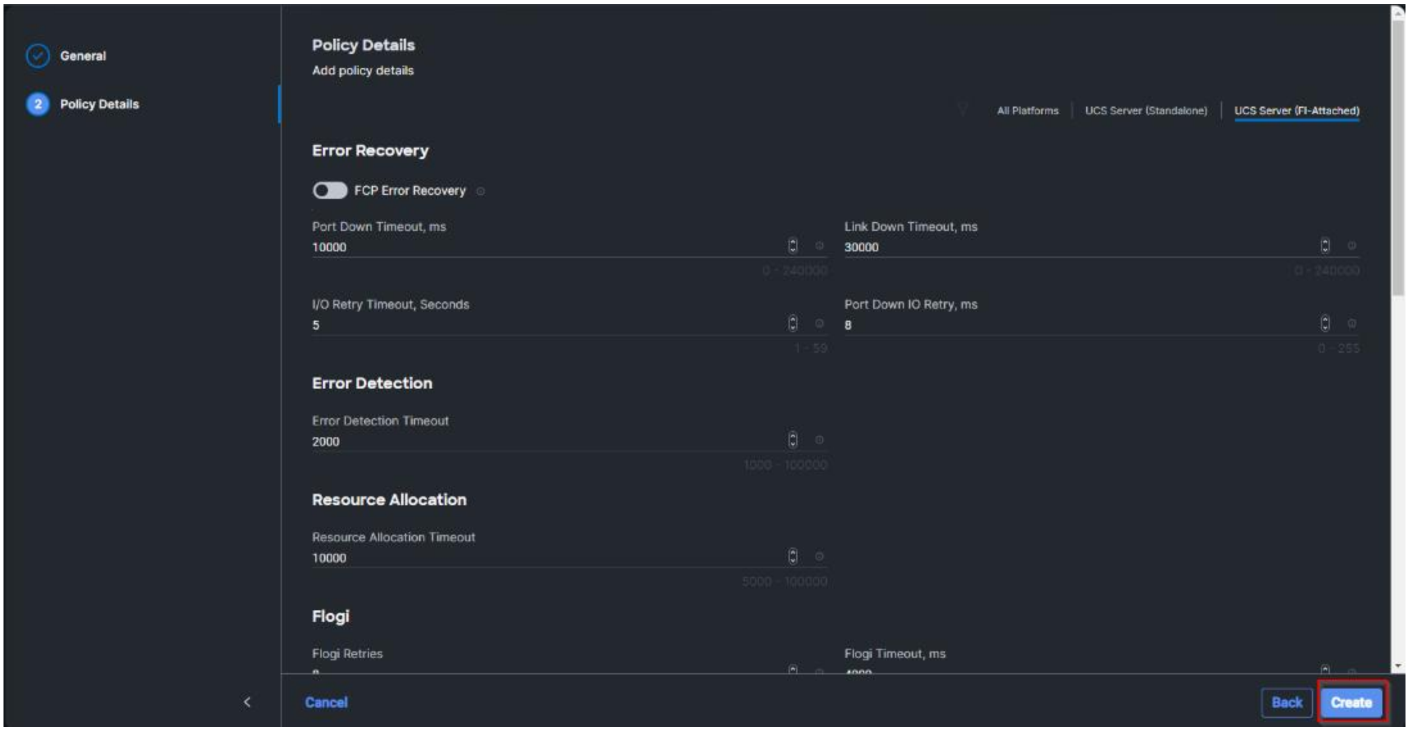

General

Policy Details

Policy Details

Add policy details

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

Error Recovery

FCP Error Recovery

Port Down Timeout, ms
10000
0 - 240000

Link Down Timeout, ms
30000
0 - 240000

I/O Retry Timeout, Seconds
5
1 - 59

Port Down IO Retry, ms
8
0 - 255

Error Detection

Error Detection Timeout
2000
1000 - 100000

Resource Allocation

Resource Allocation Timeout
10000
5000 - 100000

Flogi

Flogi Retries

Flogi Timeout, ms

Step 110. Click Add to create vHBA-A.

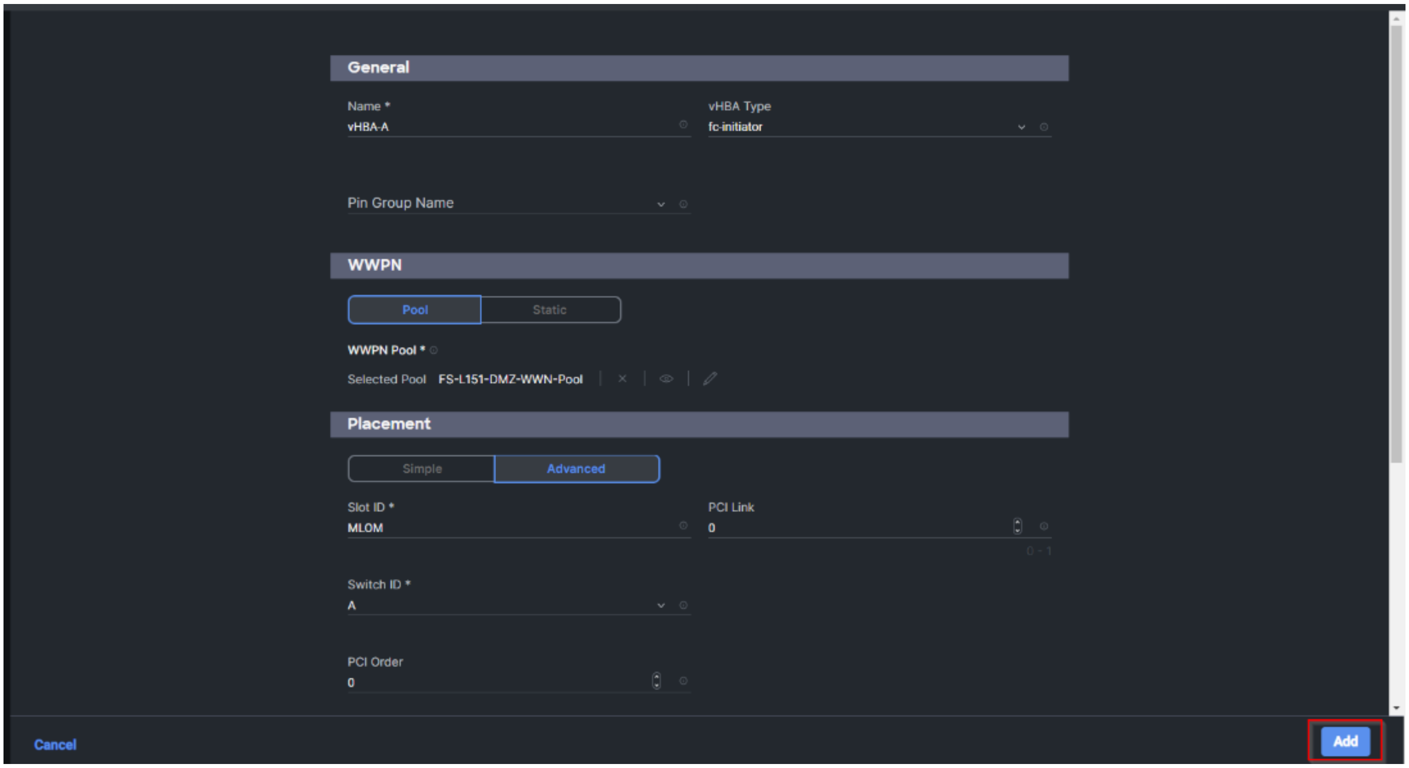

General

Name * vHBA-A

vHBA Type fc-initiator

Pin Group Name

WWPN

WWPN Pool *

Selected Pool FS-L151-DMZ-WWN-Pool

Placement

Slot ID * MLOM

PCI Link 0

0 - 1

Switch ID * A

PCI Order 0

Step 111. Create the vHBA-B using the same steps from above using pools and Fibre Channel Network policy for SAN-B.

Step 112. Verify both vHBAs are added to the SAN connectivity policy.

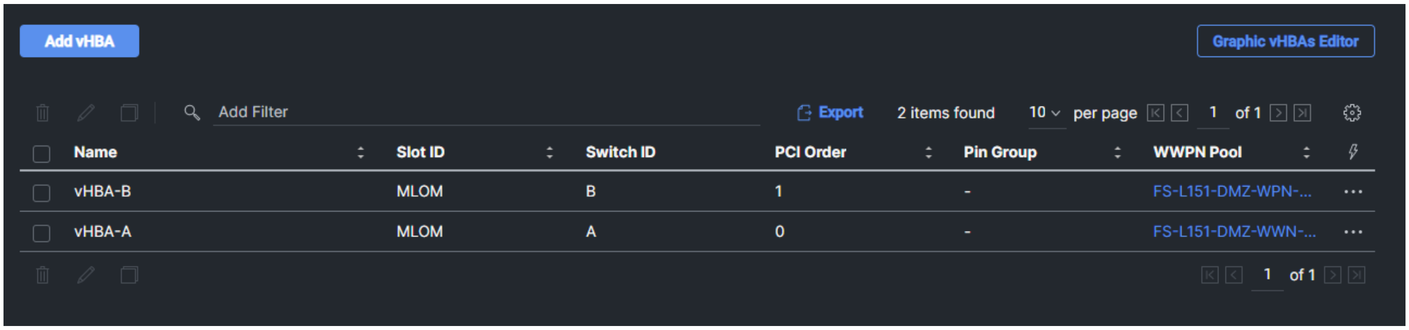

2 items found 10 per page 1 of 1

Name	Slot ID	Switch ID	PCI Order	Pin Group	WWPN Pool
vHBA-B	MLOM	B	1	-	FS-L151-DMZ-WPN-...
vHBA-A	MLOM	A	0	-	FS-L151-DMZ-WWN-...

1 of 1

Step 113. When the LAN connectivity policy and SAN connectivity policy are created and assigned, click Next to move to the Summary screen.

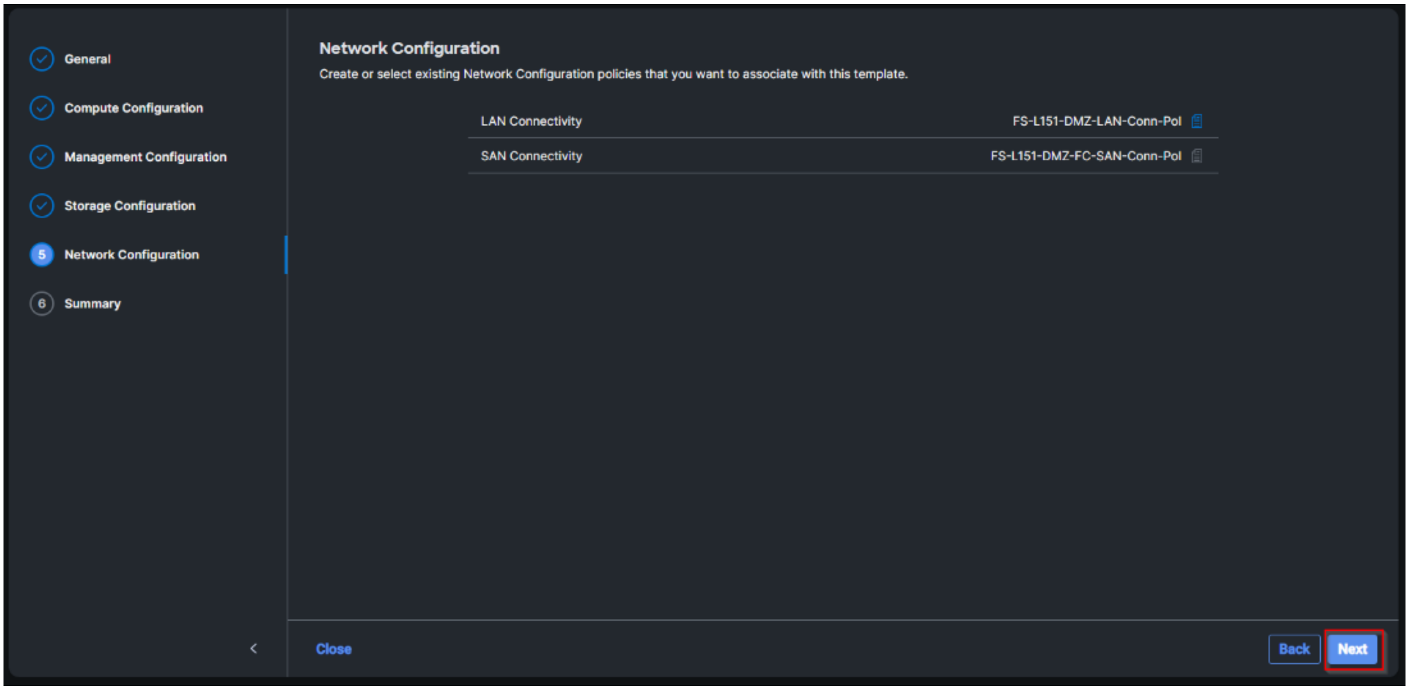

Network Configuration

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

LAN Connectivity	FS-L151-DMZ-LAN-Conn-Pol
SAN Connectivity	FS-L151-DMZ-FC-SAN-Conn-Pol

Step 114. From the Server profile template Summary screen, click Derive Profiles.

Step 115. This action can also be performed later by navigating to Templates, clicking "…" next to the template name and selecting Derive Profiles.

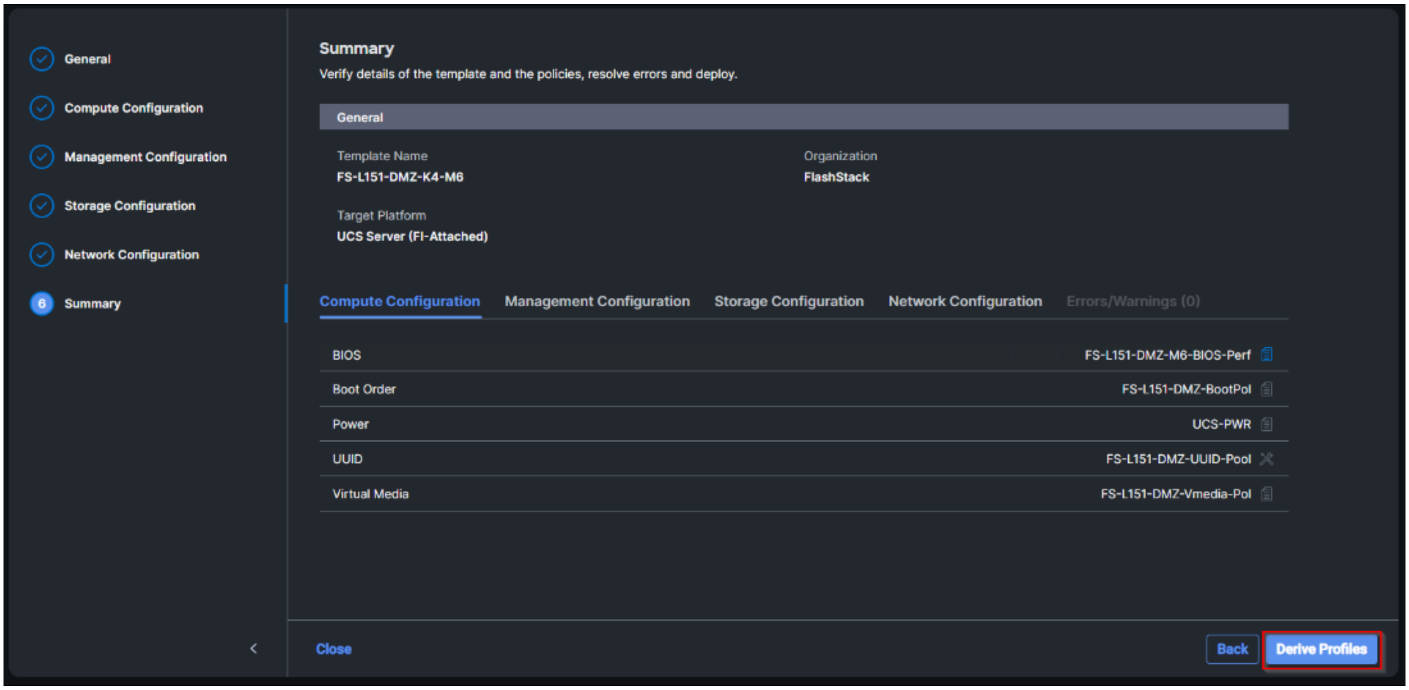

Summary

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

General

Template Name: FS-L151-DMZ-K4-M6

Organization: FlashStack

Target Platform: UCS Server (FI-Attached)

BIOS	FS-L151-DMZ-M6-BIOS-Perf
Boot Order	FS-L151-DMZ-BootPol
Power	UCS-PWR
UUID	FS-L151-DMZ-UUID-Pool
Virtual Media	FS-L151-DMZ-Vmedia-Pol

Step 116. Under the Server Assignment, select Assign Now and select Cisco UCS X210c M6 Nodes. You can select one or more servers depending on the number of profiles to be deployed. Click Next.

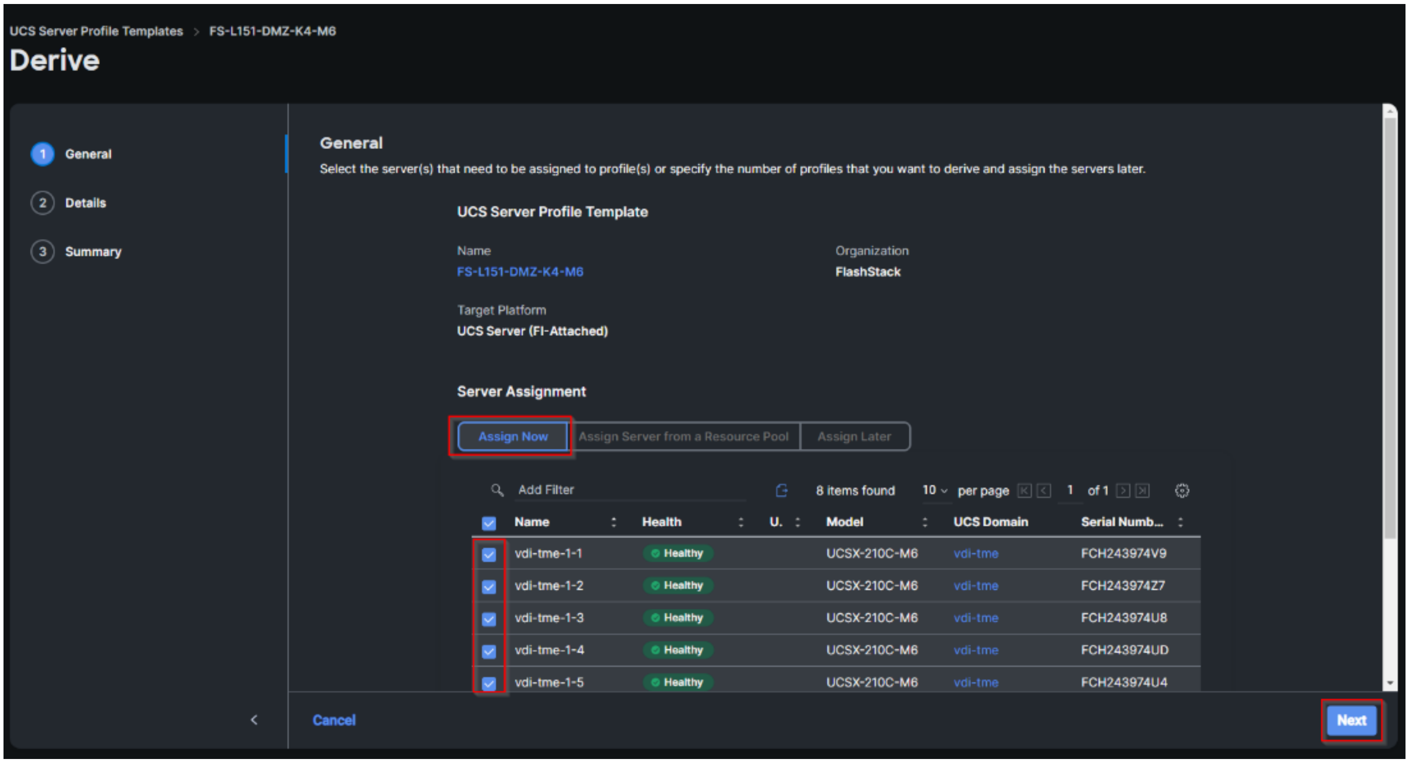

UCS Server Profile Templates > FS-L151-DMZ-K4-M6

Derive

1 General
2 Details
3 Summary

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: FS-L151-DMZ-K4-M6
Organization: FlashStack
Target Platform: UCS Server (FI-Attached)

Server Assignment

Assign Now | Assign Server from a Resource Pool | Assign Later

Add Filter — 8 items found — 10 per page — 1 of 1

☑	Name	Health	U.	Model	UCS Domain	Serial Numb...
☑	vdi-tme-1-1	Healthy		UCSX-210C-M6	vdi-tme	FCH243974V9
☑	vdi-tme-1-2	Healthy		UCSX-210C-M6	vdi-tme	FCH243974Z7
☑	vdi-tme-1-3	Healthy		UCSX-210C-M6	vdi-tme	FCH243974U8
☑	vdi-tme-1-4	Healthy		UCSX-210C-M6	vdi-tme	FCH243974UD
☑	vdi-tme-1-5	Healthy		UCSX-210C-M6	vdi-tme	FCH243974U4

Cancel | Next

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

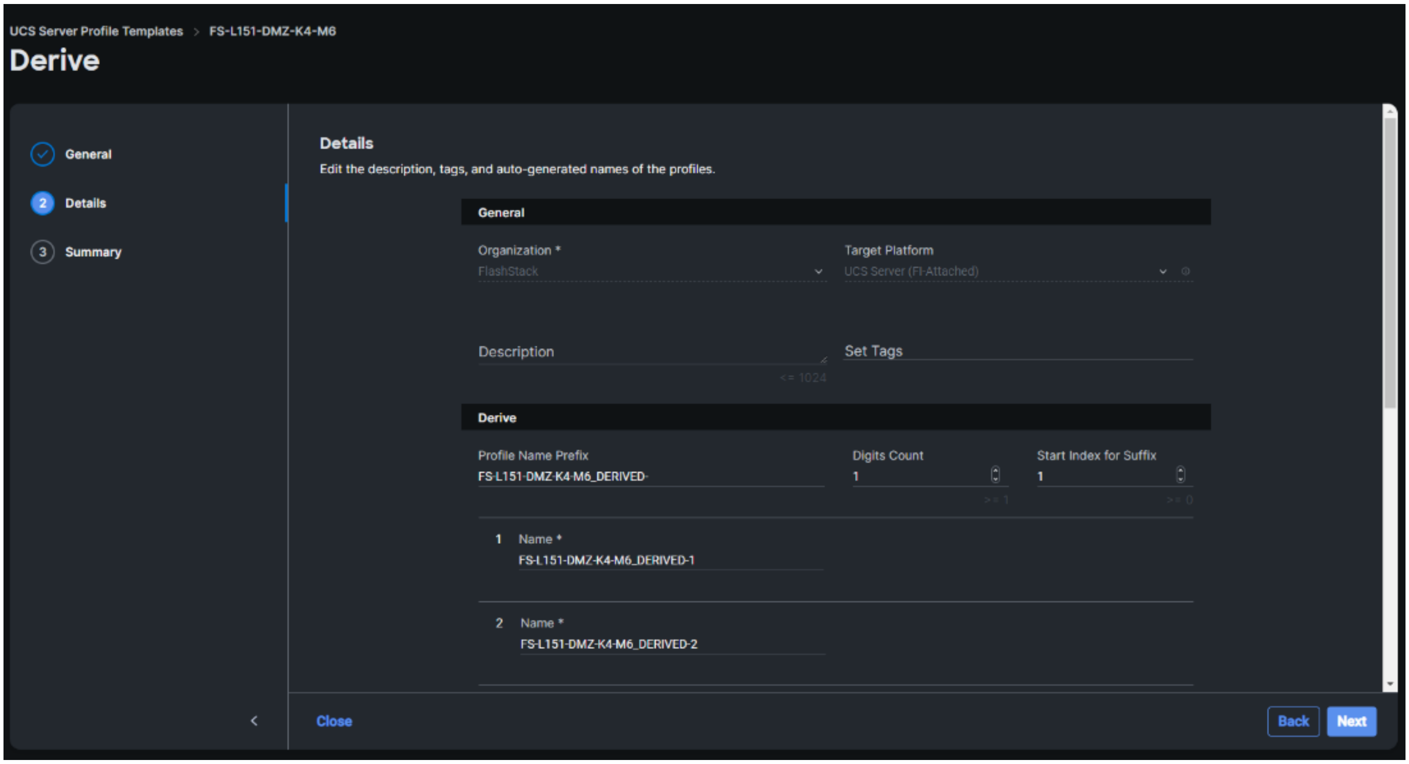

UCS Server Profile Templates > FS-L151-DMZ-K4-M6

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: FS-L151-DMZ-K4-M6_DERIVED-
Digits Count: 1 (>= 1)
Start Index for Suffix: 1 (>= 0)

1 Name *: FS-L151-DMZ-K4-M6_DERIVED-1

2 Name *: FS-L151-DMZ-K4-M6_DERIVED-2

Close | Back | Next

Step 117. Adjust the Prefix and number as needed. Click Next.

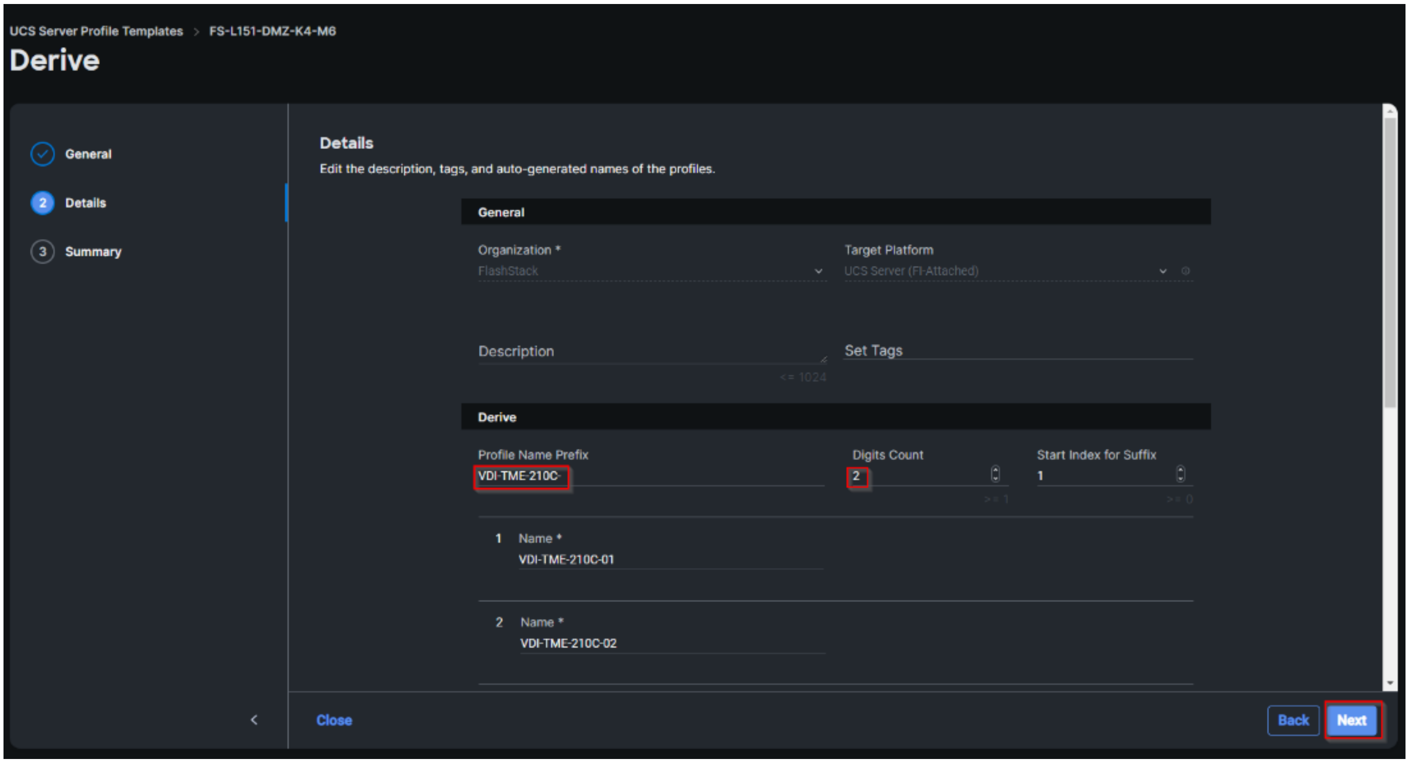

UCS Server Profile Templates > FS-L151-DMZ-K4-M6

Derive

Details

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

General

Organization *: FlashStack

Target Platform: UCS Server (FI-Attached)

Description

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 118. Verify the information and click Derive to create the Server Profiles.

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: We created VLAN 30, 31, 32, 33 and 36.

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

Step 2. Create VLAN 30:

```
config terminal
VLAN 30
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 Oracle Database Servers is listed in Table 11.

Table 11. vPC Summary

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

As listed in Table 11, a single vPC domain with Domain ID 70 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 3 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
	Eth1/53	40Gbe	Cisco Nexus 93180YC-FX B	Eth1/53
	Eth1/54	40Gbe	Cisco Nexus 93180YC-FX B	Eth1/54
	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
	Eth1/53	40Gbe	Cisco Nexus 93180YC-FX A	Eth1/53
	Eth1/54	40Gbe	Cisco Nexus 93180YC-FX A	Eth1/54
	MGMT0	Gbe	Gbe management switch	Any

Cisco UCS Fabric Interconnect 6454 Cabling

The following tables list the 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	FC 1/1	32G FC	Cisco MDS 9132T 32-Gb-A	FC 1/13
	FC 1/2	32G FC	Cisco MDS 9132T 32-Gb-A	FC 1/14
	Eth1/17-24	25Gbe	Cisco UCS 9508 Chassis IFM-A Chassis 1	IO Module 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	FC 1/1	32Gb FC	Cisco MDS 9132T 32-Gb-B	FC 1/13
	FC 1/2	32Gb FC	Cisco MDS 9132T 32-	FC 1/14

Local Device	Local Port	Connection	Remote Device	Remote Port
			Gb-B	
	Eth1/17-24	25Gbe	Cisco UCS 9508 Chassis IFM-B Chassis 1	IO Module 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 1
peer-keepalive destination 10.29.164.234 source 10.29.164.233
exit
interface port-channel 30
description VPC peer-link
switchport mode trunk
switchport trunk allowed VLAN 1,30-36
spanning-tree port type network
vpc peer-link
exit
interface Ethernet1/53
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1,30-36
channel-group 70 mode active
no shutdown
exit
interface Ethernet1/54
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1,30-36
channel-group 70 mode active
no shutdown
exit
copy running-config startup-config
```

Step 3. Log in as admin user into the Cisco Nexus Switch B and repeat steps 1 and 2 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 1,30-36
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 1,30-36
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 1,30-36
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 1,30-36
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 1,30-36
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 1,30-36
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 1,30-36
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 1,30-36
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 25 shows the verification of the vPC status on both Cisco Nexus Switches.

Figure 25. vPC Description for Cisco Nexus Switch A and B

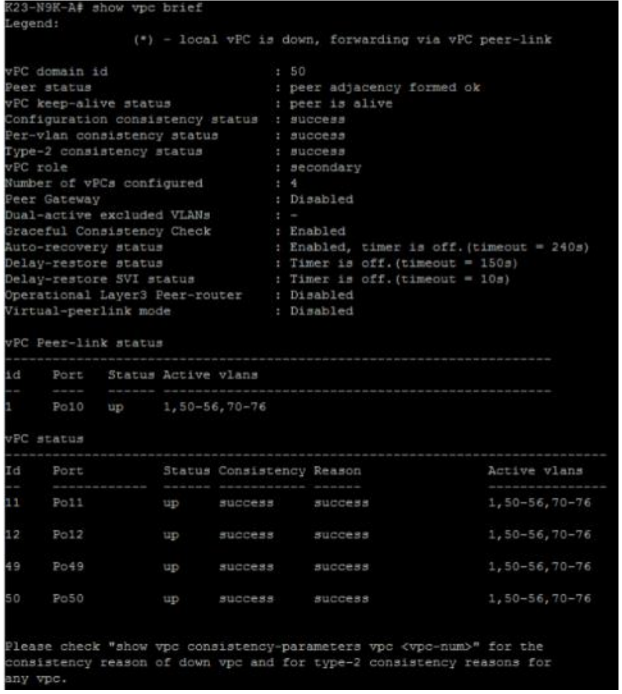

```
K23-N9K-A# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

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

49    Po49          up     success     success               1,50-56,70-76

50    Po50          up     success     success               1,50-56,70-76


Please check "show vpc consistency-parameters vpc <vpc-num>" for the
consistency reason of down vpc and for type-2 consistency reasons for
any vpc.
```

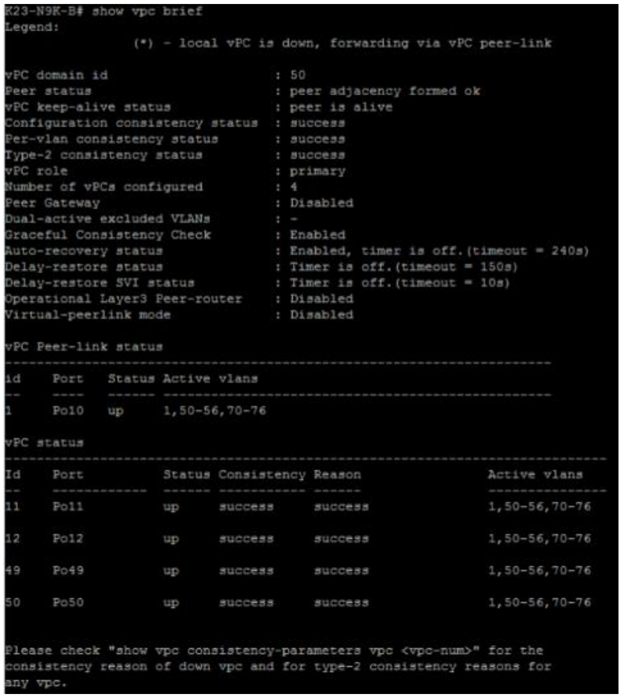

```
K23-N9K-B# show vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

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

49    Po49          up     success     success               1,50-56,70-76

50    Po50          up     success     success               1,50-56,70-76


Please check "show vpc consistency-parameters vpc <vpc-num>" for the
consistency reason of down vpc and for type-2 consistency reasons for
any vpc.
```

Cisco MDS 9132T 32-Gb FC Switch Configuration

Figure 25 illustrates the cable connectivity between the Cisco MDS 9132T 32-Gb switch and the Cisco 6454 Fabric Interconnects and Pure Storage FlashArray//X70 R3 storage.

Note: We used two 32Gb FC connections from each fabric interconnect to each MDS switch and two 32Gb FC connections from each Pure Storage FlashArray//X70 R3 array controller to each MDS switch.

Table 16. Cisco MDS 9132T-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-A	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC0
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC0
	FC1/13	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/1
	FC1/14	32Gb FC	Cisco 6454 Fabric Interconnect-A	FC1/2

Table 17. Cisco MDS 9132T-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T-B	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC2
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC2
	FC1/13	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/1
	FC1/14	32Gb FC	Cisco 6454 Fabric Interconnect-B	FC1/2

Pure Storage FlashArray//X70 R3 to MDS SAN Fabric Connectivity

Pure Storage FlashArray//X70 R3 to MDS A and B Switches using VSAN 100 for Fabric A and VSAN 101 Configured for Fabric B

In this solution, two ports (ports FC1/9 and FC1/10) of MDS Switch A and two ports (ports FC1/9 and FC1/10) of MDS Switch B are connected to Pure Storage System as listed in Table 18. All ports connected to the Pure Storage Array carry 32 Gb/s FC Traffic.

Table 18. MDS 9132T 32-Gb switch Port Connection to Pure Storage System

Local Device	Local Port	Connection	Remote Device	Remote Port
MDS Switch A	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC0
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC0
MDS Switch B	FC1/9	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 0	CT0.FC2
	FC1/10	32Gb FC	Pure Storage FlashArray//X70 R3 Controller 1	CT1.FC2

Procedure 1. Configure Feature for MDS Switch A and MDS Switch B

Follow these steps on both MDS switches.

Step 1. Log in as admin user into MDS Switch A:

```
config terminal
feature npiv
feature telnet
switchname FlashStack-MDS-A
copy running-config startup-config
```

Step 2. Log in as admin user into MDS Switch B. Repeat step 1 on MDS Switch B.

Procedure 2. Configure VSANs for MDS Switch A and MDS Switch B

Step 1. Log in as admin user into MDS Switch A. Create VSAN 100 for Storage Traffic:

```
config terminal
VSAN database
vsan 500
exit
zone smart-zoning enable vsan 500
vsan database
vsan 500 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 500
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

Step 2. Log in as admin user into MDS Switch B. Create VSAN 101 for Storage Traffic:

```
config terminal
VSAN database
vsan 501
exit
zone smart-zoning enable vsan 501
vsan database
vsan 501 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 501
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

Procedure 3. Create and Configure Fiber Channel Zoning

This procedure sets up the Fibre Channel connections between the Cisco MDS 9132T 32-Gb switches, the Cisco UCS Fabric Interconnects, and the Pure Storage FlashArray systems.

Note: Before you configure the zoning details, decide how many paths are needed for each LUN and extract the WWPN numbers for each of the HBAs from each server. We used 2 HBAs for each Server. One of the HBAs (HBA-A) is connected to MDS Switch-A and other HBAs (HBA-B) is connected to MDS Switch-B.

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

Step 2. From the Service Selector drop-down list, choose Infrastructure Service.

Step 3. Navigate to Configure > Pools. Filter WWPN type pools.

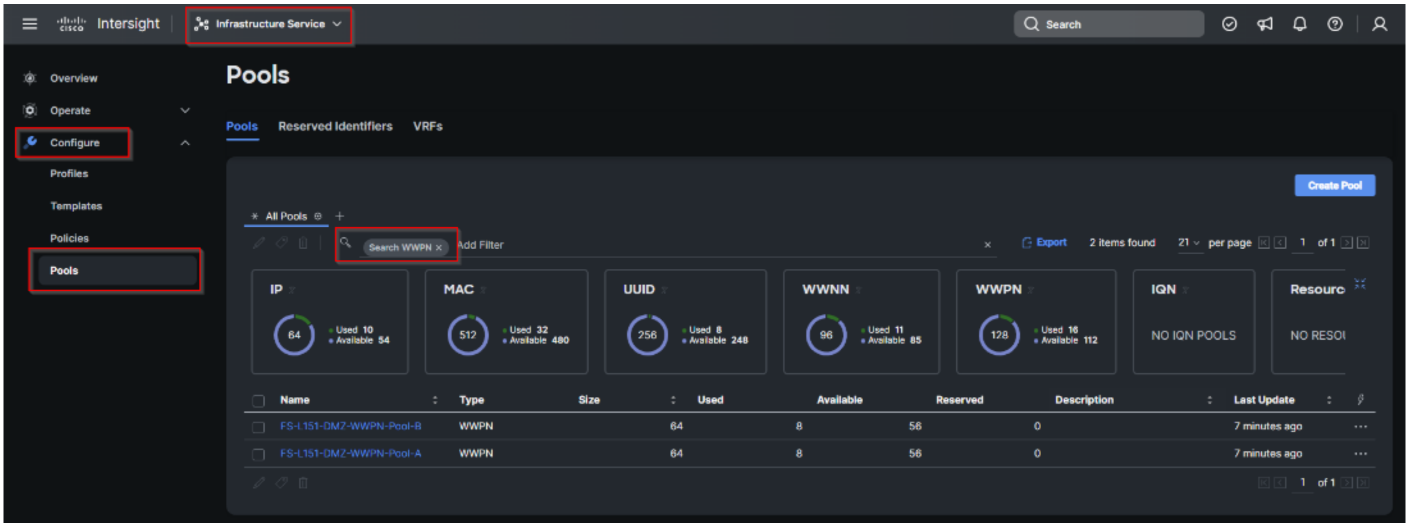

Name	Type	Size	Used	Available	Reserved	Description	Last Update
FS-L151-DMZ-WWPN-Pool-B	WWPN	64	8	56	0		7 minutes ago
FS-L151-DMZ-WWPN-Pool-A	WWPN	64	8	56	0		7 minutes ago

Step 4. Select Usage tab and collect the WWPNs and profiles to which they are assigned.

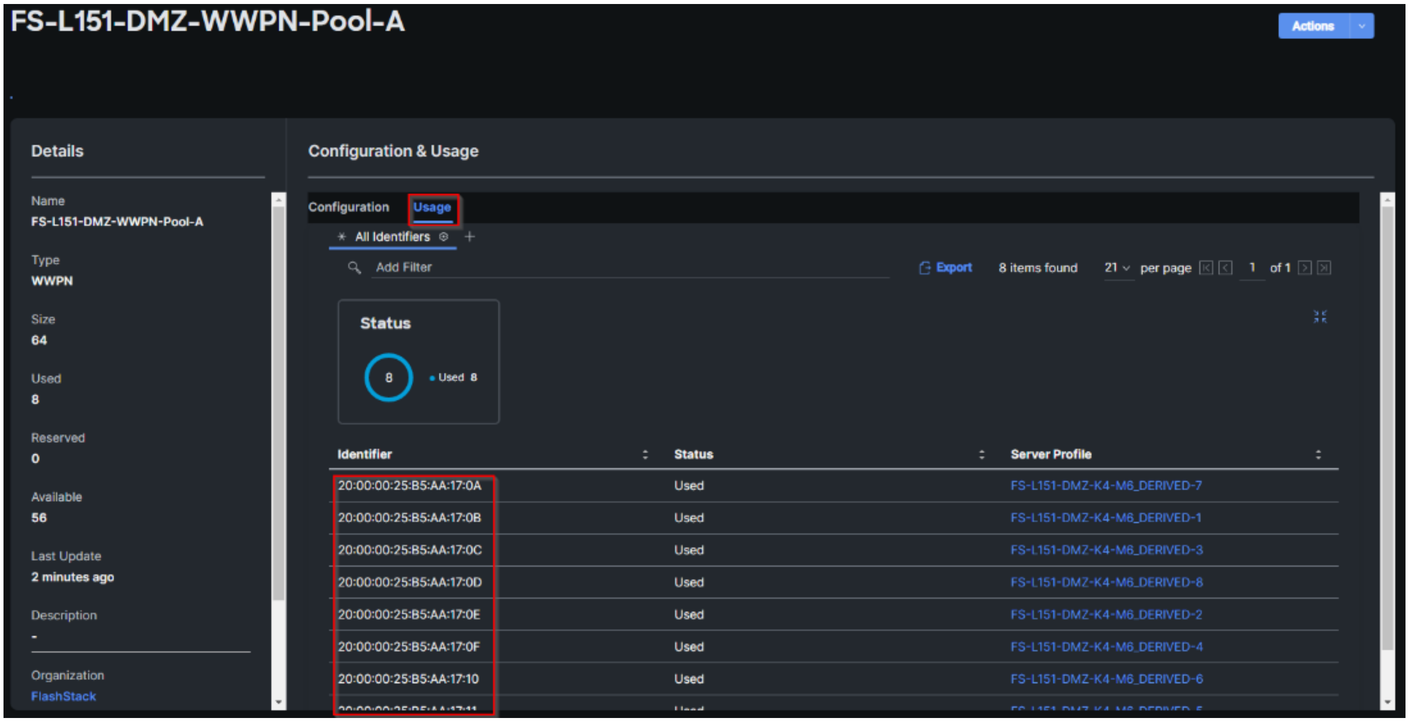

Name
FS-L151-DMZ-WWPN-Pool-A

Type
WWPN

Size
64

Used
8

Reserved
0

Available
56

Last Update
2 minutes ago

Description
-

Organization
FlashStack

Identifier	Status	Server Profile
20:00:00:25:B5:AA:17:0A	Used	FS-L151-DMZ-K4-M6_DERIVED-7
20:00:00:25:B5:AA:17:0B	Used	FS-L151-DMZ-K4-M6_DERIVED-1
20:00:00:25:B5:AA:17:0C	Used	FS-L151-DMZ-K4-M6_DERIVED-3
20:00:00:25:B5:AA:17:0D	Used	FS-L151-DMZ-K4-M6_DERIVED-8
20:00:00:25:B5:AA:17:0E	Used	FS-L151-DMZ-K4-M6_DERIVED-2
20:00:00:25:B5:AA:17:0F	Used	FS-L151-DMZ-K4-M6_DERIVED-4
20:00:00:25:B5:AA:17:10	Used	FS-L151-DMZ-K4-M6_DERIVED-6
20:00:00:25:B5:AA:17:11	Used	FS-L151-DMZ-K4-M6_DERIVED-5

Step 5. Connect to the Pure Storage System Health and go to the Connections tab and extract the WWPN of FC Ports connected to the Cisco MDS Switches from Array Ports section.

Note: We connected 4 FC ports from Pure Storage System to Cisco MDS Switches. FC ports CT0.FC0, CT1.FC0 are connected to MDS Switch-A and similarly FC ports CT1.FC2, CT0.FC2 are connected to MDS Switch-B.

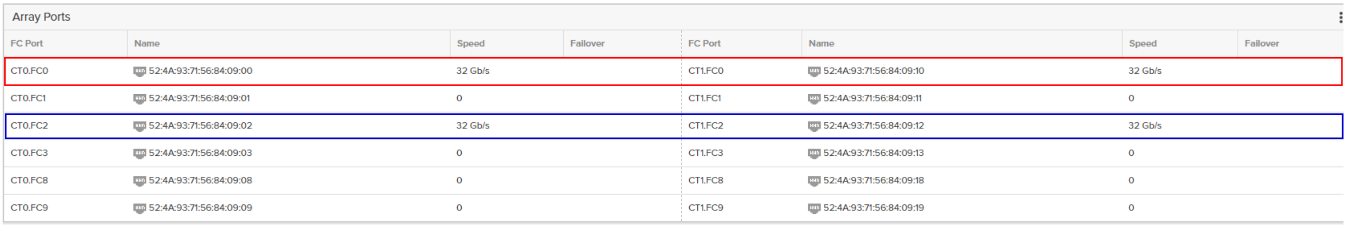

Array Ports

FC Port	Name	Speed	Failover	FC Port	Name	Speed	Failover
CT0.FC0	52:4A:93:71:56:84:09:00	32 Gb/s		CT1.FC0	52:4A:93:71:56:84:09:10	32 Gb/s	
CT0.FC1	52:4A:93:71:56:84:09:01	0		CT1.FC1	52:4A:93:71:56:84:09:11	0	
CT0.FC2	52:4A:93:71:56:84:09:02	32 Gb/s		CT1.FC2	52:4A:93:71:56:84:09:12	32 Gb/s	
CT0.FC3	52:4A:93:71:56:84:09:03	0		CT1.FC3	52:4A:93:71:56:84:09:13	0	
CT0.FC8	52:4A:93:71:56:84:09:08	0		CT1.FC8	52:4A:93:71:56:84:09:18	0	
CT0.FC9	52:4A:93:71:56:84:09:09	0		CT1.FC9	52:4A:93:71:56:84:09:19	0	

Procedure 4. Create Device Aliases for Fiber Channel Zoning for SAN Boot Paths and Datapaths on Cisco MDS Switch A

Step 1. Log in as admin user and run the following commands from the global configuration mode:

```
configure terminal
device-alias mode enhanced
device-alias database
device-alias name VDI-Host01-HBA0 pwwn 20:00:00:25:B5:AA:17:00
device-alias name X70R3-CT0-FC0 pwwn 52:4A:93:71:56:84:09:00
device-alias name X70R3-CT1-FC0 pwwn 52:4A:93:71:56:84:09:10
exit
device-alias commit
```

Procedure 5. Create Device Aliases for Fiber Channel Zoning for SAN Boot Paths and Datapaths on Cisco MDS Switch B

Step 1. Log in as admin user and run the following commands from the global configuration mode:

```
configure terminal
device-alias mode enhanced
device-alias database
device-alias name Host-FCP-1-HBA1 pwwn 20:00:00:25:b5:bb:17:03
device-alias name X70R3-CT0-FC2 pwwn 52:4A:93:71:56:84:09:02
device-alias name X70R3-CT1-FC2 pwwn 52:4A:93:71:56:84:09:12
exit
device-alias commit
```

Procedure 6. Create Fiber Channel Zoning for Cisco MDS Switch A for each Service Profile

Step 1. Log in as admin user and create the zone:

```
configure terminal
zone name FlashStack-Fabric-A vsan 500
    member device-alias X70R3-CT0-FC0 target
    member device-alias X70R3-CT1-FC0 target
    member device-alias Host-FCP-1-HBA0 init
```

Step 2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the created zones as members:

```
configure terminal
zoneset name VDI-Fabric-A vsan 500
    member FlashStack-Fabric-A
```

Step 3. Activate the zone set by running following commands:

```
zoneset activate name VDI-Fabric-A vsan 500
exit
copy running-config startup-config
```

Procedure 7. Create Fiber Channel Zoning for Cisco MDS Switch B for each Service Profile

Step 1. Log in as admin user and create the zone as shown below:

```
configure terminal zone name FlashStack-Fabric-B vsan 501
```

```
    member device-alias X70R3-CT0-FC2 target
    member device-alias X70R3-CT1-FC2 target
    member device-alias Host-FCP-1-HBA1 init
```

Step 2. After the zone for the Cisco UCS service profile has been created, create the zone set and add the necessary members:

```
zoneset name VDI-Fabric-B vsan 501
    member FlashStack-Fabric-B
```

Step 3. Activate the zone set by running following commands:

```
zoneset activate name VDI-Fabric-B vsan 501
exit
copy running-config startup-config
```

Configure Pure Storage FlashArray//X70 R3

The design goal of the reference architecture is to best represent a real-world environment as closely as possible. The approach included the features of Cisco UCS to rapidly deploy stateless servers and use Pure Storage FlashArray's boot LUNs to provision the ESXi on top of Cisco UCS. Zoning was performed on the Cisco MDS 9132T 32-Gb switches to enable the initiators discover the targets during boot process.

A Service Profile was created within Cisco UCS Manager to deploy the thirty-two servers quickly with a standard configuration. SAN boot volumes for these servers were hosted on the same Pure Storage FlashArray//X70 R3. Once the stateless servers were provisioned, following process was performed to enable rapid deployment of thirty-two blade servers.

Each Blade Server has dedicated single LUN to install operating system and all the thirty-two Blade Servers configured to boot from SAN. For this solution, we installed VMware vSphere ESXi 8.0 Cisco Custom ISO on this LUNs to create solution.

Using logical servers that are disassociated from the physical hardware removes many limiting constraints around how servers are provisioned. Cisco UCS Service Profiles contain values for a server's property settings, including virtual network interface cards (vNICs), MAC addresses, boot policies, firmware policies, fabric connectivity, external management, and HA information. The service profiles represent all the attributes of a logical server in Cisco UCS model. By abstracting these settings from the physical server into a Cisco Service Profile, the Service Profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, Service Profiles can, at any time, be migrated from one physical server to another. Furthermore, Cisco is the only hardware provider to offer a truly unified management platform, with Cisco UCS Service Profiles and hardware abstraction capabilities extending to both blade and rack servers.

In addition to the service profiles, the use of Pure Storage's FlashArray's with SAN boot policy provides the following benefits:

- Scalability - Rapid deployment of new servers to the environment in a very few steps.
- Manageability - Enables seamless hardware maintenance and upgrades without any restrictions. This is a huge benefit in comparison to another appliance model like Exadata.
- Flexibility - Easy to repurpose physical servers for different applications and services as needed.
- Availability - Hardware failures are not impactful and critical. In rare case of a server failure, it is easier to associate the logical service profile to another healthy physical server to reduce the impact.

Configure Host, WWNs, and Volume Connectivity with FlashArray Management Tools

Procedure 1. Configure Host

Note: Before using a boot volume (LUN) by a Cisco UCS Blade Server, a host representing this blade server must be defined on Pure Storage FlashArray.

Step 1. Log into Pure Storage FlashArray Management interface.

Step 2. Click the Storage tab.

Step 3. Click the + sign in the Hosts section and select Create Host.

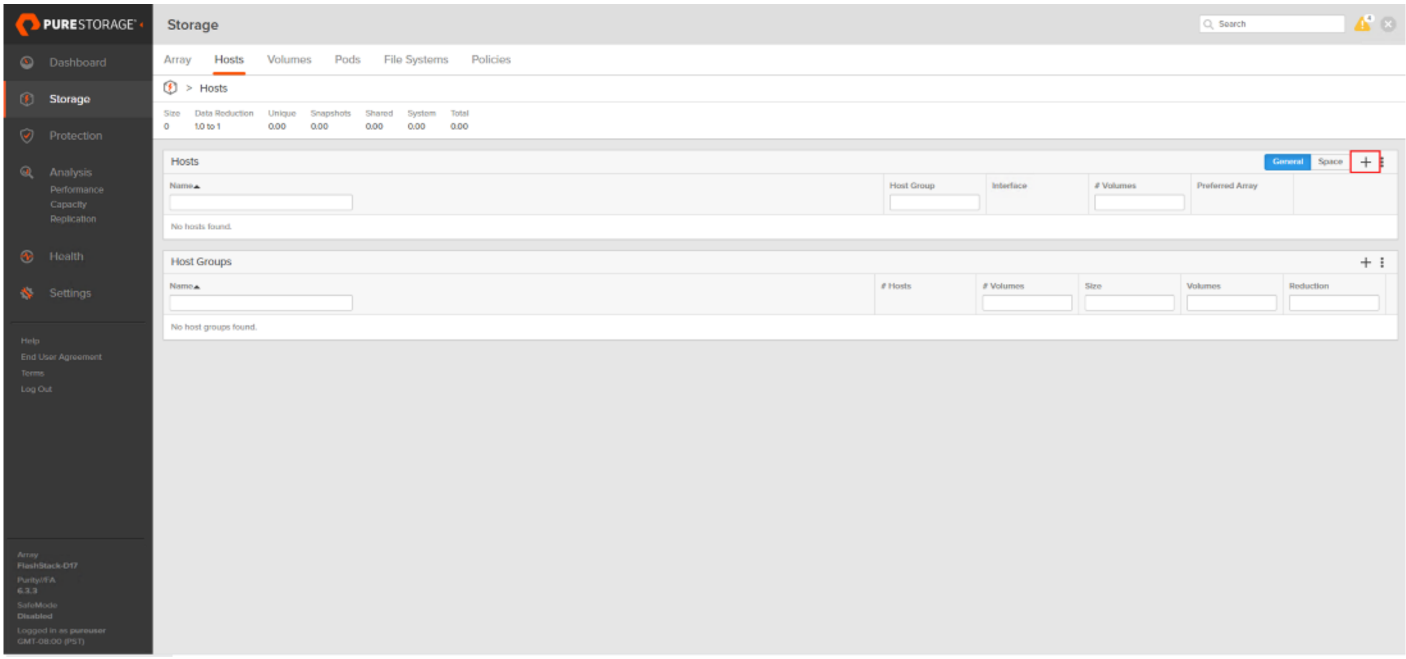

Step 4. Click Create Multiple to create a Host entries under the Hosts category.

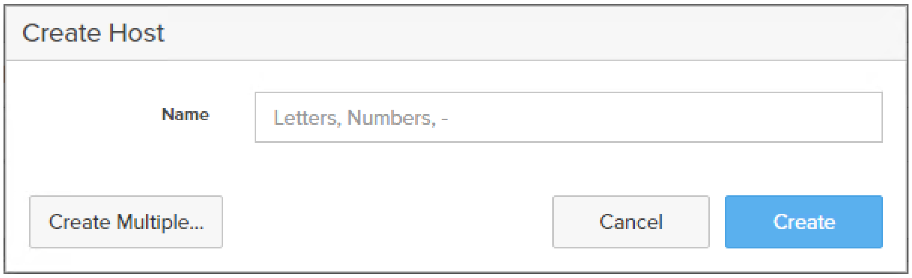
Create Host

Name: Letters, Numbers, -

Step 5. Enter the required information and click Create.

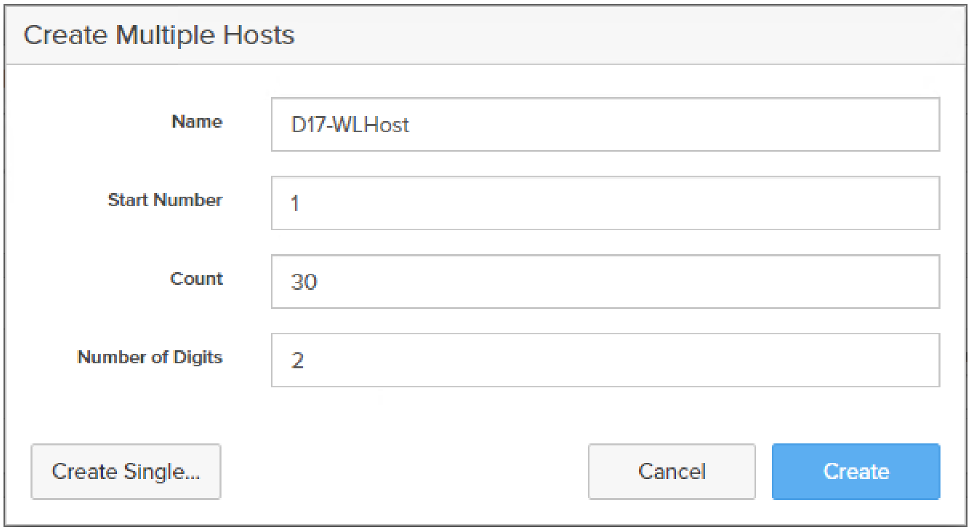

Create Multiple Hosts

Name: D17-WLHost

Start Number: 1

Count: 30

Number of Digits: 2

Step 6. Select one of the newly created hosts, in Host Ports section from the drop-down list select Configure WWNs.

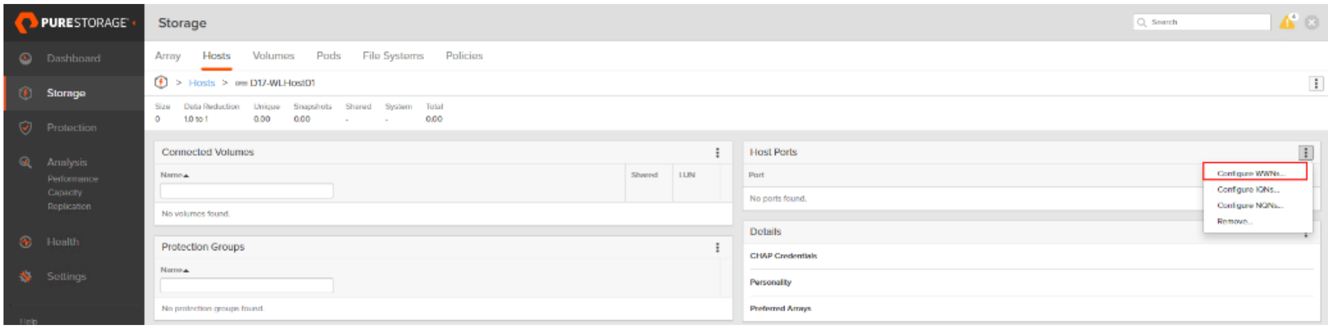

Step 7. Select the list of WWNs that belongs to the host in the next window and click Add.

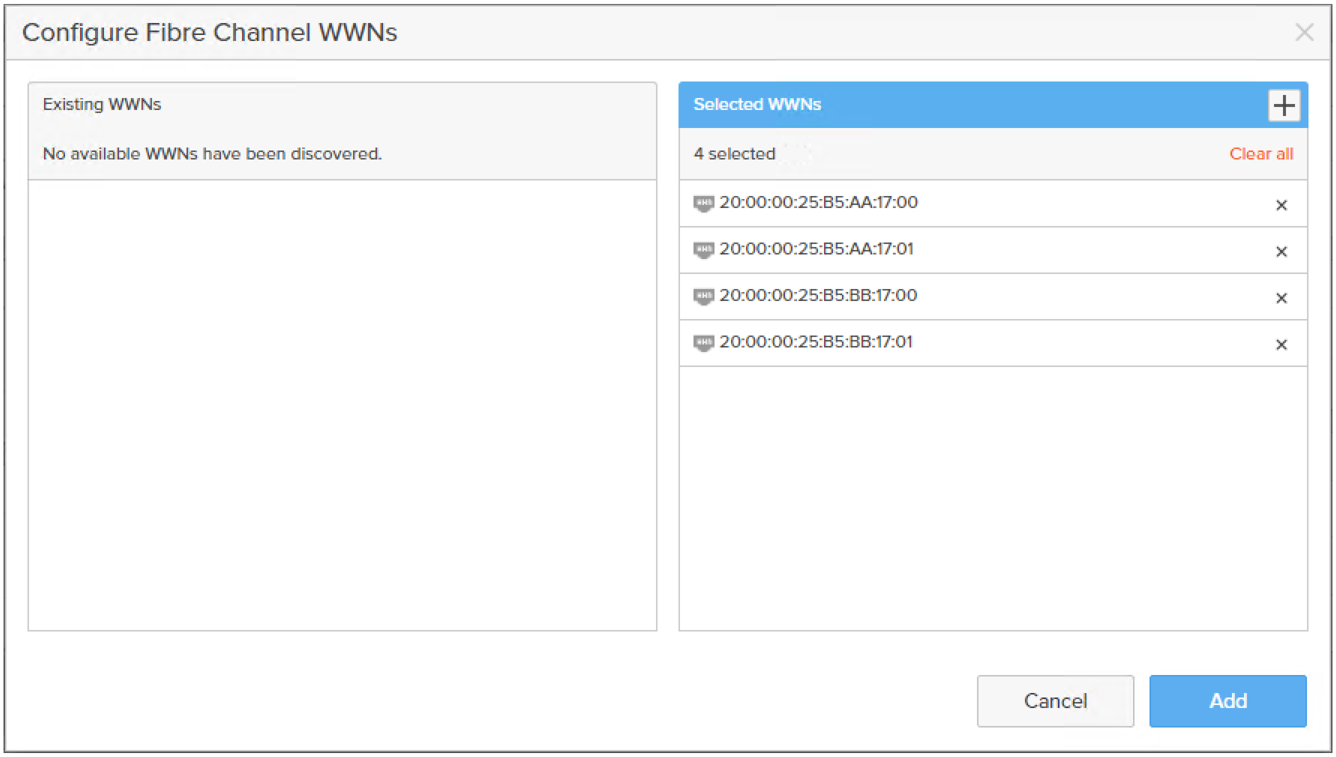

Configure Fibre Channel WWNs

Existing WWNs

No available WWNs have been discovered.

Selected WWNs

4 selected

20:00:00:25:B5:AA:17:00

20:00:00:25:B5:AA:17:01

20:00:00:25:B5:BB:17:00

20:00:00:25:B5:BB:17:01

Note: Make sure the zoning has been setup to include the WWNs details of the initiators along with the target, without which the SAN boot will not work.

Note: WWNs will appear only if the appropriate FC connections were made, and the zones were setup on the underlying FC switch.

Note: Alternatively, the WWN can be added manually by clicking the + in the Selected WWNs section and manually inputting the blade's WWNs.

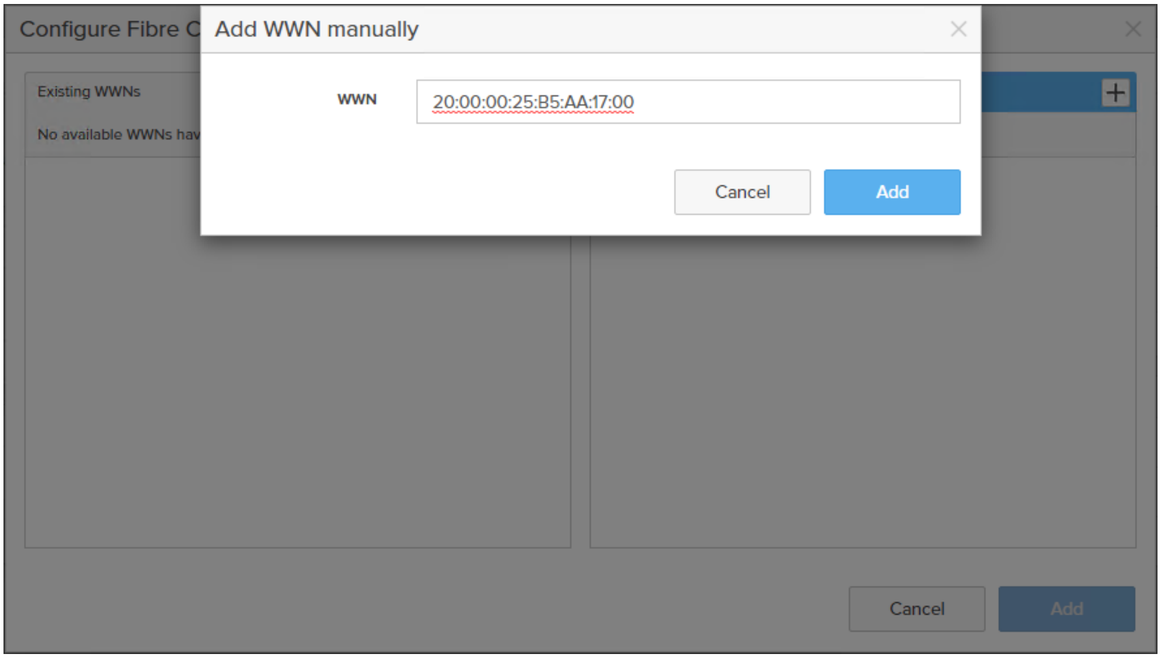

WWN: 20:00:00:25:B5:AA:17:00

Procedure 2. Configure Volume Connectivity

Step 1. Click the Storage tab.

Step 2. Click the + sign in the Volumes section and click Create Volume.

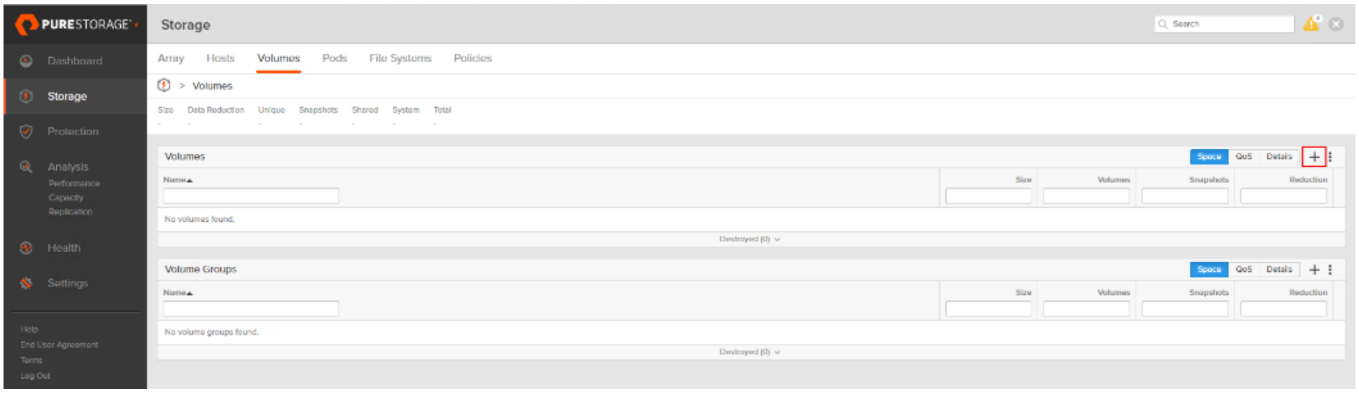

No volumes found.

No volume groups found.

Step 3. Click Create Multiple to open Create Multiple Volumes wizard.

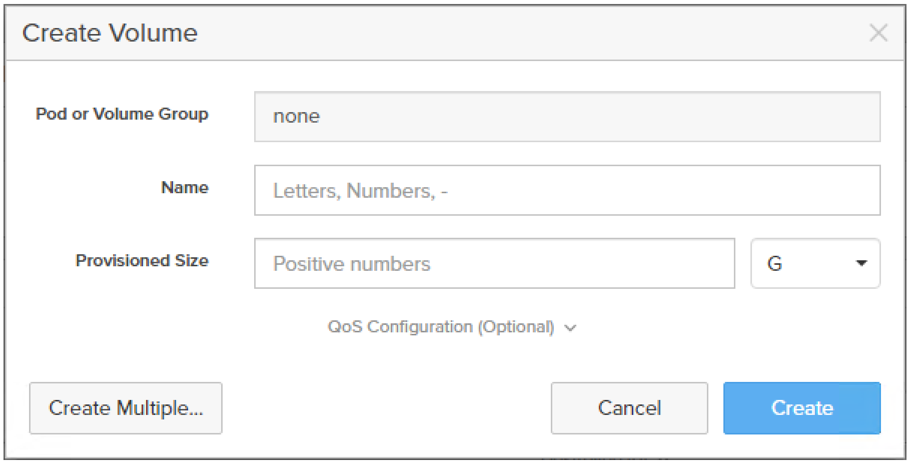

Create Volume

Pod or Volume Group	none
Name	Letters, Numbers, -
Provisioned Size	Positive numbers G

Step 4. Provide the common name of the volume, size, choose the size type (KB, MB, GB, TB, PB) and click Create to create volumes.

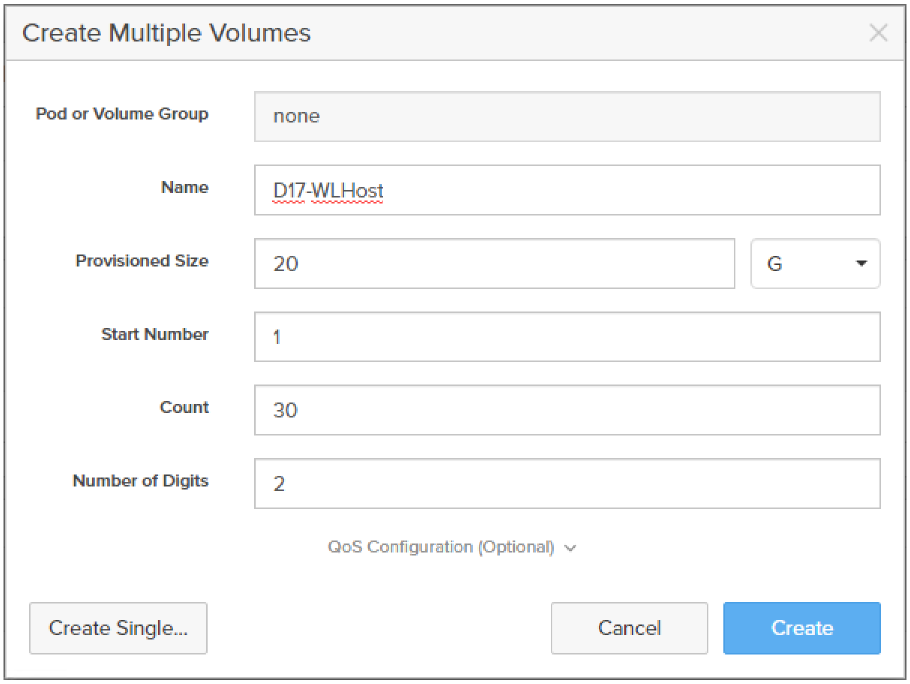

Create Multiple Volumes

Pod or Volume Group	none
Name	D17-WLHost
Provisioned Size	20 G
Start Number	1
Count	30
Number of Digits	2

Step 5. Select one of the hosts and in Connected Volumes section from the drop-down list select Connect.

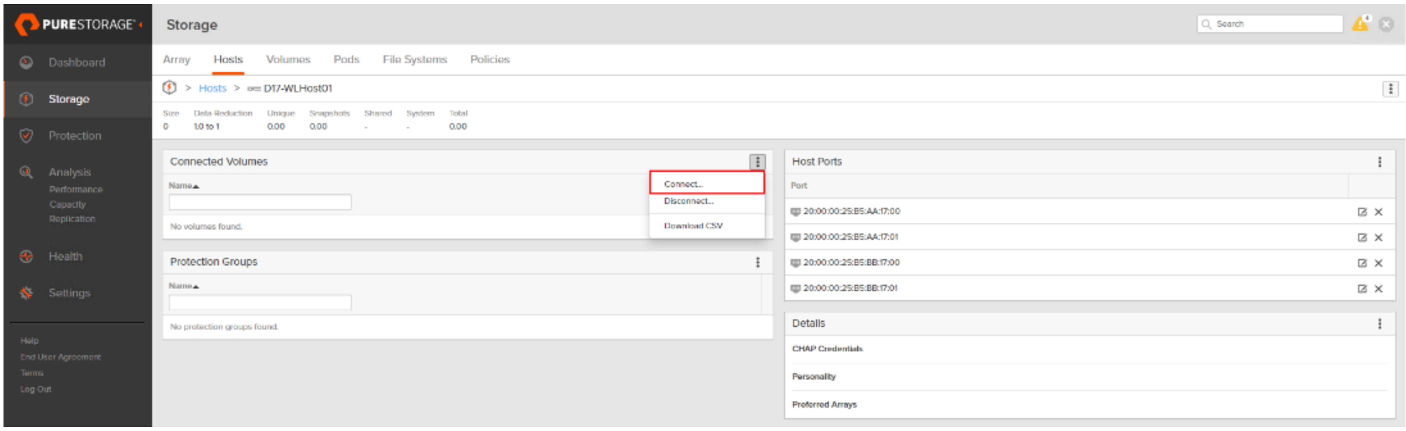

> Hosts > D17-WLHost01

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

Connected Volumes

Name▲

No volumes found.

Protection Groups

Name▲

No protection groups found.

Host Ports

Port
20:00:00:25:B5:AA:17:00
20:00:00:25:B5:AA:17:01
20:00:00:25:B5:BB:17:00
20:00:00:25:B5:BB:17:01

Details

CHAP Credentials

Personality

Preferred Arrays

Step 6. In the Connect Volumes to Host wizard select the volume configured for ESXi installation, click Connect.

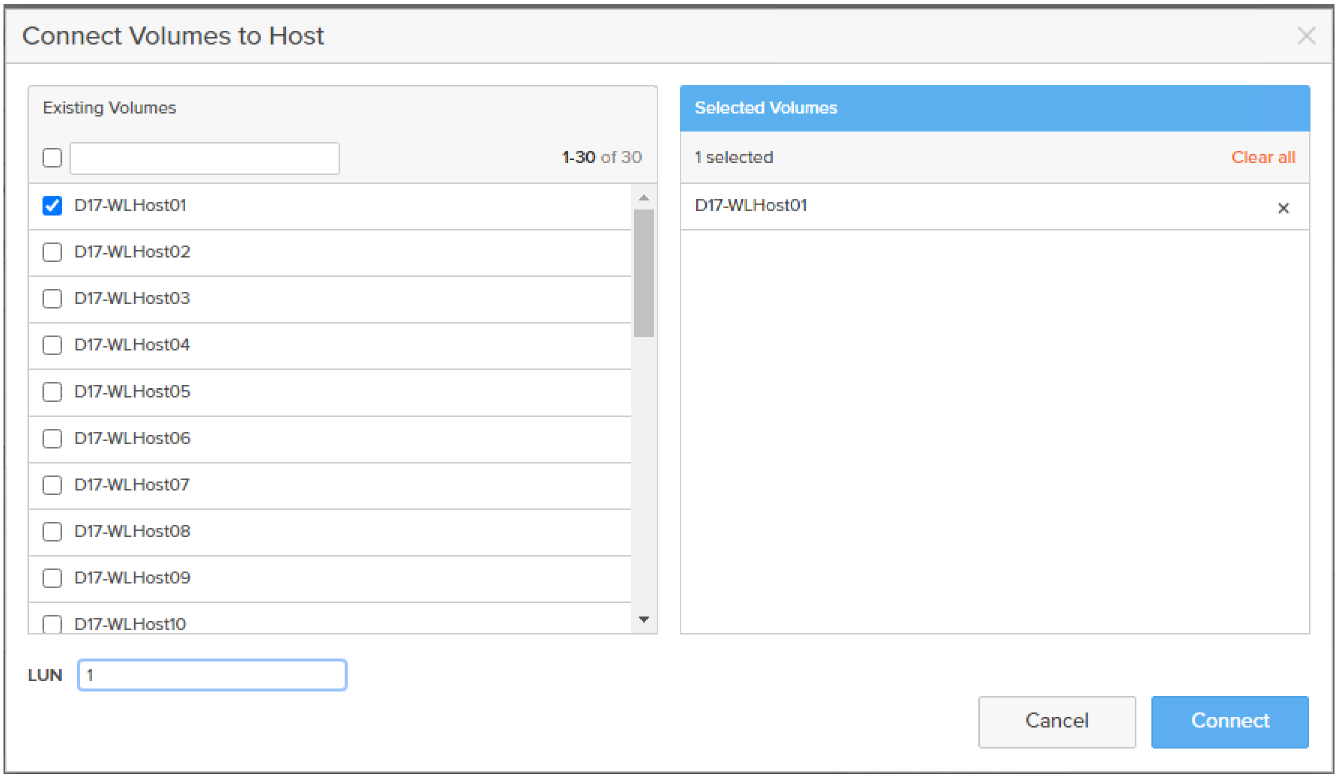

Connect Volumes to Host

Existing Volumes

□ 1-30 of 30

- ☑ D17-WLHost01
- □ D17-WLHost02
- □ D17-WLHost03
- □ D17-WLHost04
- □ D17-WLHost05
- □ D17-WLHost06
- □ D17-WLHost07
- □ D17-WLHost08
- □ D17-WLHost09
- □ D17-WLHost10

Selected Volumes

1 selected

D17-WLHost01

LUN: 1

Note: Make sure the SAN Boot Volumes has the LUN ID “1” since this is important while configuring Boot from SAN. You will also configure the LUN ID as “1” when configuring Boot from SAN policy in Cisco UCS Manager.

Note: More LUNs can be connected by adding a connection to existing or new volume(s) to an existing node.

Configure File Services

FA File services can be activated by Pure Storage Technical Services (Support). Please refer to FA File Services Support Matrix to verify that your hardware offers support for running File Services.

Currently all FA File services activations require Pure Storage Product Management approval. Customers can work with their local account representatives to obtain approval to activate File Services.

For additional information on FA File Services setup and configuration see:

- FA File Services Quick Start Guide
- FA File Services Best Practices

Procedure 1. Create Virtual Interface(s)

The VIF provides high-availability network access across 2 physical Ethernet ports per array controller. Each VIF requires 2 physical ports per controller. Any physical ethernet port can be used with the restriction that any port that is in use by management services, a bond, or subnet configuration cannot be part of a VIF. For the maximum number of VIFs supported, please see the FA File Services Limits KB.

Note: VIFs are created by CLI over SSH, configured and enabled using the Management Console. An account with administrator privileges is required.

Step 1. Connect to the array via SSH.

Step 2. Run the following syntax to create the VIF on the array:

```
purenetwork create vif --subinterfacelist ct0.ethX,ct1.ethX,ct0.ethY,ct1.ethY <name of interface>
```

Procedure 2. Configure and Enable the Virtual Interface for File Services

Step 1. Connect to the array GUI.

Step 2. Navigate to Settings > Network.

Step 3. Locate the File VIF in the interface list and click the edit icon.

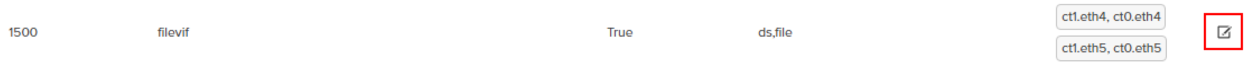

1500	filevif	True	ds,file	ct1.eth4, ct0.eth4 ct1.eth5, ct0.eth5

Step 4. In the Edit Interface dialog turn on the Enabled option, provide the IP Address, Netmask, and Gateway used by the interface. Click Save.

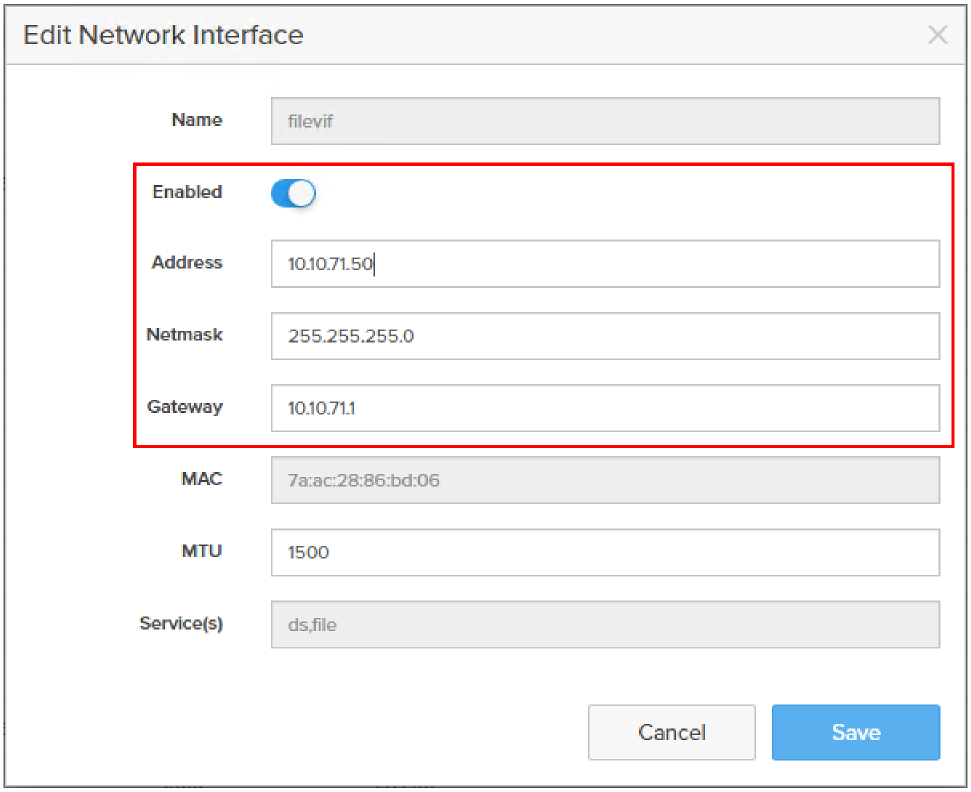

Edit Network Interface

Name	filevif
Enabled	●
Address	10.10.71.50
Netmask	255.255.255.0
Gateway	10.10.71.1
MAC	7a:ac:28:86:bd:06
MTU	1500
Service(s)	ds,file

Step 5. Scroll to the bottom of the Network tab and click the edit icon for DNS Settings.

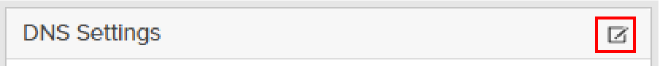

Step 6. In the Edit DNS Settings dialog, enter desired values for Domain and DNS server IPs. Click Save.

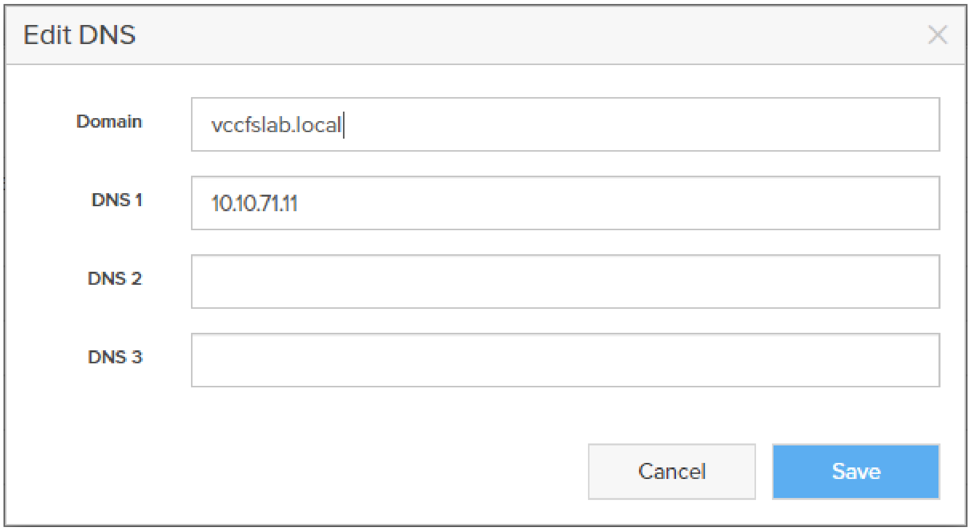
Edit DNS

Domain: vccfslab.local

DNS 1: 10.10.71.11

DNS 2:

DNS 3:

Note: More than one DNS server can be configured with the caveat that all DNS servers must have a record for Directory Service servers such as LDAP or Microsoft Active Directory.

Procedure 3. Create Active Directory Account for the Array

Step 1. Navigate to Settings > Access > Active Directory Accounts.

Step 2. To open the Create Dialog, click the + icon.

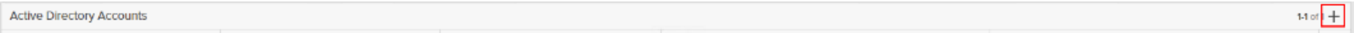

Step 3. Enter the following information:

- Name = Array management name for this AD account
- Domain = AD domain name
- Computer Name = Computer Object name within AD
- User = Domain user that can create computer objects and join to the domain.
- Password = Users password for the above domain user

Step 4. Click Create to finalize AD account creation.

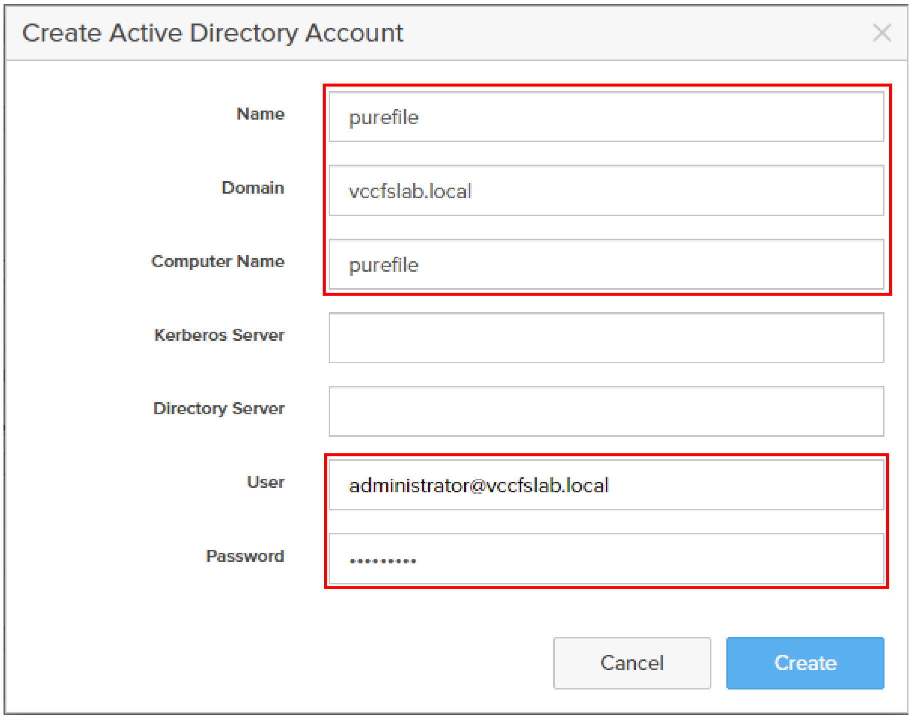

Create Active Directory Account

Name: purefile

Domain: vccfslab.local

Computer Name: purefile

Kerberos Server:

Directory Server:

User: administrator@vccfslab.local

Password: •••••••••

Procedure 4. Create a File System and Shared Directory

Step 1. Navigate to Storage > File Systems.

Step 2. Click the + icon.

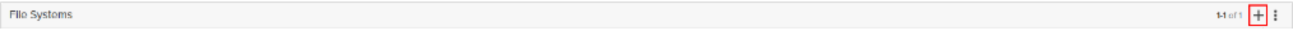

Step 3. In Create File System enter a file system name and click Create.

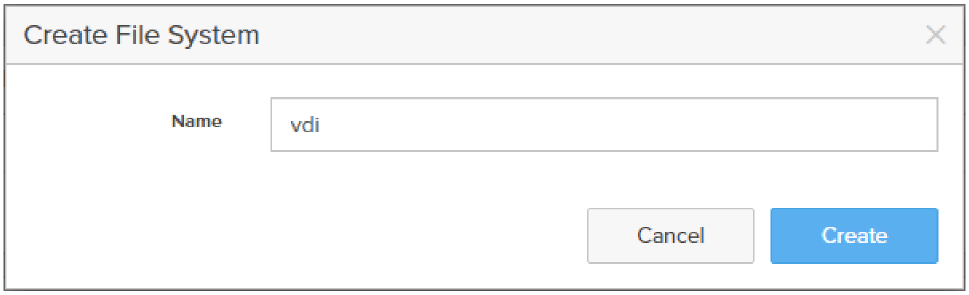

Create File System

Name: vdi

Step 4. Navigate to Storage > File Systems > Directories.

Step 5. Click the + icon.

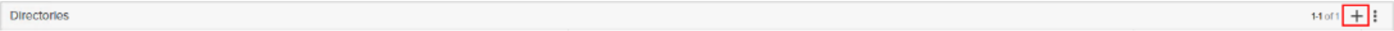

Step 6. In Create Directory, enter Select a file system from the drop-down list, enter the desired management name of the directory, and enter the directory path in the file system. (for example, dir or /dir, for sub-level directories /dir/subdir or /dir/subdir/subdir1 can be used). Click Create.

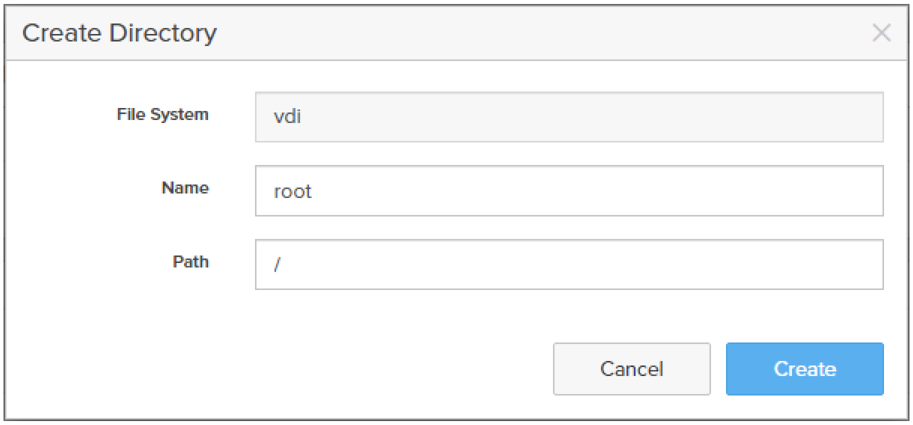

Create Directory

File System: vdi

Name: root

Path: /

Note: Polices for exports/shares/snapshots can only be attached to managed directories at the file system root or 1 level deep (/ and /dir in the example above). Space and performance metrics can be seen at all levels of managed directories.

Step 7. Navigate to Storage > Policies.

Step 8. Click the + icon.

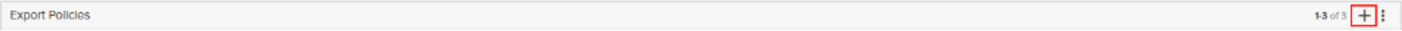

Step 9. In the Create Export Policy pop-up choose SMB from the Type drop-down list and enter a name for the policy. Click Create.

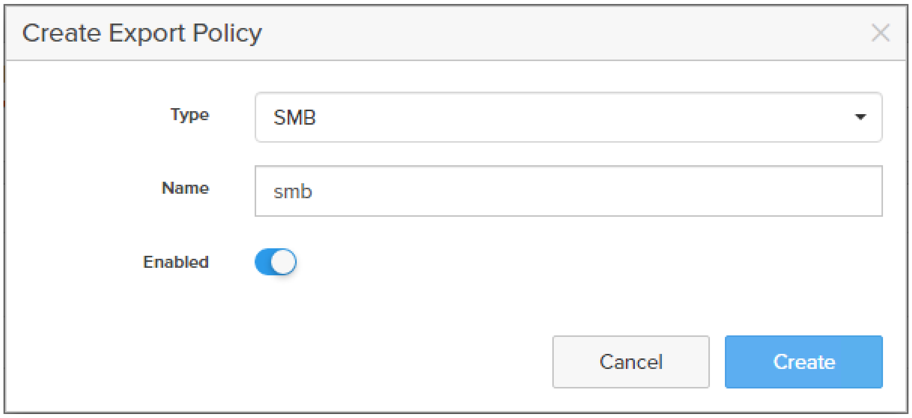

Create Export Policy

Type: SMB

Name: smb

Enabled: ●

Step 10. Click Created Policy and click the + icon.

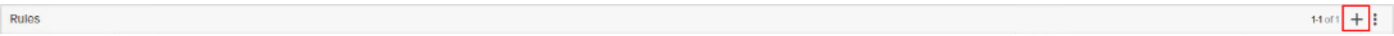

Step 11. Complete the Client filter for read-write access and click Add to complete the rule creation.

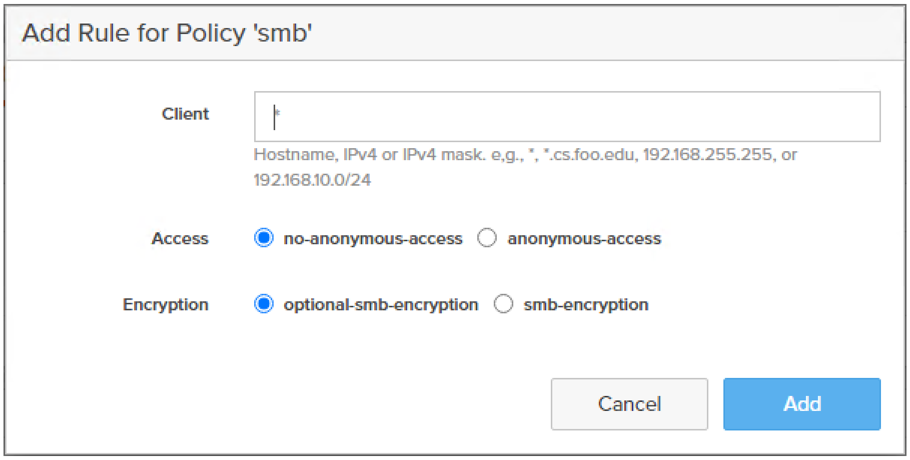

Add Rule for Policy 'smb'

Client: *

Hostname, IPv4 or IPv4 mask. e,g., *, *.cs.foo.edu, 192.168.255.255, or 192.168.10.0/24

Access: ● no-anonymous-access ○ anonymous-access

Encryption: ● optional-smb-encryption ○ smb-encryption

Step 12. Attach the export policy(s) to a managed directory. Click the + icon.

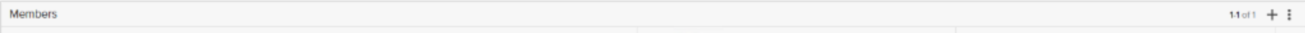

Step 13. Select a managed directory from the drop-down list, enter a share/export name, and click Create.

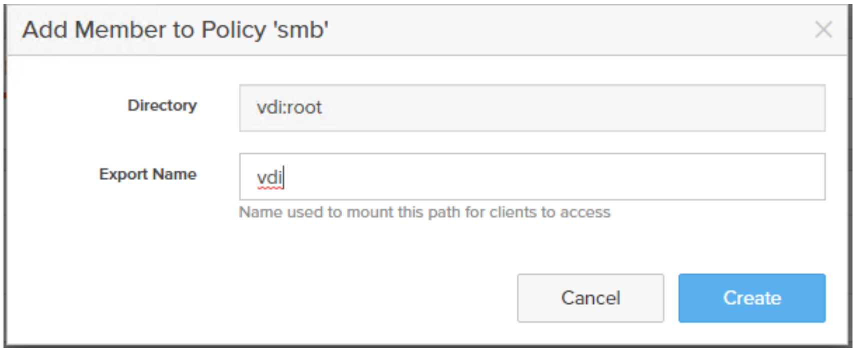

Add Member to Policy 'smb'

Directory: vdi:root

Export Name: vdi

Name used to mount this path for clients to access

Step 14. Verify access to the created share from the Windows client.

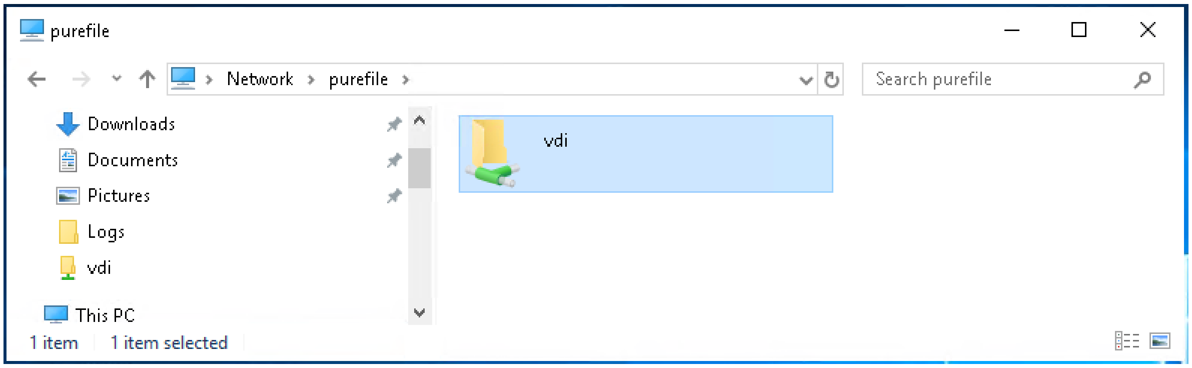

Install and Configure VMware ESXi 8.0

This section explains how to install VMware ESXi 8.0 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 boot logical unit number (LUN). Upon completion of steps outlined here, ESXi hosts will be booted from their corresponding SAN Boot LUNs.

Download Cisco Custom Image for VMware vSphere ESXi 8.0

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

Procedure 1. Install VMware vSphere ESXi 8.0

Step 1. In the Cisco UCS Manager navigation pane, click the Equipment tab.

Step 2. Under Servers > Service Profiles> VDI-Host1

Step 3. Right-click on VDI-Host1 and select KVM Console.

Step 4. Click Boot Device and then select CD/DVD.

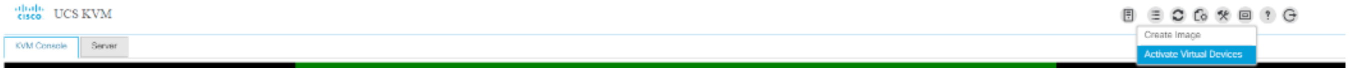

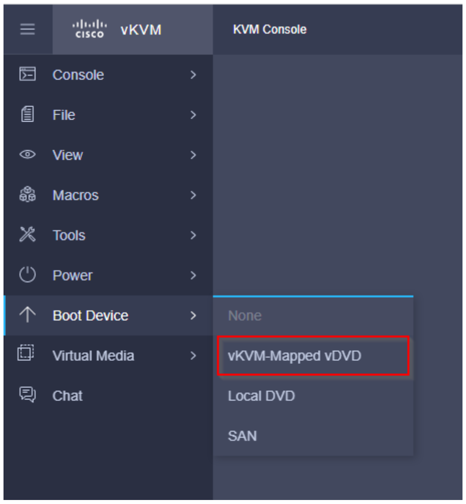

Step 5. Click Virtual Media and Mount the ESXi ISO image.

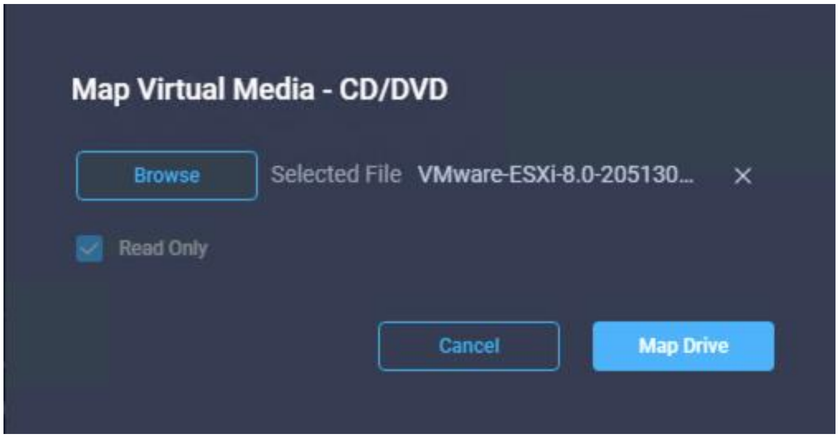
Map Virtual Media - CD/DVD

Browse Selected File VMware-ESXi-8.0-205130...

☑ Read Only

Step 6. Boot into ESXi installer and follow the prompts to complete installing VMware vSphere ESXi hypervisor.

Step 7. When selecting a storage device to install ESXi, select Remote LUN provisioned through Pure Storage Administrative console and access through FC connection.

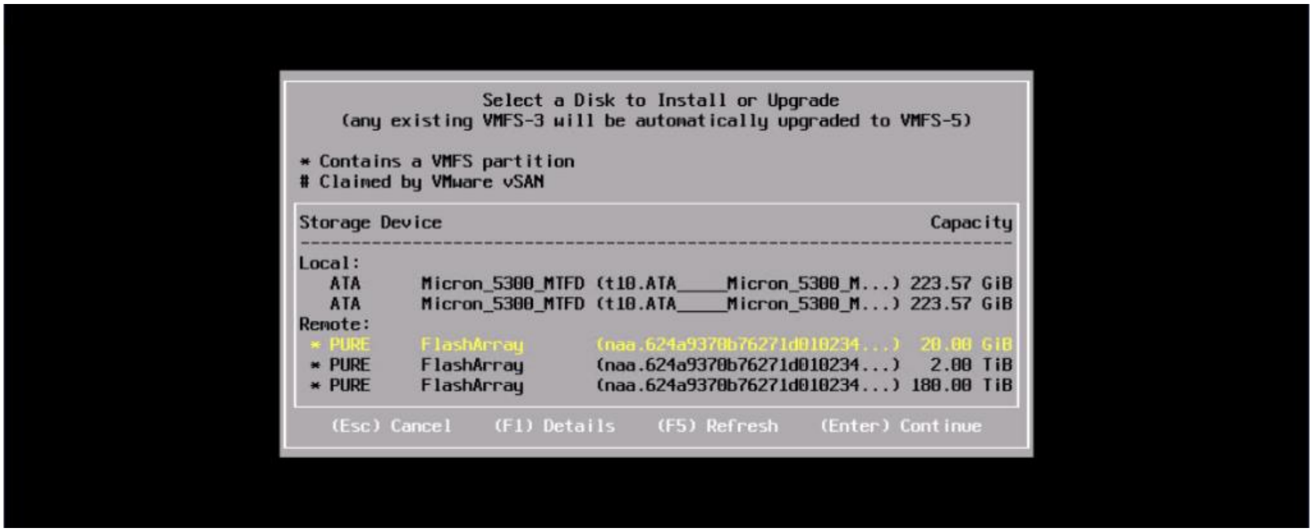

```
          Select a Disk to Install or Upgrade
  (any existing VMFS-3 will be automatically upgraded to VMFS-5)

* Contains a VMFS partition
# Claimed by VMware vSAN

Storage Device                                                   Capacity
--------------------------------------------------------------------------
Local:
   ATA      Micron_5300_MTFD (t10.ATA_____Micron_5300_M...) 223.57 GiB
   ATA      Micron_5300_MTFD (t10.ATA_____Micron_5300_M...) 223.57 GiB
Remote:
 * PURE     FlashArray       (naa.624a9370b76271d010234...)  20.00 GiB
 * PURE     FlashArray       (naa.624a9370b76271d010234...)   2.00 TiB
 * PURE     FlashArray       (naa.624a9370b76271d010234...) 180.00 TiB

   (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 26. Sample ESXi Configure Management Network

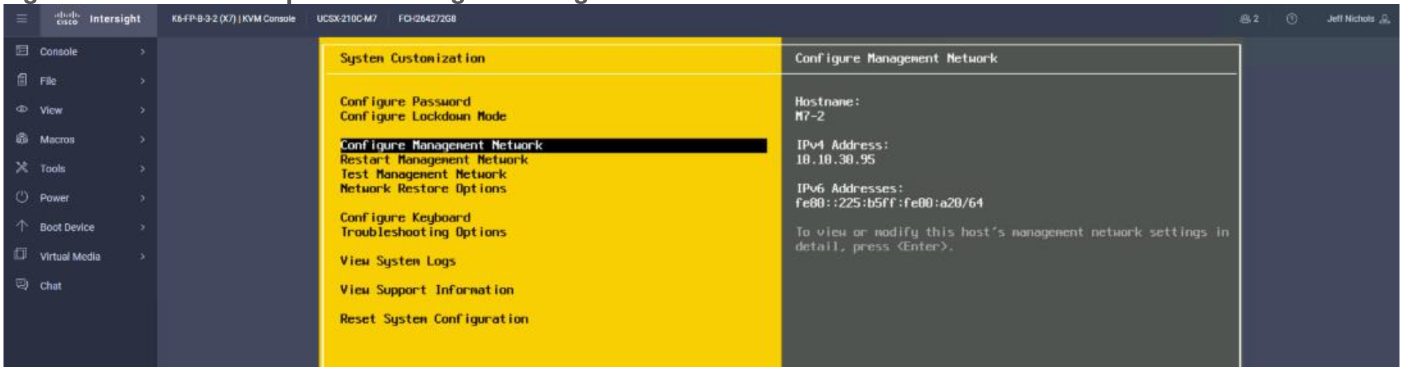

```
System Customization                         Configure Management Network

Configure Password                           Hostname:
Configure Lockdown Mode                      M7-2

Configure Management Network                 IPv4 Address:
Restart Management Network                   10.10.30.95
Test Management Network
Network Restore Options                      IPv6 Addresses:
                                             fe80::225:b5ff:fe00:a20/64
Configure Keyboard
Troubleshooting Options                      To view or modify this host's management network settings in
                                             detail, press <Enter>.
View System Logs

View Support Information

Reset System Configuration
```

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 27. Servers HCL Status in Cisco Intersight Infrastructure Services

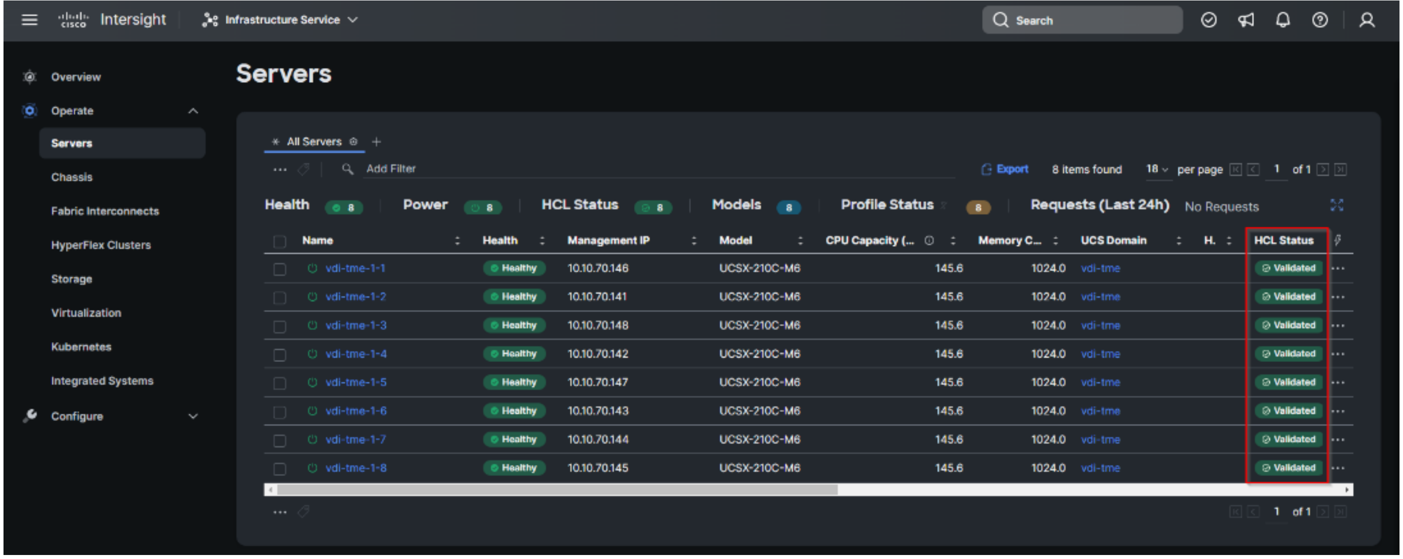

Servers

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status
vdi-tme-1-1	Healthy	10.10.70.146	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-2	Healthy	10.10.70.141	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-3	Healthy	10.10.70.148	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-4	Healthy	10.10.70.142	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-5	Healthy	10.10.70.147	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-6	Healthy	10.10.70.143	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-7	Healthy	10.10.70.144	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-8	Healthy	10.10.70.145	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated

In this Validated Design the following drivers were used (VMware-ESXi-8.0-20513097-Custom-Cisco-4.2.3-b):

- Cisco-nenic- 1.0.45.0-1OEM.700.1.0.15843807
- Cisco-nfnic- 5.0.0.37-1OEM.700.1.0.15843807

VMware Clusters

The VMware vSphere Client was configured to support the solution and testing environment as follows:

- Datacenter: FlashStack - Pure Storage FlashArray//X70 R3 with Cisco UCS
- Cluster: FlashStack-VDI - Single-session/Multi-session OS VDA workload
- Infrastructure : Infrastructure virtual machines (vCenter, Active Directory, DNS, DHCP, SQL Server, Citrix StoreFront Servers, Citrix Apps and Desktop Controllers, and other common services), Login VSI launcher infrastructure were connected using the same set of switches but hosted on separate HX 4.5.2a 4 server cluster.

Figure 28. **VMware vSphere WebUI Reporting Cluster Configuration for this Validated Design**

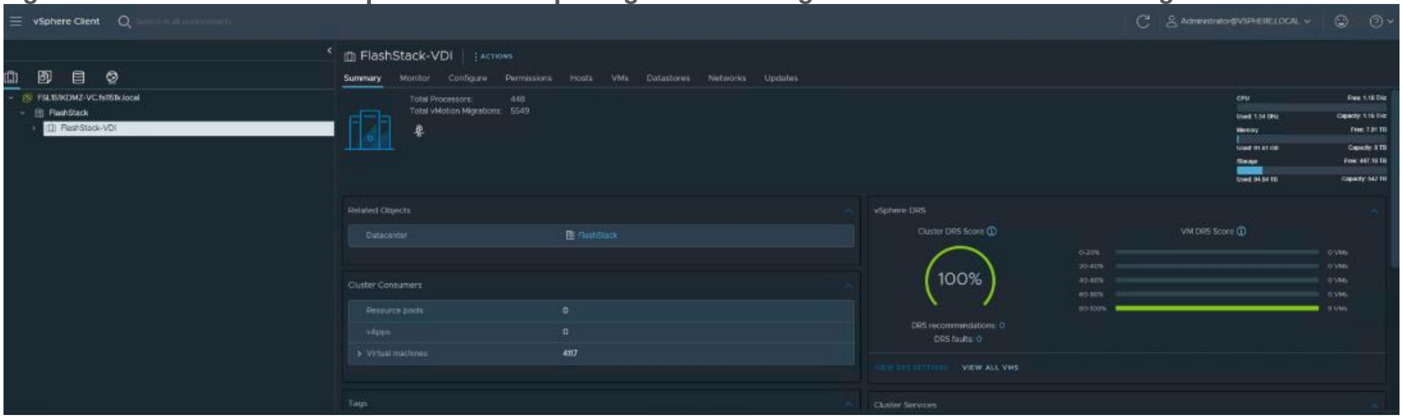

FlashStack-VDI

Total Processors: 448
Total vMotion Migrations: 5549

CPU Free: 1.16 THz
Used: 1.34 GHz Capacity: 1.16 THz
Memory Free: 7.91 TB
Used: 91.61 GB Capacity: 8 TB
Storage Free: 447.16 TB
Used: 94.84 TB Capacity: 542 TB

Related Objects

Datacenter	FlashStack

Cluster Consumers

Resource pools	0
vApps	0
Virtual machines	4117

Tags

vSphere DRS

Cluster DRS Score: 100%

DRS recommendations: 0
DRS faults: 0

VM DRS Score

0-20%	0 VMs
20-40%	0 VMs
40-60%	0 VMs
60-80%	0 VMs
80-100%	9 VMs

Cisco Intersight Orchestration

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. FlashStack 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 FlashStack-Management Cluster, first download the latest release of the OVA from: https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-230.

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 FlashStack-Management cluster and click Deploy OVF Template.

Step 4. Specify a URL or browse to the intersight-virtual-appliance-1.0.9-230.ova file. Click NEXT.

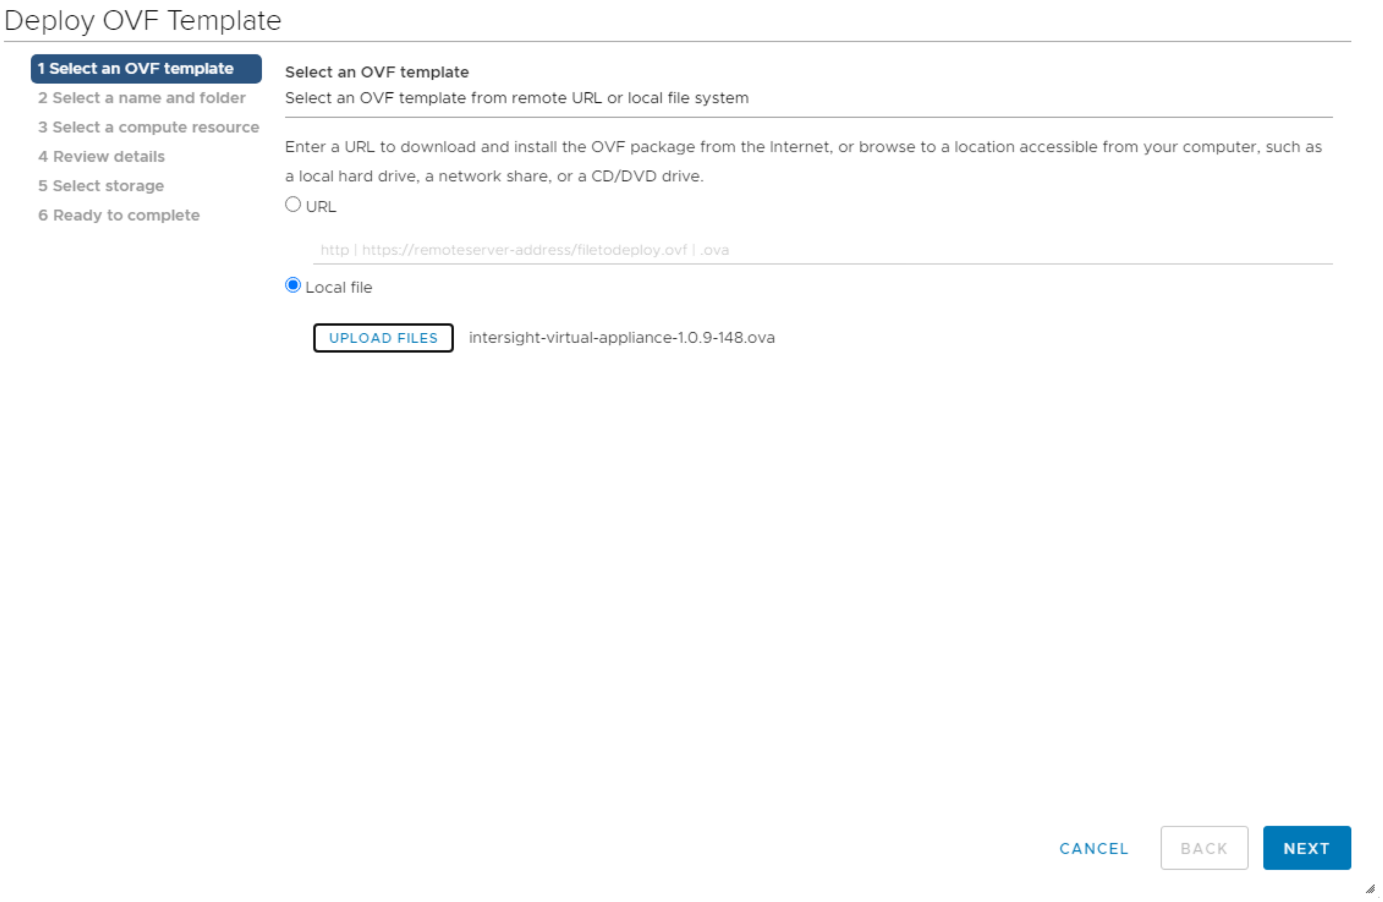

Select an OVF template

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 FlashStack-Management cluster and click NEXT.

Step 7. Review details and click NEXT.

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

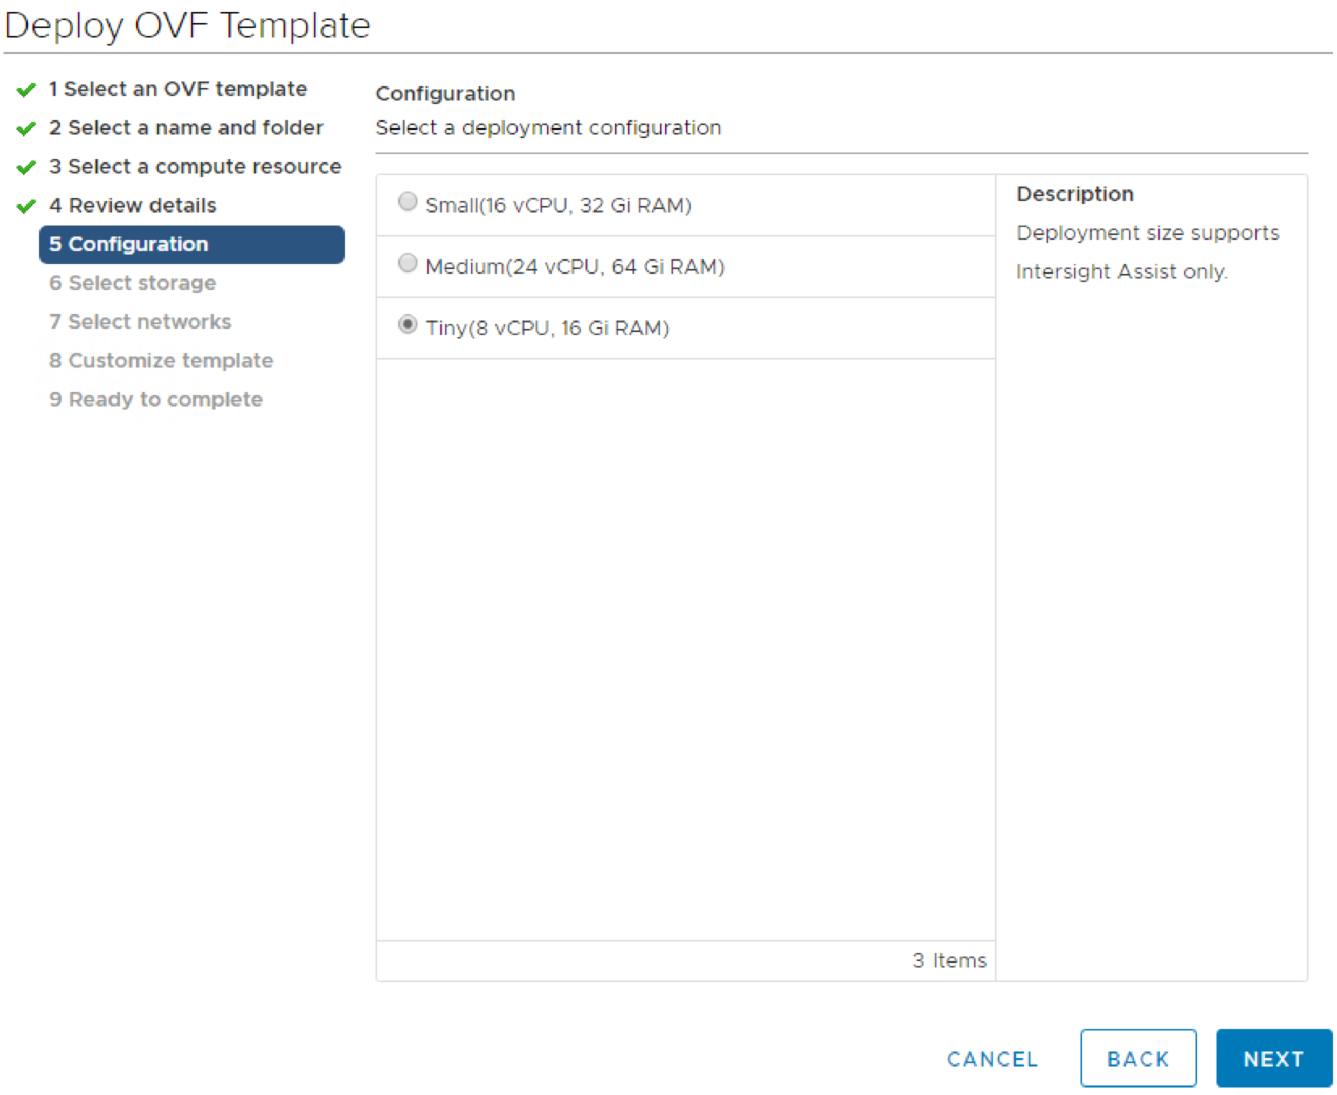

Deploy OVF Template

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

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.

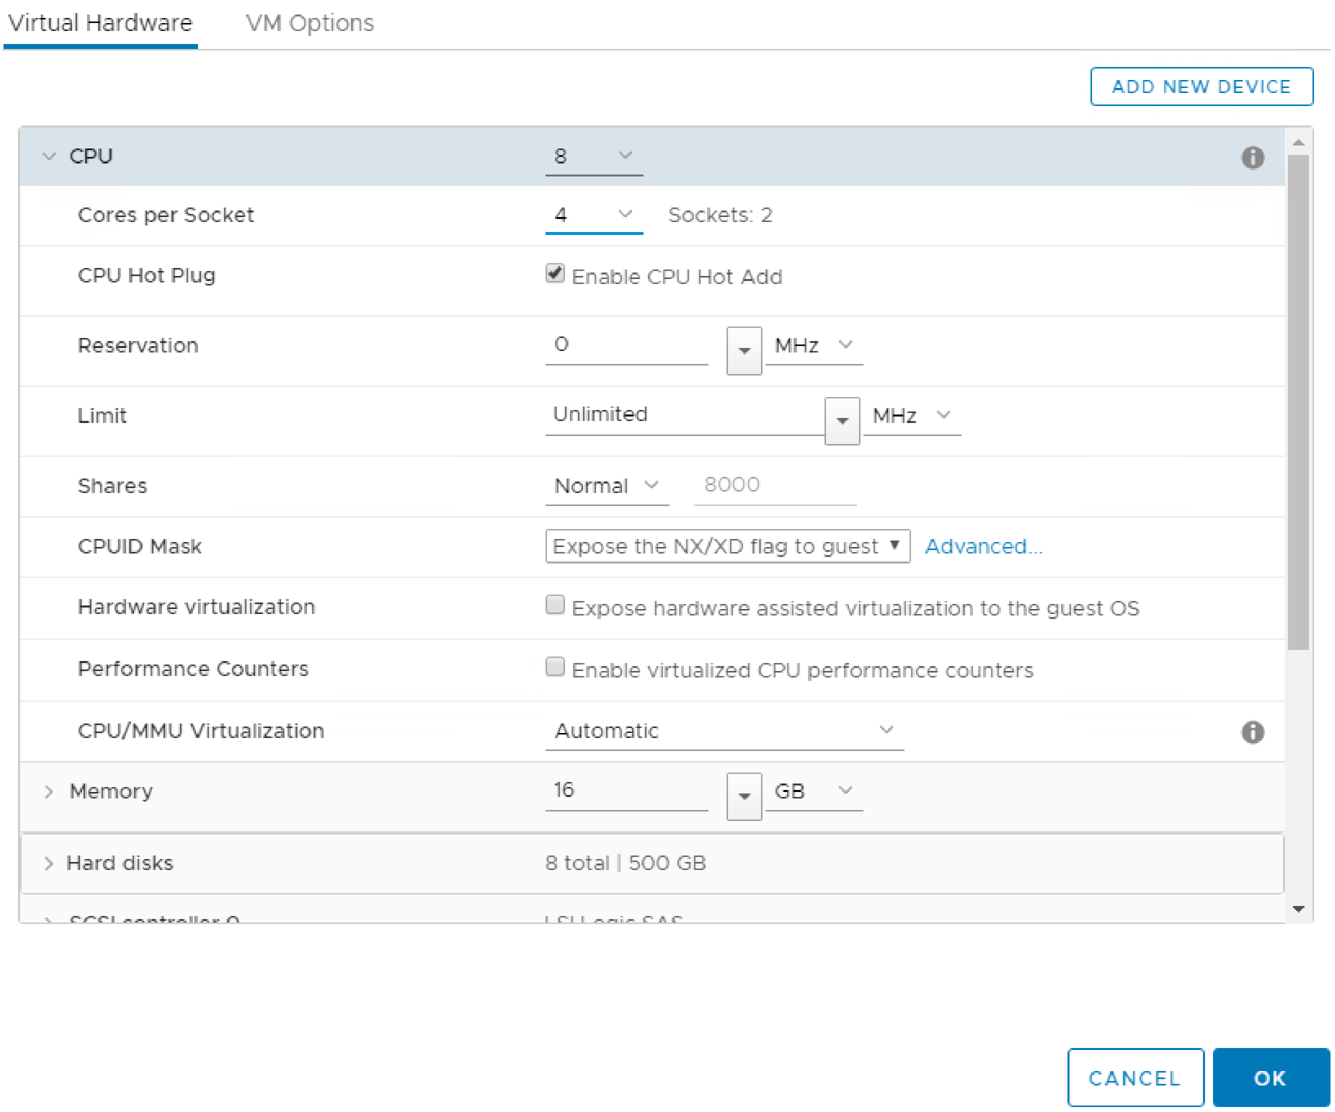

CPU	8
Cores per Socket	4 Sockets: 2
CPU Hot Plug	☑ Enable CPU Hot Add
Reservation	0 MHz
Limit	Unlimited MHz
Shares	Normal 8000
CPUID Mask	Expose the NX/XD flag to guest Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS
Performance Counters	□ Enable virtualized 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 choose 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.

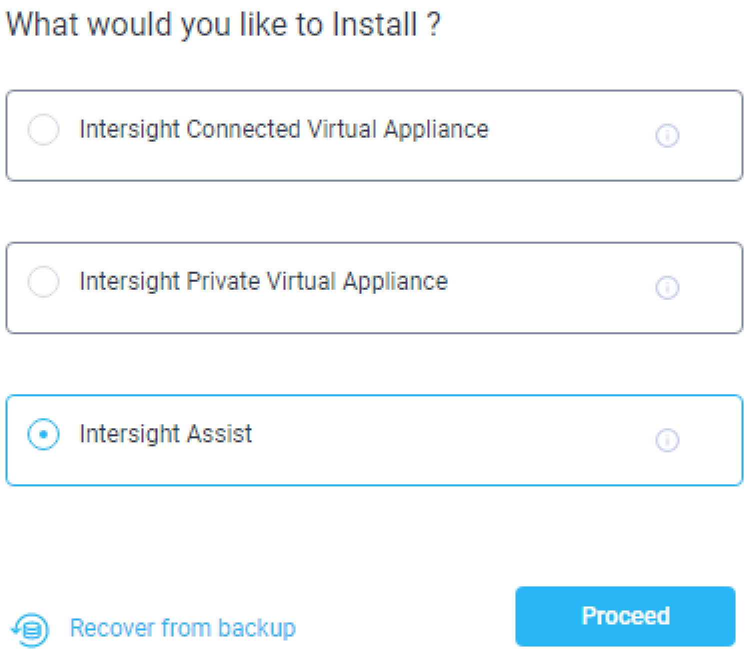

What would you like to Install ?

○ 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.

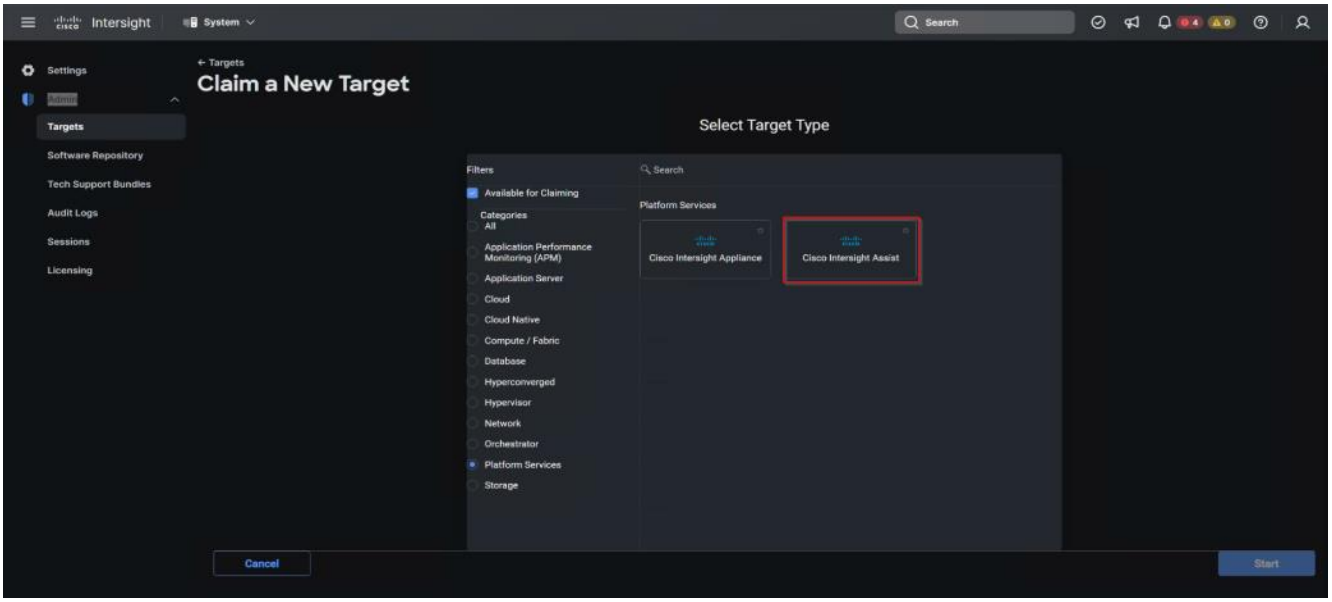

Step 3. Fill in the Intersight Assist information and click Claim.

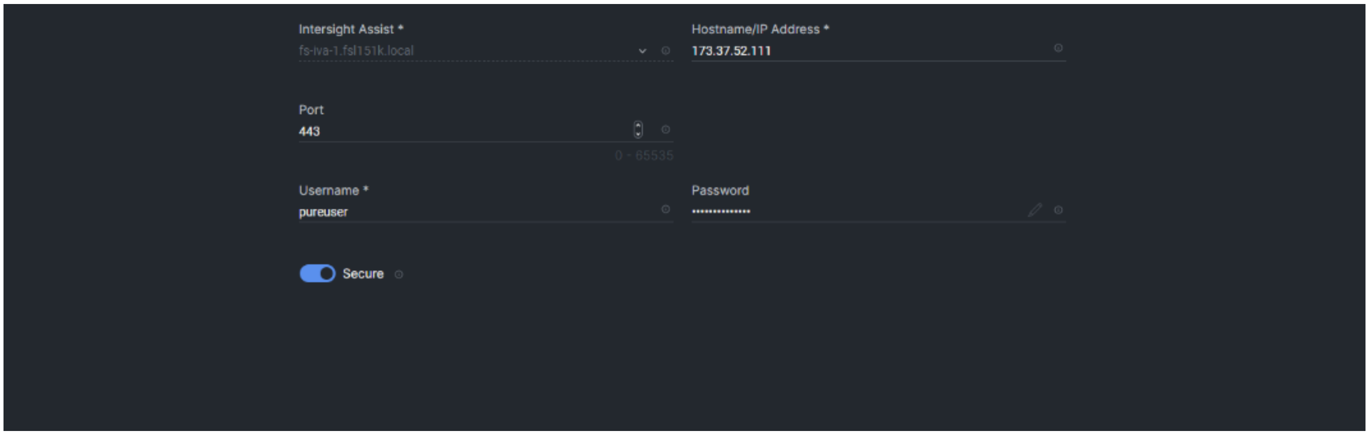

Intersight Assist *
fs-iva-1.fsl151k.local

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.

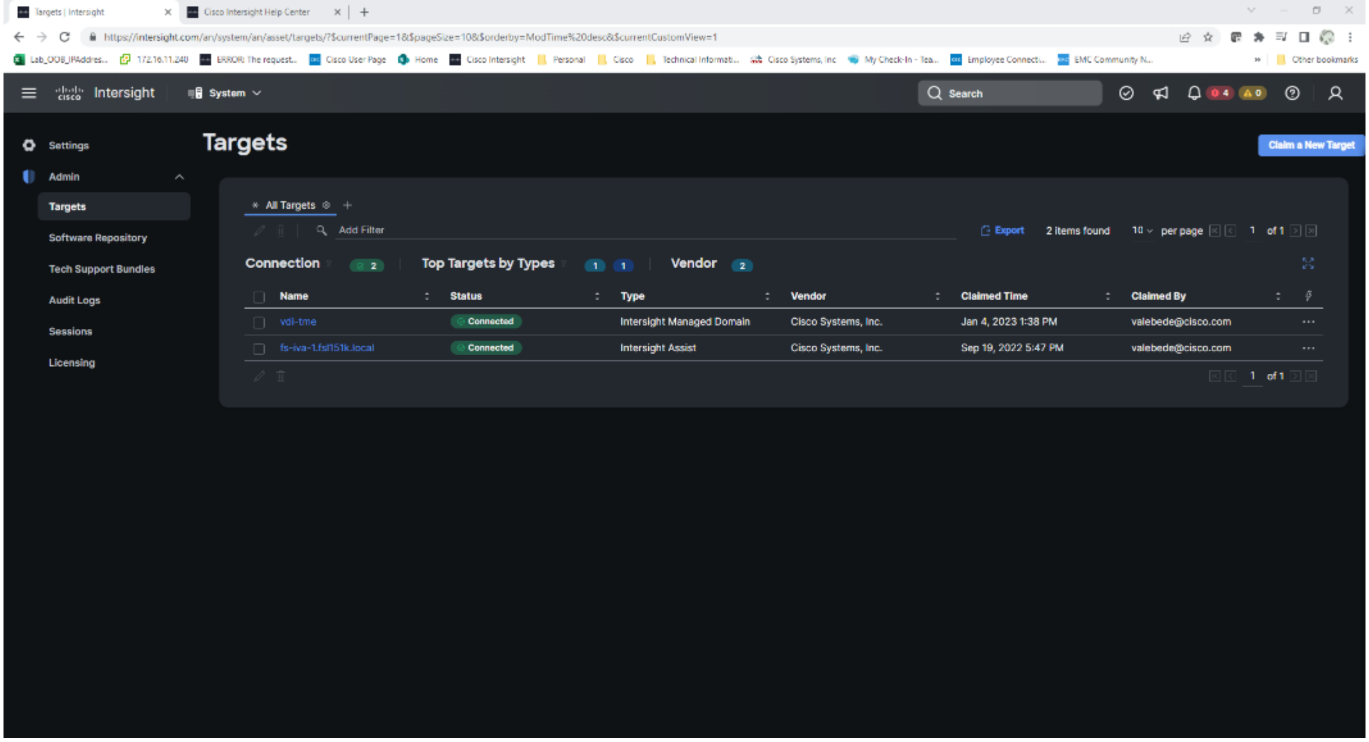

Targets

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.

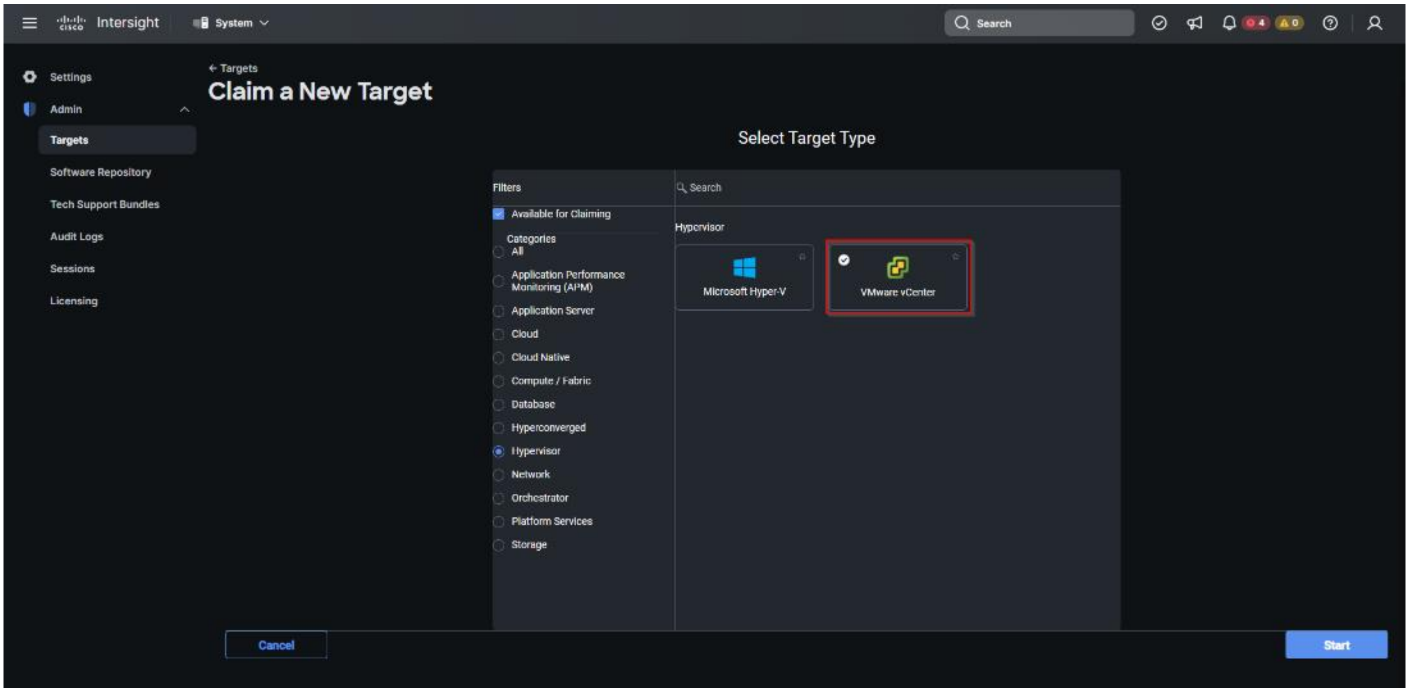

Step 2. In the VMware vCenter window, make sure the Intersight Assist is correctly selected, fill in the vCenter information, and click Claim.

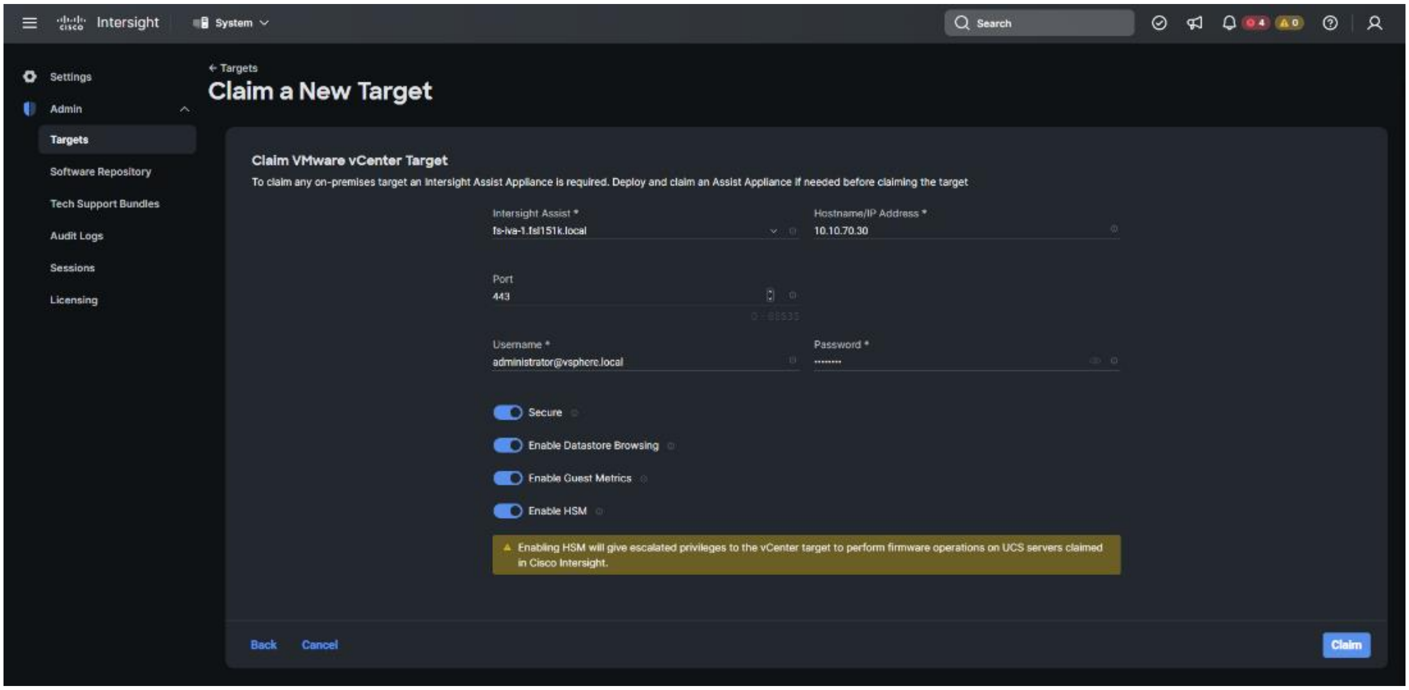

Claim a New Target

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 *: 10.10.70.30

Port: 443

Username *: administrator@vsphere.local

Password *: ••••••••

Secure

Enable Datastore Browsing

Enable Guest Metrics

Enable HSM

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

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.

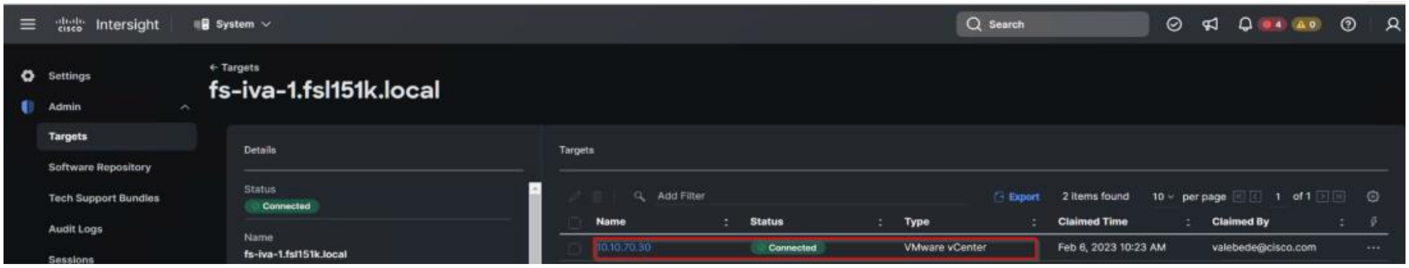

fs-iva-1.fsl151k.local

Details

Status: Connected

Name: fs-iva-1.fsl151k.local

Targets

2 items found

Name	Status	Type	Claimed Time	Claimed By
10.10.70.30	Connected	VMware vCenter	Feb 6, 2023 10:23 AM	valebede@cisco.com

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

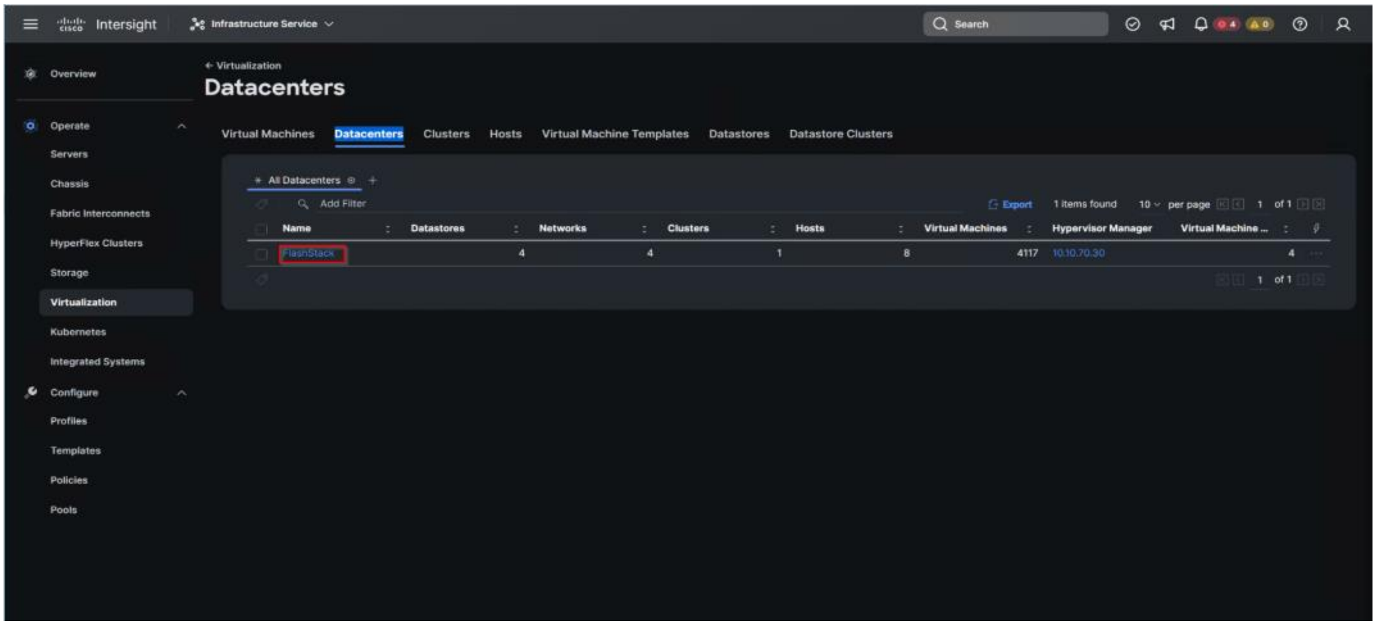

Datacenters

1 items found

Name	Datastores	Networks	Clusters	Hosts	Virtual Machines	Hypervisor Manager	Virtual Machine ...
FlashStack	4	4	1	8	4117	10.10.70.30	4

Procedure 4. Claim FlashArray//X in Cisco Intersight

Note: Claiming a Pure Storage FlashArray also requires the use of an Intersight Assist virtual machine.

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

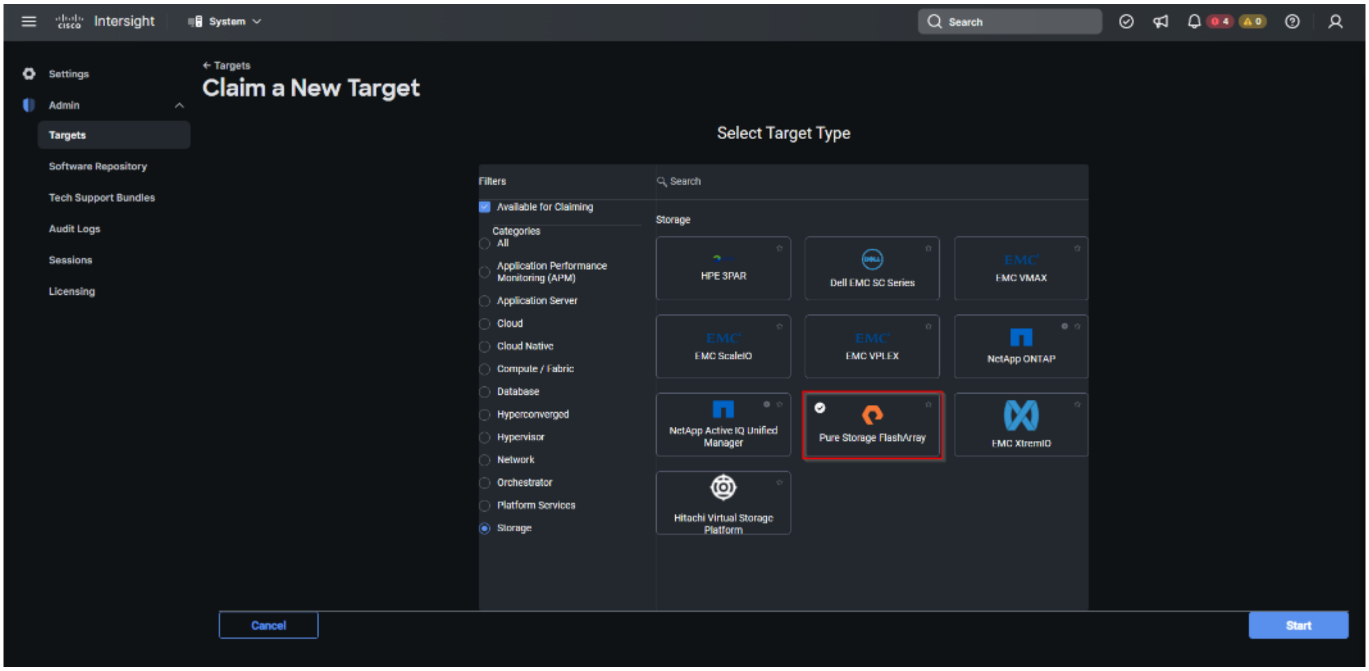

Step 2. Enter FlashArray Hostname/ IP address and credentials and click Claim.

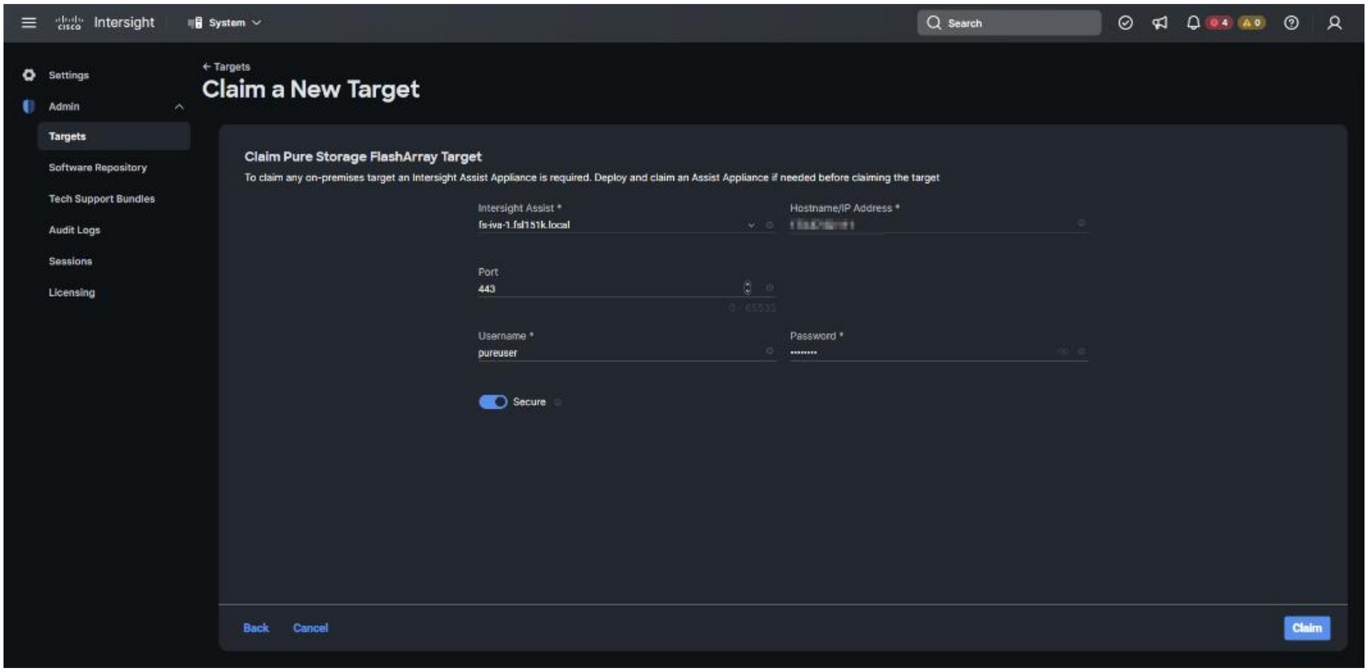

Claim Pure Storage FlashArray Target

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 * [redacted]

Port 443

Username * pureuser

Password * ••••••••

Secure

Procedure 5. FC Host Registration using Cisco Intersight

Step 1. From Cloud Orchestration service, select Configure > Workflows.

Step 2. Select New Storage Host. Click Execute.

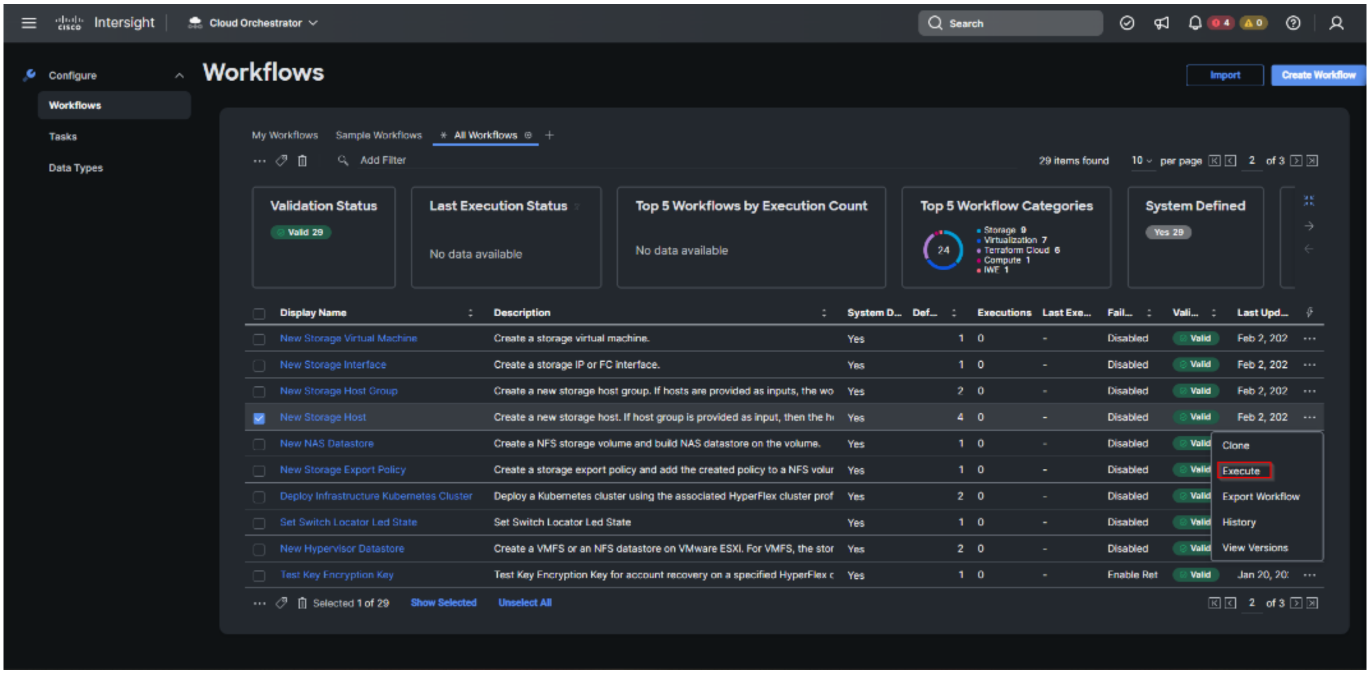

Workflows

Validation Status	Last Execution Status	Top 5 Workflows by Execution Count	Top 5 Workflow Categories	System Defined
Valid 29	No data available	No data available	24 Storage 9 Virtualization 7 Terraform Cloud 6 Compute 1 IWE 1	Yes 29

	Display Name	Description	System D...	Def...	Executions	Last Exe...	Fail...	Vali...	Last Upd...
□	New Storage Virtual Machine	Create a storage virtual machine.	Yes	1	0	-	Disabled	Valid	Feb 2, 202
□	New Storage Interface	Create a storage IP or FC interface.	Yes	1	0	-	Disabled	Valid	Feb 2, 202
□	New Storage Host Group	Create a new storage host group. If hosts are provided as inputs, the wo	Yes	2	0	-	Disabled	Valid	Feb 2, 202
☑	New Storage Host	Create a new storage host. If host group is provided as input, then the h	Yes	4	0	-	Disabled	Valid	Feb 2, 202
□	New NAS Datastore	Create a NFS storage volume and build NAS datastore on the volume.	Yes	1	0	-	Disabled	Valid	
□	New Storage Export Policy	Create a storage export policy and add the created policy to a NFS volur	Yes	1	0	-	Disabled	Valid	
□	Deploy Infrastructure Kubernetes Cluster	Deploy a Kubernetes cluster using the associated HyperFlex cluster prof	Yes	2	0	-	Disabled	Valid	
□	Set Switch Locator Led State	Set Switch Locator Led State	Yes	1	0	-	Disabled	Valid	
□	New Hypervisor Datastore	Create a VMFS or an NFS datastore on VMware ESXI. For VMFS, the stor	Yes	2	0	-	Disabled	Valid	
□	Test Key Encryption Key	Test Key Encryption Key for account recovery on a specified HyperFlex c	Yes	1	0	-	Enable Ret	Valid	Jan 20, 20:

Step 3. Select the appropriate Organization (default by default).

Step 4. Select the appropriate Pure Storage device.

Step 5. Enter the name of the Host name and WWNs for ESX host. Click Execute.

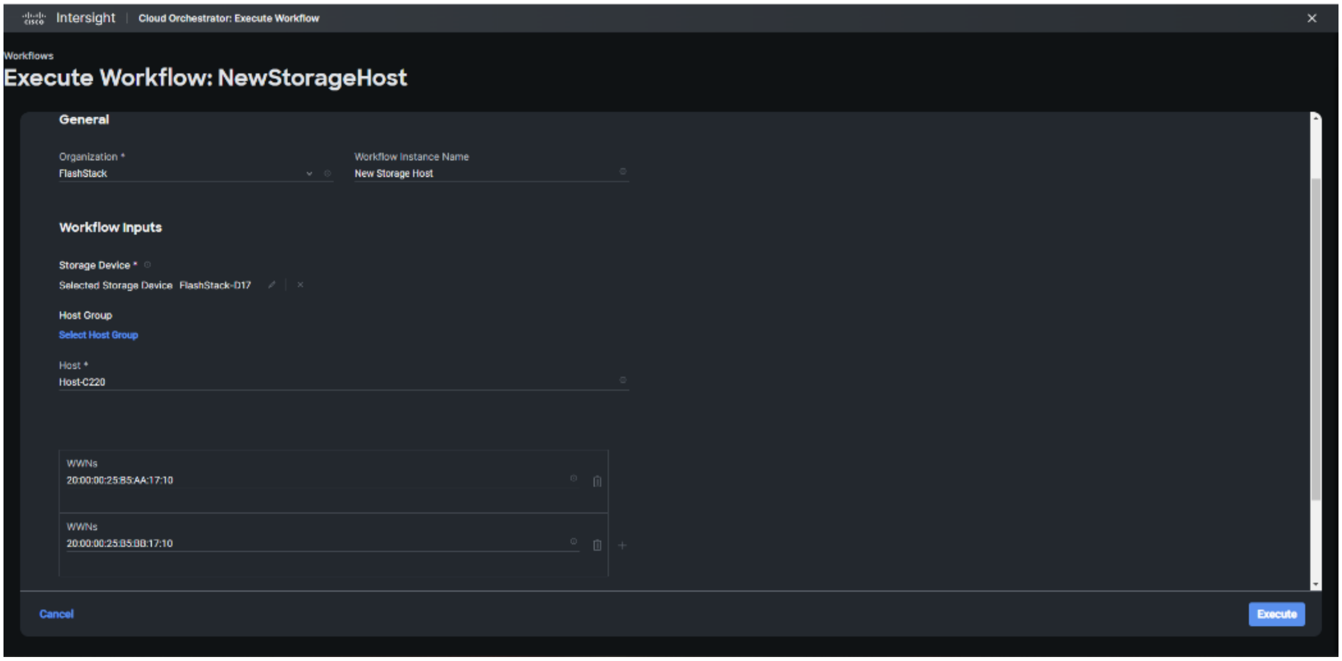

Execute Workflow: NewStorageHost

General

Organization *: FlashStack

Workflow Instance Name: New Storage Host

Workflow Inputs

Storage Device *

Selected Storage Device FlashStack-D17

Host Group

Select Host Group

Host *: Host-C220

WWNs: 20:00:00:25:B5:AA:17:10

WWNs: 20:00:00:25:B5:BB:17:10

The workflow can be monitored and rolled back.

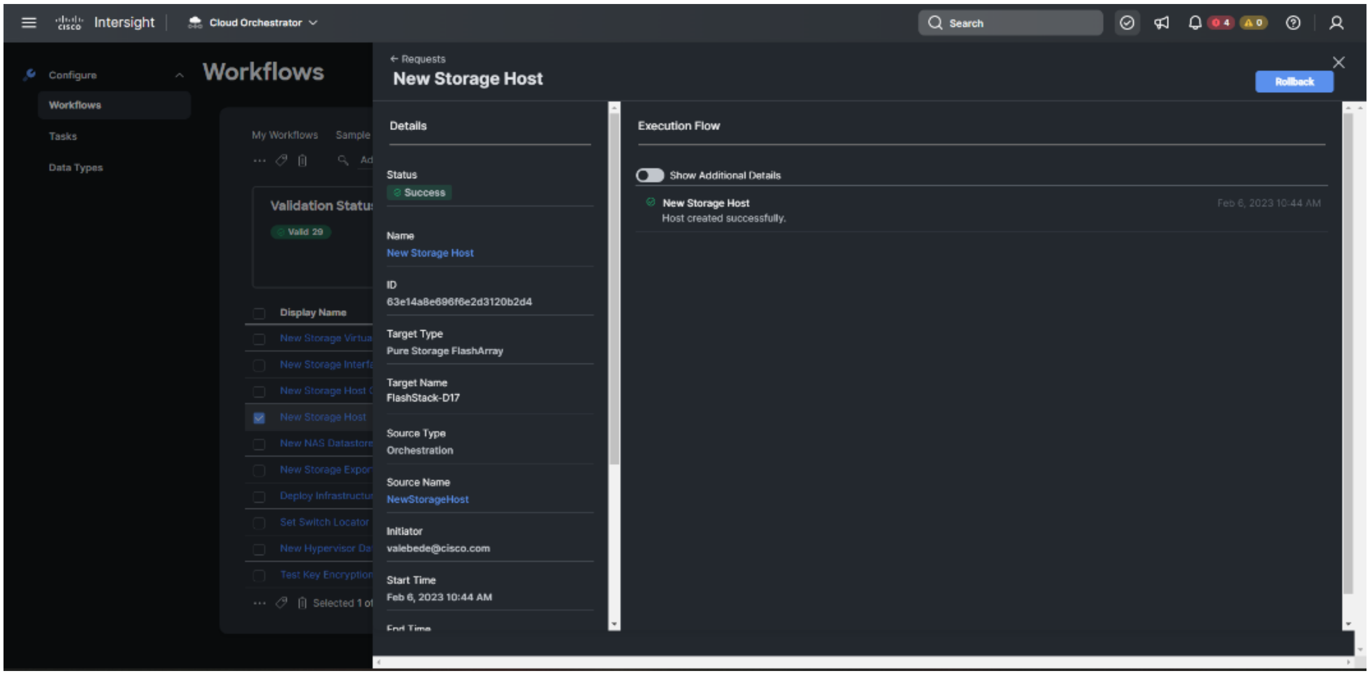

New Storage Host

Status
Success

Name
New Storage Host

ID
63e14a8e696f6e2d3120b2d4

Target Type
Pure Storage FlashArray

Target Name
FlashStack-D17

Source Type
Orchestration

Source Name
NewStorageHost

Initiator
valebede@cisco.com

Start Time
Feb 6, 2023 10:44 AM

End Time

New Storage Host
Host created successfully.
Feb 6, 2023 10:44 AM

Procedure 6. 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.

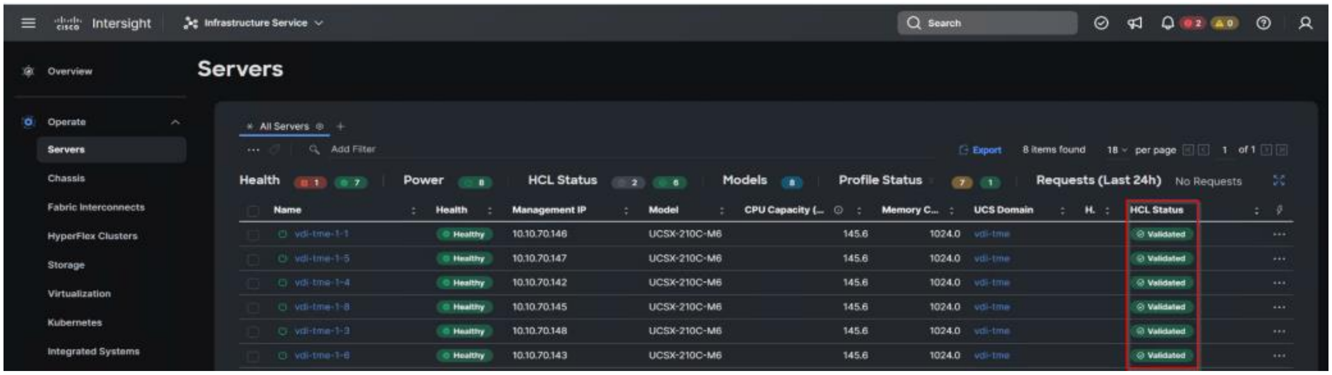

Servers

8 items found 18 per page 1 of 1

Name	Health	Management IP	Model	CPU Capacity (...	Memory C...	UCS Domain	H.	HCL Status
vdi-tme-1-1	Healthy	10.10.70.146	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-5	Healthy	10.10.70.147	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-4	Healthy	10.10.70.142	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-8	Healthy	10.10.70.145	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-3	Healthy	10.10.70.148	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated
vdi-tme-1-6	Healthy	10.10.70.143	UCSX-210C-M6	145.6	1024.0	vdi-tme		Validated

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

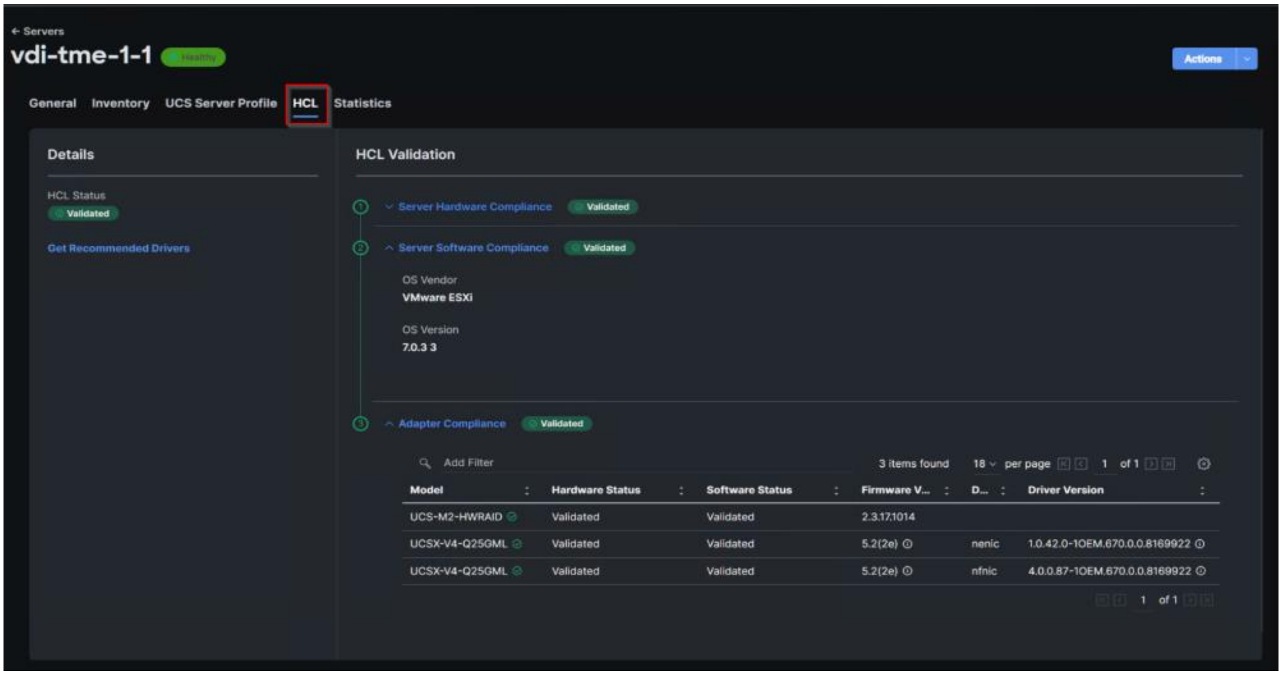

vdi-tme-1-1

Details

HCL Status
Validated

Get Recommended Drivers

HCL Validation

1 Server Hardware Compliance — Validated

2 Server Software Compliance — Validated

OS Vendor
VMware ESXi

OS Version
7.0.3 3

3 Adapter Compliance — Validated

Model	Hardware Status	Software Status	Firmware V...	D...	Driver Version
UCS-M2-HWRAID	Validated	Validated	2.3.17.1014		
UCSX-V4-Q25GML	Validated	Validated	5.2(2e)	nenic	1.0.42.0-1OEM.670.0.0.8169922
UCSX-V4-Q25GML	Validated	Validated	5.2(2e)	nfnic	4.0.0.87-1OEM.670.0.0.8169922

Pure Storage CloudSnap

Configure Pure Storage CloudSnap

Procedure 1. Create an S3 Bucket in the Customer's AWS Account

Step 1. Log in to the AWS management console and go to S3. From the AWS S3 console dashboard, select Create Bucket.

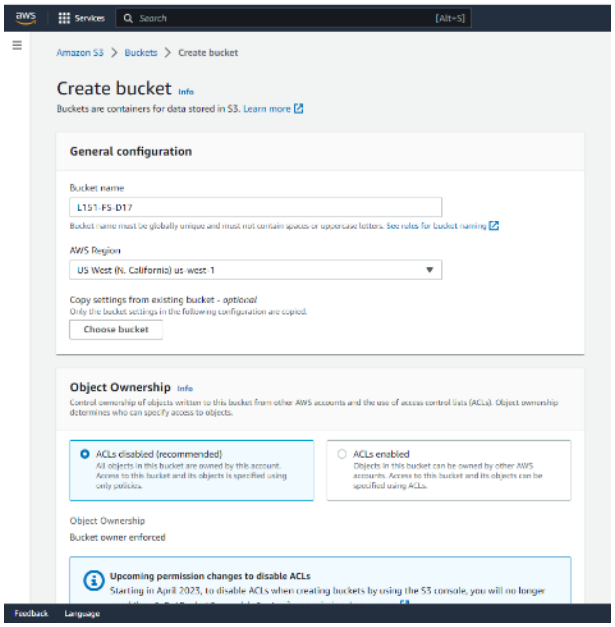

Create bucket

Buckets are containers for data stored in S3. Learn more

General configuration

Bucket name
L151-FS-D17
Bucket name must be globally unique and must not contain spaces or uppercase letters. See rules for bucket naming

AWS Region
US West (N. California) us-west-1

Copy settings from existing bucket - *optional*
Only the bucket settings in the following configuration are copied.

Object Ownership

Control ownership of objects written to this bucket from other AWS accounts and the use of access control lists (ACLs). Object ownership determines who can specify access to objects.

● ACLs disabled (recommended)
All objects in this bucket are owned by this account. Access to this bucket and its objects is specified using only policies.

○ ACLs enabled
Objects in this bucket can be owned by other AWS accounts. Access to this bucket and its objects can be specified using ACLs.

Object Ownership
Bucket owner enforced

Upcoming permission changes to disable ACLs
Starting in April 2023, to disable ACLs when creating buckets by using the S3 console, you will no longer

Step 2. Select Default encryption.

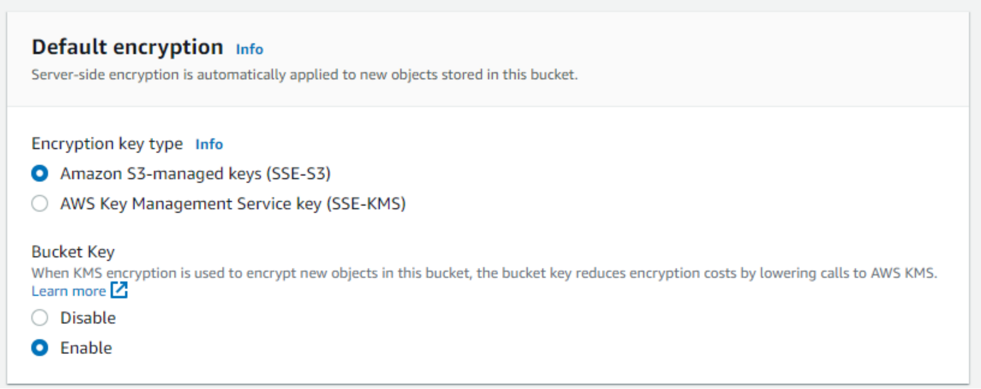

Default encryption

Info

Server-side encryption is automatically applied to new objects stored in this bucket.

Encryption key type Info

● Amazon S3-managed keys (SSE-S3)

○ AWS Key Management Service key (SSE-KMS)

Bucket Key

When KMS encryption is used to encrypt new objects in this bucket, the bucket key reduces encryption costs by lowering calls to AWS KMS. Learn more

○ Disable

● Enable

Step 3. Click Create bucket.

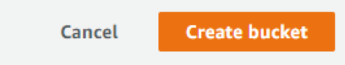

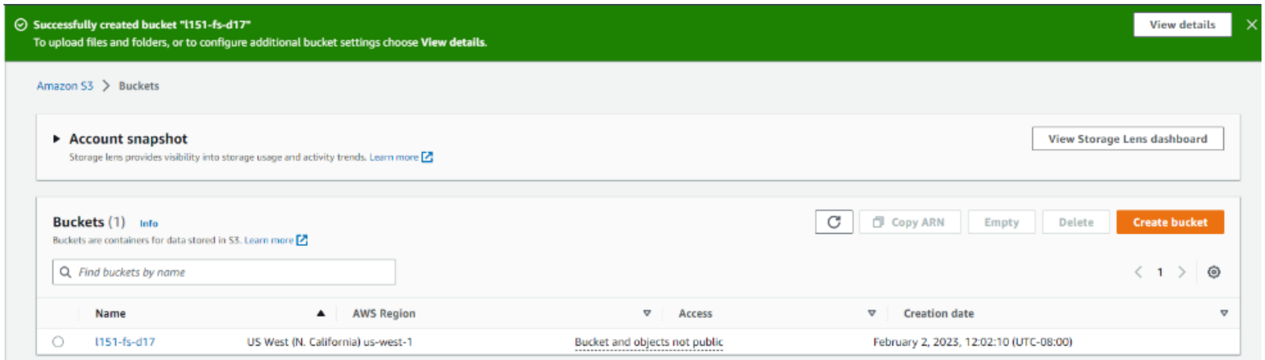

Successfully created bucket "l151-fs-d17"
To upload files and folders, or to configure additional bucket settings choose **View details**.

Account snapshot

Storage lens provides visibility into storage usage and activity trends. Learn more

Buckets (1) Info

Buckets are containers for data stored in S3. Learn more

	Name	AWS Region	Access	Creation date
○	l151-fs-d17	US West (N. California) us-west-1	Bucket and objects not public	February 2, 2023, 12:02:10 (UTC-08:00)

Step 4. From the AWS IAM console click Add users.

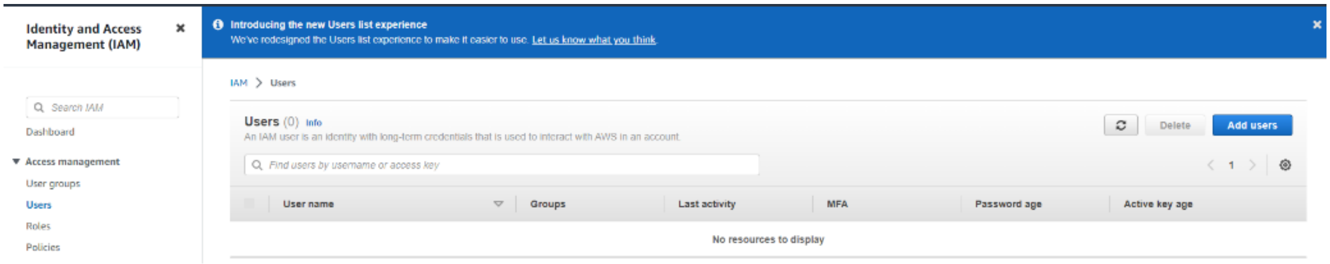

Introducing the new Users list experience
We've redesigned the Users list experience to make it easier to use. Let us know what you think.

Users (0) Info

An IAM user is an identity with long-term credentials that is used to interact with AWS in an account.

User name	Groups	Last activity	MFA	Password age	Active key age
No resources to display					

Step 5. Provide a user name and click Install.

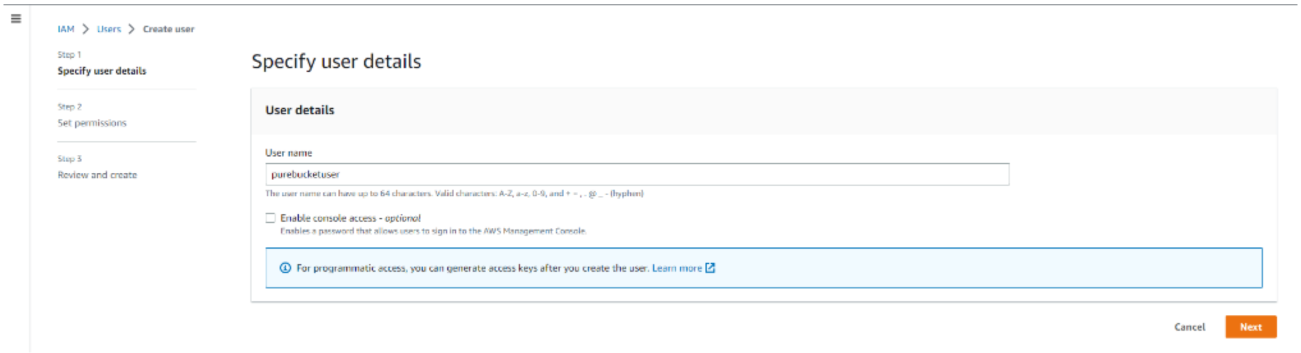

Specify user details

User details

User name: purebucketuser

The user name can have up to 64 characters. Valid characters: A-Z, a-z, 0-9, and + = , . @ _ - (hyphen)

□ Enable console access - optional
Enables a password that allows users to sign in to the AWS Management Console.

For programmatic access, you can generate access keys after you create the user. Learn more

Step 6. Attach the appropriate access policy and click Next.

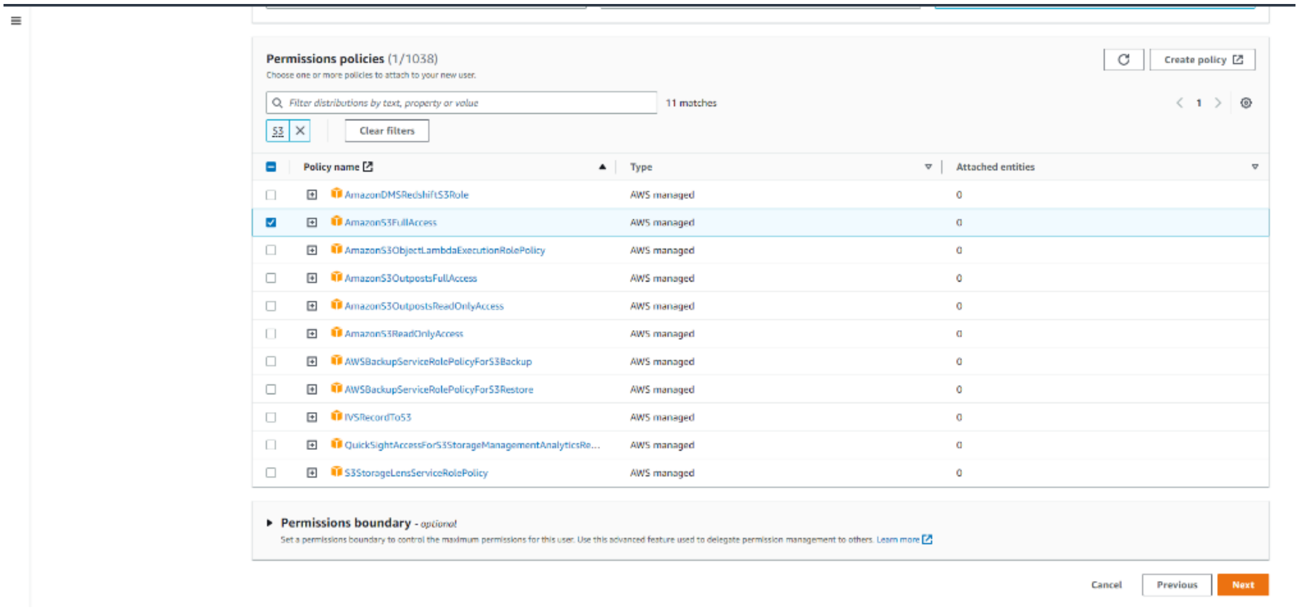

Permissions policies (1/1038)

Choose one or more policies to attach to your new user.

Filter distributions by text, property or value — 11 matches

☐	Policy name	Type	Attached entities
☐	AmazonDMSRedshiftS3Role	AWS managed	0
☑	AmazonS3FullAccess	AWS managed	0
☐	AmazonS3ObjectLambdaExecutionRolePolicy	AWS managed	0
☐	AmazonS3OutpostsFullAccess	AWS managed	0
☐	AmazonS3OutpostsReadOnlyAccess	AWS managed	0
☐	AmazonS3ReadOnlyAccess	AWS managed	0
☐	AWSBackupServiceRolePolicyForS3Backup	AWS managed	0
☐	AWSBackupServiceRolePolicyForS3Restore	AWS managed	0
☐	IVSRecordToS3	AWS managed	0
☐	QuickSightAccessForS3StorageManagementAnalyticsRe...	AWS managed	0
☐	S3StorageLensServiceRolePolicy	AWS managed	0

Permissions boundary - *optional*

Set a permissions boundary to control the maximum permissions for this user. Use this advanced feature used to delegate permission management to others. Learn more

Note: OPTION 1, THE SIMPLE METHOD – USING AN EXISTING AWS MANAGED POLICY

The easier option is to use AWS's pre-configured policy called AmazonS3FullAccess. This AWS policy grants the IAM user (Pure FlashArray) full access to all S3 buckets in the customer's AWS account.

OPTION 2, THE MORE RESTRICTIVE METHOD – CREATING A CUSTOMER MANAGED POLICY

This option is for users who want to create a Customer managed policy which would allow the IAM user (Pure FlashArray) full access to only the specific S3 bucket that will be used by CloudSnap to store offloaded data from FlashArray.

Step 7. Review the information and click Create user.

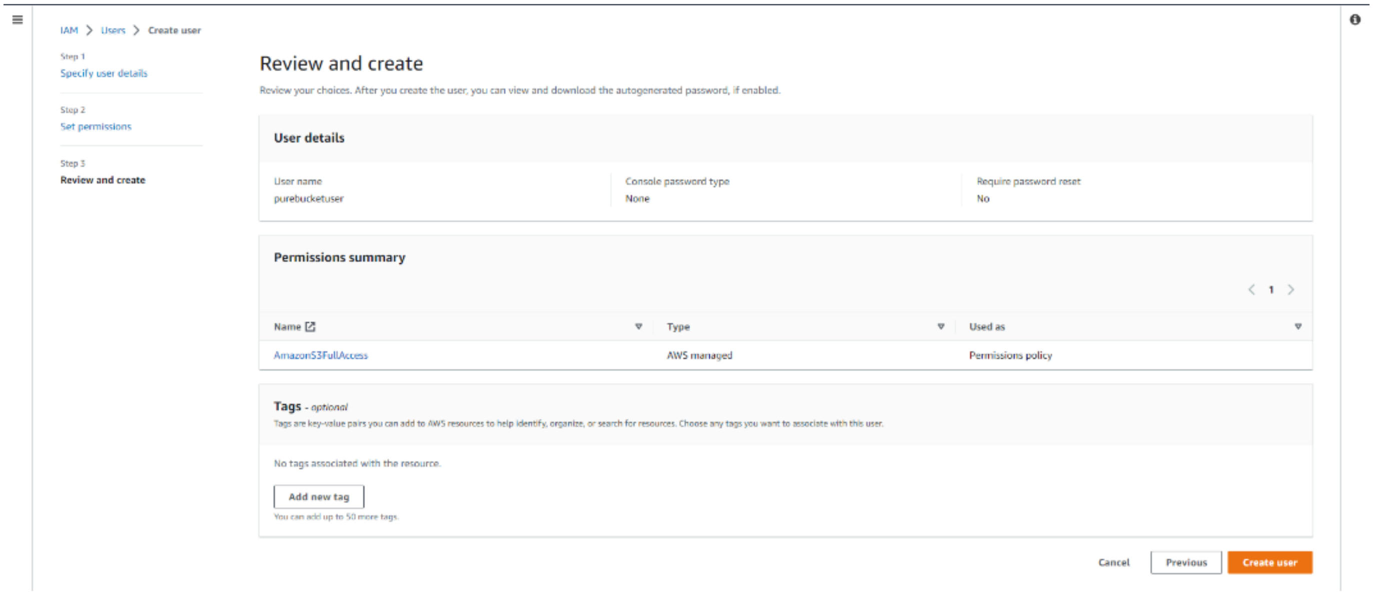

Review and create

Review your choices. After you create the user, you can view and download the autogenerated password, if enabled.

User details

User name	Console password type	Require password reset
purebucketuser	None	No

Permissions summary

Name	Type	Used as
AmazonS3FullAccess	AWS managed	Permissions policy

Tags - *optional*

Tags are key-value pairs you can add to AWS resources to help identify, organize, or search for resources. Choose any tags you want to associate with this user.

No tags associated with the resource.

You can add up to 50 more tags.

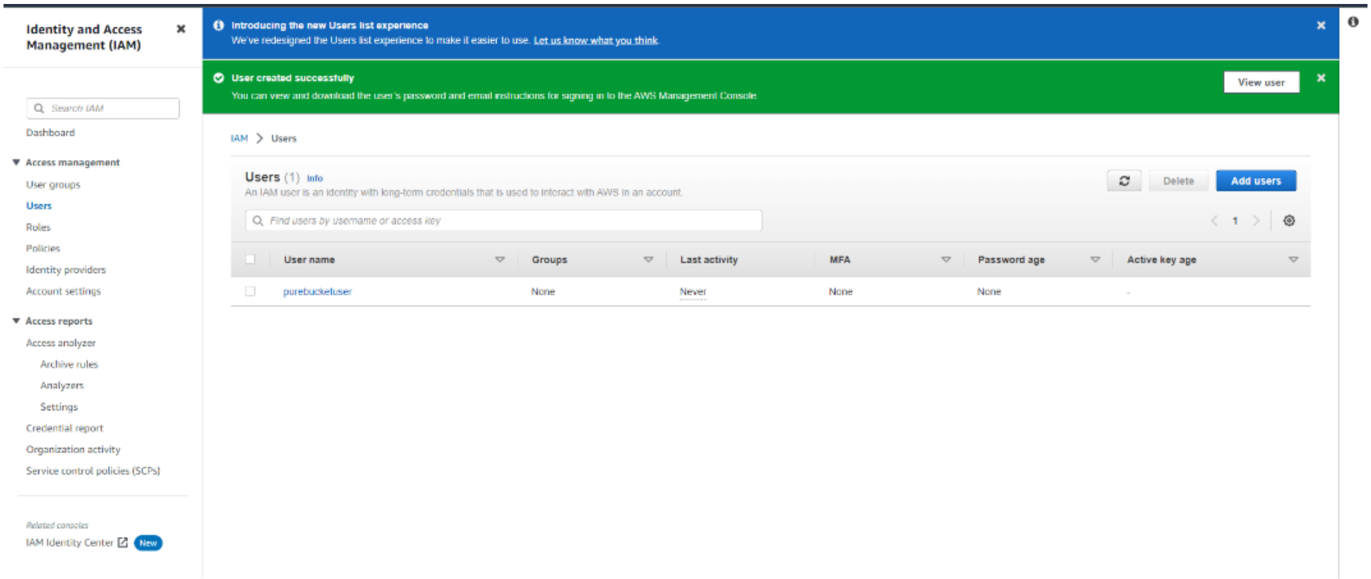

Introducing the new Users list experience
We've redesigned the Users list experience to make it easier to use. Let us know what you think.

User created successfully
You can view and download the user's password and email instructions for signing in to the AWS Management Console

Users (1) Info
An IAM user is an identity with long-term credentials that is used to interact with AWS in an account.

User name	Groups	Last activity	MFA	Password age	Active key age
purebucketuser	None	Never	None	None	-

Step 8. Create a key for the array access.

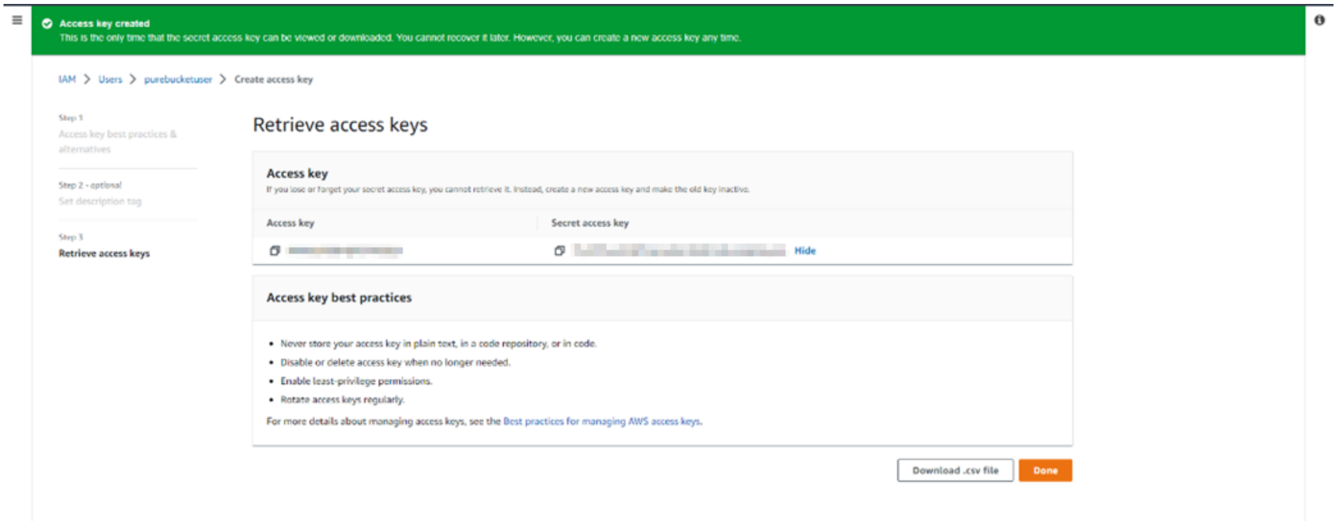

Access key created
This is the only time that the secret access key can be viewed or downloaded. You cannot recover it later. However, you can create a new access key any time.

IAM > Users > purebucketuser > Create access key
Step 1
Access key best practices & alternatives
Step 2 - optional
Set description tag
Step 3
Retrieve access keys

Retrieve access keys

Access key
If you lose or forget your secret access key, you cannot retrieve it. Instead, create a new access key and make the old key inactive.

Access key	Secret access key
[redacted]	[redacted] Hide

Access key best practices

- Never store your access key in plain text, in a code repository, or in code.
- Disable or delete access key when no longer needed.
- Enable least-privilege permissions.
- Rotate access keys regularly.

For more details about managing access keys, see the Best practices for managing AWS access keys.

Procedure 2. Configure Offload on the FlashArray//x R3

Step 1. From the Pure Storage FlashArray Management interface, go to Settings > Software > App Catalog. Select Offload and click Install.

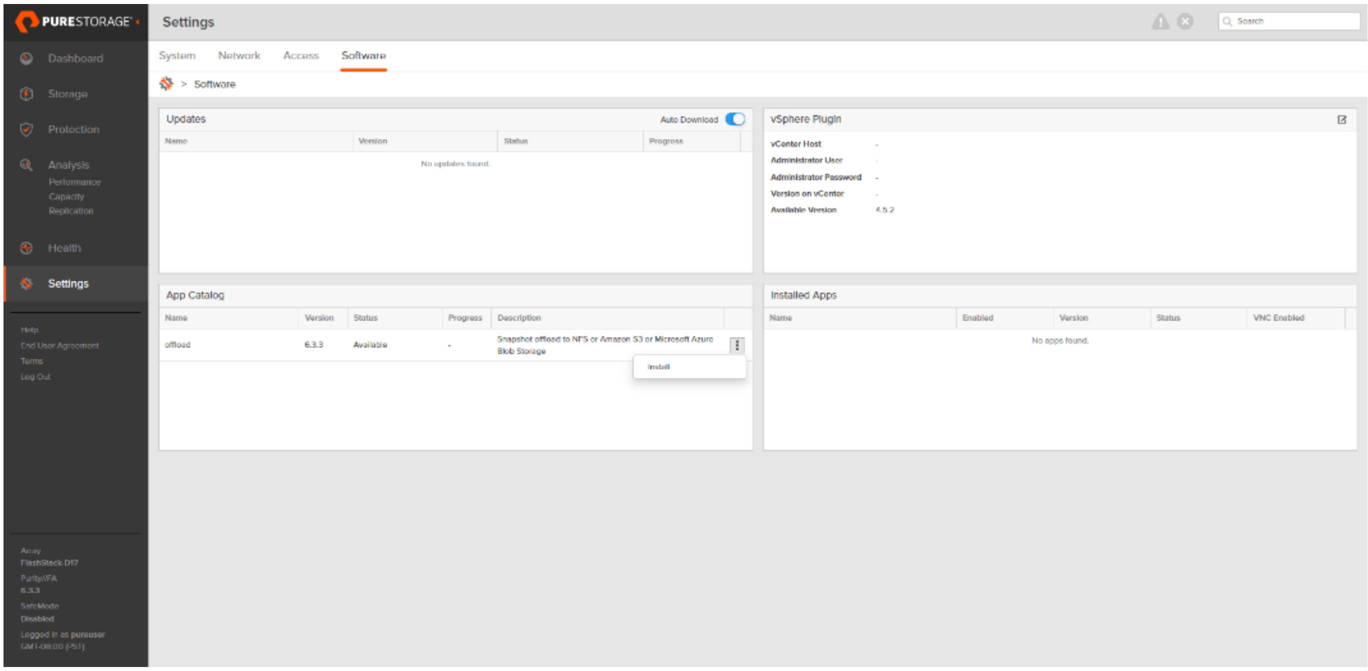

Updates

Name	Version	Status	Progress
No updates found.			

vSphere Plugin

vCenter Host	-
Administrator User	-
Administrator Password	-
Version on vCenter	-
Available Version	4.5.2

App Catalog

Name	Version	Status	Progress	Description
offload	6.3.3	Available	-	Snapshot offload to NFS or Amazon S3 or Microsoft Azure Blob Storage

Installed Apps

Name	Enabled	Version	Status	VNC Enabled
No apps found.				

Step 2. Wait for the application to finish installation.

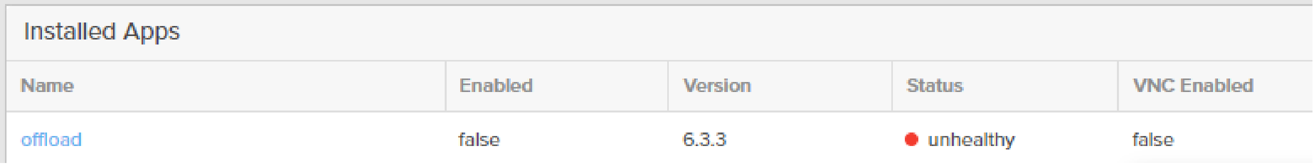

Installed Apps

Name	Enabled	Version	Status	VNC Enabled
offload	false	6.3.3	● unhealthy	false

Step 3. In the ssh array session, create the virtual offload interface.

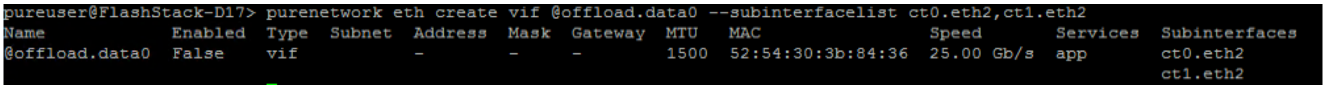

```
pureuser@FlashStack-D17> purenetwork eth create vif @offload.data0 --subinterfacelist ct0.eth2,ct1.eth2
Name           Enabled  Type  Subnet  Address  Mask  Gateway  MTU   MAC                Speed       Services  Subinterfaces
@offload.data0 False    vif           -        -     -        1500  52:54:30:3b:84:36  25.00 Gb/s  app       ct0.eth2
                                                                                                             ct1.eth2
```

Step 4. Configure the offload interface with the appropriate customer environment IP information.

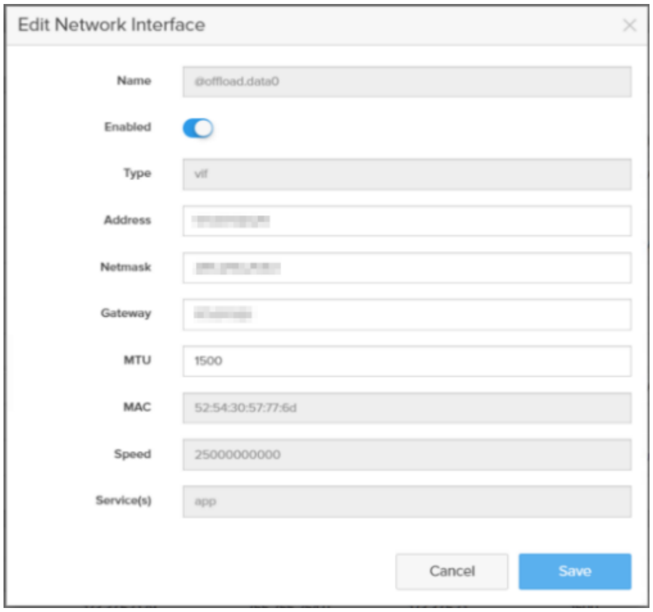

Edit Network Interface

Name	@offload.data0
Enabled	●
Type	vif
Address	[redacted]
Netmask	[redacted]
Gateway	[redacted]
MTU	1500
MAC	52:54:30:57:77:6d
Speed	25000000000
Service(s)	app

Step 5. From the Pure Storage FlashArray Management interface, go to Settings > Software > Installed Apps. Select Offload and click Enable App.

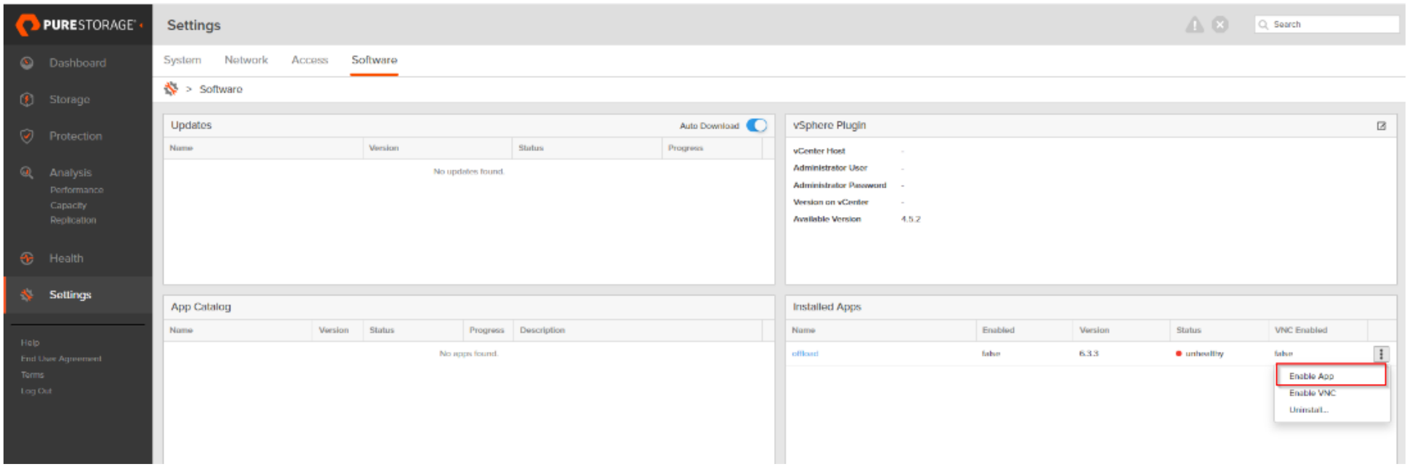

Settings > Software

Updates

Name	Version	Status	Progress

No updates found.

vSphere Plugin

vCenter Host	-
Administrator User	-
Administrator Password	-
Version on vCenter	-
Available Version	4.5.2

App Catalog

Name	Version	Status	Progress	Description

No apps found.

Installed Apps

Name	Enabled	Version	Status	VNC Enabled
offload	false	6.3.3	● unhealthy	false

Step 6. From the Pure Storage FlashArray Management interface, go to Storage > Array. Next to Offload Targets, click +.

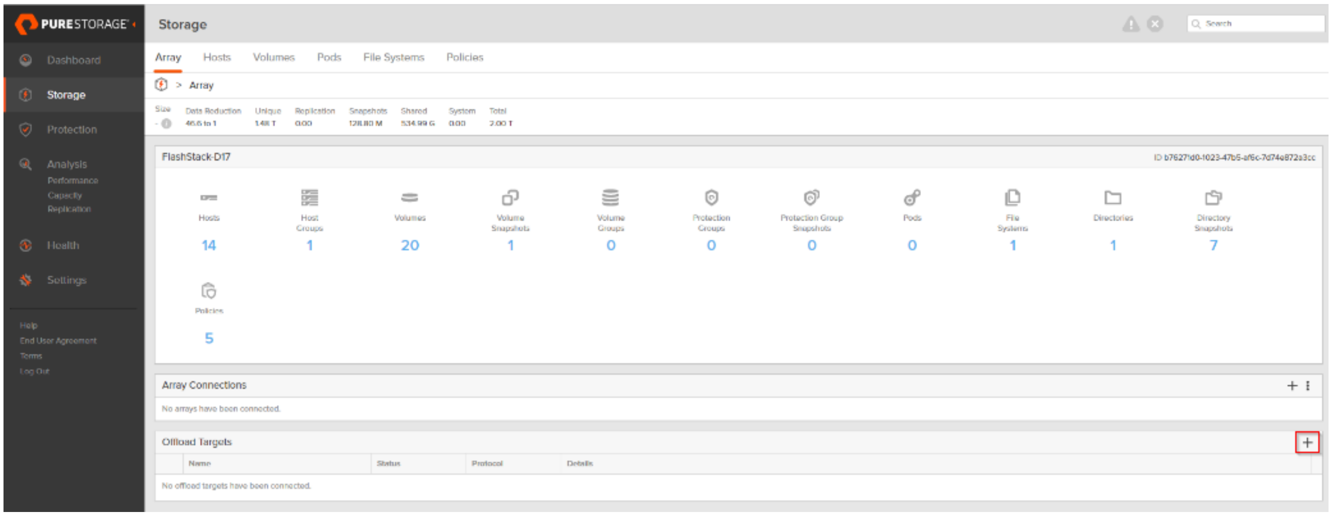

Array

Size	Data Reduction	Unique	Replication	Snapshots	Shared	System	Total
-	46.6 to 1	1.48 T	0.00	128.80 M	534.99 G	0.00	2.00 T

FlashStack-D17 ID b76271d0-1023-47b5-af6c-7d74e872a3cc

Hosts	Host Groups	Volumes	Volume Snapshots	Volume Groups	Protection Groups	Protection Group Snapshots	Pods	File Systems	Directories	Directory Snapshots	Policies
14	1	20	1	0	0	0	0	1	1	7	5

Array Connections

No arrays have been connected.

Offload Targets

Name	Status	Protocol	Details

No offload targets have been connected.

Step 7. Provide the Connection details and click Connect.

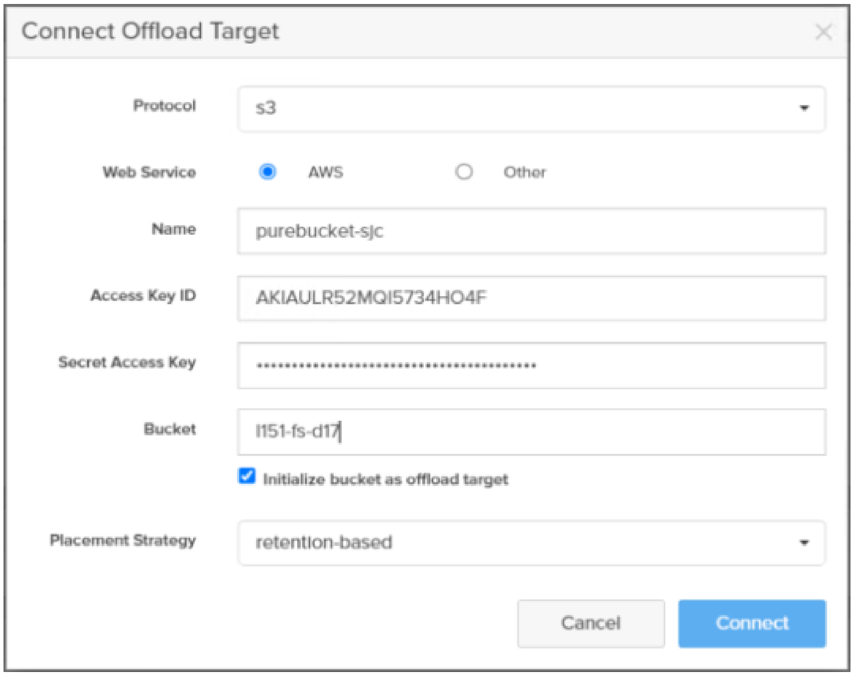

Connect Offload Target

Protocol: s3

Web Service: ● AWS ○ Other

Name: purebucket-sjc

Access Key ID: AKIAULR52MQI5734HO4F

Secret Access Key: **

Bucket: l151-fs-d17

☑ Initialize bucket as offload target

Placement Strategy: retention-based

Step 8. Select the newly added target.

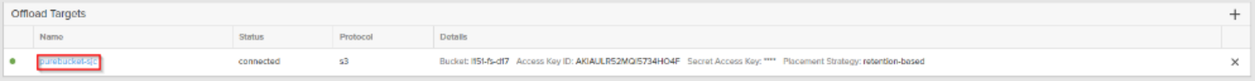

Offload Targets

	Name	Status	Protocol	Details
●	purebucket-sjc	connected	s3	Bucket: l151-fs-d17 Access Key ID: AKIAULR52MQI5734HO4F Secret Access Key: **** Placement Strategy: retention-based

Step 9. From the Pure Storage FlashArray Management interface, go to Protection > Protection Groups. Select Source Protection Groups and click Create.

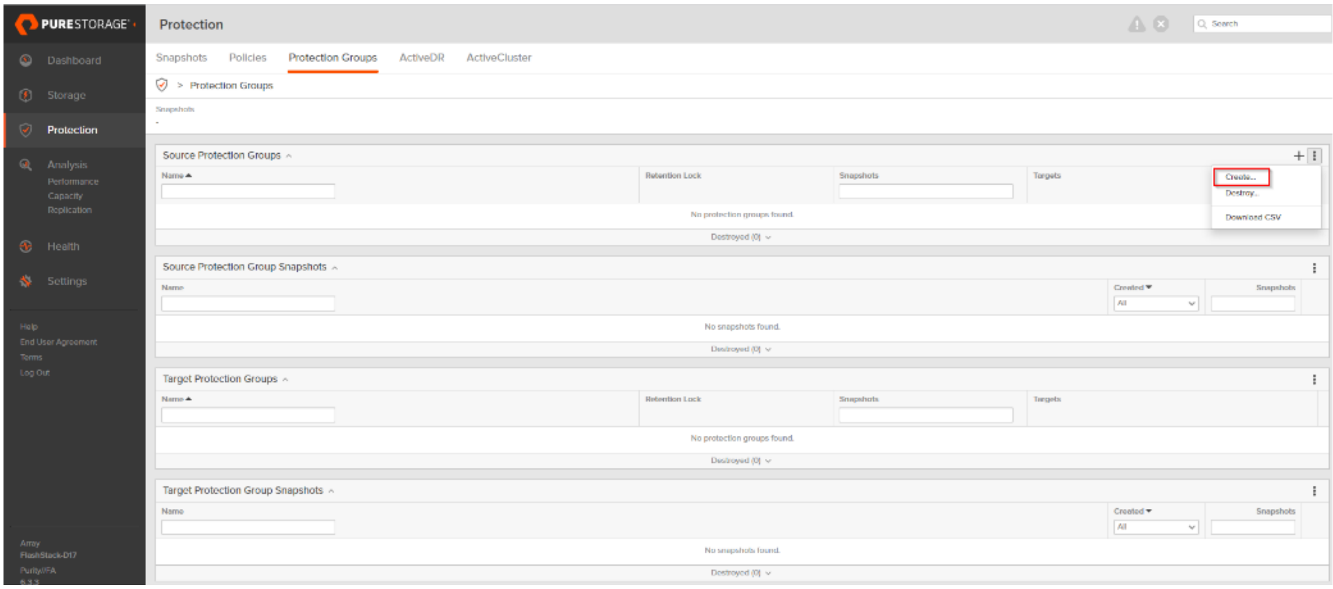

No protection groups found.

No snapshots found.

No protection groups found.

No snapshots found.

Step 10. Enter a Name and click Create.

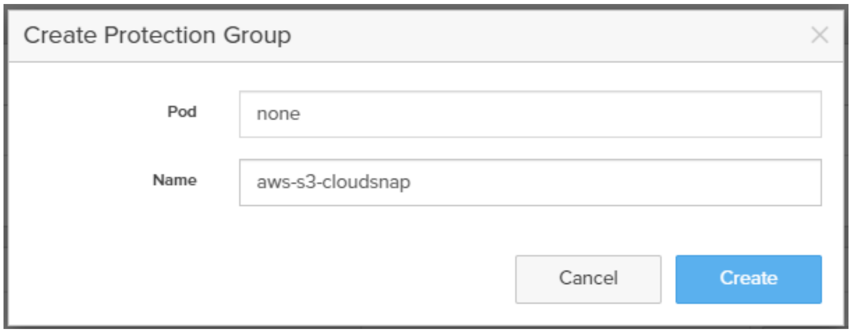

Create Protection Group

Pod: none

Name: aws-s3-cloudsnap

Step 11. From the Pure Storage FlashArray Management interface, go to Protection > Protection Groups. Select Target and click Add.

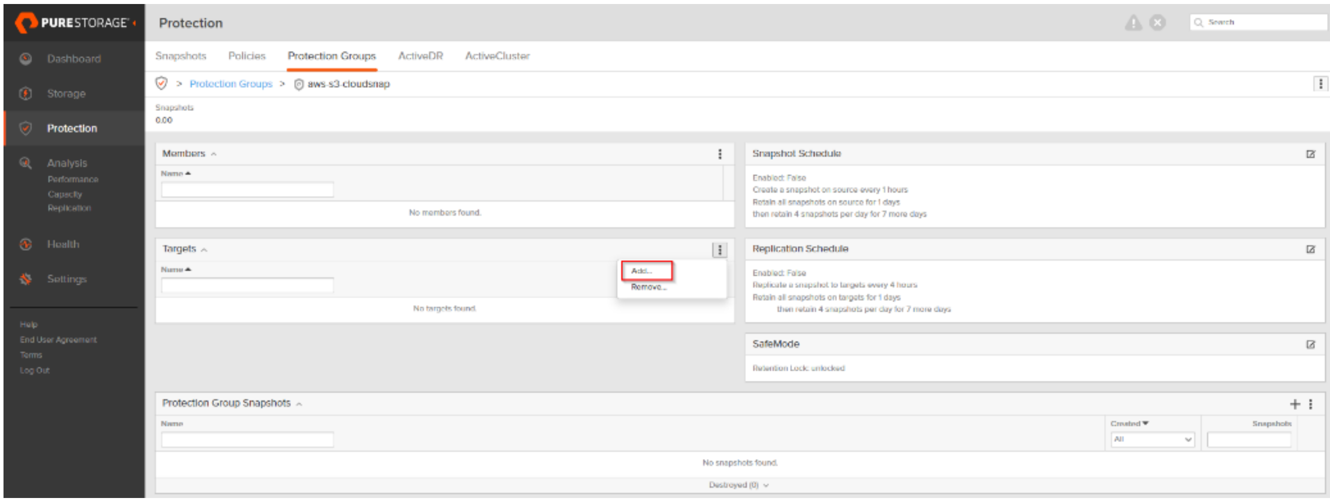

Snapshots
0.00

Members
Name
No members found.

Targets
Name
No targets found.

Snapshot Schedule
Enabled: False
Create a snapshot on source every 1 hours
Retain all snapshots on source for 1 days
then retain 4 snapshots per day for 7 more days

Replication Schedule
Enabled: False
Replicate a snapshot to targets every 4 hours
Retain all snapshots on targets for 1 days
then retain 4 snapshots per day for 7 more days

SafeMode
Retention Lock: unlocked

Protection Group Snapshots
Name | Created | Snapshots
All
No snapshots found.
Destroyed (0)

Step 12. Add the offload target.

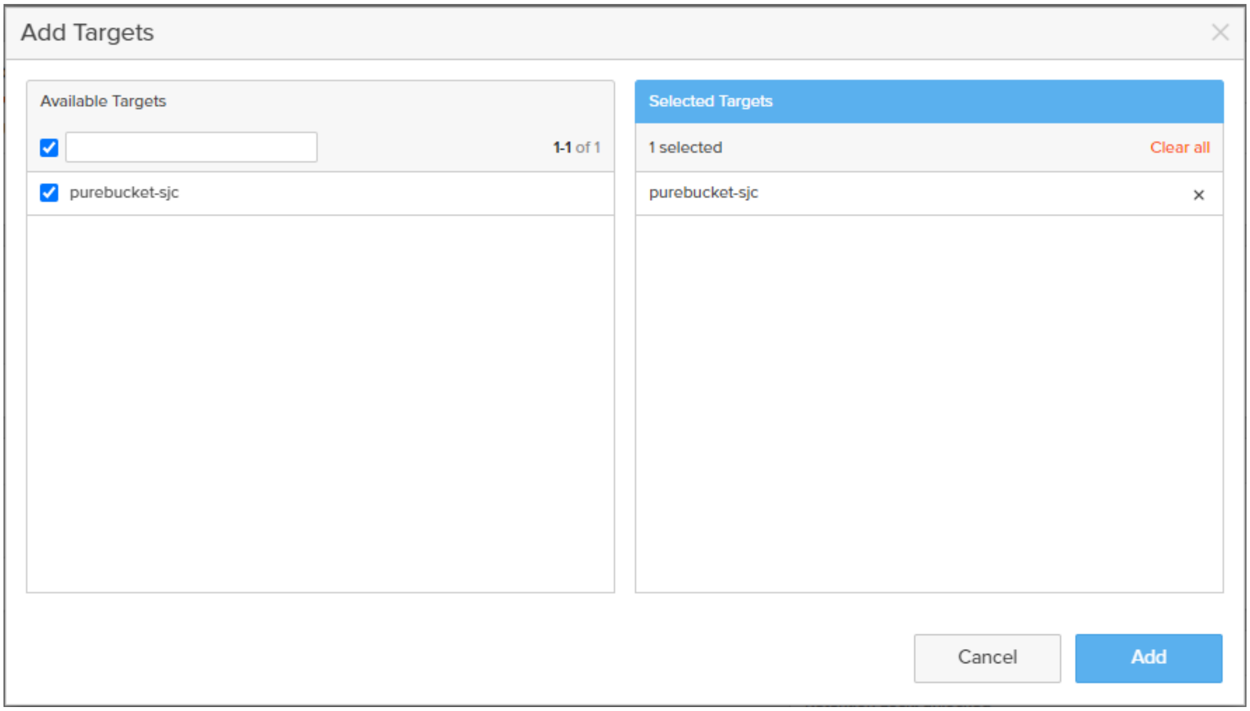

Add Targets

Available Targets
☑ 1-1 of 1
☑ purebucket-sjc

Selected Targets
1 selected
purebucket-sjc

Now your volumes are protected (such as Full Clones Desktops to AWS).

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 Citrix Virtual Apps and Desktops
- Install Citrix Virtual Apps and Desktops Delivery Controller, Citrix Licensing, and StoreFront
- Install and Configure Citrix Provisioning Server 2203
- Install the Citrix Provisioning Server Target Device Software
- Provision Virtual Desktop Machines
- Citrix Virtual Apps and Desktops Policies and Profile Management
- FSLogix for Citrix Virtual Apps & Desktops Profile Management

Prerequisites

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

Figure 29. **Example of the DHCP Scopes used in this CVD**

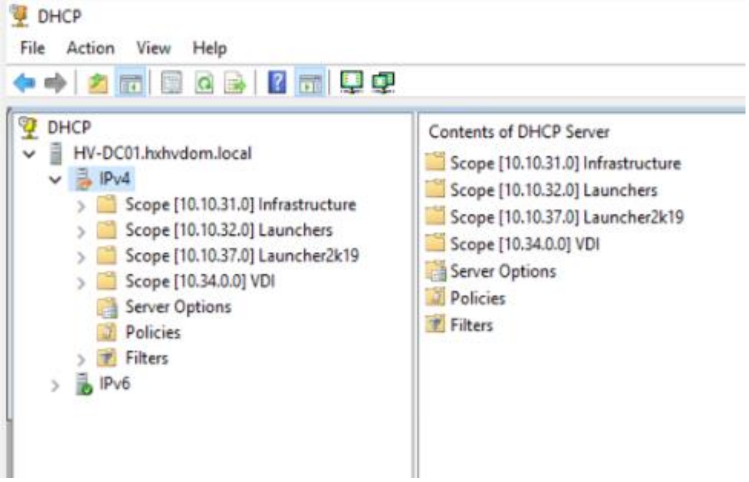

DHCP
HV-DC01.hxhvdom.local
IPv4
Scope [10.10.31.0] Infrastructure
Scope [10.10.32.0] Launchers
Scope [10.10.37.0] Launcher2k19
Scope [10.34.0.0] VDI
Server Options
Policies
Filters
IPv6

Contents of DHCP Server
Scope [10.10.31.0] Infrastructure
Scope [10.10.32.0] Launchers
Scope [10.10.37.0] Launcher2k19
Scope [10.34.0.0] VDI
Server Options
Policies
Filters

Software Infrastructure Configuration

This section explains how to configure the software infrastructure components that comprise this solution.

Install and configure the infrastructure virtual machines by following the process listed in Table 19.

Table 19. Test Infrastructure Virtual Machine Configuration

Configuration	Citrix Virtual Apps and Desktops Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Operating system	Microsoft Windows Server 2022	Microsoft Windows Server 2022
Virtual CPU amount	6	6

Configuration	Citrix Virtual Apps and Desktops Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Memory amount	12 GB	24 GB
Network	VMXNET3 Infra-Mgmt-31	VMXNET3 Infra-Mgmt-31
Disk-1 (OS) size	40 GB	40 GB
Disk-2 size	-	200 GB Disk Store

Configuration	Microsoft Active Directory DCs Virtual Machines	vCenter Server Appliance Virtual Machine
Operating system	Microsoft Windows Server 2022	VCSA – SUSE Linux
Virtual CPU amount	2	8
Memory amount	8 GB	32 GB
Network	VMXNET3 Infra-Mgmt-31	VMXNET3 IB-Mgmt-30
Disk size	40 GB	698.84 GB (across 13 VMDKs)

Configuration	Microsoft SQL Server Virtual Machine	Citrix StoreFront Controller Virtual Machine
Operating system	Microsoft Windows Server 2022 Microsoft SQL Server 2021	Microsoft Windows Server 2022
Virtual CPU amount	8	4
Memory amount	24GB	8 GB
Network	VMXNET3 Infra-Mgmt-31	VMXNET3 Infra-Mgmt-31
Disk-1 (OS) size	40 GB	40 GB
Disk-2 size		-

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 April 2023 for the Microsoft operating systems, SQL server and Microsoft Office 2021 were installed.

To prepare Single-session OS or Multi-session OS master virtual machine, there are three major steps: installing the PVS Target Device x64 software (if delivered with Citrix Provisioning Services), installing the Virtual Delivery Agents (VDAs), and installing application software.

Note: For this CVD, the images contain the basics needed to run the Login VSI workload.

The Single-session OS and Multi-session OS master target virtual machines were configured as detailed in Table 20.

Table 20. Single-session OS and Multi-session OS Virtual Machines Configurations

Configuration	Single-session OS Virtual Machines	Multi-session OS Virtual Machines
Operating system	Microsoft Windows 11 64-bit	Microsoft Windows Server 2022
Virtual CPU amount	2	8
Memory amount	4 GB	32 GB
Network	VMXNET3 DVS_VDI	VMXNET3 DVS_VDI
Citrix PVS vDisk size Citrix MCS Disk Size	48 GB (dynamic) 48 GB	90 GB (dynamic)
Write cache Disk size	6 GB	6 GB
Citrix PVS write cache RAM cache size	256 MB	1024 MB
Additional software used for testing	Microsoft Office 2021 Office Update applied Login VSI 4.1.40 Target Software (Knowledge Worker Workload)	Microsoft Office 2021 Office Update applied Login VSI 4.1.40 Target Software (Knowledge Worker Workload)
Additional Configuration	Configure DHCP Add to domain Install VMWare tool Install .Net 3.5 Activate Office Install VDA Agent Run PVS Imaging Wizard (For non-persistent Desktops only)	Configure DHCP Add to domain Install VMWare tool Install .Net 3.5 Activate Office Install VDA Agent

Install and Configure Citrix Virtual Apps and Desktops

This section explains the installation of the core components of the Citrix Virtual Apps and Desktops system. This CVD installs two Citrix Virtual Apps and Desktops Delivery Controllers to support both hosted shared desktops (HSD), non-persistent hosted virtual desktops (HVD), and persistent hosted virtual desktops (HVD).

Prerequisites

Citrix recommends that you use Secure HTTP (HTTPS) and a digital certificate to protect vSphere communications. Citrix recommends that you use a digital certificate issued by a certificate authority (CA) according to your organization's security policy. Otherwise, if the security policy allows, use the VMware-installed self-signed certificate.

Procedure 1. Install vCenter Server Self-Signed Certificate

Step 1. Add the FQDN of the computer running vCenter Server to the hosts file on that server, located at SystemRoot/
WINDOWS/system32/Drivers/etc/. This step is required only if the FQDN of the computer running vCenter Server is not already present in DNS.

Step 2. Open Internet Explorer and enter the address of the computer running vCenter Server (for example, https://FQDN as the URL).

Step 3. Accept the security warnings.

Step 4. Click the Certificate Error in the Security Status bar and select View certificates.

Step 5. Click Install certificate, select Local Machine, and then click Next.

Step 6. Select Place all certificates in the following store and then click Browse.

Step 7. Click Show physical stores.

Step 8. Click Trusted People.

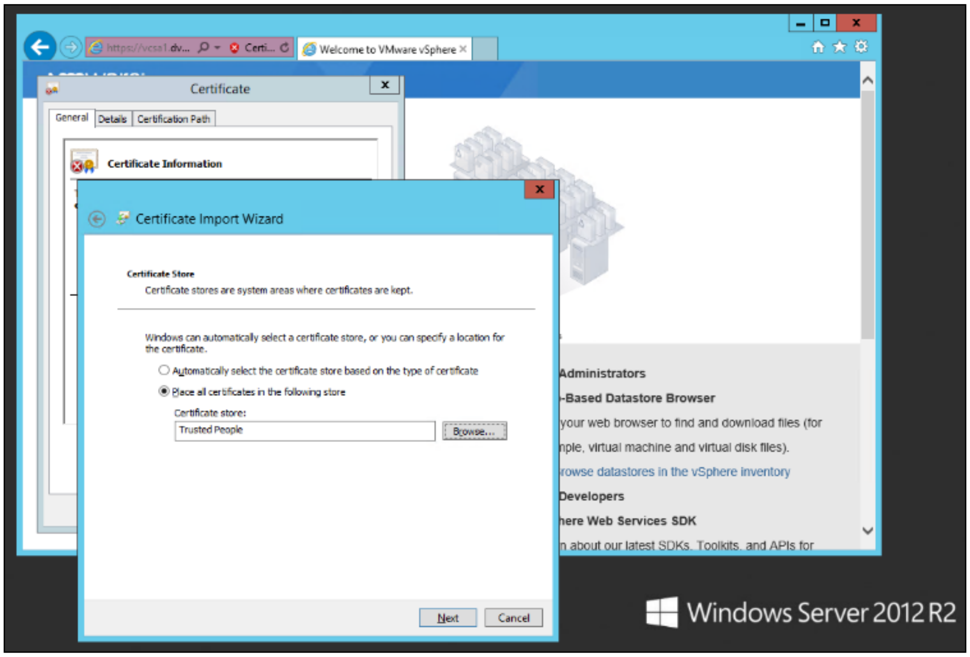

Certificate Import Wizard

Certificate Store

Certificate stores are system areas where certificates are kept.

Windows can automatically select a certificate store, or you can specify a location for the certificate.

○ Automatically select the certificate store based on the type of certificate

● Place all certificates in the following store

Certificate store: Trusted People

https://vcsa1.dv...
Certi...
Welcome to VMware vSphere
Certificate
General
Details
Certification Path
Certificate Information
Browse...
Next
Cancel
Administrators
-Based Datastore Browser
your web browser to find and download files (for
nple, virtual machine and virtual disk files).
rowse datastores in the vSphere inventory
Developers
here Web Services SDK
n about our latest SDKs, Toolkits, and APIs for
Windows Server 2012 R2

Step 9. Click Next and then click Finish.

Step 10. Repeat steps 1-9 on all Delivery Controllers and Provisioning Servers.

Install Citrix Virtual Apps and Desktops Delivery Controller, Citrix Licensing, and StoreFront

The process of installing the Citrix Virtual Apps and Desktops Delivery Controller also installs other key Citrix Virtual Apps and Desktops software components, including Studio, which is used to create and manage infrastructure components, and Director, which is used to monitor performance and troubleshoot problems.

Note: Dedicated StoreFront and License servers should be implemented for large scale deployments.

Procedure 1. Install Citrix License Server

Step 1. To begin the installation, connect to the first Citrix License server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2203 ISO.

Step 2. Click Start.

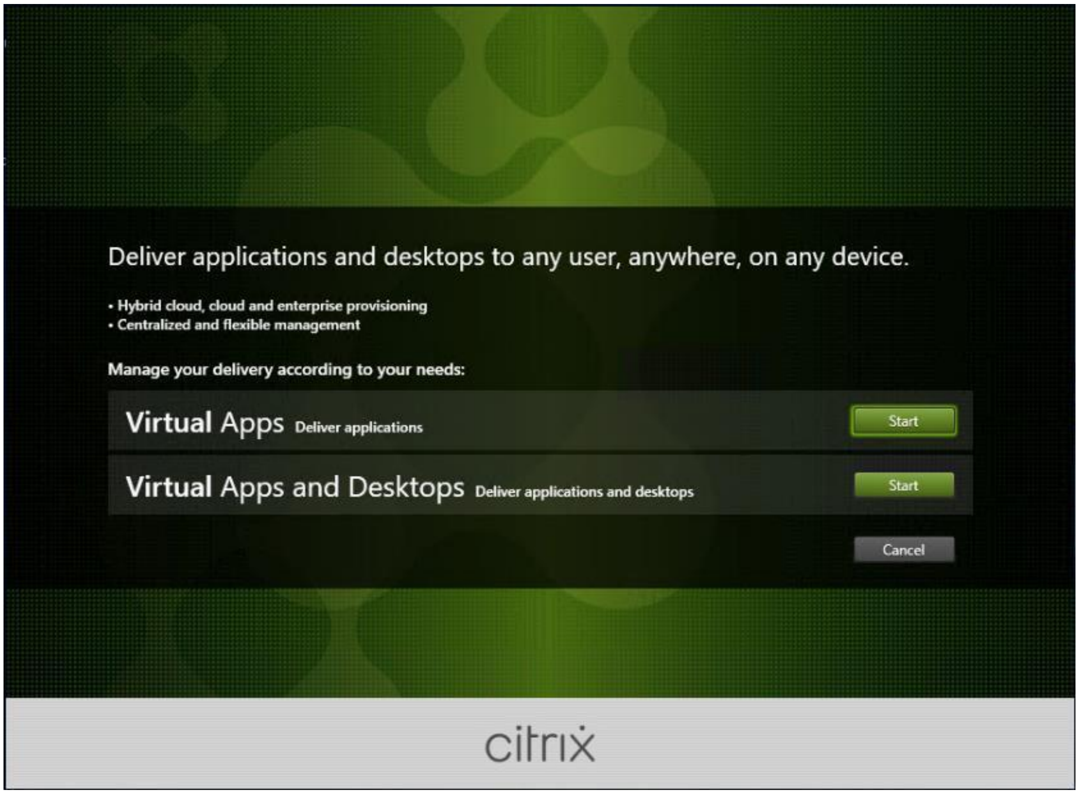

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

Virtual Apps Deliver applications
Start
Virtual Apps and Desktops Deliver applications and desktops
Start
Cancel
citrix

Step 3. Click Extend Deployment – Citrix License Server.

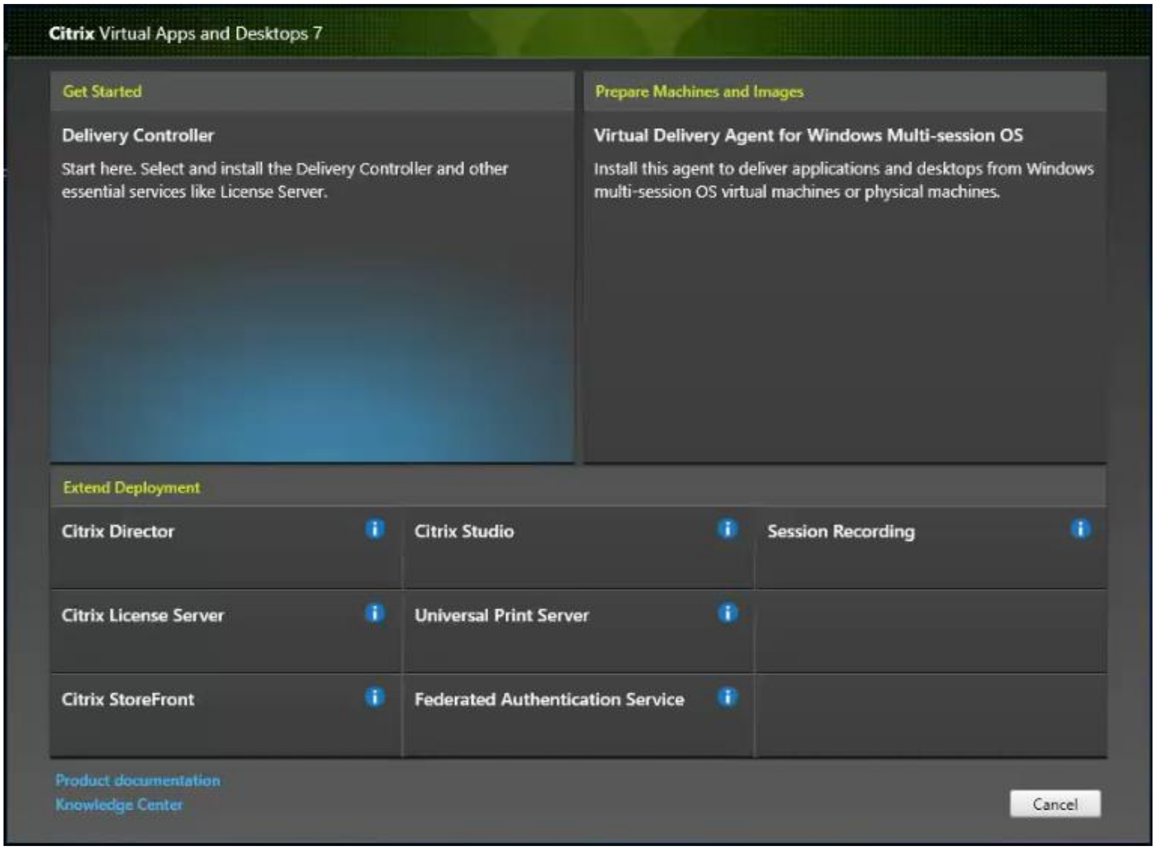

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server.

Virtual Delivery Agent for Windows Multi-session OS

Install this agent to deliver applications and desktops from Windows multi-session OS virtual machines or physical machines.

Step 4. Read the Citrix License Agreement. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

Step 5. Click Next.

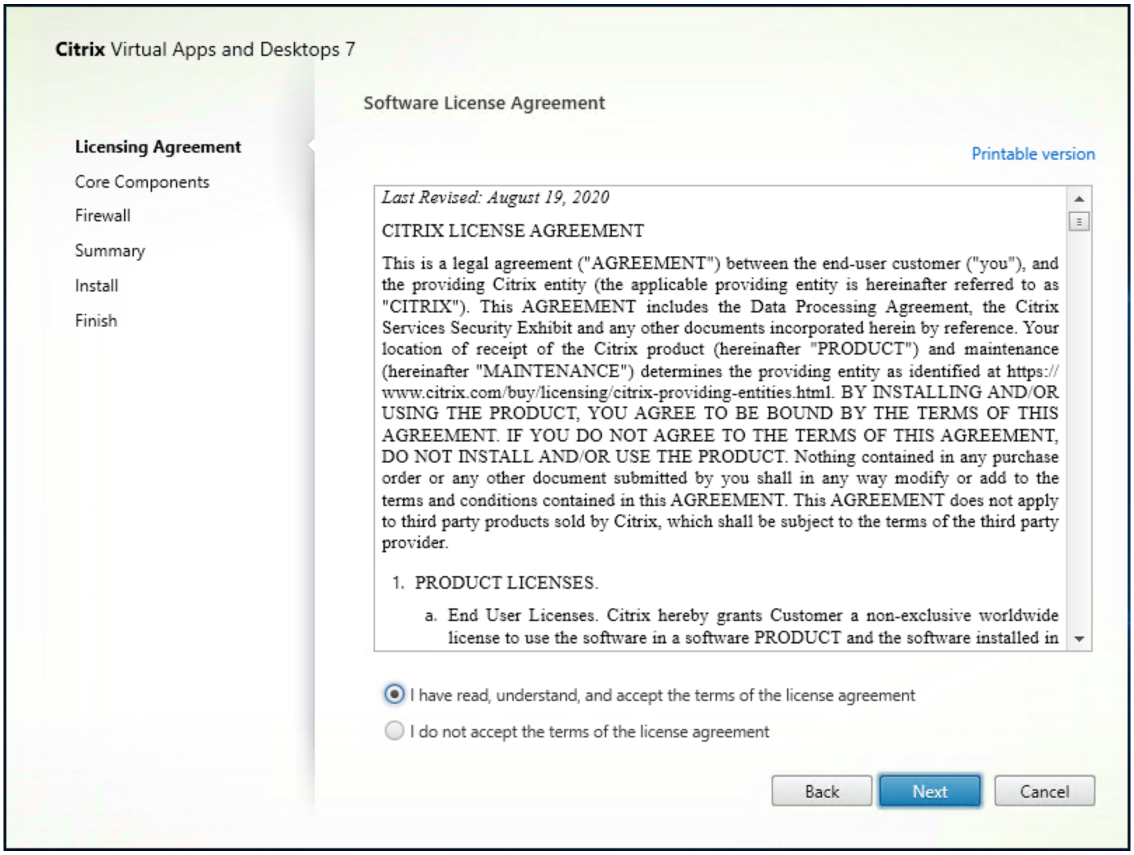

Software License Agreement

Last Revised: August 19, 2020

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). This AGREEMENT includes the Data Processing Agreement, the Citrix Services Security Exhibit and any other documents incorporated herein by reference. Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY INSTALLING AND/OR USING THE PRODUCT, YOU AGREE TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL AND/OR USE THE PRODUCT. Nothing contained in any purchase order or any other document submitted by you shall in any way modify or add to the terms and conditions contained in this AGREEMENT. This AGREEMENT does not apply to third party products sold by Citrix, which shall be subject to the terms of the third party provider.

1. PRODUCT LICENSES.
   a. End User Licenses. Citrix hereby grants Customer a non-exclusive worldwide license to use the software in a software PRODUCT and the software installed in

● I have read, understand, and accept the terms of the license agreement

○ I do not accept the terms of the license agreement

Step 6. Click Next.

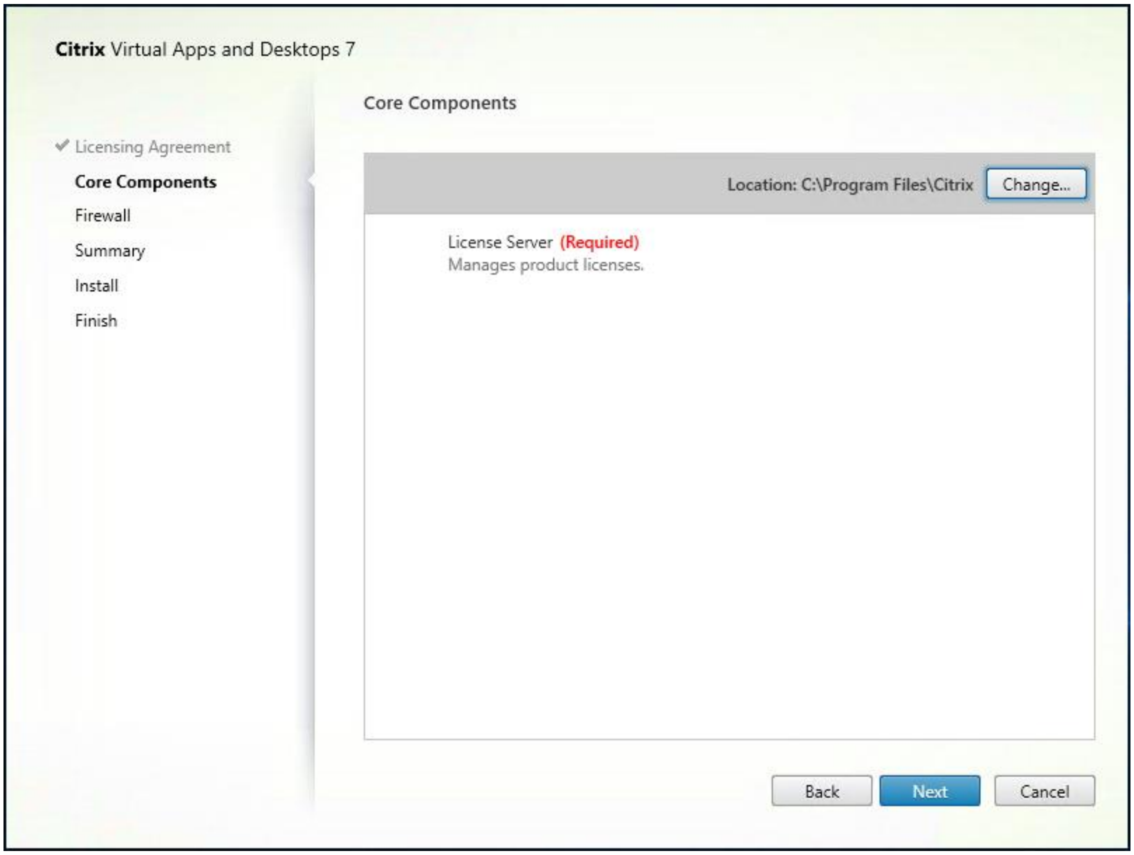

Core Components

Location: C:\Program Files\Citrix

License Server (Required)

Manages product licenses.

Step 7. Select the default ports and automatically configured firewall rules.

Step 8. Click Next.

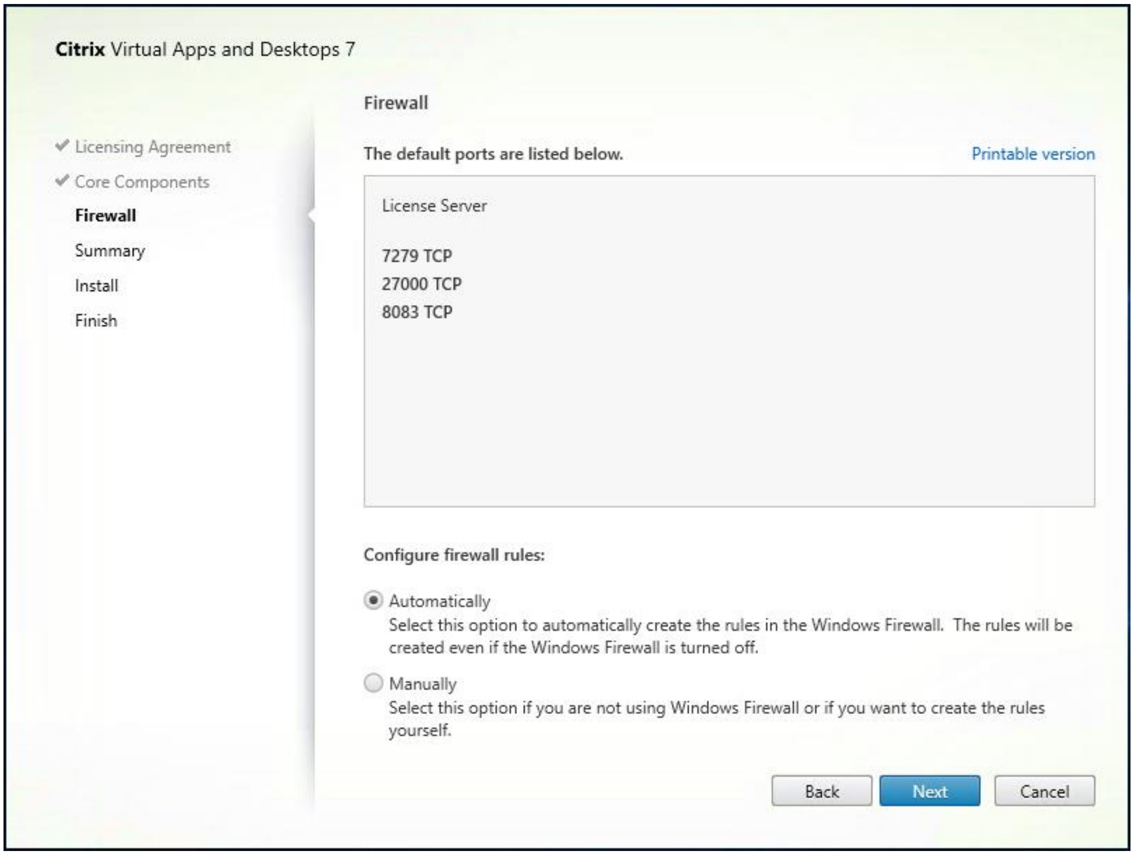

Citrix Virtual Apps and Desktops 7

Firewall

The default ports are listed below.

License Server

7279 TCP
27000 TCP
8083 TCP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

Step 9. Click Install.

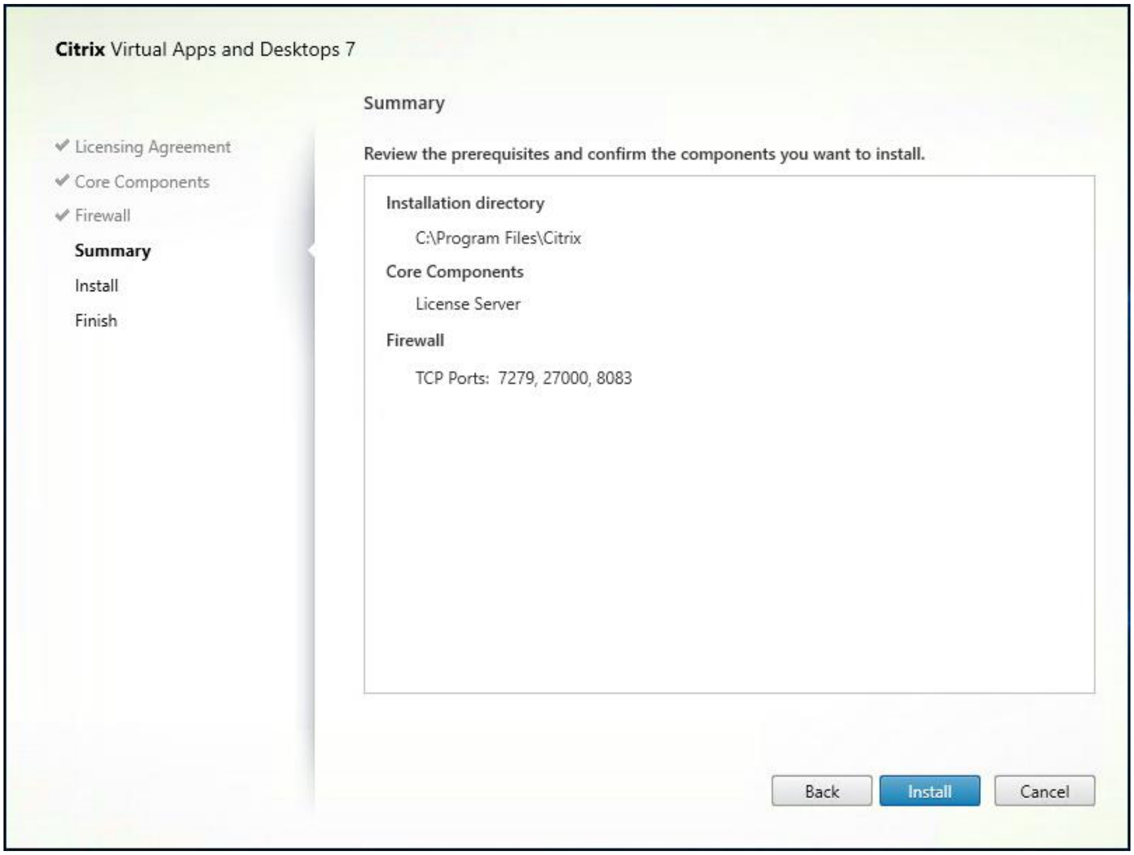

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory

C:\Program Files\Citrix

Core Components

License Server

Firewall

TCP Ports: 7279, 27000, 8083

Step 10. Click Finish to complete the installation.

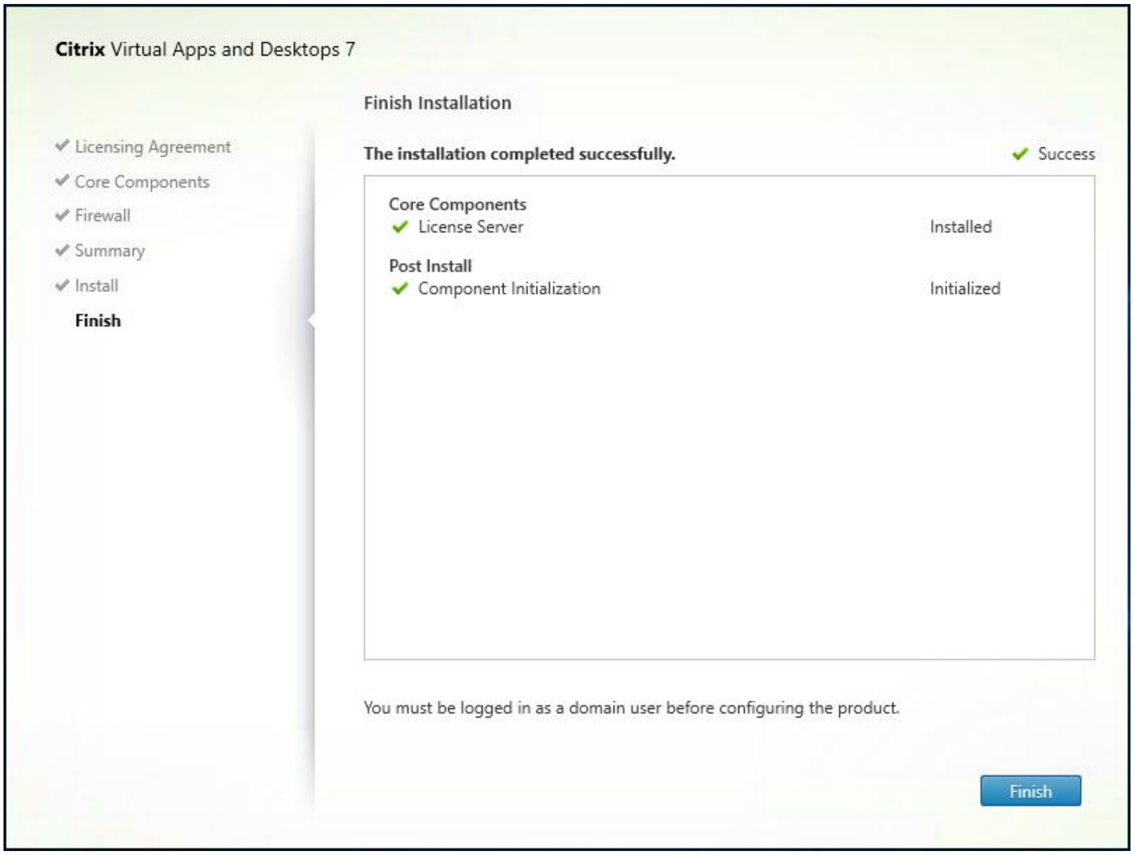

Finish Installation

The installation completed successfully. ✓ Success

Core Components	
✓ License Server	Installed
Post Install	
✓ Component Initialization	Initialized

You must be logged in as a domain user before configuring the product.

Procedure 2. Install Citrix Licenses

Step 1. Copy the license files to the default location (C:\Program Files (x86)\Citrix\Licensing\ MyFiles) on the license server.

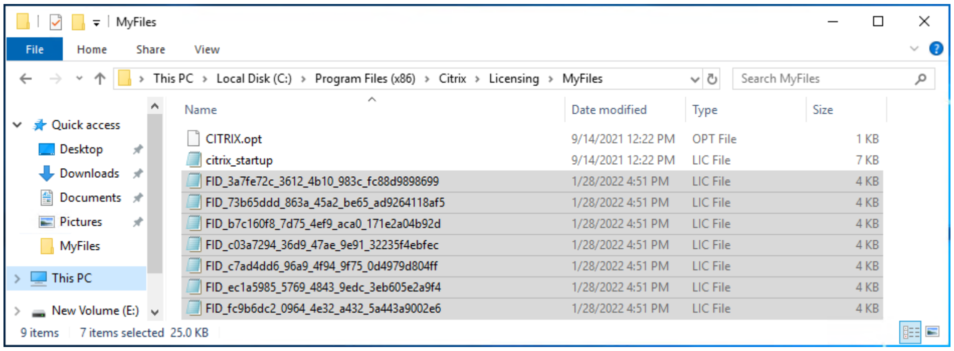

Name	Date modified	Type	Size
CITRIX.opt	9/14/2021 12:22 PM	OPT File	1 KB
citrix_startup	9/14/2021 12:22 PM	LIC File	7 KB
FID_3a7fe72c_3612_4b10_983c_fc88d9898699	1/28/2022 4:51 PM	LIC File	4 KB
FID_73b65ddd_863a_45a2_be65_ad9264118af5	1/28/2022 4:51 PM	LIC File	4 KB
FID_b7c160f8_7d75_4ef9_aca0_171e2a04b92d	1/28/2022 4:51 PM	LIC File	4 KB
FID_c03a7294_36d9_47ae_9e91_32235f4ebfec	1/28/2022 4:51 PM	LIC File	4 KB
FID_c7ad4dd6_96a9_4f94_9f75_0d4979d804ff	1/28/2022 4:51 PM	LIC File	4 KB
FID_ec1a5985_5769_4843_9edc_3eb605e2a9f4	1/28/2022 4:51 PM	LIC File	4 KB
FID_fc9b6dc2_0964_4e32_a432_5a443a9002e6	1/28/2022 4:51 PM	LIC File	4 KB

9 items 7 items selected 25.0 KB

Step 2. Restart the server or Citrix licensing services so that the licenses are activated.

Step 3. Run the application Citrix License Administration Console.

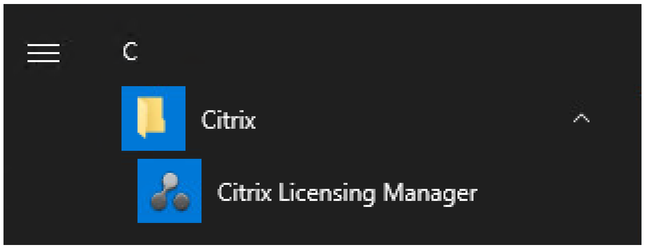

Step 4. Confirm that the license files have been read and enabled correctly.

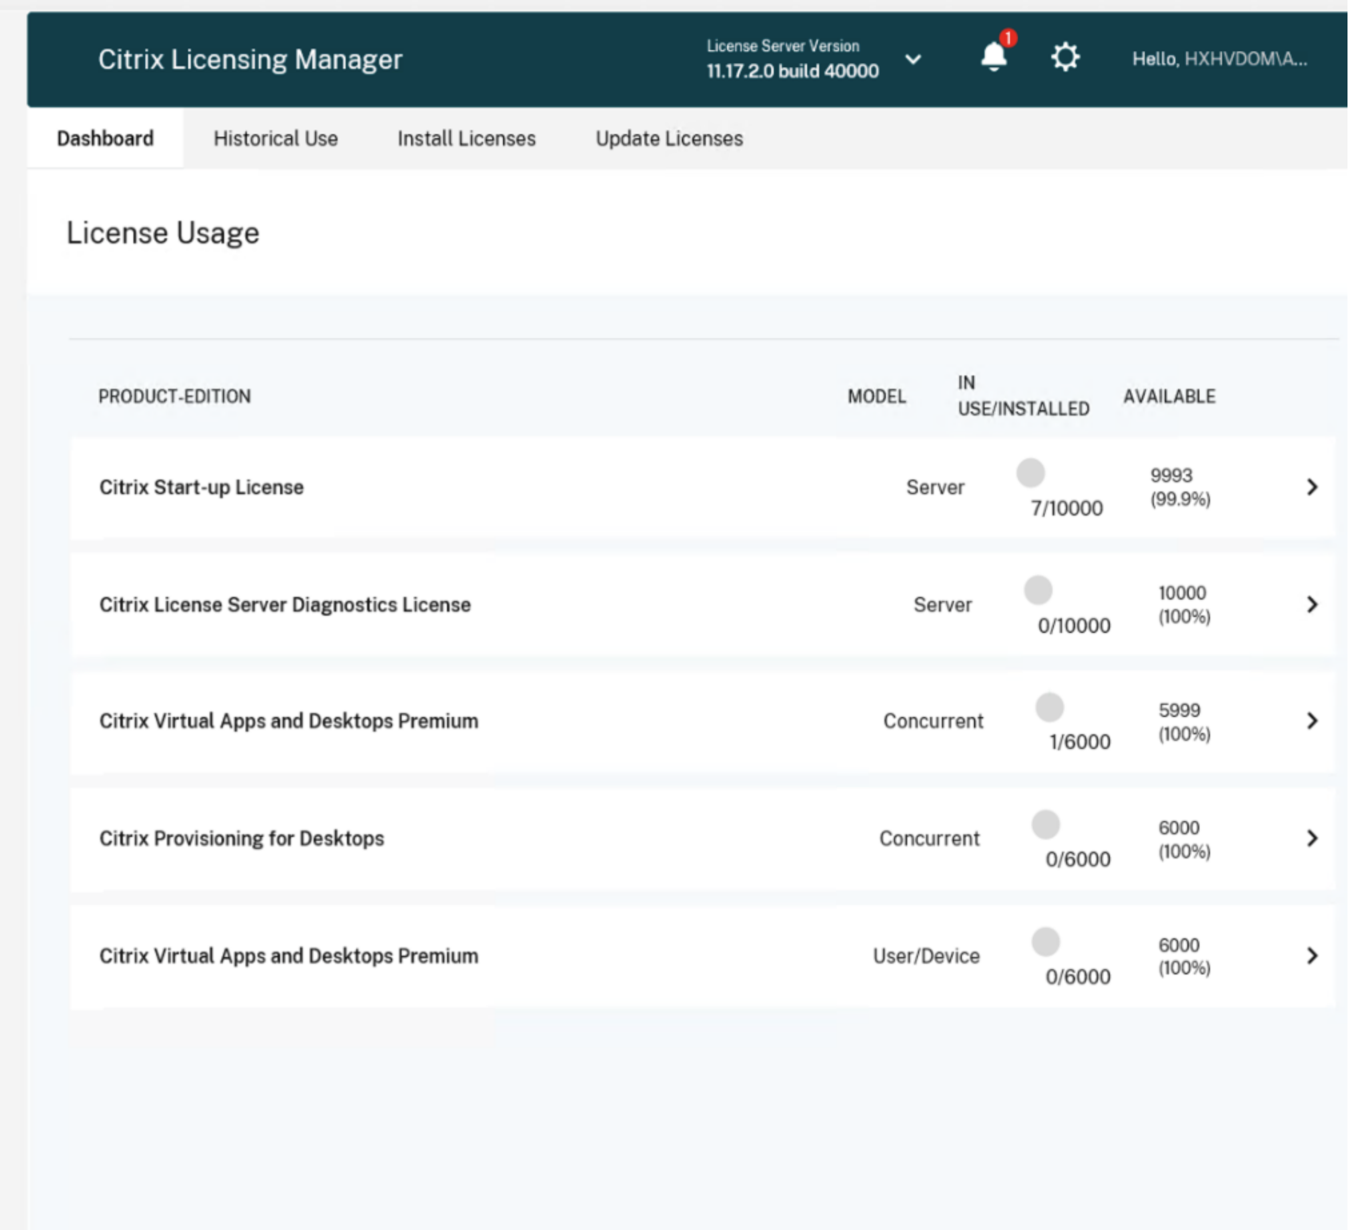

License Usage

PRODUCT-EDITION	MODEL	IN USE/INSTALLED	AVAILABLE
Citrix Start-up License	Server	7/10000	9993 (99.9%)
Citrix License Server Diagnostics License	Server	0/10000	10000 (100%)
Citrix Virtual Apps and Desktops Premium	Concurrent	1/6000	5999 (100%)
Citrix Provisioning for Desktops	Concurrent	0/6000	6000 (100%)
Citrix Virtual Apps and Desktops Premium	User/Device	0/6000	6000 (100%)

Procedure 3. Install the Citrix Virtual Apps and Desktops

Step 1. To begin the installation, connect to the first Delivery Controller server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2203 ISO.

Step 2. Click Start.

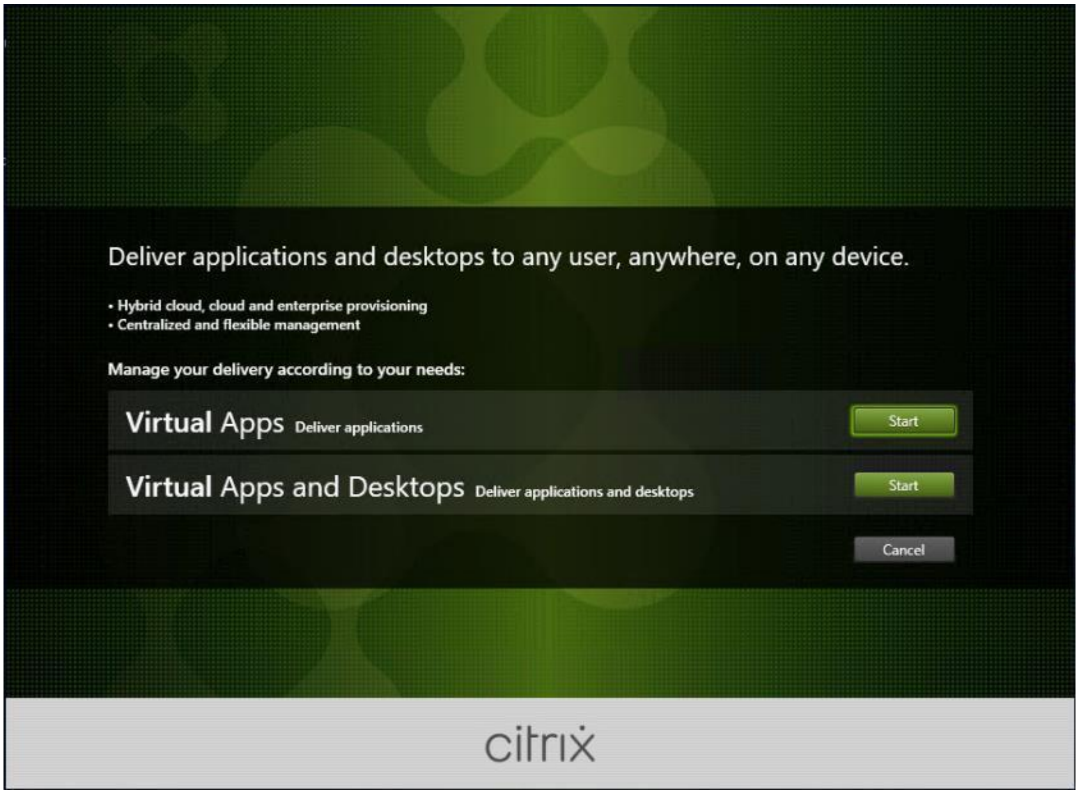

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

Virtual Apps Deliver applications
Start
Virtual Apps and Desktops Deliver applications and desktops
Start
Cancel
citrix

Step 3. The installation wizard presents a menu with three subsections. Click Get Started – Delivery Controller.

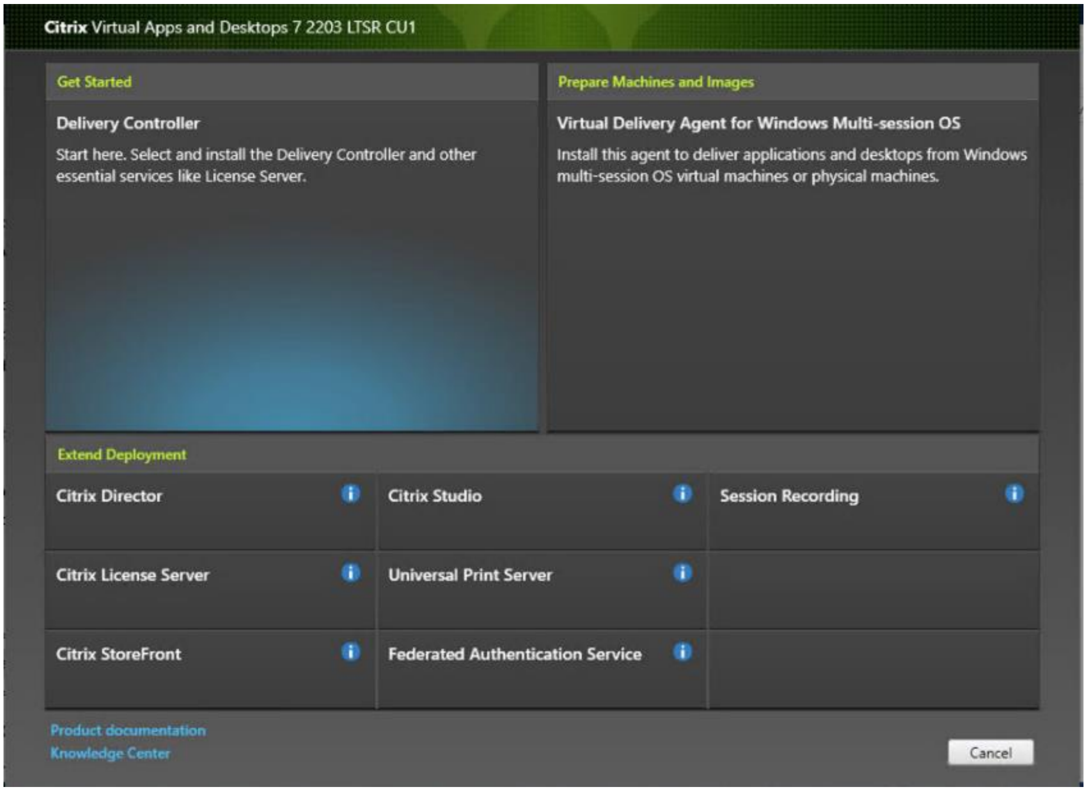

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Get Started

Delivery Controller

Start here. Select and install the Delivery Controller and other essential services like License Server.

Prepare Machines and Images

Virtual Delivery Agent for Windows Multi-session OS

Install this agent to deliver applications and desktops from Windows multi-session OS virtual machines or physical machines.

Step 4. Read the Citrix License Agreement. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

Step 5. Click Next.

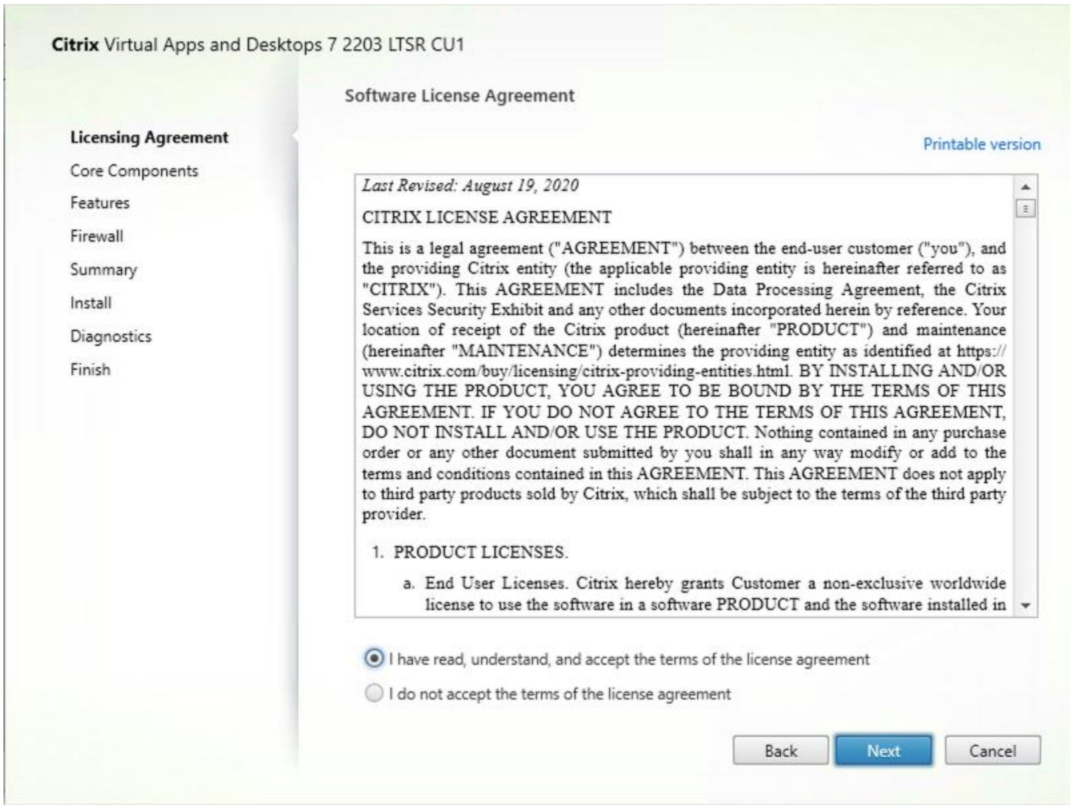

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Software License Agreement

Last Revised: August 19, 2020

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). This AGREEMENT includes the Data Processing Agreement, the Citrix Services Security Exhibit and any other documents incorporated herein by reference. Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY INSTALLING AND/OR USING THE PRODUCT, YOU AGREE TO BE BOUND BY THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, DO NOT INSTALL AND/OR USE THE PRODUCT. Nothing contained in any purchase order or any other document submitted by you shall in any way modify or add to the terms and conditions contained in this AGREEMENT. This AGREEMENT does not apply to third party products sold by Citrix, which shall be subject to the terms of the third party provider.

1. PRODUCT LICENSES.
   a. End User Licenses. Citrix hereby grants Customer a non-exclusive worldwide license to use the software in a software PRODUCT and the software installed in

● I have read, understand, and accept the terms of the license agreement

○ I do not accept the terms of the license agreement

Step 6. Select the components to be installed on the first Delivery Controller Server:

- Delivery Controller
- Studio
- Director

Step 7. Click Next.

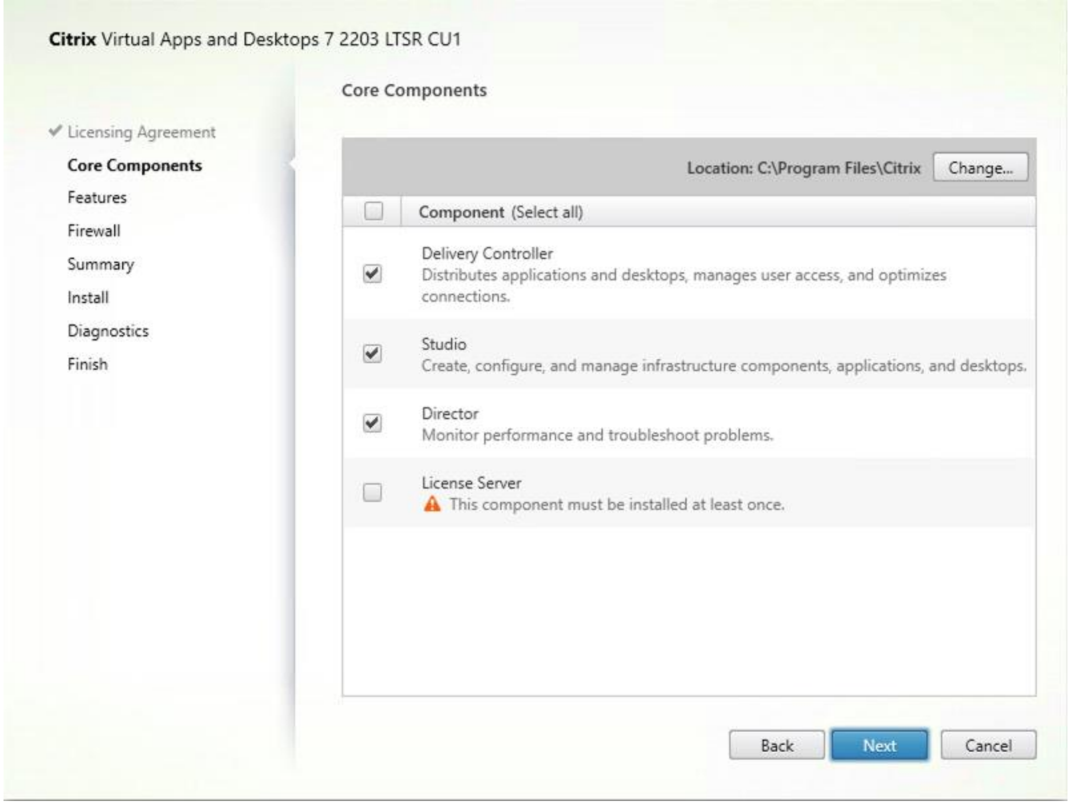

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Core Components

Location: C:\Program Files\Citrix

	Component (Select all)
☑	Delivery Controller Distributes applications and desktops, manages user access, and optimizes connections.
☑	Studio Create, configure, and manage infrastructure components, applications, and desktops.
☑	Director Monitor performance and troubleshoot problems.
□	License Server ⚠ This component must be installed at least once.

Step 8. Since a dedicated SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2014 SP2 Express” unchecked.

Step 9. Click Next.

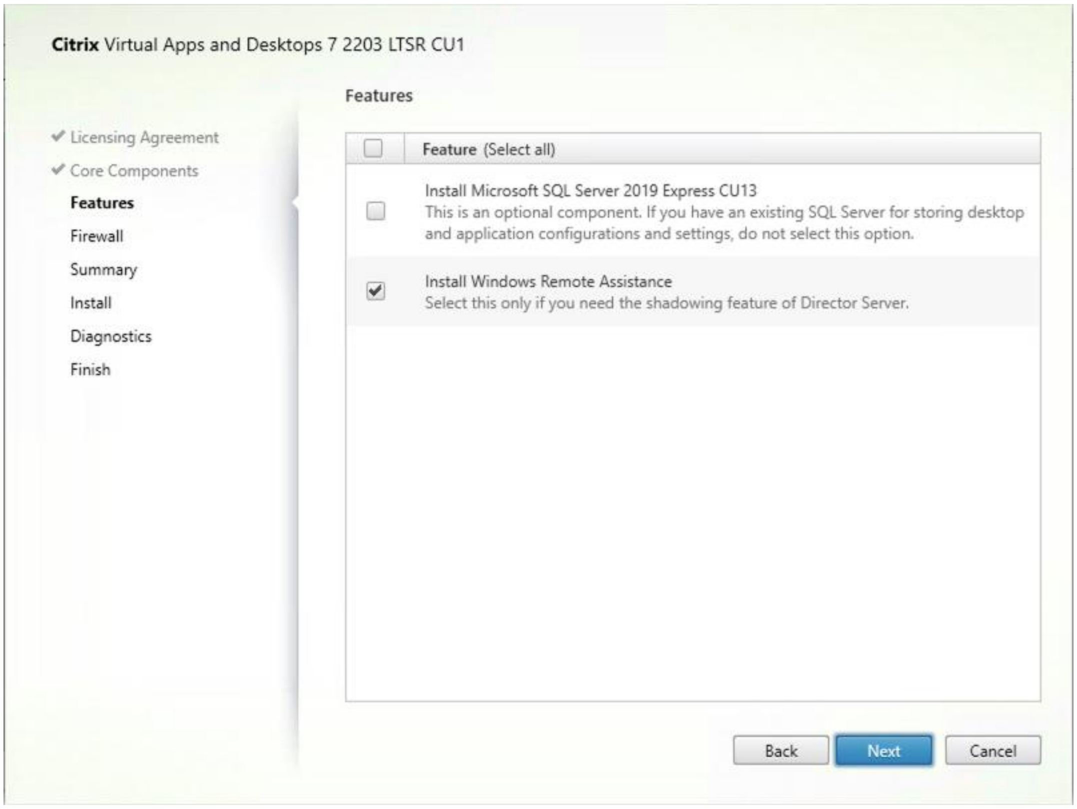

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Features

	Feature (Select all)
□	Install Microsoft SQL Server 2019 Express CU13 This is an optional component. If you have an existing SQL Server for storing desktop and application configurations and settings, do not select this option.
☑	Install Windows Remote Assistance Select this only if you need the shadowing feature of Director Server.

Step 10. Select the default ports and automatically configured firewall rules.

Step 11. Click Next.

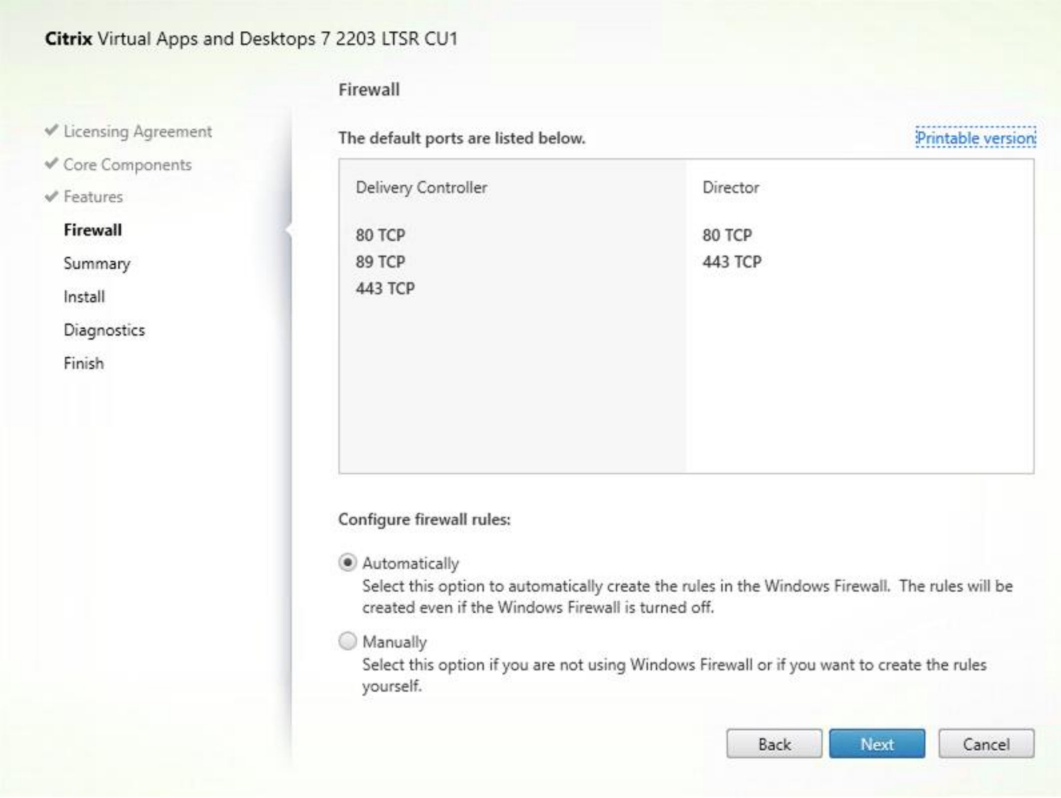

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Firewall

The default ports are listed below.

Delivery Controller	Director
80 TCP 89 TCP 443 TCP	80 TCP 443 TCP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

Step 12. Click Finish to begin the installation.

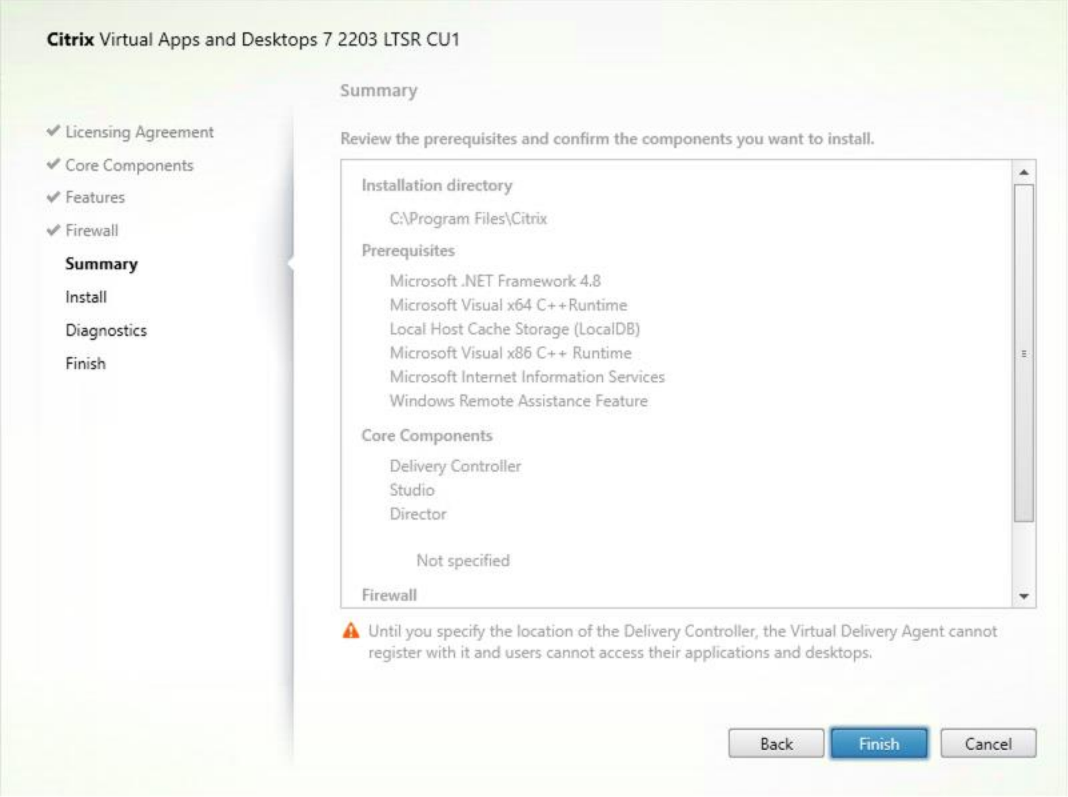

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Summary

Review the prerequisites and confirm the components you want to install.

Installation directory
C:\Program Files\Citrix

Prerequisites
Microsoft .NET Framework 4.8
Microsoft Visual x64 C++Runtime
Local Host Cache Storage (LocalDB)
Microsoft Visual x86 C++ Runtime
Microsoft Internet Information Services
Windows Remote Assistance Feature

Core Components
Delivery Controller
Studio
Director

Not specified

Firewall

⚠ Until you specify the location of the Delivery Controller, the Virtual Delivery Agent cannot register with it and users cannot access their applications and desktops.

Note: Multiple reboots may be required to finish installation.

Step 13. (Optional) Collect diagnostic information/Call Home participation.

Step 14. Click Next.

Step 15. Click Finish to complete the installation.

Step 16. (Optional) Check Launch Studio to launch Citrix Studio Console.

Procedure 4. Additional Delivery Controller Configuration

Note: After the first controller is completely configured and the Site is operational, you can add additional controllers. In this CVD, we created two Delivery Controllers.

To configure additional Delivery Controllers, repeat the steps in section Install the Citrix Virtual Apps and Desktops.

Step 1. To begin the installation of the second Delivery Controller, connect to the second Delivery Controller server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2203 ISO.

Step 2. Click Start.

Step 3. Click Delivery Controller.

Step 4. Repeat the same steps used to install the first Delivery Controller; Install the Citrix Virtual Apps and Desktops, including the step of importing an SSL certificate for HTTPS between the controller and vSphere.

Step 5. Review the Summary configuration. Click Finish.

Step 6. (Optional) Configure Collect diagnostic information /Call Home participation. Click Next.

Step 7. Verify the components installed successfully. Click Finish.

Procedure 5. Create Site

Citrix Studio is a management console that allows you to create and manage infrastructure and resources to deliver desktops and applications. Replacing Desktop Studio from earlier releases, it provides wizards to set up your environment, create workloads to host applications and desktops, and assign applications and desktops to users.

Citrix Studio launches automatically after the Delivery Controller installation, or if necessary, it can be launched manually. Studio is used to create a Site, which is the core of the Citrix Virtual Apps and Desktops environment consisting of the Delivery Controller and the Database.

Step 1. From Citrix Studio, click Deliver applications and desktops to your users.

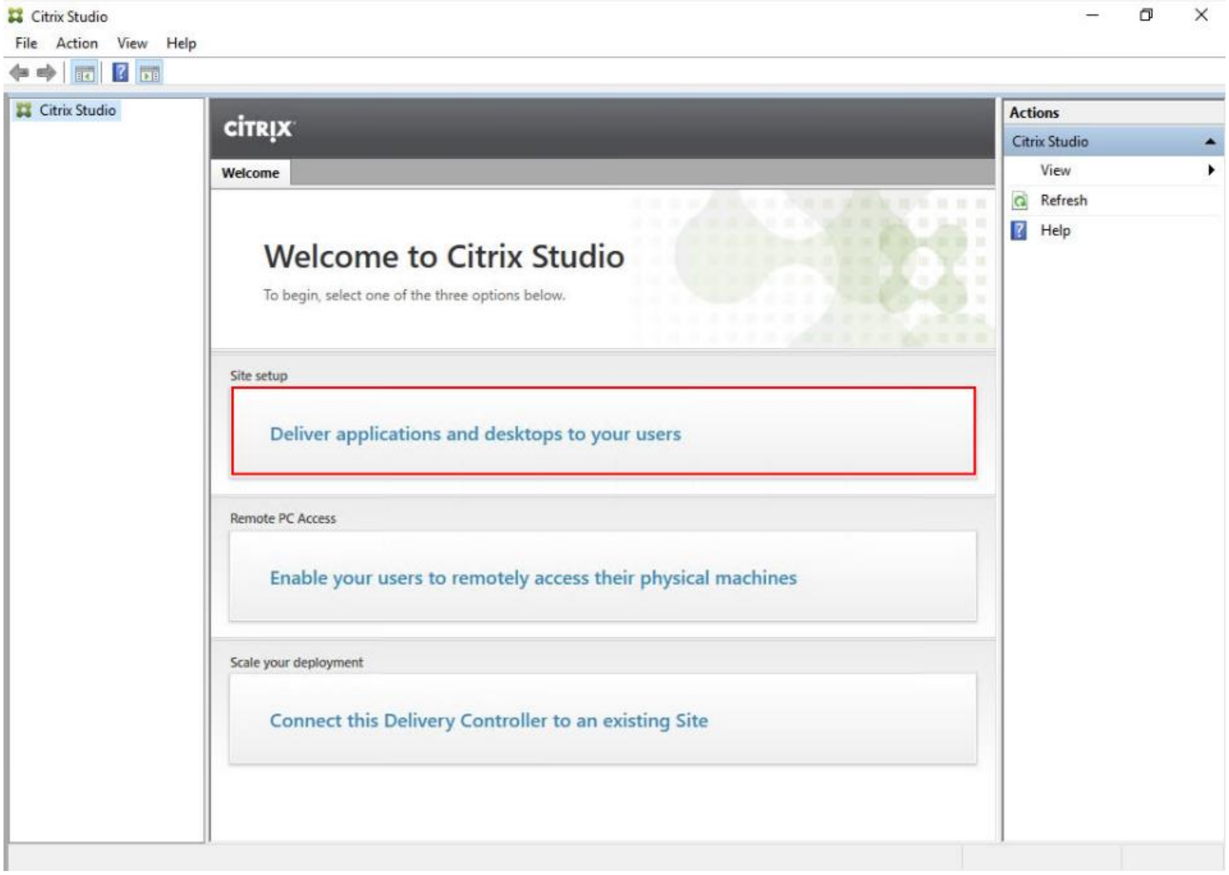

Welcome to Citrix Studio

To begin, select one of the three options below.

Site setup

Deliver applications and desktops to your users

Remote PC Access

Enable your users to remotely access their physical machines

Scale your deployment

Connect this Delivery Controller to an existing Site

Step 2. Select the “An empty, unconfigured Site” radio button.

Step 3. Enter a site name.

Step 4. Click Next.

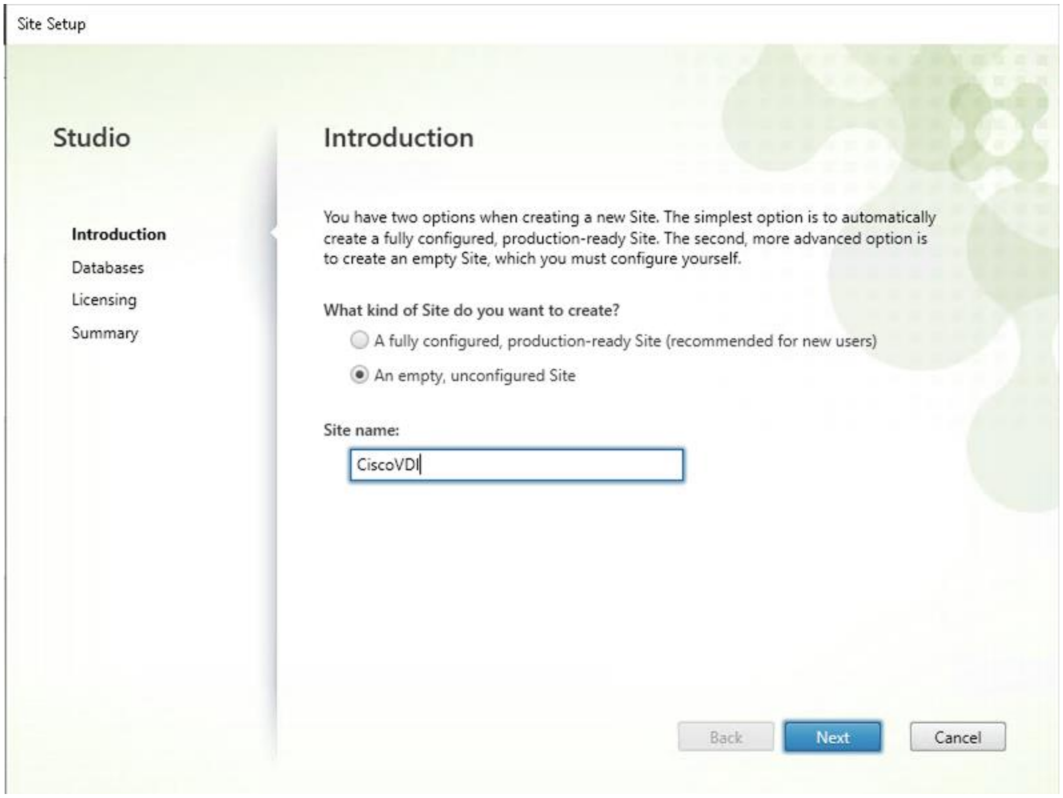
Site Setup

Studio

Introduction

You have two options when creating a new Site. The simplest option is to automatically create a fully configured, production-ready Site. The second, more advanced option is to create an empty Site, which you must configure yourself.

What kind of Site do you want to create?

○ A fully configured, production-ready Site (recommended for new users)

● An empty, unconfigured Site

Site name:

CiscoVDI

Step 5. Provide the Database Server Locations for each data type.

Note: For an SQL AlwaysOn Availability Group, use the group’s listener DNS name.

Step 6. Click Select to specify additional controllers (Optional at this time. Additional controllers can be added later).

Step 7. Click Next.

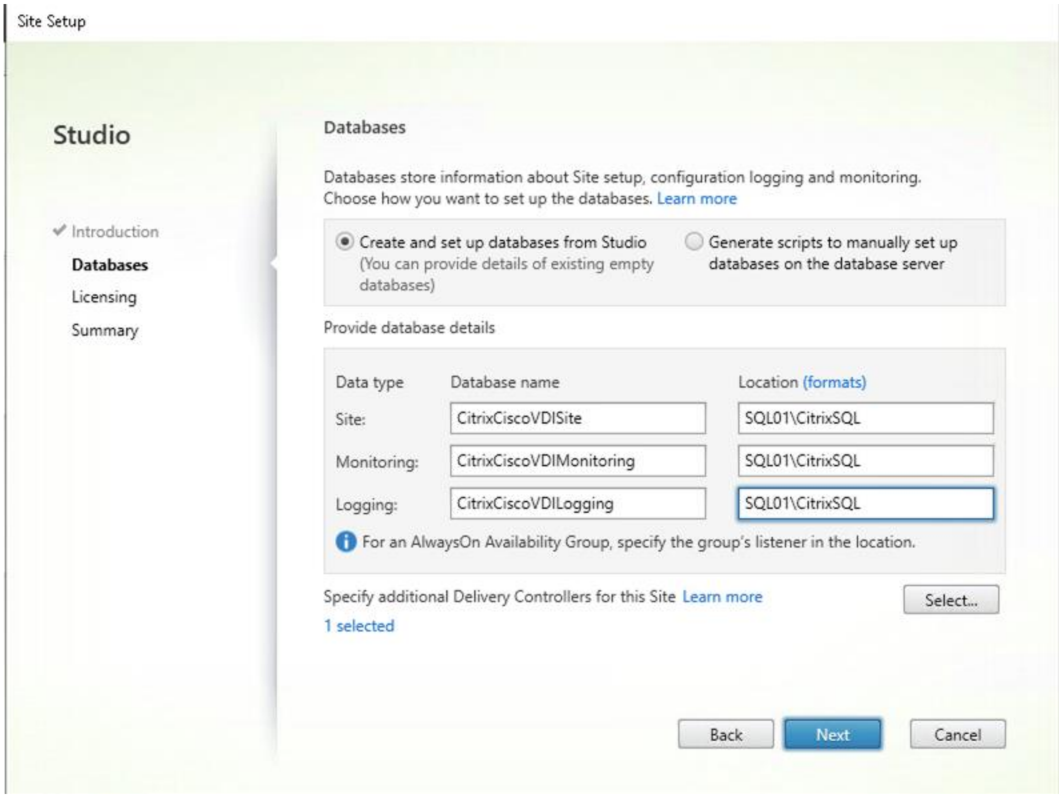

Site Setup

Studio

Databases

Databases store information about Site setup, configuration logging and monitoring. Choose how you want to set up the databases. Learn more

● Create and set up databases from Studio (You can provide details of existing empty databases)

○ Generate scripts to manually set up databases on the database server

Provide database details

Data type	Database name	Location (formats)
Site:	CitrixCiscoVDISite	SQL01\CitrixSQL
Monitoring:	CitrixCiscoVDIMonitoring	SQL01\CitrixSQL
Logging:	CitrixCiscoVDILogging	SQL01\CitrixSQL

For an AlwaysOn Availability Group, specify the group's listener in the location.

Specify additional Delivery Controllers for this Site Learn more

1 selected

Step 8. Provide the FQDN of the license server.

Step 9. Click Connect to validate and retrieve any licenses from the server.

Note: If no licenses are available, you can use the 30-day free trial or activate a license file.

Step 10. Select the appropriate product edition using the license radio button.

Step 11. Click Next.

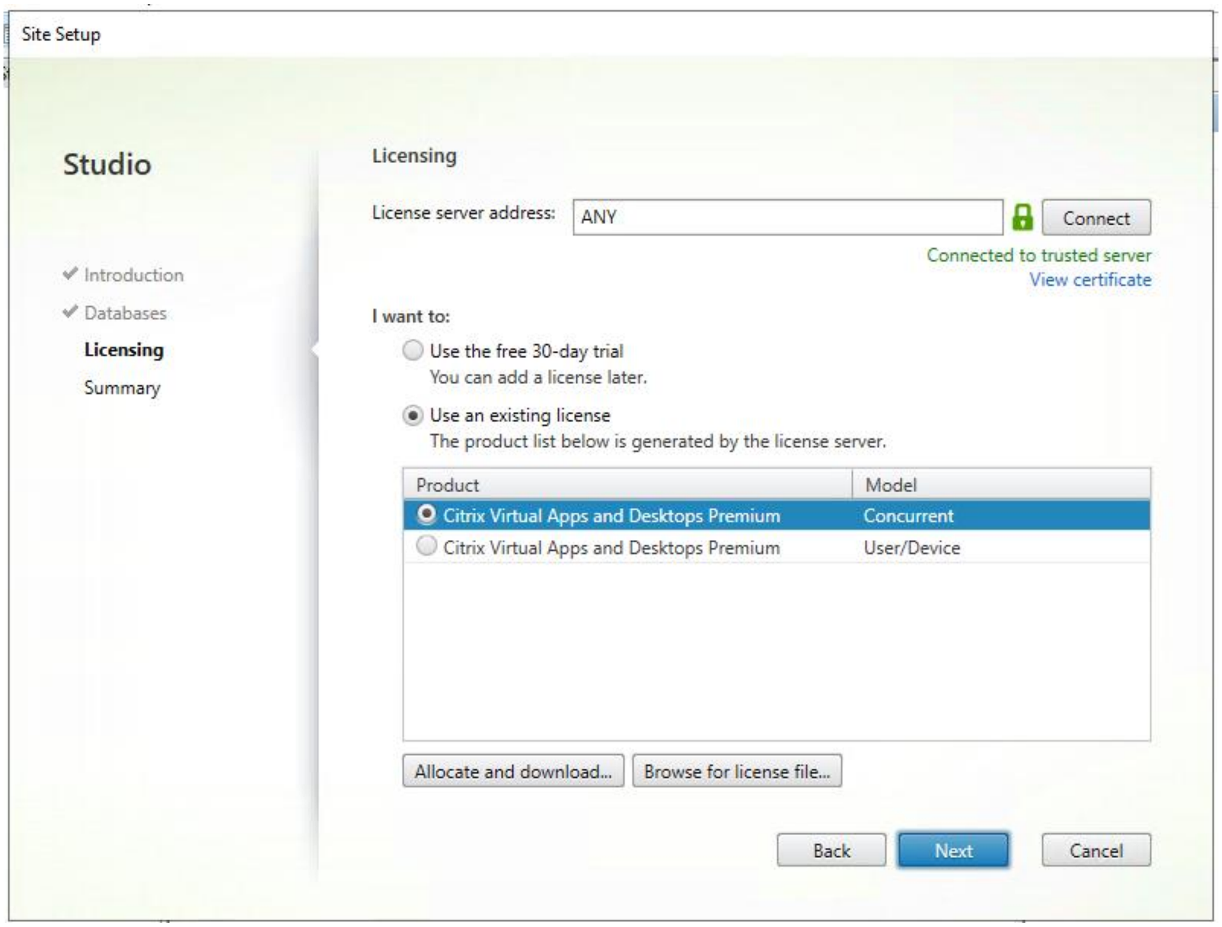

Site Setup

Studio

✔ Introduction
✔ Databases
Licensing
Summary

Licensing

License server address: ANY

Connected to trusted server
View certificate

I want to:

○ Use the free 30-day trial
You can add a license later.

● Use an existing license
The product list below is generated by the license server.

Product	Model
● Citrix Virtual Apps and Desktops Premium	Concurrent
○ Citrix Virtual Apps and Desktops Premium	User/Device

Step 12. Verify information on the Summary page.

Step 13. Click Finish.

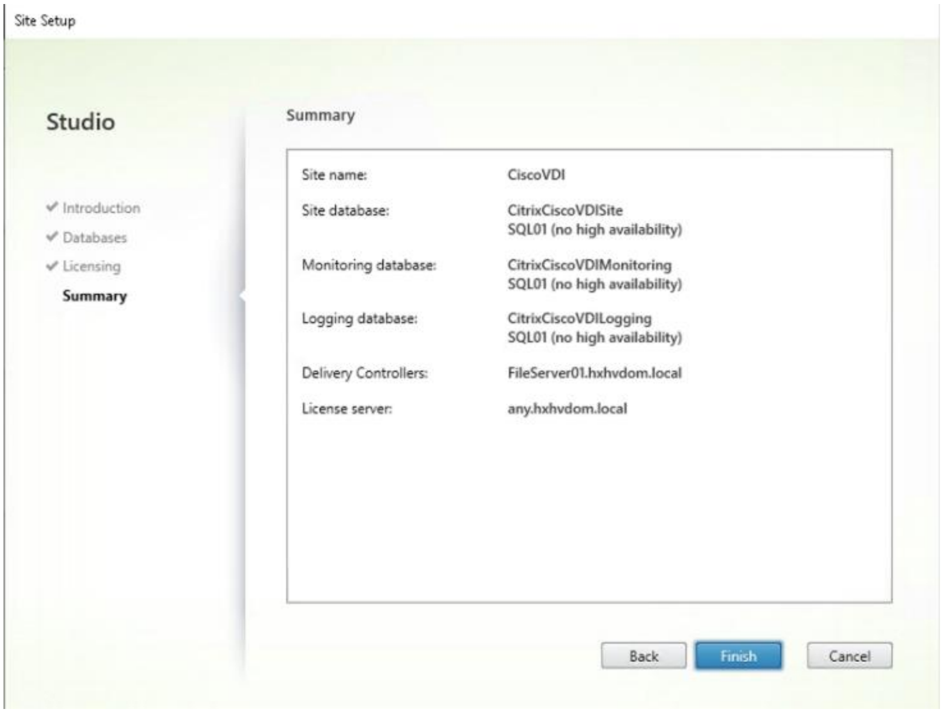

Site Setup

Studio

✔ Introduction
✔ Databases
✔ Licensing
Summary

Summary

Site name:	CiscoVDI
Site database:	CitrixCiscoVDISite SQL01 (no high availability)
Monitoring database:	CitrixCiscoVDIMonitoring SQL01 (no high availability)
Logging database:	CitrixCiscoVDILogging SQL01 (no high availability)
Delivery Controllers:	FileServer01.hxhvdom.local
License server:	any.hxhvdom.local

Procedure 6. Configure the Citrix Virtual Apps and Desktops Site Hosting Connection

Step 1. From Configuration > Hosting in Studio, click Add Connection and Resources in the right pane.

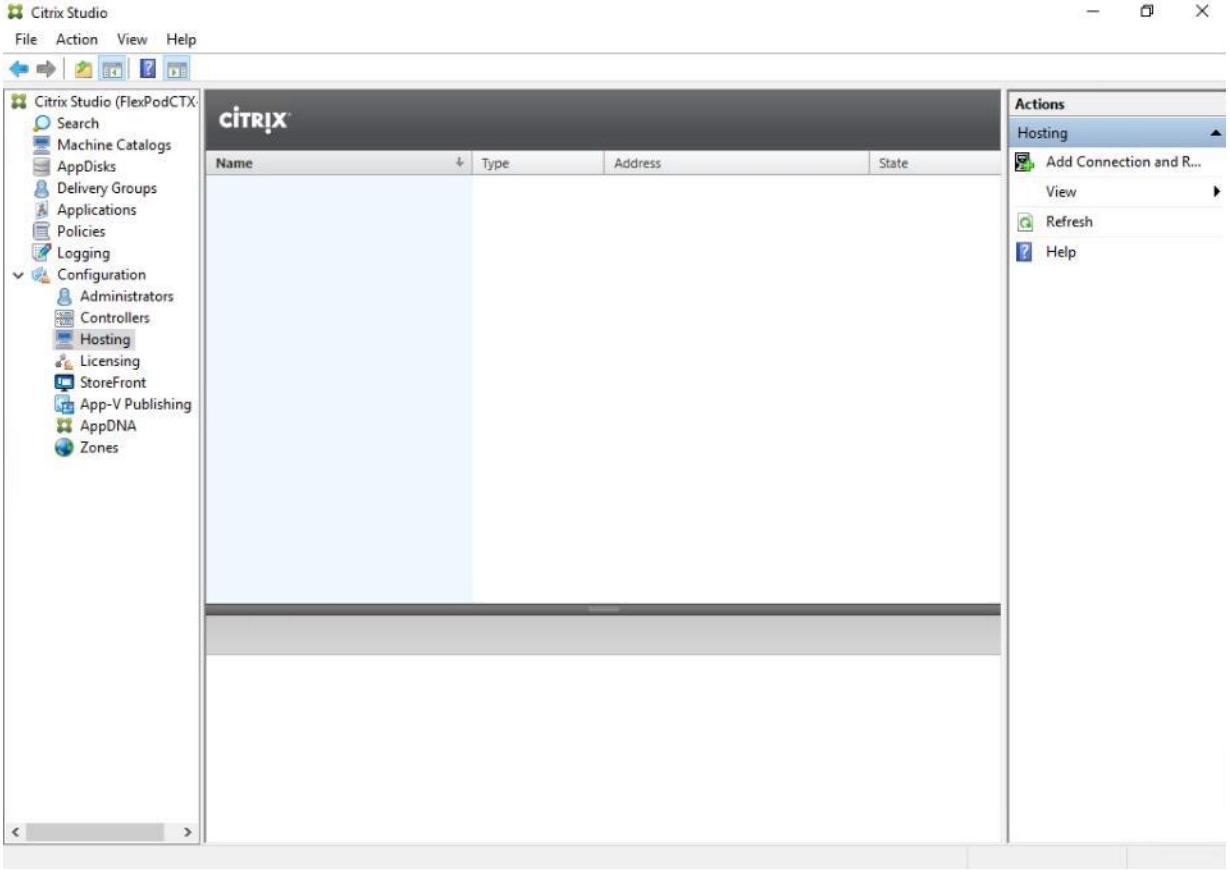

Step 2. On the Connection page:

a. Select the Connection type of VMware vSphere.

b. Enter the FQDN of the vCenter server (in Server_FQDN/sdk format).

c. Enter the username (in domain\username format) for the vSphere account.

d. Provide the password for the vSphere account.

e. Provide a connection name.

f. Choose the tool to create virtual machines: Machine Creation Services or Citrix Provisioning

Step 3. Click Next.

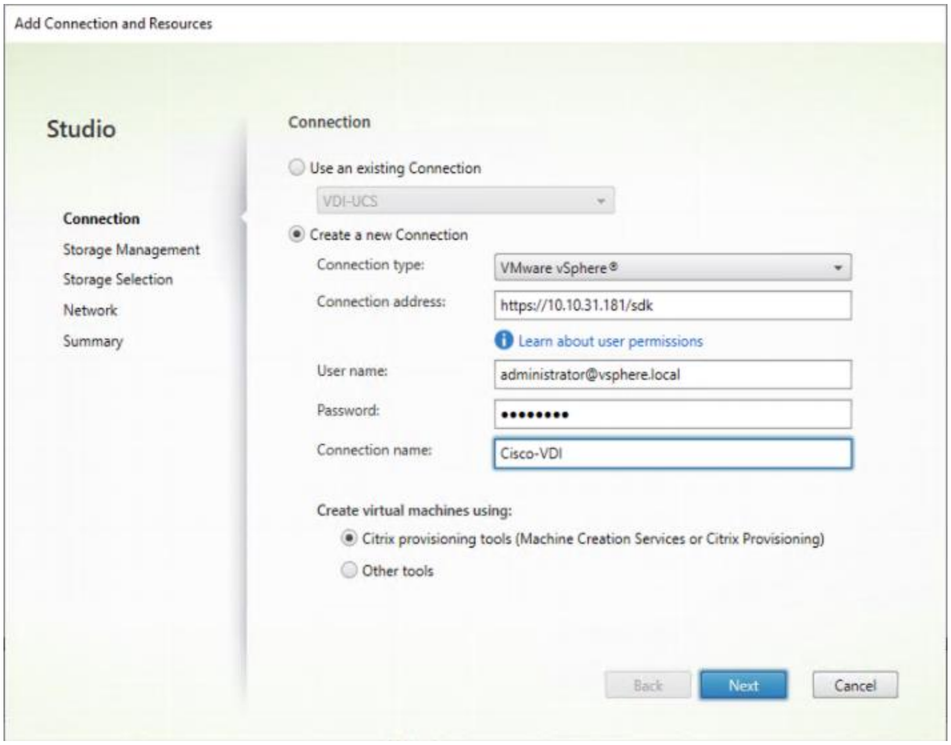
Add Connection and Resources

Studio

Connection

○ Use an existing Connection

VDI-UCS

● Create a new Connection

Connection type: VMware vSphere®

Connection address: https://10.10.31.181/sdk

Learn about user permissions

User name: administrator@vsphere.local

Password: ••••••••

Connection name: Cisco-VDI

Create virtual machines using:

● Citrix provisioning tools (Machine Creation Services or Citrix Provisioning)

○ Other tools

Step 4. Accept the certificate and click OK to trust the hypervisor connection.

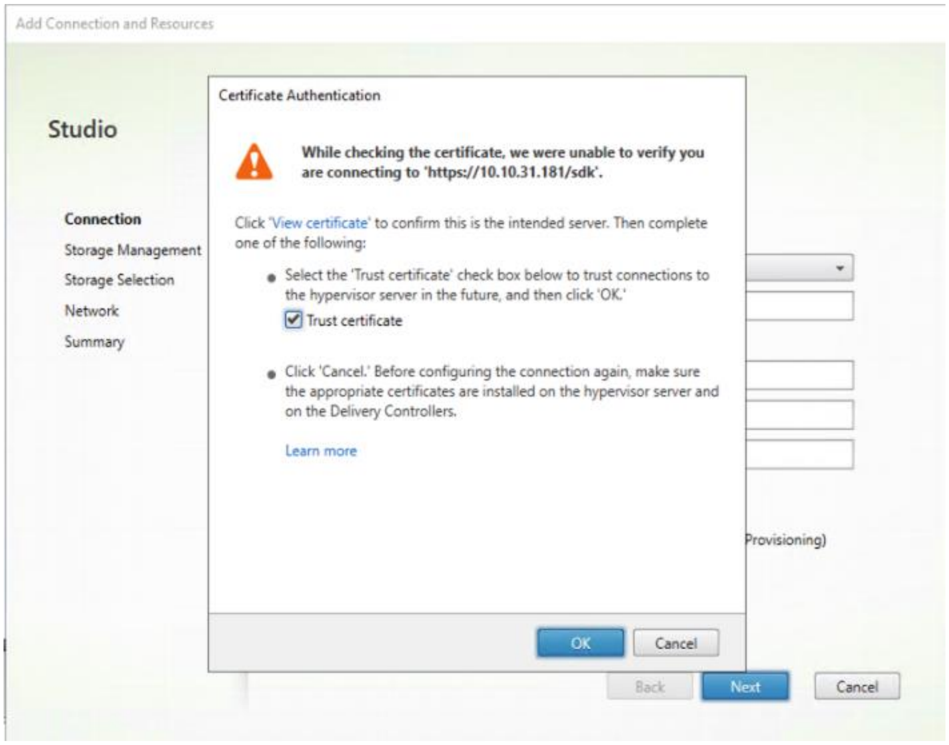
Certificate Authentication

While checking the certificate, we were unable to verify you are connecting to 'https://10.10.31.181/sdk'.

Click 'View certificate' to confirm this is the intended server. Then complete one of the following:

- Select the 'Trust certificate' check box below to trust connections to the hypervisor server in the future, and then click 'OK.'

  ☑ Trust certificate

- Click 'Cancel.' Before configuring the connection again, make sure the appropriate certificates are installed on the hypervisor server and on the Delivery Controllers.

  Learn more

Step 5. Select a storage management method:

Step 6. Select Cluster that will be used by this connection.

Step 7. Check Use storage shared by hypervisors radio button.

Step 8. Click Next.

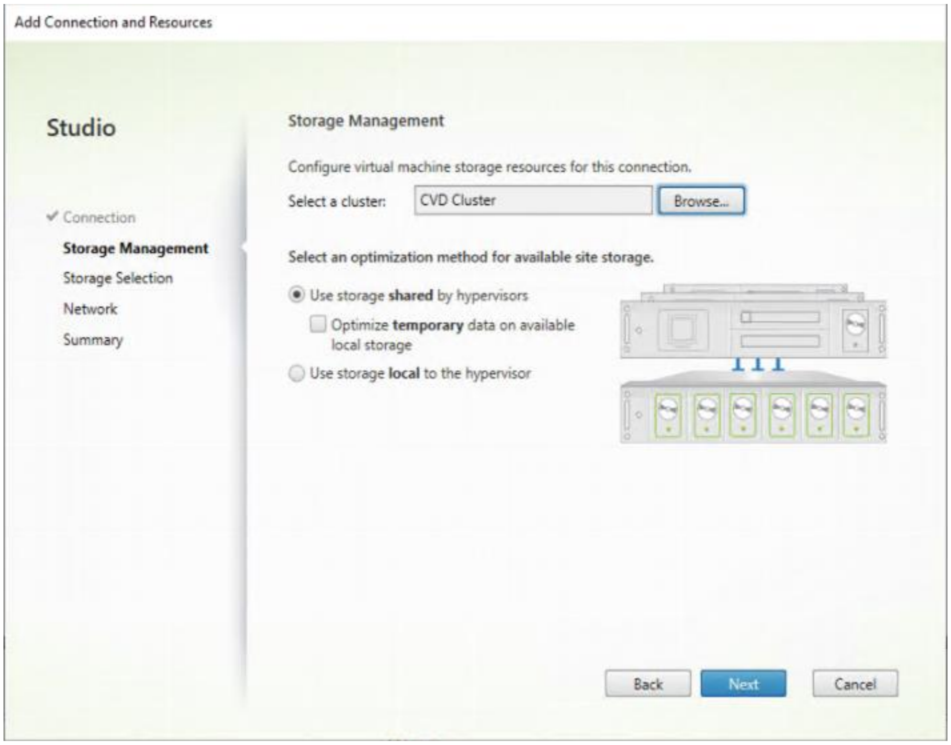

Storage Management

Configure virtual machine storage resources for this connection.

Select a cluster: CVD Cluster [Browse...]

Select an optimization method for available site storage.

● Use storage **shared** by hypervisors

□ Optimize **temporary** data on available local storage

○ Use storage **local** to the hypervisor

Step 9. Select the Storage to be used by this connection, use all provisioned for desktops datastores.

Step 10. Click Next.

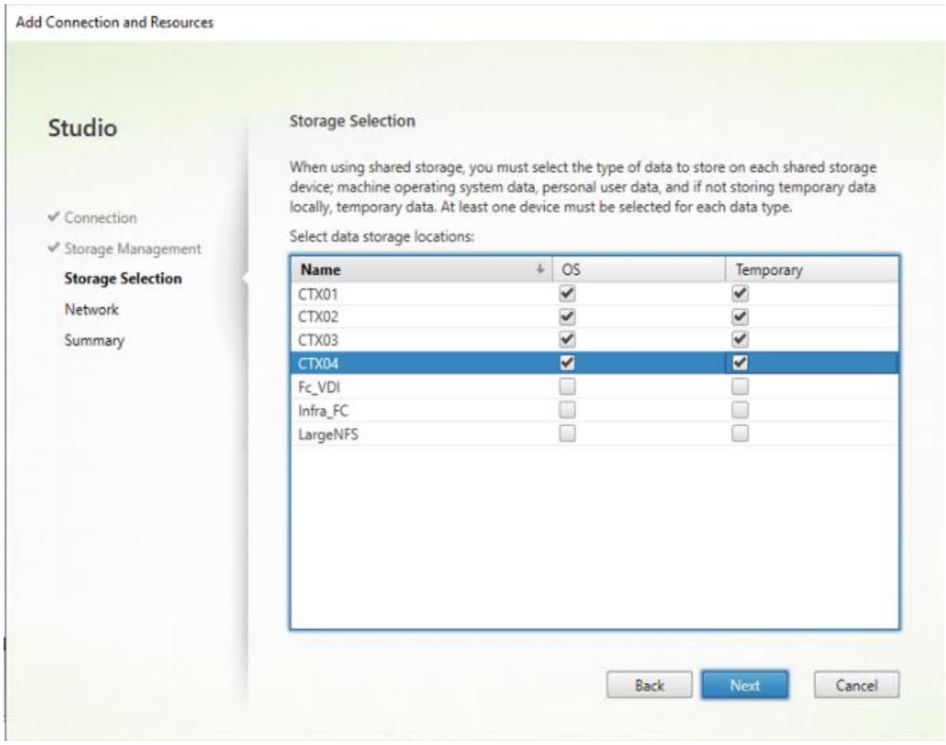

Storage Selection

When using shared storage, you must select the type of data to store on each shared storage device; machine operating system data, personal user data, and if not storing temporary data locally, temporary data. At least one device must be selected for each data type.

Select data storage locations:

Name	OS	Temporary
CTX01	☑	☑
CTX02	☑	☑
CTX03	☑	☑
CTX04	☑	☑
Fc_VDI	□	□
Infra_FC	□	□
LargeNFS	□	□

Step 11. Select the Network to be used by this connection.

Step 12. Click Next.

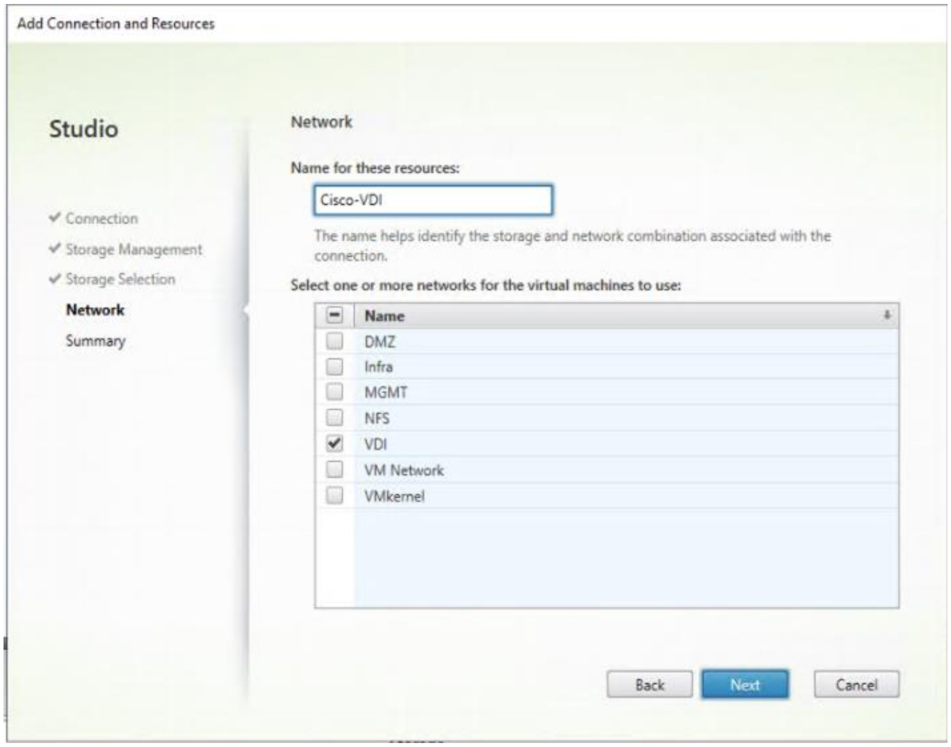

Add Connection and Resources

Studio

Network

Name for these resources: Cisco-VDI

The name helps identify the storage and network combination associated with the connection.

Select one or more networks for the virtual machines to use:

☐	Name
☐	DMZ
☐	Infra
☐	MGMT
☐	NFS
☑	VDI
☐	VM Network
☐	VMkernel

Step 13. Review Add Connection and Recourses Summary.

Step 14. Click Finish.

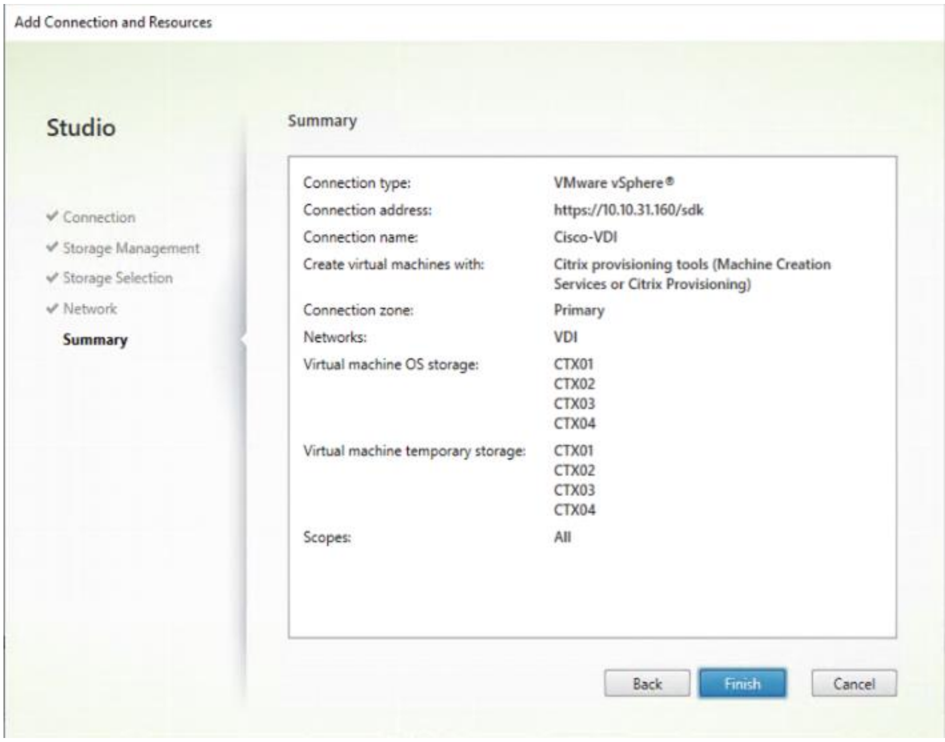

Add Connection and Resources

Studio

Summary

Connection type:	VMware vSphere®
Connection address:	https://10.10.31.160/sdk
Connection name:	Cisco-VDI
Create virtual machines with:	Citrix provisioning tools (Machine Creation Services or Citrix Provisioning)
Connection zone:	Primary
Networks:	VDI
Virtual machine OS storage:	CTX01 CTX02 CTX03 CTX04
Virtual machine temporary storage:	CTX01 CTX02 CTX03 CTX04
Scopes:	All

Procedure 7. Configure the Citrix Virtual Apps and Desktops Site Administrators

Step 1. Connect to the Citrix Virtual Apps and Desktops server and open Citrix Studio Management console.

Step 2. From the Configuration menu, right-click Administrator and select Create Administrator from the drop-down list.

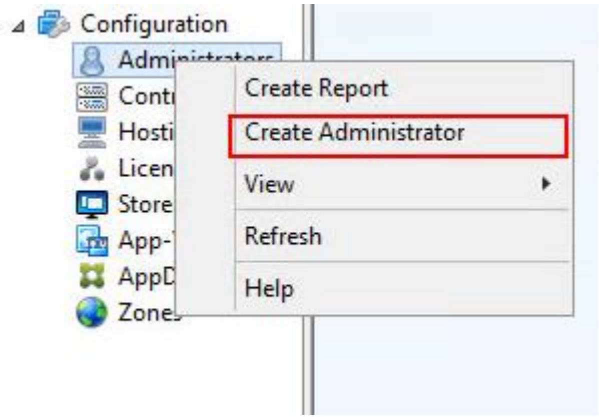

Step 3. Select or Create appropriate scope and click Next.

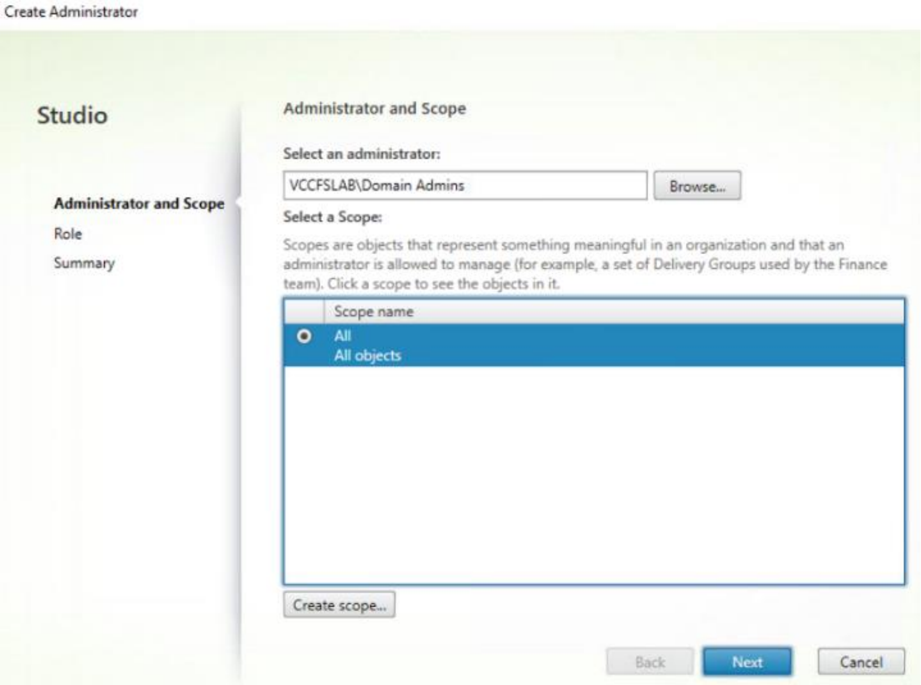
Create Administrator

Studio

Administrator and Scope

Select an administrator:

VCCFSLAB\Domain Admins

Select a Scope:

Scopes are objects that represent something meaningful in an organization and that an administrator is allowed to manage (for example, a set of Delivery Groups used by the Finance team). Click a scope to see the objects in it.

	Scope name
●	All All objects

Step 4. Select an appropriate Role.

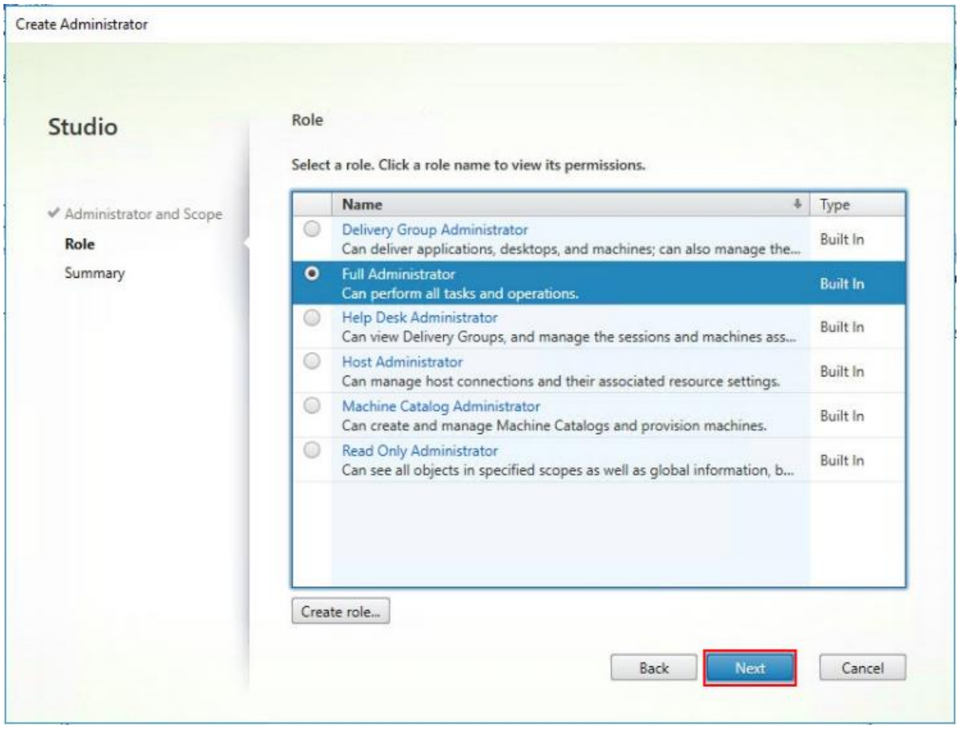
Create Administrator

Role

Select a role. Click a role name to view its permissions.

	Name	Type
○	Delivery Group Administrator Can deliver applications, desktops, and machines; can also manage the...	Built In
●	Full Administrator Can perform all tasks and operations.	Built In
○	Help Desk Administrator Can view Delivery Groups, and manage the sessions and machines ass...	Built In
○	Host Administrator Can manage host connections and their associated resource settings.	Built In
○	Machine Catalog Administrator Can create and manage Machine Catalogs and provision machines.	Built In
○	Read Only Administrator Can see all objects in specified scopes as well as global information, b...	Built In

Step 5. Review the Summary, check Enable administrator and click Finish.

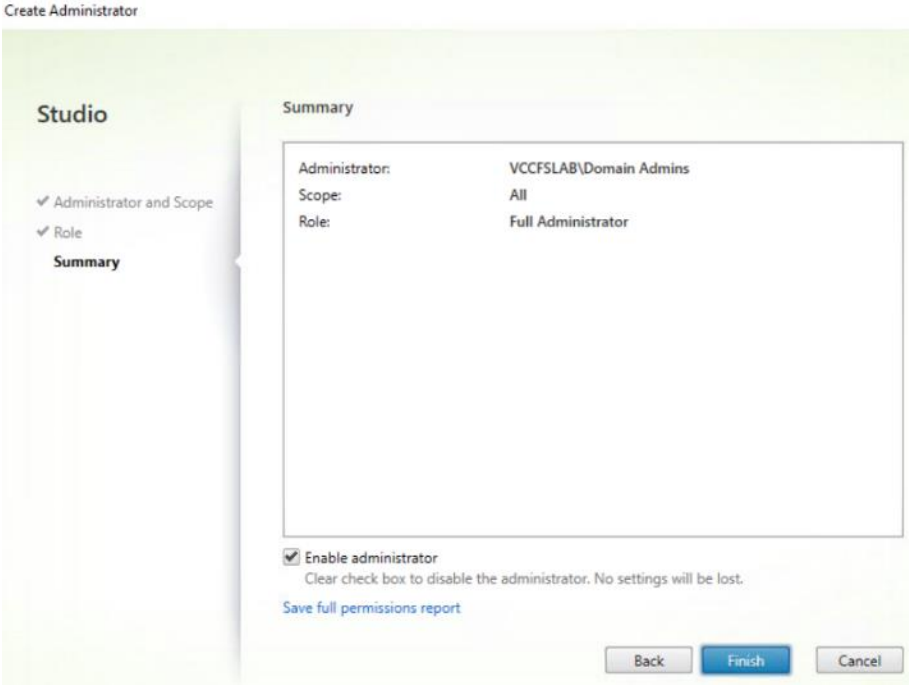
Create Administrator

Summary

Administrator:	VCCFSLAB\Domain Admins
Scope:	All
Role:	Full Administrator

☑ Enable administrator
Clear check box to disable the administrator. No settings will be lost.

Procedure 8. Install and Configure StoreFront

Note: Citrix StoreFront stores aggregate desktops and applications from Citrix Virtual Apps and Desktops sites, making resources readily available to users. In this CVD, we created two StoreFront servers on dedicated virtual machines.

Step 1. To begin the installation of the StoreFront, connect to the first StoreFront server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2203 ISO.

Step 2. Click Start.

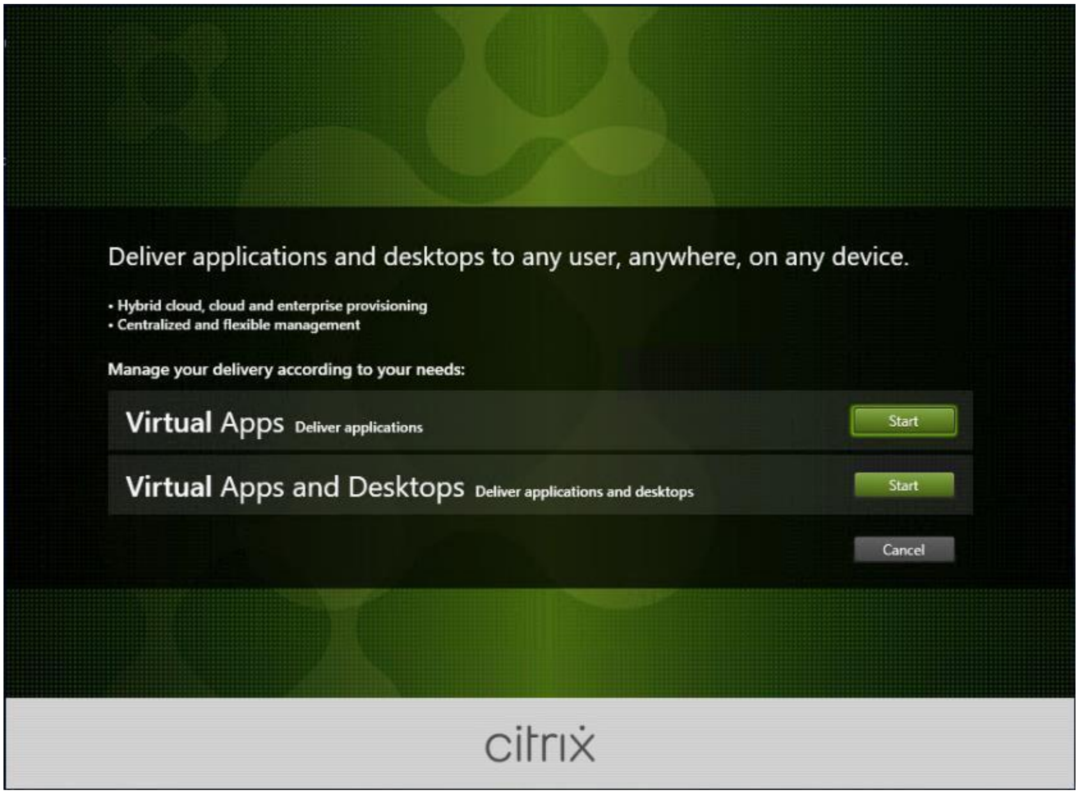

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

Virtual Apps Deliver applications
Start
Virtual Apps and Desktops Deliver applications and desktops
Start
Cancel
citrix

Step 3. Click Extend Deployment Citrix StoreFront.

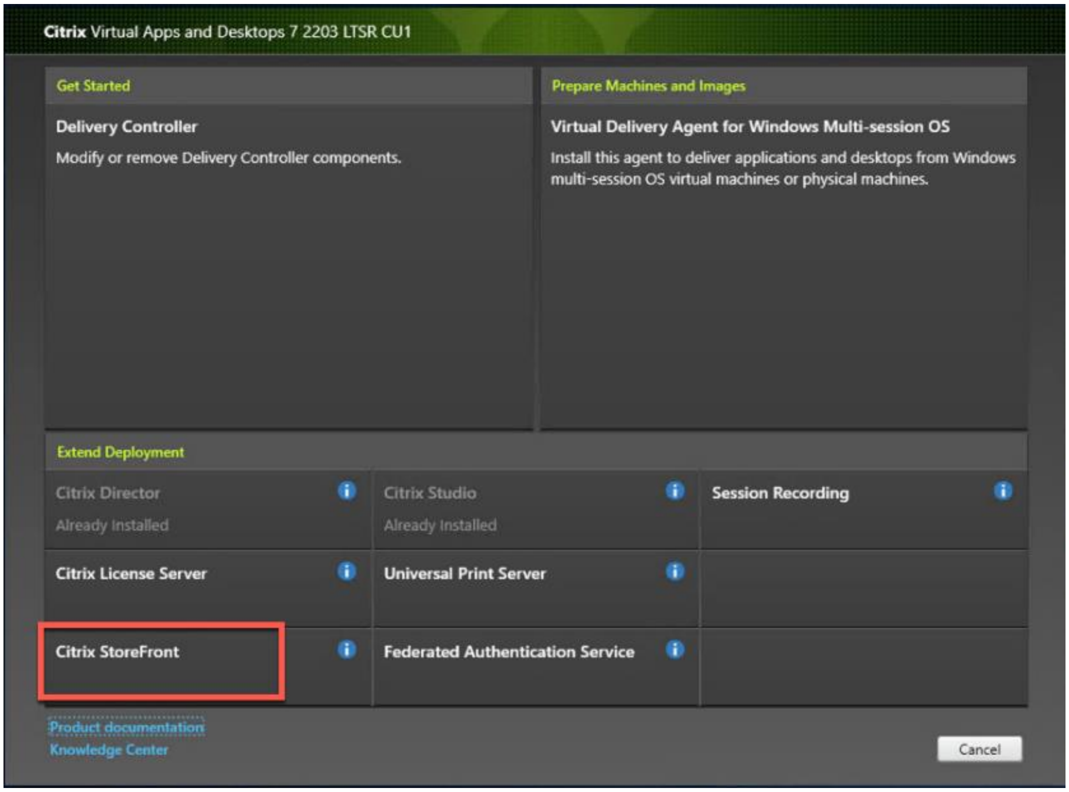

Citrix Virtual Apps and Desktops 7 2203 LTSR CU1

Get Started

Delivery Controller

Modify or remove Delivery Controller components.

Prepare Machines and Images

Virtual Delivery Agent for Windows Multi-session OS

Install this agent to deliver applications and desktops from Windows multi-session OS virtual machines or physical machines.

Extend Deployment

Citrix Director
Already Installed

Citrix Studio
Already Installed

Session Recording

Citrix License Server

Universal Print Server

Citrix StoreFront

Federated Authentication Service

Step 4. Indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement.”

Step 5. Click Next.

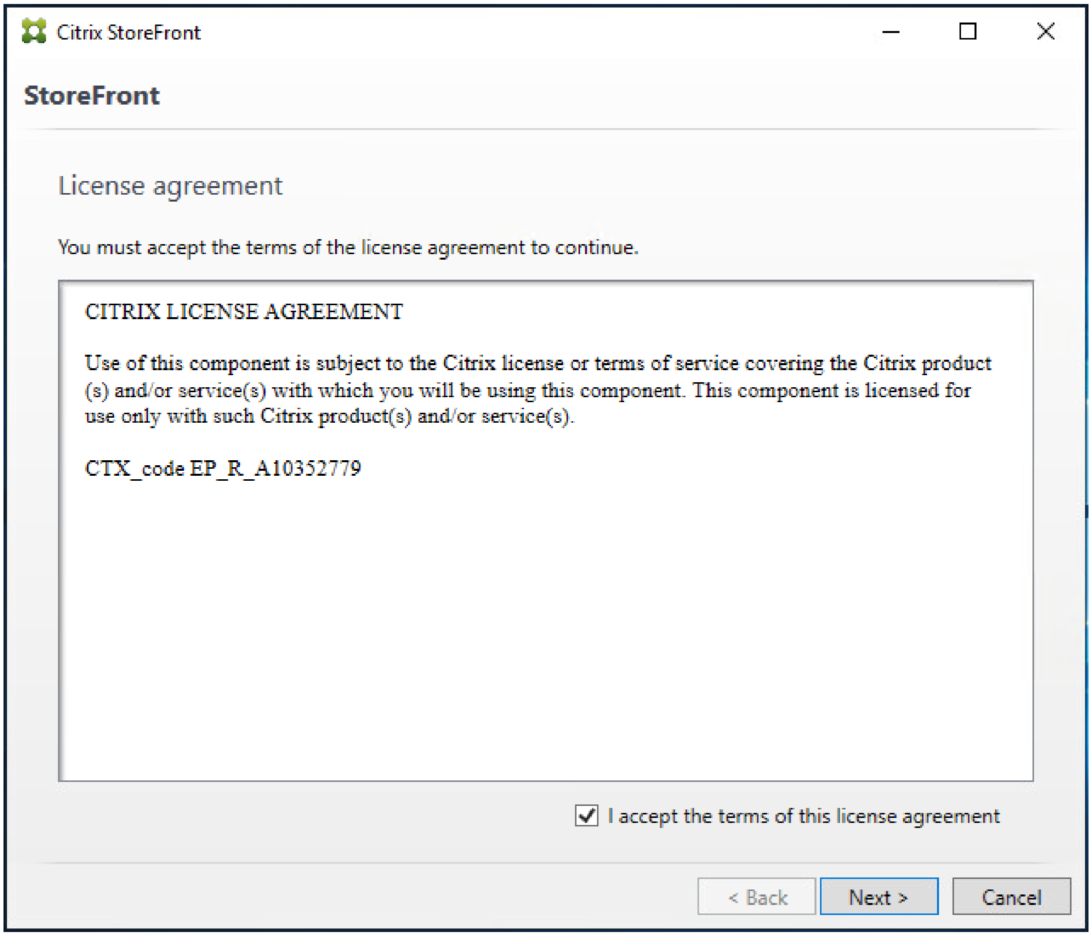

License agreement

You must accept the terms of the license agreement to continue.

CITRIX LICENSE AGREEMENT

Use of this component is subject to the Citrix license or terms of service covering the Citrix product(s) and/or service(s) with which you will be using this component. This component is licensed for use only with such Citrix product(s) and/or service(s).

CTX_code EP_R_A10352779

☑ I accept the terms of this license agreement

Step 6. On Prerequisites page click Next.

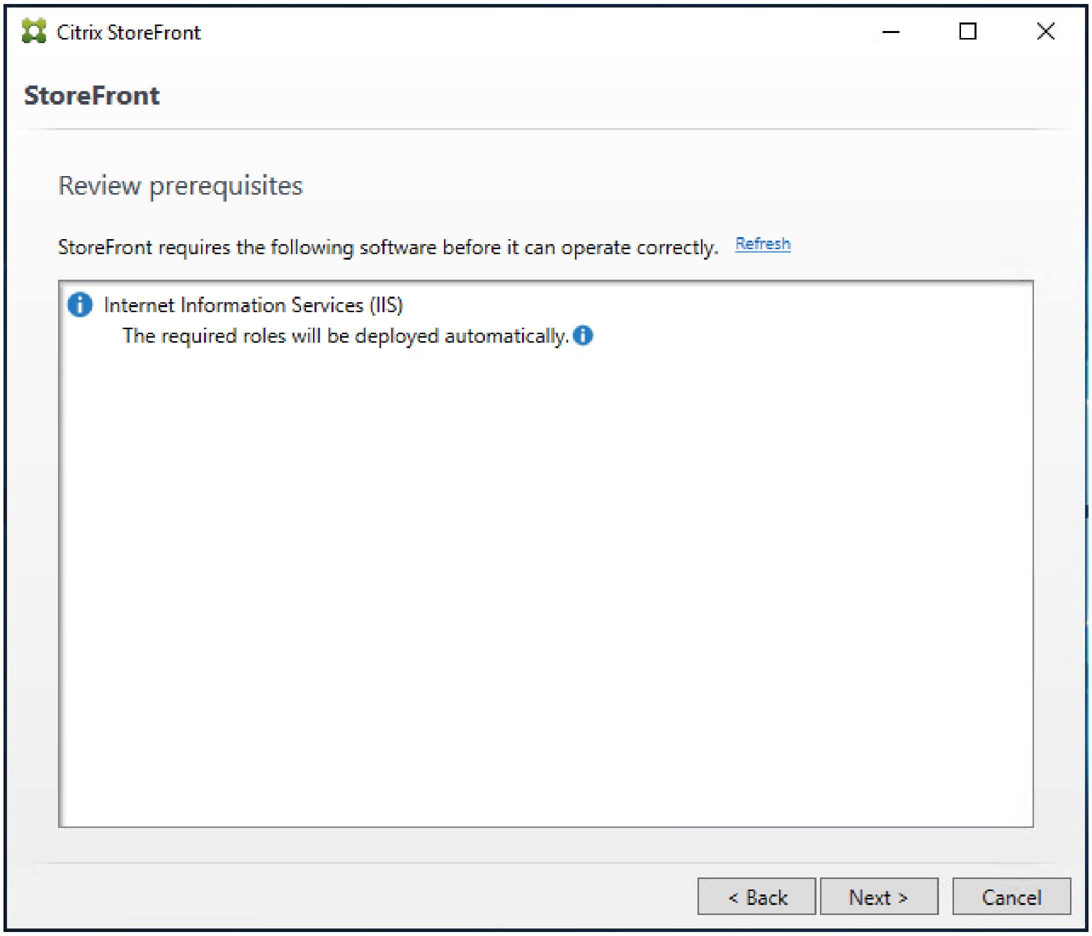

Review prerequisites

StoreFront requires the following software before it can operate correctly. Refresh

Internet Information Services (IIS)
The required roles will be deployed automatically.

Step 7. Click Install.

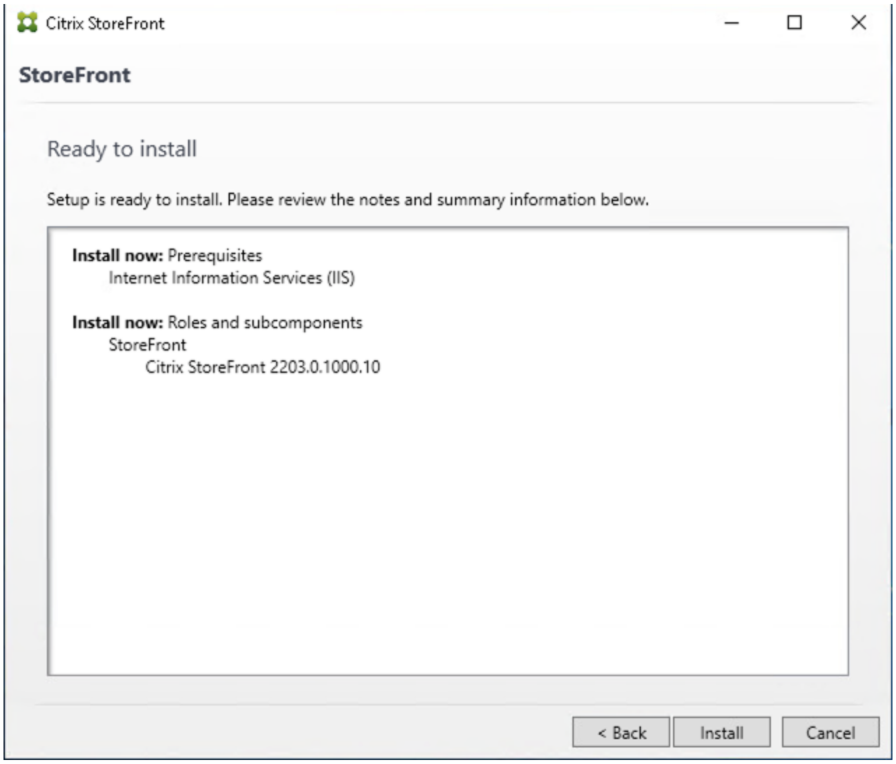

Ready to install

Setup is ready to install. Please review the notes and summary information below.

Install now: Prerequisites
Internet Information Services (IIS)

Install now: Roles and subcomponents
StoreFront
Citrix StoreFront 2203.0.1000.10

Step 8. Click Finish.

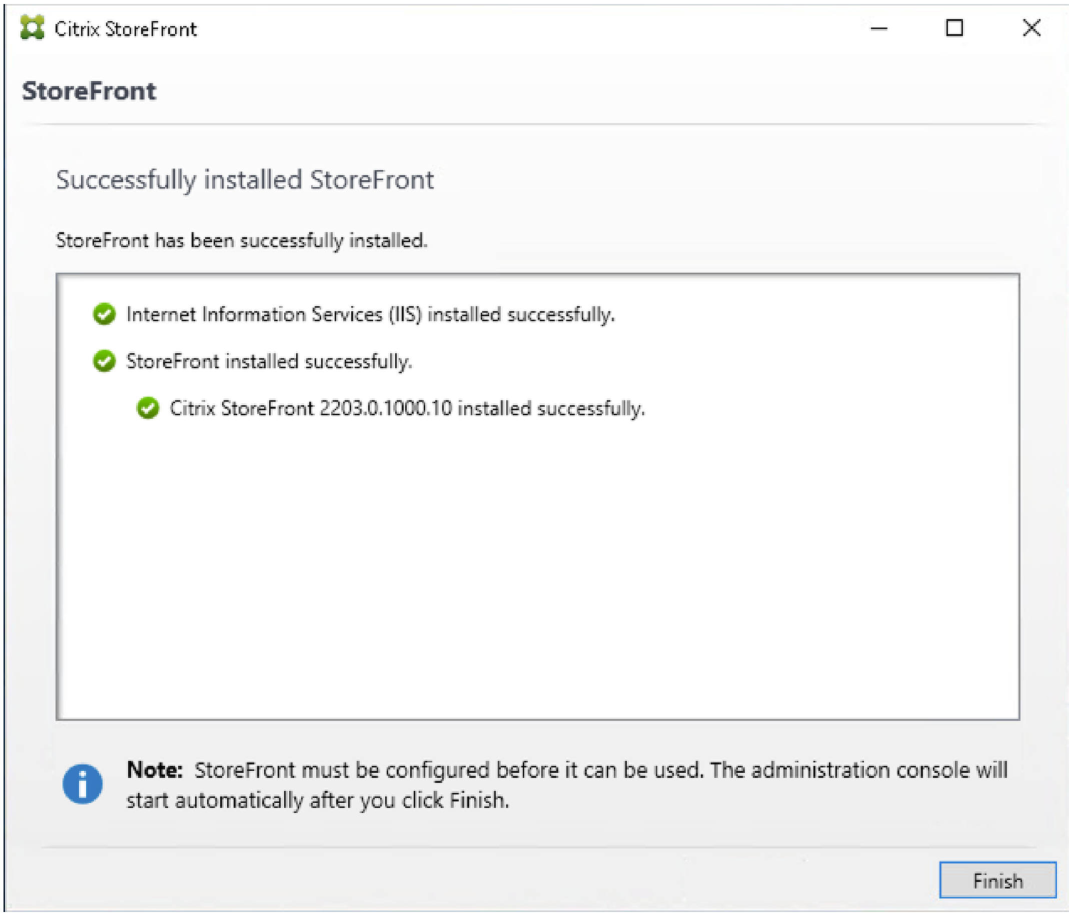

StoreFront

Successfully installed StoreFront

StoreFront has been successfully installed.

✓ Internet Information Services (IIS) installed successfully.

✓ StoreFront installed successfully.

✓ Citrix StoreFront 2203.0.1000.10 installed successfully.

Note: StoreFront must be configured before it can be used. The administration console will start automatically after you click Finish.

Step 9. Click Yes to reboot the server.

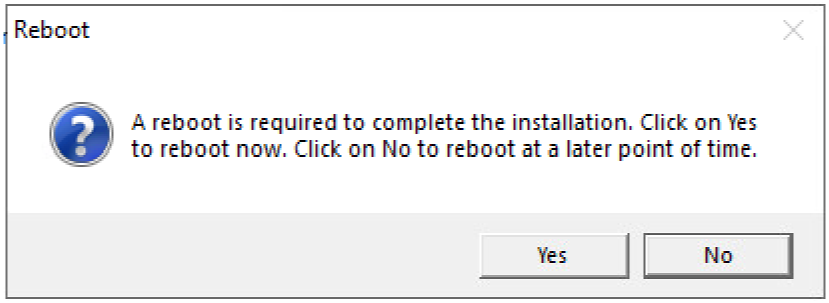

Reboot

A reboot is required to complete the installation. Click on Yes to reboot now. Click on No to reboot at a later point of time.

Step 10. Open the StoreFront Management Console.

Step 11. Click Create a new deployment.

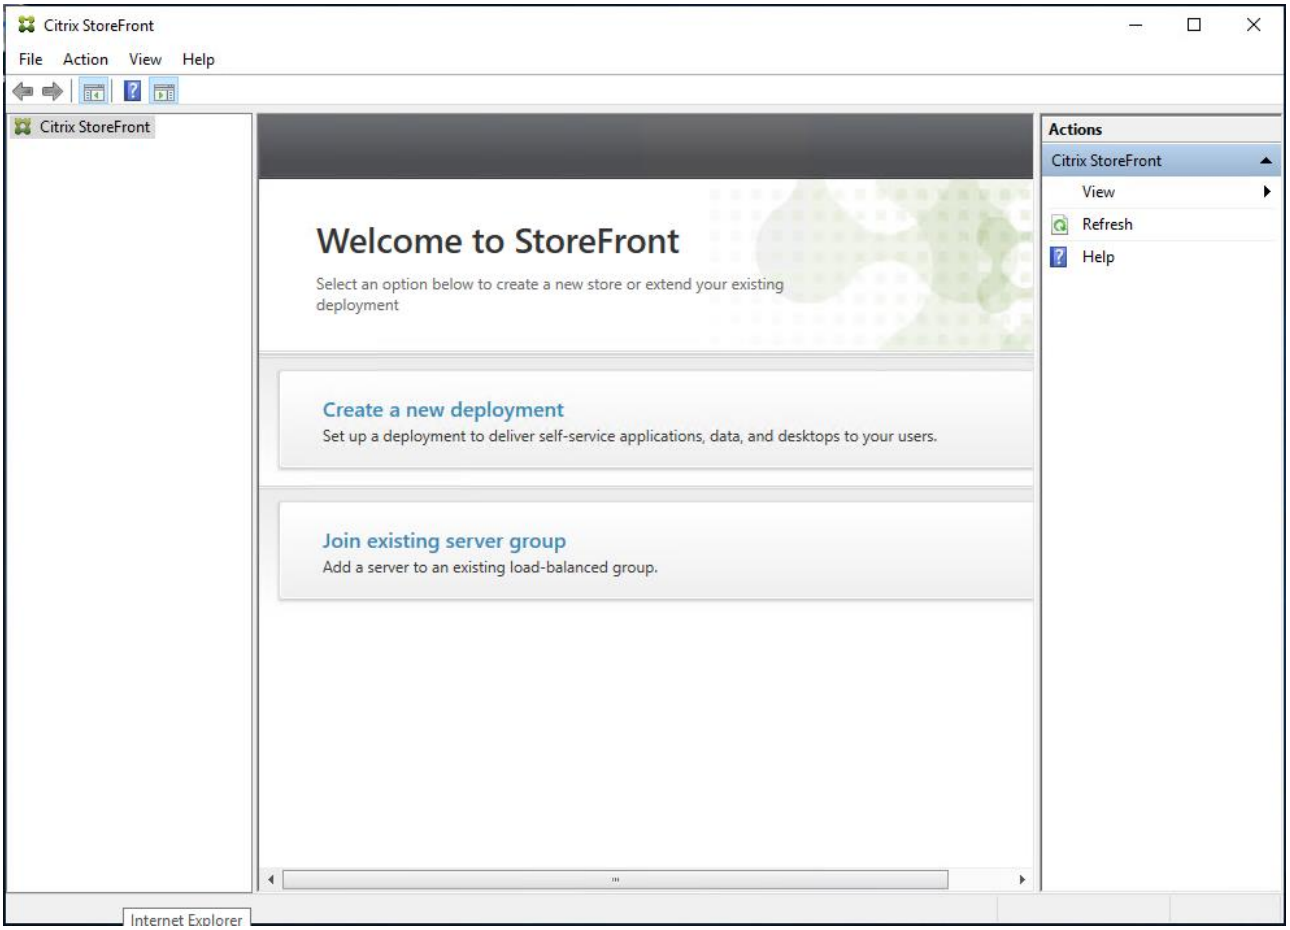

Welcome to StoreFront

Select an option below to create a new store or extend your existing deployment

Create a new deployment

Set up a deployment to deliver self-service applications, data, and desktops to your users.

Join existing server group

Add a server to an existing load-balanced group.

Step 12. Specify name for your Base URL.

Step 13. Click Next.

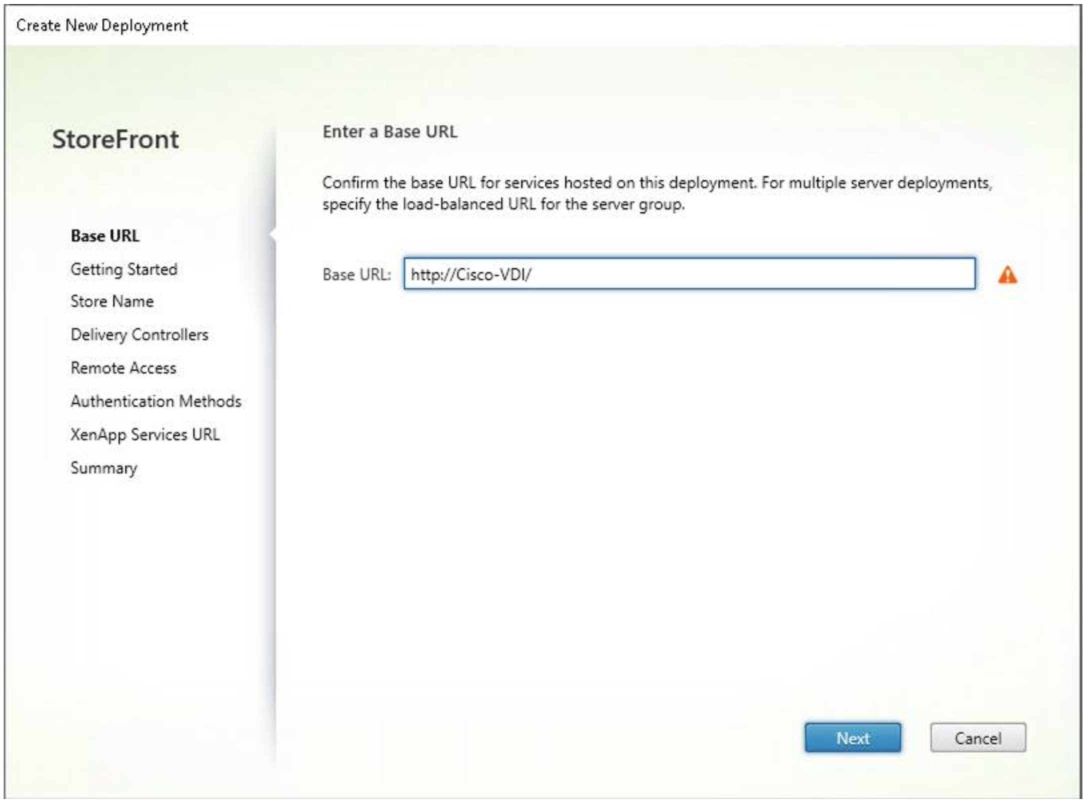

Create New Deployment

StoreFront

Enter a Base URL

Confirm the base URL for services hosted on this deployment. For multiple server deployments, specify the load-balanced URL for the server group.

Base URL: http://Cisco-VDI/

Note: For a multiple server deployment use the load balancing environment in the Base URL box.

Step 14. Click Next.

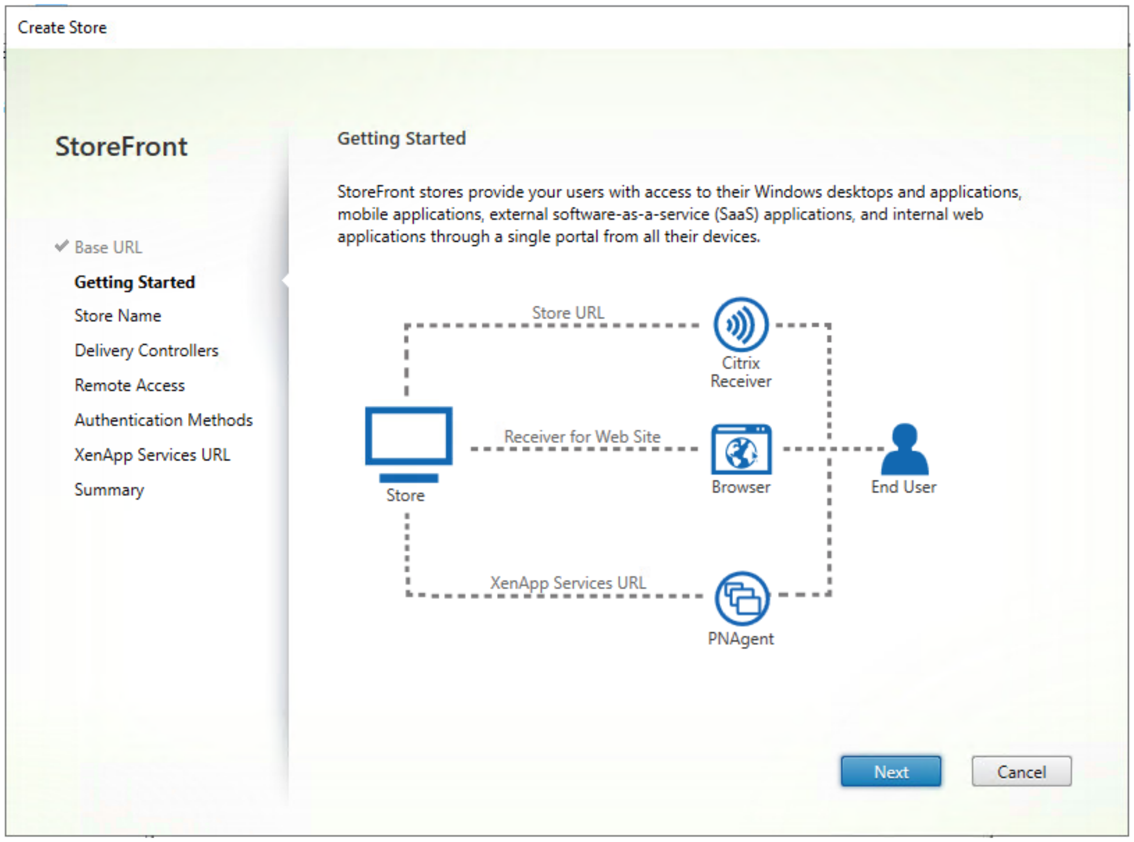

Create Store

StoreFront

- ✓ Base URL
- **Getting Started**
- Store Name
- Delivery Controllers
- Remote Access
- Authentication Methods
- XenApp Services URL
- Summary

Getting Started

StoreFront stores provide your users with access to their Windows desktops and applications, mobile applications, external software-as-a-service (SaaS) applications, and internal web applications through a single portal from all their devices.

Step 15. Specify a name for your store.

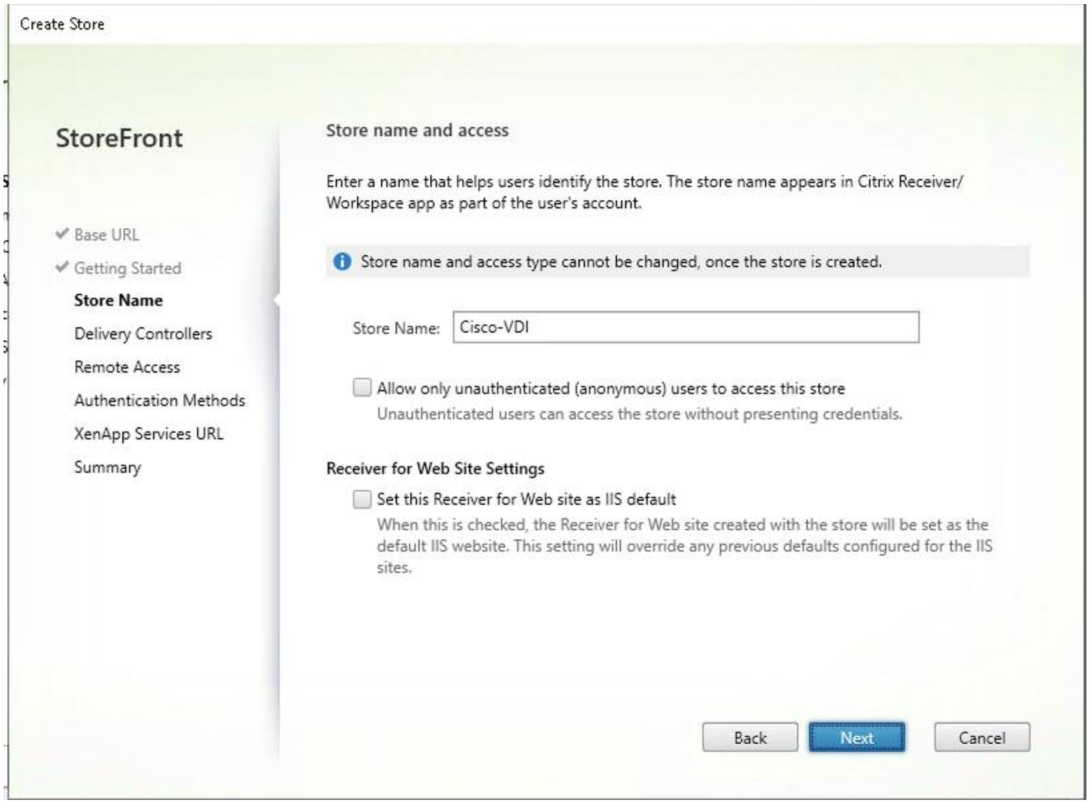

Create Store

StoreFront

- ✓ Base URL
- ✓ Getting Started
- **Store Name**
- Delivery Controllers
- Remote Access
- Authentication Methods
- XenApp Services URL
- Summary

Store name and access

Enter a name that helps users identify the store. The store name appears in Citrix Receiver/ Workspace app as part of the user's account.

Store name and access type cannot be changed, once the store is created.

Store Name: Cisco-VDI

□ Allow only unauthenticated (anonymous) users to access this store
Unauthenticated users can access the store without presenting credentials.

Receiver for Web Site Settings

□ Set this Receiver for Web site as IIS default
When this is checked, the Receiver for Web site created with the store will be set as the default IIS website. This setting will override any previous defaults configured for the IIS sites.

Step 16. Click Add to specify Delivery controllers for your new Store.

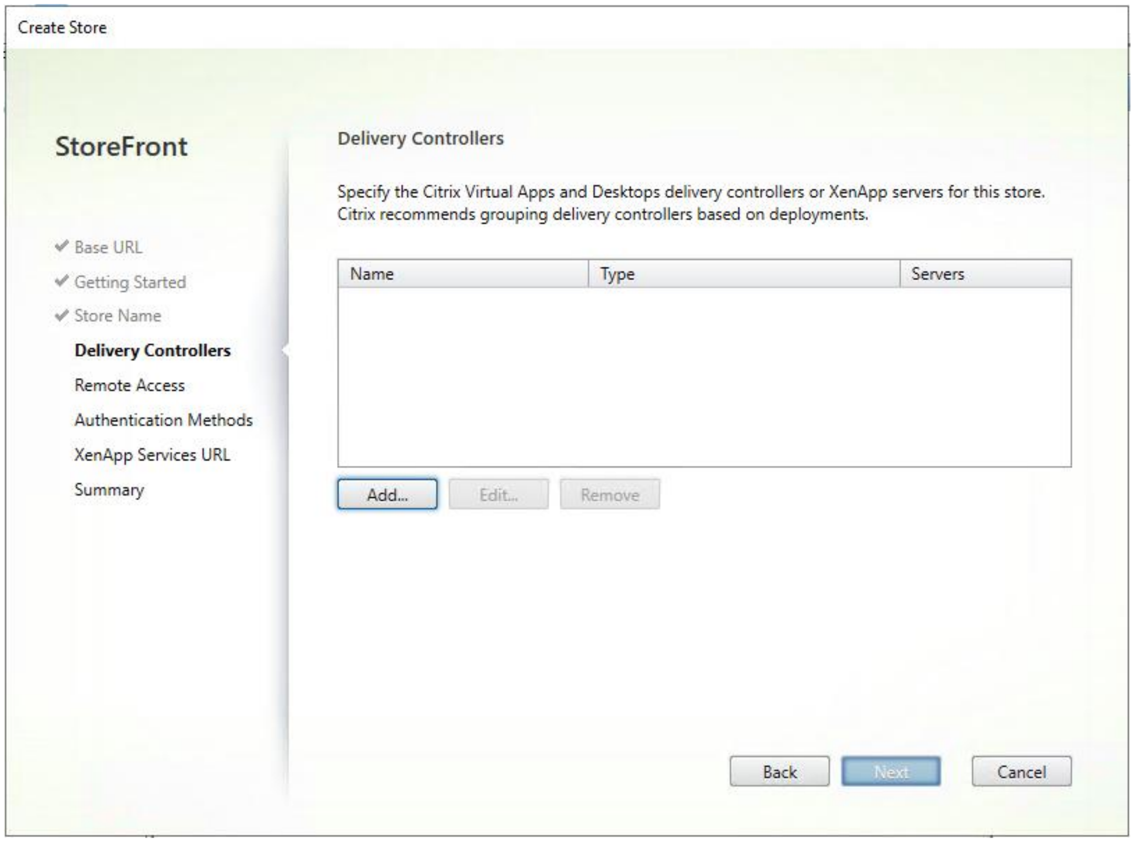

Create Store

StoreFront

Delivery Controllers

Specify the Citrix Virtual Apps and Desktops delivery controllers or XenApp servers for this store. Citrix recommends grouping delivery controllers based on deployments.

Name	Type	Servers

Step 17. Add the required Delivery Controllers to the store.

Step 18. Click OK.

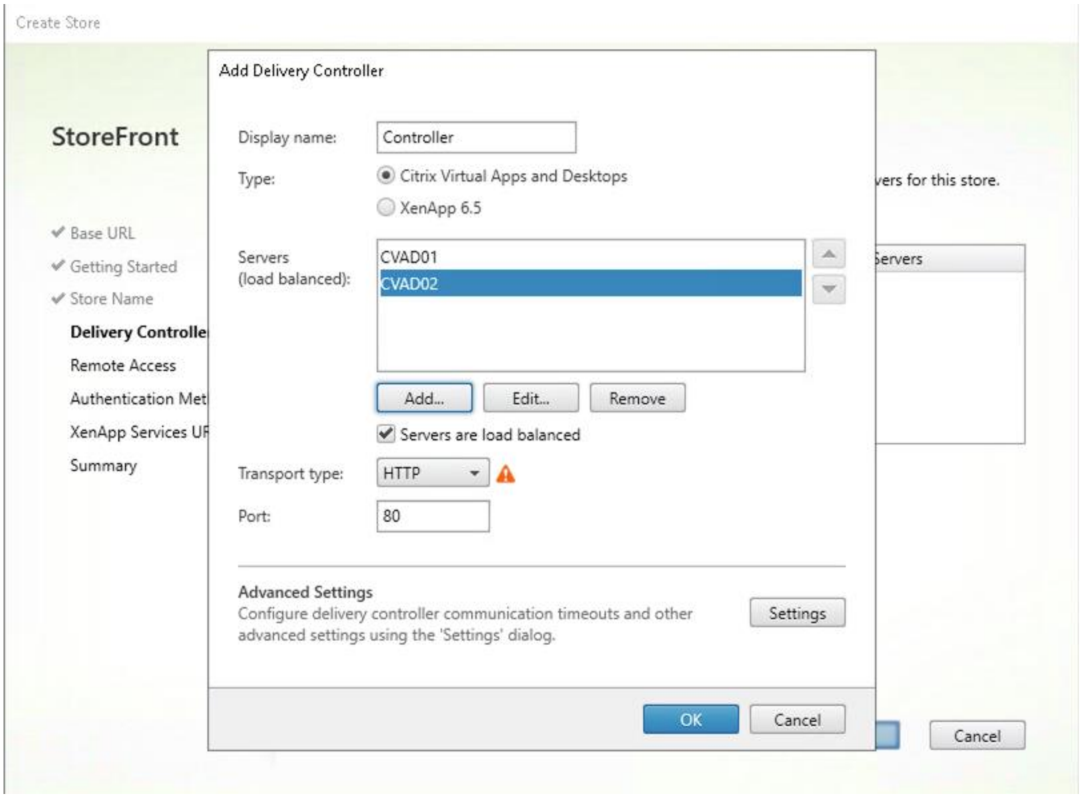

Add Delivery Controller

Display name: Controller

Type: ● Citrix Virtual Apps and Desktops ○ XenApp 6.5

Servers (load balanced):

- CVAD01
- CVAD02

☑ Servers are load balanced

Transport type: HTTP

Port: 80

Advanced Settings

Configure delivery controller communication timeouts and other advanced settings using the 'Settings' dialog.

Step 19. Click Next.

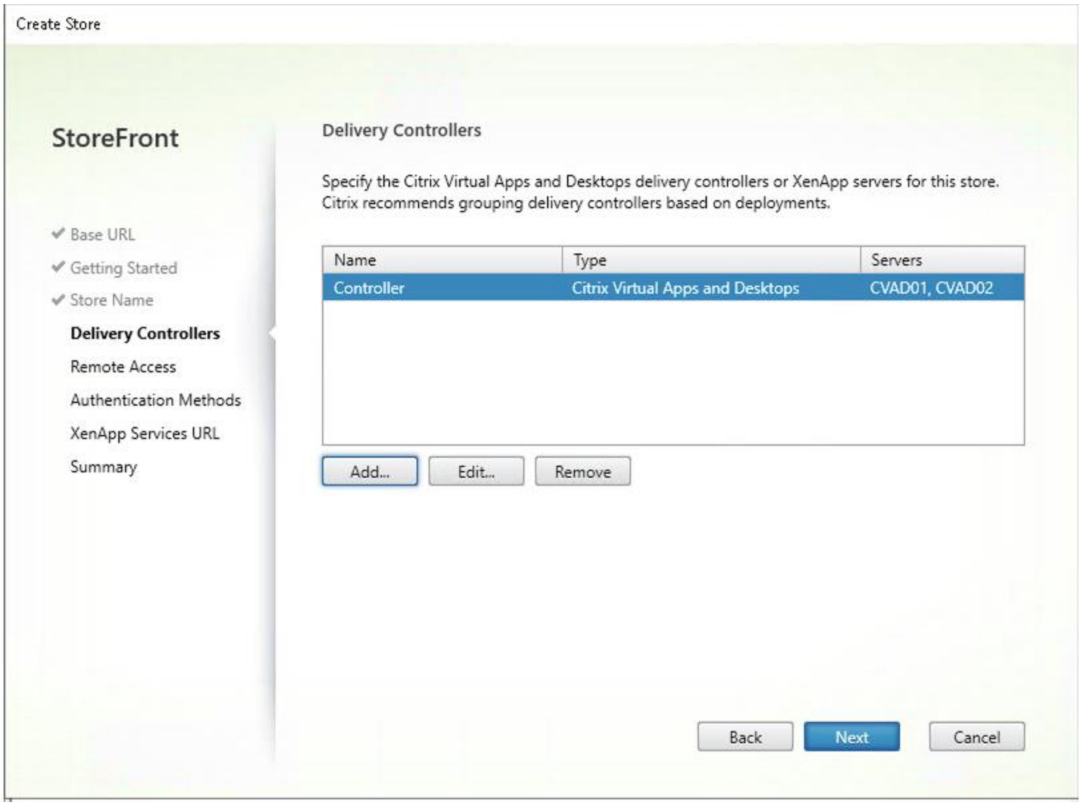

Create Store

Delivery Controllers

Specify the Citrix Virtual Apps and Desktops delivery controllers or XenApp servers for this store.
Citrix recommends grouping delivery controllers based on deployments.

Name	Type	Servers
Controller	Citrix Virtual Apps and Desktops	CVAD01, CVAD02

Step 20. Specify how connecting users can access the resources, in this environment only local users on the internal network are able to access the store.

Step 21. Click Next.

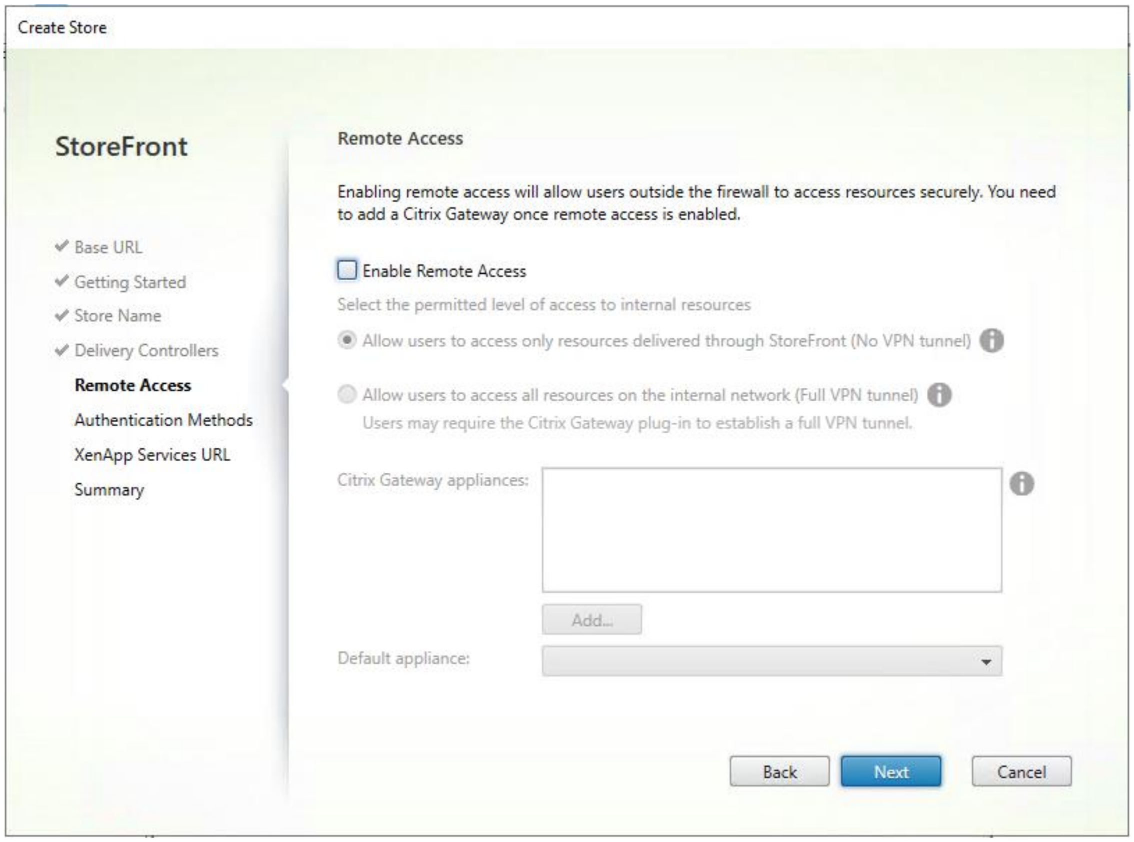

Create Store

StoreFront

- ✓ Base URL
- ✓ Getting Started
- ✓ Store Name
- ✓ Delivery Controllers
- **Remote Access**
- Authentication Methods
- XenApp Services URL
- Summary

Remote Access

Enabling remote access will allow users outside the firewall to access resources securely. You need to add a Citrix Gateway once remote access is enabled.

□ Enable Remote Access

Select the permitted level of access to internal resources

● Allow users to access only resources delivered through StoreFront (No VPN tunnel)

○ Allow users to access all resources on the internal network (Full VPN tunnel)
Users may require the Citrix Gateway plug-in to establish a full VPN tunnel.

Citrix Gateway appliances:

Default appliance:

Step 22. On the "Authentication Methods" page, select the methods your users will use to authenticate to the store. The following methods were configured in this deployment:

- Username and password: Users enter their credentials and are authenticated when they access their stores.
- Domain passthrough: Users authenticate to their domain-joined Windows computers and their credentials are used to log them on automatically when they access their stores.

Step 23. Click Next.

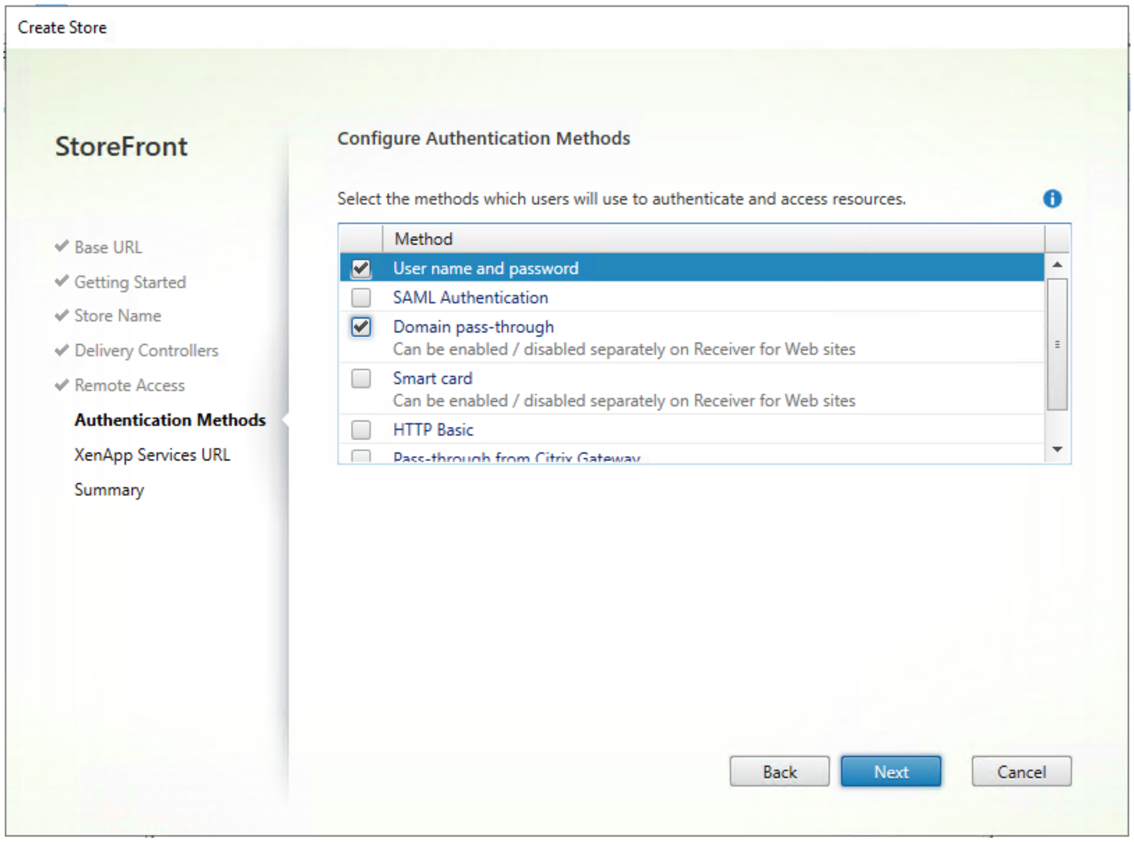

Create Store

StoreFront

Configure Authentication Methods

Select the methods which users will use to authenticate and access resources.

	Method
☑	User name and password
□	SAML Authentication
☑	Domain pass-through Can be enabled / disabled separately on Receiver for Web sites
□	Smart card Can be enabled / disabled separately on Receiver for Web sites
□	HTTP Basic
□	Pass-through from Citrix Gateway

Step 24. Configure the XenApp Service URL for users who use PNAgent to access the applications and desktops.

Step 25. Click Create.

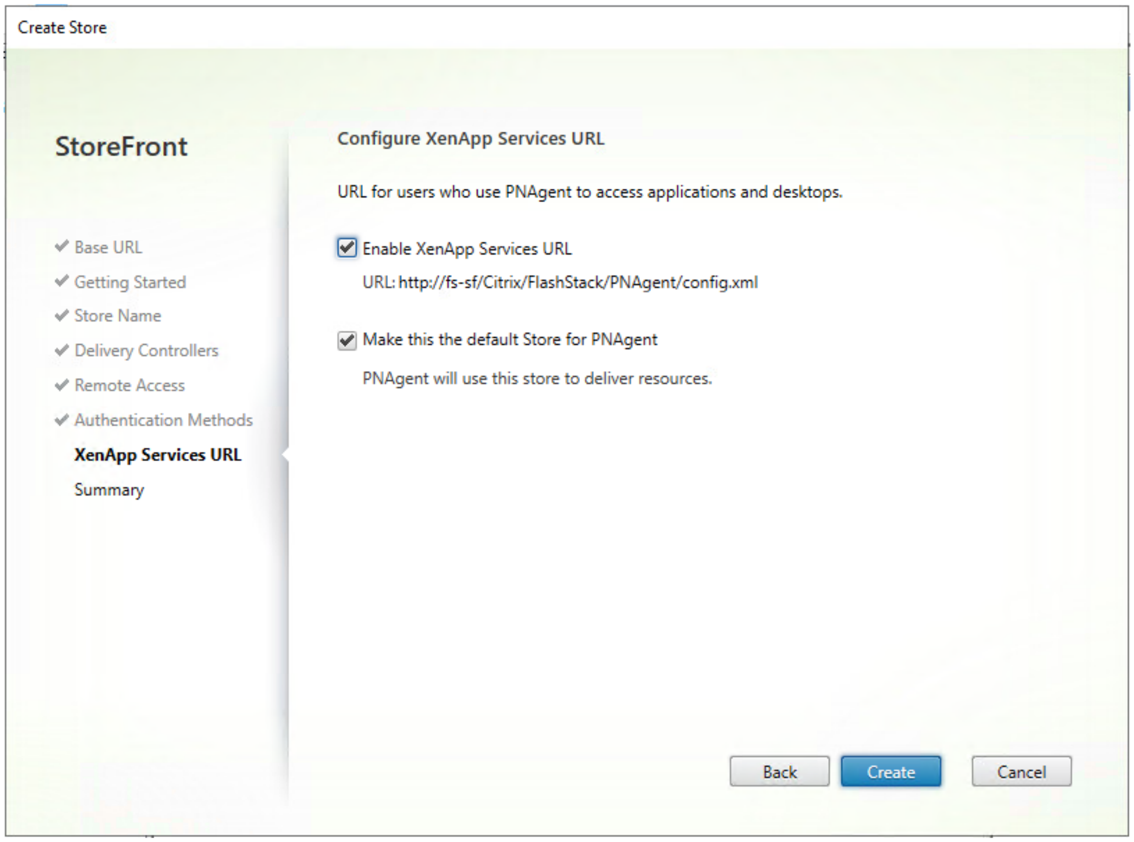
Create Store

StoreFront

✓ Base URL
✓ Getting Started
✓ Store Name
✓ Delivery Controllers
✓ Remote Access
✓ Authentication Methods
XenApp Services URL
Summary

Configure XenApp Services URL

URL for users who use PNAgent to access applications and desktops.

☑ Enable XenApp Services URL

URL: http://fs-sf/Citrix/FlashStack/PNAgent/config.xml

☑ Make this the default Store for PNAgent

PNAgent will use this store to deliver resources.

Step 26. After creating the store click Finish.

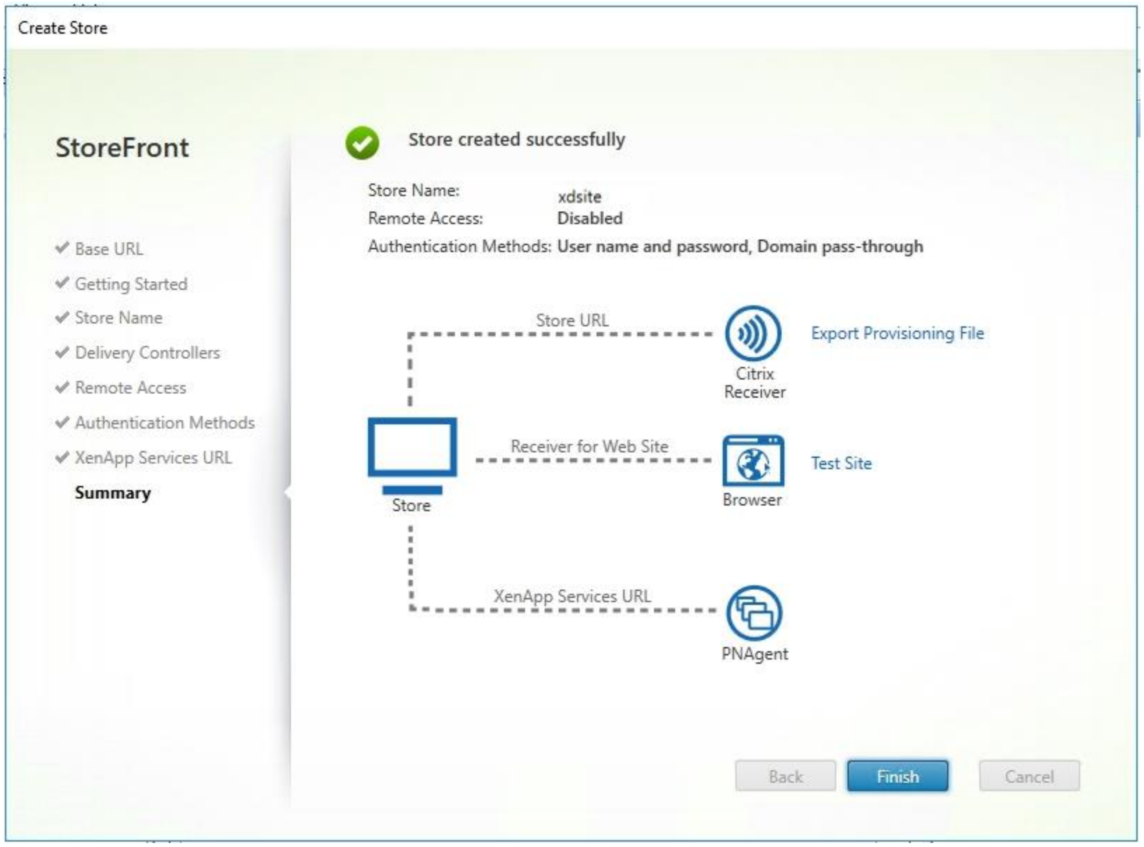
Create Store

StoreFront

✓ Base URL
✓ Getting Started
✓ Store Name
✓ Delivery Controllers
✓ Remote Access
✓ Authentication Methods
✓ XenApp Services URL
Summary

Store created successfully

Store Name:	xdsite
Remote Access:	Disabled
Authentication Methods:	User name and password, Domain pass-through

Procedure 9. Additional StoreFront Configuration

Note: After the first StoreFront server is completely configured and the Store is operational, you can add additional servers.

Step 1. Install the second StoreFront using the same installation steps outlined above.

Step 2. Connect to the first StoreFront server.

Step 3. To add the second server and generate the authorization information that allows the additional StoreFront server to join the server group, select Add Server from Actions pane in the Server Group.

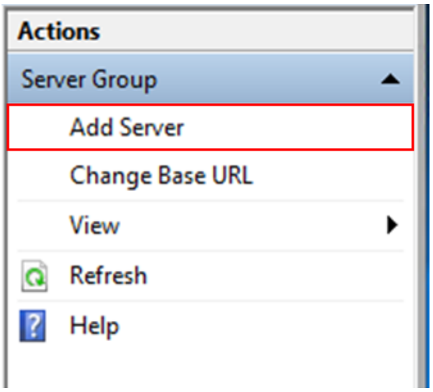

Step 4. Copy the authorization code.

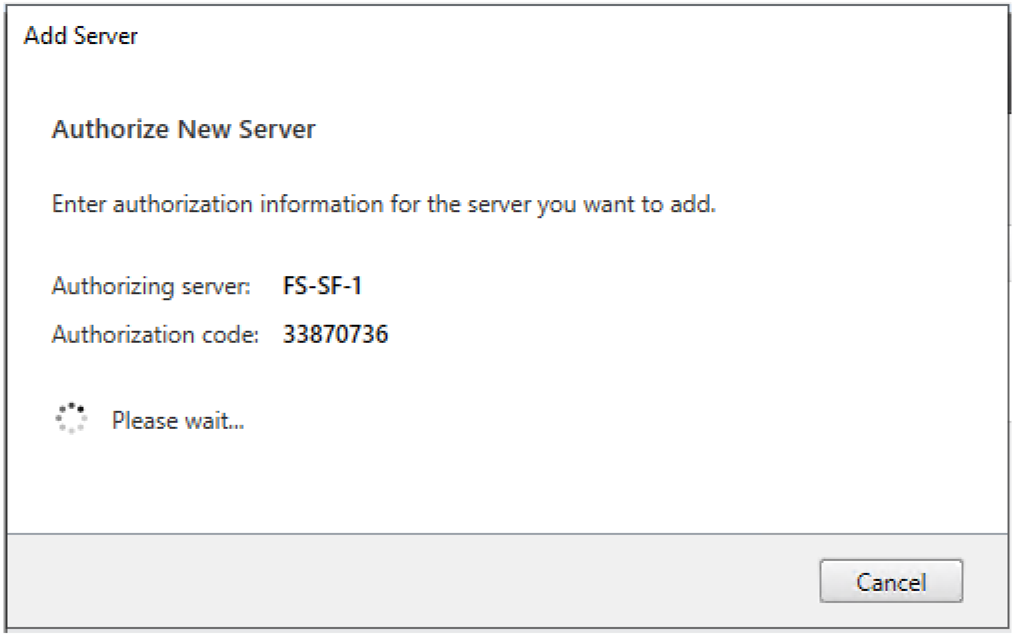
Add Server

Authorize New Server

Enter authorization information for the server you want to add.

Authorizing server: FS-SF-1

Authorization code: 33870736

Please wait...

Step 5. From the StoreFront Console on the second server select “Join existing server group.”

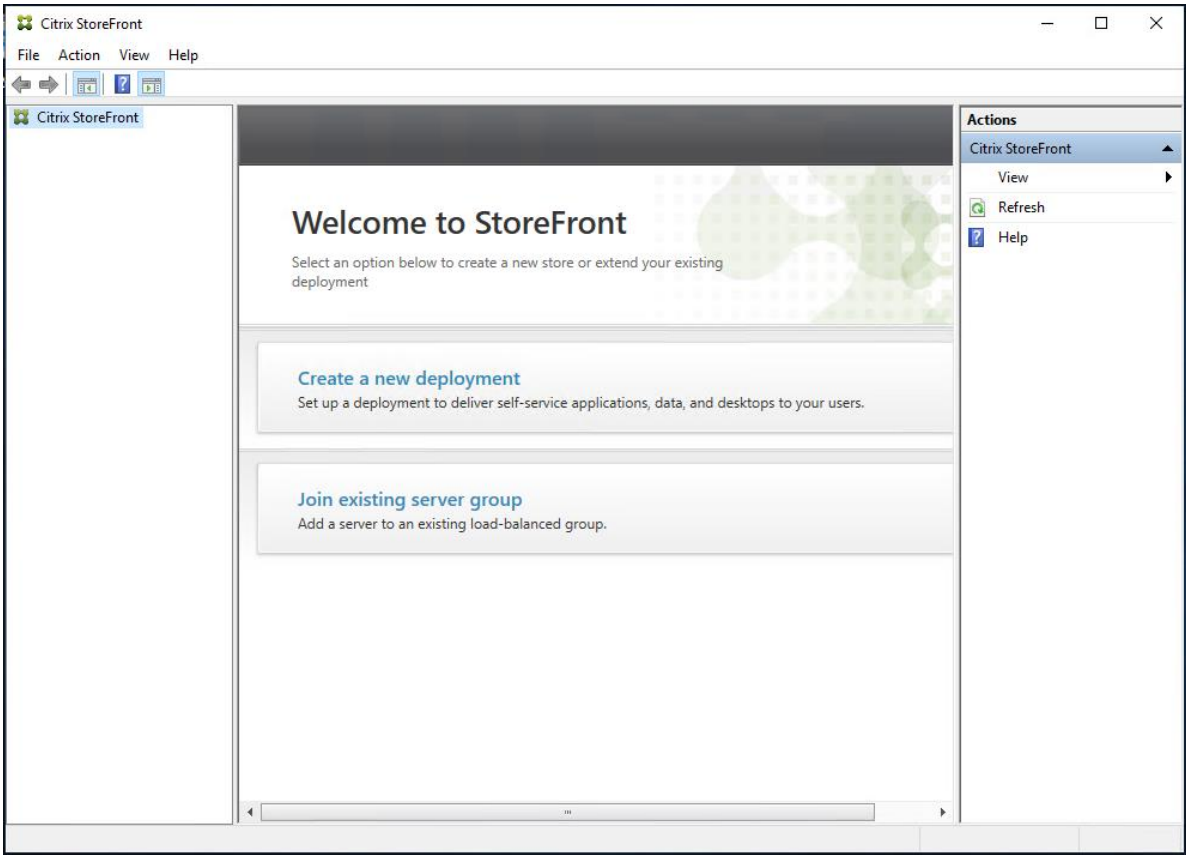

Welcome to StoreFront

Select an option below to create a new store or extend your existing deployment

Create a new deployment

Set up a deployment to deliver self-service applications, data, and desktops to your users.

Join existing server group

Add a server to an existing load-balanced group.

Step 6. In the Join Server Group dialog, enter the name of the first Storefront server and paste the Authorization code into the Join Server Group dialog.

Step 7. Click Join.

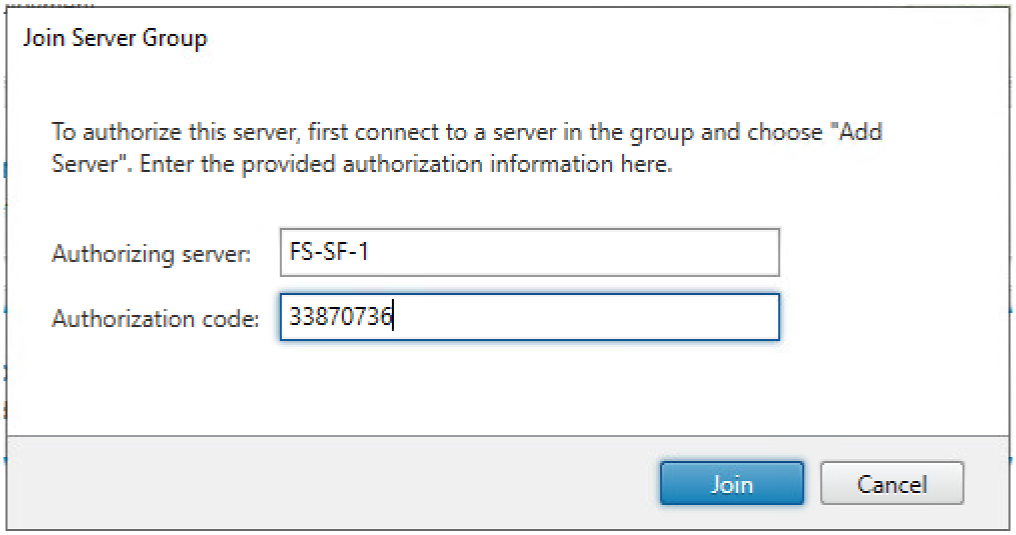

Join Server Group

To authorize this server, first connect to a server in the group and choose "Add Server". Enter the provided authorization information here.

Authorizing server: FS-SF-1

Authorization code: 33870736

Step 8. A message appears when the second server has joined successfully.

Step 9. Click OK.

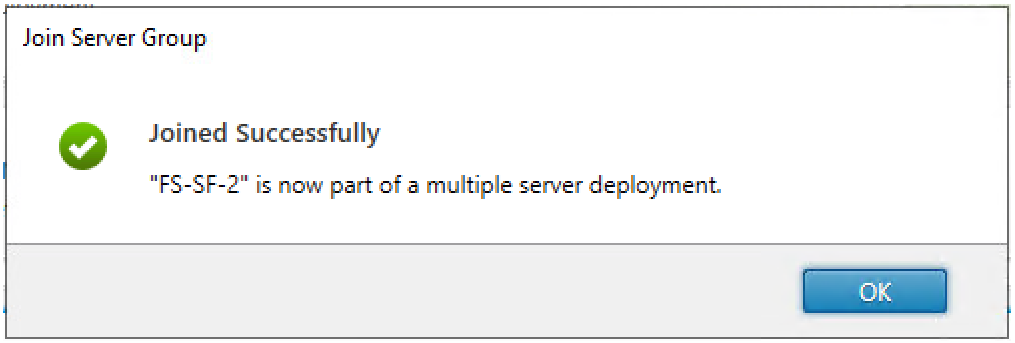

Join Server Group

Joined Successfully

"FS-SF-2" is now part of a multiple server deployment.

The second StoreFront is now in the Server Group.

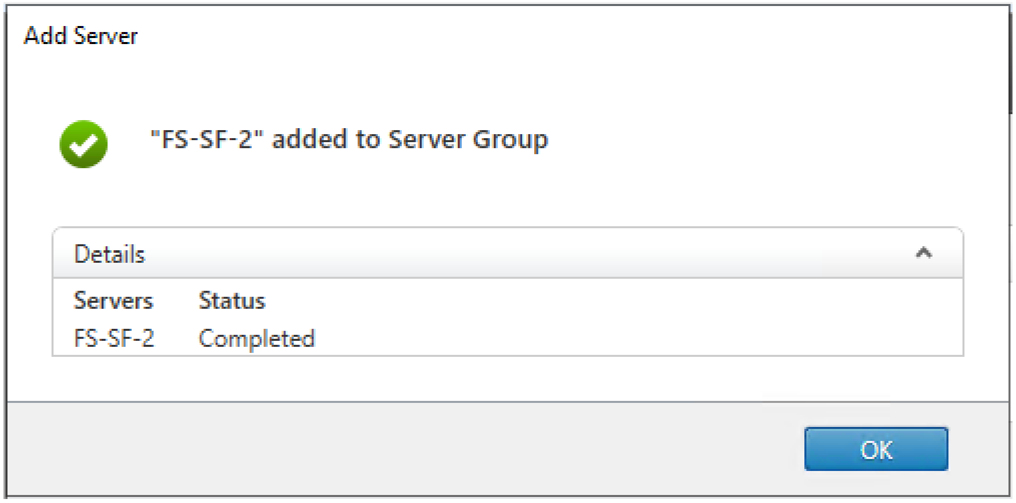

Add Server

"FS-SF-2" added to Server Group

Details

Servers	Status
FS-SF-2	Completed

Install and Configure Citrix Provisioning Server 2203

In most implementations, there is a single vDisk providing the standard image for multiple target devices. Thousands of target devices can use a single vDisk shared across multiple Provisioning Services (PVS) servers in the same farm, simplifying virtual desktop management. This section describes the installation and configuration tasks required to create a PVS implementation.

The PVS server can have many stored vDisks, and each vDisk can be several gigabytes in size. Your streaming performance and manageability can be improved using a RAID array, SAN, or NAS. PVS software and hardware requirements are available in the Provisioning Services 2203 document.

Procedure 1. Configure Prerequisites

Step 1. Set the following Scope Options on the DHCP server hosting the PVS target machines:

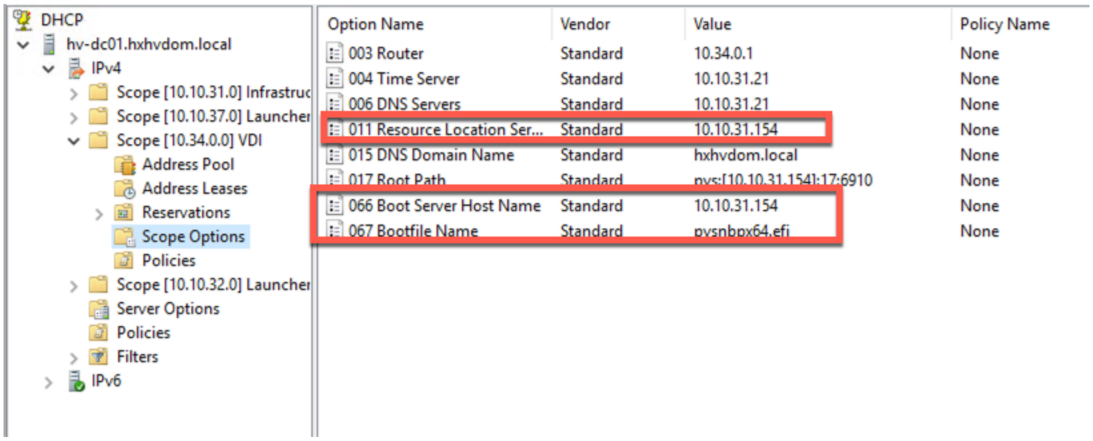

Option Name	Vendor	Value	Policy Name
003 Router	Standard	10.34.0.1	None
004 Time Server	Standard	10.10.31.21	None
006 DNS Servers	Standard	10.10.31.21	None
011 Resource Location Ser...	Standard	10.10.31.154	None
015 DNS Domain Name	Standard	hxhvdom.local	None
017 Root Path	Standard	pvs:[10.10.31.154]:17:6910	None
066 Boot Server Host Name	Standard	10.10.31.154	None
067 Bootfile Name	Standard	pvsnbpx64.efi	None

Step 2. Create a DNS host records with multiple PVS Servers IP for TFTP Load Balancing:

Step 3. As a Citrix best practice cited in this CTX article, apply the following registry setting both the PVS servers and target machines:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\TCPIP\Parameters\
Key: "DisableTaskOffload" (dword)
Value: "1"
```

Note: Only one MS SQL database is associated with a farm. You can choose to install the Provisioning Services database software on an existing SQL database, if that machine can communicate with all Provisioning Servers within the farm, or with a new SQL Express database machine, created using the SQL Express software that is free from Microsoft.

The following databases are supported: Microsoft SQL Server 2008 SP3 through 2019 (x86, x64, and Express editions). Please check Citrix documentation for further reference.

Note: Microsoft SQL 2019 was installed separately for this CVD.

Procedure 2. Install and Configure Citrix Provisioning Service 2203

Step 1. Connect to Citrix Provisioning server and launch Citrix Provisioning Services 2203 ISO and let AutoRun launch the installer.

Step 2. Click Console Installation.

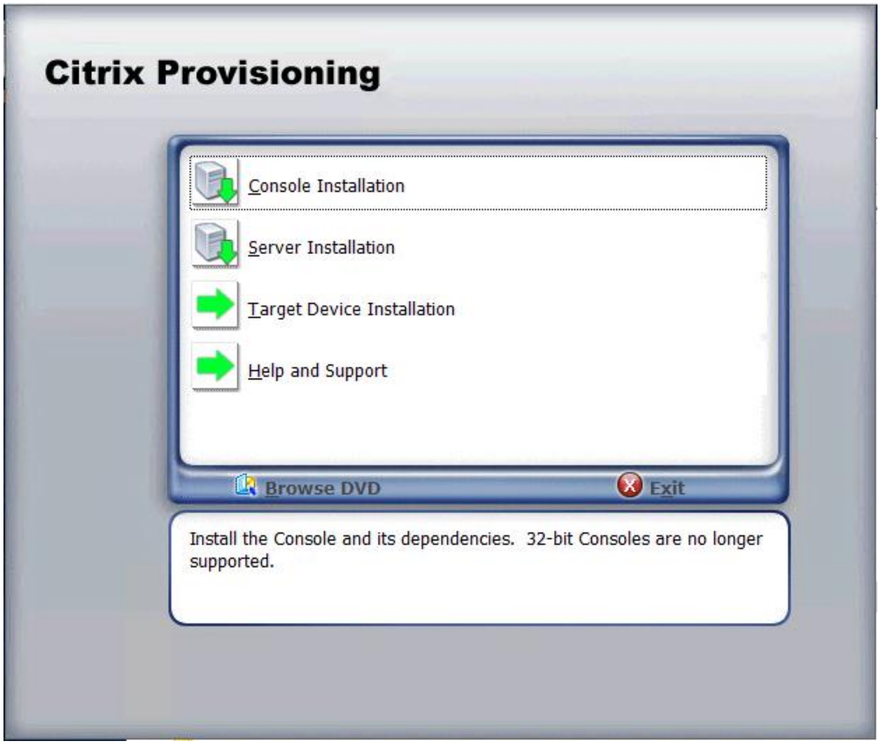

Install the Console and its dependencies. 32-bit Consoles are no longer supported.

Step 3. Click Install to start the console installation.

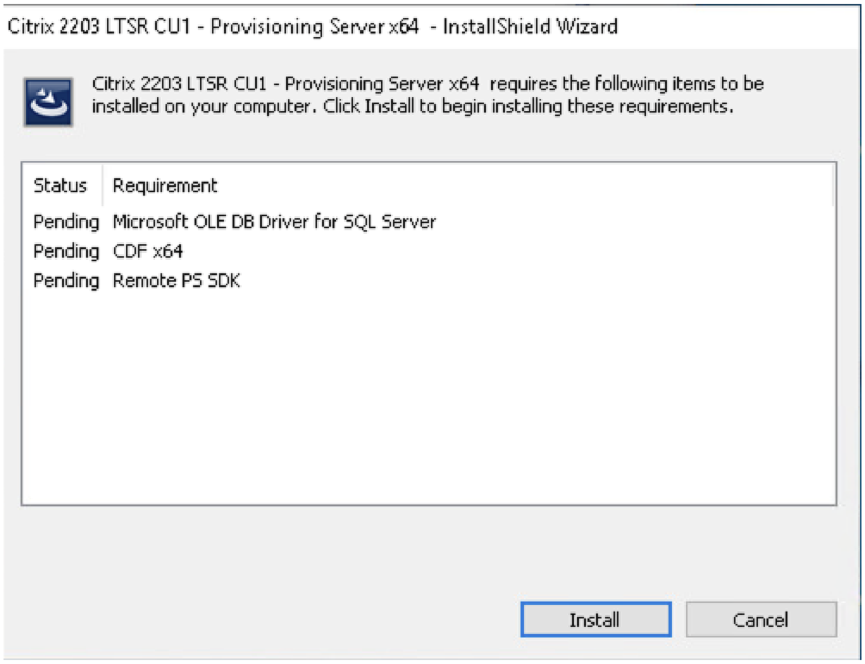
Citrix 2203 LTSR CU1 - Provisioning Server x64 - InstallShield Wizard
Install
Cancel

Citrix 2203 LTSR CU1 - Provisioning Server x64 requires the following items to be installed on your computer. Click Install to begin installing these requirements.

Status	Requirement
Pending	Microsoft OLE DB Driver for SQL Server
Pending	CDF x64
Pending	Remote PS SDK

Step 4. Read the .NET License Agreement. If acceptable, check “I have read and accept the license terms.”

Step 5. Click Next.

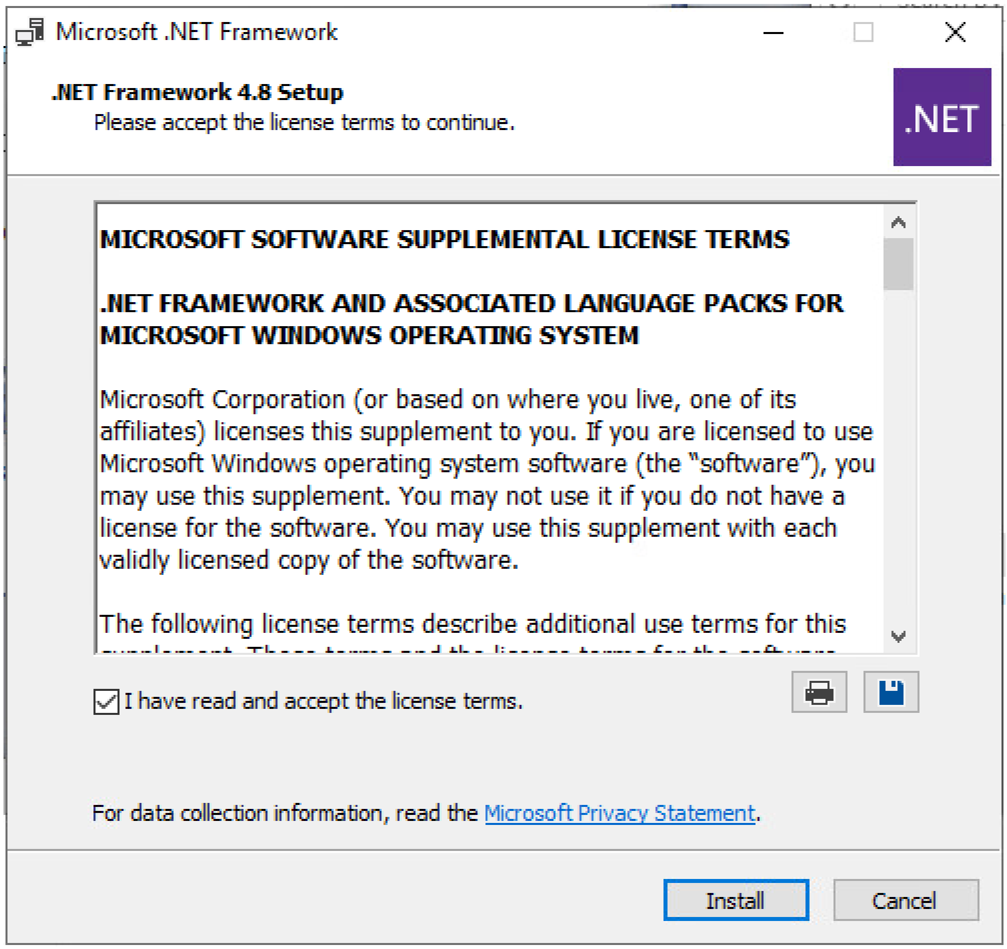

.NET Framework 4.8 Setup

Please accept the license terms to continue.

MICROSOFT SOFTWARE SUPPLEMENTAL LICENSE TERMS

.NET FRAMEWORK AND ASSOCIATED LANGUAGE PACKS FOR MICROSOFT WINDOWS OPERATING SYSTEM

Microsoft Corporation (or based on where you live, one of its affiliates) licenses this supplement to you. If you are licensed to use Microsoft Windows operating system software (the "software"), you may use this supplement. You may not use it if you do not have a license for the software. You may use this supplement with each validly licensed copy of the software.

The following license terms describe additional use terms for this

☑ I have read and accept the license terms.

For data collection information, read the Microsoft Privacy Statement.

Step 6. Click Finish.

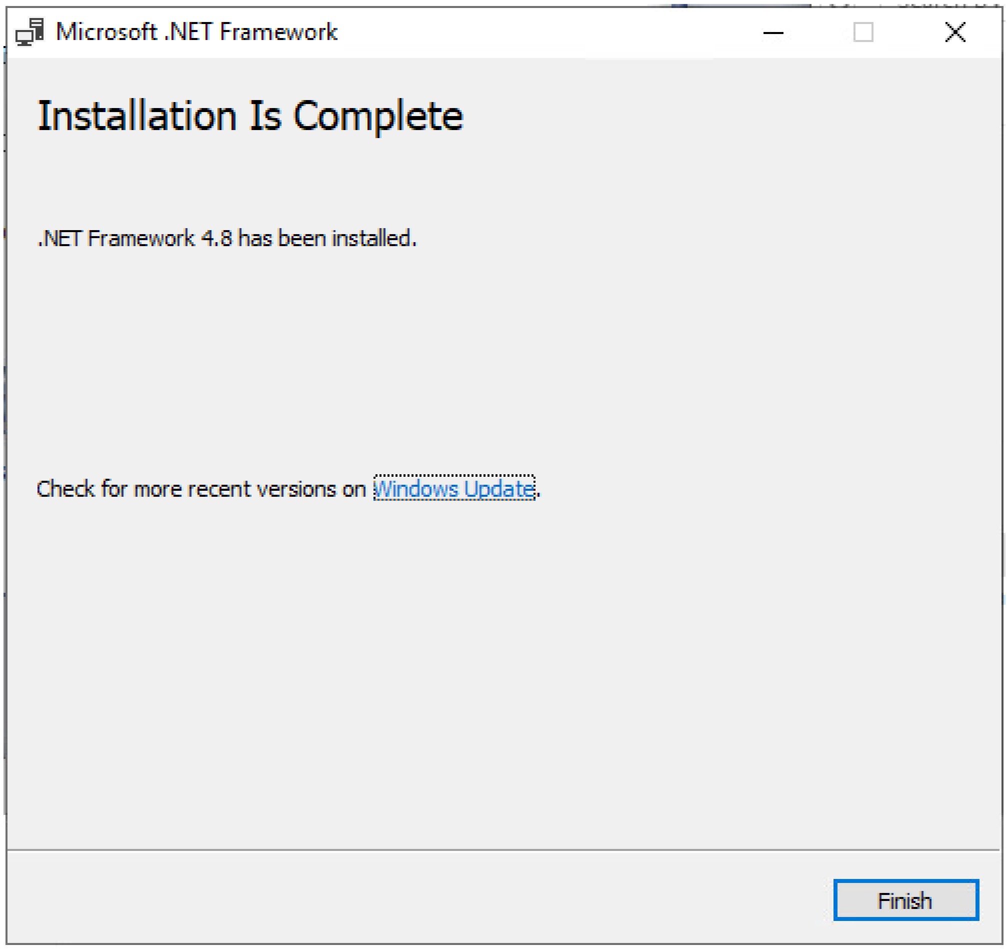

Microsoft .NET Framework

Installation Is Complete

.NET Framework 4.8 has been installed.

Check for more recent versions on Windows Update.

Step 7. Restart the Virtual Machine.

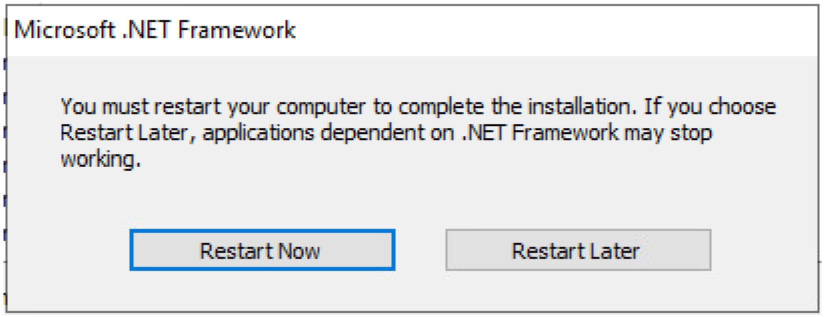

Microsoft .NET Framework

You must restart your computer to complete the installation. If you choose Restart Later, applications dependent on .NET Framework may stop working.

Step 8. Logging into the Operating system automatically launches the installation wizard.

Step 9. Click Next.

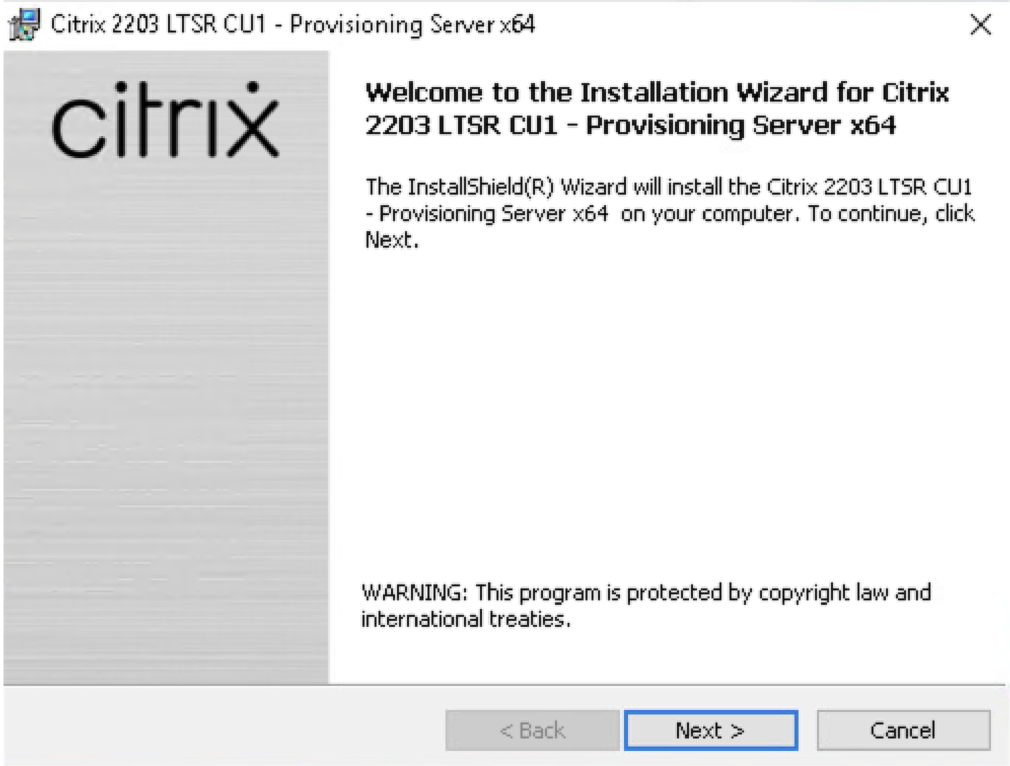
Citrix 2203 LTSR CU1 - Provisioning Server x64

Welcome to the Installation Wizard for Citrix 2203 LTSR CU1 - Provisioning Server x64

The InstallShield(R) Wizard will install the Citrix 2203 LTSR CU1 - Provisioning Server x64 on your computer. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

Step 10. Read the Citrix License Agreement. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”

Step 11. Click Next.

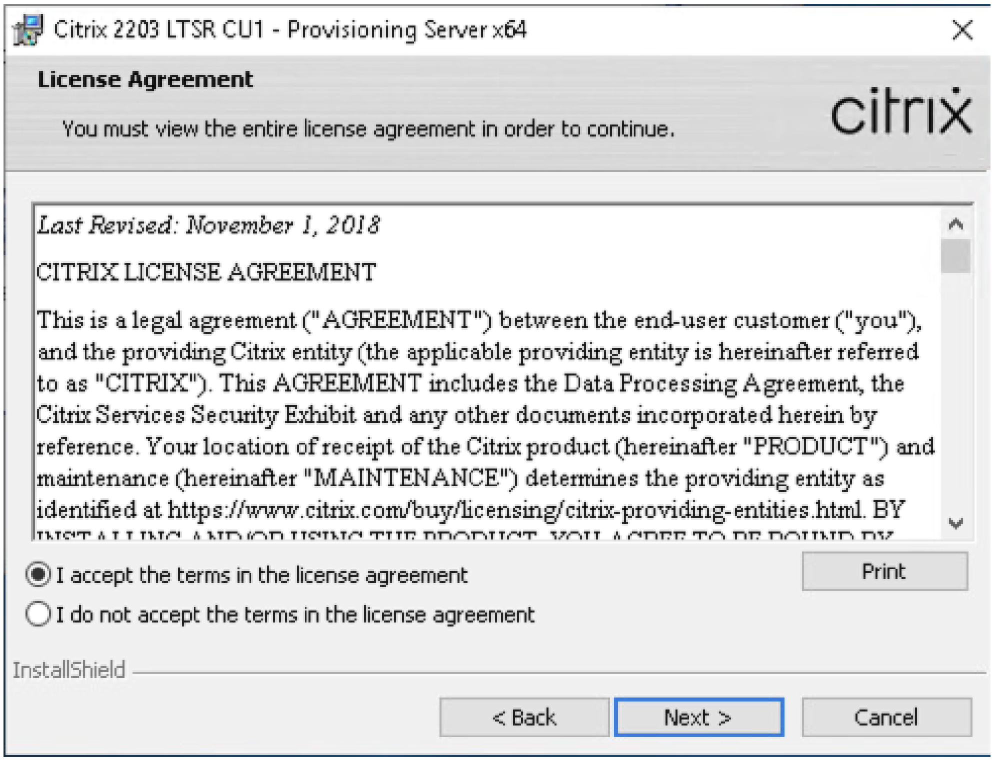
Citrix 2203 LTSR CU1 - Provisioning Server x64

License Agreement

You must view the entire license agreement in order to continue.

Last Revised: November 1, 2018

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). This AGREEMENT includes the Data Processing Agreement, the Citrix Services Security Exhibit and any other documents incorporated herein by reference. Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 12. Optionally, provide User Name and Organization.

Step 13. Click Next.

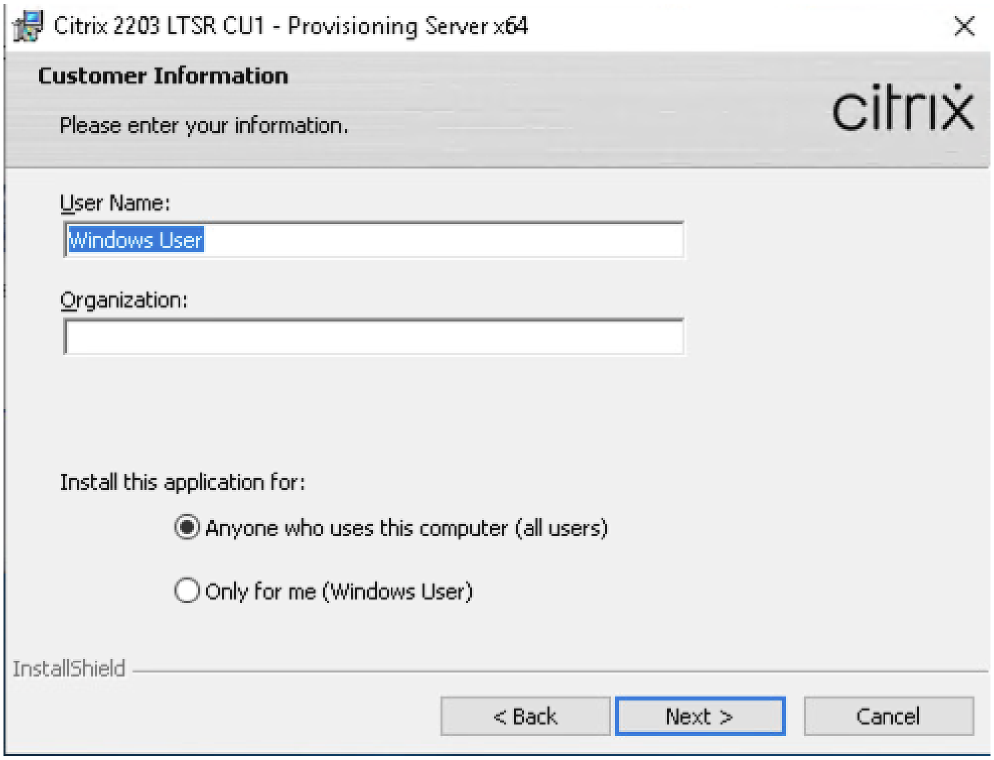

Citrix 2203 LTSR CU1 - Provisioning Server x64

Customer Information

Please enter your information.

User Name: Windows User

Organization:

Install this application for:

● Anyone who uses this computer (all users)

○ Only for me (Windows User)

Step 14. Accept the default path.

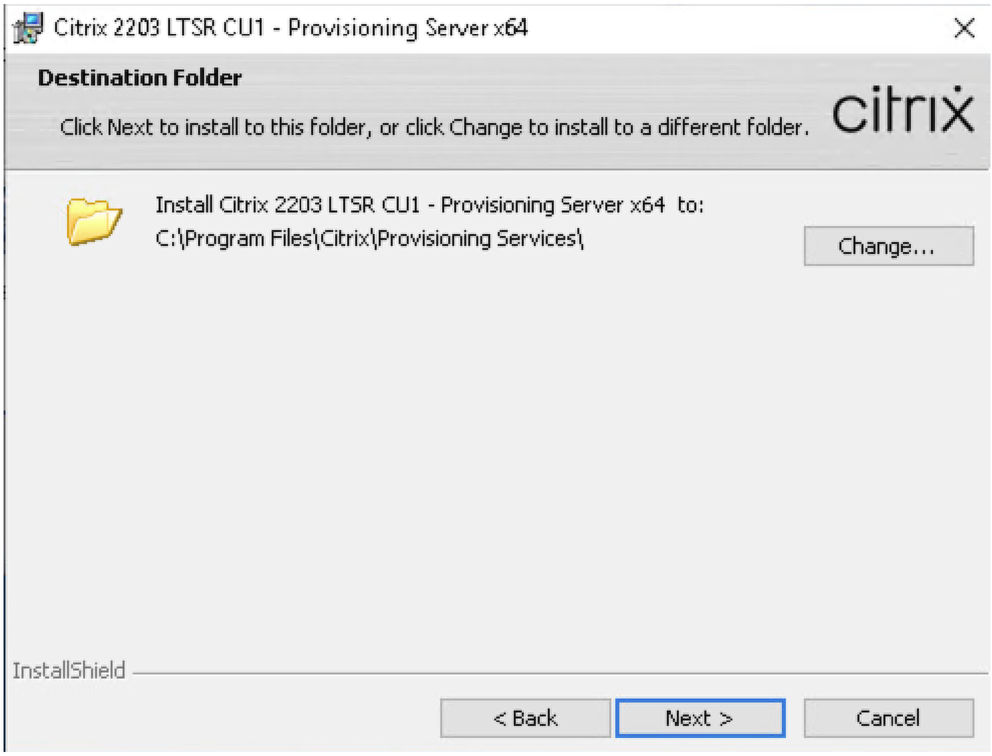

Citrix 2203 LTSR CU1 - Provisioning Server x64

Destination Folder

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

Install Citrix 2203 LTSR CU1 - Provisioning Server x64 to:
C:\Program Files\Citrix\Provisioning Services\

Step 15. Click Install.

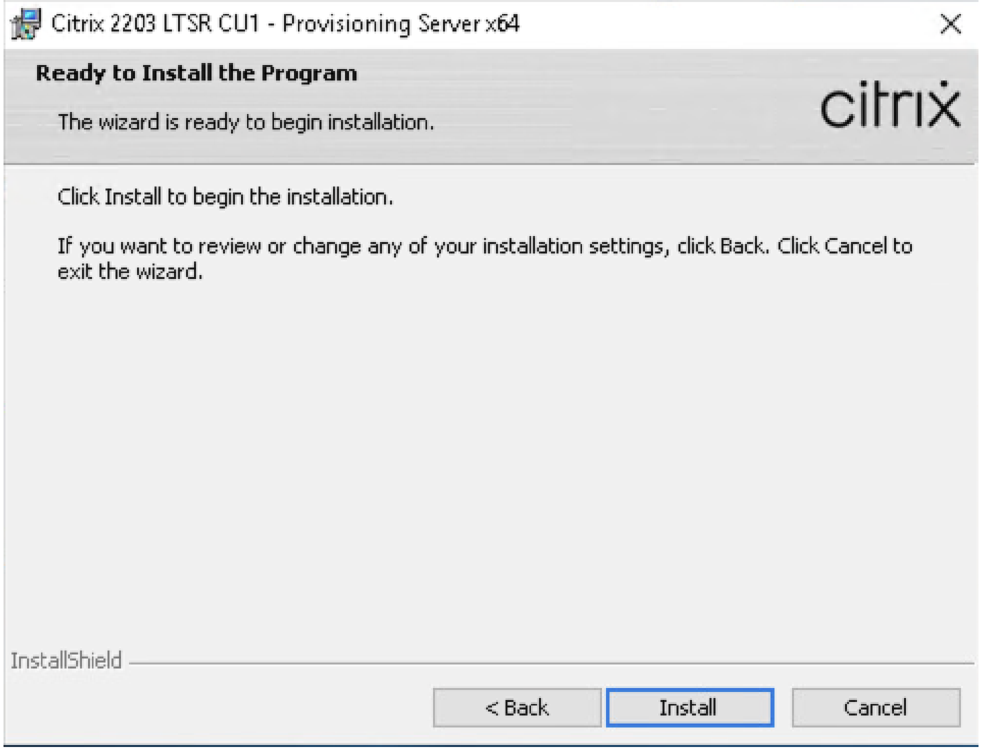
Citrix 2203 LTSR CU1 - Provisioning Server x64

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

Step 16. Click Finish after successful installation.

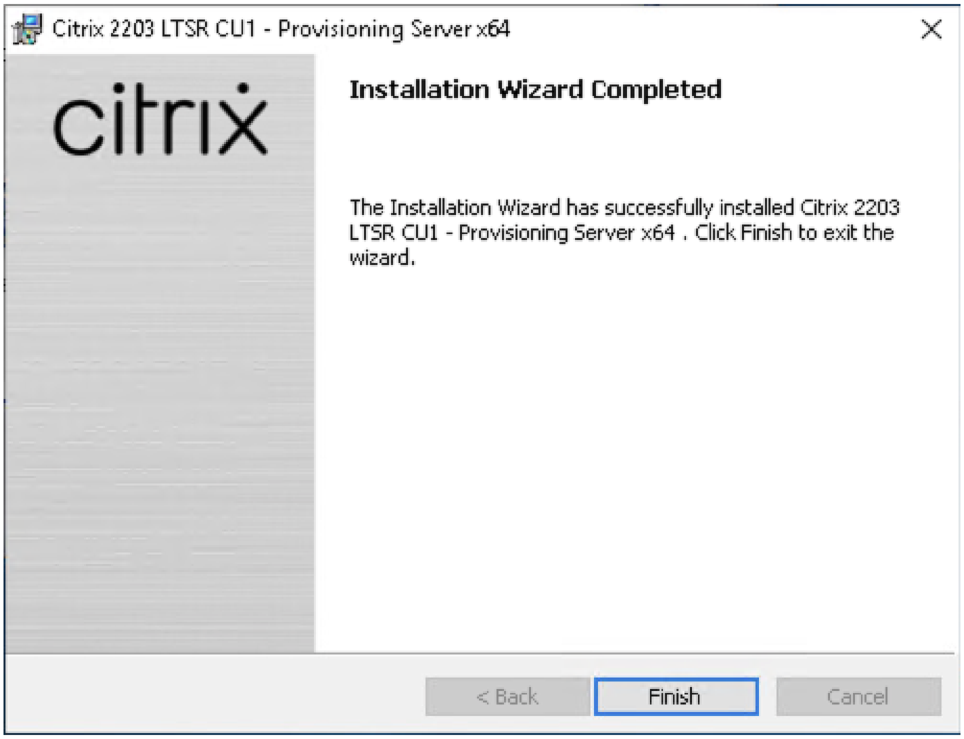
Citrix 2203 LTSR CU1 - Provisioning Server x64

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix 2203 LTSR CU1 - Provisioning Server x64 . Click Finish to exit the wizard.

Step 17. From the main installation screen, select Server Installation.

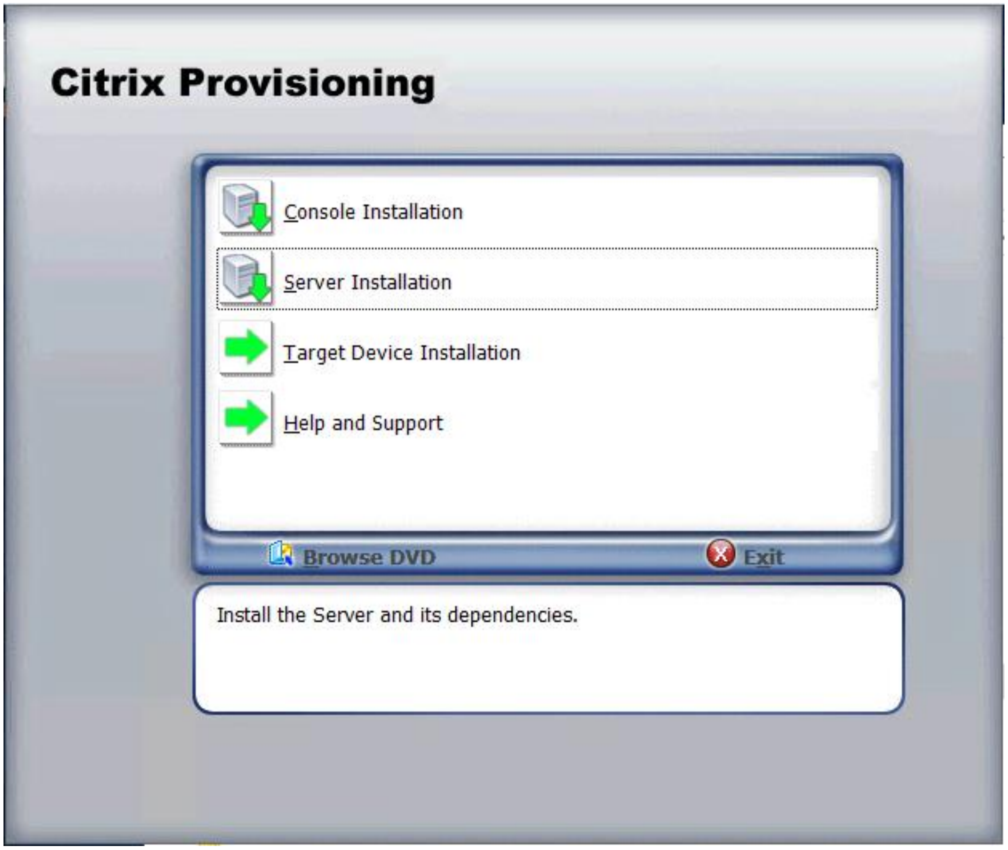

Install the Server and its dependencies.

Step 18. Click Install on the prerequisites dialog.

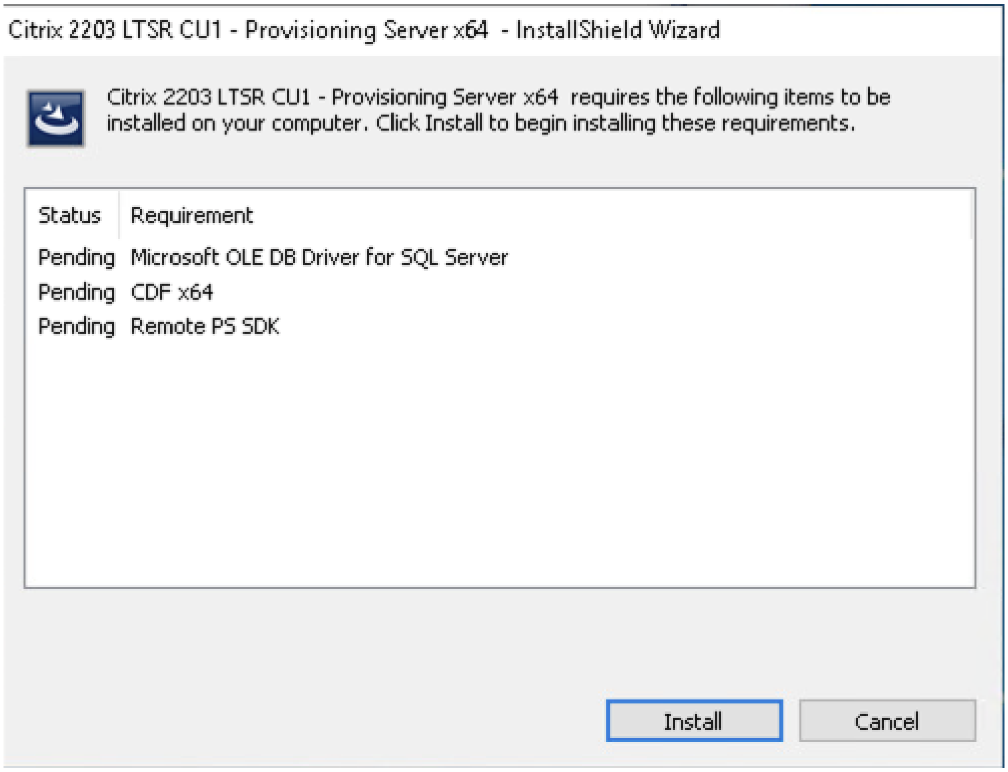

Citrix 2203 LTSR CU1 - Provisioning Server x64 - InstallShield Wizard

Citrix 2203 LTSR CU1 - Provisioning Server x64 requires the following items to be installed on your computer. Click Install to begin installing these requirements.

Status	Requirement
Pending	Microsoft OLE DB Driver for SQL Server
Pending	CDF x64
Pending	Remote PS SDK

Step 19. Click Next when the Installation wizard starts.

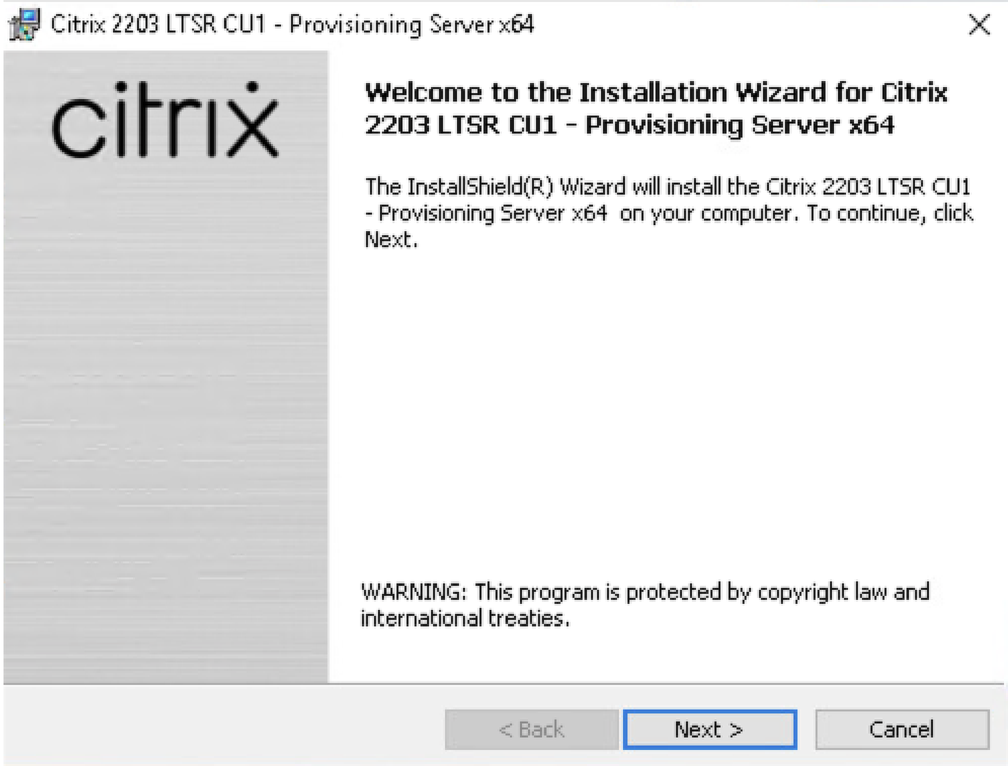

Citrix 2203 LTSR CU1 - Provisioning Server x64

Welcome to the Installation Wizard for Citrix 2203 LTSR CU1 - Provisioning Server x64

The InstallShield(R) Wizard will install the Citrix 2203 LTSR CU1 - Provisioning Server x64 on your computer. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

Step 20. Review the license agreement terms. If acceptable, select the radio button labeled "I accept the terms in the license agreement."

Step 21. Click Next.

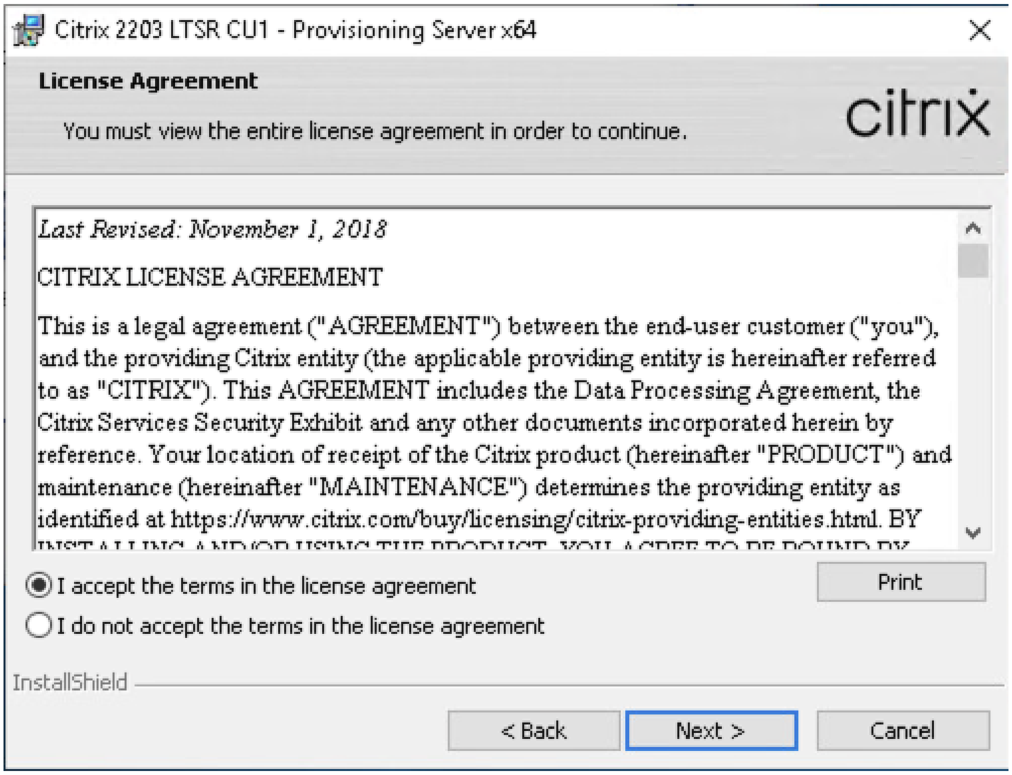

Citrix 2203 LTSR CU1 - Provisioning Server x64

License Agreement

You must view the entire license agreement in order to continue.

Last Revised: November 1, 2018

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). This AGREEMENT includes the Data Processing Agreement, the Citrix Services Security Exhibit and any other documents incorporated herein by reference. Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 22. Select Automatically open Citrix PVS Firewall Ports.

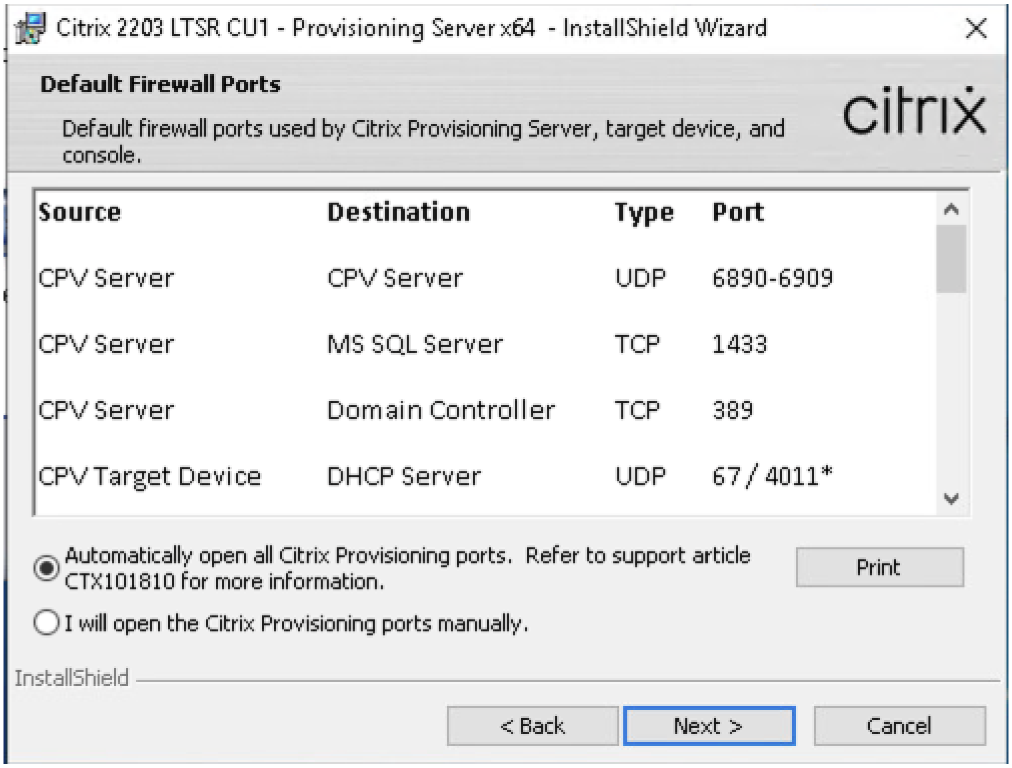

Citrix 2203 LTSR CU1 - Provisioning Server x64 - InstallShield Wizard

Default Firewall Ports

Default firewall ports used by Citrix Provisioning Server, target device, and console.

Source	Destination	Type	Port
CPV Server	CPV Server	UDP	6890-6909
CPV Server	MS SQL Server	TCP	1433
CPV Server	Domain Controller	TCP	389
CPV Target Device	DHCP Server	UDP	67 / 4011*

● Automatically open all Citrix Provisioning ports. Refer to support article CTX101810 for more information.

○ I will open the Citrix Provisioning ports manually.

Step 23. Provide User Name and Organization information. Select who will see the application.

Step 24. Click Next.

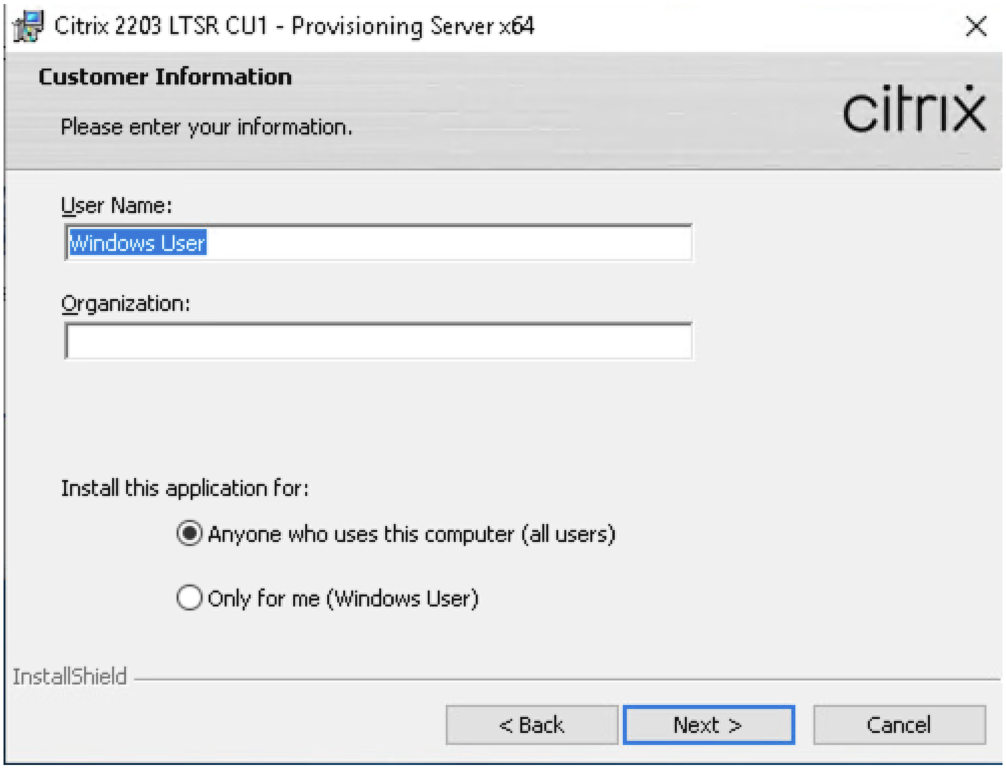

Citrix 2203 LTSR CU1 - Provisioning Server x64

Customer Information

Please enter your information.

User Name: Windows User

Organization:

Install this application for:

● Anyone who uses this computer (all users)

○ Only for me (Windows User)

Step 25. Accept the default installation location.

Step 26. Click Next.

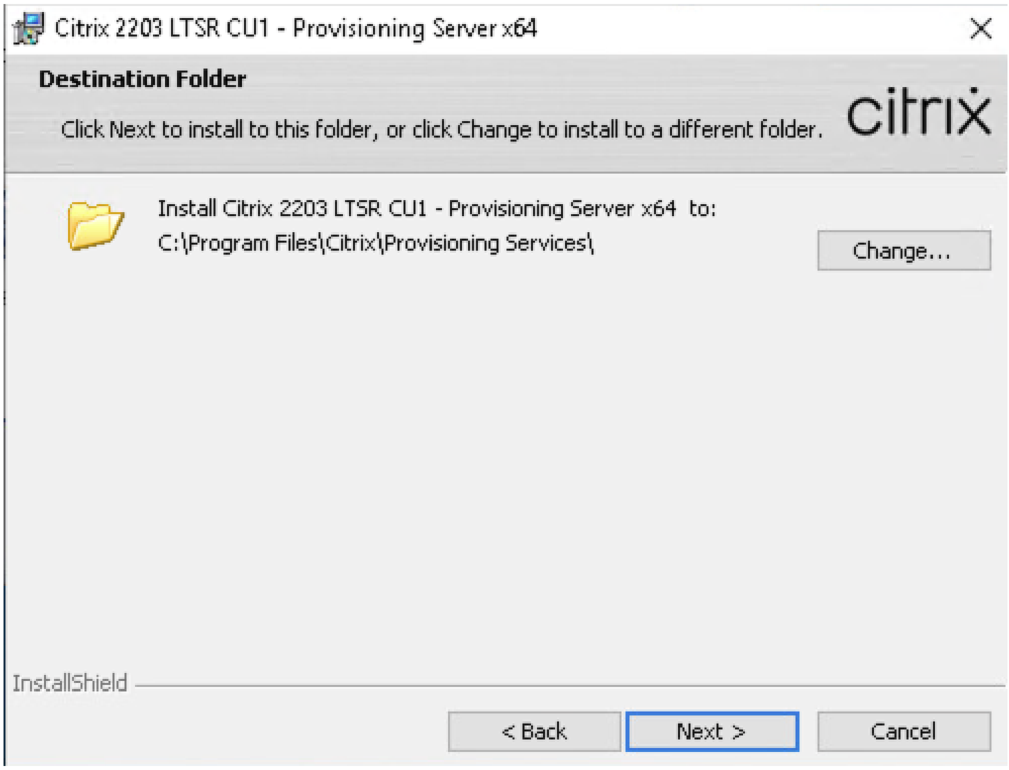

Citrix 2203 LTSR CU1 - Provisioning Server x64

Destination Folder

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

Install Citrix 2203 LTSR CU1 - Provisioning Server x64 to:
C:\Program Files\Citrix\Provisioning Services\

Step 27. Click Install to begin the installation.

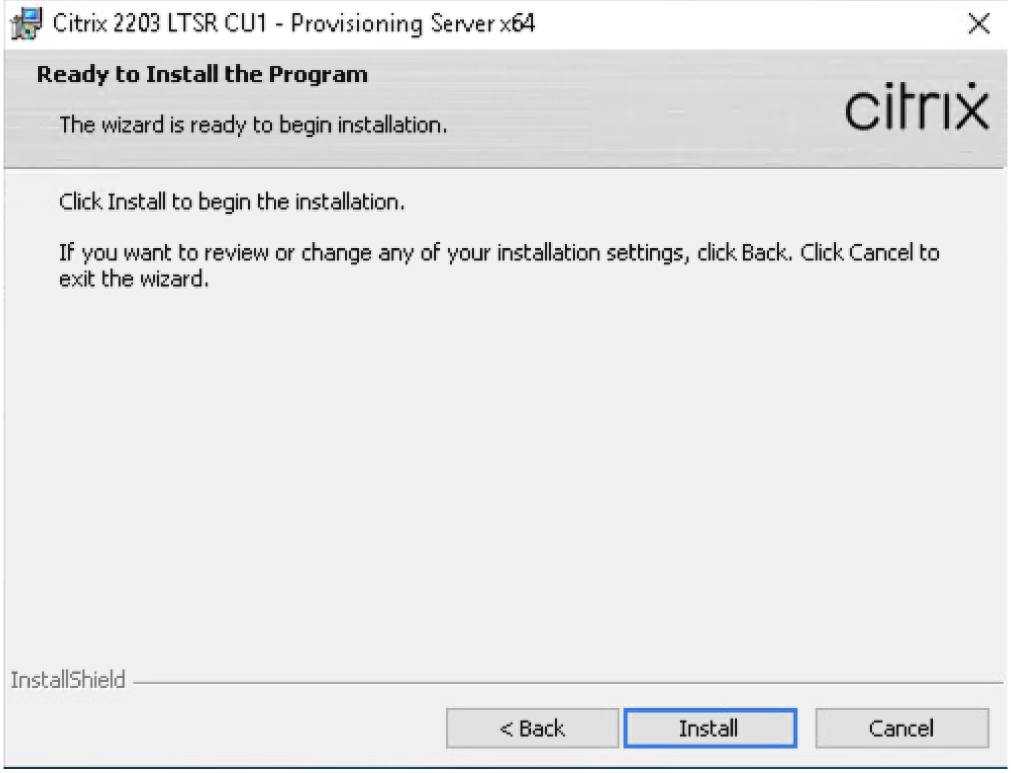

Citrix 2203 LTSR CU1 - Provisioning Server x64

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

Step 28. Click Finish when the install is complete.

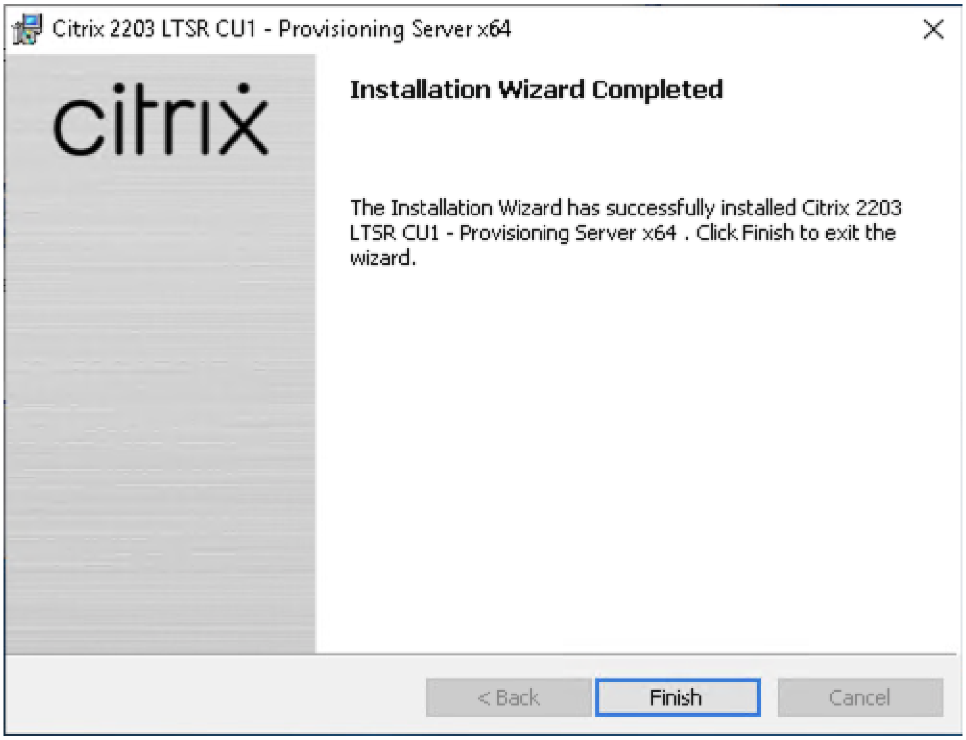

Citrix 2203 LTSR CU1 - Provisioning Server x64

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix 2203 LTSR CU1 - Provisioning Server x64 . Click Finish to exit the wizard.

Procedure 3. Configure Citrix Provisioning Services

Step 1. Start PVS Configuration Wizard.

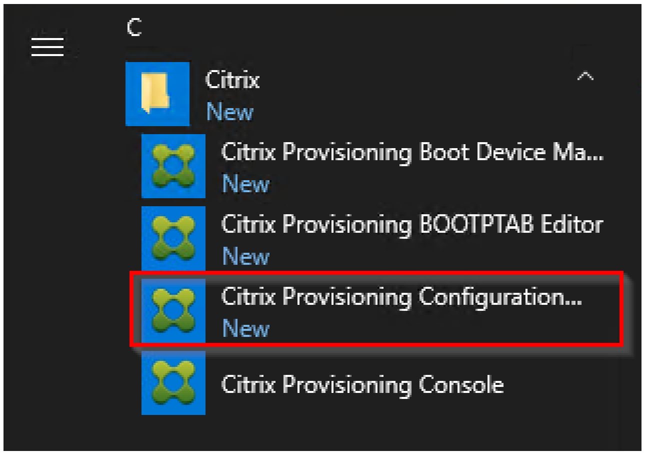

C
Citrix
New
Citrix Provisioning Boot Device Ma...
New
Citrix Provisioning BOOTPTAB Editor
New
Citrix Provisioning Configuration...
New
Citrix Provisioning Console

Step 2. Click Next.

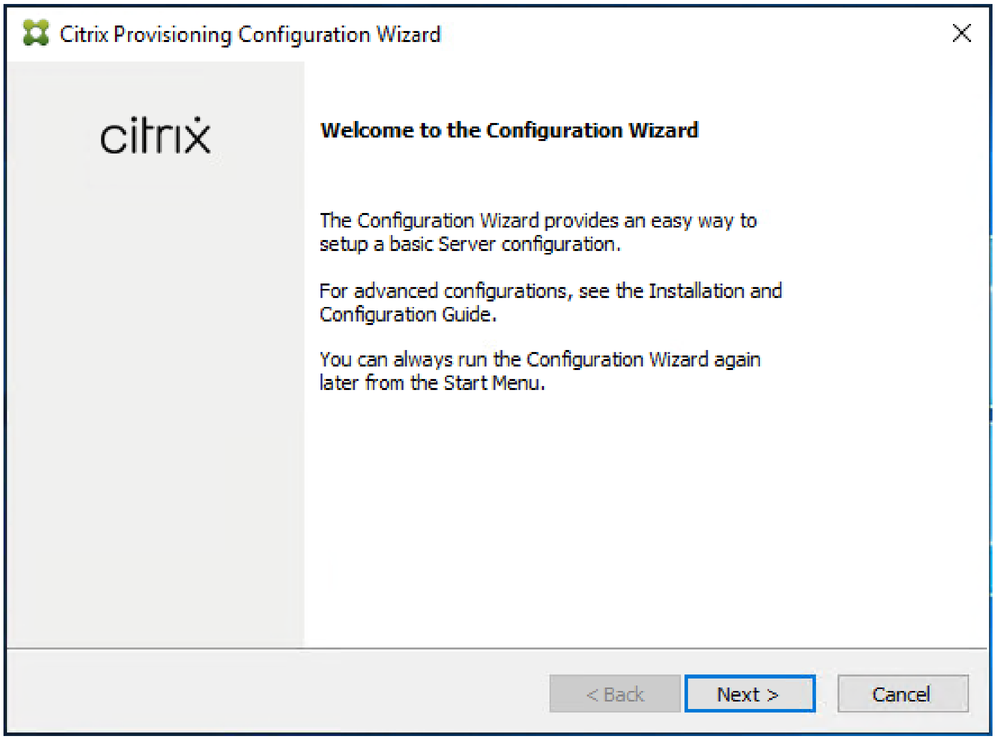

Citrix Provisioning Configuration Wizard

Welcome to the Configuration Wizard

The Configuration Wizard provides an easy way to setup a basic Server configuration.

For advanced configurations, see the Installation and Configuration Guide.

You can always run the Configuration Wizard again later from the Start Menu.

Step 3. Since the PVS server is not the DHCP server for the environment, select the radio button labeled, "The service that runs on another computer."

Step 4. Click Next.

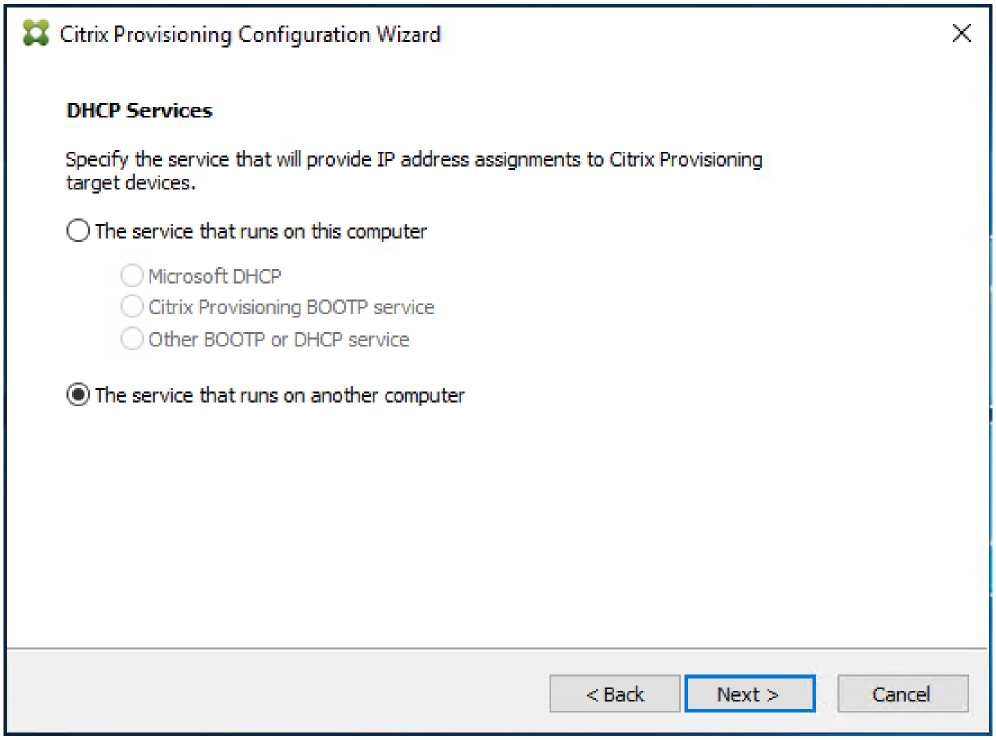

Citrix Provisioning Configuration Wizard

DHCP Services

Specify the service that will provide IP address assignments to Citrix Provisioning target devices.

○ The service that runs on this computer

- ○ Microsoft DHCP
- ○ Citrix Provisioning BOOTP service
- ○ Other BOOTP or DHCP service

● The service that runs on another computer

Step 5. Since DHCP boot options are used for TFTP services, select the radio button labeled, "The service that runs on another computer."

Step 6. Click Next.

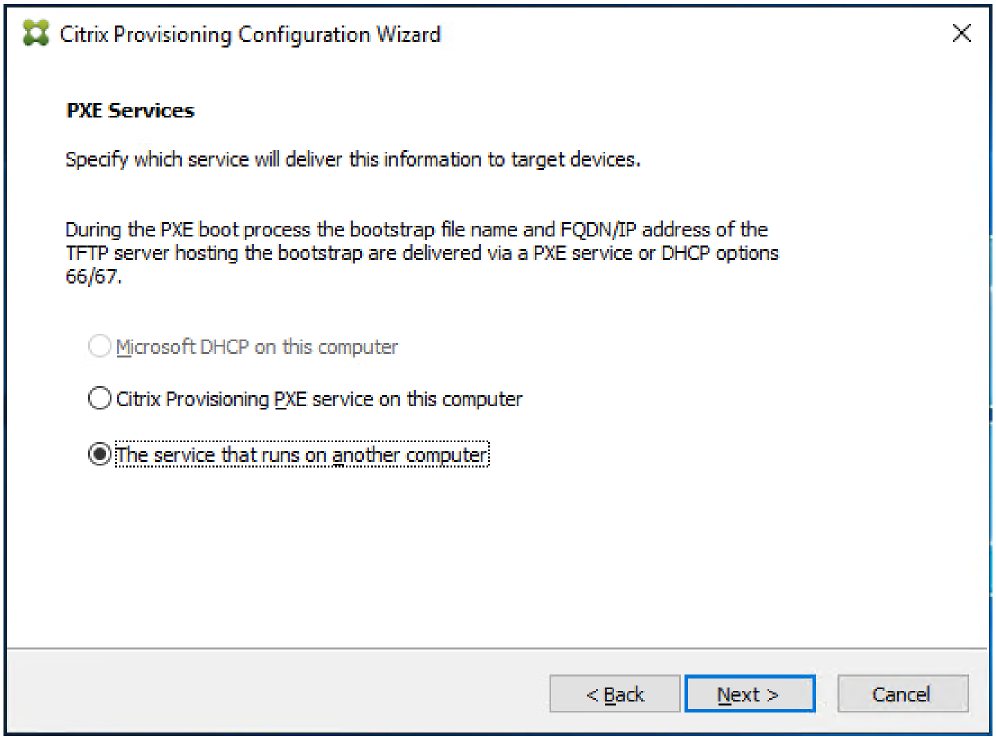

Citrix Provisioning Configuration Wizard

PXE Services

Specify which service will deliver this information to target devices.

During the PXE boot process the bootstrap file name and FQDN/IP address of the TFTP server hosting the bootstrap are delivered via a PXE service or DHCP options 66/67.

○ Microsoft DHCP on this computer

○ Citrix Provisioning PXE service on this computer

● The service that runs on another computer

Step 7. Since this is the first server in the farm, select the radio button labeled, "Create farm."

Step 8. Click Next.

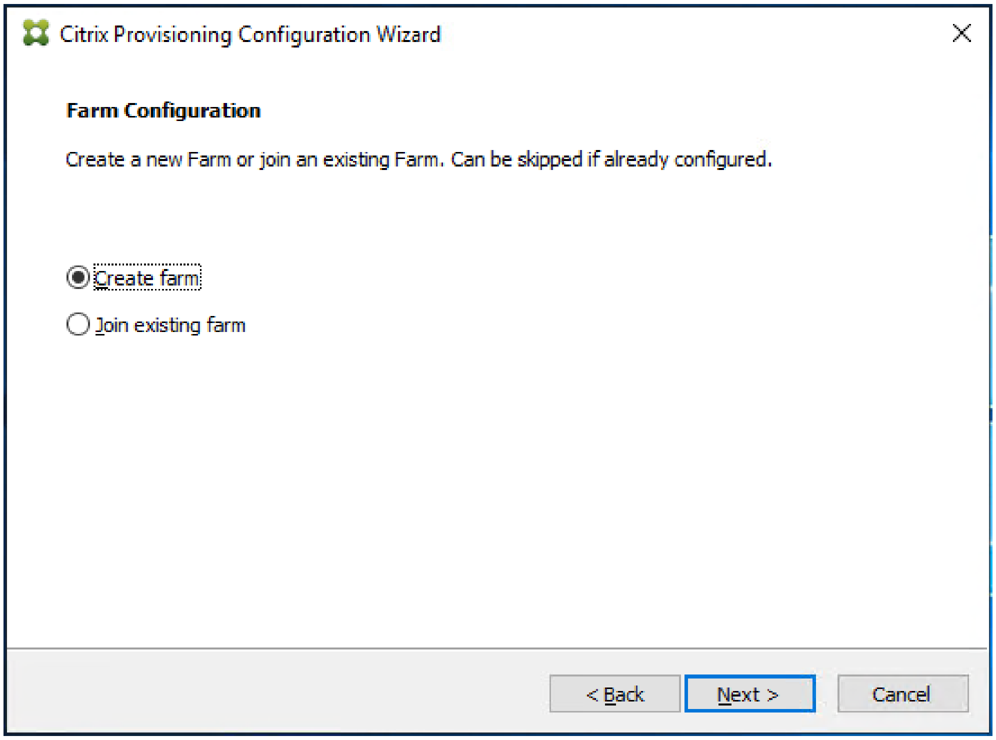

Citrix Provisioning Configuration Wizard

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

● Create farm

○ Join existing farm

Step 9. Enter the FQDN of the SQL server.

Step 10. Click Next.

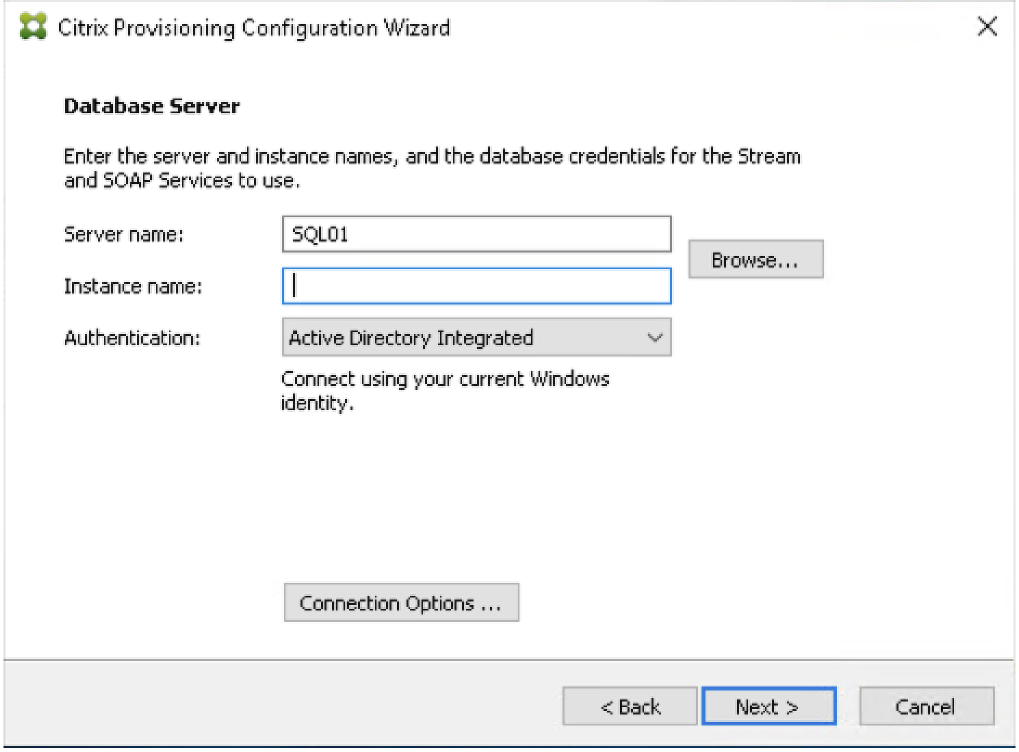

Citrix Provisioning Configuration Wizard

Database Server

Enter the server and instance names, and the database credentials for the Stream and SOAP Services to use.

Server name: SQL01

Instance name:

Authentication: Active Directory Integrated

Connect using your current Windows identity.

Step 11. Provide the Database, Farm, Site, and Collection name.

Step 12. Click Next.

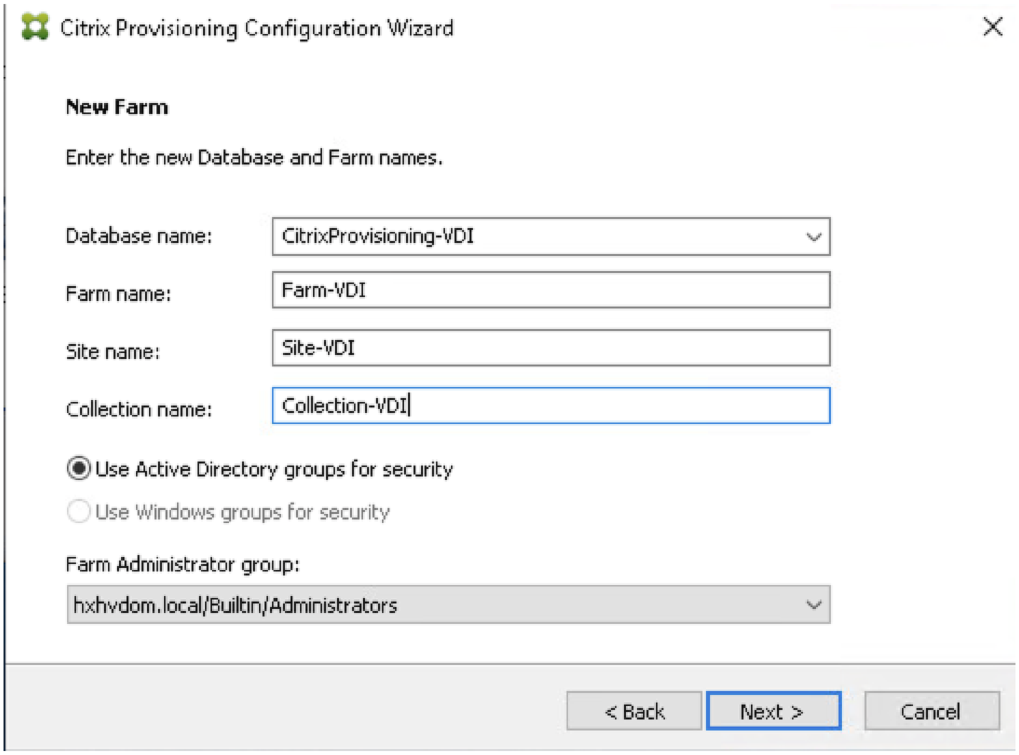

Citrix Provisioning Configuration Wizard

New Farm

Enter the new Database and Farm names.

Database name: CitrixProvisioning-VDI

Farm name: Farm-VDI

Site name: Site-VDI

Collection name: Collection-VDI

● Use Active Directory groups for security

○ Use Windows groups for security

Farm Administrator group: hxhvdom.local/Builtin/Administrators

Step 13. Provide the vDisk Store details.

Step 14. Click Next.

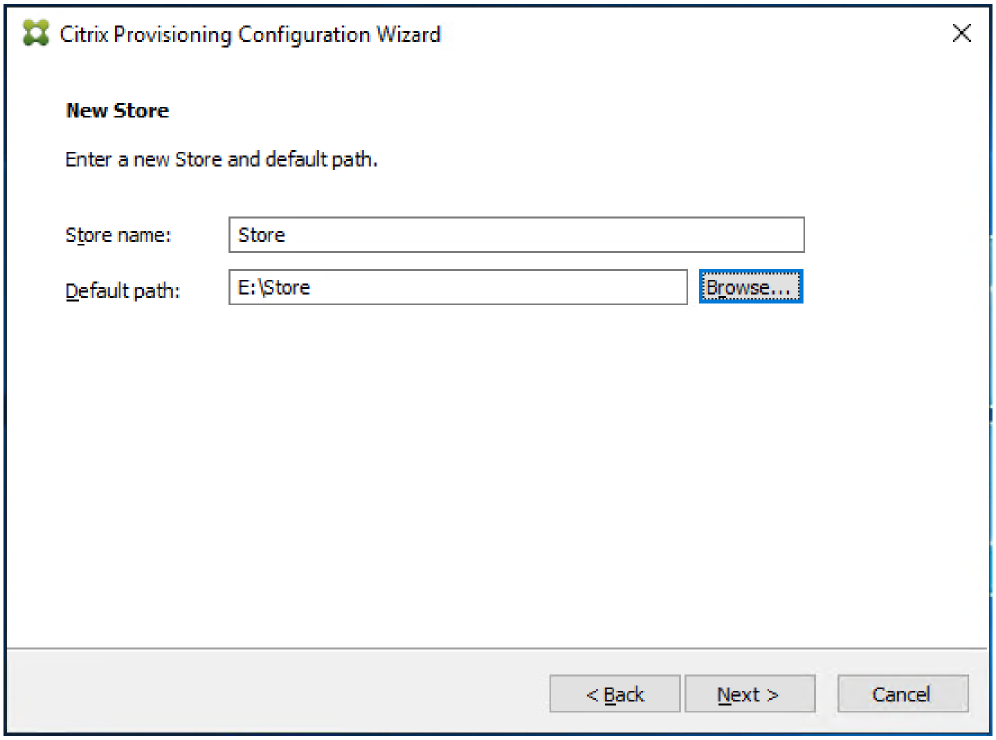

New Store

Enter a new Store and default path.

Store name: Store

Default path: E:\Store

Step 15. For large scale PVS environment, it is recommended to create the share using support for CIFS/SMB3 on an enterprise ready File Server.

Step 16. Provide the FQDN of the license server.

Step 17. Optionally, provide a port number if changed on the license server.

Step 18. Click Next.

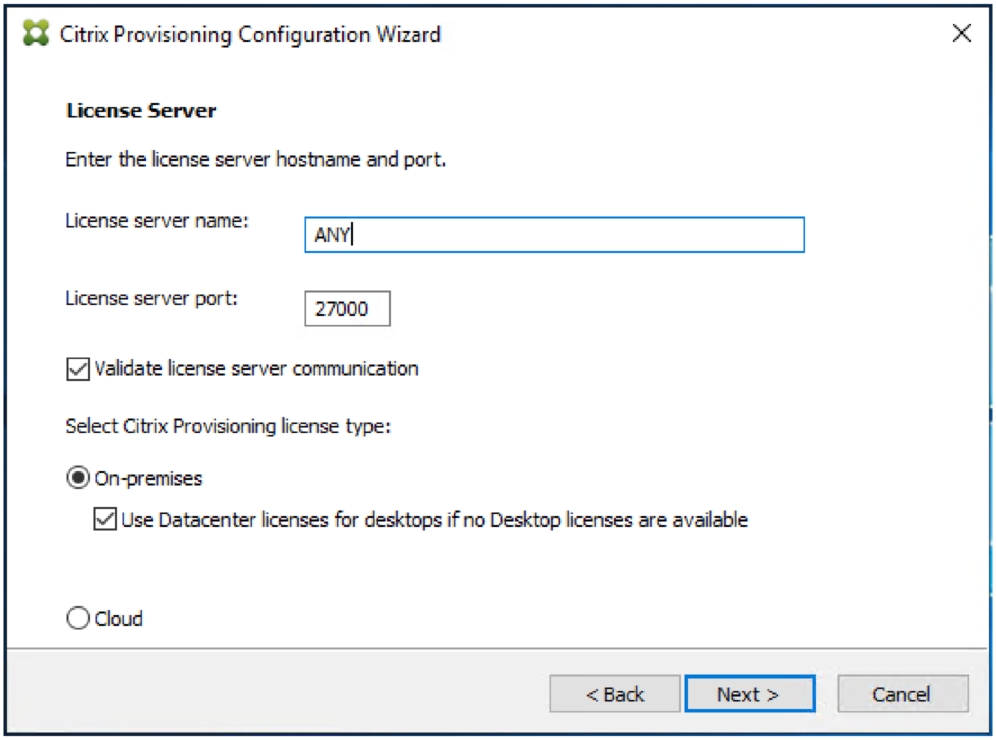

License Server

Enter the license server hostname and port.

License server name: ANY

License server port: 27000

☑ Validate license server communication

Select Citrix Provisioning license type:

● On-premises

☑ Use Datacenter licenses for desktops if no Desktop licenses are available

○ Cloud

Step 19. If an Active Directory service account is not already setup for the PVS servers, create that account prior to clicking Next on this dialog.

Step 20. Select the Specified user account radio button.

Step 21. Complete the User name, Domain, Password, and Confirm password fields, using the PVS account information created earlier.

Step 22. Click Next.

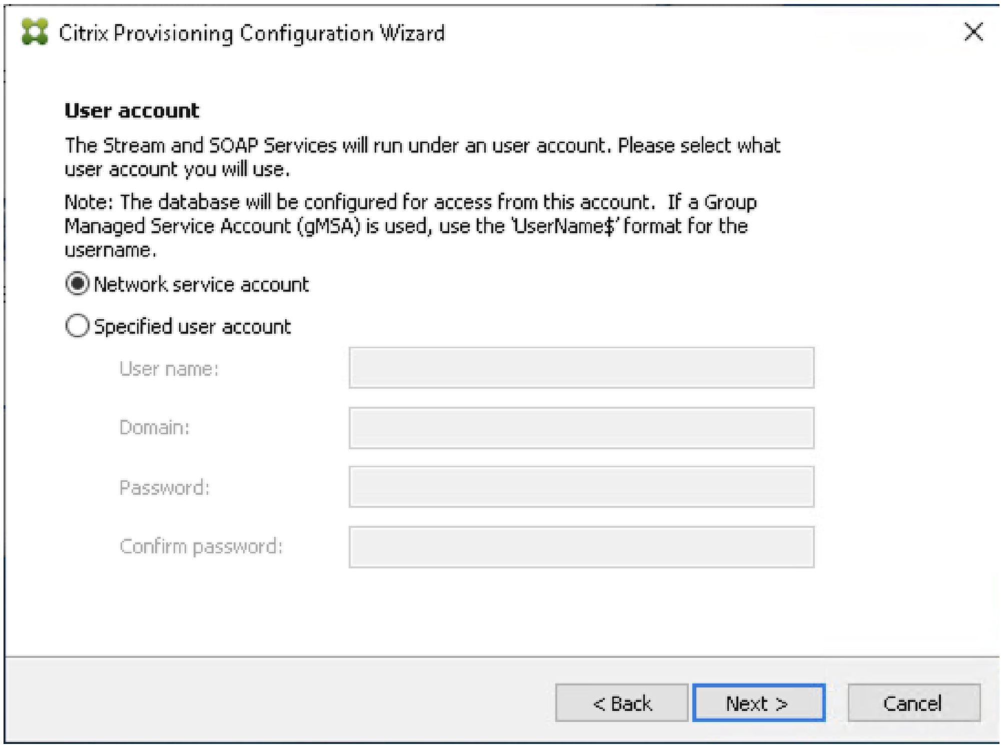

Citrix Provisioning Configuration Wizard

User account

The Stream and SOAP Services will run under an user account. Please select what user account you will use.

Note: The database will be configured for access from this account. If a Group Managed Service Account (gMSA) is used, use the 'UserName$' format for the username.

● Network service account

○ Specified user account

User name:

Domain:

Password:

Confirm password:

Step 23. Set the Days between password updates to 7.

Note: This will vary per environment. “7 days” for the configuration was appropriate for testing purposes.

Step 24. Click Next.

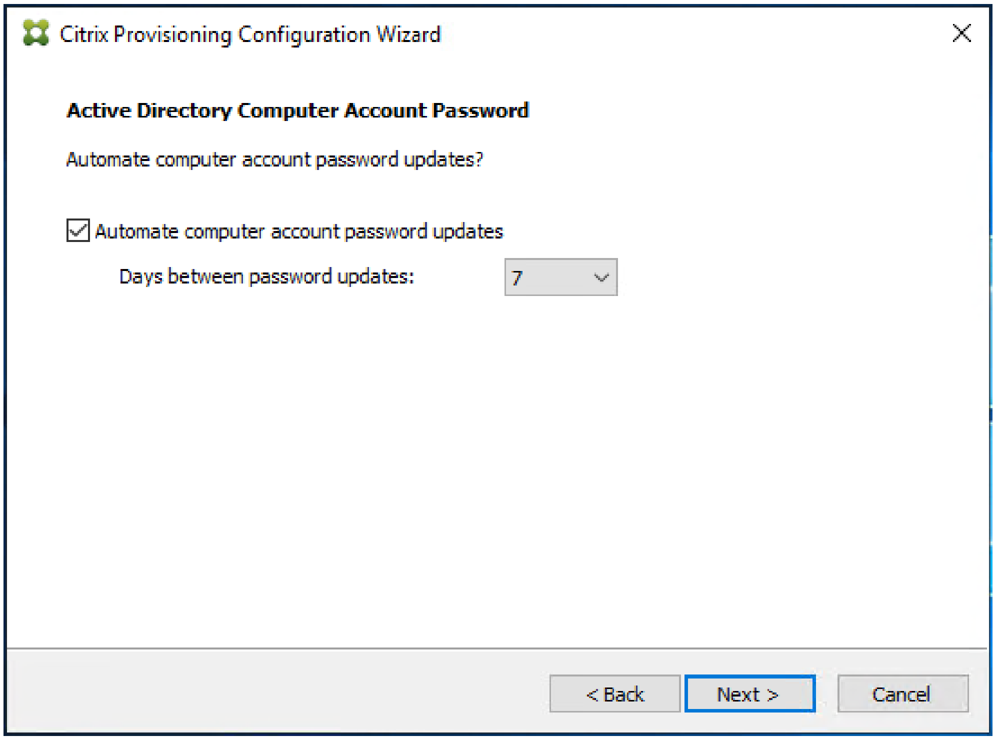

Active Directory Computer Account Password

Automate computer account password updates?

☑ Automate computer account password updates

Days between password updates: 7

Step 25. Keep the defaults for the network cards.

Step 26. Click Next.

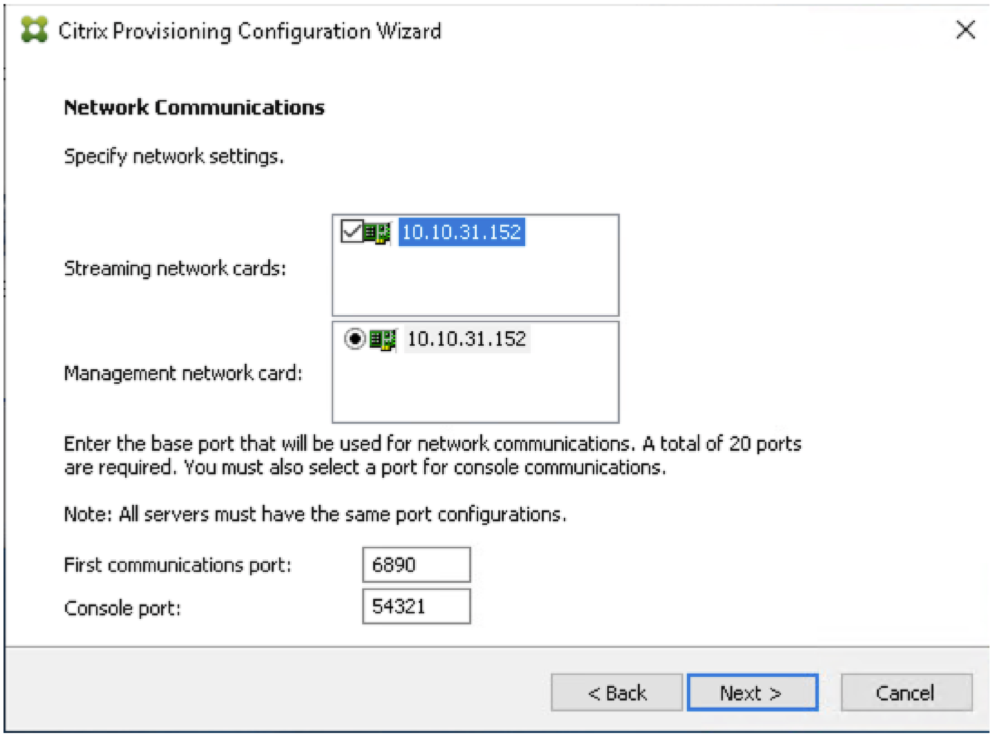

Network Communications

Specify network settings.

Streaming network cards: ☑ 10.10.31.152

Management network card: ● 10.10.31.152

Enter the base port that will be used for network communications. A total of 20 ports are required. You must also select a port for console communications.

Note: All servers must have the same port configurations.

First communications port: 6890

Console port: 54321

Step 27. Select Use the Provisioning Services TFTP service checkbox.

Step 28. Click Next.

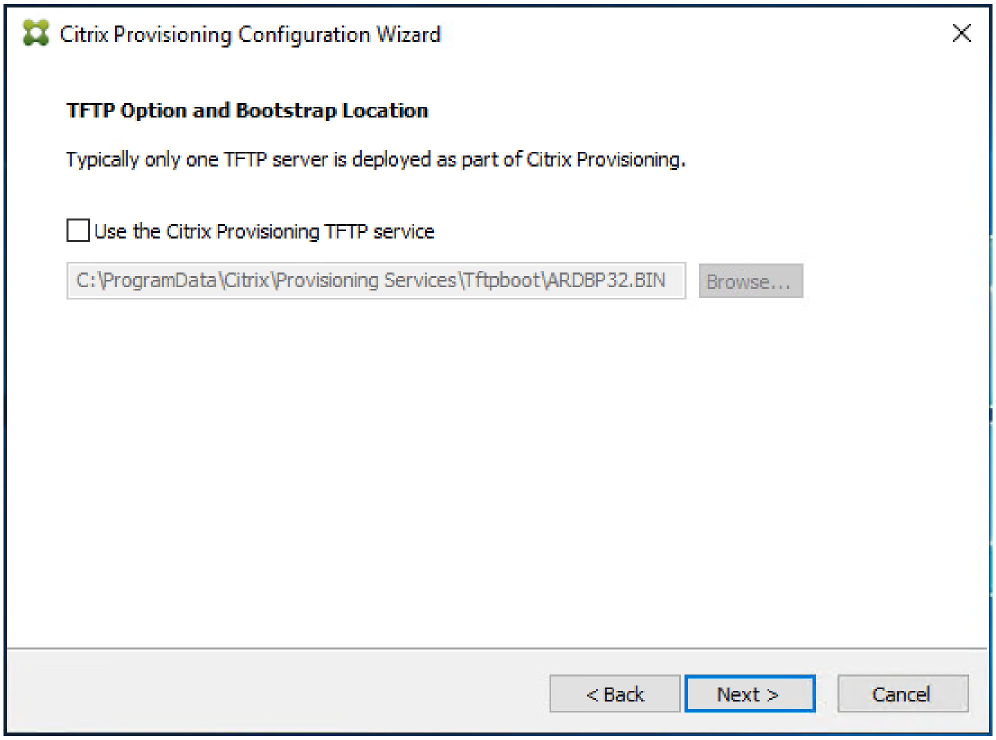

Citrix Provisioning Configuration Wizard

TFTP Option and Bootstrap Location

Typically only one TFTP server is deployed as part of Citrix Provisioning.

☐ Use the Citrix Provisioning TFTP service

C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN

Step 29. If Soap Server is used, provide details.

Step 30. Click Next.

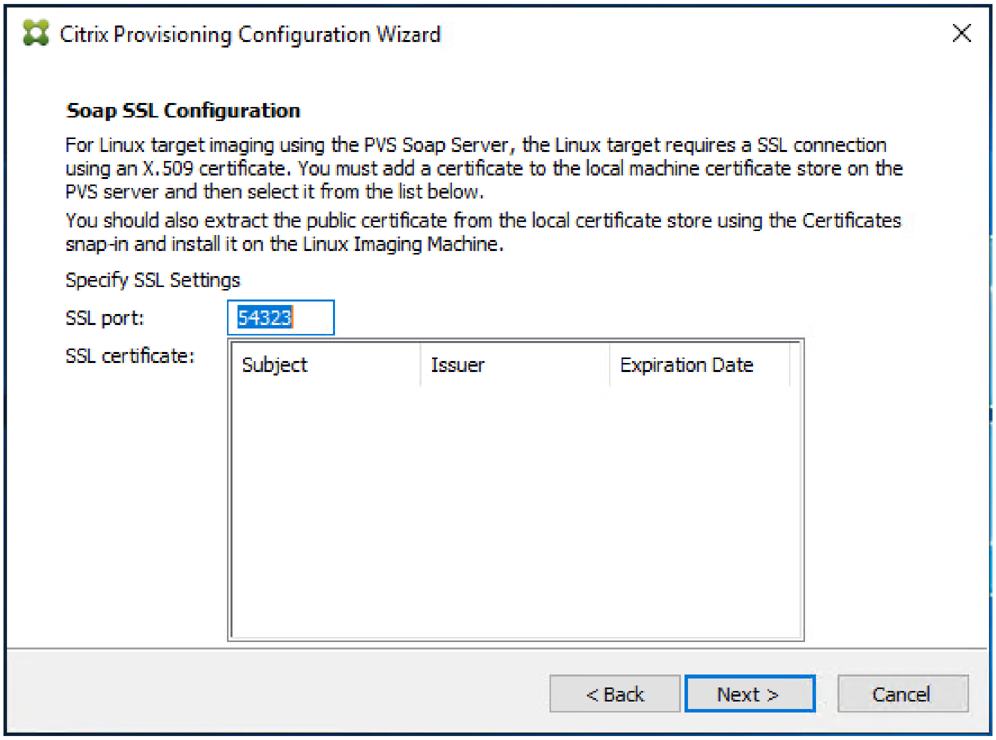

Citrix Provisioning Configuration Wizard

Soap SSL Configuration

For Linux target imaging using the PVS Soap Server, the Linux target requires a SSL connection using an X.509 certificate. You must add a certificate to the local machine certificate store on the PVS server and then select it from the list below.

You should also extract the public certificate from the local certificate store using the Certificates snap-in and install it on the Linux Imaging Machine.

Specify SSL Settings

SSL port: 54323

SSL certificate:

Subject	Issuer	Expiration Date

Step 31. If desired fill in Problem Report Configuration.

Step 32. Click Next.

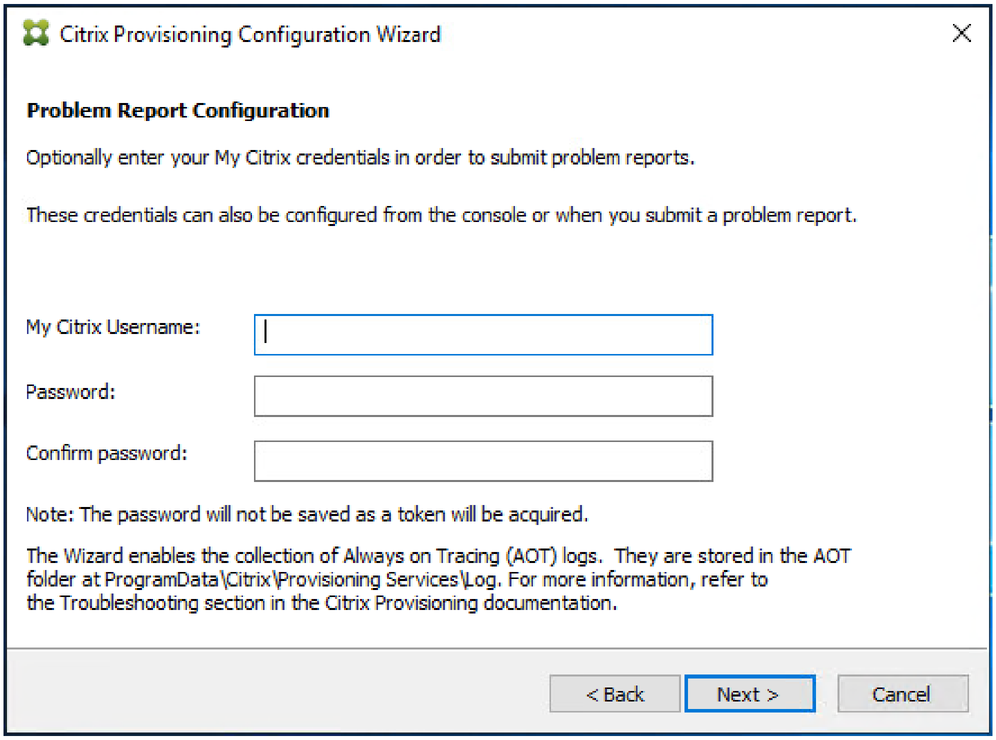

Citrix Provisioning Configuration Wizard

Problem Report Configuration

Optionally enter your My Citrix credentials in order to submit problem reports.

These credentials can also be configured from the console or when you submit a problem report.

My Citrix Username:

Password:

Confirm password:

Note: The password will not be saved as a token will be acquired.

The Wizard enables the collection of Always on Tracing (AOT) logs. They are stored in the AOT folder at ProgramData\Citrix\Provisioning Services\Log. For more information, refer to the Troubleshooting section in the Citrix Provisioning documentation.

Step 33. Click Finish to start the installation.

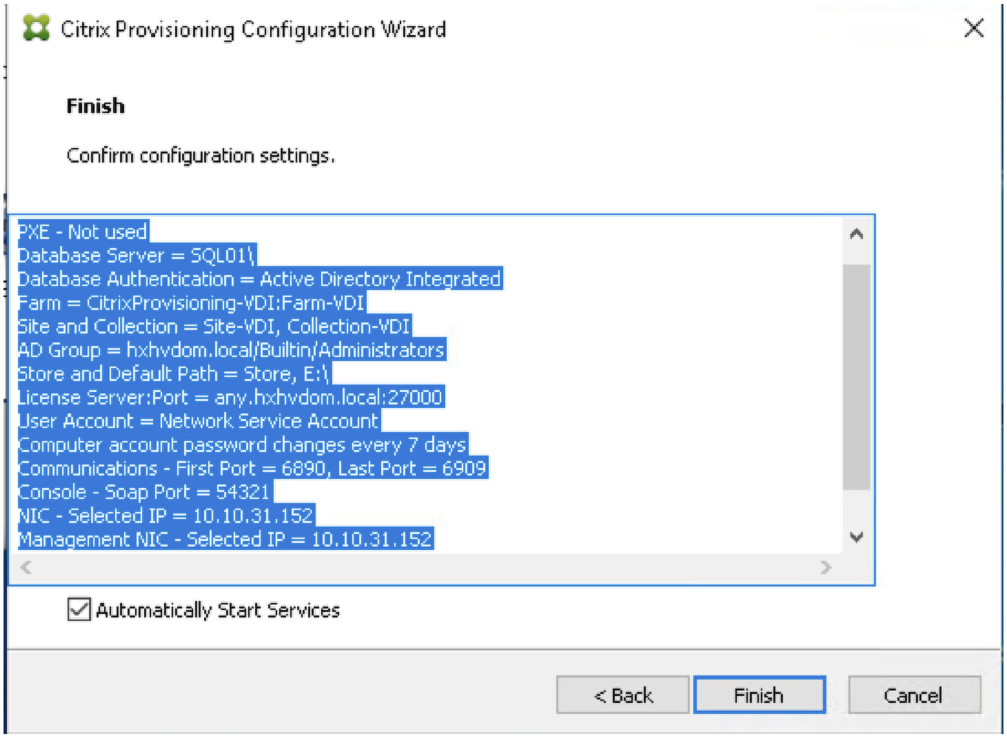

Citrix Provisioning Configuration Wizard

Finish

Confirm configuration settings.

```
PXE - Not used
Database Server = SQL01\
Database Authentication = Active Directory Integrated
Farm = CitrixProvisioning-VDI:Farm-VDI
Site and Collection = Site-VDI, Collection-VDI
AD Group = hxhvdom.local/Builtin/Administrators
Store and Default Path = Store, E:\
License Server:Port = any.hxhvdom.local:27000
User Account = Network Service Account
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.10.31.152
Management NIC - Selected IP = 10.10.31.152
```

☑ Automatically Start Services

Step 34. When the installation is completed, click Done.

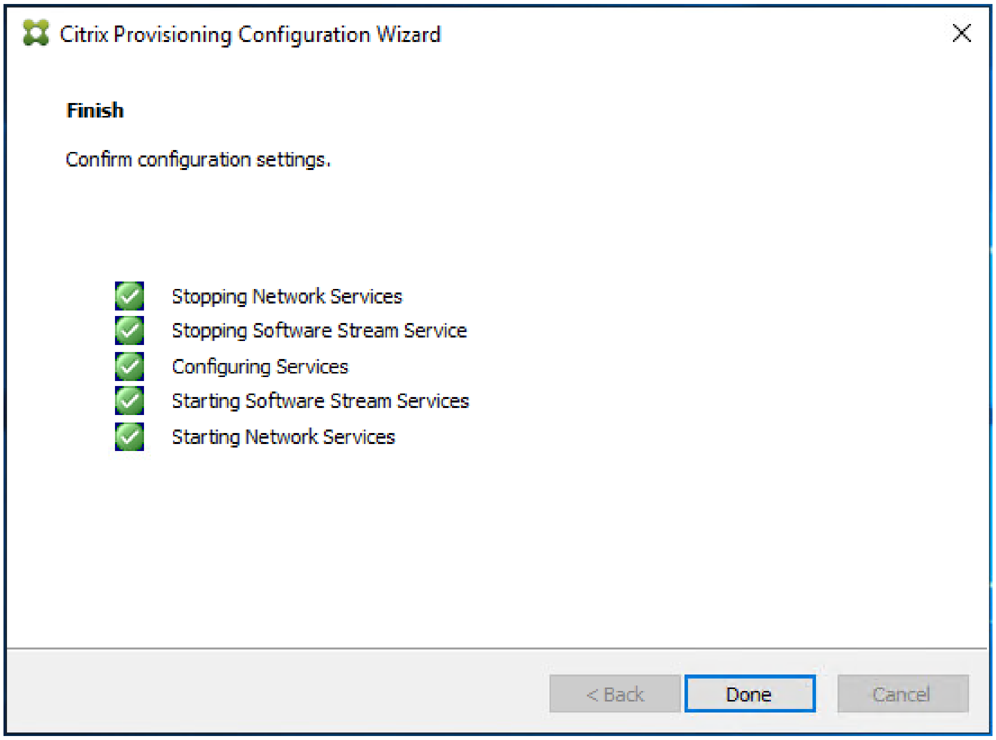

Citrix Provisioning Configuration Wizard

Finish

Confirm configuration settings.

✓ Stopping Network Services

✓ Stopping Software Stream Service

✓ Configuring Services

✓ Starting Software Stream Services

✓ Starting Network Services

Procedure 4. Install Additional PVS Servers

Note: Complete the installation steps on the additional PVS servers up to the configuration step where it asks to Create or Join a farm. In this CVD, we repeated the procedure to add a total of three PVS servers.

This procedure details how to join additional Provisioning servers to the farm already configured in the previous steps.

Step 1. On the Farm Configuration dialog, select “Join existing farm.”

Step 2. Click Next.

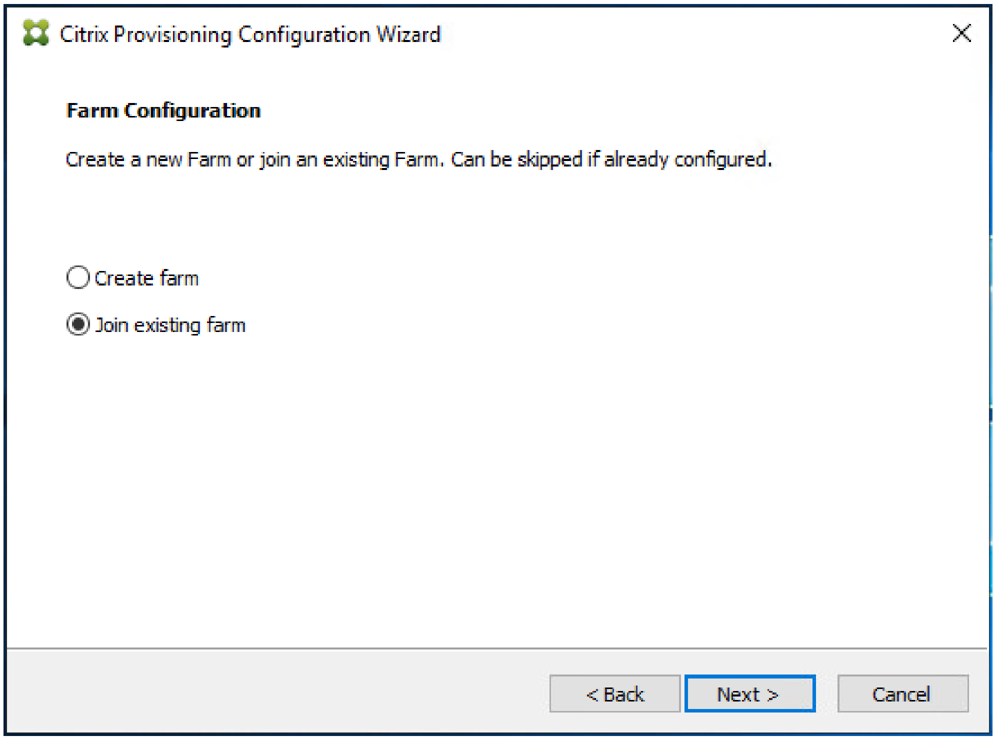

Citrix Provisioning Configuration Wizard

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

○ Create farm

● Join existing farm

Step 3. Provide the FQDN of the SQL Server.

Step 4. Click Next.

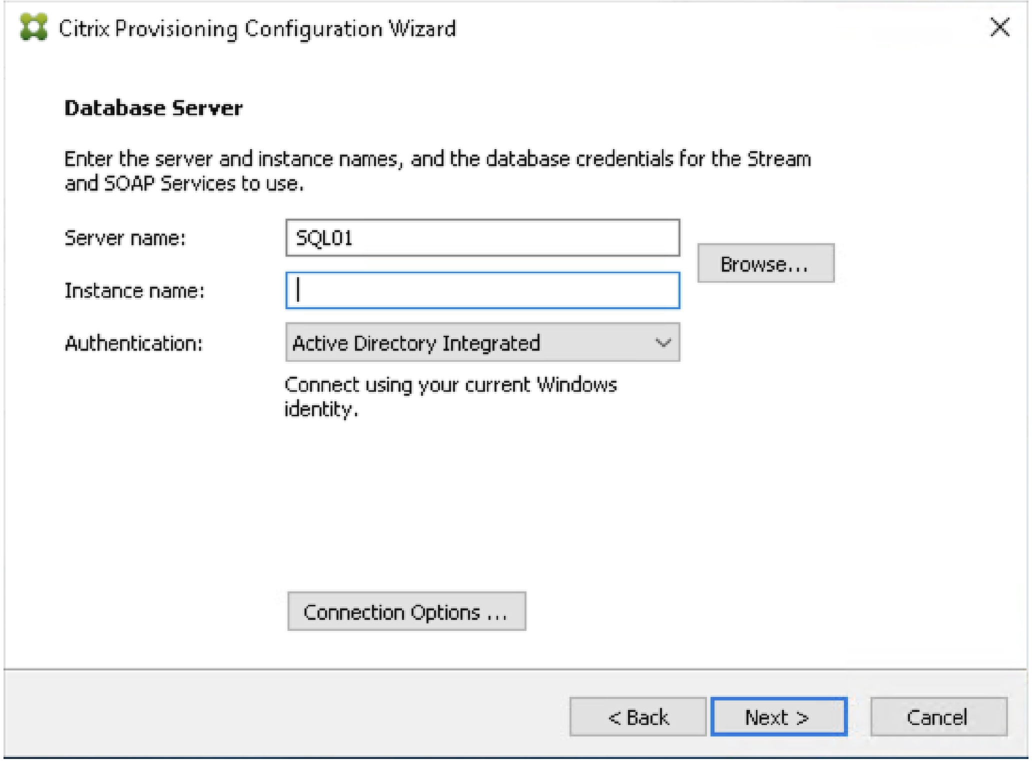

Citrix Provisioning Configuration Wizard

Database Server

Enter the server and instance names, and the database credentials for the Stream and SOAP Services to use.

Server name: SQL01

Instance name:

Authentication: Active Directory Integrated

Connect using your current Windows identity.

Step 5. Accept the Farm Name.

Step 6. Click Next.

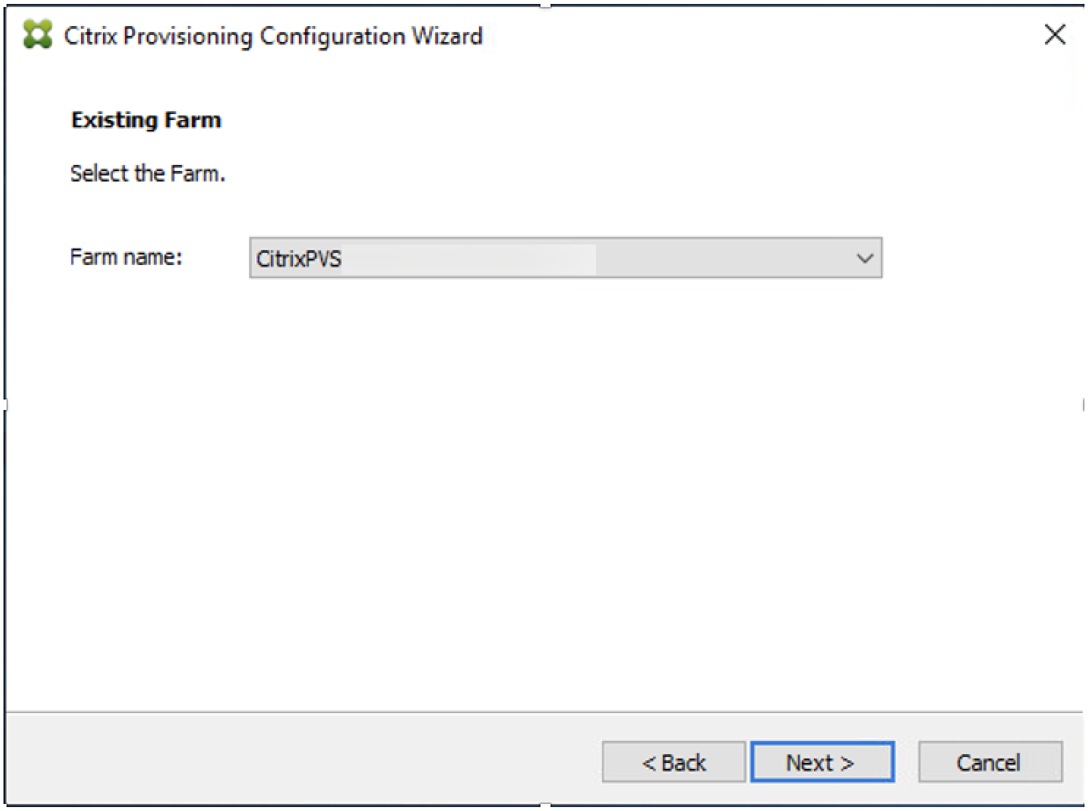

Citrix Provisioning Configuration Wizard

Existing Farm

Select the Farm.

Farm name: CitrixPVS

Step 7. Accept the Existing Site.

Step 8. Click Next.

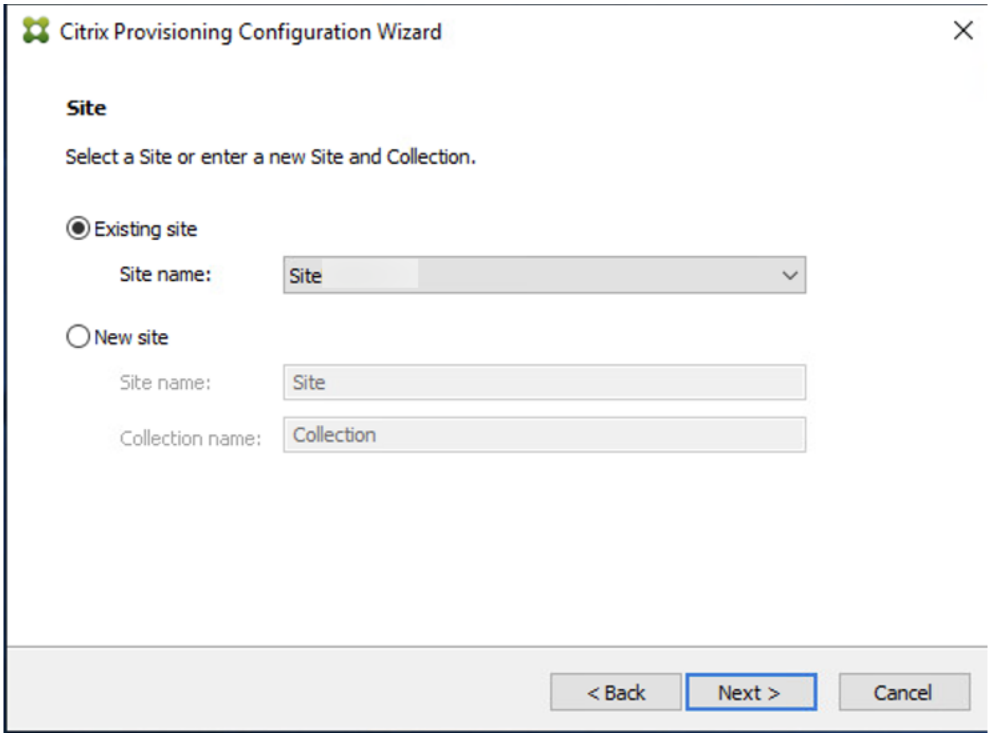

Citrix Provisioning Configuration Wizard

Site

Select a Site or enter a new Site and Collection.

● Existing site

Site name: Site

○ New site

Site name: Site

Collection name: Collection

Step 9. Accept the existing vDisk store.

Step 10. Click Next.

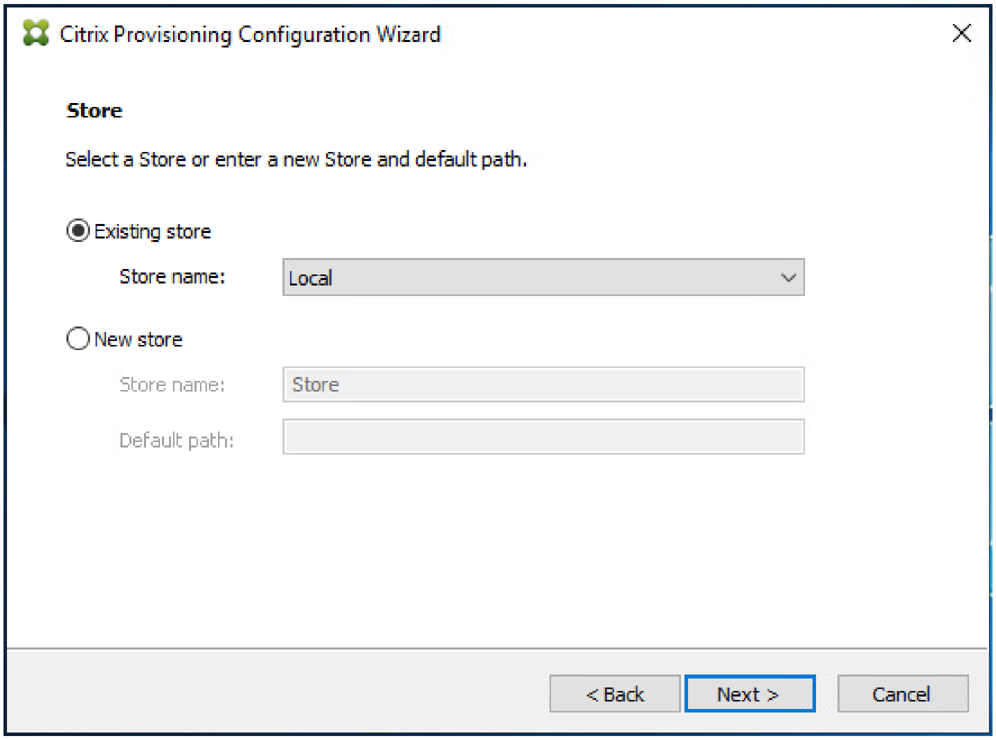

Store

Select a Store or enter a new Store and default path.

● Existing store

Store name: Local

○ New store

Store name: Store

Default path:

Step 11. Provide the FQDN of the license server.

Step 12. Optionally, provide a port number if changed on the license server.

Step 13. Click Next.

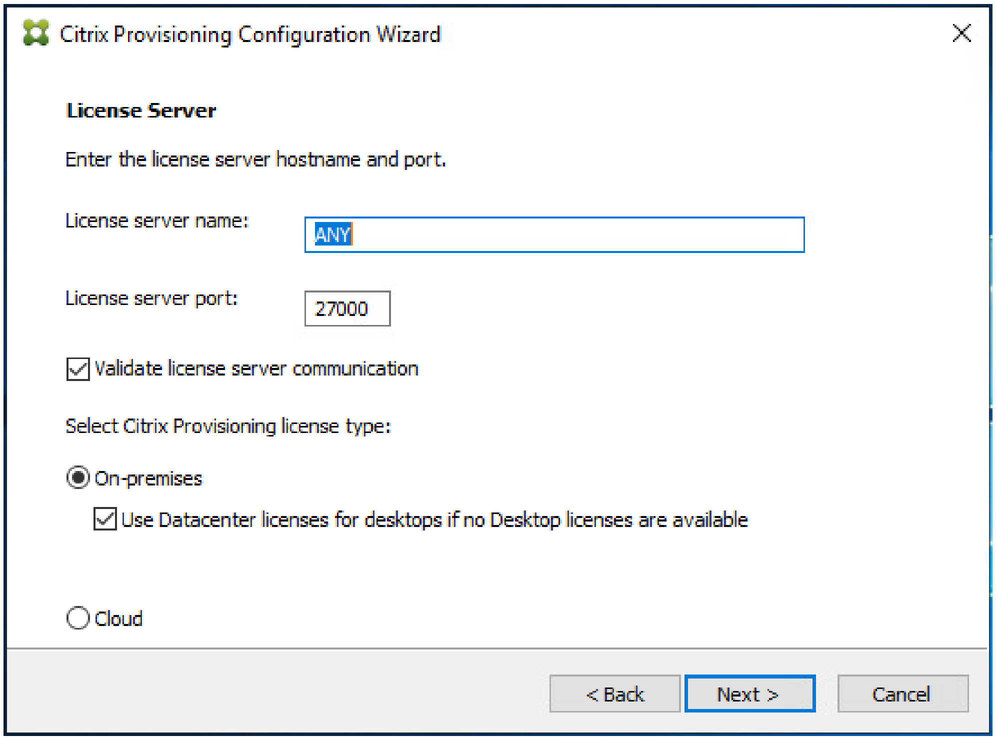

License Server

Enter the license server hostname and port.

License server name: ANY

License server port: 27000

☑ Validate license server communication

Select Citrix Provisioning license type:

● On-premises

☑ Use Datacenter licenses for desktops if no Desktop licenses are available

○ Cloud

Step 14. Provide the PVS service account information.

Step 15. Click Next.

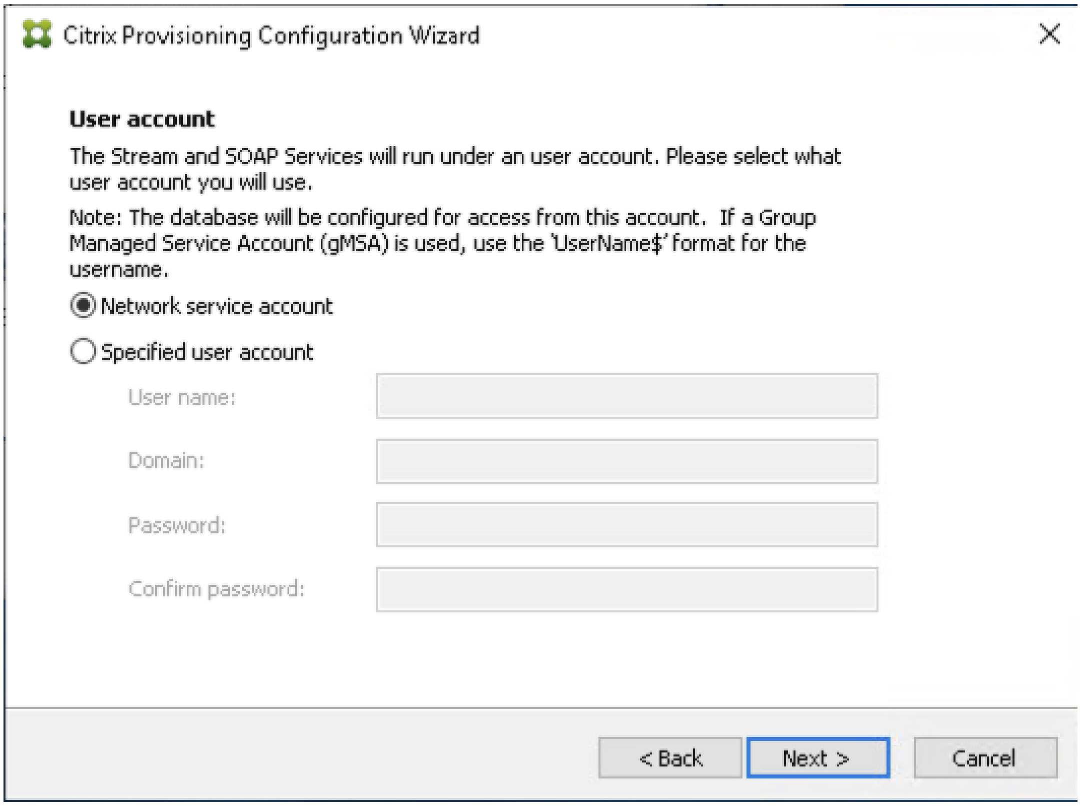

Citrix Provisioning Configuration Wizard

User account

The Stream and SOAP Services will run under an user account. Please select what user account you will use.

Note: The database will be configured for access from this account. If a Group Managed Service Account (gMSA) is used, use the 'UserName$' format for the username.

● Network service account

○ Specified user account

User name:

Domain:

Password:

Confirm password:

Step 16. Set the Days between password updates to 7.

Step 17. Click Next.

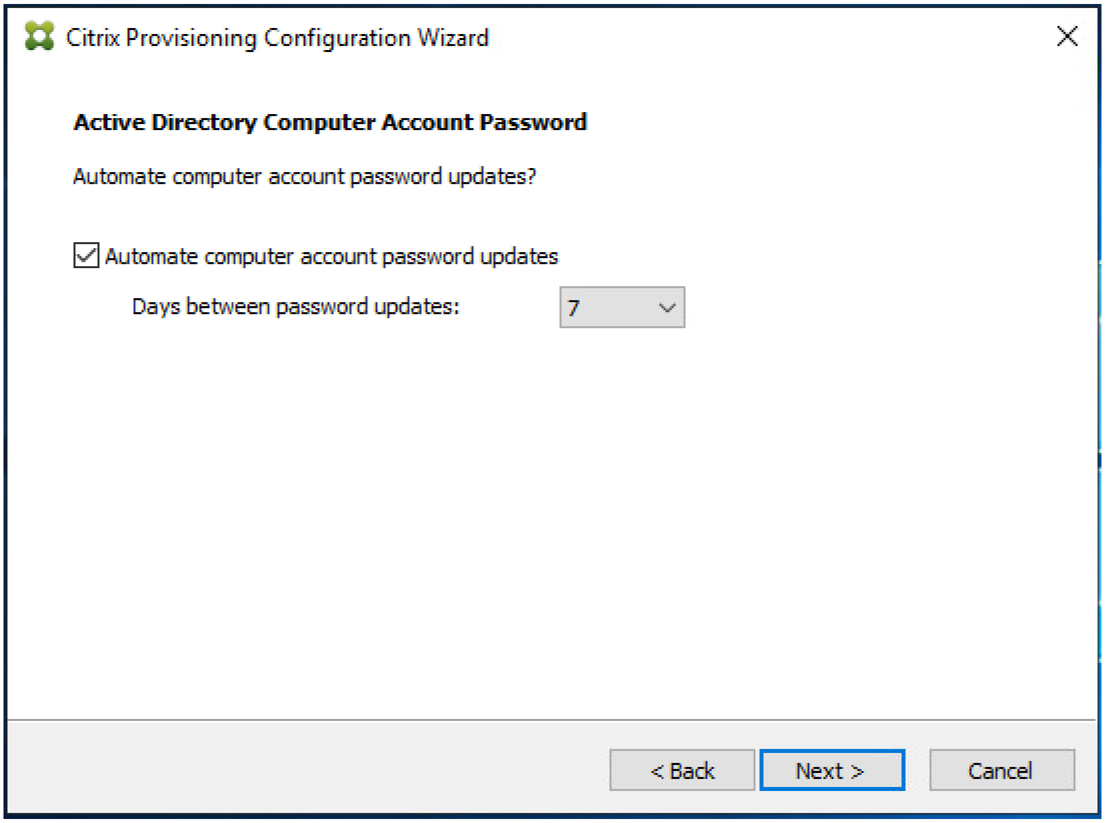

Citrix Provisioning Configuration Wizard

Active Directory Computer Account Password

Automate computer account password updates?

☑ Automate computer account password updates

Days between password updates: 7

Step 18. Accept the network card settings.

Step 19. Click Next.

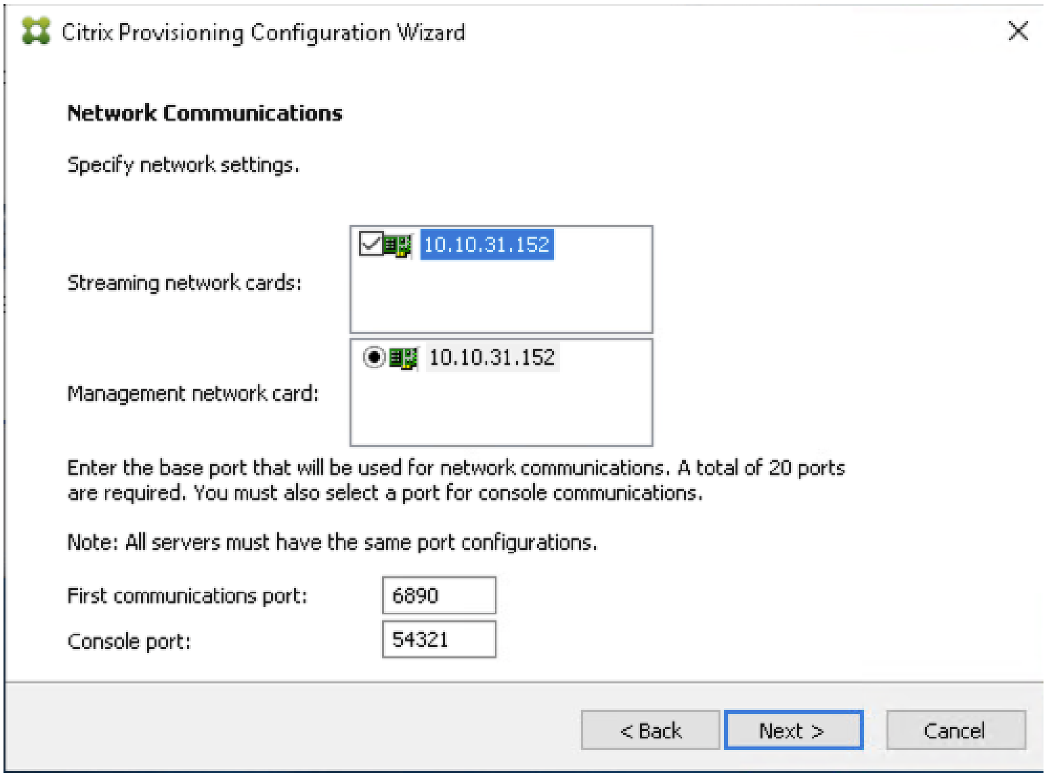

Citrix Provisioning Configuration Wizard

Network Communications

Specify network settings.

Streaming network cards: ☑ 10.10.31.152

Management network card: ● 10.10.31.152

Enter the base port that will be used for network communications. A total of 20 ports are required. You must also select a port for console communications.

Note: All servers must have the same port configurations.

First communications port: 6890

Console port: 54321

Step 20. Select Use the Provisioning Services TFTP service checkbox.

Step 21. Click Next.

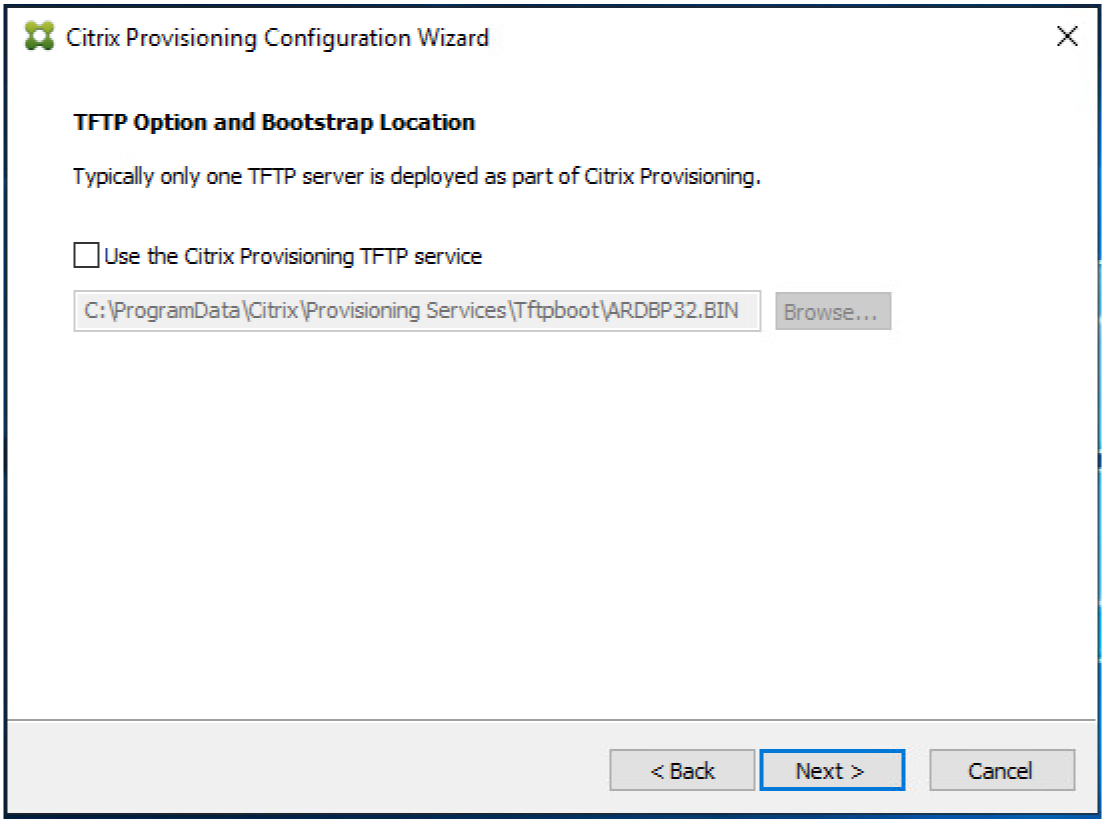

TFTP Option and Bootstrap Location

Typically only one TFTP server is deployed as part of Citrix Provisioning.

□ Use the Citrix Provisioning TFTP service

C:\ProgramData\Citrix\Provisioning Services\Tftpboot\ARDBP32.BIN

Step 22. If Soap Server is used, provide details.

Step 23. Click Next.

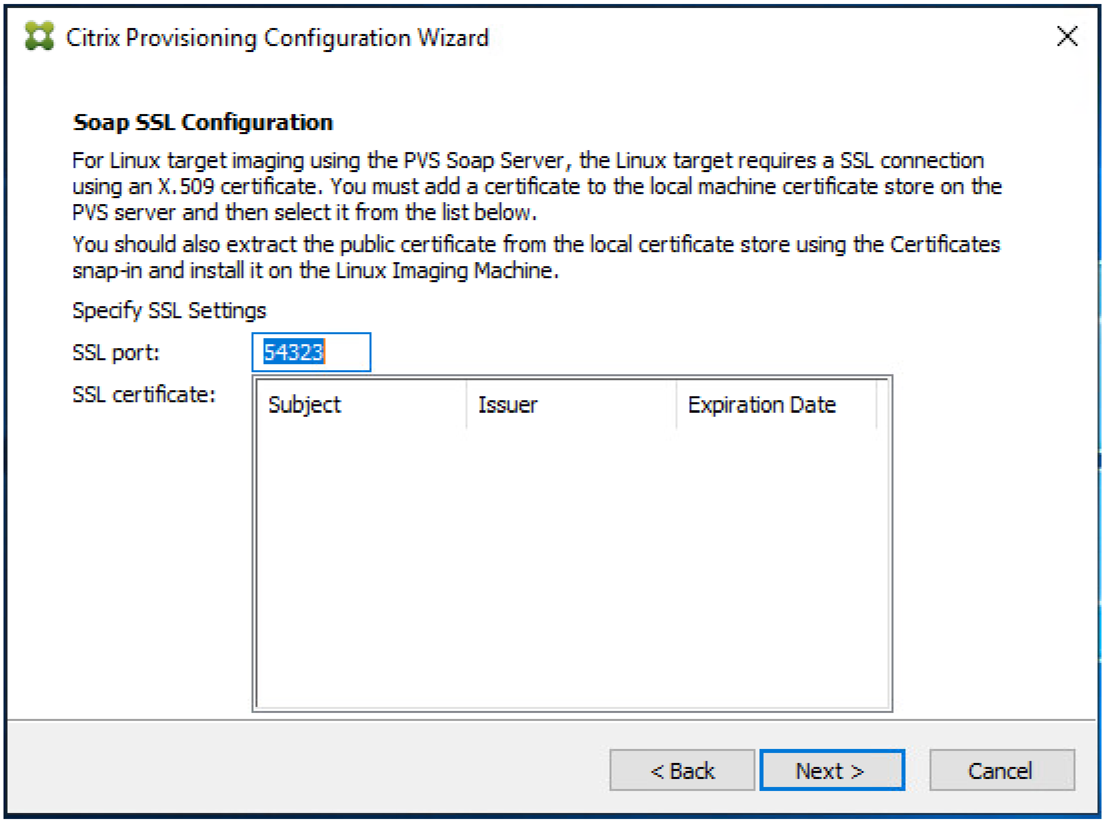

Soap SSL Configuration

For Linux target imaging using the PVS Soap Server, the Linux target requires a SSL connection using an X.509 certificate. You must add a certificate to the local machine certificate store on the PVS server and then select it from the list below.

You should also extract the public certificate from the local certificate store using the Certificates snap-in and install it on the Linux Imaging Machine.

Specify SSL Settings

SSL port: 54323

SSL certificate:

Subject	Issuer	Expiration Date

Step 24. If desired, fill in Problem Report Configuration.

Step 25. Click Next.

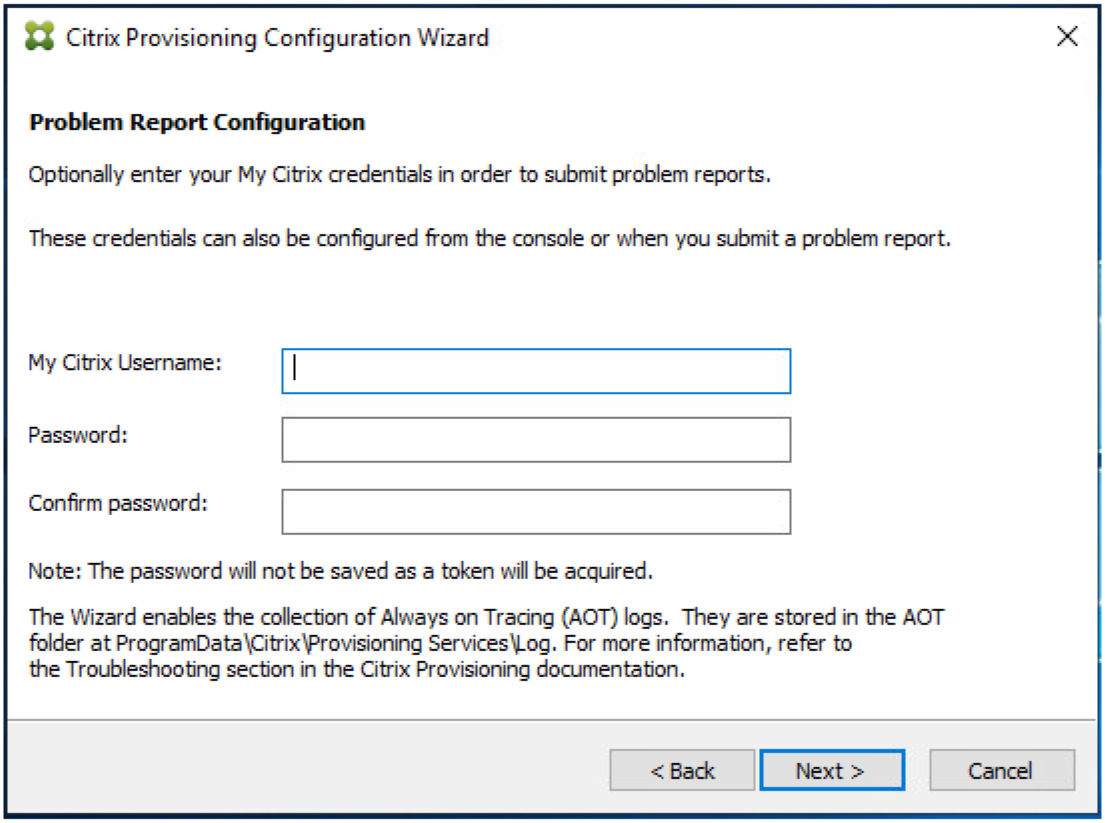

Problem Report Configuration

Optionally enter your My Citrix credentials in order to submit problem reports.

These credentials can also be configured from the console or when you submit a problem report.

My Citrix Username:

Password:

Confirm password:

Note: The password will not be saved as a token will be acquired.

The Wizard enables the collection of Always on Tracing (AOT) logs. They are stored in the AOT folder at ProgramData\Citrix\Provisioning Services\Log. For more information, refer to the Troubleshooting section in the Citrix Provisioning documentation.

Step 26. Click Finish to start the installation process.

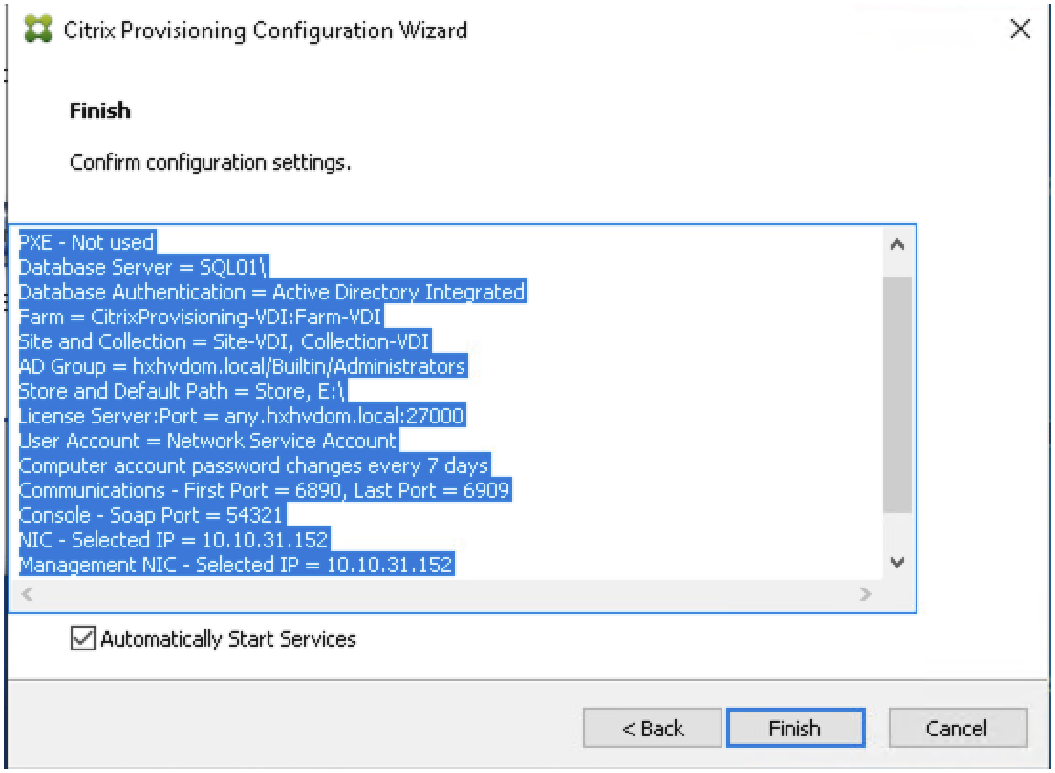

Finish

Confirm configuration settings.

```
PXE - Not used
Database Server = SQL01\
Database Authentication = Active Directory Integrated
Farm = CitrixProvisioning-VDI:Farm-VDI
Site and Collection = Site-VDI, Collection-VDI
AD Group = hxhvdom.local/Builtin/Administrators
Store and Default Path = Store, E:\
License Server:Port = any.hxhvdom.local:27000
User Account = Network Service Account
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.10.31.152
Management NIC - Selected IP = 10.10.31.152
```

☑ Automatically Start Services

Step 27. Click Done when the installation finishes.

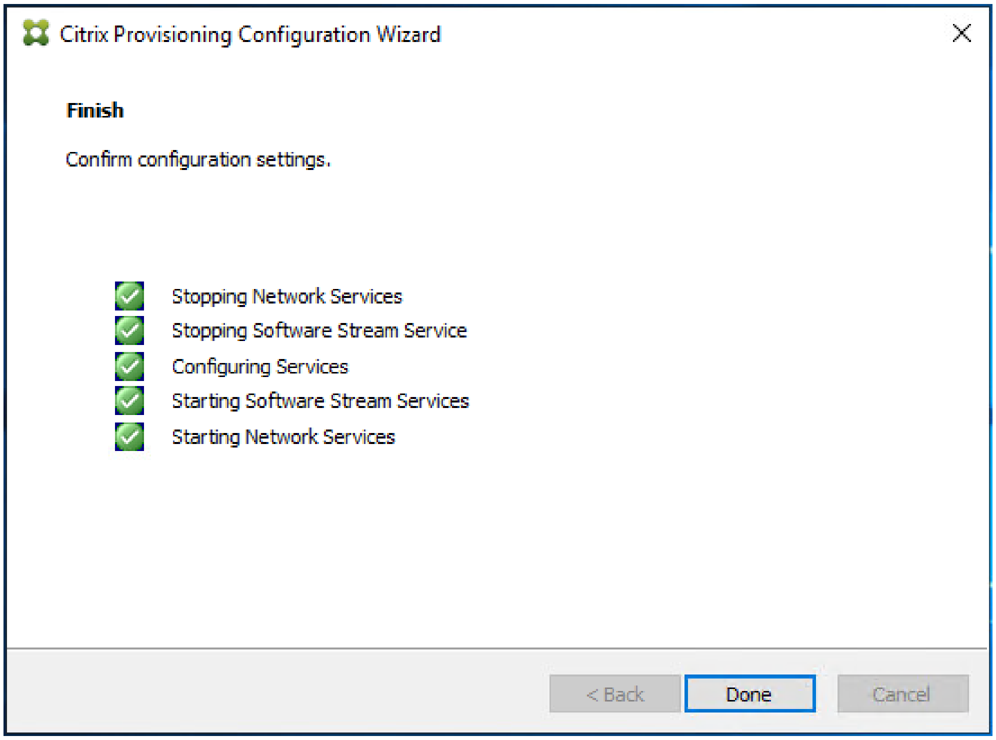

Finish

Confirm configuration settings.

✓ Stopping Network Services
✓ Stopping Software Stream Service
✓ Configuring Services
✓ Starting Software Stream Services
✓ Starting Network Services

Note: You can optionally install the Provisioning Services console on the second PVS server following the procedure in the section [Install the Citrix Provisioning Server Target Device Software](#).

Step 28. After completing the steps to install the three additional PVS servers, launch the Provisioning Services Console to verify that the PVS Servers and Stores are configured and that DHCP boot options are defined.

Step 29. Launch Provisioning Services Console and select Connect to Farm.

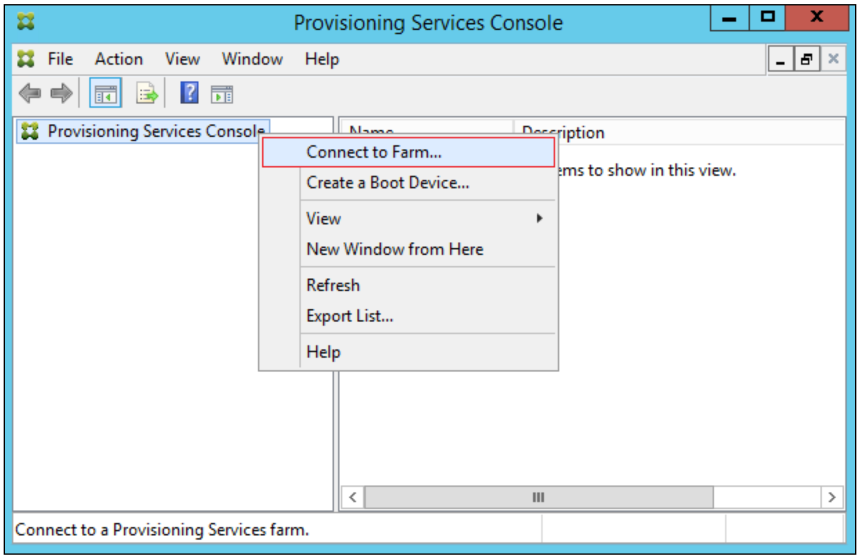

Step 30. Enter localhost for the PVS1 server.

Step 31. Click Connect.

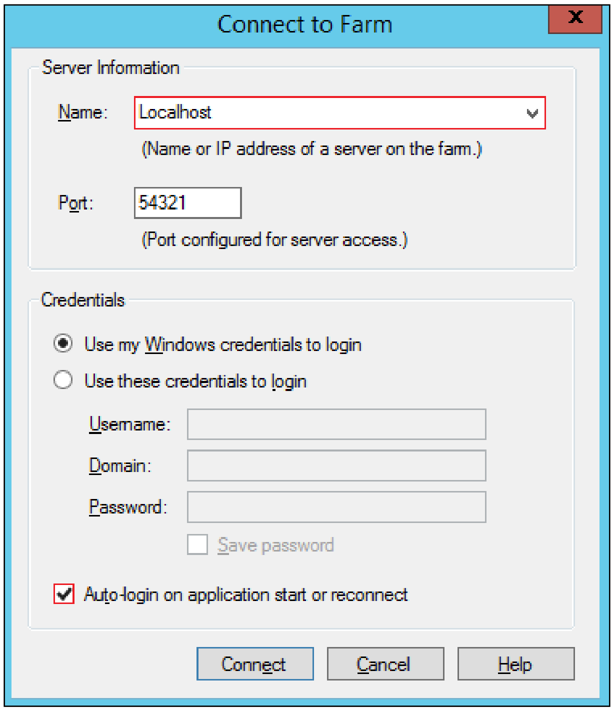

Name: Localhost

(Name or IP address of a server on the farm.)

Port: 54321

(Port configured for server access.)

● Use my Windows credentials to login

○ Use these credentials to login

Username:

Domain:

Password:

□ Save password

☑ Auto-login on application start or reconnect

Step 32. Select Store Properties from the drop-down list.

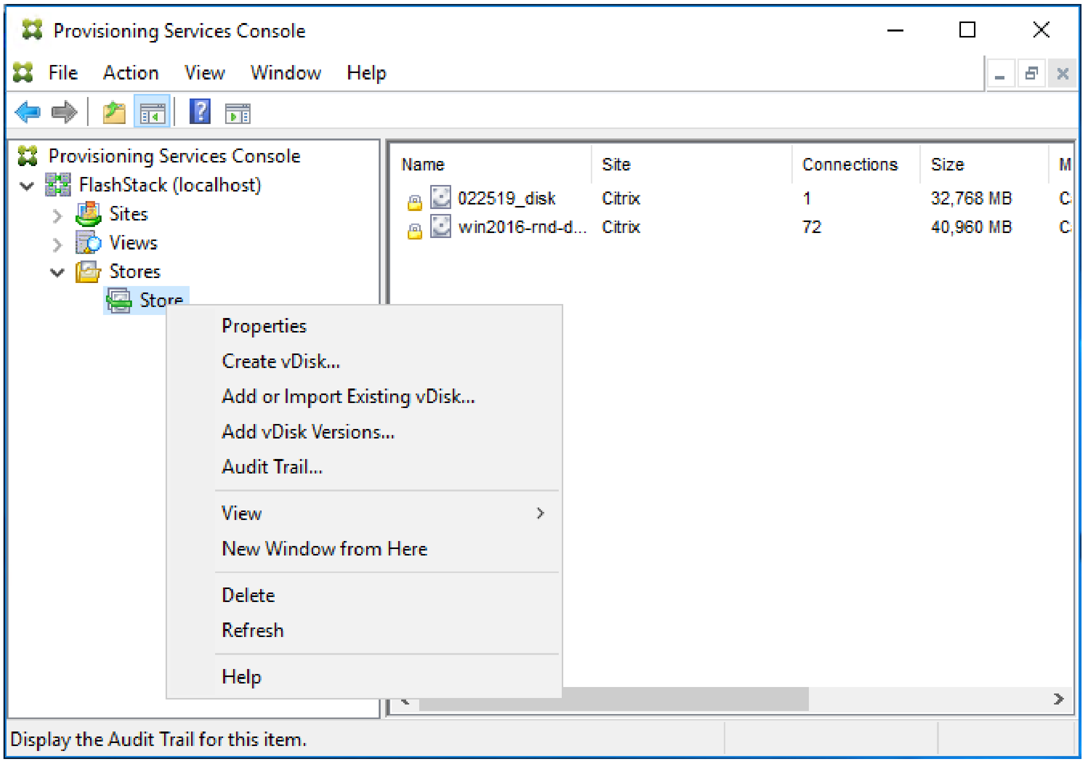

Name	Site	Connections	Size	M
022519_disk	Citrix	1	32,768 MB	C
win2016-rnd-d...	Citrix	72	40,960 MB	C

Display the Audit Trail for this item.

Step 33. In the Store Properties dialog, add the Default store path to the list of Default write cache paths.

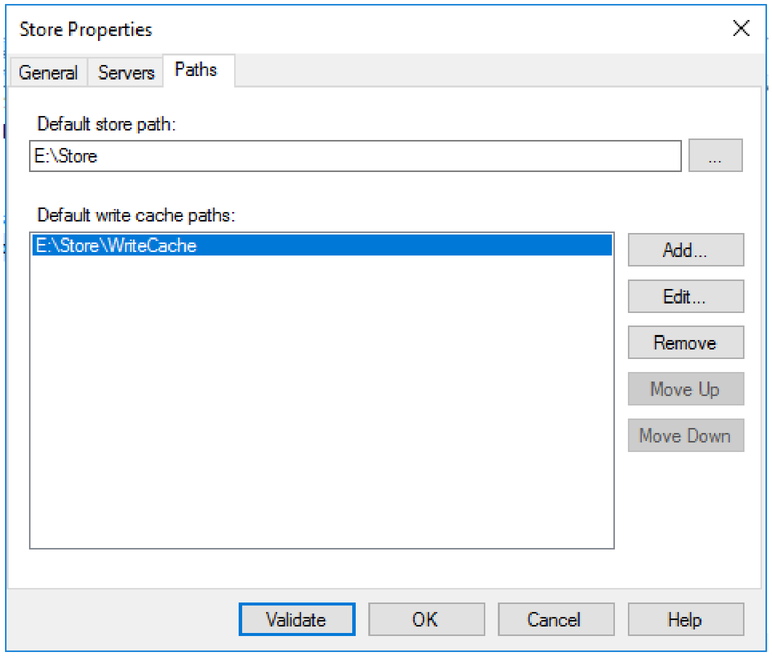

Default store path: E:\Store

Default write cache paths: E:\Store\WriteCache

Step 34. Click Validate. If the validation is successful, click Close and then click OK to continue.

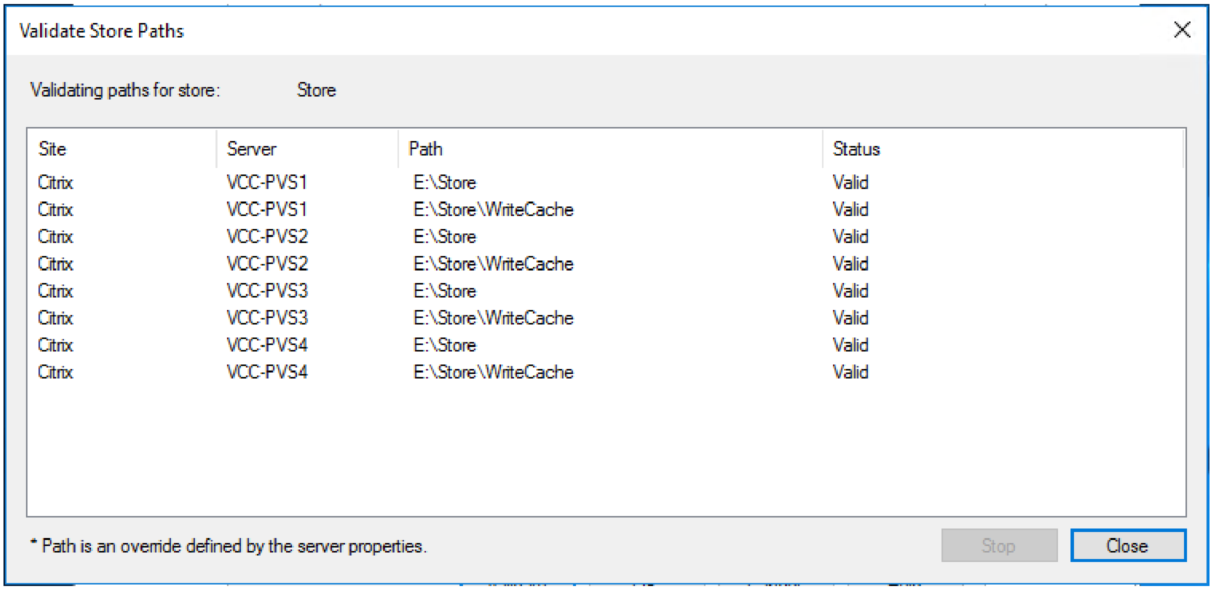

Validating paths for store: Store

Site	Server	Path	Status
Citrix	VCC-PVS1	E:\Store	Valid
Citrix	VCC-PVS1	E:\Store\WriteCache	Valid
Citrix	VCC-PVS2	E:\Store	Valid
Citrix	VCC-PVS2	E:\Store\WriteCache	Valid
Citrix	VCC-PVS3	E:\Store	Valid
Citrix	VCC-PVS3	E:\Store\WriteCache	Valid
Citrix	VCC-PVS4	E:\Store	Valid
Citrix	VCC-PVS4	E:\Store\WriteCache	Valid

* Path is an override defined by the server properties.

Procedure 5. Install Citrix Virtual Apps and Desktops Virtual Desktop Agents

Virtual Delivery Agents (VDAs) are installed on the server and workstation operating systems and enable connections for desktops and apps. The following procedure was used to install VDAs for both Single-session OS and Multi-session OS.

By default, when you install the Virtual Delivery Agent, Citrix User Profile Management is installed silently on master images. (Using profile management as a profile solution is optional, but FSLogix was used for this CVD and is described in a later section.)

Step 1. Launch the Citrix Virtual Apps and Desktops installer from the Citrix_Virtual_Apps_and_Desktops_7_2203 ISO.

Step 2. Click Start on the Welcome Screen.

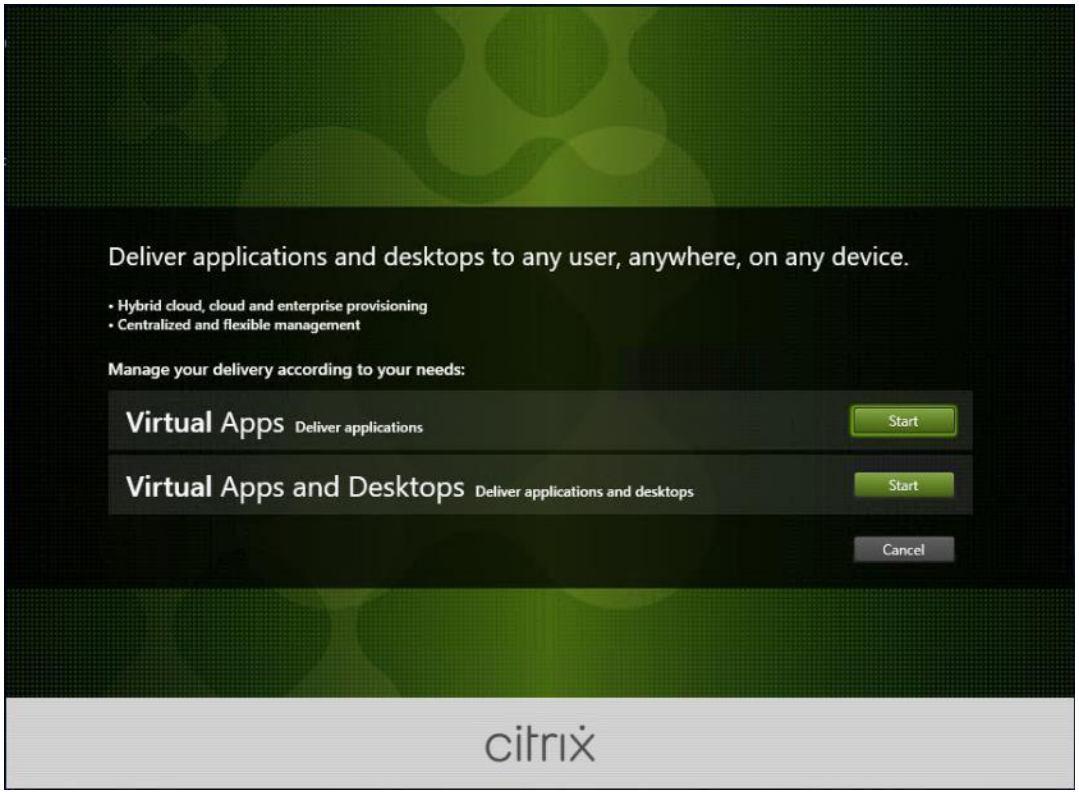

Deliver applications and desktops to any user, anywhere, on any device.

- Hybrid cloud, cloud and enterprise provisioning
- Centralized and flexible management

Manage your delivery according to your needs:

Virtual Apps Deliver applications
Start
Virtual Apps and Desktops Deliver applications and desktops
Start
Cancel
citrix

Step 3. To install the VDA for the Hosted Virtual Desktops (VDI), select Virtual Delivery Agent for Windows Single-session OS.

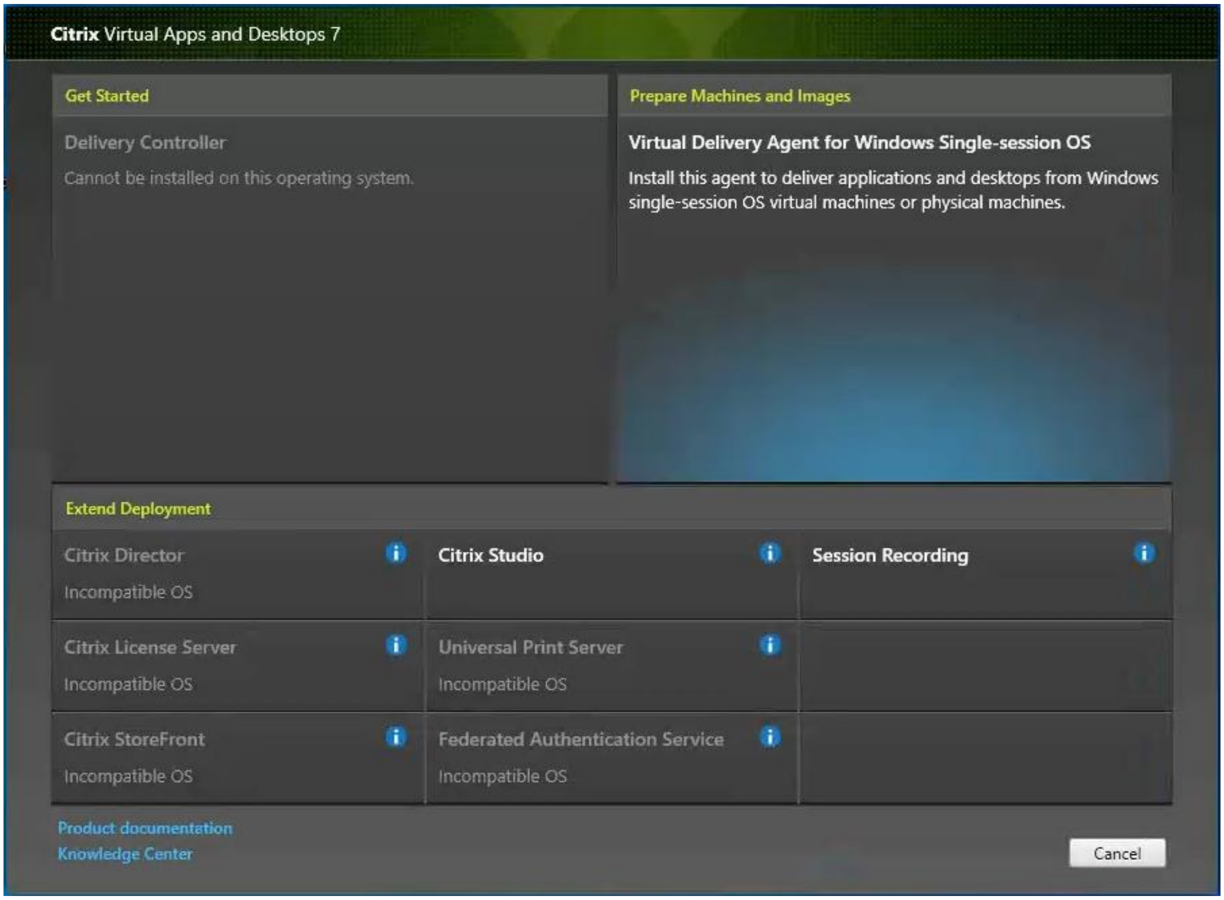

Citrix Virtual Apps and Desktops 7

Get Started

Delivery Controller

Cannot be installed on this operating system.

Prepare Machines and Images

Virtual Delivery Agent for Windows Single-session OS

Install this agent to deliver applications and desktops from Windows single-session OS virtual machines or physical machines.

Extend Deployment

Citrix Director

Incompatible OS

Citrix Studio

Session Recording

Citrix License Server

Incompatible OS

Universal Print Server

Incompatible OS

Citrix StoreFront

Incompatible OS

Federated Authentication Service

Incompatible OS

Note: When installing Virtual Delivery Agent for Windows Multi-session OS and follow the same basic steps.

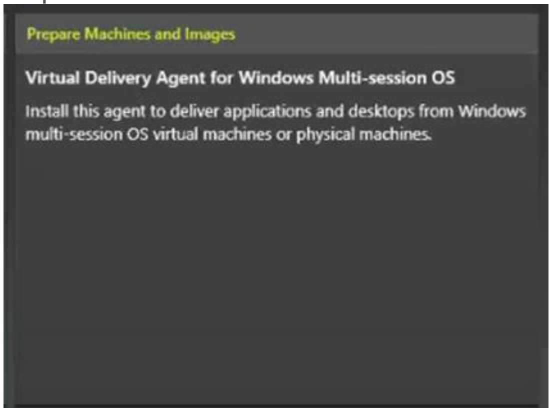

Prepare Machines and Images

Virtual Delivery Agent for Windows Multi-session OS

Install this agent to deliver applications and desktops from Windows multi-session OS virtual machines or physical machines.

Step 4. Select Create a master MCS Image.

Step 5. Click Next.

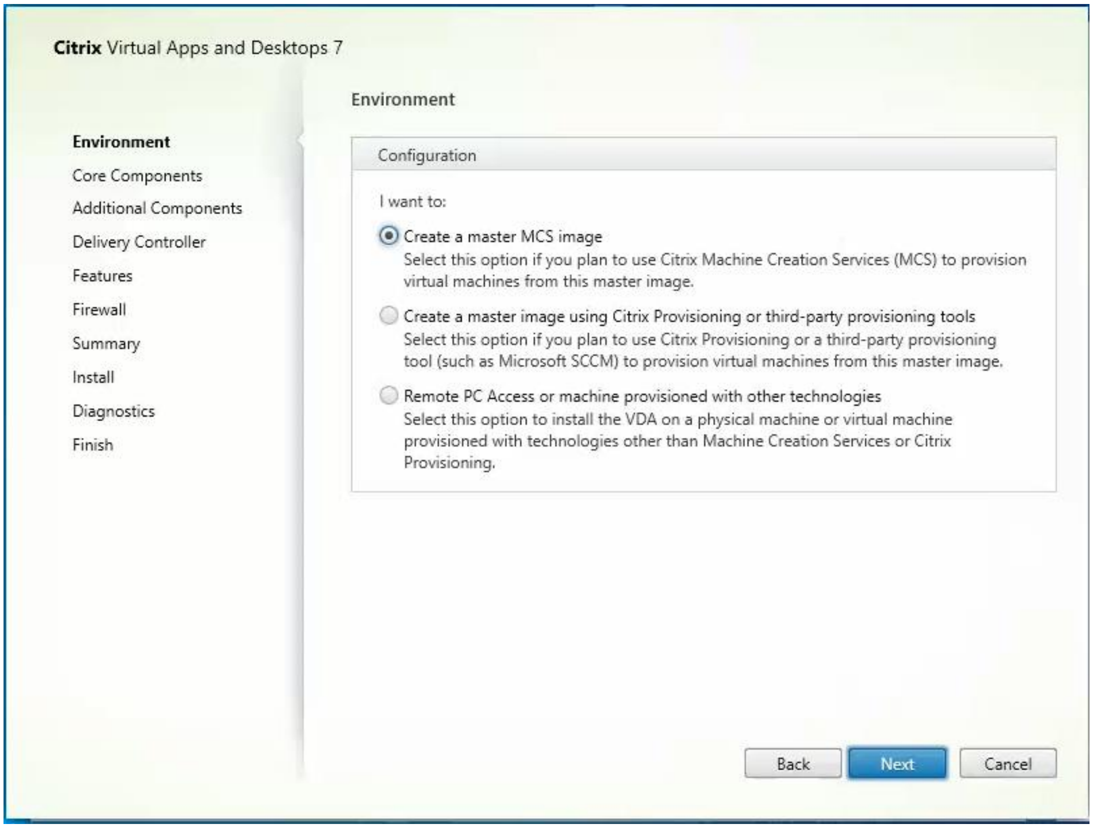

Citrix Virtual Apps and Desktops 7

Environment

Configuration

I want to:

● Create a master MCS image
Select this option if you plan to use Citrix Machine Creation Services (MCS) to provision virtual machines from this master image.

○ Create a master image using Citrix Provisioning or third-party provisioning tools
Select this option if you plan to use Citrix Provisioning or a third-party provisioning tool (such as Microsoft SCCM) to provision virtual machines from this master image.

○ Remote PC Access or machine provisioned with other technologies
Select this option to install the VDA on a physical machine or virtual machine provisioned with technologies other than Machine Creation Services or Citrix Provisioning.

Step 6. Select "Create a master image using Citrix Provisioning or third-party provisioning tools" when building image to be delivered with Citrix Provisioning tools.

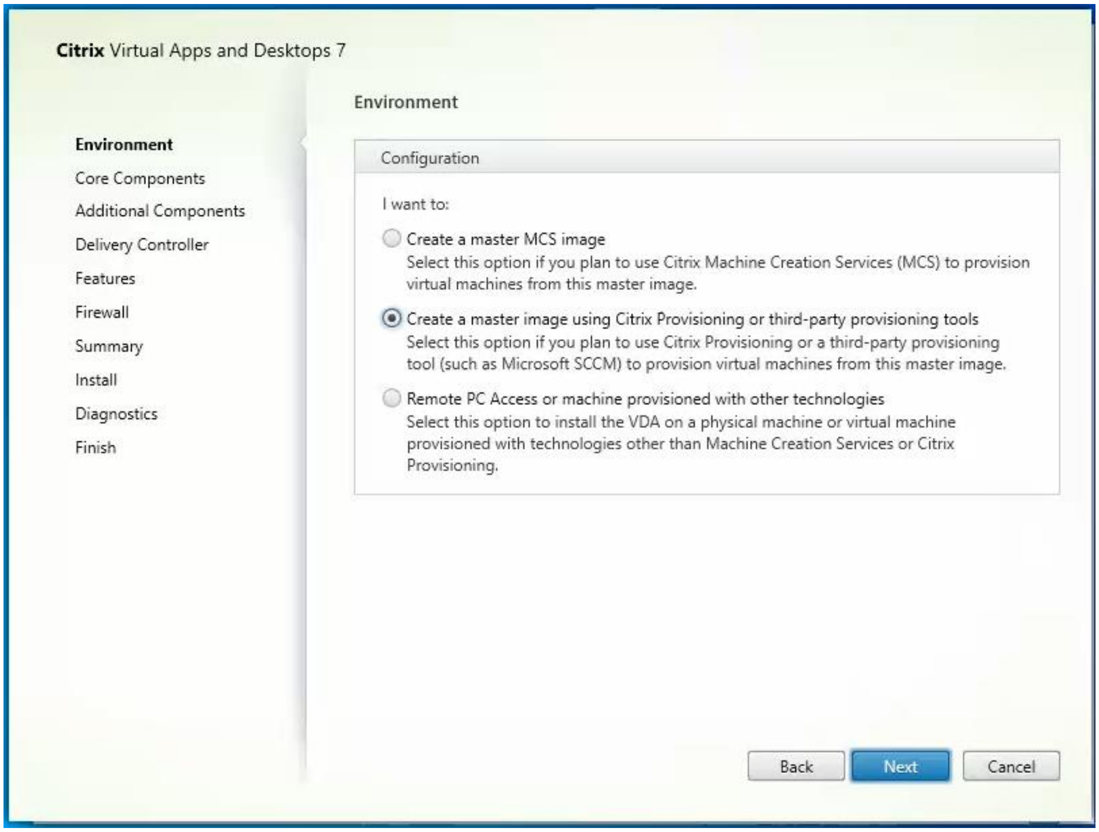

Citrix Virtual Apps and Desktops 7

Environment

Configuration

I want to:

○ Create a master MCS image
Select this option if you plan to use Citrix Machine Creation Services (MCS) to provision virtual machines from this master image.

● Create a master image using Citrix Provisioning or third-party provisioning tools
Select this option if you plan to use Citrix Provisioning or a third-party provisioning tool (such as Microsoft SCCM) to provision virtual machines from this master image.

○ Remote PC Access or machine provisioned with other technologies
Select this option to install the VDA on a physical machine or virtual machine provisioned with technologies other than Machine Creation Services or Citrix Provisioning.

Step 7. Optionally, do not select Citrix Workspace App**.**

Step 8. Click Next.

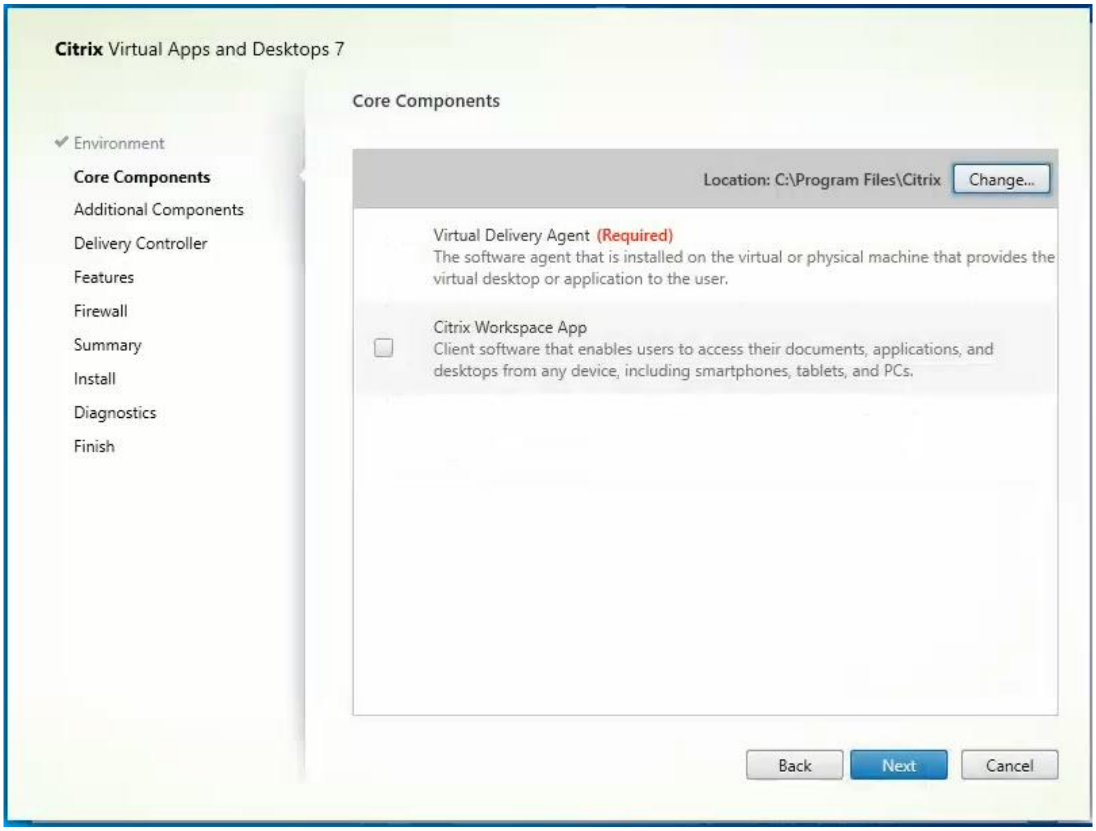

Citrix Virtual Apps and Desktops 7

Core Components

Location: C:\Program Files\Citrix

Virtual Delivery Agent (Required)
The software agent that is installed on the virtual or physical machine that provides the virtual desktop or application to the user.

□ Citrix Workspace App
Client software that enables users to access their documents, applications, and desktops from any device, including smartphones, tablets, and PCs.

Step 9. Select the additional components required for your image.

Note: In this design, only the default components were installed on the image.

Step 10. Click Next.

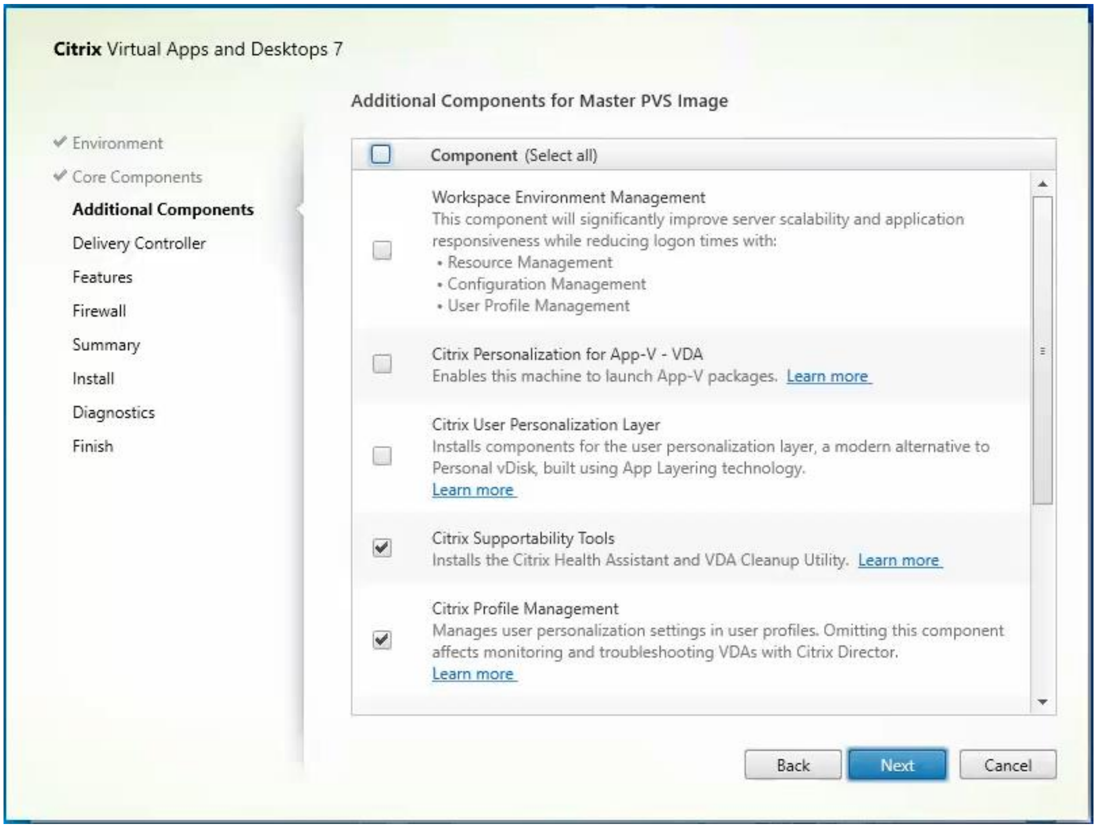

Citrix Virtual Apps and Desktops 7

Additional Components for Master PVS Image

	Component (Select all)
□	Workspace Environment Management This component will significantly improve server scalability and application responsiveness while reducing logon times with: • Resource Management • Configuration Management • User Profile Management
□	Citrix Personalization for App-V - VDA Enables this machine to launch App-V packages. Learn more
□	Citrix User Personalization Layer Installs components for the user personalization layer, a modern alternative to Personal vDisk, built using App Layering technology. Learn more
☑	Citrix Supportability Tools Installs the Citrix Health Assistant and VDA Cleanup Utility. Learn more
☑	Citrix Profile Management Manages user personalization settings in user profiles. Omitting this component affects monitoring and troubleshooting VDAs with Citrix Director. Learn more

Step 11. Configure Delivery Controllers at this time.

Step 12. Click Next.

Step 13. Optional, select additional features**.**

Step 14. Click Next.

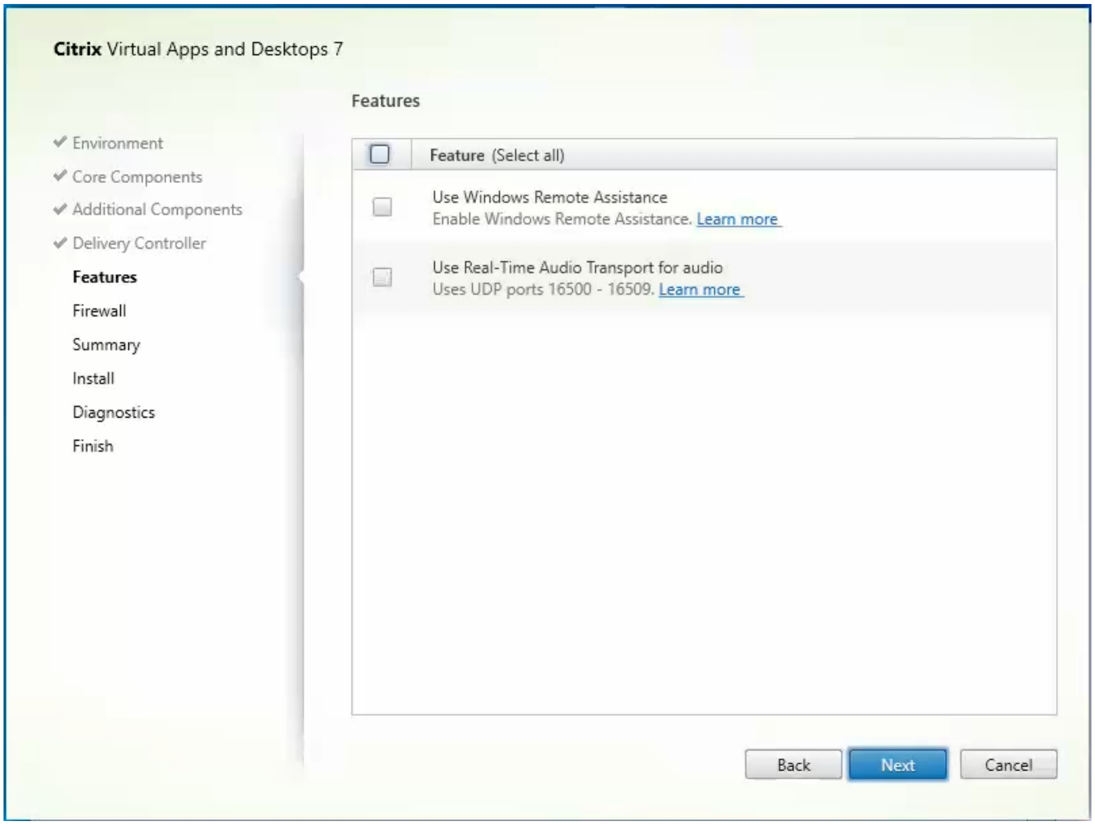

Features

□	Feature (Select all)
□	Use Windows Remote Assistance Enable Windows Remote Assistance. Learn more
□	Use Real-Time Audio Transport for audio Uses UDP ports 16500 - 16509. Learn more

Step 15. Allow the firewall rules to configure Automatically.

Step 16. Click Next.

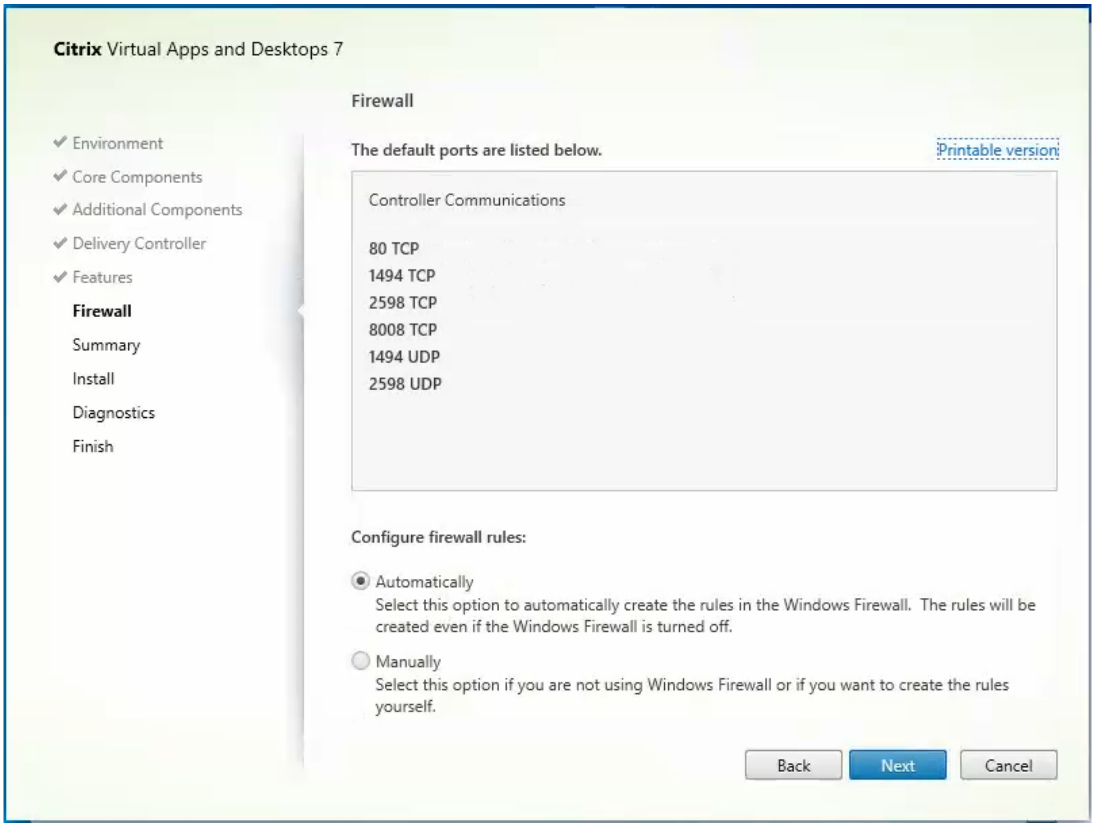

Firewall

The default ports are listed below.

Controller Communications

80 TCP
1494 TCP
2598 TCP
8008 TCP
1494 UDP
2598 UDP

Configure firewall rules:

● Automatically
Select this option to automatically create the rules in the Windows Firewall. The rules will be created even if the Windows Firewall is turned off.

○ Manually
Select this option if you are not using Windows Firewall or if you want to create the rules yourself.

Step 17. Verify the Summary and click Install**.**

Step 18. Optional, configure Citrix Call Home participation.

Step 19. Click Next.

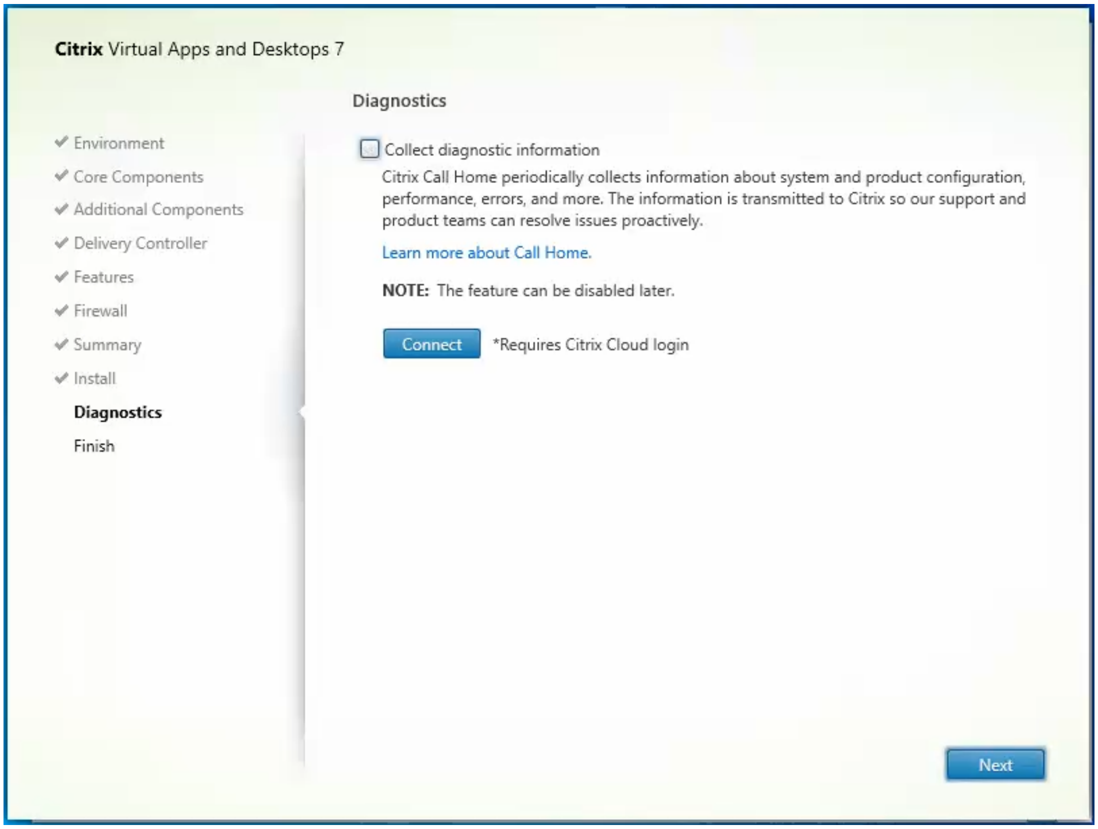

Diagnostics

□ Collect diagnostic information

Citrix Call Home periodically collects information about system and product configuration, performance, errors, and more. The information is transmitted to Citrix so our support and product teams can resolve issues proactively.

Learn more about Call Home.

NOTE: The feature can be disabled later.

Connect *Requires Citrix Cloud login

Step 20. Check Restart Machine.

Step 21. Click Finish and the machine will reboot automatically.

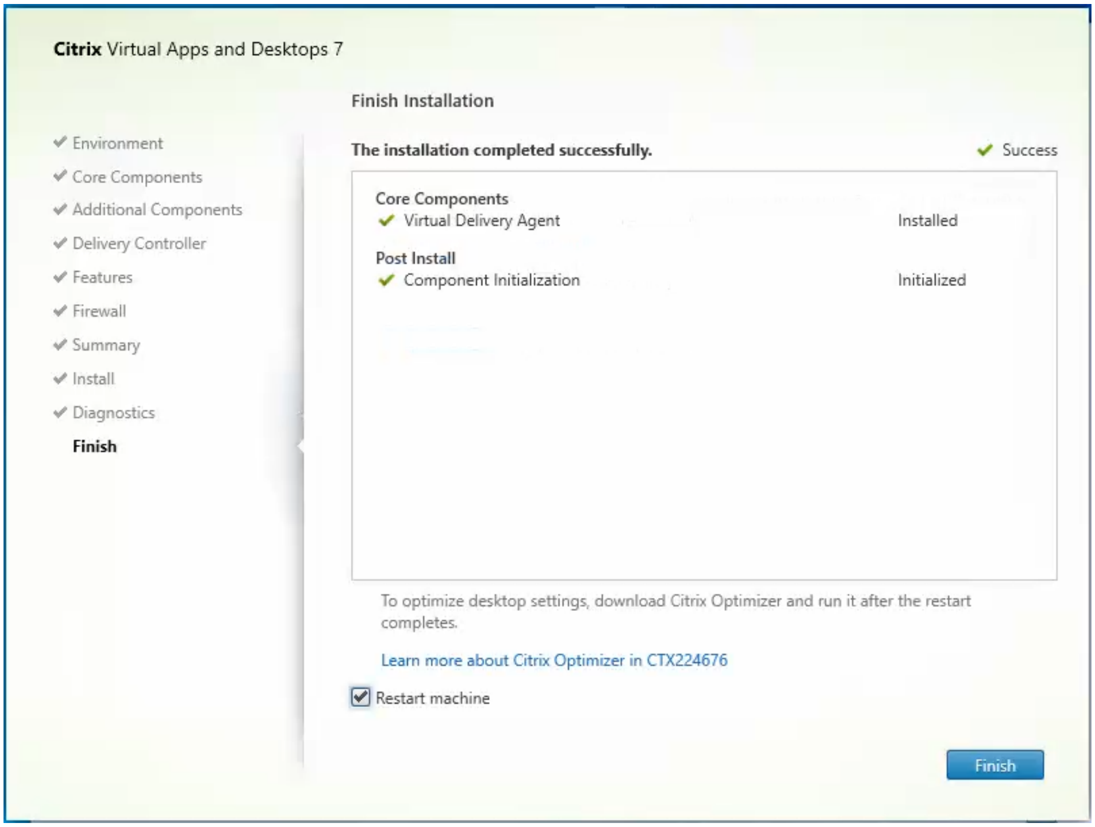

Citrix Virtual Apps and Desktops 7

Finish Installation

The installation completed successfully. ✓ Success

Core Components
✓ Virtual Delivery Agent Installed

Post Install
✓ Component Initialization Initialized

To optimize desktop settings, download Citrix Optimizer and run it after the restart completes.

Learn more about Citrix Optimizer in CTX224676

☑ Restart machine

Install the Citrix Provisioning Server Target Device Software

The Master Target Device refers to the target device from which a hard disk image is built and stored on a vDisk. Provisioning Services then streams the contents of the vDisk created to other target devices. This procedure installs the PVS Target Device software that is used to build the RDS and VDI golden images.

Procedure 1. Install the Citrix Provisioning Server Target Device software

Step 1. Launch the PVS installer from the Citrix_Provisioning_2203 ISO.

Step 2. Click Target Device Installation.

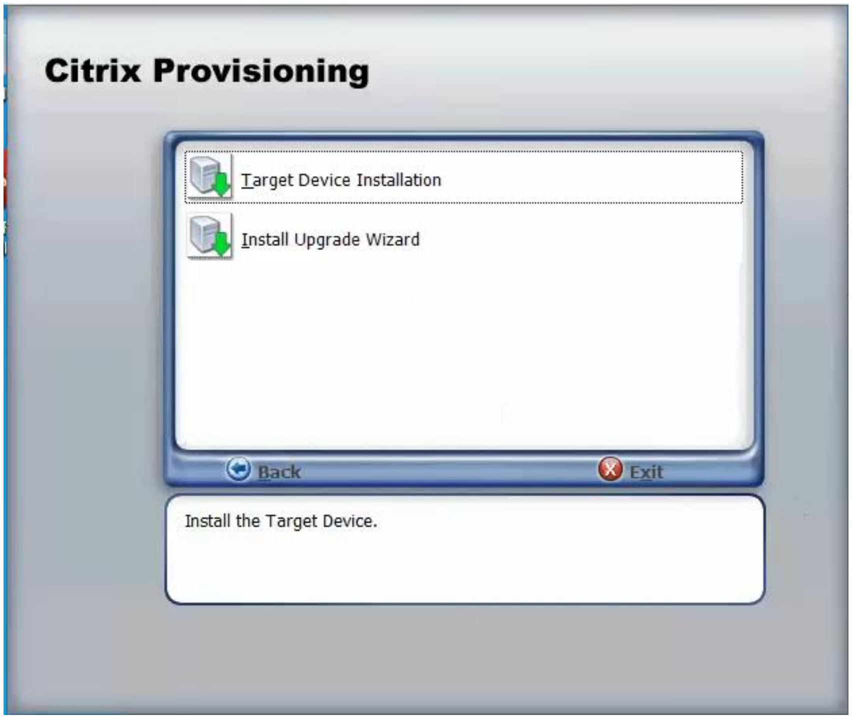

Citrix Provisioning

Target Device Installation

Install Upgrade Wizard

Install the Target Device.

Note: The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.

Step 3. Click Next.

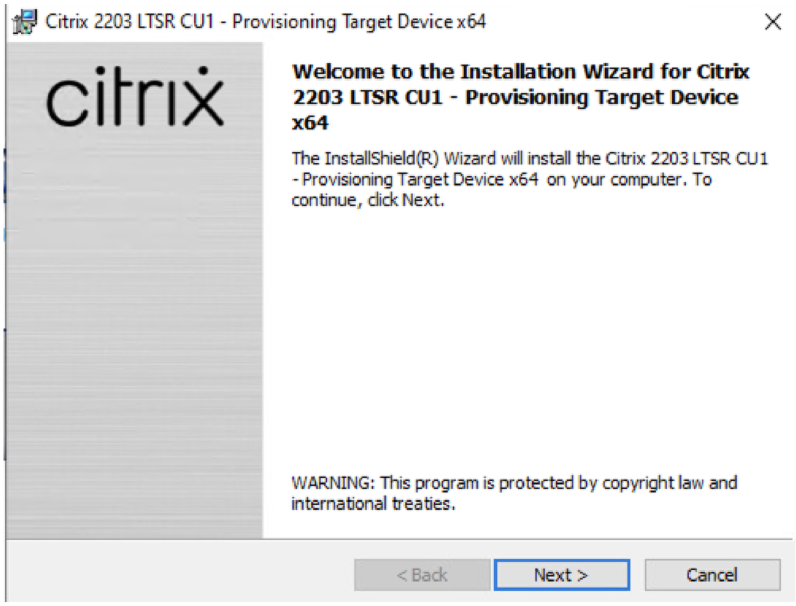

Citrix 2203 LTSR CU1 - Provisioning Target Device x64

Welcome to the Installation Wizard for Citrix 2203 LTSR CU1 - Provisioning Target Device x64

The InstallShield(R) Wizard will install the Citrix 2203 LTSR CU1 - Provisioning Target Device x64 on your computer. To continue, click Next.

WARNING: This program is protected by copyright law and international treaties.

Step 4. Indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

Step 5. Click Next.

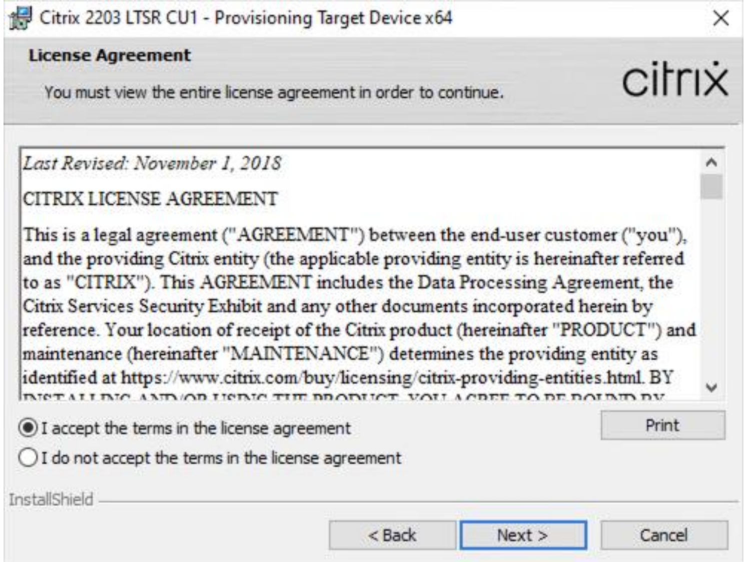

Citrix 2203 LTSR CU1 - Provisioning Target Device x64

License Agreement

You must view the entire license agreement in order to continue.

Last Revised: November 1, 2018

CITRIX LICENSE AGREEMENT

This is a legal agreement ("AGREEMENT") between the end-user customer ("you"), and the providing Citrix entity (the applicable providing entity is hereinafter referred to as "CITRIX"). This AGREEMENT includes the Data Processing Agreement, the Citrix Services Security Exhibit and any other documents incorporated herein by reference. Your location of receipt of the Citrix product (hereinafter "PRODUCT") and maintenance (hereinafter "MAINTENANCE") determines the providing entity as identified at https://www.citrix.com/buy/licensing/citrix-providing-entities.html. BY

● I accept the terms in the license agreement

○ I do not accept the terms in the license agreement

Step 6. Optionally, provide the Customer information.

Step 7. Click Next.

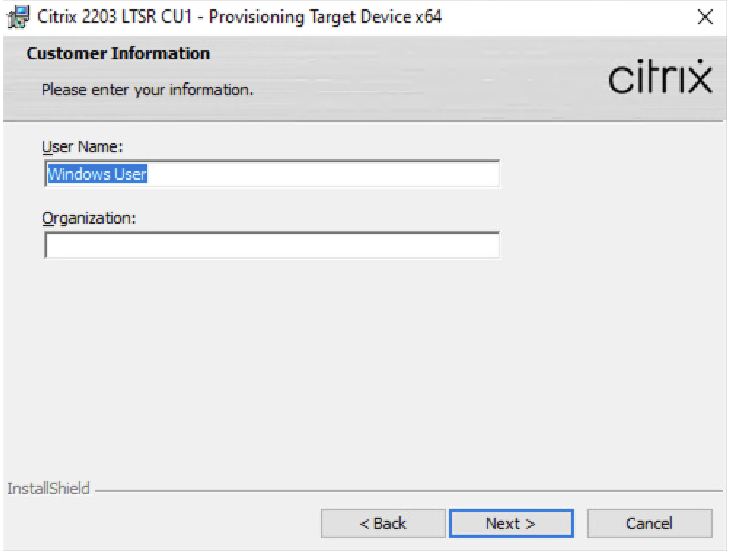

Citrix 2203 LTSR CU1 - Provisioning Target Device x64

Customer Information

Please enter your information.

User Name: Windows User

Organization:

Step 8. Accept the default installation path.

Step 9. Click Next.

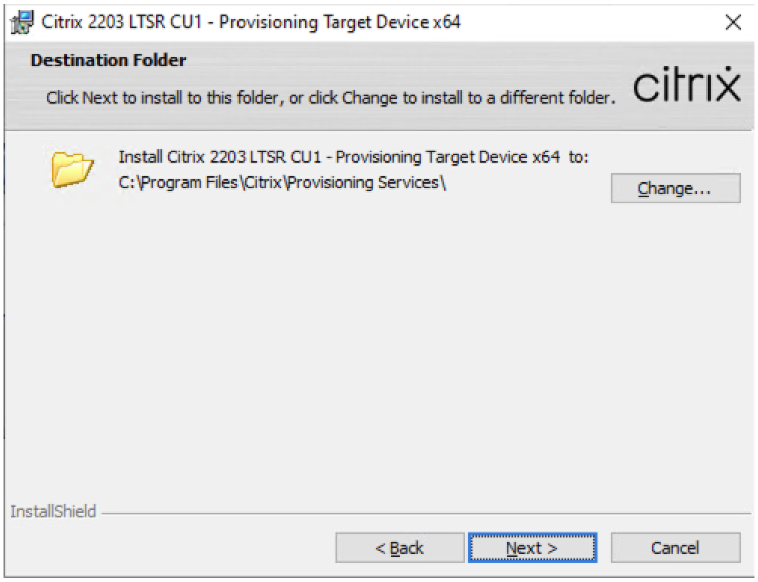

Citrix 2203 LTSR CU1 - Provisioning Target Device x64

Destination Folder

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

Install Citrix 2203 LTSR CU1 - Provisioning Target Device x64 to:
C:\Program Files\Citrix\Provisioning Services\

Step 10. Click Install.

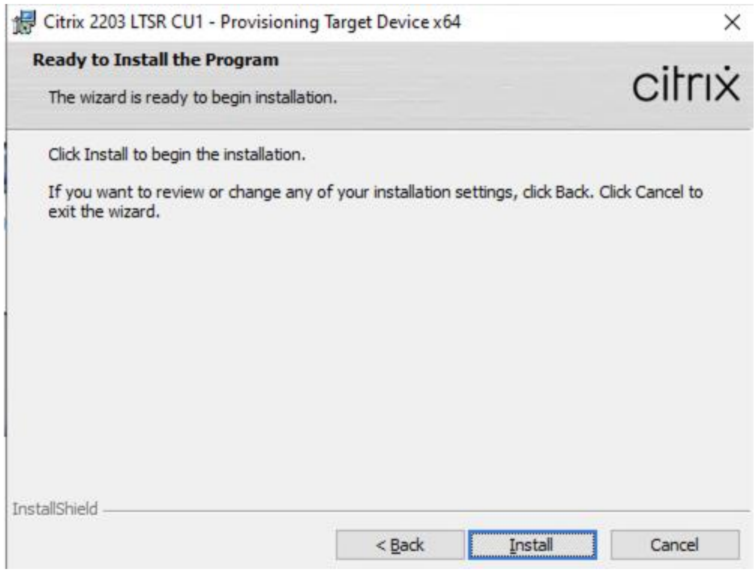

Citrix 2203 LTSR CU1 - Provisioning Target Device x64

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

Step 11. Deselect the checkbox to launch the Imaging Wizard and click Finish.

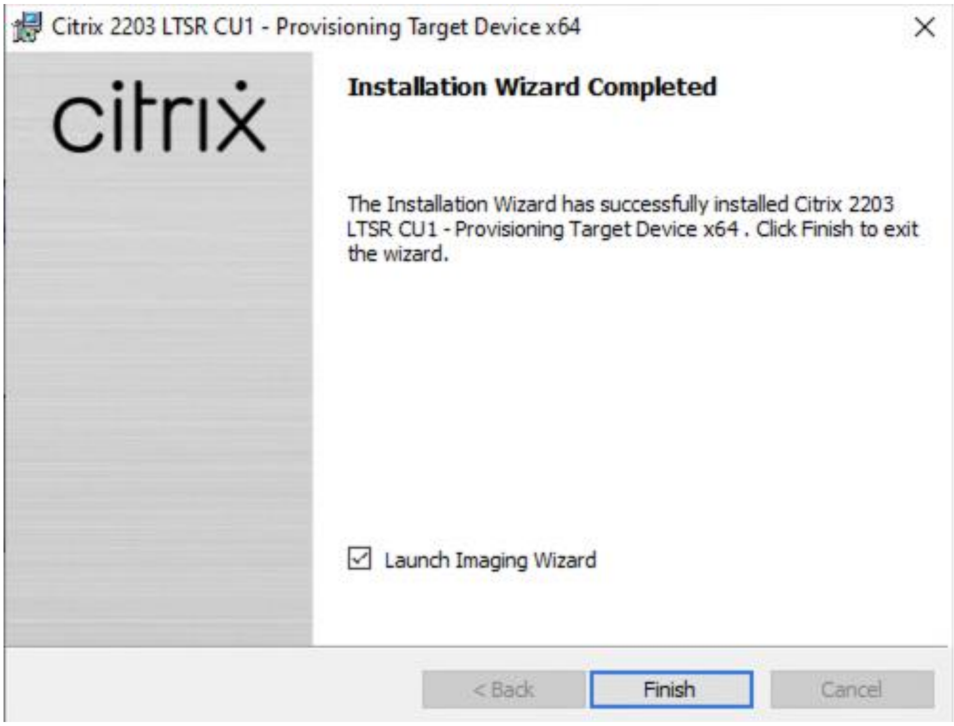

Citrix 2203 LTSR CU1 - Provisioning Target Device x64

Installation Wizard Completed

The Installation Wizard has successfully installed Citrix 2203 LTSR CU1 - Provisioning Target Device x64 . Click Finish to exit the wizard.

☑ Launch Imaging Wizard

Step 12. Click Yes to reboot the machine.

Procedure 2. Create Citrix Provisioning Server vDisks

The PVS Imaging Wizard automatically creates a base vDisk image from the master target device.

Step 1. The PVS Imaging Wizard's Welcome page appears.

Step 2. Click Next.

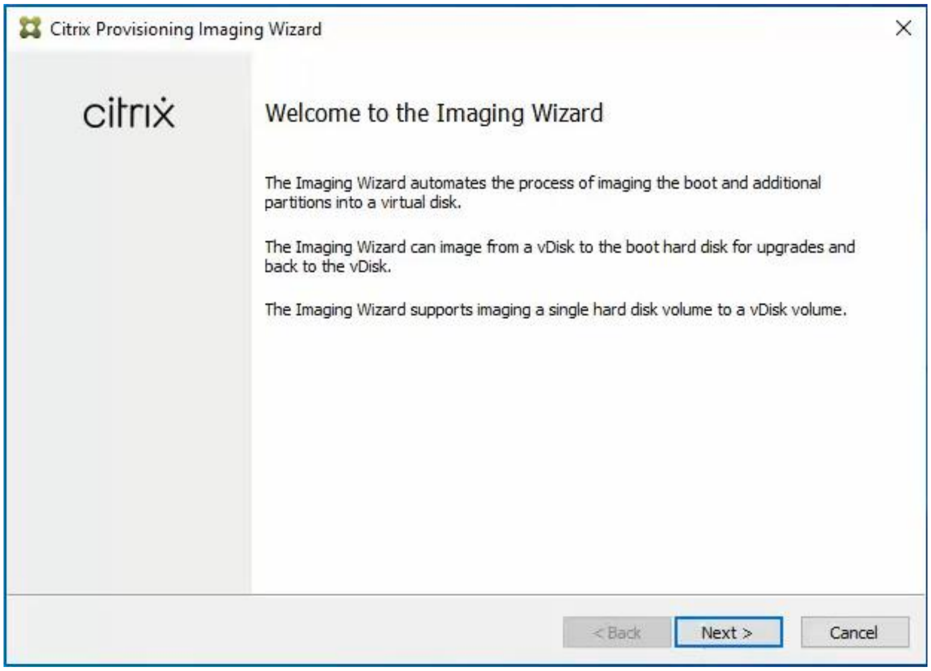

Citrix Provisioning Imaging Wizard

Welcome to the Imaging Wizard

The Imaging Wizard automates the process of imaging the boot and additional partitions into a virtual disk.

The Imaging Wizard can image from a vDisk to the boot hard disk for upgrades and back to the vDisk.

The Imaging Wizard supports imaging a single hard disk volume to a vDisk volume.

Step 3. The Connect to Farm page appears. Enter the name or IP address of a Provisioning Server within the farm to connect to and the port to use to make that connection.

Step 4. Use the Windows credentials (default) or enter different credentials.

Step 5. Click Next.

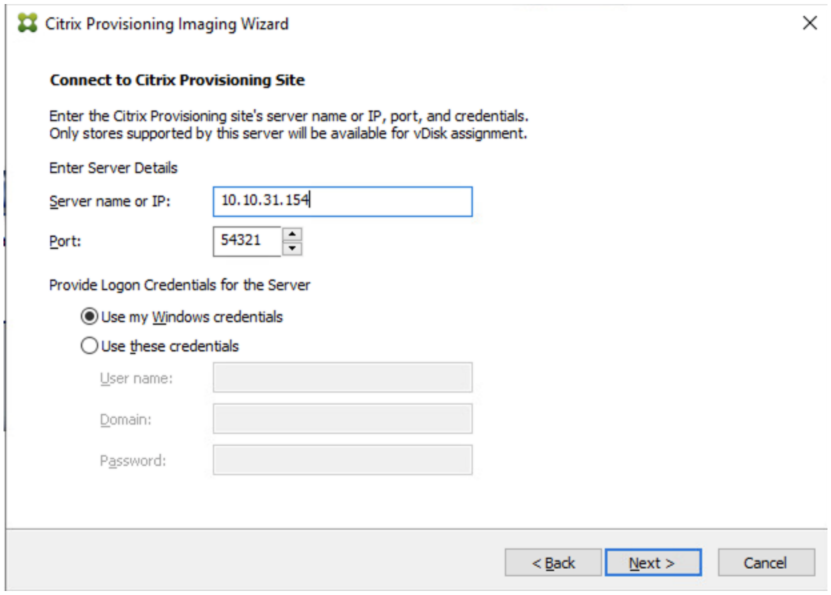

Citrix Provisioning Imaging Wizard

Connect to Citrix Provisioning Site

Enter the Citrix Provisioning site's server name or IP, port, and credentials.
Only stores supported by this server will be available for vDisk assignment.

Enter Server Details

Server name or IP: 10.10.31.154

Port: 54321

Provide Logon Credentials for the Server

● Use my Windows credentials

○ Use these credentials

User name:

Domain:

Password:

Step 6. Select Create a vDisk.

Step 7. Click Next.

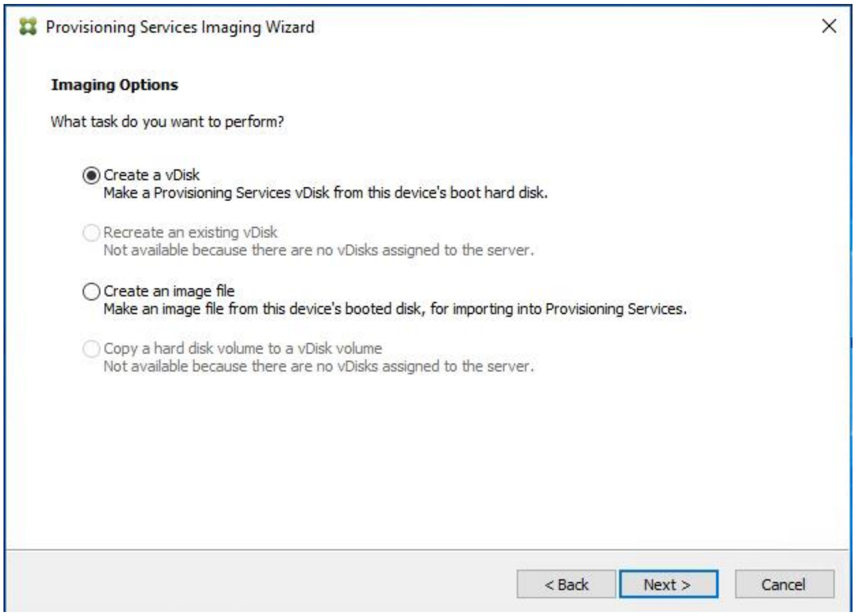

Provisioning Services Imaging Wizard

Imaging Options

What task do you want to perform?

● Create a vDisk
Make a Provisioning Services vDisk from this device's boot hard disk.

○ Recreate an existing vDisk
Not available because there are no vDisks assigned to the server.

○ Create an image file
Make an image file from this device's booted disk, for importing into Provisioning Services.

○ Copy a hard disk volume to a vDisk volume
Not available because there are no vDisks assigned to the server.

The Add Target Device page appears.

Step 8. Select the Target Device Name, the MAC address associated with one of the NICs that was selected when the target device software was installed on the master target device, and the Collection to which you are adding the device.

Step 9. Click Next.

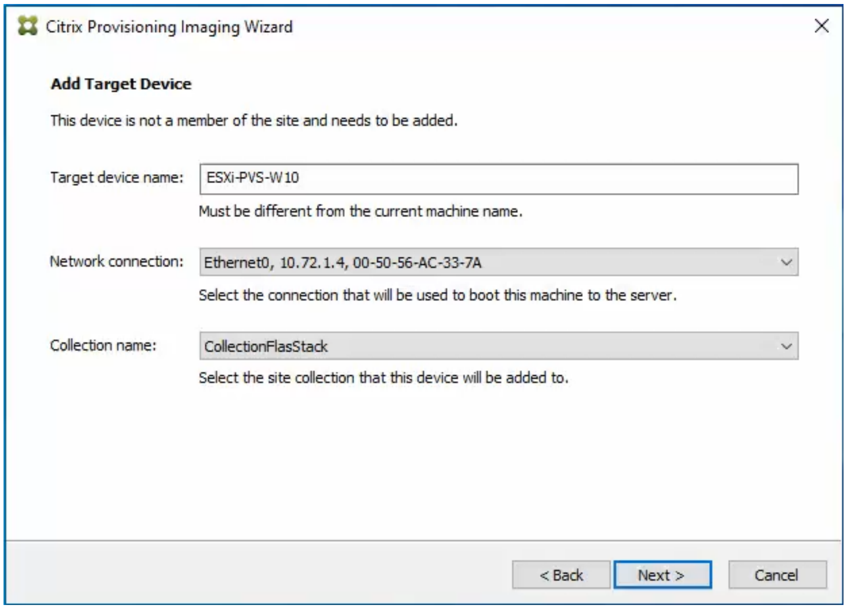

Citrix Provisioning Imaging Wizard

Add Target Device

This device is not a member of the site and needs to be added.

Target device name: ESXi-PVS-W10

Must be different from the current machine name.

Network connection: Ethernet0, 10.72.1.4, 00-50-56-AC-33-7A

Select the connection that will be used to boot this machine to the server.

Collection name: CollectionFlasStack

Select the site collection that this device will be added to.

Step 10. The New vDisk dialog displays. Enter the name of the vDisk.

Step 11. Select the Store where the vDisk will reside. Select the vDisk type, either Fixed or Dynamic, from the drop-down list.

Note: This CVD used Dynamic rather than Fixed vDisks.

Step 12. Click Next.

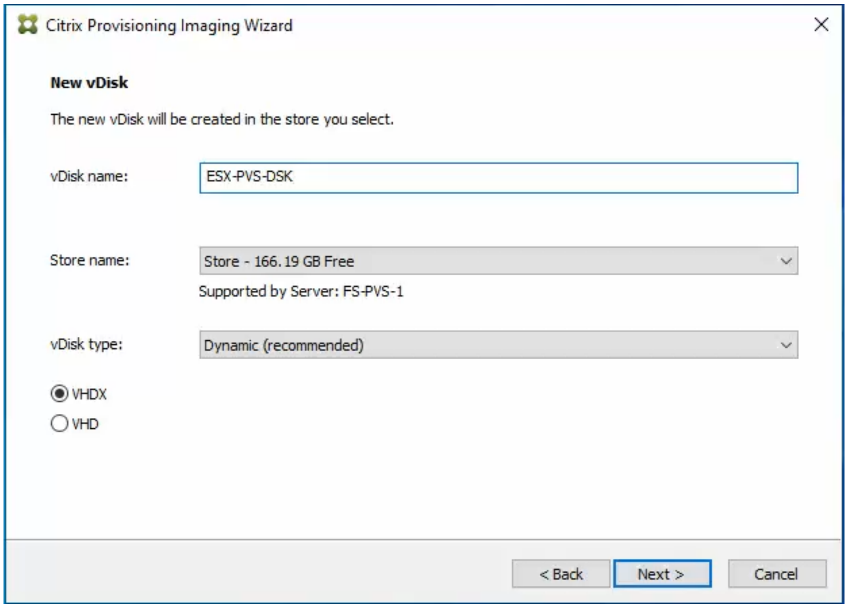

Citrix Provisioning Imaging Wizard

New vDisk

The new vDisk will be created in the store you select.

vDisk name: ESX-PVS-DSK

Store name: Store - 166.19 GB Free

Supported by Server: FS-PVS-1

vDisk type: Dynamic (recommended)

● VHDX

○ VHD

Step 13. On the Microsoft Volume Licensing page, select the volume license option to use for target devices. For this CVD, volume licensing is not used, so the None button is selected.

Step 14. Click Next.

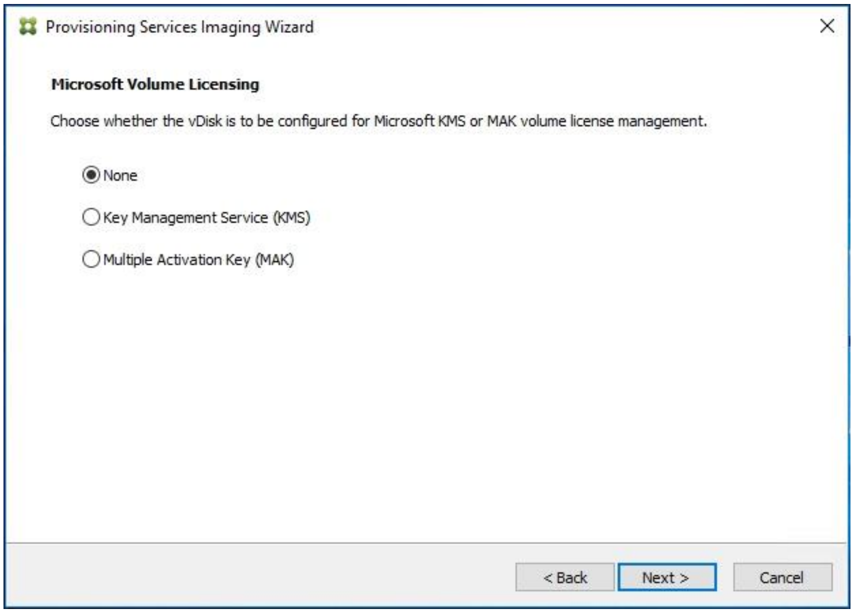

Provisioning Services Imaging Wizard

Microsoft Volume Licensing

Choose whether the vDisk is to be configured for Microsoft KMS or MAK volume license management.

● None

○ Key Management Service (KMS)

○ Multiple Activation Key (MAK)

Step 15. Select Image entire boot disk on the Configure Image Volumes page.

Step 16. Click Next.

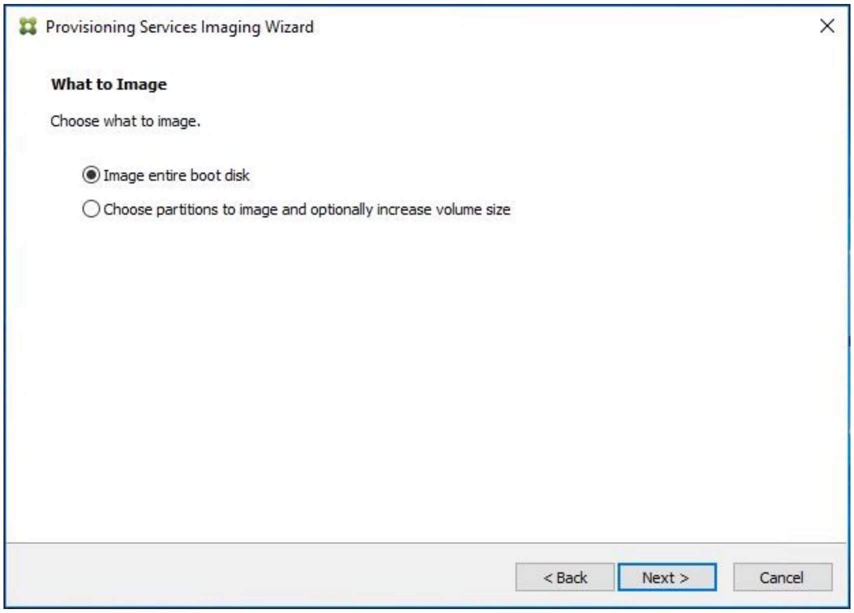

Provisioning Services Imaging Wizard

What to Image

Choose what to image.

● Image entire boot disk

○ Choose partitions to image and optionally increase volume size

Step 17. Select Optimize for hard disk again for Provisioning Services before imaging on the Optimize Hard Disk for Provisioning Services.

Step 18. Click Next.

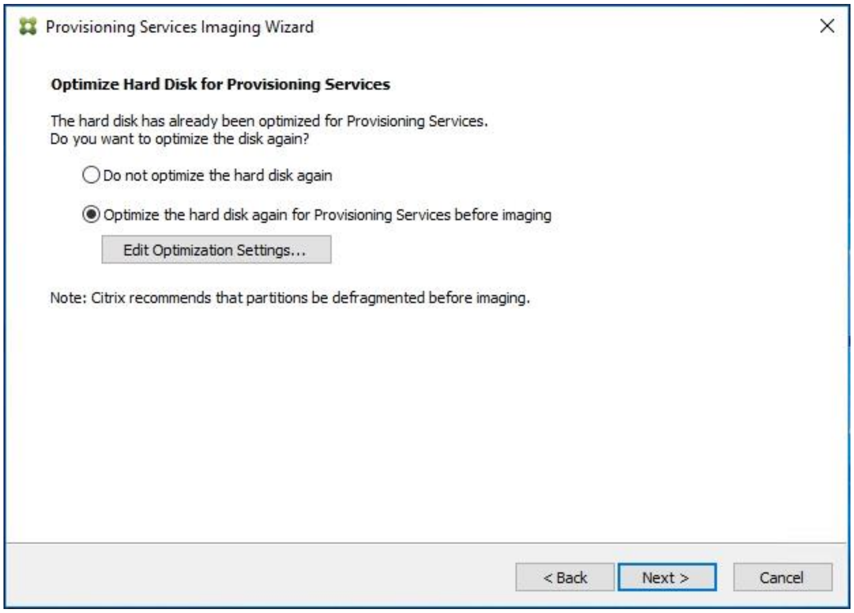
Provisioning Services Imaging Wizard

Optimize Hard Disk for Provisioning Services

The hard disk has already been optimized for Provisioning Services.
Do you want to optimize the disk again?

○ Do not optimize the hard disk again

● Optimize the hard disk again for Provisioning Services before imaging

Edit Optimization Settings...

Note: Citrix recommends that partitions be defragmented before imaging.

Step 19. Click Create on the Summary page.

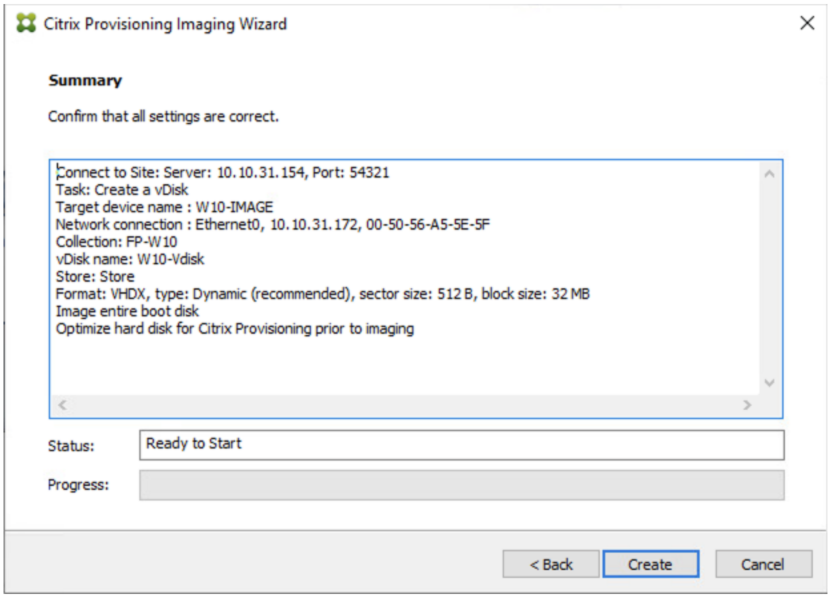

Summary

Confirm that all settings are correct.

```
Connect to Site: Server: 10.10.31.154, Port: 54321
Task: Create a vDisk
Target device name : W10-IMAGE
Network connection : Ethernet0, 10.10.31.172, 00-50-56-A5-5E-5F
Collection: FP-W10
vDisk name: W10-Vdisk
Store: Store
Format: VHDX, type: Dynamic (recommended), sector size: 512 B, block size: 32 MB
Image entire boot disk
Optimize hard disk for Citrix Provisioning prior to imaging
```

Status: Ready to Start

Progress:

Step 20. Review the configuration and click Continue.

Step 21. When prompted, click No to shut down the machine.

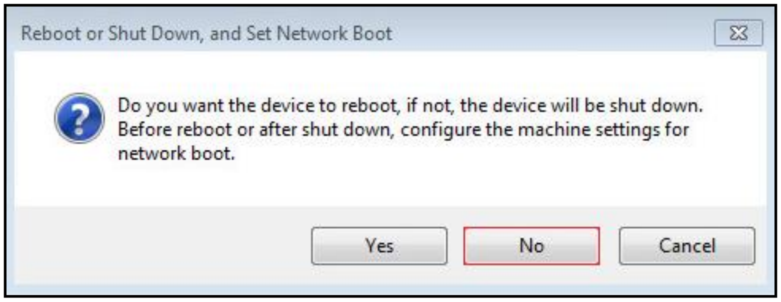

Do you want the device to reboot, if not, the device will be shut down. Before reboot or after shut down, configure the machine settings for network boot.

Step 22. Edit the VM settings and select Force EFI Setup under Boot Options.

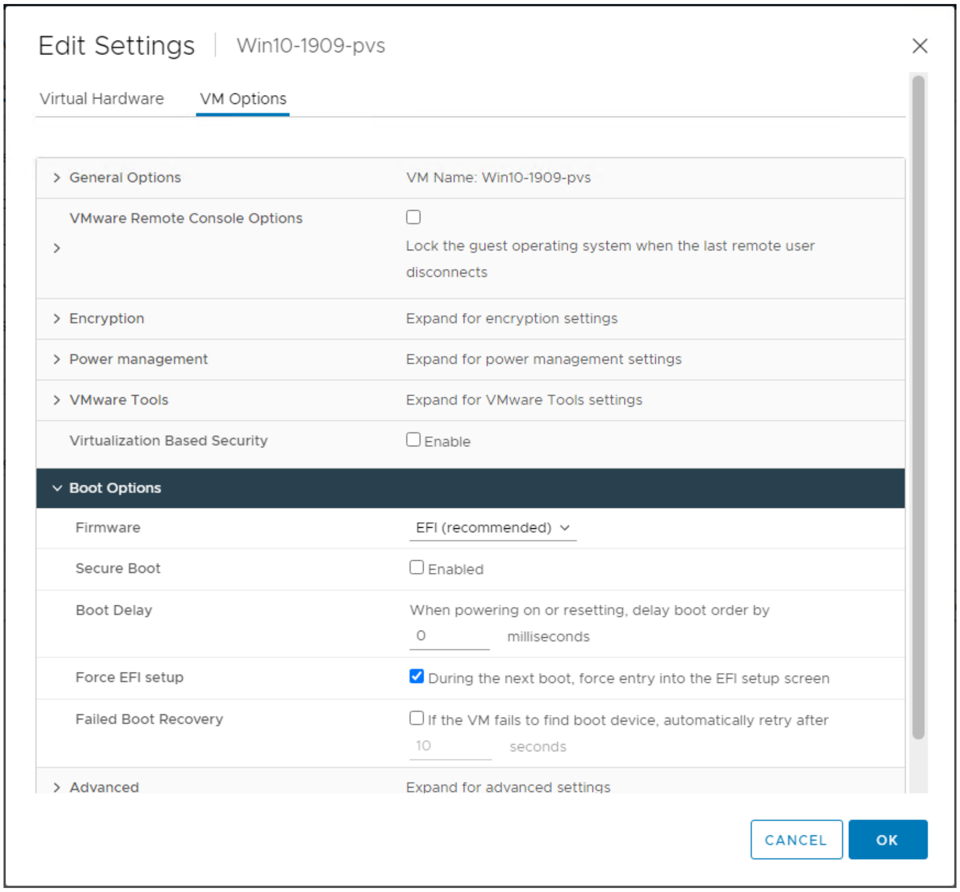

Edit Settings | Win10-1909-pvs

Virtual Hardware | VM Options

General Options	VM Name: Win10-1909-pvs
VMware Remote Console Options	□ Lock the guest operating system when the last remote user disconnects
Encryption	Expand for encryption settings
Power management	Expand for power management settings
VMware Tools	Expand for VMware Tools settings
Virtualization Based Security	□ Enable
Boot Options	
Firmware	EFI (recommended)
Secure Boot	□ Enabled
Boot Delay	When powering on or resetting, delay boot order by 0 milliseconds
Force EFI setup	☑ During the next boot, force entry into the EFI setup screen
Failed Boot Recovery	□ If the VM fails to find boot device, automatically retry after 10 seconds
Advanced	Expand for advanced settings

Step 23. Configure the VM settings for EFI network boot.

Step 24. Click Commit changes and exit.

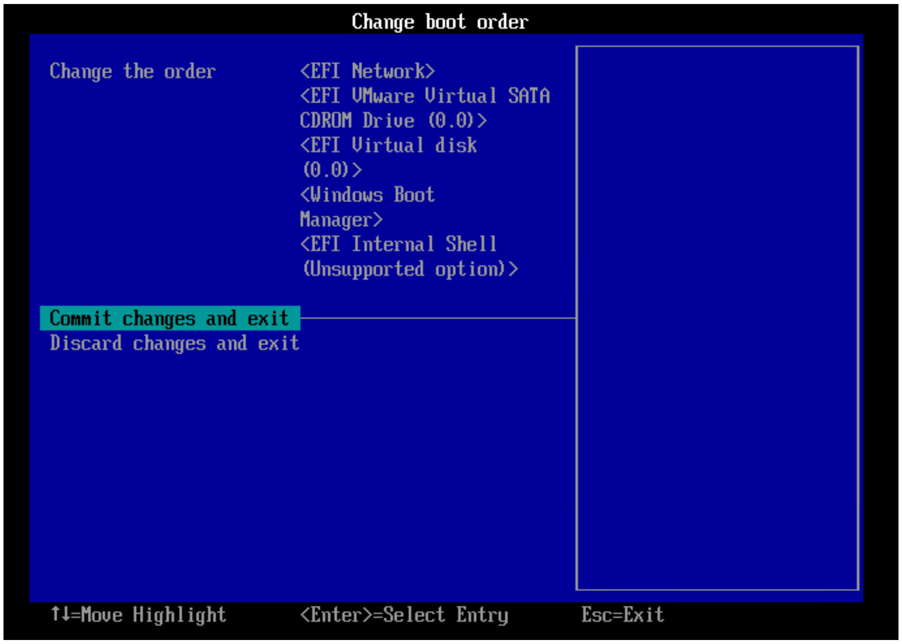

```
                         Change boot order

 Change the order          <EFI Network>
                           <EFI VMware Virtual SATA
                           CDROM Drive (0.0)>
                           <EFI Virtual disk
                           (0.0)>
                           <Windows Boot
                           Manager>
                           <EFI Internal Shell
                           (Unsupported option)>

 Commit changes and exit
 Discard changes and exit

 ↑↓=Move Highlight         <Enter>=Select Entry       Esc=Exit
```

Step 25. After restarting the virtual machine, log into the master target. The PVS imaging process begins, copying the contents of the C: drive to the PVS vDisk located on the server.

Note: If prompted to Format disk, disregard the message, and allow Provisioning Imaging Wizard to finish.

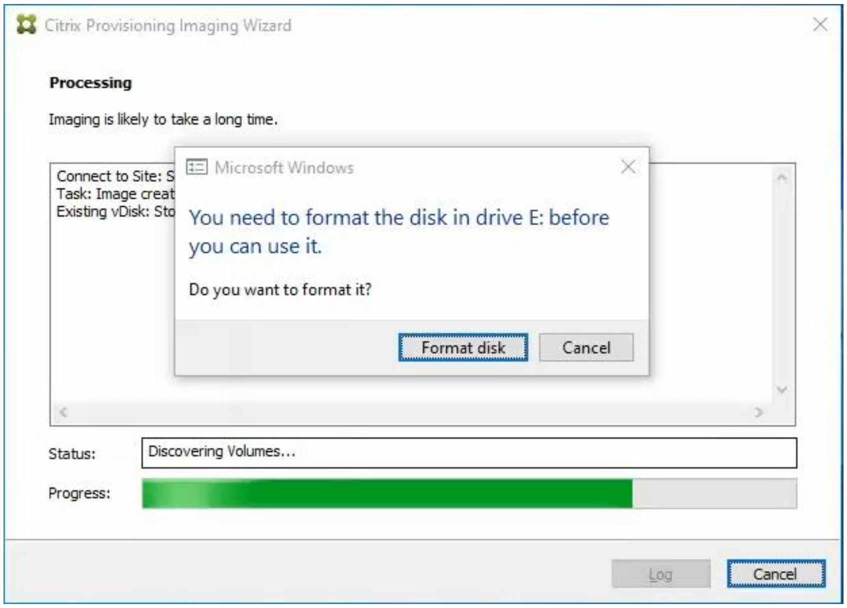

Processing

Imaging is likely to take a long time.

```
Connect to Site: S
Task: Image creat
Existing vDisk: Sto
```

You need to format the disk in drive E: before you can use it.

Do you want to format it?

Status: Discovering Volumes...

Progress:

Step 26. A message is displayed when the conversion is complete, click Done.

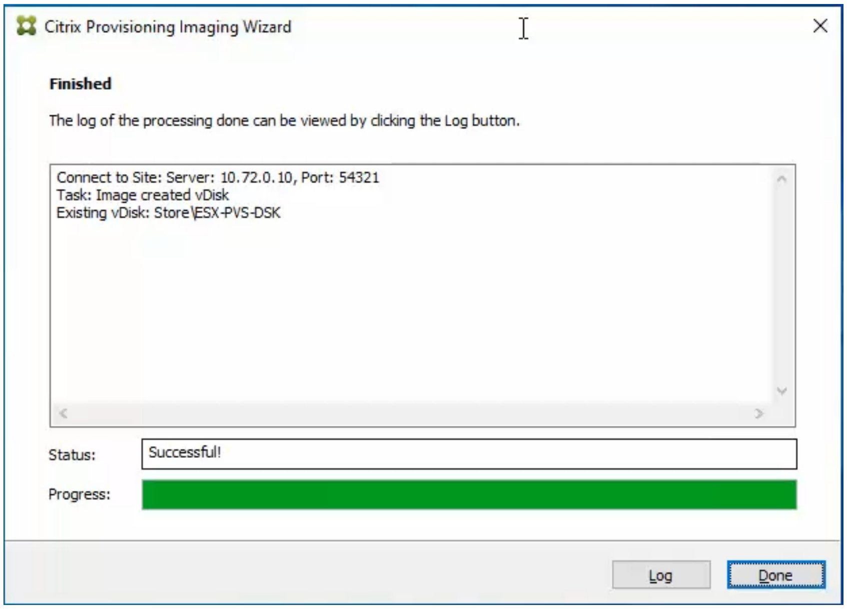

Finished

The log of the processing done can be viewed by clicking the Log button.

```
Connect to Site: Server: 10.72.0.10, Port: 54321
Task: Image created vDisk
Existing vDisk: Store\ESX-PVS-DSK
```

Status: Successful!

Progress:

Step 27. Shutdown the virtual machine used as the VDI or RDS master target.

Step 28. Connect to the PVS server and validate that the vDisk image is available in the Store.

Step 29. Right-click the newly created vDisk and select Properties.

Step 30. On the vDisk Properties dialog, change Access mode to “Standard Image (multi-device, read-only access).”

Step 31. Set the Cache Type to “Cache in device RAM with overflow on hard disk.”

Step 32. Set Maximum RAM size (MBs): 256.

Step 33. Click OK.

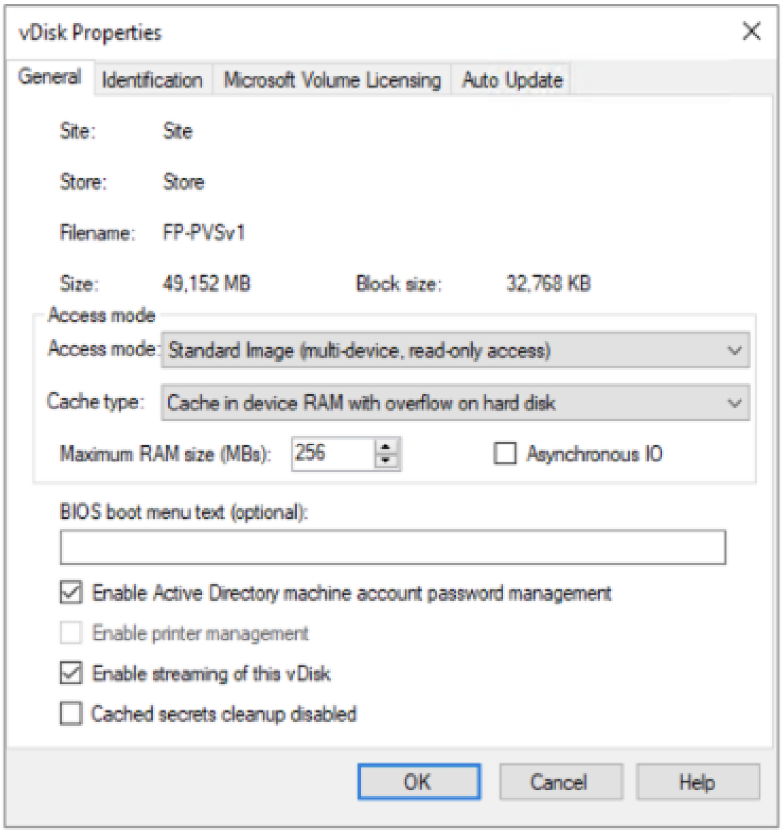

vDisk Properties

Site: Site

Store: Store

Filename: FP-PVSv1

Size: 49,152 MB Block size: 32,768 KB

Access mode

Access mode: Standard Image (multi-device, read-only access)

Cache type: Cache in device RAM with overflow on hard disk

Maximum RAM size (MBs): 256 □ Asynchronous IO

BIOS boot menu text (optional):

☑ Enable Active Directory machine account password management

□ Enable printer management

☑ Enable streaming of this vDisk

□ Cached secrets cleanup disabled

Provision Virtual Desktop Machines

Citrix Provisioning Services Citrix Virtual Desktop Setup Wizard

Procedure 1. Create PVS Streamed Virtual Desktop Machines

Step 1. Create a Master Target Virtual Machine:

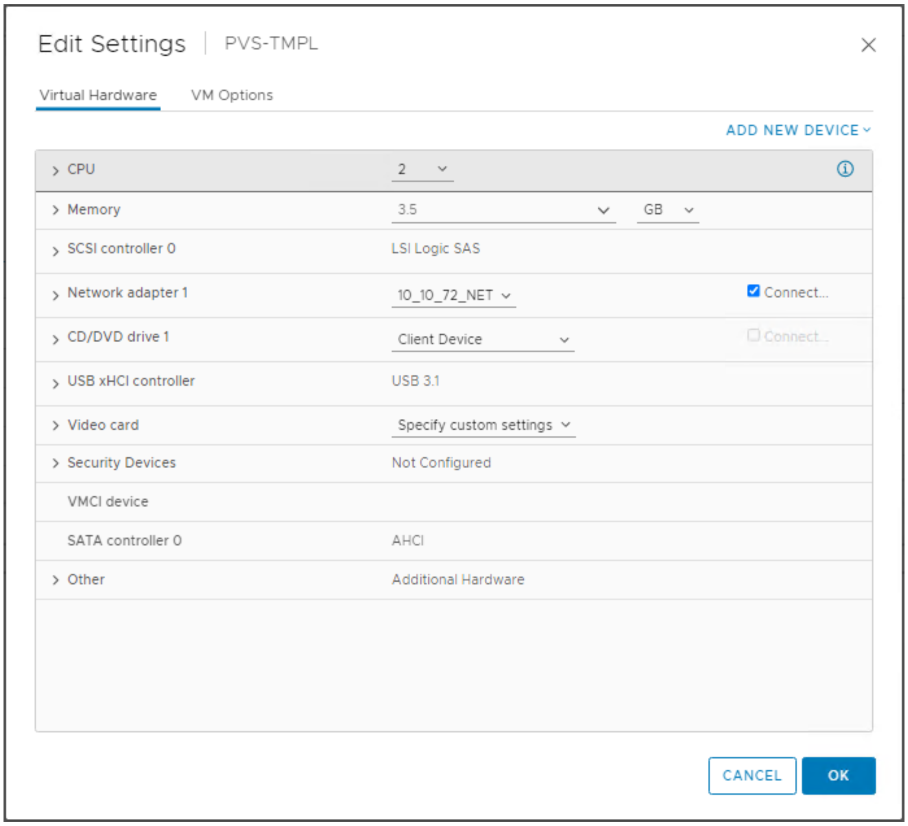

Edit Settings | PVS-TMPL

Virtual Hardware | VM Options

ADD NEW DEVICE

CPU	2	
Memory	3.5 GB	
SCSI controller 0	LSI Logic SAS	
Network adapter 1	10_10_72_NET	☑ Connect...
CD/DVD drive 1	Client Device	□ Connect...
USB xHCI controller	USB 3.1	
Video card	Specify custom settings	
Security Devices	Not Configured	
VMCI device		
SATA controller 0	AHCI	
Other	Additional Hardware	

CANCEL OK

Step 2. Right-click and clone the Master Target VM to the Template.

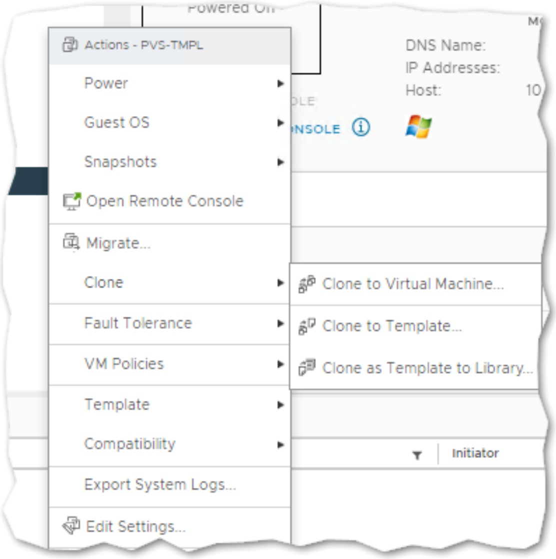

Step 3. Start the Citrix Virtual Apps and Desktops Setup Wizard from the Provisioning Services Console.

Step 4. Right-click the Site.

Step 5. Select Citrix Virtual Desktop Setup Wizard... from the context menu.

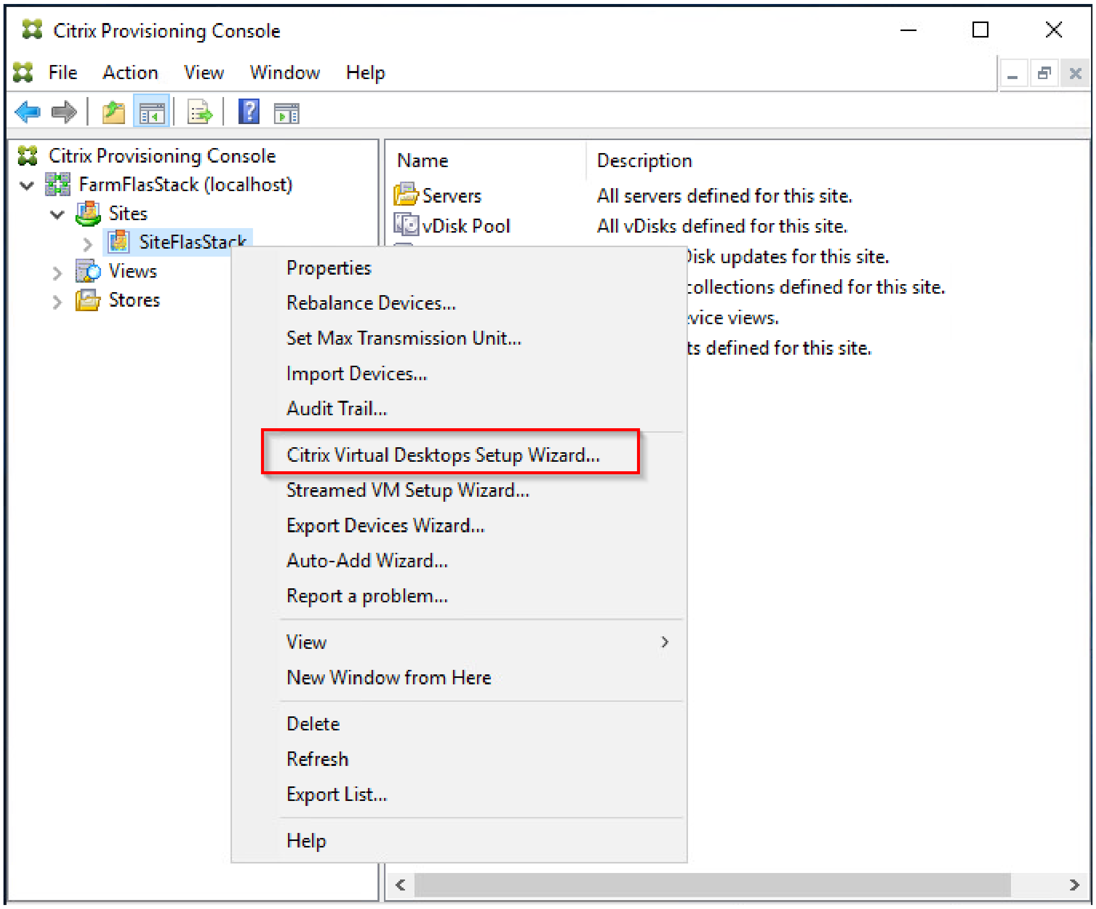
Citrix Provisioning Console
File Action View Window Help
Citrix Provisioning Console
FarmFlasStack (localhost)
Sites
SiteFlasStack
Views
Stores
Name
Description
Servers
All servers defined for this site.
vDisk Pool
All vDisks defined for this site.
Disk updates for this site.
collections defined for this site.
vice views.
ts defined for this site.
Properties
Rebalance Devices...
Set Max Transmission Unit...
Import Devices...
Audit Trail...
Citrix Virtual Desktops Setup Wizard...
Streamed VM Setup Wizard...
Export Devices Wizard...
Auto-Add Wizard...
Report a problem...
View
New Window from Here
Delete
Refresh
Export List...
Help

Step 6. Click Next.

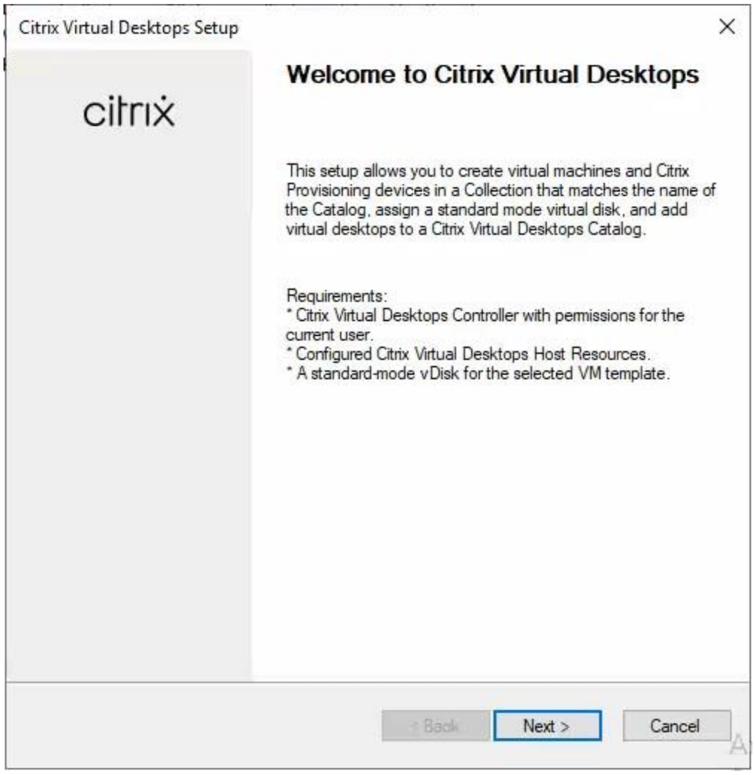

Welcome to Citrix Virtual Desktops

This setup allows you to create virtual machines and Citrix Provisioning devices in a Collection that matches the name of the Catalog, assign a standard mode virtual disk, and add virtual desktops to a Citrix Virtual Desktops Catalog.

Requirements:
* Citrix Virtual Desktops Controller with permissions for the current user.
* Configured Citrix Virtual Desktops Host Resources.
* A standard-mode vDisk for the selected VM template.

Step 7. Enter the address of the Citrix Virtual Desktop Controller that will be used for the wizard operations.

Step 8. Click Next.

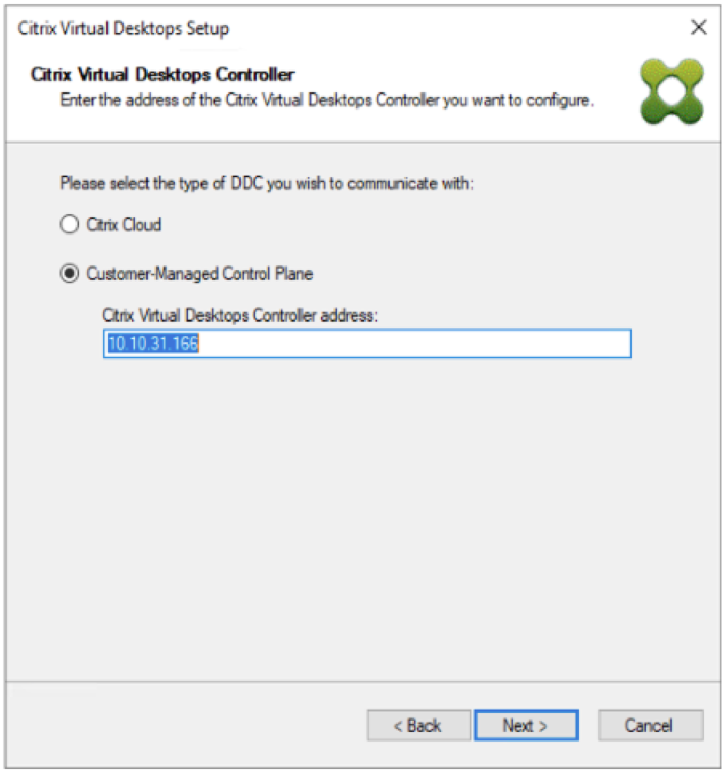

Citrix Virtual Desktops Setup

Citrix Virtual Desktops Controller

Enter the address of the Citrix Virtual Desktops Controller you want to configure.

Please select the type of DDC you wish to communicate with:

○ Citrix Cloud

● Customer-Managed Control Plane

Citrix Virtual Desktops Controller address: 10.10.31.166

Step 9. Select Host Resources that will be used for the wizard operations.

Step 10. Click Next.

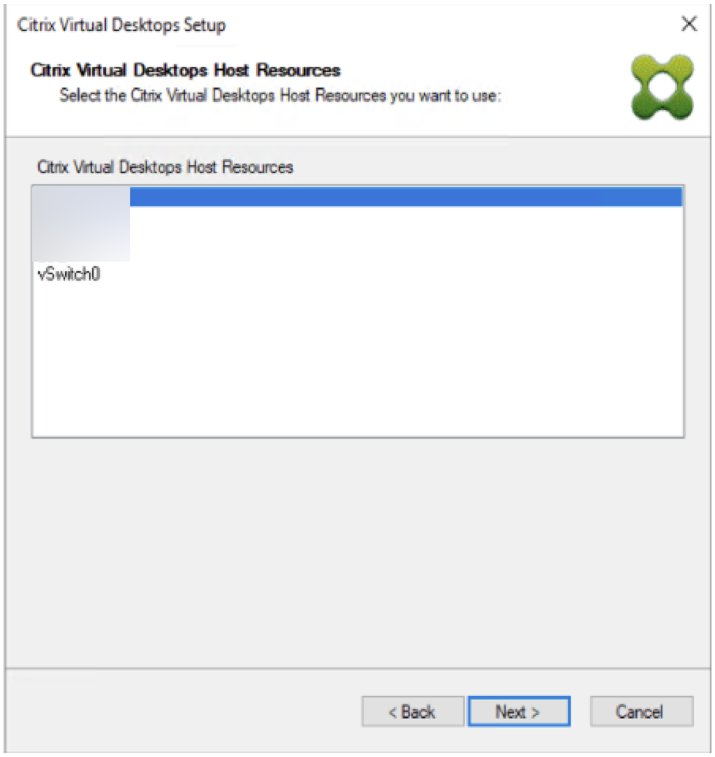

Citrix Virtual Desktops Setup

Citrix Virtual Desktops Host Resources

Select the Citrix Virtual Desktops Host Resources you want to use:

Citrix Virtual Desktops Host Resources

vSwitch0

Step 11. Provide Citrix Virtual Desktop Controller credentials.

Step 12. Click OK.

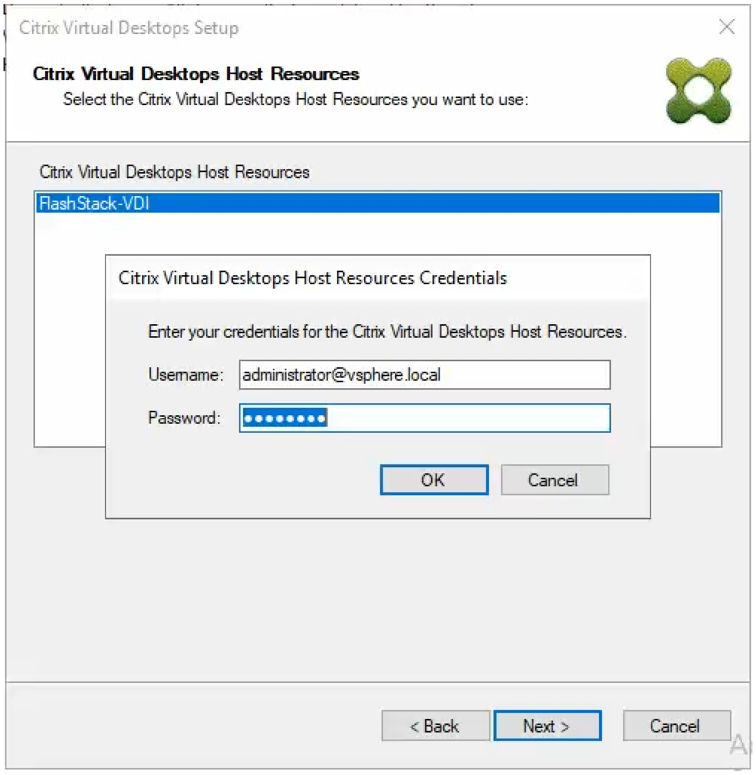

Citrix Virtual Desktops Setup

Citrix Virtual Desktops Host Resources

Select the Citrix Virtual Desktops Host Resources you want to use:

Citrix Virtual Desktops Host Resources

FlashStack-VDI

Citrix Virtual Desktops Host Resources Credentials

Enter your credentials for the Citrix Virtual Desktops Host Resources.

Username: administrator@vsphere.local

Password: ••••••••

Step 13. Select the Template created earlier.

Step 14. Click Next.

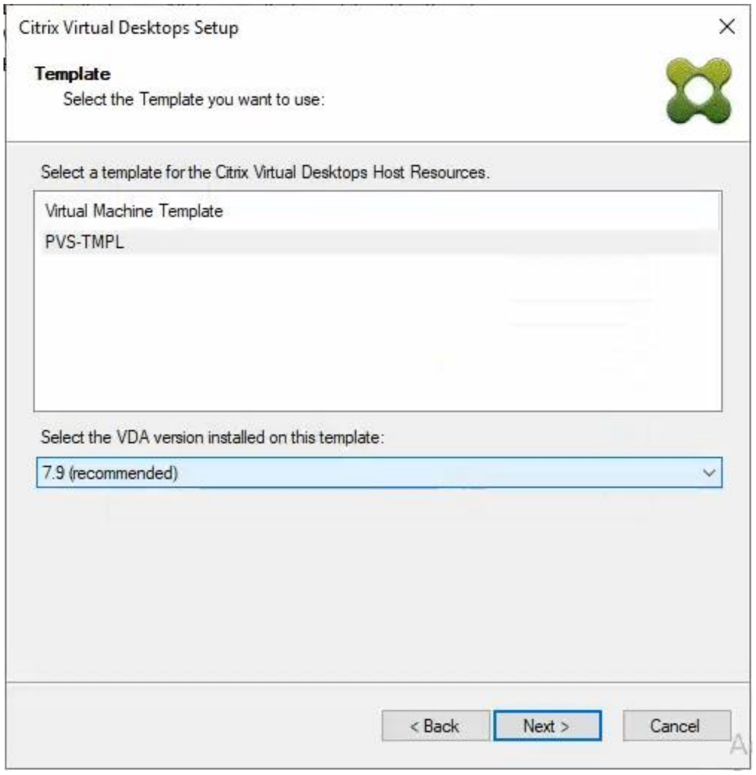

Citrix Virtual Desktops Setup

Template

Select the Template you want to use:

Select a template for the Citrix Virtual Desktops Host Resources.

Virtual Machine Template

PVS-TMPL

Select the VDA version installed on this template:

7.9 (recommended)

Step 15. Select the virtual disk (vDisk) that will be used to stream the provisioned virtual machines.

Step 16. Click Next.

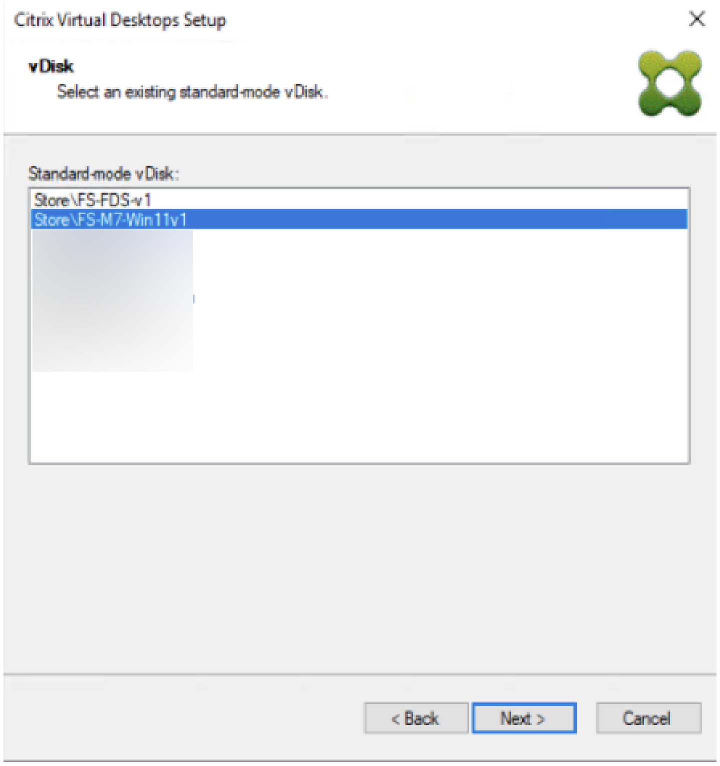

Citrix Virtual Desktops Setup

vDisk

Select an existing standard-mode vDisk.

Standard-mode vDisk:

- Store\FS-FDS-v1
- Store\FS-M7-Win11v1

Step 17. Select Create new catalog.

Step 18. Provide a catalog name.

Step 19. Click Next.

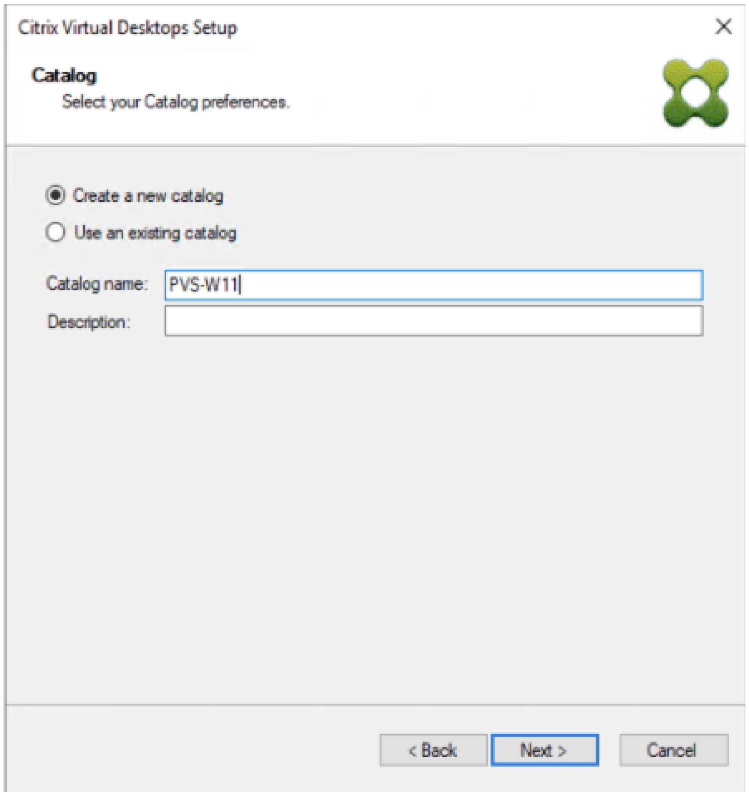

Citrix Virtual Desktops Setup

Catalog

Select your Catalog preferences.

● Create a new catalog

○ Use an existing catalog

Catalog name: PVS-W11

Description:

Step 20. Select Single-session OS for Machine catalog Operating System.

Step 21. Click Next.

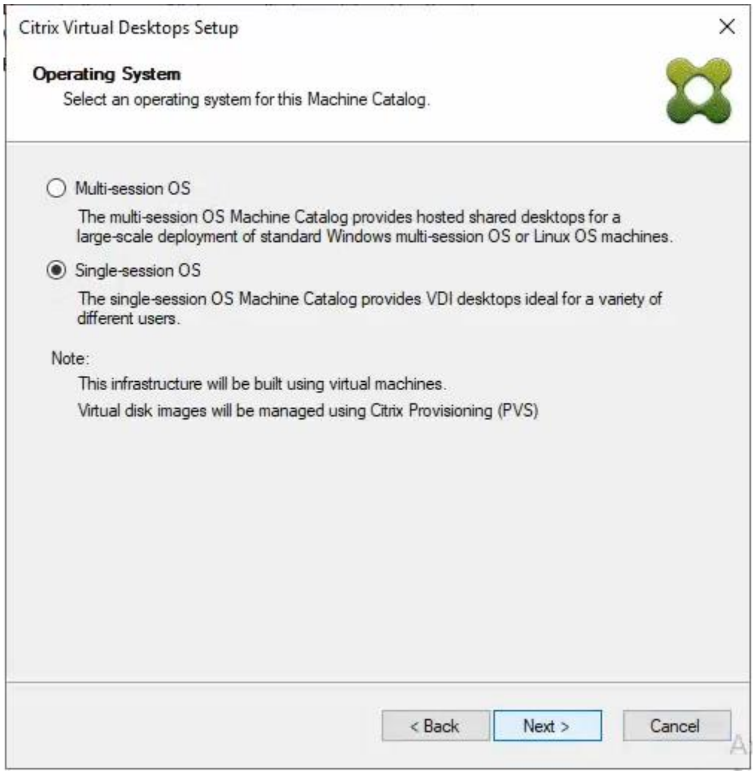

Citrix Virtual Desktops Setup

Operating System

Select an operating system for this Machine Catalog.

○ Multi-session OS

The multi-session OS Machine Catalog provides hosted shared desktops for a large-scale deployment of standard Windows multi-session OS or Linux OS machines.

● Single-session OS

The single-session OS Machine Catalog provides VDI desktops ideal for a variety of different users.

Note:

This infrastructure will be built using virtual machines.

Virtual disk images will be managed using Citrix Provisioning (PVS)

Step 22. Select random for the User Experience.

Step 23. Click Next.

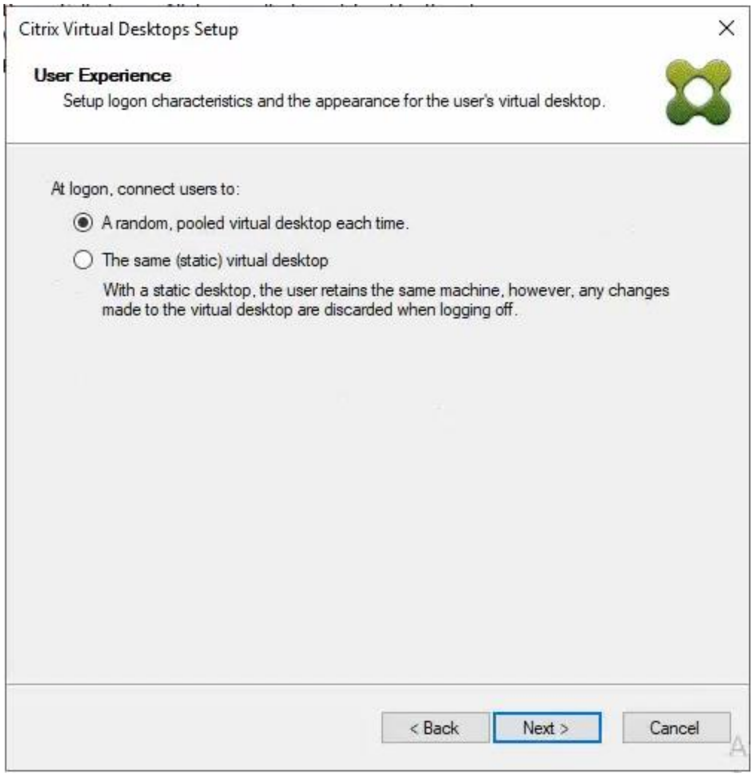

Citrix Virtual Desktops Setup

User Experience

Setup logon characteristics and the appearance for the user's virtual desktop.

At logon, connect users to:

● A random, pooled virtual desktop each time.

○ The same (static) virtual desktop

With a static desktop, the user retains the same machine, however, any changes made to the virtual desktop are discarded when logging off.

Step 24. On the Virtual machines dialog, specify the following:

- The number of virtual machines to create.
- 2 for Number of vCPUs for the virtual machine
- 4 GB for the amount of memory for the virtual machine
- 6GB for the Local write cache disk.

Step 25. Click Next.

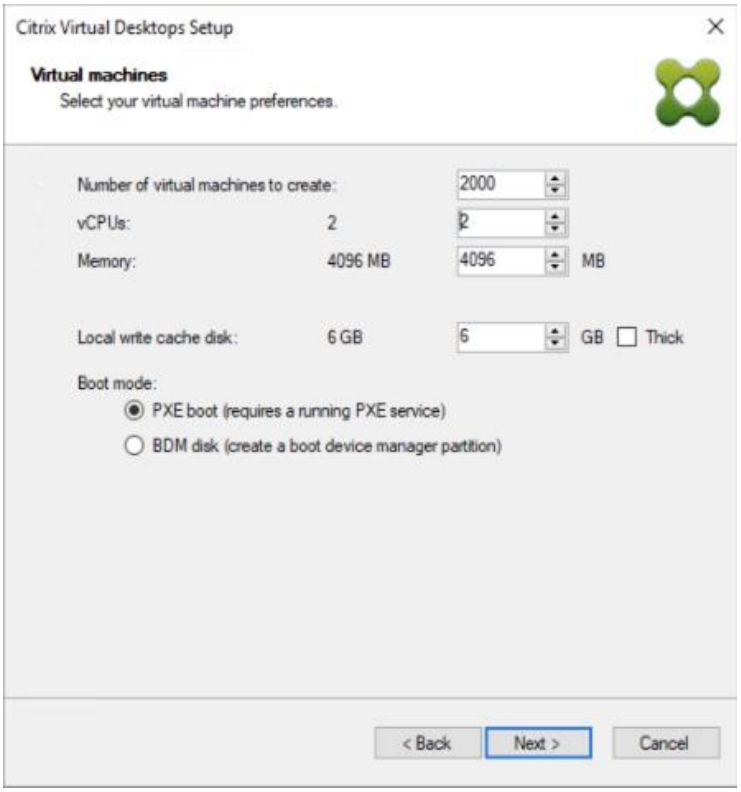

Citrix Virtual Desktops Setup
Virtual machines
Select your virtual machine preferences.
< Back
Next >
Cancel

Number of virtual machines to create: 2000

vCPUs: 2 — 2

Memory: 4096 MB — 4096 MB

Local write cache disk: 6 GB — 6 GB □ Thick

Boot mode:

● PXE boot (requires a running PXE service)

○ BDM disk (create a boot device manager partition)

Step 26. Select the Create new accounts.

Step 27. Click Next.

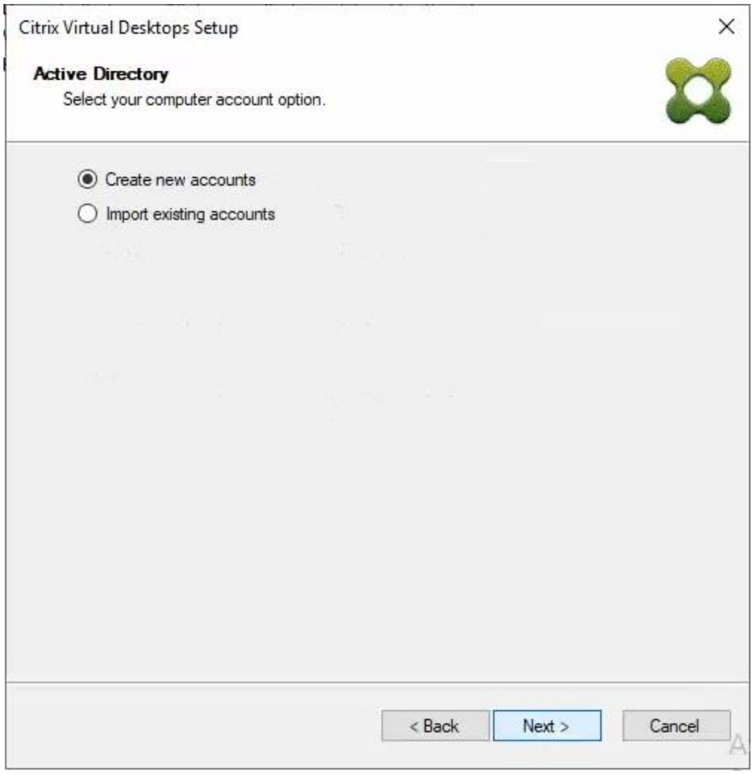

Citrix Virtual Desktops Setup
Active Directory
Select your computer account option.
< Back
Next >
Cancel

● Create new accounts

○ Import existing accounts

Step 28. Specify the Active Directory Accounts and Location. This is where the wizard should create computer accounts.

Step 29. Provide the Account naming scheme. An example name is shown in the text box below the naming scheme selection location.

Step 30. Click Next.

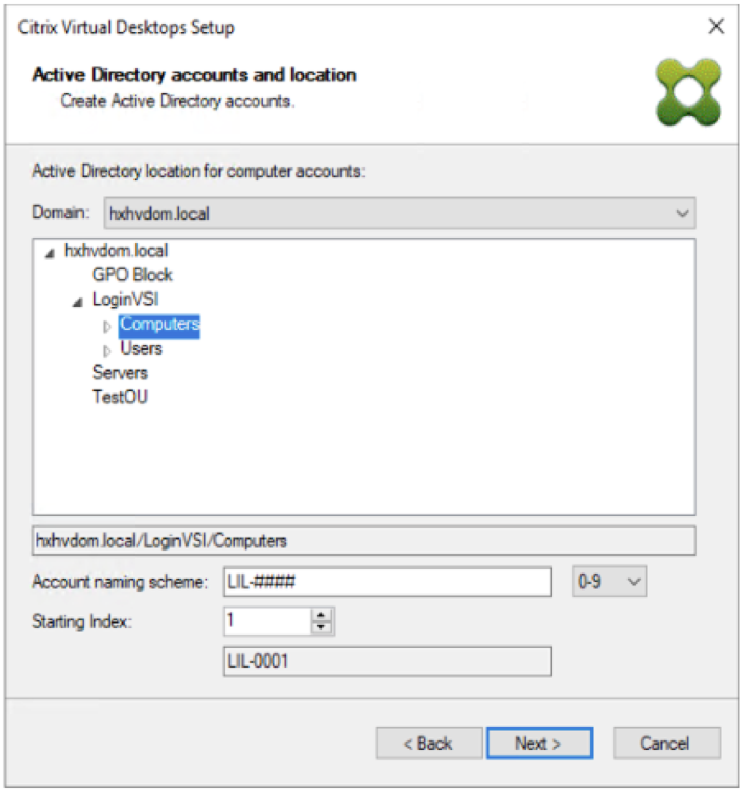

Citrix Virtual Desktops Setup

Active Directory accounts and location

Create Active Directory accounts.

Active Directory location for computer accounts:

Domain: hxhvdom.local

- hxhvdom.local
  - GPO Block
  - LoginVSI
    - Computers
    - Users
  - Servers
  - TestOU

hxhvdom.local/LoginVSI/Computers

Account naming scheme: LIL-#### 0-9

Starting Index: 1

LIL-0001

Step 31. Verify the information on the Summary screen.

Step 32. Click Finish to begin the virtual machine creation.

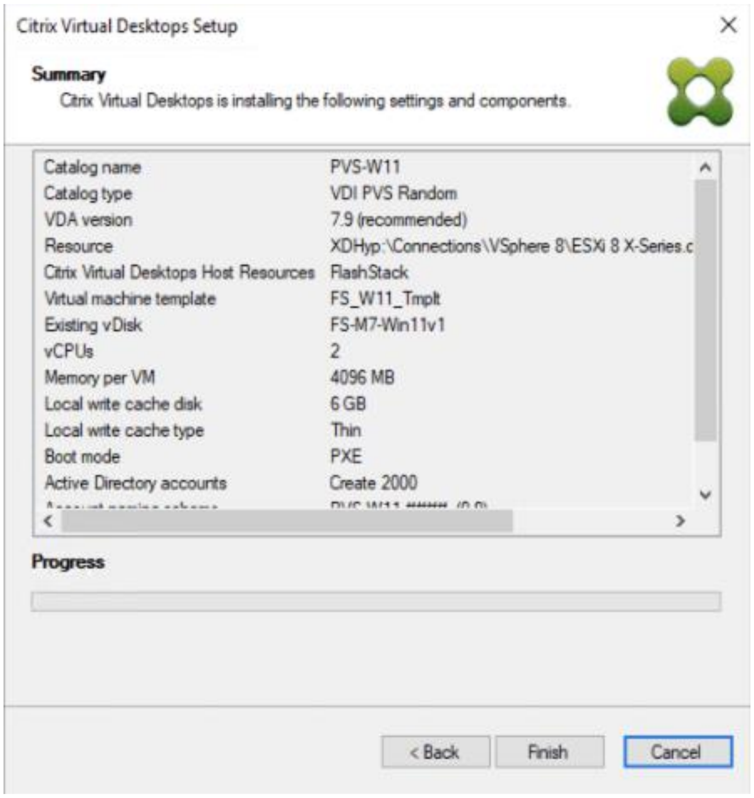

Citrix Virtual Desktops is installing the following settings and components.

Catalog name	PVS-W11
Catalog type	VDI PVS Random
VDA version	7.9 (recommended)
Resource	XDHyp:\Connections\VSphere 8\ESXi 8 X-Series.c
Citrix Virtual Desktops Host Resources	FlashStack
Virtual machine template	FS_W11_Tmplt
Existing vDisk	FS-M7-Win11v1
vCPUs	2
Memory per VM	4096 MB
Local write cache disk	6 GB
Local write cache type	Thin
Boot mode	PXE
Active Directory accounts	Create 2000
Account naming scheme	PVS-W11-###### (0-9)

Step 33. When the wizard is done provisioning the virtual machines, click Done.

Step 34. When the wizard is done provisioning the virtual machines, verify the Machine Catalog on the Citrix Virtual Apps and Desktops Controller:

- Connect to a Citrix Virtual Apps and Desktops server and launch Citrix Studio.
- Select Machine Catalogs in the Studio navigation pane.
- Select a machine catalog.

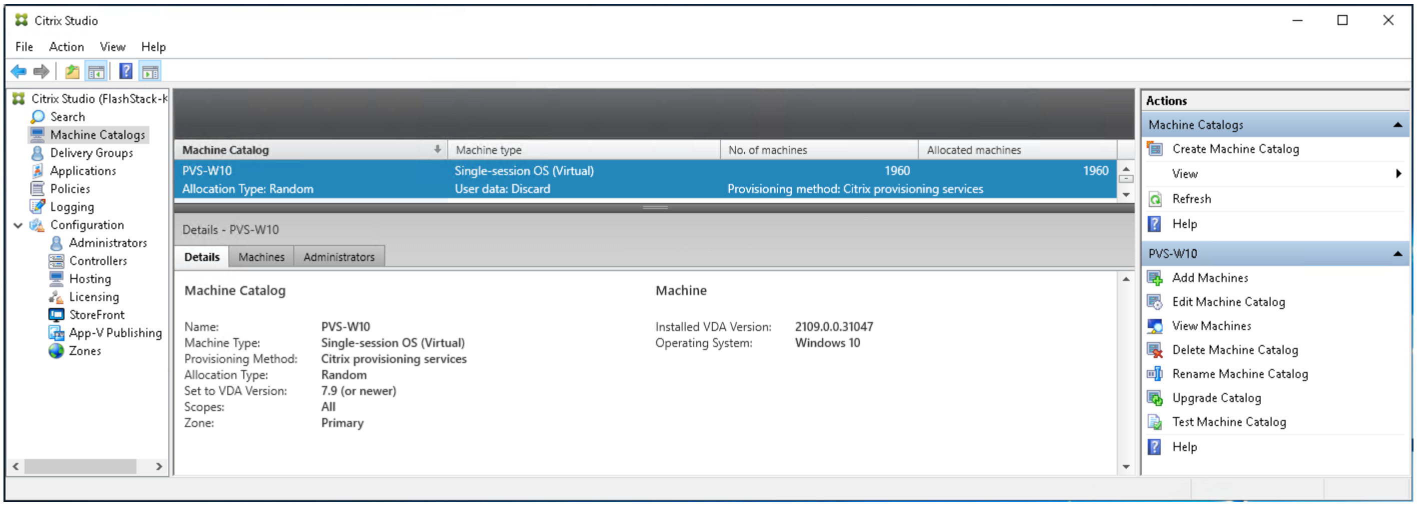

Machine Catalog	Machine type	No. of machines	Allocated machines
PVS-W10 Allocation Type: Random	Single-session OS (Virtual) User data: Discard	1960 Provisioning method: Citrix provisioning services	1960

Details - PVS-W10

Machine Catalog

Name:	PVS-W10
Machine Type:	Single-session OS (Virtual)
Provisioning Method:	Citrix provisioning services
Allocation Type:	Random
Set to VDA Version:	7.9 (or newer)
Scopes:	All
Zone:	Primary

Machine

Installed VDA Version:	2109.0.0.31047
Operating System:	Windows 10

Procedure 2. Citrix Machine Creation Services

Step 35. Connect to a Citrix Virtual Apps and Desktops server and launch Citrix Studio.

Step 36. Choose Create Machine Catalog from the Actions pane.

Step 37. Click Next.

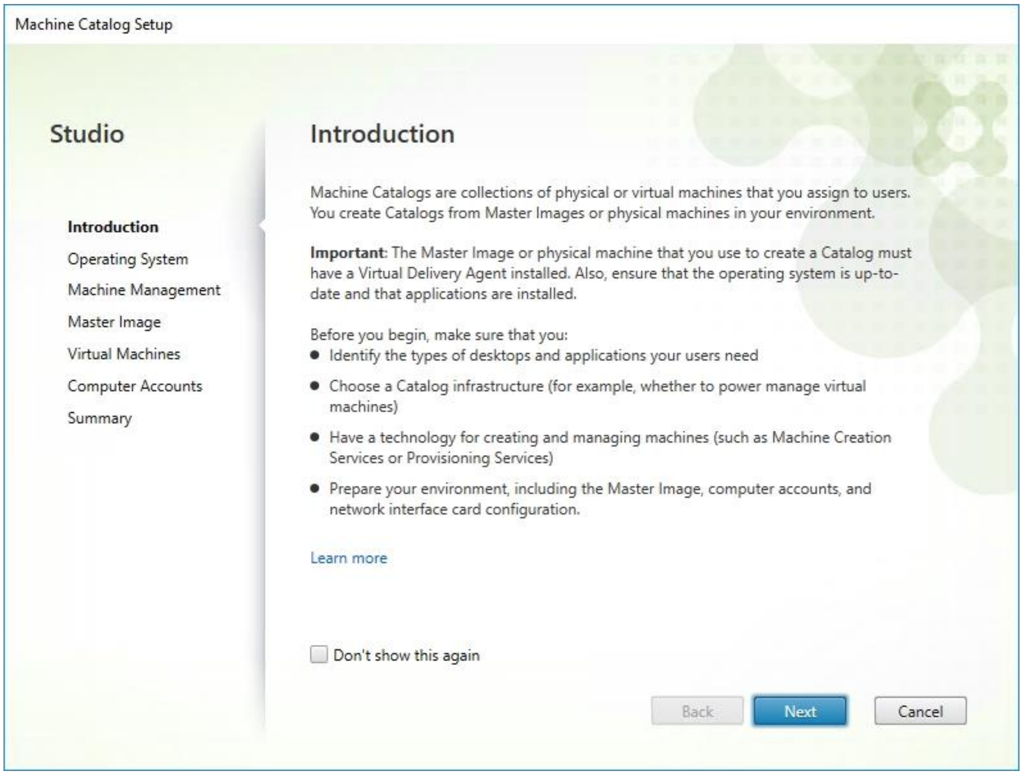

Machine Catalog Setup

Introduction

Machine Catalogs are collections of physical or virtual machines that you assign to users. You create Catalogs from Master Images or physical machines in your environment.

Important: The Master Image or physical machine that you use to create a Catalog must have a Virtual Delivery Agent installed. Also, ensure that the operating system is up-to-date and that applications are installed.

Before you begin, make sure that you:

- Identify the types of desktops and applications your users need
- Choose a Catalog infrastructure (for example, whether to power manage virtual machines)
- Have a technology for creating and managing machines (such as Machine Creation Services or Provisioning Services)
- Prepare your environment, including the Master Image, computer accounts, and network interface card configuration.

Learn more

□ Don't show this again

Step 38. Select Single-session OS.

Step 39. Click Next.

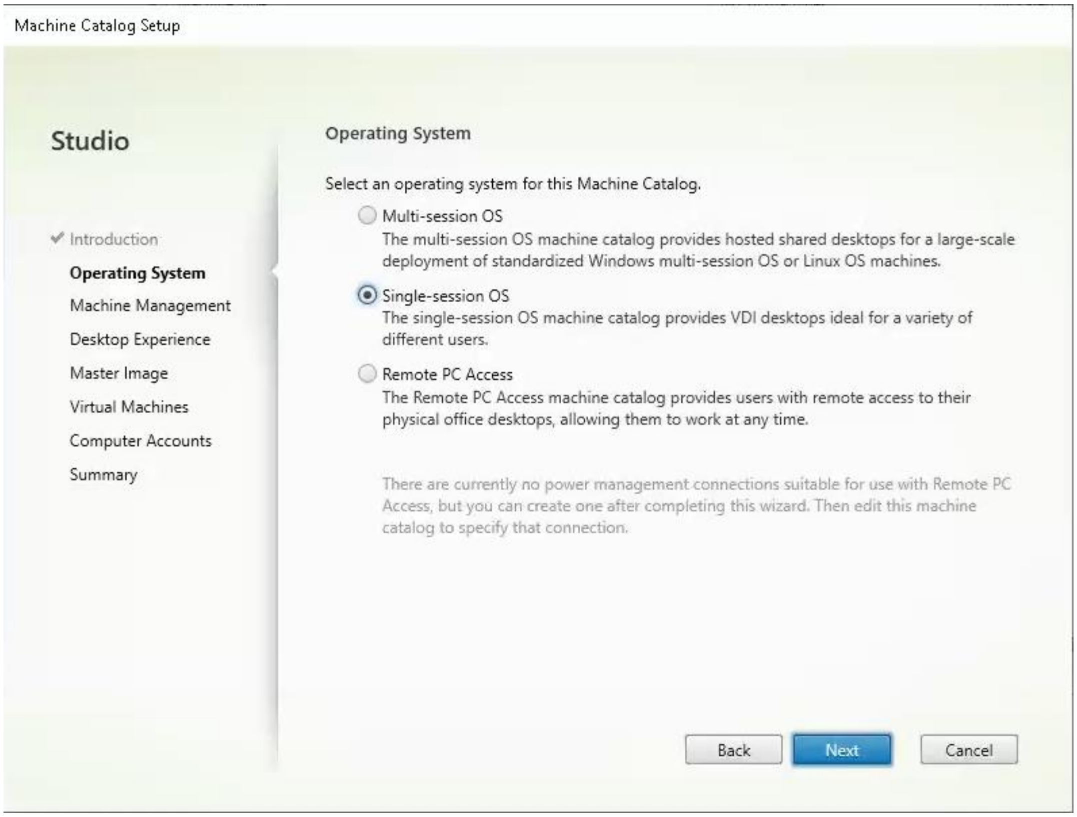

Machine Catalog Setup

Operating System

Select an operating system for this Machine Catalog.

○ Multi-session OS
The multi-session OS machine catalog provides hosted shared desktops for a large-scale deployment of standardized Windows multi-session OS or Linux OS machines.

● Single-session OS
The single-session OS machine catalog provides VDI desktops ideal for a variety of different users.

○ Remote PC Access
The Remote PC Access machine catalog provides users with remote access to their physical office desktops, allowing them to work at any time.

There are currently no power management connections suitable for use with Remote PC Access, but you can create one after completing this wizard. Then edit this machine catalog to specify that connection.

Step 40. Select Multi-session OS when using Windows Server 2022 desktops.

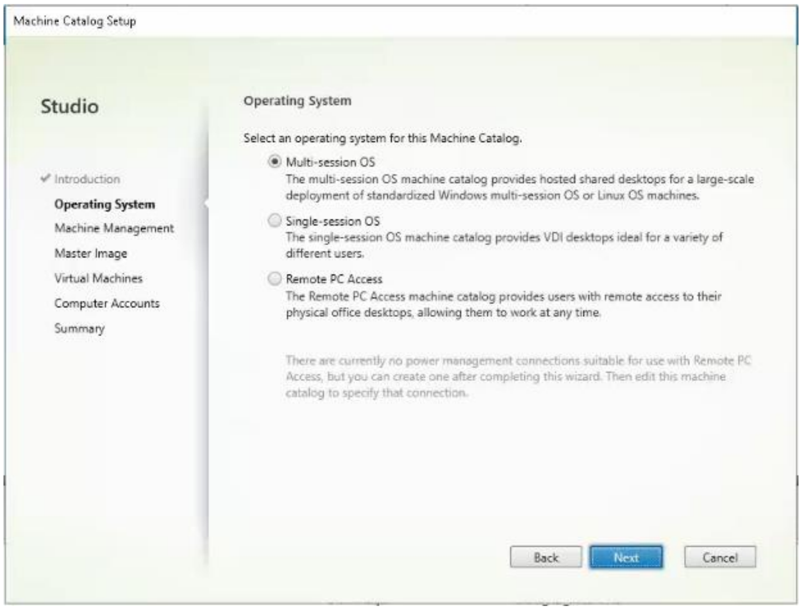

Operating System

Select an operating system for this Machine Catalog.

● Multi-session OS
The multi-session OS machine catalog provides hosted shared desktops for a large-scale deployment of standardized Windows multi-session OS or Linux OS machines.

○ Single-session OS
The single-session OS machine catalog provides VDI desktops ideal for a variety of different users.

○ Remote PC Access
The Remote PC Access machine catalog provides users with remote access to their physical office desktops, allowing them to work at any time.

There are currently no power management connections suitable for use with Remote PC Access, but you can create one after completing this wizard. Then edit this machine catalog to specify that connection.

Step 41. Select the appropriate machine management.

Step 42. Click Next.

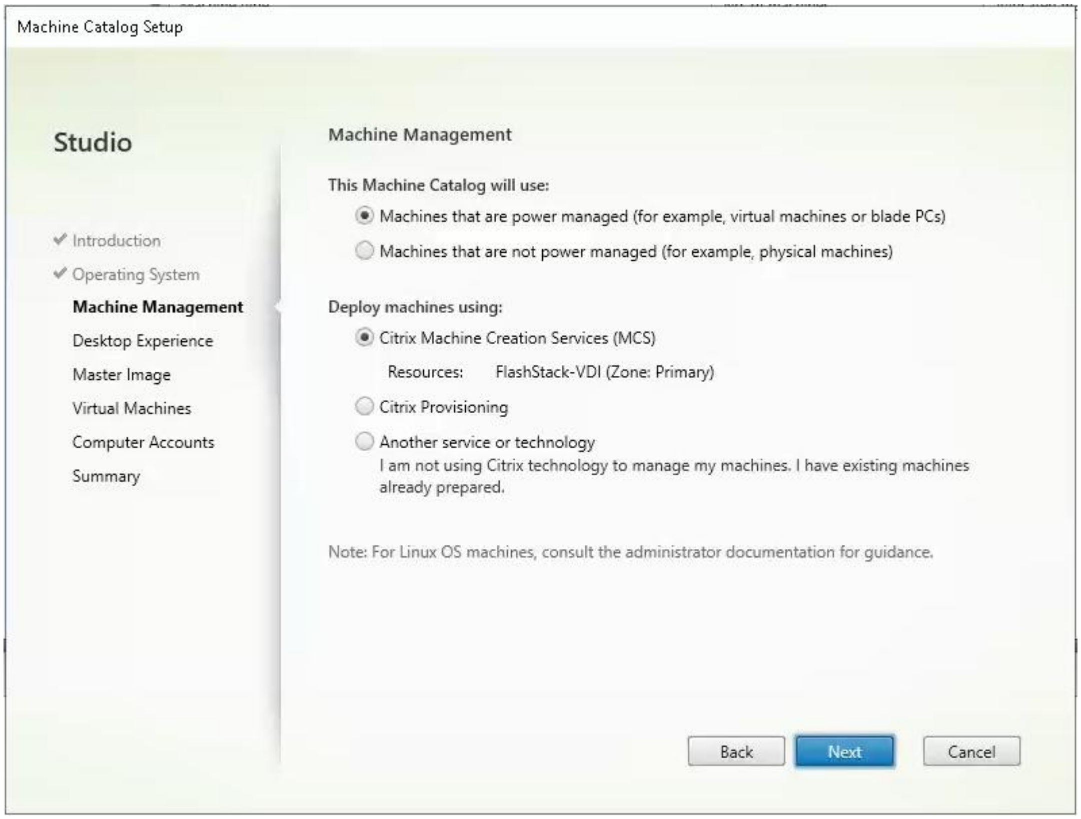

Machine Management

This Machine Catalog will use:

● Machines that are power managed (for example, virtual machines or blade PCs)

○ Machines that are not power managed (for example, physical machines)

Deploy machines using:

● Citrix Machine Creation Services (MCS)

Resources: FlashStack-VDI (Zone: Primary)

○ Citrix Provisioning

○ Another service or technology
I am not using Citrix technology to manage my machines. I have existing machines already prepared.

Note: For Linux OS machines, consult the administrator documentation for guidance.

Step 43. Select (static) for Desktop Experience.

Step 44. Click Next.

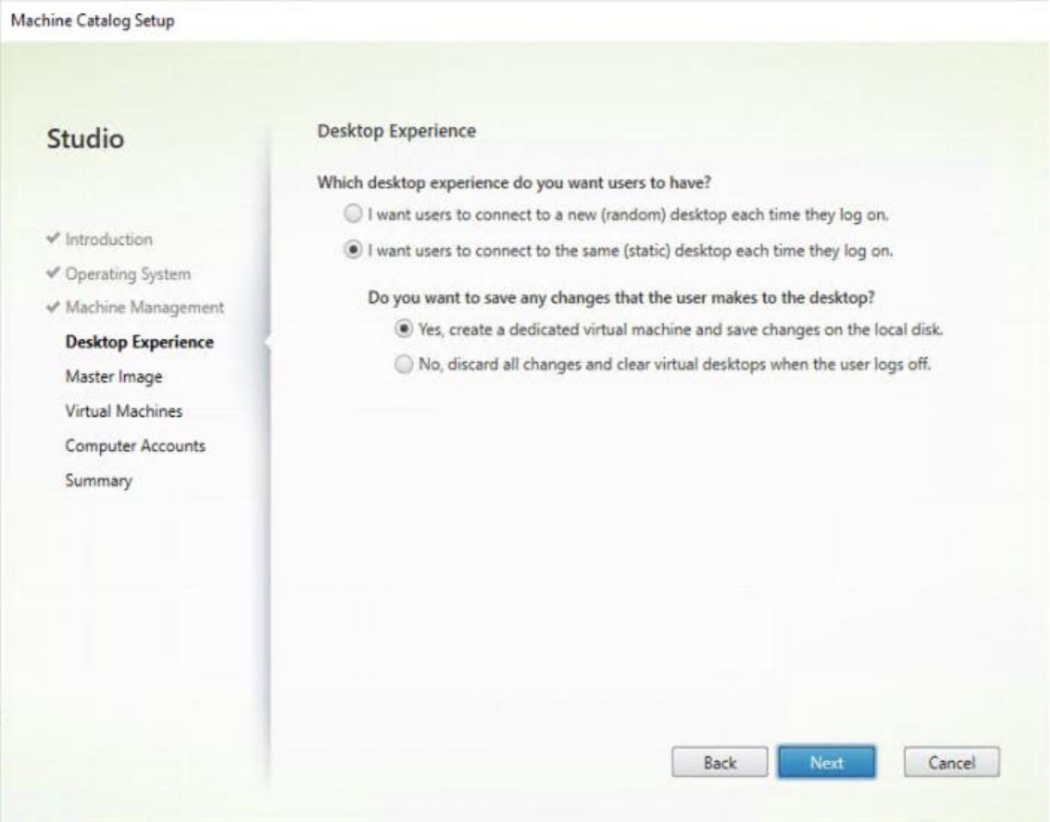

Machine Catalog Setup

Desktop Experience

Which desktop experience do you want users to have?

○ I want users to connect to a new (random) desktop each time they log on.

● I want users to connect to the same (static) desktop each time they log on.

Do you want to save any changes that the user makes to the desktop?

● Yes, create a dedicated virtual machine and save changes on the local disk.

○ No, discard all changes and clear virtual desktops when the user logs off.

Step 45. Select a Virtual Machine to be used for Catalog Master Image.

Step 46. Click Next.

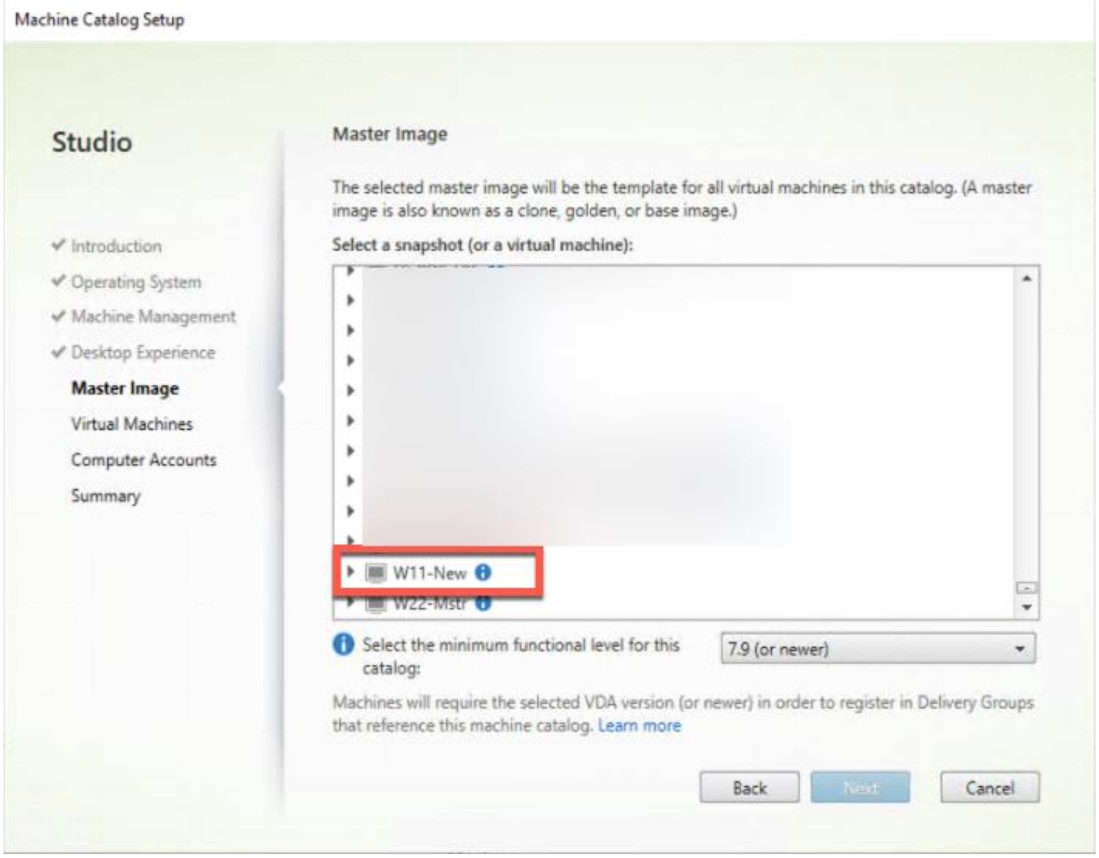

Machine Catalog Setup

Master Image

The selected master image will be the template for all virtual machines in this catalog. (A master image is also known as a clone, golden, or base image.)

Select a snapshot (or a virtual machine):

- W11-New
- W22-Mstr

Select the minimum functional level for this catalog: 7.9 (or newer)

Machines will require the selected VDA version (or newer) in order to register in Delivery Groups that reference this machine catalog. Learn more

Step 47. Specify the number of desktops to create and machine configuration.

Step 48. Set amount of memory (MB) to be used by virtual desktops.

Step 49. Select Full Copy for machine copy mode.

Step 50. Click Next.

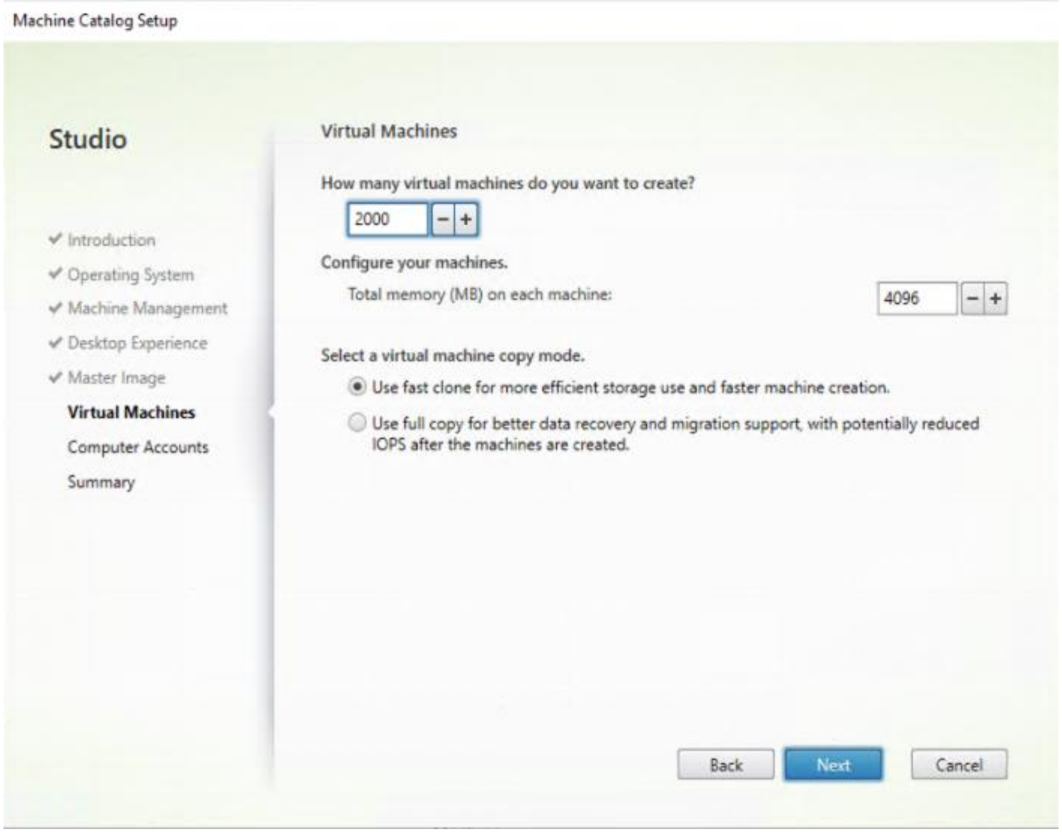

Machine Catalog Setup

Studio

✓ Introduction
✓ Operating System
✓ Machine Management
✓ Desktop Experience
✓ Master Image
Virtual Machines
Computer Accounts
Summary

Virtual Machines

How many virtual machines do you want to create? 2000

Configure your machines.

Total memory (MB) on each machine: 4096

Select a virtual machine copy mode.

● Use fast clone for more efficient storage use and faster machine creation.

○ Use full copy for better data recovery and migration support, with potentially reduced IOPS after the machines are created.

Step 51. Specify the AD account naming scheme and OU where accounts will be created.

Step 52. Click Next.

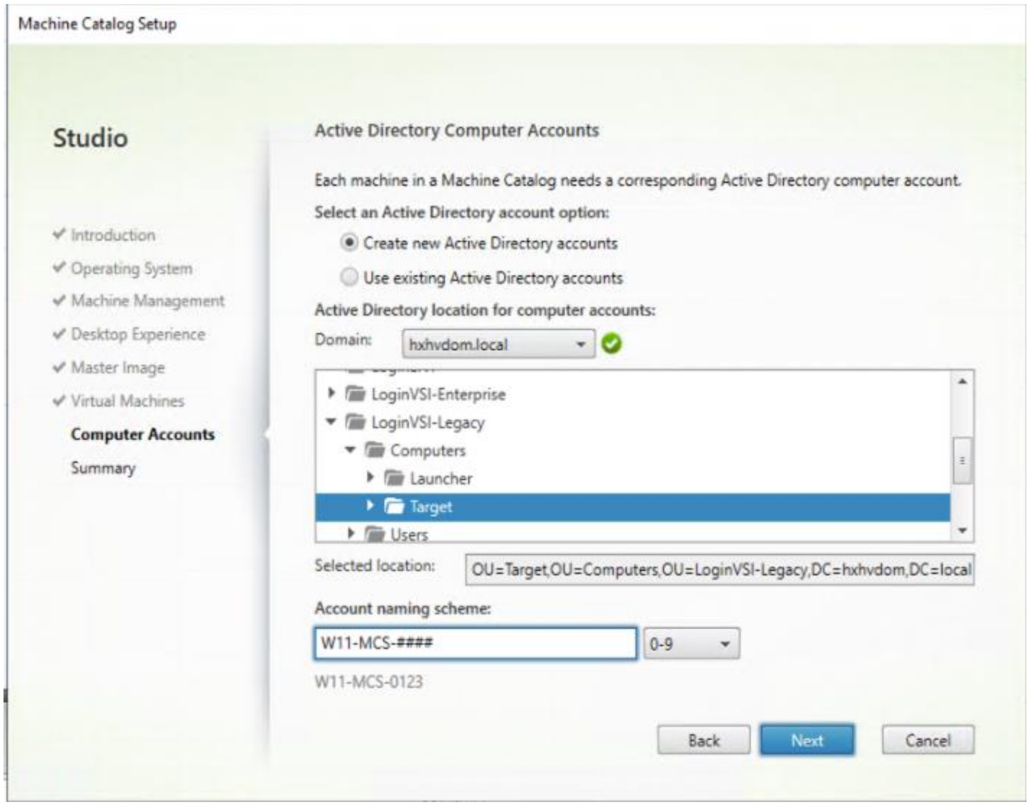

Machine Catalog Setup

Studio

✓ Introduction
✓ Operating System
✓ Machine Management
✓ Desktop Experience
✓ Master Image
✓ Virtual Machines
Computer Accounts
Summary

Active Directory Computer Accounts

Each machine in a Machine Catalog needs a corresponding Active Directory computer account.

Select an Active Directory account option:

● Create new Active Directory accounts

○ Use existing Active Directory accounts

Active Directory location for computer accounts:

Domain: hxhvdom.local

Selected location: OU=Target,OU=Computers,OU=LoginVSI-Legacy,DC=hxhvdom,DC=local

Account naming scheme: W11-MCS-#### 0-9

W11-MCS-0123

Step 53. On the Summary page specify Catalog name and click Finish to start the deployment.

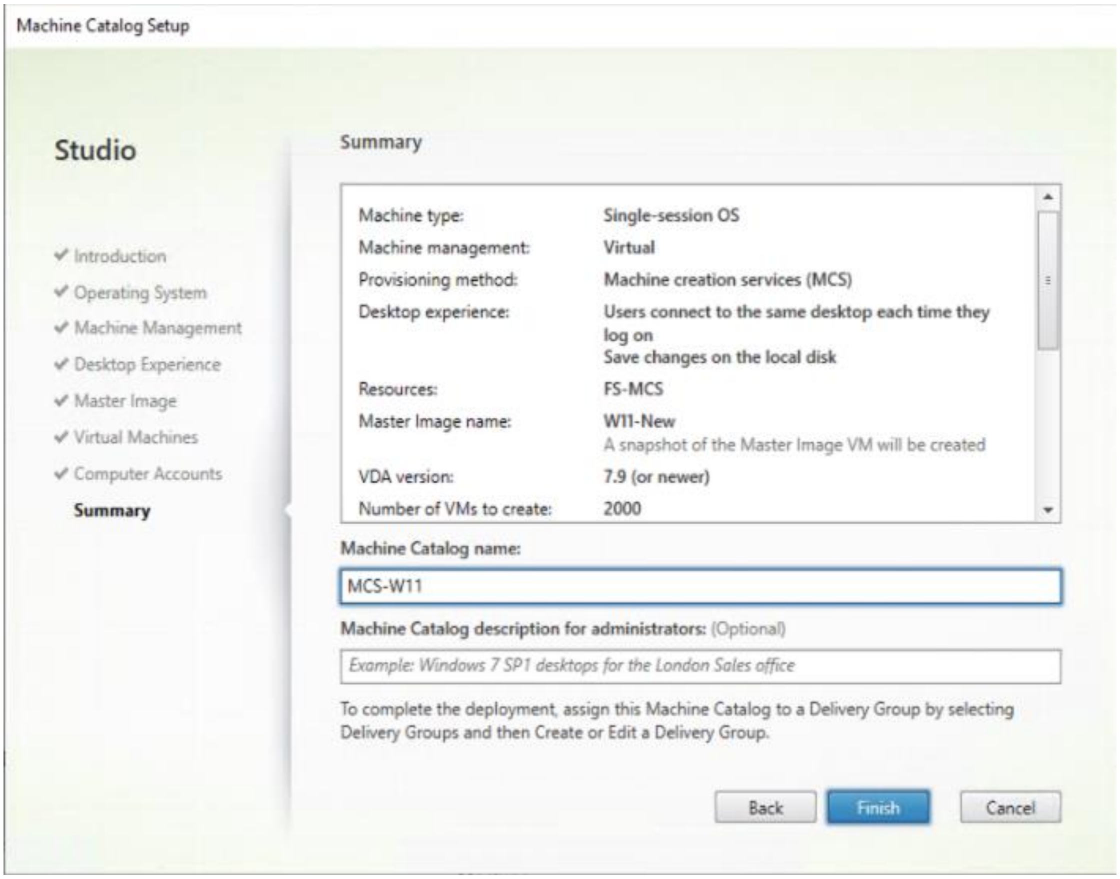

Summary

Machine type:	Single-session OS
Machine management:	Virtual
Provisioning method:	Machine creation services (MCS)
Desktop experience:	Users connect to the same desktop each time they log on Save changes on the local disk
Resources:	FS-MCS
Master Image name:	W11-New A snapshot of the Master Image VM will be created
VDA version:	7.9 (or newer)
Number of VMs to create:	2000

Machine Catalog name:

MCS-W11

Machine Catalog description for administrators: (Optional)

Example: Windows 7 SP1 desktops for the London Sales office

To complete the deployment, assign this Machine Catalog to a Delivery Group by selecting Delivery Groups and then Create or Edit a Delivery Group.

Procedure 3. Create Delivery Groups

Delivery Groups are collections of machines that control access to desktops and applications. With Delivery Groups, you can specify which users and groups can access which desktops and applications.

Note: The instructions below outline the procedure to create a Delivery Group for persistent VDI desktops. When you have completed these steps, repeat the procedure to a Delivery Group for RDS desktops.

Step 1. Connect to a Citrix Virtual Apps and Desktops server and launch Citrix Studio.

Step 2. Choose Create Delivery Group from the drop-down list.

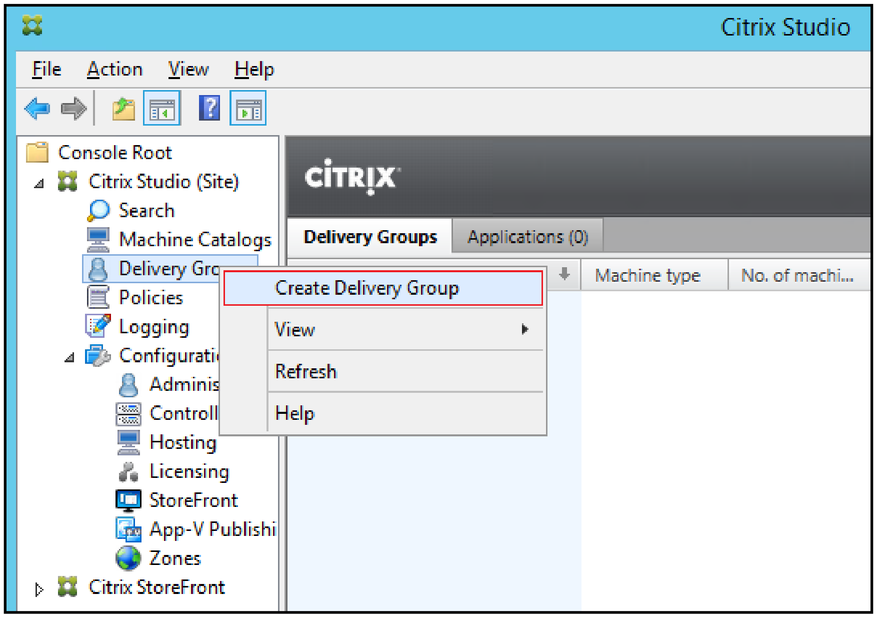

Step 3. Click Next.

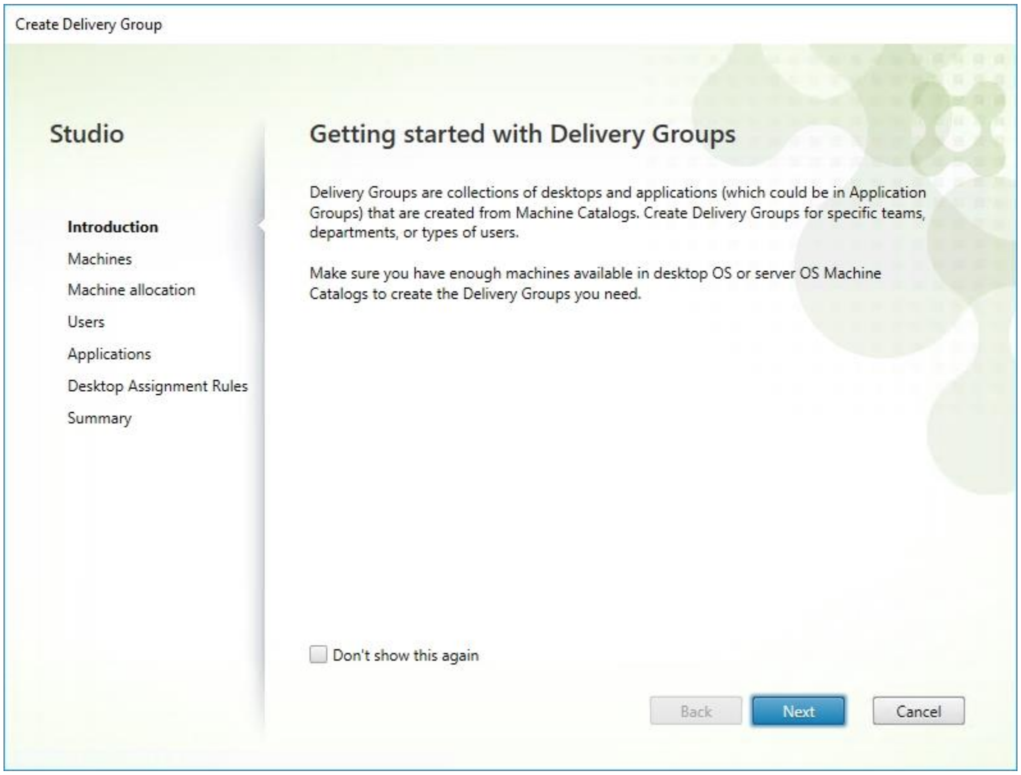
Create Delivery Group

Studio

Introduction
Machines
Machine allocation
Users
Applications
Desktop Assignment Rules
Summary

Getting started with Delivery Groups

Delivery Groups are collections of desktops and applications (which could be in Application Groups) that are created from Machine Catalogs. Create Delivery Groups for specific teams, departments, or types of users.

Make sure you have enough machines available in desktop OS or server OS Machine Catalogs to create the Delivery Groups you need.

□ Don't show this again

Step 4. Specify the Machine Catalog and increment the number of machines to add.

Step 5. Click Next.

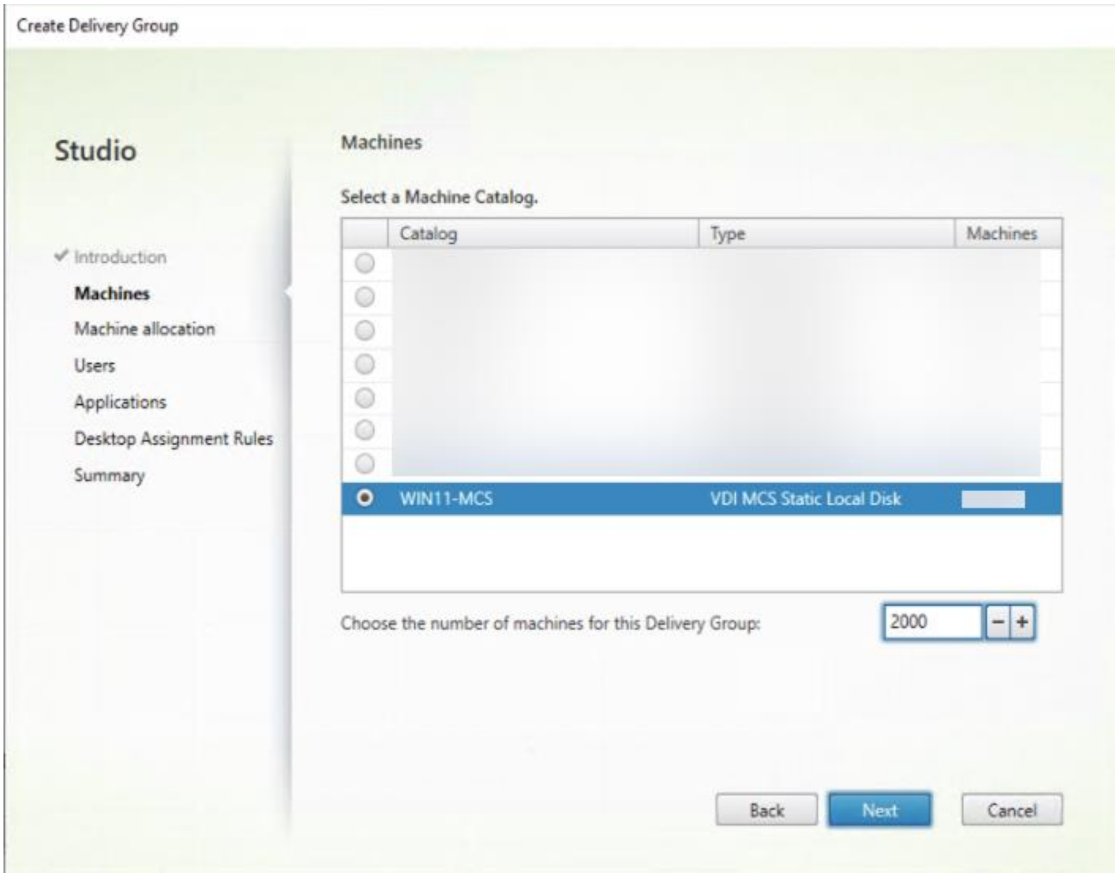

Machines

Select a Machine Catalog.

	Catalog	Type	Machines
○			
○			
○			
○			
○			
○			
○			
●	WIN11-MCS	VDI MCS Static Local Disk	

Choose the number of machines for this Delivery Group: 2000

Step 6. Specify what the machines in the catalog will deliver: Desktops, Desktops and Applications, or Applications.

Step 7. Select Desktops.

Step 8. Click Next.

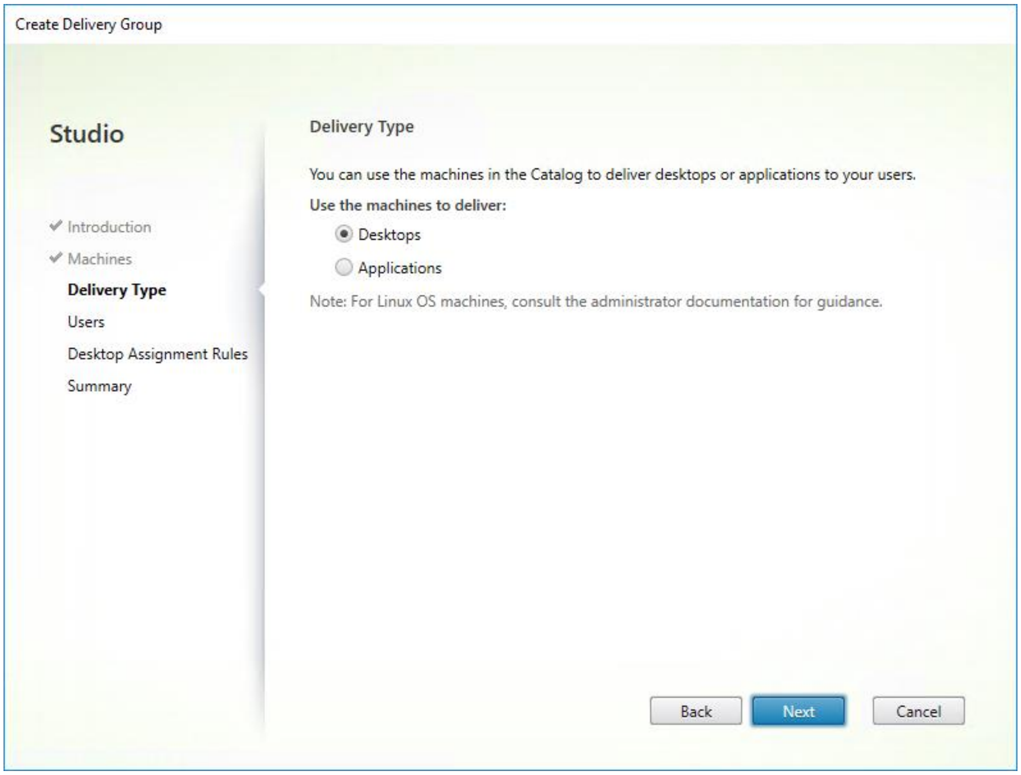

Delivery Type

You can use the machines in the Catalog to deliver desktops or applications to your users.

Use the machines to deliver:

● Desktops

○ Applications

Note: For Linux OS machines, consult the administrator documentation for guidance.

Step 9. To make the Delivery Group accessible, you must add users. Select Allow any authenticated users to use this Delivery Group.

Note: User assignment can be updated any time after Delivery group creation by accessing Delivery group properties in Desktop Studio.

Step 10. Click Next.

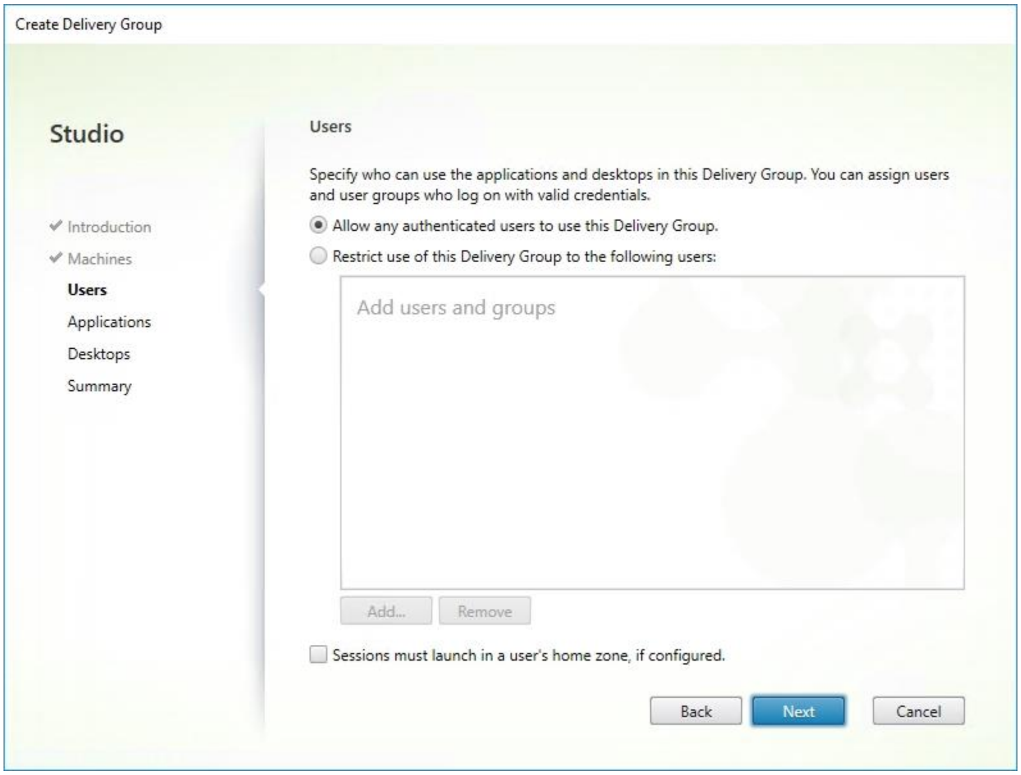

Create Delivery Group

Studio

Users

Specify who can use the applications and desktops in this Delivery Group. You can assign users and user groups who log on with valid credentials.

● Allow any authenticated users to use this Delivery Group.

○ Restrict use of this Delivery Group to the following users:

□ Sessions must launch in a user's home zone, if configured.

Step 11. Click Next (no applications are used in this design).

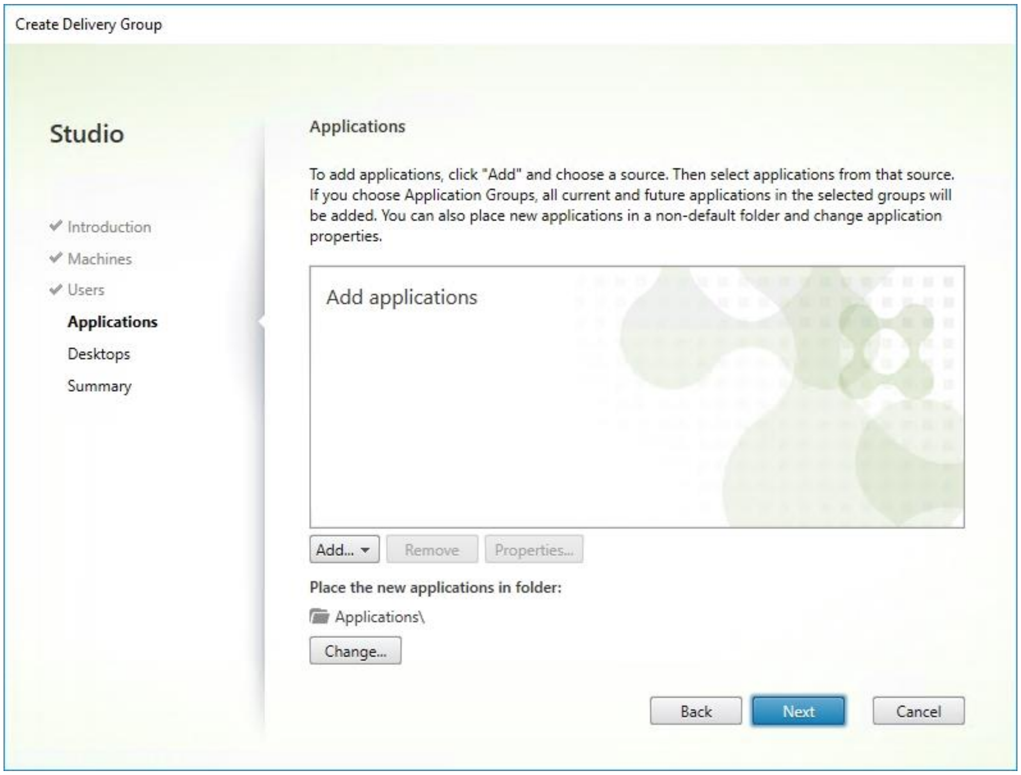

Create Delivery Group

Studio

Applications

To add applications, click "Add" and choose a source. Then select applications from that source. If you choose Application Groups, all current and future applications in the selected groups will be added. You can also place new applications in a non-default folder and change application properties.

Place the new applications in folder:

Applications\

Step 12. Enable Users to access the desktops.

Step 13. Click Next.

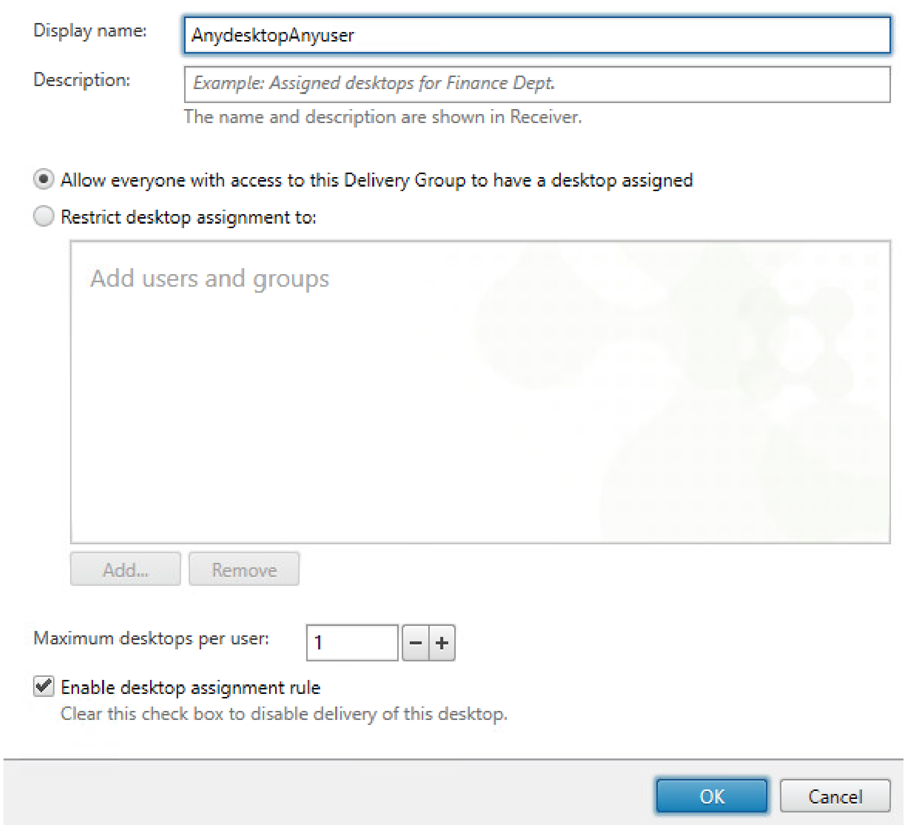

Display name: AnydesktopAnyuser

Description: *Example: Assigned desktops for Finance Dept.*

The name and description are shown in Receiver.

● Allow everyone with access to this Delivery Group to have a desktop assigned

○ Restrict desktop assignment to:

Add users and groups

Maximum desktops per user: 1

☑ Enable desktop assignment rule

Clear this check box to disable delivery of this desktop.

Step 14. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Description (Optional).

Step 15. Click Finish.

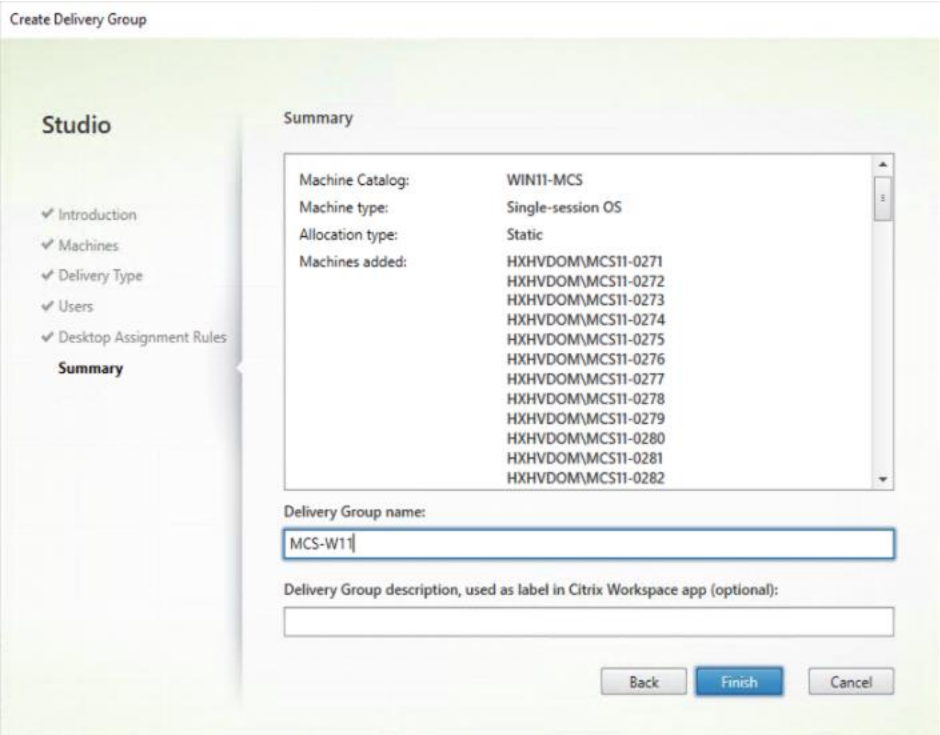

Create Delivery Group

Summary

Machine Catalog:	WIN11-MCS
Machine type:	Single-session OS
Allocation type:	Static
Machines added:	HXHVDOM\MCS11-0271 HXHVDOM\MCS11-0272 HXHVDOM\MCS11-0273 HXHVDOM\MCS11-0274 HXHVDOM\MCS11-0275 HXHVDOM\MCS11-0276 HXHVDOM\MCS11-0277 HXHVDOM\MCS11-0278 HXHVDOM\MCS11-0279 HXHVDOM\MCS11-0280 HXHVDOM\MCS11-0281 HXHVDOM\MCS11-0282

Delivery Group name: MCS-W11

Delivery Group description, used as label in Citrix Workspace app (optional):

Citrix Studio lists the created Delivery Groups as well as the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab.

Step 16. From the drop-down list, select "Turn on Maintenance Mode."

Citrix Virtual Apps and Desktops Policies and Profile Management

Policies and profiles allow the Citrix Virtual Apps and Desktops environment to be easily and efficiently customized.

Configure Citrix Virtual Apps and Desktops Policies

Citrix Virtual Apps and Desktops policies control user access and session environments, and are the most efficient method of controlling connection, security, and bandwidth settings. You can create policies for specific groups of users, devices, or connection types with each policy. Policies can contain multiple settings and are typically defined through Citrix Studio.

Note: The Windows Group Policy Management Console can also be used if the network environment includes Microsoft Active Directory and permissions are set for managing Group Policy Objects).

Figure 30 shows the policies for Login VSI testing in this CVD.

Figure 30. **Citrix Virtual Apps and Desktops Policy**

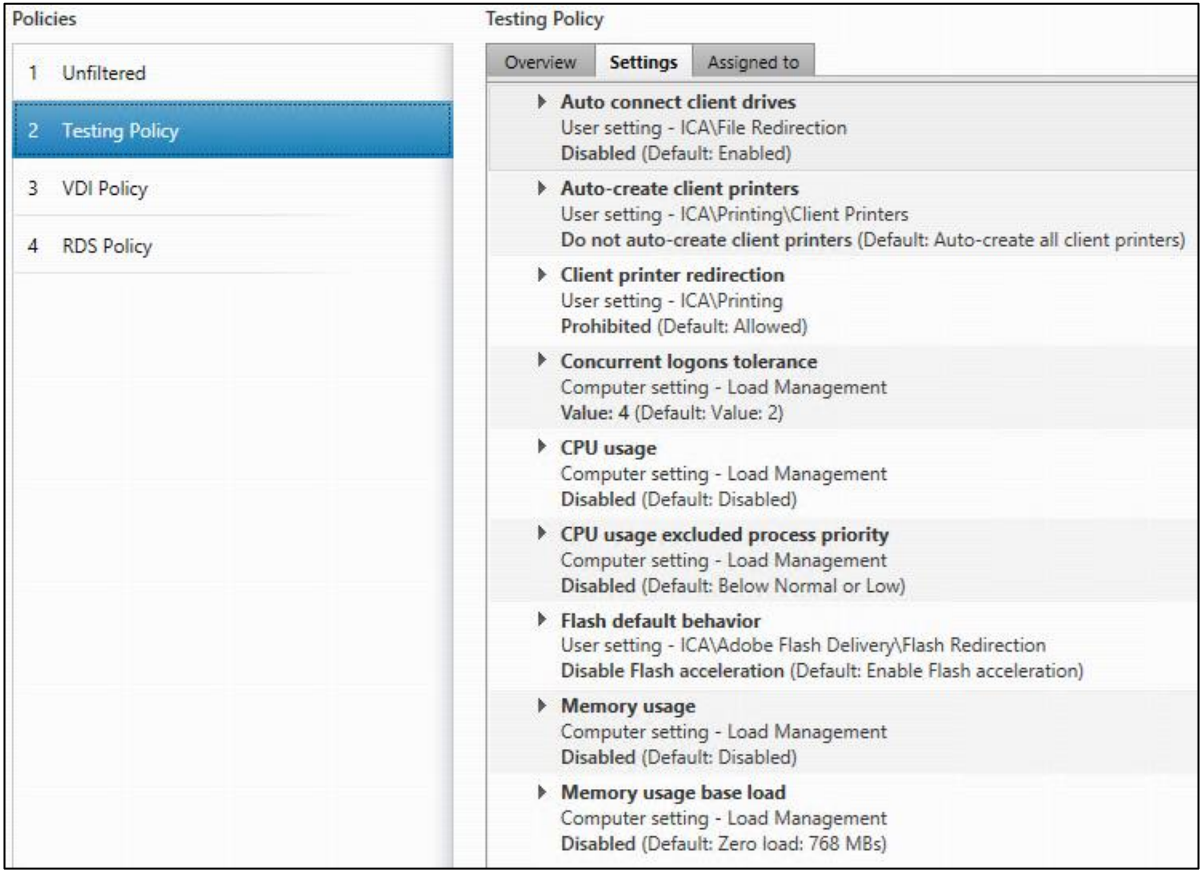

Auto connect client drives
User setting - ICA\File Redirection
Disabled (Default: Enabled)

Auto-create client printers
User setting - ICA\Printing\Client Printers
Do not auto-create client printers (Default: Auto-create all client printers)

Client printer redirection
User setting - ICA\Printing
Prohibited (Default: Allowed)

Concurrent logons tolerance
Computer setting - Load Management
Value: 4 (Default: Value: 2)

CPU usage
Computer setting - Load Management
Disabled (Default: Disabled)

CPU usage excluded process priority
Computer setting - Load Management
Disabled (Default: Below Normal or Low)

Flash default behavior
User setting - ICA\Adobe Flash Delivery\Flash Redirection
Disable Flash acceleration (Default: Enable Flash acceleration)

Memory usage
Computer setting - Load Management
Disabled (Default: Disabled)

Memory usage base load
Computer setting - Load Management
Disabled (Default: Zero load: 768 MBs)

Figure 31. Delivery Controllers Policy

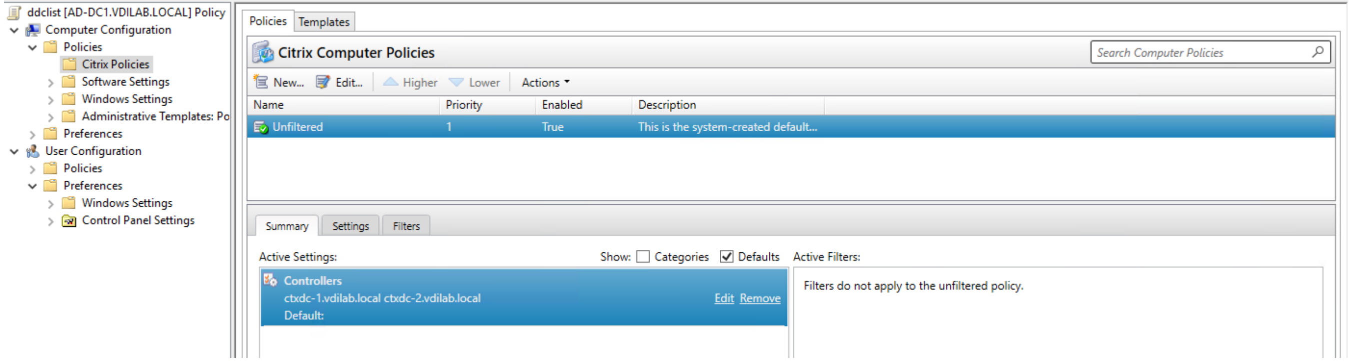

Name	Priority	Enabled	Description
Unfiltered	1	True	This is the system-created default...

Active Settings:

Controllers
ctxdc-1.vdilab.local ctxdc-2.vdilab.local
Default:

Active Filters:

Filters do not apply to the unfiltered policy.

FSLogix for Citrix Virtual Apps & Desktops Profile Management

This subject contains the following procedures:

- Configure FSLogix for Citrix Virtual Apps & Desktops Profiles Profile Container
- Configure FSLogix Profile Management

FSLogix for user profiles allows the Citrix Virtual Apps & Desktops environment to be easily and efficiently customized.

Procedure 1. Configure FSLogix for Citrix Virtual Apps & Desktops Profiles Profile Container

Profile Container is a full remote profile solution for non-persistent environments. Profile Container redirects the entire user profile to a remote location. Profile Container configuration defines how and where the profile is redirected.

Note: Profile Container is inclusive of the benefits found in Office Container.

Note: When using Profile Container, both applications and users see the profile as if it's located on the local drive.

Step 1. Verify that you meet all entitlement and configuration requirements.

Step 2. Download and install FSLogix Software

Step 3. Consider the storage and network requirements for your users' profiles (in this CVD, we used Pure File Servers to store the FSLogix Profile disks).

Step 4. Verify that your users have appropriate storage permissions where profiles will be placed.

Step 5. Profile Container is installed and configured after stopping use of other solutions used to manage remote profiles.

Step 6. Exclude the VHD(X) files for Profile Containers from Anti-Virus (AV) scanning.

Procedure 2. Configure FSLogix Profile Management

Step 1. When the FSLogix software is downloaded, copy the ‘fslogix.admx and fslogix.adml’ to the ‘PolicyDefinitions’ folder in your domain to manage the settings with Group Policy.

Step 2. On your VDI master image, install the FSLogix agent ‘FSLogixAppsSetup’ and accept all the defaults.

Step 3. Create a Group Policy object and link it to the Organizational Unit the VDI computer accounts.

Step 4. Right-click the FSLogix GPO policy.

Step 5. Enable FSLogix Profile Management.

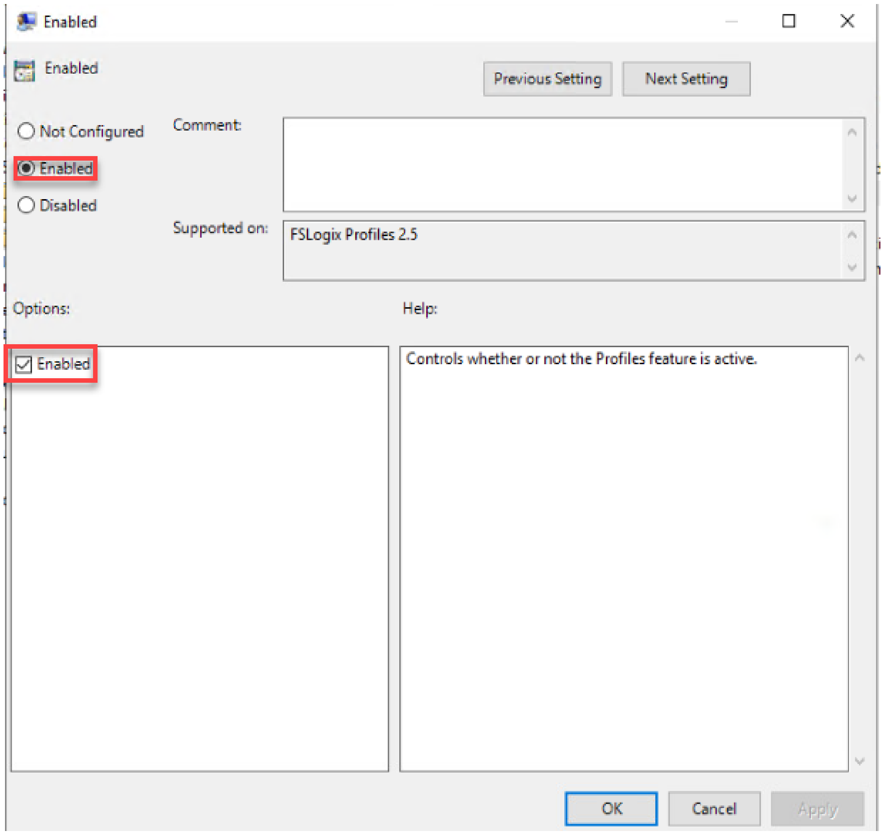

Enabled

○ Not Configured
● Enabled
○ Disabled

Comment:

Supported on: FSLogix Profiles 2.5

Options:

☑ Enabled

Help:

Controls whether or not the Profiles feature is active.

Step 6. Select Profile Type (in this solution, we used Read-Write profiles).

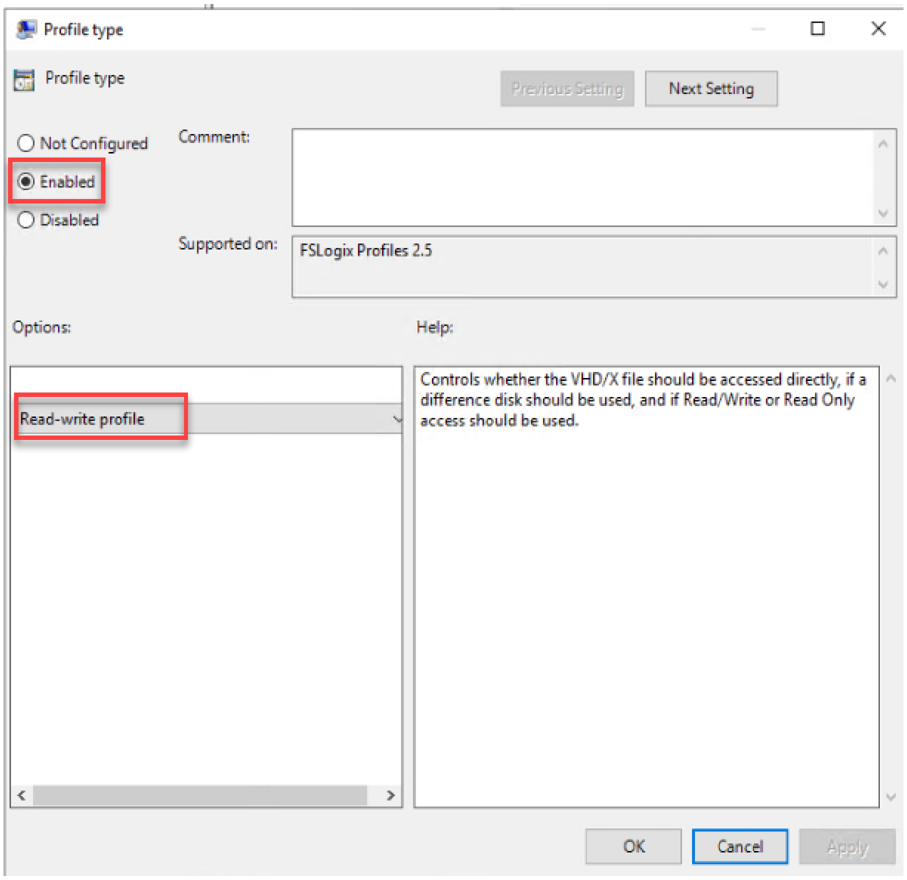

Profile type

○ Not Configured
● Enabled
○ Disabled

Supported on: FSLogix Profiles 2.5

Options: Read-write profile

Help: Controls whether the VHD/X file should be accessed directly, if a difference disk should be used, and if Read/Write or Read Only access should be used.

Step 7. Enter the location of the Profile location.

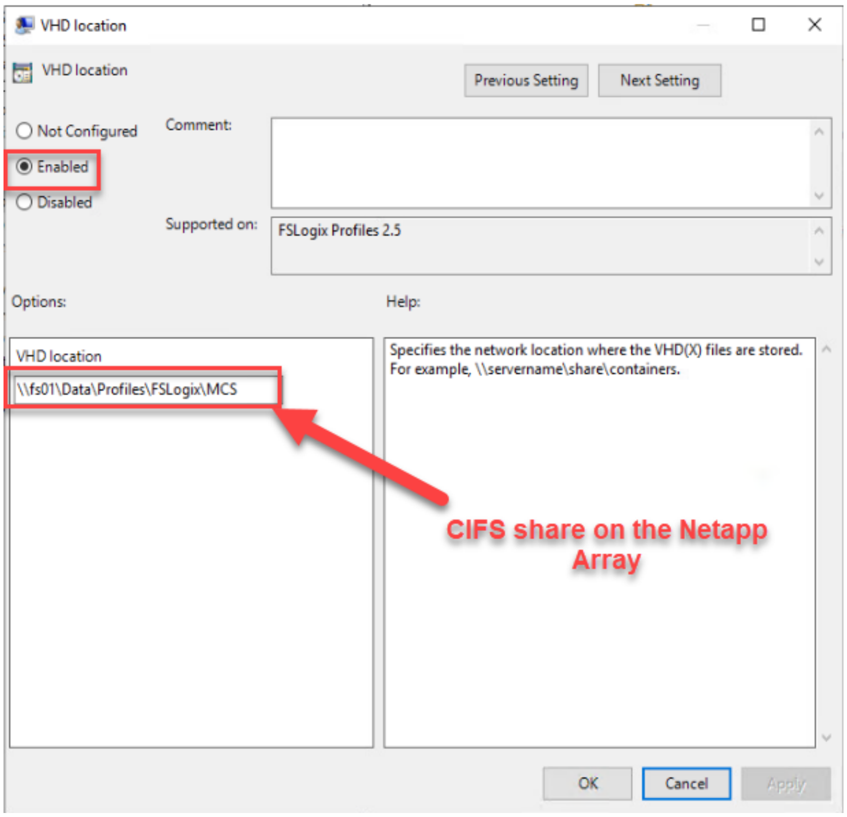

VHD location

○ Not Configured
● Enabled
○ Disabled

Supported on: FSLogix Profiles 2.5

Options: VHD location \\fs01\Data\Profiles\FSLogix\MCS

Help: Specifies the network location where the VHD(X) files are stored. For example, \\servername\share\containers.

Note: We recommend using the Dynamic VHDX setting.

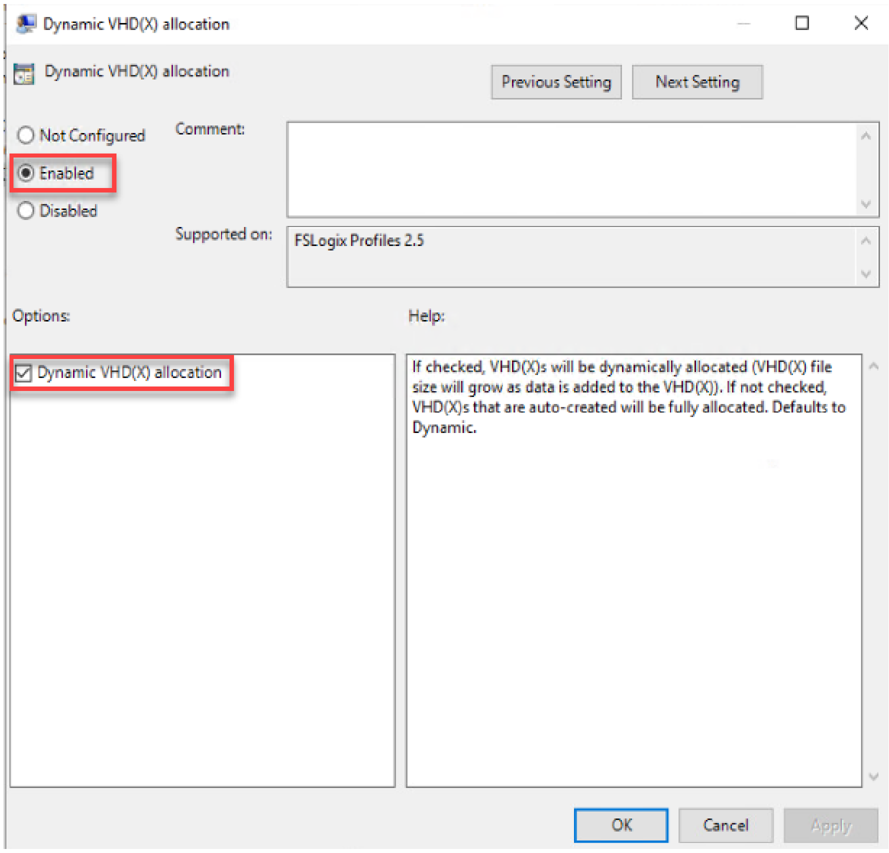

Dynamic VHD(X) allocation

○ Not Configured
● Enabled
○ Disabled

Comment:

Supported on: FSLogix Profiles 2.5

Options:

☑ Dynamic VHD(X) allocation

Help:

If checked, VHD(X)s will be dynamically allocated (VHD(X) file size will grow as data is added to the VHD(X)). If not checked, VHD(X)s that are auto-created will be fully allocated. Defaults to Dynamic.

Note: VHDX is recommended over VHD.

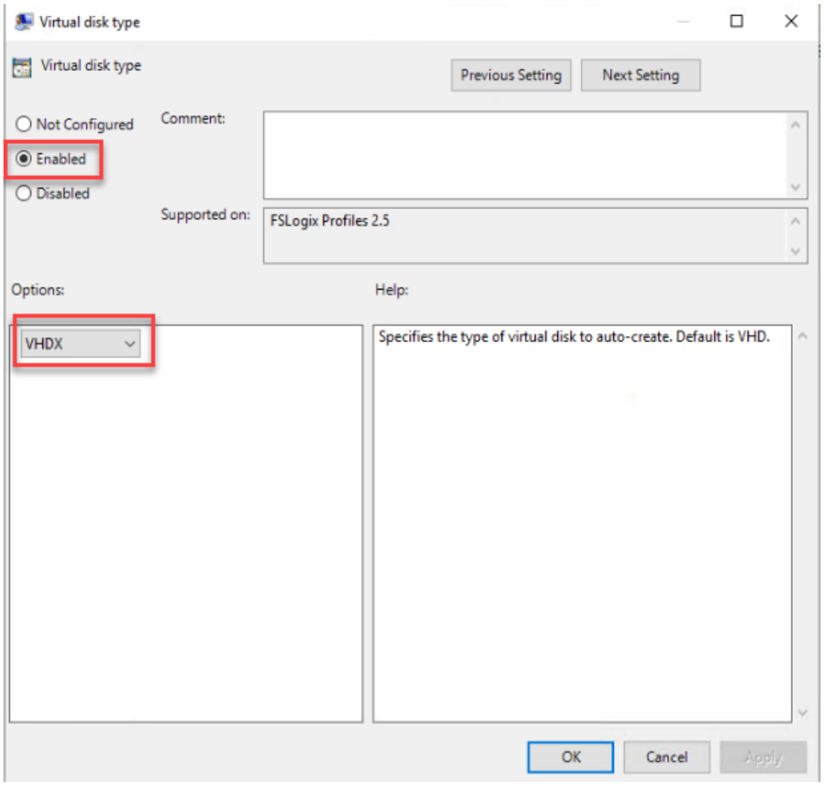
Virtual disk type

○ Not Configured
● Enabled
○ Disabled

Comment:

Supported on: FSLogix Profiles 2.5

Options:

VHDX

Help:

Specifies the type of virtual disk to auto-create. Default is VHD.

Note: We enabled the ‘Swap directory name components’ setting for an easier administration but is not necessary for improved performance.

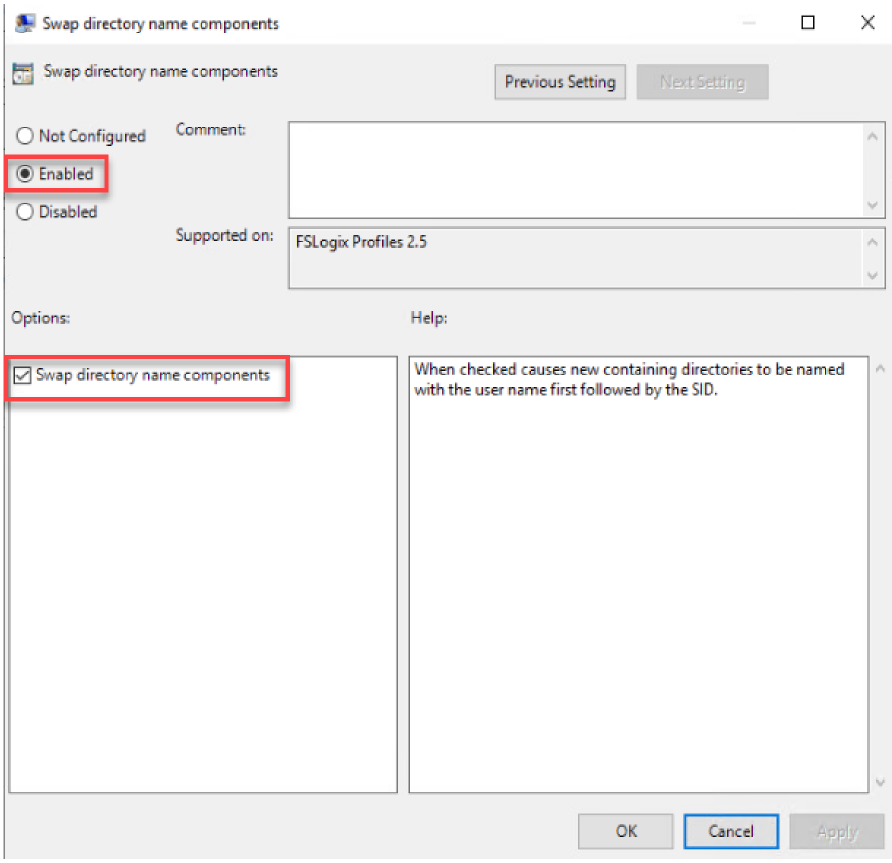
Swap directory name components

○ Not Configured
● Enabled
○ Disabled

Comment:

Supported on: FSLogix Profiles 2.5

Options:

☑ Swap directory name components

Help:

When checked causes new containing directories to be named with the user name first followed by the SID.

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 M6 Compute Node to validate against the performance of one and eight Cisco UCS X210c M6 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 Recommended Maximum Workload per host server using Citrix Virtual Apps and Desktops 2203 with 335 Multi-session OS sessions and 270 Single-session OS sessions.

Figure 32. Test Configuration for Single Server Scalability Citrix Virtual Apps and Desktops 2203 Non-persistent (PVS) Single-session OS machine VDAs

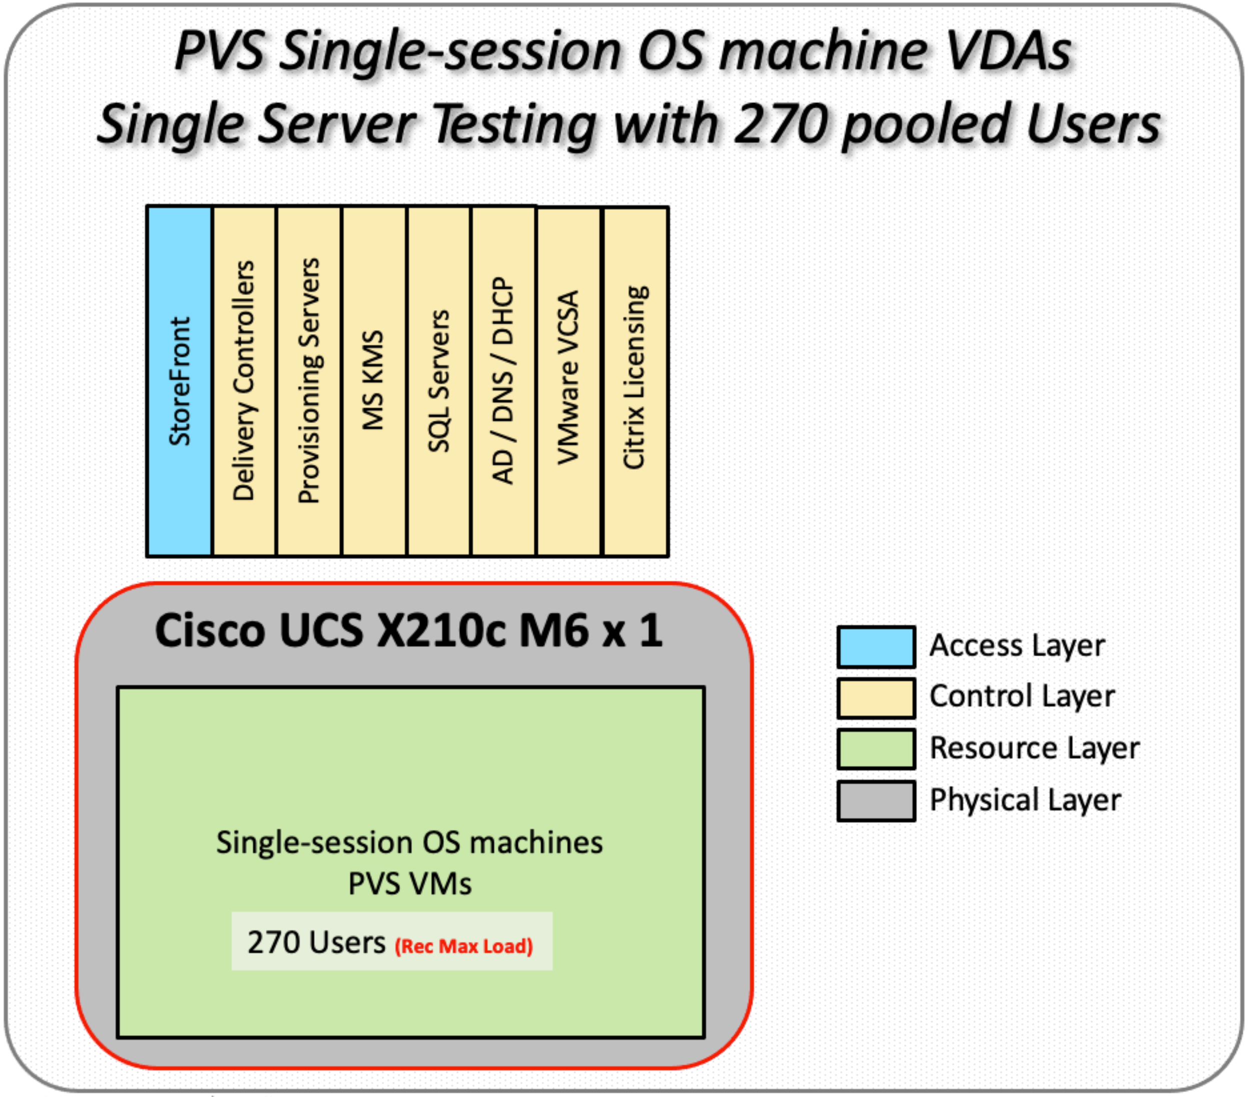

Figure 33. Test configuration for Single Server Scalability Citrix Virtual Apps and Desktops 2203 Persistent (MCS) Single-session OS machine VDAs

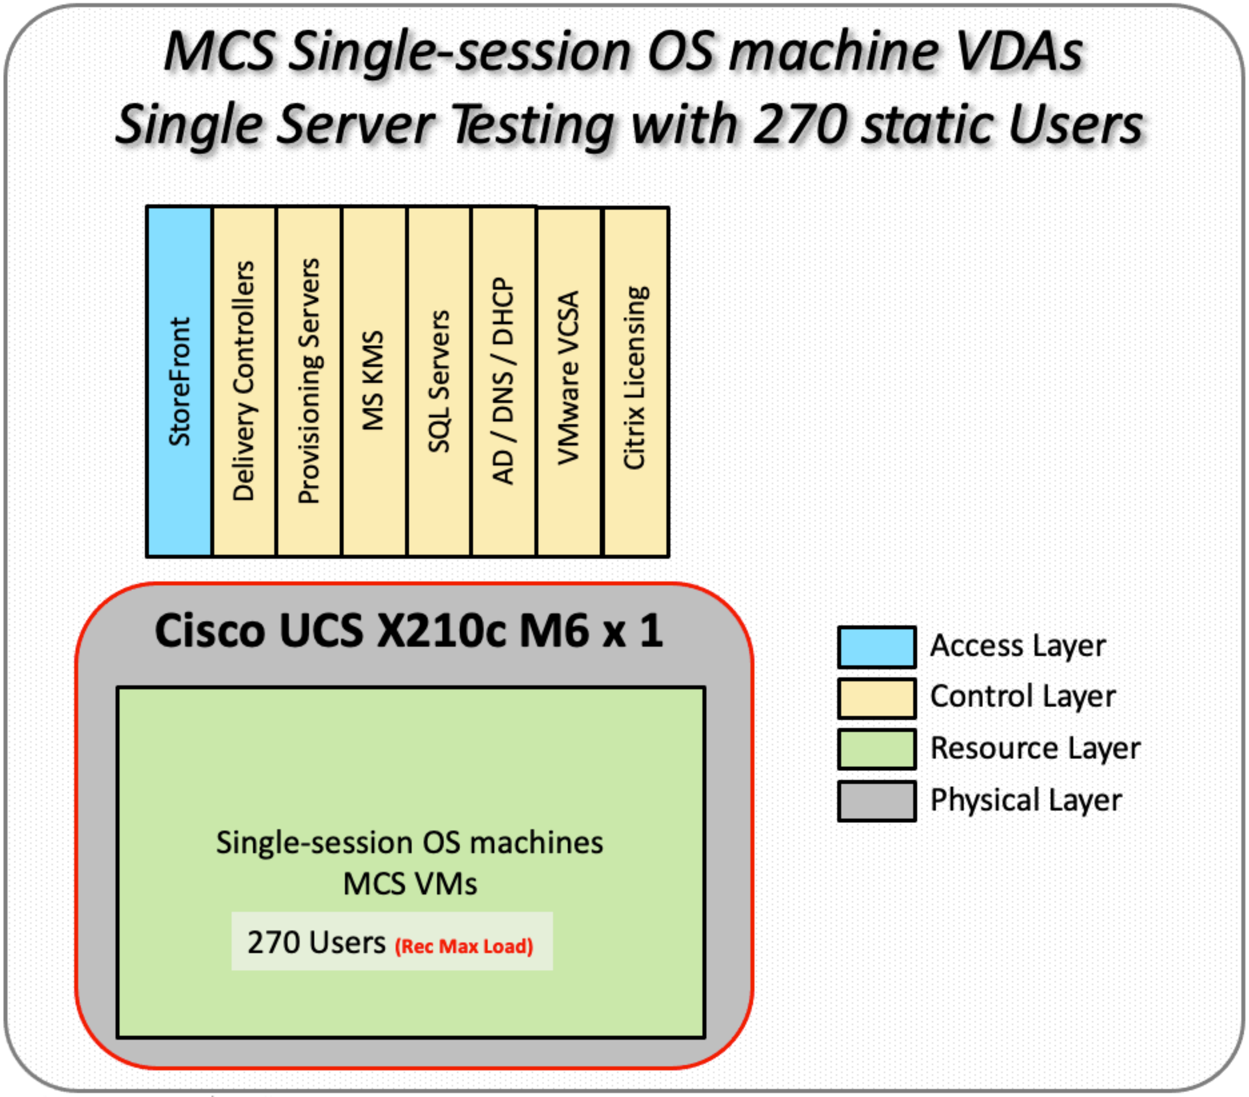

Figure 34. Test configuration for Single Server Scalability Citrix Virtual Apps and Desktops 2203 PVS Multi-session OS machine VDAs

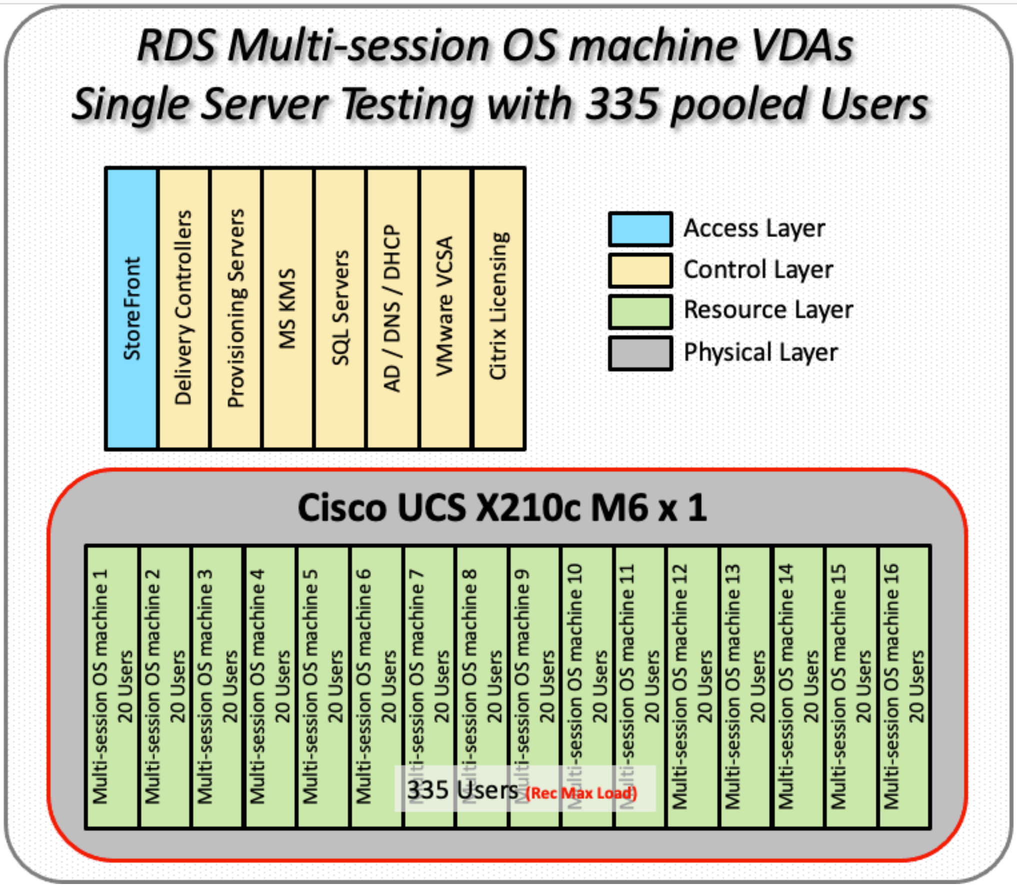
RDS Multi-session OS machine VDAs
Single Server Testing with 335 pooled Users
StoreFront
Delivery Controllers
Provisioning Servers
MS KMS
SQL Servers
AD / DNS / DHCP
VMware VCSA
Citrix Licensing
Access Layer
Control Layer
Resource Layer
Physical Layer
Cisco UCS X210c M6 x 1
Multi-session OS machine 1 20 Users
Multi-session OS machine 2 20 Users
Multi-session OS machine 3 20 Users
Multi-session OS machine 4 20 Users
Multi-session OS machine 5 20 Users
Multi-session OS machine 6 20 Users
Multi-session OS machine 7 20 Users
Multi-session OS machine 8 20 Users
Multi-session OS machine 9 20 Users
Multi-session OS machine 10 20 Users
Multi-session OS machine 11 20 Users
Multi-session OS machine 12 20 Users
Multi-session OS machine 13 20 Users
Multi-session OS machine 14 20 Users
Multi-session OS machine 15 20 Users
Multi-session OS machine 16 20 Users
335 Users (Rec Max Load)

Hardware components:

- Cisco UCS X9508 Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 1 Cisco UCS X210c M6 Compute Node Servers with Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM for all host blades
- Cisco UCS VIC 14425 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switches
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with 20 1.92TB DirectFlash NVMe drives

Software components:

- Cisco UCS firmware 5.0(2e)
- Pure Storage Purity//FA 6.3.3
- ESXi 8.0 for host blades
- Citrix Virtual Apps and Desktops 2203
- Microsoft SQL Server 2019

- Microsoft Windows 11 64 bit, 2vCPU, 4 GB RAM, 40 GB HDD (master)
- Microsoft Windows Server 2022 8vCPU, 24GB RAM, 60 GB vDisk (master)
- Microsoft Office 2021
- FSLogix 2105 HF_01
- Login VSI 4.1.39 Knowledge Worker Workload (Benchmark Mode)

Cisco UCS Test Configuration for Full Scale Testing

These test cases validate eight blades in a cluster hosting three distinct workloads using Citrix Virtual Apps and Desktops 2203 with:

- 2000 VDI-NP Single-session OS sessions (Citrix PVS)
- 2000 VDI-P Single-session OS sessions (Citrix MCS)
- 2600 Citrix PVS RDS sessions

Note: Server N+1 fault tolerance is factored into this solution for each cluster/workload.

Figure 35. **Test Configuration for Full Scale Citrix Virtual Apps and Desktops 2203 non-persistent Single-session OS machine VDAs**

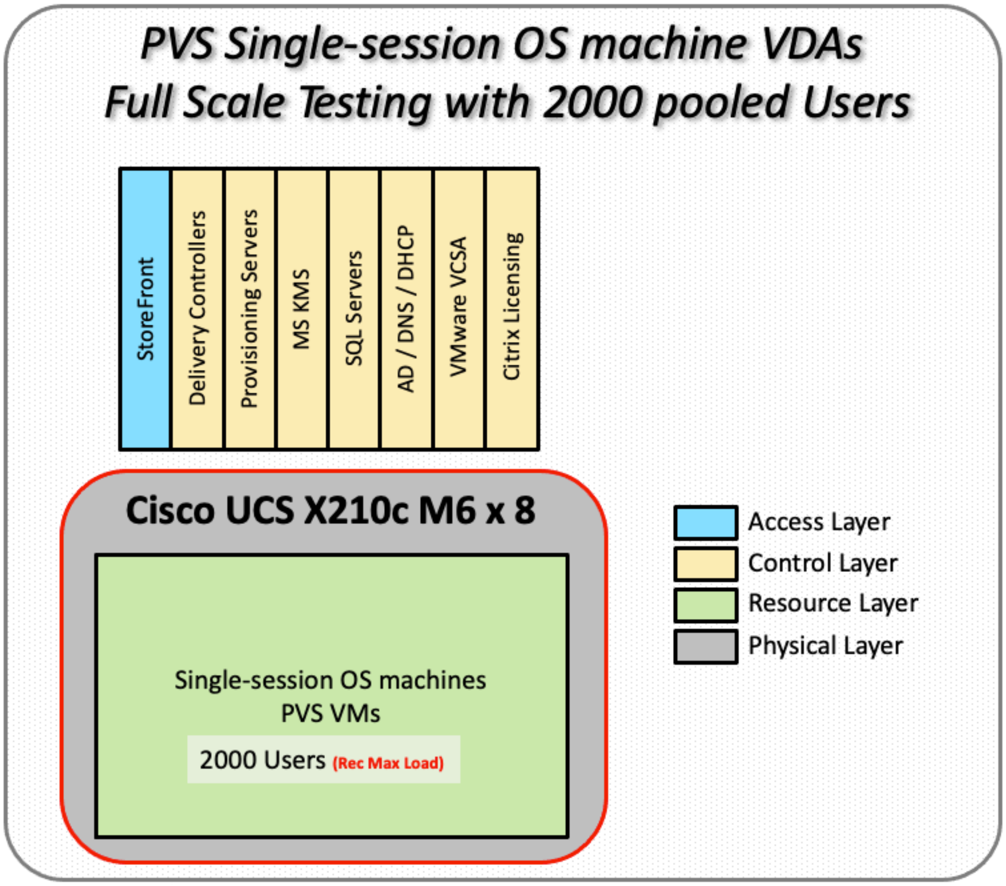

Figure 36. Test Configuration for Full Scale Citrix Virtual Apps and Desktops 2203 persistent Single-session OS machine VDAs

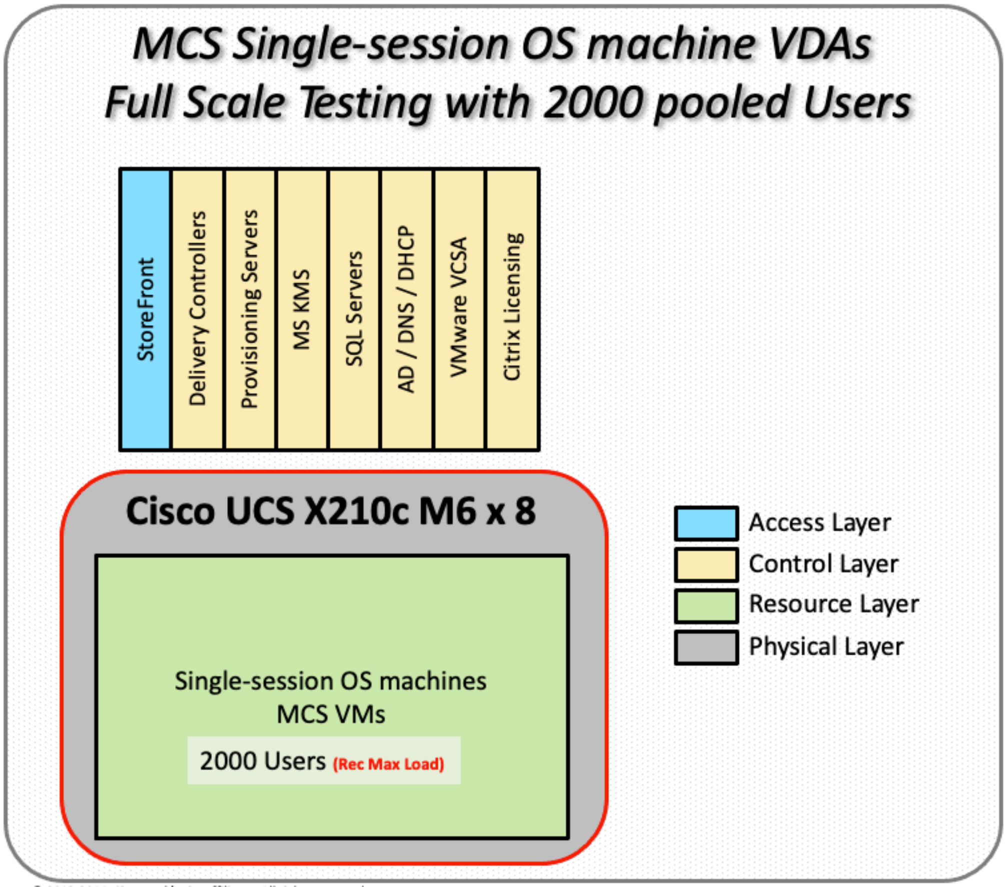

Figure 37. Test Configuration for Full Scale Citrix Virtual Apps and Desktops 2203 instant-clones Multi-session OS machine VDAs

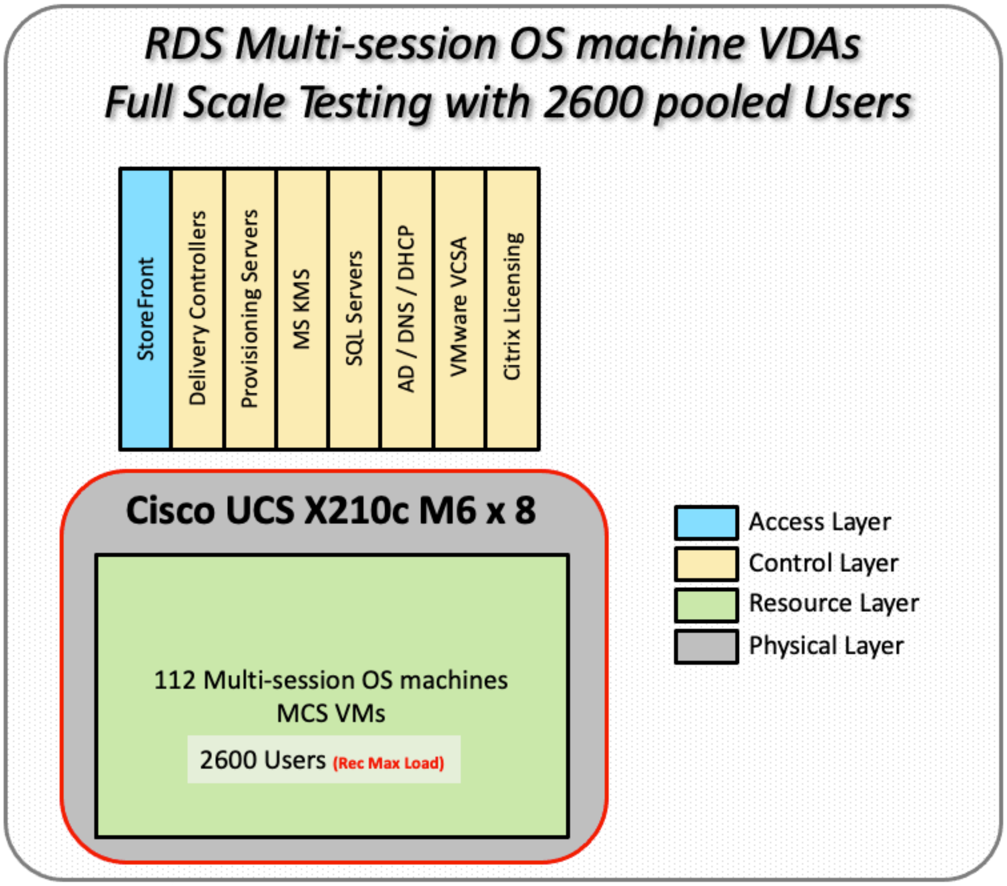

Hardware components:

- Cisco UCS X9508 Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 8 Cisco UCS X210c M6 Compute Node Servers with Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM for all host blades
- Cisco VIC 14425 CNA (1 per blade)
- 2 Cisco Nexus 93180YC-FX Access Switches
- 2 Cisco MDS 9132T 32-Gb 32-Port Fibre Channel Switches
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with 20 1.92TB DirectFlash NVMe drives

Software components:

- Cisco UCS firmware 5.0(2e)
- Pure Storage Purity//FA 6.3.3
- ESXi 8.0 host blades
- Citrix Virtual Apps and Desktops 2203
- Microsoft SQL Server 2019

- Microsoft Windows 11, 2vCPU, 4GB RAM, 40 GB HDD (master)
- Microsoft Windows Server 2022, 8vCPU, 24GB RAM, 60 GB vDisk (master)
- Microsoft Office 2021
- FSLogix 2015 HF_01
- Login VSI 4.1.39 Knowledge Worker Workload (Benchmark Mode)

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

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 Citrix Studio.

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 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 Citrix Studio 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 dynamic in our testing. FlashStack Data Center with Cisco UCS and Citrix Virtual Apps and Desktops 2203 on VMware ESXi 8.0 Test Results.

The purpose of this testing is to provide the data needed to validate Citrix Virtual Apps and Desktops Remote Desktop Sessions (RDS) and Citrix Virtual Desktop (PVS) non-persistent and Citrix Virtual Desktop (MCS) full-clones models using ESXi and vCenter to virtualize Microsoft Windows 11 desktops and Microsoft Windows Server 2022 sessions on Cisco UCS X210c M6 Compute Node Servers using the Pure Storage FlashArray//X70 R3 storage system.

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 show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what 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 38. **Sample of a VSI Max Response Time Graph, Representing a Normal Test**

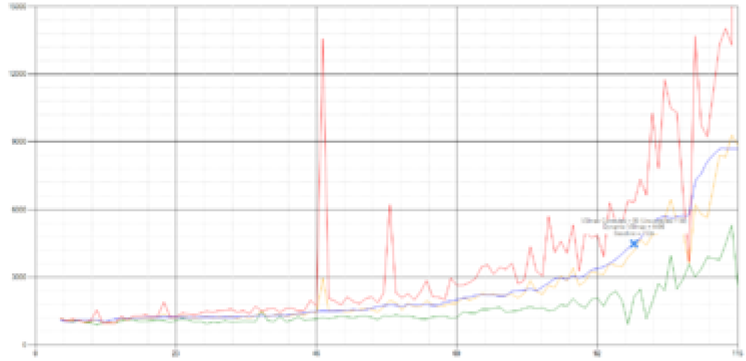

Figure 39. **Sample of a VSI Test Response Time Graph with a Performance Issue**

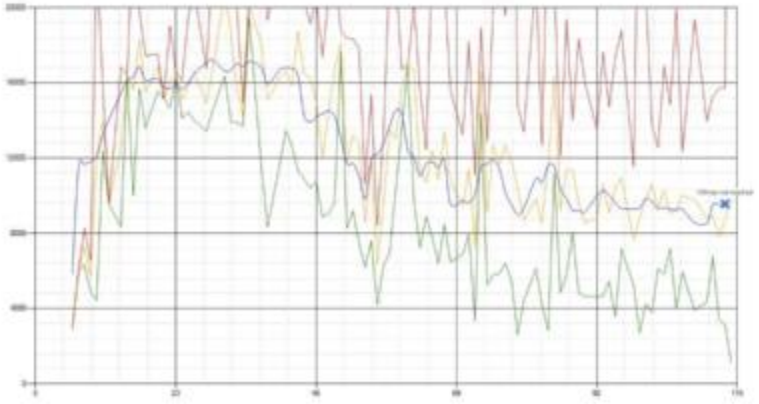

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached, and the number of sessions ran successfully.

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 2 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 Citrix Virtual Apps and Desktops 2203 Virtual Desktop and Citrix Virtual Apps and Desktops 2203 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.

Memory should never be oversubscribed for Desktop Virtualization workloads.

Table 21. 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

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:

- 270 VDI Non-Persistent sessions (Random)
- 270 VDI Persistent sessions (Static)
- 335 Multisession OS RDS sessions (Random)

Single-Server Recommended Maximum Workload for non-persistent Single-session OS Random Sessions with 270 Users

The recommended maximum workload for a Cisco UCS X210c M6 Compute Node server with dual Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM is 270 Windows 11 64-bit non-persistent virtual machines with 2 vCPU and 4 GB RAM.

Login VSI performance data is shown below:

Figure 40. **Single Server | Citrix Virtual Apps and Desktops 2203 non-persistent Single-session OS machine VDAs | VSI Score**

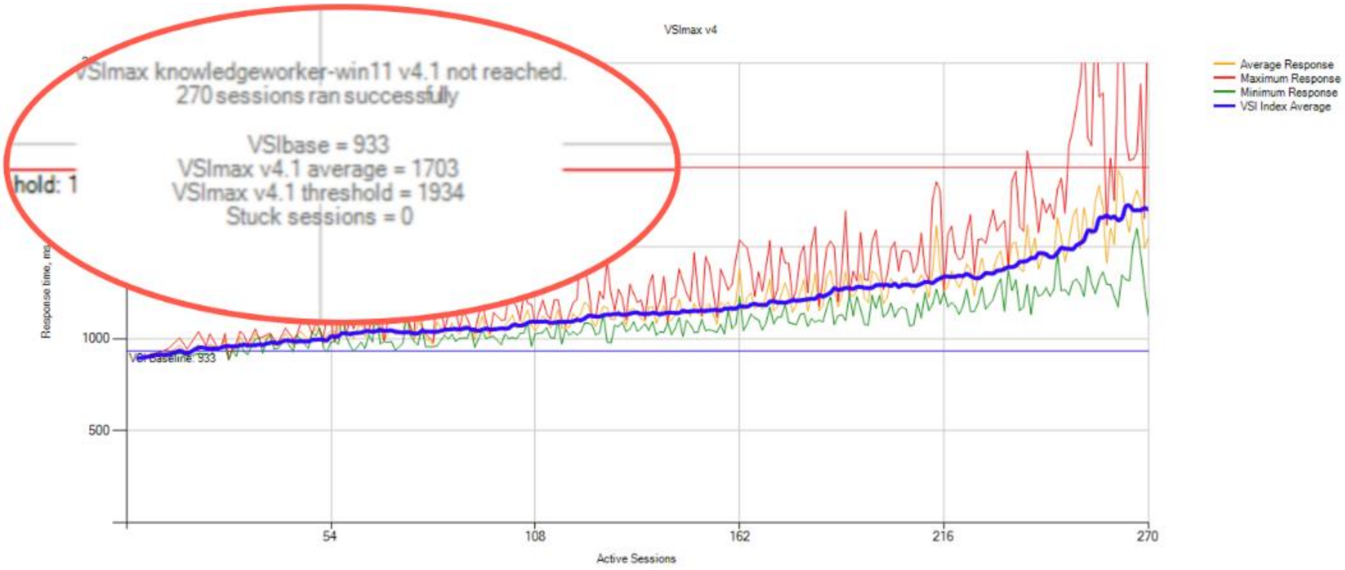

Performance data for the server running the workload is shown below:

Figure 41. Single Server | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs | Host CPU Utilization

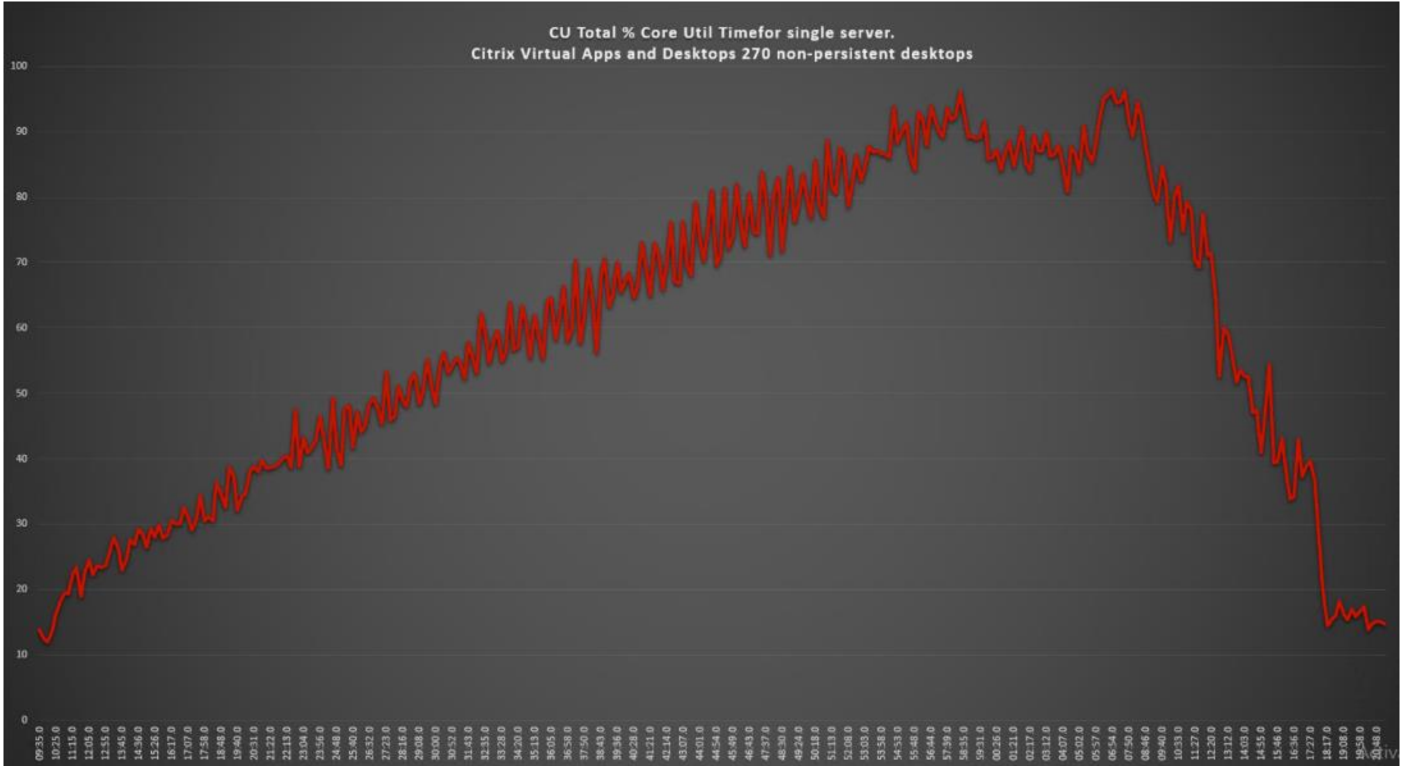

Figure 42. Single Server | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs | Host Memory Utilization

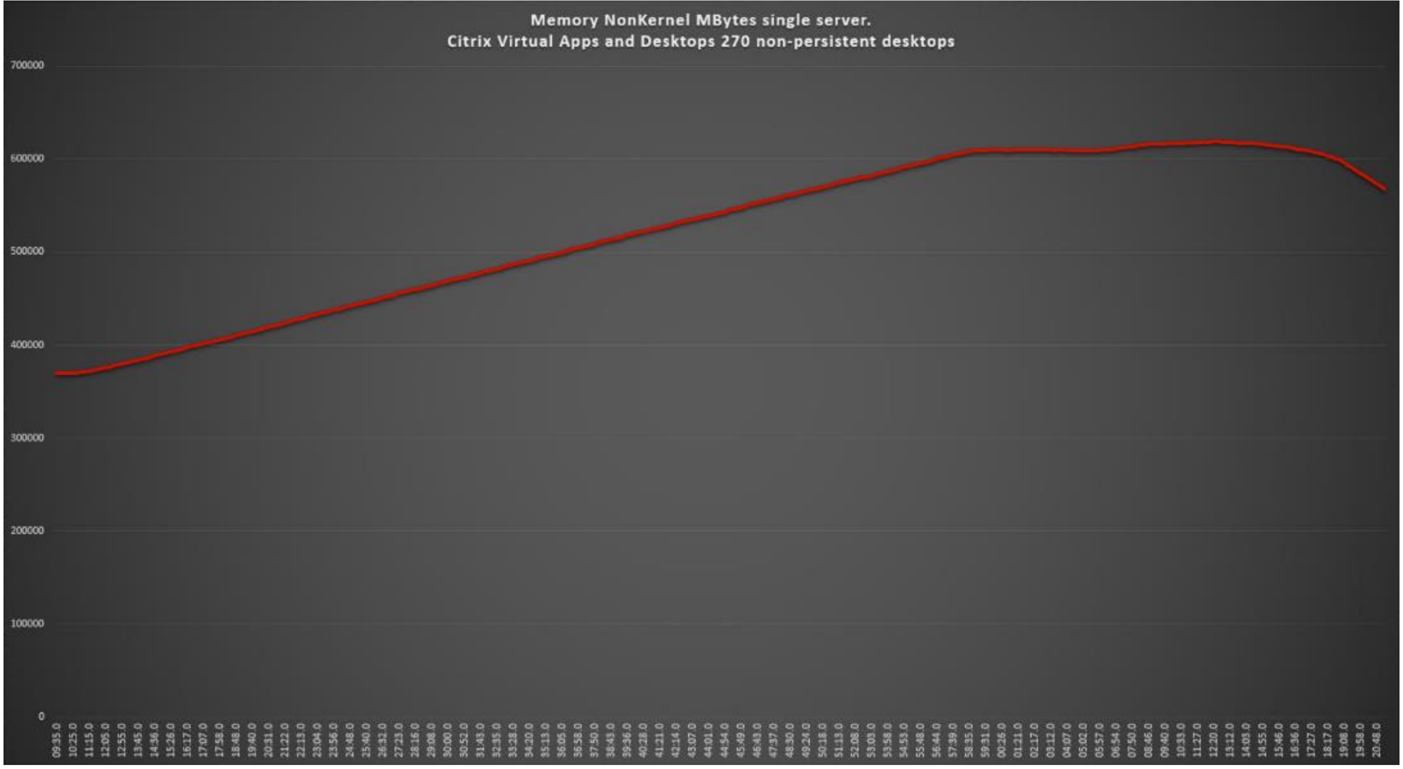

Figure 43. Single Server | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs | Host Network Utilization

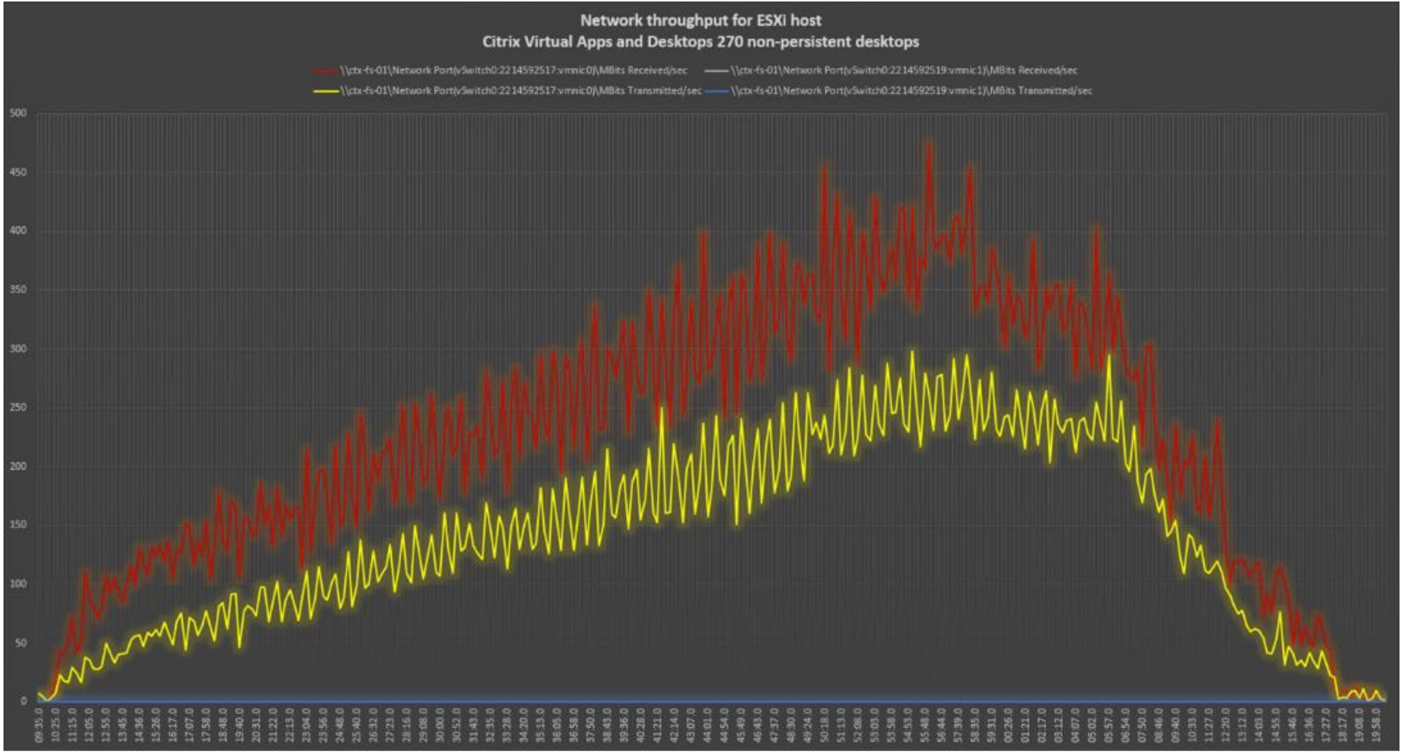

Single-Server Recommended Maximum Workload for PERSISTENT desktops with 270 Users

The recommended maximum workload for a Cisco UCS X210c M6 Compute Node server with dual Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM is 270 Windows 11 64-bit VDI PERSISTENT virtual machines with 2 vCPU and 4GB RAM.

Login VSI performance data is as shown below:

Figure 44. Single Server | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | VSI Score

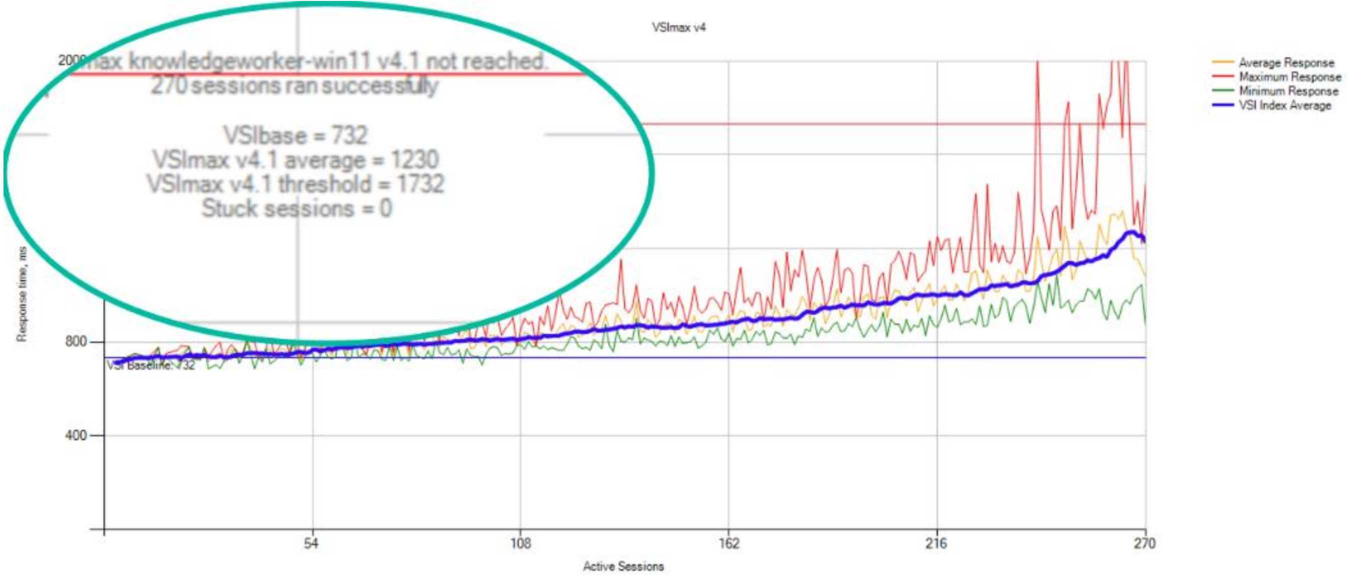

Performance data for the server running the workload is shown below:

Figure 45. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | Host CPU Utilization

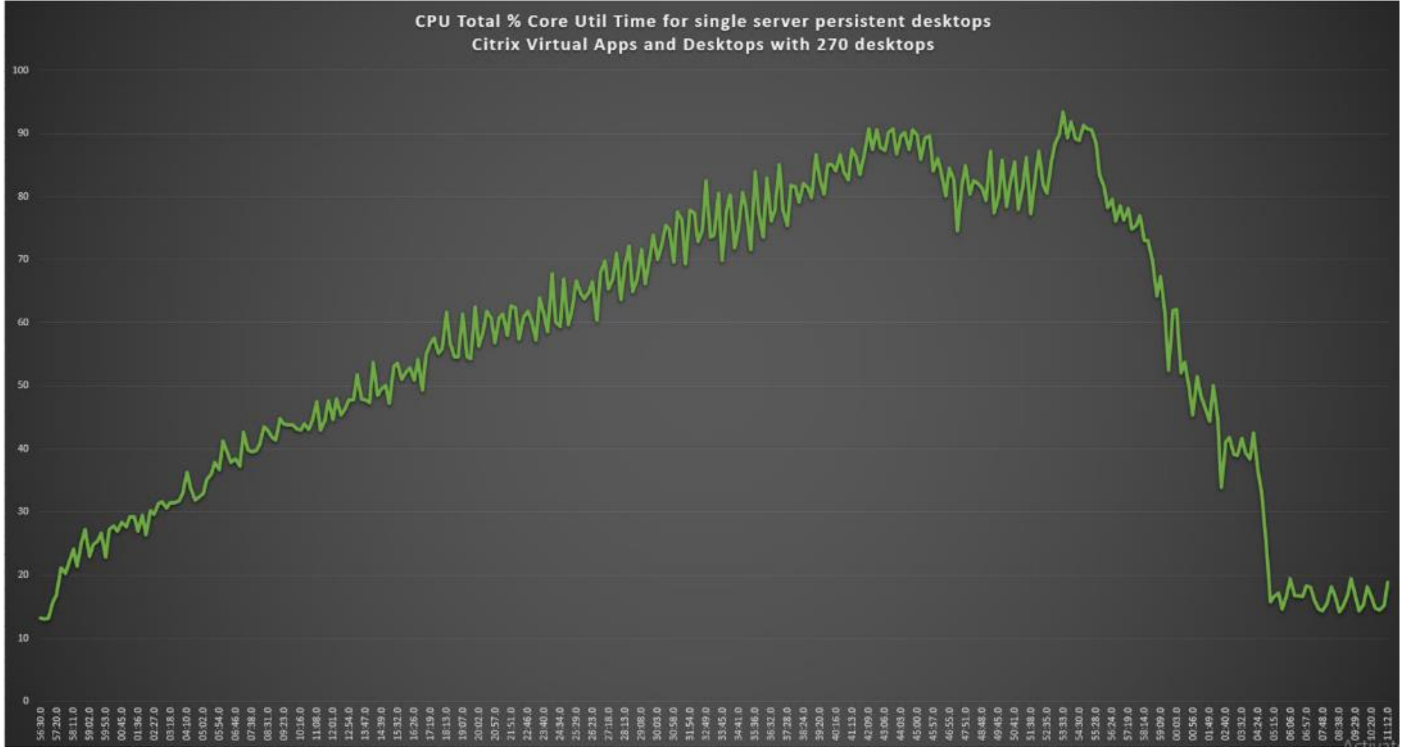

Figure 46. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | Host Memory Utilization

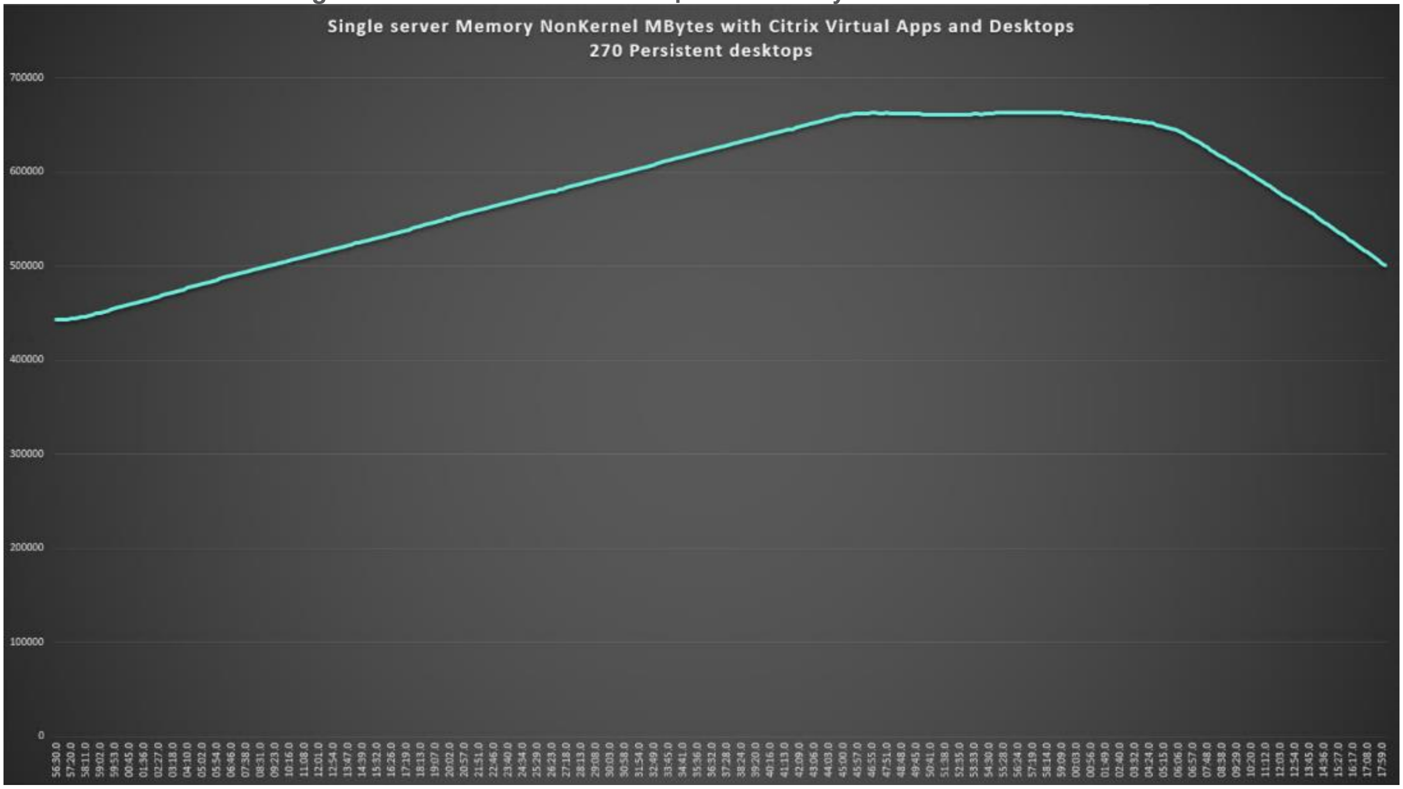

Figure 47. Single Server | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | Host Network Utilization

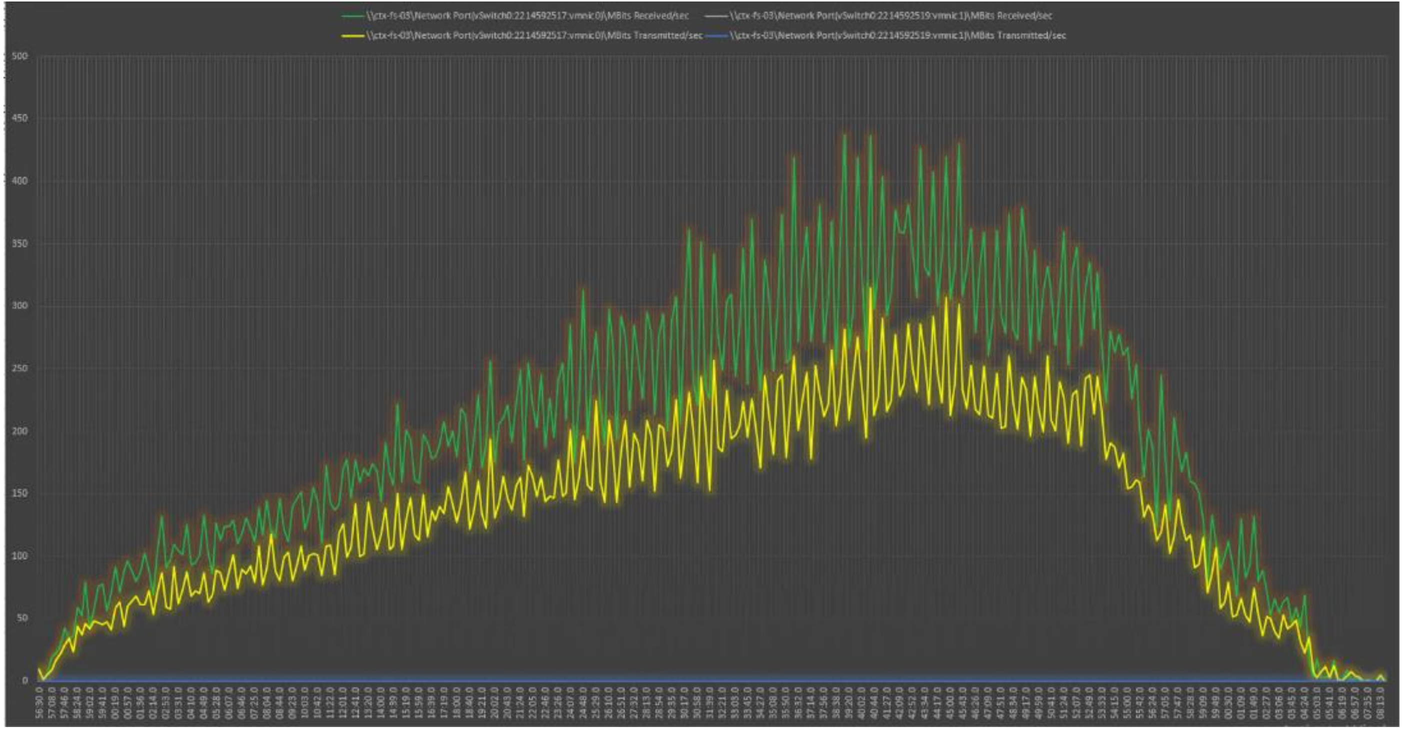

Single-Server Recommended Maximum Workload for RDS Sessions with 335 Users

The recommended maximum workload for a Cisco UCS X210c M6 Compute Node server with dual Intel(R) Xeon(R) Gold 6348 CPU 2.60GHz 28-core processors, 1TB 3200MHz RAM is 335 Windows Server 2022 sessions. The blade server ran 25 Windows Server 2022 Virtual Machines. Each virtual server was configured with 8 vCPUs and 24GB RAM.

LoginVSI data is shown below:

Figure 48. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | VSI Score

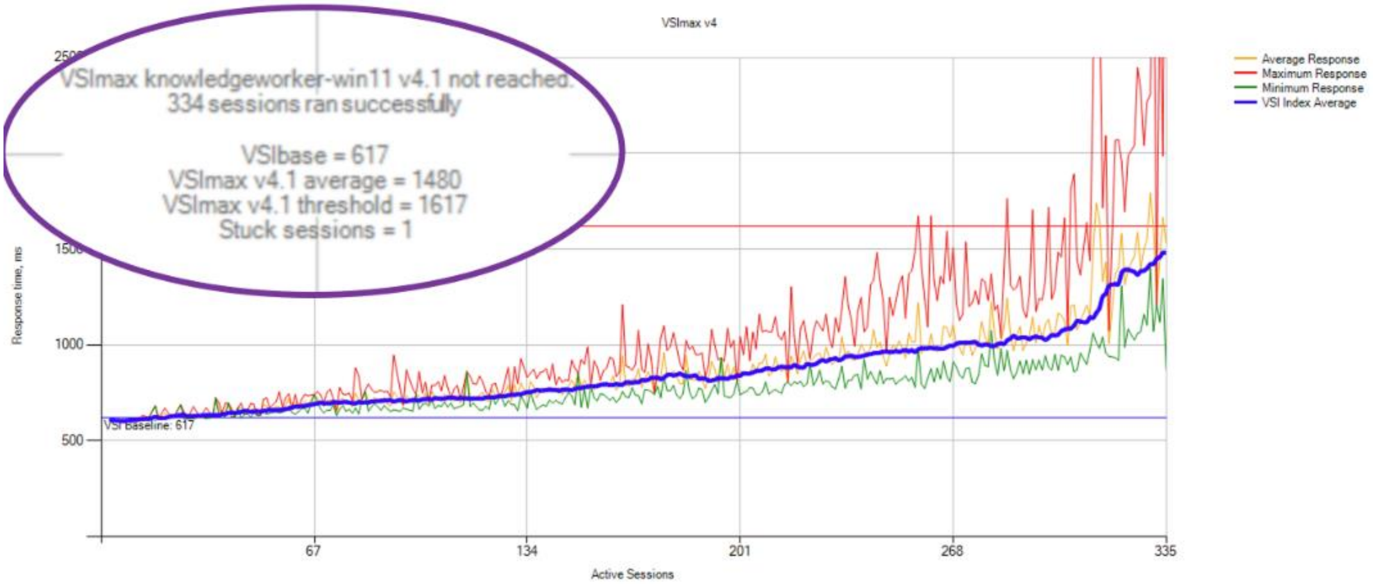

Performance data for the server running the workload is shown below:

Figure 49. Single Server Recommended Maximum Workload Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | Host CPU Utilization

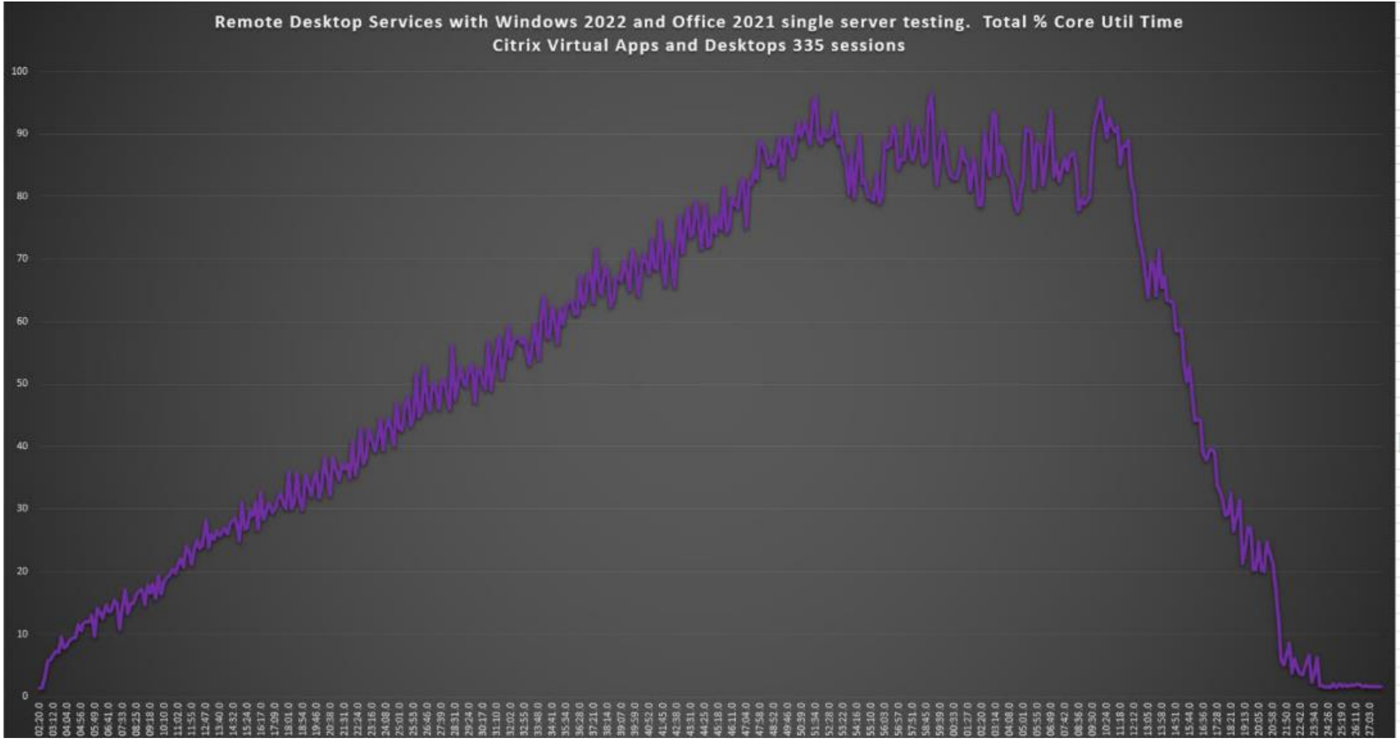

Figure 50. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | Host Memory Utilization

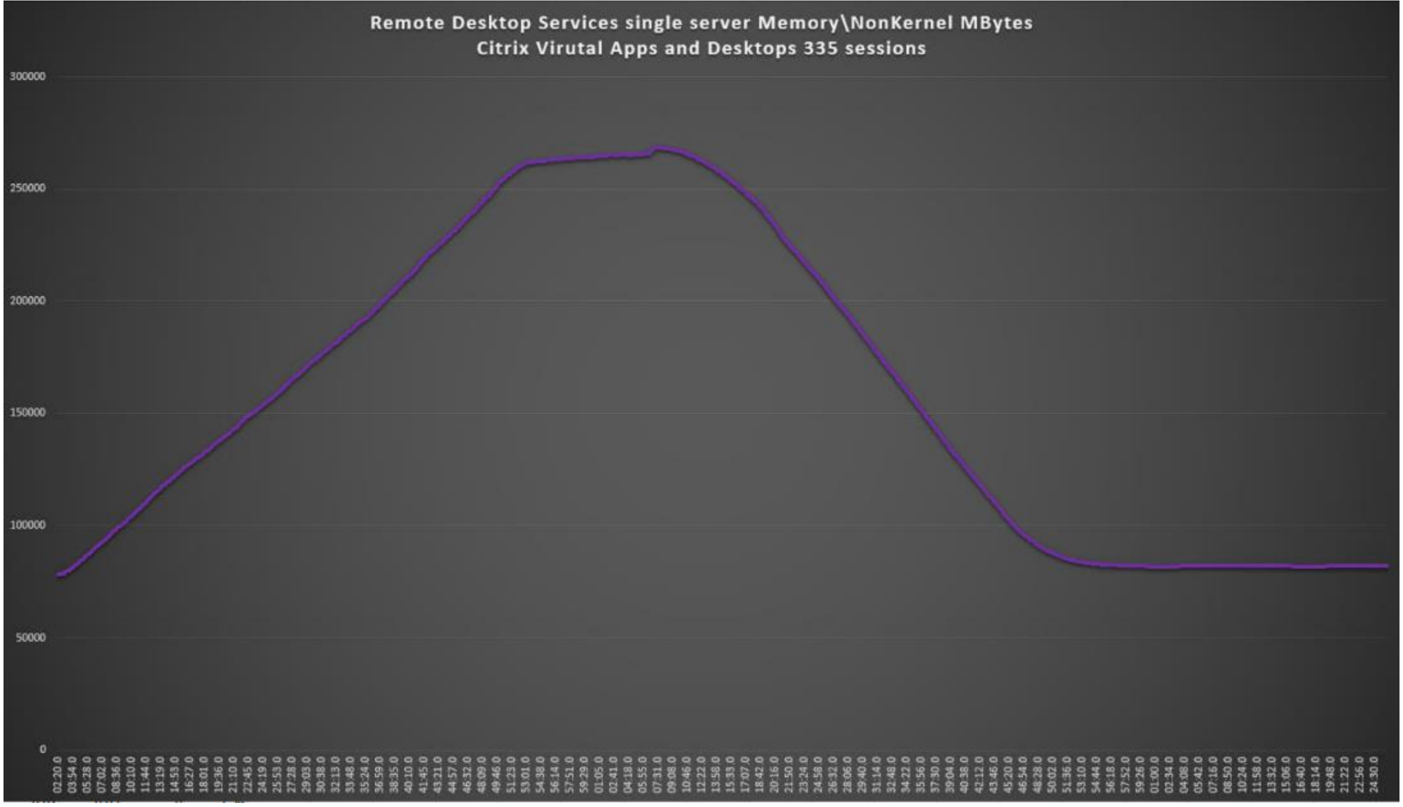

Figure 51. Single Server | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | Host Network Utilization

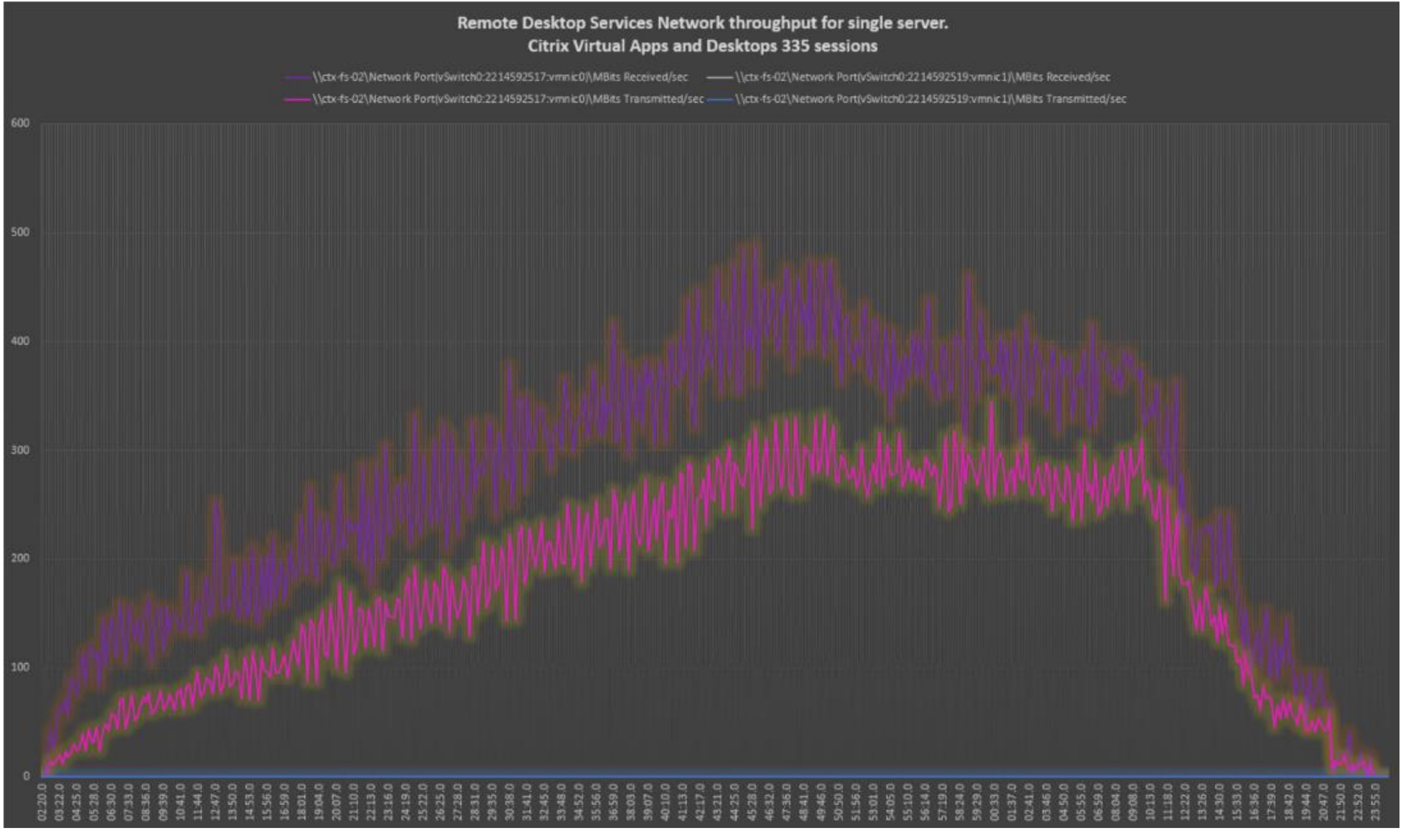

Full Scale Workload Testing

This section describes the key performance metrics that were captured on the Cisco UCS, during the full-scale testing. Full Scale testing was done with following Workloads using eight Cisco UCS X210c M6 Compute Node Servers, configured in a single ESXi Host Pool, and designed to support single Host failure (N+1 Fault tolerance):

- 2000 NON-PERSISTENT Single-session OS sessions (Citrix PVS)
- 2000 PERSISTENT Single-session OS sessions (Citrix MCS)
- 2600 NON-PERSISTENT Multisession OS sessions (RDS)

To achieve the target, sessions were launched against each workload set at a time. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

Full Scale Recommended Maximum Workload Testing for NON-PERSISTENT Single-session OS Machine VDAs with 2000 Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray//X70 R3 array during the full-scale testing with 2000 NON-PERSISTENT Single-session OS machines using 8 blades in a single pool.

The workload for the test is 2000 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 52. Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs| VSI Score

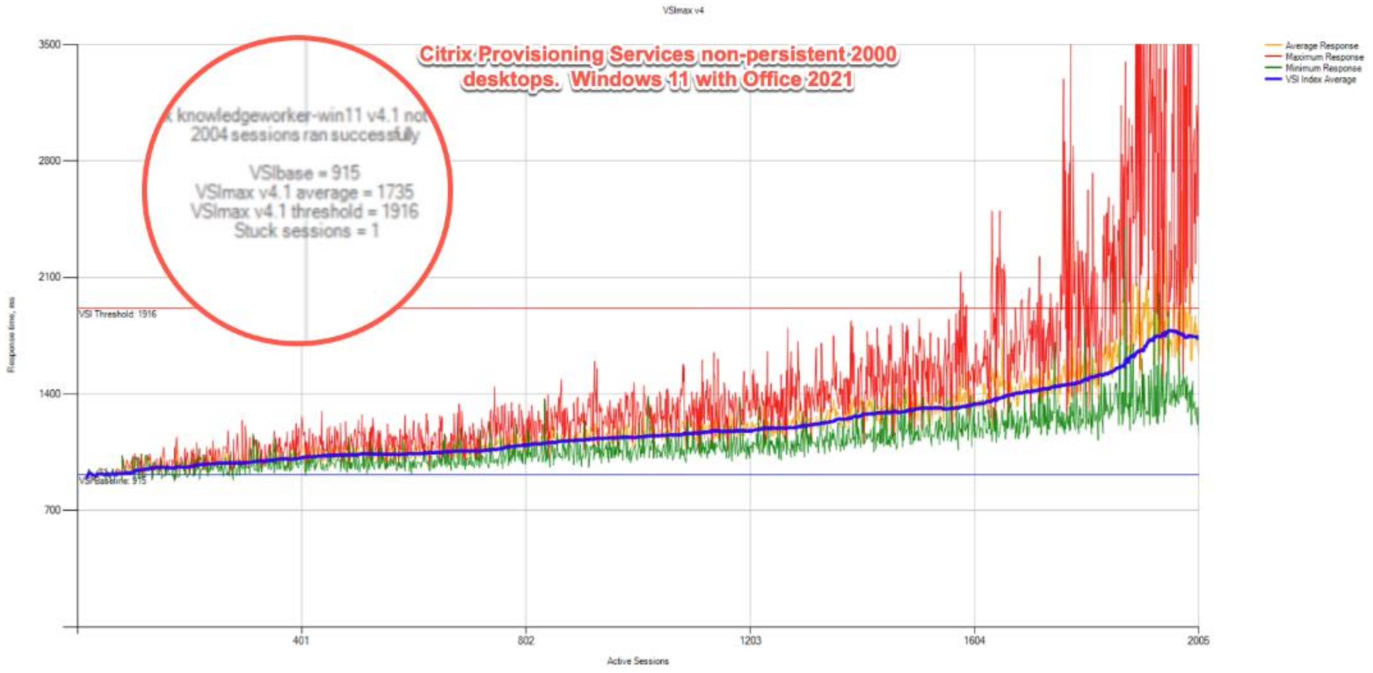

Figure 53. Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs | Test repeatability

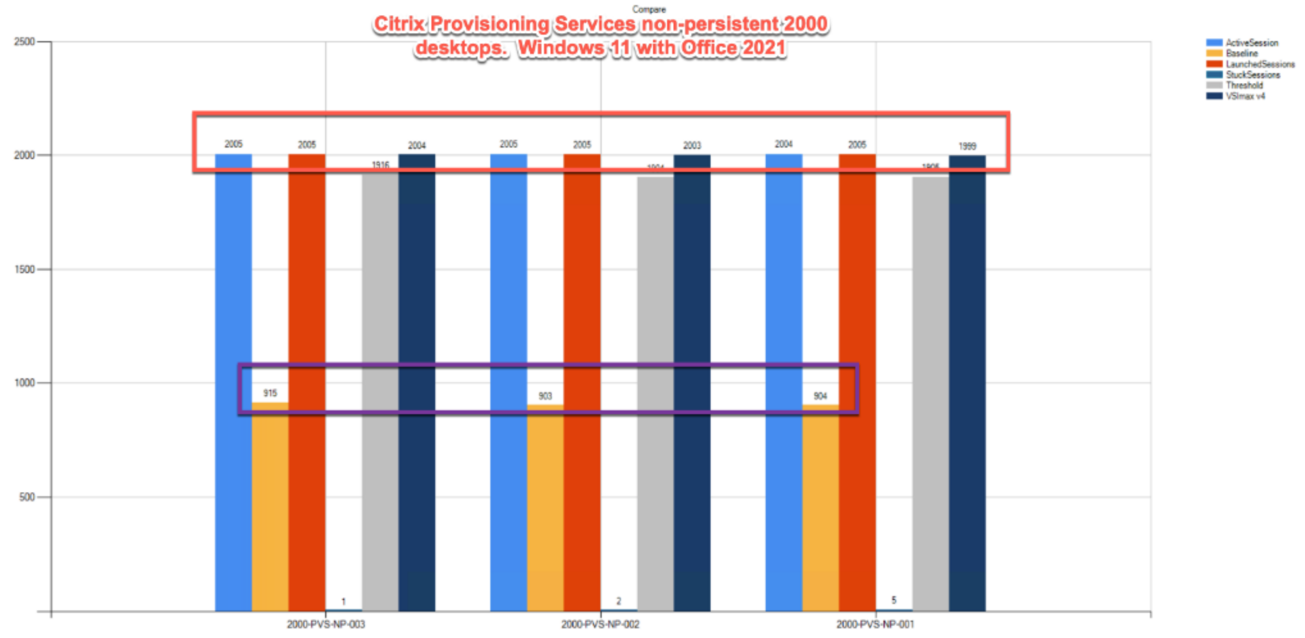

Figure 54. Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs | FlashArray//X70 R3 Performance Chart

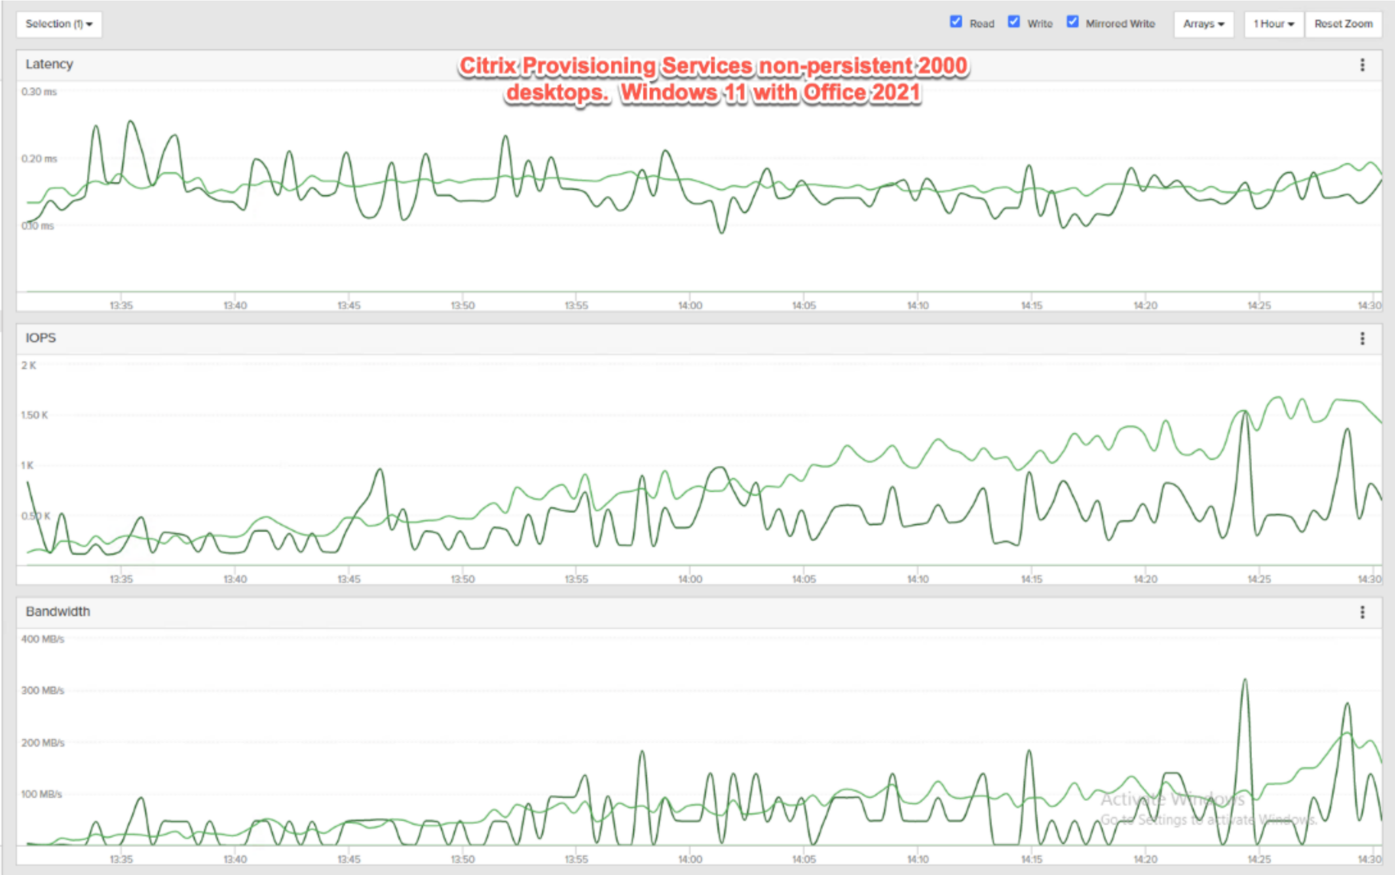

Figure 55. Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Single-session OS machine VDAs | FlashArray//X70 R3 volume data optimization

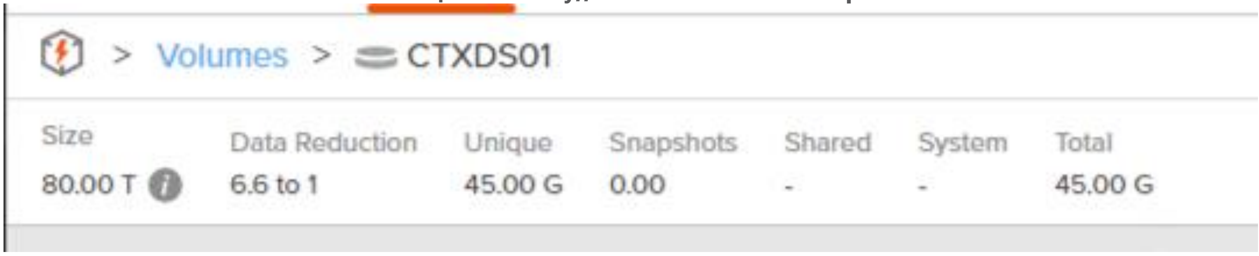

> Volumes > CTXDS01

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
80.00 T	6.6 to 1	45.00 G	0.00	-	-	45.00 G

Full Scale Recommended Maximum Workload Testing for PERSISTENT Single-session OS Machine VDAs with 2000 Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray during the persistent desktop full-scale testing with 2000 PERSISTENT Single-session OS machines using 8 blades in a single pool.

The workload for the test is 2000 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 56. **Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | VSI Score**

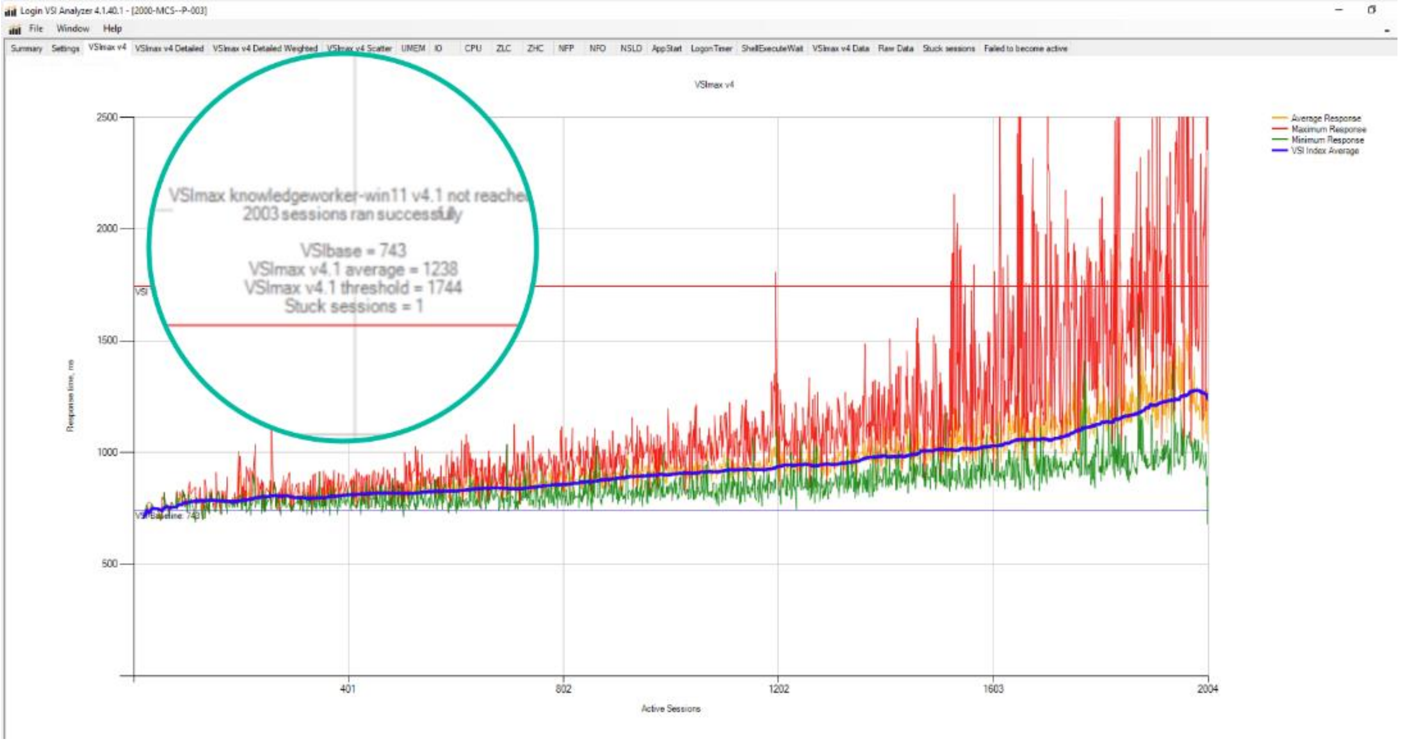

VSImax knowledgeworker-win11 v4.1 not reached
2003 sessions ran successfully

VSIbase = 743
VSImax v4.1 average = 1238
VSImax v4.1 threshold = 1744
Stuck sessions = 1

Figure 57. **Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | Test repeatability**

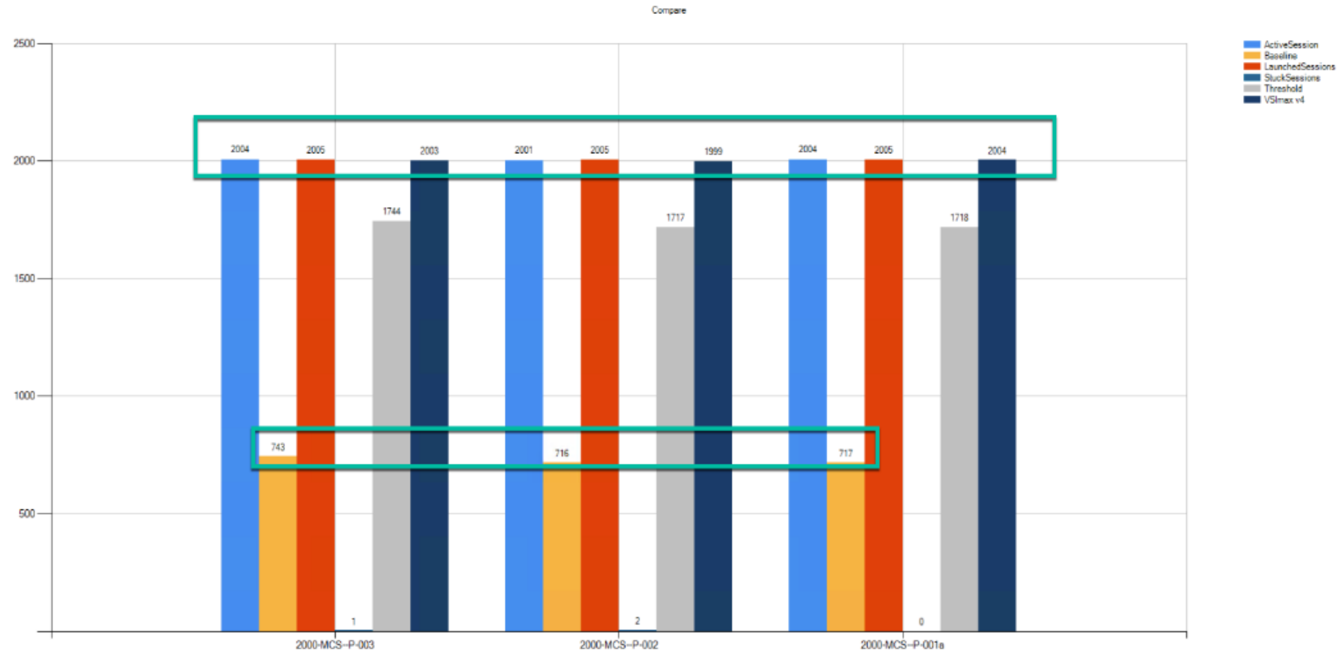

Figure 58. Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | FlashArray//X70 R3 System Performance Chart

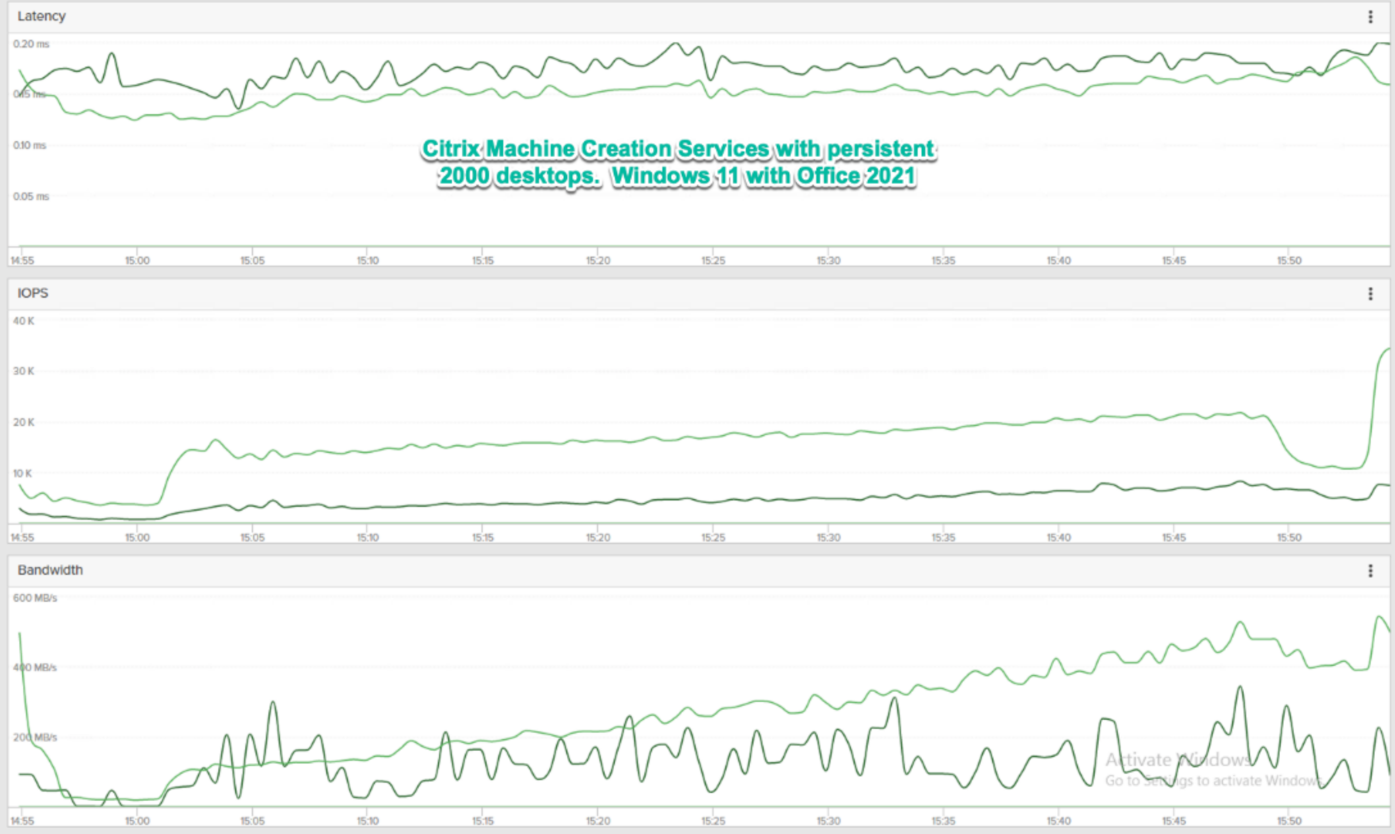

Figure 59. Full Scale | 2000 Users | Citrix Virtual Apps and Desktops 2203 PERSISTENT Single-session OS machine VDAs | FlashArray//X70 R3 volume data optimization

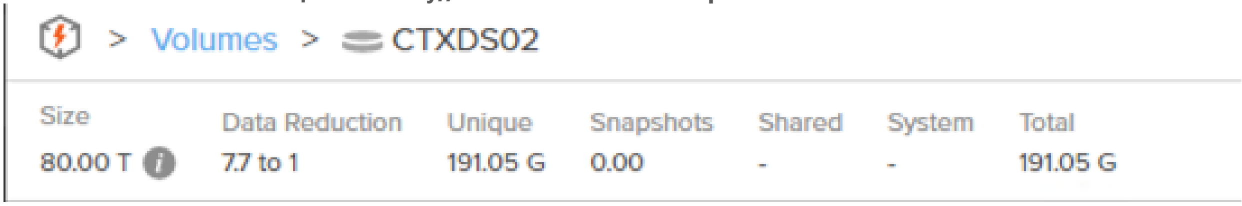

> Volumes > CTXDS02

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
80.00 T	7.7 to 1	191.05 G	0.00	-	-	191.05 G

Full Scale Recommended Maximum Workload for NON-PERSISTENT Multi-session OS Random Sessions with 2600 Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray//X70 R3 array, during the NON-PERSISTENT Multi-session OS full-scale testing with 2600 Desktop Sessions using 8 blades configured in single Host Pool.

The Multi-session OS workload for the solution is 2600 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 60. **Full Scale | 2600 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | VSI Score**

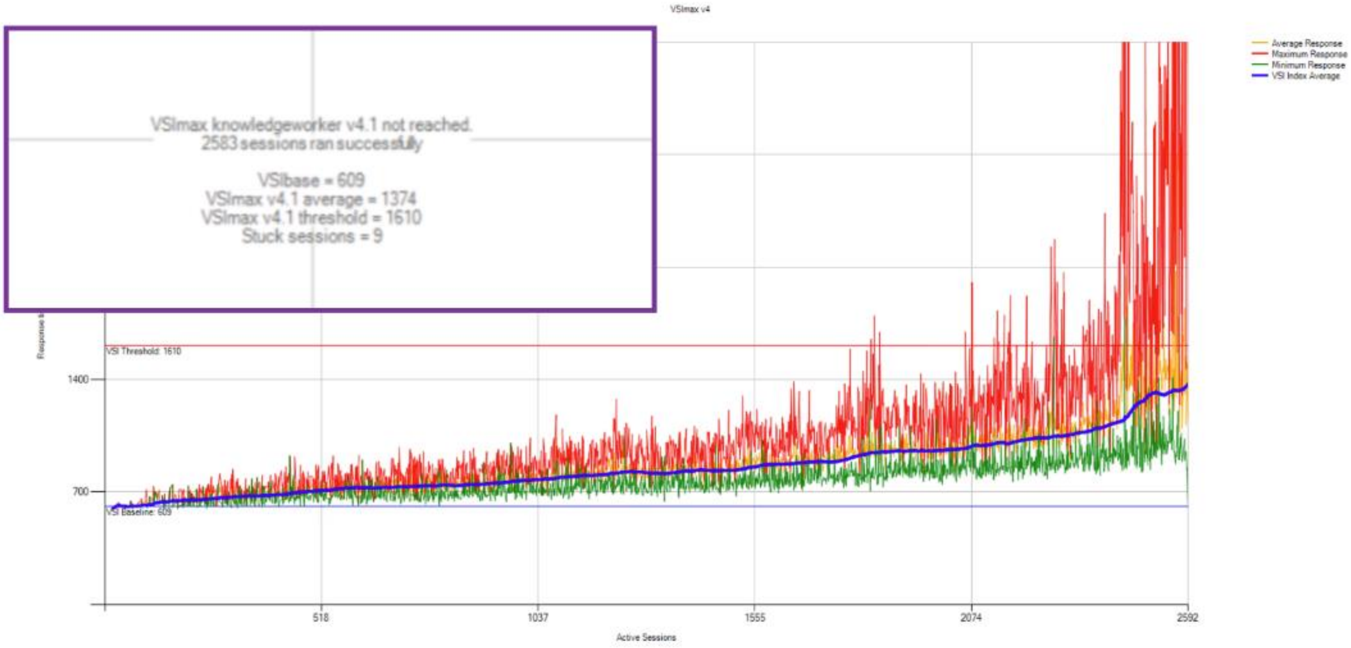

VSImax knowledgeworker v4.1 not reached.
2583 sessions ran successfully

VSIbase = 609
VSImax v4.1 average = 1374
VSImax v4.1 threshold = 1610
Stuck sessions = 9

Figure 61. **Full Scale | 2600 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | Test repeatability**

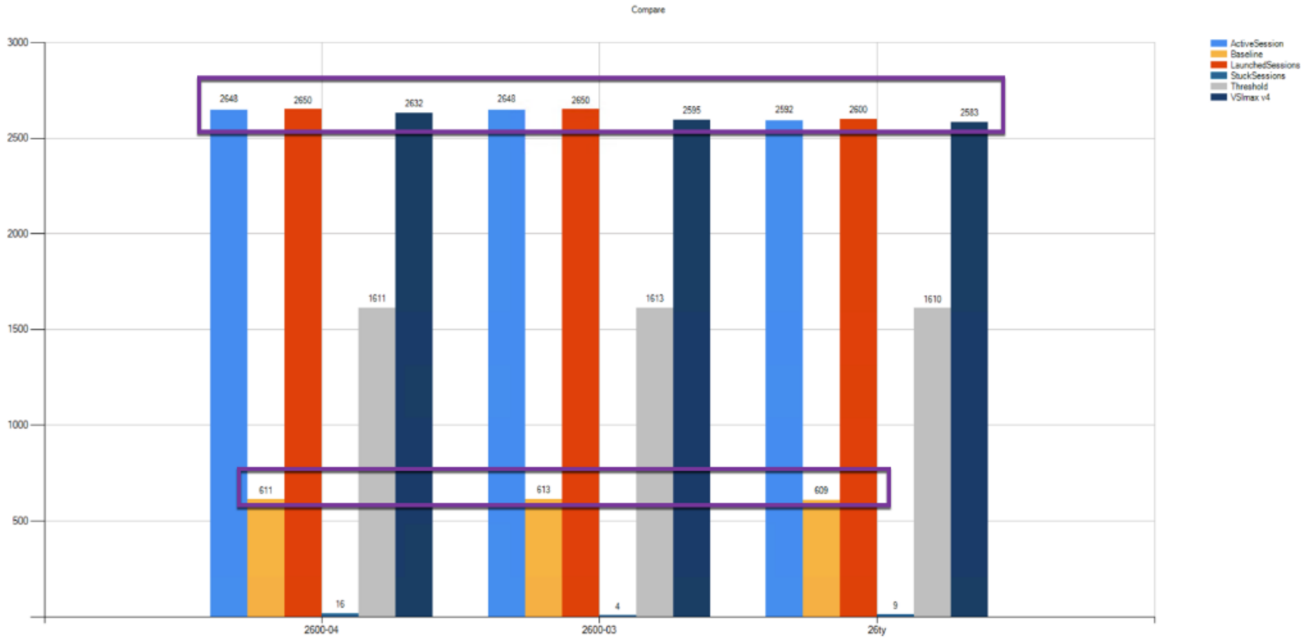

Figure 62. Full Scale | 2600 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | FlashArray//X70 R3 System Performance Chart

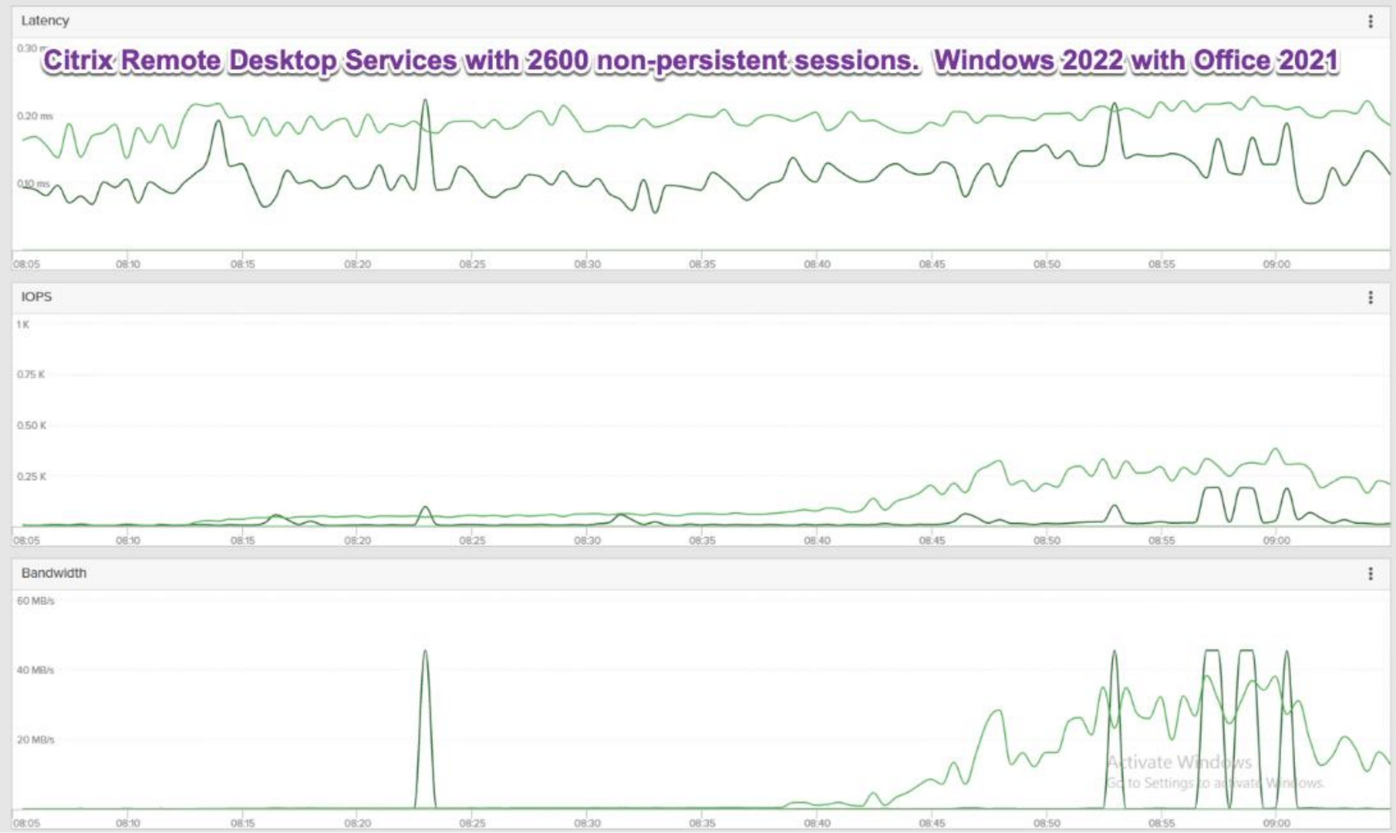

Figure 63. Full Scale | 2600 Users | Citrix Virtual Apps and Desktops 2203 NON-PERSISTENT Multi-session OS machine VDAs | FlashArray//X70 R3 volume data optimization

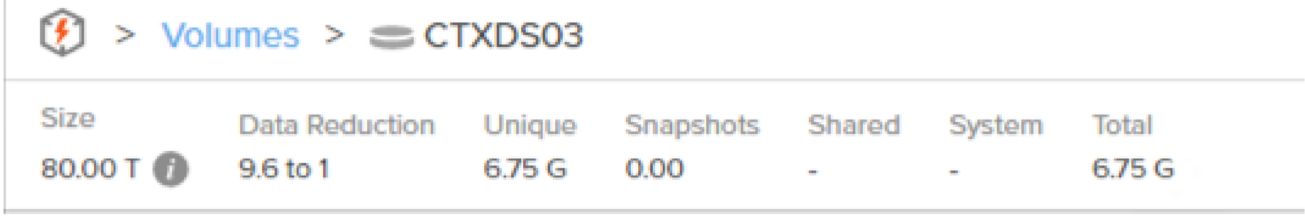
> Volumes > CTXDS03

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
80.00 T	9.6 to 1	6.75 G	0.00	-	-	6.75 G

Full Scale Server Performance Chart with LoginVSI Knowledge Worker Workload Test

This section provides a detailed performance chart for ESXi 8.0 installed on Cisco UCS X210c M6 Compute Node Server as part of the workload test with Citrix Virtual Apps and Desktops 2203 deployed on Pure Storage FlashArray//70 R3 system running LoginVSI v4.1.39 based knowledge worker workload part of the FlashStack reference architecture defined here.

The charts below are defined in the set of 8 hosts in the single performance chart.

Figure 64. Full Scale | 2000 Users| NON-PERSISTENT Single-session OS machine VDAs | Host CPU Utilization

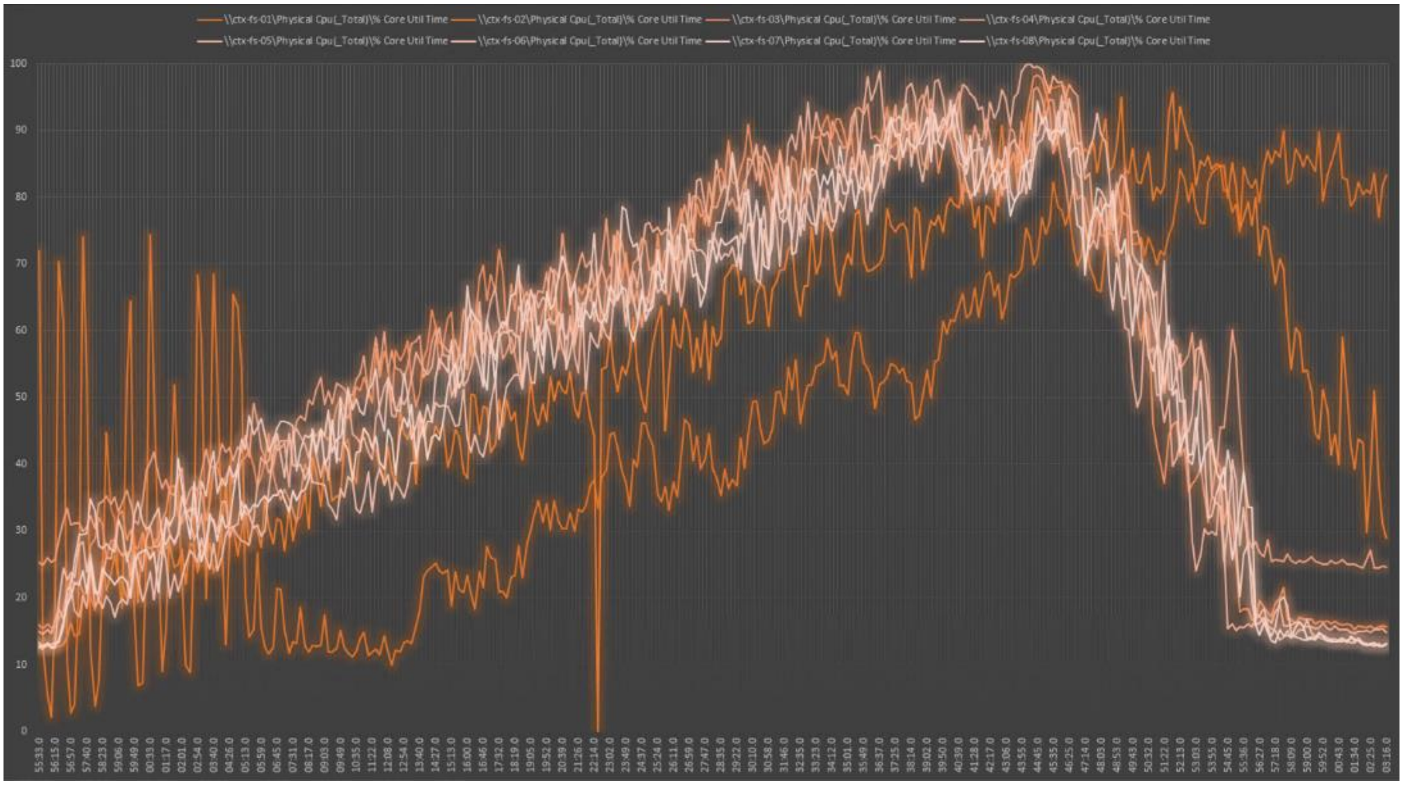

Figure 65. Full Scale | 2000 Users| NON-PERSISTENT Single-session OS machine VDAs | Host Memory Utilization

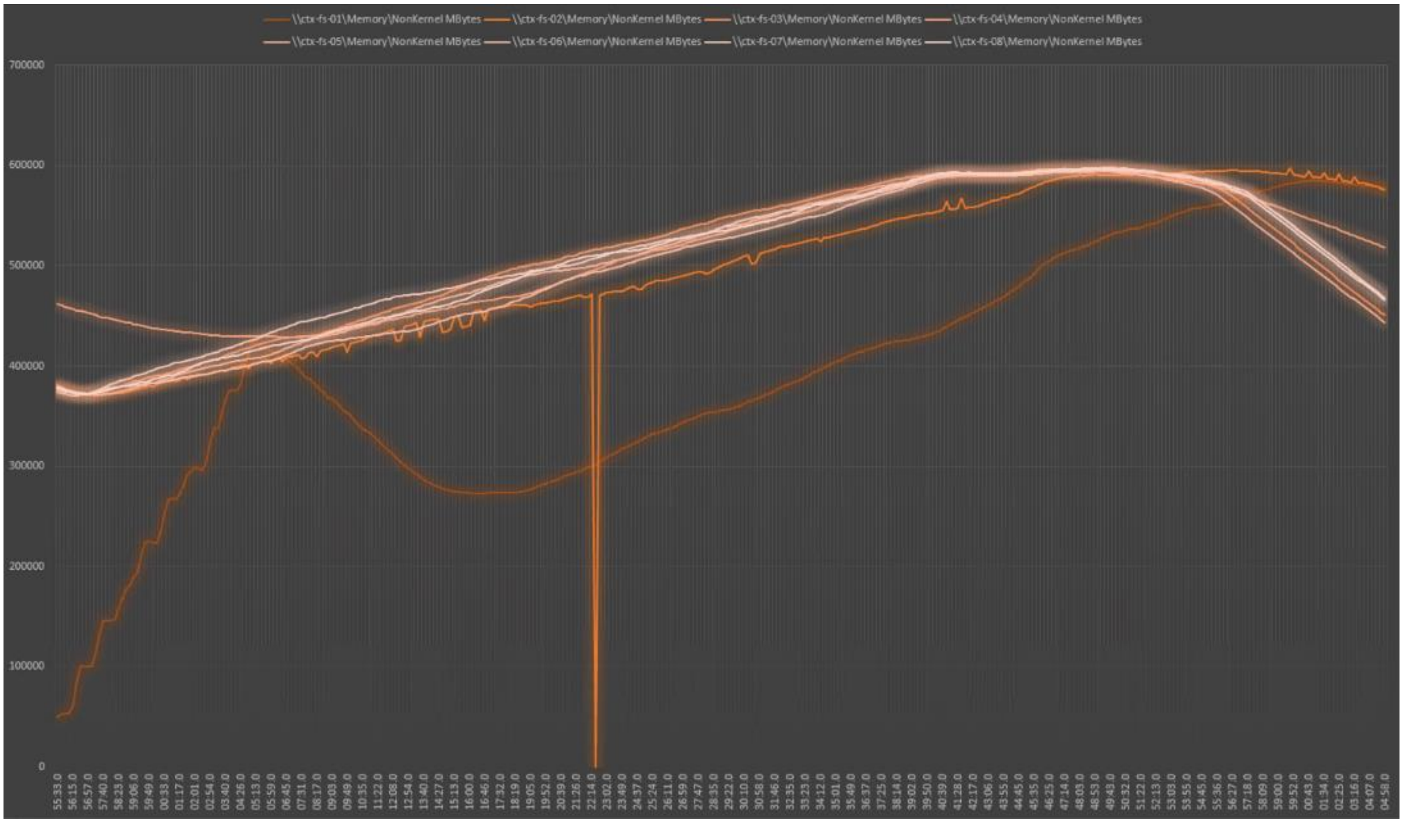

Figure 66. **Full Scale | 2000 Users| NON-PERSISTENT Single-session OS machine VDAs | Host Network Utilization**

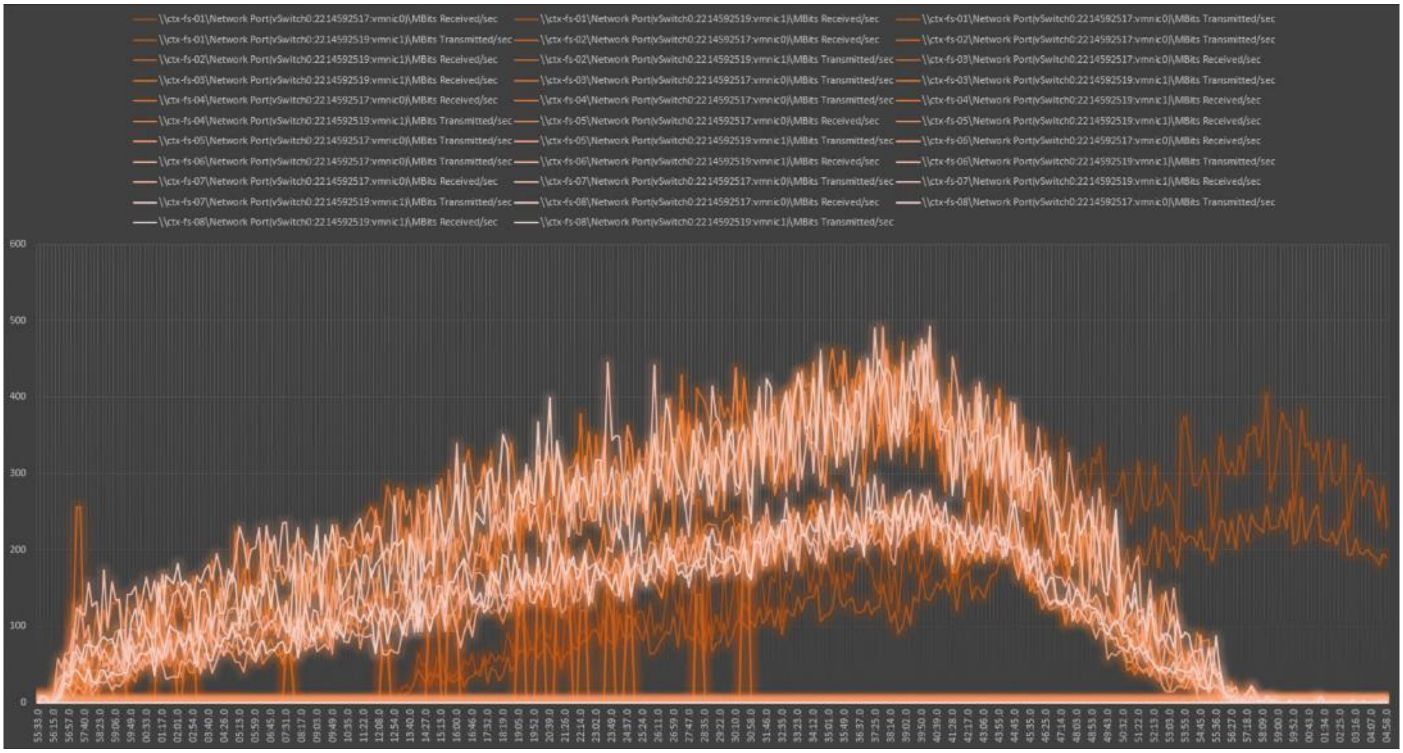

Figure 67. **Full Scale | 2000 Users| PERSISTENT Single-session OS machine VDAs | Host CPU Utilization**

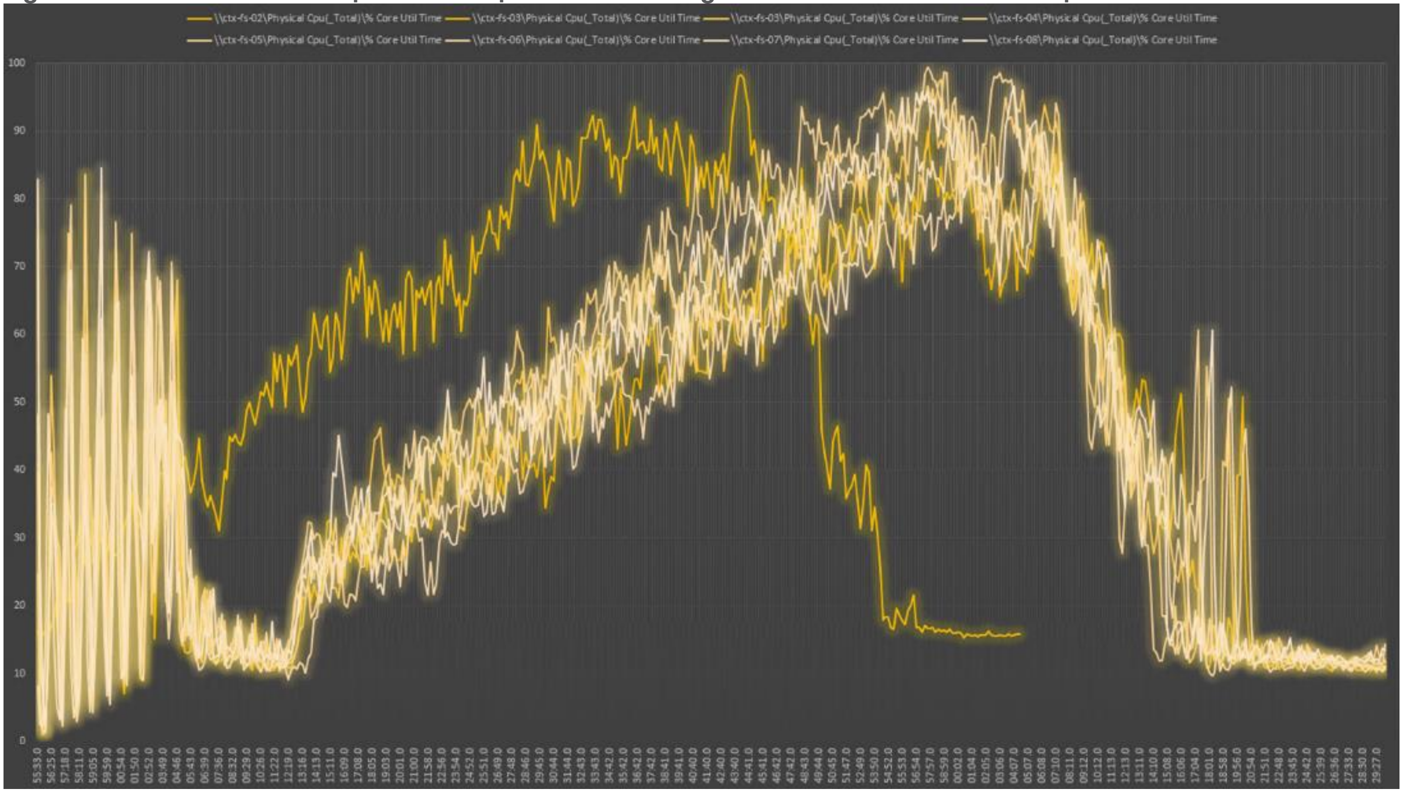

Figure 68. Full Scale | 2000 Users| PERSISTENT Single-session OS machine VDAs | Host Memory Utilization

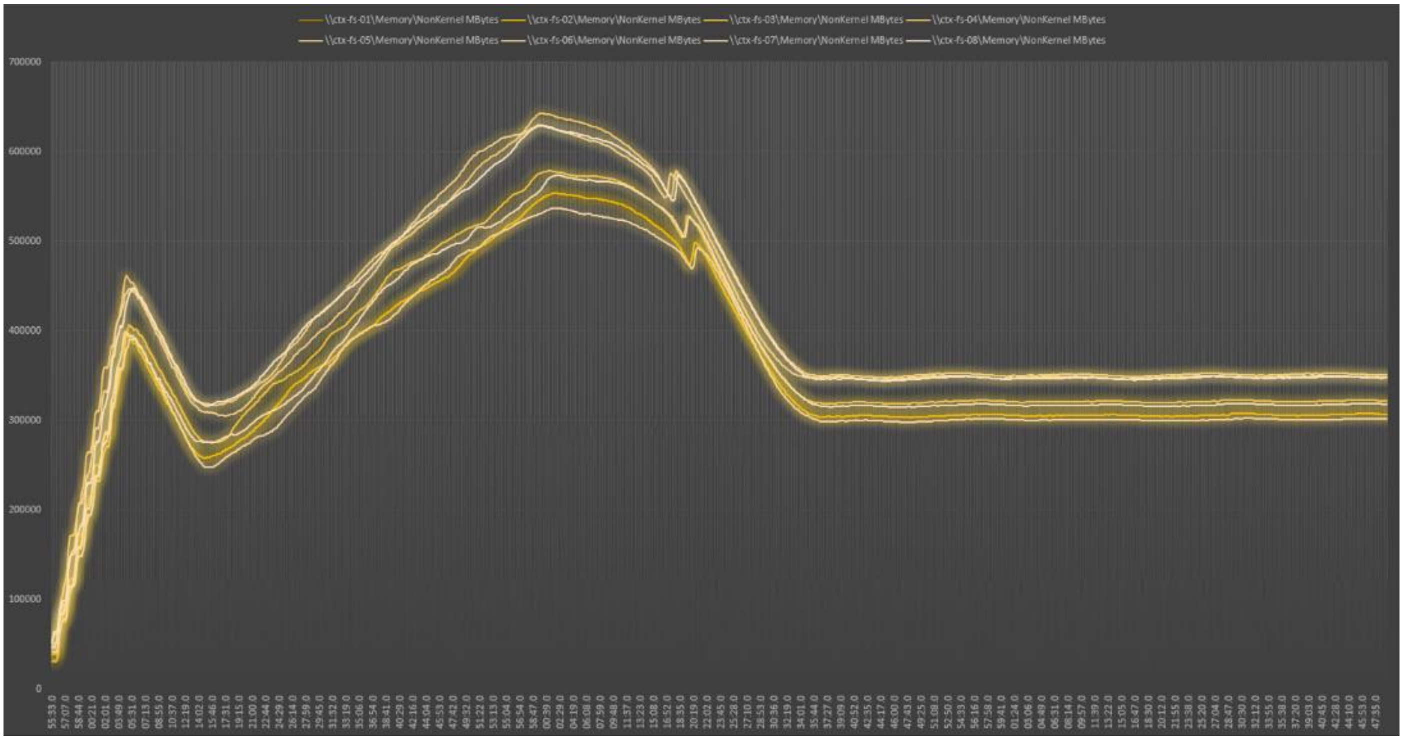

Figure 69. Full Scale | 2000 Users| PERSISTENT Single-session OS machine VDAs | Host Network Utilization

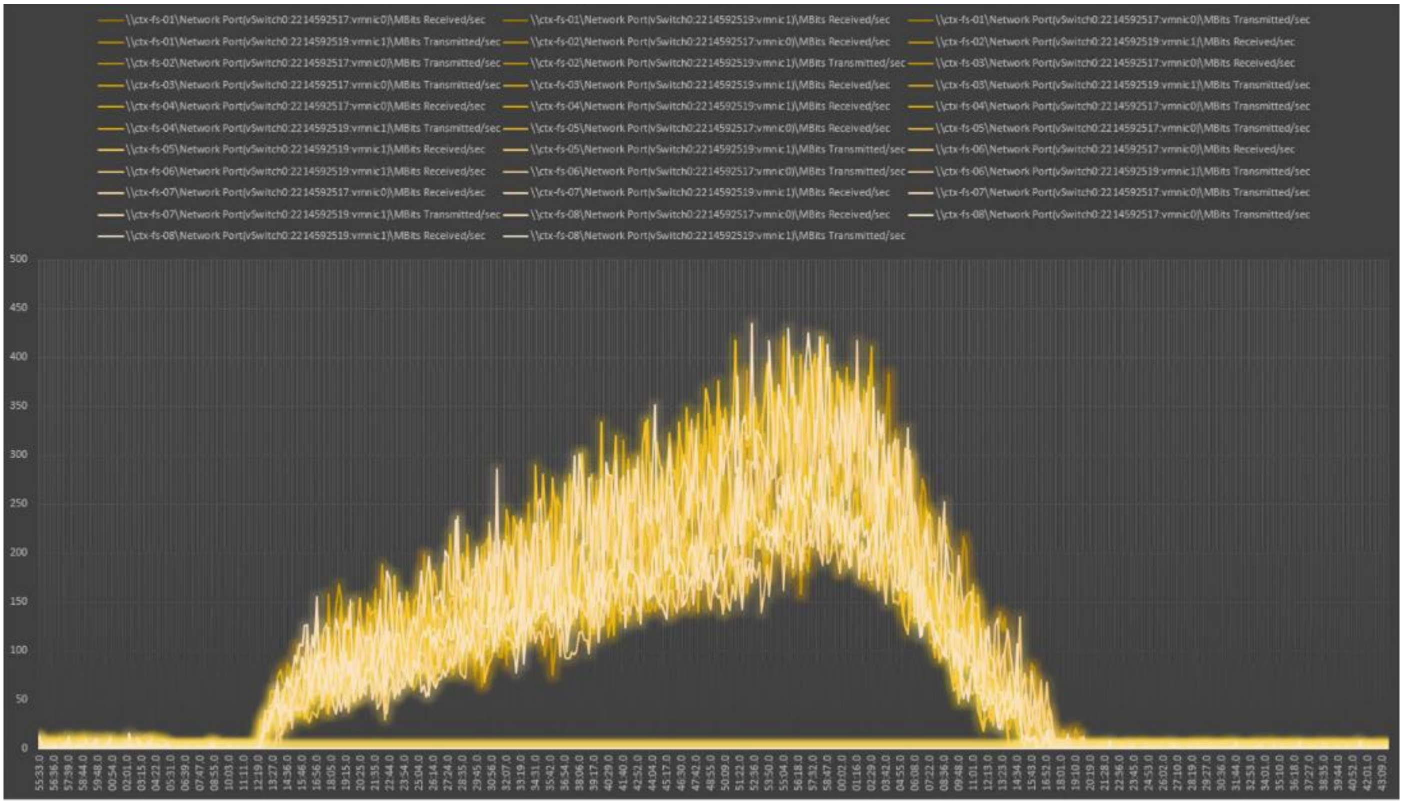

Figure 70. Full Scale | 2600 Users| Multi-session OS machine VDAs | Host CPU Utilization

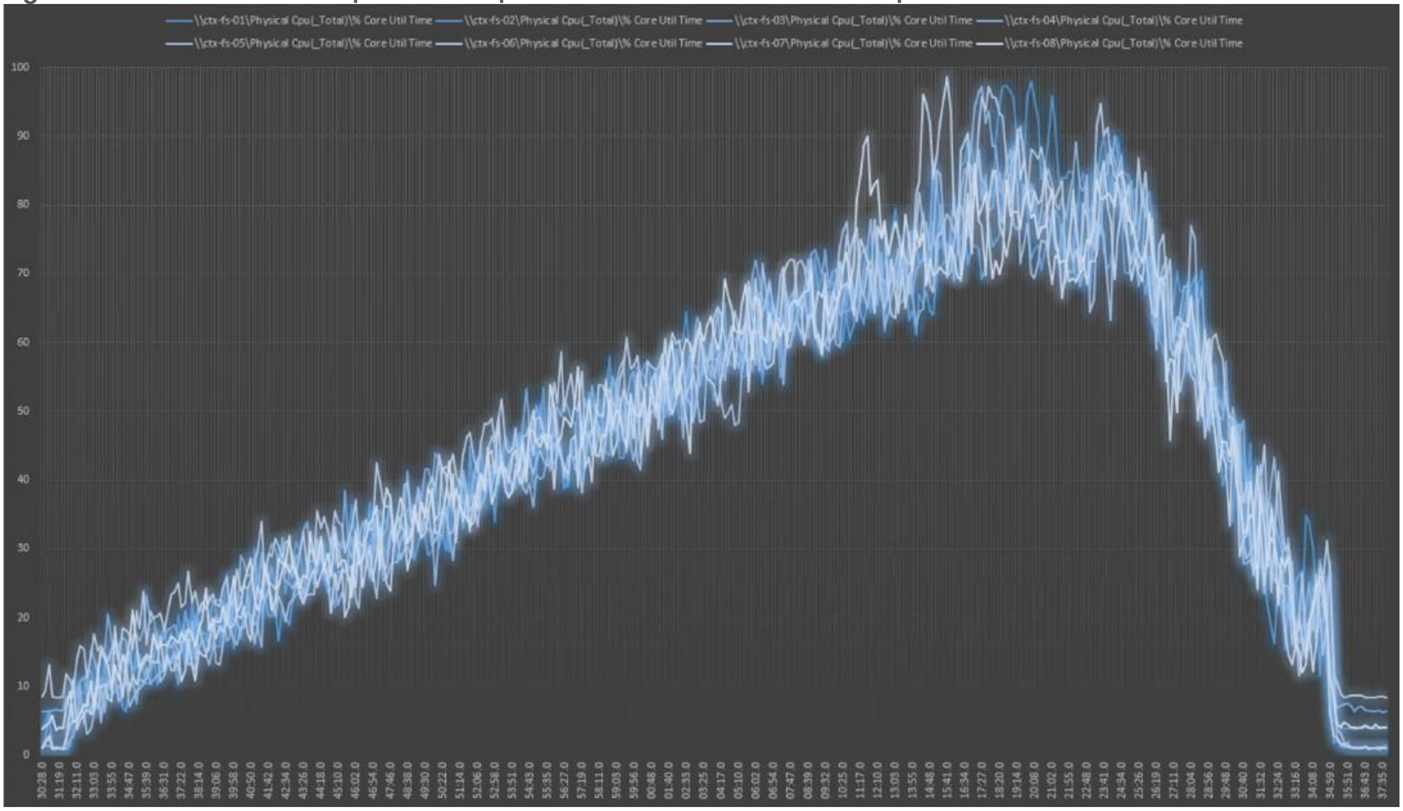

Figure 71. Full Scale | 2600 Users| Multi-session OS machine VDAs | Host Memory Utilization

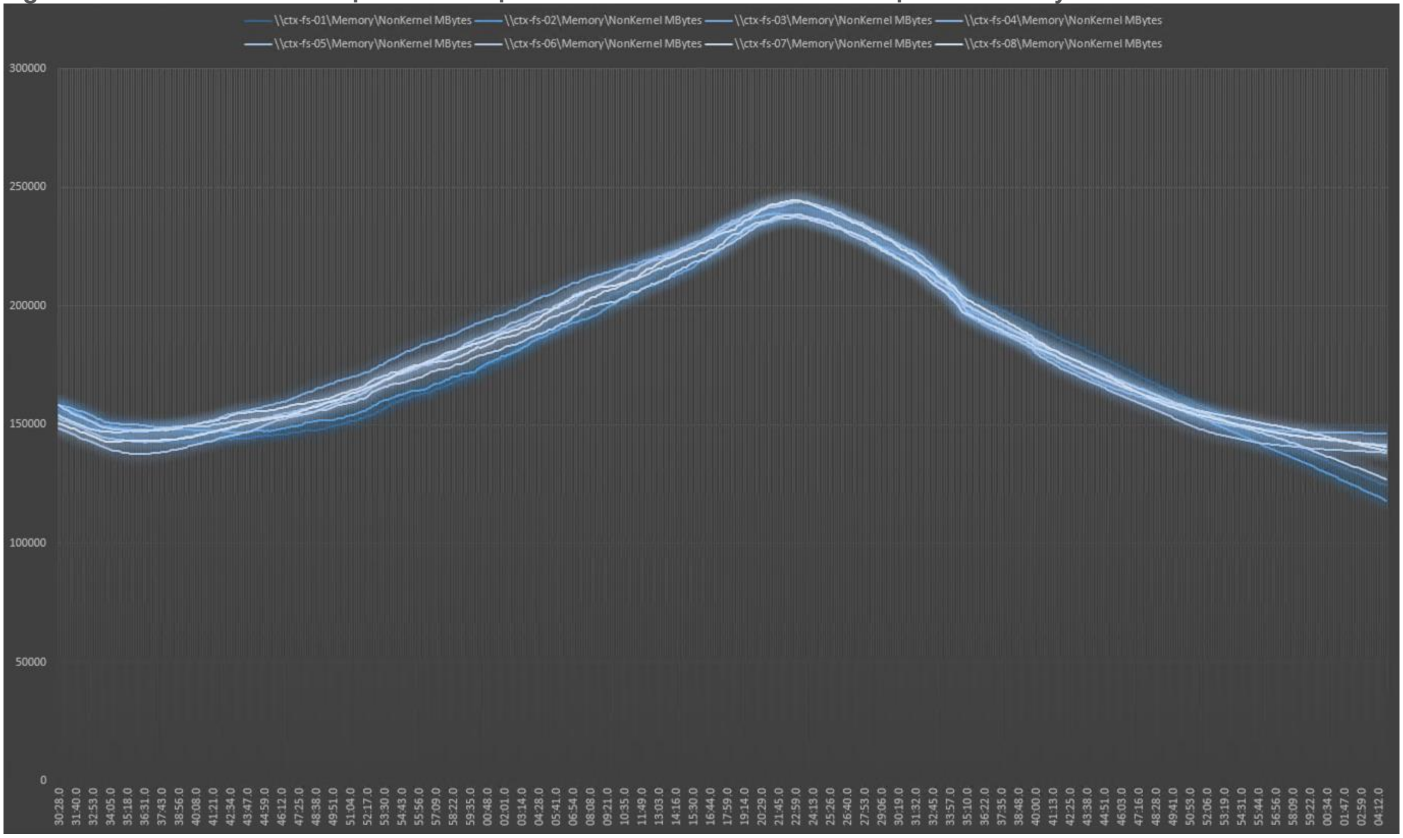

Figure 72. Full Scale | 2600 Users| Multi-session OS machine VDAs | Host Network Utilization

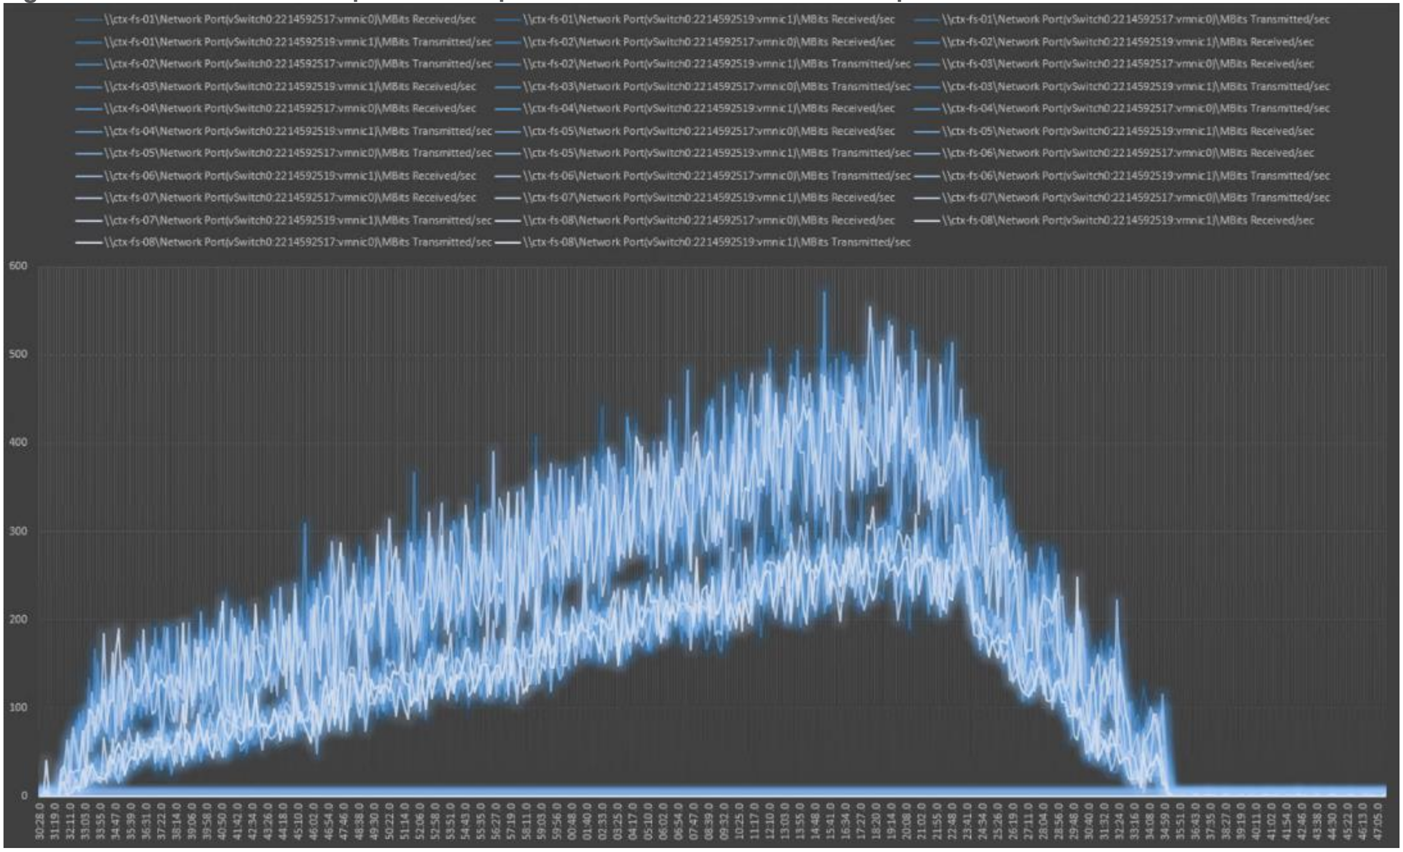

Summary

FlashStack delivers a platform for enterprise end-user computing deployments and cloud data centers using Cisco UCS Blade and Rack Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS 9100 Fibre Channel switches and Pure Storage FlashArray//X70 R3 Storage Array. The introduction of Cisco X-Series modular platform and Cisco Intersight with its services to FlashStack enhances the ability to provide complete visibility and Orchestration across all elements of FlashStack datacenter and modernize the infrastructure and operations of FlashStack datacenter.

FlashStack is designed and validated using compute, network and storage best practices and high availability to reduce deployment time, project risk and IT costs while maintaining scalability and flexibility for addressing a multitude of IT initiatives. This CVD validates the design, performance, management, scalability, and resilience that FlashStack provides to customers wishing to deploy enterprise-class VDI.

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, Pure Storage FlashArray//X70 R3 storage and our certified partners can help. We collaborate with you to enhance your IT capabilities through a full 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 and Pure Storage Support provide in-depth knowledge transfer and education services that give you access to our global technical resources and intellectual property.

About the Author

Jeff Nichols–Leader, Technical Marketing, CSPG UCS Solutions - US

Jeff Nichols is a member of the Cisco's Computing Systems Product Group team focusing on design, testing, solutions validation, technical content creation, and performance testing/benchmarking. He has years of experience in Virtual Desktop Infrastructure (VDI), Server and Desktop Virtualization using Microsoft and VMware products.

Jeff is a subject matter expert on Desktop/Server virtualization, Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA/AMD 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:

- Sreenivasa Edula, Technical Marketing Engineer, Cisco Systems, Inc.
- John George, Technical Marketing Engineer, Cisco Systems, Inc.
- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Joe Houghes, Senior Solutions Architect, Pure Storage, Inc.
- Craig Waters, Technical Director, Pure Storage, Inc.

Appendices

This appendix is organized into the following sections:

Appendix A – References used in this 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-m6-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/

- Pure Storage Reference Documents

  https://www.flashstack.com/

  https://www.purestorage.com/content/dam/purestorage/pdf/datasheets/ps_ds_flasharray_03.pdf

  https://www.purestorage.com

  https://www.purestorage.com/products/evergreen-subscriptions.html

  https://www.purestorage.com/solutions/infrastructure/vdi.html

  https://www.purestorage.com/solutions/infrastructure/vdi-calculator.html

  https://support.purestorage.com/FlashArray/PurityFA/FlashArray_File_Services/001_Getting_Started/001_FA_File_Services_Quick_Start_Guide

  https://support.purestorage.com/FlashArray/PurityFA/FlashArray_File_Services/001_Getting_Started/002_FA_File_Services_Requirements_and_Best_Practices

Appendix B – Glossary

This glossary addresses some terms used in this document, for the purposes of aiding understanding. This is not a complete list of all multicloud terminology. Some Cisco product links are supplied here also, where considered useful for the purposes of clarity, but this is by no means intended to be a complete list of all applicable Cisco products.

aaS/XaaS **(IT capability provided as a Service)**	Some IT capability, X, provided as a service (XaaS). Some benefits are: • The provider manages the design, implementation, deployment, upgrades, resiliency, scalability, and overall delivery of the service and the infrastructure that supports it. • There are very low barriers to entry, so that services can be quickly adopted and dropped in response to business demand, without the penalty of inefficiently utilized CapEx. • The service charge is an IT OpEx cost (pay-as-you-go), whereas the CapEx and the service infrastructure is the responsibility of the provider. • Costs are commensurate to usage and hence more easily controlled with respect to business demand and outcomes. Such services are typically implemented as “microservices,” which are accessed via REST APIs. This architectural style supports composition of service components into systems. Access to and management of aaS assets is via a web GUI and/or APIs, such that Infrastructure-as-code (IaC) techniques can be used for automation, for example, Ansible and Terraform. The provider can be any entity capable of implementing an aaS “cloud-native” architecture. The cloud-native architecture concept is well-documented and supported by open-source software and a rich ecosystem of services such as training and consultancy. The provider

	can be an internal IT department or any of many third-party companies using and supporting the same open-source platforms. Service access control, integrated with corporate IAM, can be mapped to specific users and business activities, enabling consistent policy controls across services, wherever they are delivered from.
Ansible	An infrastructure automation tool, used to implement processes for instantiating and configuring IT service components, such as VMs on an IaaS platform. Supports the consistent execution of processes defined in YAML “playbooks” at scale, across multiple targets. Because the Ansible artefacts (playbooks) are text-based, they can be stored in a Source Code Management (SCM) system, such as GitHub. This allows for software development like processes to be applied to infrastructure automation, such as, Infrastructure-as-code (see IaC below). https://www.ansible.com
AWS **(Amazon Web Services)**	Provider of IaaS and PaaS. https://aws.amazon.com
Azure	Microsoft IaaS and PaaS. https://azure.microsoft.com/en-gb/
Co-located data center	“A colocation center (CoLo)...is a type of data center where equipment, space, and bandwidth are available for rental to retail customers. Colocation facilities provide space, power, cooling, and physical security for the server, storage, and networking equipment of other firms and also connect them to a variety of telecommunications and network service providers with a minimum of cost and complexity.” https://en.wikipedia.org/wiki/Colocation_centre

Containers **(Docker)**	A (Docker) container is a means to create a package of code for an application and its dependencies, such that the application can run on different platforms which support the Docker environment. In the context of aaS, microservices are typically packaged within Linux containers orchestrated by Kubernetes (K8s). https://www.docker.com https://www.cisco.com/c/en/us/products/cloud-systems-management/containerplatform/index.html
DevOps	The underlying principle of DevOps is that the application development and operations teams should work closely together, ideally within the context of a toolchain that automates the stages of development, test, deployment, monitoring, and issue handling. DevOps is closely aligned with IaC, continuous integration and deployment (CI/CD), and Agile software development practices. https://en.wikipedia.org/wiki/DevOps https://en.wikipedia.org/wiki/CI/CD
Edge compute	Edge compute is the idea that it can be more efficient to process data at the edge of a network, close to the endpoints that originate that data, or to provide virtualized access services, such as at the network edge. This could be for reasons related to low latency response, reduction of the amount of unprocessed data being transported, efficiency of resource utilization, and so on. The generic label for this is Multi-access Edge Computing (MEC), or Mobile Edge Computing for mobile networks specifically. From an application experience perspective, it is important to be able to utilize, at the edge, the same operations model, processes, and tools used for any other compute node in the system. https://en.wikipedia.org/wiki/Mobile_edge_computing
IaaS **(Infrastructure as-a-Service)**	Infrastructure components provided aaS, located in data centers operated by a provider, typically accessed over the public Internet. IaaS provides a base platform for the deployment of workloads, typically with containers and Kubernetes (K8s).
IaC **(Infrastructure as-Code)**	Given the ability to automate aaS via APIs, the implementation of the automation is typically via Python code, Ansible playbooks, and similar. These automation artefacts are programming code that define how the services are consumed. As such, they can be subject to the same code management and software development regimes as any other body of code. This means that infrastructure automation can be subject to all of the quality and consistency benefits, CI/CD, traceability, automated testing, compliance checking, and so on, that could be applied to any coding project. https://en.wikipedia.org/wiki/Infrastructure_as_code
IAM **(Identity and Access Management)**	IAM is the means to control access to IT resources so that only those explicitly authorized to access given resources can do so. IAM is an essential foundation to a secure multicloud environment. https://en.wikipedia.org/wiki/Identity_management
IBM **(Cloud)**	IBM IaaS and PaaS. https://www.ibm.com/cloud
Intersight	Cisco Intersight is a Software-as-a-Service (SaaS) infrastructure lifecycle management platform that delivers simplified configuration, deployment, maintenance, and support.

	https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html
GCP **(Google Cloud Platform)**	Google IaaS and PaaS. https://cloud.google.com/gcp
Kubernetes **(K8s)**	Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. https://kubernetes.io
Microservices	A microservices architecture is characterized by processes implementing fine-grained services, typically exposed via REST APIs and which can be composed into systems. The processes are often container-based, and the instantiation of the services often managed with Kubernetes. Microservices managed in this way are intrinsically well suited for deployment into IaaS environments, and as such, are the basis of a cloud native architecture. https://en.wikipedia.org/wiki/Microservices
PaaS **(Platform-as-a-Service)**	PaaS is a layer of value-add services, typically for application development, deployment, monitoring, and general lifecycle management. The use of IaC with IaaS and PaaS is very closely associated with DevOps practices.
Private on-premises data center	A data center infrastructure housed within an environment owned by a given enterprise is distinguished from other forms of data center, with the implication that the private data center is more secure, given that access is restricted to those authorized by the enterprise. Thus, circumstances can arise where very sensitive IT assets are only deployed in a private data center, in contrast to using public IaaS. For many intents and purposes, the underlying technology can be identical, allowing for hybrid deployments where some IT assets are privately deployed but also accessible to other assets in public IaaS. IAM, VPNs, firewalls, and similar are key technologies needed to underpin the security of such an arrangement.
REST API	Representational State Transfer (REST) APIs is a generic term for APIs accessed over HTTP(S), typically transporting data encoded in JSON or XML. REST APIs have the advantage that they support distributed systems, communicating over HTTP, which is a well-understood protocol from a security management perspective. REST APIs are another element of a cloud-native applications architecture, alongside microservices. https://en.wikipedia.org/wiki/Representational_state_transfer
SaaS **(Software-as-a-Service)**	End-user applications provided “aaS” over the public Internet, with the underlying software systems and infrastructure owned and managed by the provider.
SAML **(Security Assertion Markup Language)**	Used in the context of Single-Sign-On (SSO) for exchanging authentication and authorization data between an identity provider, typically an IAM system, and a service provider (some form of SaaS). The SAML protocol exchanges XML documents that contain security assertions used by the aaS for access control decisions. https://en.wikipedia.org/wiki/Security_Assertion_Markup_Language
Terraform	An open-source IaC software tool for cloud services, based on declarative configuration files. https://www.terraform.io

Appendix C - Acronyms

AAA–Authentication, Authorization, and Accounting

ACP–Access-Control Policy

ACI–Cisco Application Centric Infrastructure

ACK–Acknowledge or Acknowledgement

ACL–Access-Control List

AD–Microsoft Active Directory

AFI–Address Family Identifier

AMP–Cisco Advanced Malware Protection

AP–Access Point

API–Application Programming Interface

APIC– Cisco Application Policy Infrastructure Controller (ACI)

ASA–Cisco Adaptative Security Appliance

ASM–Any-Source Multicast (PIM)

ASR–Aggregation Services Router

Auto-RP–Cisco Automatic Rendezvous Point protocol (multicast)

AVC–Application Visibility and Control

BFD–Bidirectional Forwarding Detection

BGP–Border Gateway Protocol

BMS–Building Management System

BSR–Bootstrap Router (multicast)

BYOD–Bring Your Own Device

CAPWAP–Control and Provisioning of Wireless Access Points Protocol

CDP–Cisco Discovery Protocol

CEF–Cisco Express Forwarding

CMD–Cisco Meta Data

CPU–Central Processing Unit

CSR–Cloud Services Routers

CTA–Cognitive Threat Analytics

CUWN–Cisco Unified Wireless Network

CVD–Cisco Validated Design

CYOD–Choose Your Own Device

DC–Data Center

DHCP–Dynamic Host Configuration Protocol

DM–Dense-Mode (multicast)

DMVPN–Dynamic Multipoint Virtual Private Network

DMZ–Demilitarized Zone (firewall/networking construct)

DNA–Cisco Digital Network Architecture

DNS–Domain Name System

DORA–Discover, Offer, Request, ACK (DHCP Process)

DWDM–Dense Wavelength Division Multiplexing

ECMP–Equal Cost Multi Path

EID–Endpoint Identifier

EIGRP–Enhanced Interior Gateway Routing Protocol

EMI–Electromagnetic Interference

ETR–Egress Tunnel Router (LISP)

EVPN–Ethernet Virtual Private Network (BGP EVPN with VXLAN data plane)

FHR–First-Hop Router (multicast)

FHRP–First-Hop Redundancy Protocol

FMC–Cisco Firepower Management Center

FTD–Cisco Firepower Threat Defense

GBAC–Group-Based Access Control

GbE–Gigabit Ethernet

Gbit/s–Gigabits Per Second (interface/port speed reference)

GRE–Generic Routing Encapsulation

GRT–Global Routing Table

HA–High-Availability

HQ–Headquarters

HSRP–Cisco Hot-Standby Routing Protocol

HTDB–Host-tracking Database (SD-Access control plane node construct)

IBNS–Identity-Based Networking Services (IBNS 2.0 is the current version)

ICMP– Internet Control Message Protocol

IDF–Intermediate Distribution Frame; essentially a wiring closet.

IEEE–Institute of Electrical and Electronics Engineers

IETF–Internet Engineering Task Force

IGP–Interior Gateway Protocol

IID–Instance-ID (LISP)

IOE–Internet of Everything

IoT–Internet of Things

IP–Internet Protocol

IPAM–IP Address Management

IPS–Intrusion Prevention System

IPSec–Internet Protocol Security

ISE–Cisco Identity Services Engine

ISR–Integrated Services Router

IS-IS–Intermediate System to Intermediate System routing protocol

ITR–Ingress Tunnel Router (LISP)

LACP–Link Aggregation Control Protocol

LAG–Link Aggregation Group

LAN–Local Area Network

L2 VNI–Layer 2 Virtual Network Identifier; as used in SD-Access Fabric, a VLAN.

L3 VNI– Layer 3 Virtual Network Identifier; as used in SD-Access Fabric, a VRF.

LHR–Last-Hop Router (multicast)

LISP–Location Identifier Separation Protocol

MAC–Media Access Control Address (OSI Layer 2 Address)

MAN–Metro Area Network

MEC–Multichassis EtherChannel, sometimes referenced as ***MCEC***

MDF–Main Distribution Frame; essentially the central wiring point of the network.

MnT–Monitoring and Troubleshooting Node (Cisco ISE persona)

MOH–Music on Hold

MPLS–Multiprotocol Label Switching

MR–Map-resolver (LISP)

MS–Map-server (LISP)

MSDP–Multicast Source Discovery Protocol (multicast)

MTU–Maximum Transmission Unit

NAC—Network Access Control

NAD—Network Access Device

NAT—Network Address Translation

NBAR—Cisco Network-Based Application Recognition (NBAR2 is the current version).

NFV—Network Functions Virtualization

NSF—Non-Stop Forwarding

OSI—Open Systems Interconnection model

OSPF—Open Shortest Path First routing protocol

OT—Operational Technology

PAgP—Port Aggregation Protocol

PAN—Primary Administration Node (Cisco ISE persona)

PCI DSS—Payment Card Industry Data Security Standard

PD—Powered Devices (PoE)

PETR—Proxy-Egress Tunnel Router (LISP)

PIM—Protocol-Independent Multicast

PITR—Proxy-Ingress Tunnel Router (LISP)

PnP—Plug-n-Play

PoE—Power over Ethernet (Generic term, may also refer to IEEE 802.3af, 15.4W at PSE)

PoE+—Power over Ethernet Plus (IEEE 802.3at, 30W at PSE)

PSE—Power Sourcing Equipment (PoE)

PSN—Policy Service Node (Cisco ISE persona)

pxGrid—Platform Exchange Grid (Cisco ISE persona and publisher/subscriber service)

PxTR—Proxy-Tunnel Router (LISP - device operating as both a PETR and PITR)

QoS—Quality of Service

RADIUS—Remote Authentication Dial-In User Service

REST—Representational State Transfer

RFC—Request for Comments Document (IETF)

RIB—Routing Information Base

RLOC—Routing Locator (LISP)

RP—Rendezvous Point (multicast)

RP—Redundancy Port (WLC)

RP–Route Processer

RPF–Reverse Path Forwarding

RR–Route Reflector (BGP)

RTT–Round-Trip Time

SA–Source Active (multicast)

SAFI–Subsequent Address Family Identifiers (BGP)

SD–Software-Defined

SDA–Cisco Software Defined-Access

SDN–Software-Defined Networking

SFP–Small Form-Factor Pluggable (1 GbE transceiver)

SFP+– Small Form-Factor Pluggable (10 GbE transceiver)

SGACL–Security-Group ACL

SGT–Scalable Group Tag, sometimes reference as Security Group Tag

SM–Spare-mode (multicast)

SNMP–Simple Network Management Protocol

SSID–Service Set Identifier (wireless)

SSM–Source-Specific Multicast (PIM)

SSO–Stateful Switchover

STP–Spanning-tree protocol

SVI–Switched Virtual Interface

SVL–Cisco StackWise Virtual

SWIM–Software Image Management

SXP–Scalable Group Tag Exchange Protocol

Syslog–System Logging Protocol

TACACS+–Terminal Access Controller Access-Control System Plus

TCP–Transmission Control Protocol (OSI Layer 4)

UCS– Cisco Unified Computing System

UDP–User Datagram Protocol (OSI Layer 4)

UPoE–Cisco Universal Power Over Ethernet (60W at PSE)

UPoE+– Cisco Universal Power Over Ethernet Plus (90W at PSE)

URL–Uniform Resource Locator

VLAN—Virtual Local Area Network

VM—Virtual Machine

VN—Virtual Network, analogous to a VRF in SD-Access

VNI—Virtual Network Identifier (VXLAN)

vPC—virtual Port Channel (Cisco Nexus)

VPLS—Virtual Private LAN Service

VPN—Virtual Private Network

VPNv4—BGP address family that consists of a Route-Distinguisher (RD) prepended to an IPv4 prefix

VPWS—Virtual Private Wire Service

VRF—Virtual Routing and Forwarding

VSL—Virtual Switch Link (Cisco VSS component)

VSS—Cisco Virtual Switching System

VXLAN—Virtual Extensible LAN

WAN—Wide-Area Network

WLAN—Wireless Local Area Network (generally synonymous with IEEE 802.11-based networks)

WoL—Wake-on-LAN

xTR—Tunnel Router (LISP – device operating as both an ETR and ITR)

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, " DE-SIGNS") IN THIS MANUAL ARE PRESENTED " AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WAR-RANTIES, 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 APPLICA-TION 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.**