

Flashstack with Citrix 7, VMware vSphere 7.0U2, and Pure Storage FlashArray//X R3 for up to 2,500 Users

Deployment Guide for Virtual Desktop Infrastructure built on Cisco UCS B200 M6 with 3rd Generation Intel Xeon Scalable Processors, Cisco UCS Manager 4.2.(1), Pure Storage FlashArray//X70 R3Array, Citrix Virtual Apps and Desktops7 2019, and VMware vSphere 7.0 U2 Hypervisor

Published: February 2022



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>.

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_UT1)

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

© 2022 Cisco Systems, Inc. All rights reserved.

Contents

Executive Summary.....	5
Solution Overview	6
Technology Overview.....	8
Solution Design	25
Deployment Hardware and Software	40
Solution Configuration	45
Configuration and Installation.....	47
Build the Virtual Machines and Environment for Workload Testing.....	148
Test Setup, Configuration, and Load Recommendation.....	307
Test Procedure	315
Test Results.....	323
Summary	362
About the Author	363
References.....	364
Appendix	367
Feedback	518

Executive Summary

Cisco Validated Designs 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 Server Infrastructure for VMware vSphere 7.0 U2 Design Guide](#), which describes a validated Converged Infrastructure (CI) jointly developed by Cisco and Pure Storage. The solution explains the deployment of a predesigned, best-practice data center architecture with Citrix Virtual Apps and Desktops and VMware vSphere built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and 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 Inter-sight 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).

Solution Overview

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility and reducing costs. Cisco, Pure Storage, Citrix 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 a large-scale Citrix Virtual Apps and Desktops 7 with Pure Storage FlashArray//X array, Cisco UCS M6 Blade Servers running VMware vSphere 7.0 U2, Cisco Nexus 9000 Series Ethernet Switches and 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 M6 Servers featuring the 3rd Gen Intel Xeon Scalable processors.

Highlights for this design include:

- Support for Cisco UCS B200 M6 blade servers with 3rd Gen Intel Xeon Scalable Family processors and 3200 MHz memory
- Support for the Cisco UCS Manager 4.2
- Support for Pure Storage FlashArray//X50 R3 with Purity version 6.1.7
- Citrix Virtual Apps and Desktops 7 2019
- Support for VMware vSphere 7.0 U2
- 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

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 Second-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 lower costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.

Virtualization: The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtual environments to support evolving business needs.

Storage access: Cisco UCS 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 resources in minutes instead of days.

Cisco UCS Differentiators

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

Embedded Management: In Cisco UCS, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating the need for any external physical or virtual devices to manage the servers.

Unified Fabric: In Cisco UCS, from blade server chassis or rack servers to FI, there is a single Ethernet cable used for LAN, SAN, and management traffic. This converged I/O results in reduced cables, SFPs and adapters - reducing capital and operational expenses of the overall solution.

Auto Discovery: By simply inserting the blade server in the chassis or connecting the rack server to the fabric interconnect, discovery and inventory of compute resources occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.

Policy Based Resource Classification: Once a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD showcases the policy-based resource classification of Cisco UCS Manager.

Combined Rack and Blade Server Management: Cisco UCS Manager can manage Cisco UCS B-series blade servers and Cisco UCS C-series rack servers under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.

Model based Management Architecture: The Cisco UCS Manager architecture and management database is model based, and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.

Policies, Pools, Templates: The management approach in Cisco UCS Manager is based on defining policies, pools, and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network, and storage resources.

Loose Referential Integrity: In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each other. This provides great flexibility where different experts from different domains, such as network, storage, security, server, and virtualization work together to accomplish a complex task.

Policy Resolution: In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to another policy by name is resolved in the organizational hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then the special policy named "de-

fault” 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 UCS Manager

Cisco UCS Manager (UCSM) provides a unified, integrated management for all software and hardware components in Cisco UCS. Using [Cisco Single Connect](#) technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive graphical user interface (GUI), a command-line interface (CLI), or a through a robust application programming interface (API).

Cisco UCS Manager is embedded into the Cisco UCS Fabric Interconnect and provides a unified management interface that integrates server, network, and storage. Cisco UCS Manager performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers a comprehensive set of XML API for third party integration, exposes thousands of integration points, and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

Cisco UCS Manager 4.0 provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System™ (Cisco UCS) across multiple chassis and Cisco UCS servers. Cisco UCS Manager 4.0 is a unified software release for all supported Cisco UCS hardware platforms. Release 4.0 enables support for UCS 6454 Fabric Interconnects, VIC 1400 series adapter cards on Cisco UCS M6 servers and third-Generation Intel® Xeon® Scalable processor refresh and Intel® Optane™ Data Center persistent memory modules on UCS Intel-based M6 servers.

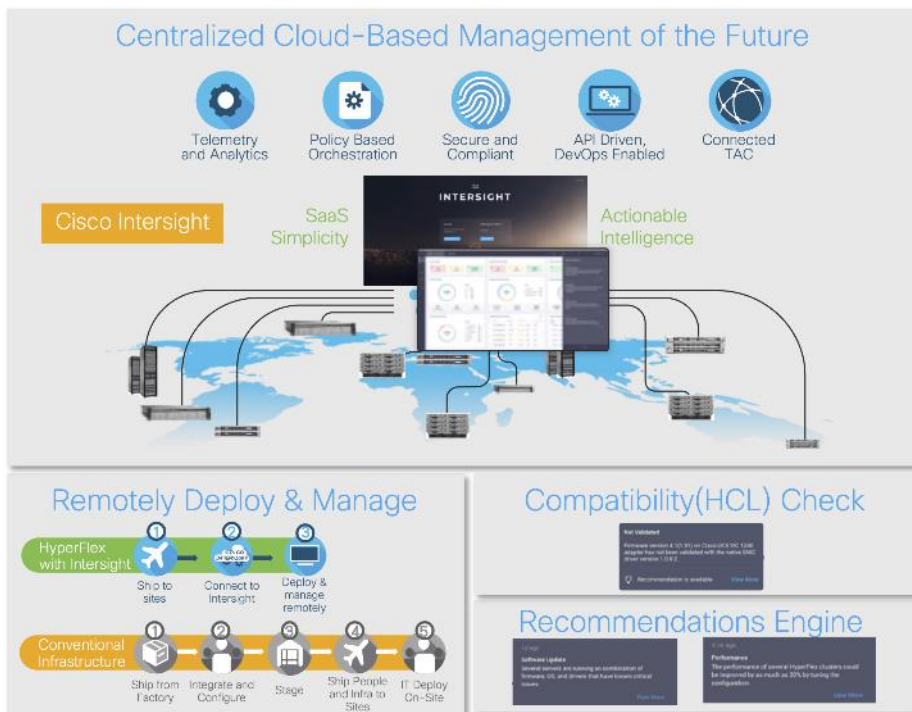
For more information on Cisco UCS Manager Release 4.0 refer to the [Release Notes page](#).

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®) and Cisco HyperFlex™ systems.

Cisco Intersight software delivers a new level of cloud-powered intelligence that supports lifecycle management with continuous improvement. It is tightly integrated with the Cisco® Technical Assistance Center (TAC). Expertise and information flow seamlessly between Cisco Intersight and IT teams, providing global management of Cisco infrastructure, anywhere. Remediation and problem resolution are supported with automated upload of error logs for rapid root-cause analysis.

Figure 1. Cisco Intersight



Automate your infrastructure

Cisco has a strong track record for management solutions that deliver policy-based automation to daily operations. Intersight SaaS is a natural evolution of our strategies. Cisco designed Cisco UCS and HyperFlex 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 HyperFlex 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.

Automate updates to Cisco HyperFlex™ Data Platform software, reducing complexity and manual efforts.

Actionable intelligence

Use best practices to enable faster, proactive IT operations.

Gain actionable insight for ongoing improvement and problem avoidance.

Manage anywhere

Deploy in the data center and at the edge with massive scale.

Get visibility into the health and inventory detail for your Intersight Managed environment on-the-go with the Cisco Intersight Mobile App.

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

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 networks. The Fabric Interconnect supports multiple traffic classes over the Ethernet fabric from the servers to the uplinks. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6454 Fabric Interconnect

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

Figure 2. Cisco UCS 6454 Fabric Interconnect



Cisco UCS B200 M6 Blade Server

The Cisco UCS B200 M6 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads, including Virtual Desktop Infrastructure (VDI), web infrastructure, distributed databases, converged infrastructure, and enterprise applications such as Oracle and SAP HANA. The B200 M6 server can quickly deploy stateless physical and virtual workloads through programmable, easy-to-use Cisco UCS Manager and Cisco Intersight™ and simplified server access through Cisco® SingleConnect technology. It includes:

- 3rd Gen Intel® Xeon® Scalable and processors with up to 40 cores per socket

- Up to 32 DDR4 DIMMs for improved performance with up to 16 DIMM slots ready for Intel Optane™ PMem
- Up to 2 Small Form-Factor (SFF) drives or up to 4 M.2 SATA drives
- Up to 80 Gbps of I/O throughput

Figure 3. Cisco UCS B200M6



Cisco UCS VIC 1457 MLOM Interface Card

The Cisco UCS VIC 1457 Card is a quad-port Enhanced Small Form-Factor Pluggable (SFP+) 10/25-Gbps Ethernet, and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS C-Series Rack Servers. The Cisco UCS VIC 1457 is used in conjunction with the Cisco UCS 6454 model Fabric Interconnects. The mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases. The card enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host, each dynamically configured as either a network interface card (NICs) or host bus adapter (HBA). The personality of the interfaces is set programmatically using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and Worldwide Name [WWN]), failover policy, adapter settings, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all specified using the service profile.

Figure 4. Cisco UCS VIC 1457 mLOM Card



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 infrastructure
- 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 Gigabit 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 5. 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 6](#)) 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 7.) 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 6. Cisco MDS 9132T 32-Gb Fibre Channel Switch



Figure 7. Cisco MDS 9132T 32-Gb 16-Port Fibre Channel Port Expansion Module



- Features
 - High performance: 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: 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 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 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 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 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 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 MDS 9132T can be provisioned, managed, monitored, and troubleshoot using Cisco Data Center Network Manager (DCNM), which currently manages the entire suite of Cisco data center products.
- Self-contained advanced anticounterfeiting technology: The 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 App and Desktops 7 2019

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.

For more information, go to: [Citrix Virtual Apps and Desktops](#)

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 – with FlashArray//X – 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 last year, and they've been busy at work 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 native files alongside blocks.

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 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 management.

Figure 8. 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)	
<ul style="list-style-type: none"> • 2 × 1/10/25Gb Ethernet • 2 × 1/10/25Gb Ethernet Replication • 2 × 1Gb Management Ports 	<ul style="list-style-type: none"> • 2-port 10GBase-T Ethernet • 2-port 1/10/25Gb Ethernet • 2-port 40Gb Ethernet 	<ul style="list-style-type: none"> • 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.

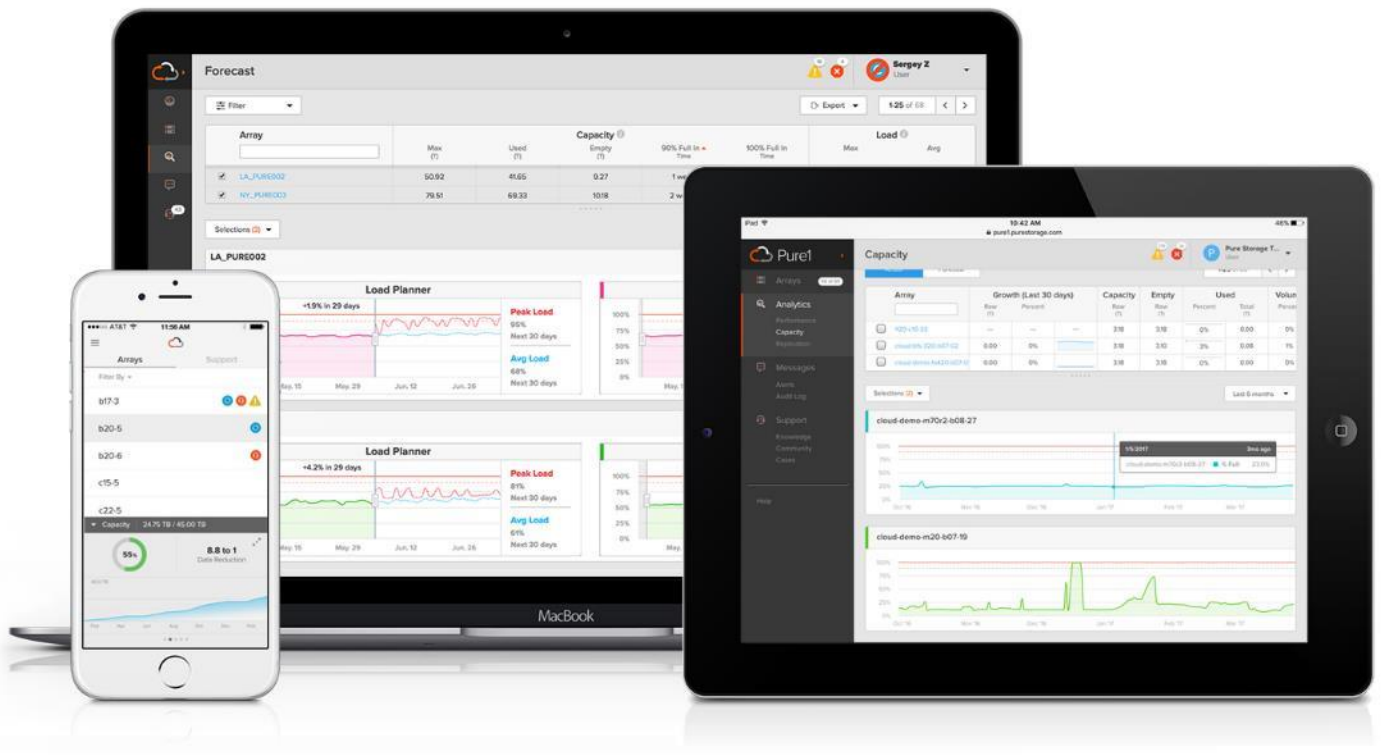
† Array accepts Pure Storage DirectFlash Shelf and/or Pure Storage SAS-based expansion shelf.

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

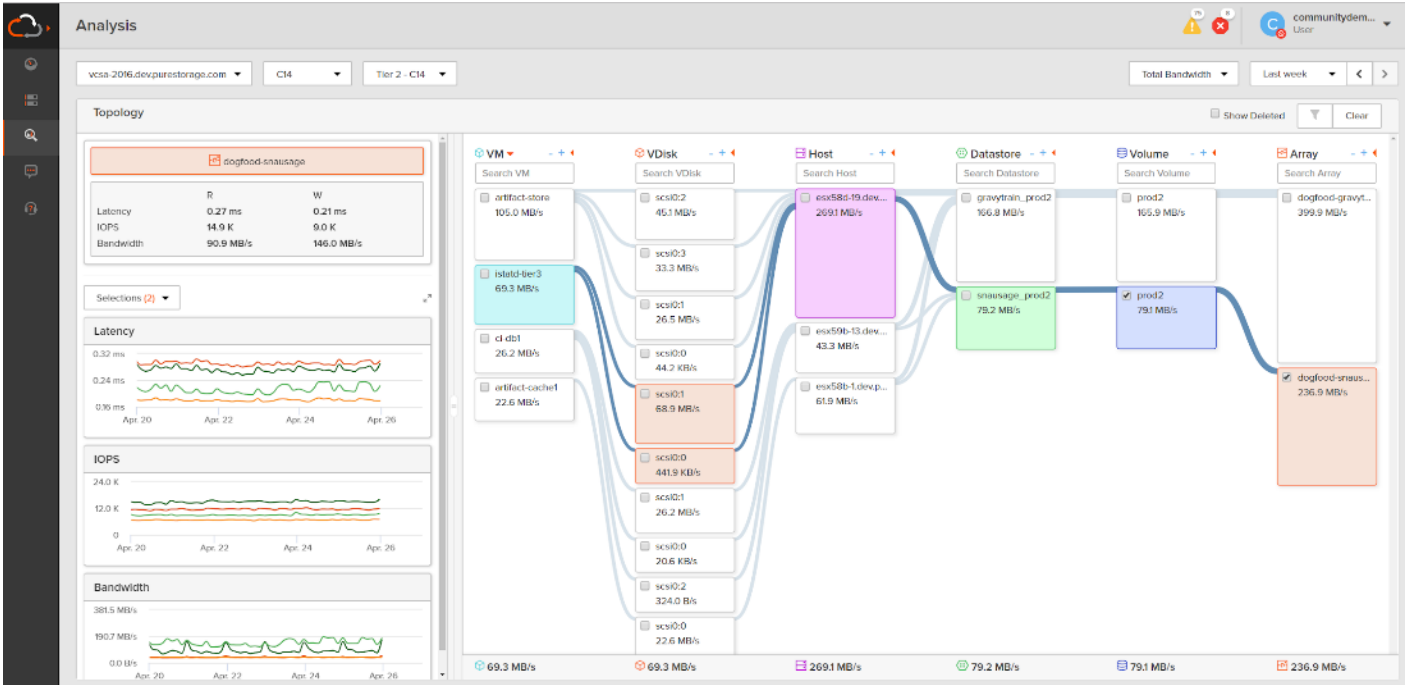
The foundation of Pure1 services, Pure1 Meta is global intelligence built from a massive collection of storage array health and performance data. By continuously scanning call-home telemetry from Pure's installed base, Pure1 Meta uses machine learning predictive analytics to help resolve potential issues 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.



Solution Design

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 desktop environments for each user:

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

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

- **Published Applications:** Published applications run entirely on the VMware RDS server virtual machines and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft Office, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.
- **Streamed Applications:** Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly, but the resources may only 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.

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 be used for virtual desktop deployment?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is there a 3rd party graphics component?
- Is anti-virus a part of the image?
- What is the SQL server version for 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 7.0 U2 has been selected as the hypervisor for this Citrix Virtual Apps and Desktop 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 multi-pathing 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.

VMware Virtual Volumes Considerations

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

Ensure that the Config vVol is either part of an existing FlashArray Protection Group, Storage Policy that includes snapshots, or manual snapshots of the Config vVol are taken. This will help with the VM recovery process if the VM is deleted.

There are some FlashArray limits on Volume Connections per Host, Volume Count, and Snapshot Count. For more information about FlashArray limits review the following:

https://support.purestorage.com/FlashArray/PurityFA/General_Troubleshooting/Pure_Storage_FlashArray_Limits

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

Pure Storage FlashArray Best Practices for VMware vSphere 7.0

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 7.0.

-
- vSphere 7.0 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

Pure Storage FlashArray Best Practices for VMware Virtual Volumes (vVols)

Along with the Pure Storage Best Practices for VMware vSphere, the following should be considered as part of a design that includes the implementation of vVols as part of the solution:

- Create a Local FlashArray Array Admin user to register the storage provider with vs using the local pureuser account, vvols-admin for example.
- Use the Round Robin pathing policy (default) for the Protocol Endpoint.
- Use the Pure Storage Plugin for the vSphere Client to register the FlashArray storage provider and mount the vVols Datastore if possible.

-
- If manually registering the storage providers, Register both controllers' storage providers with CT0.ETH0 and CT1.ETH0. It is supported to use Eth1 if a customer certificate is used.
 - If manually mounting the vVol datastore, you will need to connect the protocol endpoint.
 - A single PE should be enough for the design utilizing the default device queue depth for the PE.
 - Keep VM Templates on vVols when deploying new vVol VMs from a template.
 - When resizing a VM's VMDK that resides on a vVol, complete the task from vSphere Client and not the FlashArray GUI.
 - vCenter Server should not reside on a vVol
 - All ESXi Hosts, vCenter Server and FlashArray should have the same NTP Server synchronization configuration and be configured to send their logs to a syslog target.
 - TCP port 8084 must be open and accessible from vCenter Servers and ESXi hosts to the FlashArray that will be used for vVol.
 - The FlashArray Protocol Endpoint object 'pure-protocol-endpoint' must exist. The FlashArray admin must not rename, delete or otherwise edit the default FlashArray Protocol Endpoint.

For more information about vVols best practices, go to:

https://support.purestorage.com/Solutions/VMware_Platform_Guide/Quick_Reference_by_VMware_Product_and_Integration/Virtual_Volumes_Quick_Reference

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 7 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.

Users 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:

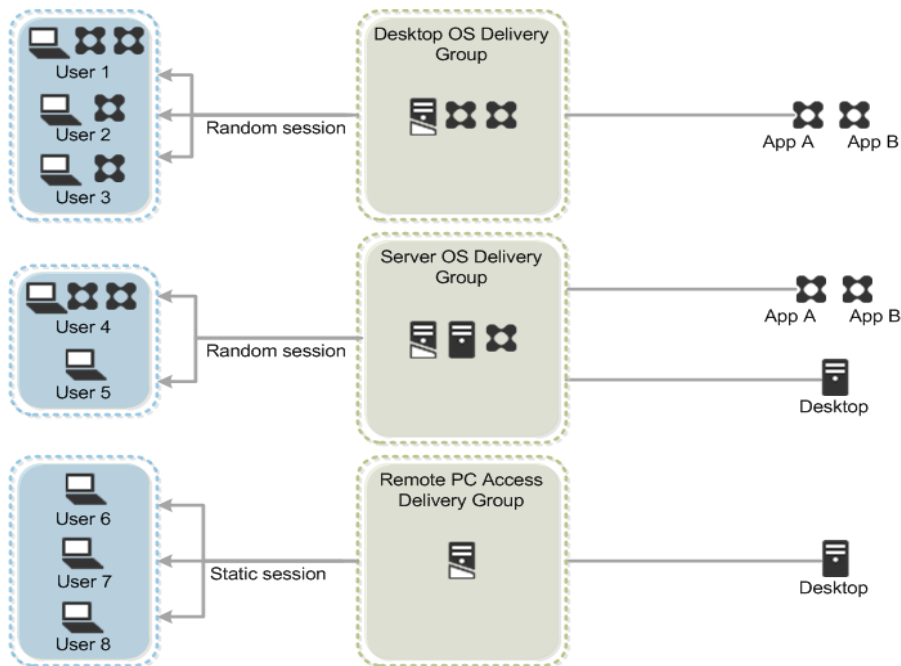
- Use machines from multiple catalogs
- Allocate a user to multiple machines
- Allocate multiple users to one machine

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 9](#) illustrates how users access desktops and applications through machine catalogs and delivery groups.

Figure 9. Access Desktops and Applications through Machine Catalogs and Delivery Groups




Citrix Provisioning Services

Citrix Virtual Apps and Desktops 7 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 10. 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 on device hard drive. Write cache exists as a file in NTFS format, located on the target device's hard drive. This option frees up the Provisioning Server since it does not have to process write requests and does not have the finite limitation of RAM.
- Cache on device hard drive persisted. (Experimental Phase) This is the same as "Cache on device hard drive", except that the cache persists. At this time, this method is an experimental feature only, and is only supported for NT6.1 or later (Windows 10 and Windows 2008 R2 and later). This method also requires a different bootstrap.
- 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 cre-

ated. 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 2019 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 2019 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 11](#) 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 11. Example of a Distributed Components Configuration

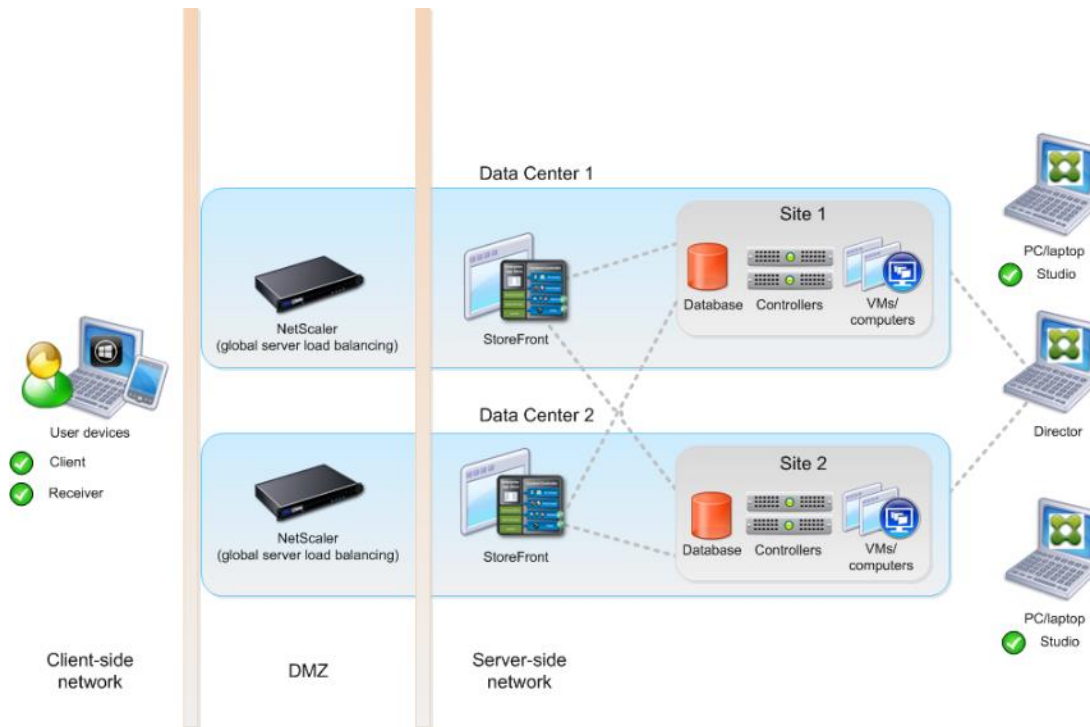


Multiple Site Configuration

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 12](#) 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 12. 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 App and Desktop 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.

Table 1. Desktop Types and User Experience

Desktop Type	User Experience
Server OS Machines	<p>You want: Inexpensive server-based delivery to minimize the cost of delivering applications to a large number of users, while providing a secure, high-definition user experience.</p> <p>Your users: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations.</p> <p>Application types: Any application.</p>
Desktop OS Machines	<p>You want: A client-based application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition.</p> <p>Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications.</p> <p>Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines.</p> <p>Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.</p>
Remote PC Access	<p>You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the data center.</p> <p>Your users: Employees or contractors that have the option to work from home but need access to specific software or data on their corporate desktops to perform their jobs remotely.</p> <p>Host: The same as Desktop OS machines.</p> <p>Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.</p>

For this Cisco Validated Design, the following designs are included:

1. Single-session OS Solution:

MCS: 1960 Windows 10 Virtual desktops random pooled were configured and tested

PVS: 1960 Windows 10 Virtual desktops random pooled were configured and tested

2. Multi-session OS Solution:

MCS: 2688 Windows Server 2019 random pooled desktops were configured and tested

Deployment Hardware and Software

Architecture

FlashStack with Cisco UCS M6 servers, Citrix Virtual Apps and Desktops 2019, and vSphere 7.0 U2 delivers a Virtual Desktop Infrastructure that is redundant, using the best practices of Cisco and Pure Storage. The solution includes VMware vSphere 7.0 U2 hypervisor installed on the Cisco UCS M6 compute nodes configured for stateless compute design using boot from SAN. Pure Storage FlashArray//X70 R3 provides the storage infrastructure required for setting up the VDI workload. Cisco UCS manager is utilized to configure and manage the Cisco UCS infrastructure with Cisco Intersight providing lifecycle management capabilities. The solution requirements and design details are covered in this section.

Physical Topology

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 covered below.

Products Deployed

This CVD details the deployment of up to 2688 Multi-session OS, 1960 Single-session OS VDI users featuring the following software:

- VMware vSphere ESXi 7.0 U2 Hypervisor
- Microsoft SQL Server 2019
- Microsoft Windows Server 2019 and Windows 10 64-bit virtual machine Operating Systems
- Citrix Virtual Apps and Desktops 2019
- Citrix Provisioning Server 2019
- FSLogix 2105 HF_01
- Citrix StoreFront 2019

Figure 13. FlashStack VDI - Physical Topology for FC

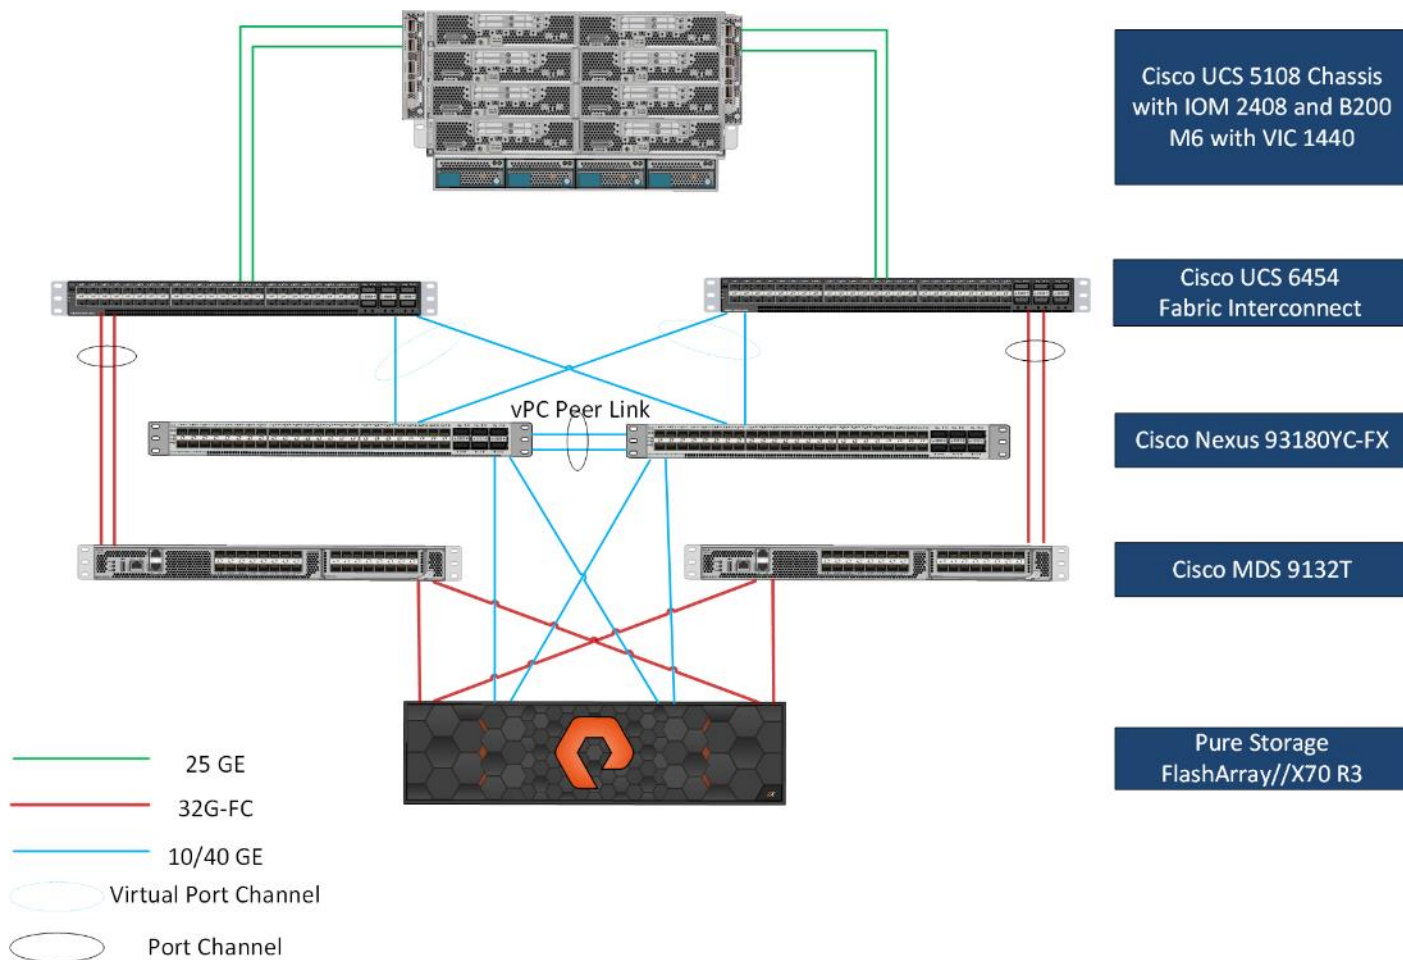


Figure 13 details the physical hardware and cabling deployed to enable this solution:

- Two Cisco Nexus 93180YC-FX Switches in NX-OS Mode.
- Two Cisco MDS 9132T 32-Gb Fibre Channel Switches.
- Four Cisco UCS 5108 Blade Server Chassis with two Cisco UCS-IOM-2408 IO Modules.
- Eight Cisco UCS B200 M6 Blade Servers with Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM, and one Cisco VIC1440 mezzanine card, providing N+1 server fault tolerance.
- Pure Storage FlashArray//X70 R3 with dual redundant controllers, with Twenty 1.92TB Direct-Flash NVMe drives.

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

[Table 2](#) lists the software versions of the primary products installed in the environment.

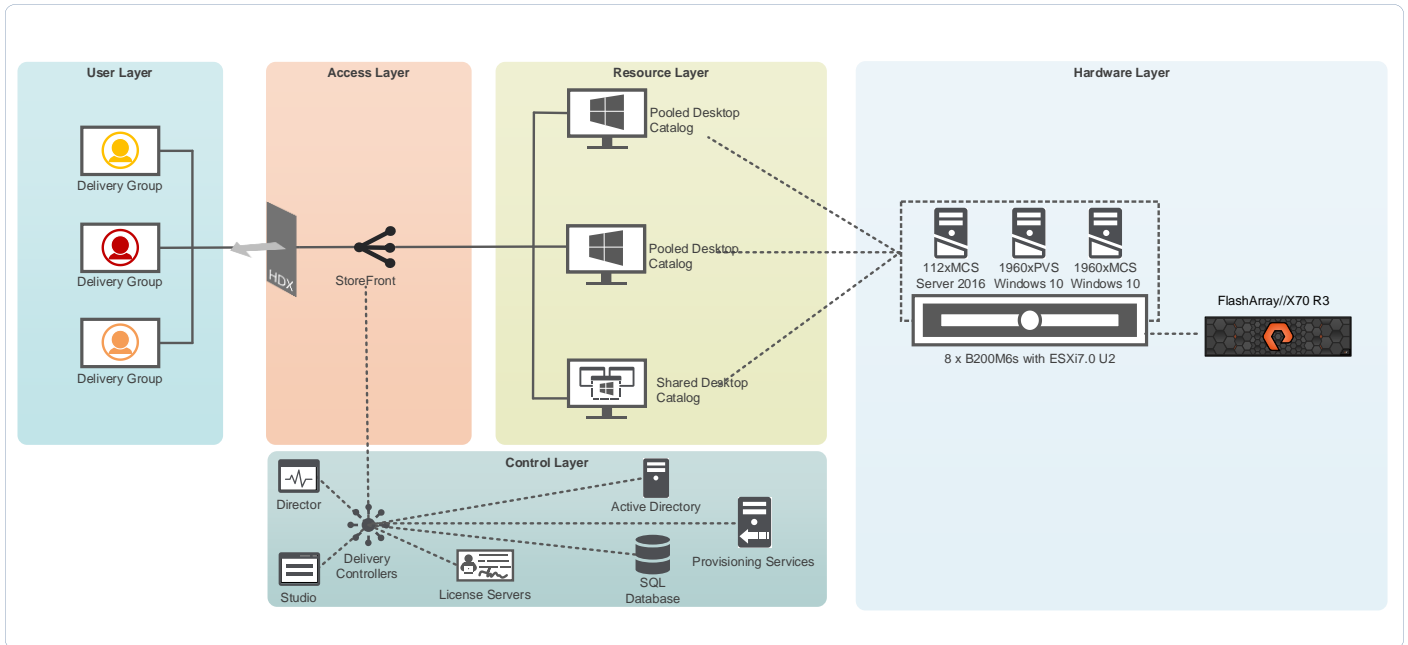
Table 2. Software and Firmware Versions

Vendor	Product / Component	Version / Build / Code
Cisco	UCS Component Firmware	4.2(1f) bundle release
Cisco	UCS Manager	4.2(1f) bundle release
Cisco	UCS B200 M5 Blades	4.2(1f) bundle release
Cisco	VIC 1440	4.2(1f) bundle release
Cisco	Cisco Nexus 93180YC-FX	9.3(7a)
Cisco	Cisco MDS 9132T	8.5(1a)
Pure Storage	FlashArray//X70 R3	Purity//FA 6.1.7
VMware	vCenter Server Appliance	7.0.2.00200 Build: 17958471
VMware	vSphere 7. 0 2U	7.0.2, 17867351
Citrix	Citrix Virtual Apps and Desktops 7 2109	2109.0.0.31047
Citrix	Provisioning Services	2109.0.0
Citrix	Store Front	2109.0.0.31047
Citrix VDA		2109.0.0.31047
Microsoft	FSLogix 2015 HF_01	2.9.7654.46150
VMware	Tools	11.2.5.17337674

Logical Architecture

The logical architecture of the validated solution which is designed to support up to 2688 users on a single chassis containing 8 blades, with physical redundancy for the blade servers for each workload type is illustrated in [Figure 14](#).

Figure 14. Logical Architecture Overview



Configuration Guidelines

The VMware Horizon solution described in this document provides details for configuring a fully redundant, highly-available configuration. Configuration guidelines are provided that refer to which redundant component is being configured with each step, whether that be A or B. For example, 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 VMware Horizon 7.12 customer environment as a stand-alone solution.

VLANs

The VLAN configuration recommended for the environment includes a total of six VLANs as outlined in [Table 3](#).

Table 3. VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
In-Band-Mgmt	70	In-Band management interfaces
Infra-Mgmt	71	Infrastructure Virtual Machines
VCC/VM-Network	72	RDSH, VDI Persistent and Non-Persistent
vMotion	73	VMware vMotion

VLAN Name	VLAN ID	VLAN Purpose
OOB-Mgmt	164	Out of Band management interfaces

VSANs

[Table 4](#) lists the two virtual SANs that were configured for communications and fault tolerance in this design.

Table 4. VSANs Configured in this study

VSAN Name	VSAN ID	Purpose
VSAN 100	100	VSAN for Primary SAN communication
VSAN 101	101	VSAN for Secondary SAN communication

Solution Configuration

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution.

Solution Cabling

The following sections detail the physical connectivity configuration of the FlashStack 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.

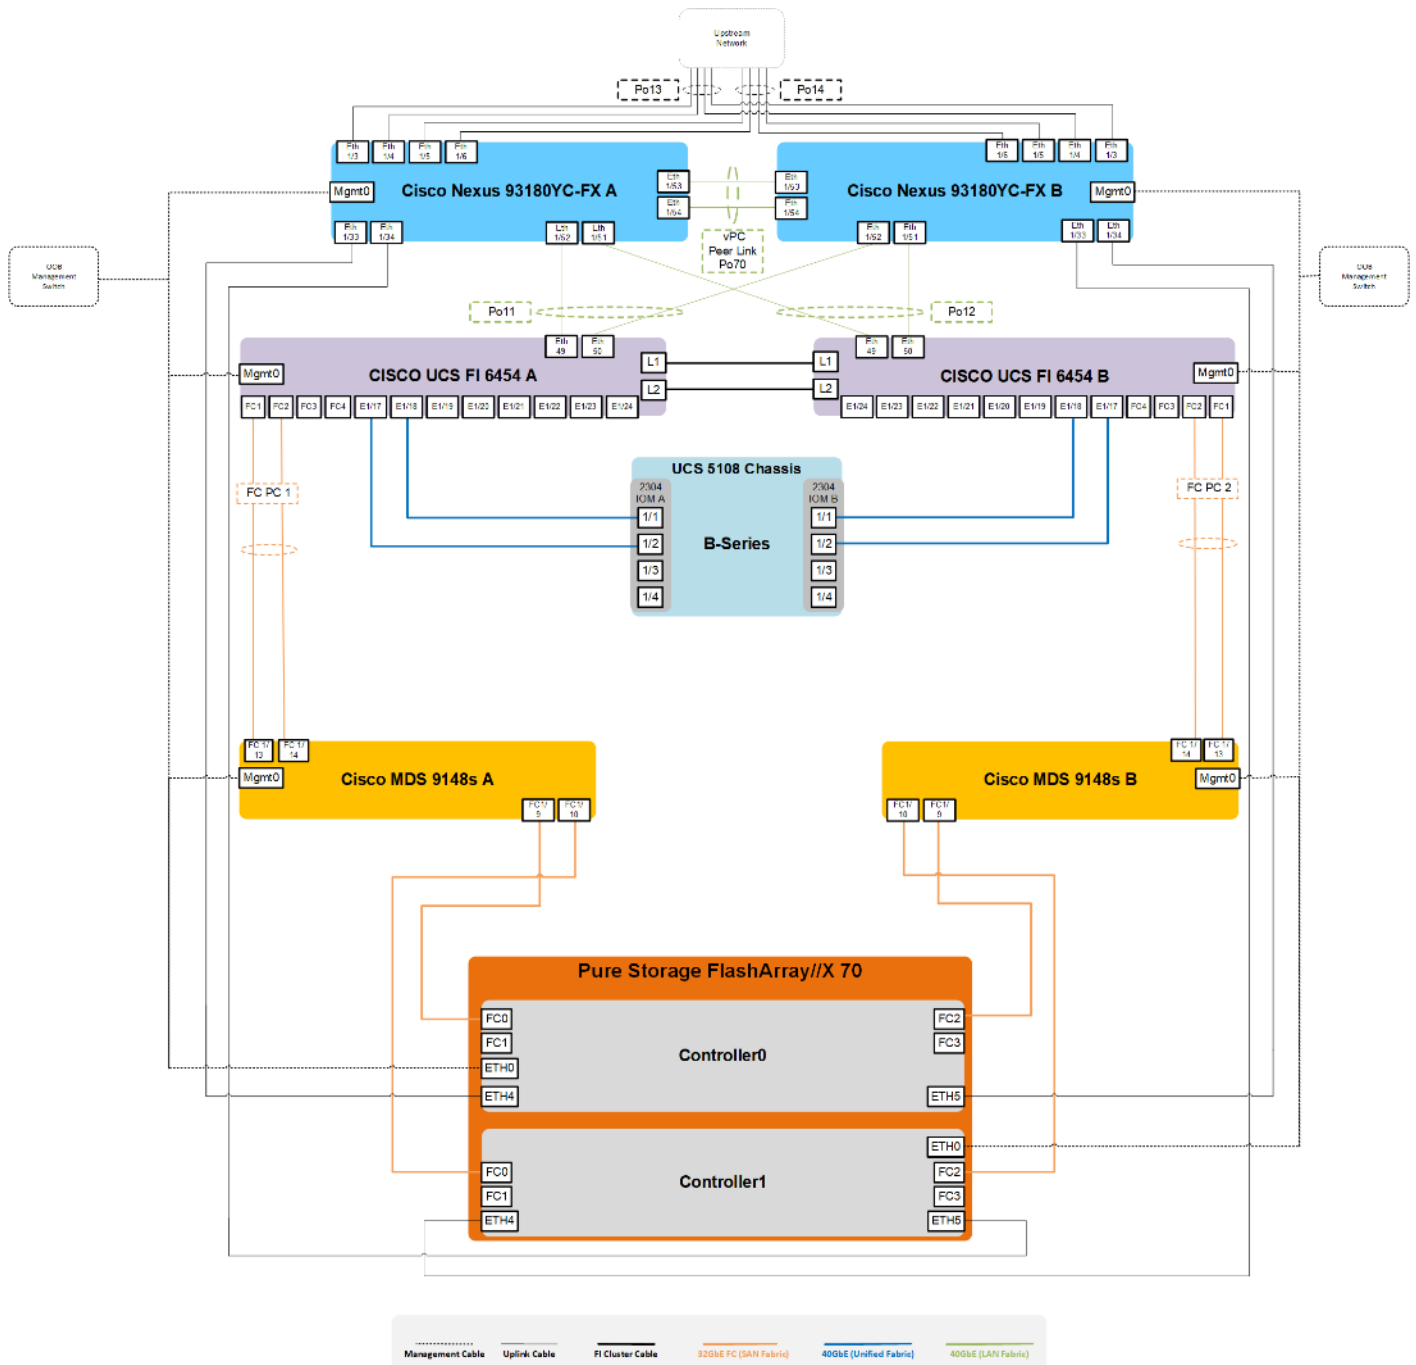
The tables in this section list the details for the prescribed and supported configuration of the Pure Storage FlashArray//X70 R3 storage array to the Cisco 6454 Fabric Interconnects through Cisco MDS 9132T 32-Gb FC switches.

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 15](#) 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 15. FlashStack Solution Cabling Diagram



Configuration and Installation

FlashStack Automated Deployment with Ansible

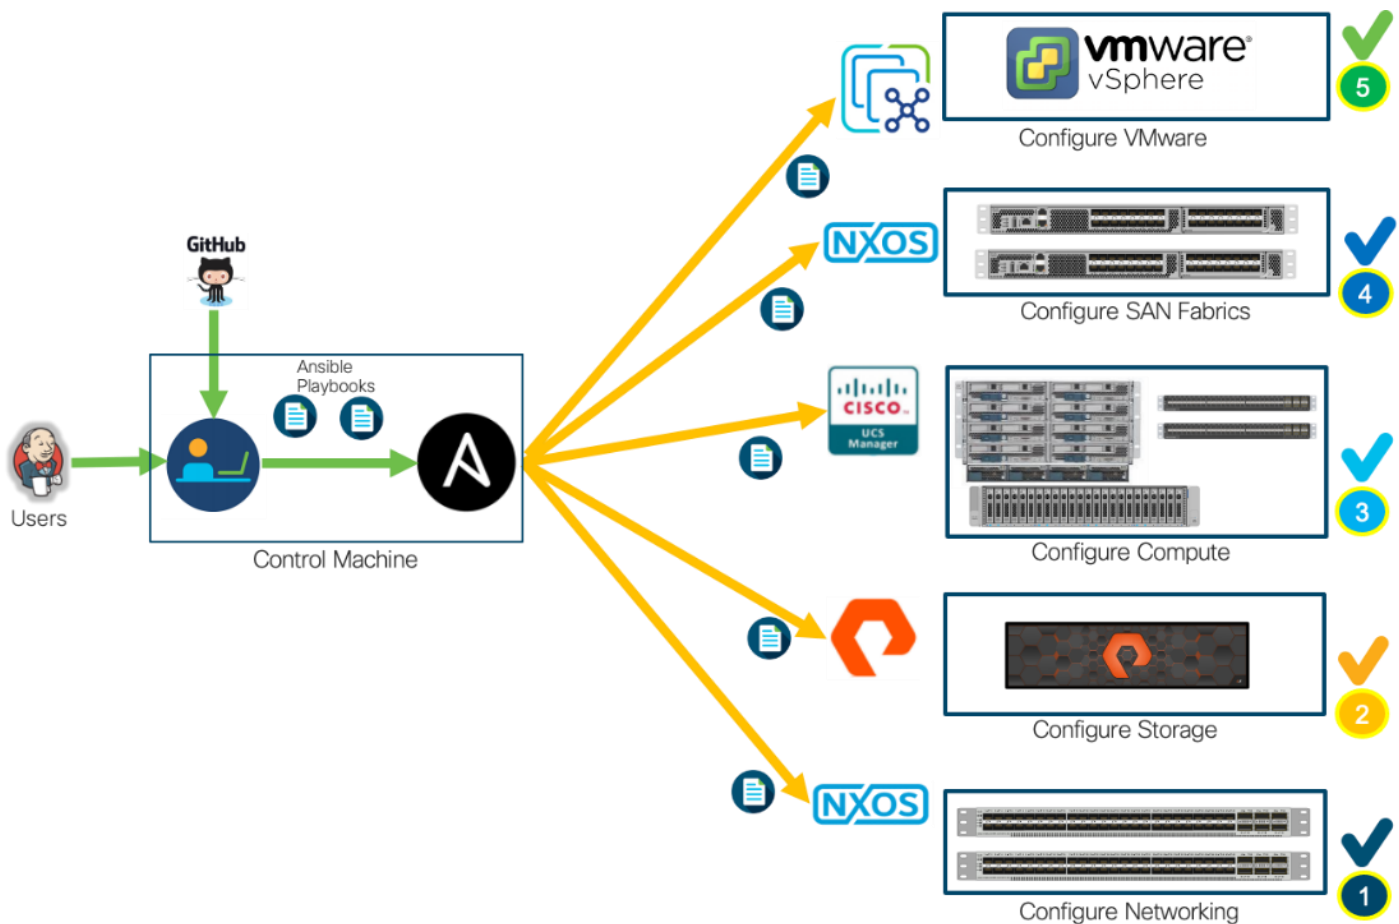
If using the published Ansible playbooks to configure the FlashStack infrastructure, complete this section of the document.

Ansible Automation Workflow and Solution Deployment

This FlashStack with vSphere 7.0 U2 and Cisco UCS M6 solution uses a management workstation (control machine) to run Ansible playbooks to configure Cisco Nexus, Cisco UCS, Pure Storage and Install VMware Cluster.

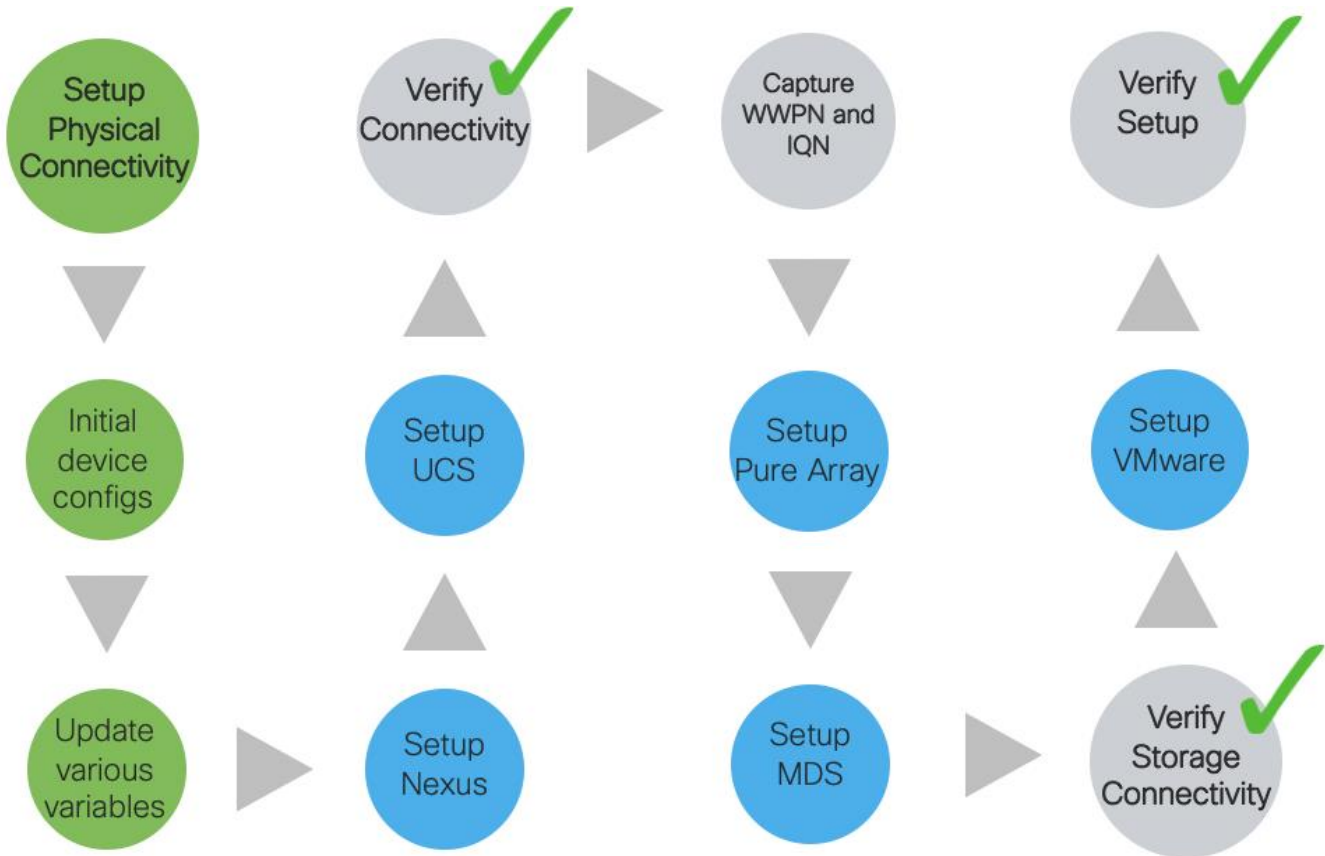
[Figure 16](#) illustrates the FlashStack with VMware vSphere 7.0 U2 and Cisco UCS solution implementation workflow, which is explained in the following sections. The FlashStack Ansible based automated deployment is shown in [Figure 17](#).

Figure 16. High-level FlashStack Automation



[Figure 17](#) illustrates the FlashStack Automated deployment workflow.

Figure 17. FlashStack Automated Deployment Workflow



Prerequisites

Setting up the solution begins with a management workstation that has access to the internet and has a working installation of Ansible. The management workstation runs a variant of Linux or MacOS for ease of use with these command-line-based tools. Instructions for installing the workstation are not included in this document, but the basic installation and configuration of Ansible is explained. For detailed information, go to: [Getting Started with Red Hat Ansible](#)

The following is a list of prerequisites:

1. To use the Ansible playbooks demonstrated in this document, the management workstation must also have a working installation of Git and access to the Cisco DevNet public GitHub repository. The Ansible playbooks used in this document are cloned from the public repositories, located at the following links:
 - Cisco DevNet: <https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/Flashstack-laC-UCSM6/>
 - GitHub repository for FlashStack infrastructure setup: <https://github.com/ucs-compute-solutions/Flashstack-laC-UCSM6>

2. The Cisco Nexus Switches, Pure Storage and Cisco UCS must be physically racked, cabled, powered, and configured with the management IP addresses before the Ansible-based installation procedure can begin as shown in the cabling diagram ([Figure 16](#)). If necessary, upgrade the Nexus Switches to release 9.3(7) and the UCS System to 4.2(1f) with the default firmware packages for both blades and rack servers set to 4.2(1f).
3. Before running each Ansible Playbook to setup the Network, Storage, Cisco UCS and VMware, various variables must be updated based on the customers environment and specific implementation with values such as the VLANs, pools and ports on Cisco UCS, IP addresses for iSCSI interfaces and values needed for the ESXi installation and configuration.

Note: Day 2 Configuration tasks, such as adding datastores or ESXi servers, were performed manually or with Cisco Intersight Cloud Orchestrator (ICO) and the information has been provided in the respective sections of this document.

Prepare Management Workstation (Control Machine)

In this section, the installation steps are performed on the CentOS management host to prepare the host for solution deployment to support the automation of Cisco UCS, Cisco Nexus, Pure Storage and VMware installation using Ansible Playbooks.

To prepare the management workstation, follow these steps:

1. Install the EPEL repository on the management host.

```
[root@FSV-Automation ~]# yum install epel-release
```

2. Install Ansible engine.

```
[root@FSV-Automation ~]# yum install ansible
```

3. Verify the Ansible version to make sure it's at least release 2.9.

```
[root@FS-Automation tasks]# ansible --version
ansible 2.10.7
  config file = None
  configured module search path = ['/root/.ansible/plugins/modules', '/usr/share/ansible/plugins/modules']
  ansible python module location = /usr/local/lib/python3.6/site-packages/ansible
  executable location = /usr/local/bin/ansible
  python version = 3.6.8 (default, Aug 24 2020, 17:57:11) [GCC 8.3.1 20191121 (Red Hat 8.3.1-5)]
```

4. Install **pip** the package installer for Python.

```
[root@FSV-Automation ~]# yum install python-pip
```

5. Install the UCS SDK.

```
[root@FSV-Automation ~]# pip3 install ucsm sdk
```

6. Install the **paramiko** package for Cisco Nexus automation.

```
[root@FSV-Automation ~]# pip3 install paramiko
```

7. SSH into each of the Cisco Nexus and Cisco MDS switches using Ansible so that the SSH keys are cached.

```
[root@FSV-Automation ~]# ssh admin@10.1.164.61
The authenticity of host '10.1.164.61 (10.1.164.61)' can't be established.
RSA key fingerprint is SHA256:mtomJluZVkcITgSLhVygocSnojlyPPDPmcJLQX2dfu4.
RSA key fingerprint is MD5:b4:e3:86:97:99:58:df:0d:5d:20:b2:5b:d5:69:aa:23.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.1.164.61' (RSA) to the list of known hosts.
User Access Verification
Password:
```

8. Install the Pure Storage SDK.

```
[root@FSV-Automation ~]# pip3 install purestorage
```

9. Install ansible-galaxy collections for Cisco UCS, Cisco Nexus/MDS switches and Pure Storage Array as follows:

```
[root@FSV-Automation ~]# ansible-galaxy collection install cisco.nxos
[root@FSV-Automation ~]# ansible-galaxy collection install cisco.ucs
[root@FSV-Automation ~]# ansible-galaxy collection install purestorage.flasharray
```

Note: We validated the Ansible automation with both Python 2.7.5 and Python 3.6 as the Python interpreter for Ansible.

FlashStack Manual Deployment

Cisco Unified Computing System Base Configuration

This section details the Cisco UCS configuration that was done as part of the infrastructure build out. The racking, power, and installation of the chassis are described in the [Cisco UCS Manager Getting Started Guide](#) and it is beyond the scope of this document. For more information about each step, refer to the following document, [Cisco UCS Manager - Configuration Guides](#).

Cisco UCS Manager Software Version 4.2(1f)

This document assumes you are using Cisco UCS Manager Software version 4.2(1f). To upgrade the Cisco UCS Manager software and the Cisco UCS 6454 Fabric Interconnect software to a higher version of the firmware,) go to [Cisco UCS Manager Install and Upgrade Guides](#).

Configure Fabric Interconnects at Console

To configure the fabric Interconnects, follow these steps:

1. Connect a console cable to the console port on what will become the primary fabric interconnect.
2. If the fabric interconnect was previously deployed and you want to erase it to redeploy, follow these steps:

- Login with the existing user name and password:

```
# connect local-mgmt
# erase config
# yes (to confirm)
```

- After the fabric interconnect restarts, the out-of-box first time installation prompt appears, type “console” and press Enter.
3. Follow the [Initial Configuration](#) steps as outlined in [Cisco UCS Manager Getting Started Guide](#). When configured, log into UCSM IP Address through the web interface to perform base Cisco UCS configuration.

Configure Fabric Interconnects for a Cluster Setup

To configure the Cisco UCS Fabric Interconnects, follow these steps:

1. Verify the following physical connections on the fabric interconnect:
 - The management Ethernet port (mgmt0) is connected to an external hub, switch, or router
2. The L1 ports on both fabric interconnects are directly connected to each other
3. The L2 ports on both fabric interconnects are directly connected to each other
4. Connect to the console port on the first Fabric Interconnect.
5. Review the settings on the console. Answer yes to Apply and Save the configuration.
6. Wait for the login prompt to make sure the configuration has been saved to Fabric Interconnect A.
7. Connect the console port on the second Fabric Interconnect, configure secondary FI.

Figure 18. Initial Setup of Cisco UCS Manager on Primary Fabric Interconnect

```
Enter the configuration method. (console/gui) ? console

Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup

You have chosen to setup a new Fabric interconnect. Continue? (y/n): y

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

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

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes

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

Enter the system name: VCC-AAD17

Physical Switch Mgmt0 IP address : 10.29.164.246

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.29.164.1

Cluster IPv4 address : 10.29.164.245

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

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

Join centralized management environment (UCS Central)? (yes/no) [n]:

Following configurations will be applied:

Switch Fabric=A
System Name=VCC-AAD17
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=10.29.164.246
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.29.164.1
Ipv6 value=0

Cluster Enabled=yes
Cluster IP Address=10.29.164.245
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized.
      UCSM will be functional only after peer FI is configured in clustering mode.

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

Configuration file - Ok

Cisco UCS 6300 Series Fabric Interconnect
VCC-AAD17-A login: █
```


Figure 19. Initial Setup of Cisco UCS Manager on Secondary Fabric Interconnect

```
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:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.164.246
Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
Cluster IPv4 address      : 10.29.164.245

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

Physical Switch Mgmt0 IP address : 10.29.164.247

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

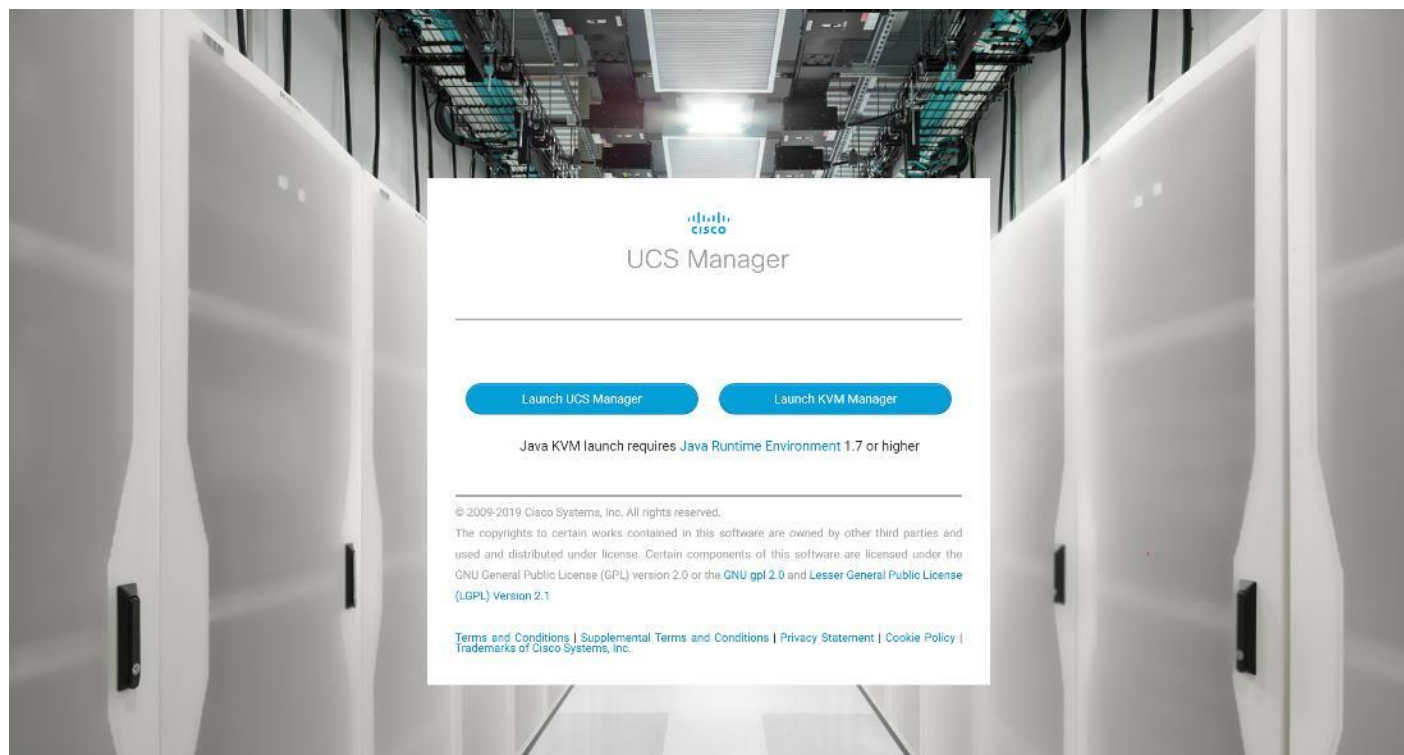
Fri Feb 16 18:53:15 UTC 2018
Configuration file - Ok

Cisco UCS 6300 Series Fabric Interconnect
VCC-AAD17-B login: █
```

To log into the Cisco Unified Computing System (Cisco UCS) environment, follow these steps:

1. Open a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address previously configured.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software. If prompted, accept the security certificates.

Figure 20. Cisco UCS Manager Web Interface



3. When prompted, enter the user name and password enter the password. Click Log In to login to Cisco UCS Manager.

Figure 21. Cisco UCS Manager Web Interface after Login



Configure Base Cisco Unified Computing System

The following are the high-level steps involved for a Cisco UCS configuration:

1. Configure Fabric Interconnects for a Cluster Setup
2. Set Fabric Interconnects to Fibre Channel End Host Mode
3. Synchronize Cisco UCS to NTP
4. Configure Fabric Interconnects for Chassis and Blade Discovery
5. Configure Global Policies
6. Configure Server Ports
7. Configure LAN and SAN on Cisco UCS Manager
8. Configure Ethernet LAN Uplink Ports
9. Create Uplink Port Channels to Cisco Nexus Switches
10. Configure FC SAN Uplink Ports
11. Configure VLAN
12. Configure VSAN
13. Configure IP, UUID, Server, MAC, WWNN and WWPN Pools
14. IP Pool Creation

15. UUID Suffix Pool Creation

16. Server Pool Creation

17. MAC Pool Creation

18. WWNN and WWPN Pool Creation

19. Set Jumbo Frames in both the Cisco Fabric Interconnect

20. Configure Server BIOS Policy

21. Create Adapter Policy

22. Configure Update Default Maintenance Policy

23. Configure vNIC and vHBA Template

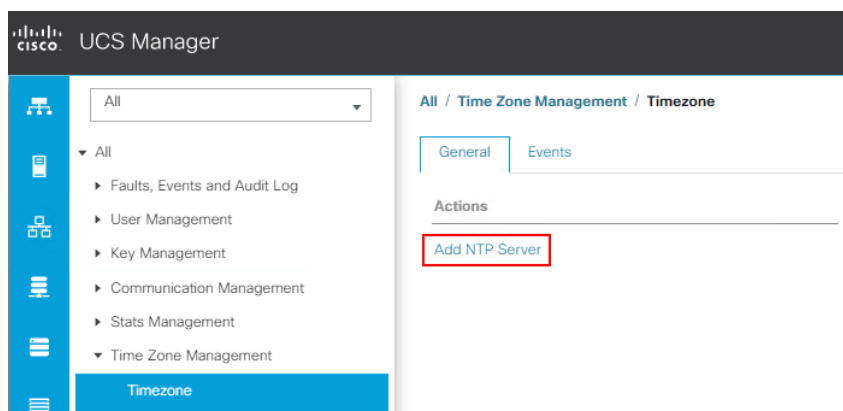
24. Create Server Boot Policy for SAN Boot

Details for each step are discussed in the following sections.

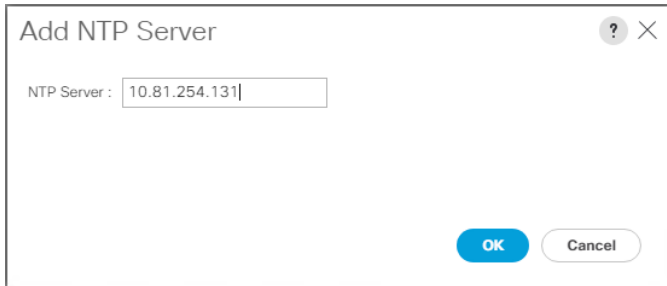
Synchronize Cisco UCSM to NTP

To synchronize the Cisco UCS environment to the NTP server, follow these steps:

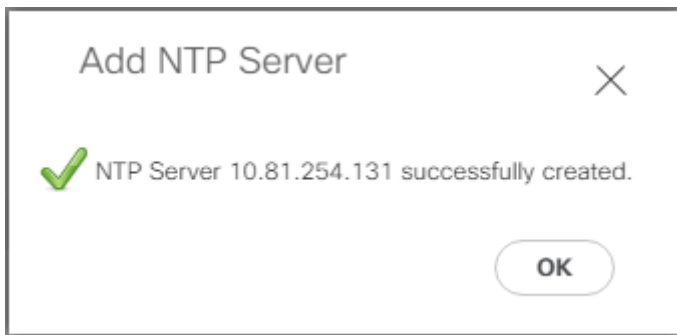
1. In Cisco UCS Manager, in the navigation pane, click the Admin tab.
2. Select All > Time zone Management.
3. In the Properties pane, select the appropriate time zone in the Time zone menu.
4. Click Save Changes and then click OK.
5. Click Add NTP Server.



6. Enter the NTP server IP address and click OK.



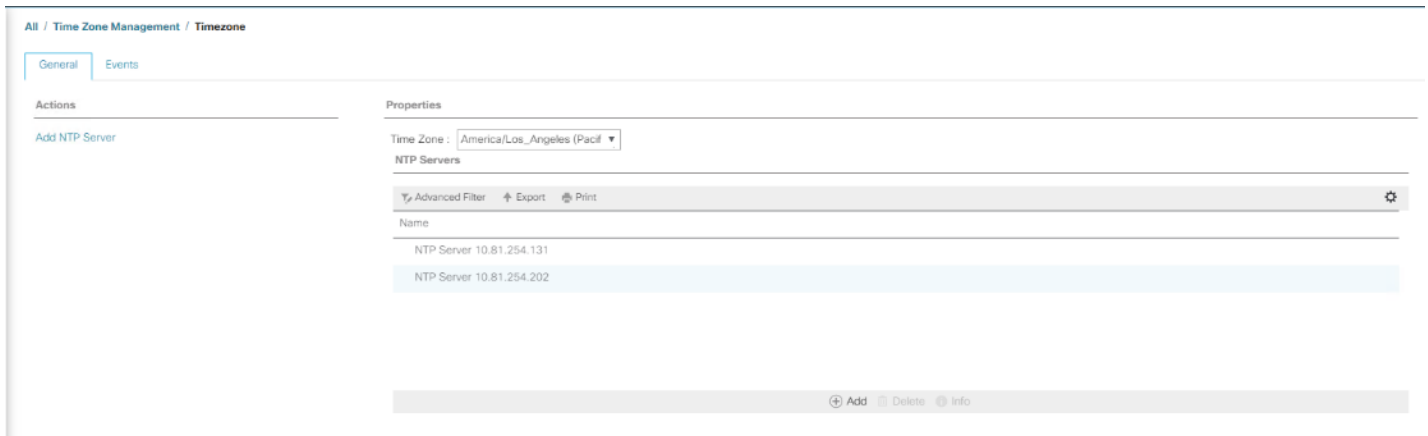
7. Click OK to finish.



8. Repeat steps 1-7 to configure additional NTP servers.

9. Click Save Changes.

Figure 22. Synchronize Cisco UCS Manager to NTP



Configure Fabric Interconnects for Chassis and Blade Discovery

Cisco UCS 6454 Fabric Interconnects are configured for redundancy, this provides resiliency in case of failures. The first step is to establish connectivity between blades and Fabric Interconnects.

Configure Global Policies

The chassis discovery policy determines how the system reacts when you add a new chassis. We recommend using the platform max value as shown. Using platform max helps ensure that Cisco UCS Manager uses the maximum number of IOM uplinks available.

To configure global policies, follow these steps:

1. In Cisco UCS Manager, go to Equipment > Policies (right pane) > Global Policies > Chassis/FEX Discovery Policies. As shown in the screenshot below, for Action select “Platform Max” from the drop-down list and set Link Grouping to Port Channel.
2. Click Save Changes.
3. Click OK.


Figure 23. Cisco UCS Global Policy

The screenshot shows the 'Equipment' section of the Cisco UCS Manager interface, specifically the 'Policies' tab. The 'Global Policies' sub-tab is active, showing the configuration for the 'Chassis/FEX Discovery Policy'. The 'Action' is set to 'Platform Max' and 'Link Grouping Preference' is set to 'Port Channel'. Other policies shown include 'Rack Server Discovery Policy' (Action: Immediate), 'Rack Management Connection Policy' (Action: Auto Acknowledged), 'Power Policy' (Redundancy: N+1), 'MAC Address Table Aging' (Aging Time: Mode Default), 'Global Power Allocation Policy' (Allocation Method: Policy Driven Chassis Group Cap), and 'Firmware Auto Sync Server Policy' (Sync State: No Actions). At the bottom right, there are 'Save Changes' and 'Reset Values' buttons.

Fabric Ports: Discrete versus Port Channel Mode

[Figure 24](#) illustrates the advantage of Discrete Vs Port-Channel mode in UCSM.

Figure 24. Port Channel versus Discrete Mode



Set Fabric Interconnects to Fibre Channel End Host Mode

In order to configure the FC Uplink ports connected to the Cisco UCS MDS 9132T 32-Gb FC switch, set the Fabric Interconnects to the Fibre Channel End Host Mode. Verify that the fabric interconnects are operating in FC End-Host Mode.

The screenshot shows the Cisco UCS Manager interface. On the left, a navigation pane lists 'Fabric Interconnects' with 'Fabric Interconnect A (primary)' selected. Underneath, the 'Set FC End-Host Mode' option is highlighted with a red box. A red arrow points from this box to the 'Actions' button at the bottom of the page. The main content area shows the configuration for 'Fabric Interconnect A (primary)' with tabs for 'General', 'Physical Ports', 'Fans', 'PSUs', and 'Physical DI'. The 'General' tab is active, displaying a 'Fault Summary' section with three status indicators (red, orange, and green) and a list of operational status items: Operable, OK, End Host, and Off.

The fabric interconnect automatically reboots if switched to operational mode; perform this task on one FI first, wait for the FI to come up and repeat this process on the second FI.

Configure FC SAN Uplink Ports

To configure Fibre Channel Uplink ports, follow these steps:

1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > General tab > Actions pane, click Configure Unified Ports.

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate)

General | Physical Ports | Fans | PSUs | Physical Display | FSM | Neighbors | Faults | Events | Statistics

Status

Overall Status : ↑ **Operable**
Thermal : ↑ **OK**
Ethernet Mode : **End Host**
FC Mode : **End Host**
Admin Evac Mode : **Off**
Oper Evac Mode : **Off**

Actions

- Configure Evacuation
- Configure Unified Ports**
- Internal Fabric Manager
- LAN Uplinks Manager
- NAS Appliance Manager
- SAN Uplinks Manager
- SAN Storage Manager
- Enable Ports ▼
- Disable Ports ▼
- Set Ethernet End-Host Mode
- Set Ethernet Switching Mode
- Set FC End-Host Mode
- Set FC Switching Mode
- Activate Firmware
- Management Interfaces
- Turn off Locator LED

Properties

Name : **A**
Product Name : **Cisco UCS 6454**
Vendor : **Cisco Systems, Inc.** PID : **UCS-FI-6454**
Revision : **0** Serial : **FDO22241ZLJ**
Available Memory : Total Memory : **62.761 (GB)**

Locator LED : ●

⊕ Part Details

⊕ Local Storage Information

⊕ Access

⊕ High Availability Details

⊕ VLAN Port Count

⊕ FC Zone Count

Firmware

Boot-loader Version : **v05.40(01/17/2020)**
Kernel Version : **7.0(3)N2(4.12a)**
System Version : **7.0(3)N2(4.12a)**
Service Pack Version : **4.1(2)SP0(Default)**
Package Version : **4.1(2a)A**
Startup Kernel Version : **7.0(3)N2(4.12a)**
Activate Status : **Ready**

2. Click Yes to confirm in the pop-up window.

Configure Unified Ports

⚠ The Configure Unified Ports wizard allows you to change the port mode from Ethernet to Fibre Channel or FC to Ethernet. Changing the port mode on either module causes an interruption in data traffic because changes to the fixed module require a reboot of the fabric interconnect and changes on an expansion module require a reboot of that module. Are you sure you want to launch this wizard and reboot the modules associated with any reconfigured ports?

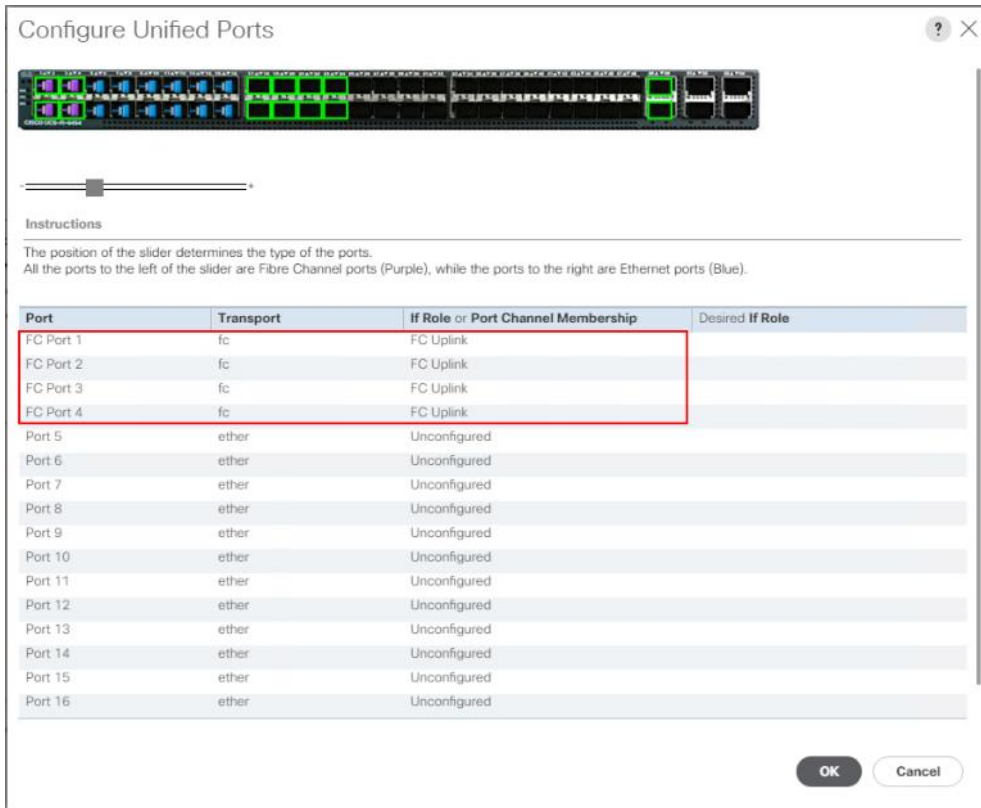
Yes No

3. Move the slider to the right.

4. Click OK.

Note: Ports to the right of the slider will become FC ports. For our study, we configured the first four ports (Ports are configured in sets of 4 ports) on the FI as FC Uplink ports.

Applying this configuration will cause the immediate reboot of the fabric interconnect and/or the expansion module(s).



Configure Unified Ports

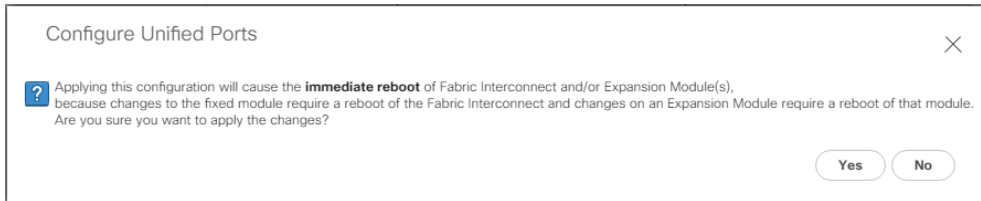
Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
FC Port 1	fc	FC Uplink	
FC Port 2	fc	FC Uplink	
FC Port 3	fc	FC Uplink	
FC Port 4	fc	FC Uplink	
Port 5	ether	Unconfigured	
Port 6	ether	Unconfigured	
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Unconfigured	

OK Cancel

5. Click Yes to apply the changes.

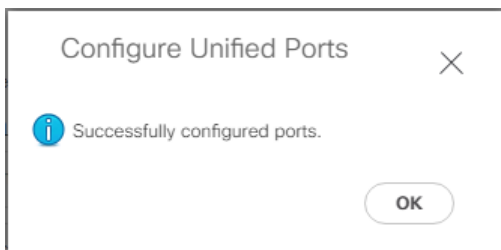


Configure Unified Ports

Applying this configuration will cause the **immediate reboot** of Fabric Interconnect and/or Expansion Module(s), because changes to the fixed module require a reboot of the Fabric Interconnect and changes on an Expansion Module require a reboot of that module. Are you sure you want to apply the changes?

Yes No

6. Click OK to proceed.



Configure Unified Ports

Successfully configured ports.

OK

After the FI reboot, your FC Ports configuration will look like [Figure 25](#).

7. Repeat steps 1-6 on Fabric Interconnect B.

Figure 25. FC Uplink Ports on Fabric Interconnect A

Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
1	1	20-01-00:3A-9C:0E:33:20	Network	Physical	Up	Enabled
1	2	20-02-00:3A-9C:0E:33:20	Network	Physical	Up	Enabled
1	3	20-03-00:3A-9C:0E:33:20	Network	Physical	Up	Enabled
1	4	20-04-00:3A-9C:0E:33:20	Network	Physical	Up	Enabled

Configure Server Ports

Configure the server ports to initiate chassis and blade discovery. To configure server ports, follow these steps:

1. Go to Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module > Ethernet Ports.
2. Select the ports (for this solution ports are 17-24) which are connected to the Cisco IO Modules of the two B-Series 5108 Chassis.
3. Right-click and select “Configure as Server Port.”

Figure 26. Configure Server Port on Cisco UCS Manager Fabric Interconnect for Chassis/Server Discovery

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	17	00:3A:9C:0E:33:38	Unconfigured	Physical	Admin Down	Disabled	
1	0	18	00:3A:9C:0E:33:39	Unconfigured	Physical	Admin Down	Disabled	
1	0	19	00:3A:9C:0E:33:3A	Unconfigured	Physical	Admin Down	Disabled	
1	0	20	00:3A:9C:0E:33:3B	Unconfigured	Physical	Admin Down	Disabled	
1	0	21	00:3A:9C:0E:33:3C	Unconfigured	Physical	Admin Down	Disabled	
1	0	22	00:3A:9C:0E:33:3D	Unconfigured	Physical	Admin Down	Disabled	
1	0	23	00:3A:9C:0E:33:3E	Unconfigured	Physical	Admin Down	Disabled	
1	0	24	00:3A:9C:0E:33:3F	Unconfigured	Physical	Admin Down	Disabled	
1	0	25	00:3A:9C:0E:33:40	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	26	00:3A:9C:0E:33:41	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	27	00:3A:9C:0E:33:42	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	28	00:3A:9C:0E:33:43	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	29	00:3A:9C:0E:33:44	Unconfigured	Physical	Slp Not Present	Disabled	

4. Click Yes to confirm and click OK.

5. Repeat steps 1-4 to configure the Server Port on Fabric Interconnect B.

When configured, the server port will look like [Figure 27](#) on both Fabric Interconnects.

Figure 27. Server Ports on Fabric Interconnect A

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State	Peer
1	0	17	00:3A:9C:0E:33:38	Server	Physical	Up	Enabled	sys/chassis-1/slot-2/fabri...
1	0	18	00:3A:9C:0E:33:39	Server	Physical	Up	Enabled	sys/chassis-2/slot-2/fabri...
1	0	19	00:3A:9C:0E:33:3A	Server	Physical	Up	Enabled	sys/chassis-3/slot-2/fabri...
1	0	20	00:3A:9C:0E:33:3B	Server	Physical	Up	Enabled	sys/chassis-4/slot-2/fabri...
1	0	21	00:3A:9C:0E:33:3C	Server	Physical	Up	Enabled	sys/chassis-1/slot-2/fabri...
1	0	22	00:3A:9C:0E:33:3D	Server	Physical	Up	Enabled	sys/chassis-2/slot-2/fabri...
1	0	23	00:3A:9C:0E:33:3E	Server	Physical	Up	Enabled	sys/chassis-3/slot-2/fabri...
1	0	24	00:3A:9C:0E:33:3F	Server	Physical	Link Up	Enabled	sys/chassis-4/slot-2/fabri...
1	0	25	00:3A:9C:0E:33:40	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	26	00:3A:9C:0E:33:41	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	27	00:3A:9C:0E:33:42	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	28	00:3A:9C:0E:33:43	Unconfigured	Physical	Slp Not Present	Disabled	
1	0	29	00:3A:9C:0E:33:44	Unconfigured	Physical	Slp Not Present	Disabled	

- After configuring Server Ports, acknowledge both the Chassis. Go to Equipment > Chassis > Chassis 1 > General > Actions > select “Acknowledge Chassis”. Similarly, acknowledge the chassis 2-4.
- After acknowledging both the chassis, re-acknowledge all the servers placed in the chassis. Go to Equipment > Chassis 1 > Servers > Server 1 > General > Actions > select Server Maintenance > select option “Re-acknowledge” and click OK. Repeat this process to re-acknowledge all eight Servers.
- When the acknowledgement of the Servers is completed, verify the Port-channel of Internal LAN. Go to the LAN tab > Internal LAN > Internal Fabric A > Port Channels as shown in [Figure 28](#).

Figure 28. Internal LAN Port Channels

Name	Slot ID	Port ID	Aggr. Port ID	Peer Slot ID	Peer Port ID	Fabric ID	Peer
Eth Interface 1/17	1	17	0	2	1	A	sys/switch-A/access-ethsp...
Eth Interface 1/18	1	18	0	2	5	A	sys/switch-A/access-ethsp...

Configure Ethernet LAN Uplink Ports

To configure network ports that are used to uplink the Fabric Interconnects to the Cisco Nexus switches, follow these steps:

- In Cisco UCS Manager, in the navigation pane, click the Equipment tab.

2. Click Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports (for this solution ports are 49-50) that are connected to the Nexus switches, right-click them, and select Configure as Network Port.

Figure 29. Network Uplink Port Configuration on Fabric Interconnect Configuration

Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 35	1	35	00:3A:9C:0E:33:4A	Unconfigured	Physical	Sfp Not Present	Disabled
Port 36	1	36	00:3A:9C:0E:33:4B	Unconfigured	Physical	Sfp Not Present	Disabled
Port 37	1	37	00:3A:9C:0E:33:4C	Unconfigured	Physical	Sfp Not Present	Disabled
Port 38	1	38	00:3A:9C:0E:33:4D	Unconfigured	Physical	Sfp Not Present	Disabled
Port 39	1	39	00:3A:9C:0E:33:4E	Unconfigured	Physical	Sfp Not Present	Disabled
Port 40		40	00:3A:9C:0E:33:4F	Unconfigured	Physical	Sfp Not Present	Disabled
Port 41		41	00:3A:9C:0E:33:50	Unconfigured	Physical	Sfp Not Present	Disabled
Port 42		42	00:3A:9C:0E:33:51	Unconfigured	Physical	Sfp Not Present	Disabled
Port 43		43	00:3A:9C:0E:33:52	Unconfigured	Physical	Sfp Not Present	Disabled
Port 44		44	00:3A:9C:0E:33:53	Unconfigured	Physical	Sfp Not Present	Disabled
Port 45		45	00:3A:9C:0E:33:54	Unconfigured	Physical	Sfp Not Present	Disabled
Port 46		46	00:3A:9C:0E:33:55	Unconfigured	Physical	Sfp Not Present	Disabled
Port 47		47	00:3A:9C:0E:33:56	Unconfigured	Physical	Sfp Not Present	Disabled
Port 48		48	00:3A:9C:0E:33:57	Unconfigured	Physical	Sfp Not Present	Disabled
Port 49		49	00:3A:9C:0E:33:58	Unconfigured	Physical	Admin Down	Disabled
Port 50		50	00:3A:9C:0E:33:59	Unconfigured	Physical	Admin Down	Disabled
Port 51	1	51	00:3A:9C:0E:33:60	Unconfigured	Physical	Sfp Not Present	Disabled
Port 52	1	52	00:3A:9C:0E:33:64	Unconfigured	Physical	Sfp Not Present	Disabled
Port 53	1	53	00:3A:9C:0E:33:68	Unconfigured	Physical	Sfp Not Present	Disabled
Port 54	1	54	00:3A:9C:0E:33:6C	Unconfigured	Physical	Sfp Not Present	Disabled

5. Click Yes to confirm ports and click OK.
6. Verify the Ports connected to Cisco Nexus upstream switches are now configured as network ports.
7. Repeat steps 1-6 for Fabric Interconnect B. The screenshot below shows the network uplink ports for Fabric A.

Figure 30. Network Uplink Port on Fabric Interconnect

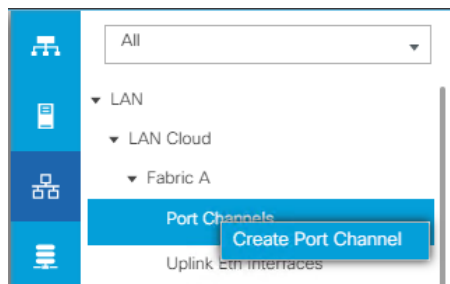
Name	Slot	Port ID	MAC	If Role	If Type	Overall Status	Admin State
Port 37	1	37	00:3A:9C:0E:33:4C	Unconfigured	Physical	Slip Not Present	Disabled
Port 38	1	38	00:3A:9C:0E:33:4D	Unconfigured	Physical	Slip Not Present	Disabled
Port 39	1	39	00:3A:9C:0E:33:4E	Unconfigured	Physical	Slip Not Present	Disabled
Port 40	1	40	00:3A:9C:0E:33:4F	Unconfigured	Physical	Slip Not Present	Disabled
Port 41	1	41	00:3A:9C:0E:33:50	Unconfigured	Physical	Slip Not Present	Disabled
Port 42	1	42	00:3A:9C:0E:33:51	Unconfigured	Physical	Slip Not Present	Disabled
Port 43	1	43	00:3A:9C:0E:33:52	Unconfigured	Physical	Slip Not Present	Disabled
Port 44	1	44	00:3A:9C:0E:33:53	Unconfigured	Physical	Slip Not Present	Disabled
Port 45	1	45	00:3A:9C:0E:33:54	Unconfigured	Physical	Slip Not Present	Disabled
Port 46	1	46	00:3A:9C:0E:33:55	Unconfigured	Physical	Slip Not Present	Disabled
Port 47	1	47	00:3A:9C:0E:33:56	Unconfigured	Physical	Slip Not Present	Disabled
Port 48	1	48	00:3A:9C:0E:33:57	Unconfigured	Physical	Slip Not Present	Disabled
Port 49	1	49	00:3A:9C:0E:33:58	Network	Physical	Up	Enabled
Port 50	1	50	00:3A:9C:0E:33:5C	Network	Physical	Up	Enabled

You have now created two uplink ports on each Fabric Interconnect as shown above. These ports will be used to create Virtual Port Channel in the next section.

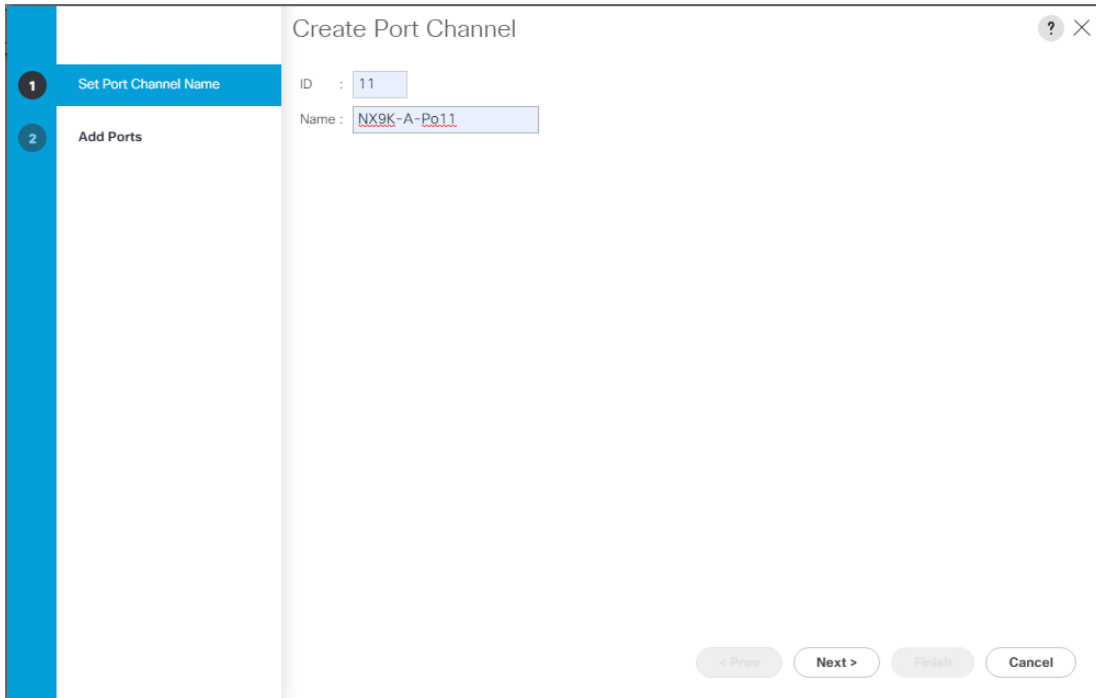
Create Uplink Port Channels to Cisco Nexus Switches

In this procedure, two port channels were created one from Fabric A to both Cisco Nexus 93180YC-FX switches and one from Fabric B to both Cisco Nexus 93180YC-FX switches. To configure the necessary port channels in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click LAN > LAN Cloud > Fabric A.
3. Right-click Port Channels.
4. Click Create Port Channel.

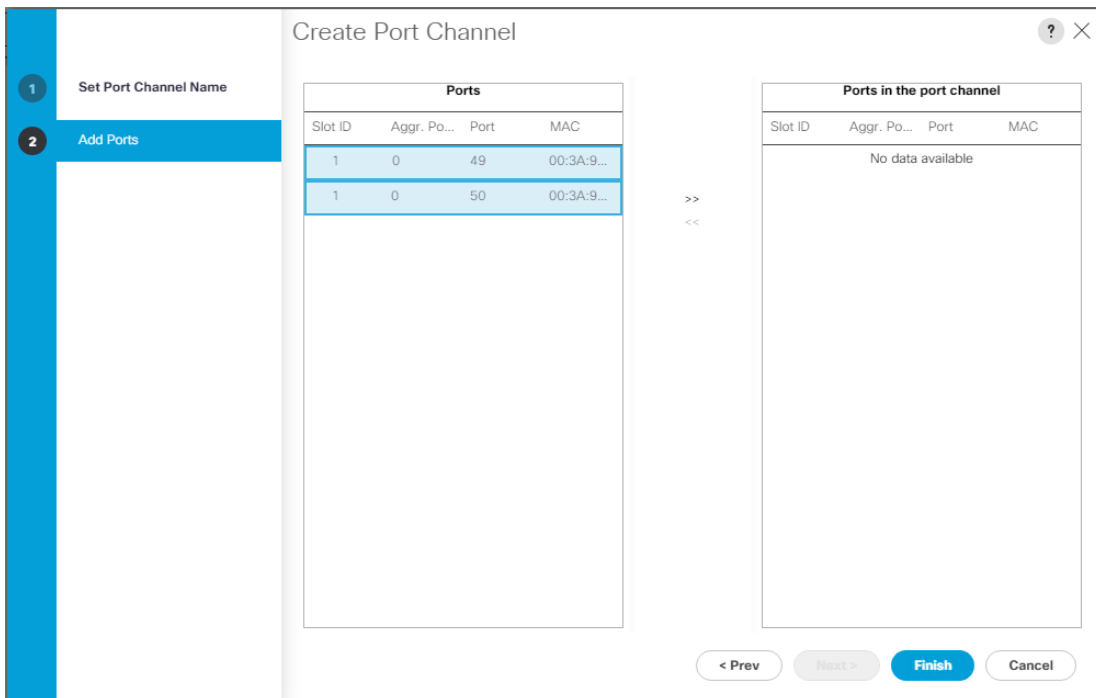


5. Enter 11 as the unique ID of the port channel and name of the port channel.

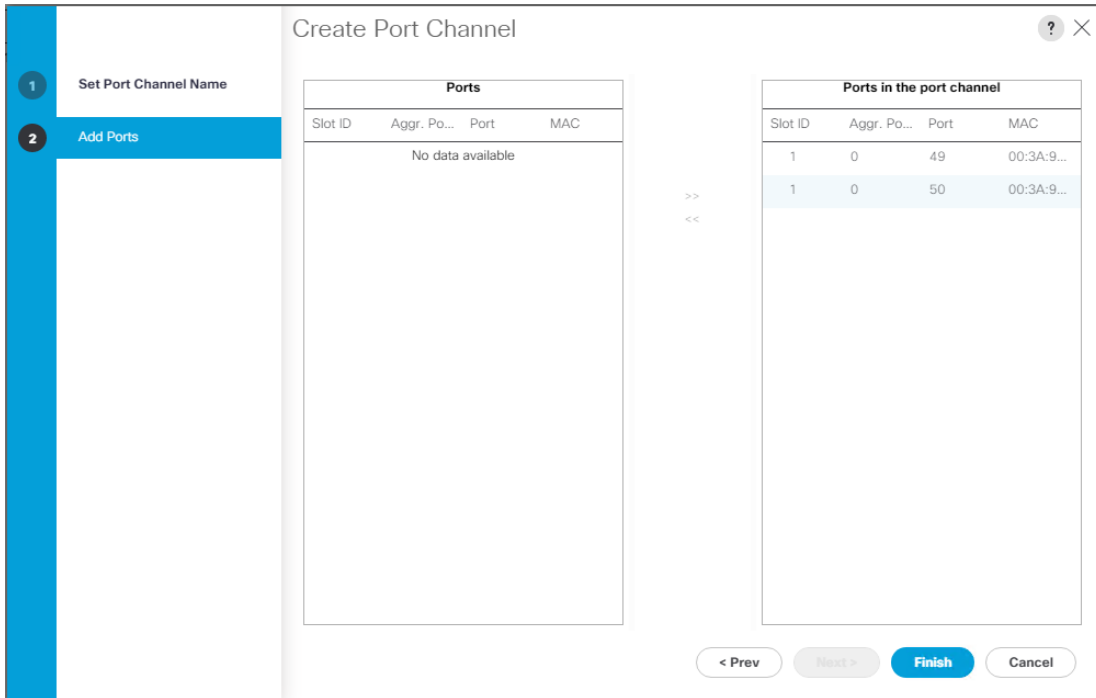


6. Click Next.


7. Select Ethernet ports 49-50 for the port channel.



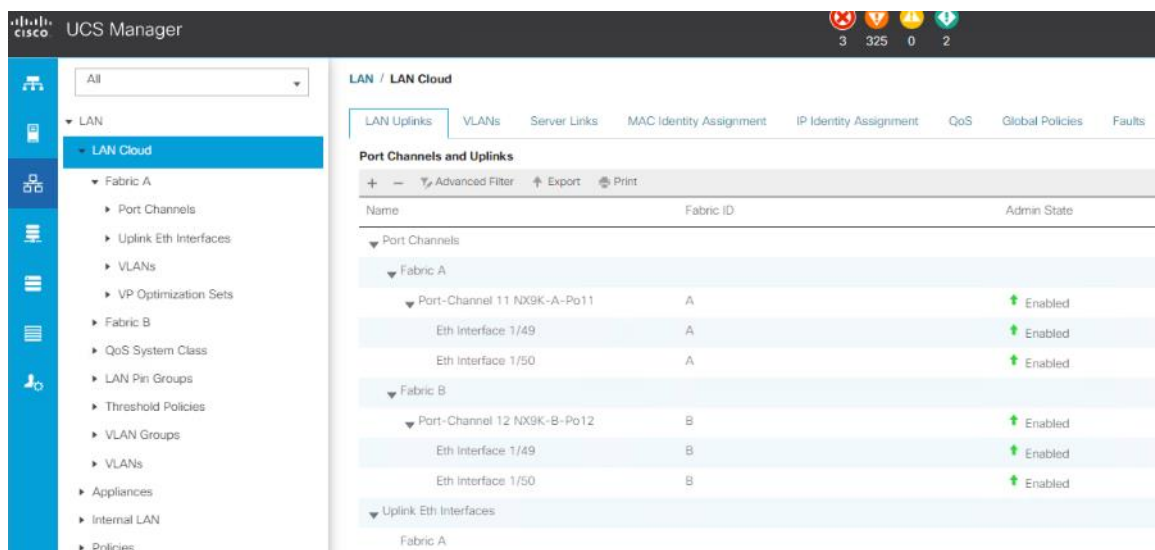
8. Click Finish.



9. Click OK.



10. Repeat steps 1-9 for the Port Channel configuration on FI-B.



Configure VLAN

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter InBand-Mgmt as the name of the VLAN to be used for Public Network Traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter 70 as the ID of the VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK.

Create VLANs

VLAN Name/Prefix :

Multicast Policy Name : [Create Multicast Policy](#)

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
 Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs :

Sharing Type : None Primary Isolated Community

10. Repeat steps 1-9 to create required VLANs. [Figure 31](#) shows the VLANs configured for this solution.

Figure 31. VLANs Configured for this Solution

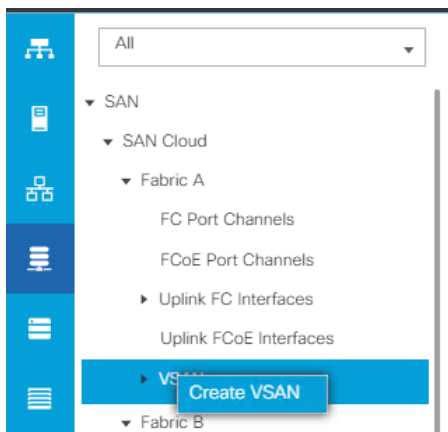
Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN InBand-Mgmt (70)	70	Lan	Ether	No	None
VLAN Infra-Mgmt (71)	71	Lan	Ether	No	None
VLAN Launcher (76)	76	Lan	Ether	No	None
VLAN VM-Network (72)	72	Lan	Ether	No	None
VLAN vMotion (73)	73	Lan	Ether	No	None

IMPORTANT! Create both VLANs with global access across both fabric interconnects. This makes sure the VLAN identity is maintained across the fabric interconnects in case of a NIC failover.

Configure VSAN

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Click SAN > SAN Cloud.
3. Under VSANs, right-click VSANs.
4. Click Create VSANs.



5. Enter the name of the VSAN, such as FlashStack-A.

Note: In this solution, we created two VSANs; VSAN FlashStack-A 100 on the Cisco UCS Fabric A and VSAN FlashStack-B 101 on the Cisco UCS Fabric B for SAN Boot and Storage Access.

6. Select Disabled for FC Zoning.

Note: In this solution we used two Cisco MDS 9132T 32-Gb switches that provide Fibre Channel zoning.

7. Select Fabric A for the scope of the VSAN:
 - Enter 100 as VSAN ID and FCoE VLAN ID.
 - Click OK.

Create VSAN ? X

Name :

FC Zoning Settings

FC Zoning : Disabled Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A. A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VSAN ID that maps to this VSAN. Enter the VLAN ID that maps to this VSAN.

VSAN ID : FCoE VLAN :

8. Repeat steps 1-7 to create the VSANs necessary for this solution.

[Figure 32](#) shows VSAN 100 and 101 configured for this solution.

Figure 32. VSANs Configured for this solution

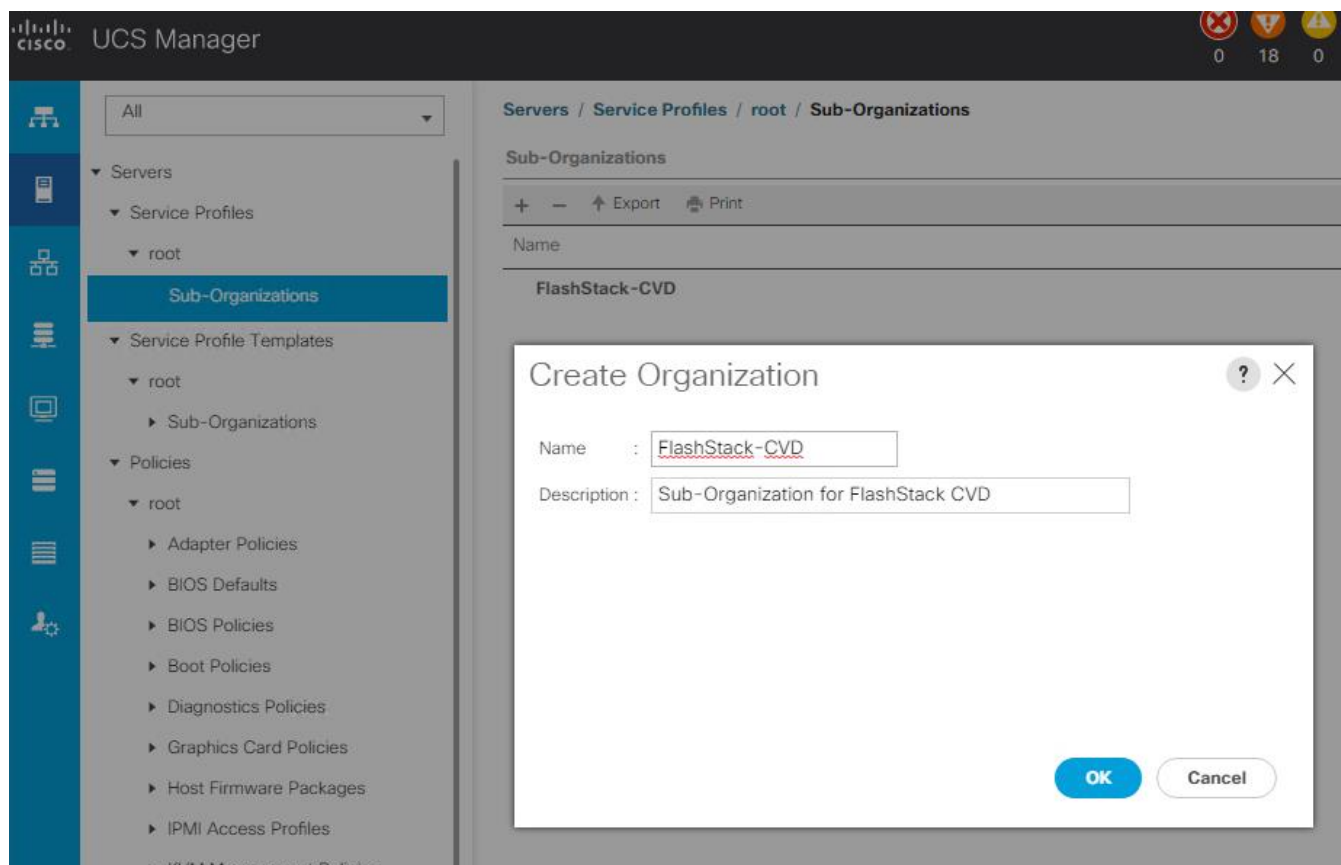
Name	ID	Fabric ID	If Type	If Role	Transport	FCoE VLAN ID	Operational State
Fabric A							
VSANs							
VSAN FlashStack-A (100)	100	A	Virtual	Network	Fc	100	OK
Fabric B							
VSANs							
VSAN FlashStack-B (101)	101	B	Virtual	Network	Fc	101	OK
VSANs							
VSAN default (1)	1	Dual	Virtual	Network	Fc	4048	OK

Create New Sub-Organization

To configure the necessary Sub-Organization for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click root > Sub-Organization.

3. Right-click Sub-Organization.
4. Enter the name of the Sub-Organization.
5. Click OK.



Note: You will create pools and policies required for this solution under the newly created “FlashStack-CVD” sub-organization.

Configure IP, UUID, Server, MAC, WWNN, and WWPN Pools

IP Pool Creation

An IP address pool on the out of band management network must be created to facilitate KVM access to each compute node in the Cisco UCS domain. To create a block of IP addresses for server KVM access in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Click Pools > root > Sub-Organizations > FlashStack-CVD > IP Pools > click Create IP Pool.
3. Select the option Sequential to assign IP in sequential order then click Next.

Create IP Pool

Name : FlashStack-KVMPool

Description :

Assignment Order : Default Sequential

4. Click Add IPv4 Block.
5. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information as shown below.

Create Block of IPv4 Addresses

From : 10.29.164.166 Size : 32

Subnet Mask : 255.255.255.0 Default Gateway : 10.29.164.1

Primary DNS : 0.0.0.0 Secondary DNS : 0.0.0.0

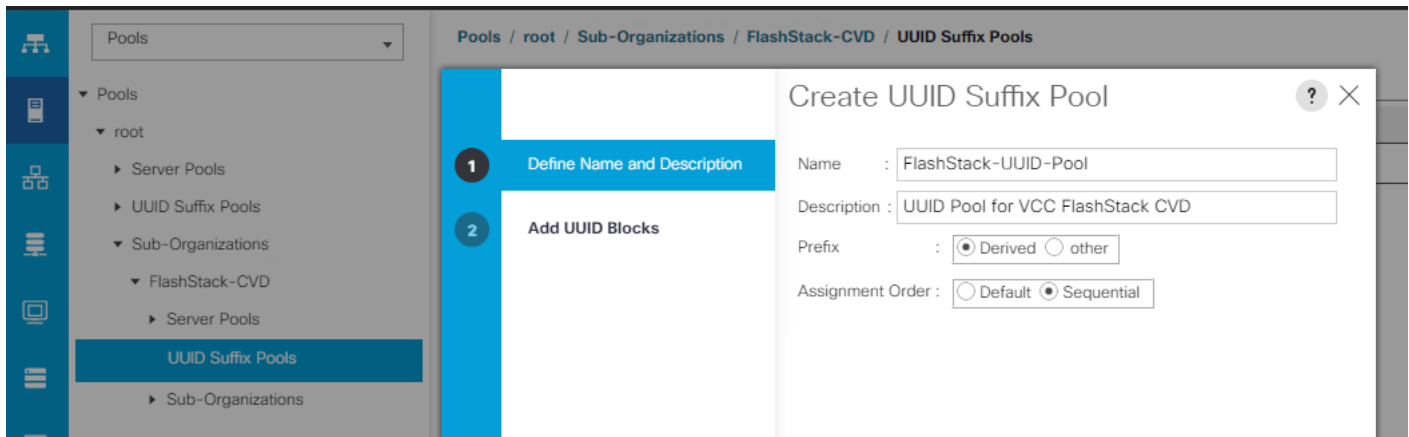
OK Cancel

➕ Add - Delete


UUID Suffix Pool Creation

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack-CVD.
3. Right-click UUID Suffix Pools and then select Create UUID Suffix Pool.
4. Enter the name of the UUID name.
5. Optional: Enter a description for the UUID pool.
6. Keep the prefix at the derived option and select Sequential in as Assignment Order then click Next.



7. Click Add to add a block of UUIDs.
8. Create a starting point UUID as per your environment.
9. Specify a size for the UUID block that is sufficient to support the available blade or server resources.

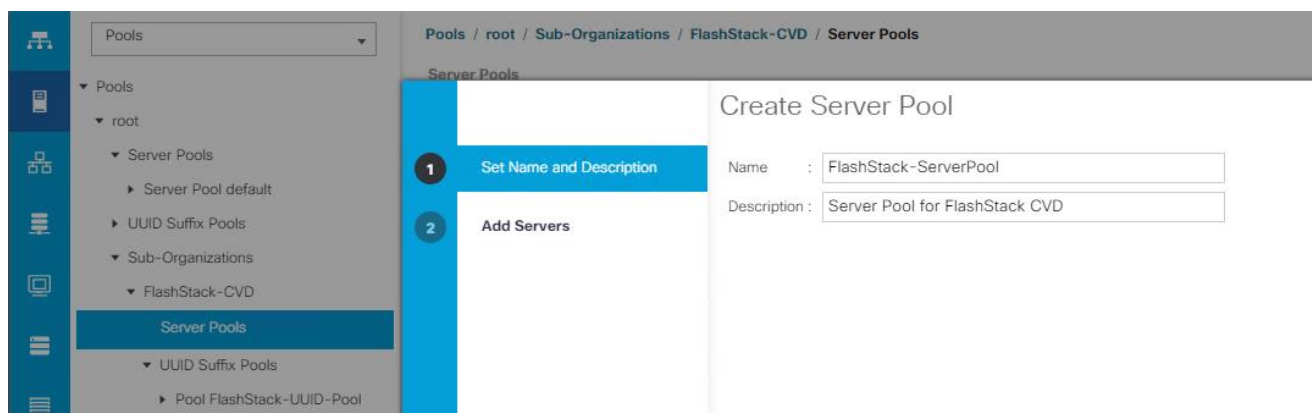


Server Pool Creation

To configure the necessary server pool for the Cisco UCS environment, follow these steps:

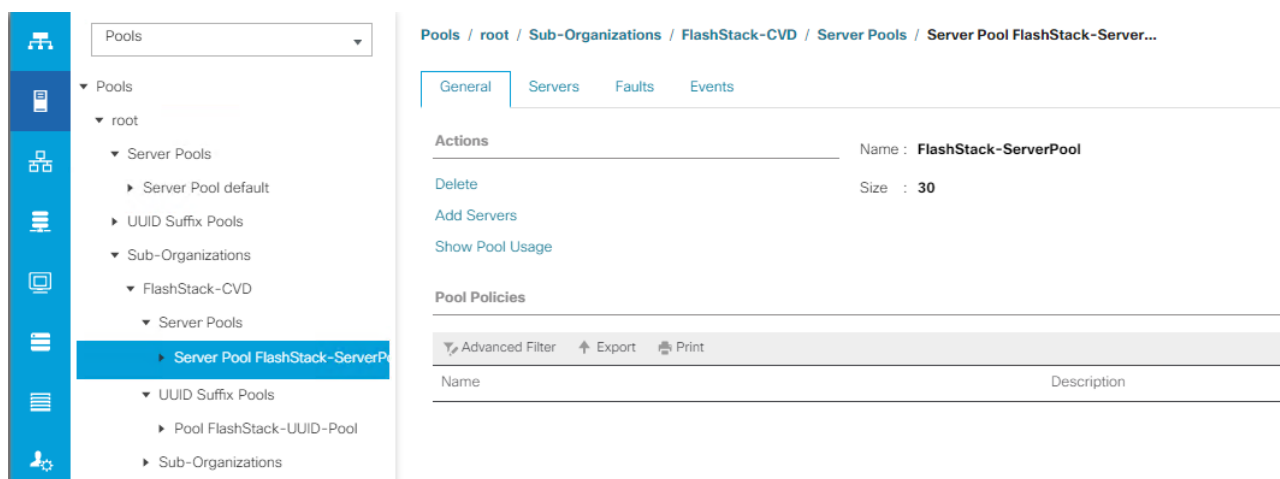
Note: Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack-CVD > right-click Server Pools > Select Create Server Pool.
3. Enter the name of the server pool.
4. Optional: Enter a description for the server pool then click Next.



5. Select the servers to be used for the deployment and click > to add them to the server pool. In our case we added thirty servers in this server pool.

6. Click Finish and then click OK.

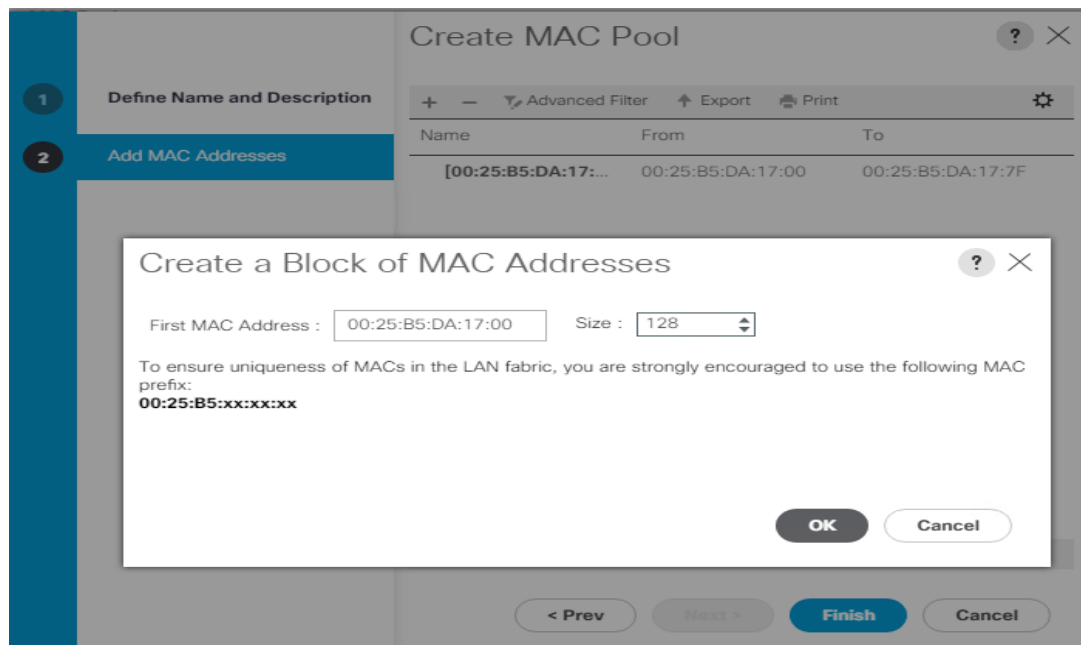


MAC Pool Creation

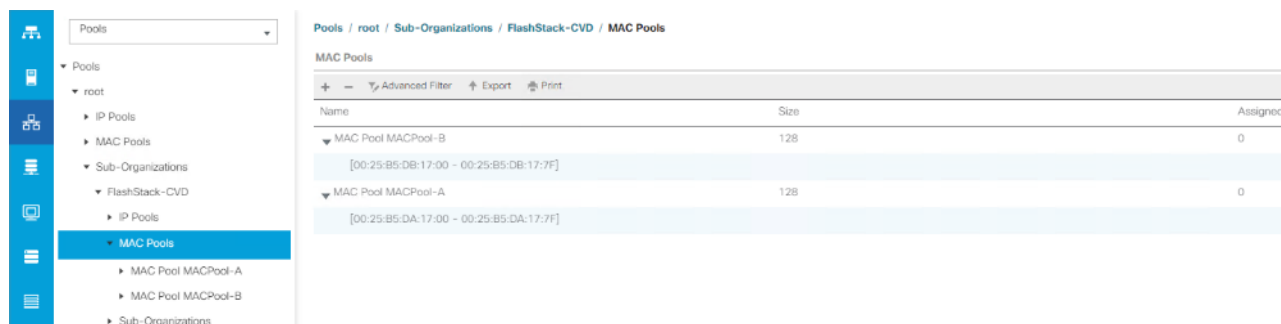
To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack > right-click MAC Pools under the root organization.
3. Click Create MAC Pool to create the MAC address pool.
4. Enter name for MAC pool. Select Assignment Order as Sequential.
5. Enter the seed MAC address and provide the number of MAC addresses to be provisioned.
6. Click OK and then click Finish.

7. In the confirmation message, click OK.



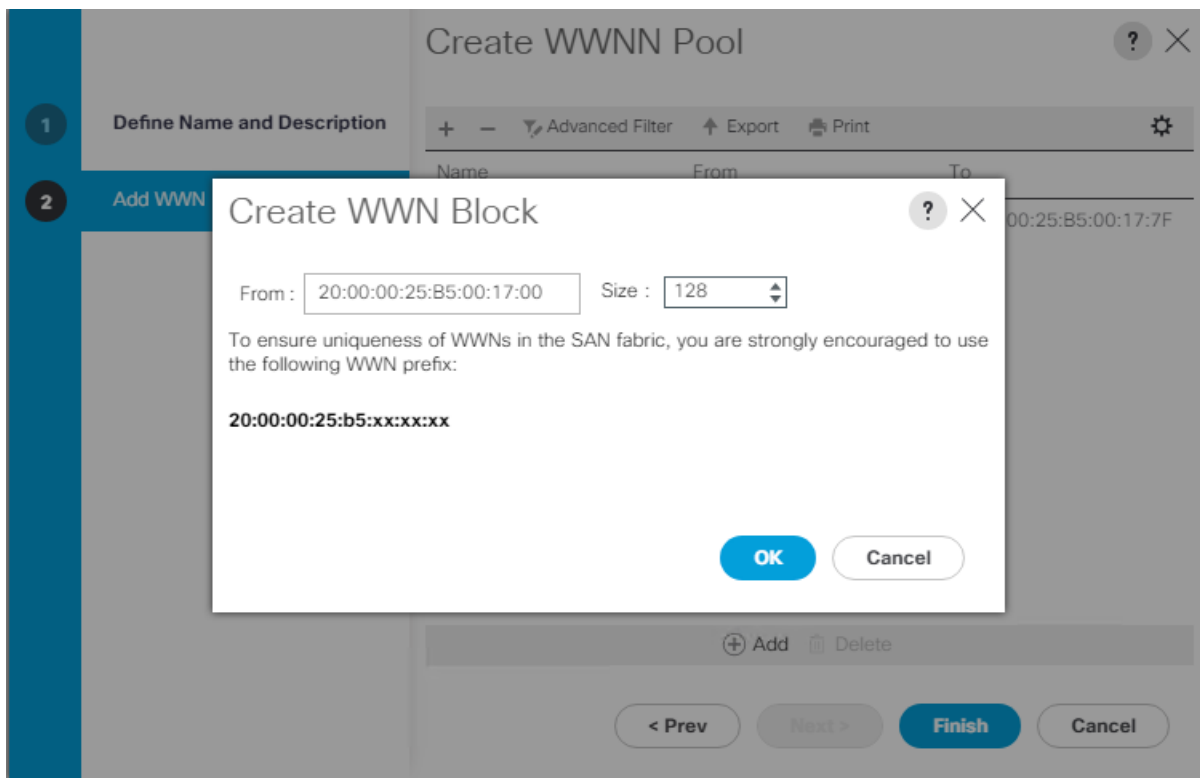
8. Create MAC Pool B and assign unique MAC Addresses as shown below.



WWNN and WWPN Pool Creation

To configure the necessary WWNN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Click Pools > Root > Sub-Organization > FlashStack-CVD > WWNN Pools > right-click WWNN Pools > select Create WWNN Pool.
3. Assign name and Assignment Order as sequential.
4. Click Next and then click Add to add block of Ports.
5. Enter Block for WWN and size of WWNN Pool as shown below.

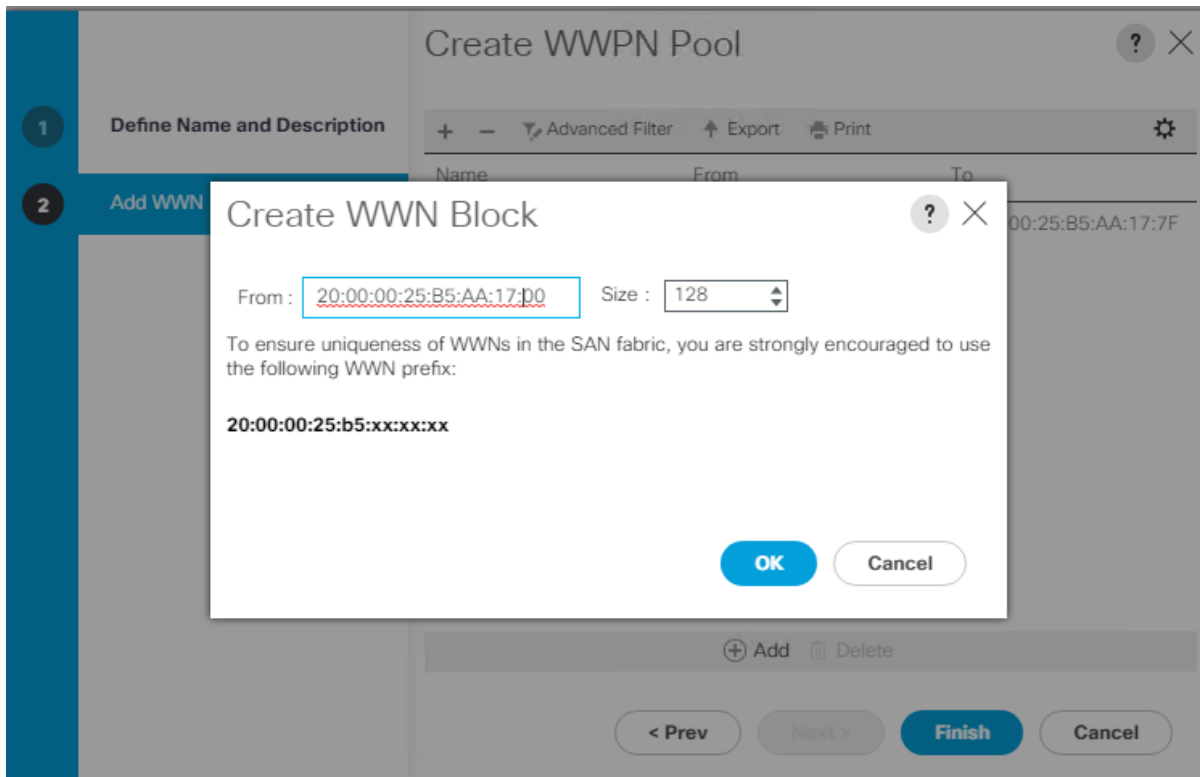


6. Click OK and then click Finish.

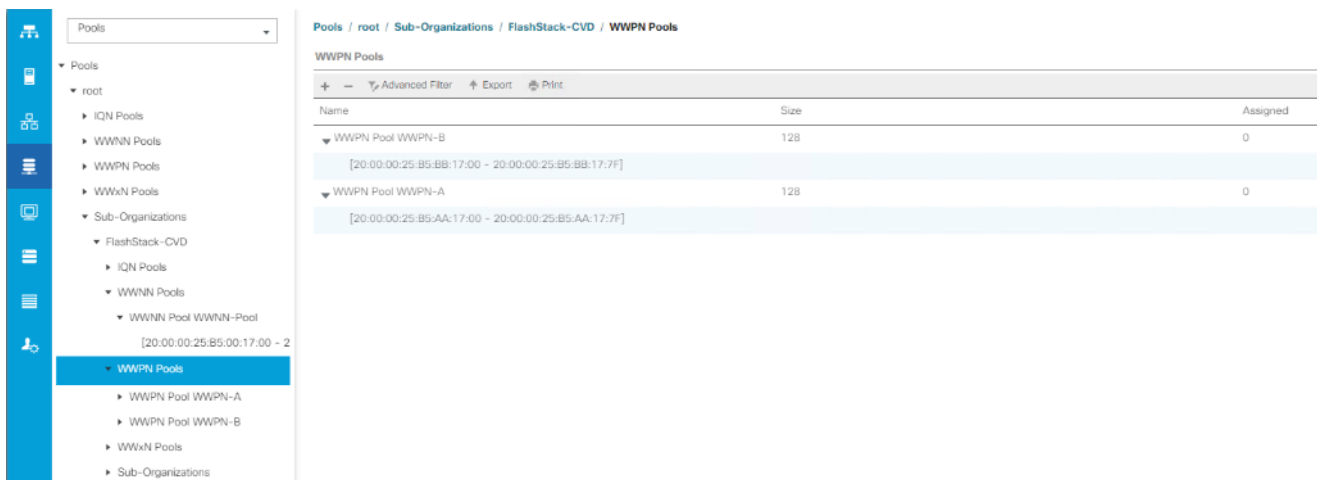
To configure the necessary WWPN pools for the Cisco UCS environment, follow these steps:

Note: We created two WWPN as WWPN-A Pool and WWPN-B as World Wide Port Name as shown below. These WWNN and WWPN entries will be used to access storage through SAN configuration.

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Pools > Root > WWPN Pools > right-click WWPN Pools > select Create WWPN Pool.
3. Assign name and Assignment Order as sequential.
4. Click Next and then click Add to add block of Ports.
5. Enter Block for WWN and size.
6. Click OK and then click Finish.



7. Configure the WWPN-B Pool and assign the unique block IDs as shown below.



Set Jumbo Frames in both the Cisco Fabric Interconnect

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click LAN > LAN Cloud > QoS System Class.

3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes.
6. Click OK.

The screenshot shows the 'LAN Cloud / QoS System Class' configuration page. The 'General' tab is selected. The 'Properties' section shows the 'Owner' as 'Local'. A table lists various QoS classes with their configurations. The 'Best Effort' class is highlighted with a red box.

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

Create Host Firmware Package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select root > Sub-Organization > FlashStack-CVD > Host Firmware Packages.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter name of the host firmware package.
6. Leave Simple selected.
7. Select the version 4.2(1f) for both the Blade Package.
8. Click OK to create the host firmware package.

Create Host Firmware Package ? X

Name :

Description :

How would you like to configure the Host Firmware Package?

Simple Advanced

Blade Package :

Rack Package :

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

<input type="checkbox"/>	Adapter
<input type="checkbox"/>	BIOS
<input type="checkbox"/>	Board Controller
<input type="checkbox"/>	CIMC
<input type="checkbox"/>	FC Adapters
<input type="checkbox"/>	Flex Flash Controller
<input type="checkbox"/>	GPUs
<input type="checkbox"/>	HBA Option ROM
<input type="checkbox"/>	Host NIC
<input type="checkbox"/>	Host NIC Option ROM
<input checked="" type="checkbox"/>	Local Disk
<input type="checkbox"/>	NVME Mswitch Firmware
<input type="checkbox"/>	PSU
<input type="checkbox"/>	PCI Switch Firmware

Create Server Pool Policy

Creating the server pool policy requires you to create the Server Pool Policy and Server Pool Qualification Policy.

To create a Server Pools Policy, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack-CVD > Server Pools.
3. Right-click Server Pools Select Create Server Pools Policy; Enter Policy name.
4. Select server from left pane to add as pooled server.

Note: In our case, we created two server pools policies. For the HOST-FCP-A policy, we added Servers as Chassis 1 Slot 1-8 and Chassis 3 Slot 1-8 and for the VDI-CVD02 policy, we added Chassis 2 Slot 1-8 and Chassis 4 Slot 1-8.

Pools / root / Sub-Organizations / FlashStack-CVD / Server Pools

Server Pools

Name	Size	Assigned
Server Pool VCC-CVD01	16	16
Server Pool VCC-CVD02	16	16

Create Server Pool Policy Qualifications

To create a Server Pool Policy Qualification Policy, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack-CVD > Server Pool Policy Qualification.
3. Right-click Server Pools Select Create Server Pool Policy Qualification; Enter Policy name.
4. Select Chassis/Server Qualification from left pane to add in Qualifications.
5. Click Add or OK to either Add more servers to existing policy to Finish creation of Policy.

Create Server Pool Policy Qualification



Naming

Name : VCC-CVD01-Qual

Description :

This server pool policy qualification will apply to new or re-discovered servers. Existing servers are not qualified until they are re-discovered

Actions

- Create Adapter Qualifications
- Create Chassis/Server Qualifications**
- Create Memory Qualifications
- Create CPU/Cores Qualifications
- Create Storage Qualifications
- Create Server PID Qualifications
- Create Power Group Qualifications
- Create Rack Qualifications

Qualifications

Name	Max	Model	From	To	Architec...	Speed	Stepping	Power G...
Chassis id range [1 - 1]								

+ Add - Delete i Info

Note: In our case, we created two server pools policies. For the HOST-FCP-A policy, we added Servers as Chassis 1 Slot 1-8 and Chassis 3 Slot 1-8 and for the “VDI-CVD02” policy, we added Chassis 2 Slot 1-8 and Chassis 4 Slot 1-8.

Policies / root / Sub-Organizations / FlashStack-CVD / Server Pool Policy Qualifications

Server Pool Policy Qualifications

Name	Max	Model	From	To
▼ VCC-CVD01-Qual				
Chassis id range [1 - 1]			1	1
Chassis id range [3 - 3]			3	3
▼ VCC-CVD02-Qual				
Chassis id range [2 - 2]			2	2
Chassis id range [4 - 4]			4	4

To create a Server Pool Policy, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Pools > root > Sub-Organization > FlashStack-CVD > Server Pool Policies.
3. Right-click Server Pool Policies and Select Create Server Pool Policy; Enter Policy name.
4. Select Target Pool and Qualification from the drop-down list.
5. Click OK.

Create Server Pool Policy ? ×

Name :

Description :

Target Pool :

Qualification :

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.

2. Click Policies > root > Sub-Organization > FlashStack-CVD > Network Control Policies.
3. Right-click Network Control Policies.
4. Click Create Network Control Policy.
5. Enter policy name.
6. Select the Enabled option for CDP.
7. Click OK to create the network control policy.

Create Network Control Policy

Name : CDP_Enabled

Description :

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

OK Cancel

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Policies > root > Sub-Organization > FlashStack-CVD > Power Control Policies.
3. Right-click Power Control Policies.
4. Click Create Power Control Policy.
5. Select Fan Speed Policy as Max Power.
6. Enter NoPowerCap as the power control policy name.
7. Change the power capping setting to No Cap.

8. Click OK to create the power control policy.

Create Power Control Policy ? X

Name :

Description :

Fan Speed Policy :

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

No Cap cap

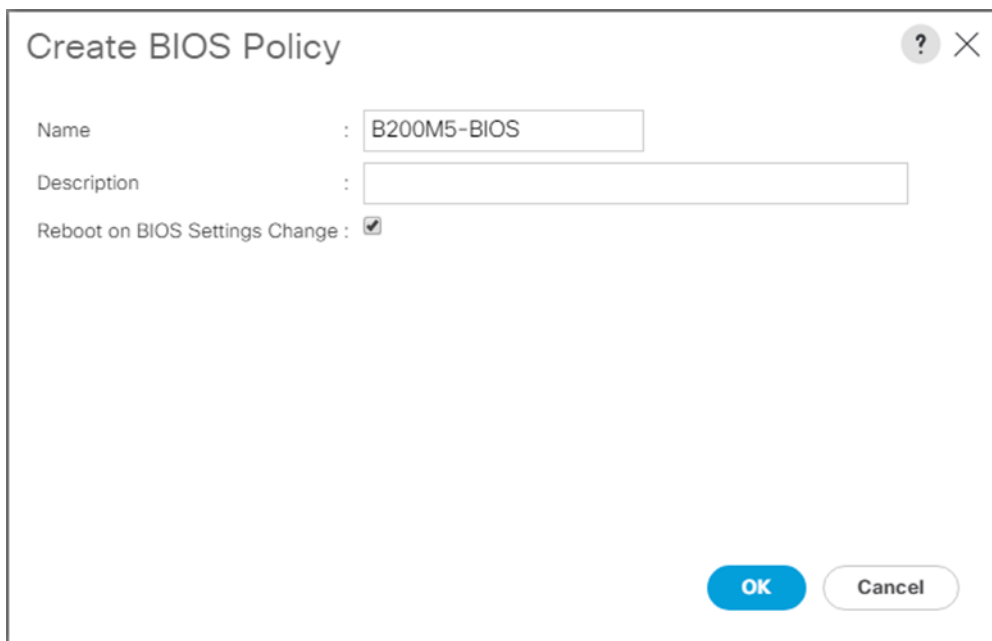
Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK **Cancel**

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Policies > root > Sub-Organization > FlashStack-CVD > BIOS Policies.
3. Right-click BIOS Policies.
4. Click Create BIOS Policy.
5. Enter B200-M6-BIOS as the BIOS policy name.
6. Click OK to create policy.



Create BIOS Policy

Name : B200M5-BIOS

Description :

Reboot on BIOS Settings Change :

OK Cancel

7. Leave all BIOS Settings as Platform Default.

Configure Maintenance Policy

To update the default Maintenance Policy, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Click Policies > root > Sub-Organization > FlashStack-CVD > Maintenance Policies.
3. Right-click Maintenance Policies to create a new policy.
4. Enter name for Maintenance Policy
5. Change the Reboot Policy to User Ack.
6. Click Save Changes.
7. Click OK to accept the change.

General Events

Actions

Delete
 Show Policy Usage
 Use Global

Properties

Name : **UserAck**
 Description :
 Owner : **Local**
 Soft Shutdown Timer : 150 Secs
 Storage Config. Deployment Policy : Immediate User Ack
 Reboot Policy : Immediate User Ack Timer Automatic
 On Next Boot (Apply pending changes at next reboot.)

Create vNIC Templates

A total of 4 vNIC Templates will be created. Two of the vNIC templates (vSwitch0-A and vSwitch0-B) will be created for vNICs to connect to VMware ESXi vSwitch0. vSwitch0 will have port groups for the IB-MGMT and OOB-MGMT. The third and fourth vNIC templates (vDS0-A and vDS0-B) will be created for vNICs to connect to the VMware Virtual Distributed Switch (vDS0). The vDS will have port groups for the vMotion and VM-Traffic VLANs.

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Click Policies > root > Sub-Organization > FlashStack-CVD > vNIC Template.
3. Right-click vNIC Templates.
4. Click Create vNIC Template.
5. Enter name for vNIC template.
6. Keep Fabric A selected. Do not select the Enable Failover checkbox.
7. For Redundancy Type, select Primary Template. Leave the Peer Redundancy Template set to <not set>
8. Select Updating Template as the Template Type.
9. Under VLANs, select the checkboxes for desired VLANs to add as part of the vNIC Template.
10. Set Native-VLAN as the native VLAN.
11. For MTU, enter 9000.

12. In the MAC Pool list, select MAC Pool configure for Fabric A.
13. In the Network Control Policy list, select CDP_Enabled.
14. Click OK to create the vNIC template.

Create vNIC Template ? X

Template Type : Initial Template Updating Template

VLANs

VLAN Groups

Advanced Filter Export Print ⚙️

Select	Name	Native VLAN	VLAN ID
<input checked="" type="checkbox"/>	K-23_CI-InBand-Mgmt	<input type="radio"/>	70
<input type="checkbox"/>	K-23_CI-Infra-Mgmt	<input type="radio"/>	71
<input type="checkbox"/>	K-23_CI-VM-Network	<input type="radio"/>	72
<input type="checkbox"/>	K-23_CI-vMotion	<input type="radio"/>	73
<input checked="" type="checkbox"/>	OOB-Mgmt	<input type="radio"/>	132
<input type="checkbox"/>	vm-network	<input type="radio"/>	54

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

15. Repeat steps 1-14 to create a vNIC Template for Fabric B. For Peer redundancy Template, select vNIC-Template-A created in the previous step.

Create vNIC Template ? X

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : vSwitch0-A ▼

Target

Adapter
 VM

Warning

If **VM** is selected, a port profile by the same name will be created.
 If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs VLAN Groups

▼ Advanced Filter
↑ Export
Print
⚙

Select	Name	Native VLAN	VLAN ID
<input checked="" type="checkbox"/>	K-23_CI-VM-Network	<input type="radio"/>	72
<input checked="" type="checkbox"/>	K-23_CI-vMotion	<input type="radio"/>	73
<input type="checkbox"/>	OOB-Mgmt	<input type="radio"/>	132
<input type="checkbox"/>	vm-network	<input type="radio"/>	54
<input type="checkbox"/>	VML-ib-mgmt	<input type="radio"/>	511

OK
Cancel

16. Verify that vNIC-Template-A Peer Redundancy Template is set to vNIC-Template-B.

17. Repeat same steps to create another set of the adapter templates

Create vHBA Templates

Two vHBAs (vHBA-A and vHBA-B) will be created for boot from SAN connectivity. To create multiple virtual host bus adapter (vHBA) templates for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Click Policies > root > Sub-Organization > FlashStack-CVD > vHBA Template.

3. Right-click vHBA Templates.
4. Click Create vHBA Template.
5. Enter vHBA-A as the vHBA template name.
6. Keep Fabric A selected.
7. Select VSAN created for Fabric A from the drop-down list.
8. Change to Updating Template.
9. For Max Data Field keep 2048.
10. Select WWPN Pool for Fabric A (created earlier) for our WWPN Pool.
11. Leave the remaining fields as is.
12. Click OK.

Create vHBA Template



Name : vHBA-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : FlashStack-A [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-A(128/128) ▼

QoS Policy : <not set> ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

OK

Cancel

13. Repeat steps 1-12 to create a vHBA Template for Fabric B.

Create Server Boot Policy for SAN Boot

All Cisco UCS B200 M6 Blade Servers for the workload and the two Infrastructure servers were set to boot from SAN for this Cisco Validated Design as part of the Service Profile template. The benefits of booting from SAN are numerous; disaster recovery, lower cooling, and power requirements for each server since a local drive is not required, and better performance, to name just a few.

Note: We strongly recommend using “Boot from SAN” to realize the full benefits of Cisco UCS stateless computing features, such as service profile mobility.

This process applies to a Cisco UCS environment in which the storage SAN ports are configured as explained in the following section.

A Local disk configuration for the Cisco UCS is necessary if the servers in the environments have a local disk.

To configure Local disk policy, follow these steps:

1. Go to tab Servers > Policies > root > Sub-Organization > FlashStack-CVD > right-click Local Disk Configuration Policy > Enter SAN-Boot for the local disk configuration policy name and change the mode to No Local Storage.
2. Click OK to create the policy.

Create Local Disk Configuration Policy
? X

Name :

Description :

Mode :

FlexFlash

FlexFlash State : Disable Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately. Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : Disable Enable

FlexFlash Removable State : Yes No No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily. Please ensure SD cards are not in use before changing the FlexFlash Removable State.

As shown in the screenshot below, the Pure Storage FlashArray have four active FC connections that pair with the Cisco MDS 9132T 32-Gb switches. Two FC ports are connected to Cisco MDS-A and the other Two FC ports are connected to Cisco MDS-B Switches. All FC ports are 32 Gb/s. The SAN Port CT0.FC0 of Pure Storage FlashArray Controller 0 is connected to Cisco MDS Switch A and SAN port CT0.FC2 is connected to MDS Switch B. The SAN Port CT1.FC0 of Pure Storage FlashArray Controller 1 is connected to Cisco MDS Switch A and SAN port CT1.FC2 connected to MDS Switch B.

FC Port	Name	Speed	Fallover	FC Port	Name	Speed	Fallover
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	

Create SAN Policy A

The FLASHSTACK-SAN-A boot policy configures the SAN Primary's primary-target to be port CT0.FC0 on the Pure Storage cluster and SAN Primary's secondary-target to be port CT1.FC0 on the Pure Storage cluster. Similarly, the SAN Secondary's primary-target should be port CT1.FC2 on the Pure Storage cluster and SAN Secondary's secondary-target should be port CT0.FC2 on the Pure Storage cluster.

To create Boot Policies for the Cisco UCS environments, follow these steps:

1. Log into the storage controller and verify all the port information is correct. This information can be found in the Pure Storage GUI under System > Connections > Target Ports.

Note: You have to create a SAN Primary (hba0) and a SAN Secondary (hba1) in SAN-A Boot Policy by entering WWPN of Pure Storage FC Ports.

2. Go to Cisco UCS Manager and then go to Servers > Policies > root > Sub Organization > FlashStack-CVD > Boot Policies. Right-click and select Create Boot Policy.
 - a. Enter FLASHSTACK-SAN-A for the name of the boot policy.
 - b. Optional: Enter a description for the boot policy.
 - c. Do not select the Reboot on Boot Order Change checkbox.
 - d. Choose the Uefi Boot Mode.
 - e. Choose the Boot Security checkbox.

Create Boot Policy

Name : FlashStack-San-A

Description : Used in Cisco Validated Design

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
No data available									

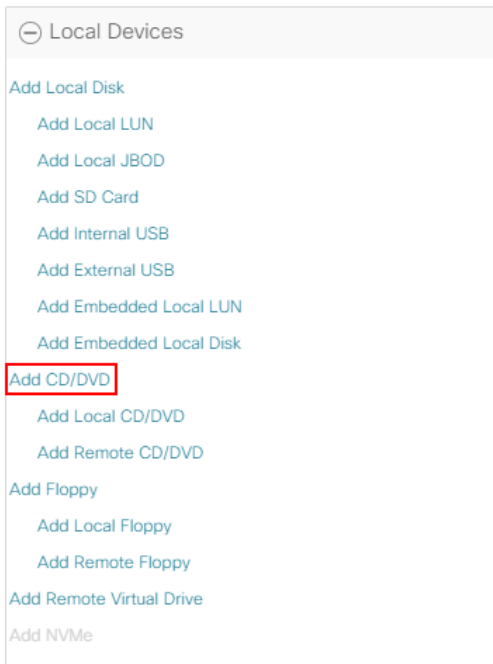
Move Up Move Down Delete

Set Uefi Boot Parameters


OK Cancel

Note: UEFI Secure Boot can be used to boot VMware ESXi 7.0 U2 with or without a TPM 2.0 module in the UCS server.

3. Expand the Local Devices drop-down list and choose Add CD/DVD.

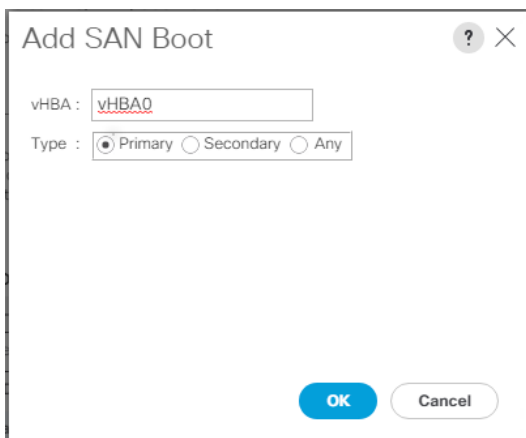


4. Expand the vHBAs drop-down list and choose Add SAN Boot.



The SAN boot paths and targets will include primary and secondary options in order to maximize resiliency and number of paths.

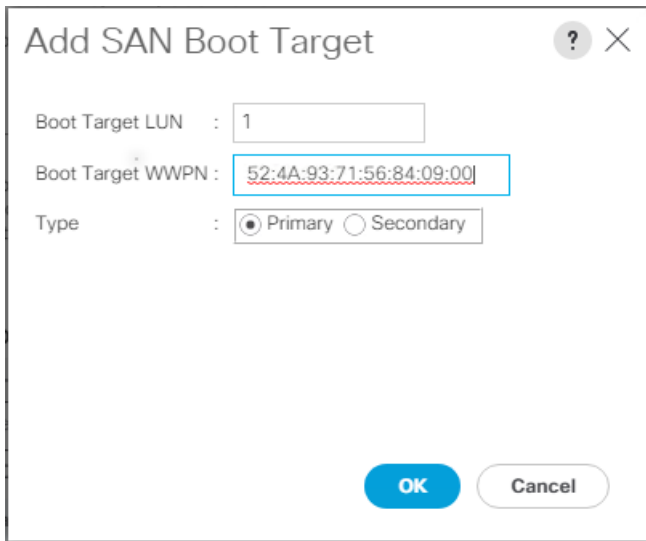
5. In the Add SAN Boot dialog box, for vHBA enter vHBA0 and for Type select Primary and. Click OK to add SAN Boot.



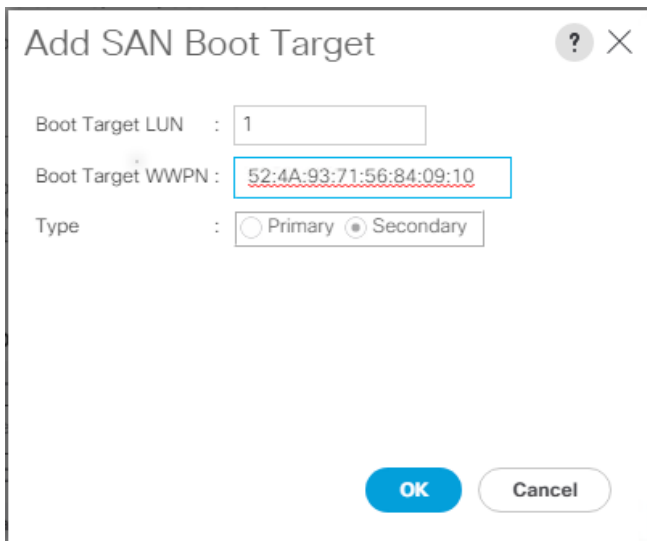
6. Click Add SAN Boot Target.



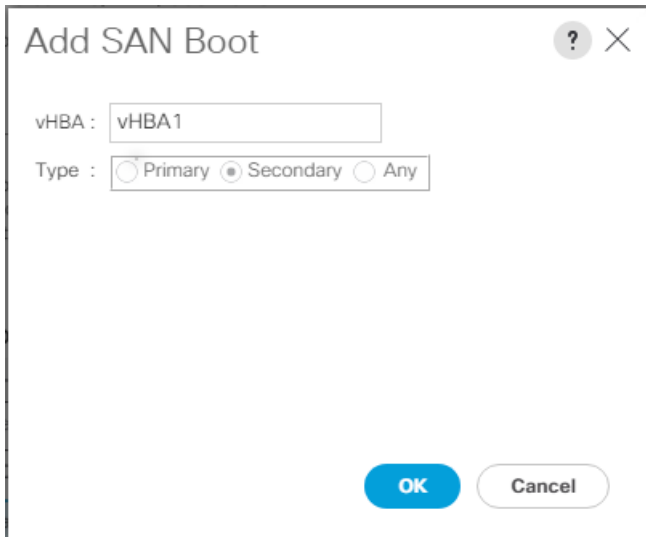
7. Keep **1** as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC0 of Pure Storage and add SAN Boot Primary Target.



8. Add a secondary SAN Boot target into same hba0, enter the boot target LUN as **1** and WWPN for FC port CT1.FC0 of Pure Storage, and add SAN Boot Secondary Target.

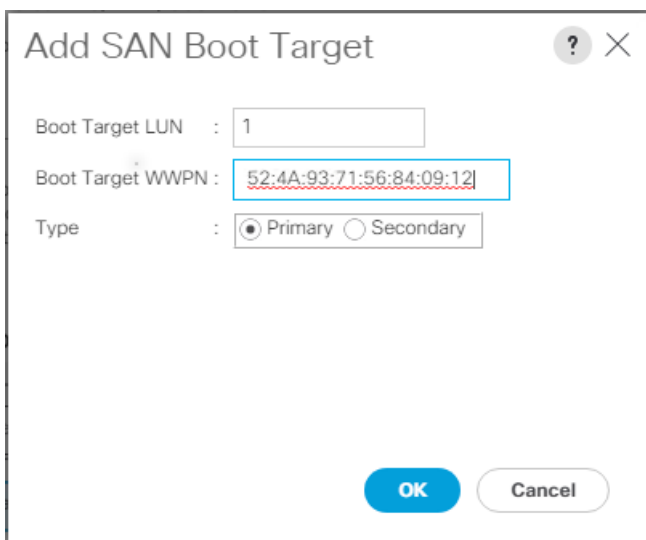


9. From the vHBA drop-down list and choose Add SAN Boot. In the Add SAN Boot dialog box, enter "vHBA1" in the vHBA field. Click OK to SAN Boot, then choose Add SAN Boot Target.



Dialog box titled "Add SAN Boot" with a help icon and close button. The "vHBA" field contains "vHBA1". The "Type" field has radio buttons for "Primary", "Secondary", and "Any", with "Secondary" selected. There are "OK" and "Cancel" buttons at the bottom.

10. Keep 1 as the value for the Boot Target LUN. Enter the WWPN for FC port CT1.FC2 of Pure Storage and add SAN Boot Primary Target.



Dialog box titled "Add SAN Boot Target" with a help icon and close button. The "Boot Target LUN" field contains "1". The "Boot Target WWPN" field contains "52:4A:93:71:56:84:09:12". The "Type" field has radio buttons for "Primary" and "Secondary", with "Primary" selected. There are "OK" and "Cancel" buttons at the bottom.

11. Add a secondary SAN Boot target into same vHBA1 and enter the boot target LUN as 1 and WWPN for FC port CT0.FC2 of Pure Storage and add SAN Boot Secondary Target.

Add SAN Boot Target

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

12. Click Save Changes.

Servers / Policies / root / Sub-Organizations / FlashStack-CVD / Boot Policies / Boot Policy FI...

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **FlashStack-San-A**

Description :

Owner : **Local**

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

Warning

The type (primary/secondary) does not indicate a boot order presence. The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order. If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported. If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
CD/DVD	1								
San	2								

13. Expand SAN > SAN Primary and select SAN Target Primary. Select Set Uefi Boot Parameters.

Servers / Policies / root / Sub-Organizations / FlashStack-CVD / Boot Policies / Boot Policy Fl...

General Events

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

Warning

The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If Enforce vNIC/vHBA/iSCSI Name is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Or...	vNIC/...	Type	LUN ...	WWN	Slot N...	Boot ...	Boot ...	Descri...
CD/DVD	1								
San	2								
SAN-Primary		vHBA0	Primary						
SAN Target Pri...			Primary	1	52:4A:93:71:56:84:09:00				
SAN Target S...			Seco...	1	52:4A:93:71:56:84:09:10				
SAN Secondary		vHBA1	Seco...						

Move Up Move Down Delete

Set Uefi Boot Parameters

Save Changes Reset Values

Note: For Cisco UCS B200 M6 and M5, and Cisco UCS C220 M6 and M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for Cisco UCS M4 and earlier servers, VMware ESXi 7.0 and above will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

14. Fill in the Set Uefi Boot Parameters exactly as shown in the following screenshot:

Set Uefi Boot Parameters



Uefi Boot Parameters

Boot Loader Name	:	<input type="text" value="BOOTX64.EFI"/>
Boot Loader Path	:	<input "="" type="text" value="\EFI\BOOT\"/>
Boot Loader Description	:	<input type="text"/>



15. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target and click OK for the confirmation.
16. Repeat this process to set Uefi Boot Parameters for each of the 4 SAN Boot Targets.
17. Click OK, then click OK again to create the boot policy.

Create SAN Policy B

The FLASHSTACK-SAN-B boot policy configures the SAN Primary's primary-target to be port CT0.FC6 on the Pure Storage cluster and SAN Primary's secondary-target to be port CT1.FC6 on the Pure Storage cluster. Similarly, the SAN Secondary's primary-target should be port CT1.FC0 on the Pure Storage cluster and SAN Secondary's secondary-target should be port CT0.FC0 on the Pure Storage cluster.

To create boot policies for the Cisco UCS environments, follow these steps:

1. Log into the storage controller and verify all the port information is correct. This information can be found in the Pure Storage GUI under System > Connections > Target Ports.

Note: You have to create SAN Primary (vHBA1) and SAN Secondary (vHBA0) in SAN-B Boot Policy by entering WWPN of Pure Storage FC Ports as explained in the following section.

2. Go to Cisco UCS Manager and then go to tab Servers > Policies > root > Sub Organization > FlashStack-CVD > Boot Policies.

3. Right-click and select Create Boot Policy. Enter FLASHSTACK-SAN-B as the name of the boot policy.

Create Boot Policy

Name : FlashStack-San-B

Description : Used in Cisco Validated Design

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
No data available									

Move Up Move Down Delete

Set Uefi Boot Parameters

OK Cancel

4. Expand the Local Devices drop-down list and Choose Add CD/DVD. Expand the vHBAs drop-down list and choose Add SAN Boot.

Note: The SAN boot paths and targets include primary and secondary options in order to maximize resiliency and number of paths.

5. In the Add SAN Boot dialog box, for Type select Primary and name vHBA as vHBA0. Click OK to add SAN Boot.

Add SAN Boot ? X

vHBA :

Type : Primary Secondary Any

OK **Cancel**

6. Select Add SAN Boot Target to enter WWPN address of storage port. Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC2 of Pure Storage and add SAN Boot Primary Target.

Add SAN Boot Target ? X

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK **Cancel**

7. Add the secondary SAN Boot target into the same hba0; enter boot target LUN as 1 and WWPN for FC port CT0.FC0 of Pure Storage and add SAN Boot Secondary Target.

Add SAN Boot Target ? X

Boot Target LUN : 1

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

Type : Primary Secondary

OK Cancel

- From the vHBA drop-down list, choose Add SAN Boot. In the Add SAN Boot dialog box, enter "hba1" in the vHBA field. Click OK to SAN Boot, then choose Add SAN Boot Target.

Add SAN Boot ? X

vHBA : vHBA0

Type : Primary Secondary Any

OK Cancel

- Keep 1 as the value for Boot Target LUN. Enter the WWPN for FC port CT0.FC1 of Pure Storage and Add SAN Boot Primary Target.

The screenshot shows a dialog box titled "Add SAN Boot Target" with a question mark icon and a close button (X) in the top right corner. The dialog contains three input fields: "Boot Target LUN" with the value "1", "Boot Target WWPN" with the value "52:4A:93:71:56:84:09:02" (highlighted with a blue border and a red dashed underline), and "Type" with radio buttons for "Primary" (selected) and "Secondary". At the bottom, there are two buttons: "OK" (blue) and "Cancel" (white with a grey border).

10. Add secondary SAN Boot target into same hba1 and enter boot target LUN as 1 and WWPN for FC port CT1.FC1 of Pure Storage and add SAN Boot Secondary Target.

The screenshot shows a dialog box titled "Add SAN Boot Target" with a question mark icon and a close button (X) in the top right corner. The dialog contains three input fields: "Boot Target LUN" with the value "1", "Boot Target WWPN" with the value "52:4A:93:71:56:84:09:12", and "Type" with radio buttons for "Primary" and "Secondary" (selected). At the bottom, there are two buttons: "OK" (blue) and "Cancel" (white with a grey border).

11. Click OK.

Create Boot Policy

Name : SAN-B

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

Add SAN Boot

Add SAN Boot Target

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Or...	vNIC/...	Type	LUN ...	WWN	Slot N...	Boot ...	Boot ...	Descri...
CD/DVD		1							
▶ San		2							

↑ Move Up ↓ Move Down Delete

Set Uefi Boot Parameters

OK Cancel

12. Expand San > SAN Primary and select SAN Target Primary. Select Set Uefi Boot Parameters.

Servers / Policies / root / Sub-Organizations / FlashStack-CVD / Boot Policies / Boot Policy Fl...

General Events

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

Warning

The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Or...	vNIC/...	Type	LUN ...	WWN	Slot N...	Boot ...	Boot ...	Descri...
CD/DVD	1								
San	2								
SAN-Primary		vHBA0	Primary						
SAN Target Pri...			Primary	1	52:4A:93:71:56:84:09:00				
SAN Target S...			Seco...	1	52:4A:93:71:56:84:09:10				
SAN Secondary		vHBA1	Seco...						

Move Up Move Down Delete

Set Uefi Boot Parameters

Save Changes Reset Values

Note: For Cisco UCS B200 M6 and M5, and Cisco UCS C220 M6 and M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for M4 and earlier servers, VMware ESXi 7.0 and above will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

13. Fill in the Set Uefi Boot Parameters exactly as shown in the following screenshot:

Set Uefi Boot Parameters



Uefi Boot Parameters

Boot Loader Name	:	<input type="text" value="BOOTX64.EFI"/>
Boot Loader Path	:	<input "="" type="text" value="\EFI\BOOT\"/>
Boot Loader Description	:	<input type="text"/>



14. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target and click OK for the confirmation.
15. Repeat this process to set Uefi Boot Parameters for each of the 4 SAN Boot Targets.
16. Click OK, then click OK again to create the boot policy.

Note: For this solution, we created two Boot Policies, “SAN-A” and “SAN-B”. For 8 Cisco UCS B200 M6 blade servers, you will assign the first 4 Service Profiles with FLASHSTACK-SAN-A to the first 4 servers and the remaining 4 Service Profiles with FLASHSTACK-SAN-B to the remaining 4 servers as explained in the following section.

Configure and Create a Service Profile Template

Service profile templates enable policy-based server management that helps ensure consistent server resource provisioning suitable to meet predefined workload needs.

You will create two Service Profile templates; the first Service profile template “Host-FCP-AHOST-FCP-A” uses the boot policy “SAN-A” and the second Service profile template “Host-FCP-B” uses the boot policy “SAN-B” to utilize all the FC ports from Pure Storage for high-availability in case any FC links go down.

HOST-FCP-A Create Service Profile Template

To create a service profile template, follow these steps:

1. In the Cisco UCS Manager, go to Servers > Service Profile Templates > root Sub Organization > FlashStack-CVD > and right-click Create Service Profile Template.

2. Enter the Service Profile Template name, select the UUID Pool that was previously created, and click Next.

Create Service Profile (expert)

You must enter a name for the service profile. You can also specify how a UUID will be assigned to this profile and enter a description of the profile.

Name :

The service profile will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-FlashStack-CVD**

Specify how the UUID will be assigned to the server associated with this service profile.
UUID

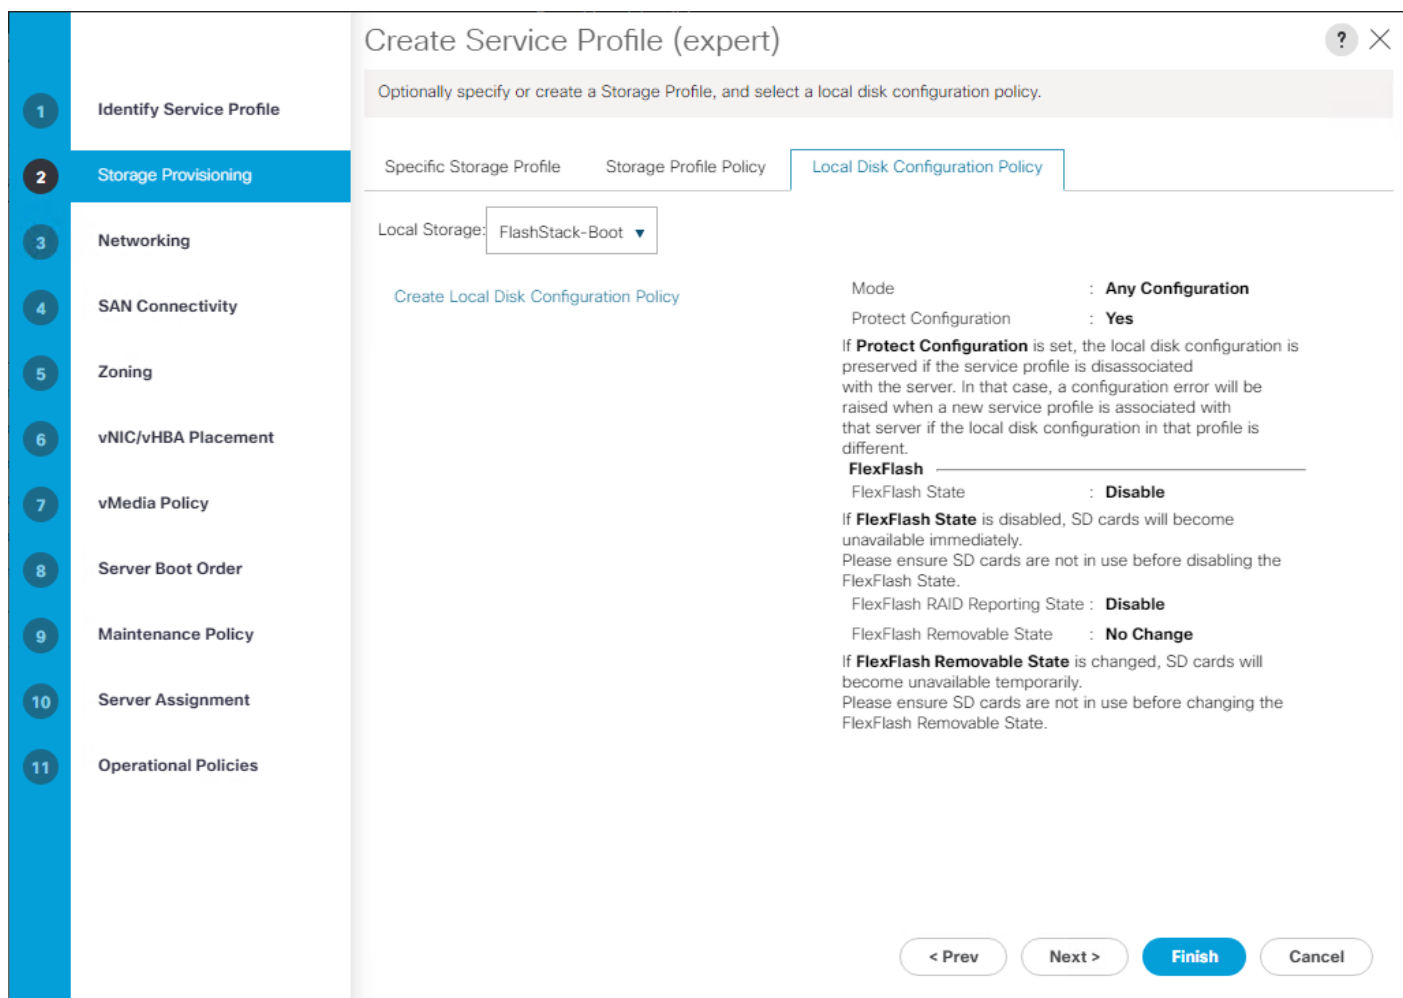
UUID Assignment:

[Create UUID Suffix Pool](#)
The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

< Prev Next > **Finish** Cancel

3. Click Local Disk Configuration Policy to SAN-Boot as No Local Storage.



4. In the networking window, select the “Use Connectivity Policy” option to configure the LAN connectivity.
5. Choose FC-Boot from the LAN Connectivity Policy drop-down list. Leave the Initiator Name Assignment as <not set>.

Create Service Profile (expert)

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Hardware Inherited Use Connectivity Policy

LAN Connectivity Policy : [Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment:

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

6. Click Next.

7. In the SAN Connectivity menu, select the Use Connectivity Policy option.

1 Identify Service Profile

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile (expert)

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

Simple Expert No vHBAs Hardware Inherited Use Connectivity Policy

SAN Connectivity Policy : [Create SAN Connectivity Policy](#)

< Prev Next > **Finish** Cancel

8. Click Next.

9. Skip zoning. For this FlashStack Configuration, the Cisco MDS 9132T 32-Gb is used for zoning.

10. Select the default option Let System Perform Placement in the Placement Selection menu.

Create Service Profile (expert) ? ×

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: [Create Placement Policy](#)

System will perform automatic placement of vNICs and vHBAs based on PCI order.

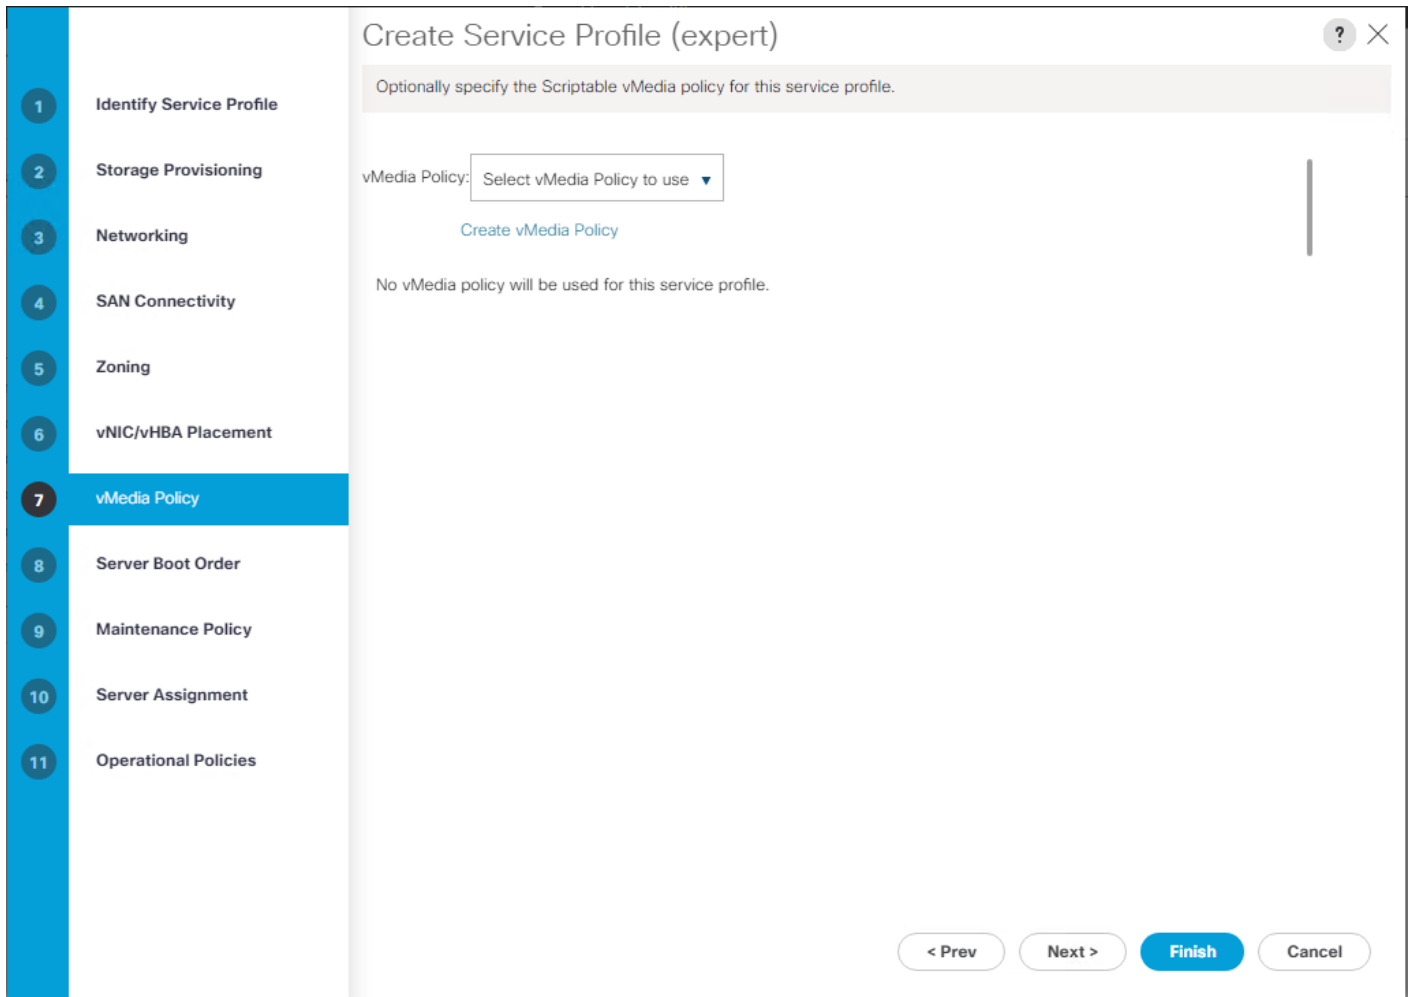
Name	Address	Order
vHBA vHBA1	Derived	1
vHBA vHBA0	Derived	2
vNIC 01-vSwitch0-B	Derived	3
vNIC 00-vSwitch0-A	Derived	4
vNIC 03-vDS0-B	Derived	5
vNIC 02-vDS0-A	Derived	6

↑ Move Up ↓ Move Down 🗑 Delete ↻ Reorder ⓘ Modify

< Prev Next > **Finish** Cancel

11. Click Next.

12. Do not select a vMedia Policy.



13. Click Next.

14. For the Server Boot Policy, select FlashStack-SAN-A, which you previously created.

Create Service Profile (expert)

Optionally specify the boot policy for this service profile.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **FlashStack-San-A**
 Description : **Used in Cisco Validated Design**
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Uefi**
 Boot Security : **Yes**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vH...	Type	LUN Name	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
CD/D...	1								
▶ San	2								

[Create ISCSI vNIC](#) [Set ISCSI Boot Parameters](#) [Set Uefi Boot Parameters](#)

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

15. Click Next.

16. Select FlashStack-UAck maintenance policy, which requires user acknowledgement prior rebooting server when making changes to policy or pool configuration tied to a service profile.

Create Service Profile (expert)

Specify how disruptive changes (such as reboot, network interruptions, firmware upgrades) should be applied to the system.

⊖ Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: FlashStack-UAck ▼ [Create Maintenance Policy](#)

Name	: FlashStack-UAck
Description	: Used in Cisco Validated Design
Soft Shutdown Timer	: 150 Secs
Storage Config. Deployment Policy	: User Ack
Reboot Policy	: User Ack

< Prev Next > **Finish** Cancel

17. Select Server Pool policy to automatically assign service profile to a server that meets the requirement for server qualification based on the pool configuration.

18. On the same page you can configure “Host firmware Package Policy” which helps to keep the firmware in sync when associated to server.

Create Service Profile (expert) ? ×

Optionally specify a server or server pool for this service profile.

You can select an existing server or server pool, or specify the physical location of the server you want to associate with this service profile.

Server Assignment: [Create Server Pool](#)

Select the power state to be applied when this profile is associated with the server.

Up Down

The service profile is not automatically associated with a server. Either select a server from the list or associate the service profile manually later.

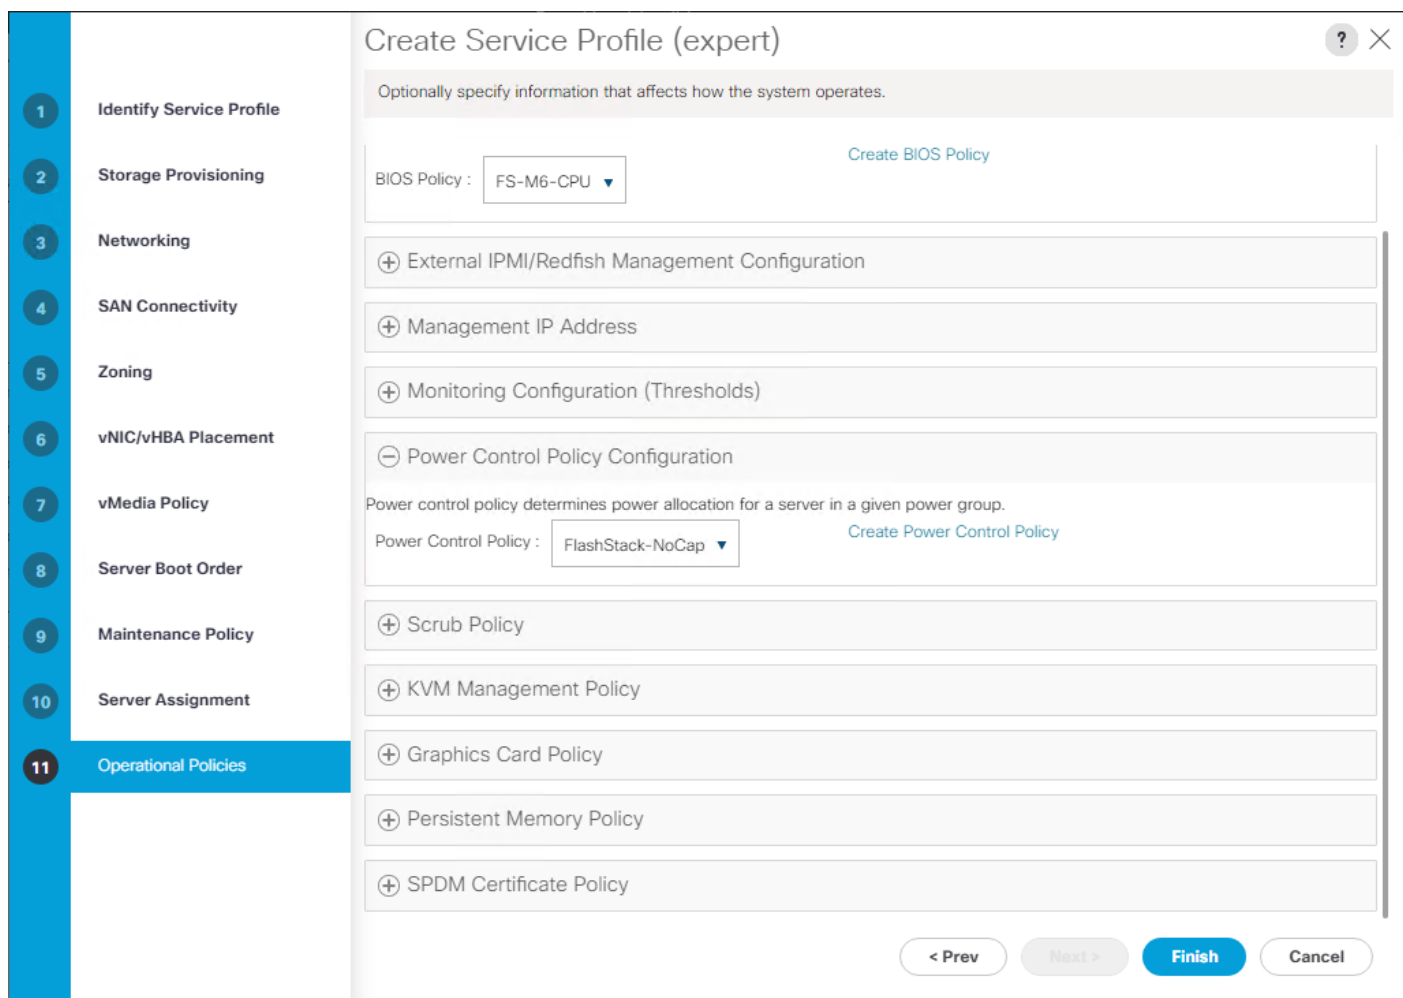
⊖ Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: [Create Host Firmware Package](#)

< Prev Next > **Finish** Cancel

Note: On the Operational Policy page, we configured the BIOS policy for the Cisco UCS B200 M6 blade server, Power Control Policy with “NoPowerCap” for maximum performance.

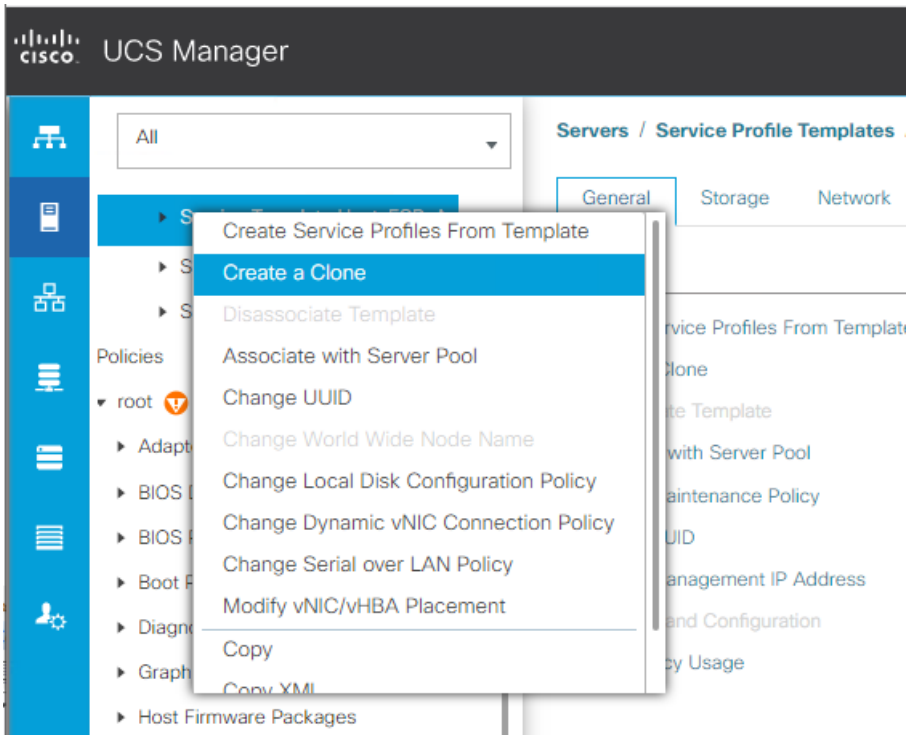


19. Finish to create the service profile template “Host-FCP-A”.

Clone Service Profile Template

To clone the Service Profile template, follow these steps:

1. In the Cisco UCS Manager, go to Servers > Service Profile Templates > root > Sub Organization > FlashStack-CVD > Service Template HOST-FCP-A and right-click Create a Clone as shown below.



2. Enter name to create Clone from existing Service Profile template. Click OK.

The image shows a dialog box titled 'Create Clone From Host-FCP-A'. It has a close button (X) in the top right corner. There are two input fields: 'Clone Name' with the value 'Host-FCP-B' and 'Org' with the value 'FlashStack-CVD'. At the bottom of the dialog, there are three buttons: 'OK', 'Cancel', and 'Help'.

Note: This HOST-FCP-B service profile template will be used to create the remaining sixteen service profiles for VDI workload and Infrastructure server02.

3. To change boot order from FLASHSTACK-SAN-A to FLASHSTACK-SAN-B for HOST-FCP-B, click Cloned Service Profile template > Select Boot Order tab. Click Modify Boot Policy.

UCS Manager

Servers / Service Profile Templates / root / Sub-Organizations / FlashStack-CVD / Service Template Host-FCP-B

General Storage Network iSCSI vNICs vMedia Policy **Boot Order** Policies Events FSM

Actions

Global Boot Policy

Modify Boot Policy

Name : **FlashStack-San-B**
 Description :
 Reboot on Boot Order Change : **No**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Uefi**
 Boot Security : **Yes**

WARNINGS:
 The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHB...	Type	LUN Name	WWN	Slot Num...	Boot Name	Boot Path	Description
CD/DVD	1								
San	2								

Create iSCSI vNIC Set iSCSI Boot Parameters Set UEFI Boot Parameters

Save Changes Reset Values

Logged in as admin@10.29.132.212 System Time: 2022-03-01T09:52

4. From the drop-down list, for the Boot Policy, select FlashStack-San-B and click OK.

Modify Boot Policy

Boot Policy: FlashStack-San-B

Name : FlashStack-San-B
 Description :
 Reboot on Boot Order Change : No
 Enforce vNIC/vHBA/iSCSI Name : Yes
 Boot Mode : Uefi
 Boot Security : Yes

WARNINGS:
 The type (primary/secondary) does not indicate the effective order of boot devices within the boot policy.
 If **Enforce vNIC/vHBA/iSCSI Name** is selected, the vNICs/vHBAs are selected by PCIe bus scan order. If it is not selected, the vNICs/vHBAs are selected by lowest PCIe bus scan order.

Boot Order

Name	Order	vNIC/vHBA/iSCSI...	Type	LUN Name	WWN	Slot Number	Boot Name	Boot Path	Description
CD/DVD	1								
▶ San	2								

Create iSCSI vNIC Set iSCSI Boot Parameters Set Uefi Boot Parameters

OK Cancel

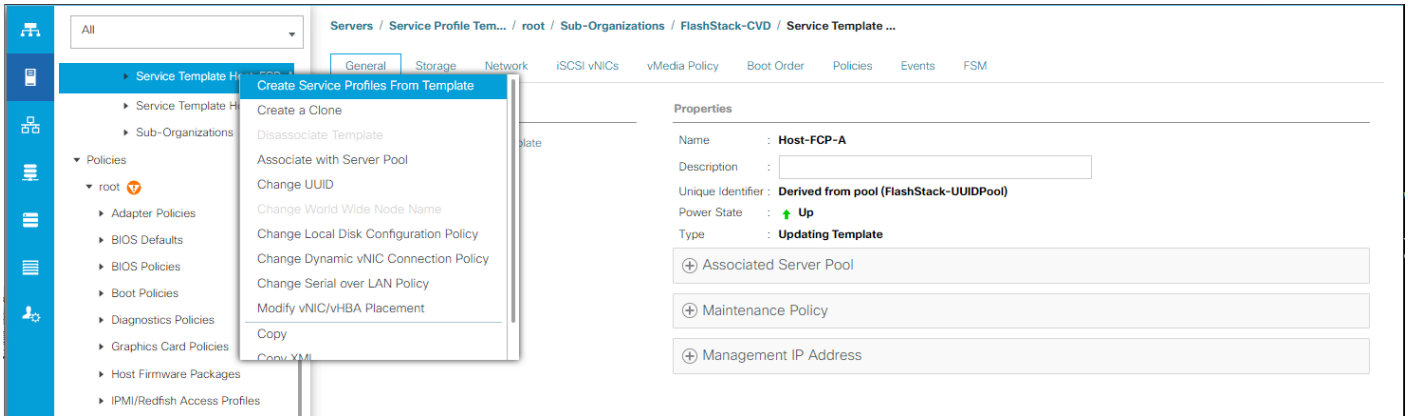
Note: You have now created the Service Profile template “HOST-FCP-A” and “HOST-FCP-B” with each having two vHBAs and four vNICs.

Create Service Profiles from Template and Associate to Servers

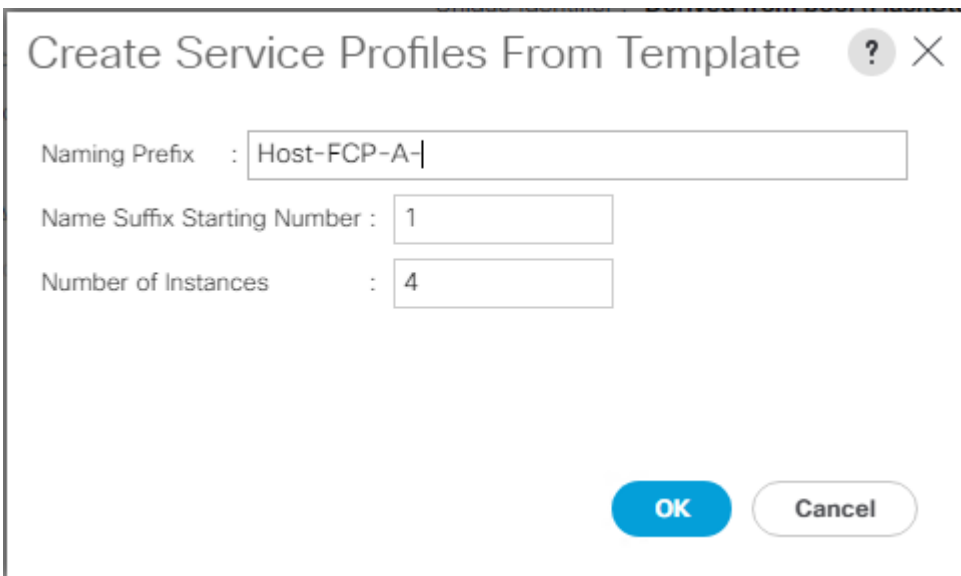
You will create 4 service profiles from the HOST-FCP-A template and 4 service profiles from the HOST-FCP-B template as explained in the following sections.

To create first four Service Profiles from Template, follow these steps:

1. Go to the Servers tab > Service Profiles > root > Sub-Organization > FlashStack-CVD and right-click Create Service Profiles from Template.



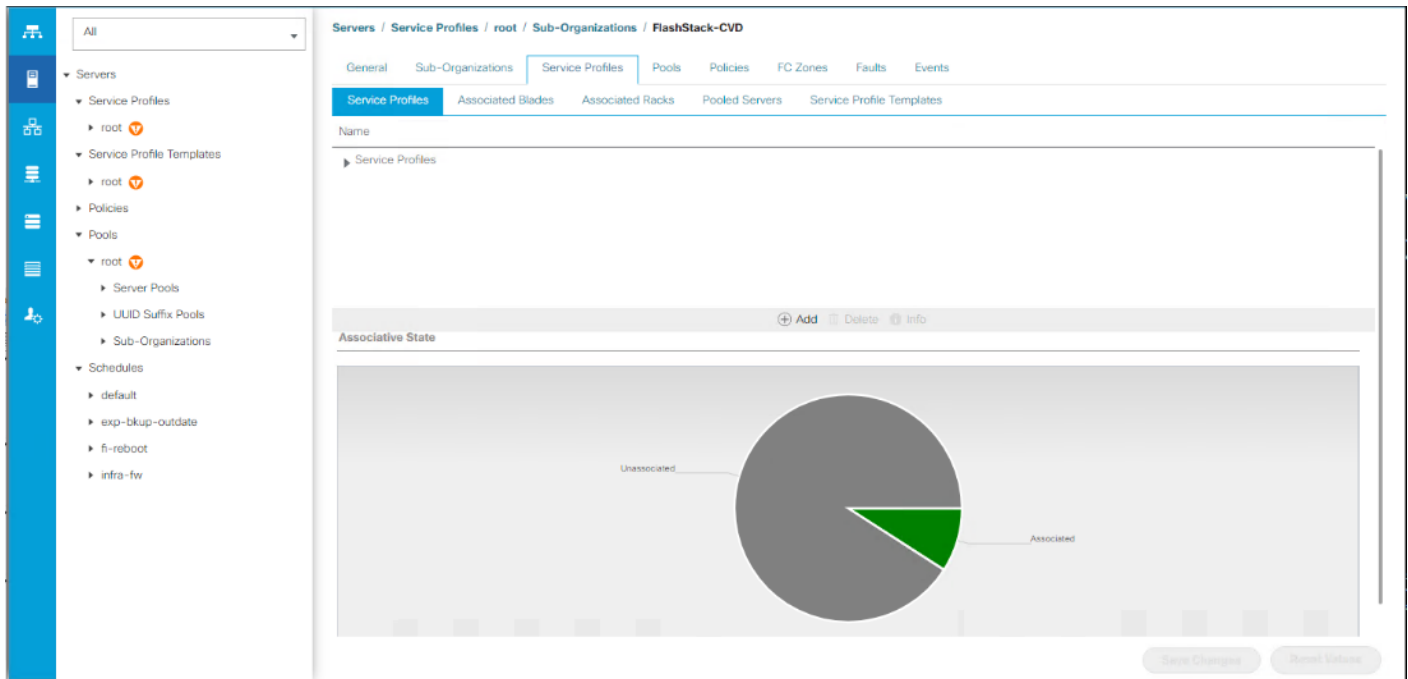
2. Select “HOST-FCP-A” for the Service profile template which you created earlier and name the service profile “Host-FCP-A-X.” To create four service profiles, enter 4 for the Number of Instances, as shown below. This process will create service profiles “Host-FCP-A-1”, “Host-FCP-A-2”, “Host-FCP-A-3” and “Host-FCP-A-4.”



3. Create the remaining four Service Profiles “Host-FCP-B-1”, “Host-FCP-B-2”, Host-FCP-B-3 and “Host-FCP-B-4” from Template “HOST-FCP-B.”

Note: When the service profiles are created, the association of Service Profile starts automatically to servers based on the Server Pool Policies if defined. Otherwise manually associate the profiles to the servers.

4. Service Profile association can be verified in Cisco UCS Manager > Servers > Service Profiles. Different tabs can provide details on Service profile association based on Server Pools Policy, Service Profile Template to which Service Profile is tied to, and so on.



Configure Cisco Nexus 93180YC-FX Switches

The following section details the steps for the Nexus 93180YC-FX switch configuration.

Configure Global Settings for Cisco Nexus A and Cisco Nexus B

To set global configuration, follow these steps on both Cisco Nexus switches:

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
```

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

Configure VLANs for Cisco Nexus A and Cisco Nexus B Switches

To create the necessary virtual local area networks (VLANs), follow these steps on both Cisco Nexus switches.

Note: We created VLAN 70, 71, 72, 73 and 76.

1. Log in as admin user into the Cisco Nexus Switch A.
2. Create VLAN 70:

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

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 5](#).

Table 5. vPC Summary

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

As listed in [Table 5](#), 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 6. 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 7. 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 8. 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	40Gbe	UCS 5108 Chassis IOM-A Chassis 1-4	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

Local Device	Local Port	Connection	Remote Device	Remote Port
	Mgmt 0	1Gbe	Management Switch	Any
	L1	1Gbe	Cisco UCS FI - A	L1
	L2	1Gbe	Cisco UCS FI - B	L2

Table 9. 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-Gb-B	FC 1/14
	Eth1/17-24	40Gbe	UCS 5108 Chassis IOM-B Chassis 1-4	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

Create vPC Peer-Link Between the Two Cisco Nexus Switches

To create the vPC Peer-Link, follow these steps:

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.

2. To create the necessary port channels between devices, follow these steps 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 70
description VPC peer-link
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type network
vpc peer-link
exit
interface Ethernet1/53
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
channel-group 70 mode active
no shutdown
exit
interface Ethernet1/54
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
channel-group 70 mode active
no shutdown
exit
copy running-config startup-config
```

3. Log in as admin user into the Nexus Switch B and repeat the above steps to configure second Cisco Nexus switch.
4. Make sure to change the peer-keepalive destination and source IP address appropriately for Cisco Nexus Switch B.

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

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

To create the necessary port channels between devices, follow these steps on both Cisco Nexus switches:

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

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

```
mtu 9216
channel-group 11 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

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

```
config Terminal
interface port-channel11
description FI-A-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 11
no shutdown
exit
interface port-channel12
description FI-B-Uplink
switchport mode trunk
switchport trunk allowed VLAN 1,70-76
spanning-tree port type edge trunk
vpc 12
no shutdown
exit
interface Ethernet1/51
description FI-A-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 11 mode active
no shutdown
exit
interface Ethernet1/52
description FI-B-Uplink
switch mode trunk
switchport trunk allowed vlan 1,70-76
spanning-tree port type edge trunk
mtu 9216
channel-group 12 mode active
no shutdown
exit
copy running-config startup-config
```

Verify all vPC Status is up on both Cisco Nexus Switches

[Figure 33](#) shows the verification of the vPC status on both Cisco Nexus Switches.

Figure 33. vPC Description for Cisco Nexus Switch A and B

```

AAD17-NX9K-A# sh vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 70
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  : Disabled
Delay-restore status   : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router : Disabled

vPC Peer-link status
-----
id  Port  Status Active vlans
--  ---  -
1   Po70  up    1,70-76

vPC status
-----
Id  Port  Status Consistency Reason      Active vlans
--  ---  -
11  Po11  up    success  success      1,70-76
12  Po12  up    success  success      1,70-76
13  Po13  up    success  success      1,70-76
14  Po14  up    success  success      1,70-76
    
```

```

AAD17-NX9K-B# sh vpc brief
Legend:
(*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 70
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  : Disabled
Delay-restore status   : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router : Disabled

vPC Peer-link status
-----
id  Port  Status Active vlans
--  ---  -
1   Po70  up    1,70-76

vPC status
-----
Id  Port  Status Consistency Reason      Active vlans
--  ---  -
11  Po11  up    success  success      1,70-76
12  Po12  up    success  success      1,70-76
13  Po13  up    success  success      1,70-76
14  Po14  up    success  success      1,70-76
    
```

Cisco MDS 9132T 32-Gb FC Switch Configuration

[Figure 13](#) 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 10. 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 11. 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

Local Device	Local Port	Connection	Remote Device	Remote Port
	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 12](#). All ports connected to the Pure Storage Array carry 32 Gb/s FC Traffic.

Table 12. 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

Configure Feature for MDS Switch A and MDS Switch B

To set feature on MDS Switches, follow these steps on both MDS switches:

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
```

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

Configure VSANs for MDS Switch A and MDS Switch B

To create VSANs, follow these steps:

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

```
config terminal
VSAN database
vsan 100
exit
zone smart-zoning enable vsan 100
vsan database
```

```
vsan 100 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 100
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

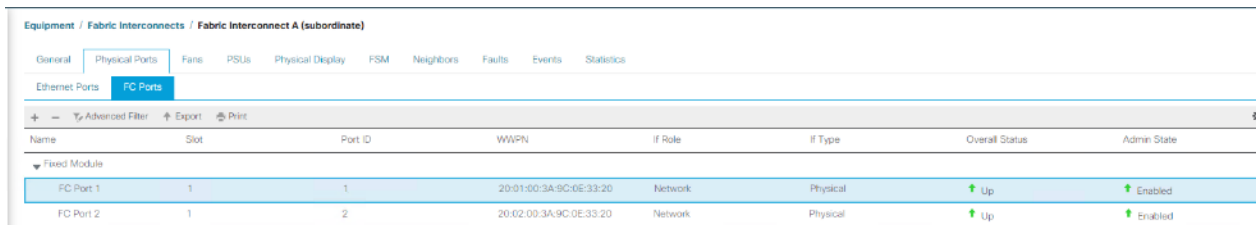
2. Log in as admin user into MDS Switch B. Create VSAN 101 for Storage Traffic:

```
config terminal
VSAN database
vsan 101
exit
zone smart-zoning enable vsan 101
vsan database
vsan 101 interface fc 1/9-16
exit
interface fc 1/9-16
switchport trunk allowed vsan 101
switchport trunk mode off
port-license acquire
no shutdown
exit
copy running-config startup-config
```

Add FC Uplink Ports to Corresponding VSAN on Fabric Interconnect

To add the FC Ports to the corresponding VSAN, follow these steps:

1. In Cisco UCS Manager, in the Equipment tab, select Fabric Interconnects > Fabric Interconnect A > Physical Ports > FC Ports.

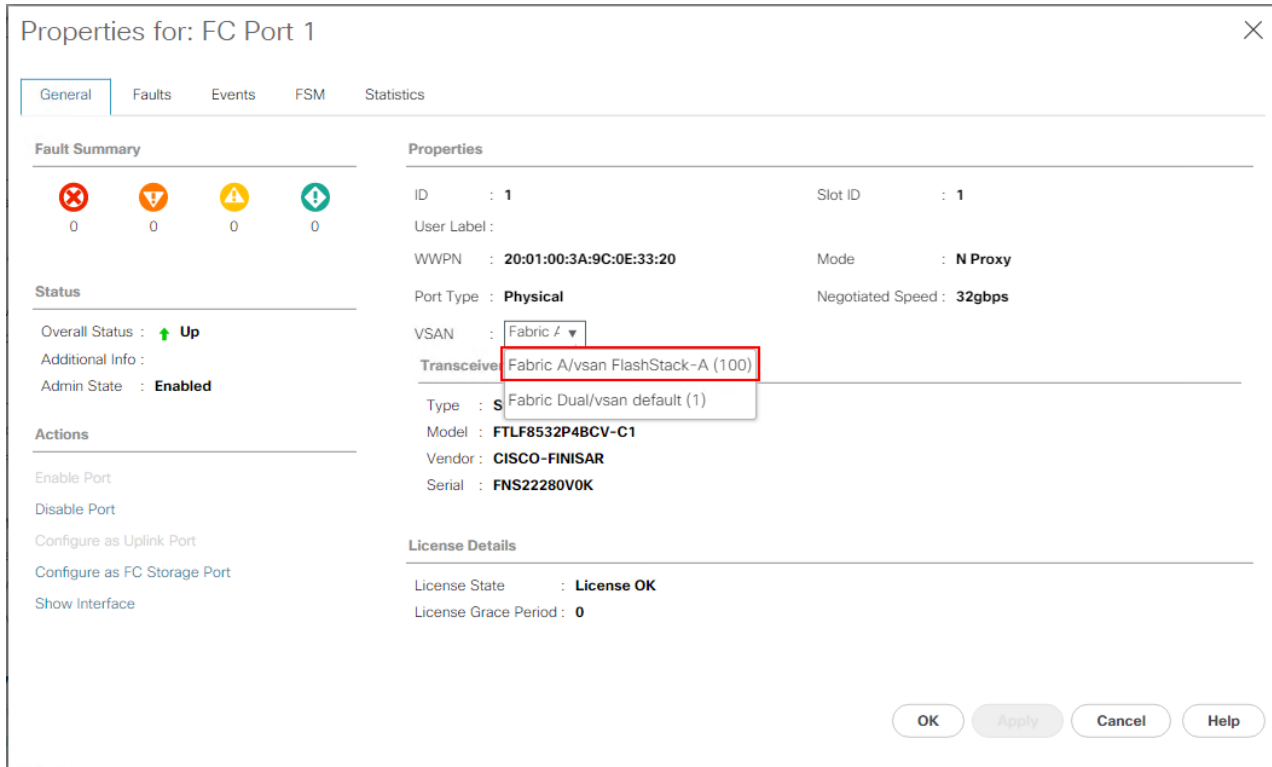


The screenshot shows the Cisco UCS Manager interface for Fabric Interconnect A. The 'Physical Ports' tab is selected, and the 'FC Ports' sub-tab is active. A table lists the configured FC ports:

Name	Slot	Port ID	WWPN	If Role	If Type	Overall Status	Admin State
FC Port 1	1	1	20:01:00:3A:9C:0E:33:20	Network	Physical	Up	Enabled
FC Port 2	1	2	20:02:00:3A:9C:0E:33:20	Network	Physical	Up	Enabled

2. From the drop-down list double-click FC Port 1 and select VSAN 100.

Figure 34. VSAN Assignment on FC Uplink Ports to MDS Switch



3. Repeat steps 1 and 2 to add the FC Port 1-4 to VSAN 100 on Fabric A and FC Port 1-4 to VSAN 101 on Fabric B.

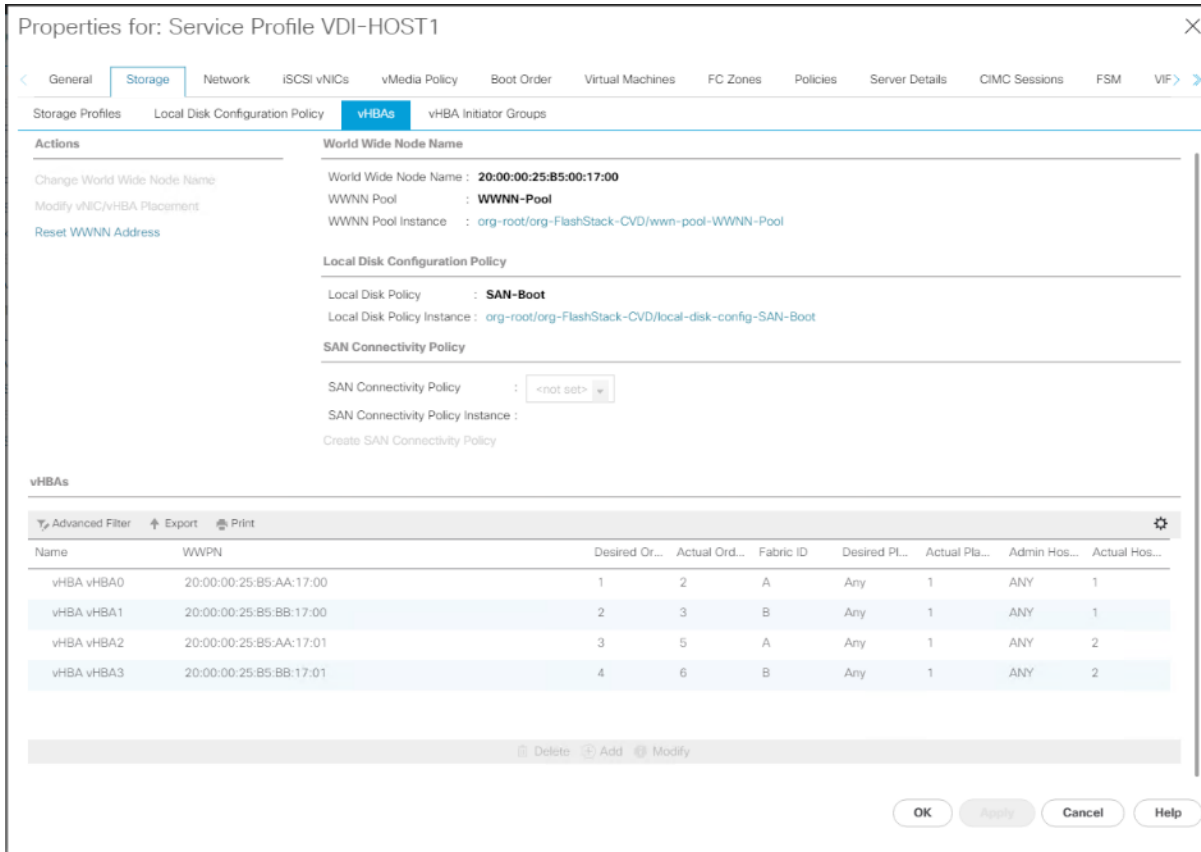
Create and Configure Fiber Channel Zoning

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 4 HBAs for each Server. Two HBAs (HBA0 and HBA2) are connected to MDS Switch-A and other two HBAs (HBA1 and HBA3) are connected to MDS Switch-B.

To create and configure the fiber channel zoning, follow these steps:

1. Log into the Cisco UCS Manager and go to Servers > Service Profiles > Sub-Organizations > FlashStack-CVD > VDI-HostX, then click the Storage tab and HBA's tab to get the WWPN of HBA's as shown in the screenshot below. Repeat for all the configured host profiles.



2. 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	Fallover	FC Port	Name	Speed	Fallover
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	

Create Device Aliases for Fiber Channel Zoning

Cisco MDS Switch A

To configure device aliases and zones for the SAN boot paths as well as the datapaths of MDS switch A, follow these steps:

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 pwnn 20:00:00:25:B5:AA:17:00
device-alias name X70R3-CT0-FC0 pwnn 52:4A:93:71:56:84:09:00
device-alias name X70R3-CT1-FC0 pwnn 52:4A:93:71:56:84:09:10
exit
device-alias commit
```

Cisco MDS Switch B

To configure device aliases and zones for the SAN boot paths as well as datapaths of MDS switch B, follow this 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 pwnn 20:00:00:25:b5:bb:17:03
device-alias name X70R3-CT0-FC2 pwnn 52:4A:93:71:56:84:09:02
device-alias name X70R3-CT1-FC2 pwnn 52:4A:93:71:56:84:09:12
exit
device-alias commit
```

Create Fiber Channel Zoning

Cisco MDS Switch A

To configure zones for the MDS switch A, follow these steps to create a zone for each server service profile:

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

```
configure terminal
zone name FlashStack-Fabric-A vsan 100
  member device-alias X70R3-CT0-FC0 target
  member device-alias X70R3-CT1-FC0 targetshow
  member device-alias Host-FCP-1-HBA0 init
```

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 100
  member FlashStack-Fabric-A
```

3. Activate the zone set by running following commands:

```
zoneset activate name VDI-Fabric-A vsan 100
exit
copy running-config startup-config
```

Cisco MDS Switch B

To configure zones for the MDS switch B, follow these steps to create a zone for each server service profile:

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

```
configure terminal zone name FlashStack-Fabric-B vsan 101
  member device-alias X70R3-CT0-FC2 target
  member device-alias X70R3-CT1-FC2 target
  member device-alias Host-FCP-1-HBA1 init
```

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 101
  member FlashStack-Fabric-B
```

3. Activate the zone set by running following commands:

```
zoneset activate name VDI-Fabric-B vsan 101
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 have installed vSphere ESXi 7.0 Update 2 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:

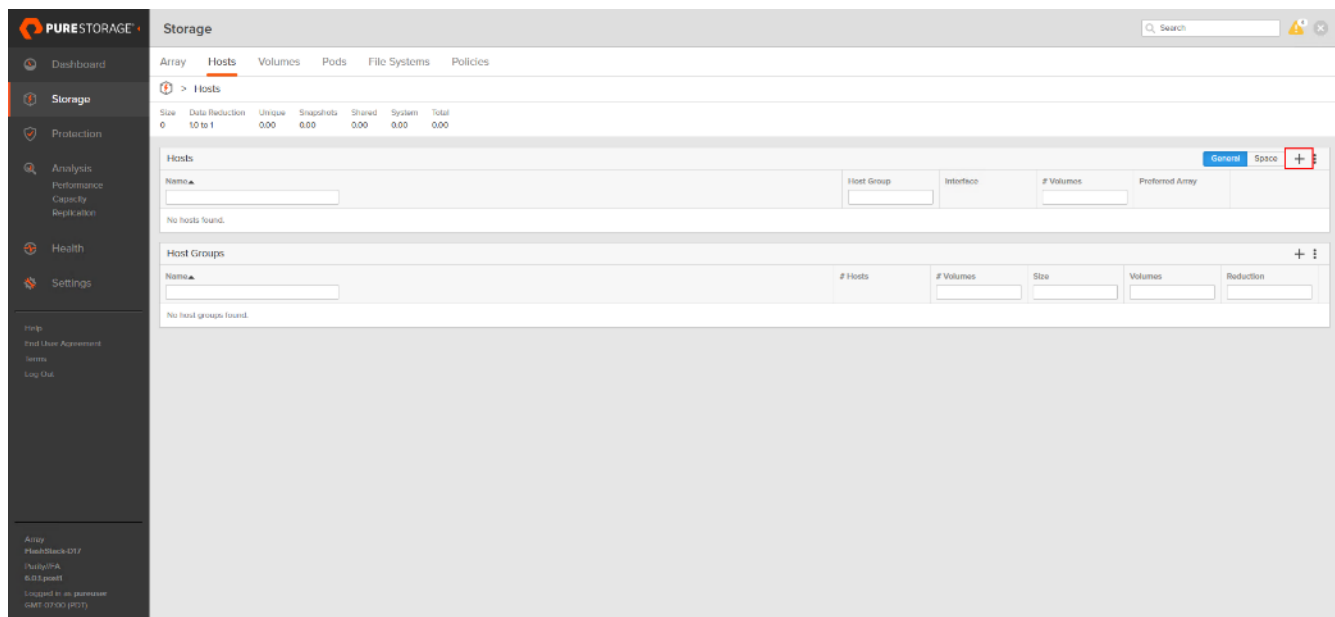
- Scalability - Rapid deployment of new servers to the environment in a very few steps.
- Manageability - Enables seamless hardware maintenance and upgrades without any restrictions. This is a huge benefit in comparison to another appliance model like Exadata.
- Flexibility - Easy to repurpose physical 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

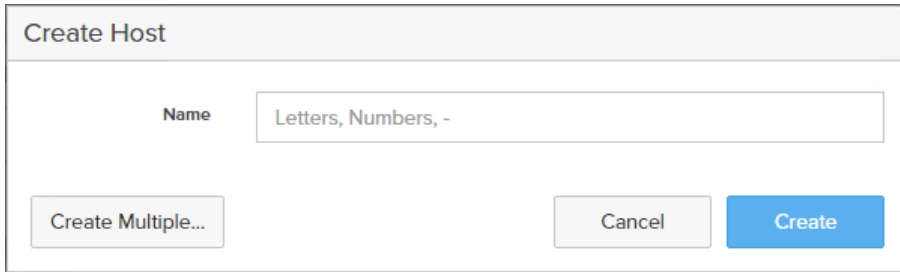
Configure Host

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. To set up a host, follow these steps:

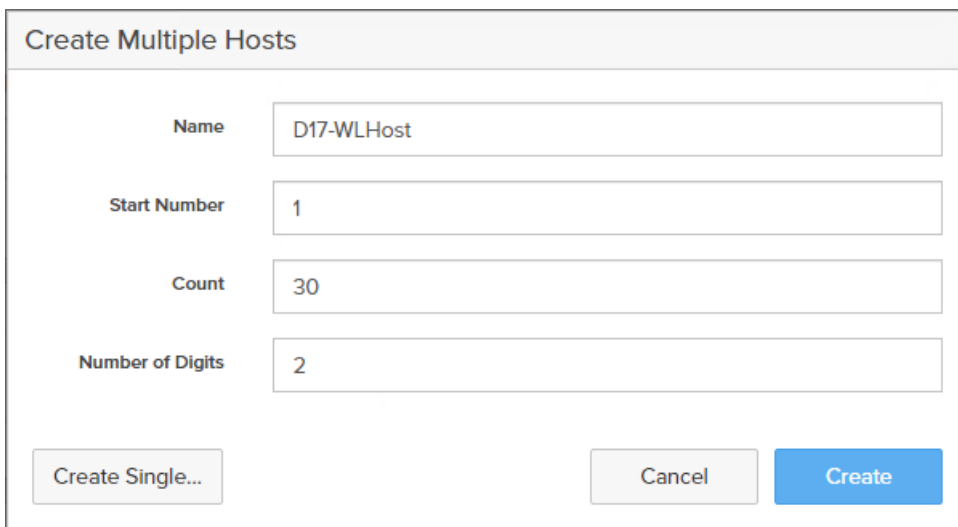
1. Log into Pure Storage FlashArray Management interface.
2. Click the Storage tab.
3. Click the + sign in the Hosts section and select Create Host.



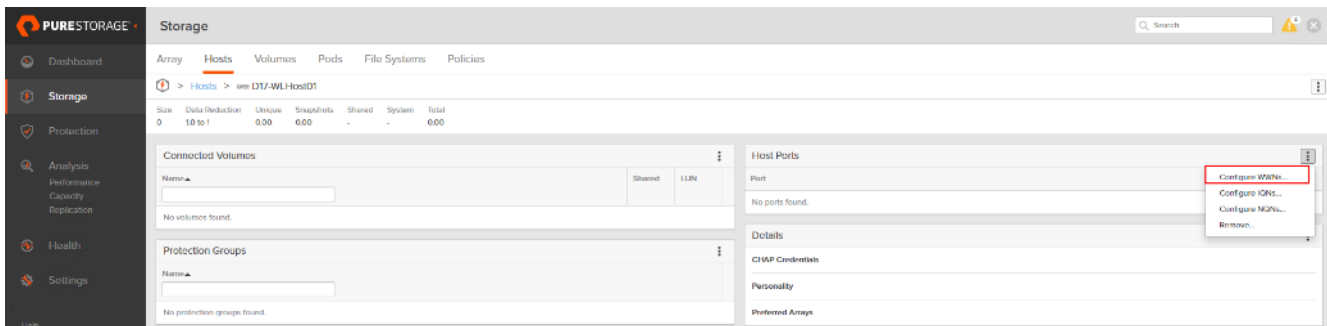
4. Click Create Multiple to create a Host entries under the Hosts category.



5. Enter the required information and click Create.

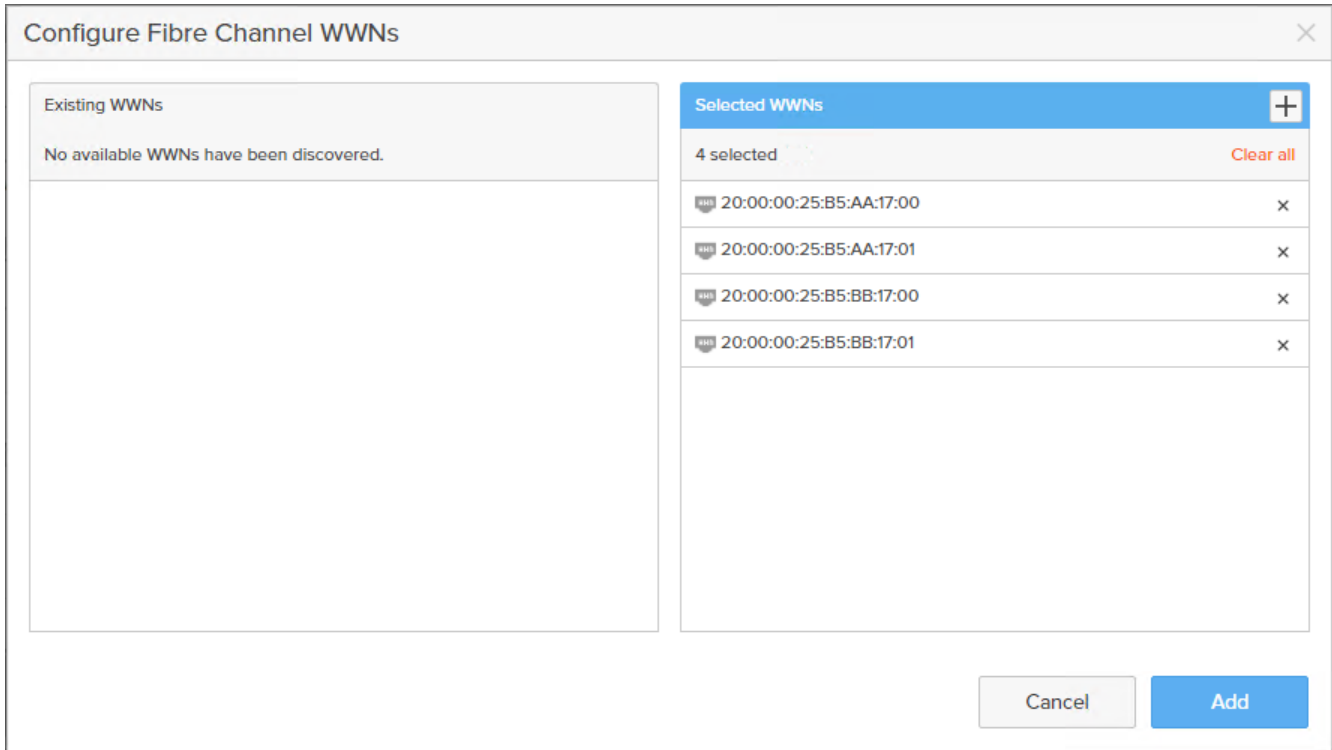


6. Select one of the newly created hosts, in Host Ports section from the drop-down list select Configure WWNs.



Size	Data Deduction	Unique	Snapshots	Shared	System	Total
0	10 to 1	0.00	0.00	-	-	0.00

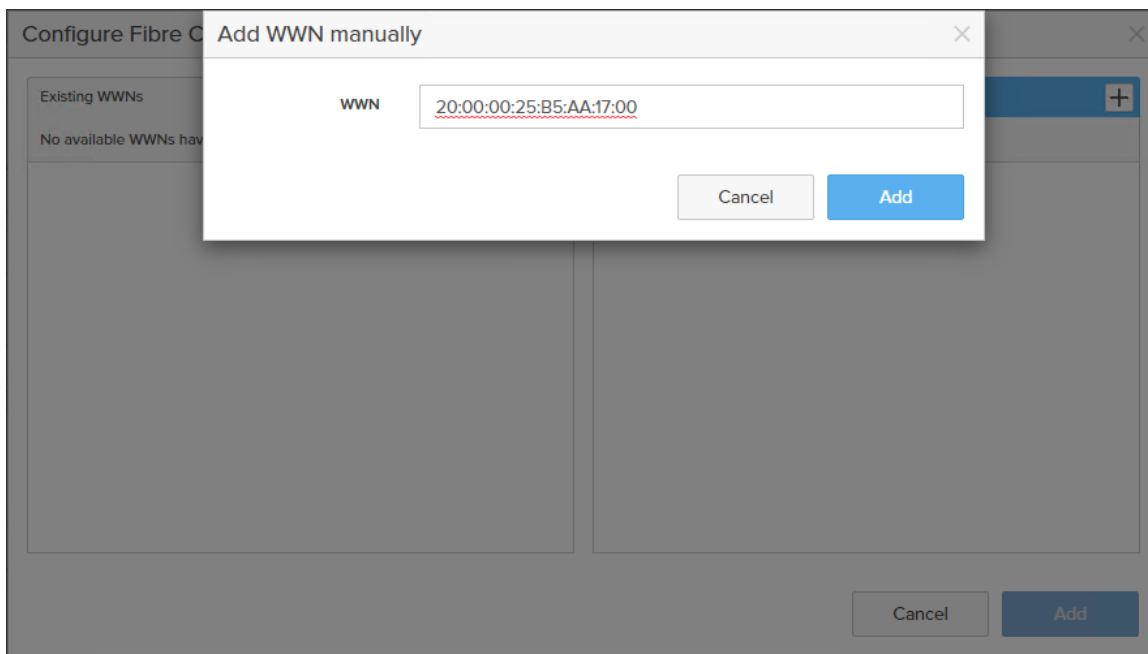
7. Select the list of WWNs that belongs to the host in the next window and click Add.



8. 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.

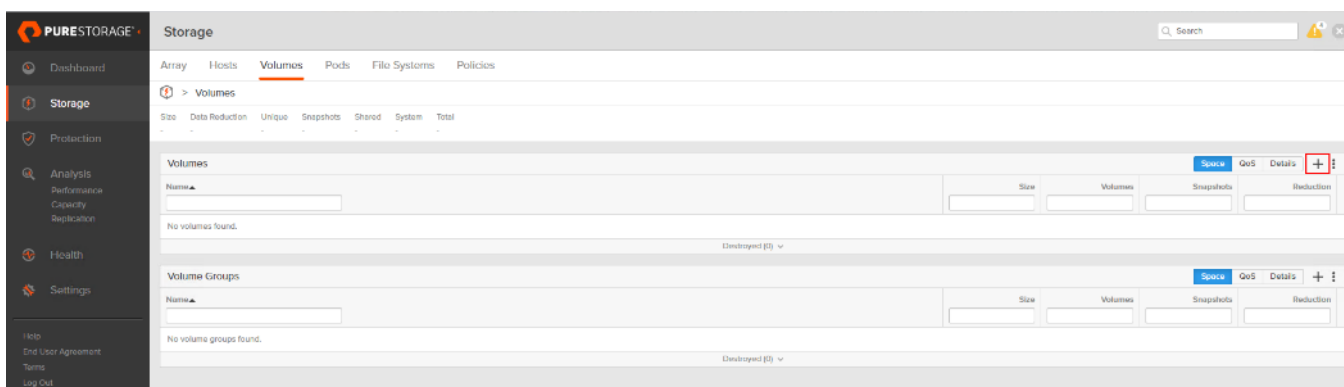
Note: Alternatively, the WWN can be added manually by clicking the + in the Selected WWNs section and manually inputting the blade's WWNs.



Configure Volume Connectivity

To configure a volume and volume connectivity, follow these steps:

1. Click the Storage tab.
2. Click the + sign in the Volumes section and click Create Volume.



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

Create Volume [X]

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

QoS Configuration (Optional) v

Buttons: Create Multiple..., Cancel, Create

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 [X]

Pod or Volume Group: none

Name: D17-WLHost

Provisioned Size: 20 G

Start Number: 1

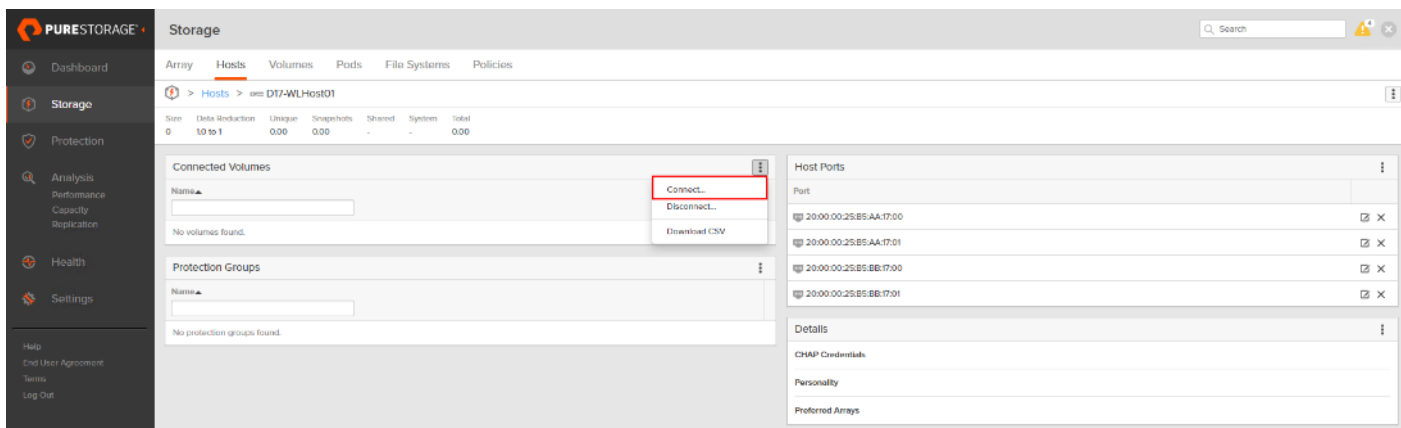
Count: 30

Number of Digits: 2

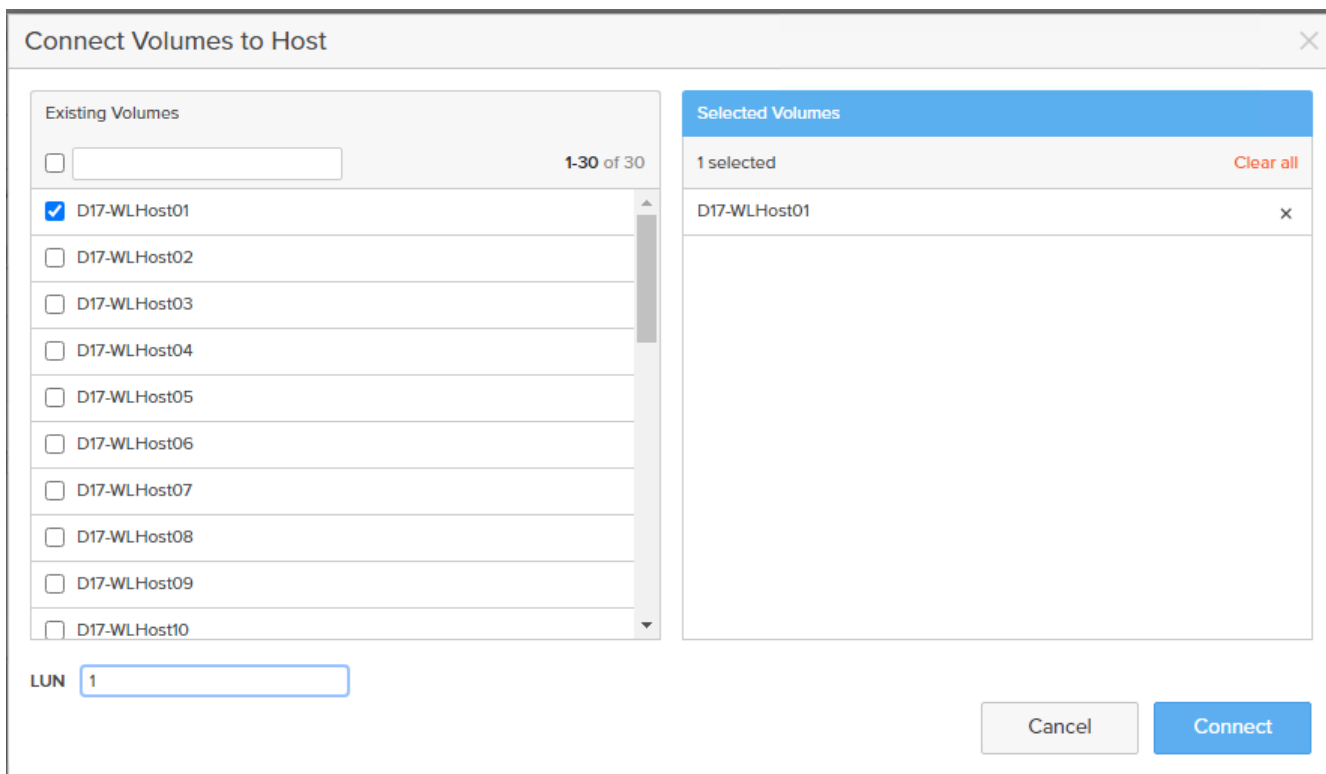
QoS Configuration (Optional) v

Buttons: Create Single..., Cancel, Create

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



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



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.

7. 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](#)

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 created by CLI over SSH, configured and enabled via Management Console. Account with administrator privileges is required.

To create File Virtual Interface, follow these steps:

1. Connect to the array via SSH.
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>
```

Configure and Enable the Virtual Interface for File Services

To configure and enable the virtual interface, follow these steps:

1. Connect to the array GUI.
2. Navigate to Settings > Network.
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



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

Cancel Save

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

DNS Settings 

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

Cancel Save

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.

Create Active Directory Account for the Array

To create the Active Directory Account, follow these steps:

1. Navigate to Settings > Access > Active Directory Accounts.
2. To open the Create Dialog, click the + icon.

Active Directory Accounts

1 of 1 +

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

4. Click Create to finalize AD account creation.


Create Active Directory Account

Name	<input type="text" value="purefile"/>
Domain	<input type="text" value="vccfslab.local"/>
Computer Name	<input type="text" value="purefile"/>
Kerberos Server	<input type="text"/>
Directory Server	<input type="text"/>
User	<input type="text" value="administrator@vccfslab.local"/>
Password	<input type="password" value="....."/>

Create a File System and Shared Directory

To create a file system and shared directory, follow these steps:

1. Navigate to Storage > File Systems.
2. Click the + icon.




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

Create File System

Name vdi

Cancel Create

4. Navigate to Storage > File Systems > Directories.
5. Click the + icon.



6. In Create Directory pop-up dialog 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 /

Cancel Create

Note: Policies 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.

7. Navigate to Storage > Policies.


8. Click the + icon.



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.

A dialog box titled 'Create Export Policy' with a close button (X) in the top right corner. It contains three fields: 'Type' is a dropdown menu with 'SMB' selected; 'Name' is a text input field containing 'smb'; 'Enabled' is a toggle switch that is turned on. At the bottom right, there are two buttons: 'Cancel' and 'Create'.


10. Click Created Policy and click the + icon.



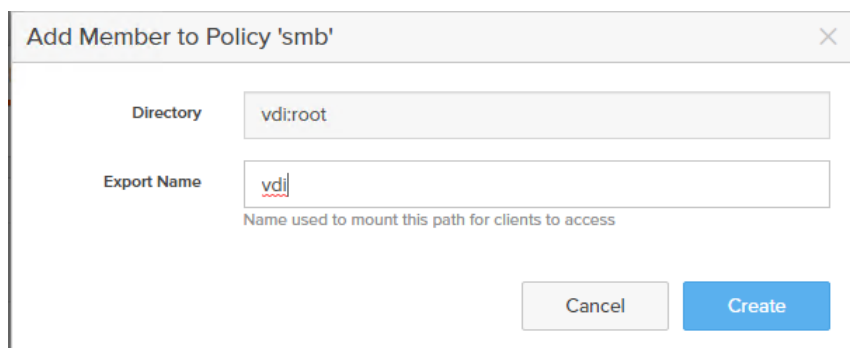
11. Complete the Client filter for read-write access and click Add to complete the rule creation.

A dialog box titled 'Add Rule for Policy 'smb'' with a close button (X) in the top right corner. It contains three sections: 'Client' is a text input field with a cursor, with a tooltip below it showing 'Hostname, IPv4 or IPv4 mask. e.g., *, *.cs.foo.edu, 192.168.255.255, or 192.168.10.0/24'; 'Access' has two radio buttons, 'no-anonymous-access' is selected; 'Encryption' has two radio buttons, 'optional-smb-encryption' is selected. At the bottom right, there are two buttons: 'Cancel' and 'Add'.

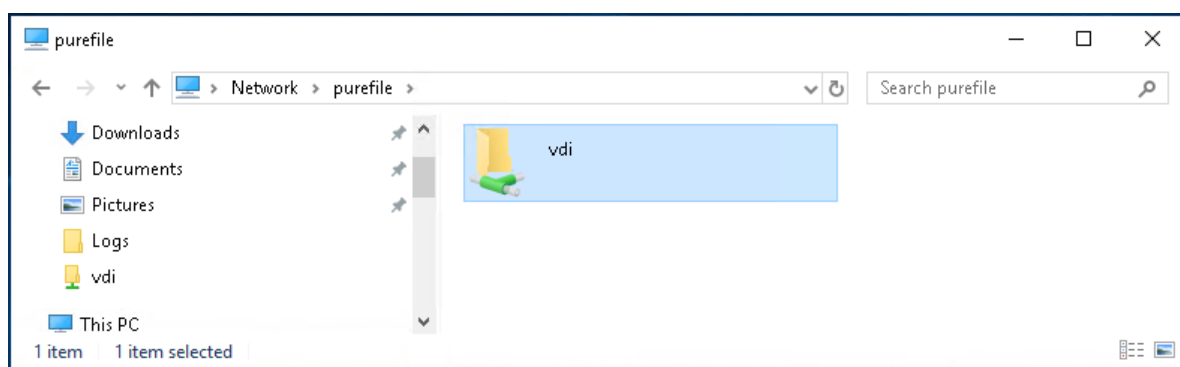
12. Attach the export policy(s) to a managed directory. Click the + icon.



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



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



Install and Configure VMware ESXi 7.0

This section explains how to install VMware ESXi 7.0 Update 2a 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 7.0

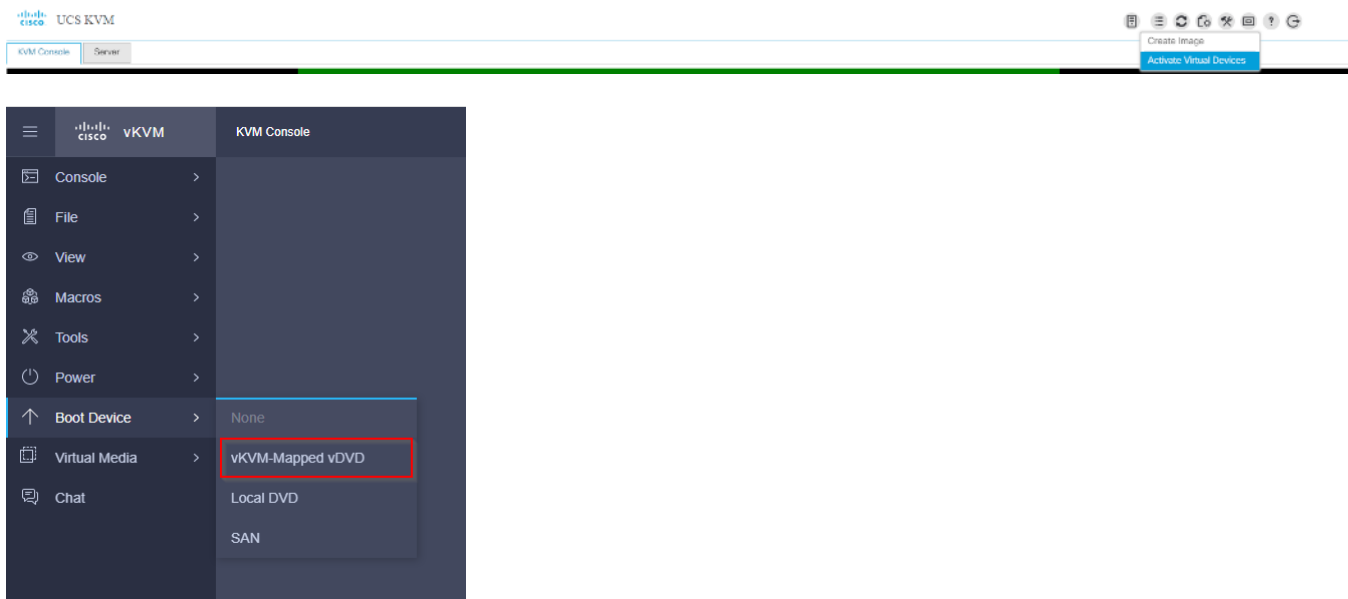
To download the Cisco Custom Image for VMware ESXi 7.0 Update 2a, from the [VMware vSphere Hypervisor 7.0 U2](#) page click the Custom ISOs tab.

Install VMware vSphere ESXi 7.0 U2


To install VMware vSphere ESXi hypervisor on Cisco UCS Server, follow these steps:

1. In the Cisco UCS Manager navigation pane, click the Equipment tab.
2. Under Servers > Service Profiles > VDI-Host1
3. Right-click on VDI-Host1 and select KVM Console.

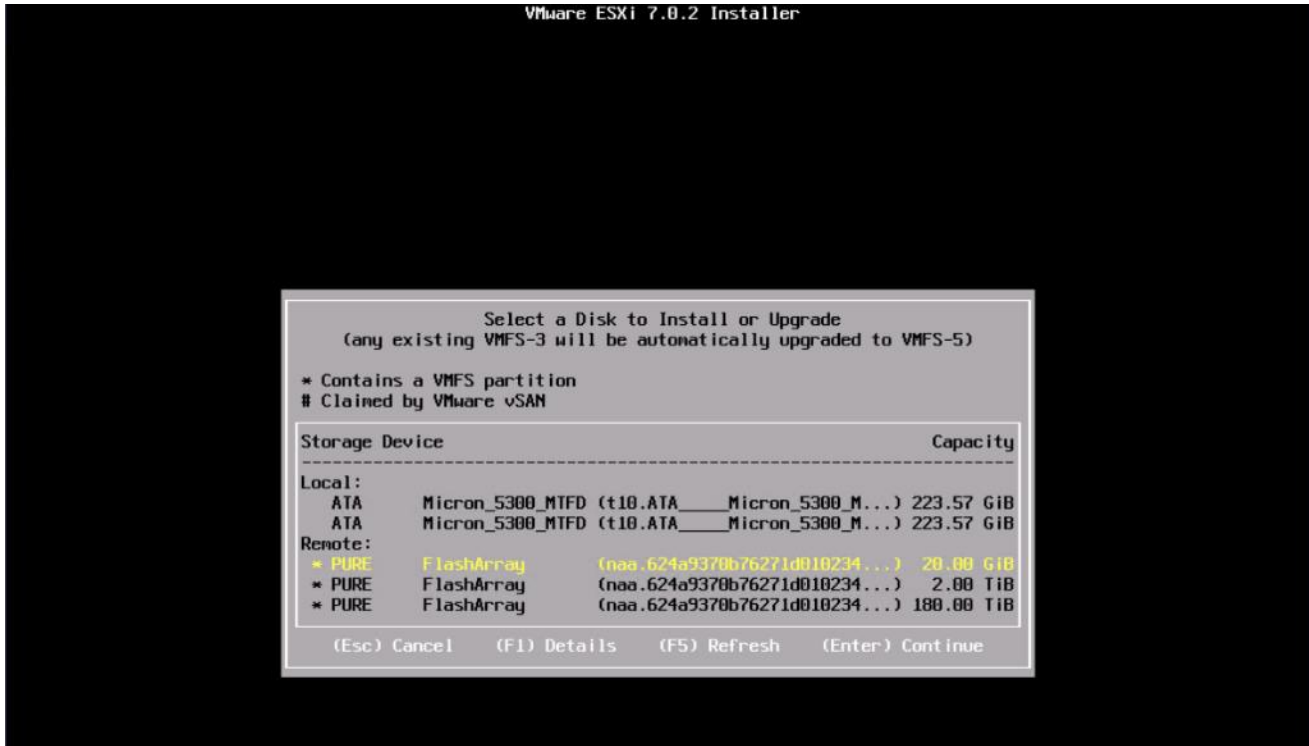
4. Click Boot Device and then select CD/DVD.



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



6. Boot into ESXi installer and follow the prompts to complete installing VMware vSphere ESXi hypervisor.
7. When selecting a storage device to install ESXi, select Remote LUN provisioned through Pure Storage Administrative console and access through FC connection.



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. Please select the IP address that can communicate with existing or new vCenter Server.

To configure the ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, press F2 to enter in to configuration wizard for ESXi Hypervisor.
2. Log in as root and enter the corresponding password.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter. Enter the VLAN In-Band management ID and press Enter.
5. From the Configure Management Network menu, select IP Configuration and press Enter.
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.

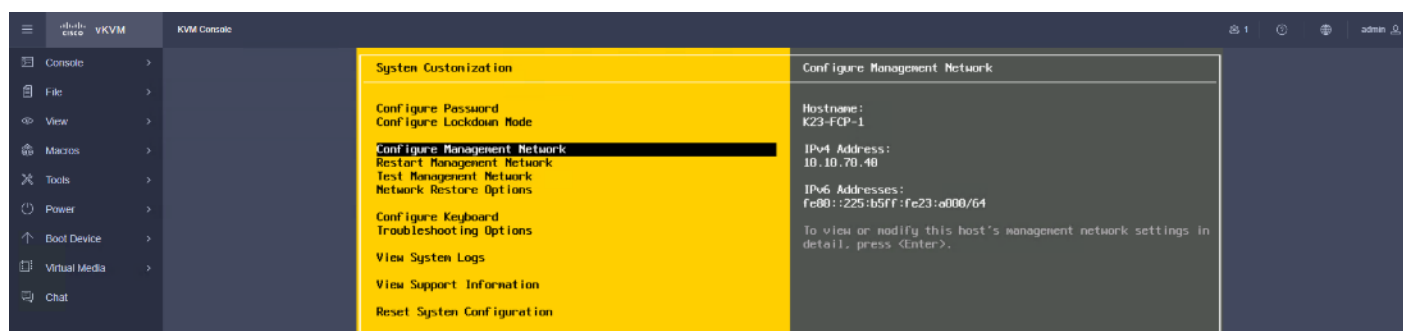
Note: IPv6 Configuration is set to automatic.

7. Select the DNS Configuration option and press Enter.
8. Enter the IP address of the primary and secondary DNS server. Enter Hostname
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 35. Sample ESXi Configure Management Network



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](#).

In this Validated Design the following drivers were used:

- Cisco-nenic- 1.0.35.0
- Cisco-nfnic- 5.0.0.15

To update the Cisco VIC drivers for ESXi, follow these steps:

1. Log into your VMware Account to download required drivers for FNIC and NENIC as per the recommendation.
2. Enable SSH on ESXi to run following commands:

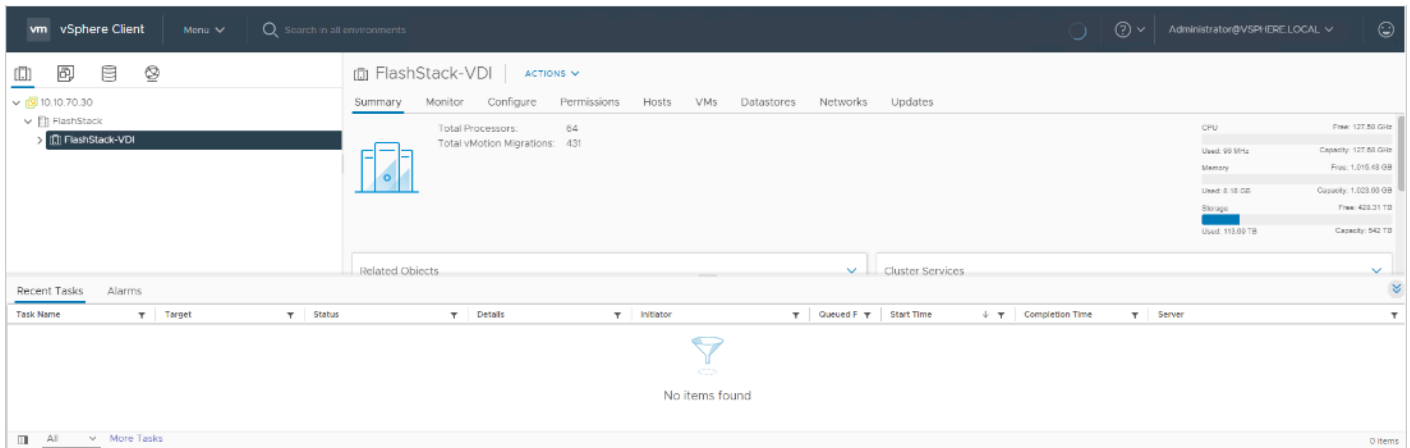
```
esxcli software vib update -d /path/offline-bundle.zip
```

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 36. VMware vSphere WebUI Reporting Cluster Configuration for this Validated Design




Build the Virtual Machines and Environment for Workload Testing

Prerequisites

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

Figure 37. Example of the DHCP Scopes used in this CVD



Software Infrastructure Configuration

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

Install and configure the infrastructure virtual machines by following the process listed in [Table 13](#).

Table 13. Test Infrastructure Virtual Machine Configuration

Configuration	Citrix Virtual Apps and Desktops Controllers Virtual Machines	Citrix Provisioning Servers Virtual Machines
Operating system	Microsoft Windows Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	6	6
Memory amount	24 GB	24 GB
Network	VMXNET3 k23-Infra-Mgmt-71	VMXNET3 k23-Infra-Mgmt-71
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 2019	VCSA - SUSE Linux
Virtual CPU amount	4	16
Memory amount	8 GB	32 GB
Network	VMXNET3 k23-Infra-Mgmt-71	VMXNET3 k23-InBand-Mgmt-70
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 2019 Microsoft SQL Server 2019	Microsoft Windows Server 2019
Virtual CPU amount	6	4
Memory amount	24GB	8 GB
Network	VMXNET3 k23-Infra-Mgmt-71	VMXNET3 k23-Infra-Mgmt-71
Disk-1 (OS) size	40 GB	40 GB
Disk-2 size	100 GB SQL Databases\Logs	-

Prepare the Master Targets

This section provides guidance regarding creating the golden (or master) images for the environment. Virtual machines for the master targets must first be installed with the software components needed to build the golden images. Additionally, all available security patches as of October 2021 for the Microsoft operating systems, SQL server and Microsoft Office 2016 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 14](#).

Table 14. Single-session OS and Multi-session OS Virtual Machines Configurations

Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2016
Virtual CPU amount	3	8
Memory amount	3 GB reserve for all guest memory	32 GB reserve for all guest memory
Network	VMXNET3 10_10_72_NET	VMXNET3 10_10_72_NET
Citrix PVS vDisk size	48 GB (dynamic)	90 GB (dynamic)
Citrix MCS Disk Size	48 GB	
write cache	6 GB	6 GB
Disk size		
Citrix PVS write cache	128 MB	1024 MB
RAM cache size		
Additional software used for testing	Microsoft Office 2016 Office Update applied Login VSI 4.1.39.6 Target Software (Knowledge Worker Workload)	Microsoft Office 2016 Office Update applied Login VSI 4.1.39.6 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

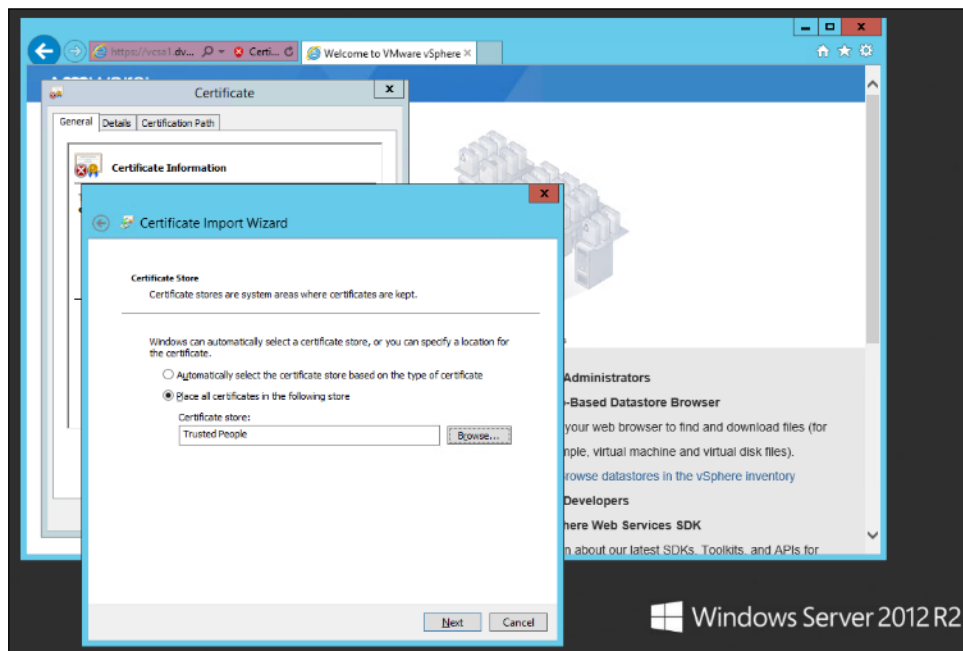
This section explains the installation of the core components of the Citrix Virtual Apps and Desktops system. This CVD installs two Citrix Virtual Apps and Desktops Delivery Controllers to support both hosted shared desktops (HSD), non-persistent hosted virtual desktops (HVD), and persistent hosted virtual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

To install vCenter Server self-signed Certificate, follow these steps:

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.
2. Open Internet Explorer and enter the address of the computer running vCenter Server (for example, https://FQDN as the URL).
3. Accept the security warnings.
4. Click the Certificate Error in the Security Status bar and select View certificates.
5. Click Install certificate, select Local Machine, and then click Next.
6. Select Place all certificates in the following store and then click Browse.
7. Click Show physical stores.
8. Click Trusted People.



9. Click Next and then click Finish.

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

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

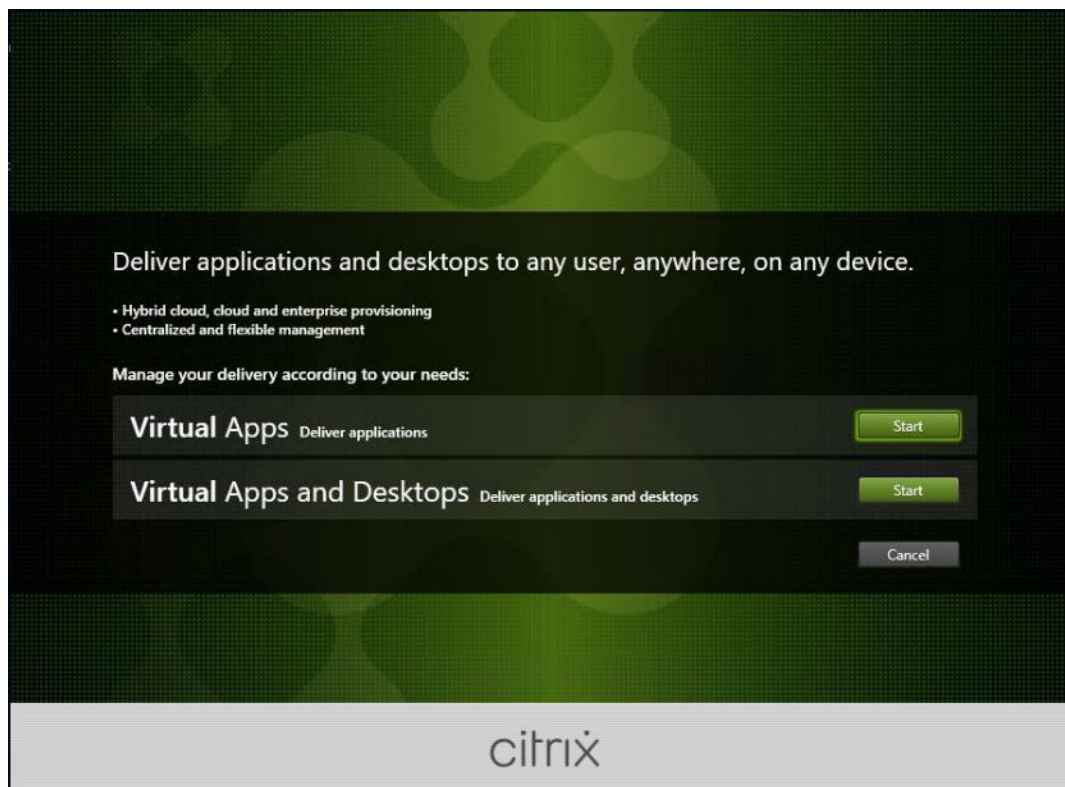
The process of installing the Citrix Virtual Apps and Desktops Delivery Controller also installs other key Citrix Virtual Apps and Desktops software components, including Studio, which is used to create and manage infrastructure components, and Director, which is used to monitor performance and troubleshoot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

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

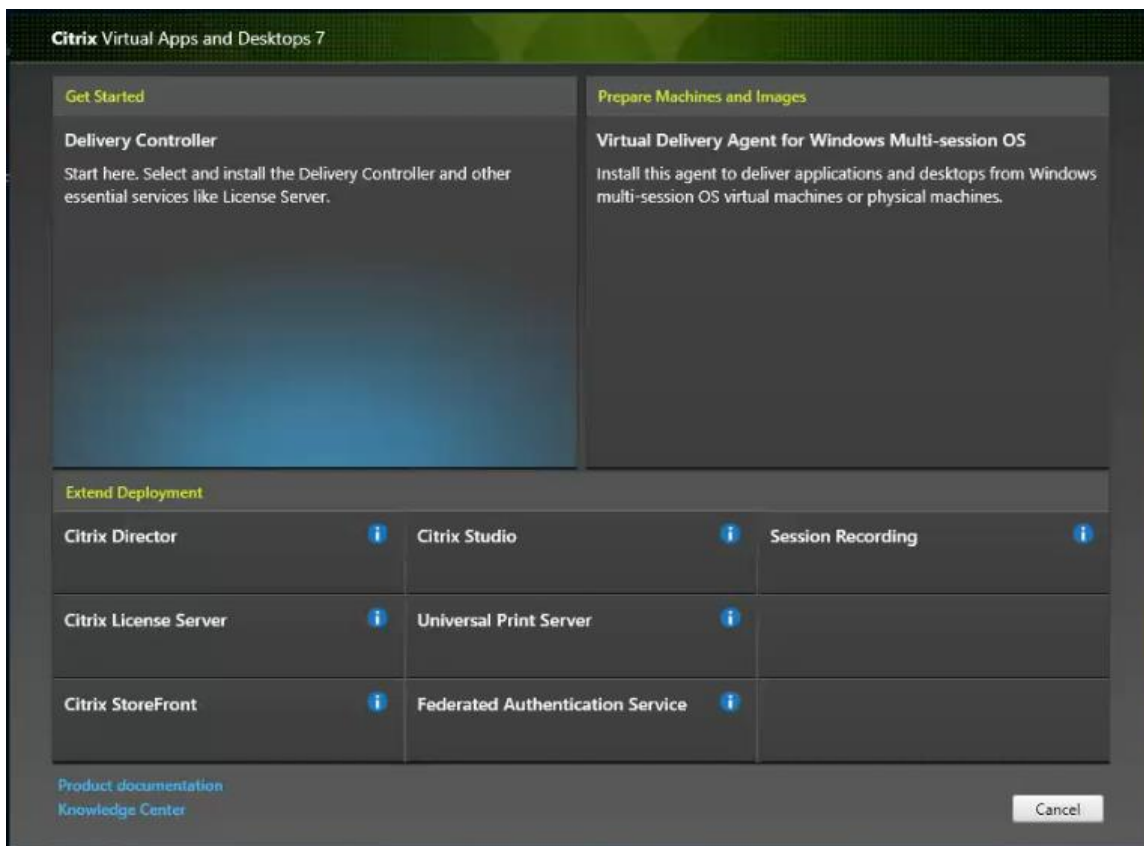
Install Citrix License Server

To install the Citrix License Server, follow these steps:

1. To begin the installation, connect to the first Citrix License server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2109 ISO.
2. Click Start.



3. Click Extend Deployment – Citrix License Server.



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.
5. Click Next.

Software License Agreement

[Printable version](#)

Licensing Agreement

Core Components

Firewall

Summary

Install

Finish

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

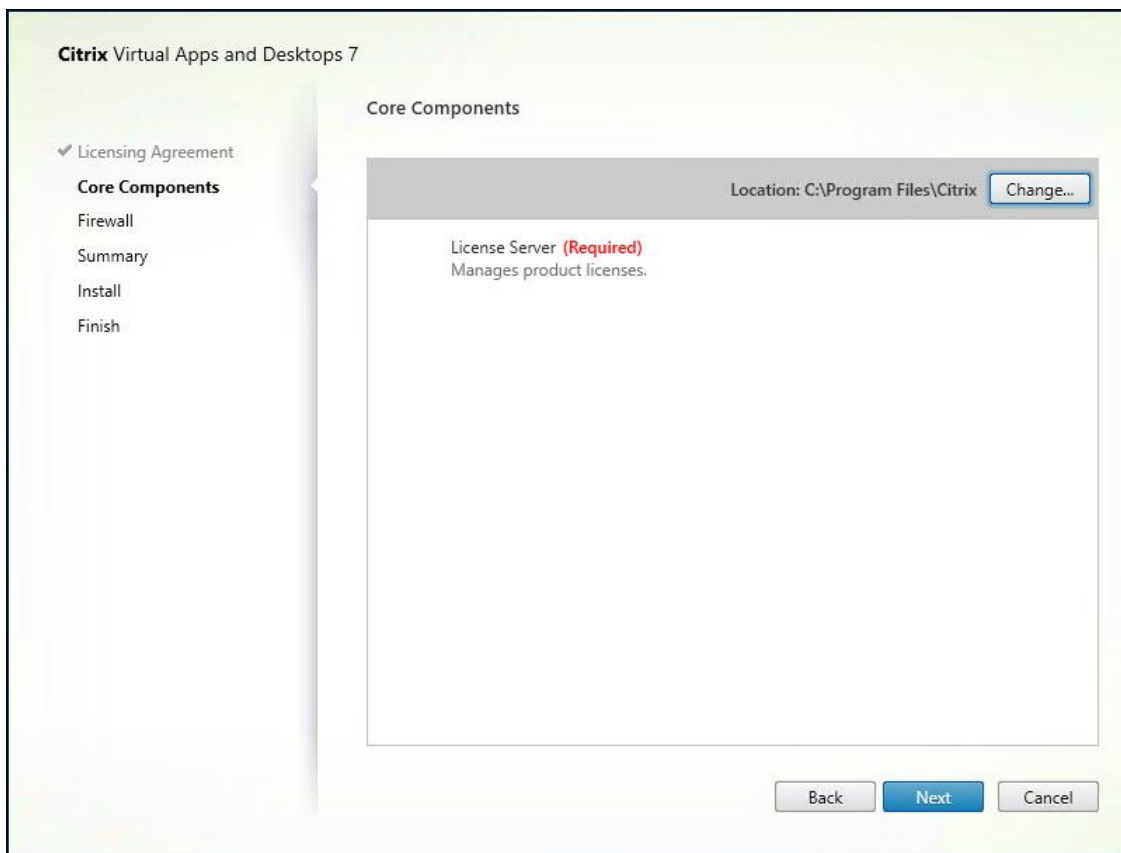
- I have read, understand, and accept the terms of the license agreement
- I do not accept the terms of the license agreement

Back

Next

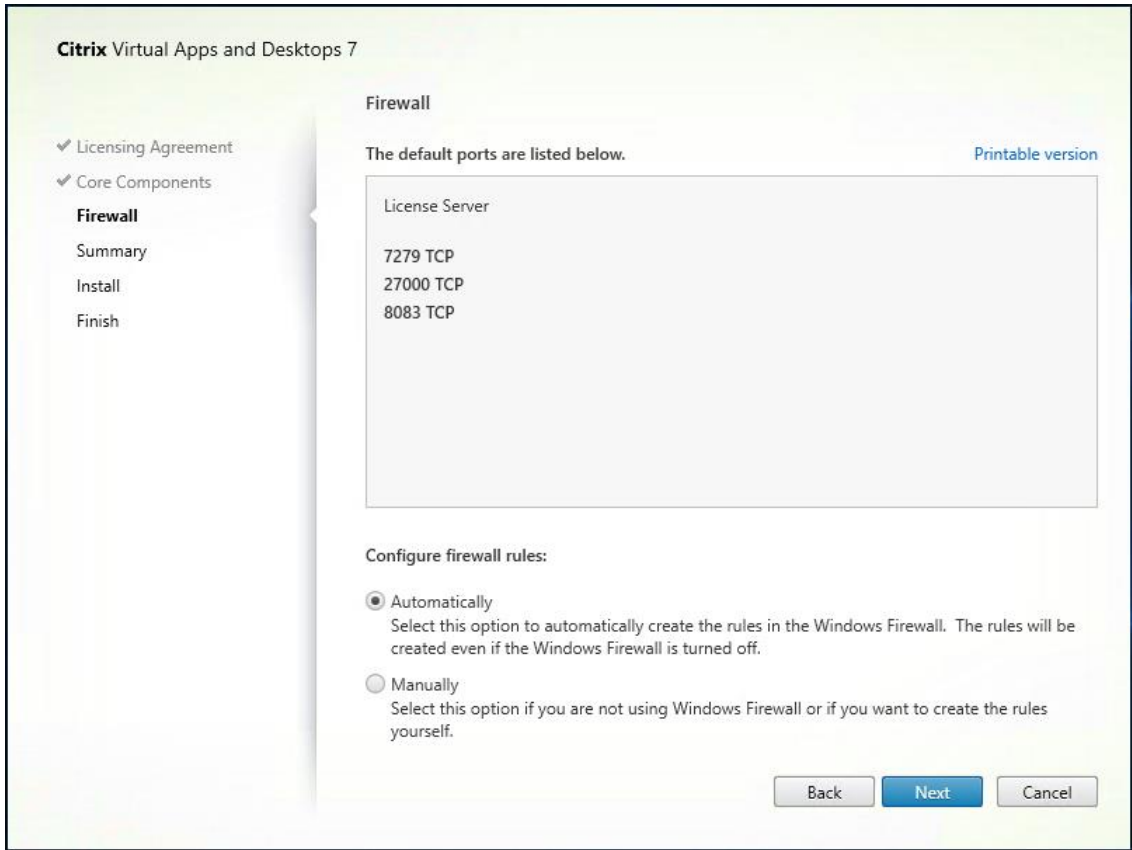
Cancel

6. Click Next.

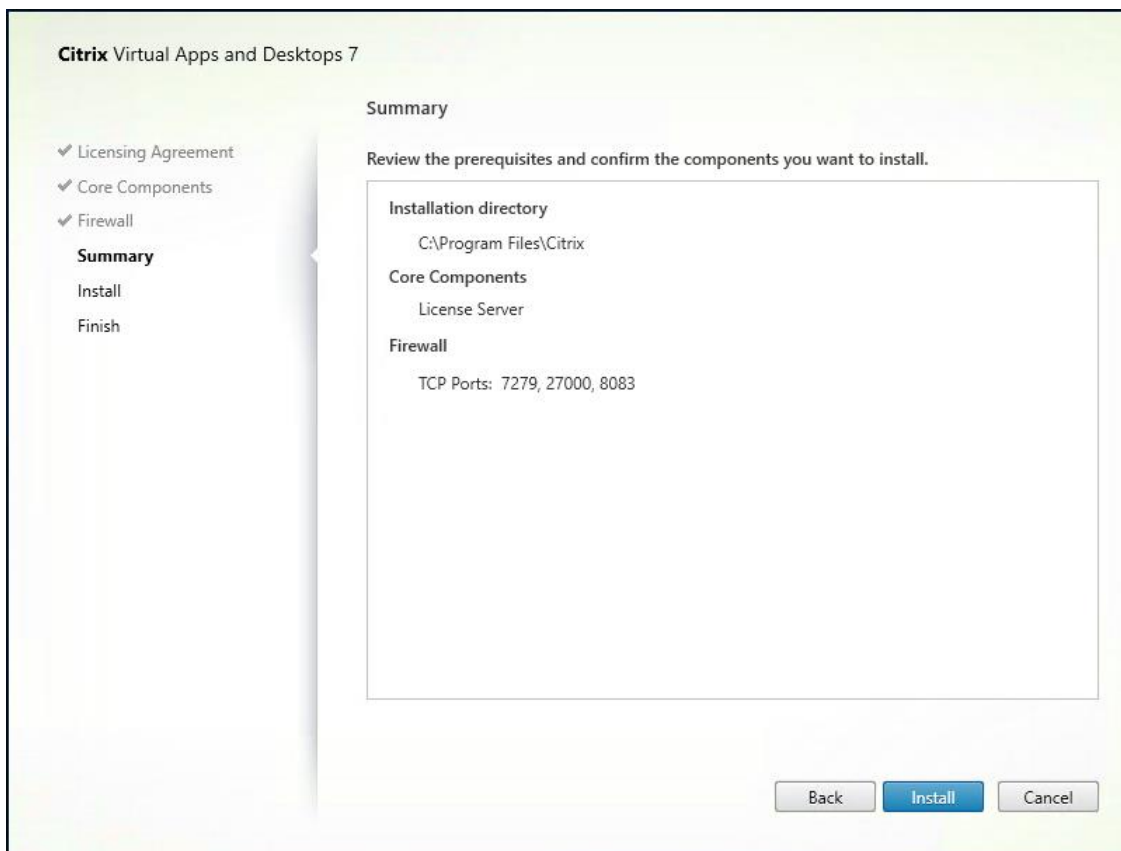


7. Select the default ports and automatically configured firewall rules.

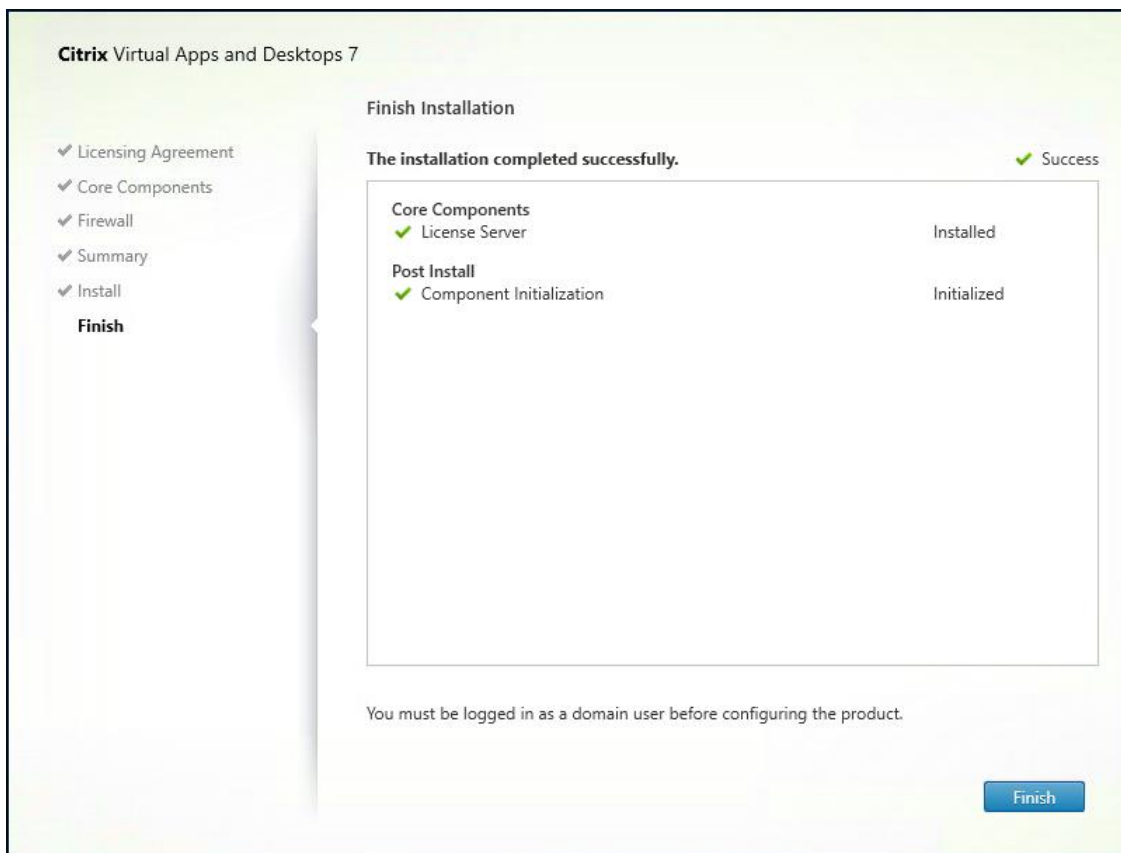
8. Click Next.



9. Click Install.



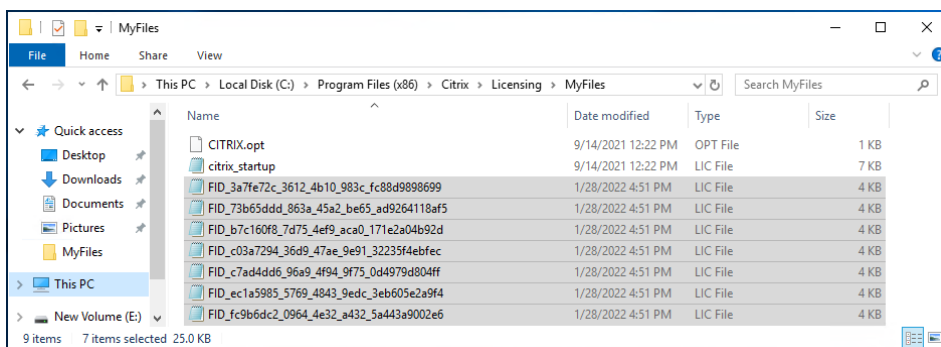
10. Click Finish to complete the installation.



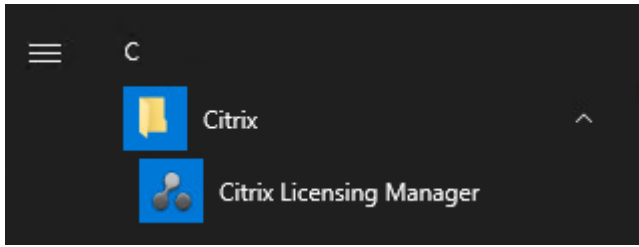
Install Citrix Licenses

To install the Citrix Licenses, follow these steps:

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



2. Restart the server or Citrix licensing services so that the licenses are activated.
3. Run the application Citrix License Administration Console.



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


The screenshot displays the Citrix Licensing Manager dashboard. The top navigation bar includes 'Dashboard', 'Historical Use', 'Install Licenses', and 'Update Licenses'. The main content area is titled 'License Usage' and contains a table with the following data:

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

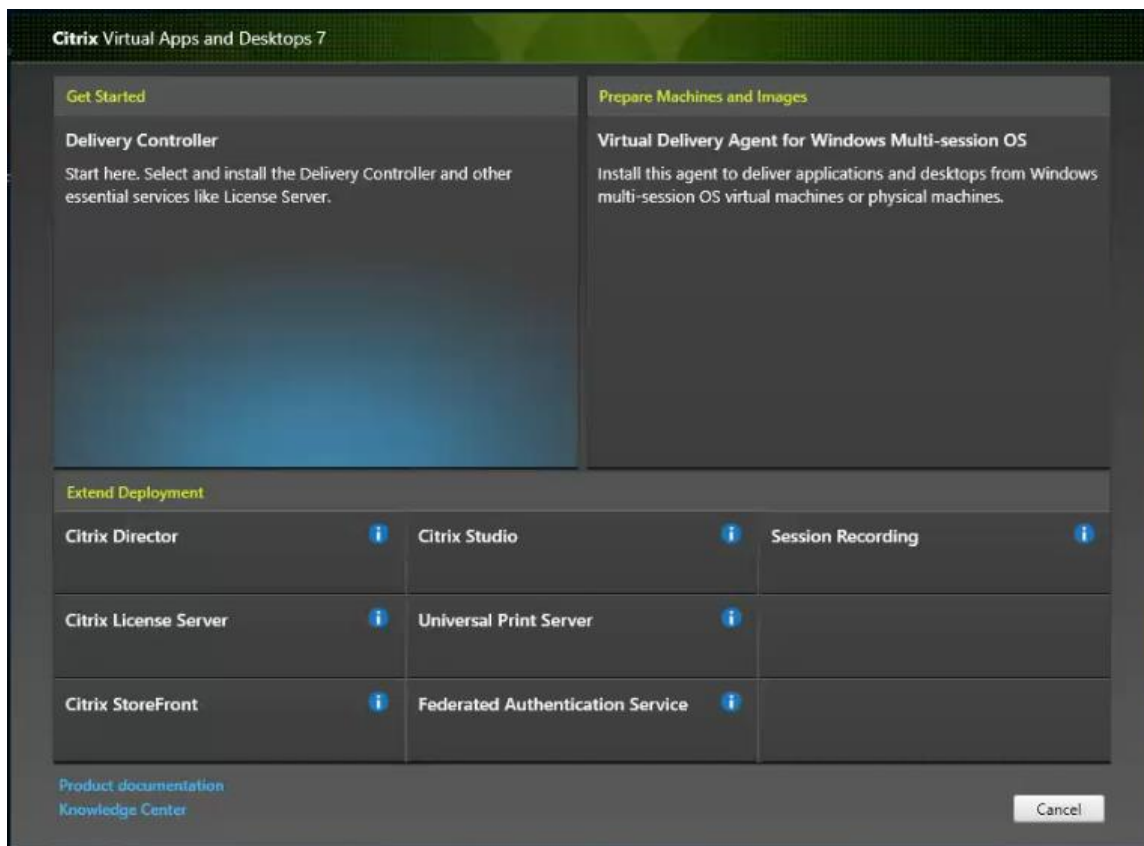
Install the Citrix Virtual Apps and Desktops

To begin the installation, connect to the first Delivery Controller server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2109 ISO, and follow these steps:

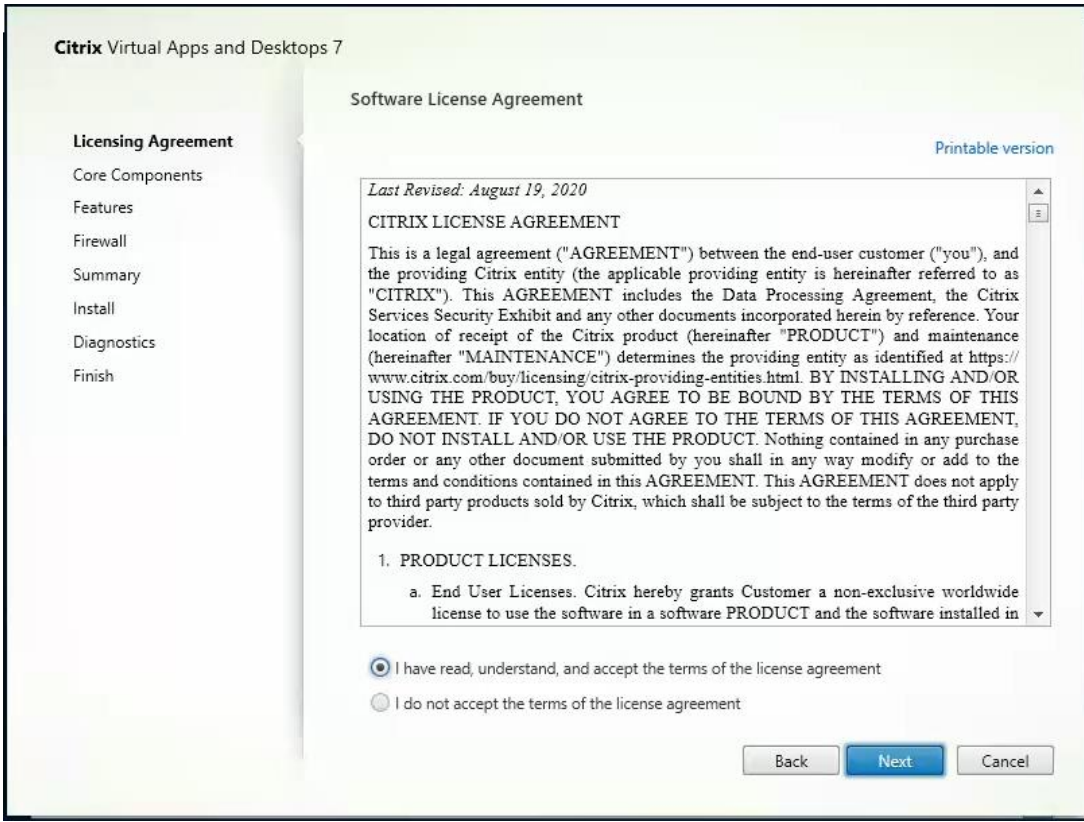
1. Click Start.



2. The installation wizard presents a menu with three subsections. Click Get Started - Delivery Controller.



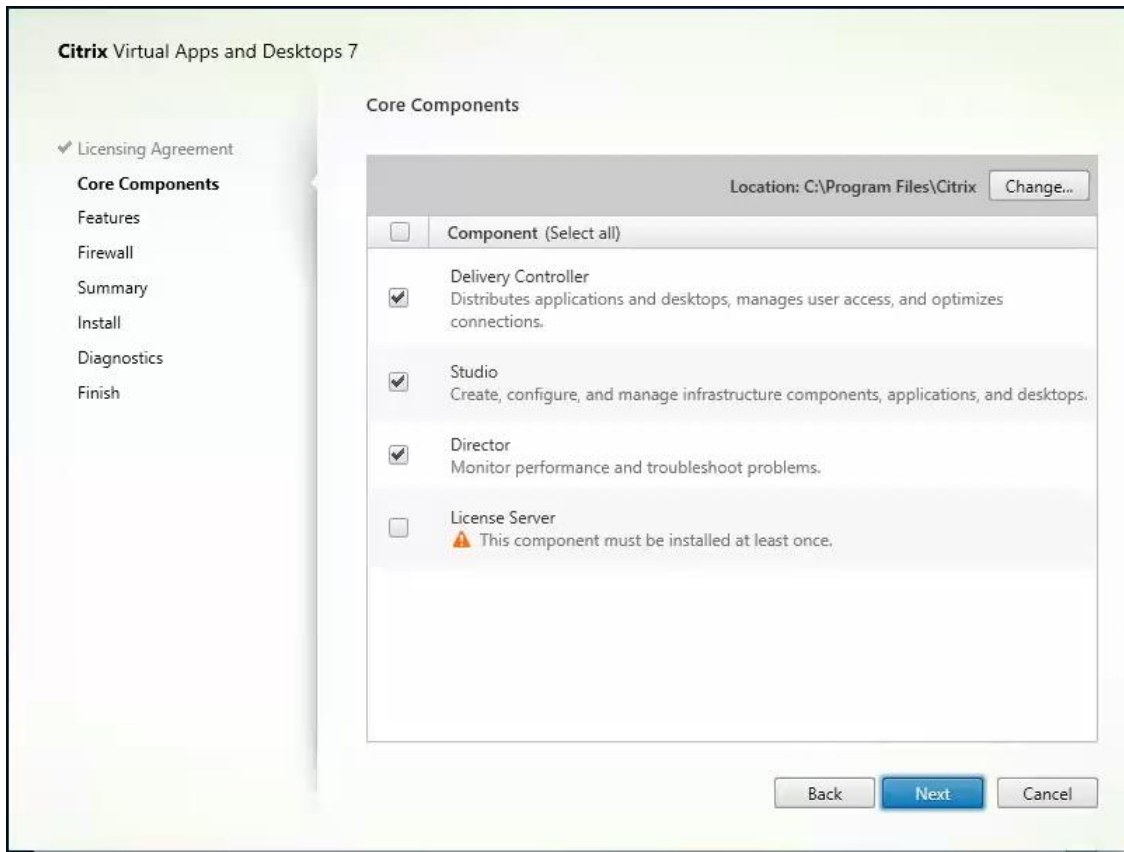
3. Read the Citrix License Agreement. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.
4. Click Next.



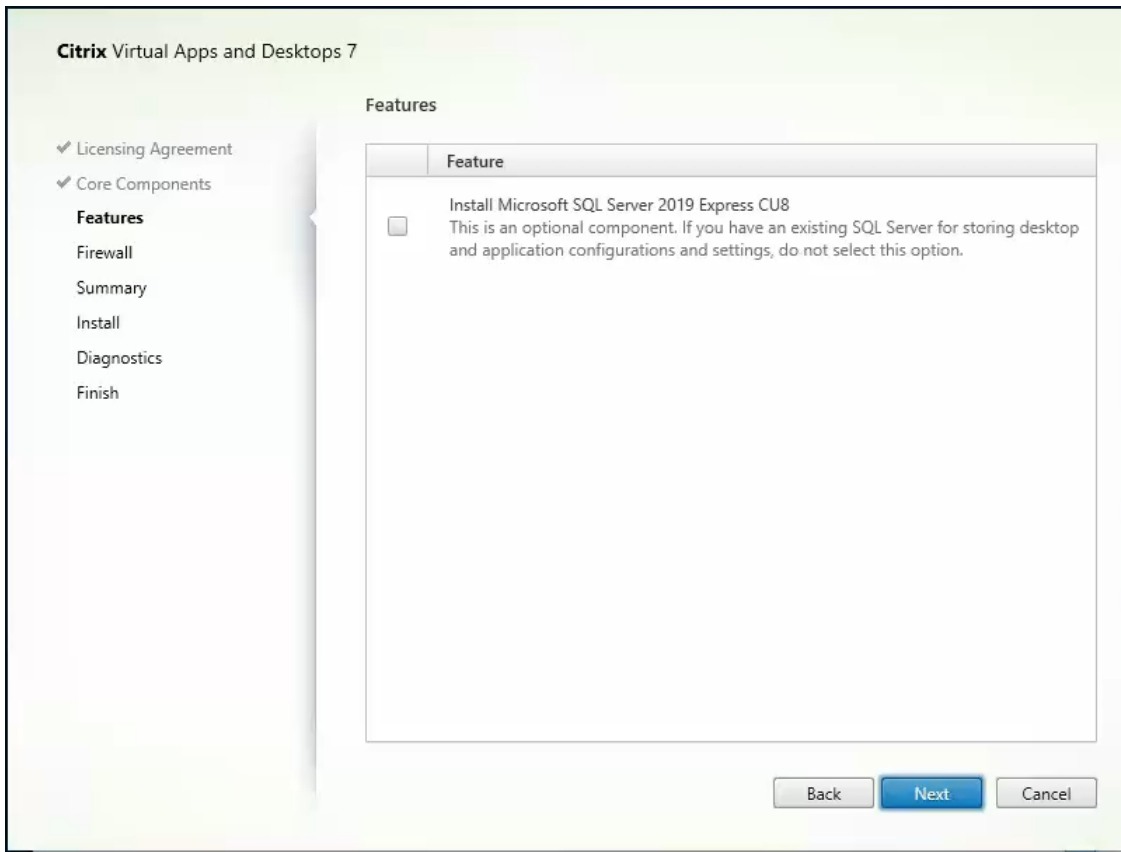
5. Select the components to be installed on the first Delivery Controller Server:

- Delivery Controller
- Studio
- Director

6. Click Next.

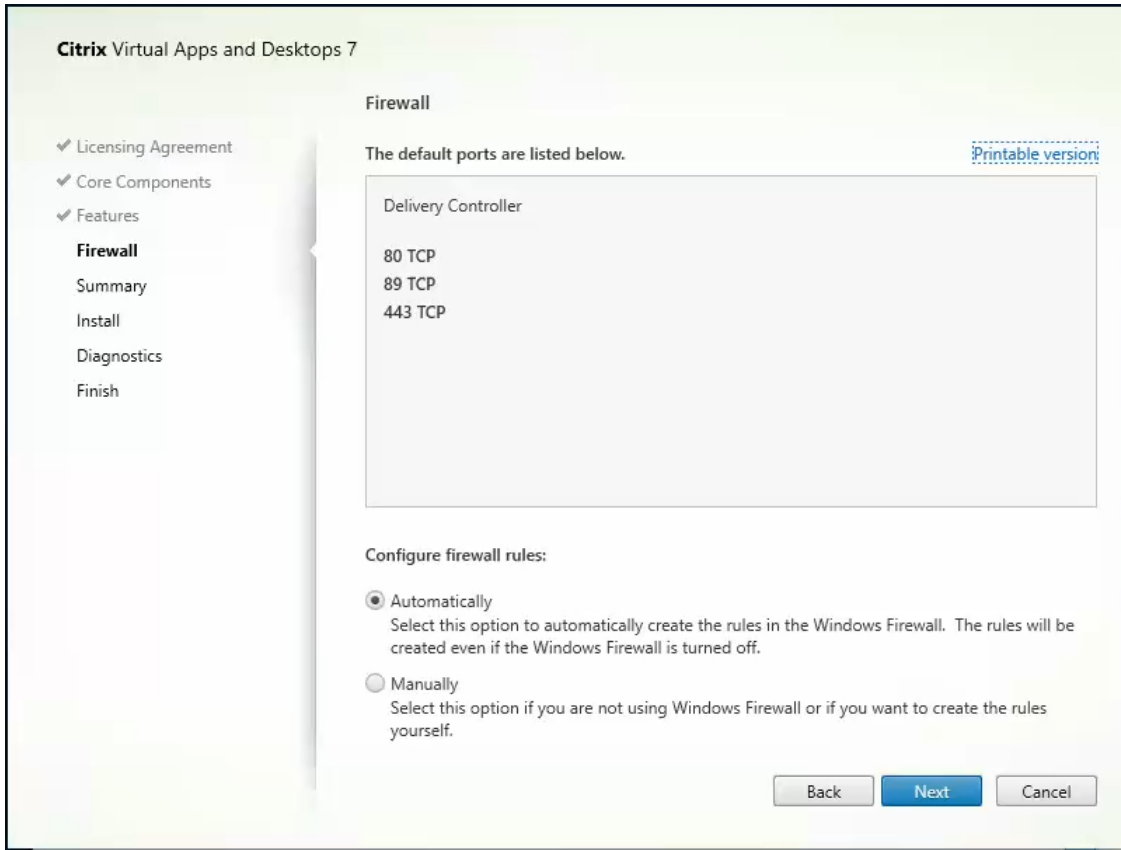


7. Since a dedicated SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2014 SP2 Express” unchecked.
8. Click Next.

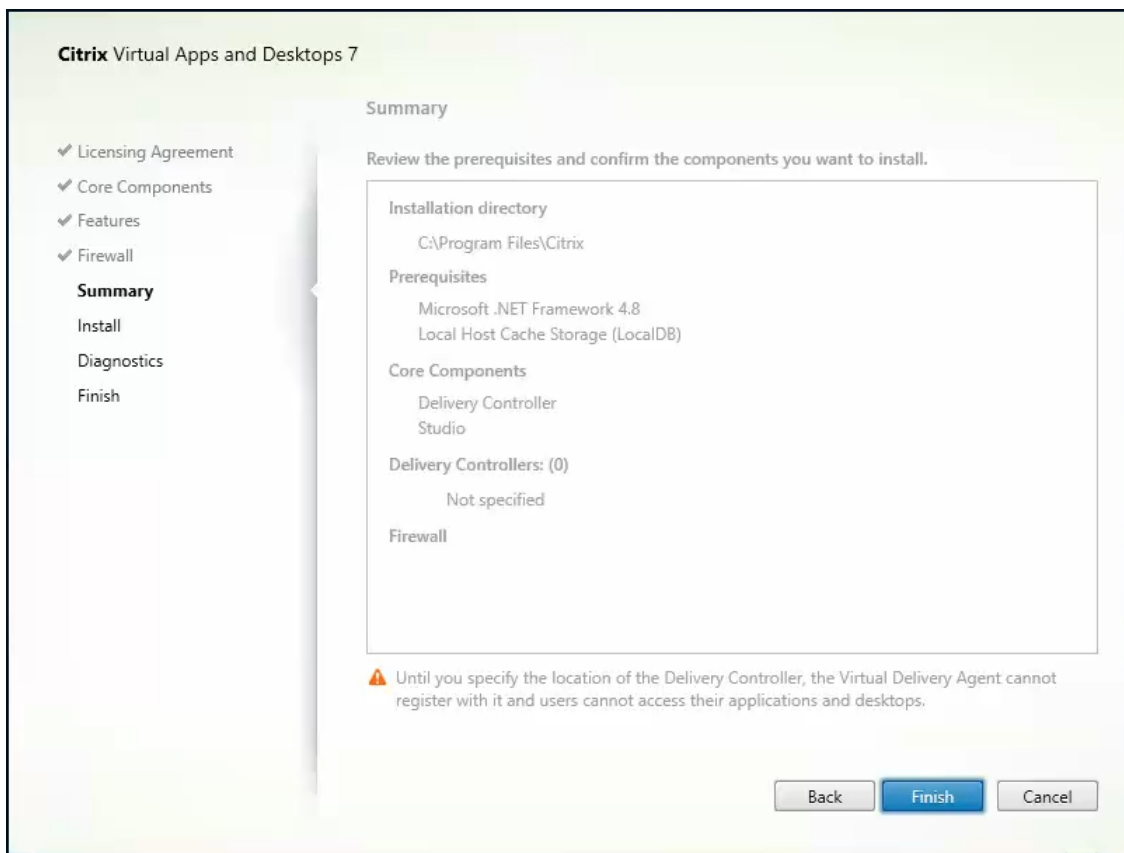


9. Select the default ports and automatically configured firewall rules.

10. Click Next.



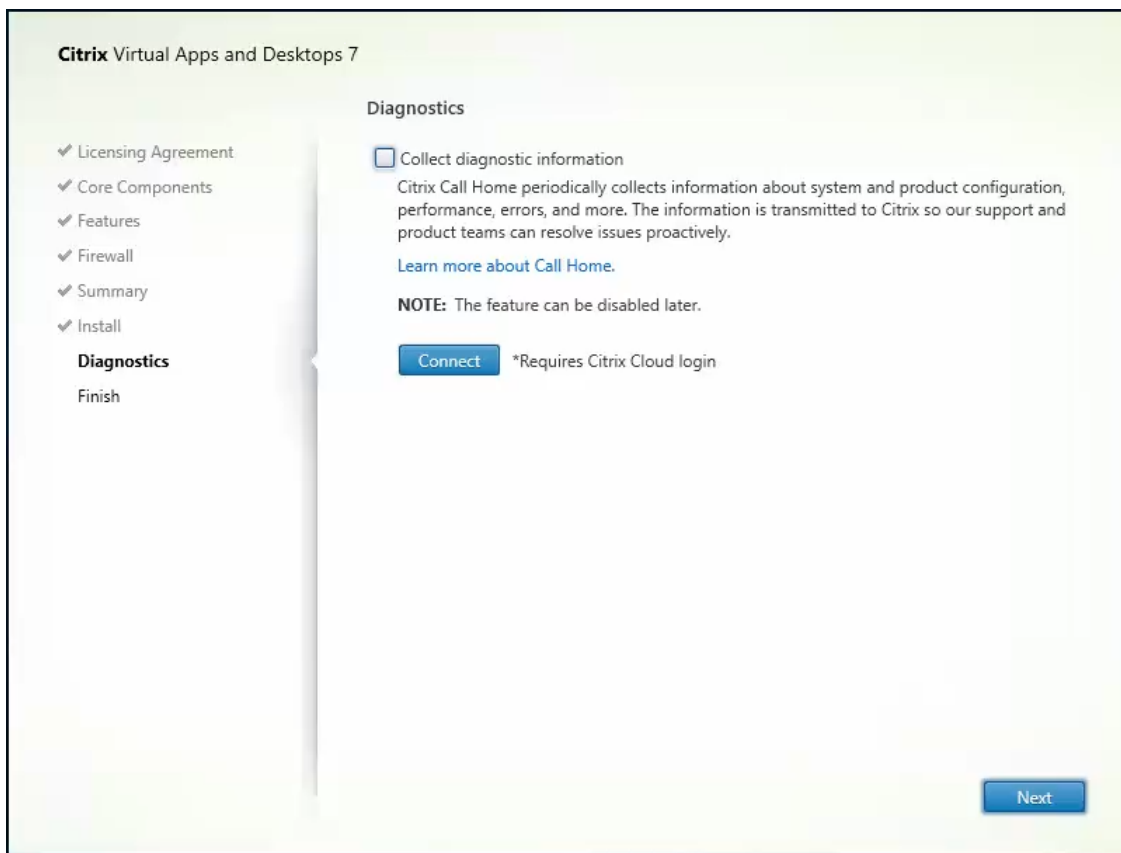
11. Click Finish to begin the installation.



Note: Multiple reboots may be required to finish installation.

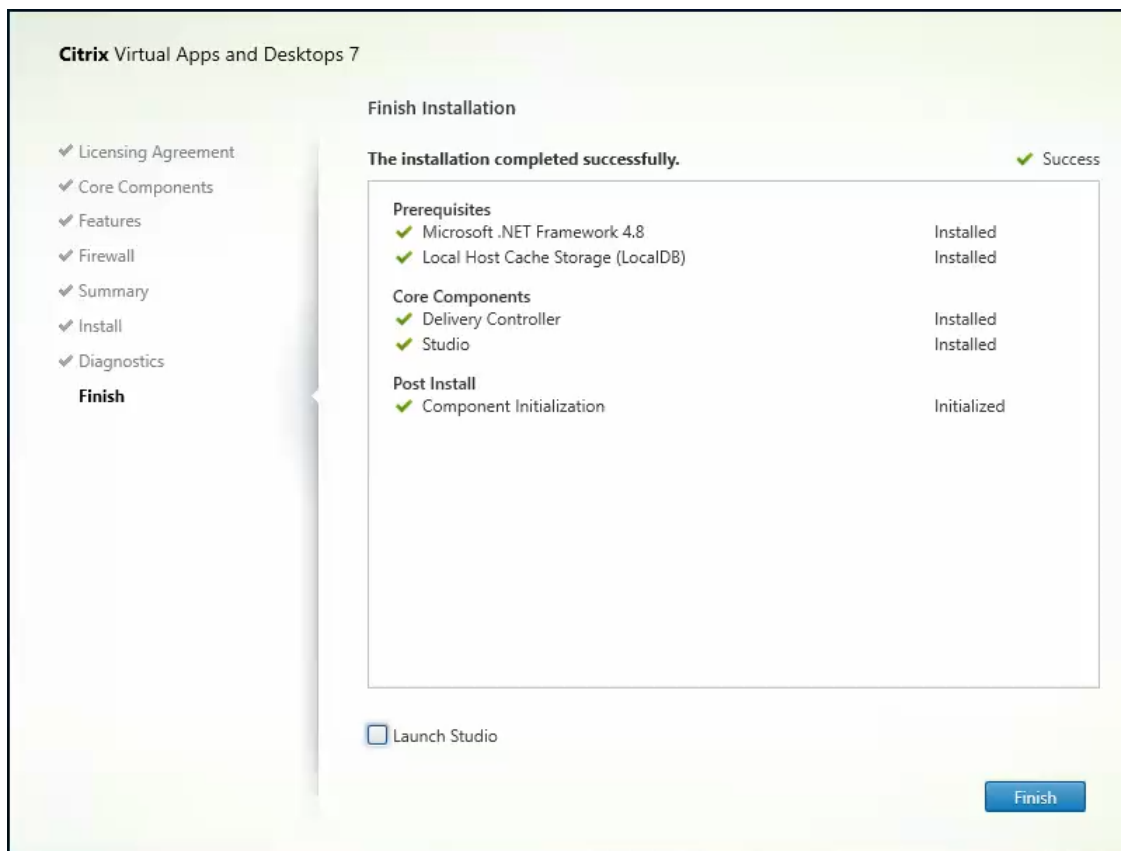
12. (Optional) Collect diagnostic information/Call Home participation.

13. Click Next.



14. Click Finish to complete the installation.

15. (Optional) Check Launch Studio to launch Citrix Studio Console.



Additional Delivery Controller Configuration

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 detailed in [Install the Citrix Virtual Apps and Desktops](#).

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_2109 ISO.

1. Click Start.
2. Click Delivery Controller.
3. 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.
4. Review the Summary configuration. Click Finish.
5. (Optional) Configure Collect diagnostic information /Call Home participation. Click Next.

6. Verify the components installed successfully. Click Finish.

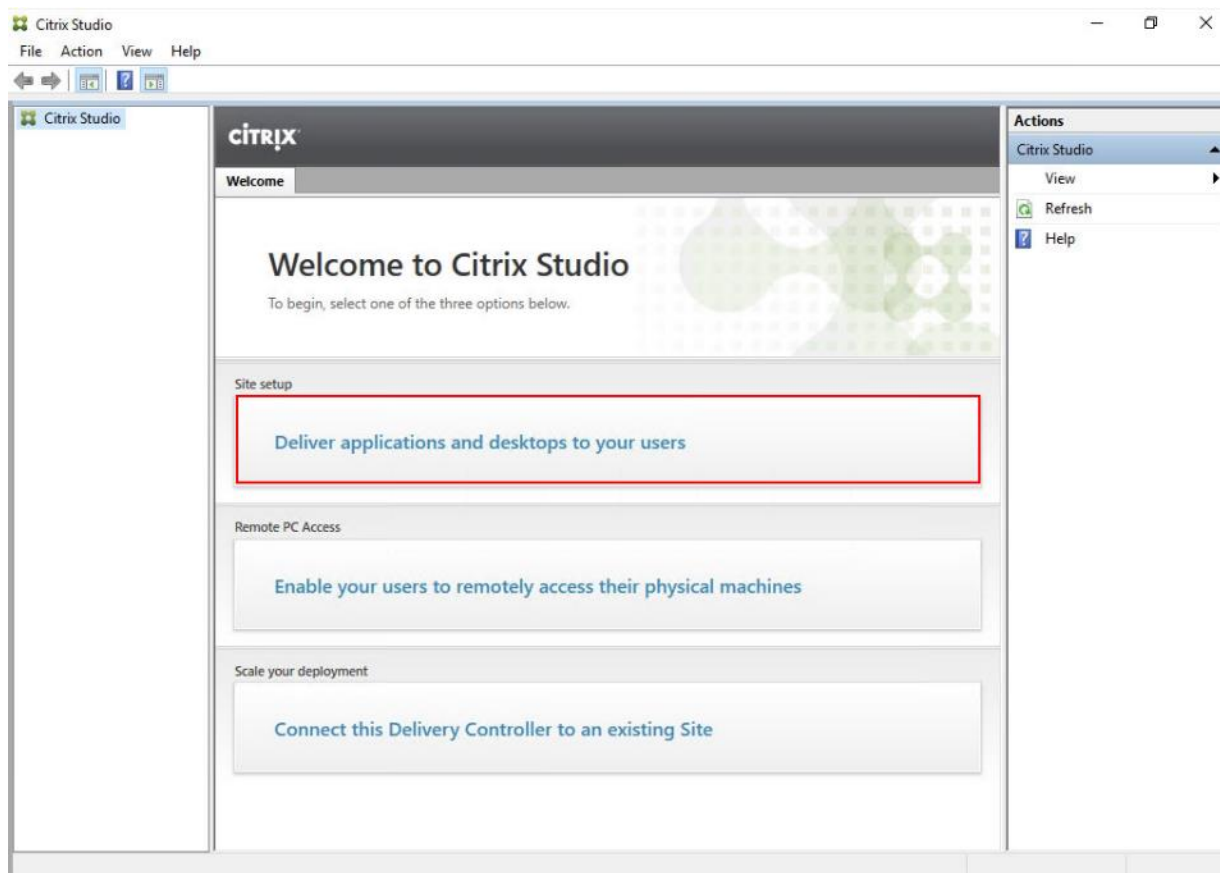
Create Site

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.

To create Site, follow these steps:

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



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

3. Enter a site name.

4. Click Next.

Site Setup

Studio

- Introduction**
- Databases
- Licensing
- Summary

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:

Back Next Cancel

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

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

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

7. Click Next.

Site Setup

Studio

- Introduction
- Databases**
- Licensing
- Summary

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:	FlashStack-k22Site	FS-SQL-1
Monitoring:	FlashStack-k22Monitoring	FS-SQL-1
Logging:	FlashStack-k22Logging	FS-SQL-1

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

Specify additional Delivery Controllers for this Site [Learn more](#)

1 selected

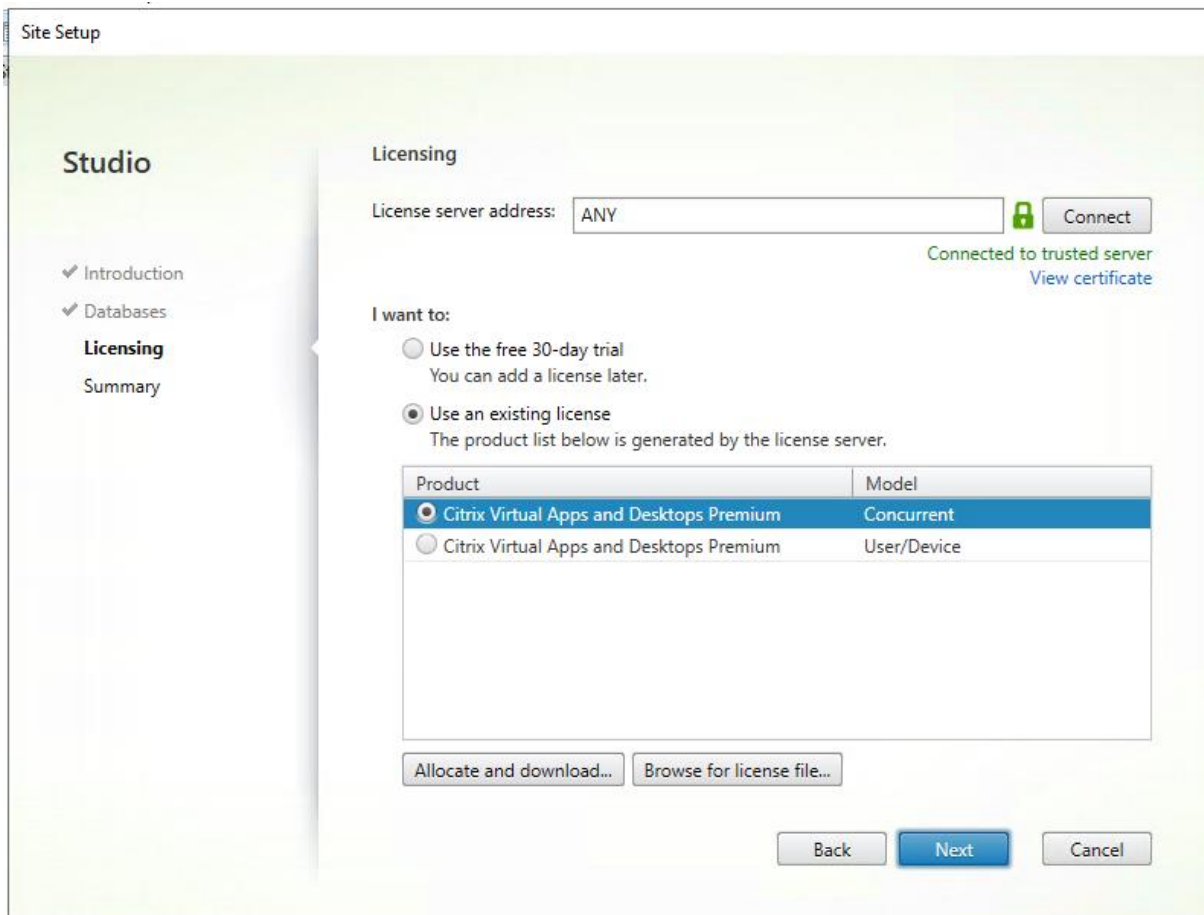
8. Provide the FQDN of the license server.

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.

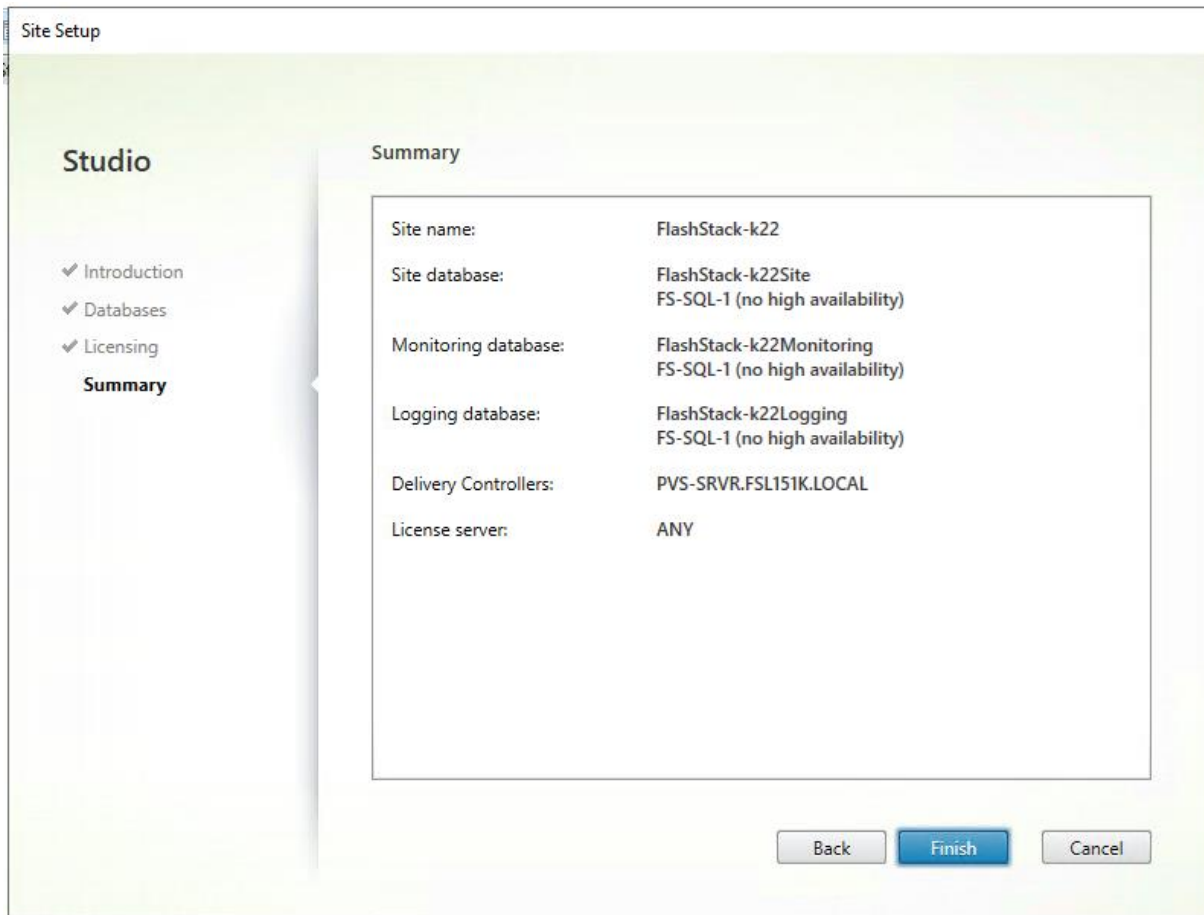
10. Select the appropriate product edition using the license radio button.

11. Click Next.



12. Verify information on the Summary page.

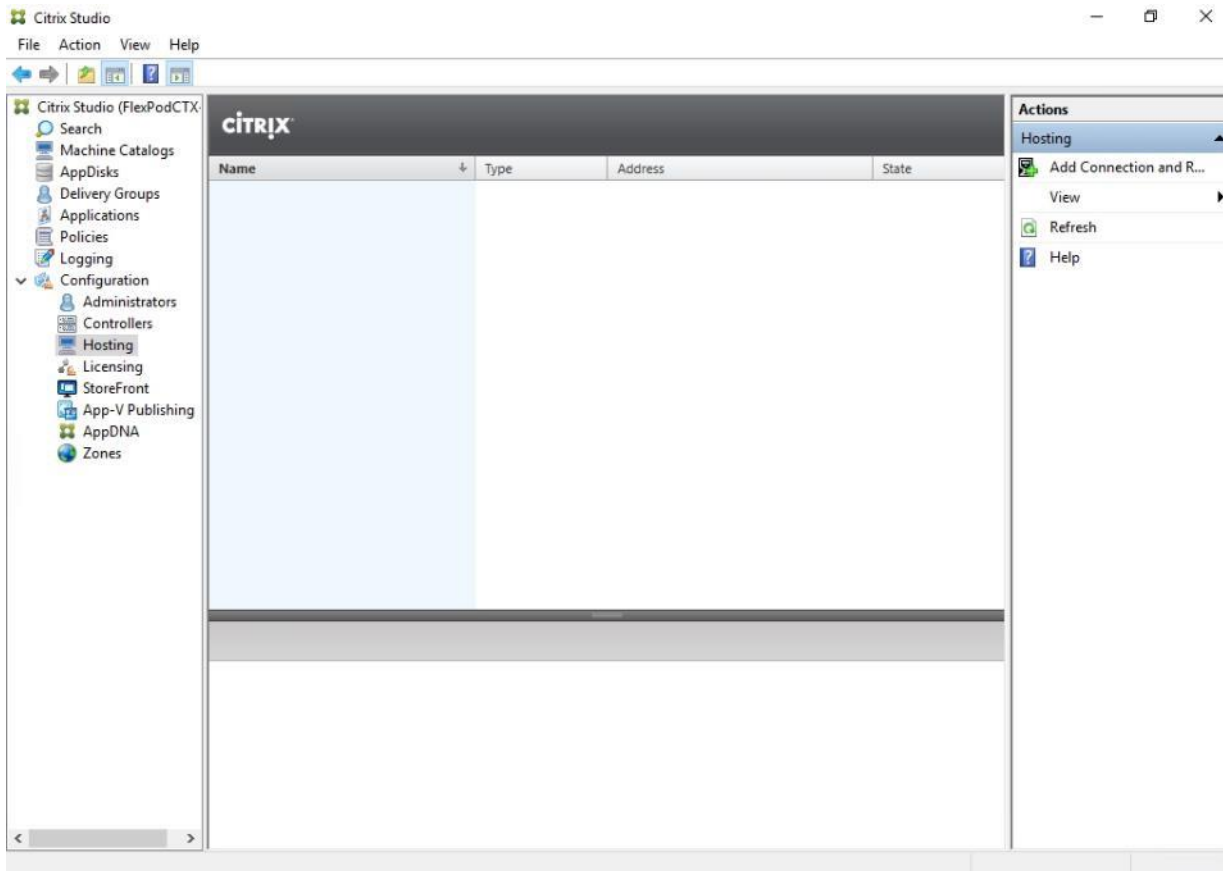
13. Click Finish.



Configure the Citrix Virtual Apps and Desktops Site Hosting Connection

To configure the Citrix Virtual Apps and Desktops site hosting connection, follow these steps:

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



2. On the Connection page:

- Select the Connection type of VMware vSphere®.
- Enter the FQDN of the vCenter server (in Server_FQDN/sdk format).
- Enter the username (in domain\username format) for the vSphere account.
- Provide the password for the vSphere account.
- Provide a connection name.
- Choose the tool to create virtual machines: Machine Creation Services or Citrix Provisioning

3. Click Next.

Add Connection and Resources

Studio

- Connection**
- Storage Management
- Storage Selection
- Network
- Summary

Connection

Connection type: VMware vSphere®

Connection address: <https://10.10.70.30/sdk>
[Learn about user permissions](#)

User name: administrator@vsphere.local

Password: ●●●●●●

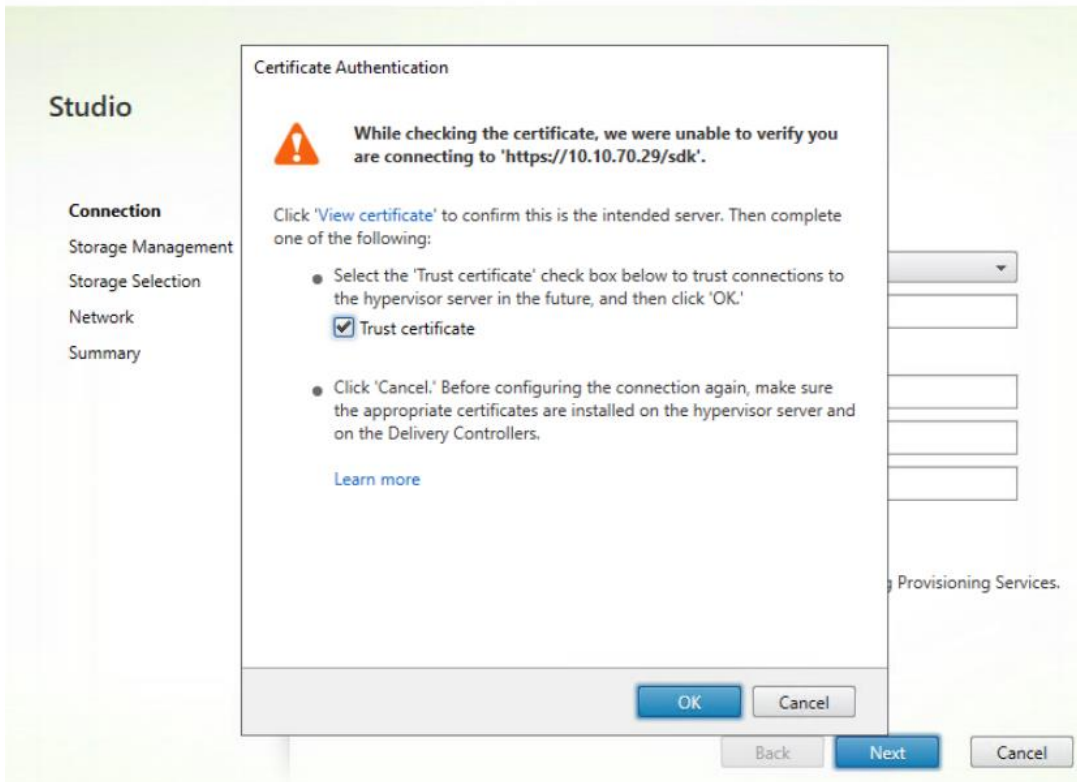
Connection name: FlashStack

Create virtual machines using:

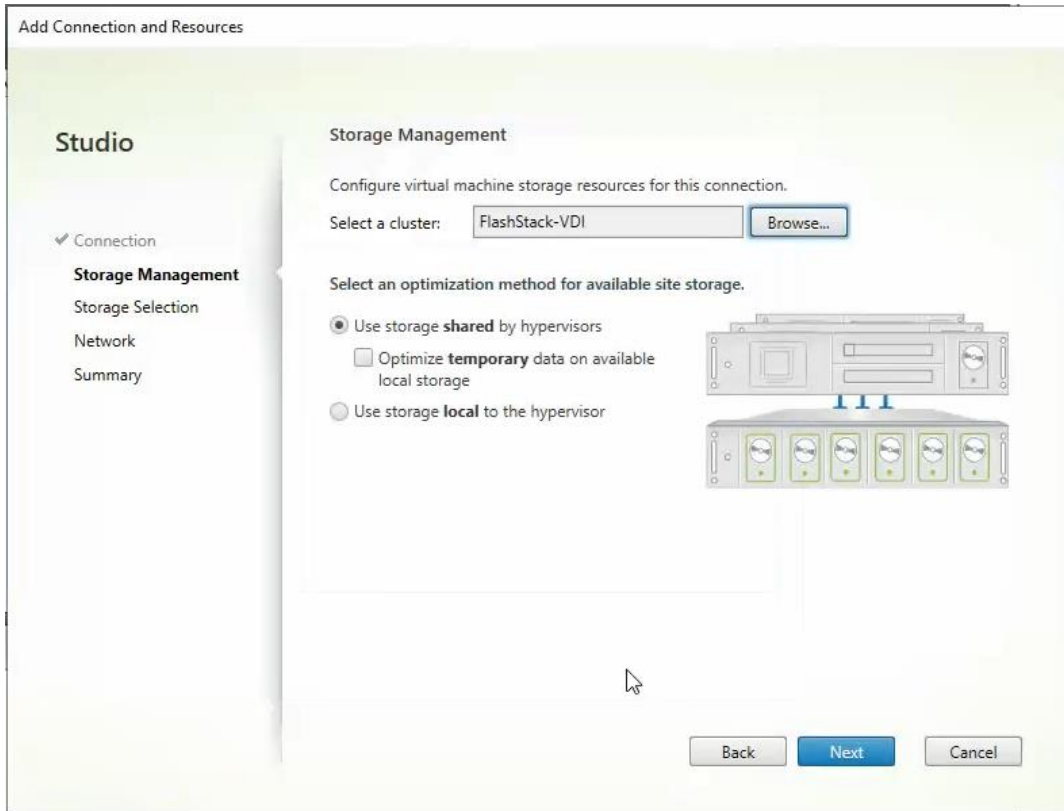
- Citrix provisioning tools (Machine Creation Services or Citrix Provisioning)
- Other tools

Back Next Cancel

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

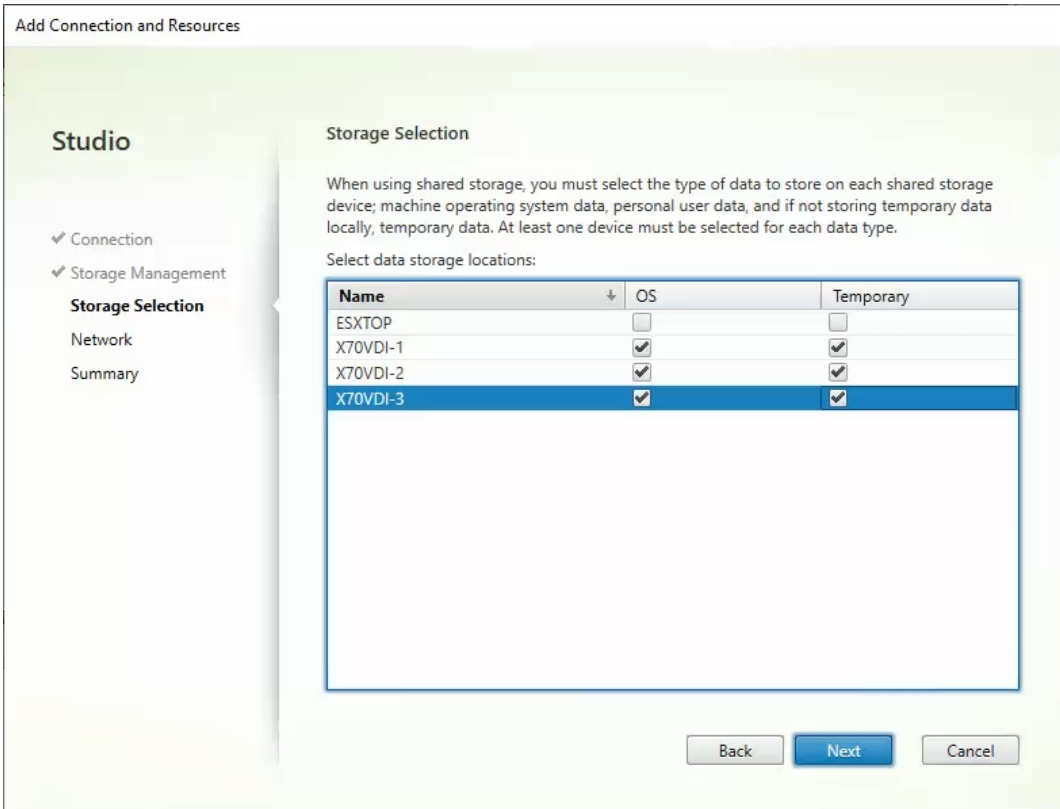


5. Select a storage management method:
6. Select Cluster that will be used by this connection.
7. Check Use storage shared by hypervisors radio button.
8. Click Next.



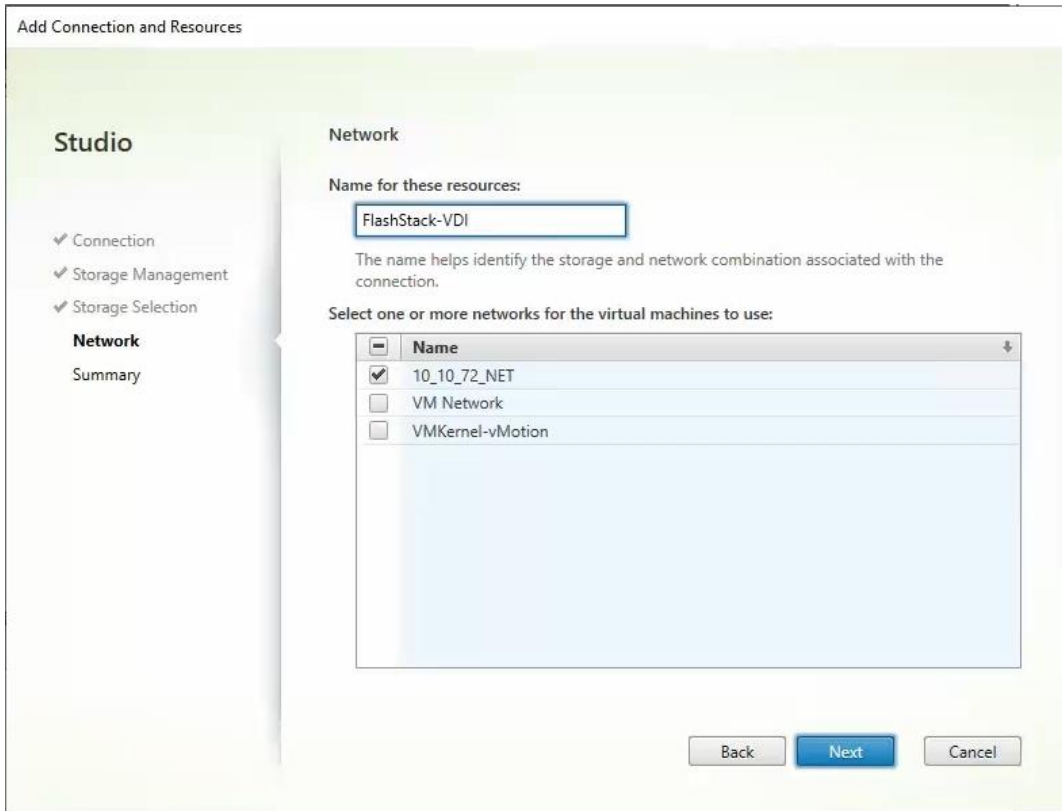
9. Select the Storage to be used by this connection, use all provisioned for desktops datastores.

10. Click Next.



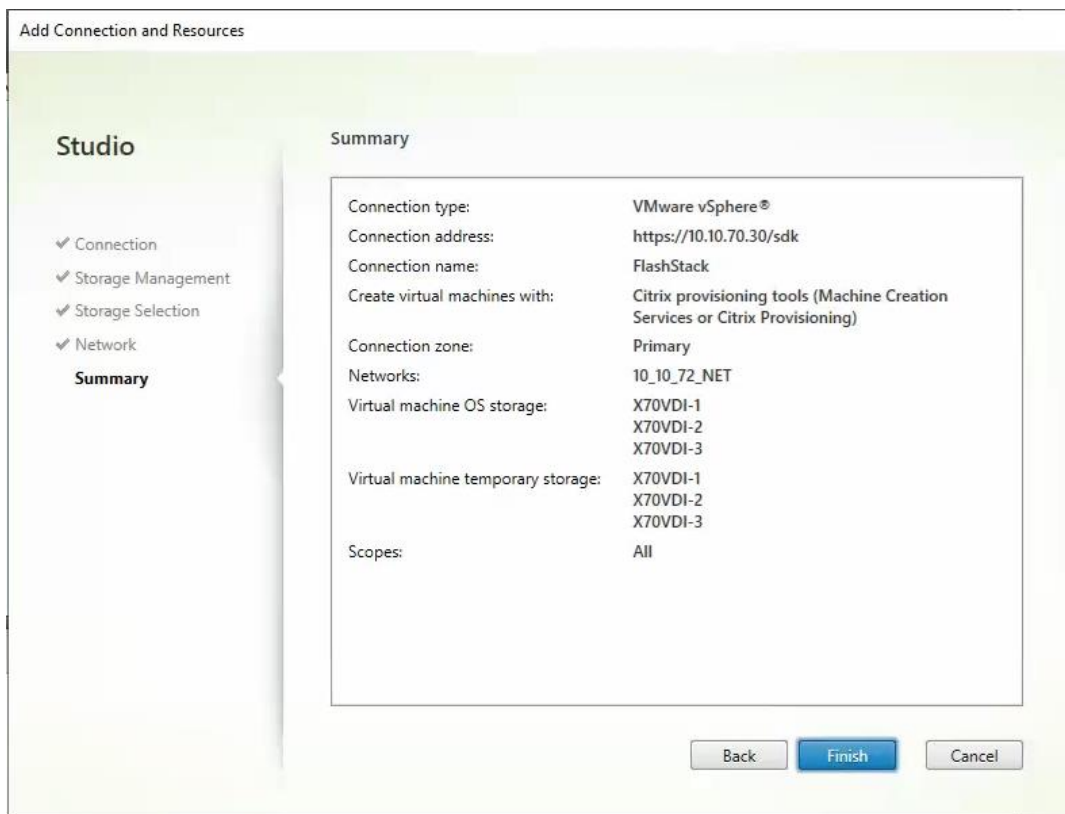
11. Select the Network to be used by this connection.

12. Click Next.



13. Review Add Connection and Resources Summary.


14. Click Finish.



Configure the Citrix Virtual Apps and Desktops Site Administrators

To configure the Citrix Virtual Apps and Desktops site administrators, follow these steps:

1. Connect to the Citrix Virtual Apps and Desktops server and open Citrix Studio Management console.
2. From the Configuration menu, right-click Administrator and select Create Administrator from the drop-down list.



3. Select or Create appropriate scope and click Next.

Create Administrator

Studio

- Administrator and Scope
- Role
- Summary

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
<input checked="" type="radio"/> All All objects

4. Select an appropriate Role.

Create Administrator

Studio

- Administrator and Scope
- Role**
- Summary

Role

Select a role. Click a role name to view its permissions.

Name	Type
<input type="radio"/> Delivery Group Administrator Can deliver applications, desktops, and machines; can also manage the...	Built In
<input checked="" type="radio"/> Full Administrator Can perform all tasks and operations.	Built In
<input type="radio"/> Help Desk Administrator Can view Delivery Groups, and manage the sessions and machines ass...	Built In
<input type="radio"/> Host Administrator Can manage host connections and their associated resource settings.	Built In
<input type="radio"/> Machine Catalog Administrator Can create and manage Machine Catalogs and provision machines.	Built In
<input type="radio"/> Read Only Administrator Can see all objects in specified scopes as well as global information, b...	Built In

5. Review the Summary, check Enable administrator and click Finish.

Create Administrator

Studio

- ✓ Administrator and Scope
- ✓ Role
- Summary**

Summary

Administrator:	VCCFSLAB\Domain Admins
Scope:	All
Role:	Full Administrator

Enable administrator
Clear check box to disable the administrator. No settings will be lost.

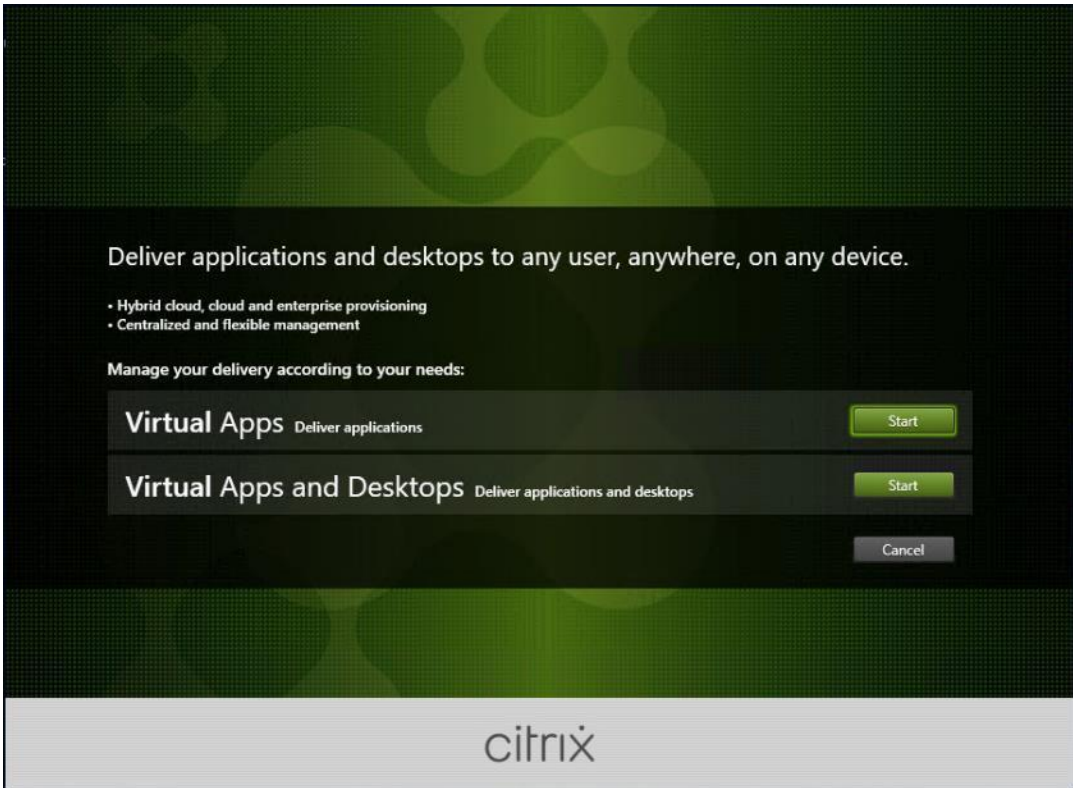
[Save full permissions report](#)

Install and Configure StoreFront

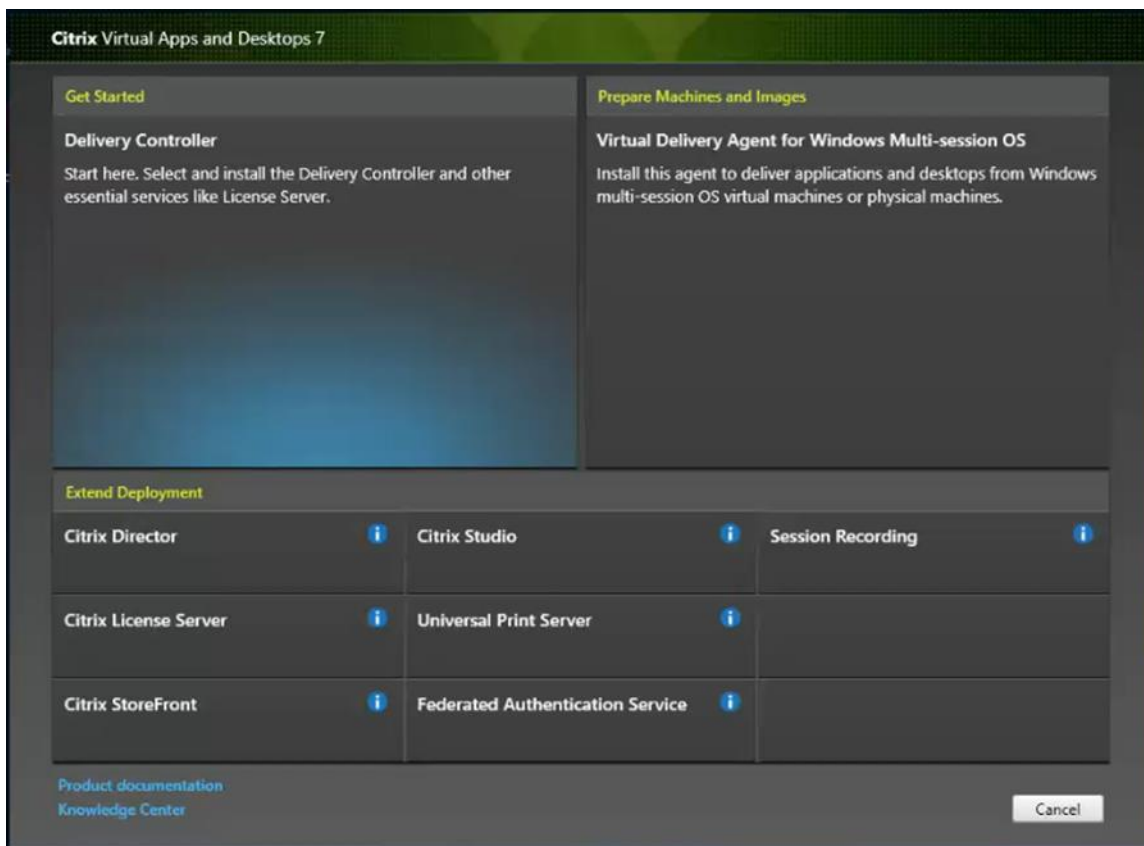
Citrix StoreFront stores aggregate desktops and applications from Citrix Virtual Apps and Desktops sites, making resources readily available to users. In this CVD, we created two StoreFront servers on dedicated virtual machines.

To install and configure StoreFront, follow these steps:

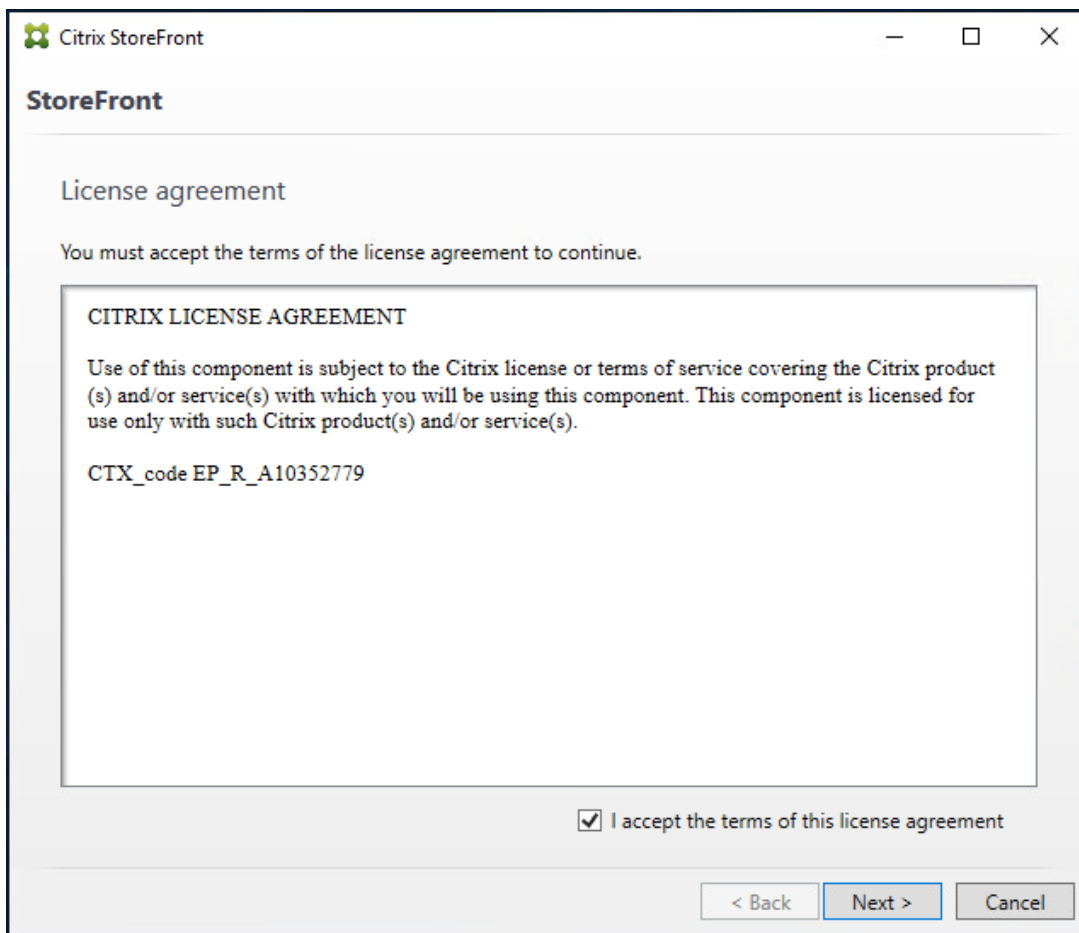
1. To begin the installation of the StoreFront, connect to the first StoreFront server and launch the installer from the Citrix_Virtual_Apps_and_Desktops_7_2109 ISO.
2. Click Start.



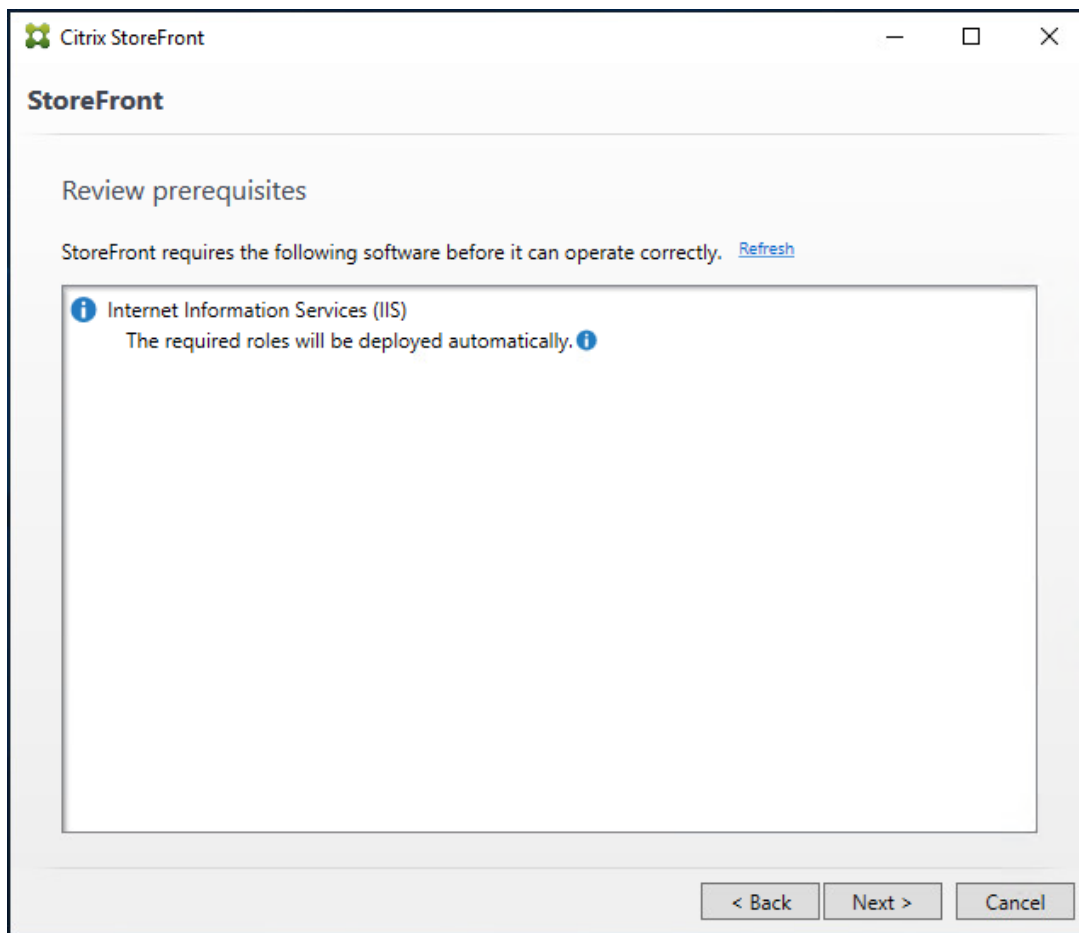
3. Click Extend Deployment Citrix StoreFront.



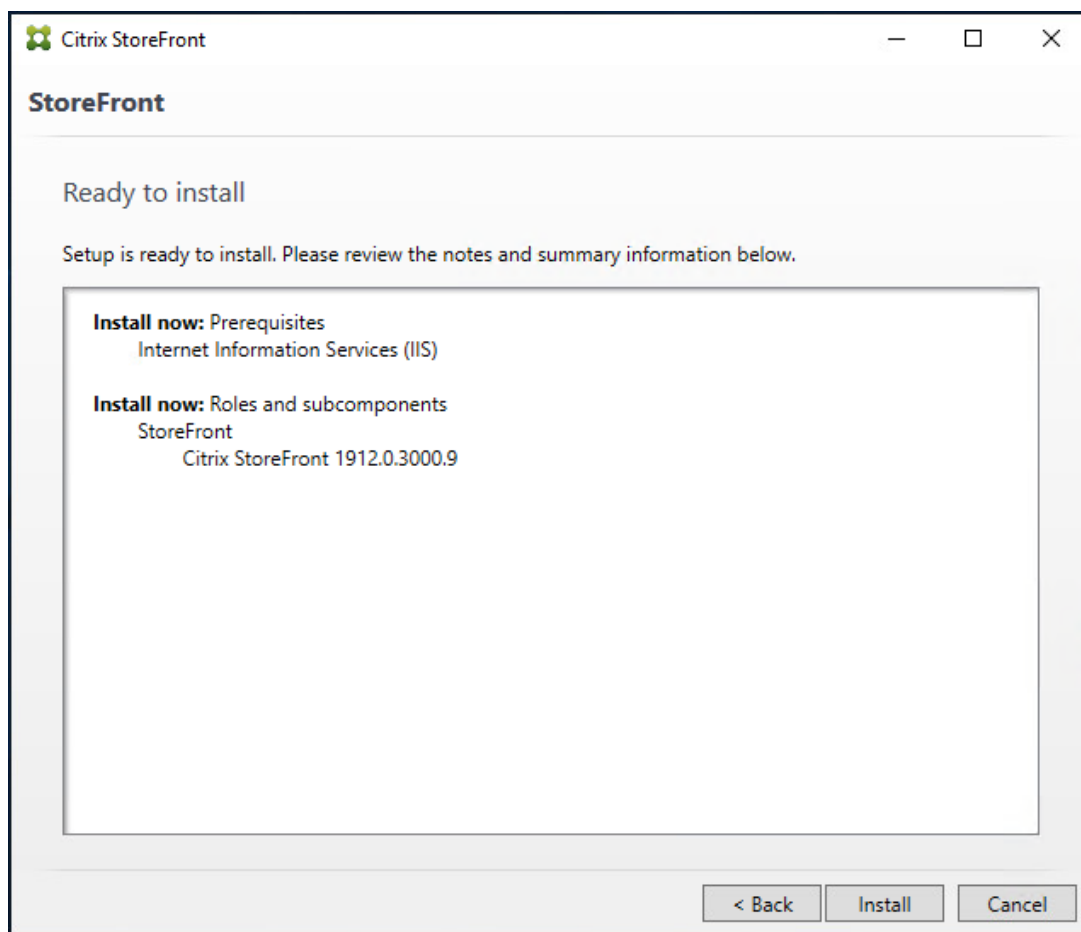
4. Indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement”.
5. Click Next.



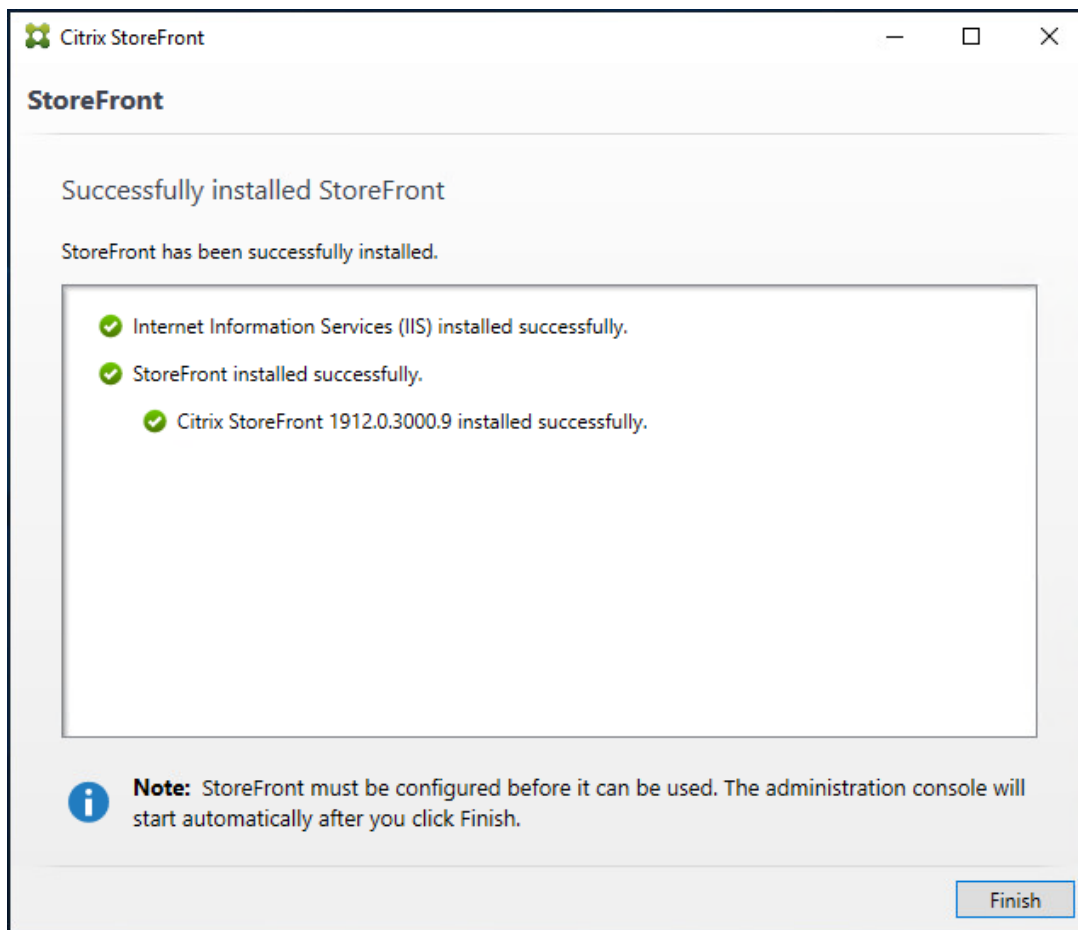
6. On Prerequisites page click Next.



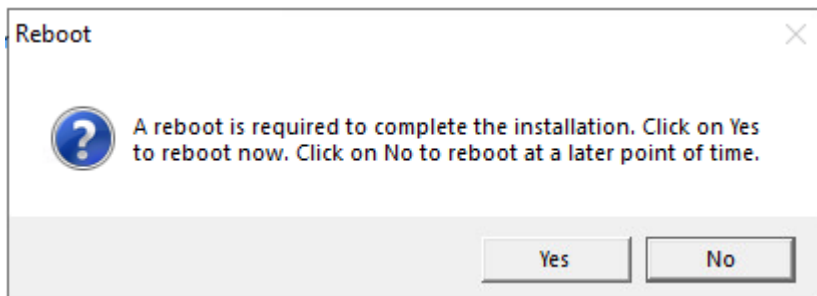
7. Click Install.



8. Click Finish.

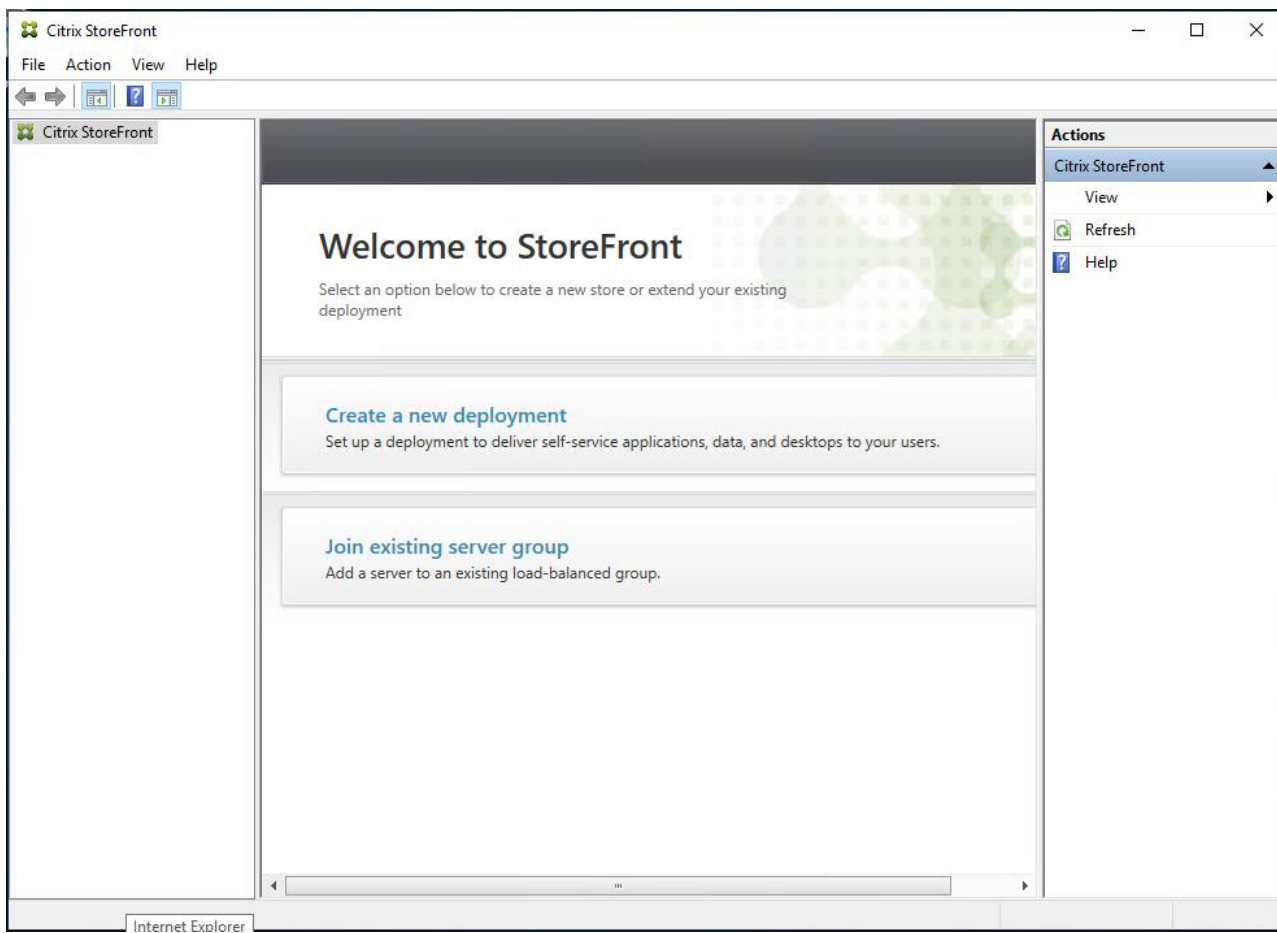


9. Click Yes to reboot the server.



10. Open the StoreFront Management Console.

11. Click Create a new deployment.



12. Specify name for your Base URL.

13. Click Next.

Create New Deployment


StoreFront

Base URL

- Getting Started
- Store Name
- Delivery Controllers
- Remote Access
- Authentication Methods
- XenApp Services URL
- Summary

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: 

Note: For a multiple server deployment use the load balancing environment in the Base URL box.

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.

```
graph LR; Store[Store] -.->|Store URL| CitrixReceiver[Citrix Receiver]; Store -.->|Receiver for Web Site| Browser[Browser]; Store -.->|XenApp Services URL| PNAgent[PNAgent]; CitrixReceiver -.-> EndUser[End User]; Browser -.-> EndUser; PNAgent -.-> EndUser;
```

Next Cancel

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.

i Store name and access type cannot be changed, once the store is created.

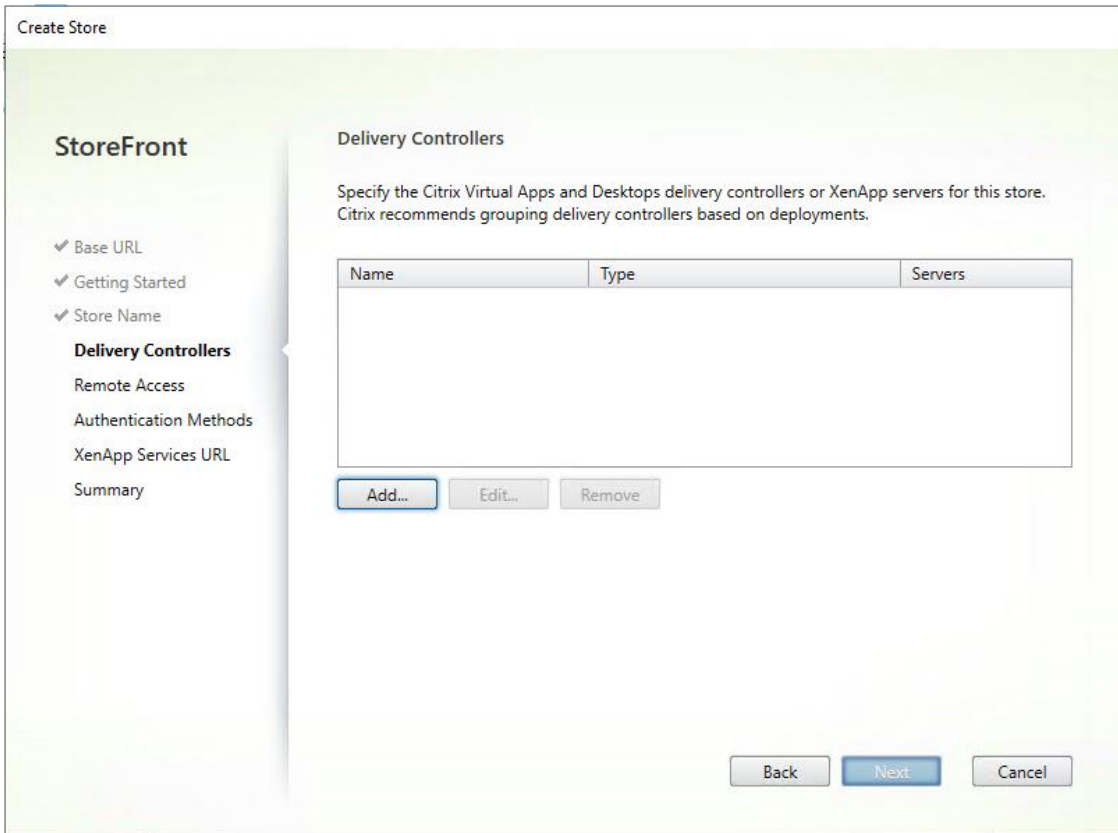
Store Name:

Allow only unauthenticated (anonymous) users to access this store
Unauthenticated users can access the store without presenting credentials.

Receiver for Web Site Settings

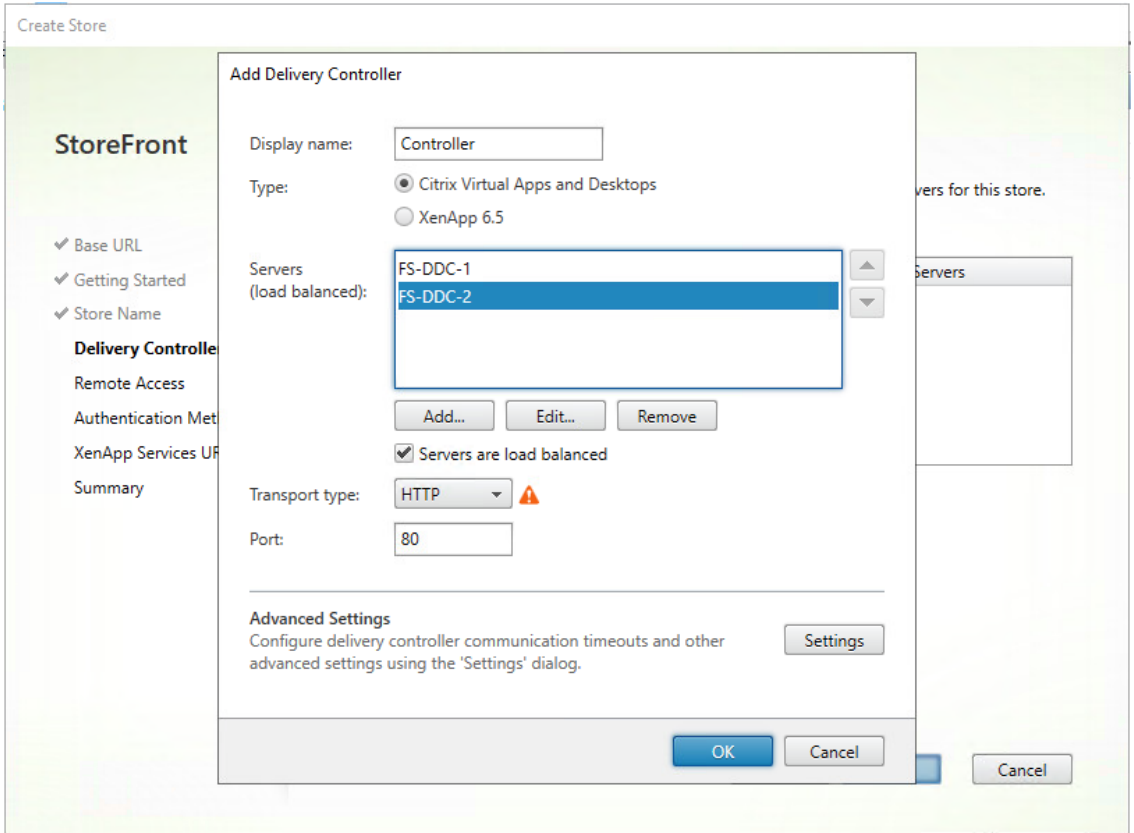
Set this Receiver for Web site as IIS default
When this is checked, the Receiver for Web site created with the store will be set as the default IIS website. This setting will override any previous defaults configured for the IIS sites.

16. Click Add to specify Delivery controllers for your new Store.

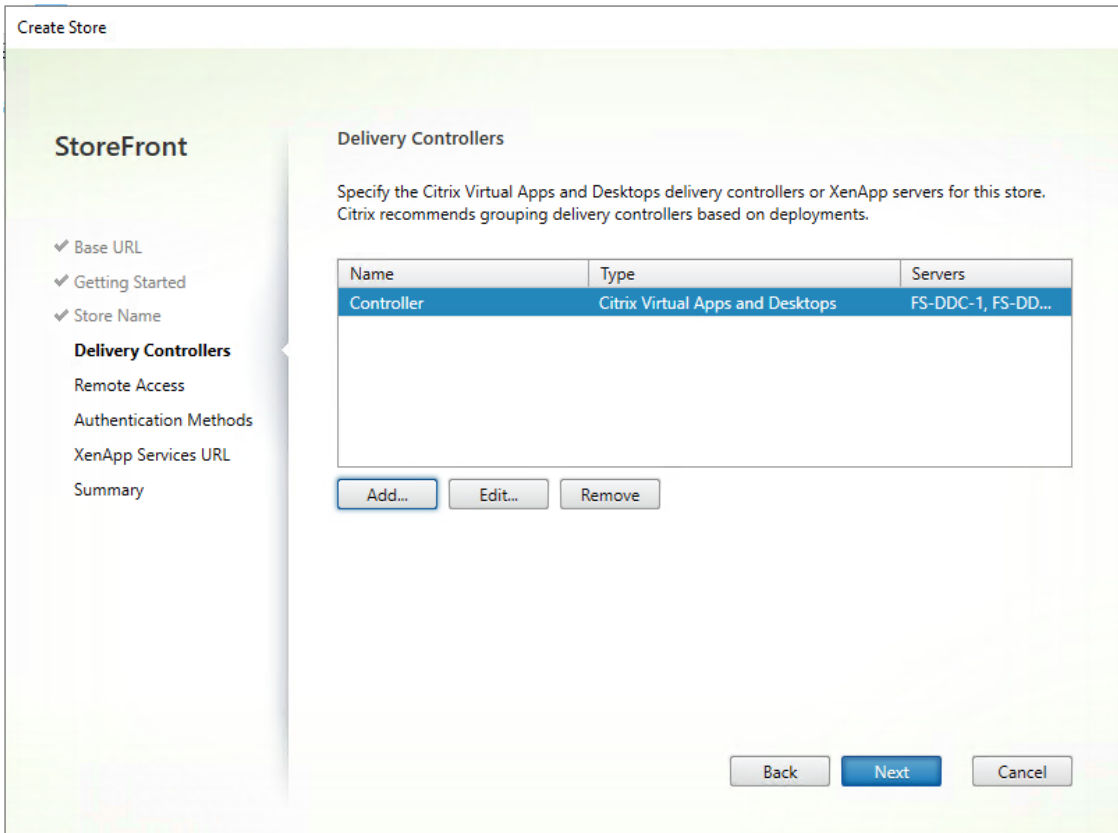


17. Add the required Delivery Controllers to the store.

18. Click OK.



19. Click Next.



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.

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.

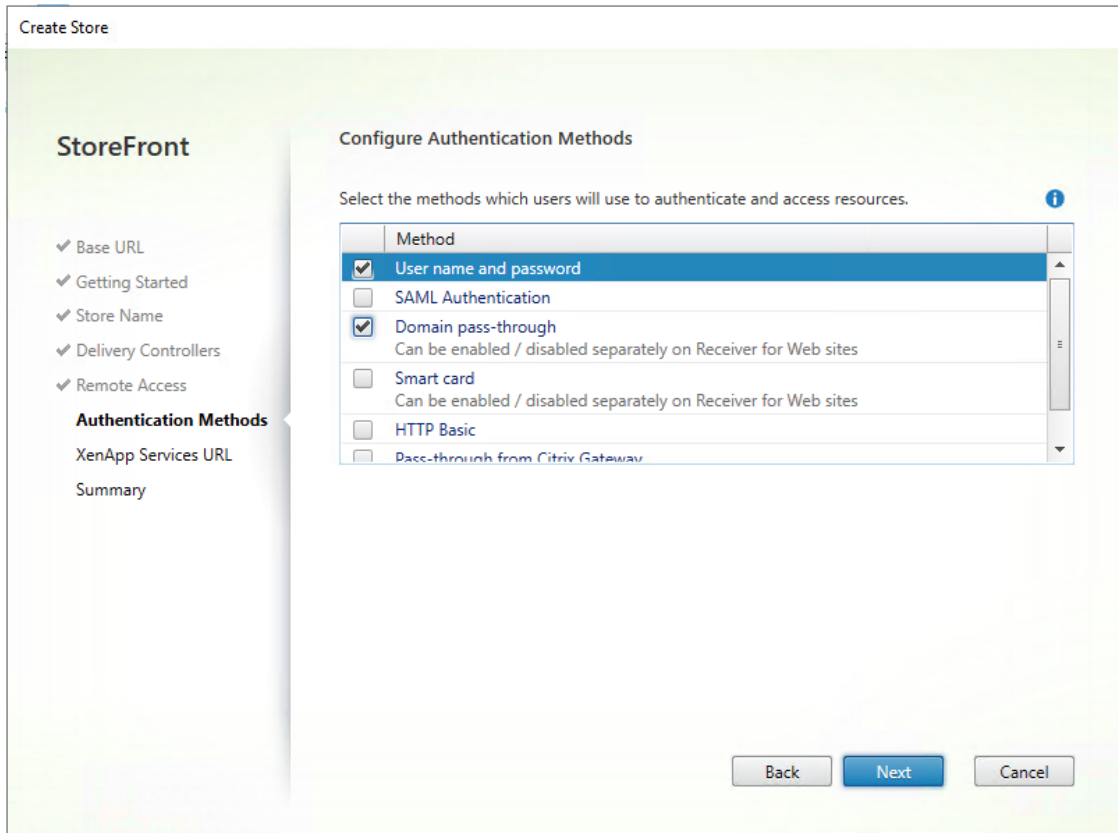
Citrix Gateway appliances: ⓘ

Default appliance:

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:

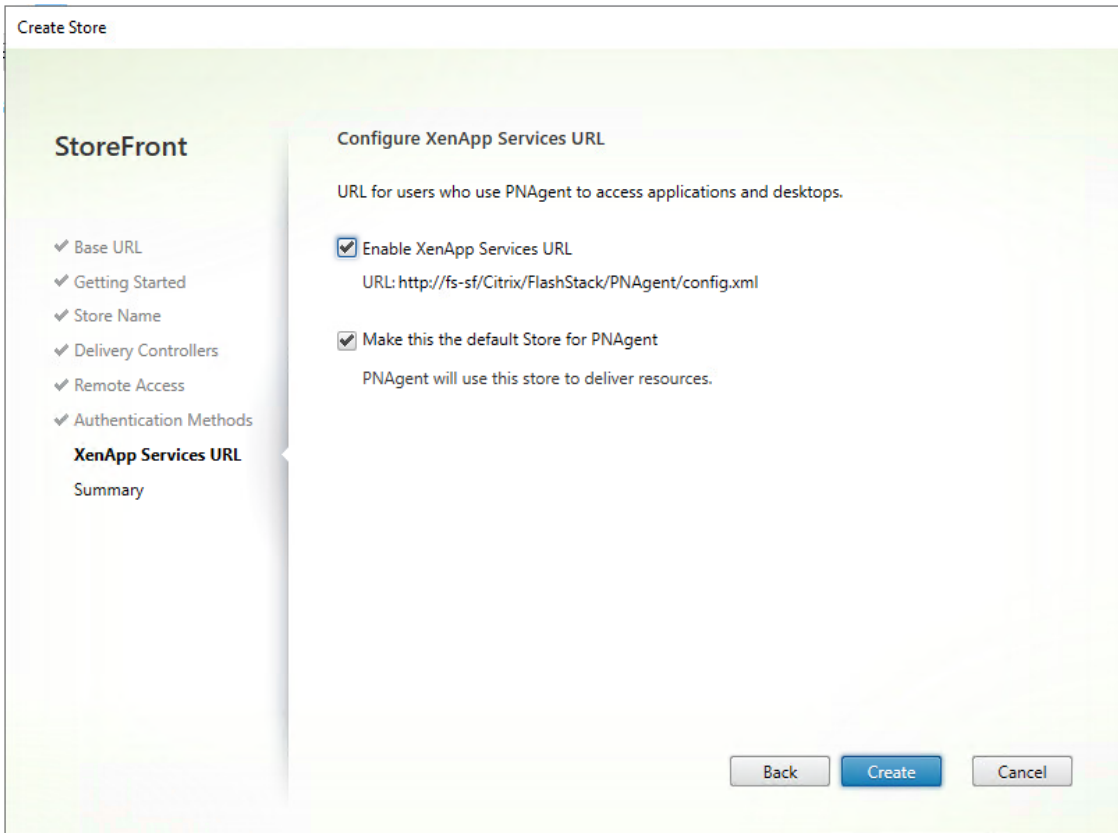
- Username and password: Users enter their credentials and are authenticated when they access their stores.
- Domain passthrough: Users authenticate to their domain-joined Windows computers and their credentials are used to log them on automatically when they access their stores.

23. Click Next.

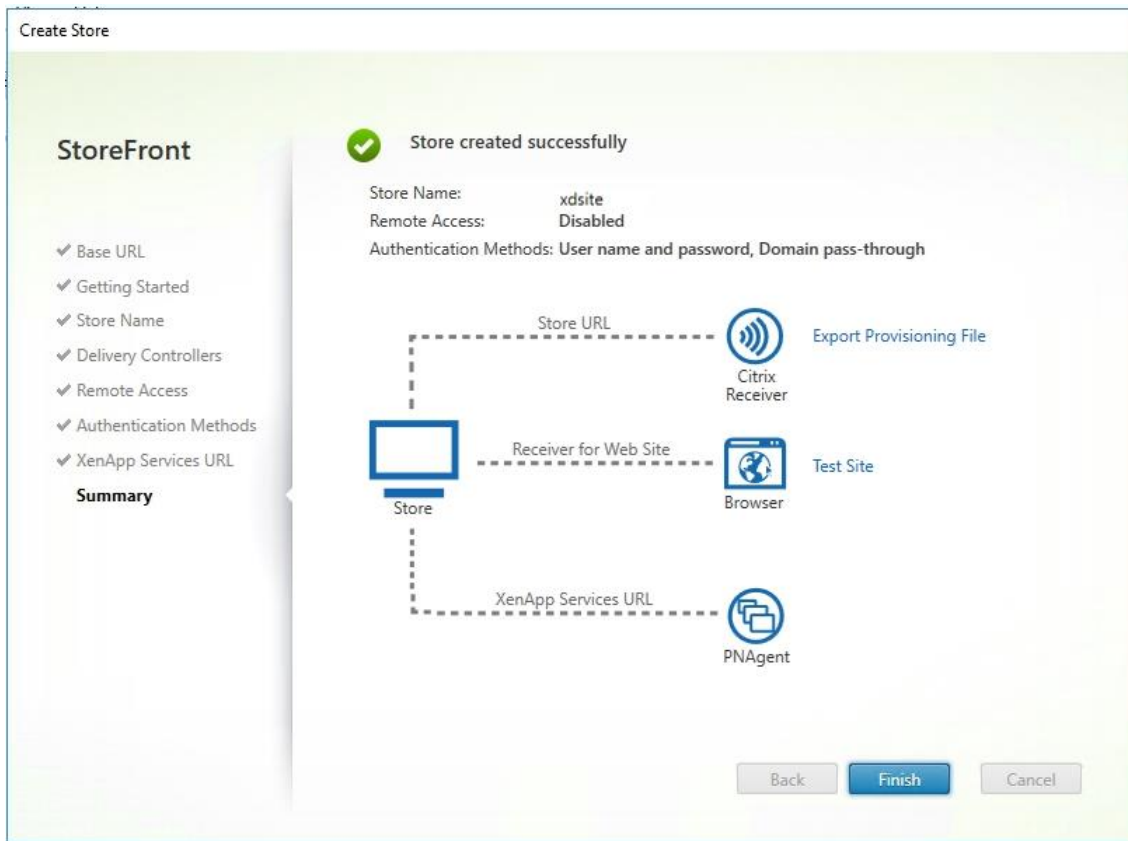


24. Configure the XenApp Service URL for users who use PNAgent to access the applications and desktops.

25. Click Create.



26. After creating the store click Finish.




Additional StoreFront Configuration

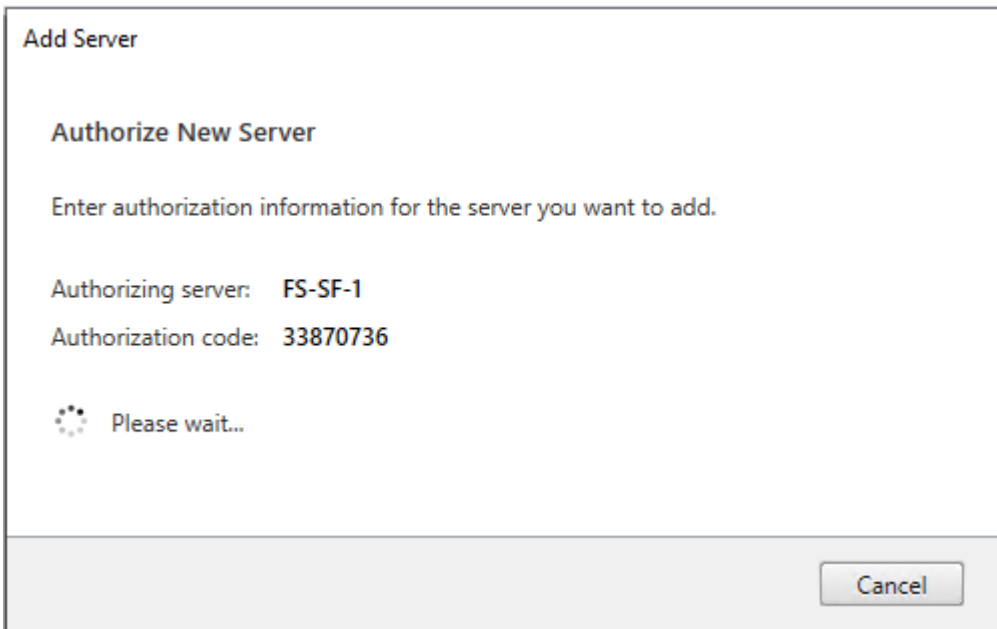
After the first StoreFront server is completely configured and the Store is operational, you can add additional servers.

To configure additional StoreFront servers, follow these steps:

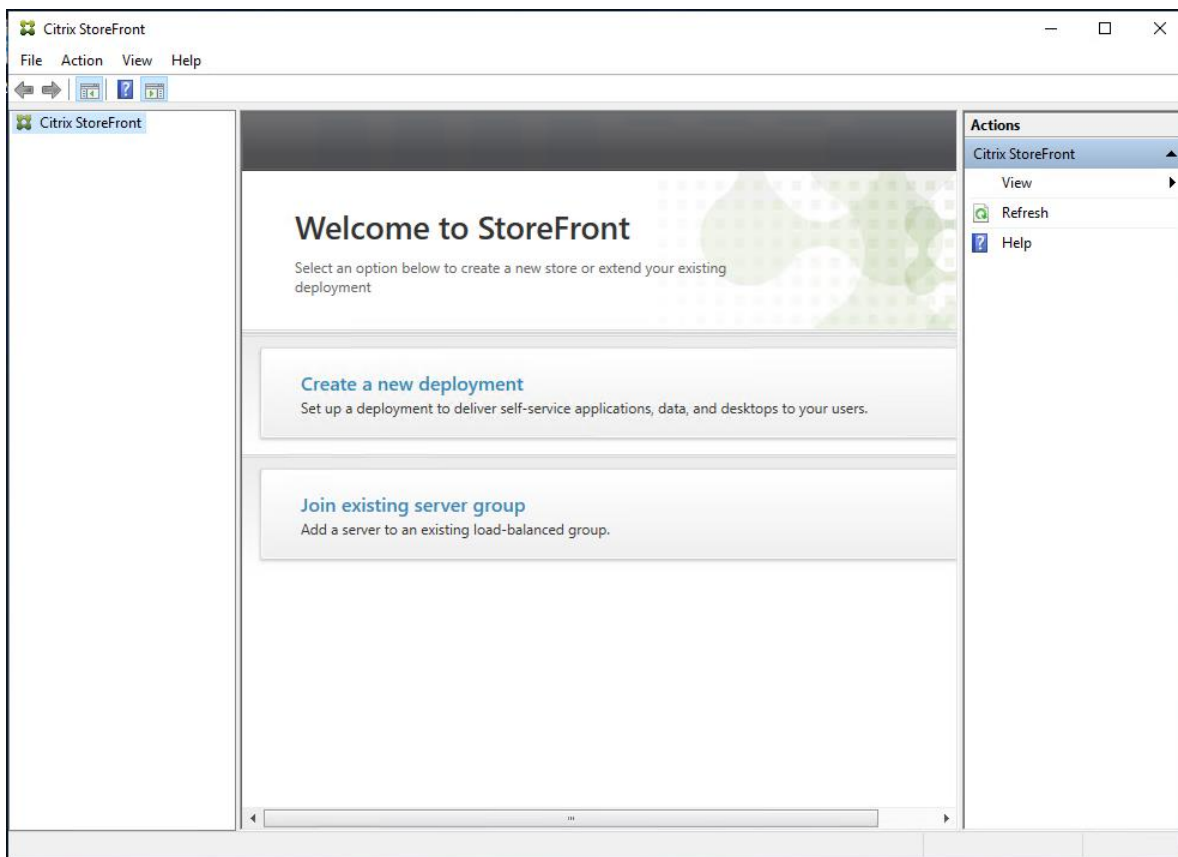
1. Install the second StoreFront using the same installation steps outlined above.
2. Connect to the first StoreFront server
3. To add the second server and generate the authorization information that allows the additional StoreFront server to join the server group, select Add Server from Actions pane in the Server Group.



4. Copy the authorization code.

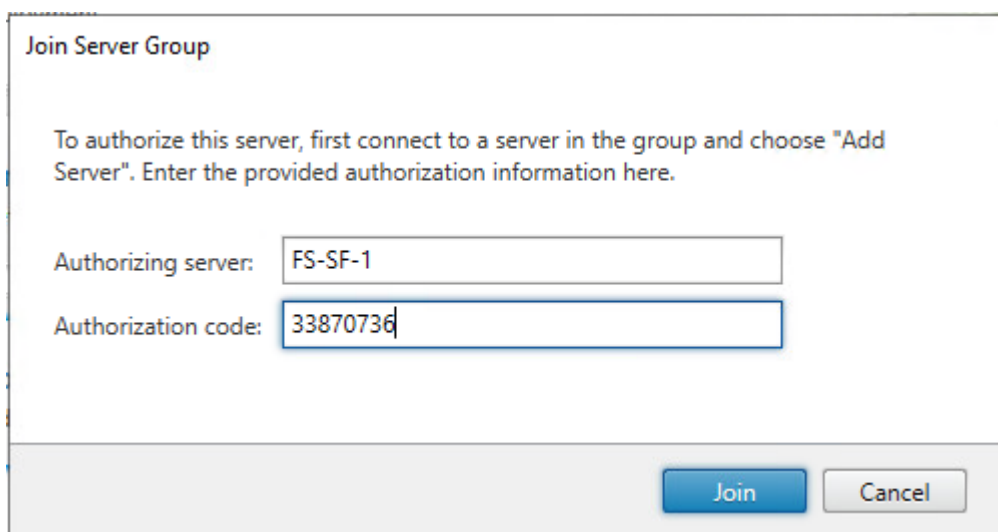


5. From the StoreFront Console on the second server select "Join existing server group."



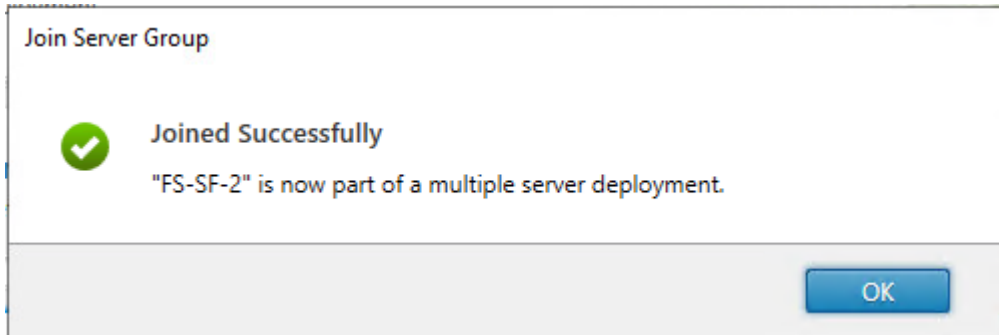
6. In the Join Server Group dialog, enter the name of the first Storefront server and paste the Authorization code into the Join Server Group dialog.

7. Click Join.

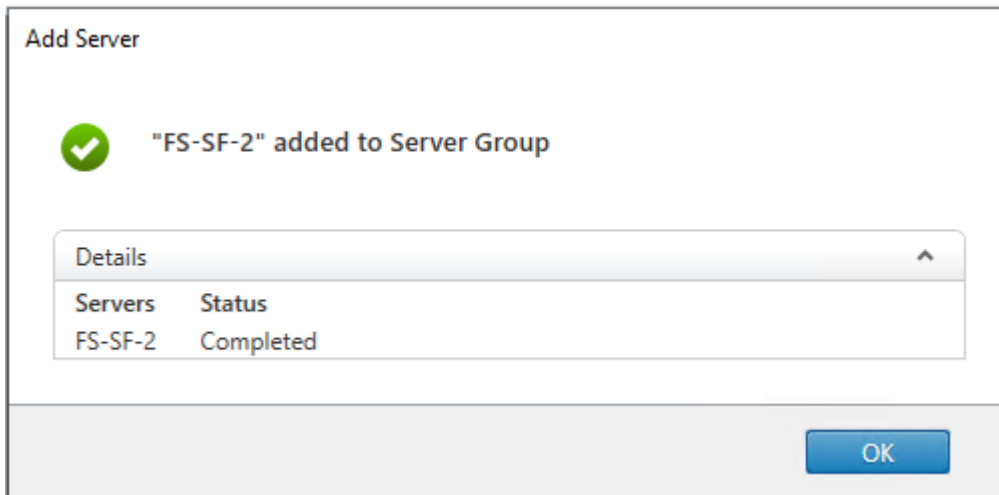


8. A message appears when the second server has joined successfully.

9. Click OK.



The second StoreFront is now in the Server Group.



Install and Configure Citrix Provisioning Server 2109

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 2109](#) document.

Prerequisites

Set the following Scope Options on the DHCP server hosting the PVS target machines:

Option Name	Vendor	Value	Policy Name
003 Router	Standard	10.72.0.1	None
006 DNS Servers	Standard	10.10.71.11	None
011 Resource Location Servers	Standard	10.72.0.10, 10.72.0.11, 10.72.0.12	None
015 DNS Domain Name	Standard	FSL151K.LOCAL	None
066 Boot Server Host Name	Standard	pvs-lb	None
067 Bootfile Name	Standard	pvsnbpx64.efi	None

Create a DNS host records with multiple PVS Servers IP for TFTP Load Balancing:

Name	Type	Data	Timestamp
W2019-MCS-Base	Host (A)	10.72.9.2	12/21/2021 12:00:00 PM
W19-MCSIMG-0105	Host (A)	10.72.9.18	1/6/2022 9:00:00 AM
pvs-lb	Host (A)	10.72.0.10	static
pvs-lb	Host (A)	10.72.0.12	static
pvs-lb	Host (A)	10.72.0.11	static
purefile	Host (A)	10.10.71.50	static
MCS-W2019-128	Host (A)	10.72.9.75	1/10/2022 10:00:00 AM

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"

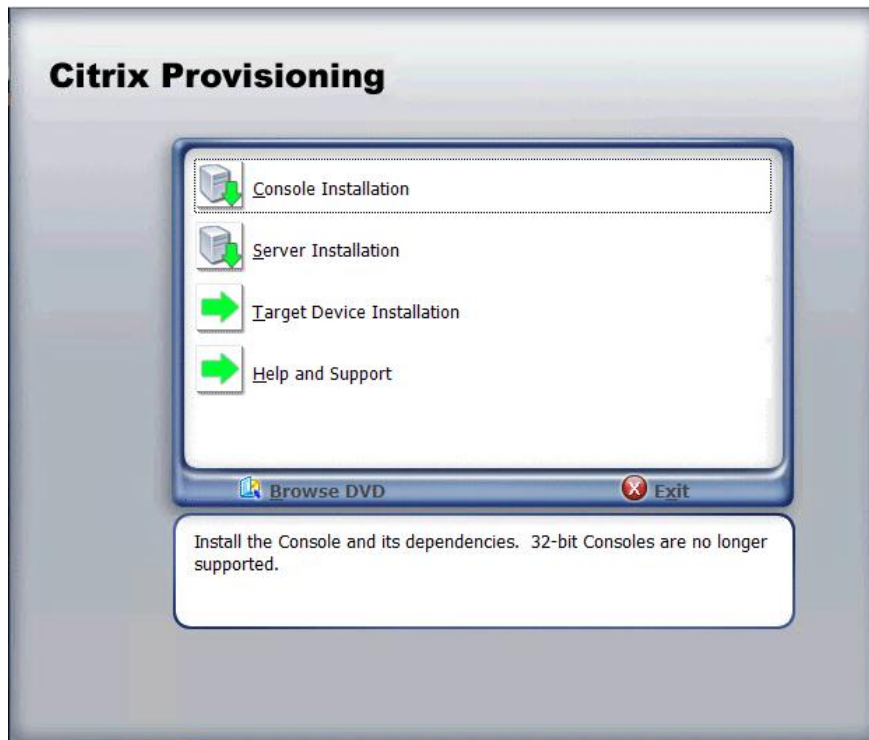
Only one MS SQL database is associated with a farm. You can choose to install the Provisioning Services database software on an existing SQL database, if that machine can communicate with all Provisioning Servers within the farm, or with a new SQL Express database machine, created using the SQL Express software that is free from Microsoft.

The following databases are supported: Microsoft SQL Server 2008 SP3 through 2016 (x86, x64, and Express editions). Please check Citrix documentation for further reference.

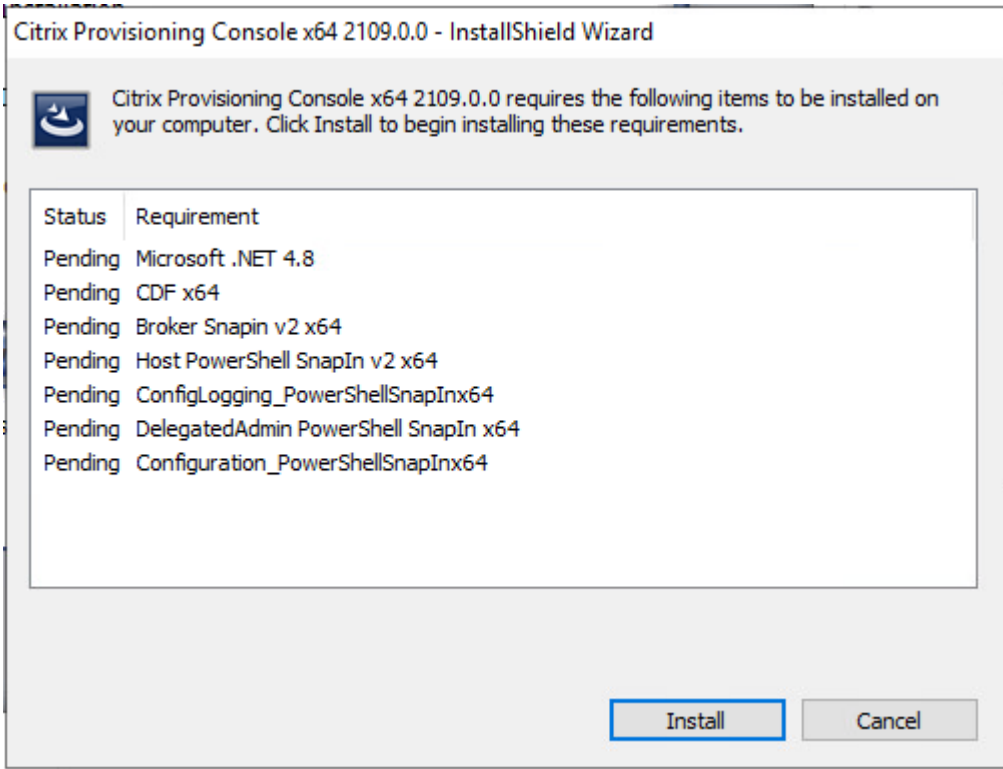
Note: Microsoft SQL 2019 was installed separately for this CVD.

To install and configure Citrix Provisioning Service 2109, follow these steps:

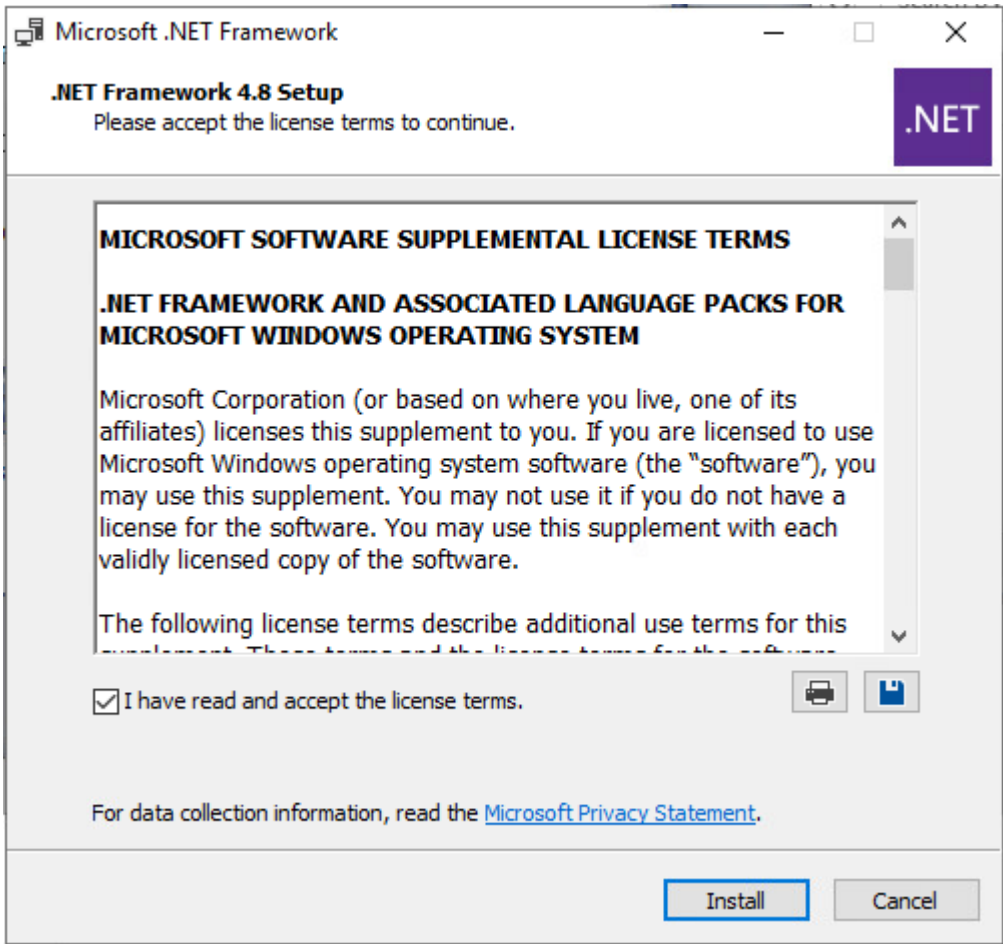
1. Connect to Citrix Provisioning server and launch Citrix Provisioning Services 2109 ISO and let AutoRun launch the installer.
2. Click Console Installation.



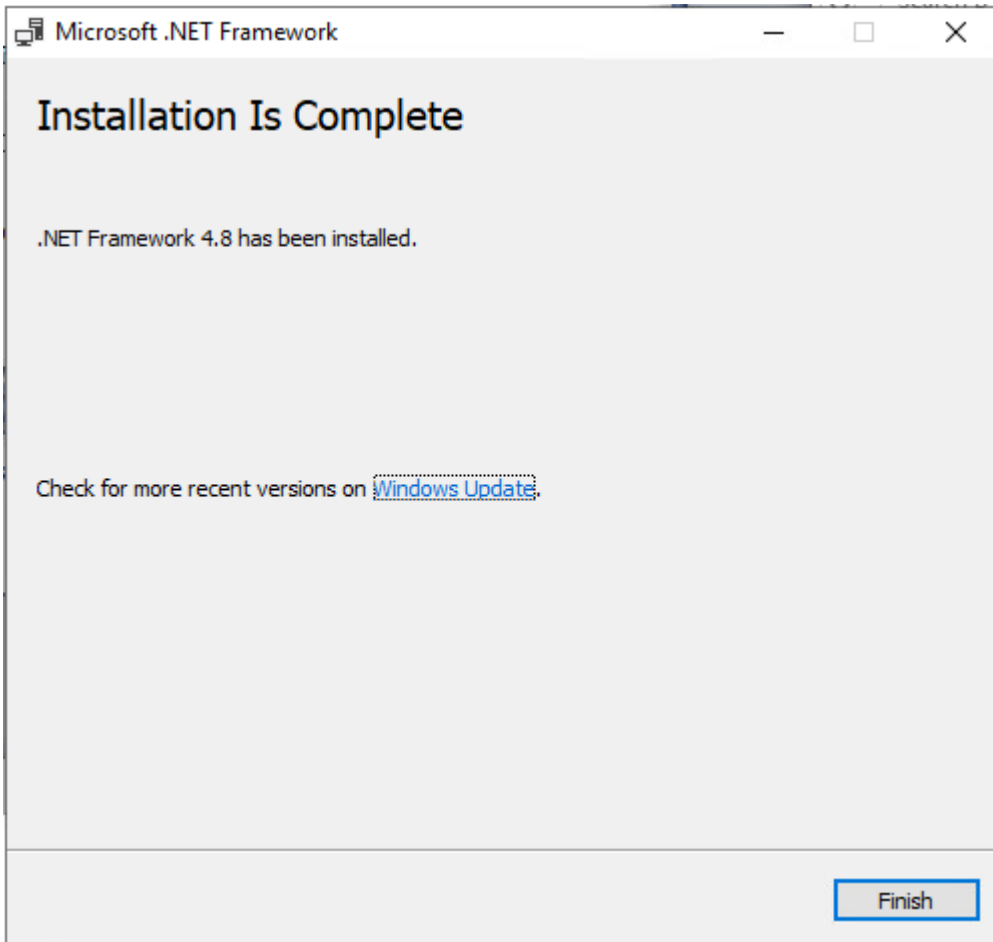
3. Click Install to start the console installation.



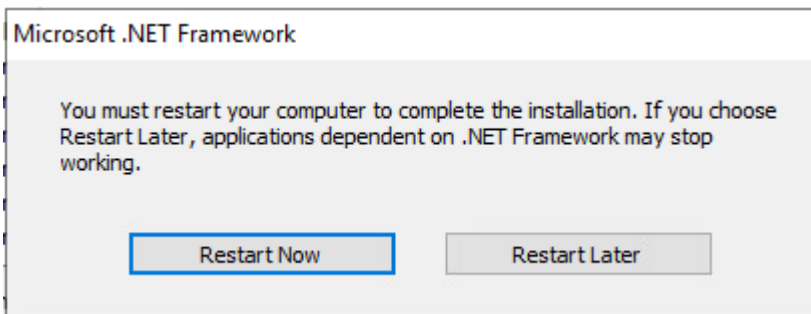
4. Read the .NET License Agreement. If acceptable, check “I have read and accept the license terms.”
5. Click Next.



6. Click Finish.

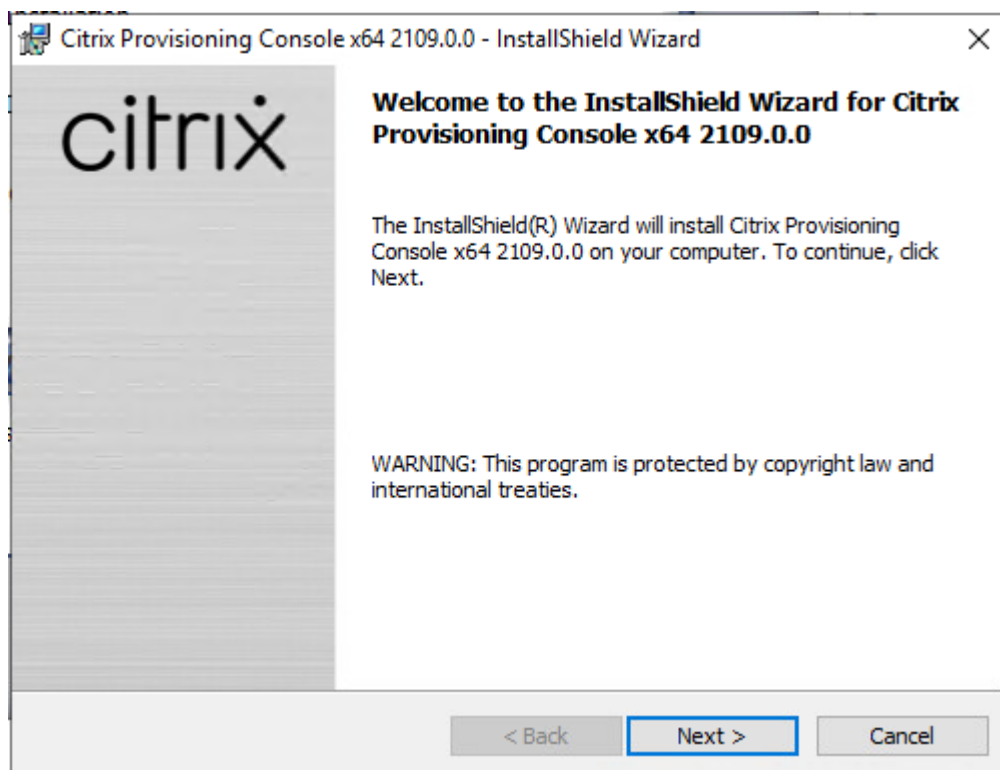


7. Restart the Virtual Machine.



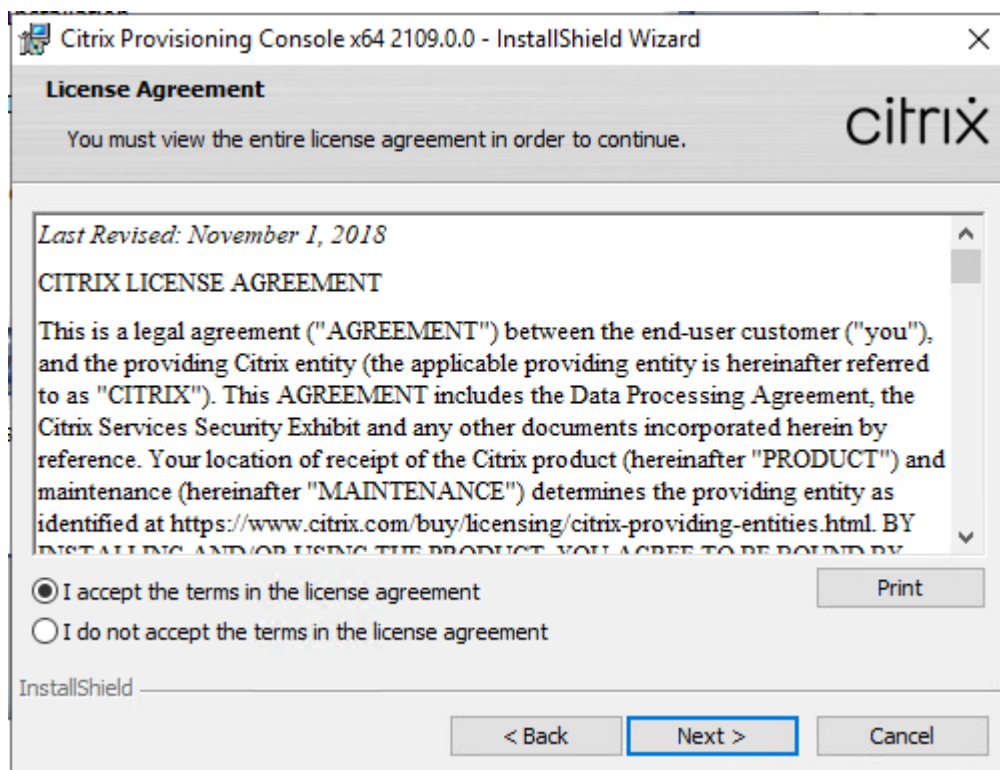
8. Logging into the Operating system automatically launches the installation wizard.

9. Click Next.



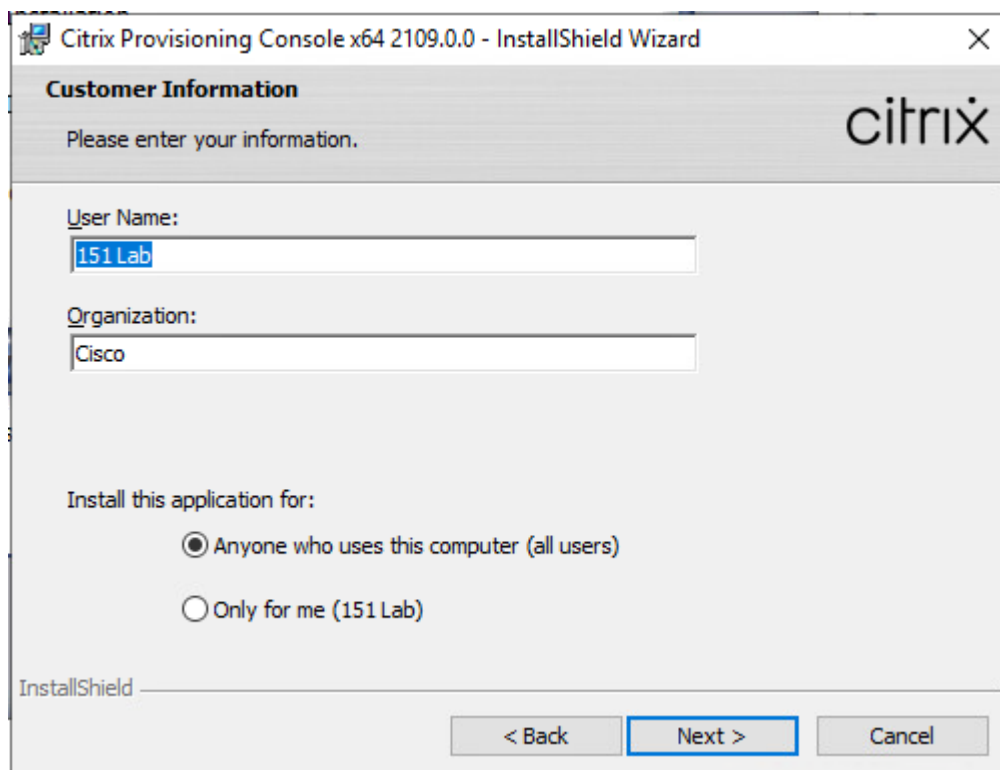
10. Read the Citrix License Agreement. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”

11. Click Next.

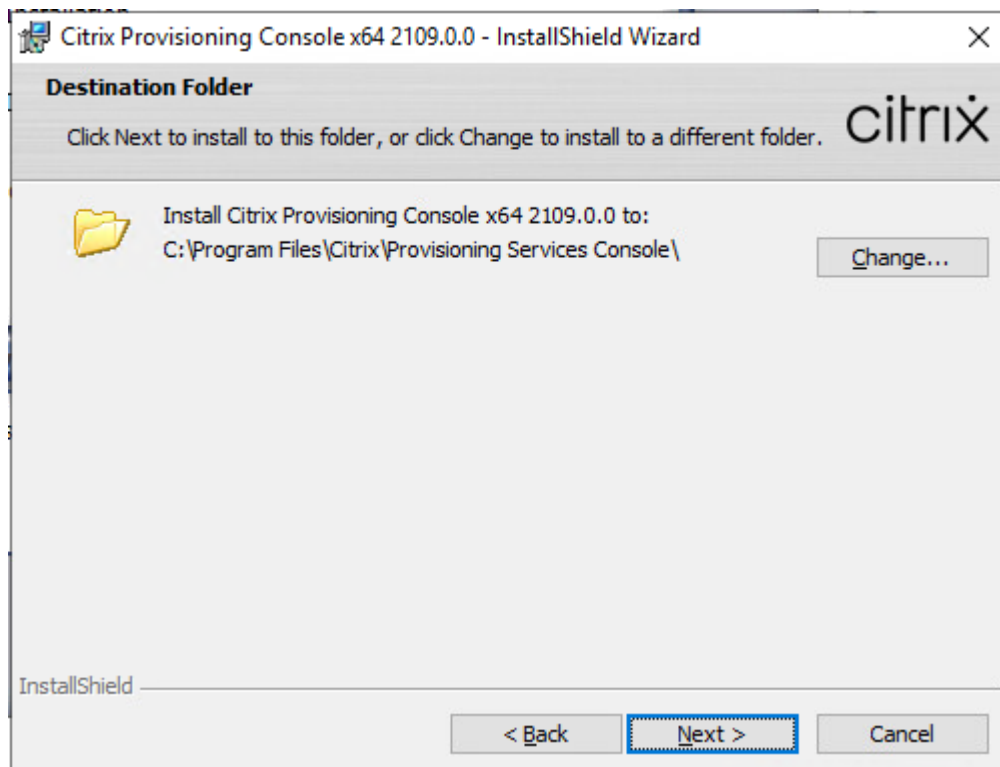


12. Optionally, provide User Name and Organization.


13. Click Next.



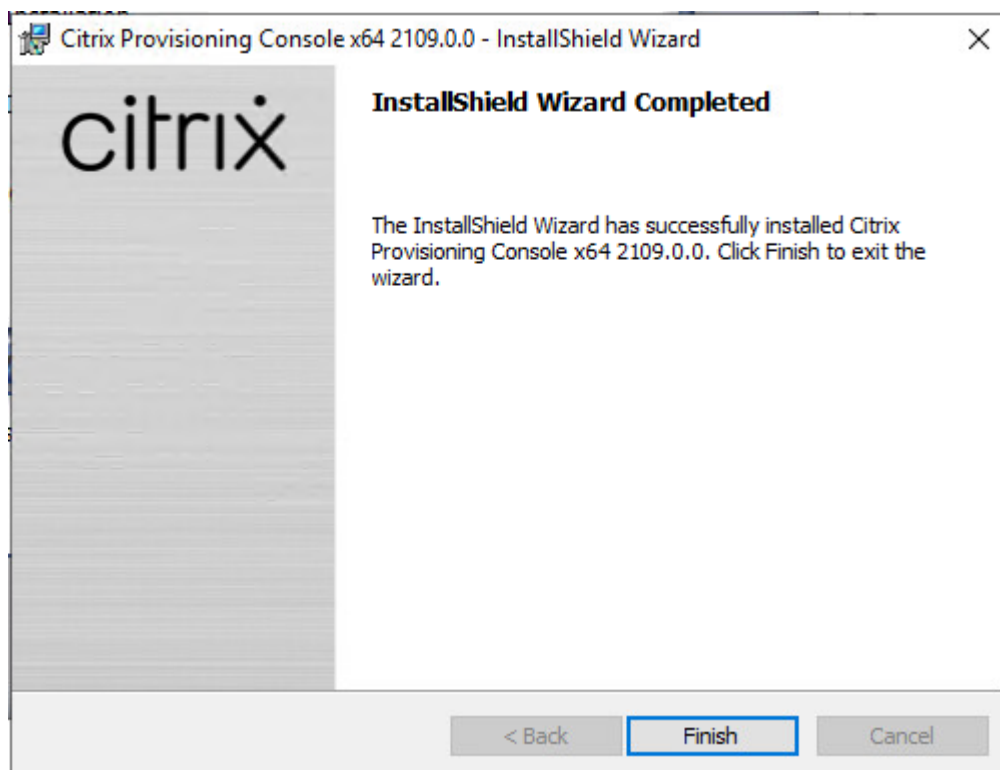
14. Accept the default path.



15. Click Install.




16. Click Finish after successful installation.




17. From the main installation screen, select Server Installation.


Citrix Provisioning

 Console Installation

 Server Installation

 Target Device Installation


 Help and Support

 Browse DVD

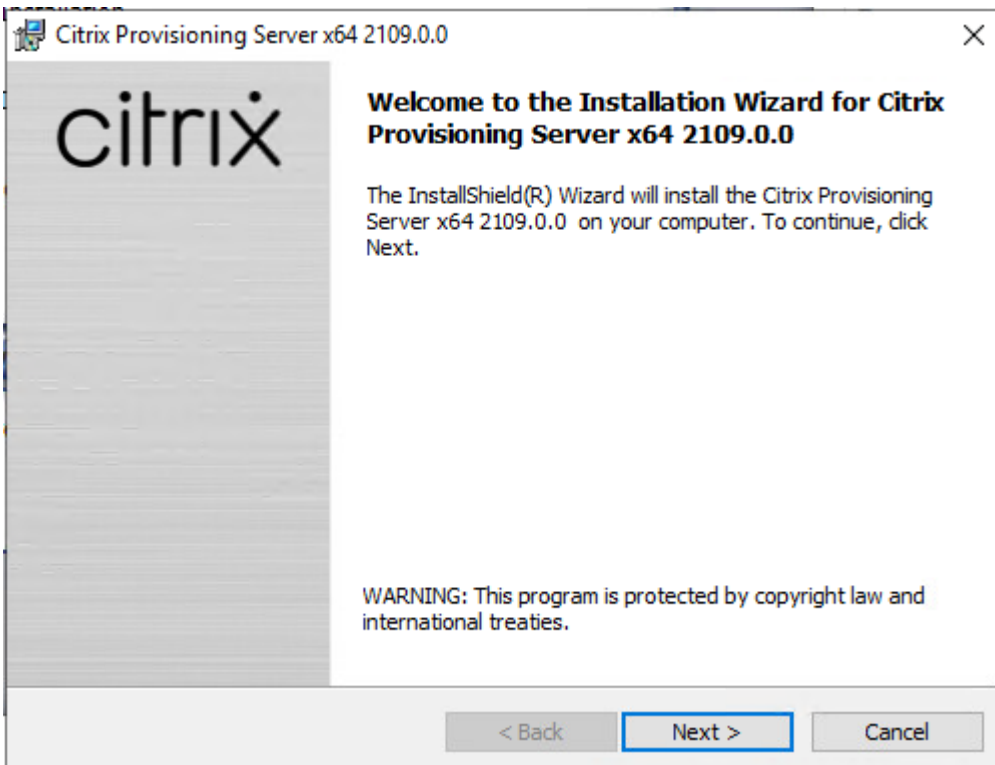
 Exit

Install the Server and its dependencies.

18. Click Install on the prerequisites dialog.

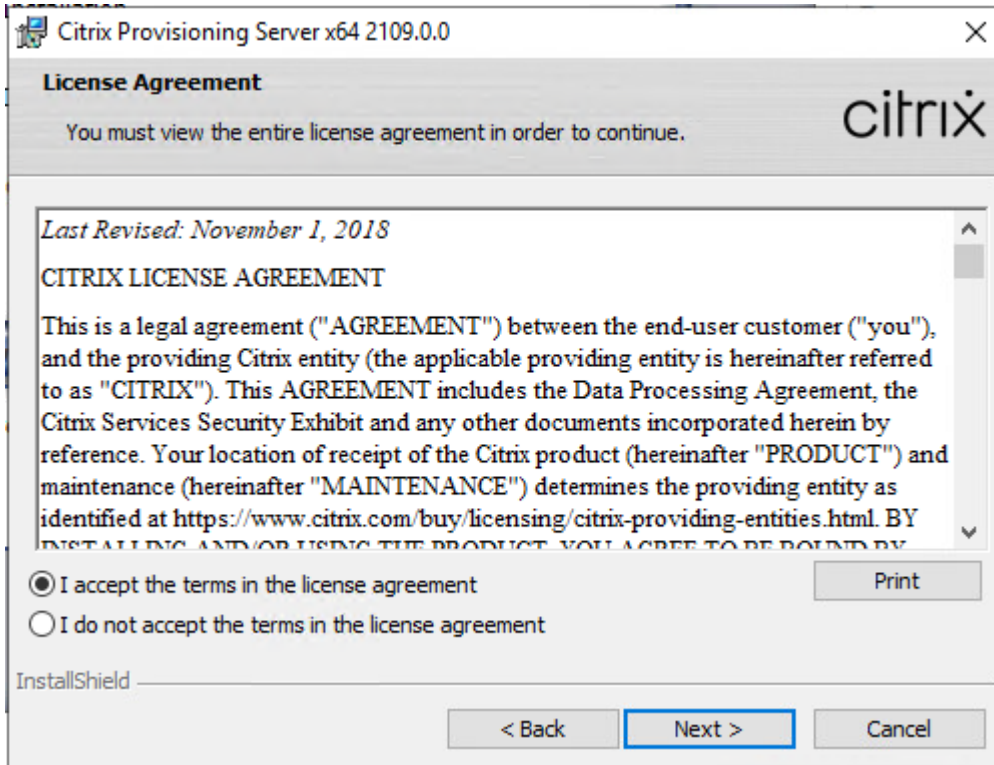


19. Click Next when the Installation wizard starts.

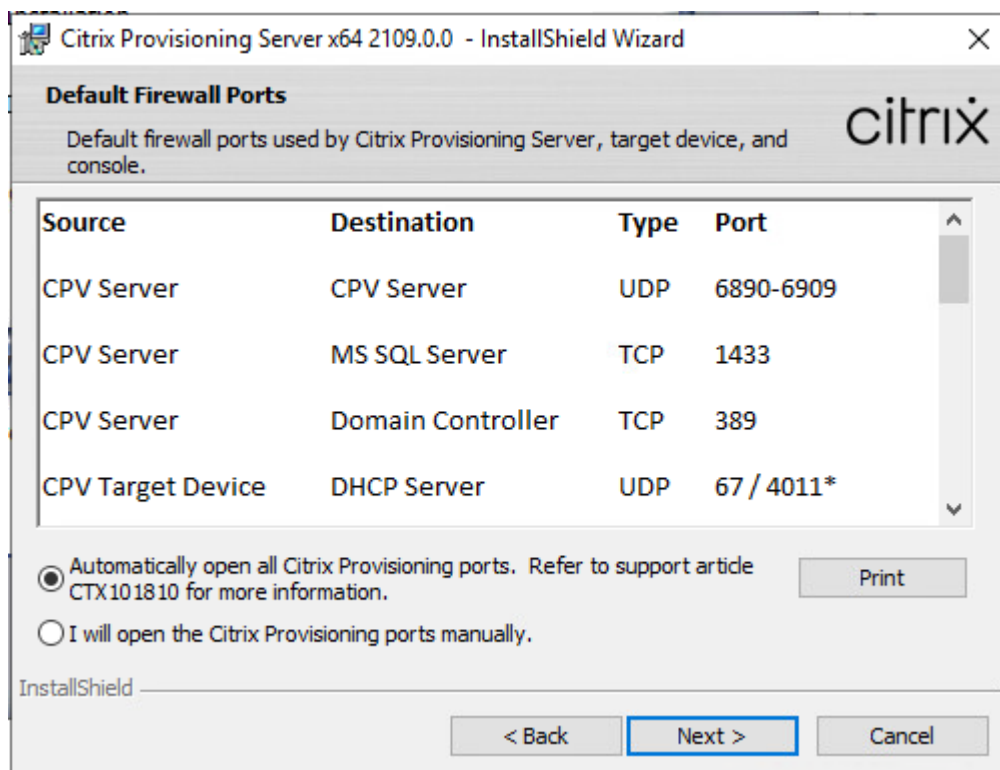


20. Review the license agreement terms. If acceptable, select the radio button labeled “I accept the terms in the license agreement.”

21. Click Next.

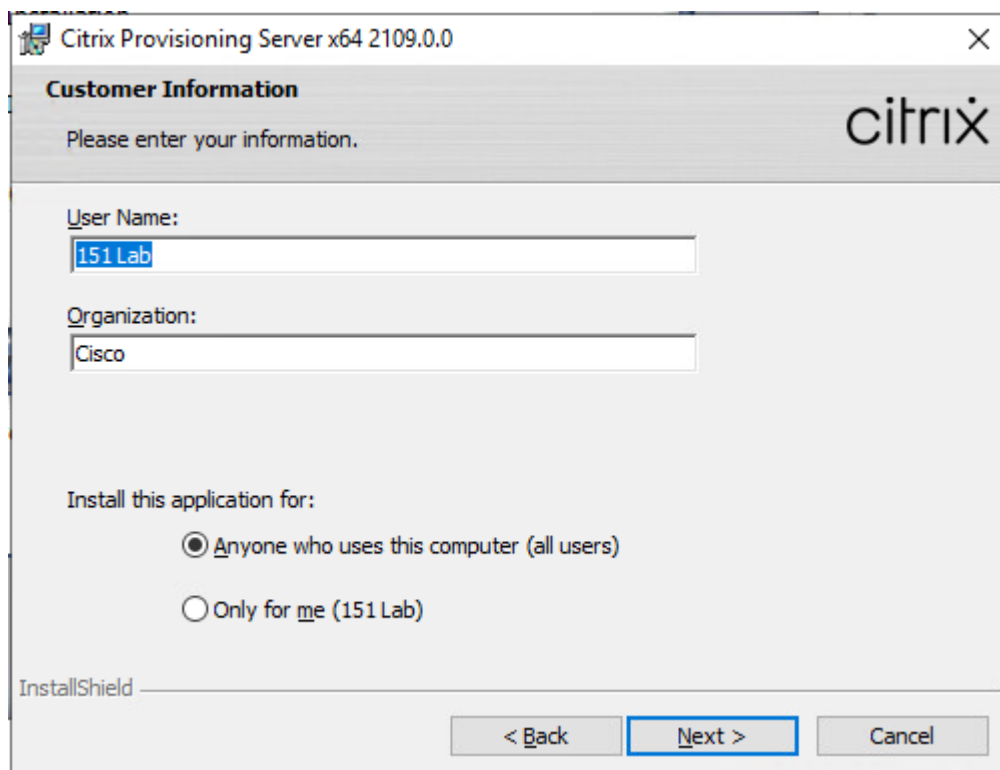


22. Select Automatically open Citrix PVS Firewall Ports.



23. Provide User Name and Organization information. Select who will see the application.

24. Click Next.



Citrix Provisioning Server x64 2109.0.0

Customer Information

Please enter your information.

User Name:
151 Lab

Organization:
Cisco

Install this application for:

Anyone who uses this computer (all users)

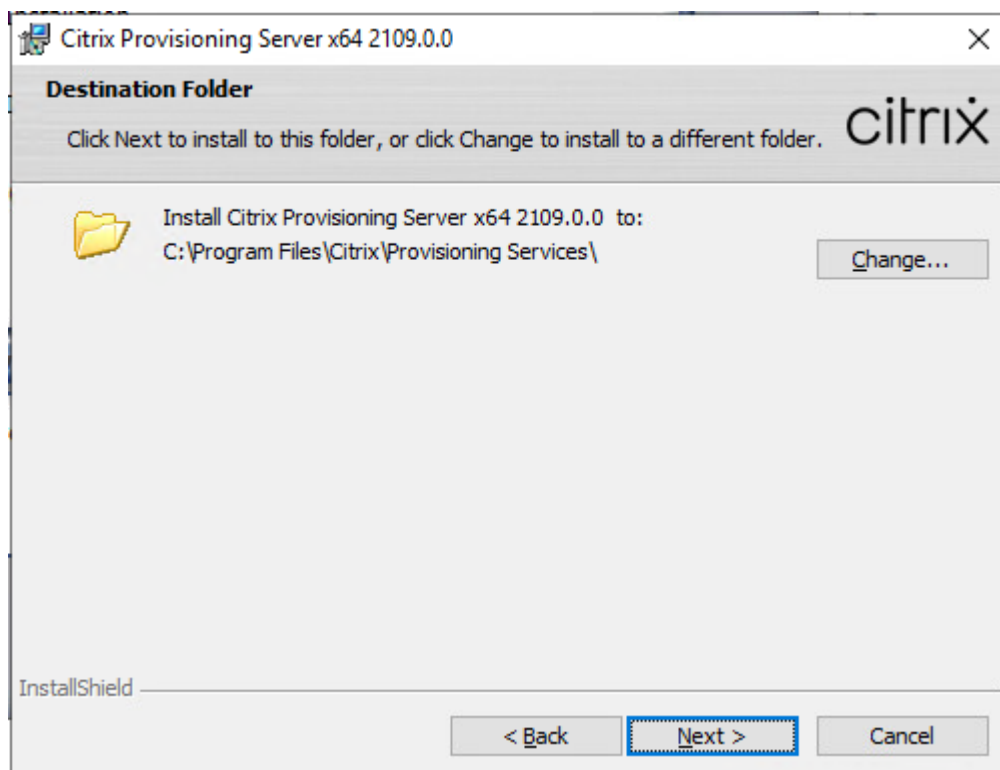
Only for me (151 Lab)

InstallShield


< Back Next > Cancel

25. Accept the default installation location.

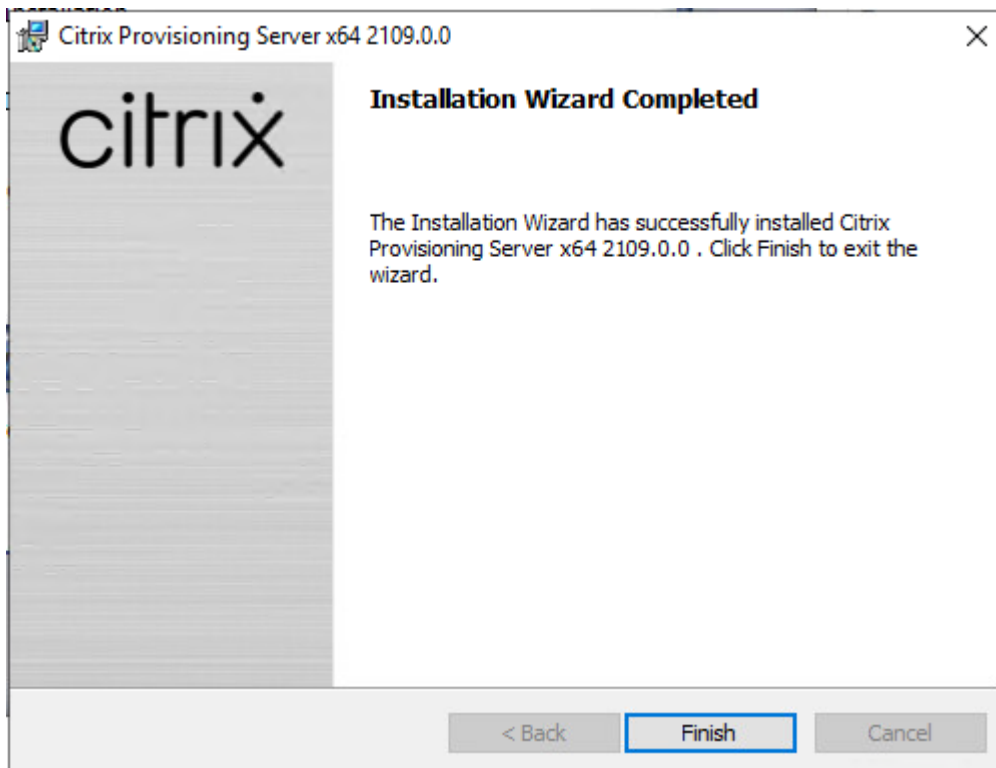
26. Click Next.



27. Click Install to begin the installation.

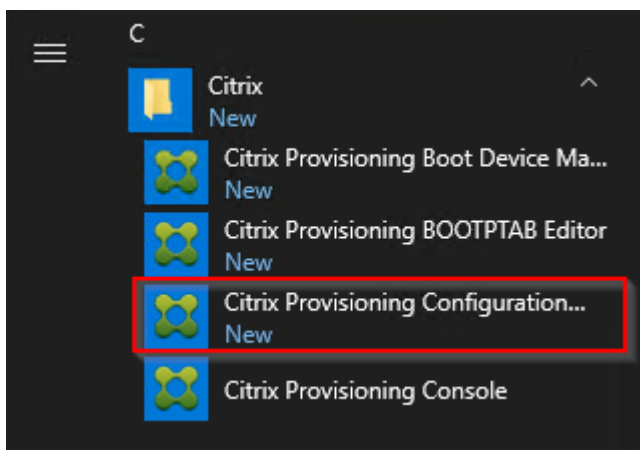


28. Click Finish when the install is complete.

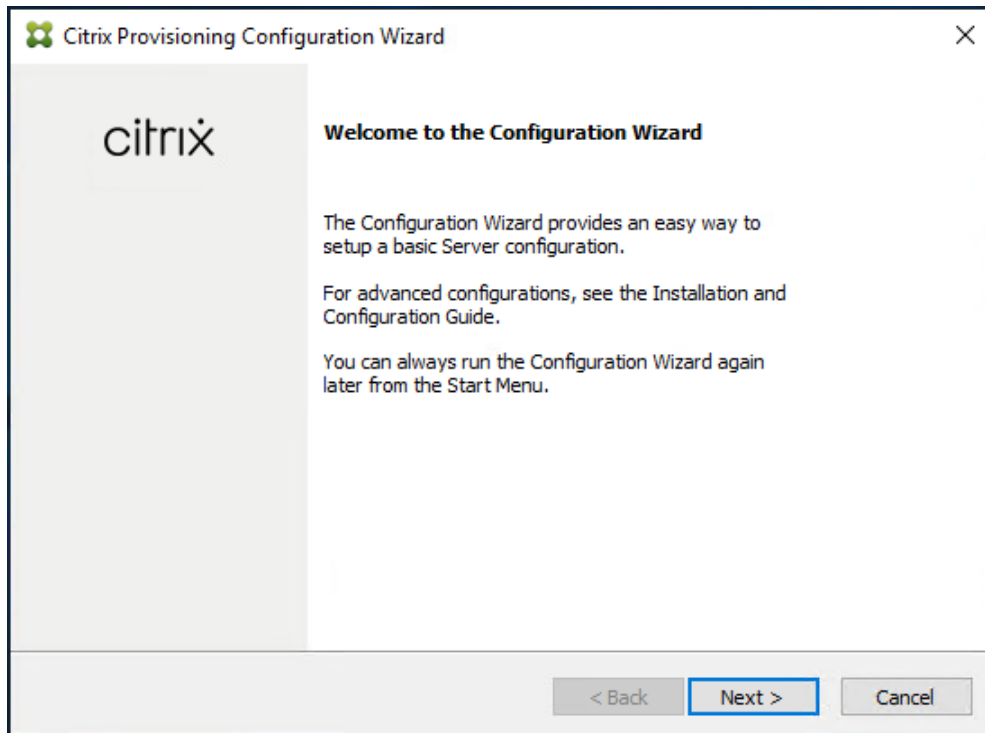


To configure Citrix Provisioning services, follow these steps:

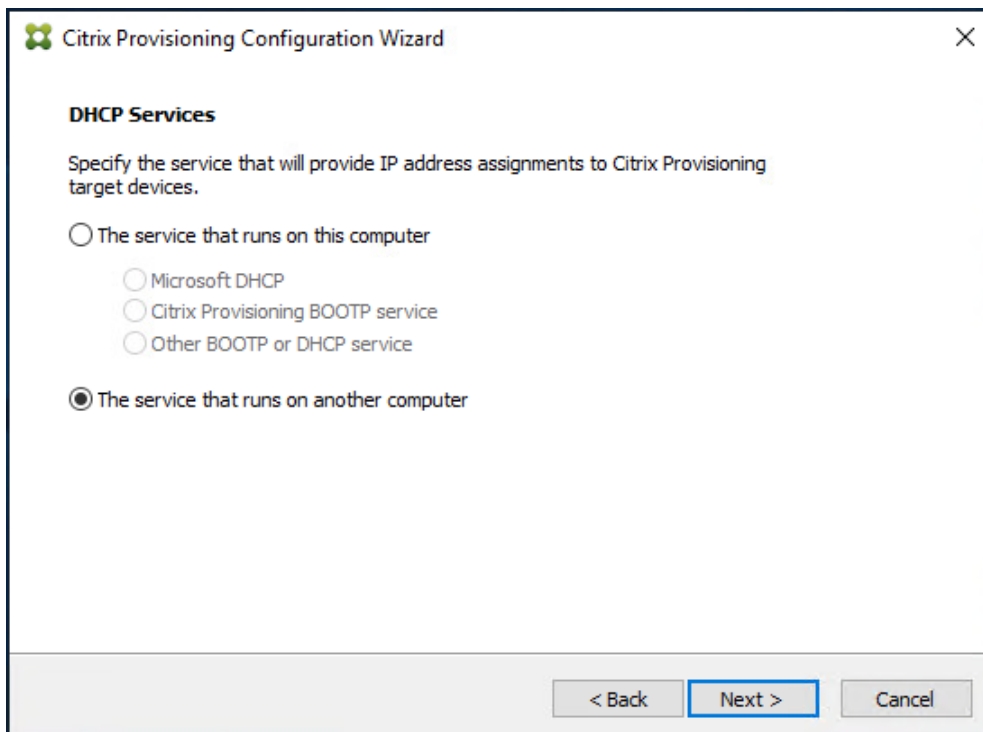
1. Start PVS Configuration Wizard.



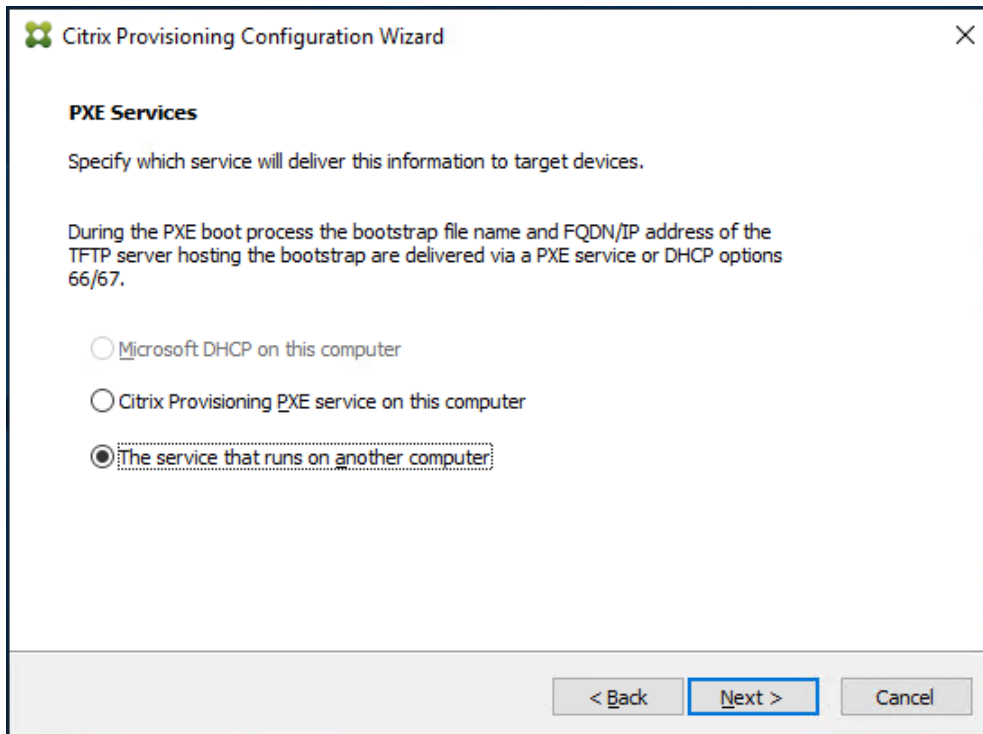
2. Click Next.



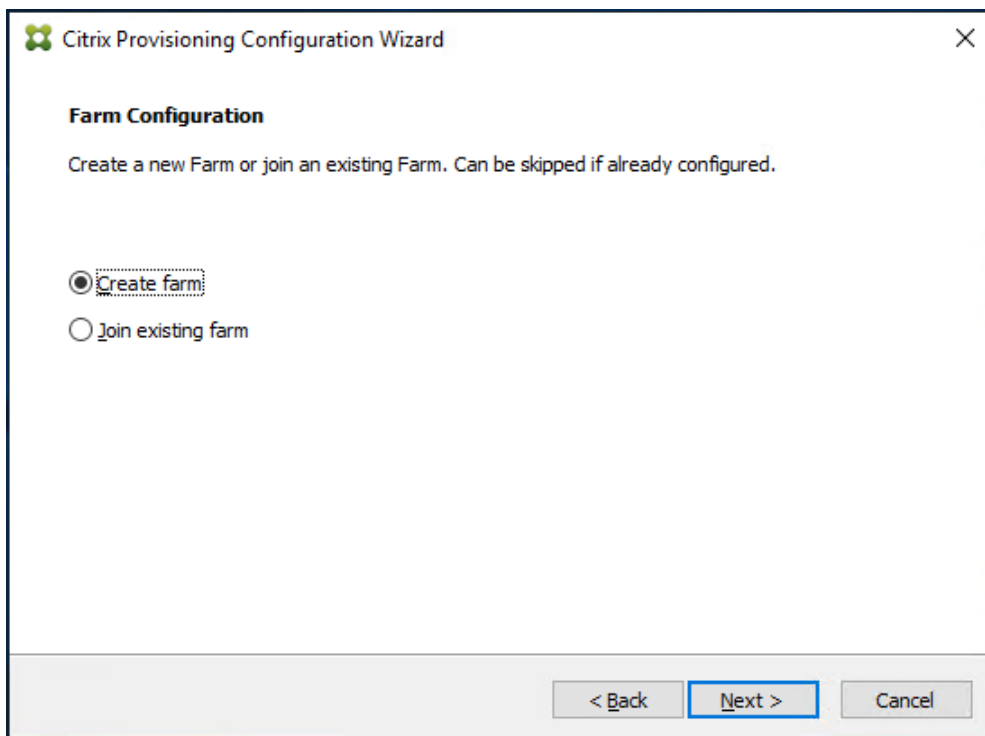
3. Since the PVS server is not the DHCP server for the environment, select the radio button labeled, "The service that runs on another computer."
4. Click Next.



5. Since DHCP boot options are used for TFTP services, select the radio button labeled, “The service that runs on another computer.”
6. Click Next.



7. Since this is the first server in the farm, select the radio button labeled, “Create farm.”
8. Click Next.



Citrix Provisioning Configuration Wizard

Farm Configuration

Create a new Farm or join an existing Farm. Can be skipped if already configured.

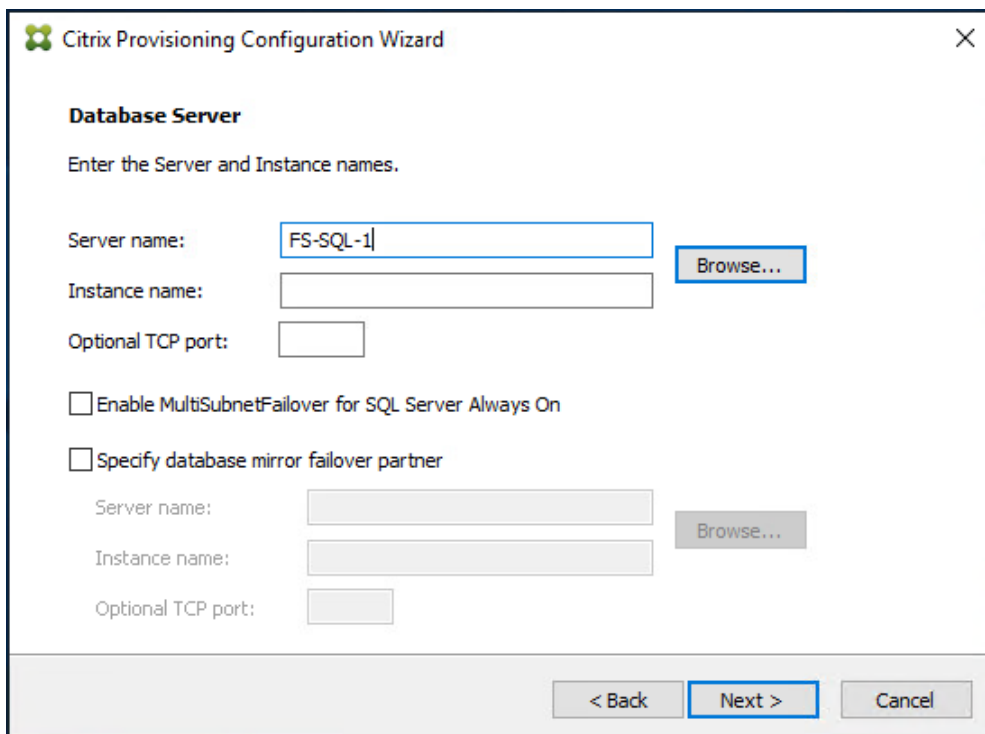
Create farm

Join existing farm

< Back Next > Cancel

9. Enter the FQDN of the SQL server.

10. Click Next.



Citrix Provisioning Configuration Wizard

Database Server

Enter the Server and Instance names.

Server name: FS-SQL-1 Browse...

Instance name:

Optional TCP port:

Enable MultiSubnetFailover for SQL Server Always On

Specify database mirror failover partner

Server name: Browse...

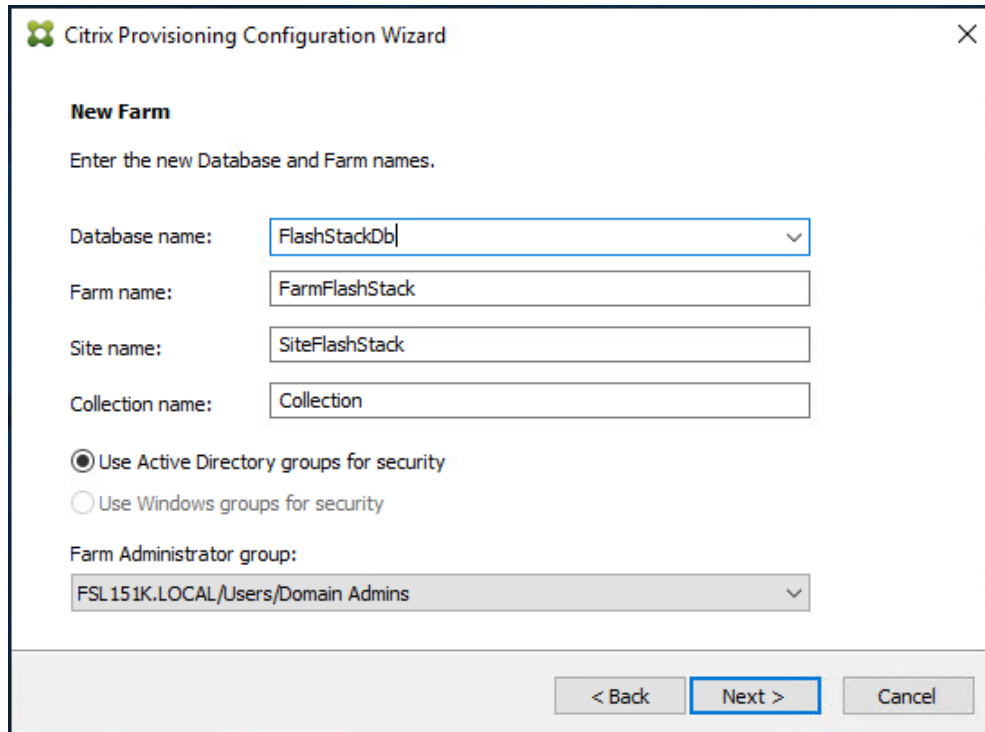
Instance name:

Optional TCP port:

< Back Next > Cancel

11. Provide the Database, Farm, Site, and Collection name.

12. Click Next.



The screenshot shows the 'Citrix Provisioning Configuration Wizard' window, specifically the 'New Farm' step. The window title is 'Citrix Provisioning Configuration Wizard' with a close button (X) in the top right corner. Below the title bar, the text 'New Farm' is displayed. A subtitle reads 'Enter the new Database and Farm names.' There are four input fields: 'Database name:' with a dropdown menu showing 'FlashStackDb'; 'Farm name:' with a text box containing 'FarmFlashStack'; 'Site name:' with a text box containing 'SiteFlashStack'; and 'Collection name:' with a text box containing 'Collection'. Below these fields are two radio button options: 'Use Active Directory groups for security' (which is selected) and 'Use Windows groups for security'. Underneath is a 'Farm Administrator group:' dropdown menu showing 'FSL151K.LOCAL/Users/Domain Admins'. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

13. Provide the vDisk Store details.

14. Click Next.

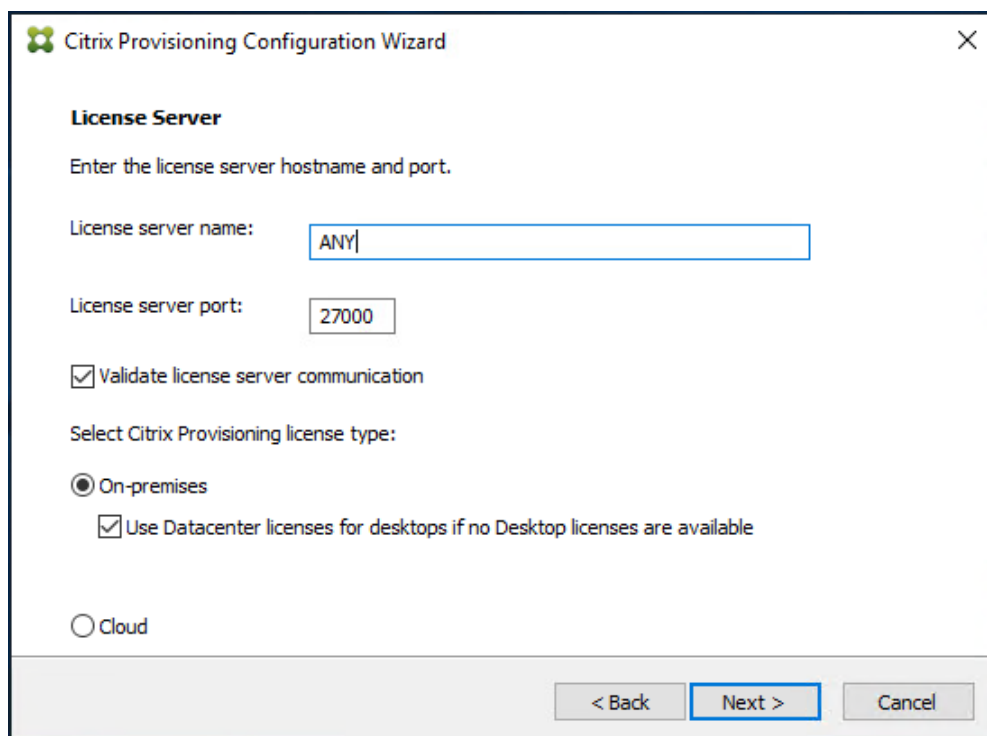
The image shows a screenshot of the 'Citrix Provisioning Configuration Wizard' dialog box, specifically the 'New Store' step. The window title is 'Citrix Provisioning Configuration Wizard' with a close button (X) in the top right corner. Below the title bar, the text 'New Store' is displayed in bold. Underneath, it says 'Enter a new Store and default path.' There are two input fields: 'Store name:' with the text 'Store' entered, and 'Default path:' with the text 'E:\Store' entered. To the right of the 'Default path' field is a 'Browse...' button. At the bottom of the dialog box, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Note: For large scale PVS environment, it is recommended to create the share using support for CIFS/SMB3 on an enterprise ready File Server.

15. Provide the FQDN of the license server.

16. Optionally, provide a port number if changed on the license server.

17. Click Next.



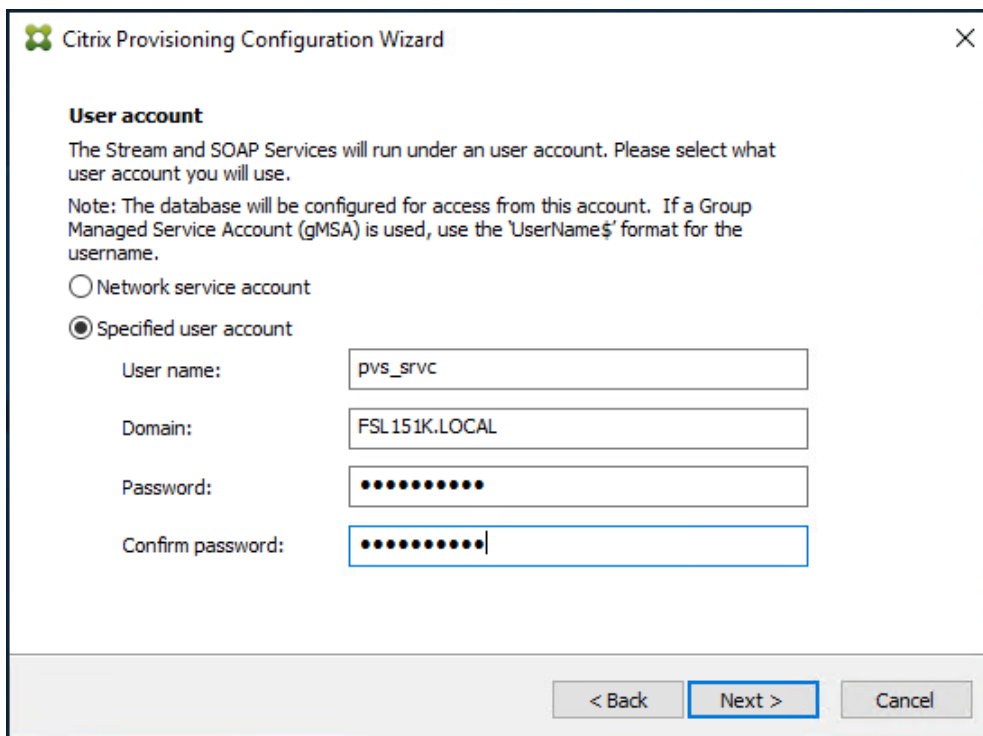
The image shows a screenshot of the 'Citrix Provisioning Configuration Wizard' dialog box, specifically the 'License Server' step. The window title is 'Citrix Provisioning Configuration Wizard' with a close button (X) in the top right corner. Below the title bar, the section is titled 'License Server'. A subtitle reads 'Enter the license server hostname and port.' There are two text input fields: 'License server name:' containing the text 'ANY' and 'License server port:' containing the text '27000'. Below these fields is a checked checkbox labeled 'Validate license server communication'. Underneath is the text 'Select Citrix Provisioning license type:'. There are two radio button options: 'On-premises' (which is selected) and 'Cloud'. Under the 'On-premises' option, there is a checked checkbox labeled 'Use Datacenter licenses for desktops if no Desktop licenses are available'. At the bottom of the dialog, there are three buttons: '< Back', 'Next >' (which is highlighted with a blue border), and 'Cancel'.

18. If an Active Directory service account is not already setup for the PVS servers, create that account prior to clicking Next on this dialog.

19. Select the Specified user account radio button.

20. Complete the User name, Domain, Password, and Confirm password fields, using the PVS account information created earlier.

21. 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:

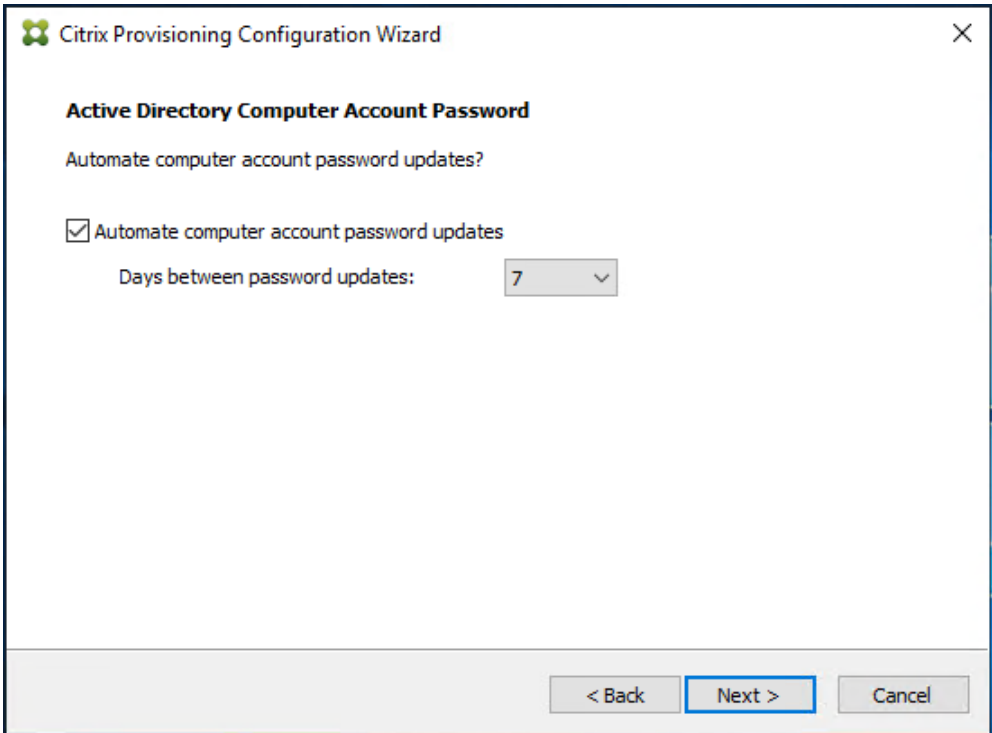
Confirm password:

< Back Next > Cancel

22. Set the Days between password updates to 7.

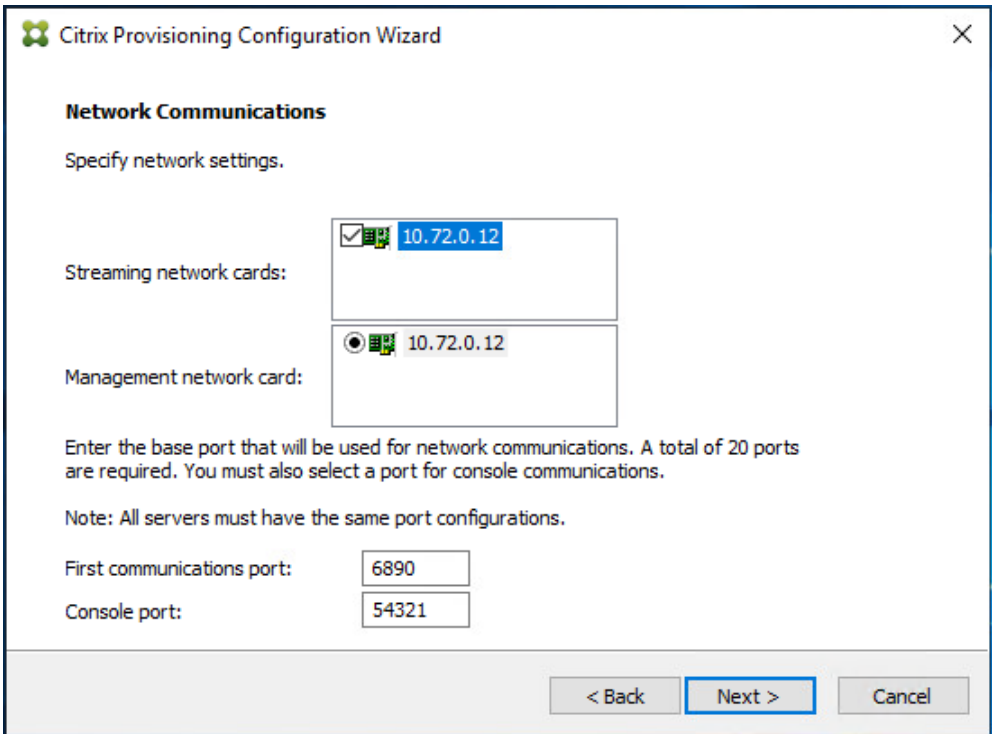
Note: This will vary per environment. “7 days” for the configuration was appropriate for testing purposes.

23. Click Next.



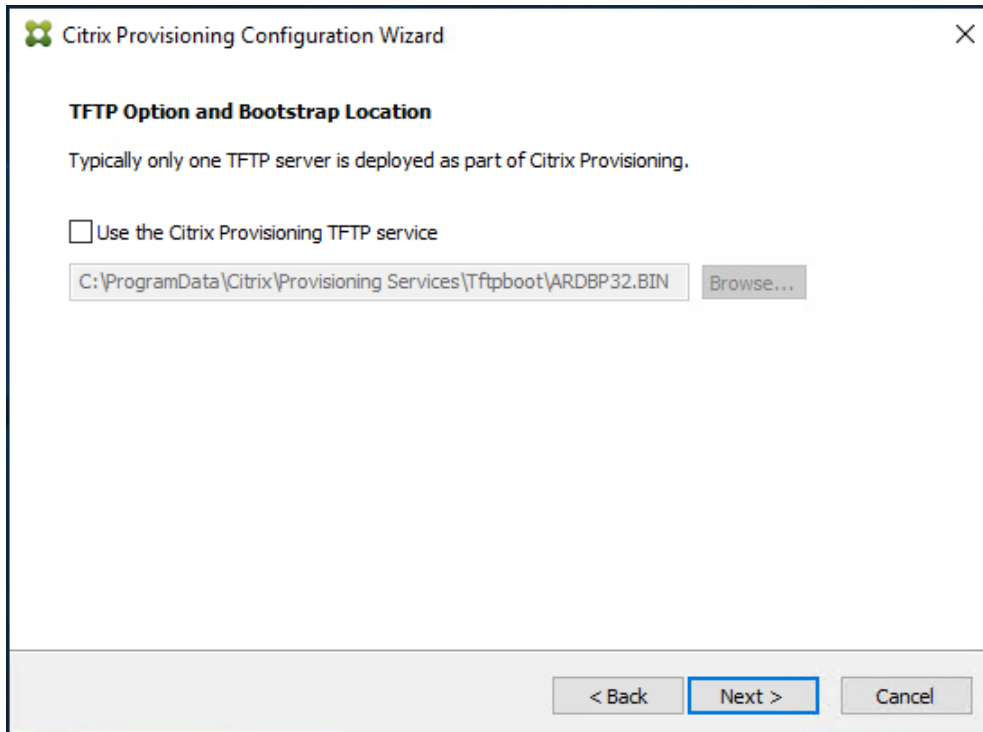
24. Keep the defaults for the network cards.

25. Click Next.



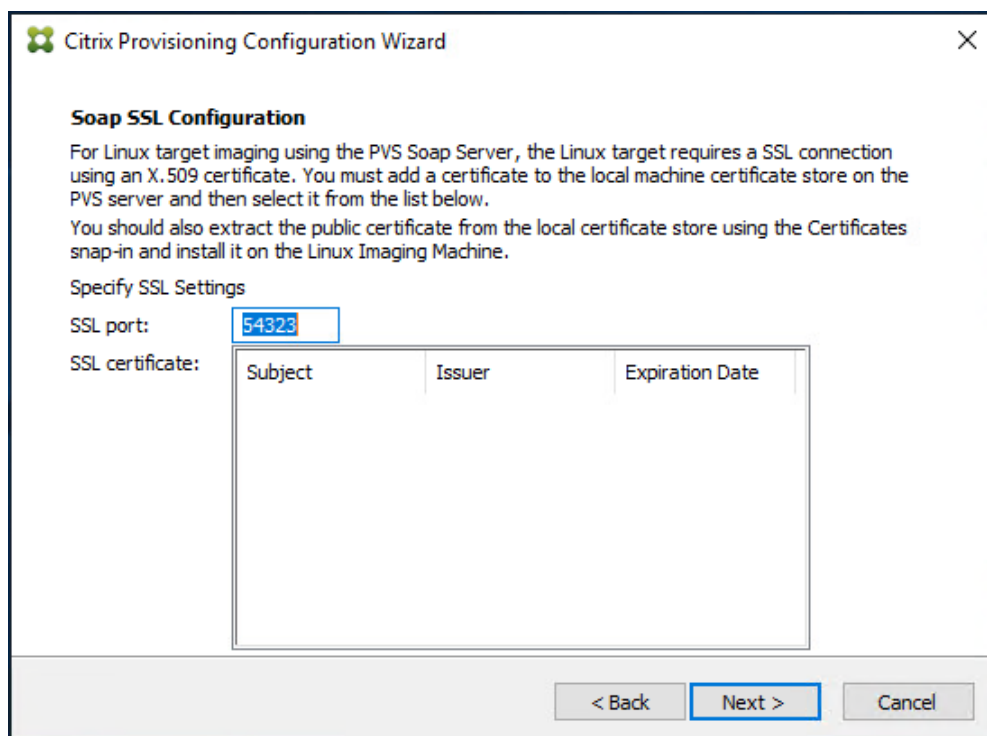
26. Select Use the Provisioning Services TFTP service checkbox.

27. Click Next.



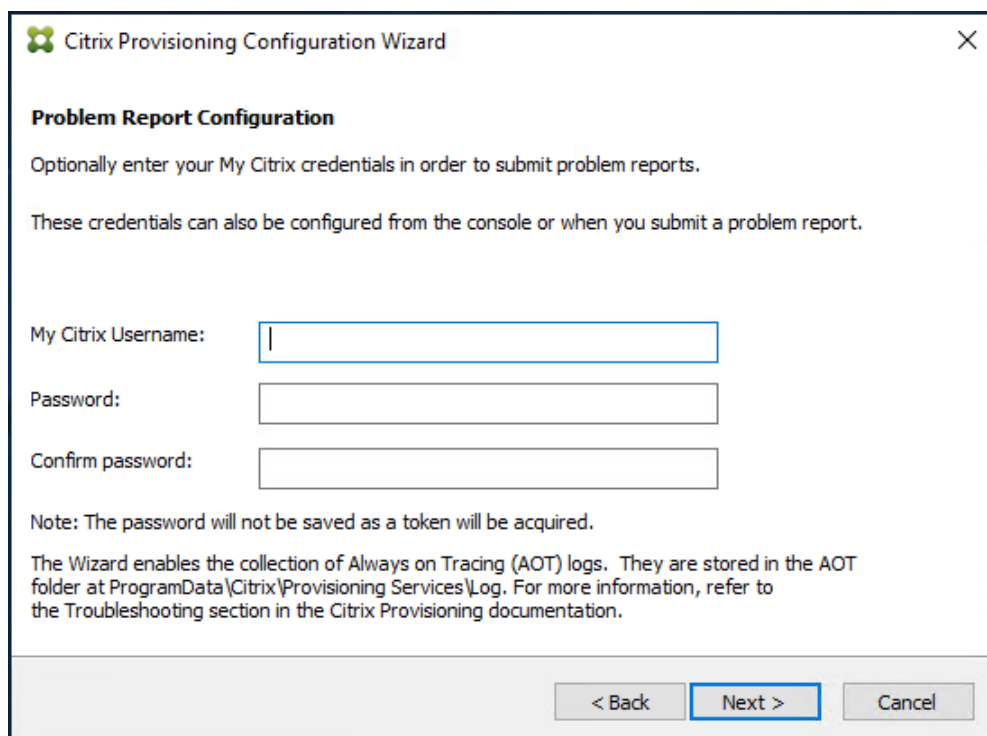
28. If Soap Server is used, provide details.

29. Click Next.

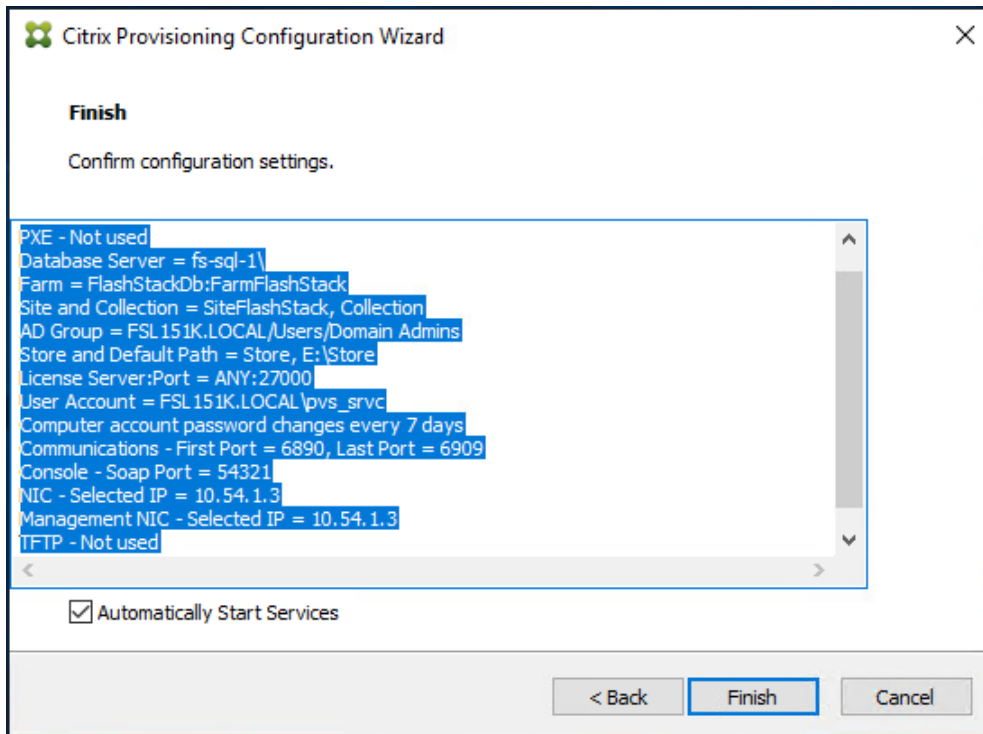


30. If desired fill in Problem Report Configuration.

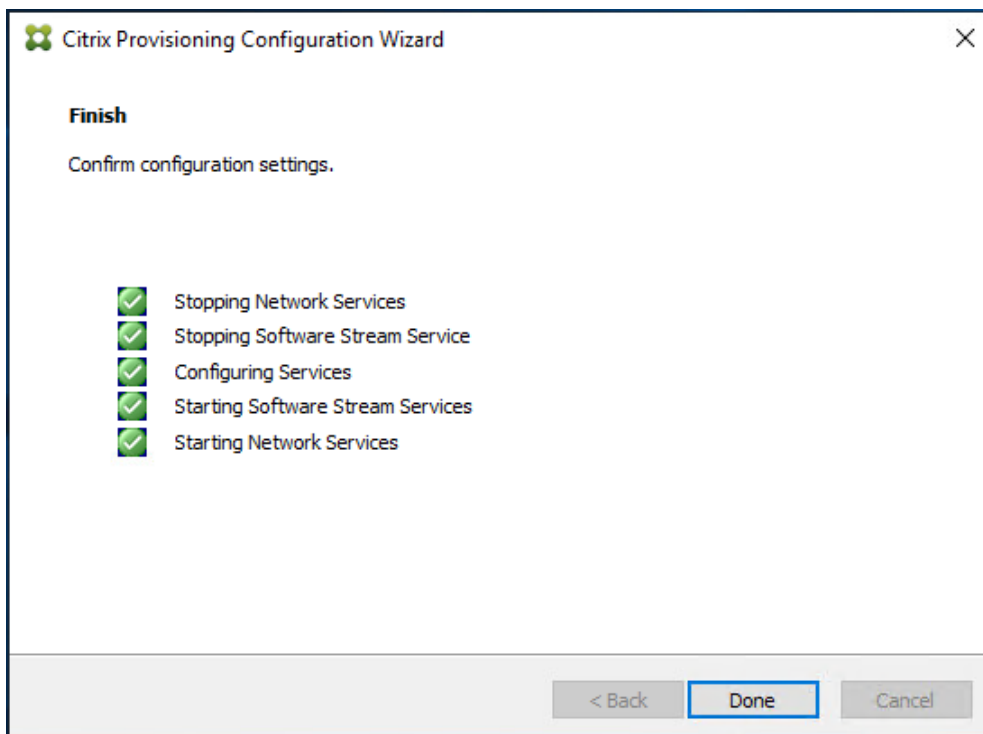
31. Click Next.



32. Click Finish to start the installation.



33. When the installation is completed, click Done.

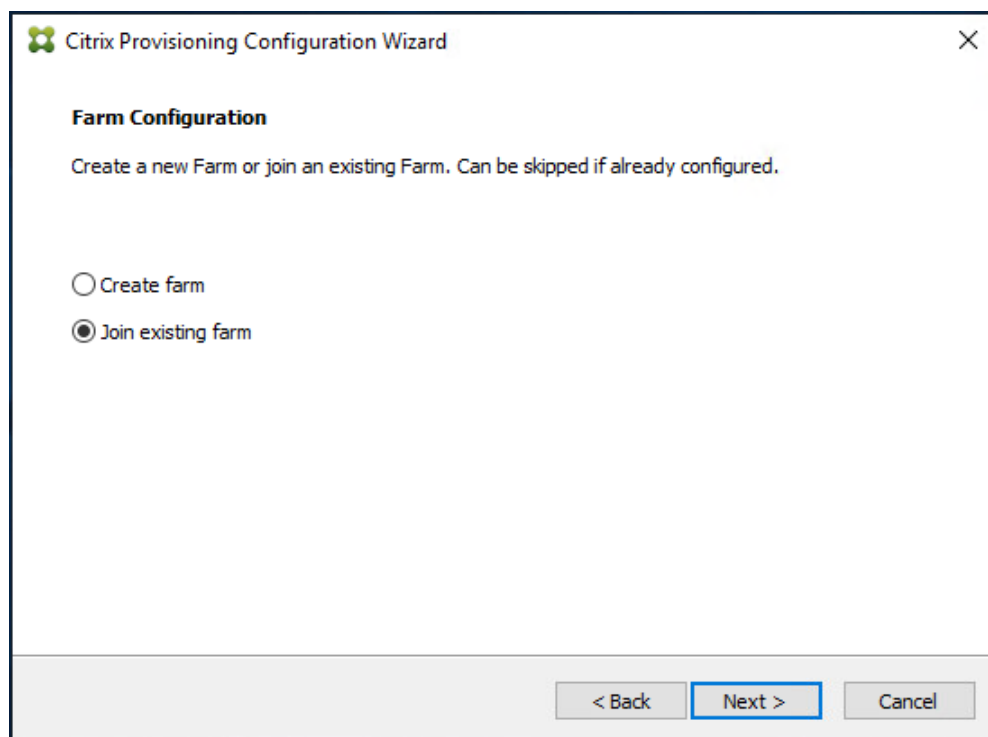


Install Additional PVS Servers

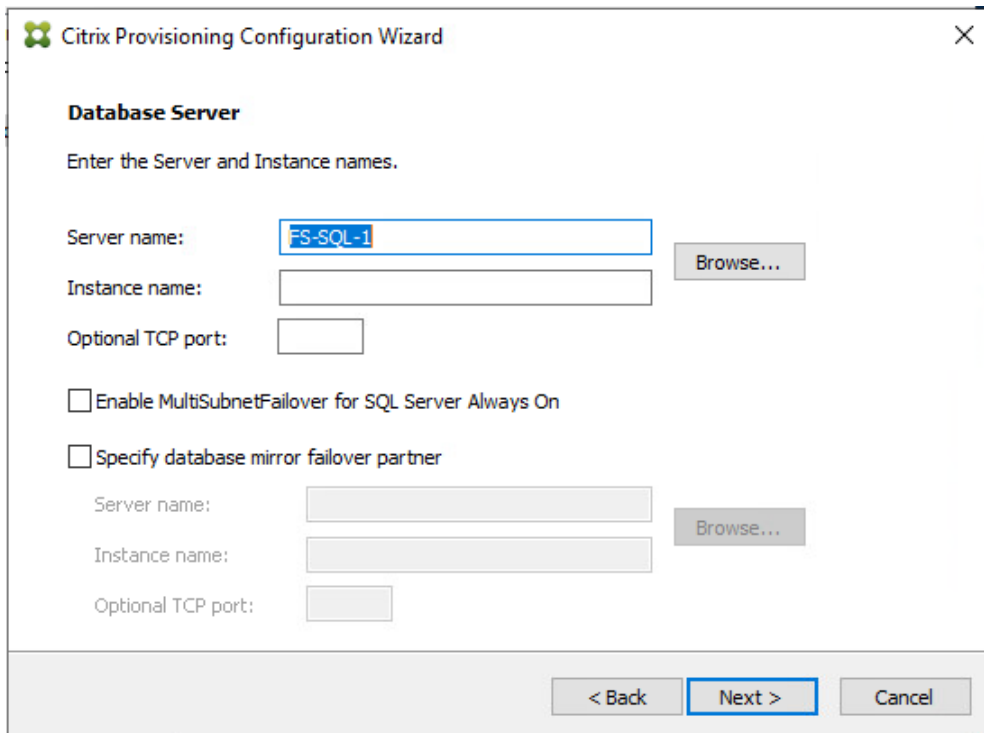
Complete the installation steps on the additional PVS servers up to the configuration step where it asks to Create or Join a farm. In this CVD, we repeated the procedure to add a total of three PVS servers.

To join additional Provisioning servers to the farm already configured in the steps above follow these steps:

1. On the Farm Configuration dialog, select “Join existing farm.”
2. Click Next.



3. Provide the FQDN of the SQL Server.
4. Click Next.



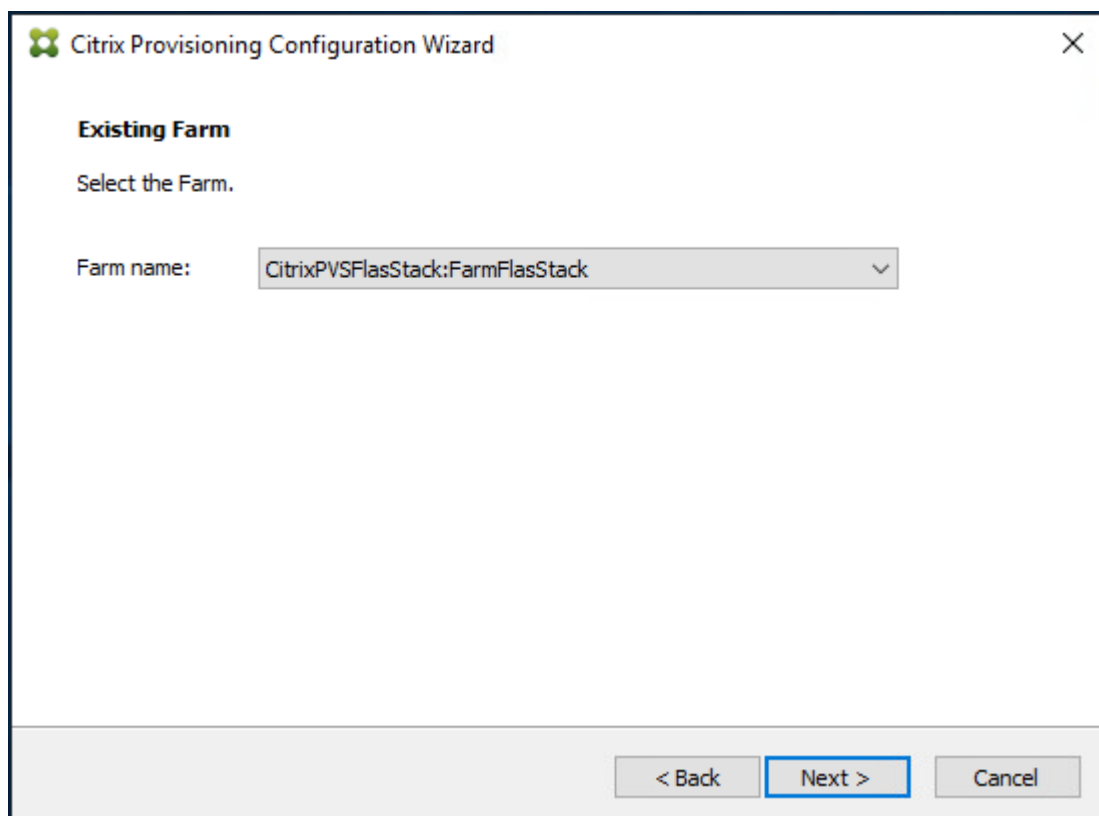
The image shows a screenshot of the Citrix Provisioning Configuration Wizard, specifically the 'Database Server' step. The window title is 'Citrix Provisioning Configuration Wizard' with a close button (X) in the top right corner. The main heading is 'Database Server' followed by the instruction 'Enter the Server and Instance names.'

The form contains the following fields and options:

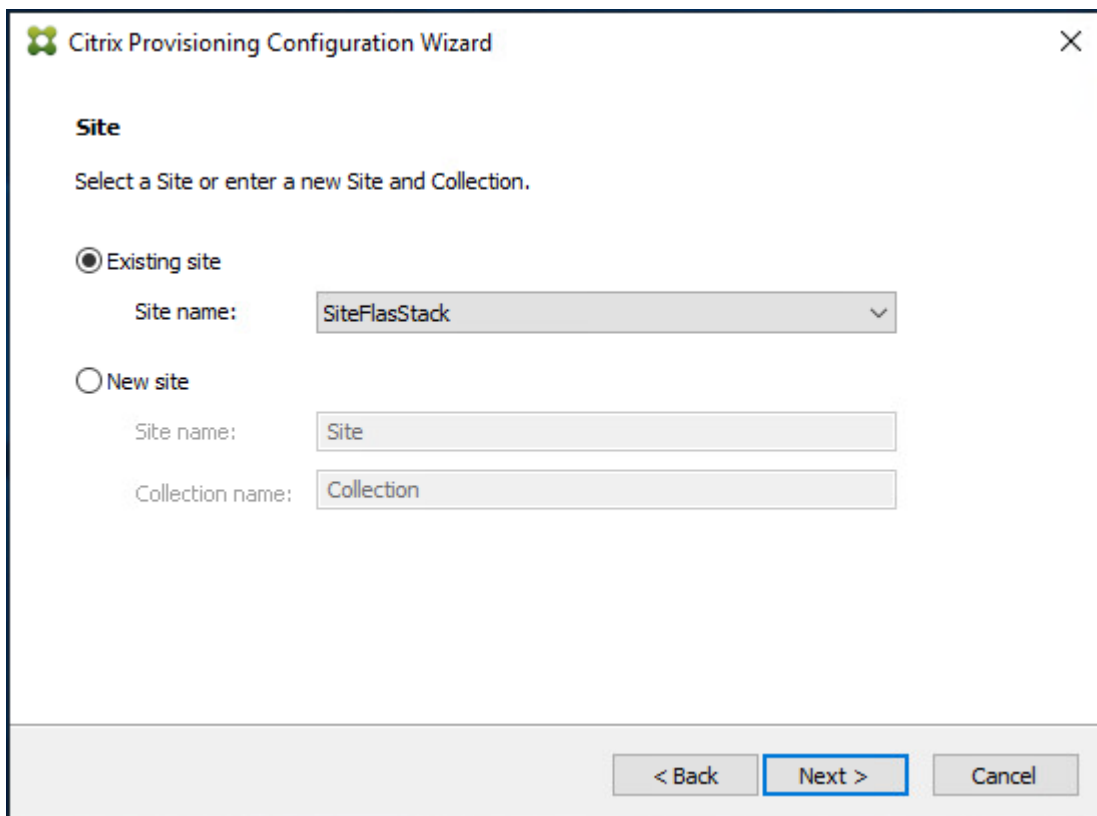
- Server name:** A text box containing 'FS-SQL-1' with a blue selection highlight. To its right is a 'Browse...' button.
- Instance name:** An empty text box.
- Optional TCP port:** An empty text box.
- Enable MultiSubnetFailover for SQL Server Always On
- Specify database mirror failover partner
- Server name:** A disabled (greyed out) text box.
- Instance name:** A disabled (greyed out) text box.
- Optional TCP port:** A disabled (greyed out) text box.
- To the right of the disabled 'Server name' field is a disabled 'Browse...' button.

At the bottom of the window, there are three buttons: '< Back', 'Next >' (which is highlighted with a blue border), and 'Cancel'.

5. Accept the Farm Name.
6. Click Next.



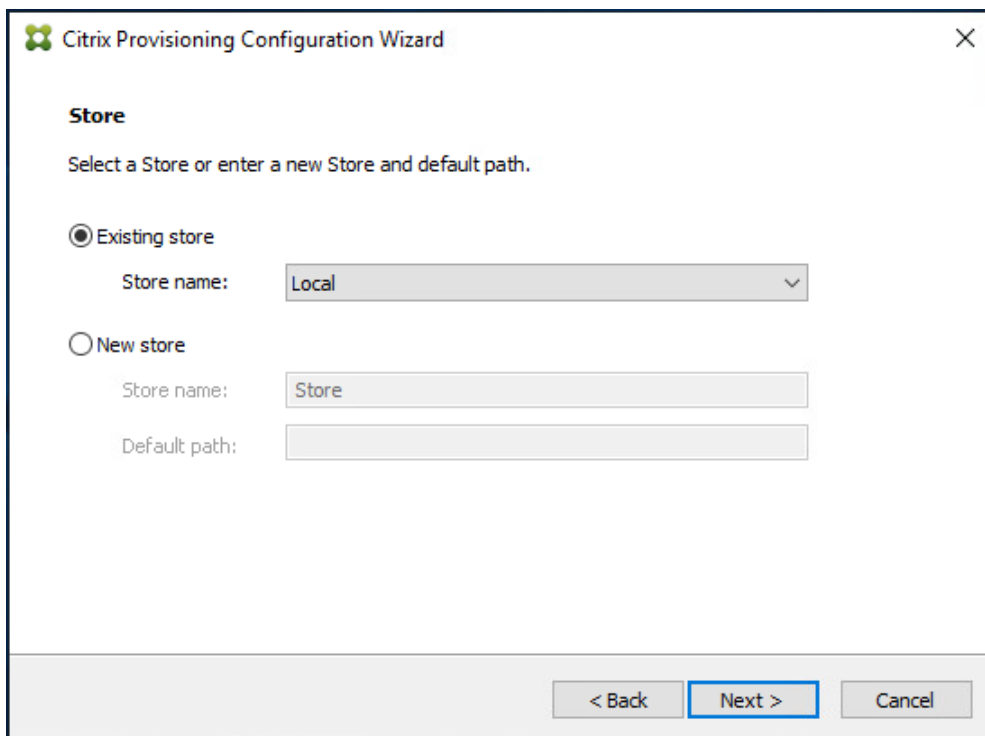
7. Accept the Existing Site.
8. Click Next.



The image shows a screenshot of the Citrix Provisioning Configuration Wizard. The window title is "Citrix Provisioning Configuration Wizard" with a close button (X) in the top right corner. The main heading is "Site". Below the heading, the instruction reads "Select a Site or enter a new Site and Collection." There are two radio button options: "Existing site" (which is selected) and "New site". Under "Existing site", there is a "Site name:" label and a dropdown menu containing the text "SiteFlasStack". Under "New site", there are two text input fields: "Site name:" with the text "Site" and "Collection name:" with the text "Collection". At the bottom of the dialog, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

9. Accept the existing vDisk store.

10. Click Next.

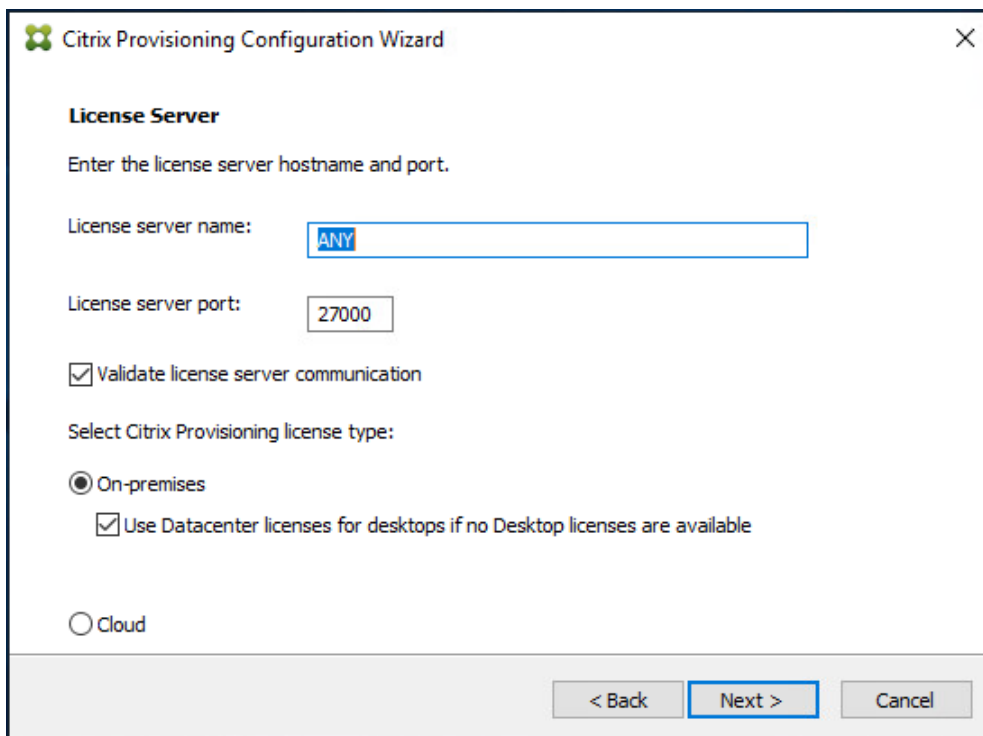


The image shows a screenshot of the Citrix Provisioning Configuration Wizard. The window title is "Citrix Provisioning Configuration Wizard" with a close button (X) in the top right corner. The main heading is "Store". Below the heading, the instruction reads "Select a Store or enter a new Store and default path." There are two radio button options: "Existing store" (which is selected) and "New store". Under "Existing store", there is a "Store name:" label followed by a dropdown menu showing "Local". Under "New store", there are two text input fields: "Store name:" containing the text "Store" and "Default path:" which is currently empty. At the bottom of the window, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

11. Provide the FQDN of the license server.

12. Optionally, provide a port number if changed on the license server.


13. Click Next.



The image shows a screenshot of the 'Citrix Provisioning Configuration Wizard' window, specifically the 'License Server' step. The window title is 'Citrix Provisioning Configuration Wizard' with a close button (X) in the top right corner. Below the title bar, the section is titled 'License Server'. A sub-instruction reads 'Enter the license server hostname and port.' There are two input fields: 'License server name:' with the text 'ANY' entered, and 'License server port:' with the text '27000' entered. Below these fields is a checked checkbox labeled 'Validate license server communication'. The next section is 'Select Citrix Provisioning license type:', which has two radio button options: 'On-premises' (selected) and 'Cloud'. Under the 'On-premises' option, there is a checked checkbox labeled 'Use Datacenter licenses for desktops if no Desktop licenses are available'. At the bottom of the window, there are three buttons: '< Back', 'Next >', and 'Cancel'. The 'Next >' button is highlighted with a blue border.

14. Provide the PVS service account information.

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:

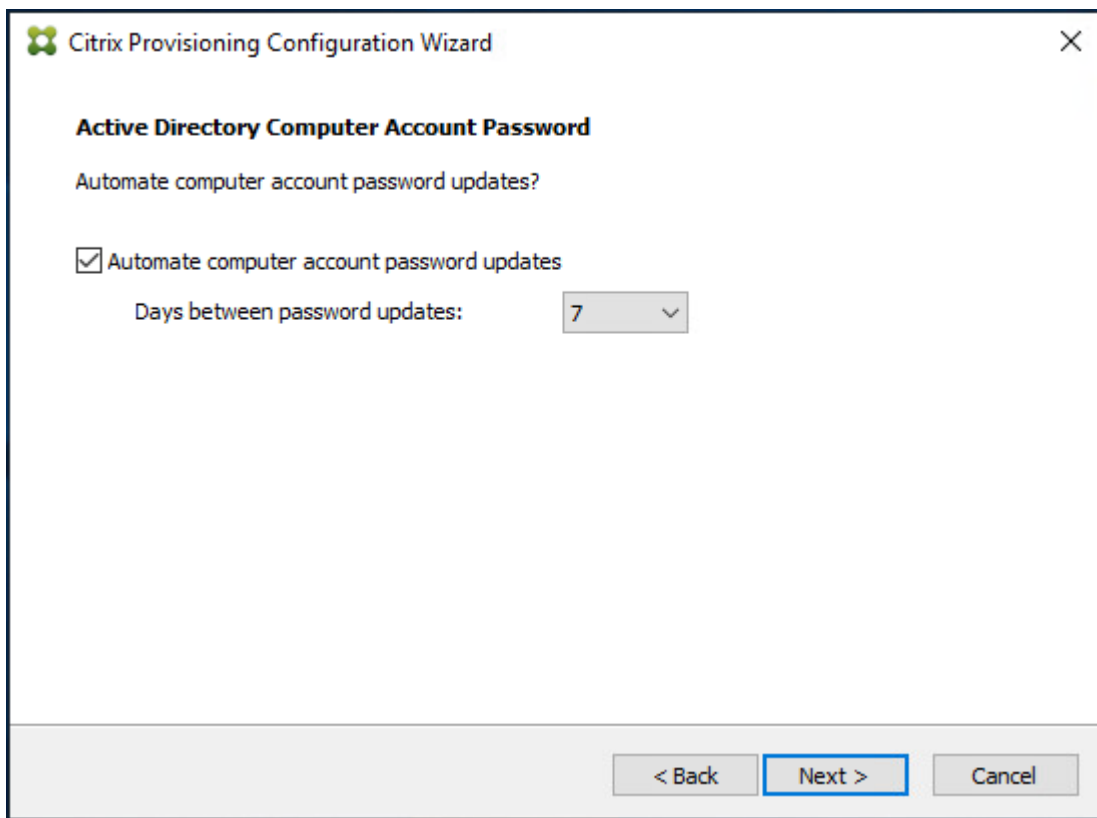
Domain:

Password:

Confirm password:


16. Set the Days between password updates to 7.

17. Click Next.




18. Accept the network card settings.


19. Click Next.

 Citrix Provisioning Configuration Wizard ✕

Network Communications

Specify network settings.

Streaming network cards:  10.54.1.3

Management network card:  10.54.1.3

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.

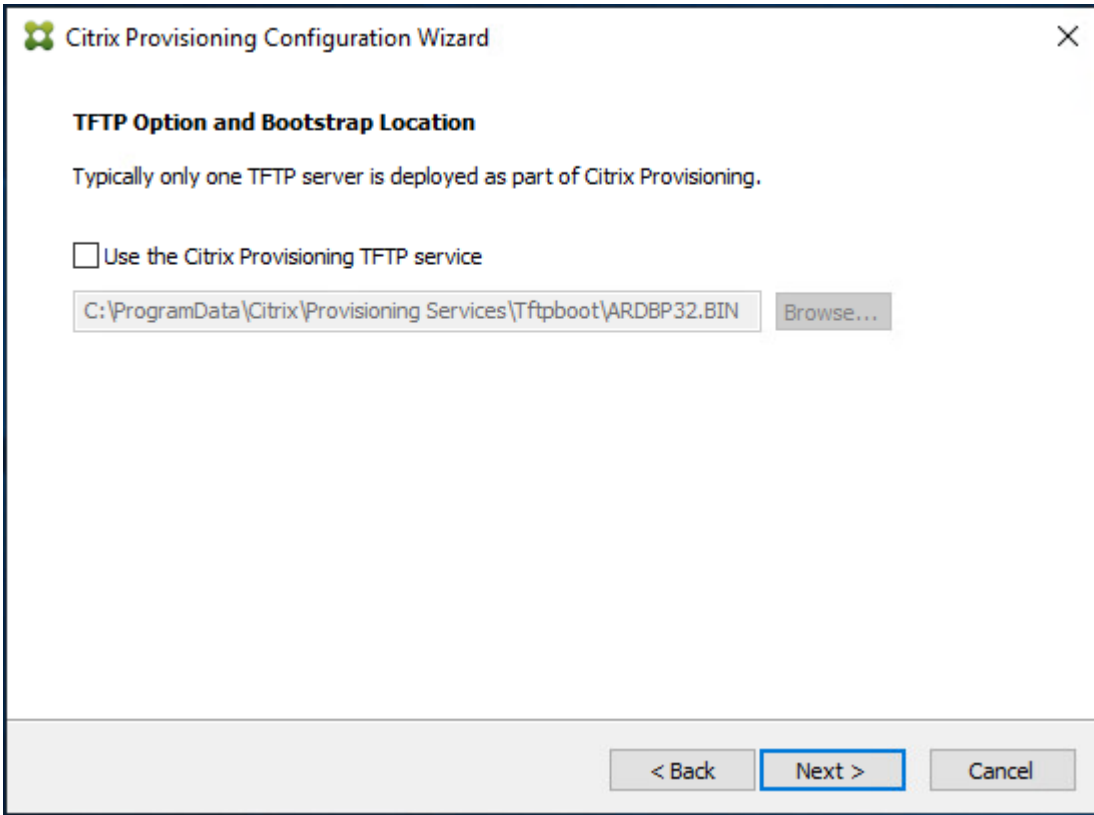
Note: All servers must have the same port configurations.

First communications port:

Console port:


20. Select Use the Provisioning Services TFTP service checkbox.

21. Click Next.



22. If Soap Server is used, provide details.

23. 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:

SSL certificate:

Subject	Issuer	Expiration Date
---------	--------	-----------------

24. If desired, fill in Problem Report Configuration.

25. Click Next.

Citrix Provisioning Configuration Wizard [X]

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.

< Back **Next >** Cancel

26. Click Finish to start the installation process.

Citrix Provisioning Configuration Wizard [X]

Finish

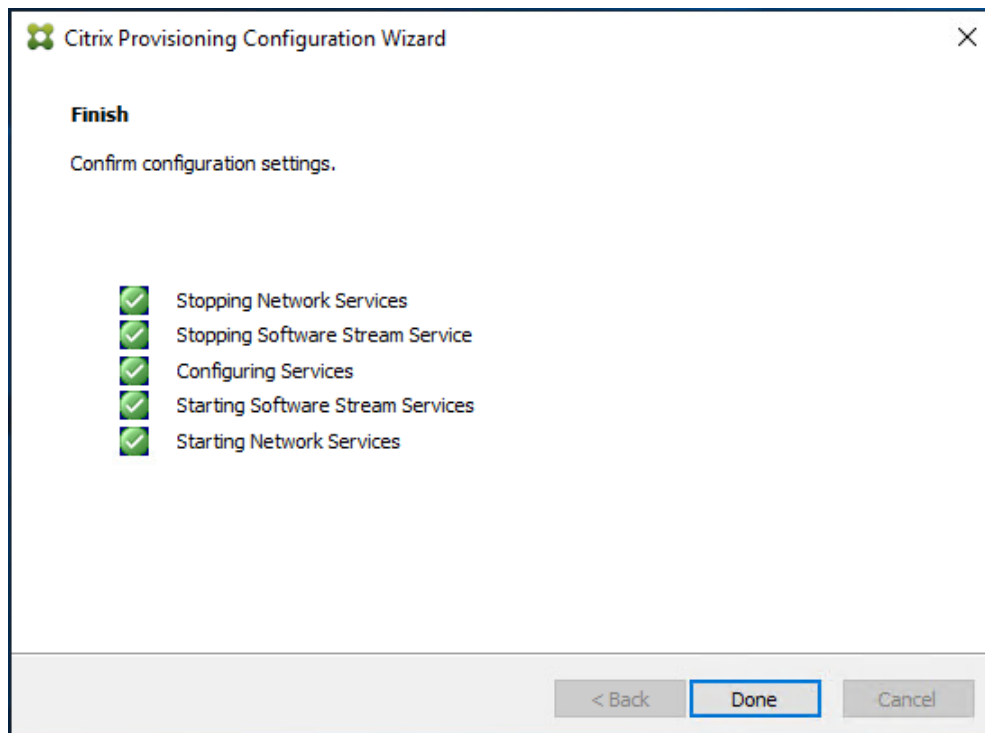
Confirm configuration settings.

PXE - Not used
Database Server = fs-sql-1\
Farm = FlashStackDb:FarmFlashStack
Site and Collection = SiteFlashStack, Collection
AD Group = FSL151K.LOCAL/Users/Domain Admins
Store and Default Path = Store, E:\Store
License Server:Port = ANY:27000
User Account = FSL151K.LOCAL\pvs_svc
Computer account password changes every 7 days
Communications - First Port = 6890, Last Port = 6909
Console - Soap Port = 54321
NIC - Selected IP = 10.54.1.3
Management NIC - Selected IP = 10.54.1.3
TFTP - Not used

Automatically Start Services

< Back **Finish** Cancel

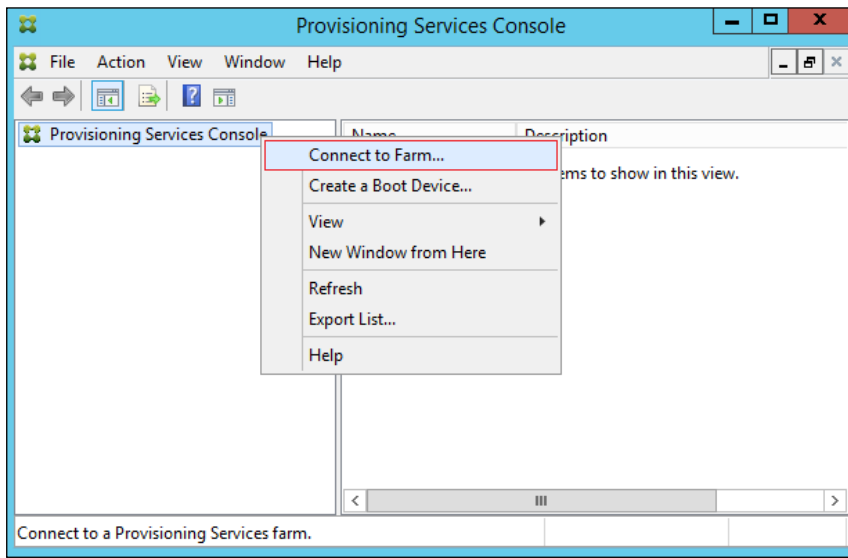
27. Click Done when the installation finishes.



Note: You can optionally install the Provisioning Services console on the second PVS server following the procedure in the section Installing Provisioning Services.

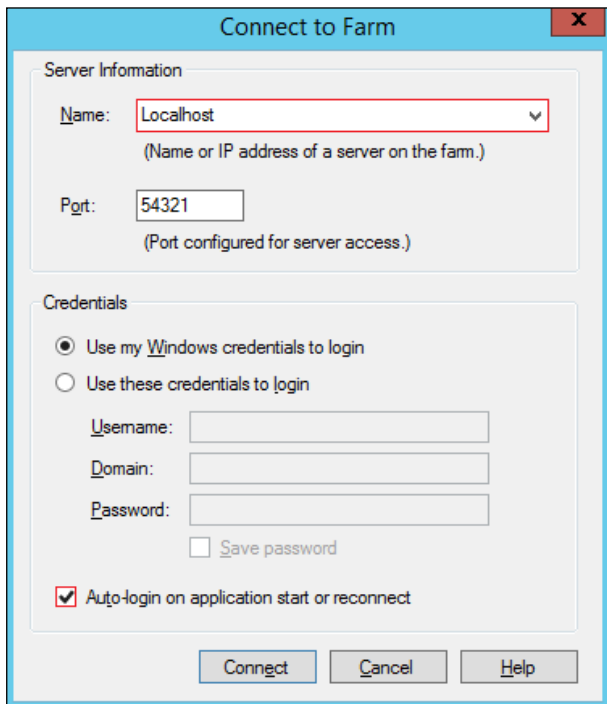
28. After completing the steps to install the three additional PVS servers, launch the Provisioning Services Console to verify that the PVS Servers and Stores are configured and that DHCP boot options are defined.

29. Launch Provisioning Services Console and select Connect to Farm.

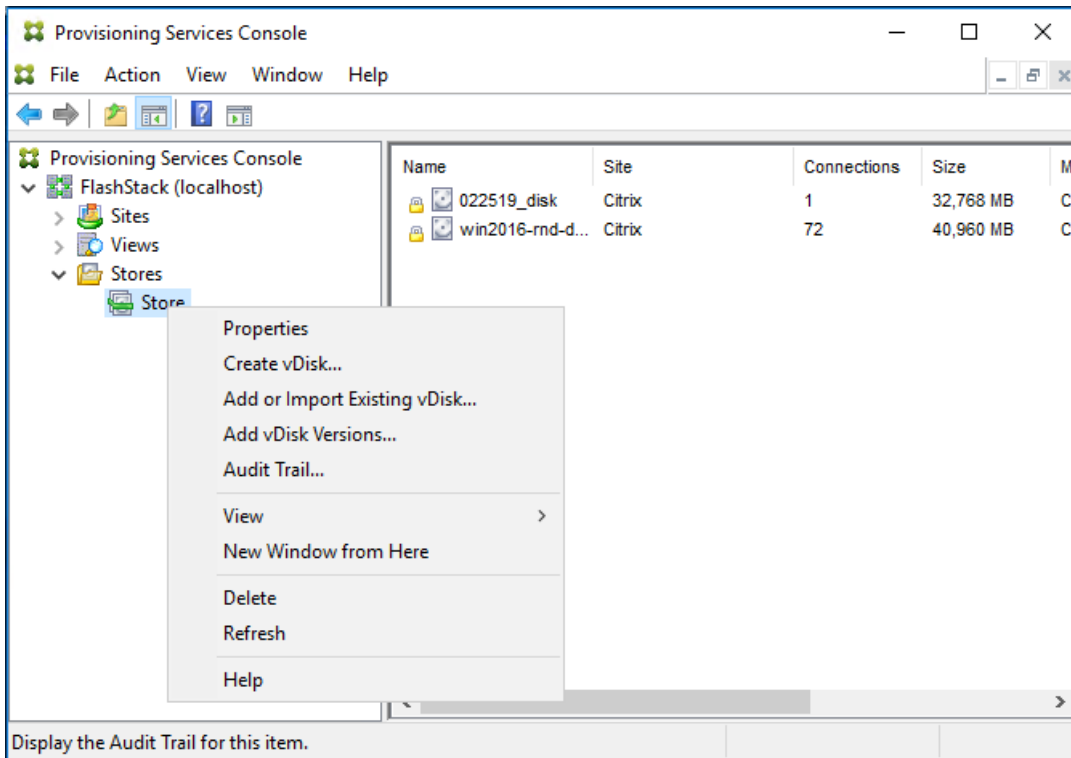


30. Enter localhost for the PVS1 server.

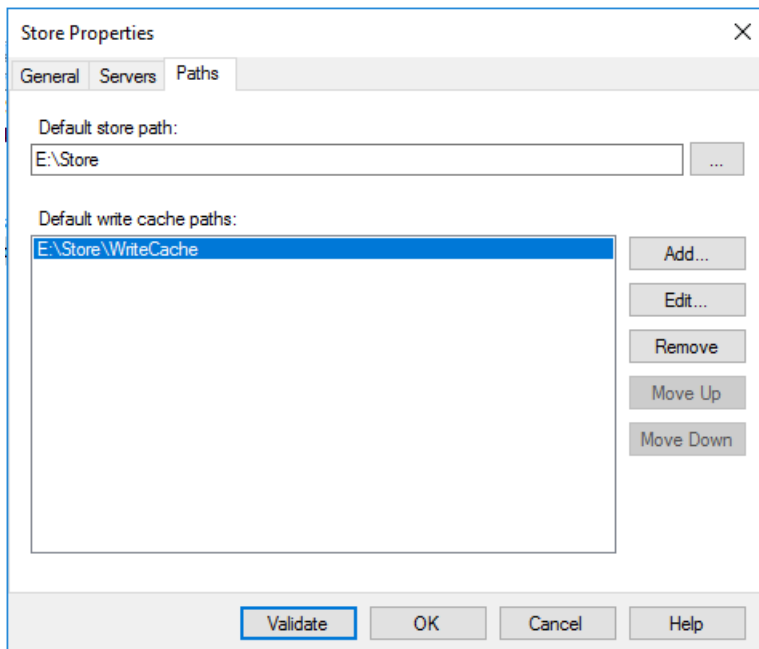
31. Click Connect.



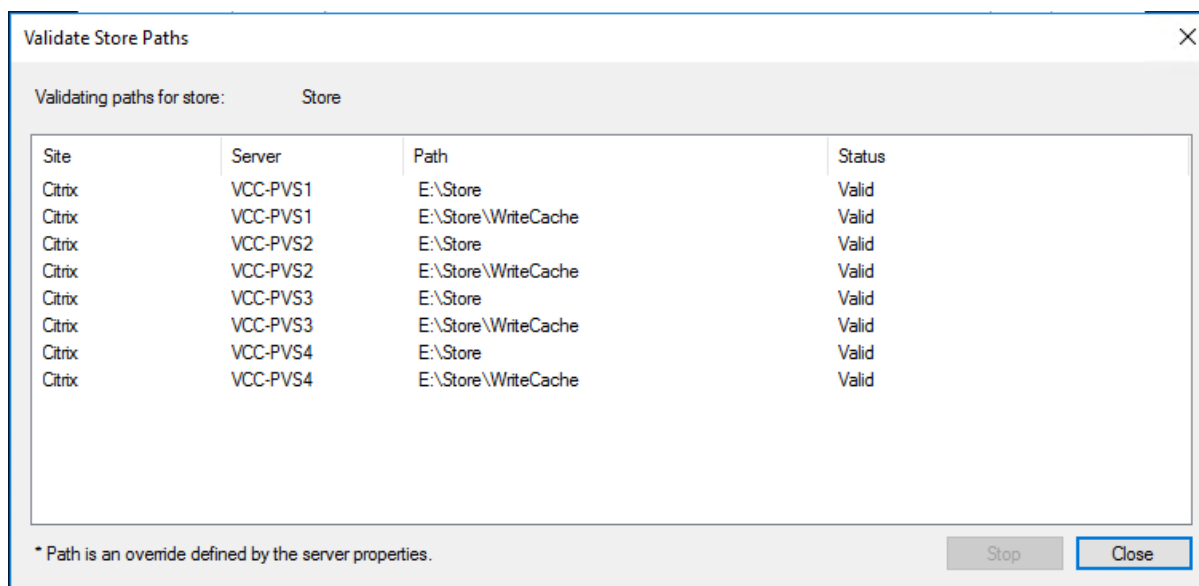
32. Select Store Properties from the drop-down list.



33. In the Store Properties dialog, add the Default store path to the list of Default write cache paths.



34. Click Validate. If the validation is successful, click Close and then click OK to continue.



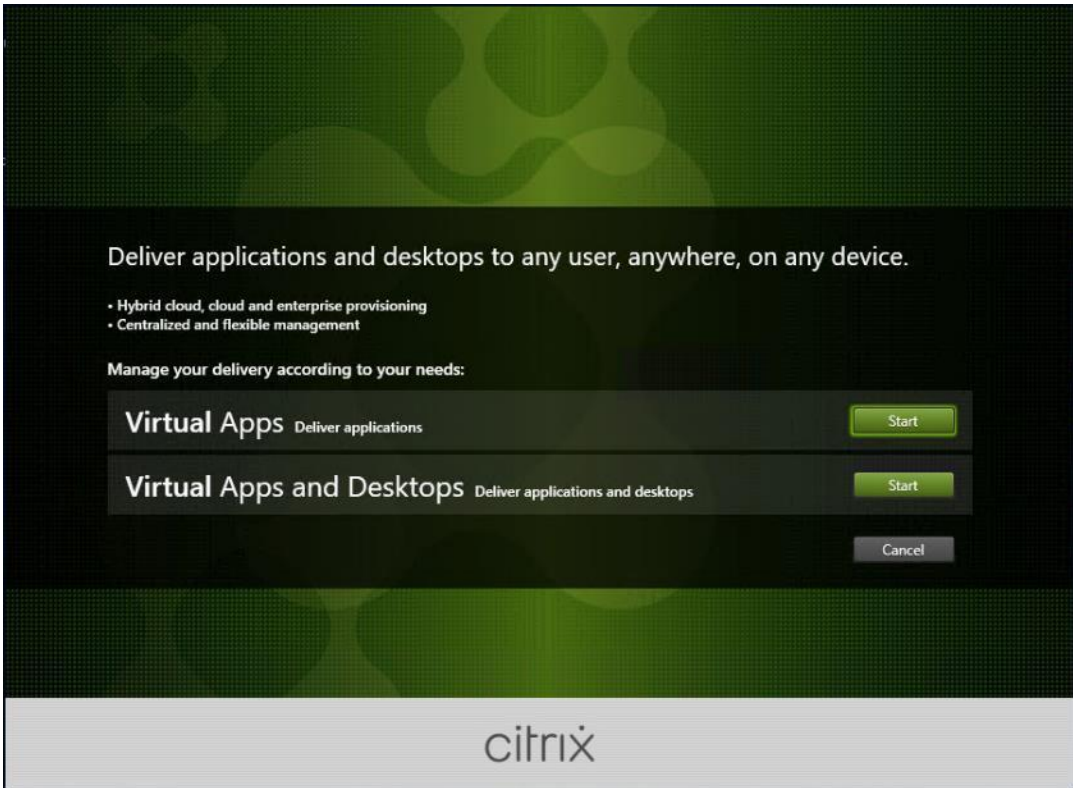
Install Citrix Virtual Apps and Desktops Virtual Desktop Agents

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.

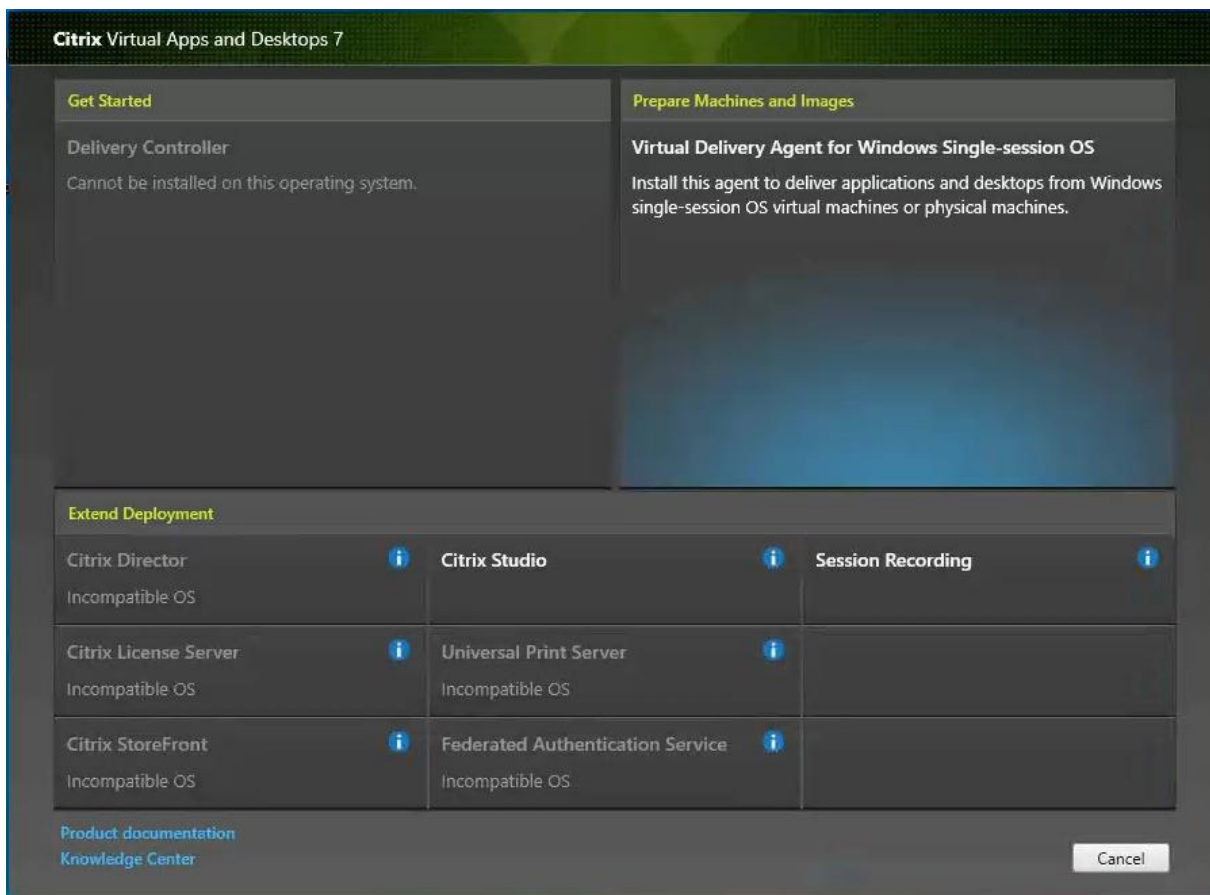
By default, when you install the Virtual Delivery Agent, Citrix User Profile Management is installed silently on master images. (Using profile management as a profile solution is optional but FSLogix was used for this CVD and is described in a later section.)

To install Citrix Virtual Apps and Desktops Virtual Desktop Agents, follow these steps:

1. Launch the Citrix Virtual Apps and Desktops installer from the Citrix_Virtual_Apps_and_Desktops_7_2109 ISO.
2. Click Start on the Welcome Screen.



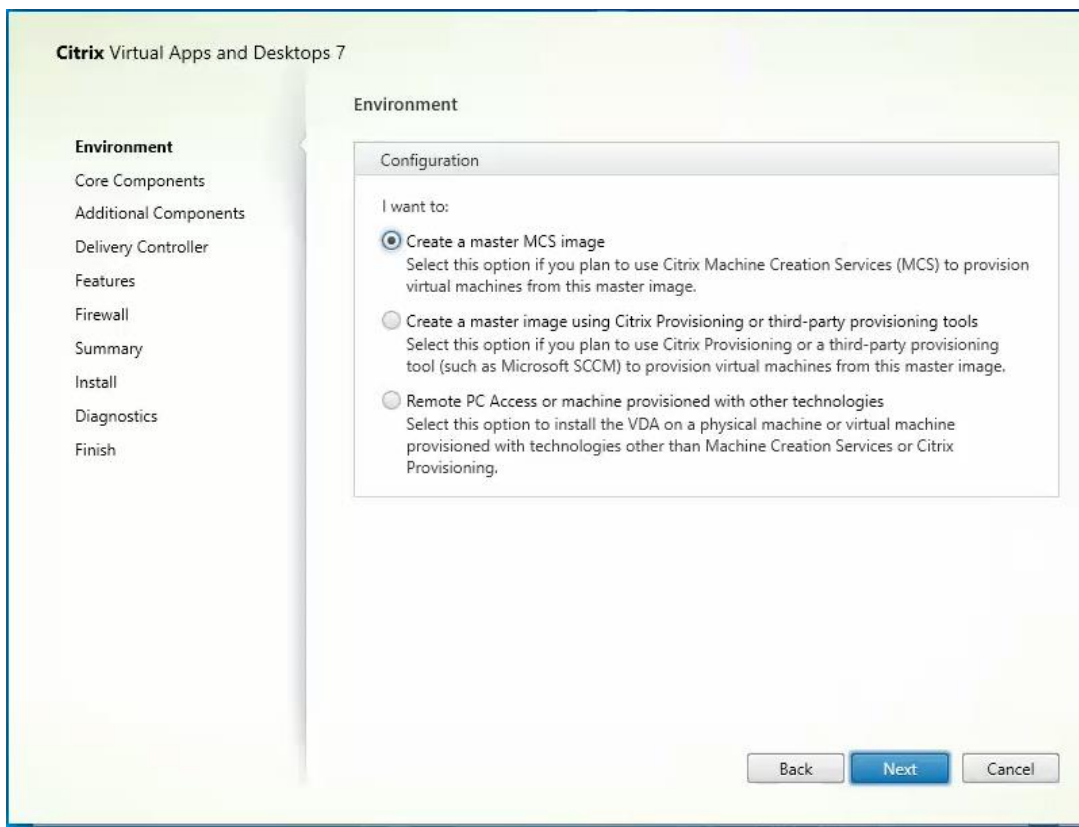
3. To install the VDA for the Hosted Virtual Desktops (VDI), select Virtual Delivery Agent for Windows Single-session OS.



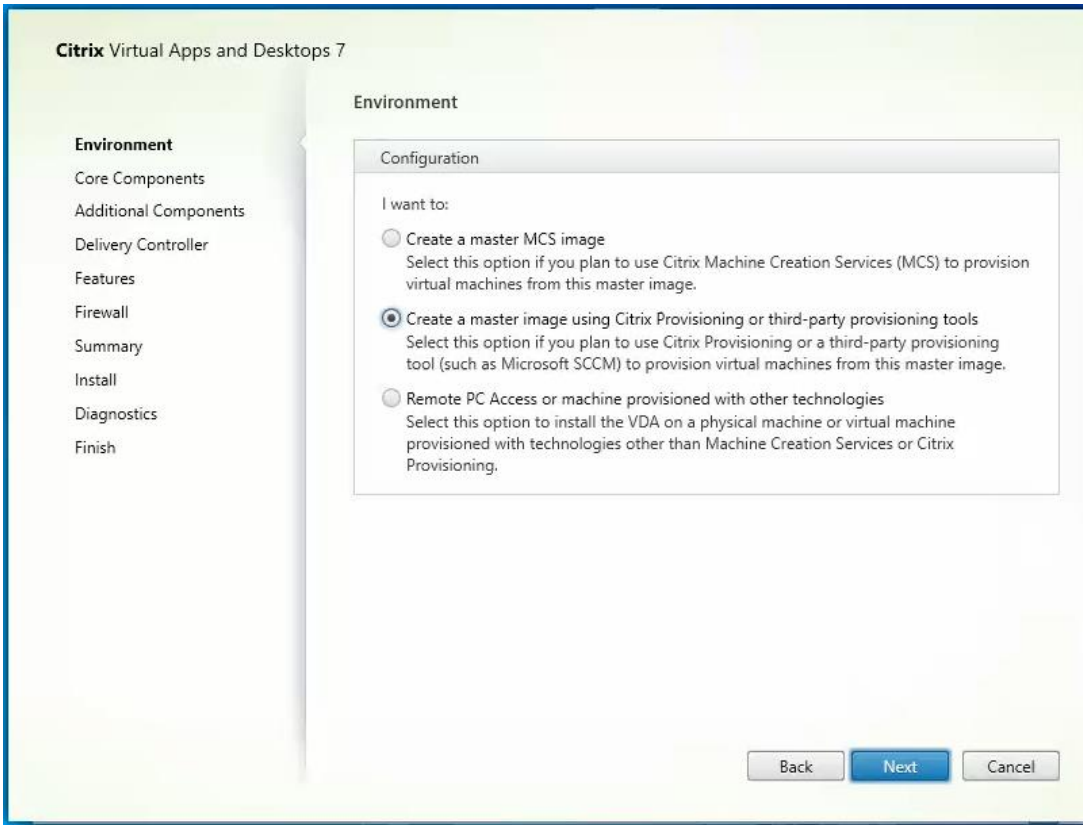
Note: When installing Virtual Delivery Agent for Windows Multi-session OS and follow the same basic steps.



4. Select Create a master MCS Image.
5. Click Next.

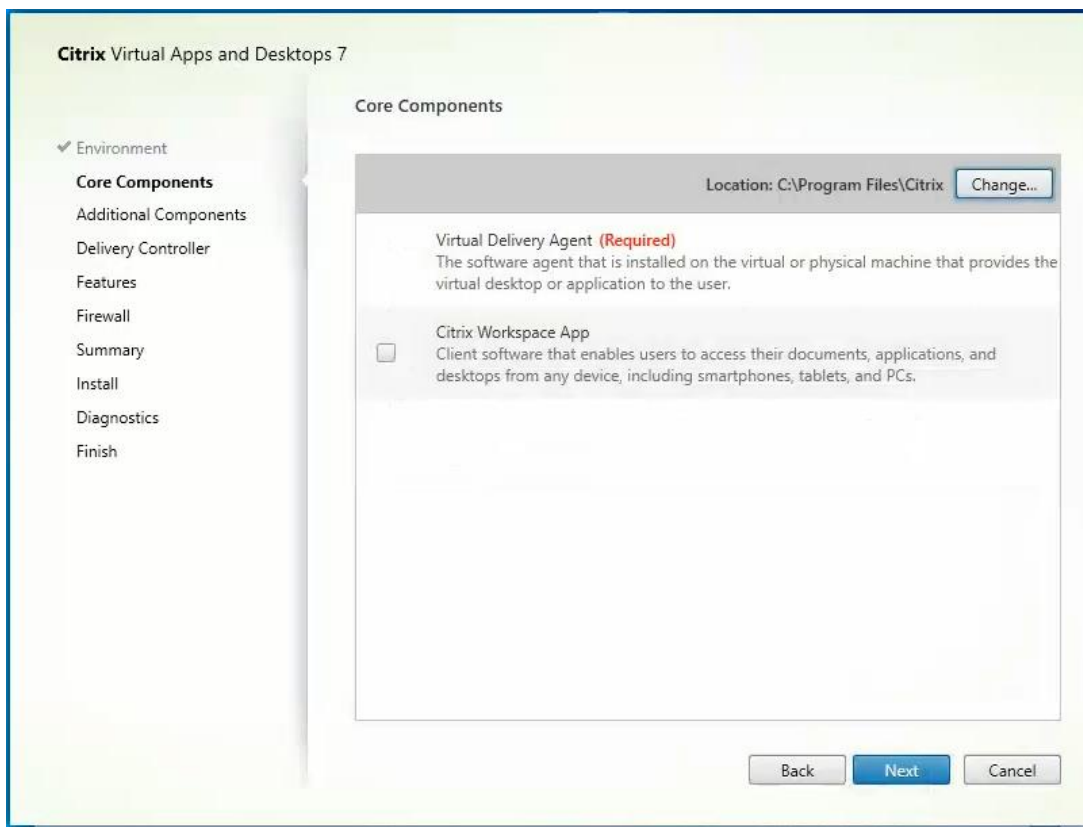


6. Select “Create a master image using Citrix Provisioning or third-party provisioning tools” when building image to be delivered with Citrix Provisioning tools.



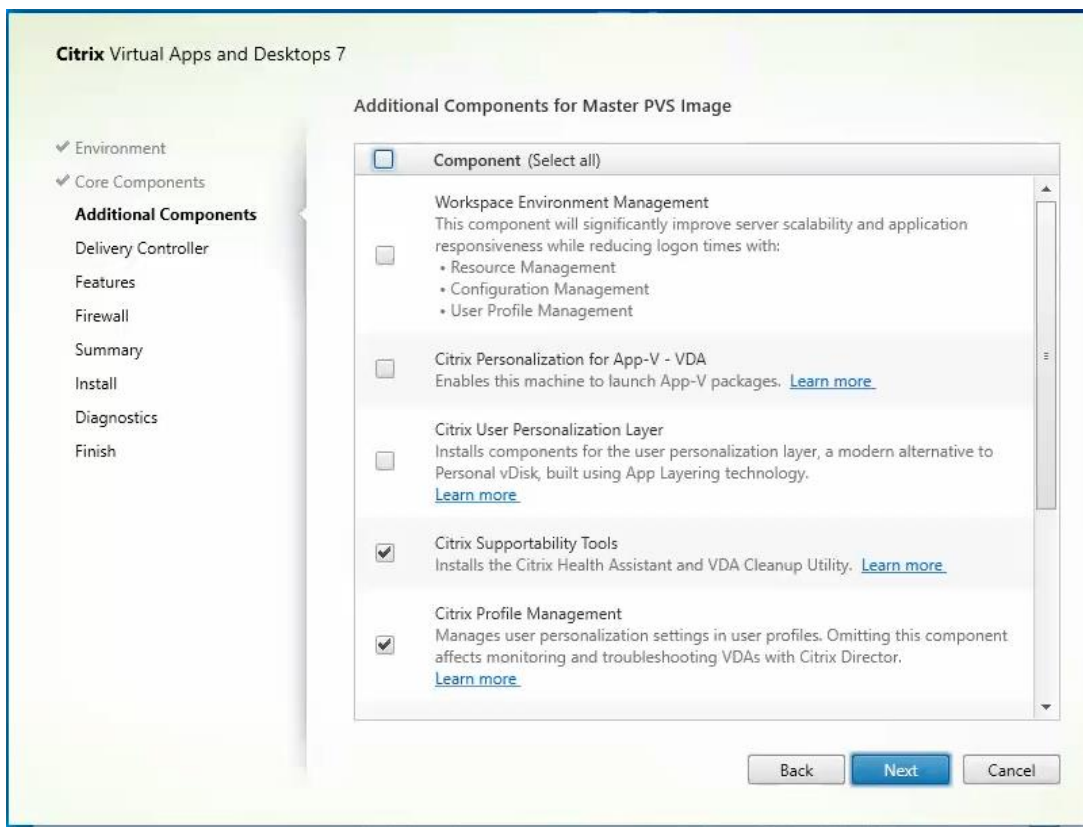
7. Optional; do not select Citrix Workspace App.

8. Click Next.



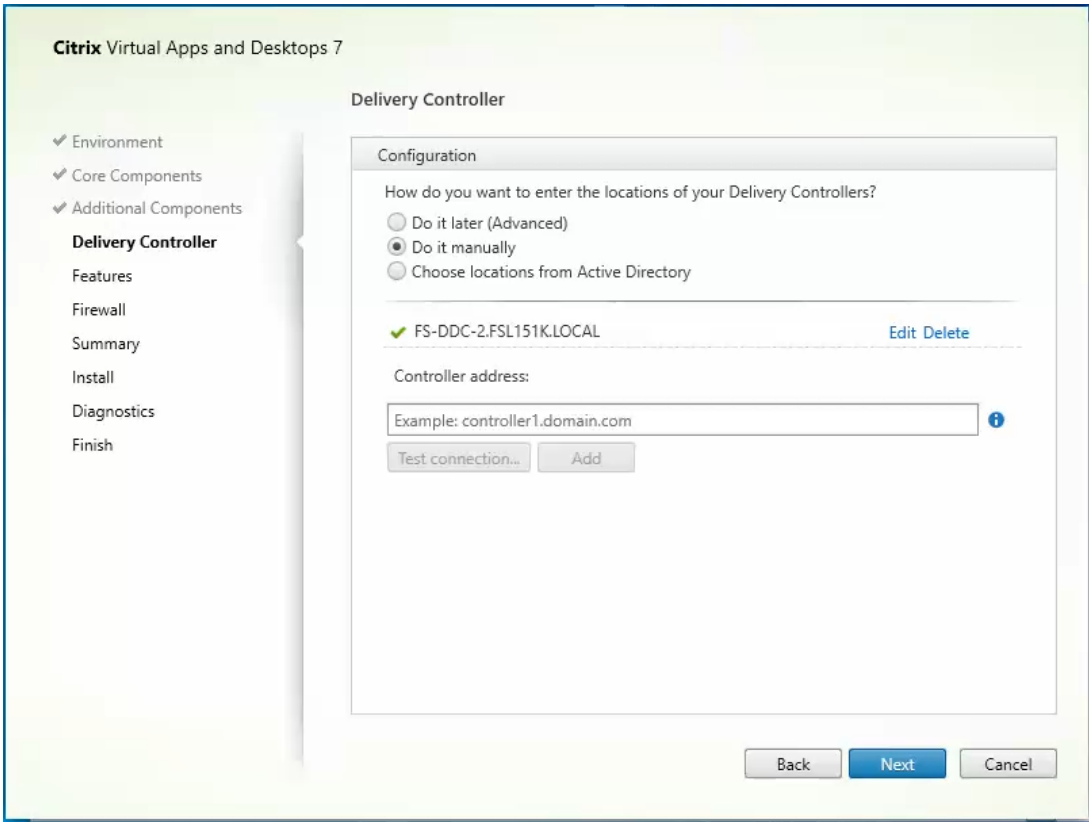
9. Select the additional components required for your image. In this design, only default components were installed on the image.

10. Click Next.



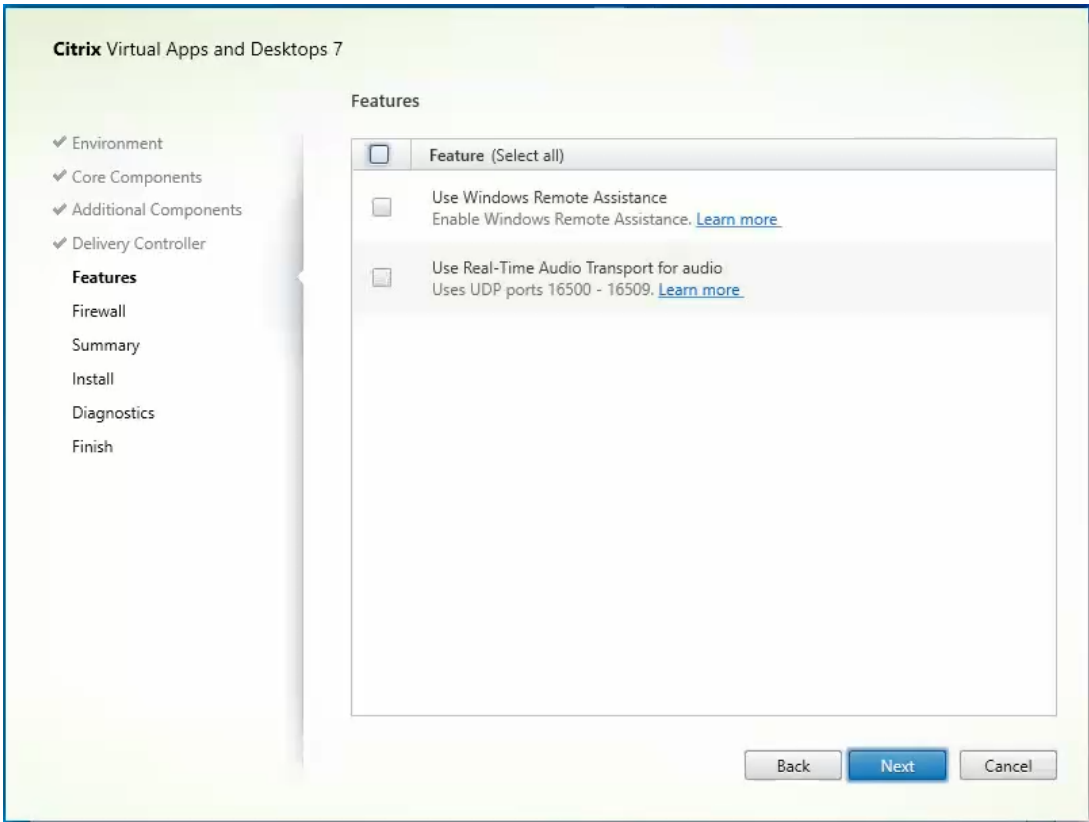
11. Configure Delivery Controllers at this time.

12. Click Next.



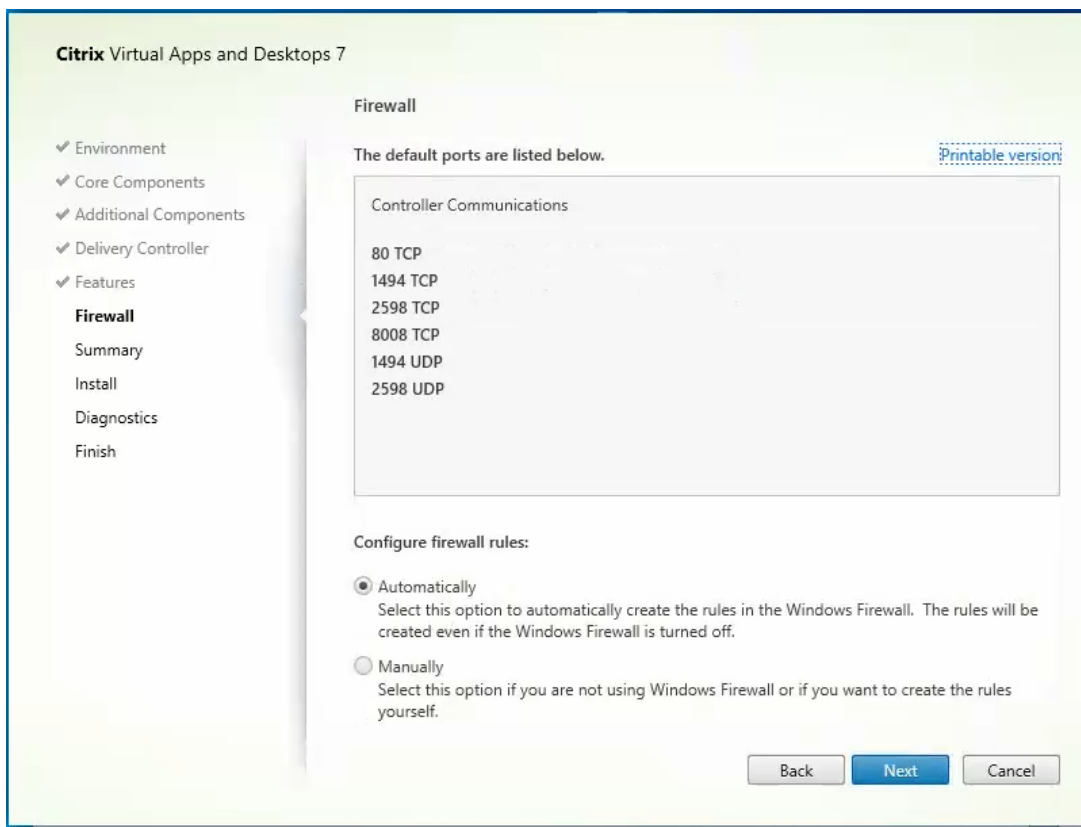
13. Optional: select additional features.

14. Click Next.

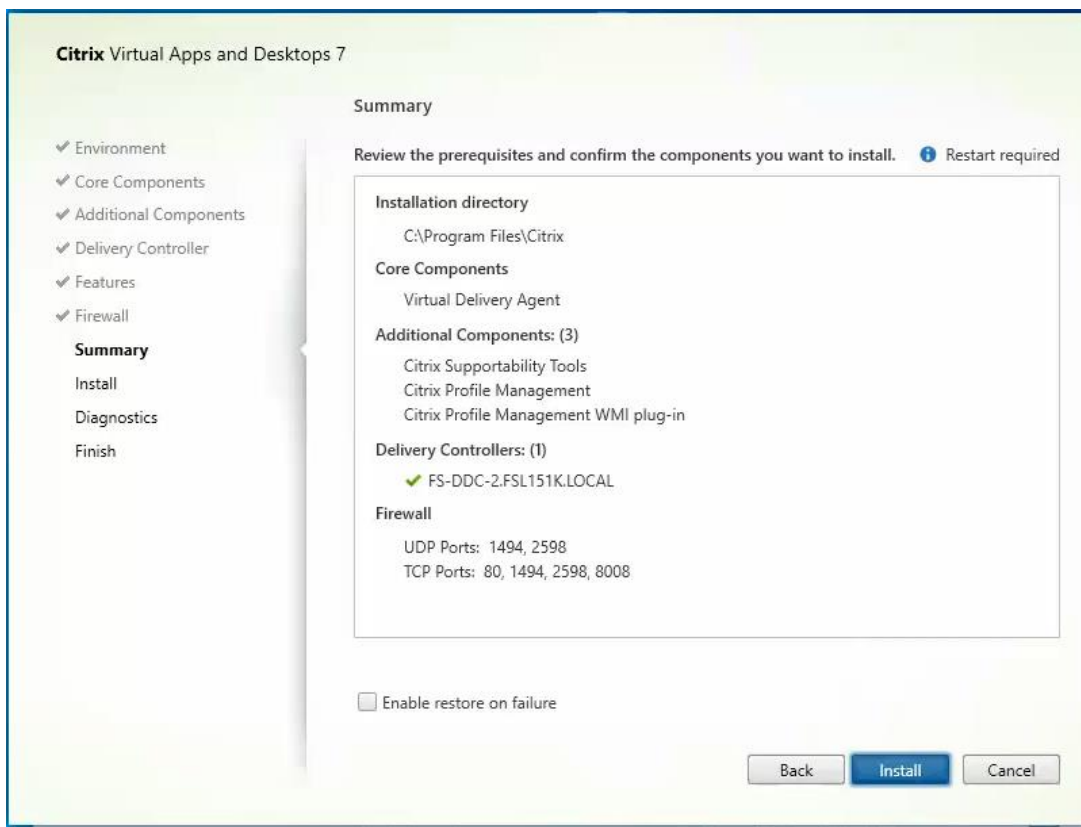


Note: Allow the firewall rules to be configured Automatically.

15. Click Next.

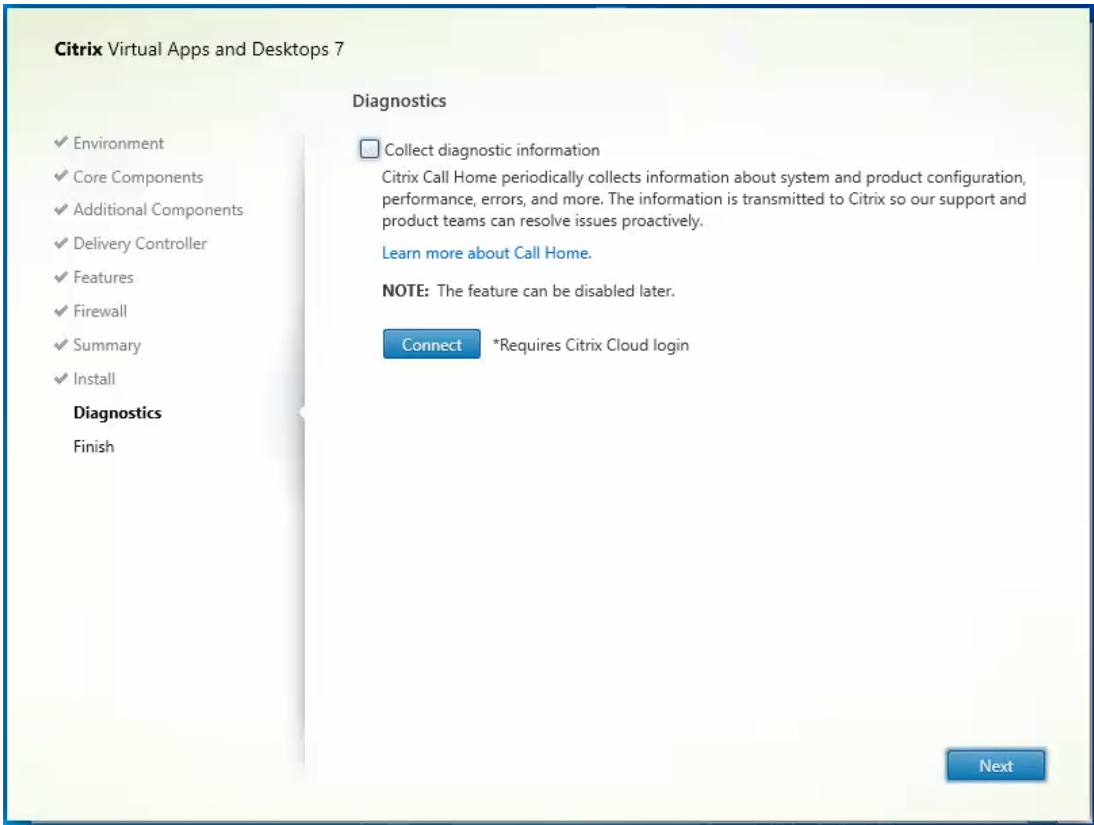


16. Verify the Summary and click Install.



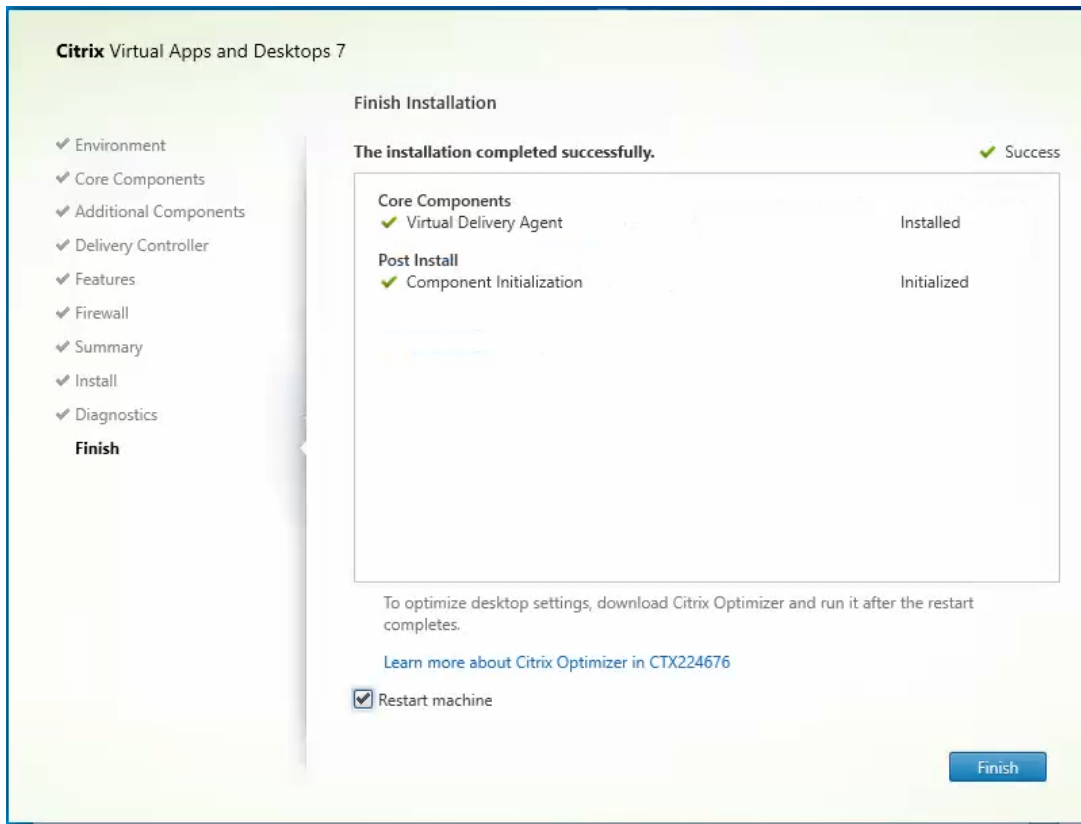
17. Optional: configure Citrix Call Home participation.

18. Click Next.



19. Check Restart Machine.

20. Click Finish and the machine will reboot automatically.

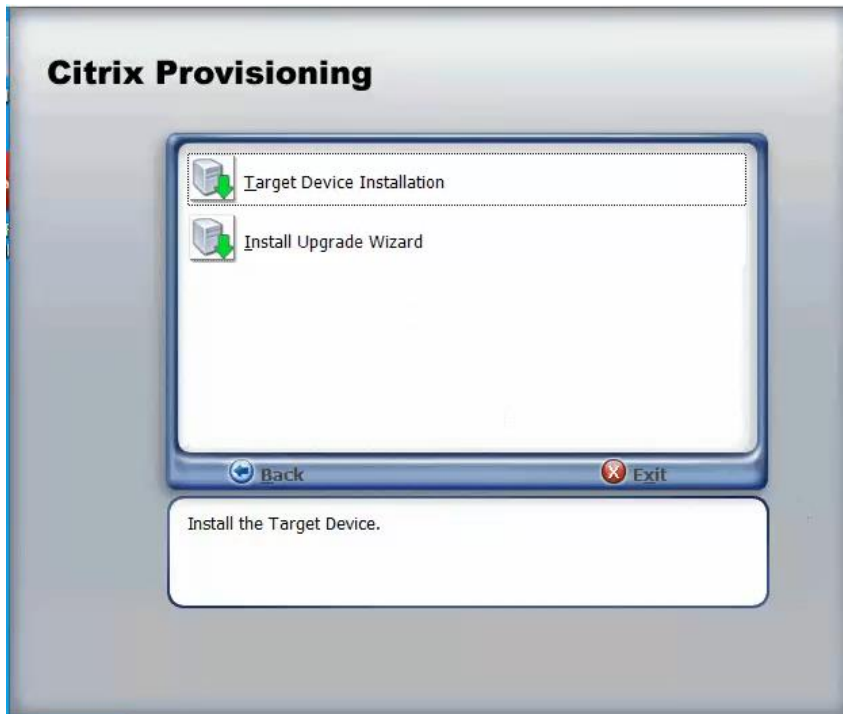


Install the Citrix Provisioning Server Target Device Software

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.

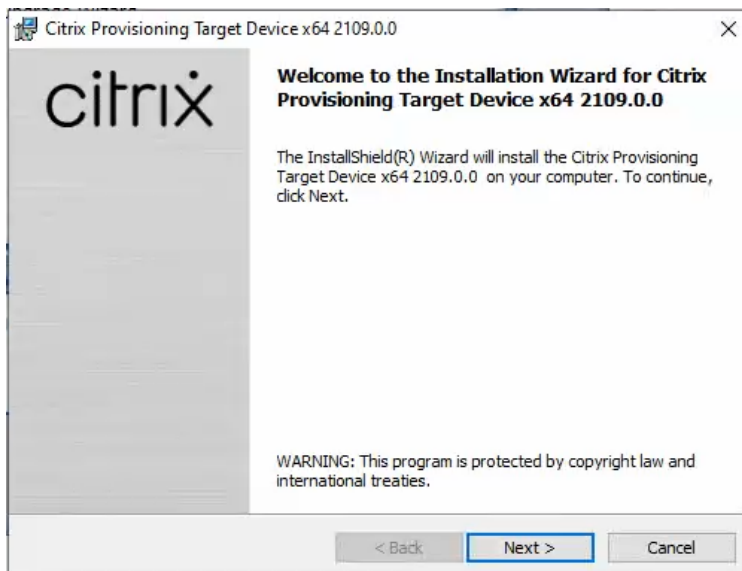
To install the Citrix Provisioning Server Target Device software, follow these steps:

1. Launch the PVS installer from the Citrix_Provisioning_2109 ISO.
2. Click Target Device Installation.



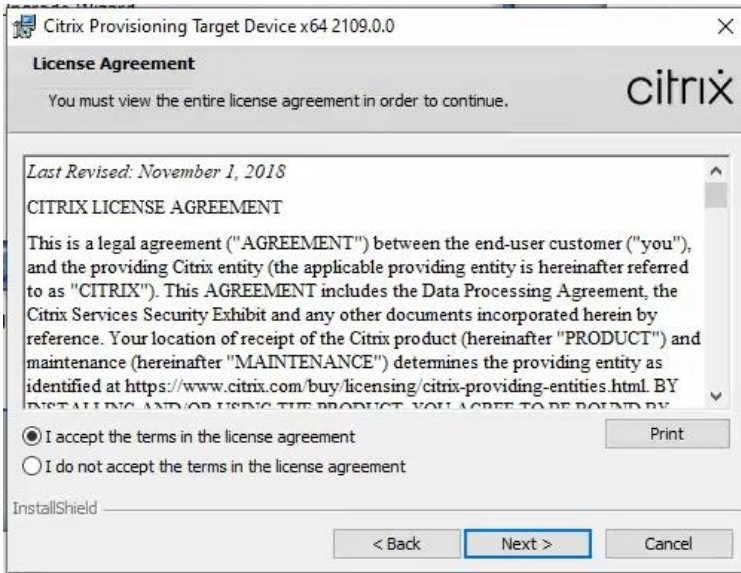
Note: The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.

3. Click Next.



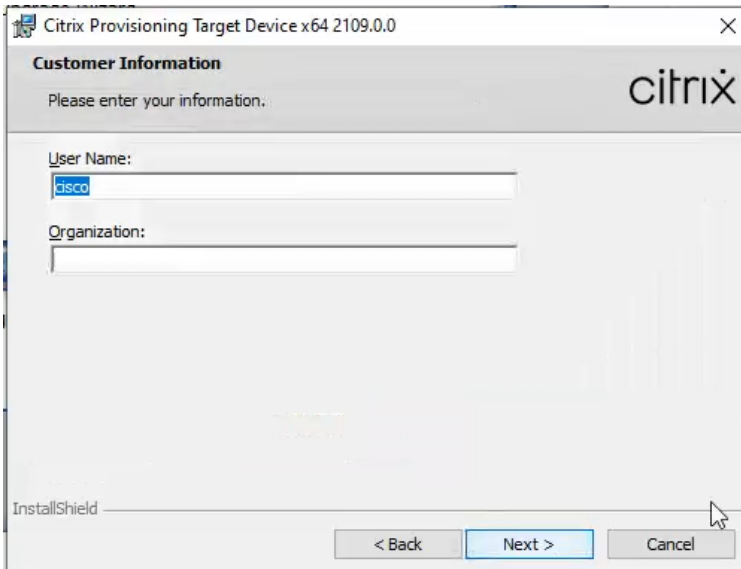
4. Indicate your acceptance of the license by selecting the "I have read, understand, and accept the terms of the license agreement" radio button.

5. Click Next.



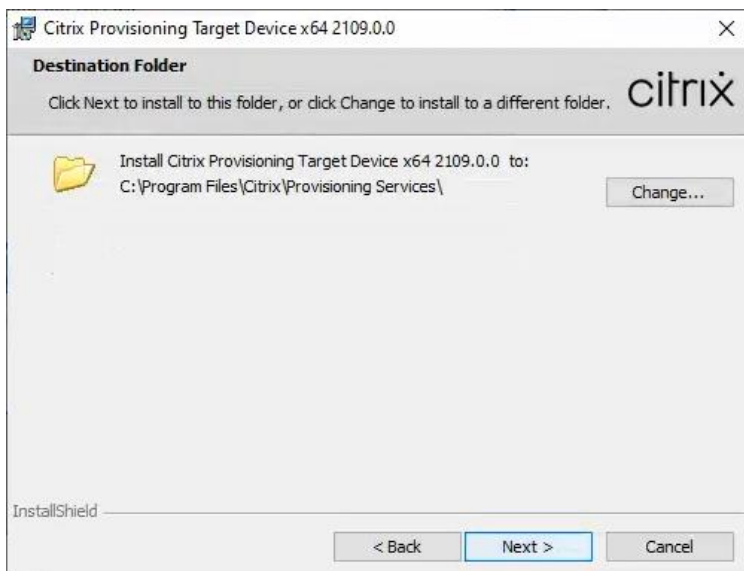
6. Optional: provide the Customer information.

7. Click Next.

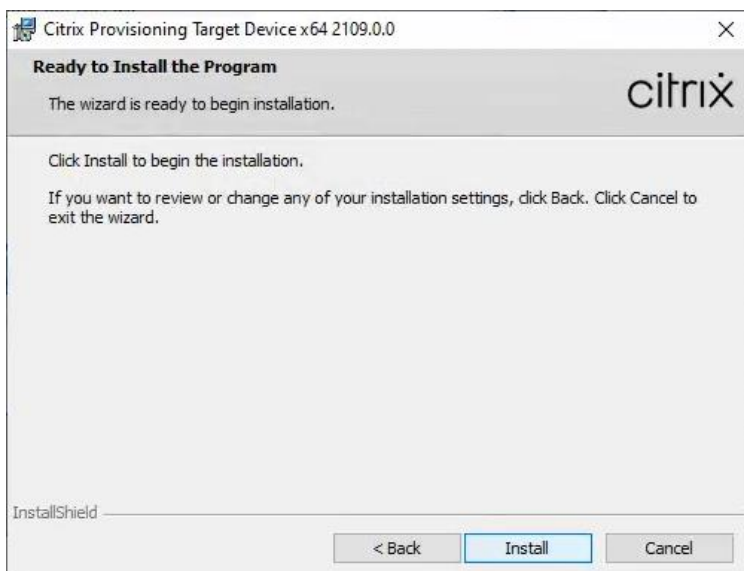


8. Accept the default installation path.

9. Click Next.



10. Click Install.



11. Deselect the checkbox to launch the Imaging Wizard and click Finish.



12. Click Yes to reboot the machine.

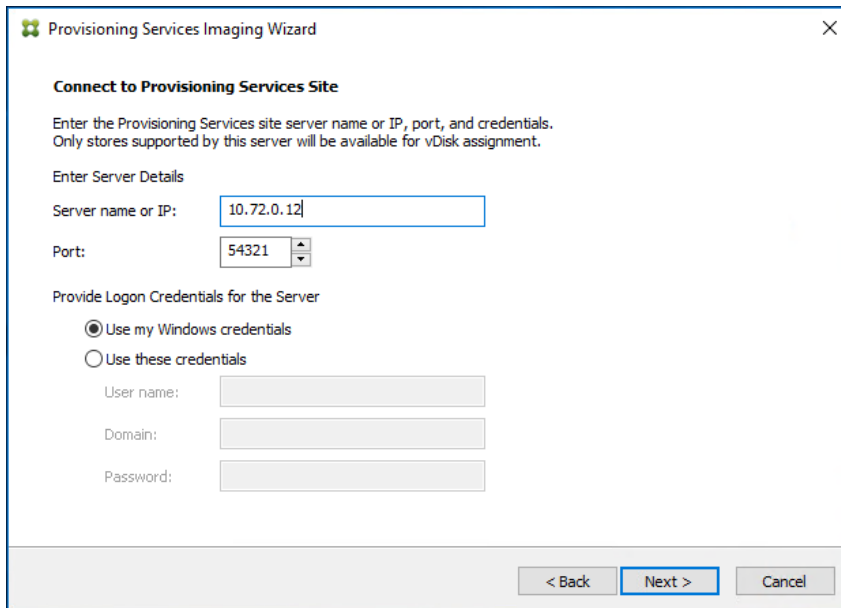
Create Citrix Provisioning Server vDisks

The PVS Imaging Wizard automatically creates a base vDisk image from the master target device. To create the Citrix Provisioning Server vDisks, follow these steps:

1. The PVS Imaging Wizard's Welcome page appears.
2. Click Next.

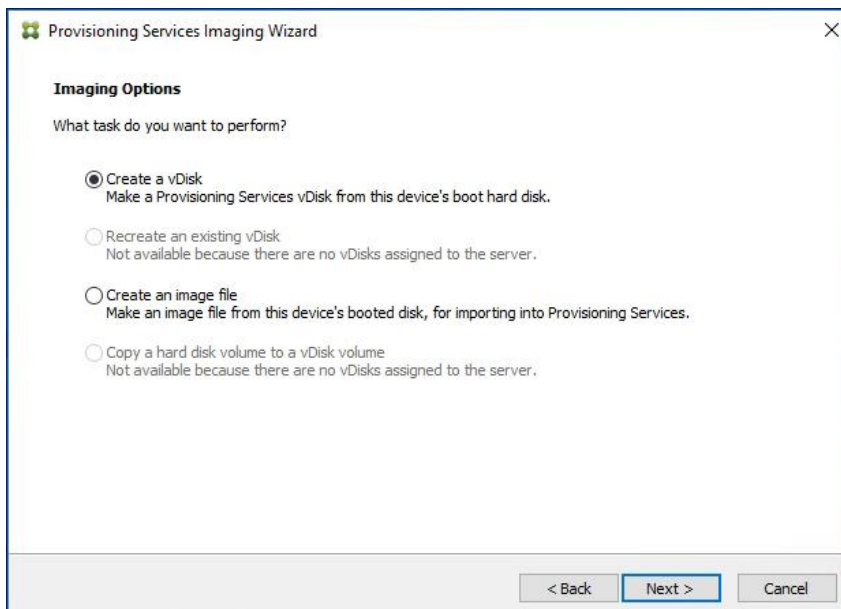


3. The Connect to Farm page appears. Enter the name or IP address of a Provisioning Services Server within the farm to connect to and the port to use to make that connection.
4. Use the Windows credentials (default) or enter different credentials.
5. Click Next.



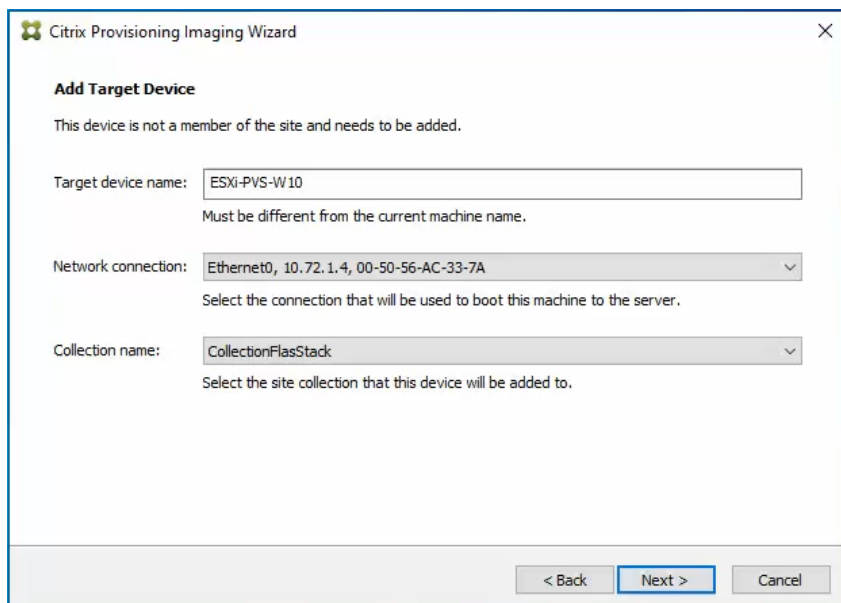
The screenshot shows the 'Provisioning Services Imaging Wizard' window. The title bar reads 'Provisioning Services Imaging Wizard'. The main heading is 'Connect to Provisioning Services Site'. Below this, there is a sub-heading 'Enter Server Details' and a text box for 'Server name or IP:' containing '10.72.0.12'. A 'Port:' spinner control is set to '54321'. Under 'Provide Logon Credentials for the Server', the 'Use my Windows credentials' radio button is selected. Below it are three empty text boxes for 'User name:', 'Domain:', and 'Password:'. At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

6. Select Create a vDisk.
7. Click Next.



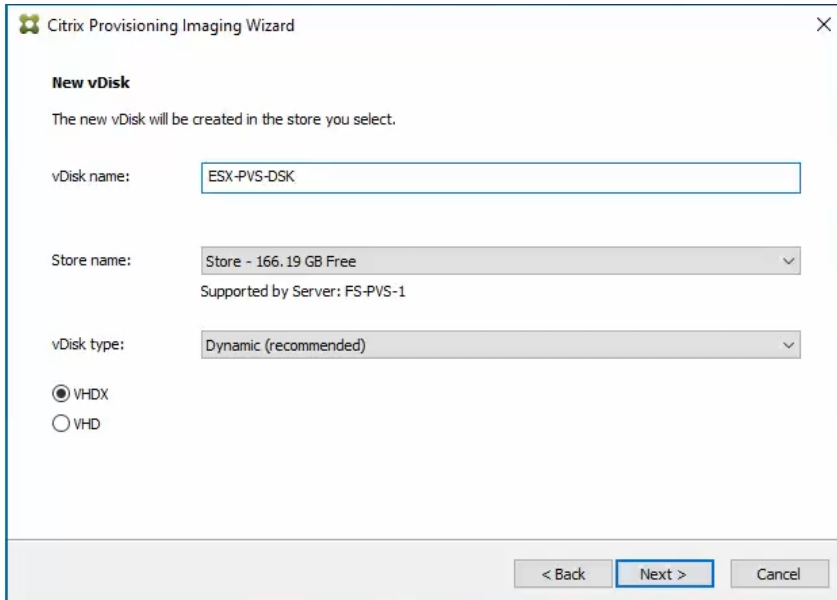
The screenshot shows the 'Provisioning Services Imaging Wizard' window. The title bar reads 'Provisioning Services Imaging Wizard'. The main heading is 'Imaging Options'. Below this, there is a sub-heading 'What task do you want to perform?'. There are four radio button options: 'Create a vDisk' (selected), 'Recreate an existing vDisk', 'Create an image file', and 'Copy a hard disk volume to a vDisk volume'. The 'Recreate an existing vDisk' and 'Copy a hard disk volume to a vDisk volume' options have a note below them: 'Not available because there are no vDisks assigned to the server.' At the bottom right, there are three buttons: '< Back', 'Next >', and 'Cancel'.

8. The Add Target Device page appears.
9. 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.
10. Click Next.




The screenshot shows the 'Add Target Device' dialog box in the Citrix Provisioning Imaging Wizard. The dialog has a title bar with the Citrix logo and the text 'Citrix Provisioning Imaging Wizard'. Below the title bar, the text 'Add Target Device' is displayed. A message states: 'This device is not a member of the site and needs to be added.' There are three input fields: 'Target device name:' with the value 'ESXi-PVS-W10' and a note 'Must be different from the current machine name.'; 'Network connection:' with a dropdown menu showing 'Ethernet0, 10.72.1.4, 00-50-56-AC-33-7A' and a note 'Select the connection that will be used to boot this machine to the server.'; and 'Collection name:' with a dropdown menu showing 'CollectionFlasStack' and a note 'Select the site collection that this device will be added to.' At the bottom of the dialog, there are three buttons: '< Back', 'Next >', and 'Cancel'.

11. The New vDisk dialog displays. Enter the name of the vDisk.
 12. 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.
13. Click Next.




14. 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.

15. Click Next.



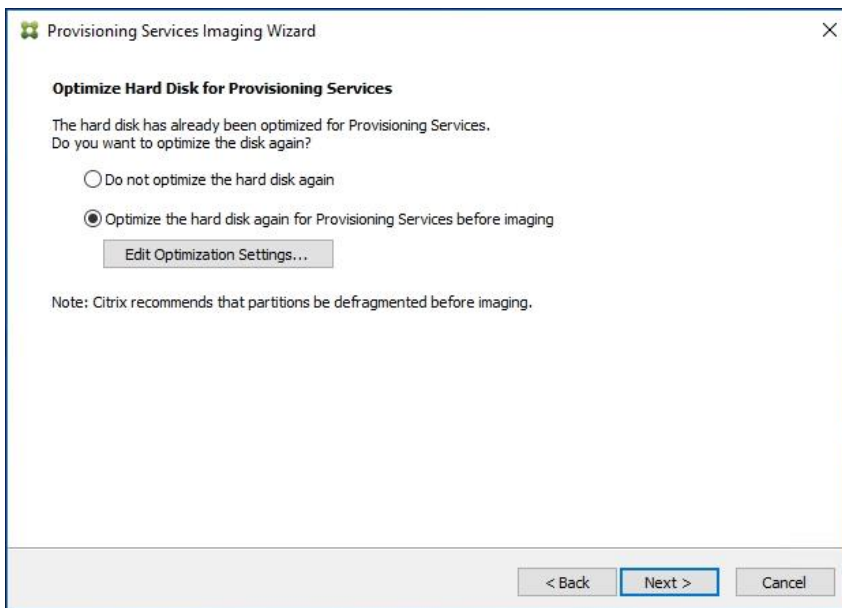
16. Select Image entire boot disk on the Configure Image Volumes page.

17. Click Next.

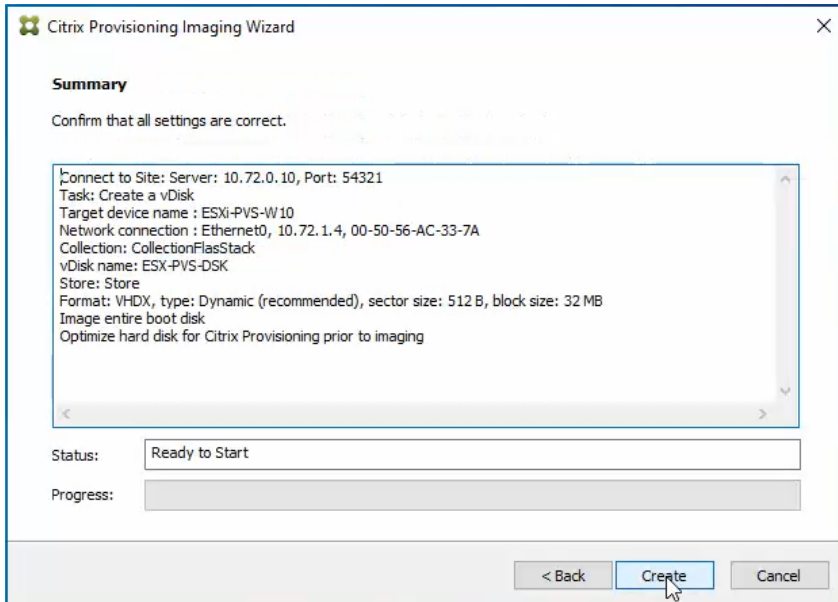


18. Select Optimize for hard disk again for Provisioning Services before imaging on the Optimize Hard Disk for Provisioning Services.

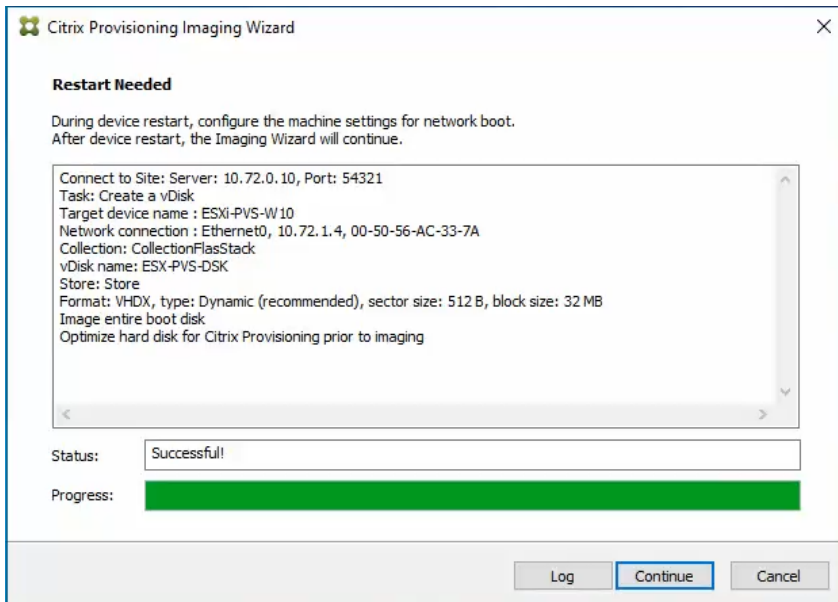
19. Click Next.




20. Click Create on the Summary page.



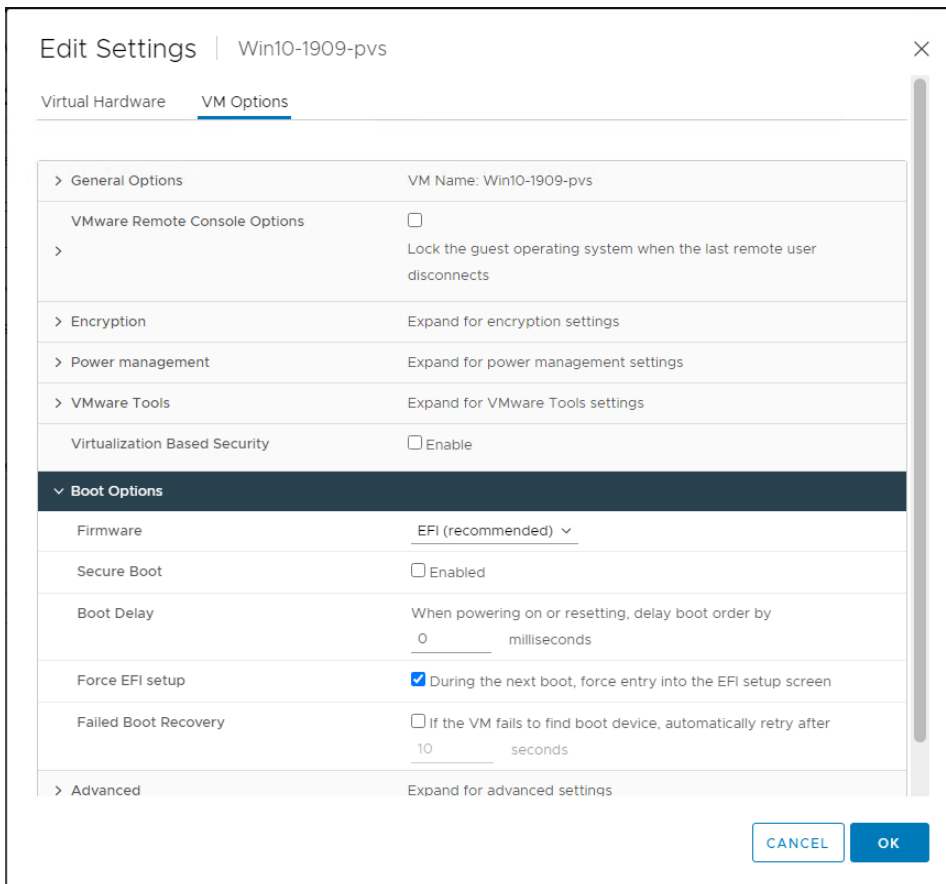
21. Review the configuration and click Continue.



22. When prompted, click No to shut down the machine.

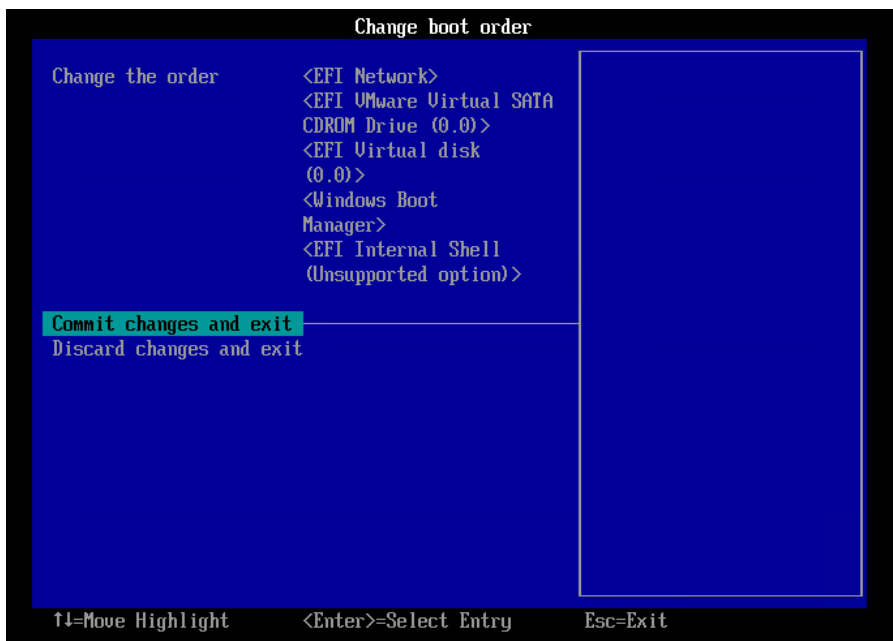


23. Edit the VM settings and select Force EFI Setup under Boot Options.



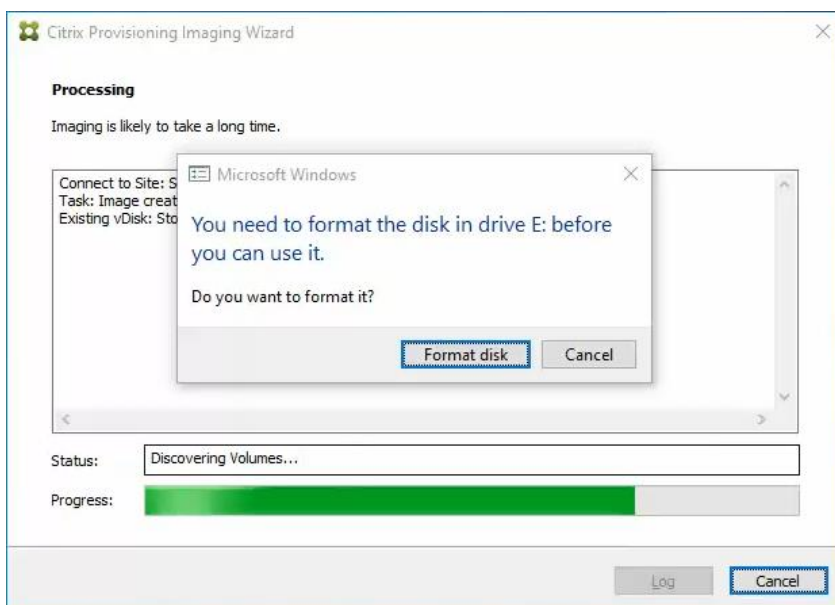
24. Configure the VM settings for EFI network boot.

25. Click Commit changes and exit.




26. 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.

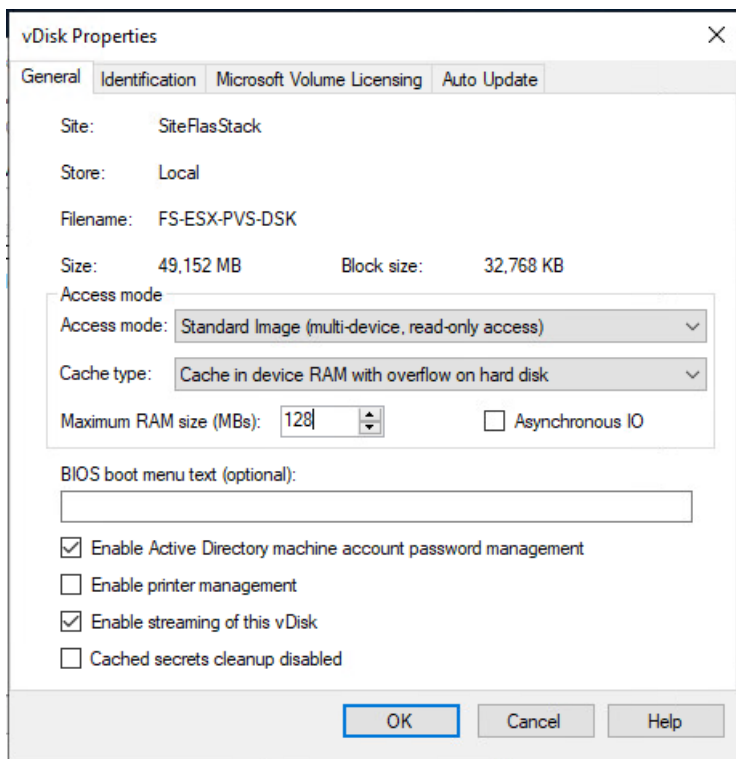
Note: If prompted to Format disk, disregard the message and allow Provisioning Imaging Wizard to finish.



27. A message is displayed when the conversion is complete, click Done.



28. Shutdown the virtual machine used as the VDI or RDS master target.
29. Connect to the PVS server and validate that the vDisk image is available in the Store.
30. Right-click the newly created vDisk and select Properties.
31. On the vDisk Properties dialog, change Access mode to “Standard Image (multi-device, read-only access).”
32. Set the Cache Type to “Cache in device RAM with overflow on hard disk.”
33. Set Maximum RAM size (MBs): 128.
34. Click OK.

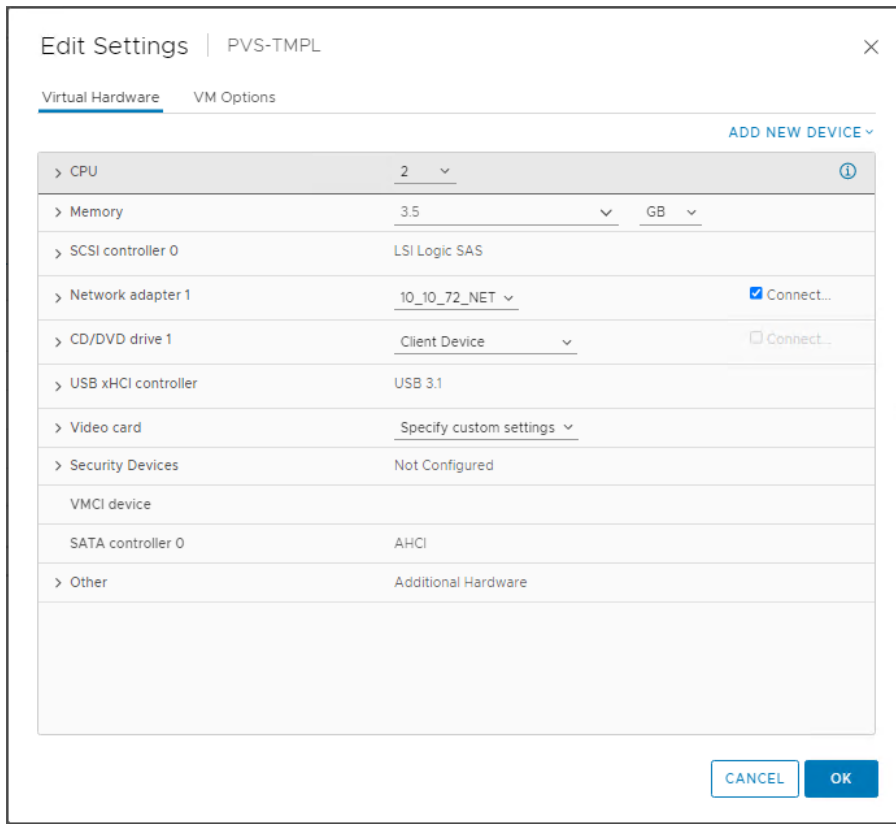


Provision Virtual Desktop Machines

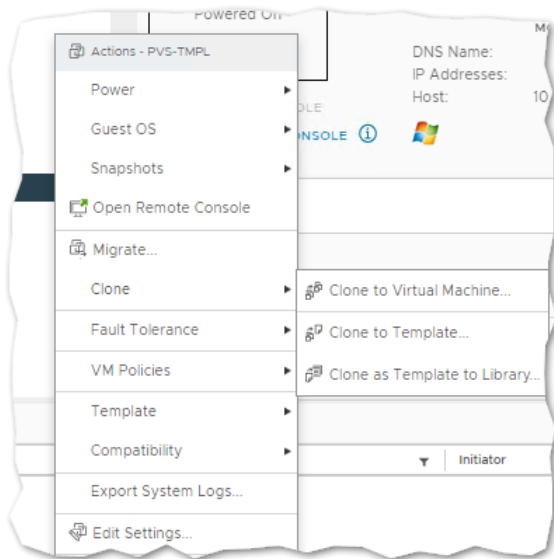
Citrix Provisioning Services Citrix Virtual Desktop Setup Wizard

To create PVS streamed virtual desktop machines, follow these steps:

1. Create a Master Target Virtual Machine:



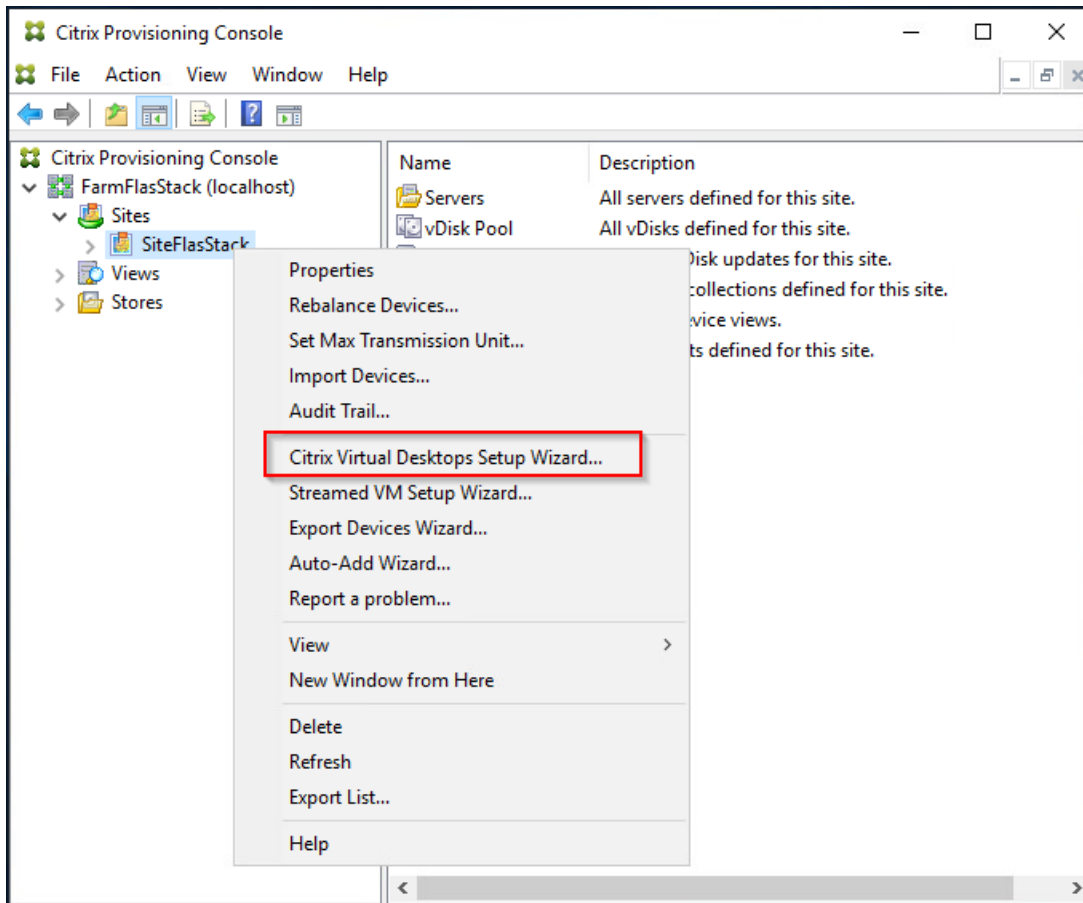
2. Right-click and clone the Master Target VM to the Template.



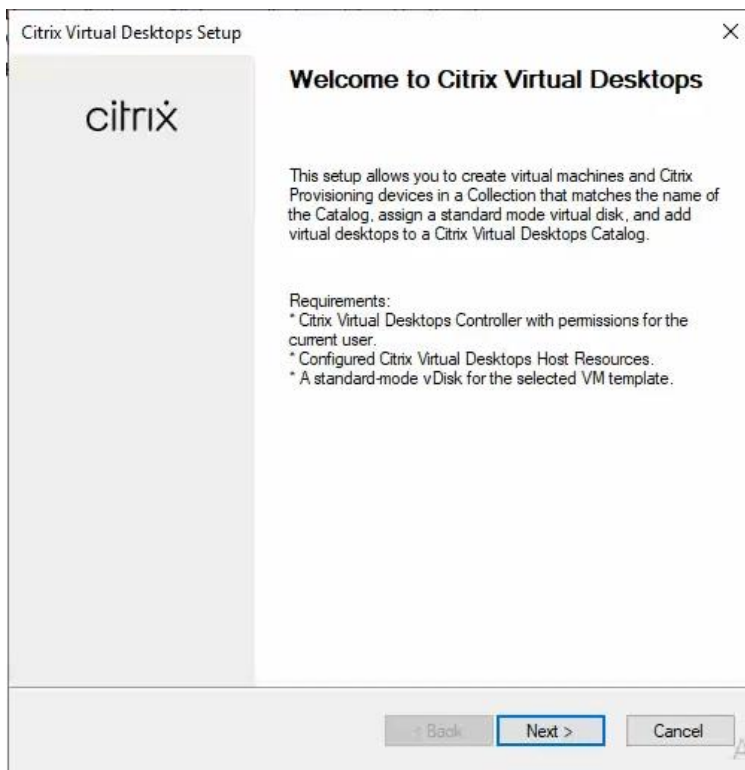
3. Start the Citrix Virtual Apps and Desktops Setup Wizard from the Provisioning Services Console.

4. Right-click the Site.

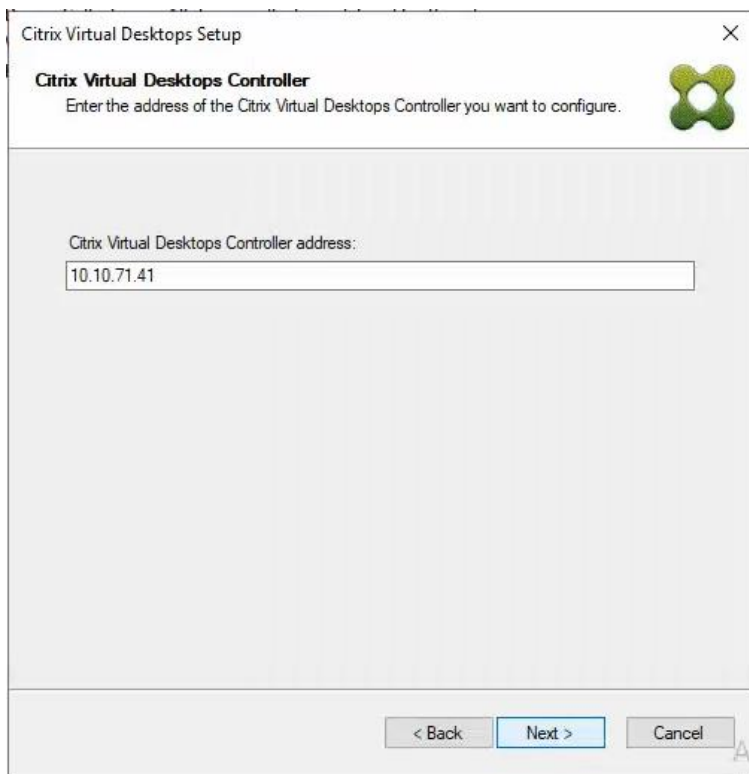
5. Select Citrix Virtual Desktop Setup Wizard... from the context menu.



6. Click Next.

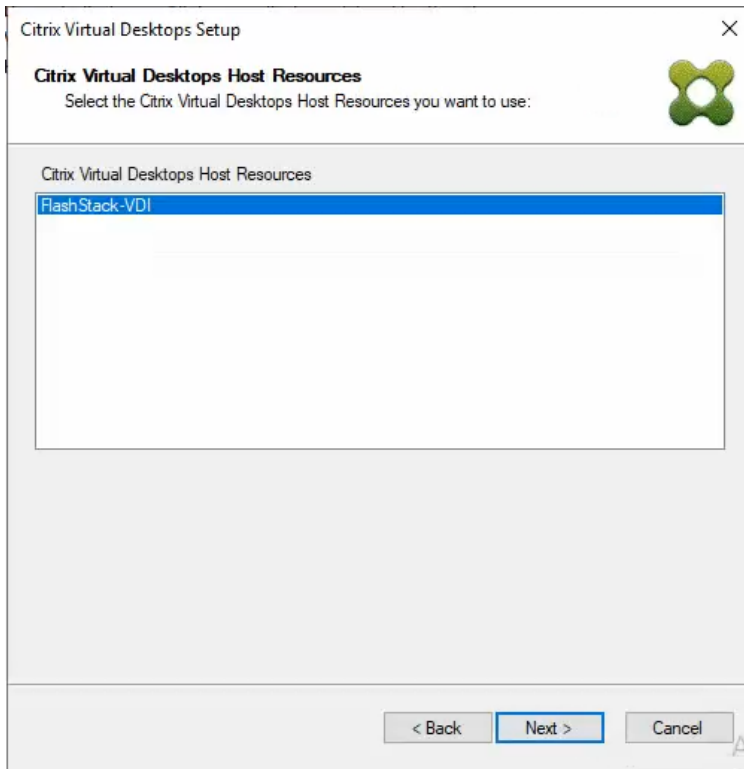


7. Enter the address of the Citrix Virtual Desktop Controller that will be used for the wizard operations.
8. Click Next.



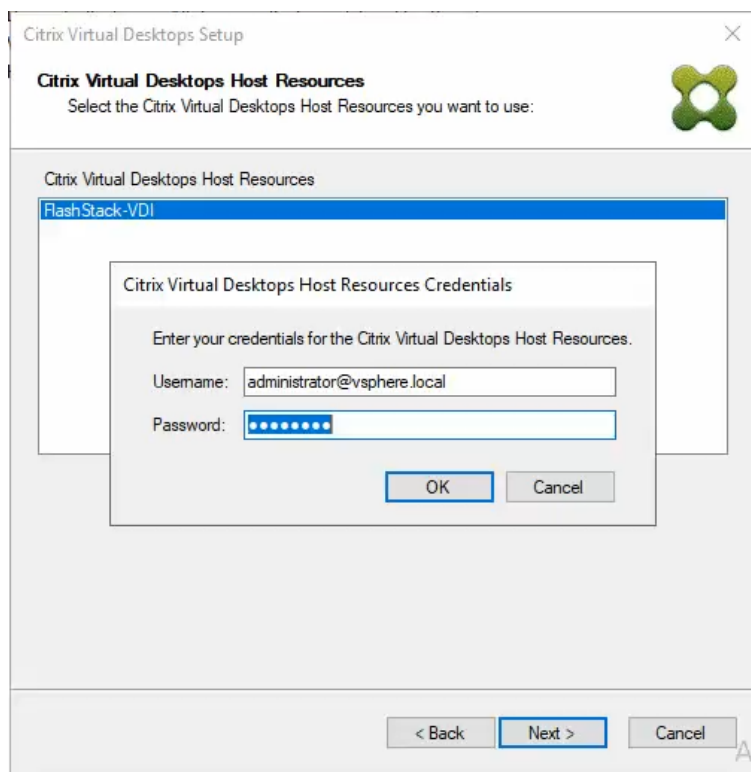
9. Select Host Resources that will be used for the wizard operations

10. Click Next.



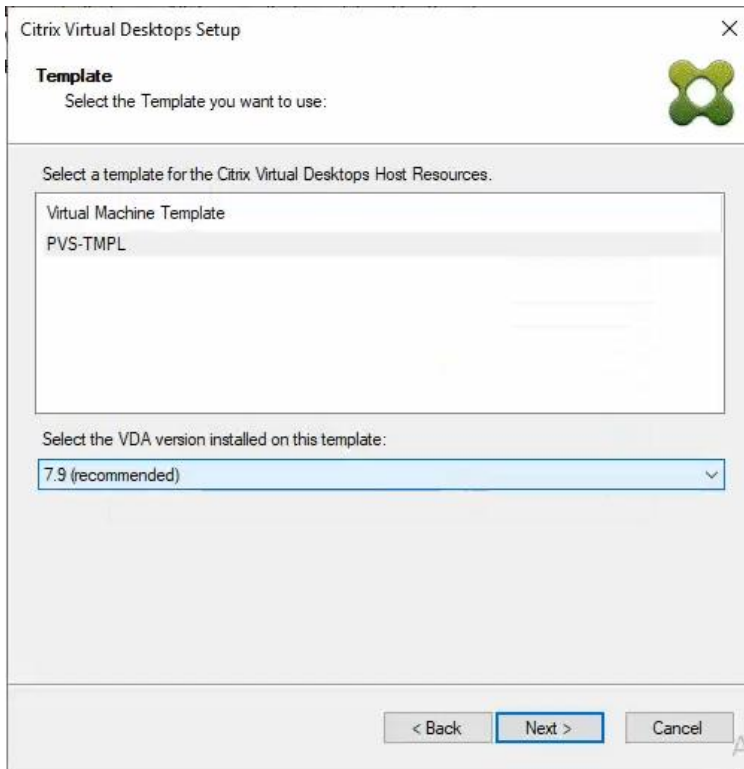
11. Provide Citrix Virtual Desktop Controller credentials.

12. Click Ok

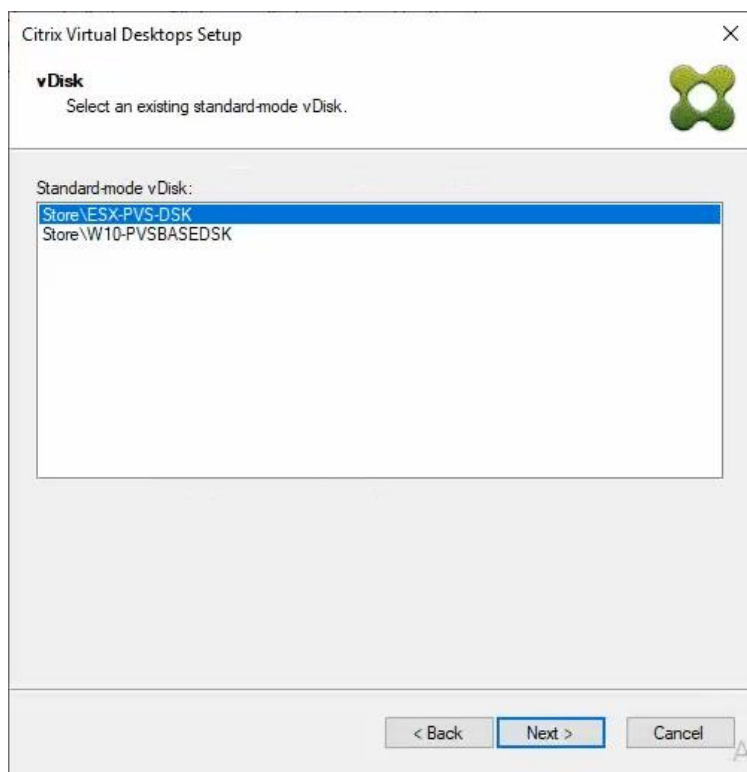


13. Select the Template created earlier.

14. Click Next.



15. Select the virtual disk (vDisk) that will be used to stream the provisioned virtual machines.
16. Click Next.



17. Select Create new catalog.

18. Provide a catalog name.

19. Click Next.

Citrix Virtual Desktops Setup

Catalog
Select your Catalog preferences.

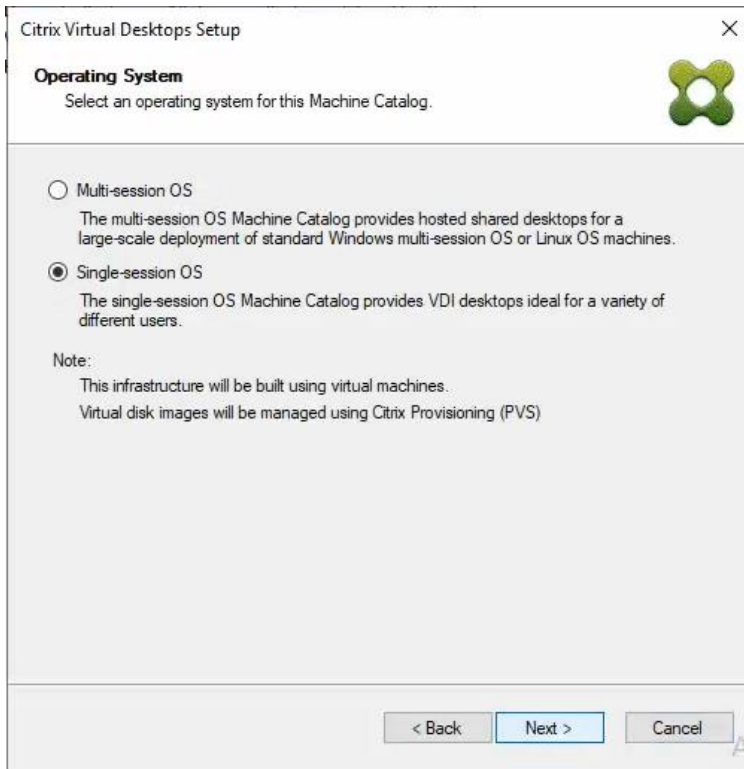
Create a new catalog
 Use an existing catalog

Catalog name:
Description:

< Back Next > Cancel

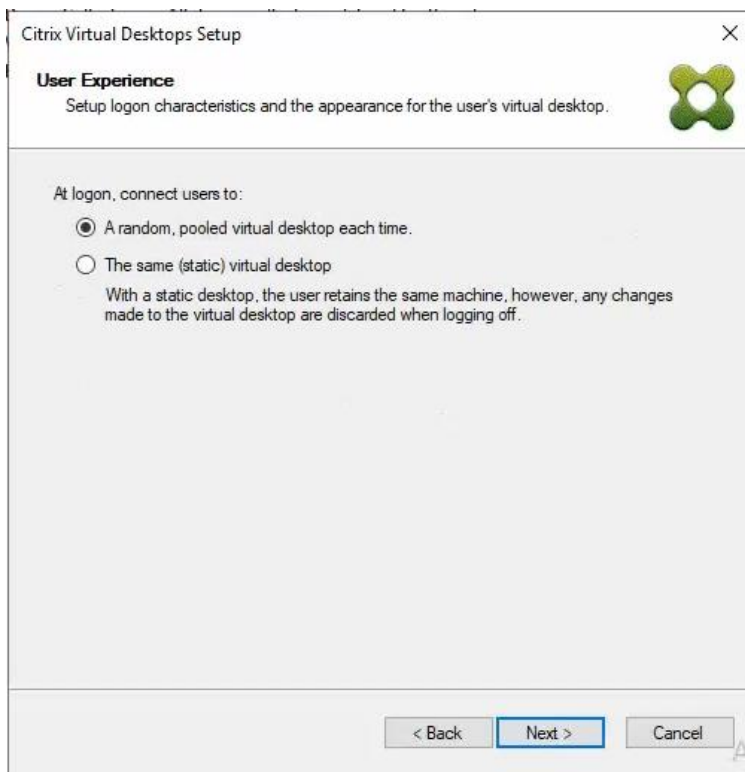
20. Select Single-session OS for Machine catalog Operating System.

21. Click Next.



22. Select random for the User Experience.

23. Click Next.



24. On the Virtual machines dialog, specify the following:

- The number of virtual machines to create.

Note: It is recommended to create 200 or less per provisioning run. Create a single virtual machine at first to verify the procedure.

- 2 as Number of vCPUs for the virtual machine
- 3584 MB as the amount of memory for the virtual machine
- 6GB as the Local write cache disk.

25. Click Next.

Citrix Virtual Desktops Setup

Virtual machines
Select your virtual machine preferences.

Number of virtual machines to create: 1960

vCPUs: 2

Memory: 3584 MB

Local write cache disk: 6 GB Thick

Boot mode:

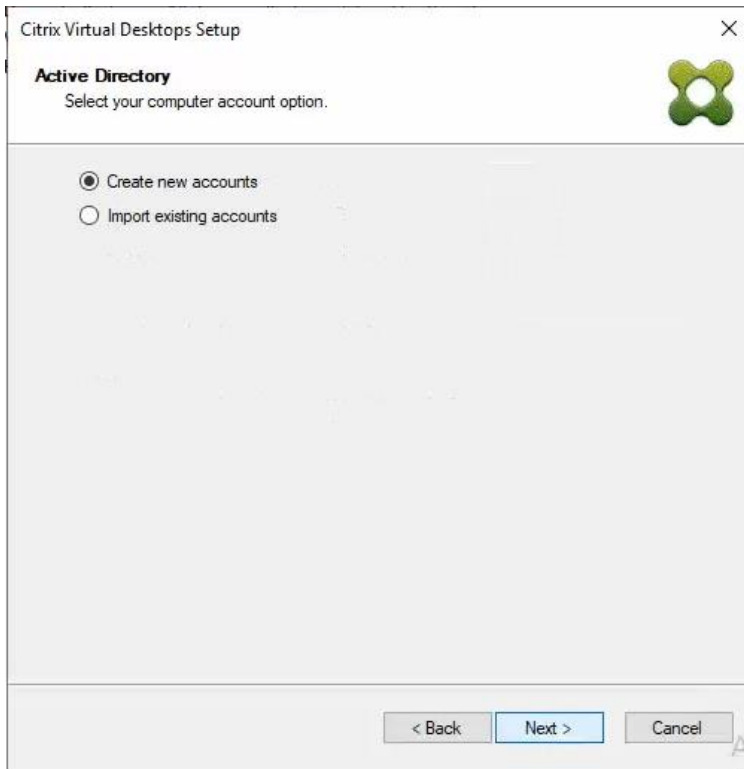
PXE boot (requires a running PXE service)

BDM disk (create a boot device manager partition)

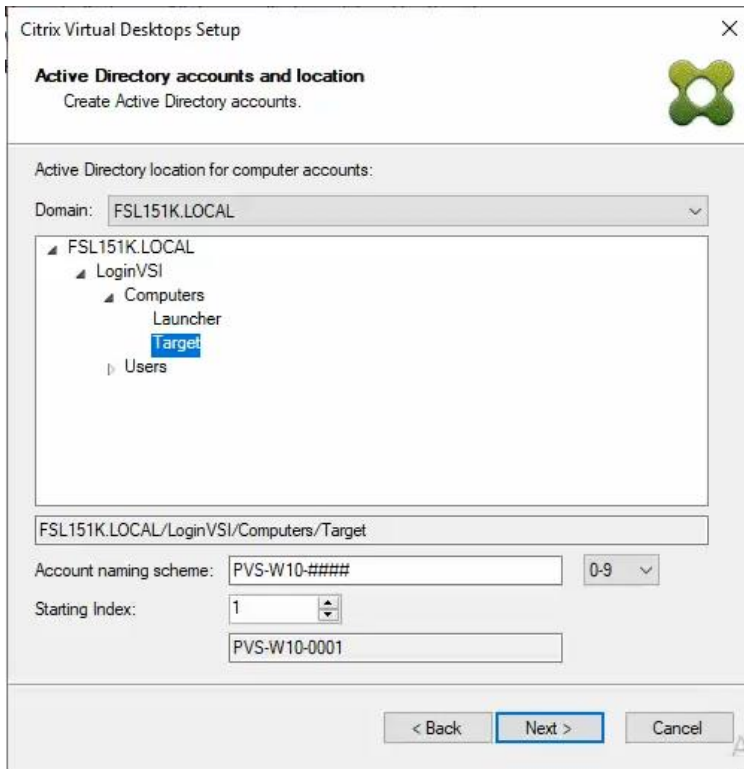
< Back Next > Cancel

26. Select the Create new accounts.

27. Click Next.

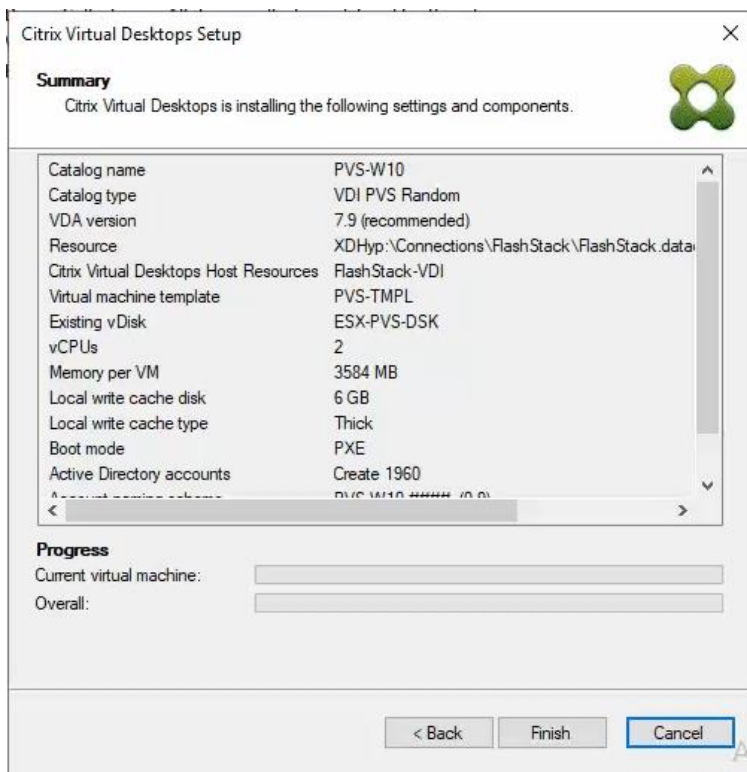


28. Specify the Active Directory Accounts and Location. This is where the wizard should create computer accounts.
29. Provide the Account naming scheme. An example name is shown in the text box below the naming scheme selection location.
30. Click Next.



31. Verify the information on the Summary screen.

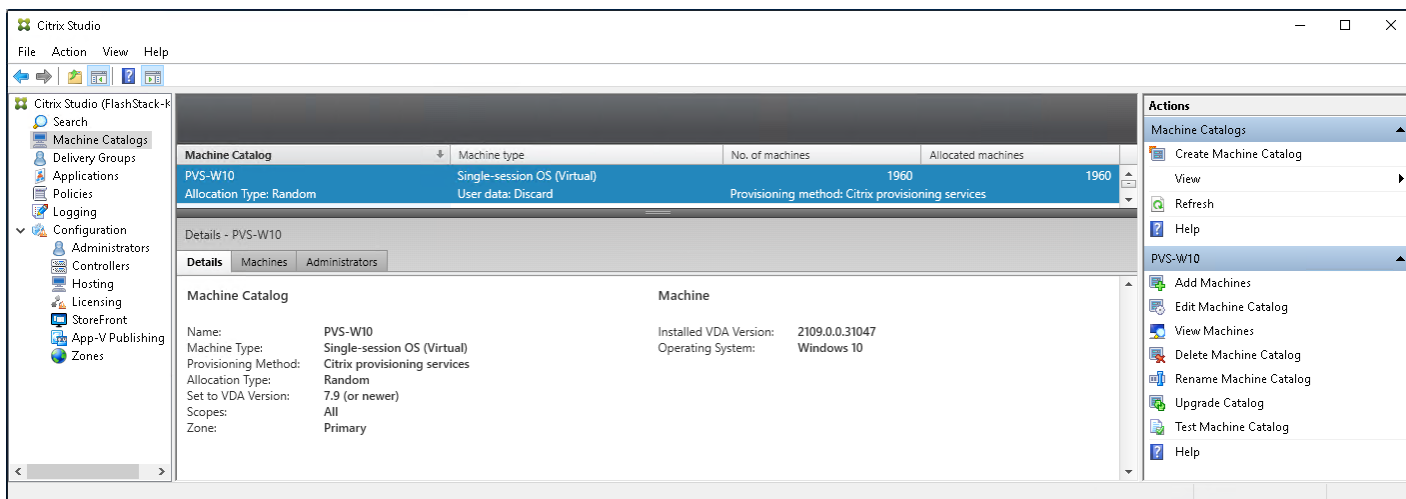
32. Click Finish to begin the virtual machine creation.



33. When the wizard is done provisioning the virtual machines, click Done.

34. When the wizard is done provisioning the virtual machines, verify the Machine Catalog on the Citrix Virtual Apps and Desktops Controller:

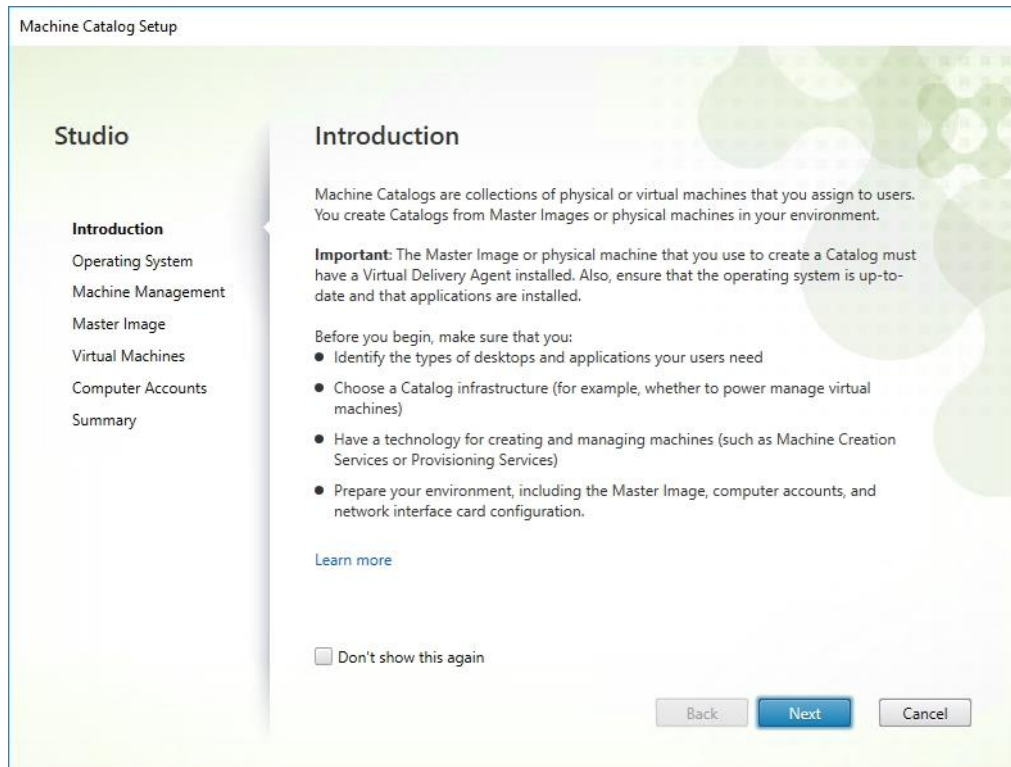
- Connect to a Citrix Virtual Apps and Desktops server and launch Citrix Studio.
- Select Machine Catalogs in the Studio navigation pane.
- Select a machine catalog.



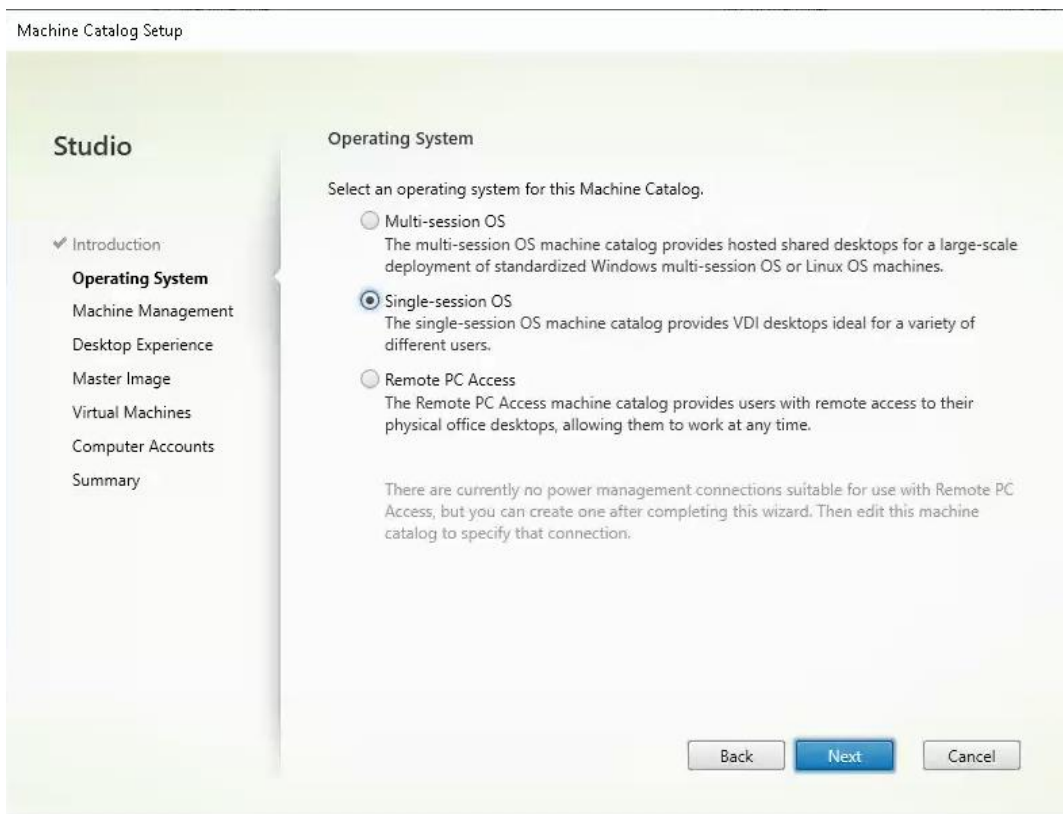
Citrix Machine Creation Services

To configure the Machine Catalog Setup, follow these steps:

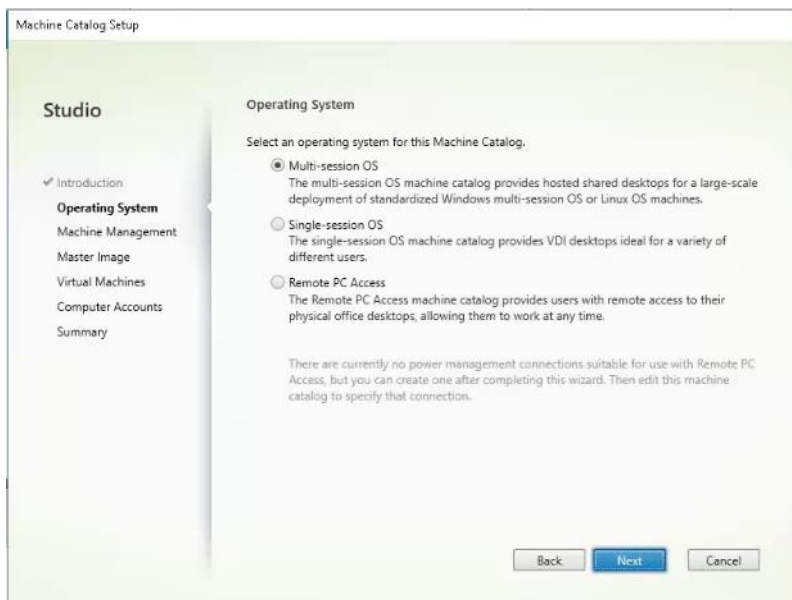
1. Connect to a Citrix Virtual Apps and Desktops server and launch Citrix Studio.
2. Choose Create Machine Catalog from the Actions pane.
3. Click Next.



4. Select Single-session OS.
5. Click Next.

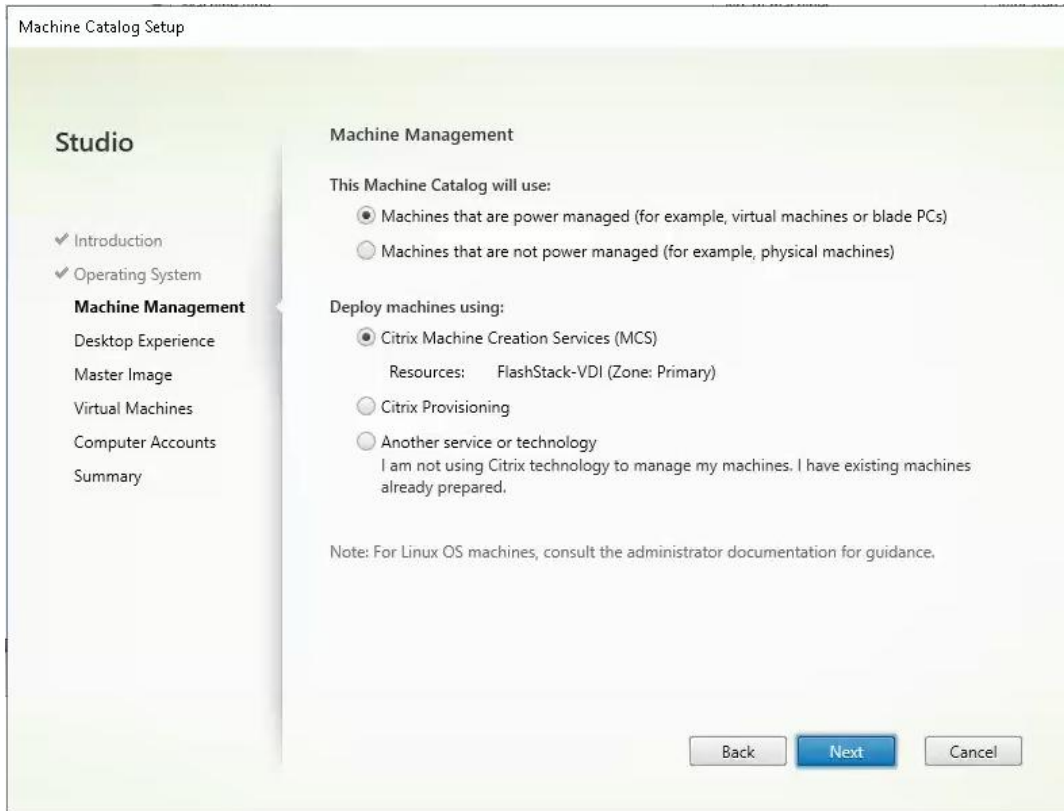


6. Select Multi-session OS when using Windows Server 2019 desktops.



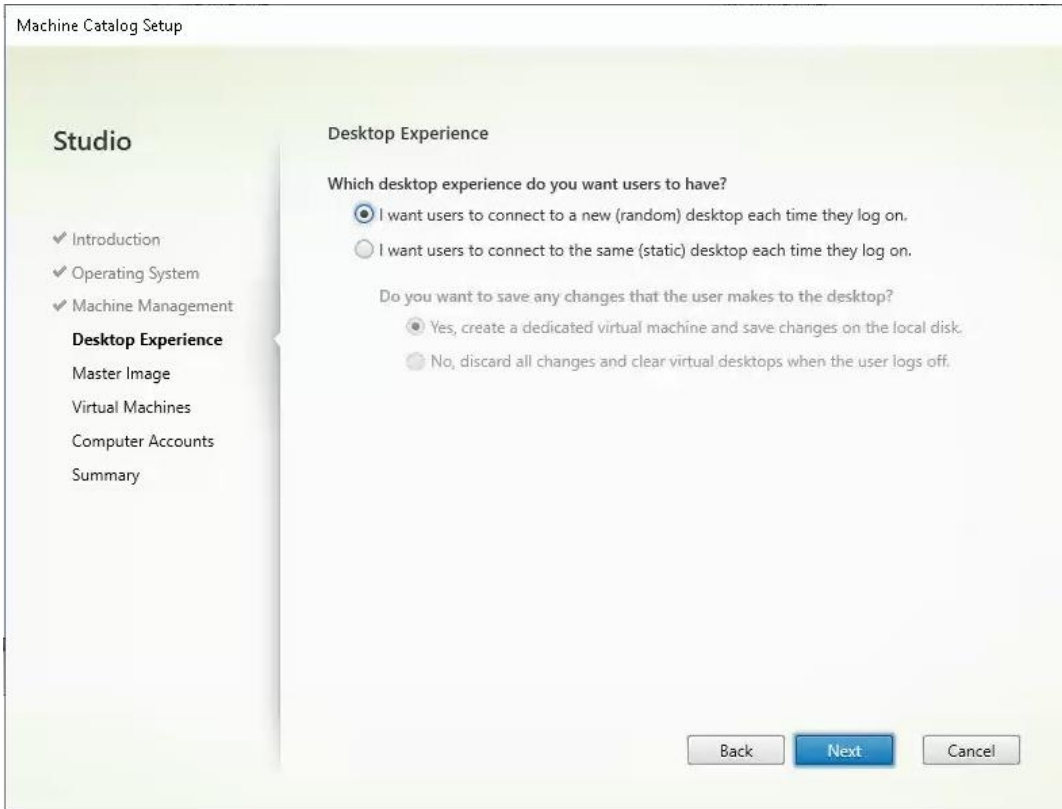
7. Select the appropriate machine management.

8. Click Next.



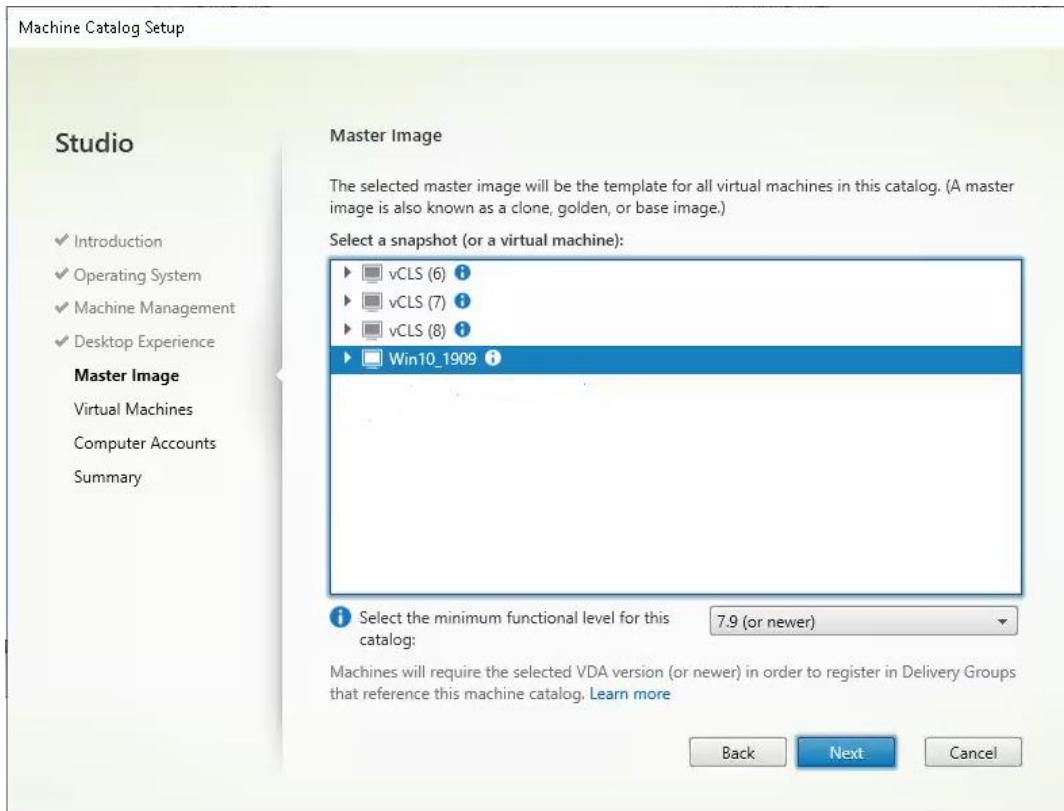
9. Select (random) for Desktop Experience.

10. Click Next.



11. Select a Virtual Machine to be used for Catalog Master Image.

12. Click Next.



13. Specify the number of desktops to create and machine configuration.

14. Set amount of memory (MB) to be used by virtual desktops.

15. Select Full Copy for machine copy mode.

16. 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?

1960 - +

Configure your machines.

Total memory (MB) on each machine: 3584 - +

Configure a cache for temporary data on each machine.

Memory allocated to cache (MB): 256 - +

Disk cache size (GB): 10 - +

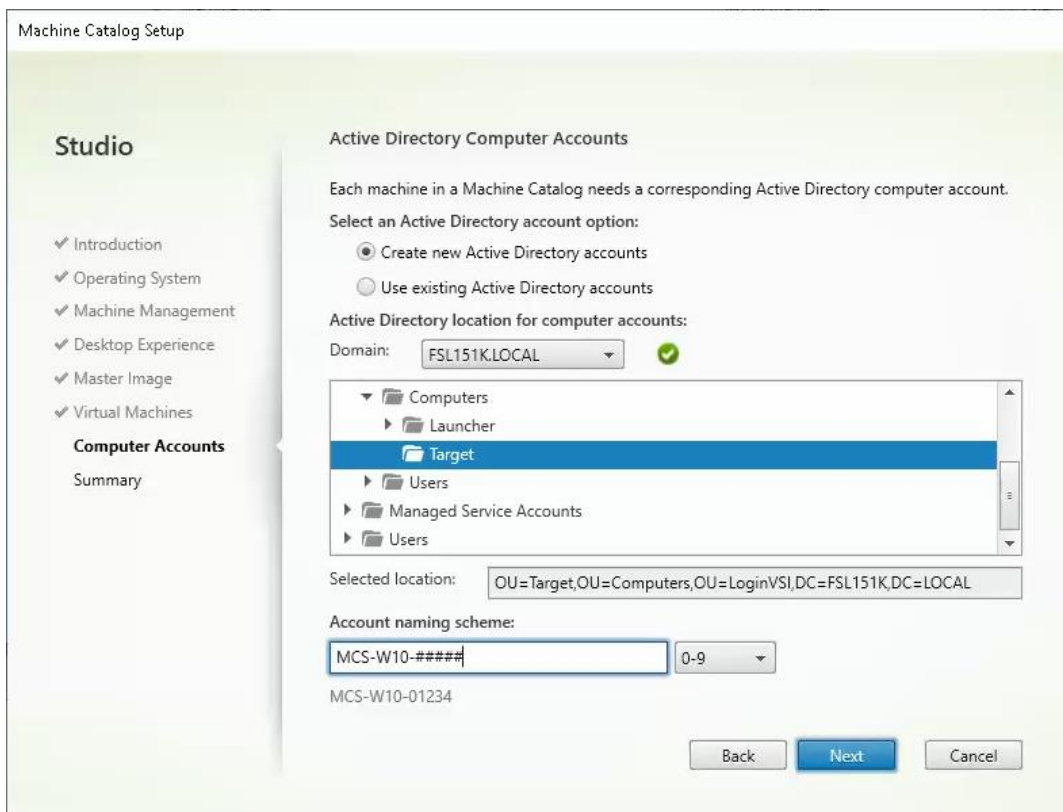
By default, both check boxes are cleared. (Temporary data is written to OS storage for each VM.) To cache temporary data, a current MCSIO driver must be installed on the VM, in addition to selecting one or both check boxes and values above.

[Learn more](#)

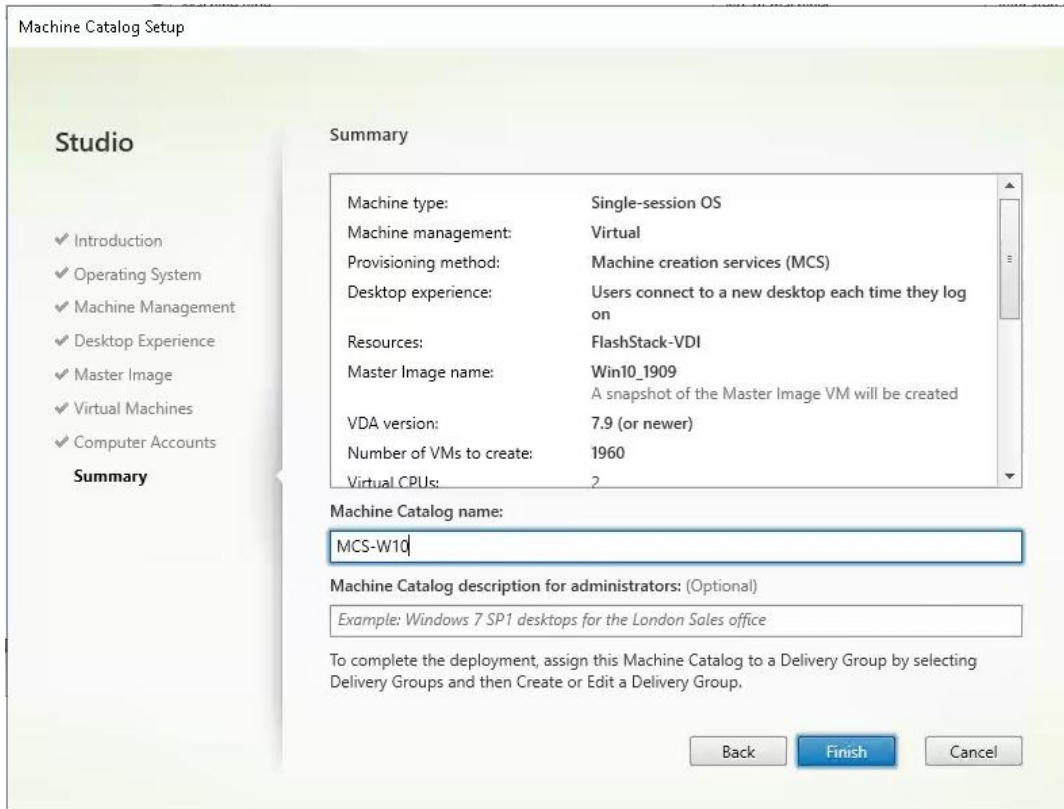
Back Next Cancel

17. Specify the AD account naming scheme and OU where accounts will be created.

18. Click Next.



19. On the Summary page specify Catalog name and click Finish to start the deployment.



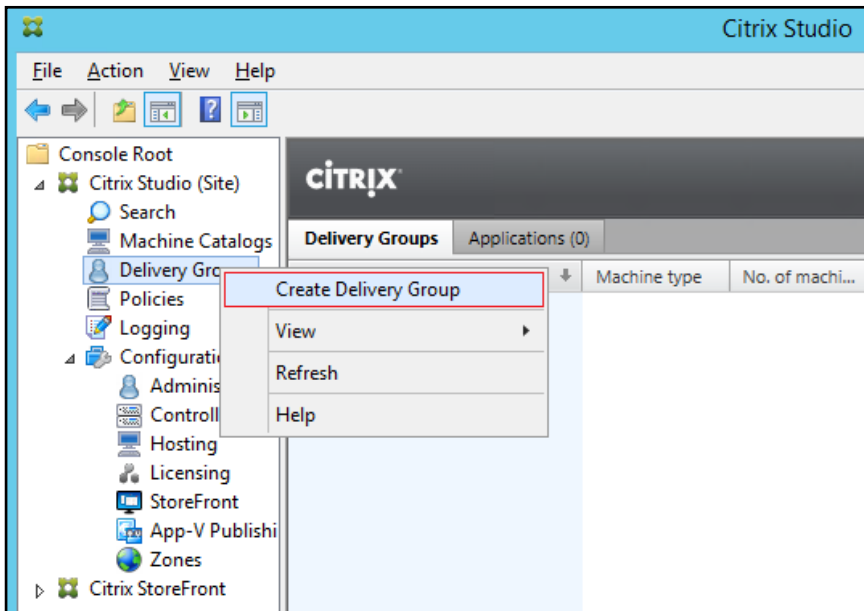
Create Delivery Groups

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.

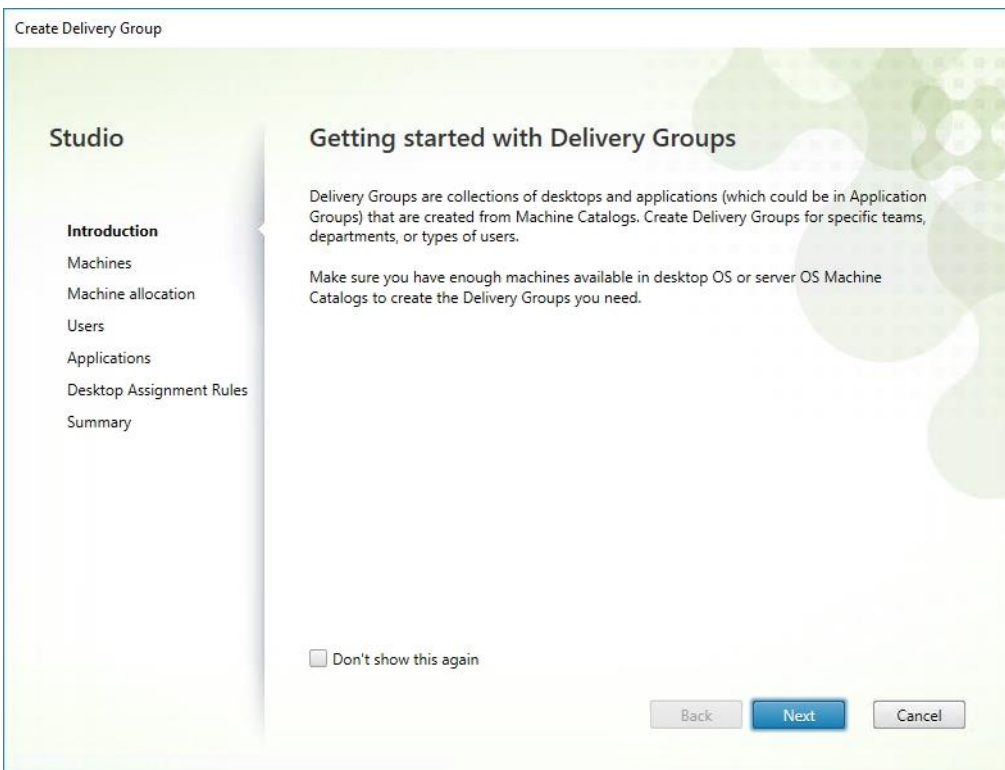
To create delivery groups, follow these steps:

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.

1. Connect to a Citrix Virtual Apps and Desktops server and launch Citrix Studio.
2. Choose Create Delivery Group from the drop-down list.



3. Click Next.



4. Specify the Machine Catalog and increment the number of machines to add.

5. Click Next.

Create Delivery Group

Studio

- Introduction
- Machines**
- Delivery Type
- Users
- Desktop Assignment Rules
- Summary

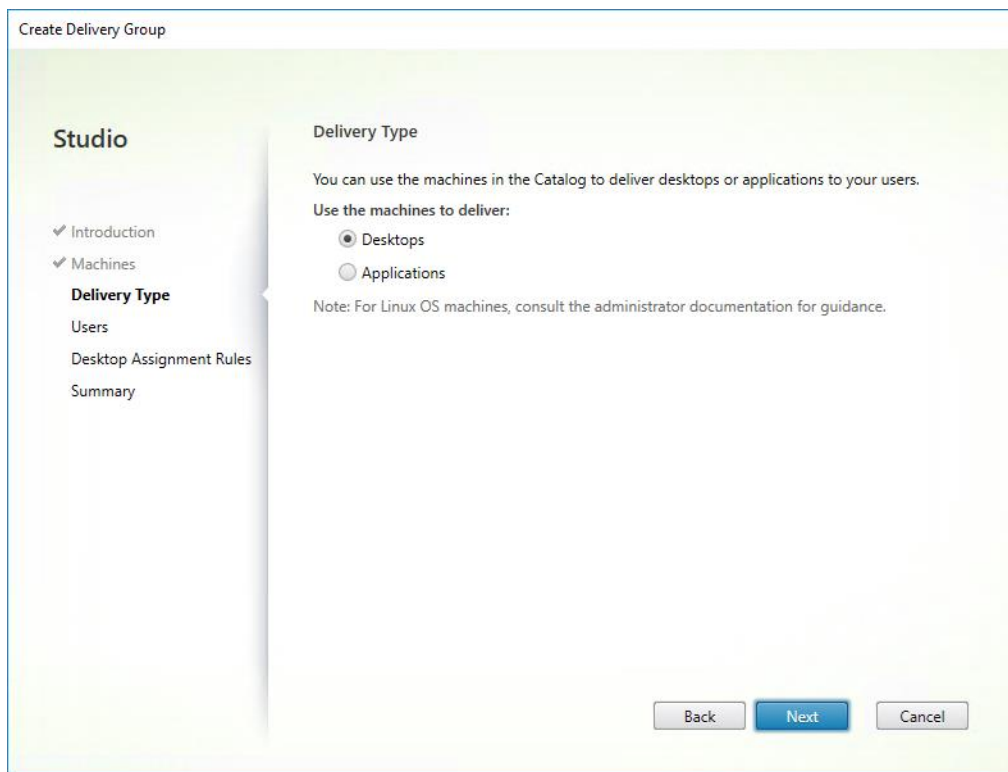
Machines

Select a Machine Catalog.

Catalog	Type	Machines
<input checked="" type="radio"/> WIN10-FS-MCS MCS WIN10 1809	VDI MCS Static Local Disk	210

Choose the number of machines for this Delivery Group: - +

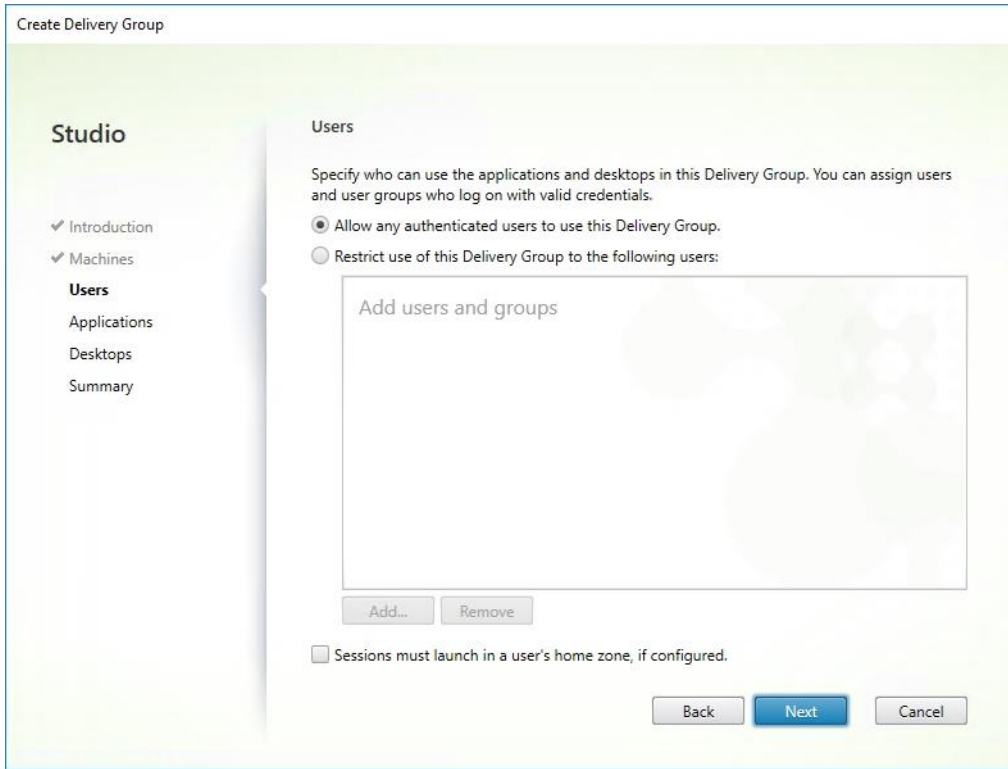
- Specify what the machines in the catalog will deliver: Desktops, Desktops and Applications, or Applications.
- Select Desktops.
- Click Next.



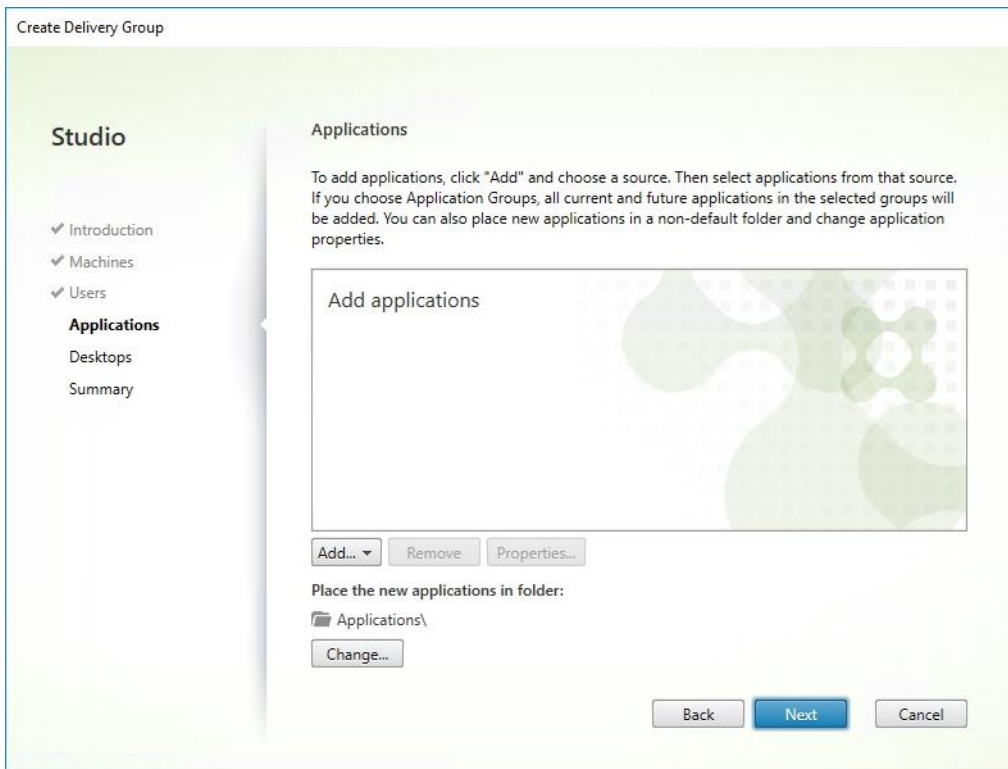
9. To make the Delivery Group accessible, you must add users. Select Allow any authenticated users to use this Delivery Group.

Note: User assignment can be updated any time after Delivery group creation by accessing Delivery group properties in Desktop Studio.

10. Click Next.



11. Click Next (no applications are used in this design).



12. Enable Users to access the desktops.

13. Click Next.

Display name:

Description:

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:

Enable desktop assignment rule
Clear this check box to disable delivery of this desktop.

14. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Description (Optional).

15. Click Finish.

Create Delivery Group

The screenshot shows the 'Summary' page of the 'Create Delivery Group' wizard in Citrix Studio. On the left, a 'Studio' sidebar lists navigation options: Introduction, Machines, Delivery Type, Users, Desktop Assignment Rules, and Summary (which is selected). The main area displays the following configuration:

Machine Catalog:	WIN10-FS-MCS
Machine type:	Desktop OS
Allocation type:	Static
Machines added:	VCCFSLAB\w10-mcs-001 VCCFSLAB\w10-mcs-002 VCCFSLAB\w10-mcs-003 VCCFSLAB\w10-mcs-004 VCCFSLAB\w10-mcs-005 VCCFSLAB\w10-mcs-006 VCCFSLAB\w10-mcs-007 VCCFSLAB\w10-mcs-008 VCCFSLAB\w10-mcs-009 VCCFSLAB\w10-mcs-010 VCCFSLAB\w10-mcs-011 VCCFSLAB\w10-mcs-012

Below the table, the 'Delivery Group name' field contains 'WIN10-DG-MCS-STATIC'. The 'Delivery Group description, used as label in Receiver (optional)' field is empty. At the bottom, there are 'Back', 'Finish', and 'Cancel' buttons.

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.

16. From the drop-down list, select “Turn on Maintenance Mode.”

Citrix Virtual Apps, 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 38](#) shows the policies for Login VSI testing in this CVD.

Figure 38. Citrix Virtual Apps and Desktops Policy

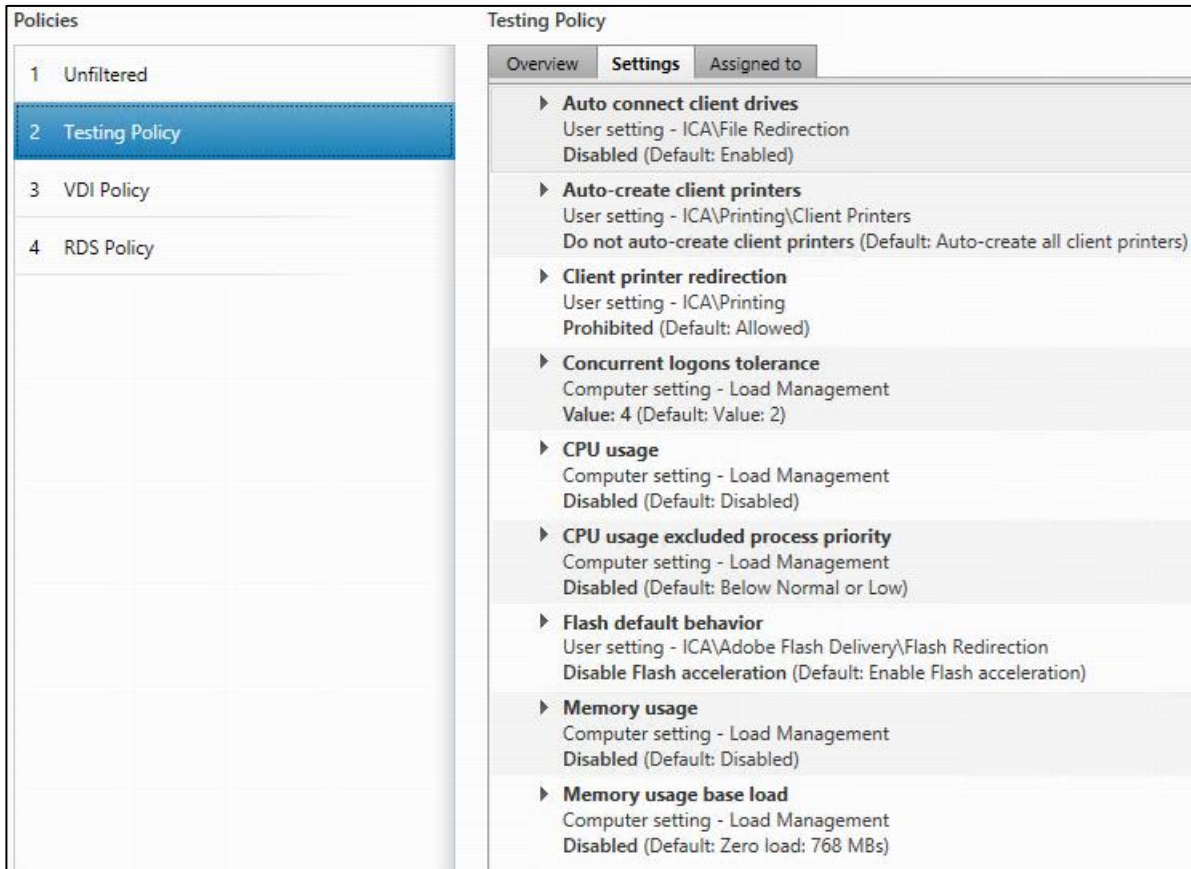
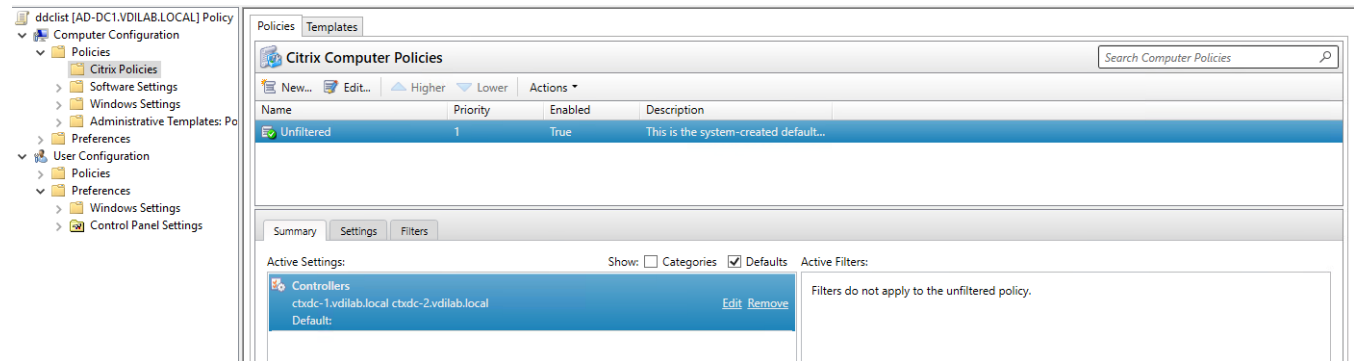


Figure 39. Delivery Controllers Policy



Configure FSLogix

A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Profile management in VDI environments is an integral part of the user experience.

Note: FSLogix, a Microsoft tool, was used to manage user profiles in this validated design.

FSLogix allows you to:


- Roam user data between remote computing session hosts
- Minimize sign in times for virtual desktop environments
- Optimize file IO between host/client and remote profile store
- Provide a local profile experience, eliminating the need for roaming profiles.
- Simplify the management of applications and 'Gold Images'

Additional documentation about the tool can be found [here](#).

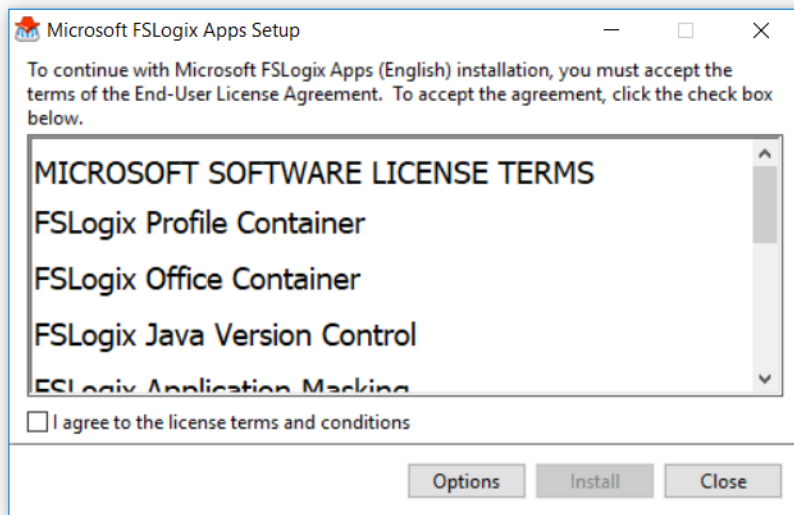
FSLogix Apps Installation

To install the FSLogix Apps, follow these steps:

1. FSLogix download file [here](#).
2. Run FSLogixAppSetup.exe on VDI master image (32 bit or 64 bit depending on your environment).
3. Click OK to proceed with default installation folder.



4. Review and accept the license agreement.
5. Click Install.



6. Reboot.

Configure Profile Container Group Policy

1. Copy "fslogix.admx" to C:\Windows\PolicyDefinitions, and "fslogix.adml" to C:\Windows\PolicyDefinitions\en-US on Active Directory Domain Controllers.
2. Create FSLogix GPO as follows and apply to the desktops OU.
3. Navigate to Computer Configuration > Administrative Templates > FSLogix > Profile Containers.
4. Configure the following settings:
 - Enabled - Enabled
 - VHD location - Enabled, with the path set to \\<FileServer>\<Profiles Directory>

Note: Consider enabling and configuring FSLogix logging as well as limiting the size of the profiles and excluding additional directories.

Figure 40. Example of FSLogix Policy

FSLogix		
Scope	Details	Settings
FSLogix/Profile Containers show		
Policy	Setting	Comment
Delete local profile when FSLogix Profile should apply	Enabled	
Delete local profile when FSLogix Profile should apply		
Enabled		
Policy	Setting	Comment
Dynamic VHD(X) allocation	Enabled	
Dynamic VHD(X) allocation		
Enabled		
Policy	Setting	Comment
Enabled	Enabled	
Enabled		
Enabled		
Policy	Setting	Comment
Profile type	Enabled	
Try for read-write profile and fallback to read-only		
Policy	Setting	Comment
Size in MBs	Enabled	
Size in MBs		
2048		
Policy	Setting	Comment
VHD location	Enabled	
VHD location		
\\purefile\wdi\RDS		
FSLogix/Profile Containers/Advanced show		
FSLogix/Profile Containers/Container and Directory Naming hide		
Policy	Setting	Comment
Virtual disk type	Enabled	
VHDX		

Cisco Intersight Cloud Based Management

[Cisco Intersight](#) is Cisco’s new systems management platform that delivers intuitive computing through cloud-powered intelligence. This platform offers a more intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in ways that were not possible with prior generations of tools. This capability empowers organizations to achieve significant savings in Total Cost of Ownership (TCO) and to deliver applications faster, so they can support new business initiatives. The advantages of the model-based management of the Cisco UCS platform plus Cisco Intersight are extended to Cisco UCS servers.

The Cisco UCS platform uses model-based management to provision servers and the associated storage and fabric automatically, regardless of form factor. Cisco Intersight works in conjunction with Cisco UCS Manager and the Cisco® Integrated Management Controller (IMC). By simply associating a model-based configuration with a resource through service profiles, your IT staff can consistently align policy, server personality, and workloads. These policies can be created once and used by IT staff with minimal effort to deploy servers. The result is improved productivity and compliance and lower risk of failures due to inconsistent configuration.

Cisco Intersight will be integrated with data center, hybrid cloud platforms, and services to securely deploy and manage infrastructure resources across data center and edge environments. In addition,

Cisco will provide future integrations to third-party operations tools to allow customers to use their existing solutions more effectively.

Figure 41. Example of User-Customizable Cisco Intersight Dashboard for FlashStack UCS Domain

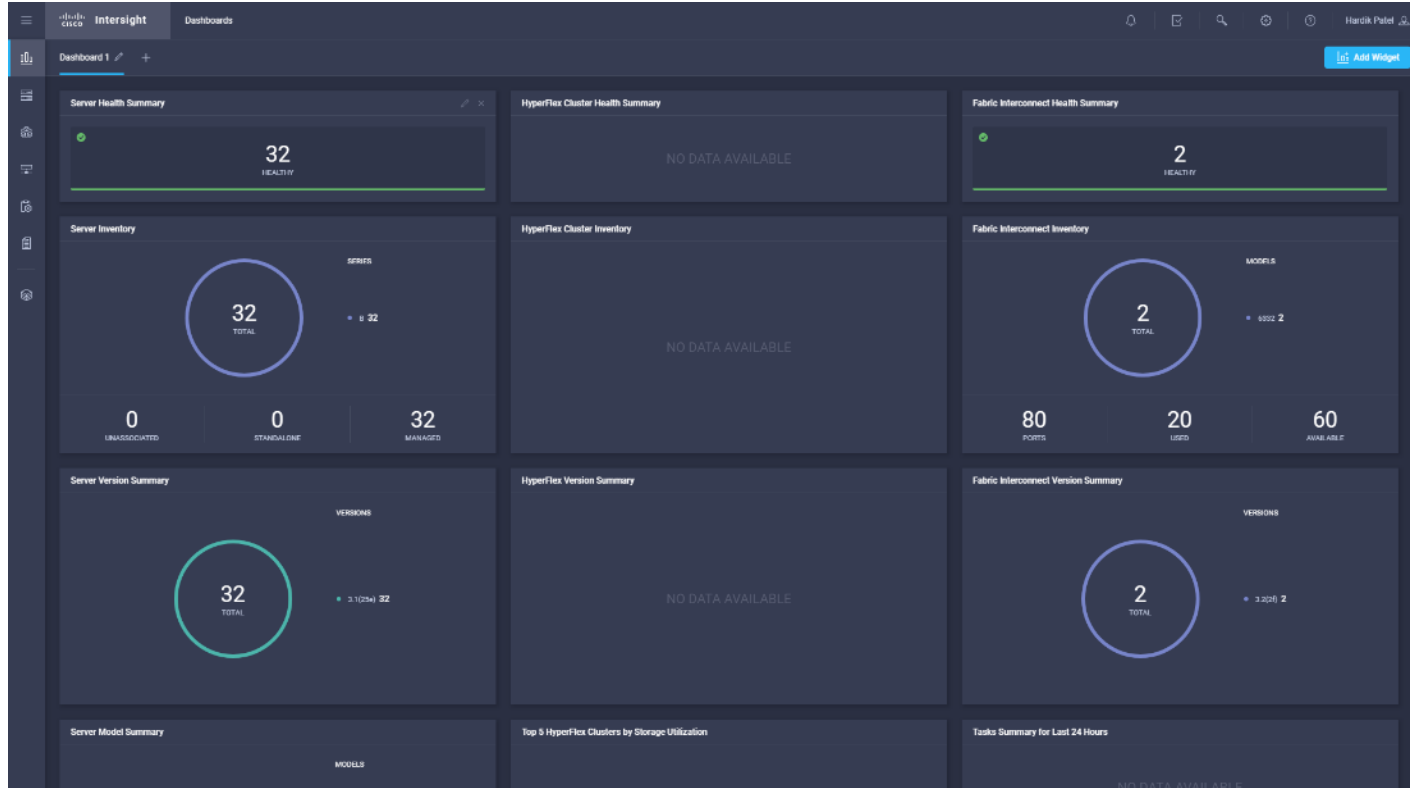


Figure 42. Cisco UCS Manager Device Connector Example

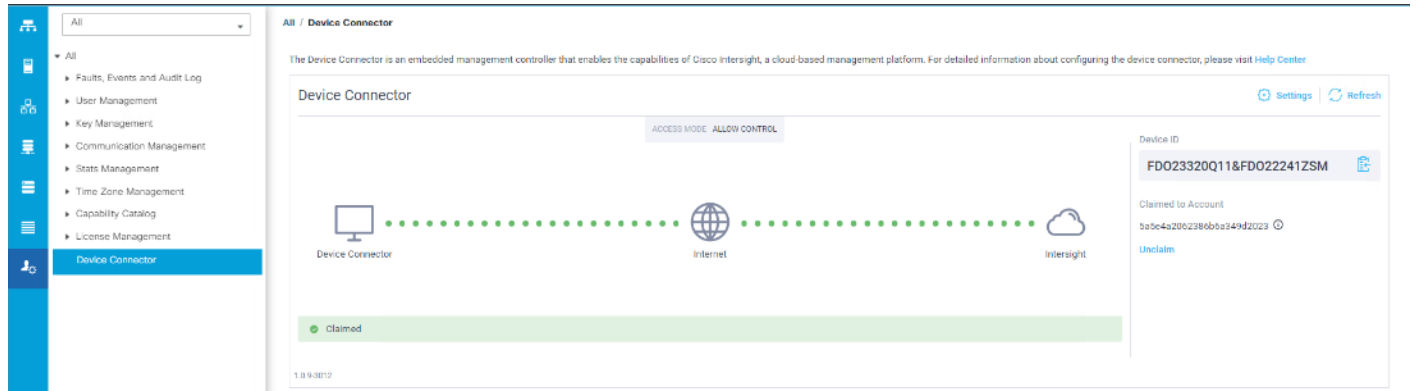


Figure 43. Cisco Intersight License

Cisco Intersight Licensing Tiers - Features i

Base	Essentials	Advantage	Premier	
<ul style="list-style-type: none"> - SaaS only - No cost for UCS and HyperFlex systems - Global monitoring of health and inventory status - User customizable dashboard - Tagging and basic search - Context launch of element managers (UCS Manager, IMC, HyperFlex Connect, and UCS Director) - Simplified Cisco HyperFlex installation and upgrades - Connected TAC: Log Collection, Open Case, Contract Status - Role Based Access Control, Single Sign-On (SAML), Multi-Factor Authentication 	<ul style="list-style-type: none"> - All the features of Base - SaaS and Virtual Appliance - Advanced global search and detailed inventory - Server HCL compliance check with driver Recommendations - Virtual Keyboard-Video-Mouse (vKVM) - ServiceNow Integration - Cisco Intersight Mobile App - HX Storage Capacity Planning (in Tech Preview) - Cisco Standalone UCS C-Series management (M4 and later) <ul style="list-style-type: none"> - Policy-based configuration with Profiles - Firmware and server actions (Power On/Off, reboot, etc) - Includes UCS Central and IMC Supervisor 	<ul style="list-style-type: none"> - All the features of Essentials - SaaS and Virtual Appliance - Tunneled Virtual Keyboard-Video-Mouse (vKVM) (Target Q1,CY2020) - Storage Widget for Pure Storage (Target Q1,CY2020) - Storage Inventory Status for Pure: Capacity and Utilization Storage (Target Q1,CY2020) - Multi-Domain Inventory correlation: Server, Virtualization, Storage (GA: Target Q1,CY2020) - OS Install (in Tech Preview , GA: Target Q1,CY2020) - HX Edge • SD-WAN (Tech Preview Target Q1,CY2020) 	<ul style="list-style-type: none"> - All the features of Advantage - SaaS and Virtual Appliance - Includes UCS Director - Storage Automation with Pure Storage (Target Q1,CY2020) - VM Automation (Target Q1,CY2020) - Workflow Designer (Tech Preview Target Q1,CY2020) 	<p>Pure Storage Integration Requirements:</p> <p>Advantage Storage Widgets and Inventory Status (Capacity/Utilization).</p> <p>Premier Storage Automation.</p>

Test Setup, Configuration, and Load Recommendation

We tested a single Cisco UCS B200 M6 blade to validate against the performance of one and eight Cisco UCS B200 M6 blades 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 7 2109 with 384 Multi-session OS sessions and 280 Single-session OS sessions.

Figure 44. Test Configuration for Single Server Scalability Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs

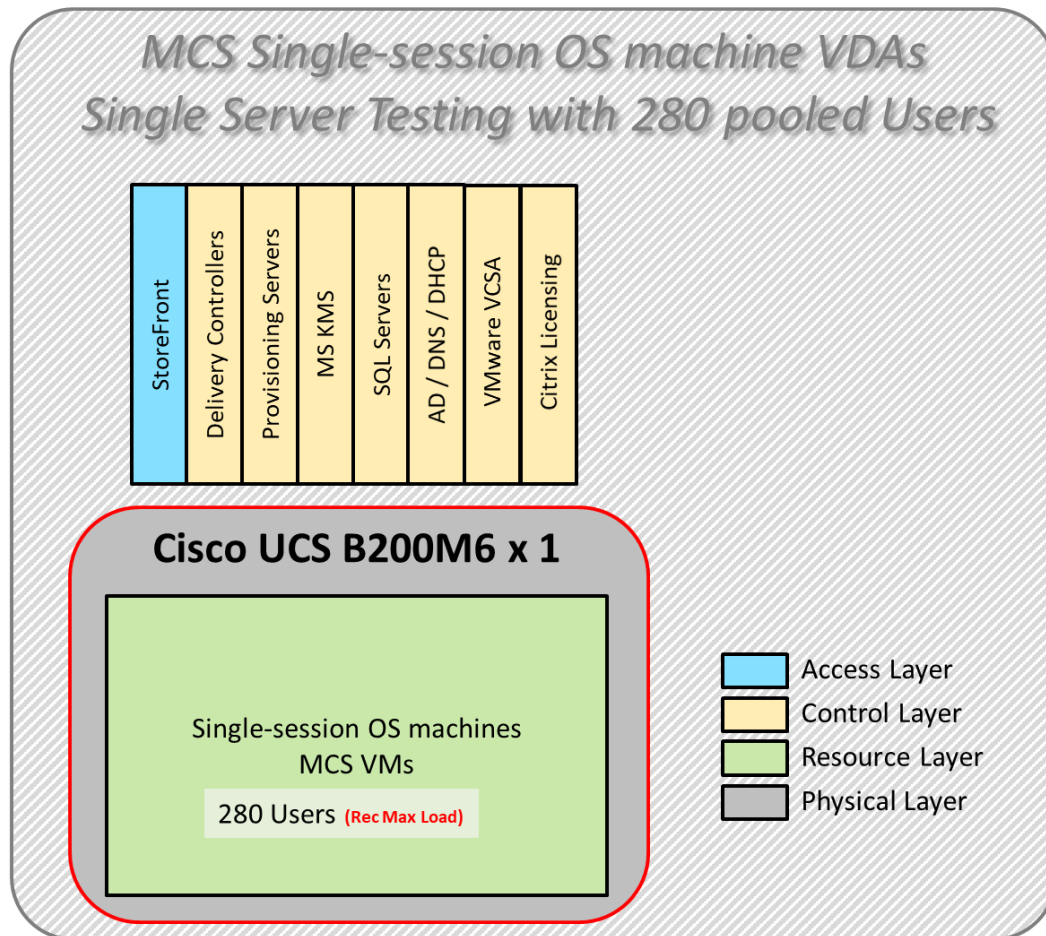


Figure 45. Test configuration for Single Server Scalability Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs

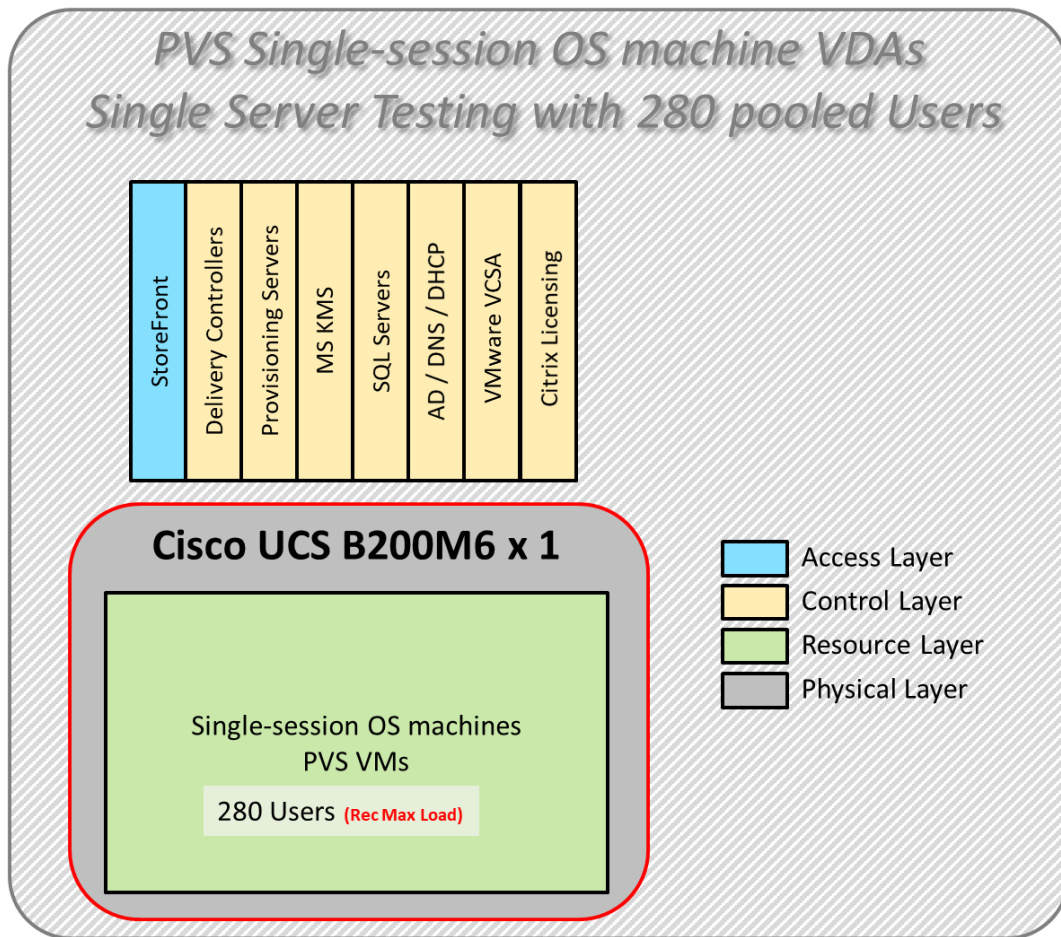
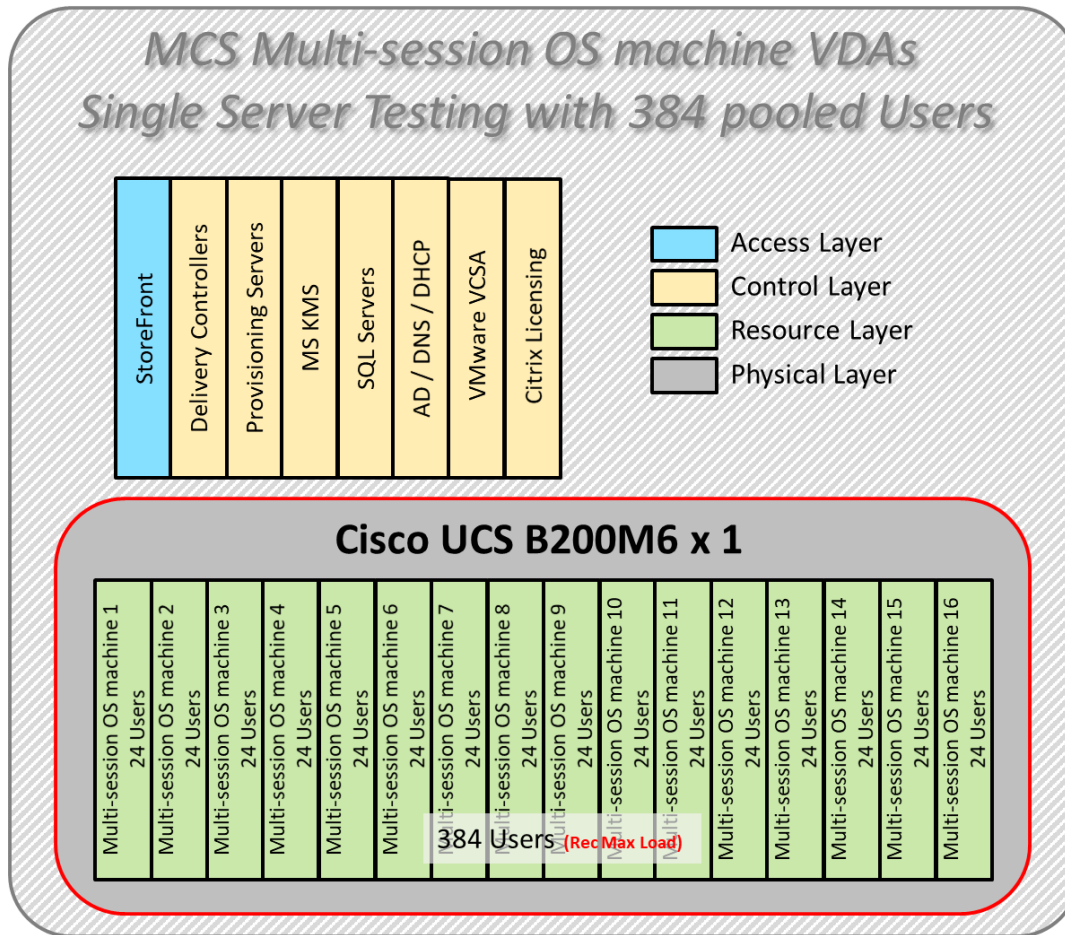


Figure 46. Test configuration for Single Server Scalability Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs



Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 1 Cisco UCS B200 M6 Blade Servers with Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM for all host blades
- Cisco UCS VIC 1440 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 4.2(1f)
- Pure Storage Purity//FA 6.1.7
- ESXi 7.0 Update 2a for host blades
- Citrix Virtual Apps and Desktops 7 2109
- Microsoft SQL Server 2019
- Microsoft Windows 10 64 bit (1909), 2vCPU, 3 GB RAM, 40 GB HDD (master)
- Microsoft Windows Server 2019 (1809), 8vCPU, 32GB RAM, 60 GB vDisk (master)
- Microsoft Office 2019 32-bit
- 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 7 2109 with:

- 1960 MCS Single-session OS sessions
- 1960 PVS Single-session OS sessions
- 2688 MCS Multi-session OS sessions

Note: Server N+1 fault tolerance is factored into this solution for each cluster/workload.

Figure 47. Test Configuration for Full Scale Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs

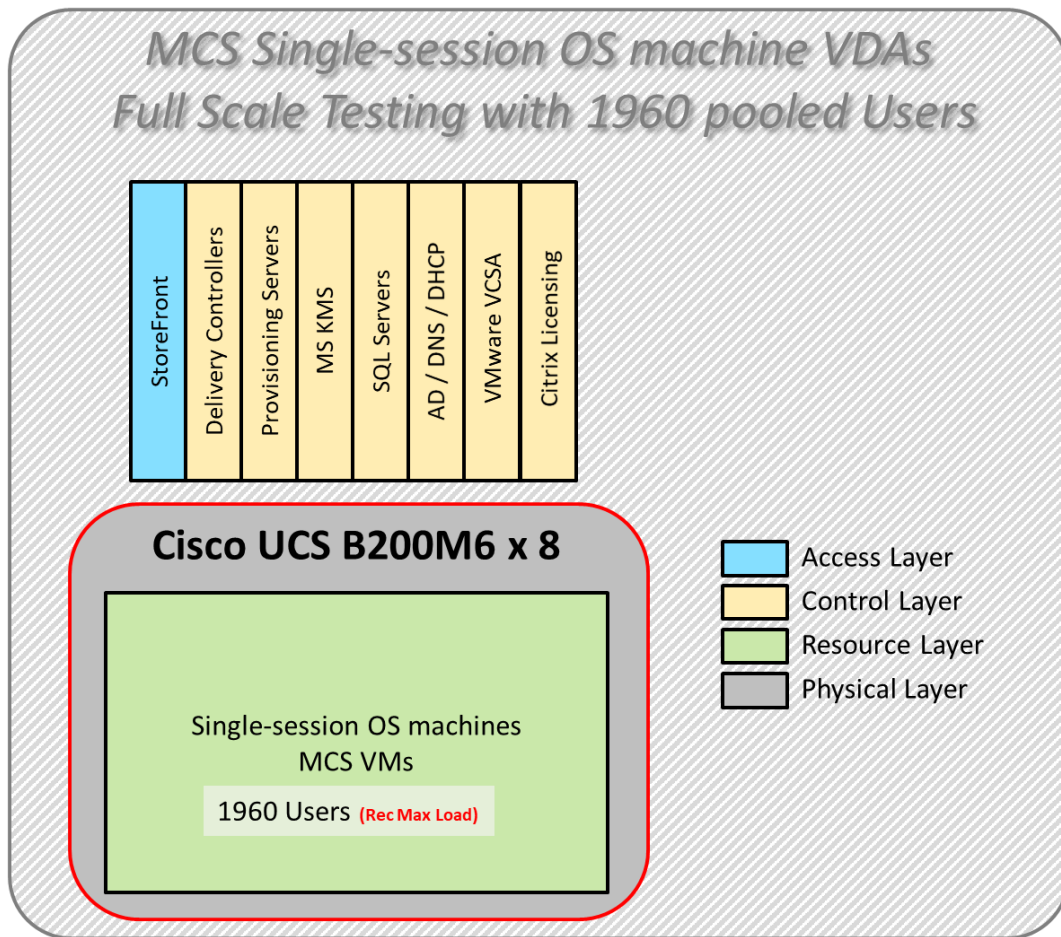


Figure 48. Test Configuration for Full Scale Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs

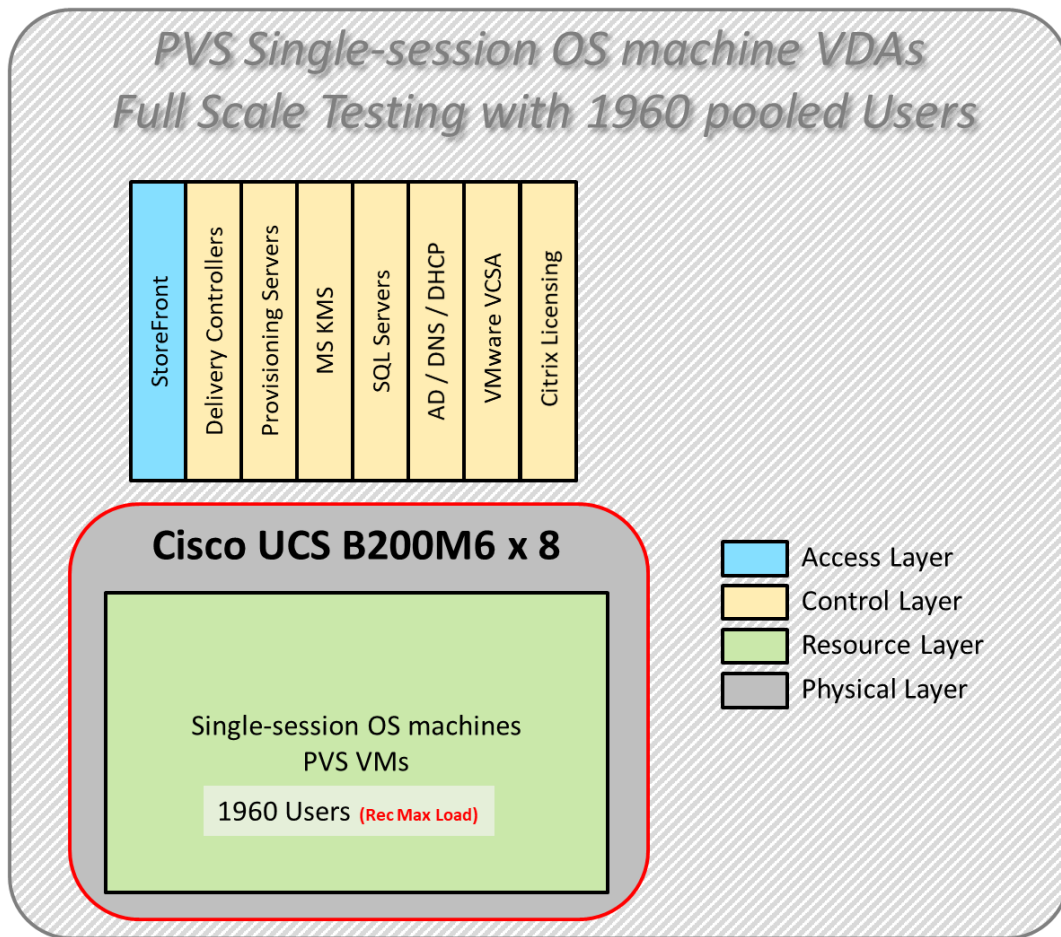
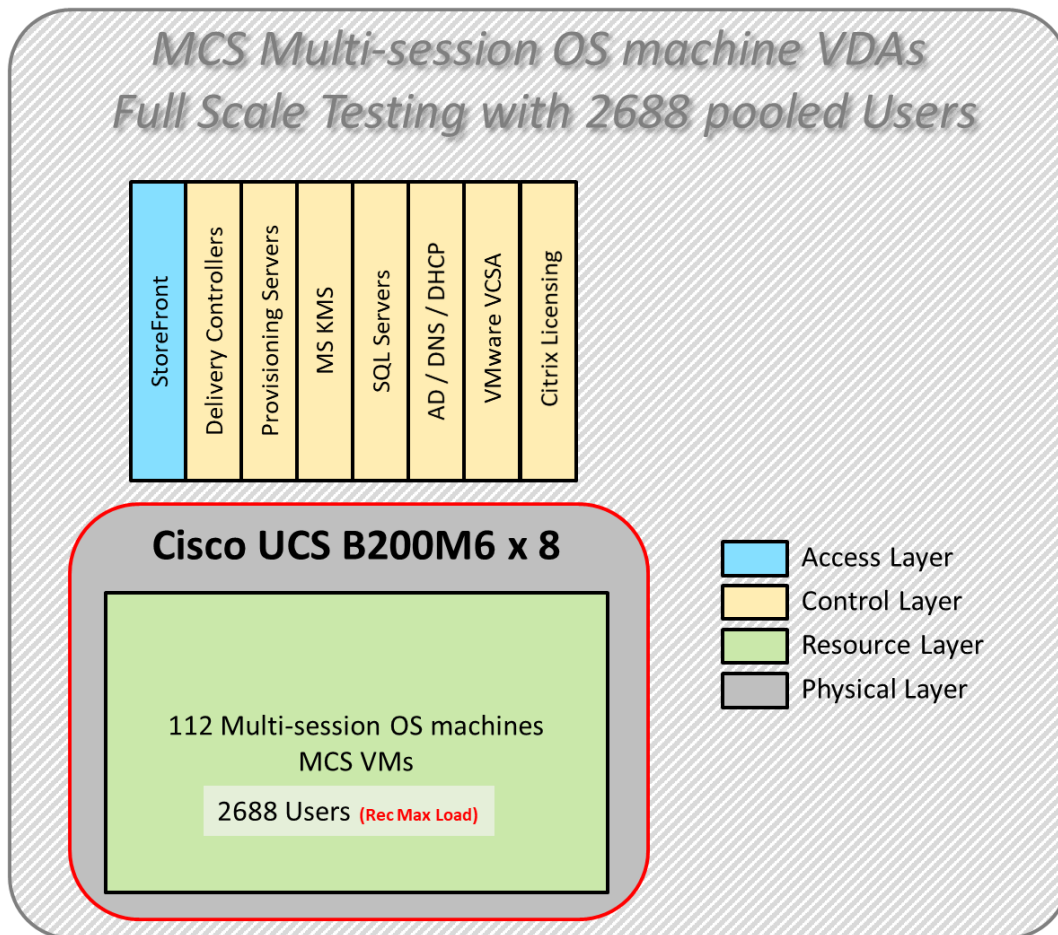


Figure 49. Test Configuration for Full Scale Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs



Hardware components:

- Cisco UCS 5108 Blade Server Chassis
- 2 Cisco UCS 6454 4th Gen Fabric Interconnects
- 8 Cisco UCS B200 M6 Blade Servers with Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM for all host blades
- Cisco VIC 1440 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 4.2(1f)
- Pure Storage Purity//FA 6.1.7
- ESXi 7.0 Update 2a for host blades
- Citrix Virtual Apps and Desktops 7 2109
- Microsoft SQL Server 2019
- Microsoft Windows 10 64 bit (1909), 2vCPU, 3 GB RAM, 40 GB HDD (master)
- Microsoft Windows Server 2019 (1809), 8vCPU, 32GB RAM, 60 GB vDisk (master)
- Microsoft Office 2019 32-bit
- 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

The following protocol was used for each test cycle in this study to ensure consistent results.

Pre-Test Setup for Single and Multi-Blade Testing

All virtual machines were shut down utilizing the VMware Horizon Console and vCenter.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.

All VMware ESXi VDI host blades to be tested were restarted prior to each test cycle.

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. For testing where the user session count exceeds 1000 users, we will now deem the test run successful with up to 1% session failure rate.

In addition, Cisco requires that the Login VSI Benchmark method is used for all single server and scale testing. This assures that our tests represent real-world scenarios. For each of the three consecutive runs on single server tests, the same process was followed.

Time 0:00:00 Start PerfMon/Esxtop Logging on the following system:

Infrastructure and VDI Host Blades used in the test run

vCenter used in the test run.

All Infrastructure virtual machines used in test run (AD, SQL, brokers, image mgmt., and so on)

Time 0:00:10 Start Storage Partner Performance Logging on Storage System.

Time 0:05: Boot Virtual Desktops/RDS Virtual Machines using View Connection server.

The boot rate should be around 10-12 virtual machines per minute per server.

Time 0:06 First machines boot.

Time 0:30 Single Server or Scale target number of desktop virtual machines booted on 1 or more blades.

No more than 30 minutes for boot up of all virtual desktops is allowed.

Time 0:35 Single Server or Scale target number of desktop virtual machines desktops available on View Connection Server.

Virtual machine settling time.

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.

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).

Time 2:23 Single Server or Scale target number of desktop virtual machines desktops launched (48 minute benchmark launch rate).

Time 2:25 All launched sessions must become active. id test run within this window.

Time 2:40 Login VSI Test Ends (based on Auto Logoff 15 minutes period designated above).

Time 2:55 All active sessions logged off.

Time 2:57 All logging terminated; Test complete.

Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows machines.

Time 3:30 Reboot all hypervisor hosts.

Time 3:45 Ready for the new test sequence.

Success Criteria

Our pass criteria for this testing is as follows:

- Cisco will run tests at a session count level that effectively utilizes the blade capacity measured by CPU utilization, memory utilization, storage utilization, and network utilization. We will use Login VSI to launch version 4.1.x Office Worker workloads. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The VMware Horizon Console be monitored throughout the steady state to make sure of the following:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered, unavailable or available state at any time during steady state
- Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Stuck sessions define a test failure condition.

- Cisco requires three consecutive runs with results within +/-1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with the proposed white paper.)

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing. FlashStack Data Center with Cisco UCS and Citrix Virtual Apps and Desktops 7 2109 on VMware ESXi 7.0 Update 2a Test Results.

The purpose of this testing is to provide the data needed to validate VMware Horizon Remote Desktop Sessions (RDS) and VMware Horizon Virtual Desktop (VDI) instant-clones and VMware Horizon Virtual Desktop (VDI) full-clones models using ESXi and vCenter to virtualize Microsoft Windows 10 desktops and Microsoft Windows Server 2019 sessions on Cisco UCS B200 M6 Blade 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 amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what 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 50. Sample of a VSI Max Response Time Graph, Representing a Normal Test



Figure 51. Sample of a VSI Test Response Time Graph with a Performance Issue



When the test is finished, VSI_{max} can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSI_{max} is not reached, and the amount of sessions ran successfully.

The response times are very different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. These response 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 samples 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.

VSI_{max} v4.1.x is reached when the VSI_{base} + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSI_{max} 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 VSI_{max} 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 VSI_{max} v4.1.x, the performance of the system is not decided by the total average response time, but by the latency it 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, VSI_{max} is not hit, and the number of sessions ran successfully. This approach is fundamentally different in comparison to previous VSI_{max} methods, as it was always required to saturate the system beyond VSI_{max} threshold.

Lastly, VSI_{max} v4.1.x is now always reported with the average baseline VSI response time result. For example: "The VSI_{max} 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. VSI_{max} indicates what the total user capacity is for the system. These two are not automatically connected and related.

When a server with a very fast dual core CPU, running at 3.6 GHz, is compared to a 10 core CPU, running at 2.26 GHz, the dual core machine will give an individual user better performance than the 10 core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

However, the server with the slower 10 core CPU will easily have a larger capacity than the faster dual core system. This is indicated by VSI_{max} v4.1.x, and the higher VSI_{max} is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSI_{max} method is introduced: VSI_{max} v4.1.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 7 2109 Virtual Desktop and Citrix Virtual Apps and Desktops 7 2109 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 15. 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

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:

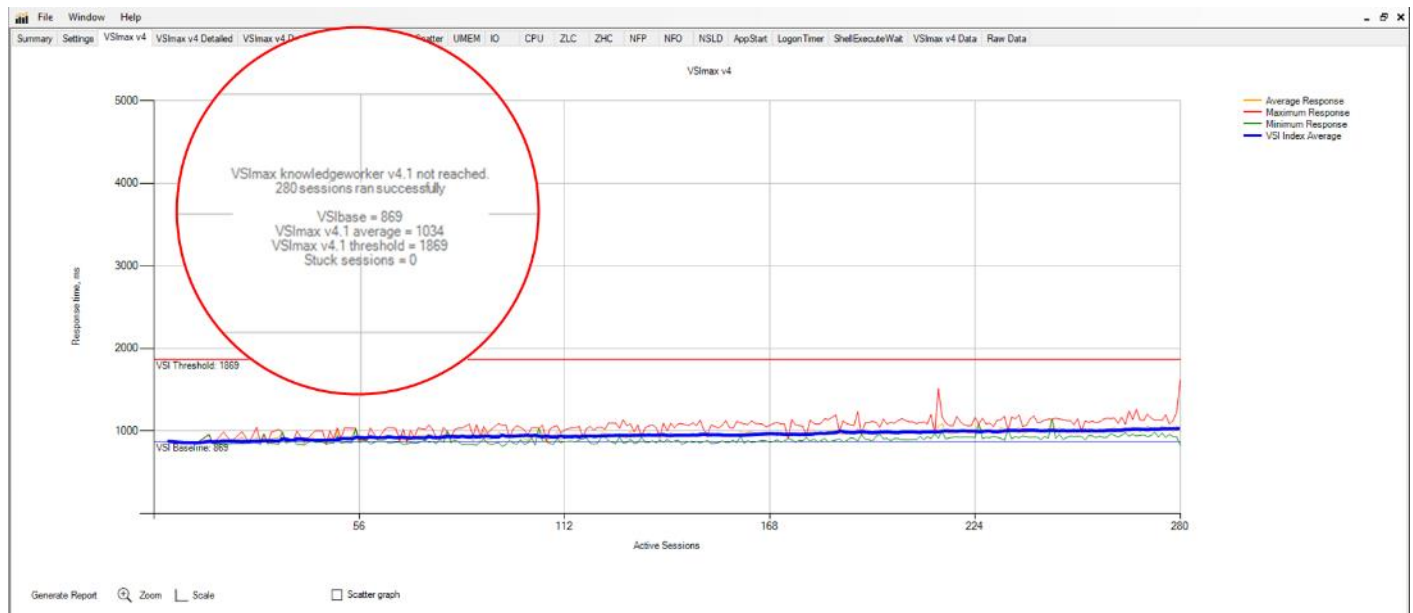
- 280 MCS Single-session OS sessions (Random)
- 280 PVS Single-session OS sessions (Random)
- 384 MCS Multi-session OS sessions (Random)

Single-Server Recommended Maximum Workload for MCS Single-session OS Random Sessions with 280 Users

The recommended maximum workload for a Cisco UCS B200 M6 blade server with dual Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM is 280 Windows 10 64-bit non-persistent MCS virtual machines with 2 vCPU and 3.5 GB RAM.

Login VSI performance data is shown below:

Figure 52. Single Server | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | VSI Score



Performance data for the server running the workload is shown below:

Figure 53. Single Server | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Host CPU Utilization




Figure 54. Single Server | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Host Memory Utilization





Figure 55. Single Server | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Host Network Utilization




Single-Server Recommended Maximum Workload for PVS Single-session OS Random Sessions with 280 Users

The recommended maximum workload for a Cisco UCS B200 M6 blade server with dual Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM is 280 Windows 10 64-bit VDI non-persistent PVS virtual machines with 2 vCPU and 3.5GB RAM.

Login VSI performance data is as shown below:

Figure 56. Single Server | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | VSI Score



Performance data for the server running the workload is shown below:

Figure 57. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | Host CPU Utilization




Figure 58. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | Host Memory Utilization





Figure 59. Single Server | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | Host Network Utilization




Single-Server Recommended Maximum Workload for MCS Multiple-session OS Random Sessions with 384 Users

The recommended maximum workload for a Cisco UCS B200 M6 blade server with dual Intel(R) Xeon(R) Gold 6338 CPU 2.00GHz 32-core processors, 1TB 3200MHz RAM is 384 Windows Server 2019 sessions. The blade server ran 16 Windows Server 2019 Virtual Machines. Each virtual server was configured with 8 vCPUs and 32GB RAM.

LoginVSI data is shown below:

Figure 60. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | VSI Score



Performance data for the server running the workload is shown below:

Figure 61. Single Server Recommended Maximum Workload Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Host CPU Utilization




Figure 62. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Host Memory Utilization





Figure 63. Single Server | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Host Network Utilization



Performance data for the RDS Virtual Machine running the workload is shown below:

Figure 64. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Virtual Machine CPU Utilization




Figure 65. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Virtual Machine Memory Utilization





Figure 66. Single Server Recommended Maximum Workload | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Network Utilization



Full Scale Workload Testing

This section describes the key performance metrics that were captured on the Cisco UCS, during the full-scale testing. Full Scale testing was done with following Workloads using 8 Cisco UCS B200M6 Blade Servers, configured in a single ESXi Host Pool and designed to support single Host failure (N+1 Fault tolerance):

- 1960 MCS Single-session OS sessions
- 1960 PVS Single-session OS sessions
- 2688 MCS Multi-session OS sessions

To achieve the target, sessions were launched against each workload set at a time. As per the Cisco Test Protocol for VDI solutions, all sessions were launched within 48 minutes (using the official Knowledge Worker Workload in VSI Benchmark Mode) and all launched sessions became active within two minutes subsequent to the last logged in session.

Full Scale Recommended Maximum Workload Testing for MCS Single-session OS Machine VDAs with 1960 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 1960 MCS Single-session OS machines using 8 blades in a single pool.

The workload for the test is 1960 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 67. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs| VSI Score




Figure 68. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Host CPU Utilization




Figure 69. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Host Memory Utilization




Figure 70. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Host Network Utilization




Figure 71. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | Pure Storage FlashArray//X70 R3 System Latency Chart




Figure 72. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | FlashArray//X70 R3 System IOPS Chart




Figure 73. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | FlashArray//X70 R3 System Bandwidth Chart


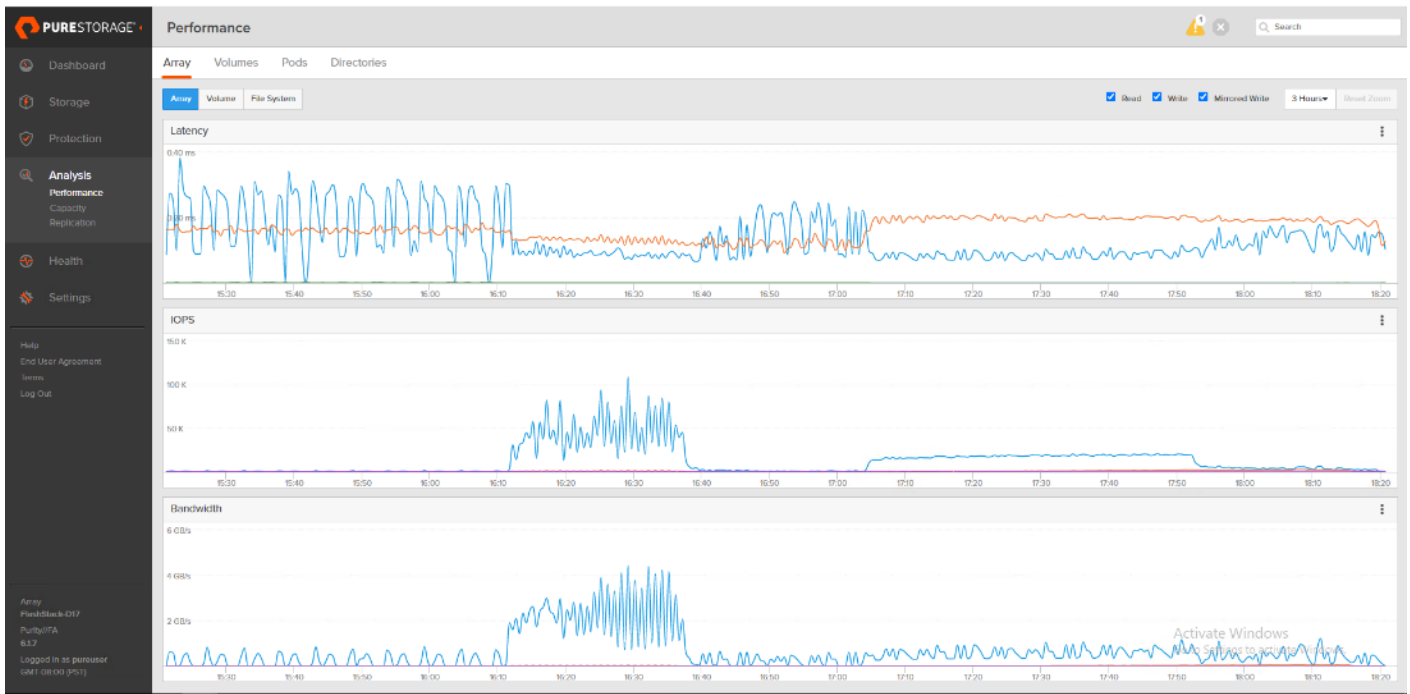


Figure 74. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Single-session OS machine VDAs | FlashArray//X70 R3 Performance Chart



Full Scale Recommended Maximum Workload Testing for PVS Single-session OS Machine VDAs with 1960 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 1960 PVS Single-session OS machines using 8 blades in a single pool.

The workload for the test is 1960 Non-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 75. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | VSI Score




Figure 76. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | Host CPU Utilization




Figure 77. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | Host Memory Utilization




Figure 78. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | Host Network Utilization




Figure 79. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | FlashArray//X70 R3 System Latency Chart




Figure 80. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | FlashArray//X70 R3 System IOPS Chart




Figure 81. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | FlashArray//X70 R3 System Bandwidth Chart


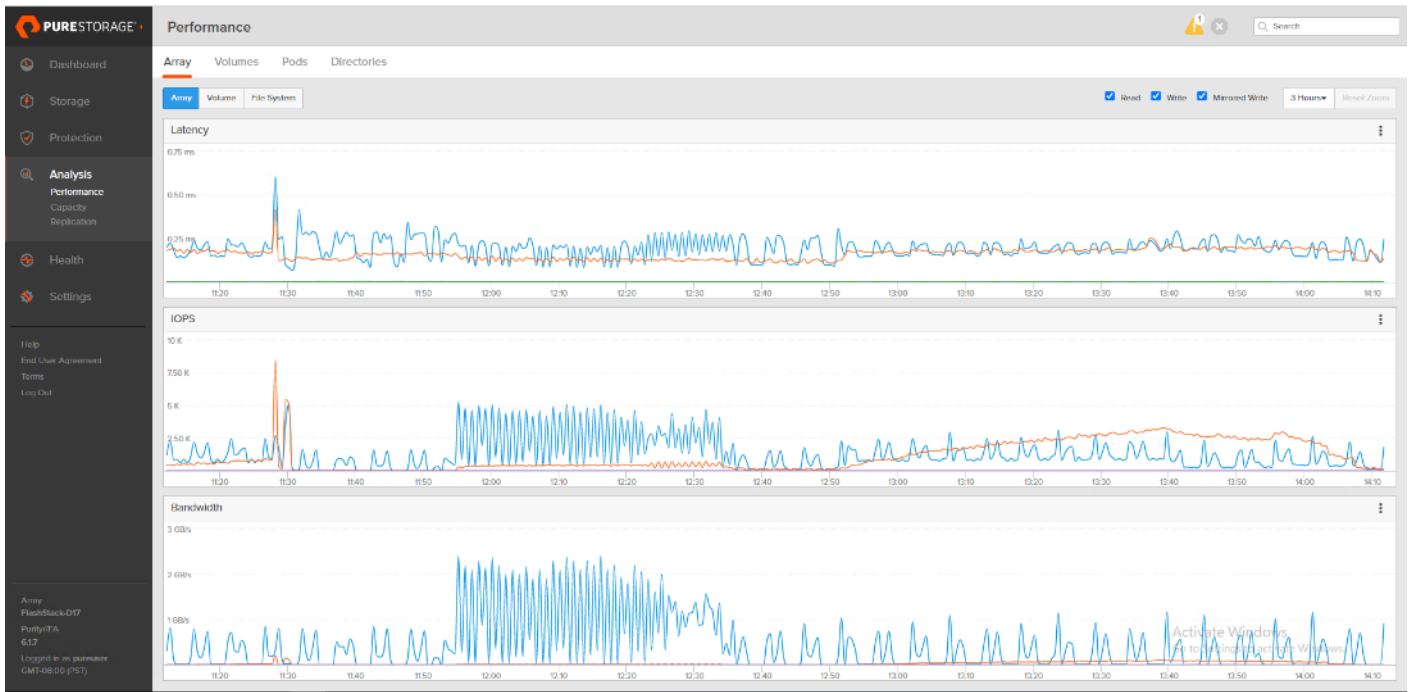


Figure 82. Full Scale | 1960 Users | Citrix Virtual Apps and Desktops 7 2109 PVS Single-session OS machine VDAs | FlashArray//X70 R3 System Performance Chart



Full Scale Recommended Maximum Workload for MCS Multi-session OS Random Sessions with 2688 Users

This section describes the key performance metrics that were captured on the Cisco UCS and Pure Storage FlashArray//X70 R3 array, during the MCS Multi-session OS full-scale testing with 2688 Desktop Sessions using 8 blades configured in single Host Pool.

The Multi-session OS workload for the solution is 2688 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 83. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | VSI Score




Figure 84. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Host CPU Utilization




Figure 85. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Host Memory Utilization




Figure 86. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Host Network Utilization




Figure 87. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Pure Storage FlashArray//X70 R3 System Latency Chart




Figure 88. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Pure Storage FlashArray//X70 R3 System IOPS Chart




Figure 89. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | Pure Storage FlashArray//X70 R3 System Bandwidth Chart





Figure 90. Full Scale | 2688 Users | Citrix Virtual Apps and Desktops 7 2109 MCS Multi-session OS machine VDAs | FlashArray//X70 R3 System Performance Chart



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. 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

Vadim Lebedev, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Vadim Lebedev for the last five years is a member of the Cisco's Computing Systems Product Group team focusing on design, testing, and solutions validation, technical content creation, and performance testing/benchmarking. He has years of experience in server and desktop virtualization. Vadim is a subject matter expert on Desktop/Server virtualization, Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, and NVIDIA Graphics.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the following for their contribution and expertise that resulted in developing this document:

- Cisco Systems, Inc

Sreenivasa Edula, Technical Marketing Engineer

John George, Technical Marketing Engineer

Haseeb Niazi, Technical Marketing Engineer

- Pure Storage, Inc.

Joe Houghes, Senior Solutions Architect

Craig Waters, Technical Director

References

This section provides links to additional information for each partner's solution component of this document.

Cisco UCS B-Series Servers

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b200M6-specsheet.pdf>

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/datasheet-listing.html>

<https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-b200-M6-blade-server/model.html>

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/hw/blade-servers/B200M6.pdf

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/collateral/interfaces-modules/unified-computing-system-adapters/datasheet-c78-741130.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>

Citrix References

<https://docs.citrix.com/en-us/citrix-virtual-apps-desktops/2109/whats-new.html>

<https://docs.citrix.com/en-us/provisioning/2109.html>

VMware References

<https://docs.microsoft.com/en-us/fslogix/>

Microsoft References

<https://docs.vmware.com/en/VMware-vSphere/index.html>

Login VSI Documentation

https://www.loginvsi.com/documentation/Main_Page

https://www.loginvsi.com/documentation/Start_your_first_test

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/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

Ethernet Network Configuration

The following section provides a detailed procedure for configuring the Cisco Nexus 9000 Switches used in this study.

Cisco Nexus 93180YC-A Configuration

```
version 9.3(7)

switchname AAD17-NX9K-A

class-map type network-qos class-fcoe
match qos-group 1

class-map type network-qos class-all-flood
match qos-group 2

class-map type network-qos class-ip-multicast
match qos-group 2

policy-map type network-qos jumbo
  class type network-qos class-fcoe
    mtu 2158
  class type network-qos class-default
    mtu 9216

install feature-set fcoe-npv

vdc AAD17-NX9K-A id 1
  allow feature-set fcoe-npv
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
```

limit-resource u6route-mem minimum 96 maximum 96

limit-resource m4route-mem minimum 58 maximum 58

limit-resource m6route-mem minimum 8 maximum 8

feature-set fcoe-npv

feature telnet

cfs eth distribute

feature interface-vlan

feature hsrp

feature lacp

feature dhcp

feature vpc

feature lldp

no password strength-check

username admin password 5 \$5\$d3vc8gvD\$hmf.YoRRPcqZ2dDGV2laVKYZsPSPIs8E9bpUzMciMZ0
role network-admin

ip domain-lookup

system default switchport

class-map type qos match-all class-fcoe

policy-map type qos jumbo

class class-default

set qos-group 0

system qos

service-policy type network-qos jumbo

copp profile lenient

snmp-server user admin network-admin auth md5 0xc9a73d344387b8db2dc0f3fc624240ac priv
0xc9a73d344387b8db2dc0f3fc624240ac localizedkey

snmp-server host 10.24.66.169 traps version 2c public udp-port 1165

snmp-server host 10.24.72.119 traps version 2c public udp-port 1163

rmon event 1 description FATAL(1) owner PMON@FATAL

rmon event 2 description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 description ERROR(3) owner PMON@ERROR

rmon event 4 description WARNING(4) owner PMON@WARNING

rmon event 5 description INFORMATION(5) owner PMON@INFO

ntp server 10.10.70.2 use-vrf default

ntp peer 10.10.70.3 use-vrf default

ntp server 72.163.32.44 use-vrf management

ntp logging

ntp master 8

vlan 1,70-76

vlan 70

name InBand-Mgmt-SP

vlan 71

name Infra-Mgmt-SP

vlan 72

name VM-Network-SP

vlan 73

name vMotion-SP

vlan 74

name Storage_A-SP

vlan 75

name Storage_B-SP

vlan 76

name Launcher-SP

service dhcp

ip dhcp relay

ip dhcp relay information option

ipv6 dhcp relay

vrf context management

ip route 0.0.0.0/0 10.29.164.1

hardware access-list tcam region ing-racl 1536

hardware access-list tcam region ing-redirect 256

vpc domain 70

role priority 1000

peer-keepalive destination 10.29.164.234 source 10.29.164.233

interface Vlan1

no shutdown

ip address 10.29.164.241/24

```
interface Vlan70
  no shutdown
  ip address 10.10.70.2/24
  hsrp version 2
  hsrp 70
    preempt
    priority 110
  ip 10.10.70.1
```

```
interface Vlan71
  no shutdown
  ip address 10.10.71.2/24
  hsrp version 2
  hsrp 71
    preempt
    priority 110
  ip 10.10.71.1
```

```
interface Vlan72
  no shutdown
  ip address 10.72.0.2/19
  hsrp version 2
  hsrp 72
    preempt
```

priority 110

ip 10.72.0.1

ip dhcp relay address 10.10.71.11

ip dhcp relay address 10.10.71.12

interface Vlan73

no shutdown

ip address 10.10.73.2/24

hsrp version 2

hsrp 73

preempt

priority 110

ip 10.10.73.1

interface Vlan74

no shutdown

ip address 10.10.74.2/24

hsrp version 2

hsrp 74

preempt

priority 110

ip 10.10.74.1

interface Vlan75

no shutdown

ip address 10.10.75.2/24

hsrp version 2

hsrp 75

preempt

priority 110

ip 10.10.75.1

interface Vlan76

no shutdown

ip address 10.10.76.2/23

hsrp version 2

hsrp 76

preempt

priority 110

ip 10.10.76.1

ip dhcp relay address 10.10.71.11

ip dhcp relay address 10.10.71.12

interface port-channel10

interface port-channel11

description FI-Uplink-D17

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 11

interface port-channel12

description FI-Uplink-D17

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 12

interface port-channel13

description FI-Uplink-D16

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 13

```
interface port-channel14
  description FI-Uplink-D16
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 14
```

```
interface port-channel70
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  spanning-tree port type network
  service-policy type qos input jumbo
  vpc peer-link
```

```
interface port-channel101
  description to PureStorage ethernet port eth2
  shutdown
  switchport access vlan 72
  spanning-tree port type edge
  mtu 9216
  service-policy type qos input jumbo
```

vpc 101

interface Ethernet1/1

interface Ethernet1/2

switchport mode trunk

switchport trunk allowed vlan 1,70-76

interface Ethernet1/3

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 13 mode active

interface Ethernet1/4

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 13 mode active

interface Ethernet1/5

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 14 mode active

interface Ethernet1/6

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 14 mode active

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

switchport access vlan 71

spanning-tree port type edge

interface Ethernet1/34

switchport access vlan 71

spanning-tree port type edge

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 11 mode active

interface Ethernet1/52

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 12 mode active

interface Ethernet1/53

switchport mode trunk

switchport trunk allowed vlan 1,70-76

channel-group 70 mode active

interface Ethernet1/54

switchport mode trunk

switchport trunk allowed vlan 1,70-76

channel-group 70 mode active

interface mgmt0

vrf member management

ip address 10.29.164.233/24

line console

line vty

boot nxos bootflash:/nxos.7.0.3.I7.2.bin

no system default switchport shutdown

Cisco Nexus 93180YC -B Configuration

version 9.3(7)

switchname AAD17-NX9K-B

class-map type network-qos class-fcoe

match qos-group 1

class-map type network-qos class-all-flood

match qos-group 2

class-map type network-qos class-ip-multicast

match qos-group 2

policy-map type network-qos jumbo

class type network-qos class-fcoe

mtu 2158

class type network-qos class-default

mtu 9216

install feature-set fcoe-npv

vdc AAD17-NX9K-B id 1

allow feature-set fcoe-npv

limit-resource vlan minimum 16 maximum 4094

limit-resource vrf minimum 2 maximum 4096

limit-resource port-channel minimum 0 maximum 511

limit-resource u4route-mem minimum 248 maximum 248

limit-resource u6route-mem minimum 96 maximum 96

limit-resource m4route-mem minimum 58 maximum 58

limit-resource m6route-mem minimum 8 maximum 8

feature-set fcoe-npv

feature telnet

cfs eth distribute

feature interface-vlan

feature hsrp

feature lacp

feature dhcp

feature vpc

feature lldp

no password strength-check

username admin password 5 \$5\$/48.OHa8\$g6pOMLlwrzqxJesMYoP5CNphujBksPPRjn4l3iFfOp. role
network-admin

ip domain-lookup

system default switchport

class-map type qos match-all class-fcoe

policy-map type qos jumbo

class class-default

set qos-group 0

system qos

service-policy type network-qos jumbo

copp profile lenient

snmp-server user admin network-admin auth md5 0x6d450e3d5a3927ddee1dadd30e5f616f priv
0x6d450e3d5a3927ddee1dadd30e5f616f localizedkey

snmp-server host 10.24.66.169 traps version 2c public udp-port 1166

snmp-server host 10.24.72.119 traps version 2c public udp-port 1164

rmon event 1 description FATAL(1) owner PMON@FATAL

rmon event 2 description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 description ERROR(3) owner PMON@ERROR

rmon event 4 description WARNING(4) owner PMON@WARNING

rmon event 5 description INFORMATION(5) owner PMON@INFO

ntp peer 10.10.70.2 use-vrf default

ntp server 10.10.70.3 use-vrf default

ntp server 72.163.32.44 use-vrf management

ntp logging

ntp master 8

vlan 1,70-76

vlan 70

```
name InBand-Mgmt-SP
vlan 71
name Infra-Mgmt-SP
vlan 72
name VM-Network-SP
vlan 73
name vMotion-SP
vlan 74
name Storage_A-SP
vlan 75
name Storage_B-SP
vlan 76
name Launcher-SP

service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
ip route 0.0.0.0/0 10.29.164.1
hardware access-list tcam region ing-racl 1536
hardware access-list tcam region ing-redirect 256
vpc domain 70
role priority 2000
```

peer-keepalive destination 10.29.164.233 source 10.29.164.234

interface Vlan1

no shutdown

ip address 10.29.164.240/24

interface Vlan70

no shutdown

ip address 10.10.70.3/24

hsrp version 2

hsrp 70

preempt

priority 110

ip 10.10.70.1

interface Vlan71

no shutdown

ip address 10.10.71.3/24

hsrp version 2

hsrp 71

preempt

priority 110

ip 10.10.71.1

interface Vlan72

no shutdown

ip address 10.72.0.2/19

hsrp version 2

hsrp 72

preempt

priority 110

ip 10.72.0.1

ip dhcp relay address 10.10.71.11

ip dhcp relay address 10.10.71.12

interface Vlan73

no shutdown

ip address 10.10.73.3/24

hsrp version 2

hsrp 73

preempt

priority 110

ip 10.10.73.1

interface Vlan74

no shutdown

ip address 10.10.74.3/24

hsrp version 2

hsrp 74

preempt

priority 110

ip 10.10.74.1

interface Vlan75

no shutdown

ip address 10.10.75.3/24

hsrp version 2

hsrp 75

preempt

priority 110

ip 10.10.75.1

interface Vlan76

no shutdown

ip address 10.10.76.3/23

hsrp version 2

hsrp 76

preempt

priority 110

ip 10.10.76.1

ip dhcp relay address 10.10.71.11

ip dhcp relay address 10.10.71.12

interface port-channel10

interface port-channel11

description FI-Uplink-D17

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 11

interface port-channel12

description FI-Uplink-D17

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 12

interface port-channel13

description FI-Uplink-D16

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 13

interface port-channel14

description FI-Uplink-D16

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type edge trunk

mtu 9216

service-policy type qos input jumbo

vpc 14

interface port-channel70

description vPC-PeerLink

switchport mode trunk

switchport trunk allowed vlan 1,70-76

spanning-tree port type network

service-policy type qos input jumbo

vpc peer-link

```
interface port-channel101
  description to PureStorage ethernet port eth2
  shutdown
  switchport access vlan 72
  mtu 9216
  service-policy type qos input jumbo
  vpc 101
```

```
interface Ethernet1/1
  switchport access vlan 70
  speed 1000
```

```
interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
```

```
interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  mtu 9216
  channel-group 13 mode active
```

```
interface Ethernet1/4
  switchport mode trunk
```

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 13 mode active

interface Ethernet1/5

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 14 mode active

interface Ethernet1/6

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 14 mode active

interface Ethernet1/7

interface Ethernet1/8

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25

interface Ethernet1/26

interface Ethernet1/27

interface Ethernet1/28

interface Ethernet1/29

interface Ethernet1/30

interface Ethernet1/31

interface Ethernet1/32

interface Ethernet1/33

switchport access vlan 71

spanning-tree port type edge

interface Ethernet1/34

switchport access vlan 71

spanning-tree port type edge

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 11 mode active

interface Ethernet1/52

switchport mode trunk

switchport trunk allowed vlan 1,70-76

mtu 9216

channel-group 12 mode active

```
interface Ethernet1/53
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  channel-group 70 mode active
```

```
interface Ethernet1/54
  switchport mode trunk
  switchport trunk allowed vlan 1,70-76
  channel-group 70 mode active
```

```
interface mgmt0
  vrf member management
  ip address 10.29.164.234/24

line console

line vty

boot nxos bootflash:/nxos.7.0.3.17.2.bin

no system default switchport shutdown
```

Fibre Channel Network Configuration

The following section provides a detailed procedure for configuring the Cisco MDS 9100 Switches used in this study.

Cisco MDS 9132T-A Configuration

```
version 8.3(1)

power redundancy-mode redundant

feature npiv

feature fport-channel-trunk
```

role name default-role

description This is a system defined role and applies to all users.

rule 5 permit show feature environment

rule 4 permit show feature hardware

rule 3 permit show feature module

rule 2 permit show feature snmp

rule 1 permit show feature system

no password strength-check

username admin password 5 \$5\$Dcs72Ao/\$8lHyVrotTm4skqb/84BC793tgdly/yWf9loMx2OEg6C role network-admin

ip domain-lookup

ip name-server 10.10.61.30

ip host ADD16-MDS-A 10.29.164.238

aaa group server radius radius

snmp-server user admin network-admin auth md5 0x616758aed4f07bab2d24f3d594ebd649 priv 0x616758aed4f07bab2d24f3d594ebd649 localizedkey

snmp-server host 10.24.30.91 traps version 2c public udp-port 1163

snmp-server host 10.24.46.67 traps version 2c public udp-port 1163

snmp-server host 10.24.66.169 traps version 2c public udp-port 1163

snmp-server host 10.24.72.119 traps version 2c public udp-port 1165

rmon event 1 log trap public description FATAL(1) owner PMON@FATAL

rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 log trap public description ERROR(3) owner PMON@ERROR

rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

ntp server 10.81.254.131

ntp server 10.81.254.202

vsan database

vsan 100 name "FlashStack-VCC-CVD-Fabric-A"

device-alias database

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

device-alias name VCC-Infra01-HBA0 pwwn 20:00:00:25:b5:aa:17:1e

device-alias name VCC-Infra01-HBA2 pwwn 20:00:00:25:b5:aa:17:1f

device-alias name VCC-Infra02-HBA0 pwwn 20:00:00:25:b5:aa:17:3e

device-alias name VCC-Infra02-HBA2 pwwn 20:00:00:25:b5:aa:17:3f

device-alias name VCC-WLHost01-HBA0 pwwn 20:00:00:25:b5:aa:17:00

device-alias name VCC-WLHost01-HBA2 pwwn 20:00:00:25:b5:aa:17:01

device-alias name VCC-WLHost02-HBA0 pwwn 20:00:00:25:b5:aa:17:02

device-alias name VCC-WLHost02-HBA2 pwwn 20:00:00:25:b5:aa:17:03

device-alias name VCC-WLHost03-HBA0 pwwn 20:00:00:25:b5:aa:17:04

device-alias name VCC-WLHost03-HBA2 pwwn 20:00:00:25:b5:aa:17:05

device-alias name VCC-WLHost04-HBA0 pwwn 20:00:00:25:b5:aa:17:06

device-alias name VCC-WLHost04-HBA2 pwwn 20:00:00:25:b5:aa:17:07

device-alias name VCC-WLHost05-HBA0 pwwn 20:00:00:25:b5:aa:17:08

device-alias name VCC-WLHost05-HBA2 pwwn 20:00:00:25:b5:aa:17:09

device-alias name VCC-WLHost06-HBA0 pwwn 20:00:00:25:b5:aa:17:0a

device-alias name VCC-WLHost06-HBA2 pwwn 20:00:00:25:b5:aa:17:0b

device-alias name VCC-WLHost07-HBA0 pwwn 20:00:00:25:b5:aa:17:0c

device-alias name VCC-WLHost07-HBA2 pwwn 20:00:00:25:b5:aa:17:0d
device-alias name VCC-WLHost08-HBA0 pwwn 20:00:00:25:b5:aa:17:0e
device-alias name VCC-WLHost08-HBA2 pwwn 20:00:00:25:b5:aa:17:0f
device-alias name VCC-WLHost09-HBA0 pwwn 20:00:00:25:b5:aa:17:10
device-alias name VCC-WLHost09-HBA2 pwwn 20:00:00:25:b5:aa:17:11
device-alias name VCC-WLHost10-HBA0 pwwn 20:00:00:25:b5:aa:17:12
device-alias name VCC-WLHost10-HBA2 pwwn 20:00:00:25:b5:aa:17:13
device-alias name VCC-WLHost11-HBA0 pwwn 20:00:00:25:b5:aa:17:14
device-alias name VCC-WLHost11-HBA2 pwwn 20:00:00:25:b5:aa:17:15
device-alias name VCC-WLHost12-HBA0 pwwn 20:00:00:25:b5:aa:17:16
device-alias name VCC-WLHost12-HBA2 pwwn 20:00:00:25:b5:aa:17:17
device-alias name VCC-WLHost13-HBA0 pwwn 20:00:00:25:b5:aa:17:18
device-alias name VCC-WLHost13-HBA2 pwwn 20:00:00:25:b5:aa:17:19
device-alias name VCC-WLHost14-HBA0 pwwn 20:00:00:25:b5:aa:17:1a
device-alias name VCC-WLHost14-HBA2 pwwn 20:00:00:25:b5:aa:17:1b
device-alias name VCC-WLHost15-HBA0 pwwn 20:00:00:25:b5:aa:17:1c
device-alias name VCC-WLHost15-HBA2 pwwn 20:00:00:25:b5:aa:17:1d
device-alias name VCC-WLHost16-HBA0 pwwn 20:00:00:25:b5:aa:17:20
device-alias name VCC-WLHost16-HBA2 pwwn 20:00:00:25:b5:aa:17:21
device-alias name VCC-WLHost17-HBA0 pwwn 20:00:00:25:b5:aa:17:22
device-alias name VCC-WLHost17-HBA2 pwwn 20:00:00:25:b5:aa:17:23
device-alias name VCC-WLHost18-HBA0 pwwn 20:00:00:25:b5:aa:17:24
device-alias name VCC-WLHost18-HBA2 pwwn 20:00:00:25:b5:aa:17:25
device-alias name VCC-WLHost19-HBA0 pwwn 20:00:00:25:b5:aa:17:26

device-alias name VCC-WLHost19-HBA2 pwwn 20:00:00:25:b5:aa:17:27
device-alias name VCC-WLHost20-HBA0 pwwn 20:00:00:25:b5:aa:17:28
device-alias name VCC-WLHost20-HBA2 pwwn 20:00:00:25:b5:aa:17:29
device-alias name VCC-WLHost21-HBA0 pwwn 20:00:00:25:b5:aa:17:2a
device-alias name VCC-WLHost21-HBA2 pwwn 20:00:00:25:b5:aa:17:2b
device-alias name VCC-WLHost22-HBA0 pwwn 20:00:00:25:b5:aa:17:2c
device-alias name VCC-WLHost22-HBA2 pwwn 20:00:00:25:b5:aa:17:2d
device-alias name VCC-WLHost23-HBA0 pwwn 20:00:00:25:b5:aa:17:2e
device-alias name VCC-WLHost23-HBA2 pwwn 20:00:00:25:b5:aa:17:2f
device-alias name VCC-WLHost24-HBA0 pwwn 20:00:00:25:b5:aa:17:30
device-alias name VCC-WLHost24-HBA2 pwwn 20:00:00:25:b5:aa:17:31
device-alias name VCC-WLHost25-HBA0 pwwn 20:00:00:25:b5:aa:17:32
device-alias name VCC-WLHost25-HBA2 pwwn 20:00:00:25:b5:aa:17:33
device-alias name VCC-WLHost26-HBA0 pwwn 20:00:00:25:b5:aa:17:34
device-alias name VCC-WLHost26-HBA2 pwwn 20:00:00:25:b5:aa:17:35
device-alias name VCC-WLHost27-HBA0 pwwn 20:00:00:25:b5:aa:17:36
device-alias name VCC-WLHost27-HBA2 pwwn 20:00:00:25:b5:aa:17:37
device-alias name VCC-WLHost28-HBA0 pwwn 20:00:00:25:b5:aa:17:38
device-alias name VCC-WLHost28-HBA2 pwwn 20:00:00:25:b5:aa:17:39
device-alias name VCC-WLHost29-HBA0 pwwn 20:00:00:25:b5:aa:17:3a
device-alias name VCC-WLHost29-HBA2 pwwn 20:00:00:25:b5:aa:17:3b
device-alias name VCC-WLHost30-HBA0 pwwn 20:00:00:25:b5:aa:17:3c
device-alias name VCC-WLHost30-HBA2 pwwn 20:00:00:25:b5:aa:17:3d

device-alias commit

fcdomain fcid database

vsan 100 wwn 20:03:00:de:fb:92:8d:00 fcid 0x300000 dynamic

vsan 100 wwn 52:4a:93:75:dd:91:0a:02 fcid 0x300020 dynamic

! [X70-CT0-FC2]

vsan 100 wwn 52:4a:93:75:dd:91:0a:17 fcid 0x300040 dynamic

vsan 100 wwn 52:4a:93:75:dd:91:0a:06 fcid 0x300041 dynamic

! [X70-CT0-FC8]

vsan 100 wwn 52:4a:93:75:dd:91:0a:07 fcid 0x300042 dynamic

vsan 100 wwn 52:4a:93:75:dd:91:0a:16 fcid 0x300043 dynamic

! [X70-CT1-FC8]

vsan 100 wwn 20:00:00:25:b5:aa:17:3e fcid 0x300060 dynamic

! [VCC-Infra02-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:07 fcid 0x300061 dynamic

! [VCC-WLHost04-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:06 fcid 0x300062 dynamic

! [VCC-WLHost04-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:3a fcid 0x300063 dynamic

! [VCC-WLHost29-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:29 fcid 0x300064 dynamic

! [VCC-WLHost20-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:13 fcid 0x300065 dynamic

! [VCC-WLHost10-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:1c fcid 0x300066 dynamic
! [VCC-WLHost15-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:32 fcid 0x300067 dynamic
! [VCC-WLHost25-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:17 fcid 0x300068 dynamic
! [VCC-WLHost12-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:2e fcid 0x300069 dynamic
! [VCC-WLHost23-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:1f fcid 0x30006a dynamic
! [VCC-Infra01-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:1b fcid 0x30006b dynamic
! [VCC-WLHost14-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:1a fcid 0x30006c dynamic
! [VCC-WLHost14-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:0a fcid 0x30006d dynamic
! [VCC-WLHost06-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:34 fcid 0x30006e dynamic
! [VCC-WLHost26-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:19 fcid 0x30006f dynamic
! [VCC-WLHost13-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:36 fcid 0x300070 dynamic
! [VCC-WLHost27-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:01 fcid 0x300071 dynamic
! [VCC-WLHost01-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:12 fcid 0x300072 dynamic
! [VCC-WLHost10-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:16 fcid 0x300073 dynamic
! [VCC-WLHost12-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:2b fcid 0x300074 dynamic
! [VCC-WLHost21-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:25 fcid 0x300075 dynamic
! [VCC-WLHost18-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:27 fcid 0x300076 dynamic
! [VCC-WLHost19-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:3d fcid 0x300077 dynamic
! [VCC-WLHost30-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:15 fcid 0x300078 dynamic
! [VCC-WLHost11-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:38 fcid 0x300079 dynamic
! [VCC-WLHost28-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:23 fcid 0x30007a dynamic
! [VCC-WLHost17-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:00 fcid 0x30007b dynamic
! [VCC-WLHost01-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:04 fcid 0x30007c dynamic
! [VCC-WLHost03-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:03 fcid 0x30007d dynamic
! [VCC-WLHost02-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:0f fcid 0x30007e dynamic

! [VCC-WLHost08-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:1d fcid 0x30007f dynamic

! [VCC-WLHost15-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:31 fcid 0x300080 dynamic

! [VCC-WLHost24-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:30 fcid 0x300081 dynamic

! [VCC-WLHost24-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:02 fcid 0x300082 dynamic

! [VCC-WLHost02-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:08 fcid 0x300083 dynamic

! [VCC-WLHost05-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:26 fcid 0x300084 dynamic

! [VCC-WLHost19-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:22 fcid 0x300085 dynamic

! [VCC-WLHost17-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:2c fcid 0x300086 dynamic

! [VCC-WLHost22-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:33 fcid 0x300087 dynamic

! [VCC-WLHost25-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:21 fcid 0x300088 dynamic

! [VCC-WLHost16-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:2d fcid 0x300089 dynamic

! [VCC-WLHost22-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:24 fcid 0x30008a dynamic
! [VCC-WLHost18-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:3f fcid 0x30008b dynamic
! [VCC-Infra02-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:39 fcid 0x30008c dynamic
! [VCC-WLHost28-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:3c fcid 0x30008d dynamic
! [VCC-WLHost30-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:14 fcid 0x30008e dynamic
! [VCC-WLHost11-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:11 fcid 0x30008f dynamic
! [VCC-WLHost09-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:10 fcid 0x300090 dynamic
! [VCC-WLHost09-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:05 fcid 0x300091 dynamic
! [VCC-WLHost03-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:0e fcid 0x300092 dynamic
! [VCC-WLHost08-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:0d fcid 0x300093 dynamic
! [VCC-WLHost07-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:0c fcid 0x300094 dynamic
! [VCC-WLHost07-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:1e fcid 0x300095 dynamic
! [VCC-Infra01-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:0b fcid 0x300096 dynamic
! [VCC-WLHost06-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:28 fcid 0x300097 dynamic
! [VCC-WLHost20-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:37 fcid 0x300098 dynamic
! [VCC-WLHost27-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:3b fcid 0x300099 dynamic
! [VCC-WLHost29-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:09 fcid 0x30009a dynamic
! [VCC-WLHost05-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:2a fcid 0x30009b dynamic
! [VCC-WLHost21-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:2f fcid 0x30009c dynamic
! [VCC-WLHost23-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:20 fcid 0x30009d dynamic
! [VCC-WLHost16-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:35 fcid 0x30009e dynamic
! [VCC-WLHost26-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:18 fcid 0x30009f dynamic
! [VCC-WLHost13-HBA0]

vsan 100 wwn 20:02:00:de:fb:92:8d:00 fcid 0x3000a0 dynamic

vsan 100 wwn 20:04:00:de:fb:92:8d:00 fcid 0x3000c0 dynamic

vsan 100 wwn 20:01:00:de:fb:92:8d:00 fcid 0x3000e0 dynamic

vsan 100 wwn 52:4a:93:75:dd:91:0a:00 fcid 0x300044 dynamic

! [X70-CT0-FC0]

vsan 100 wwn 20:01:00:3a:9c:0e:33:20 fcid 0x3000e1 dynamic

vsan 100 wwn 20:02:00:3a:9c:0e:33:20 fcid 0x3000a1 dynamic

vsan 100 wwn 20:04:00:3a:9c:0e:33:20 fcid 0x3000c1 dynamic

vsan 100 wwn 20:03:00:3a:9c:0e:33:20 fcid 0x300100 dynamic

vsan 100 wwn 52:4a:93:75:dd:91:0a:10 fcid 0x300021 dynamic

! [X70-CT1-FC0]

vsan 100 wwn 52:4a:93:71:56:84:09:12 fcid 0x300022 dynamic

vsan 100 wwn 52:4a:93:71:56:84:09:10 fcid 0x300045 dynamic

! [X70R3-CT1-FC0]

vsan 100 wwn 52:4a:93:71:56:84:09:02 fcid 0x300046 dynamic

vsan 100 wwn 52:4a:93:71:56:84:09:00 fcid 0x300023 dynamic

! [X70R3-CT0-FC0]

vsan 100 wwn 20:00:00:25:b5:aa:17:40 fcid 0x3000e2 dynamic

! [AMD-VMHost70-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:41 fcid 0x3000a2 dynamic

! [AMD-VMHost70-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:44 fcid 0x3000e3 dynamic

! [AMD-VMHost72-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:45 fcid 0x3000a3 dynamic

! [AMD-VMHost72-HBA2]

vsan 100 wwn 20:00:00:25:b5:aa:17:4e fcid 0x3000e4 dynamic

! [AMD-VMHost73-HBA0]

vsan 100 wwn 20:00:00:25:b5:aa:17:4f fcid 0x3000a4 dynamic

! [AMD-VMHost73-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:42 fcid 0x3000e5 dynamic

! [AMD-VMHost71-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:43 fcid 0x3000a5 dynamic

! [AMD-VMHost71-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:46 fcid 0x3000e6 dynamic

! [AMD-VMHost74-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:47 fcid 0x3000a6 dynamic

! [AMD-VMHost74-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:48 fcid 0x3000e7 dynamic

! [AMD-VMHost75-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:49 fcid 0x3000a7 dynamic

! [AMD-VMHost75-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:4a fcid 0x3000e8 dynamic

! [AMD-VMHost76-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:4b fcid 0x3000a8 dynamic

! [AMD-VMHost76-HBA2]
vsan 100 wwn 20:00:00:25:b5:aa:17:4c fcid 0x3000e9 dynamic

! [AMD-VMHost77-HBA0]
vsan 100 wwn 20:00:00:25:b5:aa:17:4d fcid 0x3000a9 dynamic

! [AMD-VMHost77-HBA2]

!Active Zone Database Section for vsan 100

zone name FlaskStack-VCC-CVD-WLHost01 vsan 100

member pwwn 20:00:00:25:b5:aa:17:00

! [VCC-WLHost01-HBA0]

member pwwn 20:00:00:25:b5:aa:17:01

! [VCC-WLHost01-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 100

member pwwn 20:00:00:25:b5:aa:17:02

! [VCC-WLHost02-HBA0]

member pwwn 20:00:00:25:b5:aa:17:03

! [VCC-WLHost02-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 100

member pwwn 20:00:00:25:b5:aa:17:04

! [VCC-WLHost03-HBA0]

member pwwn 20:00:00:25:b5:aa:17:05

! [VCC-WLHost03-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 100

member pwwn 20:00:00:25:b5:aa:17:06

! [VCC-WLHost04-HBA0]

member pwwn 20:00:00:25:b5:aa:17:07

! [VCC-WLHost04-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 100

member pwwn 20:00:00:25:b5:aa:17:08

! [VCC-WLHost05-HBA0]

member pwwn 20:00:00:25:b5:aa:17:09

! [VCC-WLHost05-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 100

member pwwn 20:00:00:25:b5:aa:17:0a

! [VCC-WLHost06-HBA0]

member pwwn 20:00:00:25:b5:aa:17:0b

! [VCC-WLHost06-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 100

member pwwn 20:00:00:25:b5:aa:17:0c

! [VCC-WLHost07-HBA0]

member pwwn 20:00:00:25:b5:aa:17:0d

! [VCC-WLHost07-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 100

member pwwn 20:00:00:25:b5:aa:17:0e

! [VCC-WLHost08-HBA0]

member pwwn 20:00:00:25:b5:aa:17:0f

! [VCC-WLHost08-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 100

member pwwn 20:00:00:25:b5:aa:17:10

! [VCC-WLHost09-HBA0]

member pwwn 20:00:00:25:b5:aa:17:11

! [VCC-WLHost09-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 100

member pwwn 20:00:00:25:b5:aa:17:12

! [VCC-WLHost10-HBA0]

member pwwn 20:00:00:25:b5:aa:17:13

! [VCC-WLHost10-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 100

member pwwn 20:00:00:25:b5:aa:17:14

! [VCC-WLHost11-HBA0]

member pwwn 20:00:00:25:b5:aa:17:15

! [VCC-WLHost11-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 100

member pwwn 20:00:00:25:b5:aa:17:16

! [VCC-WLHost12-HBA0]

member pwwn 20:00:00:25:b5:aa:17:17

! [VCC-WLHost12-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 100

member pwwn 20:00:00:25:b5:aa:17:18

! [VCC-WLHost13-HBA0]

member pwwn 20:00:00:25:b5:aa:17:19

! [VCC-WLHost13-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 100

member pwwn 20:00:00:25:b5:aa:17:1a

! [VCC-WLHost14-HBA0]

member pwwn 20:00:00:25:b5:aa:17:1b

! [VCC-WLHost14-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 100

member pwwn 20:00:00:25:b5:aa:17:1c

! [VCC-WLHost15-HBA0]

member pwwn 20:00:00:25:b5:aa:17:1d

! [VCC-WLHost15-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra01 vsan 100

member pwwn 20:00:00:25:b5:aa:17:1e

! [VCC-Infra01-HBA0]

member pwwn 20:00:00:25:b5:aa:17:1f

! [VCC-Infra01-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 100

member pwwn 20:00:00:25:b5:aa:17:20

! [VCC-WLHost16-HBA0]

member pwwn 20:00:00:25:b5:aa:17:21

! [VCC-WLHost16-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 100

member pwwn 20:00:00:25:b5:aa:17:22

! [VCC-WLHost17-HBA0]

member pwwn 20:00:00:25:b5:aa:17:23

! [VCC-WLHost17-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 100

member pwwn 20:00:00:25:b5:aa:17:24

! [VCC-WLHost18-HBA0]

member pwwn 20:00:00:25:b5:aa:17:25

! [VCC-WLHost18-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 100

member pwwn 20:00:00:25:b5:aa:17:26

! [VCC-WLHost19-HBA0]

member pwwn 20:00:00:25:b5:aa:17:27

! [VCC-WLHost19-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 100

member pwwn 20:00:00:25:b5:aa:17:28

! [VCC-WLHost20-HBA0]

member pwwn 20:00:00:25:b5:aa:17:29

! [VCC-WLHost20-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 100

member pwwn 20:00:00:25:b5:aa:17:2a

! [VCC-WLHost21-HBA0]

member pwwn 20:00:00:25:b5:aa:17:2b

! [VCC-WLHost21-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 100

member pwwn 20:00:00:25:b5:aa:17:2c

! [VCC-WLHost22-HBA0]

member pwwn 20:00:00:25:b5:aa:17:2d

! [VCC-WLHost22-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 100

member pwwn 20:00:00:25:b5:aa:17:2e

! [VCC-WLHost23-HBA0]

member pwwn 20:00:00:25:b5:aa:17:2f

! [VCC-WLHost23-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 100

member pwwn 20:00:00:25:b5:aa:17:30

! [VCC-WLHost24-HBA0]

member pwwn 20:00:00:25:b5:aa:17:31

! [VCC-WLHost24-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 100

member pwwn 20:00:00:25:b5:aa:17:32

! [VCC-WLHost25-HBA0]

member pwwn 20:00:00:25:b5:aa:17:33

! [VCC-WLHost25-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 100

member pwwn 20:00:00:25:b5:aa:17:34

! [VCC-WLHost26-HBA0]

member pwwn 20:00:00:25:b5:aa:17:35

! [VCC-WLHost26-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 100

member pwwn 20:00:00:25:b5:aa:17:36

! [VCC-WLHost27-HBA0]

member pwwn 20:00:00:25:b5:aa:17:37

! [VCC-WLHost27-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 100

member pwwn 20:00:00:25:b5:aa:17:38

! [VCC-WLHost28-HBA0]

member pwwn 20:00:00:25:b5:aa:17:39

! [VCC-WLHost28-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 100

member pwwn 20:00:00:25:b5:aa:17:3a

! [VCC-WLHost29-HBA0]

member pwwn 20:00:00:25:b5:aa:17:3b

! [VCC-WLHost29-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 100

member pwwn 20:00:00:25:b5:aa:17:3c

! [VCC-WLHost30-HBA0]

member pwwn 20:00:00:25:b5:aa:17:3d

! [VCC-WLHost30-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra02 vsan 100

member pwwn 20:00:00:25:b5:aa:17:3e

! [VCC-Infra02-HBA0]

member pwwn 20:00:00:25:b5:aa:17:3f

! [VCC-Infra02-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost70 vsan 100

member pwwn 20:00:00:25:b5:aa:17:40

! [AMD-VMHost70-HBA0]

member pwwn 20:00:00:25:b5:aa:17:41

! [AMD-VMHost70-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost71 vsan 100

member pwwn 20:00:00:25:b5:aa:17:42

! [AMD-VMHost71-HBA0]

member pwwn 20:00:00:25:b5:aa:17:43

! [AMD-VMHost71-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost72 vsan 100

member pwwn 20:00:00:25:b5:aa:17:44

! [AMD-VMHost72-HBA0]

member pwwn 20:00:00:25:b5:aa:17:45

! [AMD-VMHost72-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost73 vsan 100

member pwwn 20:00:00:25:b5:aa:17:4e

! [AMD-VMHost73-HBA0]

member pwwn 20:00:00:25:b5:aa:17:4f

! [AMD-VMHost73-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost74 vsan 100

member pwwn 20:00:00:25:b5:aa:17:46

! [AMD-VMHost74-HBA0]

member pwwn 20:00:00:25:b5:aa:17:47

! [AMD-VMHost74-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost75 vsan 100

member pwwn 20:00:00:25:b5:aa:17:48

! [AMD-VMHost75-HBA0]

member pwwn 20:00:00:25:b5:aa:17:49

! [AMD-VMHost75-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost76 vsan 100

member pwwn 20:00:00:25:b5:aa:17:4a

! [AMD-VMHost76-HBA0]

member pwwn 20:00:00:25:b5:aa:17:4b

! [AMD-VMHost76-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost77 vsan 100

member pwwn 20:00:00:25:b5:aa:17:4c

! [AMD-VMHost77-HBA0]

member pwwn 20:00:00:25:b5:aa:17:4d

! [AMD-VMHost77-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zoneset name FlashStack-VCC-CVD vsan 100

member FlaskStack-VCC-CVD-WLHost01

member FlaskStack-VCC-CVD-WLHost02

member FlaskStack-VCC-CVD-WLHost03

member FlaskStack-VCC-CVD-WLHost04

member FlaskStack-VCC-CVD-WLHost05

member FlaskStack-VCC-CVD-WLHost06

member FlaskStack-VCC-CVD-WLHost07

member FlaskStack-VCC-CVD-WLHost08

member FlaskStack-VCC-CVD-WLHost09
member FlaskStack-VCC-CVD-WLHost10
member FlaskStack-VCC-CVD-WLHost11
member FlaskStack-VCC-CVD-WLHost12
member FlaskStack-VCC-CVD-WLHost13
member FlaskStack-VCC-CVD-WLHost14
member FlaskStack-VCC-CVD-WLHost15
member FlaskStack-VCC-CVD-Infra01
member FlaskStack-VCC-CVD-WLHost16
member FlaskStack-VCC-CVD-WLHost17
member FlaskStack-VCC-CVD-WLHost18
member FlaskStack-VCC-CVD-WLHost19
member FlaskStack-VCC-CVD-WLHost20
member FlaskStack-VCC-CVD-WLHost21
member FlaskStack-VCC-CVD-WLHost22
member FlaskStack-VCC-CVD-WLHost23
member FlaskStack-VCC-CVD-WLHost24
member FlaskStack-VCC-CVD-WLHost25
member FlaskStack-VCC-CVD-WLHost26
member FlaskStack-VCC-CVD-WLHost27
member FlaskStack-VCC-CVD-WLHost28
member FlaskStack-VCC-CVD-WLHost29
member FlaskStack-VCC-CVD-WLHost30
member FlaskStack-VCC-CVD-Infra02

member FlaskStack-AMD-VMHost70

member FlaskStack-AMD-VMHost71

member FlaskStack-AMD-VMHost72

member FlaskStack-AMD-VMHost73

member FlaskStack-AMD-VMHost74

member FlaskStack-AMD-VMHost75

member FlaskStack-AMD-VMHost76

member FlaskStack-AMD-VMHost77

zoneset activate name FlashStack-VCC-CVD vsan 100

do clear zone database vsan 100

!Full Zone Database Section for vsan 100

zone name FlaskStack-VCC-CVD-WLHost01 vsan 100

member pwwn 20:00:00:25:b5:aa:17:00

! [VCC-WLHost01-HBA0]

member pwwn 20:00:00:25:b5:aa:17:01

! [VCC-WLHost01-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 100

member pwwn 20:00:00:25:b5:aa:17:02

! [VCC-WLHost02-HBA0]
member pwwn 20:00:00:25:b5:aa:17:03
!
! [VCC-WLHost02-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 100

member pwwn 20:00:00:25:b5:aa:17:04
!
! [VCC-WLHost03-HBA0]
member pwwn 20:00:00:25:b5:aa:17:05
!
! [VCC-WLHost03-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 100

member pwwn 20:00:00:25:b5:aa:17:06
!
! [VCC-WLHost04-HBA0]
member pwwn 20:00:00:25:b5:aa:17:07
!
! [VCC-WLHost04-HBA2]
member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 100

member pwwn 20:00:00:25:b5:aa:17:08

! [VCC-WLHost05-HBA0]

member pwwn 20:00:00:25:b5:aa:17:09

! [VCC-WLHost05-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 100

member pwwn 20:00:00:25:b5:aa:17:0a

! [VCC-WLHost06-HBA0]

member pwwn 20:00:00:25:b5:aa:17:0b

! [VCC-WLHost06-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 100

member pwwn 20:00:00:25:b5:aa:17:0c

! [VCC-WLHost07-HBA0]

member pwwn 20:00:00:25:b5:aa:17:0d

! [VCC-WLHost07-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 100

member pwwn 20:00:00:25:b5:aa:17:0e

! [VCC-WLHost08-HBA0]

member pwwn 20:00:00:25:b5:aa:17:0f

! [VCC-WLHost08-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 100

member pwwn 20:00:00:25:b5:aa:17:10

! [VCC-WLHost09-HBA0]

member pwwn 20:00:00:25:b5:aa:17:11

! [VCC-WLHost09-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 100

member pwwn 20:00:00:25:b5:aa:17:12
!
! [VCC-WLHost10-HBA0]
member pwwn 20:00:00:25:b5:aa:17:13
!
! [VCC-WLHost10-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 100

member pwwn 20:00:00:25:b5:aa:17:14
!
! [VCC-WLHost11-HBA0]
member pwwn 20:00:00:25:b5:aa:17:15
!
! [VCC-WLHost11-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 100

member pwwn 20:00:00:25:b5:aa:17:16

! [VCC-WLHost12-HBA0]

member pwwn 20:00:00:25:b5:aa:17:17

! [VCC-WLHost12-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 100

member pwwn 20:00:00:25:b5:aa:17:18

! [VCC-WLHost13-HBA0]

member pwwn 20:00:00:25:b5:aa:17:19

! [VCC-WLHost13-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 100

member pwwn 20:00:00:25:b5:aa:17:1a

! [VCC-WLHost14-HBA0]
member pwwn 20:00:00:25:b5:aa:17:1b
!
! [VCC-WLHost14-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 100

member pwwn 20:00:00:25:b5:aa:17:1c
!
! [VCC-WLHost15-HBA0]
member pwwn 20:00:00:25:b5:aa:17:1d
!
! [VCC-WLHost15-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra01 vsan 100

member pwwn 20:00:00:25:b5:aa:17:1e
!
! [VCC-Infra01-HBA0]
member pwwn 20:00:00:25:b5:aa:17:1f
!
! [VCC-Infra01-HBA2]
member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 100

member pwwn 20:00:00:25:b5:aa:17:20

! [VCC-WLHost16-HBA0]

member pwwn 20:00:00:25:b5:aa:17:21

! [VCC-WLHost16-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 100

member pwwn 20:00:00:25:b5:aa:17:22

! [VCC-WLHost17-HBA0]

member pwwn 20:00:00:25:b5:aa:17:23

! [VCC-WLHost17-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 100

member pwwn 20:00:00:25:b5:aa:17:24

! [VCC-WLHost18-HBA0]

member pwwn 20:00:00:25:b5:aa:17:25

! [VCC-WLHost18-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 100

member pwwn 20:00:00:25:b5:aa:17:26

! [VCC-WLHost19-HBA0]

member pwwn 20:00:00:25:b5:aa:17:27

! [VCC-WLHost19-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 100

member pwwn 20:00:00:25:b5:aa:17:28

! [VCC-WLHost20-HBA0]

member pwwn 20:00:00:25:b5:aa:17:29

! [VCC-WLHost20-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 100

member pwwn 20:00:00:25:b5:aa:17:2a
!
! [VCC-WLHost21-HBA0]
member pwwn 20:00:00:25:b5:aa:17:2b
!
! [VCC-WLHost21-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 100

member pwwn 20:00:00:25:b5:aa:17:2c
!
! [VCC-WLHost22-HBA0]
member pwwn 20:00:00:25:b5:aa:17:2d
!
! [VCC-WLHost22-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 100

member pwwn 20:00:00:25:b5:aa:17:2e

! [VCC-WLHost23-HBA0]

member pwwn 20:00:00:25:b5:aa:17:2f

! [VCC-WLHost23-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 100

member pwwn 20:00:00:25:b5:aa:17:30

! [VCC-WLHost24-HBA0]

member pwwn 20:00:00:25:b5:aa:17:31

! [VCC-WLHost24-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 100

member pwwn 20:00:00:25:b5:aa:17:32

! [VCC-WLHost25-HBA0]
member pwwn 20:00:00:25:b5:aa:17:33
!
! [VCC-WLHost25-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 100

member pwwn 20:00:00:25:b5:aa:17:34
!
! [VCC-WLHost26-HBA0]
member pwwn 20:00:00:25:b5:aa:17:35
!
! [VCC-WLHost26-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 100

member pwwn 20:00:00:25:b5:aa:17:36
!
! [VCC-WLHost27-HBA0]
member pwwn 20:00:00:25:b5:aa:17:37
!
! [VCC-WLHost27-HBA2]
member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 100

member pwwn 20:00:00:25:b5:aa:17:38

! [VCC-WLHost28-HBA0]

member pwwn 20:00:00:25:b5:aa:17:39

! [VCC-WLHost28-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 100

member pwwn 20:00:00:25:b5:aa:17:3a

! [VCC-WLHost29-HBA0]

member pwwn 20:00:00:25:b5:aa:17:3b

! [VCC-WLHost29-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 100

member pwwn 20:00:00:25:b5:aa:17:3c

! [VCC-WLHost30-HBA0]

member pwwn 20:00:00:25:b5:aa:17:3d

! [VCC-WLHost30-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-VCC-CVD-Infra02 vsan 100

member pwwn 20:00:00:25:b5:aa:17:3e

! [VCC-Infra02-HBA0]

member pwwn 20:00:00:25:b5:aa:17:3f

! [VCC-Infra02-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost70 vsan 100

member pwwn 20:00:00:25:b5:aa:17:40

! [AMD-VMHost70-HBA0]

member pwwn 20:00:00:25:b5:aa:17:41

! [AMD-VMHost70-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost71 vsan 100

member pwwn 20:00:00:25:b5:aa:17:42
!
! [AMD-VMHost71-HBA0]
member pwwn 20:00:00:25:b5:aa:17:43
!
! [AMD-VMHost71-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
!
! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost72 vsan 100

member pwwn 20:00:00:25:b5:aa:17:44
!
! [AMD-VMHost72-HBA0]
member pwwn 20:00:00:25:b5:aa:17:45
!
! [AMD-VMHost72-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
!
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost73 vsan 100

member pwwn 20:00:00:25:b5:aa:17:4e

! [AMD-VMHost73-HBA0]

member pwwn 20:00:00:25:b5:aa:17:4f

! [AMD-VMHost73-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost74 vsan 100

member pwwn 20:00:00:25:b5:aa:17:46

! [AMD-VMHost74-HBA0]

member pwwn 20:00:00:25:b5:aa:17:47

! [AMD-VMHost74-HBA2]

member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost75 vsan 100

member pwwn 20:00:00:25:b5:aa:17:48

! [AMD-VMHost75-HBA0]
member pwwn 20:00:00:25:b5:aa:17:49
! [AMD-VMHost75-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost76 vsan 100

member pwwn 20:00:00:25:b5:aa:17:4a
! [AMD-VMHost76-HBA0]
member pwwn 20:00:00:25:b5:aa:17:4b
! [AMD-VMHost76-HBA2]
member pwwn 52:4a:93:71:56:84:09:00
! [X70R3-CT0-FC0]
member pwwn 52:4a:93:71:56:84:09:10
! [X70R3-CT1-FC0]

zone name FlaskStack-AMD-VMHost77 vsan 100

member pwwn 20:00:00:25:b5:aa:17:4c
! [AMD-VMHost77-HBA0]
member pwwn 20:00:00:25:b5:aa:17:4d
! [AMD-VMHost77-HBA2]
member pwwn 52:4a:93:71:56:84:09:00

! [X70R3-CT0-FC0]

member pwwn 52:4a:93:71:56:84:09:10

! [X70R3-CT1-FC0]

zoneset name FlashStack-VCC-CVD vsan 100

member FlaskStack-VCC-CVD-WLHost01

member FlaskStack-VCC-CVD-WLHost02

member FlaskStack-VCC-CVD-WLHost03

member FlaskStack-VCC-CVD-WLHost04

member FlaskStack-VCC-CVD-WLHost05

member FlaskStack-VCC-CVD-WLHost06

member FlaskStack-VCC-CVD-WLHost07

member FlaskStack-VCC-CVD-WLHost08

member FlaskStack-VCC-CVD-WLHost09

member FlaskStack-VCC-CVD-WLHost10

member FlaskStack-VCC-CVD-WLHost11

member FlaskStack-VCC-CVD-WLHost12

member FlaskStack-VCC-CVD-WLHost13

member FlaskStack-VCC-CVD-WLHost14

member FlaskStack-VCC-CVD-WLHost15

member FlaskStack-VCC-CVD-Infra01

member FlaskStack-VCC-CVD-WLHost16

member FlaskStack-VCC-CVD-WLHost17

member FlaskStack-VCC-CVD-WLHost18

member FlaskStack-VCC-CVD-WLHost19
member FlaskStack-VCC-CVD-WLHost20
member FlaskStack-VCC-CVD-WLHost21
member FlaskStack-VCC-CVD-WLHost22
member FlaskStack-VCC-CVD-WLHost23
member FlaskStack-VCC-CVD-WLHost24
member FlaskStack-VCC-CVD-WLHost25
member FlaskStack-VCC-CVD-WLHost26
member FlaskStack-VCC-CVD-WLHost27
member FlaskStack-VCC-CVD-WLHost28
member FlaskStack-VCC-CVD-WLHost29
member FlaskStack-VCC-CVD-WLHost30
member FlaskStack-VCC-CVD-Infra02
member FlaskStack-AMD-VMHost70
member FlaskStack-AMD-VMHost71
member FlaskStack-AMD-VMHost72
member FlaskStack-AMD-VMHost73
member FlaskStack-AMD-VMHost74
member FlaskStack-AMD-VMHost75
member FlaskStack-AMD-VMHost76
member FlaskStack-AMD-VMHost77

interface mgmt0

ip address 10.29.164.238 255.255.255.0

vsan database

vsan 400 interface fc1/1

vsan 400 interface fc1/2

vsan 400 interface fc1/3

vsan 400 interface fc1/4

vsan 400 interface fc1/5

vsan 400 interface fc1/6

vsan 400 interface fc1/7

vsan 400 interface fc1/8

vsan 100 interface fc1/9

vsan 100 interface fc1/10

vsan 100 interface fc1/11

vsan 100 interface fc1/12

vsan 100 interface fc1/13

vsan 100 interface fc1/14

vsan 100 interface fc1/15

vsan 100 interface fc1/16

clock timezone PST 0 0

clock summer-time PDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60

switchname ADD16-MDS-A

cli alias name autozone source sys/autozone.py

line console

line vty

boot kickstart bootflash:/m9100-s6ek9-kickstart-mz.8.3.1.bin

boot system bootflash:/m9100-s6ek9-mz.8.3.1.bin

interface fc1/4

switchport speed auto

interface fc1/1

interface fc1/2

interface fc1/3

interface fc1/5

interface fc1/6

interface fc1/7

interface fc1/8

interface fc1/9

interface fc1/10

interface fc1/11

interface fc1/12

interface fc1/13

interface fc1/14

interface fc1/15

interface fc1/16

interface fc1/4

interface fc1/1

port-license acquire

no shutdown

interface fc1/2

port-license acquire

no shutdown

interface fc1/3

port-license acquire

no shutdown

interface fc1/4

port-license acquire

no shutdown

interface fc1/5

no port-license

interface fc1/6

no port-license

interface fc1/7

no port-license

interface fc1/8

no port-license

interface fc1/9

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/10

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/11

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/12

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/13

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/14

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/15

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/16

switchport trunk allowed vsan 100

switchport trunk mode off

port-license acquire

no shutdown

ip default-gateway 10.29.164.1

Cisco MDS 9132T-B Configuration

```
version 8.3(1)

power redundancy-mode redundant

feature npiv

feature fport-channel-trunk

role name default-role

    description This is a system defined role and applies to all users.

    rule 5 permit show feature environment

    rule 4 permit show feature hardware

    rule 3 permit show feature module

    rule 2 permit show feature snmp

    rule 1 permit show feature system

no password strength-check

username admin password 5 $5$1qs42bIH$hp2kMO3FA/4Zzg6EekVHWpA8IA7Mc/kBsFZVU8q1uU7
role network-admin

ip domain-lookup

ip host ADD16-MDS-B 10.29.164.239

aaa group server radius radius

snmp-server user admin network-admin auth md5 0x6fa97f514b0cdf3638e31dfd0bd19c71 priv
0x6fa97f514b0cdf3638e31dfd0bd19c71 localizedkey

snmp-server host 10.155.160.97 traps version 2c public udp-port 1164

snmp-server host 10.24.66.169 traps version 2c public udp-port 1164

snmp-server host 10.24.72.119 traps version 2c public udp-port 1166

snmp-server host 10.29.164.250 traps version 2c public udp-port 1163

rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
```

rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 log trap public description ERROR(3) owner PMON@ERROR

rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

ntp server 10.81.254.131

ntp server 10.81.254.202

vsan database

vsan 101 name "FlashStack-VCC-CVD-Fabric-B"

device-alias database

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

device-alias name VCC-Infra01-HBA1 pwwn 20:00:00:25:b5:bb:17:1e

device-alias name VCC-Infra01-HBA3 pwwn 20:00:00:25:b5:bb:17:1f

device-alias name VCC-Infra02-HBA1 pwwn 20:00:00:25:b5:bb:17:3e

device-alias name VCC-Infra02-HBA3 pwwn 20:00:00:25:b5:bb:17:3f

device-alias name VCC-WLHost01-HBA1 pwwn 20:00:00:25:b5:bb:17:00

device-alias name VCC-WLHost01-HBA3 pwwn 20:00:00:25:b5:bb:17:01

device-alias name VCC-WLHost02-HBA1 pwwn 20:00:00:25:b5:bb:17:02

device-alias name VCC-WLHost02-HBA3 pwwn 20:00:00:25:b5:bb:17:03

device-alias name VCC-WLHost03-HBA1 pwwn 20:00:00:25:b5:bb:17:04

device-alias name VCC-WLHost03-HBA3 pwwn 20:00:00:25:b5:bb:17:05

device-alias name VCC-WLHost04-HBA1 pwwn 20:00:00:25:b5:bb:17:06

device-alias name VCC-WLHost04-HBA3 pwwn 20:00:00:25:b5:bb:17:07

device-alias name VCC-WLHost05-HBA1 pwwn 20:00:00:25:b5:bb:17:08

device-alias name VCC-WLHost05-HBA3 pwwn 20:00:00:25:b5:bb:17:09
device-alias name VCC-WLHost06-HBA1 pwwn 20:00:00:25:b5:bb:17:0a
device-alias name VCC-WLHost06-HBA3 pwwn 20:00:00:25:b5:bb:17:0b
device-alias name VCC-WLHost07-HBA1 pwwn 20:00:00:25:b5:bb:17:0c
device-alias name VCC-WLHost07-HBA3 pwwn 20:00:00:25:b5:bb:17:0d
device-alias name VCC-WLHost08-HBA1 pwwn 20:00:00:25:b5:bb:17:0e
device-alias name VCC-WLHost08-HBA3 pwwn 20:00:00:25:b5:bb:17:0f
device-alias name VCC-WLHost09-HBA1 pwwn 20:00:00:25:b5:bb:17:10
device-alias name VCC-WLHost09-HBA3 pwwn 20:00:00:25:b5:bb:17:11
device-alias name VCC-WLHost10-HBA1 pwwn 20:00:00:25:b5:bb:17:12
device-alias name VCC-WLHost10-HBA3 pwwn 20:00:00:25:b5:bb:17:13
device-alias name VCC-WLHost11-HBA1 pwwn 20:00:00:25:b5:bb:17:14
device-alias name VCC-WLHost11-HBA3 pwwn 20:00:00:25:b5:bb:17:15
device-alias name VCC-WLHost12-HBA1 pwwn 20:00:00:25:b5:bb:17:16
device-alias name VCC-WLHost12-HBA3 pwwn 20:00:00:25:b5:bb:17:17
device-alias name VCC-WLHost13-HBA1 pwwn 20:00:00:25:b5:bb:17:18
device-alias name VCC-WLHost13-HBA3 pwwn 20:00:00:25:b5:bb:17:19
device-alias name VCC-WLHost14-HBA1 pwwn 20:00:00:25:b5:bb:17:1a
device-alias name VCC-WLHost14-HBA3 pwwn 20:00:00:25:b5:bb:17:1b
device-alias name VCC-WLHost15-HBA1 pwwn 20:00:00:25:b5:bb:17:1c
device-alias name VCC-WLHost15-HBA3 pwwn 20:00:00:25:b5:bb:17:1d
device-alias name VCC-WLHost16-HBA1 pwwn 20:00:00:25:b5:bb:17:20
device-alias name VCC-WLHost16-HBA3 pwwn 20:00:00:25:b5:bb:17:21
device-alias name VCC-WLHost17-HBA1 pwwn 20:00:00:25:b5:bb:17:22

device-alias name VCC-WLHost17-HBA3 pwwn 20:00:00:25:b5:bb:17:23
device-alias name VCC-WLHost18-HBA1 pwwn 20:00:00:25:b5:bb:17:24
device-alias name VCC-WLHost18-HBA3 pwwn 20:00:00:25:b5:bb:17:25
device-alias name VCC-WLHost19-HBA1 pwwn 20:00:00:25:b5:bb:17:26
device-alias name VCC-WLHost19-HBA3 pwwn 20:00:00:25:b5:bb:17:27
device-alias name VCC-WLHost20-HBA1 pwwn 20:00:00:25:b5:bb:17:28
device-alias name VCC-WLHost20-HBA3 pwwn 20:00:00:25:b5:bb:17:29
device-alias name VCC-WLHost21-HBA1 pwwn 20:00:00:25:b5:bb:17:2a
device-alias name VCC-WLHost21-HBA3 pwwn 20:00:00:25:b5:bb:17:2b
device-alias name VCC-WLHost22-HBA1 pwwn 20:00:00:25:b5:bb:17:2c
device-alias name VCC-WLHost22-HBA3 pwwn 20:00:00:25:b5:bb:17:2d
device-alias name VCC-WLHost23-HBA1 pwwn 20:00:00:25:b5:bb:17:2e
device-alias name VCC-WLHost23-HBA3 pwwn 20:00:00:25:b5:bb:17:2f
device-alias name VCC-WLHost24-HBA1 pwwn 20:00:00:25:b5:bb:17:30
device-alias name VCC-WLHost24-HBA3 pwwn 20:00:00:25:b5:bb:17:31
device-alias name VCC-WLHost25-HBA1 pwwn 20:00:00:25:b5:bb:17:32
device-alias name VCC-WLHost25-HBA3 pwwn 20:00:00:25:b5:bb:17:33
device-alias name VCC-WLHost26-HBA1 pwwn 20:00:00:25:b5:bb:17:34
device-alias name VCC-WLHost26-HBA3 pwwn 20:00:00:25:b5:bb:17:35
device-alias name VCC-WLHost27-HBA1 pwwn 20:00:00:25:b5:bb:17:36
device-alias name VCC-WLHost27-HBA3 pwwn 20:00:00:25:b5:bb:17:37
device-alias name VCC-WLHost28-HBA1 pwwn 20:00:00:25:b5:bb:17:38
device-alias name VCC-WLHost28-HBA3 pwwn 20:00:00:25:b5:bb:17:39
device-alias name VCC-WLHost29-HBA1 pwwn 20:00:00:25:b5:bb:17:3a

device-alias name VCC-WLHost29-HBA3 pwwn 20:00:00:25:b5:bb:17:3b

device-alias name VCC-WLHost30-HBA1 pwwn 20:00:00:25:b5:bb:17:3c

device-alias name VCC-WLHost30-HBA3 pwwn 20:00:00:25:b5:bb:17:3d

device-alias commit

fcdomain fcid database

vsan 101 wwn 20:03:00:de:fb:90:a4:40 fcid 0xc40000 dynamic

vsan 101 wwn 52:4a:93:75:dd:91:0a:17 fcid 0xc40020 dynamic

! [X70-CT1-FC9]

vsan 101 wwn 52:4a:93:75:dd:91:0a:07 fcid 0xc40040 dynamic

! [X70-CT0-FC9]

vsan 101 wwn 52:4a:93:75:dd:91:0a:16 fcid 0xc40021 dynamic

vsan 101 wwn 52:4a:93:75:dd:91:0a:13 fcid 0xc40041 dynamic

! [X70-CT1-FC3]

vsan 101 wwn 20:00:00:25:b5:bb:17:3e fcid 0xc40060 dynamic

! [VCC-Infra02-HBA1]

vsan 101 wwn 20:00:00:25:b5:bb:17:07 fcid 0xc40061 dynamic

! [VCC-WLHost04-HBA3]

vsan 101 wwn 20:00:00:25:b5:bb:17:3c fcid 0xc40062 dynamic

! [VCC-WLHost30-HBA1]

vsan 101 wwn 20:00:00:25:b5:bb:17:11 fcid 0xc40063 dynamic

! [VCC-WLHost09-HBA3]

vsan 101 wwn 20:00:00:25:b5:bb:17:01 fcid 0xc40064 dynamic

! [VCC-WLHost01-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:00 fcid 0xc40065 dynamic

! [VCC-WLHost01-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:13 fcid 0xc40066 dynamic

! [VCC-WLHost10-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:04 fcid 0xc40067 dynamic

! [VCC-WLHost03-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:17 fcid 0xc40068 dynamic

! [VCC-WLHost12-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:16 fcid 0xc40069 dynamic

! [VCC-WLHost12-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:30 fcid 0xc4006a dynamic

! [VCC-WLHost24-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:21 fcid 0xc4006b dynamic

! [VCC-WLHost16-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:1f fcid 0xc4006c dynamic

! [VCC-Infra01-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:1a fcid 0xc4006d dynamic

! [VCC-WLHost14-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:3f fcid 0xc4006e dynamic

! [VCC-Infra02-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:0a fcid 0xc4006f dynamic

! [VCC-WLHost06-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:38 fcid 0xc40070 dynamic

! [VCC-WLHost28-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:19 fcid 0xc40071 dynamic

! [VCC-WLHost13-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:22 fcid 0xc40072 dynamic

! [VCC-WLHost17-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:2f fcid 0xc40073 dynamic

! [VCC-WLHost23-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:1b fcid 0xc40074 dynamic

! [VCC-WLHost14-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:3b fcid 0xc40075 dynamic

! [VCC-WLHost29-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:2a fcid 0xc40076 dynamic

! [VCC-WLHost21-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:29 fcid 0xc40077 dynamic

! [VCC-WLHost20-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:1c fcid 0xc40078 dynamic

! [VCC-WLHost15-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:0b fcid 0xc40079 dynamic

! [VCC-WLHost06-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:0d fcid 0xc4007a dynamic

! [VCC-WLHost07-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:37 fcid 0xc4007b dynamic

! [VCC-WLHost27-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:31 fcid 0xc4007c dynamic

! [VCC-WLHost24-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:08 fcid 0xc4007d dynamic

! [VCC-WLHost05-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:10 fcid 0xc4007e dynamic

! [VCC-WLHost09-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:34 fcid 0xc4007f dynamic

! [VCC-WLHost26-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:25 fcid 0xc40080 dynamic

! [VCC-WLHost18-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:3d fcid 0xc40081 dynamic

! [VCC-WLHost30-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:15 fcid 0xc40082 dynamic

! [VCC-WLHost11-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:23 fcid 0xc40083 dynamic

! [VCC-WLHost17-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:3a fcid 0xc40084 dynamic

! [VCC-WLHost29-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:28 fcid 0xc40085 dynamic

! [VCC-WLHost20-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:32 fcid 0xc40086 dynamic

! [VCC-WLHost25-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:0f fcid 0xc40087 dynamic

! [VCC-WLHost08-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:0c fcid 0xc40088 dynamic

! [VCC-WLHost07-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:2e fcid 0xc40089 dynamic

! [VCC-WLHost23-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:03 fcid 0xc4008a dynamic

! [VCC-WLHost02-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:02 fcid 0xc4008b dynamic

! [VCC-WLHost02-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:2b fcid 0xc4008c dynamic

! [VCC-WLHost21-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:35 fcid 0xc4008d dynamic

! [VCC-WLHost26-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:2c fcid 0xc4008e dynamic

! [VCC-WLHost22-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:27 fcid 0xc4008f dynamic

! [VCC-WLHost19-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:18 fcid 0xc40090 dynamic

! [VCC-WLHost13-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:14 fcid 0xc40091 dynamic

! [VCC-WLHost11-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:0e fcid 0xc40092 dynamic

! [VCC-WLHost08-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:1e fcid 0xc40093 dynamic

! [VCC-Infra01-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:06 fcid 0xc40094 dynamic

! [VCC-WLHost04-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:09 fcid 0xc40095 dynamic

! [VCC-WLHost05-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:26 fcid 0xc40096 dynamic

! [VCC-WLHost19-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:24 fcid 0xc40097 dynamic

! [VCC-WLHost18-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:20 fcid 0xc40098 dynamic

! [VCC-WLHost16-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:1d fcid 0xc40099 dynamic

! [VCC-WLHost15-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:33 fcid 0xc4009a dynamic

! [VCC-WLHost25-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:36 fcid 0xc4009b dynamic

! [VCC-WLHost27-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:39 fcid 0xc4009c dynamic

! [VCC-WLHost28-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:2d fcid 0xc4009d dynamic

! [VCC-WLHost22-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:12 fcid 0xc4009e dynamic

! [VCC-WLHost10-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:05 fcid 0xc4009f dynamic

! [VCC-WLHost03-HBA3]
vsan 101 wwn 20:02:00:de:fb:90:a4:40 fcid 0xc400a0 dynamic

vsan 101 wwn 20:01:00:de:fb:90:a4:40 fcid 0xc400c0 dynamic
vsan 101 wwn 20:04:00:de:fb:90:a4:40 fcid 0xc400e0 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:00 fcid 0xc40022 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:12 fcid 0xc40042 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:11 fcid 0xc40023 dynamic
! [X70-CT1-FC1]
vsan 101 wwn 20:01:00:3a:9c:a4:fd:20 fcid 0xc400c1 dynamic
vsan 101 wwn 20:02:00:3a:9c:a4:fd:20 fcid 0xc400a1 dynamic
vsan 101 wwn 20:03:00:3a:9c:a4:fd:20 fcid 0xc40100 dynamic
vsan 101 wwn 20:04:00:3a:9c:a4:fd:20 fcid 0xc400e1 dynamic
vsan 101 wwn 52:4a:93:75:dd:91:0a:01 fcid 0xc40043 dynamic
! [X70-CT0-FC1]
vsan 101 wwn 52:4a:93:71:56:84:09:02 fcid 0xc40044 dynamic
! [X70R3-CT0-FC2]
vsan 101 wwn 52:4a:93:71:56:84:09:00 fcid 0xc40024 dynamic
vsan 101 wwn 52:4a:93:71:56:84:09:12 fcid 0xc40045 dynamic
! [X70R3-CT1-FC2]
vsan 101 wwn 20:00:00:25:b5:bb:17:40 fcid 0xc400c2 dynamic
! [AMD-VMHost70-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:41 fcid 0xc400a2 dynamic
! [AMD-VMHost70-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:44 fcid 0xc400c3 dynamic
! [AMD-VMHost72-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:45 fcid 0xc400a3 dynamic

! [AMD-VMHost72-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:4e fcid 0xc400c4 dynamic

! [AMD-VMHost73-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:4f fcid 0xc400a4 dynamic

! [AMD-VMHost73-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:42 fcid 0xc400c5 dynamic

! [AMD-VMHost71-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:43 fcid 0xc400a5 dynamic

! [AMD-VMHost71-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:46 fcid 0xc400c6 dynamic

! [AMD-VMHost74-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:47 fcid 0xc400a6 dynamic

! [AMD-VMHost74-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:48 fcid 0xc400c7 dynamic

! [AMD-VMHost75-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:49 fcid 0xc400a7 dynamic

! [AMD-VMHost75-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:4a fcid 0xc400c8 dynamic

! [AMD-VMHost76-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:4b fcid 0xc400a8 dynamic

! [AMD-VMHost76-HBA3]
vsan 101 wwn 20:00:00:25:b5:bb:17:4c fcid 0xc400c9 dynamic

! [AMD-VMHost77-HBA1]
vsan 101 wwn 20:00:00:25:b5:bb:17:4d fcid 0xc400a9 dynamic

! [AMD-VMHost77-HBA3]

!Active Zone Database Section for vsan 101

zone name FlaskStack-VCC-CVD-WLHost01 vsan 101

member pwwn 20:00:00:25:b5:bb:17:00

! [VCC-WLHost01-HBA1]

member pwwn 20:00:00:25:b5:bb:17:01

! [VCC-WLHost01-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 101

member pwwn 20:00:00:25:b5:bb:17:02

! [VCC-WLHost02-HBA1]

member pwwn 20:00:00:25:b5:bb:17:03

! [VCC-WLHost02-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 101

member pwwn 20:00:00:25:b5:bb:17:04

! [VCC-WLHost03-HBA1]

member pwwn 20:00:00:25:b5:bb:17:05

! [VCC-WLHost03-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 101

member pwwn 20:00:00:25:b5:bb:17:06

! [VCC-WLHost04-HBA1]

member pwwn 20:00:00:25:b5:bb:17:07

! [VCC-WLHost04-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 101

member pwwn 20:00:00:25:b5:bb:17:08

! [VCC-WLHost05-HBA1]

member pwwn 20:00:00:25:b5:bb:17:09

! [VCC-WLHost05-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 101

member pwwn 20:00:00:25:b5:bb:17:0a

! [VCC-WLHost06-HBA1]

member pwwn 20:00:00:25:b5:bb:17:0b

! [VCC-WLHost06-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 101

member pwwn 20:00:00:25:b5:bb:17:0c

! [VCC-WLHost07-HBA1]

member pwwn 20:00:00:25:b5:bb:17:0d

! [VCC-WLHost07-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 101

member pwwn 20:00:00:25:b5:bb:17:0e

! [VCC-WLHost08-HBA1]

member pwwn 20:00:00:25:b5:bb:17:0f

! [VCC-WLHost08-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 101

member pwwn 20:00:00:25:b5:bb:17:10

! [VCC-WLHost09-HBA1]

member pwwn 20:00:00:25:b5:bb:17:11

! [VCC-WLHost09-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 101

member pwwn 20:00:00:25:b5:bb:17:12

! [VCC-WLHost10-HBA1]

member pwwn 20:00:00:25:b5:bb:17:13

! [VCC-WLHost10-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 101

member pwwn 20:00:00:25:b5:bb:17:14

! [VCC-WLHost11-HBA1]

member pwwn 20:00:00:25:b5:bb:17:15

! [VCC-WLHost11-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 101

member pwwn 20:00:00:25:b5:bb:17:16

! [VCC-WLHost12-HBA1]

member pwwn 20:00:00:25:b5:bb:17:17

! [VCC-WLHost12-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 101

member pwwn 20:00:00:25:b5:bb:17:18

! [VCC-WLHost13-HBA1]

member pwwn 20:00:00:25:b5:bb:17:19

! [VCC-WLHost13-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 101

member pwwn 20:00:00:25:b5:bb:17:1a

! [VCC-WLHost14-HBA1]

member pwwn 20:00:00:25:b5:bb:17:1b

! [VCC-WLHost14-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 101

member pwwn 20:00:00:25:b5:bb:17:1c

! [VCC-WLHost15-HBA1]

member pwwn 20:00:00:25:b5:bb:17:1d

! [VCC-WLHost15-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra01 vsan 101

member pwwn 20:00:00:25:b5:bb:17:1e

! [VCC-Infra01-HBA1]

member pwwn 20:00:00:25:b5:bb:17:1f

! [VCC-Infra01-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 101

member pwwn 20:00:00:25:b5:bb:17:20

! [VCC-WLHost16-HBA1]

member pwwn 20:00:00:25:b5:bb:17:21

! [VCC-WLHost16-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 101

member pwwn 20:00:00:25:b5:bb:17:22

! [VCC-WLHost17-HBA1]

member pwwn 20:00:00:25:b5:bb:17:23

! [VCC-WLHost17-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 101

member pwwn 20:00:00:25:b5:bb:17:24

! [VCC-WLHost18-HBA1]

member pwwn 20:00:00:25:b5:bb:17:25

! [VCC-WLHost18-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 101

member pwwn 20:00:00:25:b5:bb:17:26

! [VCC-WLHost19-HBA1]

member pwwn 20:00:00:25:b5:bb:17:27

! [VCC-WLHost19-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 101

member pwwn 20:00:00:25:b5:bb:17:28

! [VCC-WLHost20-HBA1]

member pwwn 20:00:00:25:b5:bb:17:29

! [VCC-WLHost20-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 101

member pwwn 20:00:00:25:b5:bb:17:2a

! [VCC-WLHost21-HBA1]

member pwwn 20:00:00:25:b5:bb:17:2b

! [VCC-WLHost21-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 101

member pwwn 20:00:00:25:b5:bb:17:2c

! [VCC-WLHost22-HBA1]

member pwwn 20:00:00:25:b5:bb:17:2d

! [VCC-WLHost22-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 101

member pwwn 20:00:00:25:b5:bb:17:2e

! [VCC-WLHost23-HBA1]

member pwwn 20:00:00:25:b5:bb:17:2f

! [VCC-WLHost23-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 101

member pwwn 20:00:00:25:b5:bb:17:30

! [VCC-WLHost24-HBA1]

member pwwn 20:00:00:25:b5:bb:17:31

! [VCC-WLHost24-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 101

member pwwn 20:00:00:25:b5:bb:17:32

! [VCC-WLHost25-HBA1]

member pwwn 20:00:00:25:b5:bb:17:33

! [VCC-WLHost25-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 101

member pwwn 20:00:00:25:b5:bb:17:34

! [VCC-WLHost26-HBA1]

member pwwn 20:00:00:25:b5:bb:17:35

! [VCC-WLHost26-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 101

member pwwn 20:00:00:25:b5:bb:17:36

! [VCC-WLHost27-HBA1]

member pwwn 20:00:00:25:b5:bb:17:37

! [VCC-WLHost27-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 101

member pwwn 20:00:00:25:b5:bb:17:38

! [VCC-WLHost28-HBA1]

member pwwn 20:00:00:25:b5:bb:17:39

! [VCC-WLHost28-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 101

member pwwn 20:00:00:25:b5:bb:17:3a

! [VCC-WLHost29-HBA1]

member pwwn 20:00:00:25:b5:bb:17:3b

! [VCC-WLHost29-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 101

member pwwn 20:00:00:25:b5:bb:17:3c

! [VCC-WLHost30-HBA1]

member pwwn 20:00:00:25:b5:bb:17:3d

! [VCC-WLHost30-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra02 vsan 101

member pwwn 20:00:00:25:b5:bb:17:3e

! [VCC-Infra02-HBA1]

member pwwn 20:00:00:25:b5:bb:17:3f

! [VCC-Infra02-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost70 vsan 101

member pwwn 20:00:00:25:b5:bb:17:40

! [AMD-VMHost70-HBA1]

member pwwn 20:00:00:25:b5:bb:17:41

! [AMD-VMHost70-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost71 vsan 101

member pwwn 20:00:00:25:b5:bb:17:42

! [AMD-VMHost71-HBA1]

member pwwn 20:00:00:25:b5:bb:17:43

! [AMD-VMHost71-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost72 vsan 101

member pwwn 20:00:00:25:b5:bb:17:44

! [AMD-VMHost72-HBA1]

member pwwn 20:00:00:25:b5:bb:17:45

! [AMD-VMHost72-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost73 vsan 101

member pwwn 20:00:00:25:b5:bb:17:4e

! [AMD-VMHost73-HBA1]

member pwwn 20:00:00:25:b5:bb:17:4f

! [AMD-VMHost73-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost74 vsan 101

member pwwn 20:00:00:25:b5:bb:17:46

! [AMD-VMHost74-HBA1]

member pwwn 20:00:00:25:b5:bb:17:47

! [AMD-VMHost74-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost75 vsan 101

member pwwn 20:00:00:25:b5:bb:17:48

! [AMD-VMHost75-HBA1]

member pwwn 20:00:00:25:b5:bb:17:49

! [AMD-VMHost75-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost76 vsan 101

member pwwn 20:00:00:25:b5:bb:17:4a

! [AMD-VMHost76-HBA1]

member pwwn 20:00:00:25:b5:bb:17:4b

! [AMD-VMHost76-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost77 vsan 101

member pwwn 20:00:00:25:b5:bb:17:4c

! [AMD-VMHost77-HBA1]

member pwwn 20:00:00:25:b5:bb:17:4d

! [AMD-VMHost77-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zoneset name FlashStack-VCC-CVD vsan 101

member FlaskStack-VCC-CVD-WLHost01

member FlaskStack-VCC-CVD-WLHost02

member FlaskStack-VCC-CVD-WLHost03

member FlaskStack-VCC-CVD-WLHost04

member FlaskStack-VCC-CVD-WLHost05
member FlaskStack-VCC-CVD-WLHost06
member FlaskStack-VCC-CVD-WLHost07
member FlaskStack-VCC-CVD-WLHost08
member FlaskStack-VCC-CVD-WLHost09
member FlaskStack-VCC-CVD-WLHost10
member FlaskStack-VCC-CVD-WLHost11
member FlaskStack-VCC-CVD-WLHost12
member FlaskStack-VCC-CVD-WLHost13
member FlaskStack-VCC-CVD-WLHost14
member FlaskStack-VCC-CVD-WLHost15
member FlaskStack-VCC-CVD-Infra01
member FlaskStack-VCC-CVD-WLHost16
member FlaskStack-VCC-CVD-WLHost17
member FlaskStack-VCC-CVD-WLHost18
member FlaskStack-VCC-CVD-WLHost19
member FlaskStack-VCC-CVD-WLHost20
member FlaskStack-VCC-CVD-WLHost21
member FlaskStack-VCC-CVD-WLHost22
member FlaskStack-VCC-CVD-WLHost23
member FlaskStack-VCC-CVD-WLHost24
member FlaskStack-VCC-CVD-WLHost25
member FlaskStack-VCC-CVD-WLHost26
member FlaskStack-VCC-CVD-WLHost27

member FlaskStack-VCC-CVD-WLHost28
member FlaskStack-VCC-CVD-WLHost29
member FlaskStack-VCC-CVD-WLHost30
member FlaskStack-VCC-CVD-Infra02
member FlaskStack-AMD-VMHost70
member FlaskStack-AMD-VMHost71
member FlaskStack-AMD-VMHost72
member FlaskStack-AMD-VMHost73
member FlaskStack-AMD-VMHost74
member FlaskStack-AMD-VMHost75
member FlaskStack-AMD-VMHost76
member FlaskStack-AMD-VMHost77

zoneset activate name FlashStack-VCC-CVD vsan 101
do clear zone database vsan 101
!Full Zone Database Section for vsan 101
zone name FlaskStack-VCC-CVD-WLHost01 vsan 101
member pwwn 20:00:00:25:b5:bb:17:00
! [VCC-WLHost01-HBA1]
member pwwn 20:00:00:25:b5:bb:17:01
! [VCC-WLHost01-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost02 vsan 101

member pwwn 20:00:00:25:b5:bb:17:02

! [VCC-WLHost02-HBA1]

member pwwn 20:00:00:25:b5:bb:17:03

! [VCC-WLHost02-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost03 vsan 101

member pwwn 20:00:00:25:b5:bb:17:04

! [VCC-WLHost03-HBA1]

member pwwn 20:00:00:25:b5:bb:17:05

! [VCC-WLHost03-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost04 vsan 101

member pwwn 20:00:00:25:b5:bb:17:06

! [VCC-WLHost04-HBA1]
member pwwn 20:00:00:25:b5:bb:17:07
!
! [VCC-WLHost04-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost05 vsan 101

member pwwn 20:00:00:25:b5:bb:17:08
!
! [VCC-WLHost05-HBA1]
member pwwn 20:00:00:25:b5:bb:17:09
!
! [VCC-WLHost05-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost06 vsan 101

member pwwn 20:00:00:25:b5:bb:17:0a
!
! [VCC-WLHost06-HBA1]
member pwwn 20:00:00:25:b5:bb:17:0b
!
! [VCC-WLHost06-HBA3]
member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost07 vsan 101

member pwwn 20:00:00:25:b5:bb:17:0c

! [VCC-WLHost07-HBA1]

member pwwn 20:00:00:25:b5:bb:17:0d

! [VCC-WLHost07-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost08 vsan 101

member pwwn 20:00:00:25:b5:bb:17:0e

! [VCC-WLHost08-HBA1]

member pwwn 20:00:00:25:b5:bb:17:0f

! [VCC-WLHost08-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost09 vsan 101

member pwwn 20:00:00:25:b5:bb:17:10

! [VCC-WLHost09-HBA1]

member pwwn 20:00:00:25:b5:bb:17:11

! [VCC-WLHost09-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost10 vsan 101

member pwwn 20:00:00:25:b5:bb:17:12

! [VCC-WLHost10-HBA1]

member pwwn 20:00:00:25:b5:bb:17:13

! [VCC-WLHost10-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost11 vsan 101

member pwwn 20:00:00:25:b5:bb:17:14

! [VCC-WLHost11-HBA1]

member pwwn 20:00:00:25:b5:bb:17:15

! [VCC-WLHost11-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost12 vsan 101

member pwwn 20:00:00:25:b5:bb:17:16
!
! [VCC-WLHost12-HBA1]
member pwwn 20:00:00:25:b5:bb:17:17
!
! [VCC-WLHost12-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost13 vsan 101

member pwwn 20:00:00:25:b5:bb:17:18
!
! [VCC-WLHost13-HBA1]
member pwwn 20:00:00:25:b5:bb:17:19
!
! [VCC-WLHost13-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost14 vsan 101

member pwwn 20:00:00:25:b5:bb:17:1a

! [VCC-WLHost14-HBA1]

member pwwn 20:00:00:25:b5:bb:17:1b

! [VCC-WLHost14-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost15 vsan 101

member pwwn 20:00:00:25:b5:bb:17:1c

! [VCC-WLHost15-HBA1]

member pwwn 20:00:00:25:b5:bb:17:1d

! [VCC-WLHost15-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra01 vsan 101

member pwwn 20:00:00:25:b5:bb:17:1e

! [VCC-Infra01-HBA1]
member pwwn 20:00:00:25:b5:bb:17:1f
!
! [VCC-Infra01-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost16 vsan 101

member pwwn 20:00:00:25:b5:bb:17:20
!
! [VCC-WLHost16-HBA1]
member pwwn 20:00:00:25:b5:bb:17:21
!
! [VCC-WLHost16-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost17 vsan 101

member pwwn 20:00:00:25:b5:bb:17:22
!
! [VCC-WLHost17-HBA1]
member pwwn 20:00:00:25:b5:bb:17:23
!
! [VCC-WLHost17-HBA3]
member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost18 vsan 101

member pwwn 20:00:00:25:b5:bb:17:24

! [VCC-WLHost18-HBA1]

member pwwn 20:00:00:25:b5:bb:17:25

! [VCC-WLHost18-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost19 vsan 101

member pwwn 20:00:00:25:b5:bb:17:26

! [VCC-WLHost19-HBA1]

member pwwn 20:00:00:25:b5:bb:17:27

! [VCC-WLHost19-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost20 vsan 101

member pwwn 20:00:00:25:b5:bb:17:28

! [VCC-WLHost20-HBA1]

member pwwn 20:00:00:25:b5:bb:17:29

! [VCC-WLHost20-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost21 vsan 101

member pwwn 20:00:00:25:b5:bb:17:2a

! [VCC-WLHost21-HBA1]

member pwwn 20:00:00:25:b5:bb:17:2b

! [VCC-WLHost21-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost22 vsan 101

member pwwn 20:00:00:25:b5:bb:17:2c

! [VCC-WLHost22-HBA1]

member pwwn 20:00:00:25:b5:bb:17:2d

! [VCC-WLHost22-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost23 vsan 101

member pwwn 20:00:00:25:b5:bb:17:2e
!
! [VCC-WLHost23-HBA1]
member pwwn 20:00:00:25:b5:bb:17:2f
!
! [VCC-WLHost23-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost24 vsan 101

member pwwn 20:00:00:25:b5:bb:17:30
!
! [VCC-WLHost24-HBA1]
member pwwn 20:00:00:25:b5:bb:17:31
!
! [VCC-WLHost24-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost25 vsan 101

member pwwn 20:00:00:25:b5:bb:17:32

! [VCC-WLHost25-HBA1]

member pwwn 20:00:00:25:b5:bb:17:33

! [VCC-WLHost25-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost26 vsan 101

member pwwn 20:00:00:25:b5:bb:17:34

! [VCC-WLHost26-HBA1]

member pwwn 20:00:00:25:b5:bb:17:35

! [VCC-WLHost26-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost27 vsan 101

member pwwn 20:00:00:25:b5:bb:17:36

! [VCC-WLHost27-HBA1]
member pwwn 20:00:00:25:b5:bb:17:37
!
! [VCC-WLHost27-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost28 vsan 101

member pwwn 20:00:00:25:b5:bb:17:38
!
! [VCC-WLHost28-HBA1]
member pwwn 20:00:00:25:b5:bb:17:39
!
! [VCC-WLHost28-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost29 vsan 101

member pwwn 20:00:00:25:b5:bb:17:3a
!
! [VCC-WLHost29-HBA1]
member pwwn 20:00:00:25:b5:bb:17:3b
!
! [VCC-WLHost29-HBA3]
member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-WLHost30 vsan 101

member pwwn 20:00:00:25:b5:bb:17:3c

! [VCC-WLHost30-HBA1]

member pwwn 20:00:00:25:b5:bb:17:3d

! [VCC-WLHost30-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-VCC-CVD-Infra02 vsan 101

member pwwn 20:00:00:25:b5:bb:17:3e

! [VCC-Infra02-HBA1]

member pwwn 20:00:00:25:b5:bb:17:3f

! [VCC-Infra02-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost70 vsan 101

member pwwn 20:00:00:25:b5:bb:17:40

! [AMD-VMHost70-HBA1]

member pwwn 20:00:00:25:b5:bb:17:41

! [AMD-VMHost70-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost71 vsan 101

member pwwn 20:00:00:25:b5:bb:17:42

! [AMD-VMHost71-HBA1]

member pwwn 20:00:00:25:b5:bb:17:43

! [AMD-VMHost71-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost72 vsan 101

member pwwn 20:00:00:25:b5:bb:17:44

! [AMD-VMHost72-HBA1]

member pwwn 20:00:00:25:b5:bb:17:45

! [AMD-VMHost72-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost73 vsan 101

member pwwn 20:00:00:25:b5:bb:17:4e
!
! [AMD-VMHost73-HBA1]
member pwwn 20:00:00:25:b5:bb:17:4f
!
! [AMD-VMHost73-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12
!
! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost74 vsan 101

member pwwn 20:00:00:25:b5:bb:17:46
!
! [AMD-VMHost74-HBA1]
member pwwn 20:00:00:25:b5:bb:17:47
!
! [AMD-VMHost74-HBA3]
member pwwn 52:4a:93:71:56:84:09:02
!
! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost75 vsan 101

member pwwn 20:00:00:25:b5:bb:17:48

! [AMD-VMHost75-HBA1]

member pwwn 20:00:00:25:b5:bb:17:49

! [AMD-VMHost75-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost76 vsan 101

member pwwn 20:00:00:25:b5:bb:17:4a

! [AMD-VMHost76-HBA1]

member pwwn 20:00:00:25:b5:bb:17:4b

! [AMD-VMHost76-HBA3]

member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]

member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zone name FlaskStack-AMD-VMHost77 vsan 101

member pwwn 20:00:00:25:b5:bb:17:4c

! [AMD-VMHost77-HBA1]
member pwwn 20:00:00:25:b5:bb:17:4d

! [AMD-VMHost77-HBA3]
member pwwn 52:4a:93:71:56:84:09:02

! [X70R3-CT0-FC2]
member pwwn 52:4a:93:71:56:84:09:12

! [X70R3-CT1-FC2]

zoneset name FlashStack-VCC-CVD vsan 101

member FlaskStack-VCC-CVD-WLHost01

member FlaskStack-VCC-CVD-WLHost02

member FlaskStack-VCC-CVD-WLHost03

member FlaskStack-VCC-CVD-WLHost04

member FlaskStack-VCC-CVD-WLHost05

member FlaskStack-VCC-CVD-WLHost06

member FlaskStack-VCC-CVD-WLHost07

member FlaskStack-VCC-CVD-WLHost08

member FlaskStack-VCC-CVD-WLHost09

member FlaskStack-VCC-CVD-WLHost10

member FlaskStack-VCC-CVD-WLHost11

member FlaskStack-VCC-CVD-WLHost12

member FlaskStack-VCC-CVD-WLHost13

member FlaskStack-VCC-CVD-WLHost14

member FlaskStack-VCC-CVD-WLHost15

member FlaskStack-VCC-CVD-Infra01
member FlaskStack-VCC-CVD-WLHost16
member FlaskStack-VCC-CVD-WLHost17
member FlaskStack-VCC-CVD-WLHost18
member FlaskStack-VCC-CVD-WLHost19
member FlaskStack-VCC-CVD-WLHost20
member FlaskStack-VCC-CVD-WLHost21
member FlaskStack-VCC-CVD-WLHost22
member FlaskStack-VCC-CVD-WLHost23
member FlaskStack-VCC-CVD-WLHost24
member FlaskStack-VCC-CVD-WLHost25
member FlaskStack-VCC-CVD-WLHost26
member FlaskStack-VCC-CVD-WLHost27
member FlaskStack-VCC-CVD-WLHost28
member FlaskStack-VCC-CVD-WLHost29
member FlaskStack-VCC-CVD-WLHost30
member FlaskStack-VCC-CVD-Infra02
member FlaskStack-AMD-VMHost70
member FlaskStack-AMD-VMHost71
member FlaskStack-AMD-VMHost72
member FlaskStack-AMD-VMHost73
member FlaskStack-AMD-VMHost74
member FlaskStack-AMD-VMHost75
member FlaskStack-AMD-VMHost76

member FlaskStack-AMD-VMHost77

interface mgmt0

ip address 10.29.164.239 255.255.255.0

vsan database

vsan 101 interface fc1/9

vsan 101 interface fc1/10

vsan 101 interface fc1/11

vsan 101 interface fc1/12

vsan 101 interface fc1/13

vsan 101 interface fc1/14

vsan 101 interface fc1/15

vsan 101 interface fc1/16

clock timezone PST 0 0

clock summer-time PDT 2 Sun Mar 02:00 1 Sun Nov 02:00 60

switchname ADD16-MDS-B

cli alias name autozone source sys/autozone.py

line console

line vty

boot kickstart bootflash:/m9100-s6ek9-kickstart-mz.8.3.1.bin

boot system bootflash:/m9100-s6ek9-mz.8.3.1.bin

interface fc1/1

interface fc1/2

interface fc1/3

interface fc1/4

interface fc1/5

interface fc1/6

interface fc1/7

interface fc1/8

interface fc1/9

interface fc1/10

interface fc1/11

interface fc1/12

interface fc1/13

interface fc1/14

interface fc1/15

interface fc1/16

interface fc1/1

no port-license

interface fc1/2

no port-license

interface fc1/3

no port-license

interface fc1/4

no port-license

interface fc1/5

no port-license

interface fc1/6

no port-license

interface fc1/7

no port-license

interface fc1/8

no port-license

interface fc1/9

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/10

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/11

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/12

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/13

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/14

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/15

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

interface fc1/16

switchport trunk allowed vsan 101

switchport trunk mode off

port-license acquire

no shutdown

ip default-gateway 10.29.164.1

Full Scale Server Performance Chart with Boot and LoginVSI Knowledge Worker Workload Test

This section provides a detailed performance chart for ESXi 7.0 Update 2 installed on Cisco UCS B200 M6 Blade Server as part of the workload test with Citrix Virtual Apps and Desktops 7 2109 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 92. Full Scale | 1960 Users| MCS Single-session OS machine VDAs | Host Memory Utilization





Figure 94. Full Scale | 1960 Users| MCS Single-session OS machine VDAs | Host Network Utilization | Transmitted



VDI Server Performance Monitor Data for One Sample Test: 1960 Users PVS Single-session OS machine VDAs Scale Testing

Figure 95. Full Scale | 1960 Users | PVS Single-session OS machine VDAs | Host CPU Utilization





Figure 98. Full Scale | 1960 Users| PVS Single-session OS machine VDAs | Host Network Utilization | Received



VDI Server Performance Monitor Data for One Sample Test: 2688 Users MCS Multi-session OS machine VDAs Scale Testing

Figure 99. Full Scale | 2688 Users| MCS Multi-session OS machine VDAs | Host CPU Utilization





Figure 102. Full Scale | 2688 Users| MCS Multi-session OS machine VDAs | Host Network Utilization | Received



Feedback

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

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

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