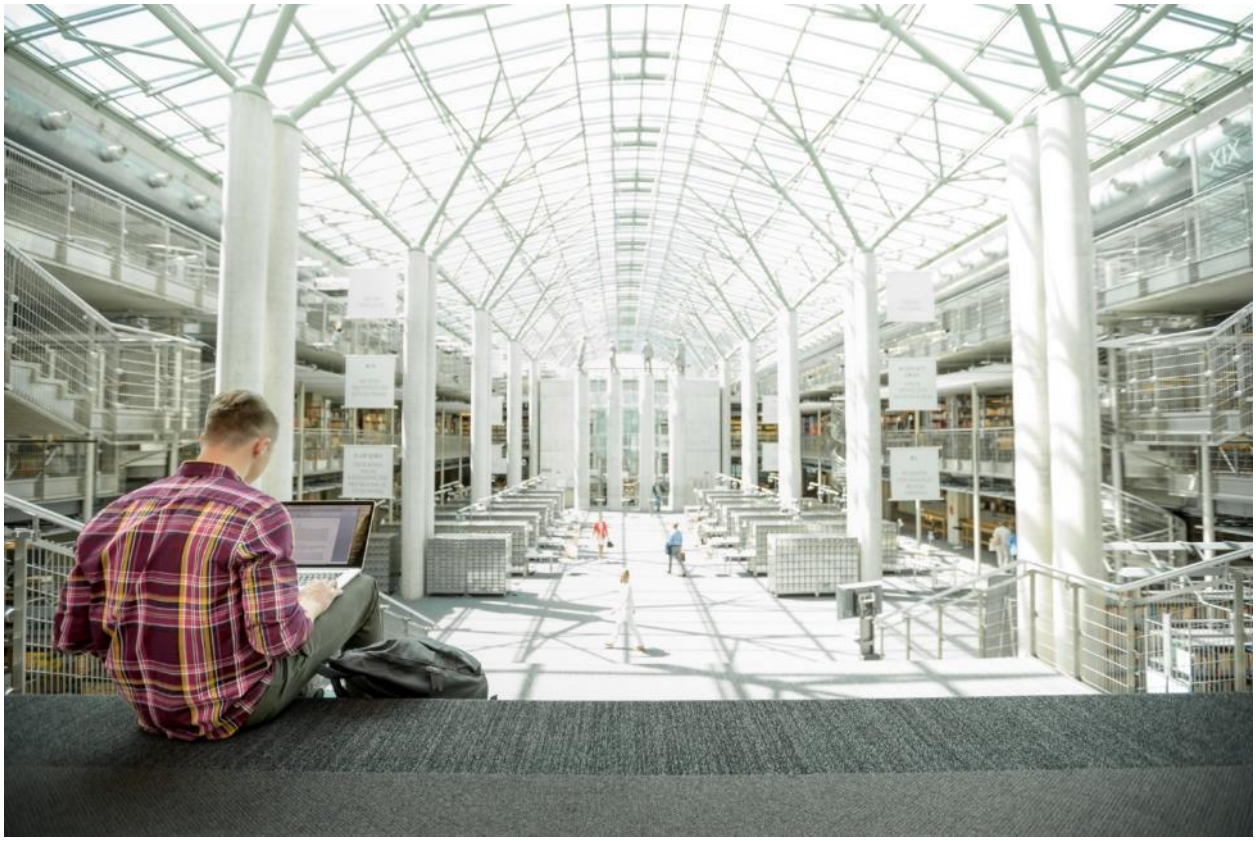

Cisco All-Flash HyperFlex 3.5 Hyperconverged System with up to 2000 Citrix Virtual Apps and Desktops Users

Design and Deployment of Cisco HyperFlex for Virtual Desktop Infrastructure with Citrix Virtual Apps and Desktops 1808

Last Updated: February 11, 2019

CISCO VALIDATED DESIGN

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.

Table of Contents

Executive Summary

To keep pace with the market, you need systems that support rapid, agile development processes. Cisco HyperFlex Systems let you unlock the full potential of hyper-convergence and adapt IT to the needs of your workloads. The systems use an end-to-end software-defined infrastructure approach, combining software-defined computing in the form of Cisco HyperFlex HX-Series Nodes, software-defined storage with the powerful Cisco HyperFlex HX Data Platform, and software-defined networking with the Cisco UCS fabric that integrates smoothly with Cisco Application Centric Infrastructure (Cisco ACI).

Together with a single point of connectivity and management, these technologies deliver a pre-integrated and adaptable cluster with a unified pool of resources that you can quickly deploy, adapt, scale, and manage to power efficiently your applications and your business.

This document provides an architectural reference and design guide for up to a 2000 user mixed workload on a 16-node (8 Cisco HyperFlex HXAF220C-M5SX servers plus 8 B200 M5 blade servers) Cisco HyperFlex system. We provide deployment guidance and performance data for Citrix Virtual Apps and Desktops 1808 virtual desktop sessions running Windows Server 2016 HSD server-based Remote Desktop Server sessions and Microsoft Windows 10 with Office 2016, highlighting provisioning through Citrix PVS and MCS Persistent virtual desktops on vSphere 6.5.

This solution is a pre-integrated, best-practice data center architecture built on Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches and Cisco HyperFlex Data Platform software version 3.5(1a.)

The solution payload is 100 percent virtualized on Cisco HyperFlex HXAF220c-M5SX hyperconverged nodes and Cisco UCS B200 M5 blade servers booting through on-board Flex-Flash controller SD cards running VMware vSphere 6.5 U2 hypervisor. The virtual desktops are configured with Citrix Virtual Apps and Desktops 1808 which incorporates both traditional persistent and non-persistent virtual Windows 10 desktops, hosted applications and remote desktop service (RDS) server 2008 R2, server 2012 R2 or server 2016 based desktops. The solution provides unparalleled scale and management simplicity. Citrix HSD server based desktop sessions (800,) PVS non-persistent Windows 10 desktops (800), Citrix Machine Creation Services Persistent Windows 10 desktops (400) are provisioned on a sixteen node Cisco HyperFlex cluster. Where applicable, this document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

The solution boots 2000 HSD virtual server and virtual desktops machines in 15 minutes or less, insuring that users will not experience delays in accessing their virtual workspace on HyperFlex.

The solution is fully capable of supporting hardware accelerated graphic workloads. Each Cisco HyperFlex HXAF240c M5 node and each Cisco UCS C240 M5 compute only server can support up to two NVIDIA M10 or P40 cards and up to six NVIDIA P4 cards. The Cisco UCS B200 M5 server supports up to two NVIDIA P6 cards for high density, high performance graphics workload support. See our Cisco Graphics White Paper for our fifth generation servers with NVIDIA GPUs and software for details on how to integrate this capability with VMware and Citrix.

This solution provides outstanding virtual desktop end user experience as measured by the Login VSI 4.1x Knowledge Worker workload running in benchmark mode. Average end-user response times for all tested delivery methods is under one (1) second, representing the best performance in the industry.

Solution Overview

Introduction

A current industry trend in data center design is towards small, granularly expandable hyperconverged infrastructures. By using virtualization along with pre-validated IT platforms, customers of all sizes have embarked on the journey to "just-in-time capacity" using this new technology. The Cisco Hyper Converged Solution can be quickly deployed, thereby increasing agility and reducing costs. Cisco HyperFlex 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 and scaled out.

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers deploying the Cisco HyperFlex System. External references are provided wherever applicable, but readers are expected to be familiar with Citrix and VMware specific technologies, infrastructure concepts, networking connectivity, and security policies of the customer installation.

Purpose of this Document

This document provides a systematic design, configuration, and implementation guide for the Cisco Validated Design for a Cisco HyperFlex All-Flash system running three different Citrix Virtual Apps and Desktops workloads with Cisco UCS 6300 series Fabric Interconnects and Cisco Nexus 93000 series networking switches.

What's New in this Release?

This is the first Cisco Validated Design with Cisco HyperFlex All-Flash system running Virtual Desktop Infrastructure on Intel Xeon Scalable Family processor-based, fifth generation Cisco UCS HyperFlex system with these features:

- Validation of Cisco Nexus 93000 YC series with Cisco HyperFlex Support
- Support for the Cisco UCS 4.0(1b) release and Cisco HyperFlex v 3.5(1a)
- VMware vSphere 6.5 U2 Hypervisor
- Citrix Virtual Apps and Desktops 1808 Remote Desktop Sever Hosted Sessions with PVS
- Citrix Virtual Apps and Desktops 1808 PVS non persistent, and MCS Persistent Desktops

Enhancements for Version 3.5

Some of the new features in Version 3.5 are as follows:

- Native Disaster Recovery Enhancements– A simple to use Planned Migration workflow for Disaster Recovery, virtual machine migration, and resuming replication. In addition, native support for Replication, Planned Migration and Disaster Recovery in Stretched Cluster deployments. For more information, see Cisco HyperFlex Systems Administration Guide, Release 3.5.

- HX Data Platform Installer Enhancements—New extended capabilities in hardening and reliability in the HX Data Platform Installer.
- Cluster Expansion for Hyper-V converged nodes.
- Cluster expansion for Stretched Cluster compute-only and converged nodes.
- Integrated Hyper-V and Windows Server OS bare metal installation included as part of cluster creation workflow.
- Networking Enhancements—Support for multi-VIC network designs and third-party NIC for HX converged and compute-only nodes. For more information, see Cisco HyperFlex Systems – Networking Topologies
- Upgrade Enhancements—Support for orchestrated ESXi hypervisor upgrades. Combined with existing support for HXDP and server firmware upgrades, this release provides the ability to perform seamless full stack upgrades all orchestrated through HX Connect.
- Starting with release 3.5(1a) and later, all future upgrades can be completed in HX Connect UI. For all future upgrades, this functionality will take affect for all clusters on release 3.5. To upgrade to the 3.5 release from older versions, continue to run the bootstrap script as outlined in the documentation. This new end-to-end UI-based upgrade capability will be utilized on all subsequent upgrades. For more information, see Cisco HyperFlex Systems Upgrade Guide, Release 3.5
- ESXi Lockdown Mode—Support for VMware ESXi lockdown mode to increase security of an ESXi host by limiting access allowed for the host. When enabled, the ESXi host can only be accessed through vCenter Server or Direct Console User Interface (DCUI). For more information, Cisco HyperFlex Systems Installation Guide, Release 3.5
- HX Edge 10GbE Edge Network option-—New 10GbE Edge support provides an additional fully redundant, high speed connectivity option for HyperFlex Edge clusters. For more information, see Cisco HyperFlex Systems Edge Deployment Guide, Release 3.5
- Cisco Container Platform (CCP) and Open Shift Platform integration (OpenShift)—Storage integration with Kubernetes that enables dynamic (on-demand) persistent volumes from HyperFlex. This feature is supported with OpenShift (version 3.10) and Cisco Container Platform (CCP). For more information, see Cisco HyperFlex Systems Kubernetes Administration Guide, Release 3.5
- Artificial intelligence and machine learning (AI/ML) Workloads on HyperFlex with NVIDIA V100 GPUs—Ability to create applications for AI/ML with NVIDIA Tesla V100 GPUs integration within HyperFlex nodes. For more information, see Cisco HyperFlex HX-Series Spec Sheets
- Permanent License Reservation (PLR) – This feature is designed for highly secure intelligence, air-gapped and military environments where external communication may be limited. For more information, see Cisco HyperFlex Systems Ordering and Licensing Guide
- DISA STIG Automation—Enhance the security posture of HyperFlex converged and compute-only nodes by automating the implementation of the Defense Information Systems Agency's(DISA) recommended Security Technical Implementation Guides (STIGs), pertaining to VMware vSphere
- Tech Support Mode –Enhance the security posture of HyperFlex converged nodes by disabling Tech Support Mode, which disables remote access to Controller virtual machines over SSH.
- Multi-hypervisor support allows HyperFlex to be installed with either the VMware ESXi hypervisor, or Microsoft Hyper-V. This document focuses on installation and support of HyperFlex with the VMware ESXi hypervisor.

- Installation of a Cisco HyperFlex cluster can span two physical locations, creating a stretched cluster. A third location is required for running a witness virtual machine to prevent a “split brain” situation.
- HyperFlex clusters can be configured with logical availability zones, which subdivide the nodes into groups and evenly distribute the data across all zones, in order to better tolerate node failures.
- Support for 64 node clusters; up to 32 converged nodes and 32 compute-only nodes can be used per cluster.
- Support for using Intel Optane based NVMe based SSDs as the caching disk in the Cisco HyperFlex all-flash nodes.
- Support for large form factor hard drives in the HyperFlex HX240-M5SL model server for higher storage capacity.
- Enhancements and customizations available for the HyperFlex Connect native HTML5 management GUI.
- Kubernetes support with automated storage and networking deployment through a new FlexVolume driver, creating a fully integrated container platform.

Documentation Roadmap

For the comprehensive documentation suite, refer to the Cisco HyperFlex Systems Documentation Roadmap.

A login is required for the Documentation Roadmap.

Cisco HX Data Platform requires specific software and hardware versions, and networking settings for successful installation. See the Cisco HyperFlex Systems Getting Started Guide for a complete list of requirements.

For a complete list of hardware and software inter-dependencies, refer to the Cisco UCS Manager release version of Hardware and Software Interoperability for Cisco HyperFlex HX-Series.

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

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

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

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both Microsoft Windows 10 virtual desktops and HSD server desktop sessions based on Microsoft Server 2016. The mixed workload solution includes Cisco HyperFlex hardware and Data Platform software, Cisco

Nexus switches, Cisco Unified Computing System, Citrix, and VMware vSphere software in a single package. The design is efficient such that the networking, computing, and storage components occupy 18-rack units footprint in an industry standard 42U rack. Port density on the Cisco Nexus switches and Cisco UCS Fabric Interconnects enables the networking components to accommodate multiple HyperFlex clusters in a single Cisco UCS domain.

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

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

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

- More power, same size. Cisco HX-series nodes, dual 18-core 2.3 GHz Intel Xeon (Gold 6140) Scalable Family processors with 768GB of 2666Mhz memory with Citrix support more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel Xeon Gold 6140 18-core scalable family processors used in this study provided a balance between increased per-server capacity and cost
- Fault-tolerance with high availability built into the design. The various designs are based on multiple Cisco HX-Series nodes, Cisco UCS rack servers and Cisco UCS blade servers for virtual desktop and infrastructure workloads. The design provides N+1 server fault tolerance for every payload type tested
- Stress-tested to the limits during aggressive boot scenario. The 2000 user mixed hosted virtual desktop and hosted shared desktop environment booted and registered with the Citrix Studio in under 15 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.
- Stress-tested to the limits during simulated login storms. All 2000 users logged in and started running workloads up to steady state in 48-minutes without overwhelming the processors, exhausting memory or exhausting the storage subsystems, providing customers with a desktop virtualization system that can easily handle the most demanding login and startup storms.
- Ultra-condensed computing for the datacenter. The rack space required to support the initial 2000 user system is 20 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects. Incremental seat Cisco HyperFlex clusters can be added one at a time to a total of 64 nodes.
- 100 percent virtualized. This CVD presents a validated design that is 100 percent virtualized on VMware ESXi 6.5. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, Citrix components, Citrix VDI desktops and HSD servers were hosted as virtual machines. This provides customers with complete flexibility for maintenance and capacity additions because the entire system runs on the Cisco HyperFlex hyper-converged infrastructure with stateless Cisco UCS HX-series servers. (Infrastructure virtual machines were hosted on two Cisco UCS C220 M4 Rack Servers outside of the HX cluster to deliver the highest capacity and best economics for the solution.)
- Cisco data center management: Cisco maintains industry leadership with the new Cisco UCS Manager 4.0(1b) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director insure that customer environments are consistent locally, across Cisco UCS Domains and across the globe. Cisco

UCS software suite offers increasingly simplified operational and deployment management, and it continues to widen the span of control for customer organizations' subject matter experts in compute, storage and network.

- Cisco 40G Fabric: Our 40G unified fabric story gets additional validation on 6300 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- Cisco HyperFlex Connect (HX Connect): An all-new HTML 5 based Web UI Introduced with HyperFlex v2.5 or later is available for use as the primary management tool for Cisco HyperFlex. Through this centralized point of control for the cluster, administrators can create volumes, monitor the data platform health, and manage resource use. Administrators can also use this data to predict when the cluster will need to be scaled.
- Cisco HyperFlex storage performance: Cisco HyperFlex provides industry-leading hyper converged storage performance that efficiently handles the most demanding I/O bursts (for example, login storms), high write throughput at low latency, delivers simple and flexible business continuity and helps reduce storage cost per desktop.
- Cisco HyperFlex agility: Cisco HyperFlex System enables users to seamlessly add, upgrade or remove storage from the infrastructure to meet the needs of the virtual desktops.
- Cisco HyperFlex vCenter integration: Cisco HyperFlex plugin for VMware vSphere provides easy-button automation for key storage tasks such as storage provisioning and storage resize, cluster health status and performance monitoring directly from the vCenter web client in a single pane of glass. Experienced vCenter administrators have a near zero learning curve when HyperFlex is introduced into the environment.
- Citrix Virtual Apps and Desktops advantage: Citrix VAD follows a new unified product architecture that supports both hosted-shared desktops and applications (RDS) and complete virtual desktops (VDI). This new Citrix release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of user increase. In addition, HDX enhancements help to optimize performance and improve the user experience across a variety of endpoint device types, from workstations to mobile devices including laptops, tablets, and smartphones.
- Optimized for performance and scale. For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the Citrix HSD virtual machines did not exceed the number of hyper-threaded (logical) cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for the virtual machines running virtualized RDS systems.
- Provisioning desktop machines made easy: Citrix Provisioning Services provisions hosted virtual desktops as well as hosted shared desktop virtual machines for this solution using a single method for both, the non-persistent HSD and VDI desktops. Dedicated user persistent desktops were provisioned with Citrix Machine Creation Services.

All-Flash Versus Hybrid

The initial HyperFlex product release featured hybrid converged nodes, which use a combination of solid-state disks (SSDs) for the short-term storage caching layer, and hard disk drives (HDDs) for the long-term storage capacity layer. The hybrid HyperFlex system is an excellent choice for entry-level or midrange storage solutions, and hybrid solutions have been successfully deployed in many non-performance sensitive virtual environments. Meanwhile, there is significant growth in deployment of highly performance sensitive and mission critical applications. The primary challenge to the hybrid HyperFlex system from these highly performance sensitive applications, is their increased sensitivity to high storage latency. Due to the characteristics of the spinning hard

disks, it is unavoidable that their higher latency becomes the bottleneck in the hybrid system. Ideally, if all of the storage operations were to occur in the caching SSD layer, the hybrid system's performance will be excellent. However, in several scenarios, the amount of data being written and read exceeds the caching layer capacity, placing larger loads on the HDD capacity layer, and the subsequent increases in latency will naturally result in reduced performance.

Cisco All-Flash HyperFlex systems are an excellent option for customers with a requirement to support high performance, latency sensitive workloads. With a purpose built, flash-optimized and high-performance log based filesystem, the Cisco All-Flash HyperFlex system provides:

- Predictable high performance across all the virtual machines on HyperFlex All-Flash and compute-only nodes in the cluster.
- Highly consistent and low latency, which benefits data-intensive applications and databases such as Microsoft SQL and Oracle.
- Support for NVMe caching SSDs, offering an even higher level of performance.
- Future ready architecture that is well suited for flash-memory configuration:
  - Cluster-wide SSD pooling maximizes performance and balances SSD usage so as to spread the wear.
  - A fully distributed log-structured filesystem optimizes the data path to help reduce write amplification.
  - Large sequential writing reduces flash wear and increases component longevity.
  - Inline space optimization, such as deduplication and compression, minimizes data operations and reduces wear.
- Lower operating cost with the higher density drives for increased capacity of the system.
- Cloud scale solution with easy scale-out and distributed infrastructure and the flexibility of scaling out independent resources separately.

Cisco HyperFlex support for hybrid and all-flash models now allows customers to choose the right platform configuration based on their capacity, applications, performance, and budget requirements. All-flash configurations offer repeatable and sustainable high performance, especially for scenarios with a larger working set of data, in other words, a large amount of data in motion. Hybrid configurations are a good option for customers who want the simplicity of the Cisco HyperFlex solution, but their needs focus on capacity-sensitive solutions, lower budgets, and fewer performance-sensitive applications.

Cisco HyperFlex Compute-Only Nodes

All current model Cisco UCS M4 and M5 generation servers, except the Cisco UCS C880 M4 and Cisco UCS C880 M5, may be used as compute-only nodes connected to a Cisco HyperFlex cluster, along with a limited number of previous M3 generation servers. Any valid CPU and memory configuration is allowed in the compute-only nodes, and the servers can be configured to boot from SAN, local disks, or internal SD cards. The following servers may be used as compute-only nodes:

- Cisco UCS B200 M3 Blade Server
- Cisco UCS B200 M4 Blade Server
- Cisco UCS B200 M5 Blade Server
- Cisco UCS B260 M4 Blade Server

- Cisco UCS B420 M4 Blade Server
- Cisco UCS B460 M4 Blade Server
- Cisco UCS B480 M5 Blade Server
- Cisco UCS C220 M3 Rack-Mount Servers
- Cisco UCS C220 M4 Rack-Mount Servers
- Cisco UCS C220 M5 Rack-Mount Servers
- Cisco UCS C240 M3 Rack-Mount Servers
- Cisco UCS C240 M4 Rack-Mount Servers
- Cisco UCS C240 M5 Rack-Mount Servers
- Cisco UCS C460 M4 Rack-Mount Servers
- Cisco UCS C480 M5 Rack-Mount Servers

Cisco HyperFlex Data Platform Software

The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology, exceeding the boundaries of first-generation hyperconverged infrastructures. The data platform has all the features expected in an enterprise shared storage system, eliminating the need to configure and maintain complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:

- Data protection creates multiple copies of the data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- Stretched clusters allow nodes to be evenly split between two physical locations, keeping a duplicate copy of all data in both locations, thereby providing protection in case of an entire site failure.
- Logical availability zones provide multiple logical grouping of nodes and distributes the data across these groups in such a way that no single group has more than one copy of the data. This enables enhanced protection from node failures, allowing for more nodes to fail while the overall cluster remains online.
- Deduplication is always on, helping reduce storage requirements in virtualization clusters in which multiple operating system instances in guest virtual machines result in large amounts of replicated data.
- Compression further reduces storage requirements, reducing costs, and the log-structured file system is designed to store variable-sized blocks, reducing internal fragmentation.
- Replication copies virtual machine level snapshots from one Cisco HyperFlex cluster to another, to facilitate recovery from a cluster or site failure, through a failover to the secondary site of all virtual machines.
- Encryption stores all data on the caching and capacity disks in an encrypted format, to prevent accidental data loss or data theft. Key management can be done using local Cisco UCS Manager managed keys, or third-party Key Management Systems (KMS) through the Key Management Interoperability Protocol (KMIP).
- Thin provisioning allows large volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a "pay as you grow" proposition.

- Fast, space-efficient clones rapidly duplicate virtual storage volumes so that virtual machines can be cloned simply through metadata operations, with actual data copied only for write operations.
- Snapshots help facilitate backup and remote-replication operations, which are needed in enterprises that require always-on data availability.

Cisco Desktop Virtualization Solutions: Data Center

The Evolving Workplace

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

This workforce is also increasingly mobile, conducting business in traditional offices, conference rooms across the enterprise campus, home offices, on the road, in hotels, and at the local coffee shop. This workforce wants to use a growing array of client computing and mobile devices that they can choose based on personal preference. These trends are increasing pressure on IT to ensure protection of corporate data and prevent data leakage or loss through any combination of user, endpoint device, and desktop access scenarios (Figure 1).

These challenges are compounded by desktop refresh cycles to accommodate aging PCs and bounded local storage and migration to new operating systems, specifically Microsoft Windows 10 and productivity tools, specifically Microsoft Office 2016.

Figure 1 Cisco Data Center Partner Collaboration

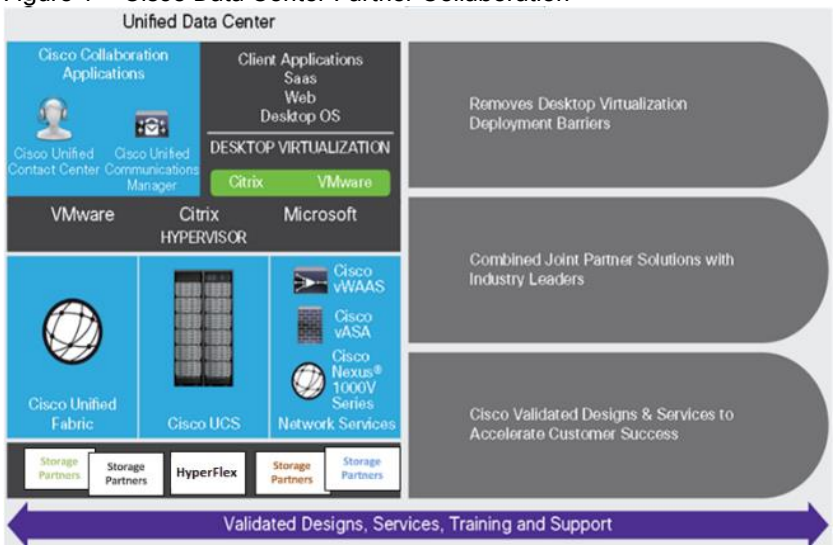

Some of the key drivers for desktop virtualization are increased data security, the ability to expand and contract capacity and reduced TCO through increased control and reduced management costs.

Cisco Desktop Virtualization Focus

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

Simplified

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

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

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

Secure

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

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

Scalable

Growth of a desktop virtualization solution is accelerating, so a solution must be able to scale, and scale predictably, with that growth. The Cisco Desktop Virtualization Solutions support high virtual-desktop density (desktops per server) and additional servers scale with near-linear performance. Cisco data center infrastructure provides a flexible platform for growth and improves business agility. Cisco UCS Manager service profiles allow on-demand desktop provisioning and make it just as easy to deploy dozens of desktops as it is to deploy thousands of desktops.

Cisco HyperFlex servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 3.0 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco HyperFlex helps maintain data availability and optimal performance during boot and login

storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs based on VMware, Citrix, Cisco HyperFlex solutions have demonstrated scalability and performance, with up to 2000 hosted virtual desktops and hosted shared desktops up and running in ~10 minutes.

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

Savings and Success

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

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

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

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

Use Cases

- Healthcare: Mobility between desktops and terminals, compliance, and cost
- Federal government: Teleworking initiatives, business continuance, continuity of operations (COOP), and training centers
- Financial: Retail banks reducing IT costs, insurance agents, compliance, and privacy
- Education: K-12 student access, higher education, and remote learning
- State and local governments: IT and service consolidation across agencies and interagency security
- Retail: Branch-office IT cost reduction and remote vendors
- Manufacturing: Task and knowledge workers and offshore contractors
- Microsoft Windows 10 migration

- Graphic intense applications
- Security and compliance initiatives
- Opening of remote and branch offices or offshore facilities
- Mergers and acquisitions

Figure 2 illustrates the Citrix Virtual Apps and Desktops on vSphere 6.5 built on Cisco UCS components and the network connections. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Figure 2 Citrix Virtual Apps and Desktops 1808 on vSphere 6.5 Built on Cisco Unified Computing System

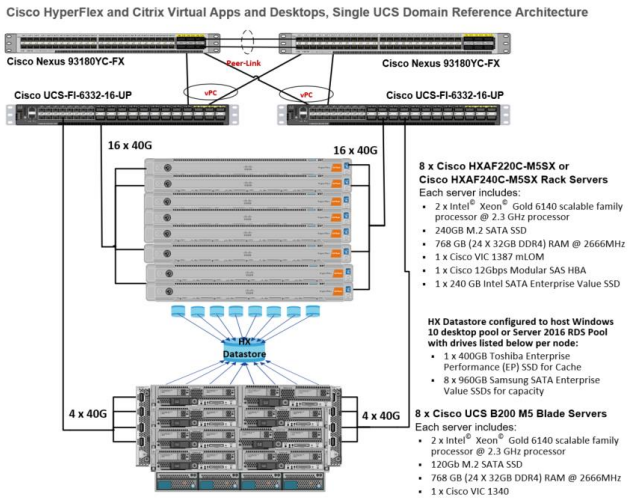

Physical Topology

The Cisco HyperFlex system is composed of a pair of Cisco UCS 6200/6300 series Fabric Interconnects, along with up to 32 HXAF-Series rack mount servers per cluster. In addition, up to 32 compute only servers can be added per cluster. Adding Cisco UCS 5108 blade chassis allows use of Cisco UCS B200 M5 blade servers for additional compute resources in a hybrid cluster design. Cisco UCS C240 and C220 servers can also be used for additional compute resources. The Fabric Interconnects both connect to every HX-Series rack mount server and both connect to every Cisco UCS 5108-blade chassis. Upstream network connections, also referred to as “northbound” network connections are made from the Fabric Interconnects to the customer datacenter network at the time of installation.

For this study, we uplinked the Cisco 6332 UP Fabric Interconnects to Cisco Nexus 93180YC-FX switches.

Figure 3 and Figure 4 illustrate the hyperconverged and hybrid hyperconverged, plus compute only topologies.

Figure 3 Cisco HyperFlex Standard Topology

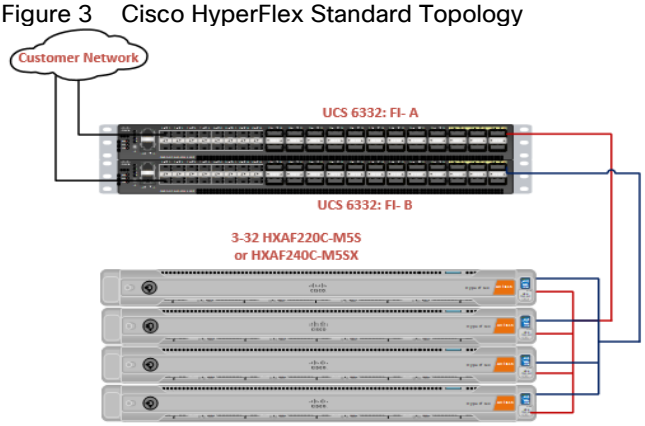

Figure 4 Cisco HyperFlex Hyperconverged Plus Compute-Only Node Topology

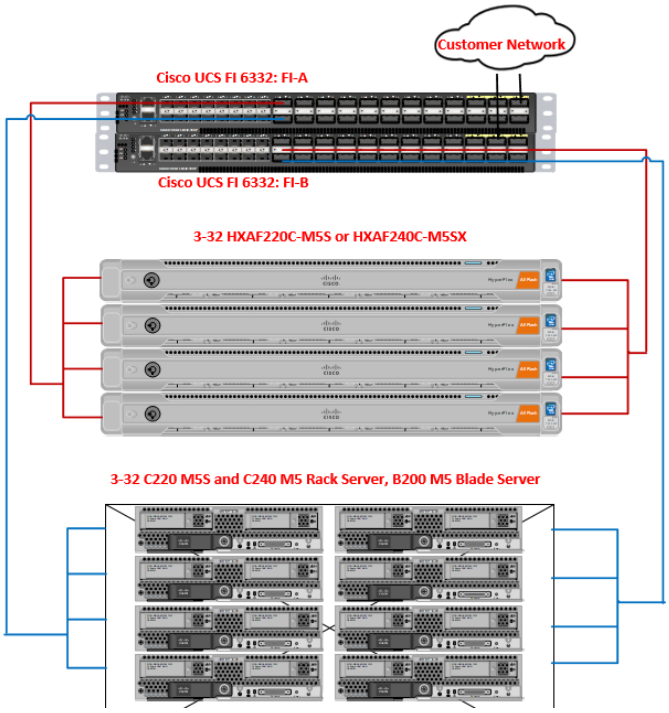

Fabric Interconnects

Fabric Interconnects (FI) are deployed in pairs, wherein the two units operate as a management cluster, while forming two separate network fabrics, referred to as the A side and B side fabrics. Therefore, many design elements will refer to FI A or FI B, alternatively called fabric A or fabric B. Both Fabric Interconnects are active at all times, passing data on both network fabrics for a redundant and highly available configuration. Management services, including Cisco UCS Manager, are provided by the two FIs but in a clustered manner, where one FI is the primary, and one is secondary, with a roaming clustered IP address. This primary/secondary relationship is only for the management cluster and has no effect on data transmission.

Fabric Interconnects have the following ports, which must be connected for proper management of the Cisco UCS domain:

- Mgmt: A 10/100/1000 Mbps port for managing the Fabric Interconnect and the Cisco UCS domain through GUI and CLI tools. Also used by remote KVM, IPMI and SoL sessions to the managed servers within the domain. This is typically connected to the customer management network.

- L1: A cross connect port for forming the Cisco UCS management cluster. This is connected directly to the L1 port of the paired Fabric Interconnect using a standard CAT5 or CAT6 Ethernet cable with RJ45 plugs. It is not necessary to connect this to a switch or hub.
- L2: A cross connect port for forming the Cisco UCS management cluster. This is connected directly to the L2 port of the paired Fabric Interconnect using a standard CAT5 or CAT6 Ethernet cable with RJ45 plugs. It is not necessary to connect this to a switch or hub.
- Console: An RJ45 serial port for direct console access to the Fabric Interconnect. Typically used during the initial FI setup process with the included serial to RJ45 adapter cable. This can also be plugged into a terminal aggregator or remote console server device.

HX-Series and C-Series Rack-Mount Servers

The HX-Series converged servers and optional Cisco UCS C-Series compute only servers are connected directly to the Cisco UCS Fabric Interconnects in Direct Connect mode. This option enables Cisco UCS Manager to manage the HX-Series rack-mount Servers and Cisco UCS C-Series servers using a single cable for both management traffic and data traffic. The HXAF220C-M5SX.HXAF240C-M5SX, C240-M5 and C220 M5 servers are configured with the Cisco VIC 1387 network interface card (NIC) installed in a modular LAN on motherboard (MLOM) slot, which has dual 40 Gigabit Ethernet (GbE) ports. The standard and redundant connection practice is to connect port 1 of the VIC 1387 to a port on FI A, and port 2 of the VIC 1387 to a port on FI B (Figure 5).

Failure to follow this cabling practice can lead to errors, discovery failures, and loss of redundant connectivity.

Figure 5 HX-Series and Cisco UCS C-Series Server Connectivity

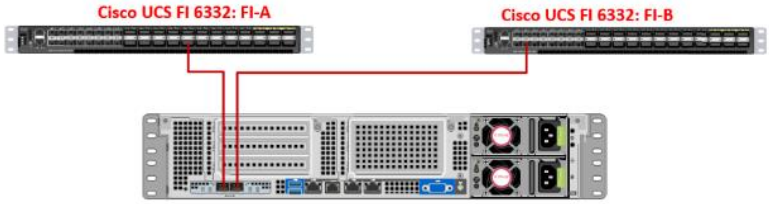

Cisco UCS B-Series Blade Servers

Hybrid HyperFlex clusters also incorporate 1-16 Cisco UCS B200 M5 blade servers for additional compute capacity. Like all other Cisco UCS B-series blade servers, the Cisco UCS B200 M5 must be installed within a Cisco UCS 5108 blade chassis. The blade chassis comes populated with 1-4 power supplies, and 8 modular cooling fans. In the rear of the chassis are two bays for installation of Cisco Fabric Extenders. The Fabric Extenders (also commonly called IO Modules, or IOMs) connect the chassis to the Fabric Interconnects. Internally, the Fabric Extenders connect to the Cisco VIC 1340 card installed in each blade server across the chassis backplane. The standard connection practice is to connect 1-4 40 GbE or 2 x 40 (native) GbE links from the left side IOM, or IOM 1, to FI A, and to connect the same number of 40 GbE links from the right side IOM, or IOM 2, to FI B (Figure 6). All other cabling configurations are invalid, and can lead to errors, discovery failures, and loss of redundant connectivity.

Figure 6 Cisco UCS 5108 Chassis Connectivity

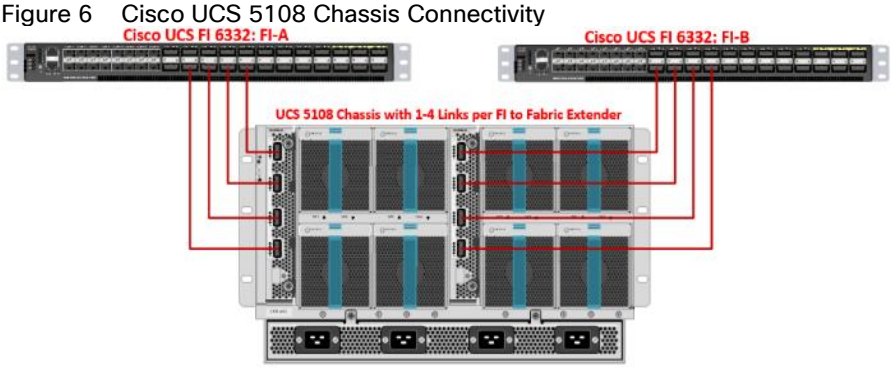

Logical Network Design

The Cisco HyperFlex system has communication pathways that fall into four defined zones (Figure 6):

- Management Zone: This zone comprises the connections needed to manage the physical hardware, the hypervisor hosts, and the storage platform controller virtual machines (SCVM). These interfaces and IP addresses need to be available to all staff who will administer the HX system, throughout the LAN/WAN. This zone must provide access to Domain Name System (DNS) and Network Time Protocol (NTP) services and allow Secure Shell (SSH) communication. In this zone are multiple physical and virtual components:
  - Fabric Interconnect management ports.
  - Cisco UCS external management interfaces used by the servers and blades, which answer through the FI management ports.
  - ESXi host management interfaces.
  - Storage Controller virtual machine management interfaces.
  - A roaming HX cluster management interface.
- VM Zone: This zone comprises the connections needed to service network IO to the guest virtual machines that will run inside the HyperFlex hyperconverged system. This zone typically contains multiple VLANs that are trunked to the Cisco UCS Fabric Interconnects through the network uplinks and tagged with 802.1Q VLAN IDs. These interfaces and IP addresses need to be available to all staff and other computer endpoints which need to communicate with the guest virtual machines in the HX system, throughout the LAN/WAN.
- Storage Zone: This zone comprises the connections used by the Cisco HX Data Platform software, ESXi hosts, and the storage controller virtual machines to service the HX Distributed Data Filesystem. These interfaces and IP addresses need to be able to communicate with each other at all times for proper operation. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B, and vice-versa. This zone is primarily jumbo frame traffic therefore; jumbo frames must be enabled on the Cisco UCS uplinks. In this zone are multiple components:
  - A vmkernel interface used for storage traffic for each ESXi host in the HX cluster.
  - Storage Controller VM storage interfaces.
  - A roaming HX cluster storage interface.

- VMotion Zone: This zone comprises the connections used by the ESXi hosts to enable vMotion of the guest virtual machines from host to host. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B, and vice-versa.

Figure 7 illustrates the logical network design.

Figure 7 Logical Network Design

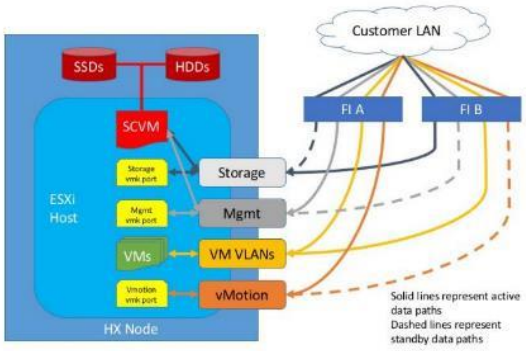

Logical Availability Zones

Larger scale HyperFlex clusters are subject to higher failure risks, simply due to the number of nodes in the cluster. While any individual node's risk of failure is the same no matter how many nodes there are, with clusters up to 64 nodes in size, there is a logically higher probability that a single node could fail, when compared to a cluster with fewer nodes. To mitigate these risks in larger scale clusters, a HyperFlex cluster of eight nodes or more, can be configured with a feature called Logical Availability Zones (LAZ). The Logical Availability Zones feature groups 2 or more HyperFlex nodes together into a logically defined zone, a minimum of 3 zones are created, and the data in the cluster is distributed in such a way that no blocks are written to the nodes within a single zone more than once. Due to this enhanced distribution pattern of data across zones, wherein each zone has multiple servers, clusters with LAZ enabled can typically withstand more failures than clusters that operate without this feature enabled. The number of failures that can tolerated varies depending on the number of zones in the cluster, and the number of servers in each of the zones. Generally speaking, multiple node failures across one or two zones will be tolerated better, and with less risk than multiple nodes failing across three or more zones. Note that the failure tolerance shown in the HyperFlex Connect dashboard will always present a "worst case scenario" view, meaning that even though the dashboard may state that two failures can be tolerated, in fact two servers could fail and the cluster can remain online, and the failure tolerance may still remain at two.

Logical availability zones should not be confused with the concept of fault domains. An example of a fault domain would be a subset of the nodes in a single HyperFlex cluster being powered by one uninterruptable power supply (UPS) or connected to one power distribution unit (PDU), meanwhile the remaining nodes would be connected to another UPS or PDU. If one of the UPS' or PDUs were to fail, then there would be a simultaneous failure of multiple nodes. While LAZ may actually prevent the cluster from failing in this scenario, to guarantee it would require that the zone membership be manually controlled, so that a failure of all of the servers protected by a single UPS or PDU, would be distributed in such a way that it would not cause an outage. The LAZ feature is not designed to be manually configured in this way, instead the zone membership is determined automatically by the system. If a HyperFlex cluster needs to be physically split in half due to a physical limitation, such as the UPS example above, or a distance requirement for fault tolerance, then the cluster should be built as a stretched cluster instead of using LAZ.

Figure 8 illustrates an example of the data distribution method for clusters with Logical Availability Zones enabled, set to replication factor 3, where each zone only contains one of the three copies of the data in the cluster. This cluster consists of eight nodes, which the system configures into four zones.

Figure 8 Data Distribution Method for Clusters with Logical Availability Zones Enabled

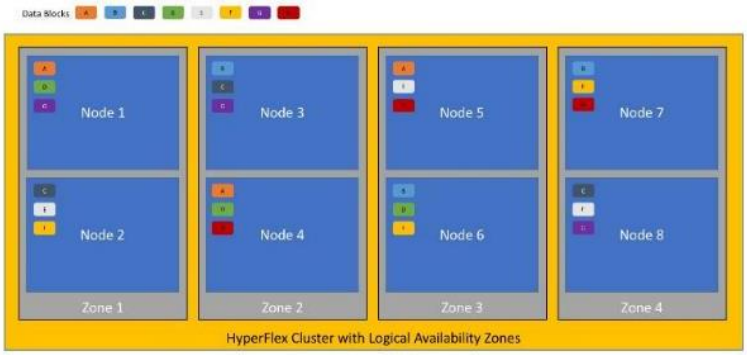

Logical availability zones are subject to the following requirements and limitations:

- Only HyperFlex clusters with 8 nodes or more can be configured with logical availability zones during the installation process.
- Logical Availability Zones can be enabled during the HyperFlex cluster installation, or it can be enabled through the command line at a later time. It is recommended to enable this feature during installation, in order to avoid a large migration and reorganization of data across the cluster, which would be necessary to comply with the data distribution rules if LAZ is turned on in a cluster already containing data.
- The number of zones can be manually specified as 3, 4, 5, or you can allow the installer to automatically choose, which is the recommended setting.
- The HyperFlex cluster determines which nodes participate in each zone, and this configuration cannot be modified.
- To maintain the most balanced consumption of space and data distribution, it is recommended that the number of nodes in a cluster are whole multiples of 3, 4, 5, or 7. For example, 8 nodes would evenly divide into 4 zones of 2 servers each, and 9 nodes would divide evenly into 3 zones of 3 servers each. Eleven nodes would create an unbalanced number of nodes across the zones, leading to unbalanced space consumption on the nodes.
- In addition to the previous point, expansion of a cluster should be done in multiples of the number of zones, when the cluster is operating with LAZ enabled. Expanding in such a way preserves a matched number of nodes in each zone and prevents any unbalance of space consumption. For example, a cluster with 3 zones should be expanded by adding 3 more nodes, because adding only 1 or 2 nodes would lead to an imbalance, as would adding 4 nodes.

The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-FX switches
- Two Cisco UCS 6332 fabric interconnects
- Eight Cisco HXAF C220 M5 rack servers running HyperFlex data platform version 3.5.1a
- Eight Cisco UCS B200 M5 blade server running HyperFlex data platform version 3.5.1a as compute-only nodes.

For desktop virtualization, the deployment includes Citrix Virtual Apps and Desktops 1808 running on VMware vSphere 6.5. The design is intended to provide a large-scale building block for both HSD and persistent/non-persistent desktops with following density per 16-node configuration:

- 800 Citrix HSD server desktop sessions
- 800 Windows 10 non-persistent desktops with Citrix PVS
- 400 Windows 10 persistent full clone virtual desktops with Citrix MCS

All of the Windows 10 virtual desktops have been provisioned with 4GB of memory for this validated design. Typically, persistent desktop users may desire more memory. If 4GB or more of memory is needed, additional memory channels on the Cisco HXAF220c-M5S HX-Series rack server and Cisco UCS B200 M5 servers should be populated.

Data provided here will allow customers to run HSD server sessions and VDI desktops to suit their environment. For example, additional Cisco HX server can be deployed in compute-only manner to increase compute capacity or additional drives can be added in existing server to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features. This document guides you through the low-level steps for deploying the base architecture, as shown in Figure 2. These procedures cover everything from physical cabling to network, compute and storage device configurations.

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a Cisco Validated Design for various type of Virtual Desktop workloads on Cisco HyperFlex. Configuration guidelines are provided that refer to which redundant component is being configured with each step. For example, Cisco Nexus A or Cisco Nexus B identifies a member in the pair of Cisco Nexus switches that are configured. Cisco UCS 6332 UP Fabric Interconnects are similarly identified. Additionally, this document details the steps for provisioning multiple Cisco UCS and HyperFlex hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, VM-Host-HSD-01, VM-Host-VDI-01 and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure.

Solution Design

This section describes the infrastructure components used in the solution outlined in this study.

Cisco Unified Computing System

Cisco UCS Manager (UCSM) provides unified, embedded management of all software and hardware components of Cisco Unified Computing System and Cisco HyperFlex through an intuitive GUI, a CLI, and an XML API. The manager provides a unified management domain with centralized management capabilities and can control multiple chassis and thousands of virtual machines.

Cisco UCS is a next-generation data center platform that unites computing, networking, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency; lossless 40 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. It is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain.

Cisco Unified Computing System Components

The main components of Cisco UCS are:

- Compute: The system is based on an entirely new class of computing system that incorporates blade, rack and hyperconverged servers based on Intel Xeon scalable family processors.
- Network: The system is integrated on a low-latency, lossless, 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables needed, and by decreasing the power and cooling requirements.
- 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 virtualized environments to better support changing business and IT requirements.
- Storage: The Cisco HyperFlex rack servers provide high performance, resilient storage using the powerful HX Data Platform software. Customers can deploy as few as three nodes (replication factor 2/3) depending on their fault tolerance requirements. These nodes form a HyperFlex storage and compute cluster. The onboard storage of each node is aggregated at the cluster level and automatically shared with all of the nodes. Storage resources are managed from the familiar VMware vCenter web client, extending the capability of vCenter administrators.
- Management: Cisco UCS uniquely integrates all system components, enabling the entire solution to be managed as a single entity by Cisco UCS Manager. The manager has an intuitive GUI, a CLI, and a robust API for managing all system configuration processes and operations.

Figure 9 Cisco HyperFlex Family Overview

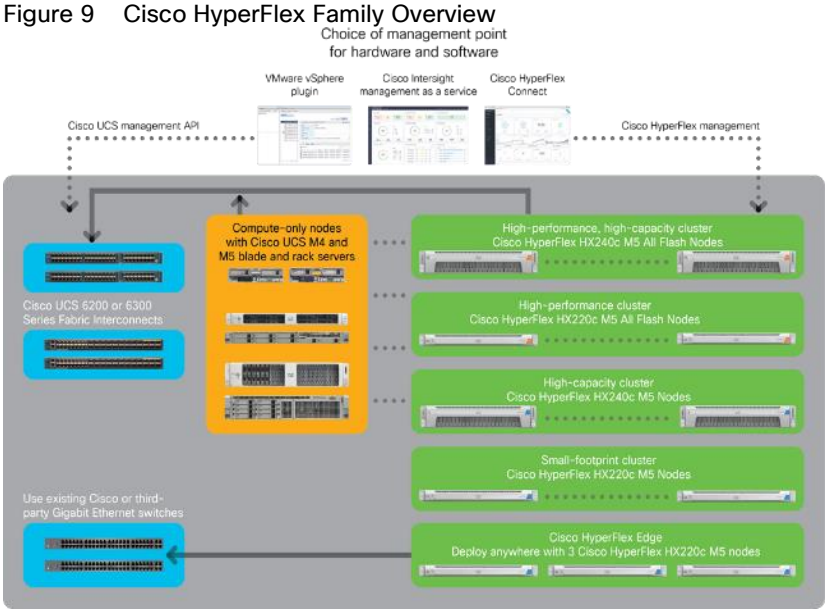

Cisco UCS and Cisco HyperFlex are designed to deliver:

- Reduced TCO and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system that unifies the technology in the data center; the system is managed, serviced and tested as a whole.
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand.
- Industry standards supported by a partner ecosystem of industry leaders.

Cisco UCS Manager provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System across multiple chassis, rack servers, and thousands of virtual machines. Cisco UCS Manager manages Cisco UCS as a single entity through an intuitive GUI, a command-line interface (CLI), or an XML API for comprehensive access to all Cisco UCS Manager Functions.

The Cisco HyperFlex system provides a fully contained virtual server platform, with compute and memory resources, integrated networking connectivity, a distributed high performance log-structured file system for virtual machine storage, and the hypervisor software for running the virtualized servers, all within a single Cisco UCS management domain.

Figure 10 Cisco HyperFlex System Overview

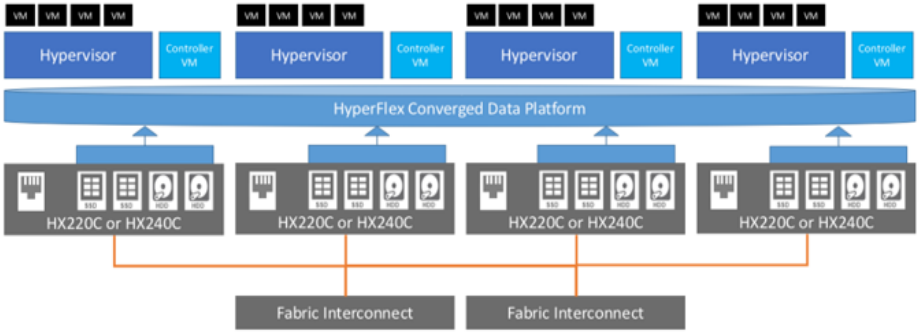

Enhancements for Version 3.5.1a

The Cisco HyperFlex system has several new capabilities and enhancements in version 3.5.1a (see Figure 11)

Figure 11 Addition of HX All-Flash Nodes in 3.5.1a

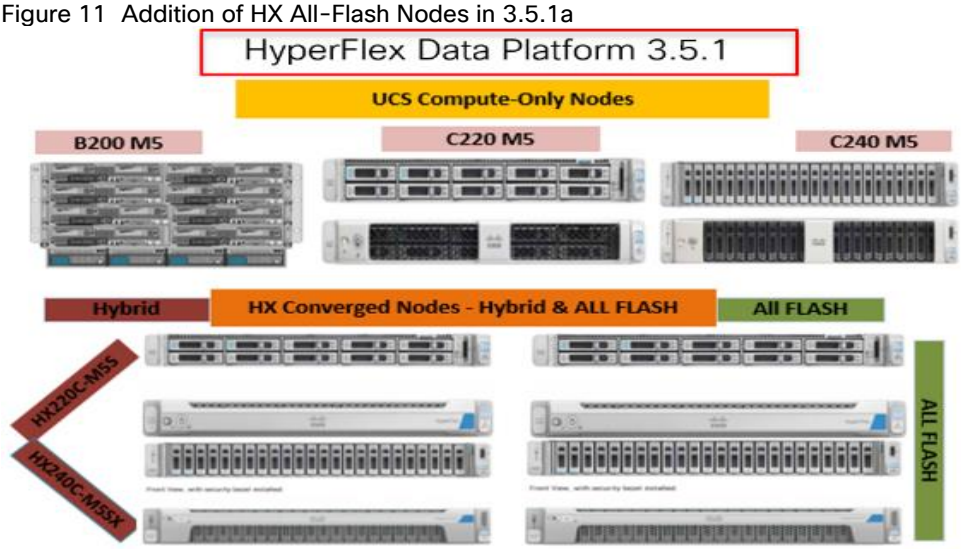

- New All-Flash HX server models are added to the Cisco HyperFlex product family that offer all flash storage using SSDs for persistent storage devices.
- Cisco HyperFlex now support the latest generation of Cisco UCS software, Cisco UCS Manager 4.0.(1b) and beyond. For new All-Flash deployments, verify that Cisco UCS Manager 4.0.(1b) or later is installed.
- Support for adding external storage (iSCSI or Fibre Channel) adapters to HX nodes during HX Data Platform software installation, which simplifies the process to connect external storage arrays to the HX domain.
- Support for adding HX nodes to an existing Cisco UCS-FI domain.
- Support for Cisco HyperFlex Sizer – A new end to end sizing tool for compute, capacity and performance.
- Multiple Hypervisors - Support for Microsoft Hyper-V in addition to already supported VMware ESXi
- Stretched cluster - for High Availability across Datacenter locations
- Kubernetes FlexVolume driver - Turnkey Kubernetes persistent storage for enterprises & foundation for Cisco Container Platform.
- Higher Scale (32 Converged + 32 Compute-Only) and Enhanced resiliency through Logical Availability Zones (LAZ)
- Intel Optane NVMe support for higher drive level performance and higher endurance
- Large Form Factor - HX M5 240 LFF chassis with 6TB, 8TB drives options
- Advanced Disaster Recovery workflows
- Cisco Intersight support across hypervisor platforms
- Expanded HyperFlex Edge configuration options
- Linked mode - HyperFlex Plugin Support for vCenter's enhanced linked mode feature

- REST APIs - Cisco HyperFlex Systems REST API Getting Started Guide on Cisco DevNet
- New All-Flash and Hybrid HX M5 server models are added to the Cisco HyperFlex product family
- Cisco Smart Licensing—Support for Cisco Smart Software Manager satellite. Please refer to the https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/Installation_VMWare_ESXi/3_5/b_HyperFlexSystems_Installation_Guide_for_VMware_ESXi_3_5.html for more details.
- M5 Servers
- Key release highlights:
  - Same software feature set as HX 3.5.1a
  - Support for M5 servers in HyperFlex.
  - Enablement for Cisco HX240c M5 and HXAF240c M5 servers:
  - Dual CPU—Intel Xeon processor scalable family
  - Up to 3TB DRAM—Recommended minimum of 256 GB DRAM
  - M.2 Drive—For ESX Boot and for Storage Controller Virtual Machine
  - Up to 2 GPUs—M10, P40, AMD 7150 x 2
  - Dedicated rear slots for caching
- Enablement for Cisco HX220c M5 and HXAF220c M5 servers:
  - Dual CPU (Except Edge)—Intel Xeon processor scalable family
  - Up to 3TB DRAM—Recommended minimum of 256 GB DRAM
  - 8 x Data Drives (SATA/SAS)
  - M.2 Drive—For ESX Boot and for Storage Controller Virtual Machine
- M4/M5 support in the same cluster.
  - A mixed cluster is defined by having both M4 and M5 HX converged nodes within the same storage cluster.
  - HyperFlex Edge does not support mixed clusters.
  - SED SKUs do not support mixed clusters.
- Peripherals
  - Option for 6-8 drives in HX220C-M5S and HXAF220C-M5S nodes.
  - Up to two GPUs for HX240C-M5SX and HXAF240C-M5SX nodes

Cisco UCS Fabric Interconnect

The Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6300 Series offers line-rate, low-latency, lossless 40 Gigabit Ethernet, FCoE, and Fibre Channel functions.

The fabric interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers, Cisco UCS C-Series and HX-Series rack servers and Cisco UCS 5100 Series Blade Server Chassis. All servers, attached to the fabric interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6300 Series provides both LAN and SAN connectivity for all blades in the domain.

For networking, the Cisco UCS 6300 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 40 Gigabit Ethernet on all ports, 2.56-terabit (Tb) switching capacity, and 320 Gbps of bandwidth per chassis, independent of packet size and enabled services. The product series supports Cisco low-latency, lossless, 40 Gigabit Ethernet unified network fabric capabilities, increasing the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnects support multiple traffic classes over a lossless Ethernet fabric, from the blade server through the interconnect. 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.

Figure 12 Cisco UCS 6332 Series Fabric Interconnect

Figure 13 Cisco UCS 6332-16UP Fabric Interconnect

Cisco HyperFlex HX-Series Nodes

Cisco HyperFlex systems are based on an end-to-end software-defined infrastructure, combining software-defined computing in the form of Cisco Unified Computing System servers; software-defined storage with the powerful Cisco HX Data Platform and software-defined networking with the Cisco UCS fabric that will integrate smoothly with Cisco Application Centric Infrastructure. Together with a single point of connectivity and hardware management, these technologies deliver a pre-integrated and adaptable cluster that is ready to provide a unified pool of resources to power applications as your business needs dictate.

A Cisco HyperFlex cluster requires a minimum of three HX-Series nodes (with disk storage). Data is being replicated across at least two of these nodes, and a third node is required for continuous operation in the event of a single-node failure. Each node that has disk storage is equipped with at least one high-performance SSD drive for data caching and rapid acknowledgment of write requests. Each node is also equipped with the platform's physical capacity of either spinning disks or enterprise-value SSDs for maximum data capacity.

Cisco UCS HXAF220c-M5S Rack Server

The HXAF220c M5 servers extend the capabilities of Cisco's HyperFlex portfolio in a 1U form factor with the addition of the Intel Xeon Processor Scalable Family, 24 DIMM slots for 2666MHz DIMMs, up to 128GB individual DIMM capacities and up to 3.0TB of total DRAM capacities.

This small footprint configuration of Cisco HyperFlex all-flash nodes contains one M.2 SATA SSD drive that act as the boot drives, a single 240-GB solid-state disk (SSD) data-logging drive, a single 400-GB SSD write-log drive, and up to eight 3.8-terabyte (TB) or 960-GB SATA SSD drives for storage capacity. A minimum of three nodes and a maximum of sixteen nodes can be configured in one HX cluster. For detailed information, see the Cisco HyperFlex HXAF220c-M5S specsheet.

Table 1 HXAF220c-M5SX Server Options

HXAF220c-M5SX options		Hardware Required
Processors		Chose a matching pair of Intel Xeon Processor Scalable Family CPUs
Memory		192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 MHz 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA
SSDs	Standard	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One 400 GB 2.5 Inch Enterprise Performance 12G SAS SSD, or one 1.6 TB 2.5 Inch Enterprise Performance NVMe SSD, or one 375 GB 2.5 Inch Optane Extreme Performance SSD Six to eight 3.8 TB 2.5 Inch Enterprise Value 6G SATA SSDs, or six to eight 960 GB 2.5 Inch Enterprise Value 6G SATA SSDs
	SED	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One 800 GB 2.5 Inch Enterprise Performance 12G SAS SED SSD Six to eight 3.8 TB 2.5 Inch Enterprise Value 6G SATA SED SSDs, or six to eight 960 GB 2.5 Inch Enterprise Value 6G SATA SED SSDs, or six to eight 800 GB 2.5 Inch Enterprise Performance 12G SAS SED SSDs
Network		Cisco UCS VIC1387 VIC MLOM
Boot Device		One 240 GB M.2 form factor SATA SSD
microSD Card		One 32GB microSD card for local host utilities storage
Optional		Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

Figure 14 Cisco UCS HXAF220c-M5SX Rack Server Front View

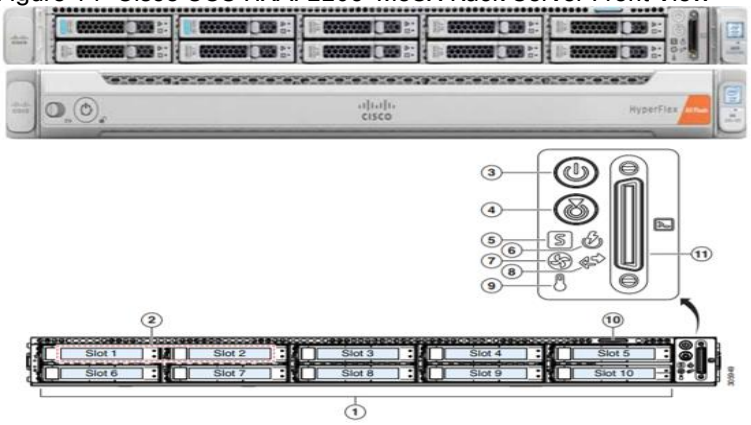

Figure 15 Cisco UCS HXAF220c-M5SX Rack Server Rear View

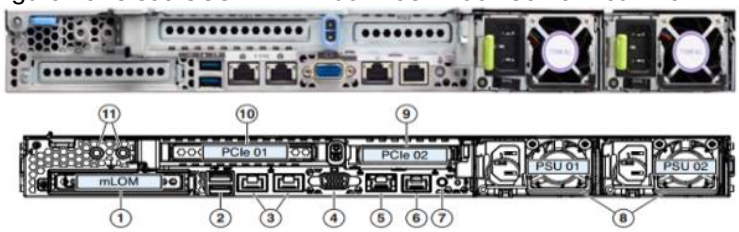

1	Modular LAN-on-motherboard (mLOM) card bay (x16)	7	Rear unit identification button/LED
2	USB 3.0 ports (two)	8	Power supplies (two, redundant as 1+1)
3	Dual 1/10-Gb Ethernet ports (LAN1 and LAN2). LAN1 is left connector and LAN2 is right connector	9	PCIe riser 2 (slot 2) (half-height, x16);
4	VGA video port (DB-15)	10	PCIe riser 1 (slot 1) (full-height, x16)
5	1-Gb Ethernet dedicated management port	11	Threaded holes for dual-hole grounding lug
6	Serial port (RJ-45 connector)	—	—

The Cisco UCS HXAF220c-M5S delivers performance, flexibility, and optimization for data centers and remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads ranging from web infrastructure to distributed databases. The Cisco UCS HXAF220c-M5SX can quickly deploy stateless physical and virtual workloads with the programmable ease of use of the Cisco UCS Manager software and simplified server access with Cisco Single Connect technology. Based on the Intel Xeon scalable family processor product family, it offers up to 1.5TB of memory using 64-GB DIMMs, up to ten disk drives, and up to 40 Gbps of I/O throughput. The Cisco UCS HXAF220c-M5Soffers exceptional levels of performance, flexibility, and I/O throughput to run your most demanding applications.

The Cisco UCS HXAF220c-M5S provides:

- Up to two multicore Intel Xeon scalable family processor for up to 56 processing cores
- 24 DIMM slots for industry-standard DDR4 memory at speeds 2666 MHz, and up to 1.5TB of total memory when using 64-GB DIMMs
- Ten hot-pluggable SAS and SATA HDDs or SSDs
- Cisco UCS VIC 1387, a 2-port, 80 Gigabit Ethernet and FCoE-capable modular (mLOM) mezzanine adapter
- Cisco FlexStorage local drive storage subsystem, with flexible boot and local storage capabilities that allow you to install and boot Hypervisor from
- Enterprise-class pass-through RAID controller
- Easily add, change, and remove Cisco FlexStorage modules

Cisco VIC 1387 MLOM Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1387 is a dual-port Enhanced Small Form-Factor Pluggable (QSFP+) 40-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)) in a modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS HX-Series Rack Servers (Figure 5). 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 that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile.

Figure 16 Cisco VIC 1387 mLOM Card

Table 2 Supported Physical Connectivity

Fabric Interconnect Model	6248	6296	6332		6332-16UP		
Port Type	10GbE	10GbE	40GbE	10GbE Breakout	40GbE	10GbE Breakout	10GbE onboard
M4 with VIC 1227	✓	✓	✗	✗	✗	✗	✗
M4 with VIC 1387	✗	✗	✓	✗	✓	✗	✗
M4 with VIC 1387 + QSA	✗	✗	✗	✗	✗	✗	✗
M5 with VIC 1387	✗	✗	✓	✗	✓	✗	✗
M5 with VIC 1387 + QSA	✓	✓	✗	✗	✗	✗	✗

Cisco HyperFlex Compute Nodes

Cisco UCS B200 M5 Blade Server

For workloads that require additional computing and memory resources, but not additional storage capacity, a compute-intensive hybrid cluster configuration is allowed. This configuration requires a minimum of three (up to sixteen) HyperFlex converged nodes with one to sixteen Cisco UCS B200 M5 Blade Servers for additional computing capacity. The HX-series Nodes are configured as described previously, and the Cisco UCS B200 M5 servers are equipped with boot drives. Using the Cisco UCS B200 M5 compute nodes also requires the Cisco UCS 5108 blade server chassis, and a pair of Cisco UCS 2300/2200 series Fabric Extenders. For detailed information, see the Cisco UCS B200 M5 Blade Server Spec Sheet.

Figure 17 Cisco UCS B200 M5 Blade Server

Cisco VIC1340 Converged Network Adapter

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 18) is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet.

The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1340 supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.

Figure 18 Cisco UCS VIC 1340

Figure 19 Cisco UCS VIC 1340 Virtual Interface Cards Deployed in the Cisco UCS B Series B200 M5 Blade Servers

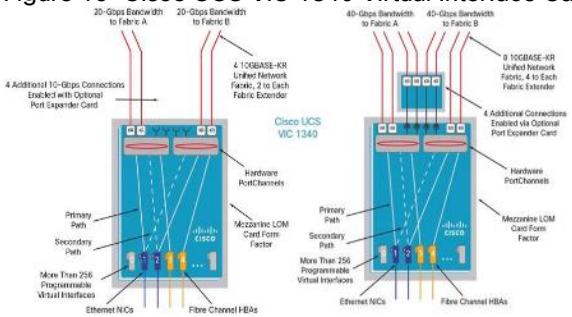

Cisco UCS 5108 Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis for today's and tomorrow's data center while helping reduce TCO.

The Cisco UCS 5108 Blade Server Chassis (Figure 20) is six Rack Units (6RU) high and can mount in an industry-standard 19-inch rack. A chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors.

Four hot-swappable power supplies are accessible from the front of the chassis, and single-phase 2500 W AC, 2500 W –48 VDC, and 2500 W 200 – 380 VDC power supplies and chassis are available. These power supplies are up to 94 percent efficient and meet the requirements for the 80 Plus Platinum rating. The power subsystem can be configured to support nonredundant, N+1 redundant, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays that can support either Cisco UCS 2000 Series Fabric Extenders or the Cisco UCS 6324 Fabric Interconnect. A passive midplane provides up to 80 Gbps of I/O bandwidth per server slot and up to 160 Gbps of I/O bandwidth for two slots. The chassis supports 40 Gigabit Ethernet standards with the 2304 Fabric Extender.

Figure 20 Cisco UCS 5108 Blade Chassis Front and Rear Views

Features and Benefits

The Cisco UCS 5108 Blade Server Chassis revolutionizes the use and deployment of blade-based systems. By incorporating unified fabric, integrated, embedded management, and fabric extender technology, the chassis uses fewer physical components, has no need for independent management, and enables greater energy efficiency than traditional blade server chassis. This simplicity eliminates the need for dedicated chassis management and blade switches, reduces cabling, and enables Cisco UCS to scale to 20 chassis without adding complexity. The Cisco UCS 5108 chassis is a critical component in delivering the Cisco UCS benefits of data center simplicity and IT responsiveness.

In addition, the Cisco UCS 5108 chassis has the architectural advantage of not having to power and cool excess switches in each chassis. With a larger power budget per blade server, Cisco can design uncompromised expandability and capabilities in its blade servers, as evidenced by the new Cisco UCS B200 M5 and B480 M5 Blade Servers. For more information, see the Cisco UCS 5100 Series Blade Server Chassis Data Sheet.

Cisco UCS 2304XP Fabric Extender

Cisco UCS 2304 Fabric Extender brings the unified fabric into the blade server enclosure, providing multiple 40 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management. It is a third-generation I/O Module (IOM) that shares the same form factor as the second-generation Cisco UCS 2200/2300 Series Fabric Extenders and is backward compatible with the shipping Cisco UCS 5108 Blade Server Chassis.

The Cisco UCS 2304 connects the I/O fabric between the Cisco UCS 6300 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic Fibre Channel over Ethernet (FCoE) fabric to connect all blades and chassis together. Fabric extender is similar to a distributed line card, it does not perform any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity and enabling Cisco UCS to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

The Cisco UCS 2304 also manages the chassis environment (power supply, fans, and blades) in conjunction with the fabric interconnect. Therefore, separate chassis management modules are not required.

Cisco UCS 2304 Fabric Extenders fit into the back of the Cisco UCS 5100 Series chassis. Each Cisco UCS 5100 Series chassis can support up to two fabric extenders, allowing increased capacity and redundancy (Figure 21).

The Cisco UCS 2304 Fabric Extender has four 40 Gigabit Ethernet, FCoE-capable, Quad Small Form-Factor Pluggable (QSFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2304 can provide one 40 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis, giving it a total eight 40G interfaces to the compute. Typically configured in pairs for redundancy, two fabric extenders provide up to 320 Gbps of I/O to the chassis.

Figure 21 Cisco UCS 2304XP Fabric Extender

Cisco UCS C220 M5 Rack Server

The Cisco UCS C220 M5 Rack Server is an enterprise-class infrastructure server in an 1RU form factor. It incorporates the Intel Xeon processor E5-2600 v4 and v3 product family, next-generation DDR4 memory, and 12-Gbps SAS throughput, delivering significant performance and efficiency gains. Cisco UCS C220 M5 Rack Server can be used to build a compute-intensive hybrid HX cluster, for an environment where the workloads require additional computing and memory resources but not additional storage capacity, along with the HX-series converged nodes. This configuration contains a minimum of three (up to eight) HX-series converged nodes with one to eight Cisco UCS C220 M5 Rack Servers for additional computing capacity.

Figure 22 Cisco UCS C220 M5 Rack Server

Cisco HyperFlex HX Data Platform Administration Plug-in

The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology, exceeding the boundaries of first-generation hyperconverged infrastructures. The data platform has all the features that you would expect of an enterprise shared storage system, eliminating the need to configure and maintain complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:

- Replication replicates data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- Deduplication is always on, helping reduce storage requirements in virtualization clusters in which multiple operating system instances in client virtual machines result in large amounts of replicated data.
- Compression further reduces storage requirements, reducing costs, and the log- structured file system is designed to store variable-sized blocks, reducing internal fragmentation.
- Thin provisioning allows large volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a "pay as you grow" proposition.
- Fast, space-efficient clones rapidly replicate storage volumes so that virtual machines can be replicated simply through metadata operations, with actual data copied only for write operations.
- Snapshots help facilitate backup and remote-replication operations: needed in enterprises that require always-on data availability.

The Cisco HyperFlex HX Data Platform is administered through a VMware vSphere web client plug-in. Through this centralized point of control for the cluster, administrators can create volumes, monitor the data platform health, and manage resource use. Administrators can also use this data to predict when the cluster will need to be scaled. For customers who prefer a lightweight web interface, there is a tech preview URL management interface available by opening a browser to the IP address of the HX cluster interface. Additionally, there is an interface to assist in running cli commands through a web browser.

Figure 23 HyperFlex Web Client Plug-in

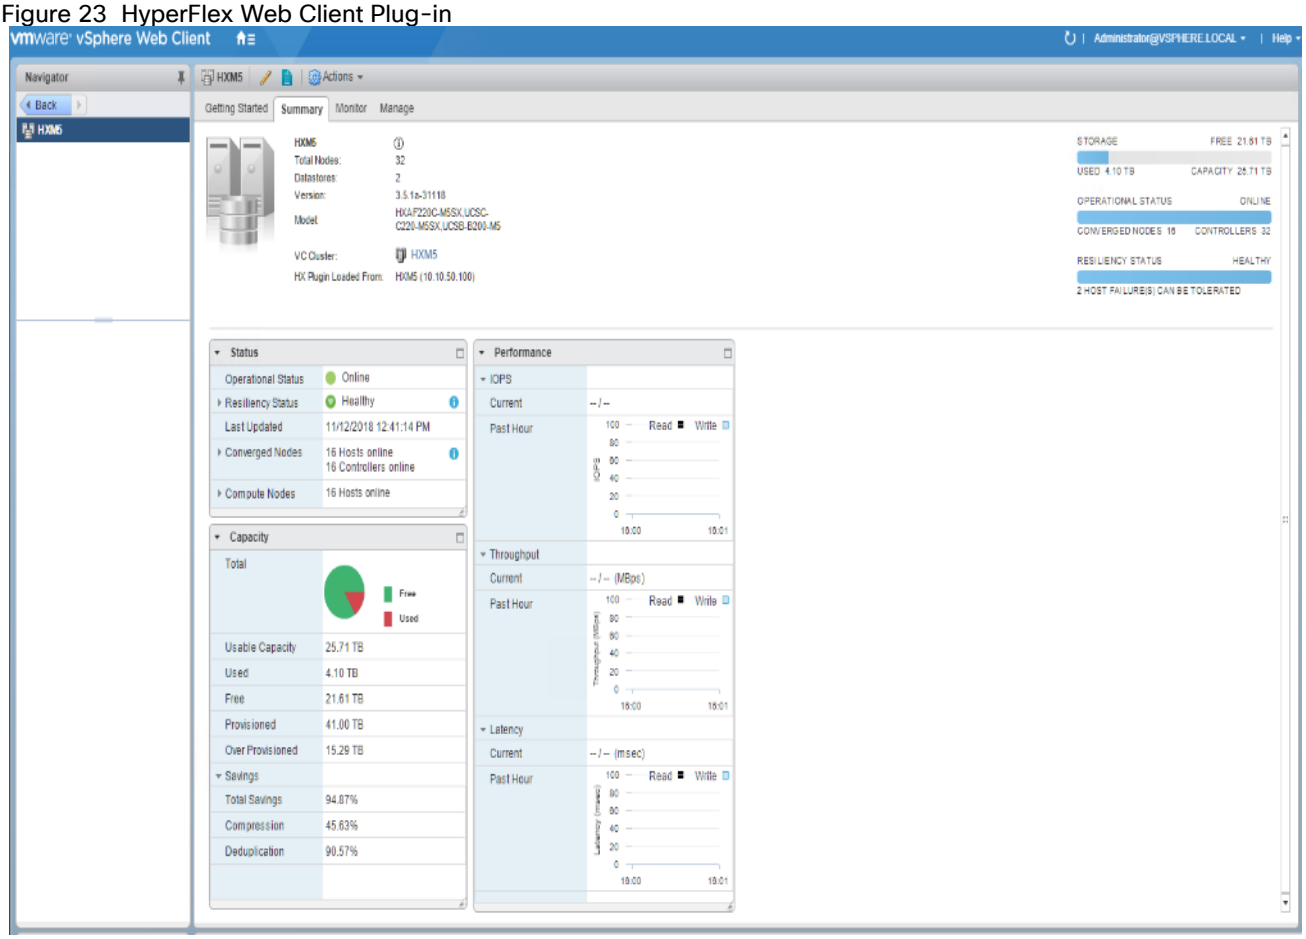

Cisco HyperFlex Connect HTML5 Management Web Page

An all-new HTML 5 based Web UI is available for use as the primary management tool for Cisco HyperFlex. Through this centralized point of control for the cluster, administrators can create volumes, monitor the data platform health, and manage resource use. Administrators can also use this data to predict when the cluster will need to be scaled. To use the HyperFlex Connect UI, connect using a web browser to the HyperFlex cluster IP address: http://<hx controller cluster ip>.

For the Tech Preview Web UI, connect to HX controller cluster IP: http://hx controller cluster ip/ui

Figure 24 HyperFlex Web GUI Preview

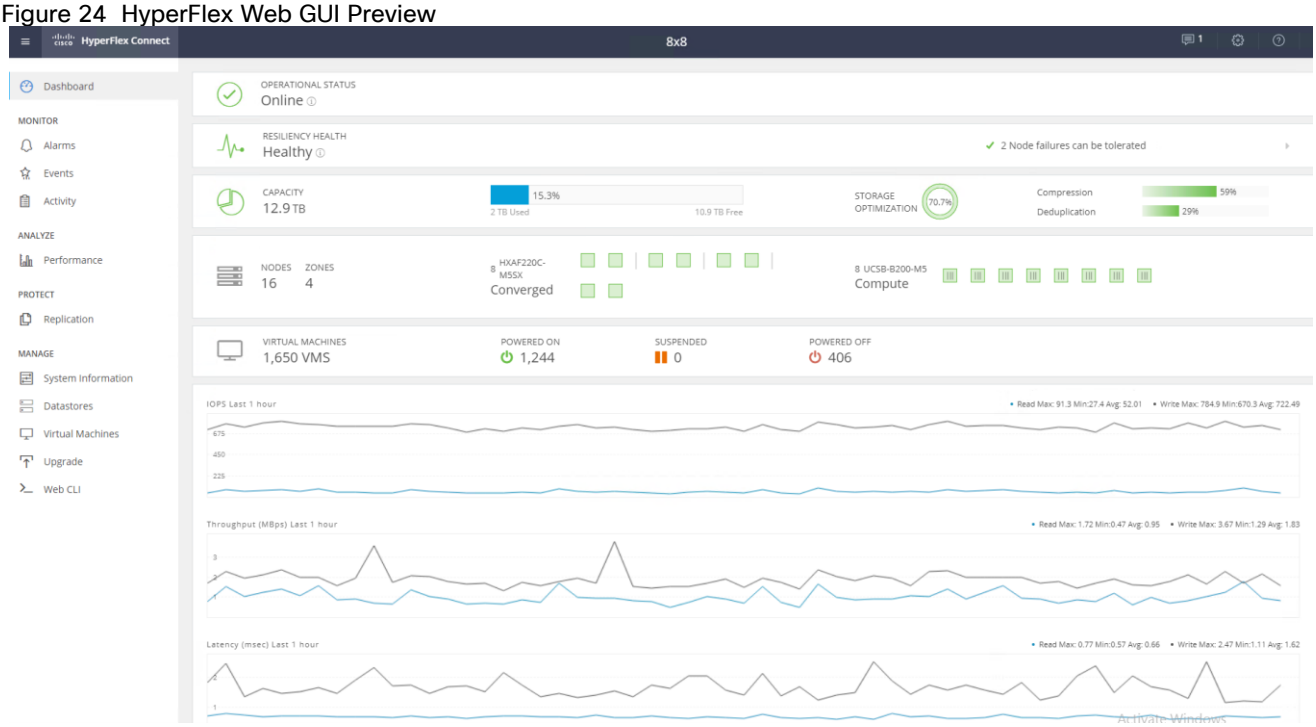

Figure 25 Web CLI

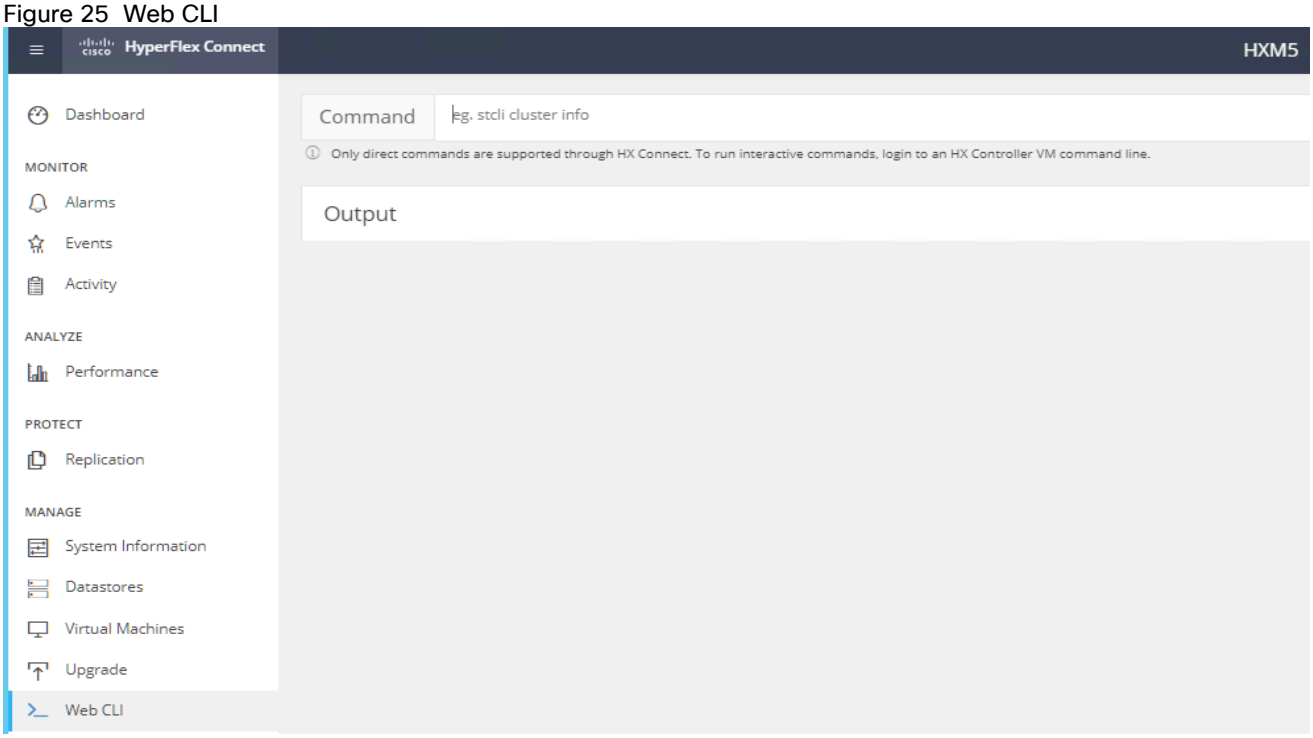

Cisco HyperFlex HX Data Platform Controller

A Cisco HyperFlex HX Data Platform controller resides on each node and implements the distributed file system. The controller runs in user space within a virtual machine and intercepts and handles all I/O from guest virtual machines. The platform controller VM uses the VMDirectPath I/O feature to provide PCI pass-through control of the physical server's SAS disk controller. This method gives the controller VM full control of the physical disk resources, utilizing the SSD drives as a read/write caching layer, and the HDDs as a capacity layer for distributed

storage. The controller integrates the data platform into VMware software through the use of two preinstalled VMware ESXi vSphere Installation Bundles (VIBs):

- IO Visor: This VIB provides a network file system (NFS) mount point so that the ESXi hypervisor can access the virtual disks that are attached to individual virtual machines. From the hypervisor's perspective, it is simply attached to a network file system.
- VMware API for Array Integration (VAAI): This storage offload API allows vSphere to request advanced file system operations such as snapshots and cloning. The controller implements these operations through manipulation of metadata rather than actual data copying, providing rapid response, and thus rapid deployment of new environments.

Replication Factor

The policy for the number of duplicate copies of each storage block is chosen during cluster setup and is referred to as the replication factor (RF).

- Replication Factor 3: For every I/O write committed to the storage layer, 2 additional copies of the blocks written will be created and stored in separate locations, for a total of 3 copies of the blocks. Blocks are distributed in such a way as to ensure multiple copies of the blocks are not stored on the same disks, nor on the same nodes of the cluster. This setting can tolerate simultaneous failures 2 entire nodes without losing data and resorting to restore from backup or other recovery processes.
- Replication Factor 2: For every I/O write committed to the storage layer, 1 additional copy of the blocks written will be created and stored in separate locations, for a total of 2 copies of the blocks. Blocks are distributed in such a way as to ensure multiple copies of the blocks are not stored on the same disks, nor on the same nodes of the cluster. This setting can tolerate a failure 1 entire node without losing data and resorting to restore from backup or other recovery processes.

Data Distribution

Incoming data is distributed across all nodes in the cluster to optimize performance using the caching tier (Figure 26). Effective data distribution is achieved by mapping incoming data to stripe units that are stored evenly across all nodes, with the number of data replicas determined by the policies you set. When an application writes data, the data is sent to the appropriate node based on the stripe unit, which includes the relevant block of information. This data distribution approach in combination with the capability to have multiple streams writing at the same time avoids both network and storage hot spots, delivers the same I/O performance regardless of virtual machine location, and gives you more flexibility in workload placement. This contrasts with other architectures that use a data locality approach that does not fully use available networking and I/O resources and is vulnerable to hot spots.

Figure 26 Data is Striped Across Nodes in the Cluster

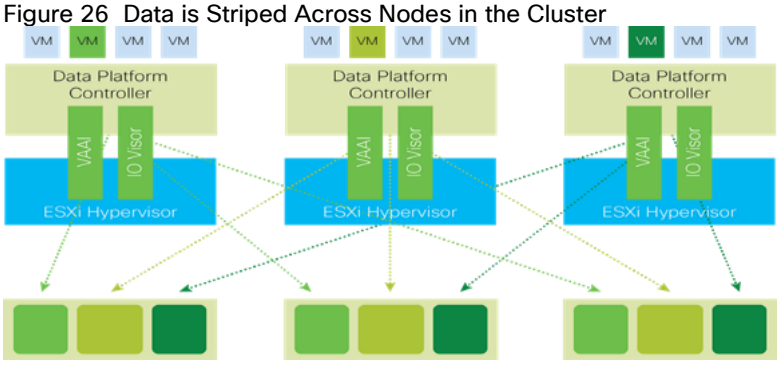

When moving a virtual machine to a new location using tools such as VMware Dynamic Resource Scheduling (DRS), the Cisco HyperFlex HX Data Platform does not require data to be moved. This approach significantly reduces the impact and cost of moving virtual machines among systems.

Data Operations

The data platform implements a distributed, log-structured file system that changes how it handles caching and storage capacity depending on the node configuration.

In the all-flash-memory configuration, the data platform uses a caching layer in SSDs to accelerate write responses, and it implements the capacity layer in SSDs. Read requests are fulfilled directly from data obtained from the SSDs in the capacity layer. A dedicated read cache is not required to accelerate read operations.

Incoming data is striped across the number of nodes required to satisfy availability requirements—usually two or three nodes. Based on policies you set, incoming write operations are acknowledged as persistent after they are replicated to the SSD drives in other nodes in the cluster. This approach reduces the likelihood of data loss due to SSD or node failures. The write operations are then de-staged to SSDs in the capacity layer in the all-flash memory configuration for long-term storage.

The log-structured file system writes sequentially to one of two write logs (three in case of RF=3) until it is full. It then switches to the other write log while de-staging data from the first to the capacity tier. When existing data is (logically) overwritten, the log-structured approach simply appends a new block and updates the metadata. This layout benefits SSD configurations in which seek operations are not time consuming. It reduces the write amplification levels of SSDs and the total number of writes the flash media experiences due to incoming writes and random overwrite operations of the data.

When data is de-staged to the capacity tier in each node, the data is deduplicated and compressed. This process occurs after the write operation is acknowledged, so no performance penalty is incurred for these operations. A small deduplication block size helps increase the deduplication rate. Compression further reduces the data footprint. Data is then moved to the capacity tier as write cache segments are released for reuse (Figure 27).

Figure 27 Data Write Operation Flow through the Cisco HyperFlex HX Data Platform

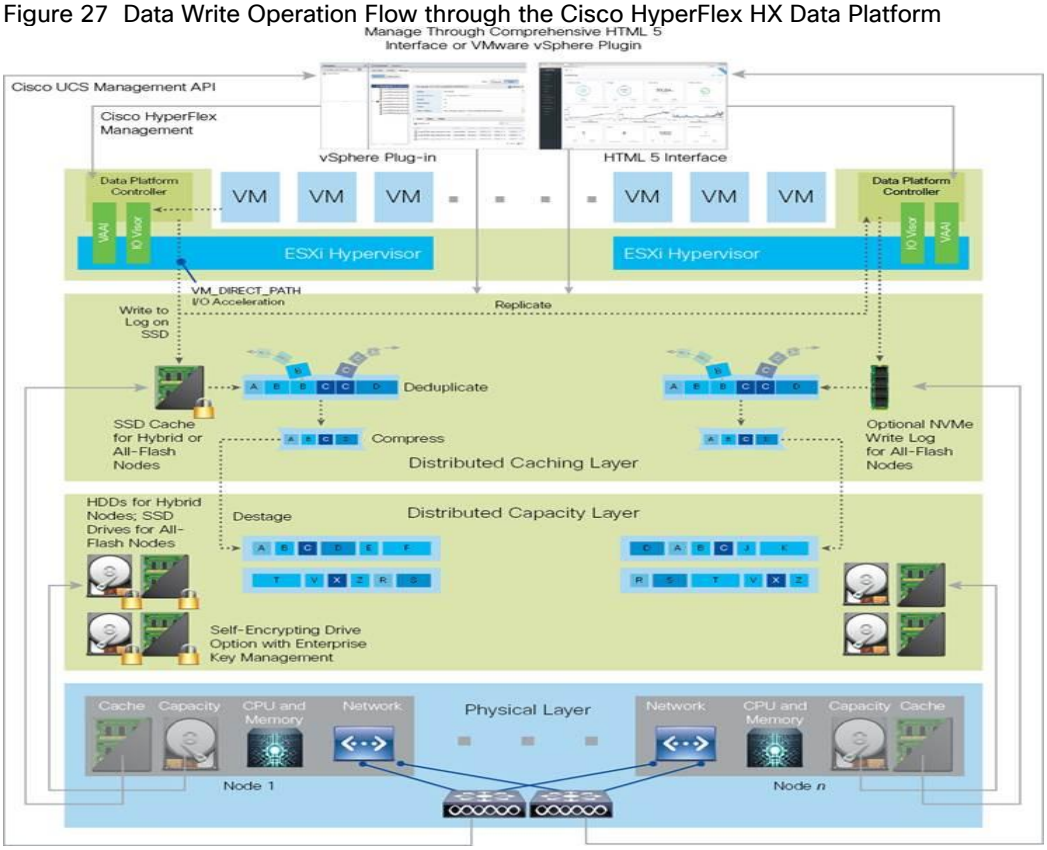

Hot data sets—data that is frequently or recently read from the capacity tier—are cached in memory. All-Flash configurations, however, do not use an SSD read cache since there is no performance benefit of such a cache; the persistent data copy already resides on high-performance SSDs. In these configurations, a read cache implemented with SSDs could become a bottleneck and prevent the system from using the aggregate bandwidth of the entire set of SSDs.

Data Optimization

The Cisco HyperFlex HX Data Platform provides finely detailed inline deduplication and variable block inline compression that is always on for objects in the cache (SSD and memory) and capacity (SSD or HDD) layers. Unlike other solutions, which require you to turn off these features to maintain performance, the deduplication and compression capabilities in the Cisco data platform are designed to sustain and enhance performance and significantly reduce physical storage capacity requirements.

Data Deduplication

Data deduplication is used on all storage in the cluster, including memory and SSD drives. Based on a patent-pending Top-K Majority algorithm, the platform uses conclusions from empirical research that show that most data, when sliced into small data blocks, has significant deduplication potential based on a minority of the data blocks. By fingerprinting and indexing just these frequently used blocks, high rates of deduplication can be achieved with only a small amount of memory, which is a high-value resource in cluster nodes (Figure 28).

Figure 28 Cisco HyperFlex HX Data Platform Optimizes Data Storage with No Performance Impact

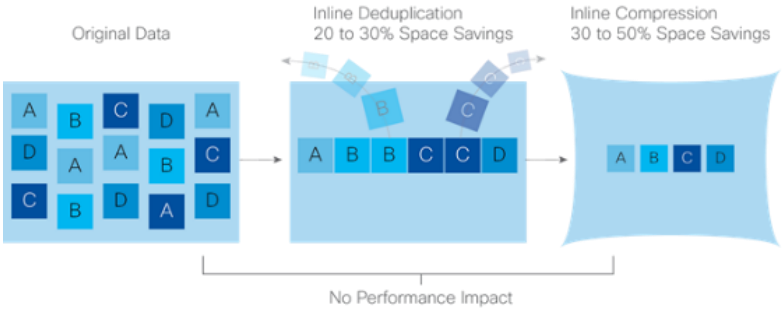

Inline Compression

The Cisco HyperFlex HX Data Platform uses high-performance inline compression on data sets to save storage capacity. Although other products offer compression capabilities, many negatively affect performance. In contrast, the Cisco data platform uses CPU-offload instructions to reduce the performance impact of compression operations. In addition, the log-structured distributed-objects layer has no effect on modifications (write operations) to previously compressed data. Instead, incoming modifications are compressed and written to a new location, and the existing (old) data is marked for deletion, unless the data needs to be retained in a snapshot.

The data that is being modified does not need to be read prior to the write operation. This feature avoids typical read-modify-write penalties and significantly improves write performance.

Log-Structured Distributed Objects

In the Cisco HyperFlex HX Data Platform, the log-structured distributed-object store layer groups and compresses data that filters through the deduplication engine into self-addressable objects. These objects are written to disk in a log-structured, sequential manner. All incoming I/O—including random I/O—is written sequentially to both the caching (SSD and memory) and persistent (SSD or HDD) tiers. The objects are distributed across all nodes in the cluster to make uniform use of storage capacity.

By using a sequential layout, the platform helps increase flash-memory endurance. Because read-modify-write operations are not used, there is little or no performance impact of compression, snapshot operations, and cloning on overall performance.

Data blocks are compressed into objects and sequentially laid out in fixed-size segments, which in turn are sequentially laid out in a log-structured manner (Figure 29). Each compressed object in the log-structured segment is uniquely addressable using a key, with each key fingerprinted and stored with a checksum to provide high levels of data integrity. In addition, the chronological writing of objects helps the platform quickly recover from media or node failures by rewriting only the data that came into the system after it was truncated due to a failure.

Figure 29 Cisco HyperFlex HX Data Platform Optimizes Data Storage with No Performance Impact

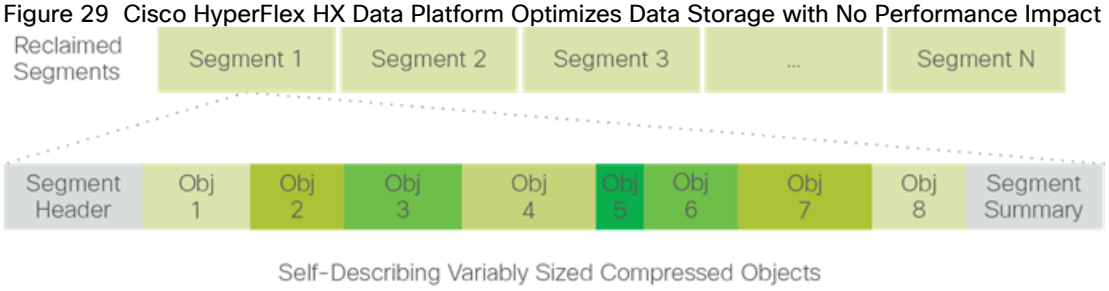

Encryption

Securely encrypted storage optionally encrypts both the caching and persistent layers of the data platform. Integrated with enterprise key management software, or with passphrase-protected keys, encrypting data at rest helps you comply with HIPAA, PCI-DSS, FISMA, and SOX regulations. The platform itself is hardened to Federal Information Processing Standard (FIPS) 140-1 and the encrypted drives with key management comply with the FIPS 140-2 standard.

Data Services

The Cisco HyperFlex HX Data Platform provides a scalable implementation of space-efficient data services, including thin provisioning, space reclamation, pointer-based snapshots, and clones—without affecting performance.

Thin Provisioning

The platform makes efficient use of storage by eliminating the need to forecast, purchase, and install disk capacity that may remain unused for a long time. Virtual data containers can present any amount of logical space to applications, whereas the amount of physical storage space that is needed is determined by the data that is written. You can expand storage on existing nodes and expand your cluster by adding more storage-intensive nodes as your business requirements dictate, eliminating the need to purchase large amounts of storage before you need it.

Snapshots

The Cisco HyperFlex HX Data Platform uses metadata-based, zero-copy snapshots to facilitate backup operations and remote replication: critical capabilities in enterprises that require always-on data availability. Space-efficient snapshots allow you to perform frequent online data backups without worrying about the consumption of physical storage capacity. Data can be moved offline or restored from these snapshots instantaneously.

- Fast snapshot updates: When modified-data is contained in a snapshot, it is written to a new location, and the metadata is updated, without the need for read-modify-write operations.
- Rapid snapshot deletions: You can quickly delete snapshots. The platform simply deletes a small amount of metadata that is located on an SSD, rather than performing a long consolidation process as needed by solutions that use a delta-disk technique.
- Highly specific snapshots: With the Cisco HyperFlex HX Data Platform, you can take snapshots on an individual file basis. In virtual environments, these files map to drives in a virtual machine. This flexible specificity allows you to apply different snapshot policies on different virtual machines.

Many basic backup applications, read the entire dataset, or the changed blocks since the last backup at a rate that is usually as fast as the storage, or the operating system can handle. This can cause performance implications since HyperFlex is built on Cisco UCS with 40GbE that could result in multiple gigabytes per second of backup throughput. These basic backup applications, such as Windows Server Backup, should be scheduled during off-peak hours, particularly the initial backup if the application lacks some form of change block tracking.

Full featured backup applications, such as Veeam Backup and Replication v9.5, have the ability to limit the amount of throughput the backup application can consume which can protect latency sensitive applications during the production hours. With the release of v9.5 update 2, Veeam is the first partner to integrate HX native

snapshots into the product. HX Native snapshots do not suffer the performance penalty of delta-disk snapshots, and do not require heavy disk IO impacting consolidation during snapshot deletion.

Particularly important for SQL administrators is the Veeam Explorer for SQL which can provide transaction level recovery within the Microsoft VSS framework. The three ways Veeam Explorer for SQL Server works to restore SQL Server databases include; from the backup restore point, from a log replay to a point in time, and from a log replay to a specific transaction – all without taking the VM or SQL Server offline.

Fast, Space-Efficient Clones

In the Cisco HyperFlex HX Data Platform, clones are writable snapshots that can be used to rapidly provision items such as virtual desktops and applications for test and development environments. These fast, space-efficient clones rapidly replicate storage volumes so that virtual machines can be replicated through just metadata operations, with actual data copying performed only for write operations. With this approach, hundreds of clones can be created and deleted in minutes. Compared to full-copy methods, this approach can save a significant amount of time, increase IT agility, and improve IT productivity.

Clones are deduplicated when they are created. When clones start diverging from one another, data that is common between them is shared, with only unique data occupying new storage space. The deduplication engine eliminates data duplicates in the diverged clones to further reduce the clone's storage footprint.

Data Replication and Availability

In the Cisco HyperFlex HX Data Platform, the log-structured distributed-object layer replicates incoming data, improving data availability. Based on policies that you set, data that is written to the write cache is synchronously replicated to one or two other SSD drives located in different nodes before the write operation is acknowledged to the application. This approach allows incoming writes to be acknowledged quickly while protecting data from SSD or node failures. If an SSD or node fails, the replica is quickly re-created on other SSD drives or nodes using the available copies of the data.

The log-structured distributed-object layer also replicates data that is moved from the write cache to the capacity layer. This replicated data is likewise protected from SSD or node failures. With two replicas, or a total of three data copies, the cluster can survive uncorrelated failures of two SSD drives or two nodes without the risk of data loss. Uncorrelated failures are failures that occur on different physical nodes. Failures that occur on the same node affect the same copy of data and are treated as a single failure. For example, if one disk in a node fails and subsequently another disk on the same node fails, these correlated failures count as one failure in the system. In this case, the cluster could withstand another uncorrelated failure on a different node. See the Cisco HyperFlex HX Data Platform system administrator's guide for a complete list of fault-tolerant configurations and settings.

If a problem occurs in the Cisco HyperFlex HX controller software, data requests from the applications residing in that node are automatically routed to other controllers in the cluster. This same capability can be used to upgrade or perform maintenance on the controller software on a rolling basis without affecting the availability of the cluster or data. This self-healing capability is one of the reasons that the Cisco HyperFlex HX Data Platform is well suited for production applications.

In addition, native replication transfers consistent cluster data to local or remote clusters. With native replication, you can snapshot and store point-in-time copies of your environment in local or remote environments for backup and disaster recovery purposes.

Data Rebalancing

A distributed file system requires a robust data rebalancing capability. In the Cisco HyperFlex HX Data Platform, no overhead is associated with metadata access, and rebalancing is extremely efficient. Rebalancing is a non-disruptive online process that occurs in both the caching and persistent layers, and data is moved at a fine level of specificity to improve the use of storage capacity. The platform automatically rebalances existing data when nodes and drives are added or removed or when they fail. When a new node is added to the cluster, its capacity and performance is made available to new and existing data. The rebalancing engine distributes existing data to the new node and helps ensure that all nodes in the cluster are used uniformly from capacity and performance perspectives. If a node fails or is removed from the cluster, the rebalancing engine rebuilds and distributes copies of the data from the failed or removed node to available nodes in the clusters.

Online Upgrades

Cisco HyperFlex HX-Series systems and the HX Data Platform support online upgrades so that you can expand and update your environment without business disruption. You can easily expand your physical resources; add processing capacity; and download and install BIOS, driver, hypervisor, firmware, and Cisco UCS Manager updates, enhancements, and bug fixes.

Cisco Nexus 93180 Switches

The Cisco Nexus 93180YC-FX Switches has 48 10/25-Gbps Small Form Pluggable Plus (SFP+) ports and 6 Quad 40/100-Gbps SFP+ (QSFP+) uplink ports. All the ports are line rate, delivering 3.6 Tbps of throughput in a 1-rack-unit (1RU) form factor. Cisco Nexus 93180-YC-FX benefits are listed below:

Specifications at-a-Glance

- 1 rack unit (1RU)
- 48 x 1/10/25-Gbps fiber ports
- 6 x 40/100-Gbps QSFP28 ports
- Up to 3.6 Tbps of bandwidth

Architectural Flexibility

- Leaf-node support for Cisco ACI architecture with flexible port configuration
- Seamless convergence thanks to 48 downlink ports that can work as 1/10/25-Gbps Ethernet or FCoE ports or as 8/16/32-Gbps Fibre Channel ports
- Easy migration with 6 uplink ports that can be configured as 40/100-Gbps Ethernet or FCoE ports

Feature Rich

- Automated policy-based systems management with Cisco ACI
- Open APIs enable third-party integration with our partners
- Better management of speed mismatch between access and uplink ports with 40 MB of shared buffer space
- Support for Fibre Channel interfaces for back-end storage connectivity

Highly Available and Efficient Design

- High-performance, non-blocking architecture
- Easily deployed into either a hot-aisle or a cold-aisle configuration
- Redundant, hot-swappable power supplies and fan trays

Simplified Operations

- Automate IT work flows and shorten app deployment from weeks to minutes

Top-notch Security

- Whitelist model, policy enforcement and application security with Cisco ACI micro-segmentation
- Wire-rate MACsec encryption on all ports

Real-time Visibility and Telemetry

- Built-in Cisco Tetration sensors for rich traffic-flow telemetry and line-rate data collection
- Get actionable insights in less than 1 second
- Get visibility into everything in your data center

Investment Protection

- Flexible migration options with support for 10-Gbps and 25-Gbps access connectivity and 40-Gbps and 100-Gbps uplinks
- Cisco's 40-Gbps bidirectional transceiver allows for reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet

Resources

- Cisco Nexus 9300-EX and 9300-FX Platform Leaf
- Switches for Cisco Application Centric Infrastructure Data Sheet

Figure 30 Cisco Nexus 93180YC-FX Switch

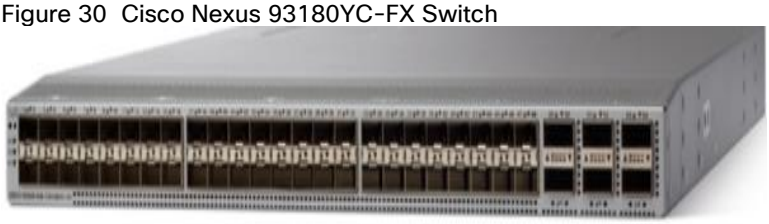

VMware vSphere 6.5

VMware provides virtualization software. VMware's enterprise software hypervisors for servers–VMware vSphere ESX, vSphere ESXi, and VSphere–are bare-metal hypervisors that run directly on server hardware without requiring an additional underlying operating system. VMware vCenter Server for vSphere provides central management and complete control and visibility into clusters, hosts, virtual machines, storage, networking, and other critical elements of your virtual infrastructure.

VMware vSphere 6.5 introduces many enhancements to vSphere Hypervisor, VMware virtual machines, vCenter Server, virtual storage, and virtual networking, further extending the core capabilities of the vSphere platform.

VMware vCenter Server

- Migration Tool
- Improved appliance management
- Native high availability
- Native backup and restore
- There are also general improvements to vCenter Server 6.5, including the vSphere Web Client and the fully supported HTML5-based vSphere Client.

VMware ESXi 6.5 Hypervisor

- With vSphere 6.5, administrators can find significant improvement in patching, upgrading and managing configuration of ESXi hosts through vSphere Update Manager that is enabled by default.
- VMware tool and virtual hardware upgrade
- Improvement in Host Profile, as well as in day to day operations
- Improvement in manageability and configuration rules for Auto-Deploy
- Enhanced monitoring, added option to monitor GPU usage.
- Dedicated Gateways for VMkernel Network Adapter
- VMware vSphere Storage I/O Control Using Storage Policy Based Management

Citrix Virtual Apps and Desktops 1808

Enterprise IT organizations are tasked with the challenge of provisioning Microsoft Windows apps and desktops while managing cost, centralizing control, and enforcing corporate security policy. Deploying Windows apps to users in any location, regardless of the device type and available network bandwidth, enables a mobile workforce that can improve productivity. With Citrix Virtual Apps and Desktops 1808, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

The Citrix Virtual Apps and Desktops 1808 release offers these benefits:

- Comprehensive virtual desktop delivery for any use case. The Citrix Virtual Apps and Desktops 1808 release incorporates the full power of Virtual Apps, delivering full desktops or just applications to users. Administrators can deploy both Virtual Apps published applications and desktops (to maximize IT control at low cost) or personalized VDI desktops (with simplified image management) from the same management console. Citrix Virtual Apps and Desktops 1808 leverages common policies and cohesive tools to govern both infrastructure resources and user access.
- Simplified support and choice of BYO (Bring Your Own) devices. Citrix Virtual Apps and Desktops 1808 brings thousands of corporate Microsoft Windows-based applications to mobile devices with a native-touch experience and optimized performance. HDX technologies create a “high definition” user experience, even for graphics intensive design and engineering applications.
- Lower cost and complexity of application and desktop management. Citrix Virtual Apps and Desktops 1808 helps IT organizations take advantage of agile and cost-effective cloud offerings, allowing the virtualized infrastructure to flex and meet seasonal demands or the need for sudden capacity changes. IT organizations can deploy Virtual Desktops application and desktop workloads to private or public clouds.

- Protection of sensitive information through centralization. Virtual Desktops decreases the risk of corporate data loss, enabling access while securing intellectual property and centralizing applications since assets reside in the datacenter.
- Virtual Delivery Agent improvements. Universal print server and driver enhancements and support for the HDX 3D Pro graphics acceleration for Windows 10 are key additions in Citrix Virtual Apps and Desktops 1808
- Improved high-definition user experience. Citrix Virtual Apps and Desktops 1808 continues the evolutionary display protocol leadership with enhanced Thinwire display remoting protocol and Framehawk support for HDX 3D Pro.

Citrix Virtual Apps and Virtual Desktops are application and desktop virtualization solutions built on a unified architecture so they're simple to manage and flexible enough to meet the needs of all your organization's users. Virtual Apps and Virtual Desktops have a common set of management tools that simplify and automate IT tasks. You use the same architecture and management tools to manage public, private, and hybrid cloud deployments as you do for on premises deployments.

Citrix Virtual Apps delivers:

- Virtual Apps published apps, also known as server-based hosted applications: These are applications hosted from Microsoft Windows servers to any type of device, including Windows PCs, Macs, smartphones, and tablets. Some Virtual Apps editions include technologies that further optimize the experience of using Windows applications on a mobile device by automatically translating native mobile-device display, navigation, and controls to Windows applications; enhancing performance over mobile networks; and enabling developers to optimize any custom Windows application for any mobile environment.
- Virtual Apps published desktops, also known as server-hosted desktops: These are inexpensive, locked-down Windows virtual desktops hosted from Windows server operating systems. They are well suited for users, such as call center employees, who perform a standard set of tasks.
- Virtual machine–hosted apps: These are applications hosted from machines running Windows desktop operating systems for applications that can’t be hosted in a server environment.
- Windows applications delivered with Microsoft App-V: These applications use the same management tools that you use for the rest of your Virtual Apps deployment.
- Citrix Virtual Desktops: Includes significant enhancements to help customers deliver Windows apps and desktops as mobile services while addressing management complexity and associated costs. Enhancements in this release include:
- Unified product architecture for Virtual Apps and Virtual Desktops: The FlexCast Management Architecture (FMA). This release supplies a single set of administrative interfaces to deliver both hosted-shared applications (RDS) and complete virtual desktops (VDI). Unlike earlier releases that separately provisioned Citrix Virtual Apps and Virtual Desktops farms, the Citrix Virtual Apps and Desktops 1808 release allows administrators to deploy a single infrastructure and use a consistent set of tools to manage mixed application and desktop workloads.
- Support for extending deployments to the cloud. This release provides the ability for hybrid cloud provisioning from Microsoft Azure, Amazon Web Services (AWS) or any Cloud Platform-powered public or private cloud. Cloud deployments are configured, managed, and monitored through the same administrative consoles as deployments on traditional on-premises infrastructure.

Citrix Virtual Desktops delivers:

- VDI desktops: These virtual desktops each run a Microsoft Windows desktop operating system rather than running in a shared, server-based environment. They can provide users with their own desktops that they can fully personalize.
- Hosted physical desktops: This solution is well suited for providing secure access powerful physical machines, such as blade servers, from within your data center.
- Remote PC access: This solution allows users to log in to their physical Windows PC from anywhere over a secure Virtual Desktops connection.
- Server VDI: This solution is designed to provide hosted desktops in multitenant, cloud environments.
- Capabilities that allow users to continue to use their virtual desktops: These capabilities let users continue to work while not connected to your network.

This product release includes the following new and enhanced features:

 **Some Virtual Desktops editions include the features available in Virtual Apps.**

Zones

Deployments that span widely-dispersed locations connected by a WAN can face challenges due to network latency and reliability. Configuring zones can help users in remote regions connect to local resources without forcing connections to traverse large segments of the WAN. Using zones allows effective Site management from a single Citrix Studio console, Citrix Director, and the Site database. This saves the costs of deploying, staffing, licensing, and maintaining additional Sites containing separate databases in remote locations.

Zones can be helpful in deployments of all sizes. You can use zones to keep applications and desktops closer to end users, which improves performance.

For more information, see the Zones article.

Improved Database Flow and Configuration

When you configure the databases during Site creation, you can now specify separate locations for the Site, Logging, and Monitoring databases. Later, you can specify different locations for all three databases. In previous releases, all three databases were created at the same address, and you could not specify a different address for the Site database later.

You can now add more Delivery Controllers when you create a Site, as well as later. In previous releases, you could add more Controllers only after you created the Site.

For more information, see the Databases and Controllers articles.

Application Limits

Configure application limits to help manage application use. For example, you can use application limits to manage the number of users accessing an application simultaneously. Similarly, application limits can be used to manage the number of simultaneous instances of resource-intensive applications, this can help maintain server performance and prevent deterioration in service.

For more information, see the Manage applications article.

Multiple Notifications Before Machine Updates or Scheduled Restarts

You can now choose to repeat a notification message that is sent to affected machines before the following types of actions begin:

- Updating machines in a Machine Catalog using a new master image
- Restarting machines in a Delivery Group according to a configured schedule

If you indicate that the first message should be sent to each affected machine 15 minutes before the update or restart begins, you can also specify that the message be repeated every five minutes until the update/restart begins.

For more information, see the Manage Machine Catalogs and Manage machines in Delivery Groups articles.

API Support for Managing Session Roaming

By default, sessions roam between client devices with the user. When the user launches a session and then moves to another device, the same session is used and applications are available on both devices. The applications follow, regardless of the device or whether current sessions exist. Similarly, printers and other resources assigned to the application follow.

You can now use the PowerShell SDK to tailor session roaming. This was an experimental feature in the previous release.

For more information, see the Sessions article.

API Support for Provisioning VMs from Hypervisor Templates

When using the PowerShell SDK to create or update a Machine Catalog, you can now select a template from other hypervisor connections. This is in addition to the currently-available choices of virtual machine images and snapshots.

Support for New and Additional Platforms

See the System requirements article for full support information. Information about support for third-party product versions is updated periodically.

By default, SQL Server 2012 Express SP2 is installed when you install the Delivery Controller. SP1 is no longer installed.

The component installers now automatically deploy newer Microsoft Visual C++ runtime versions: 32-bit and 64-bit Microsoft Visual C++ 2013, 2010 SP1, and 2008 SP1. Visual C++ 2005 is no longer deployed.

You can install Studio or VDAs for Windows Desktop OS on machines running Windows 10.

You can create connections to Microsoft Azure virtualization resources.

Figure 31 Logical Architecture of Citrix Virtual Desktops

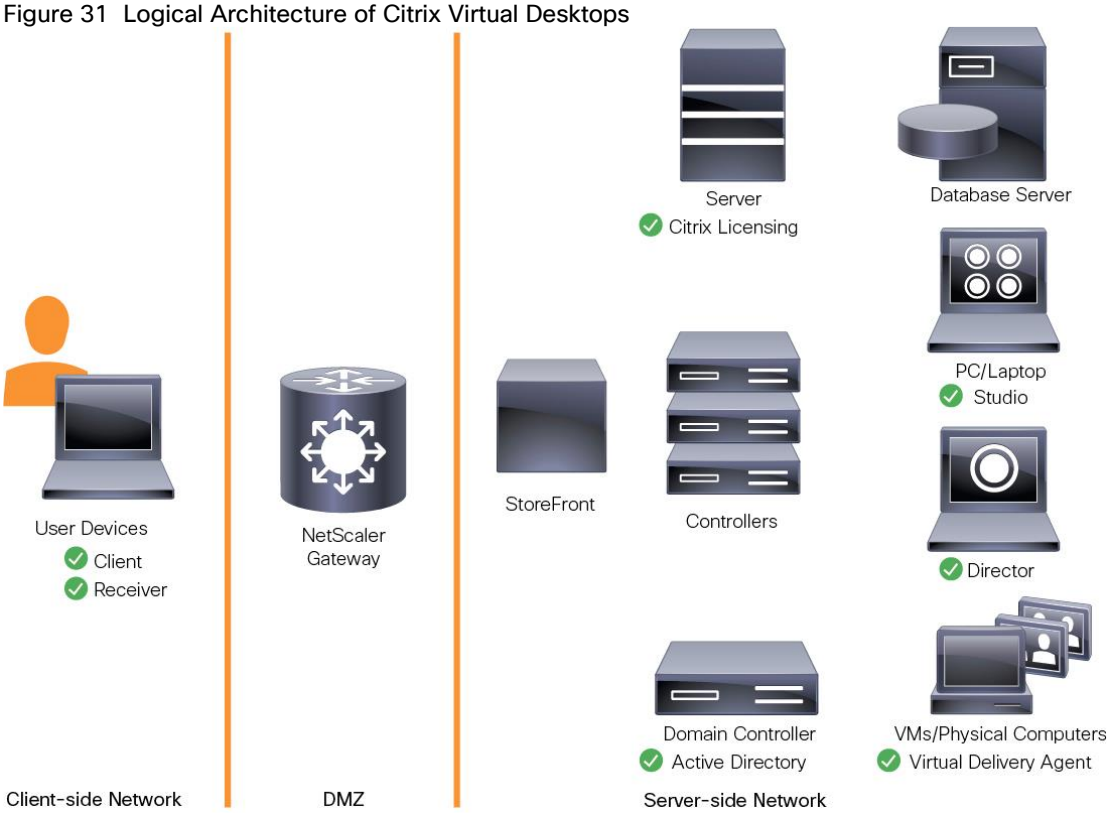

Citrix Provisioning Services 1808

Most enterprises struggle to keep up with the proliferation and management of computers in their environments. Each computer, whether it is a desktop PC, a server in a data center, or a kiosk-type device, must be managed as an individual entity. The benefits of distributed processing come at the cost of distributed management. It costs time and money to set up, update, support, and ultimately decommission each computer. The initial cost of the machine is often dwarfed by operating costs.

Citrix PVS takes a very different approach from traditional imaging solutions by fundamentally changing the relationship between hardware and the software that runs on it. By streaming a single shared disk image (vDisk) rather than copying images to individual machines, PVS enables organizations to reduce the number of disk images that they manage, even as the number of machines continues to grow, simultaneously providing the efficiency of centralized management and the benefits of distributed processing.

In addition, because machines are streaming disk data dynamically and in real time from a single shared image, machine image consistency is essentially ensured. At the same time, the configuration, applications, and even the OS of large pools of machines can be completed changed in the time it takes the machines to reboot.

Using PVS, any vDisk can be configured in standard-image mode. A vDisk in standard-image mode allows many computers to boot from it simultaneously, greatly reducing the number of images that must be maintained and the amount of storage that is required. The vDisk is in read-only format, and the image cannot be changed by target devices.

Benefits for Citrix Virtual Apps and Other Server Farm Administrators

If you manage a pool of servers that work as a farm, such as Citrix Virtual Apps servers or web servers, maintaining a uniform patch level on your servers can be difficult and time consuming. With traditional imaging

solutions, you start with a clean golden master image, but as soon as a server is built with the master image, you must patch that individual server along with all the other individual servers. Rolling out patches to individual servers in your farm is not only inefficient, but the results can also be unreliable. Patches often fail on an individual server, and you may not realize you have a problem until users start complaining or the server has an outage. After that happens, getting the server resynchronized with the rest of the farm can be challenging, and sometimes a full reimaging of the machine is required.

With Citrix PVS, patch management for server farms is simple and reliable. You start by managing your golden image, and you continue to manage that single golden image. All patching is performed in one place and then streamed to your servers when they boot. Server build consistency is assured because all your servers use a single shared copy of the disk image. If a server becomes corrupted, simply reboot it, and it is instantly back to the known good state of your master image. Upgrades are extremely fast to implement. After you have your updated image ready for production, you simply assign the new image version to the servers and reboot them. You can deploy the new image to any number of servers in the time it takes them to reboot. Just as important, rollback can be performed in the same way, so problems with new images do not need to take your servers or your users out of commission for an extended period of time.

Benefits for Desktop Administrators

Because Citrix PVS is part of Citrix Virtual Desktops, desktop administrators can use PVS's streaming technology to simplify, consolidate, and reduce the costs of both physical and virtual desktop delivery. Many organizations are beginning to explore desktop virtualization. Although virtualization addresses many of IT's needs for consolidation and simplified management, deploying it also requires deployment of supporting infrastructure. Without PVS, storage costs can make desktop virtualization too costly for the IT budget. However, with PVS, IT can reduce the amount of storage required for VDI by as much as 90 percent. And with a single image to manage instead of hundreds or thousands of desktops, PVS significantly reduces the cost, effort, and complexity for desktop administration.

Different types of workers across the enterprise need different types of desktops. Some require simplicity and standardization, and others require high performance and personalization. Virtual Desktops can meet these requirements in a single solution using Citrix FlexCast delivery technology. With FlexCast, IT can deliver every type of virtual desktop, each specifically tailored to meet the performance, security, and flexibility requirements of each individual user.

Not all desktops applications can be supported by virtual desktops. For these scenarios, IT can still reap the benefits of consolidation and single-image management. Desktop images are stored and managed centrally in the data center and streamed to physical desktops on demand. This model works particularly well for standardized desktops such as those in lab and training environments and call centers and thin-client devices used to access virtual desktops.

Citrix Provisioning Services Solution

Citrix PVS streaming technology allows computers to be provisioned and re-provisioned in real time from a single shared disk image. With this approach, administrators can completely eliminate the need to manage and patch individual systems. Instead, all image management is performed on the master image. The local hard drive of each system can be used for runtime data caching or, in some scenarios, removed from the system entirely, which reduces power use, system failure rate, and security risk.

The PVS solution's infrastructure is based on software-streaming technology. After PVS components are installed and configured, a 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 used for this process is referred to as a master target device. The devices that use the vDisks are called target devices. vDisks can exist on a PVS, file share, or in larger deployments, on a storage system with which PVS can communicate (iSCSI, SAN, network-

attached storage [NAS], and Common Internet File System [CIFS]). vDisks can be assigned to a single target device in private-image mode, or to multiple target devices in standard-image mode.

Citrix Provisioning Services Infrastructure

The Citrix PVS infrastructure design directly relates to administrative roles within a PVS farm. The PVS administrator role determines which components that administrator can manage or view in the console.

A PVS farm contains several components. Figure 32 provides a high-level view of a basic PVS infrastructure and shows how PVS components might appear within that implementation.

Figure 32 Logical Architecture of Citrix Provisioning Services

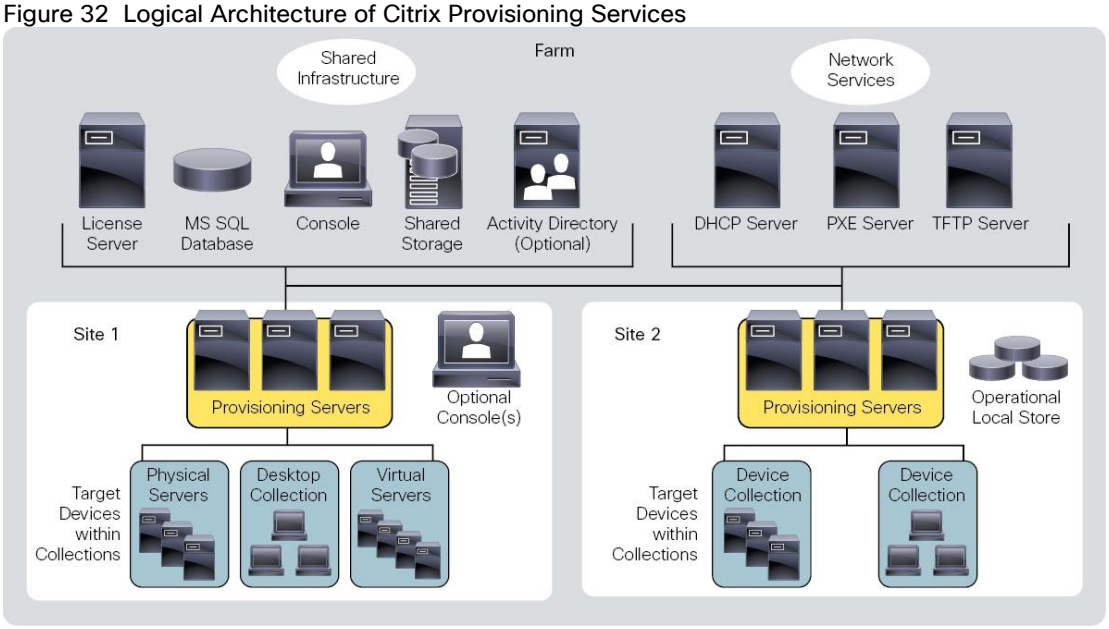

The following new features are available with Provisioning Services:

- Linux streaming
- XenServer proxy using PVS-Accelerator

Architecture and Design of Citrix Virtual Desktops on Cisco Unified Computing System and Cisco HyperFlex Storage Design Fundamentals

There are many reasons to consider a virtual desktop solution such as an ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Computer (BYOC) 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 universities and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

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

- Traditional PC: A traditional PC is what typically constituted a desktop environment; physical device with a locally installed operating system.
- Hosted Shared Desktop: A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2012, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Changes made by one user could impact the other users.
- Hosted Virtual Desktop: A hosted virtual desktop is a virtual desktop running either on virtualization layer (ESX) or on bare metal hardware. 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 Microsoft Session Hosts and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.
- Streamed Applications: Streamed desktops and applications run entirely on the user‘s local client device and are sent from a server on demand. The user interacts with the application or desktop directly but the resources may only available while they are connected to the network.
- Local Virtual Desktop: A local virtual desktop is a desktop running entirely on the user‘s local device and continues to operate when disconnected from the network. In this case, the user’s local device is used as a type 1 hypervisor and is synced with the data center when the device is connected to the network.

For the purposes of the validation represented in this document, both Virtual Desktops Virtual Desktops and Virtual Apps Hosted Shared Desktop server sessions were validated. Each of the sections provides some fundamental design decisions for this environment.

Understanding Applications and Data

When the desktop user groups and sub-groups have been identified, the next task is to catalog group application and data requirements. This can be one of the most time-consuming processes in the VDI planning exercise, but

is essential for the VDI project's success. If the applications and data are not identified and co-located, performance will be negatively affected.

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

Project Planning and Solution Sizing Sample Questions

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

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

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

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

- What is the desktop OS planned? Windows 7, Windows 8, or Windows 10?
- 32-bit or 64-bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 7/8/10?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, or graphics-intensive workloads?
- What is the end point graphics processing capability?
- Will Citrix Virtual Apps for Remote Desktop Server Hosted Sessions used?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- 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?

Citrix Virtual 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 1808 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 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, VM provisioning relies on Citrix Provisioning 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 Windows Server OS (for RDS hosted shared desktops) or a Windows Desktop OS (for hosted pooled VDI desktops).

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

Figure 33 Access Desktops and Applications through Machine Catalogs and Delivery Groups

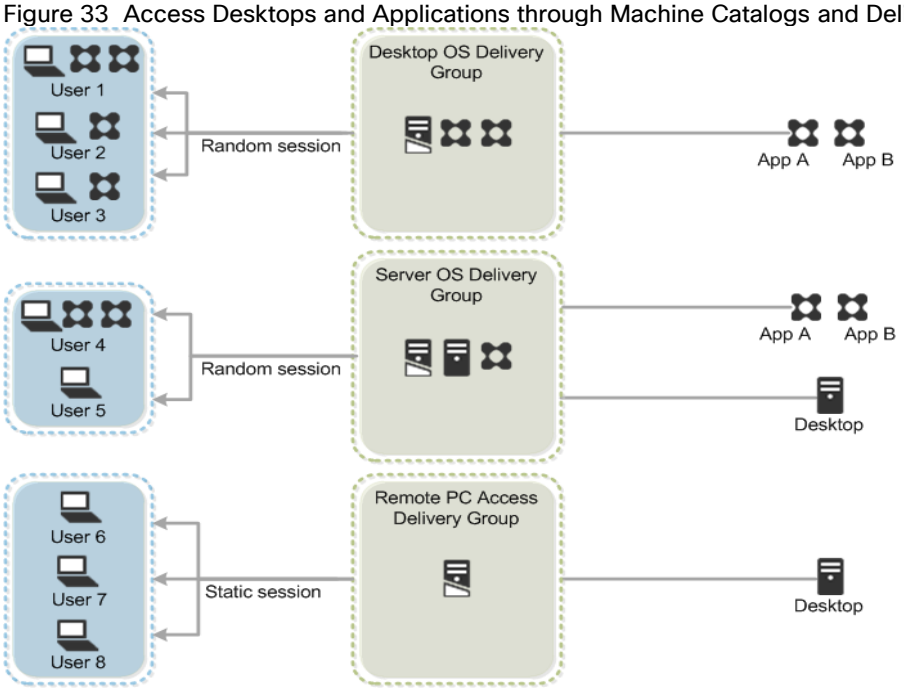

Example Virtual Desktops Deployments

Two examples of typical Virtual Desktops deployments are:

- A distributed components configuration
- A multiple site configuration

Since Virtual Apps and Citrix Virtual Apps and Desktops 1808 are based on a unified architecture, combined they can deliver a combination of Hosted Shared Desktops (HSDs, using a Server OS machine) and Hosted Virtual Desktops (HVDs, using a Desktop OS).

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 34 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 Desktops in a configuration that resembles this distributed components configuration shown. Two Cisco C220 rack servers host the required infrastructure services (AD, DNS, DHCP, Profile, SQL, Citrix Virtual Desktops management, and StoreFront servers).

Figure 34 Example of a Distributed Components Configuration

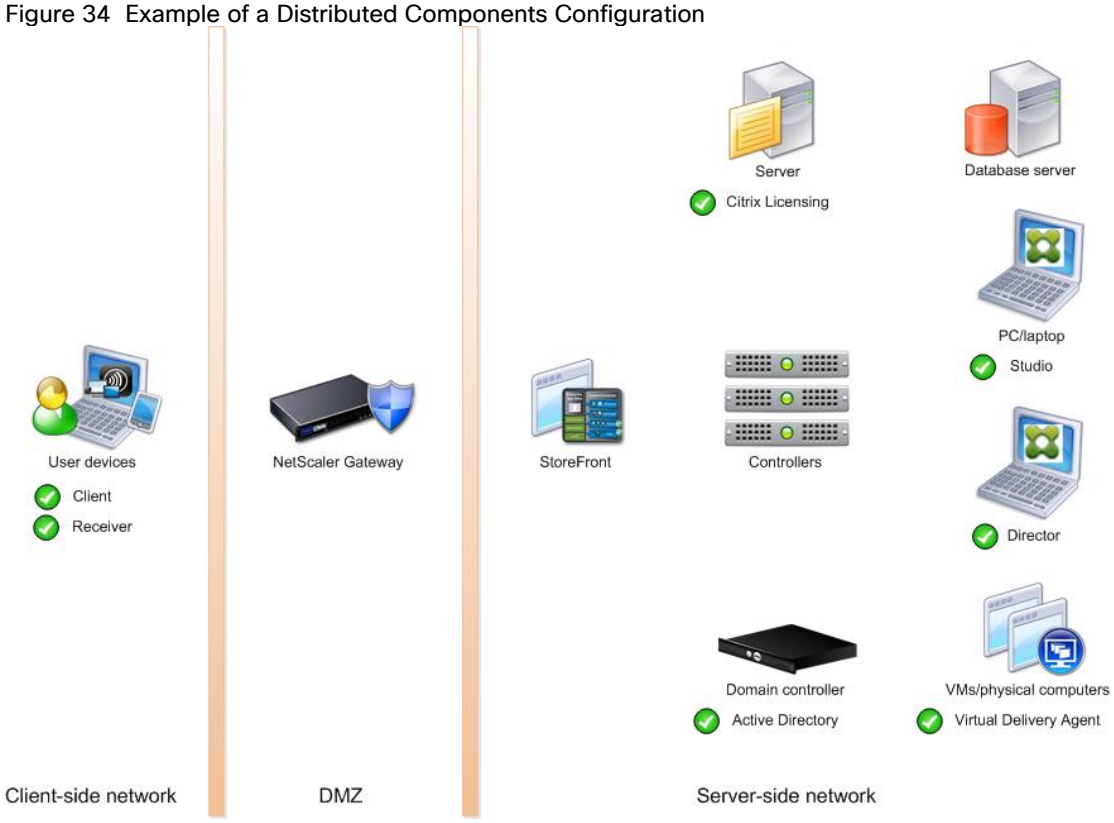

Multiple Site Configuration

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.

In Figure 35 depicting 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 35 Multiple Sites

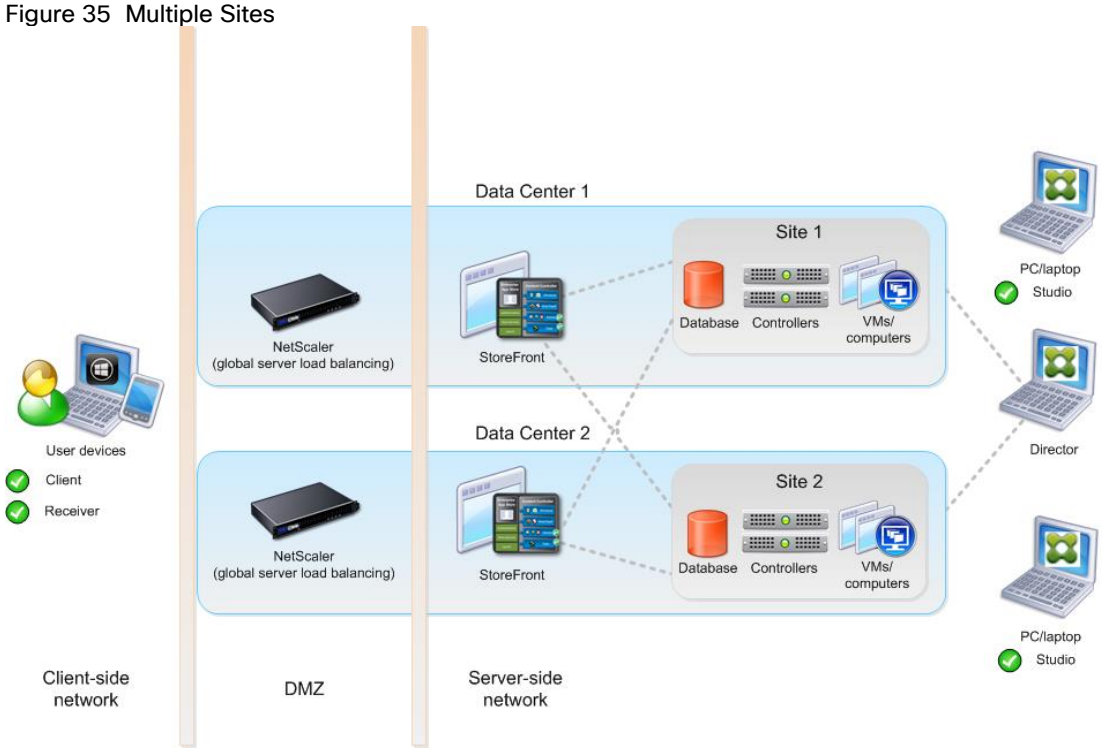

You can use StoreFront to aggregate resources from multiple sites to provide users with a single point of access with NetScaler. A separate Studio console is required to manage each site; sites cannot be managed as a single entity. You can use Director to support users across sites.

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.

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 Desktops Environment for a Mixed Workload

With Citrix Virtual Apps and Desktops 1808, 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.

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

Deployment Hardware and Software

Products Deployed

The architecture deployed is modular. While each customer's environment might vary in its exact configuration, the reference architecture contained in this document once built, can easily be scaled as requirements and demands change. This includes scaling both up (adding additional resources within existing Cisco HyperFlex system) and out (adding additional Cisco UCS HX-series nodes).

The solution includes Cisco networking, Cisco UCS and Cisco HyperFlex hyper-converged storage, which efficiently fits into a single data center rack, including the access layer network switches.

This validated design document details the deployment of the multiple configurations extending to 2000 users for Citrix virtual desktop or Citrix HSD published desktop workload respectively featuring the following software:

- Citrix Virtual Apps and Desktops Shared Remote Desktop Server Hosted (HSD) sessions on Cisco HyperFlex
- Citrix Virtual Apps and Desktops Non-Persistent and persistent Virtual Desktops (VDI) on Cisco HyperFlex
- Microsoft Windows Server 2016 for Citrix User Profile Manager
- Microsoft Windows 2016 Server for Login VSI Management and data servers to simulate real world VDI workload
- VMware vSphere ESXi 6.5.2 (Update 2) Hypervisor
- Windows Server 2016 for HSD Servers & Windows 10 64-bit Operating Systems for VDI virtual machines
- Microsoft SQL Server 2016
- Cisco HyperFlex data platform v3.5(1a)
- Citrix Virtual Apps and Desktops Studio Server and Director Servers for redundancy and support up to 2000 seat scale
- Citrix Virtual Apps and Desktops Storefront Server with redundancy

Figure 36 Detailed Reference Architecture with Physical Hardware Cabling Configured to Enable the Solution

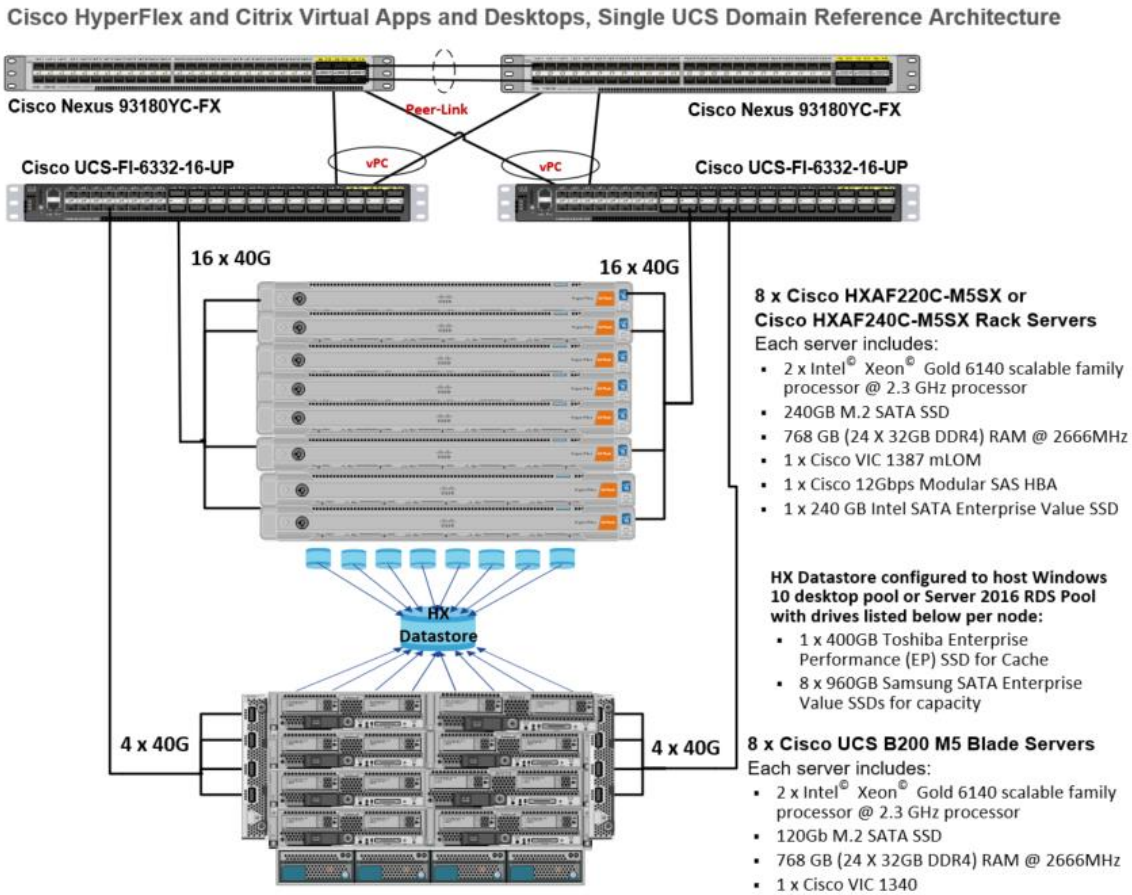

Hardware Deployed

The solution contains the following hardware as shown in Figure 36:

- Two Cisco Nexus 93180YC-FX Layer 2 Access Switches
- Two Cisco Fabric Interconnects 6332 UP
- Two Cisco UCS C220 M4 Rack servers with dual socket Intel Xeon E5-2620v4 2.1-GHz 8-core processors, 128GB RAM 2133-MHz and VIC1227 mLOM card for the hosted infrastructure with N+1 server fault tolerance. (Not show in the diagram).
- Eight Cisco UCS HXAF220c-M5S Rack servers with Intel Xeon Gold 6140 scalable family 2.3-GHz 18-core processors, 768GB RAM 2666-MHz and VIC1387 mLOM cards running Cisco HyperFlex data platform v3.5(1a) for the virtual desktop workloads with N+1 server fault tolerance.
- Eight Cisco UCS B200 M5 blade servers with Intel Xeon Gold 6140 scalable family 2.3-GHz 18-core processors, 768GB RAM 2666-MHz and VIC1340 mLOM cards running Cisco HyperFlex data platform v3.5(1a) for the virtual desktop workloads with N+1 server fault tolerance.

Software Deployed

Table 3 lists the software and firmware version used in the study.

Table 3 Software and Firmware Versions

Vendor	Product	Version
Cisco	UCS Component Firmware	4.0(1b) bundle release
Cisco	UCS Manager	4.0(1b) bundle release
Cisco	UCS HXAF220c-M5S rack server	4.0(1b) bundle release
Cisco	VIC 1387	4.2(2b)
Cisco	UCS B200 M5 blade server	4.0(1b) bundle release
Cisco	VIC 1340	4.2(2d)
Cisco	HyperFlex Data Platform	3.5.1a
Cisco	Cisco NENIC	2.1.2.71
Cisco	Cisco fNIC	1.6.0.37
Network	Cisco Nexus 9000 NX-OS	7.0(3)I7(2)
Citrix	Citrix Studio Server	1808
Citrix	Citrix Director Server	1808
Citrix	Citrix Provisioning Server	1808
Citrix	Citrix Client	4.9.0-9539668
VMware	vCenter Server Appliance	6.5.0-5973321
VMware	vSphere ESXi 6.5 Update 2	6.5.2.U2-8935087

Logical Architecture

The logical architecture of this solution has designed to support up to 2000 HSD hosted shared server desktop users and Hosted Virtual Microsoft Windows 10 Desktops within an eight node Cisco UCS HXAF220c-M5S, and eight Cisco UCS B200 M5 HyperFlex cluster, which provides physical redundancy for each workload type.

Figure 37 Logical Architecture Design

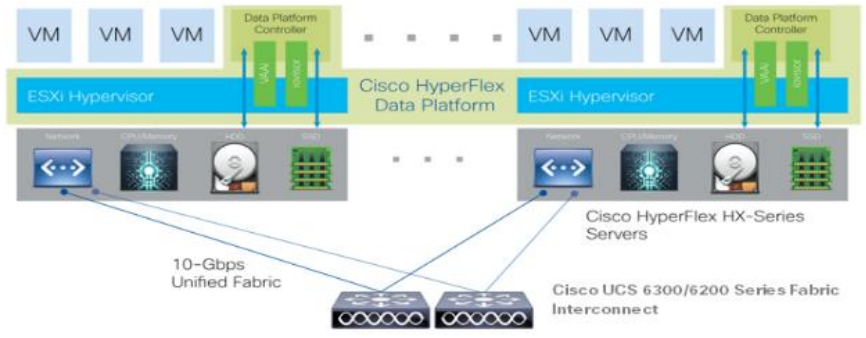

Table 3 lists the software revisions for this solution.

This document is intended to allow you to fully configure your environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well

as to record appropriate MAC addresses. Table 4 through Table 8 lists the information you need to configure your environment.

VLANs

The VLAN configuration recommended for the environment includes a total of seven VLANs as outlined in Table 4

Table 4 Table 2 VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
Hx-in-Band-Mgmt	50	VLAN for in-band management interfaces
Infra-Mgmt	51	VLAN for Virtual Infrastructure
Hx-storage-data	52	VLAN for HyperFlex Storage
Hx-vmotion	53	VLAN for VMware vMotion
Vm-network	54	VLAN for VDI Traffic
OOB-Mgmt	132	VLAN for out-of-band management interfaces

A dedicated network or subnet for physical device management is often used in datacenters. In this scenario, the mgmt0 interfaces of the two Fabric Interconnects would be connected to that dedicated network or subnet. This is a valid configuration for HyperFlex installations with the following caveat; wherever the HyperFlex installer is deployed it must have IP connectivity to the subnet of the mgmt0 interfaces of the Fabric Interconnects, and also have IP connectivity to the subnets used by the hx-inband-mgmt VLANs listed above.

Jumbo Frames

All HyperFlex storage traffic traversing the hx-storage-data VLAN and subnet is configured to use jumbo frames, or to be precise all communication is configured to send IP packets with a Maximum Transmission Unit (MTU) size of 9000 bytes. Using a larger MTU value means that each IP packet sent carries a larger payload, therefore transmitting more data per packet, and consequently sending and receiving data faster. This requirement also means that the Cisco UCS uplinks must be configured to pass jumbo frames. Failure to configure the Cisco UCS uplink switches to allow jumbo frames can lead to service interruptions during some failure scenarios, particularly when cable or port failures would cause storage traffic to traverse the northbound Cisco UCS uplink switches.

VMware Clusters

Three VMware Clusters were configured in one vCenter datacenter instance to support the solution and testing environment:

- Infrastructure Cluster: Infrastructure VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, Citrix Studio Server, Citrix Provisioning Servers, Citrix Storefront Server and HyperFlex Data Platform Installer, and so on).
- HyperFlex Cluster: Citrix Virtual App VMs (Windows Server 2016) or Persistent/Non-Persistent VDI VM Pools (Windows 10 64-bit).

HyperFlex release v3.0 or later supports 64 nodes in a single VMware cluster with 32 HXAF series HXAF220 or HXAF240 and 32 compute-only node. For more information, see: https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/Cisco_HXDataPlatform_RN_3_0.html.

- VSI Launcher Cluster: Login VSI Cluster (the Login VSI launcher infrastructure was connected using the same set of switches and vCenter instance but was hosted on separate local storage and servers).

Figure 38 VMware vSphere Clusters on vSphere Web GUI

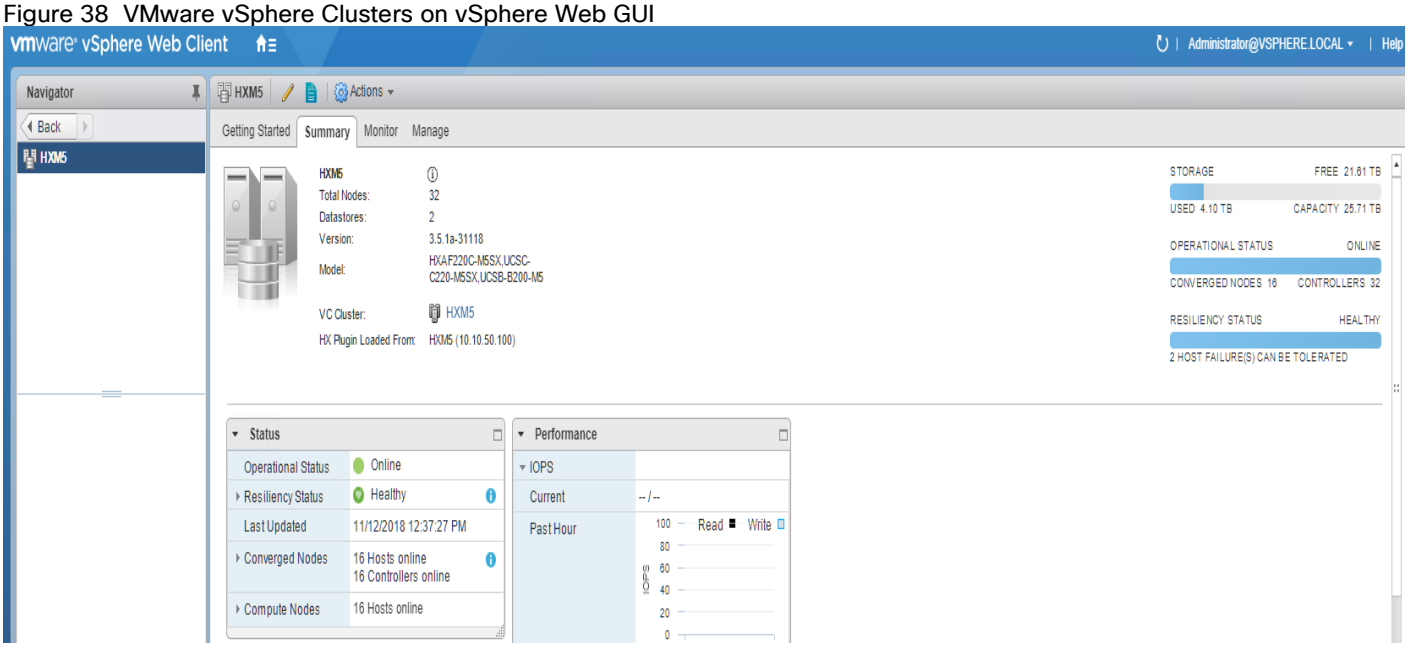

ESXi Host Design

The following sections detail the design of the elements within the VMware ESXi hypervisors, system requirements, virtual networking and the configuration of ESXi for the Cisco HyperFlex HX Distributed Data Platform.

Virtual Networking Design

The Cisco HyperFlex system has a pre-defined virtual network design at the ESXi hypervisor level. Four different virtual switches are created by the HyperFlex installer, each using two uplinks, which are each serviced by a vNIC defined in the UCS service profile. The vSwitches created are:

- vswitch-hx-inband-mgmt: This is the default vSwitch0 which is renamed by the ESXi kickstart file as part of the automated installation. The default vmkernel port, vmk0, is configured in the standard Management Network port group. The switch has two uplinks, active on fabric A and standby on fabric B, without jumbo frames. A second port group is created for the Storage Platform Controller VMs to connect to with their individual management interfaces. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere

- vswitch-hx-storage-data: This vSwitch is created as part of the automated installation. A vmkernel port, vmk1, is configured in the Storage Hypervisor Data Network port group, which is the interface used for connectivity to the HX Datastores through NFS. The switch has two uplinks, active on fabric B and standby on fabric A, with jumbo frames required. A second port group is created for the Storage Platform Controller

virtual machines to connect to with their individual storage interfaces. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere

- vswitch-hx-vm-network: This vSwitch is created as part of the automated installation. The switch has two uplinks, active on both fabrics A and B, and without jumbo frames. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere
- vmotion: This vSwitch is created as part of the automated installation. The switch has two uplinks, active on fabric A and standby on fabric B, with jumbo frames required. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere

The following table and figures help give more details into the ESXi virtual networking design as built by the HyperFlex installer:

Table 5 Table ESXi Host Virtual Switch Configuration

Virtual Switch	Port Groups	Active vmnic(s)	Passive vmnic(s)	VLAN IDs	Jumbo
vswitch-hx-inband-mgmt	Management Network Storage Controller Management Network	vmnic0	vmnic4	hx-inband-mgmt	no
vswitch-hx-storage-data	Storage Controller Data Network Storage Hypervisor Data Network	vmnic5	vmnic1	hx-storage-data	yes
vswitch-hx-vm-network	none	vmnic2 vmnic6	none	vm-network	no
vmotion	none	vmnic3	vmnic7	hx-vmotion	yes

Figure 39 ESXi Network Design

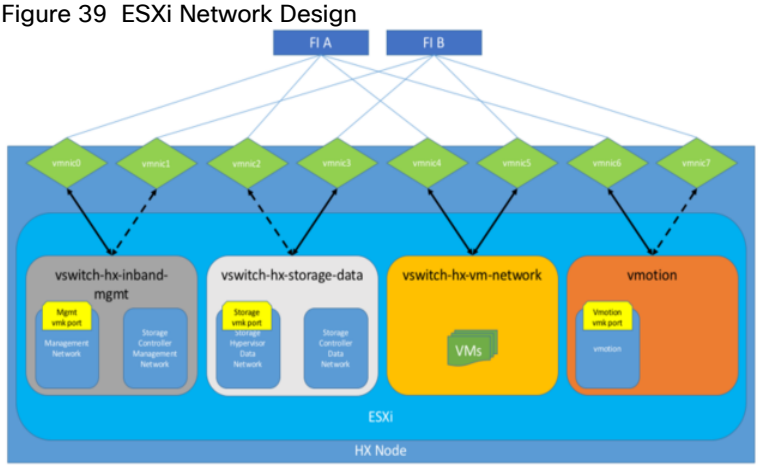

VMDirectPath I/O Pass-through

VMDirectPath I/O allows a guest VM to directly access PCI and PCIe devices in an ESXi host as though they were physical devices belonging to the VM itself, also referred to as PCI pass-through. With the appropriate driver for the hardware device, the guest VM sends all I/O requests directly to the physical device, bypassing the hypervisor. In the Cisco HyperFlex system, the Storage Platform Controller VMs use this feature to gain full control of the Cisco 12Gbps SAS HBA cards in the Cisco HX-series rack-mount servers. This gives the controller VMs

direct hardware level access to the physical disks installed in the servers, which they consume to construct the Cisco HX Distributed Filesystem. Only the disks connected directly to the Cisco SAS HBA or to a SAS extender, in turn connected to the SAS HBA are controlled by the controller VMs. Other disks, connected to different controllers, such as the SD cards, remain under the control of the ESXi hypervisor. The configuration of the VMDirectPath I/O feature is done by the Cisco HyperFlex installer and requires no manual steps.

Storage Platform Controller Virtual Machines

A key component of the Cisco HyperFlex system is the Storage Platform Controller Virtual Machine running on each of the nodes in the HyperFlex cluster. The controller virtual machines cooperate to form and coordinate the Cisco HX Distributed Filesystem, and service all the guest virtual machine IO requests. The controller VMs are deployed as a vSphere ESXi agent, which is similar in concept to that of a Linux or Windows service. ESXi agents are tied to a specific host, they start and stop along with the ESXi hypervisor, and the system is not considered to be online and ready until both the hypervisor and the agents have started. Each ESXi hypervisor host has a single ESXi agent deployed, which is the controller virtual machine for that node, and it cannot be moved or migrated to another host. The collective ESXi agents are managed through an ESXi agency in the vSphere cluster.

The storage controller virtual machine runs custom software and services that manage and maintain the Cisco HX Distributed Filesystem. The services and processes that run within the controller virtual machines are not exposed as part of the ESXi agents to the agency, therefore the ESXi hypervisors nor vCenter server have any direct knowledge of the storage services provided by the controller virtual machines. Management and visibility into the function of the controller virtual machines, and the Cisco HX Distributed Filesystem is done through a plug-in installed to the vCenter server or appliance managing the vSphere cluster. The plugin communicates directly with the controller virtual machines to display the information requested, or make the configuration changes directed, all while operating within the same web-based interface of the vSphere Web Client. The deployment of the controller virtual machines, agents, agency, and vCenter plugin are all done by the Cisco HyperFlex installer and requires no manual steps.

Controller Virtual Machine Locations

The physical storage location of the controller virtual machine is similar between the Cisco HXAF220c-M5S and HXAF240c-M5SX model servers. The storage controller virtual machine is operationally no different from any other typical virtual machines in an ESXi environment. The virtual machine must have a virtual disk with the bootable root filesystem available in a location separate from the SAS HBA that the virtual machine is controlling through VMDirectPath I/O. The configuration details of the models are as follows:

The Cisco UCS compute-only Nodes also place a lightweight storage controller virtual machine on a 3.5 GB VMFS datastore, provisioned from the M.2 SATA SSD drive.

Cisco HyperFlex Datastores

The new HyperFlex cluster has no default datastores configured for virtual machine storage, therefore the datastores must be created using the vCenter Web Client plugin or HyperFlex Connect GUI. A minimum of two datastores is recommended to satisfy vSphere High Availability datastore heartbeat requirements, although one of the two datastores can be very small. It is important to recognize that all HyperFlex datastores are thinly provisioned, meaning that their configured size can far exceed the actual space available in the HyperFlex cluster. Alerts will be raised by the HyperFlex system in the vCenter plugin when actual space consumption results in low amounts of free space, and alerts will be sent through auto support email alerts. Overall space consumption in the HyperFlex clustered filesystem is optimized by the default deduplication and compression features.

Cisco HyperFlex Datastores

The new HyperFlex cluster has no default datastores configured for virtual machine storage, therefore the datastores must be created using the vCenter Web Client plug-in. A minimum of two datastores is recommended

to satisfy vSphere High Availability datastore heartbeat requirements, although one of the two datastores can be very small. It is important to recognize that all HyperFlex datastores are thinly provisioned, meaning that their configured size can far exceed the actual space available in the HyperFlex cluster. Alerts will be raised by the HyperFlex system in the vCenter plugin when actual space consumption results in low amounts of free space, and alerts will be sent through auto support email alerts. Overall space consumption in the HyperFlex clustered filesystem is optimized by the default deduplication and compression features.

Figure 40 Datastore Example

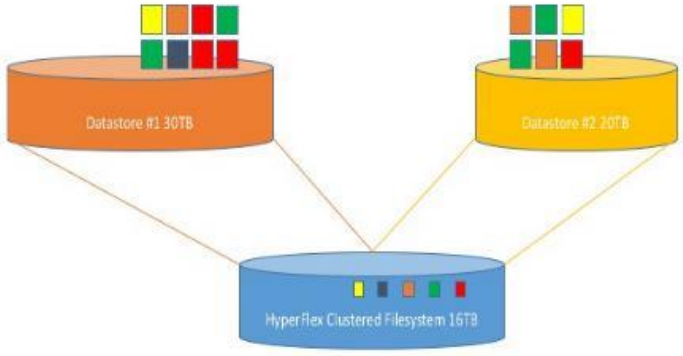

CPU Resource Reservations

Since the storage controller virtual machines provide critical functionality of the Cisco HX Distributed Data Platform, the HyperFlex installer will configure CPU resource reservations for the controller virtual machines. This reservation guarantees that the controller virtual machines will have CPU resources at a minimum level, in situations where the physical CPU resources of the ESXi hypervisor host are being heavily consumed by the guest virtual machines. Table 6 details the CPU resource reservation of the storage controller virtual machines:

Table 6 Controller Virtual Machine CPU Reservations

Number of vCPU	Shares	Reservation	Limit
8	Low	10800 MHz	unlimited

Memory Resource Reservations

Since the storage controller virtual machines provide critical functionality of the Cisco HX Distributed Data Platform, the HyperFlex installer will configure memory resource reservations for the controller virtual machines. This reservation guarantees that the controller virtual machines will have memory resources at a minimum level, in situations where the physical memory resources of the ESXi hypervisor host are being heavily consumed by the guest virtual machines.

Table 7 lists the memory resource reservation of the storage controller virtual machines.

Table 7 Controller Virtual Machine Memory Reservations

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HX220c-M5S HXAF220c-M5S	48 GB	Yes
HX240c-M5SX HXAF240c-M5SX	72 GB	Yes

The Cisco UCS compute-only Nodes have a lightweight storage controller virtual machine; it is configured with only 1 vCPU and 512 MB of memory reservation.

Solution Configuration

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution. Figure 41 illustrates the configuration topology for this solution.

Figure 41 Configuration Topology for Scalable Citrix VDI Workload with HyperFlex

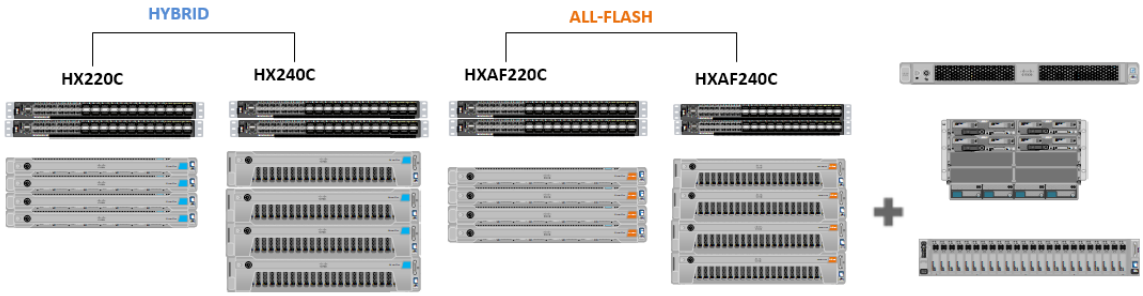

Cisco UCS Compute Platform

The following subsections detail the physical connectivity configuration of the Citrix VDI environment.

Physical Infrastructure

Solution Cabling

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

The tables in this section contain the details for the prescribed and supported configuration.

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.

Be sure to follow the cabling directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.

Figure 36 shows a cabling diagram for a Citrix configuration using the Cisco Nexus 9000 and Cisco UCS Fabric Interconnect.

Table 8 Cisco Nexus 93180 A-Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 A	Eth1/1	10GbE	Cisco Nexus 93180 B	Eth1/1
	Eth1/2	10GbE	Cisco Nexus 93180 B	Eth1/2
	Eth1/3	10GbE	Cisco Nexus 93180 B	Eth1/3
	Eth1/4	10GbE	Cisco Nexus 93180 B	Eth1/4
	Eth1/11	10GbE	Launcher-FI-A	Eth/29

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/12	10GbE	Launcher-FI-A	Eth/30
	Eth1/13	10GbE	Launcher-FI-B	Eth/29
	Eth1/14	10GbE	Launcher-FI-B	Eth/30
	MGMT0	GbE	GbE management switch	Any
	Eth1/15	10GbE	Infra-host-01	Port01
	Eth1/16	10GbE	Infra-host-02	Port01
	Eth1/49	40GbE	Cisco UCS fabric interconnect A	Eth1/35
	Eth1/50	40GbE	Cisco UCS fabric interconnect B	Eth1/35

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

Table 9 Cisco Nexus 93180 B -Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 B	Eth1/1	10GbE	Cisco Nexus 93180 A	Eth1/1
	Eth1/2	10GbE	Cisco Nexus 93180 A	Eth1/2
	Eth1/3	10GbE	Cisco Nexus 93180 A	Eth1/3
	Eth1/4	10GbE	Cisco Nexus 93180 A	Eth1/4
	Eth1/11	10GbE	Launcher-FI-A	Eth/31
	Eth1/12	10GbE	Launcher-FI-A	Eth/32
	Eth1/13	10GbE	Launcher-FI-B	Eth/31
	Eth1/14	10GbE	Launcher-FI-B	Eth/32
	MGMT0	GbE	GbE management switch	Any
	Eth1/15	10GbE	Infra-host-01	Port02
	Eth1/16	10GbE	Infra-host-02	Port02
	Eth1/49	40GbE	Cisco UCS fabric interconnect B	Eth1/36
	Eth1/50	40GbE	Cisco UCS fabric interconnect B	Eth1/36

Table 10 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric interconnect A	Eth1/1	40GbE	Server 1	Port01
	Eth1/2	40GbE	Server 2	Port01

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/3	40GbE	Server 3	Port01
	Eth1/4	40GbE	Server 4	Port01
	Eth1/5	40GbE	Server 5	Port01
	Eth1/6	40GbE	Server 6	Port01
	Eth1/7	40GbE	Server 7	Port01
	Eth1/8	40GbE	Server 8	Port01
	Eth1/9	40GbE	Server 9	Port01
	Eth1/10	40GbE	Server 10	Port01
	Eth1/11	40GbE	Server 11	Port01
	Eth1/12	40GbE	Server 12	Port01
	Eth1/13	40GbE	Server 13	Port01
	Eth1/14	40GbE	Server 14	Port01
	Eth1/15	40GbE	Server 15	Port01
	Eth1/16	40GbE	Server 16	Port01
	Eth1/17	40GbE	Server 17	Port01
	Eth1/18	40GbE	Server 18	Port01
	Eth1/19	40GbE	Server 19	Port01
	Eth1/20	40GbE	Server 20	Port01
	Eth1/21	40GbE	Server 21	Port01
	Eth1/22	40GbE	Server 22	Port01
	Eth1/23	40GbE	Server 23	Port01
	Eth1/24	40GbE	Server 24	Port01
	Eth1/25	40GbE	5108 IOM Chassis-A	IOM-A/1/1
	Eth1/26	40GbE	5108 IOM Chassis-A	IOM-A/1/2
	Eth1/27	40GbE	Cisco Nexus 93180 A	Eth1/49
	Eth1/28	40GbE	Cisco Nexus 93180 B	Eth1/49
	MGMT0	GbE	GbE management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 11 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric interconnect B	Eth1/1	40GbE	Server 1	Port02
	Eth1/2	40GbE	Server 2	Port02
	Eth1/3	40GbE	Server 3	Port02
	Eth1/4	40GbE	Server 4	Port02
	Eth1/5	40GbE	Server 5	Port02
	Eth1/6	40GbE	Server 6	Port02
	Eth1/7	40GbE	Server 7	Port02
	Eth1/8	40GbE	Server 8	Port02
	Eth1/9	40GbE	Server 9	Port02
	Eth1/10	40GbE	Server 10	Port02
	Eth1/11	40GbE	Server 11	Port02
	Eth1/12	40GbE	Server 12	Port02
	Eth1/13	40GbE	Server 13	Port02
	Eth1/14	40GbE	Server 14	Port02
	Eth1/15	40GbE	Server 15	Port02
	Eth1/16	40GbE	Server 16	Port02
	Eth1/17	40GbE	Server 17	Port02
	Eth1/18	40GbE	Server 18	Port02
	Eth1/19	40GbE	Server 19	Port02
	Eth1/20	40GbE	Server 20	Port02
	Eth1/21	40GbE	Server 21	Port02
	Eth1/22	40GbE	Server 22	Port02
	Eth1/23	40GbE	Server 23	Port02
	Eth1/24	40GbE	Server 24	Port02
	Eth1/25	40GbE	5108 IOM Chassis-B	IOM-B/1/1

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/26	40GbE	5108 IOM Chassis-B	IOM-B/1/2
	Eth1/35	40GbE	Cisco Nexus 93180 A	Eth1/50
	Eth1/36	40GbE	Cisco Nexus 93180 B	Eth1/50
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Figure 42 Cable Connectivity between Cisco Nexus 93180 A and B to Cisco UCS 6332 Fabric A and B

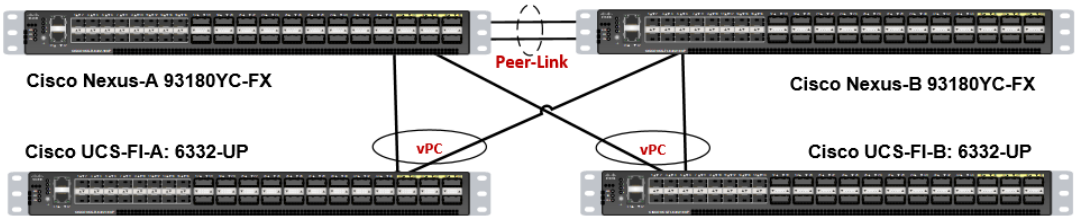

Cisco Unified Computing System Configuration

This section details the Cisco UCS configuration performed as part of the infrastructure build out by the Cisco HyperFlex installer. Many of the configuration elements are fixed in nature, meanwhile the HyperFlex installer does allow for some items to be specified at the time of creation, for example VLAN names and IDs, IP pools and more. Where the elements can be manually set during the installation, those items will be noted in << >> brackets.

For more information about racking, power, and installation of the chassis is described in the install guide (see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html) and it is beyond the scope of this document. For more information about each step, refer to the following documents: Cisco UCS Manager Configuration Guides – GUI and Command Line Interface (CLI) Cisco UCS Manager – Configuration Guides – Cisco

During the HyperFlex Installation, a Cisco UCS Sub-Organization is created named "hx-cluster." The sub-organization is created below the root level of the Cisco UCS hierarchy, and is used to contain all policies, pools, templates and service profiles used by HyperFlex. This arrangement allows for organizational control using Role-Based Access Control (RBAC) and administrative locales at a later time if desired. In this way, control can be granted to administrators of only the HyperFlex specific elements of the Cisco UCS domain, separate from control of root level elements or elements in other sub-organizations.

Figure 43 Cisco UCS Manager Configuration: HyperFlex Sub-organization

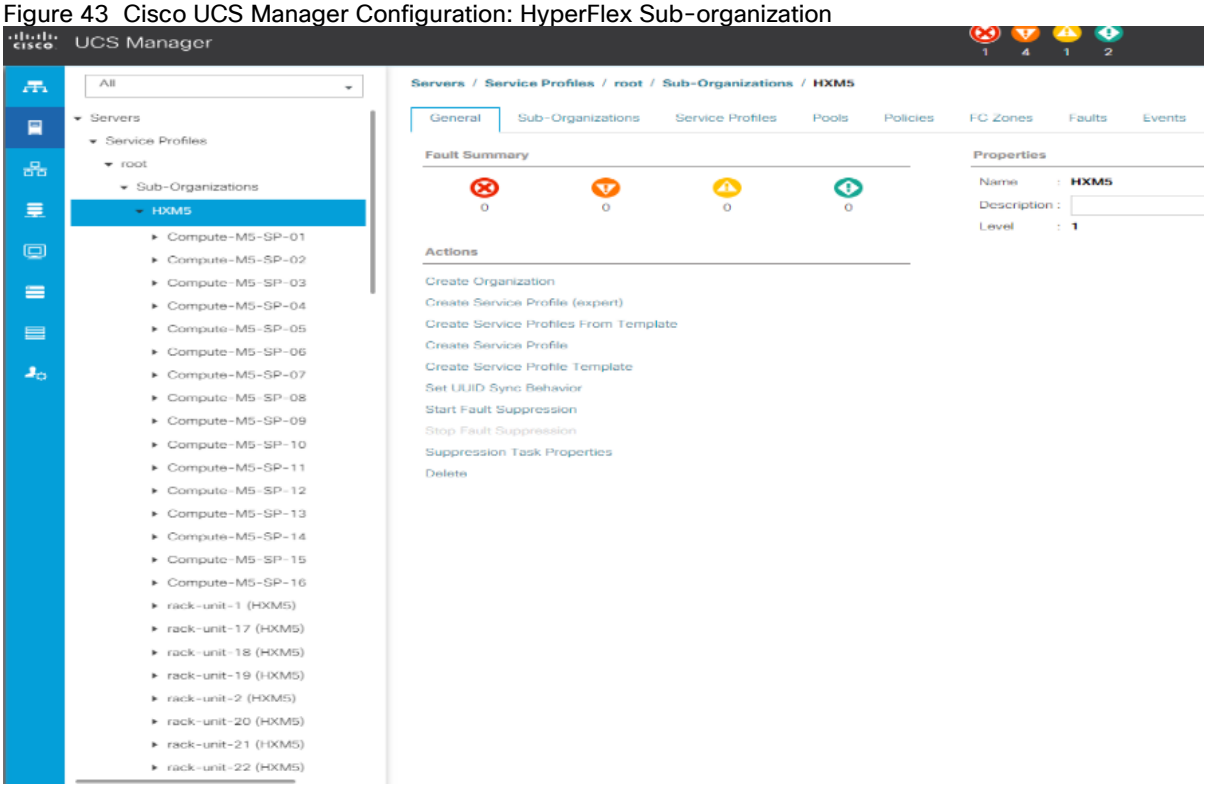

Deploy and Configure HyperFlex Data Platform

Prerequisites

To deploy and configure the HyperFlex Data Platform, complete the following steps:

1. Set Time Zone and NTP: From the Cisco UCS Manager, from the Admin tab, Configure TimeZone and add NTP server. Save changes.

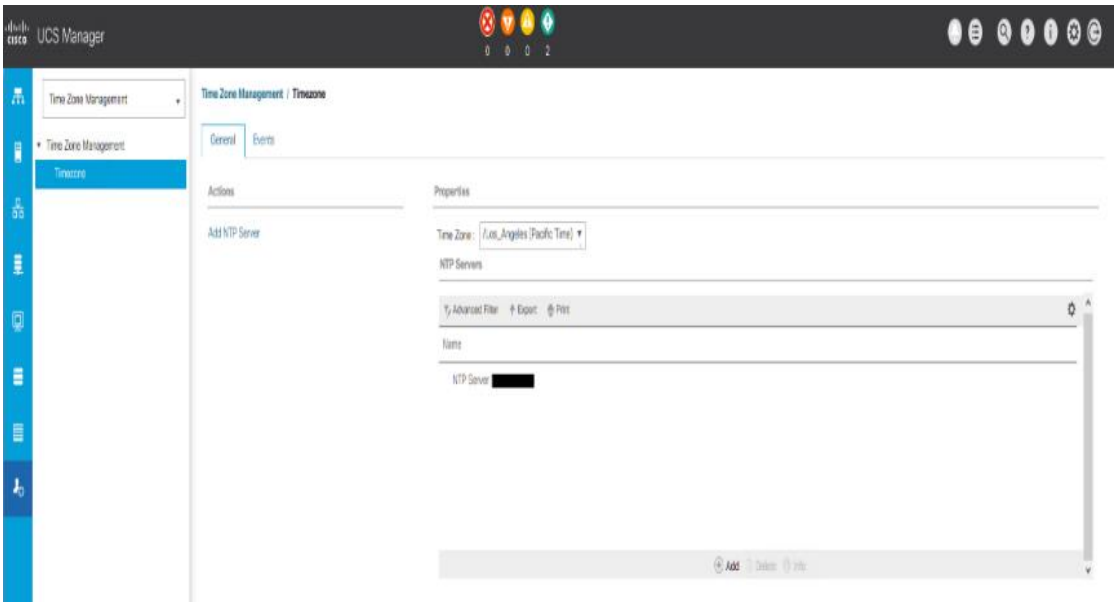

2. Configure Server Ports: Under the Equipment tab, Select Fabric A, select port to be configured as server port to manager HyperFlex rack server through Cisco UCS Manager.

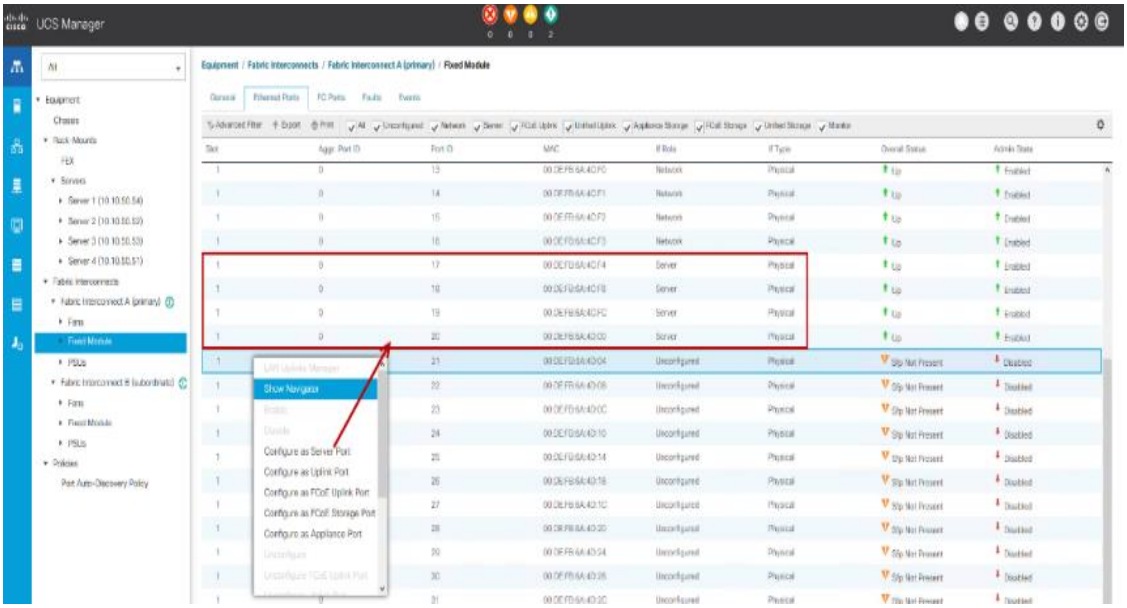

3. Repeat this step to configure server port on Fabric B.

4. Configure Uplink Ports: On Fabric A, Select port to be configured as uplink port for network connectivity to north bound switch.

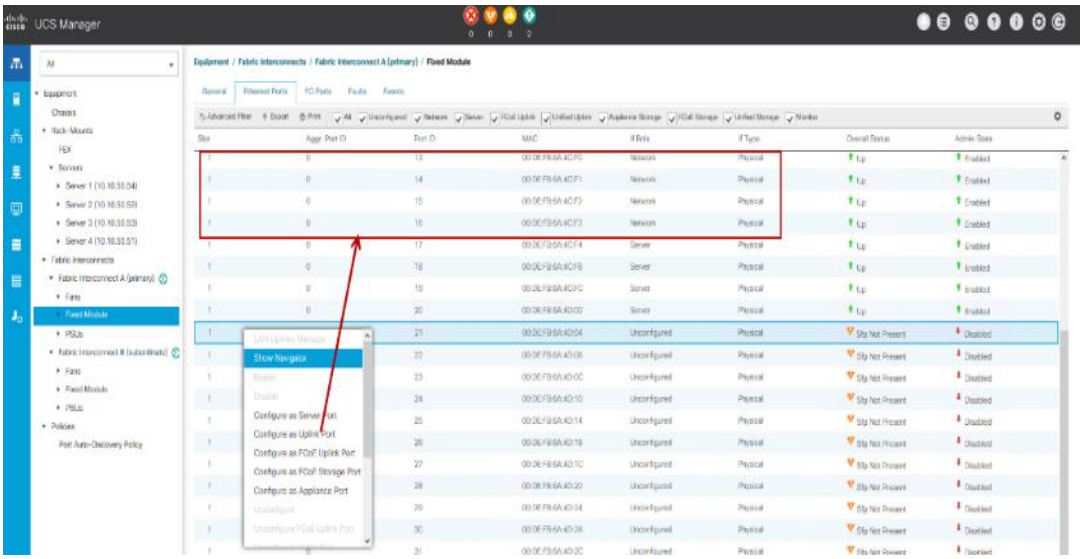

5. Repeat this same on Fabric B.

6. Create Port Channels: Under LAN tab, select expand LAN > LAN cloud > Fabric A. Right-click Port Channel.

7. Select Create a port channel to connect with the upstream switch as per Cisco UCS best practices. For our reference architecture, we connected a pair of Nexus 93180 switches.

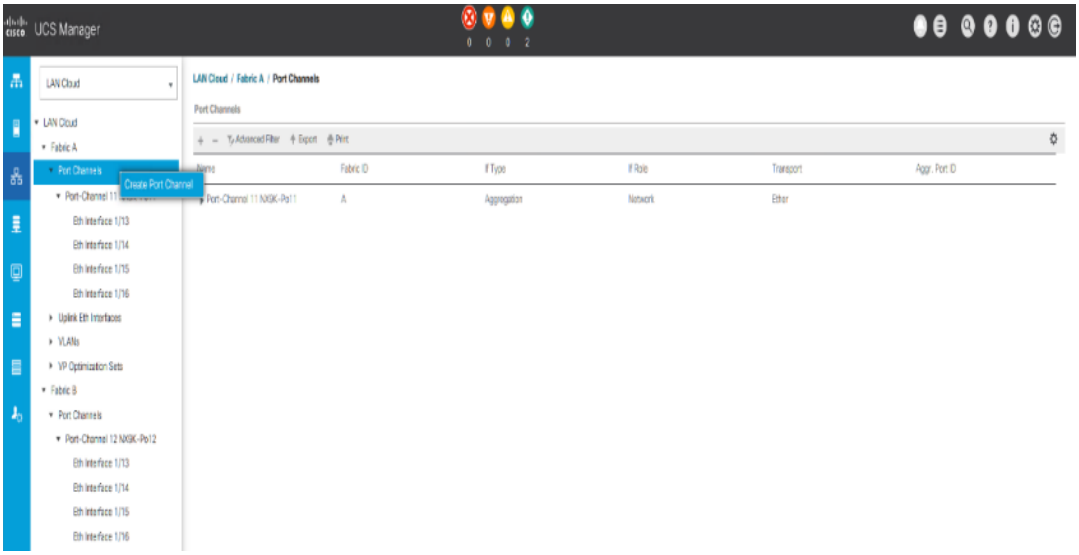

8. Enter the port channel ID number and name to be created, click Next.

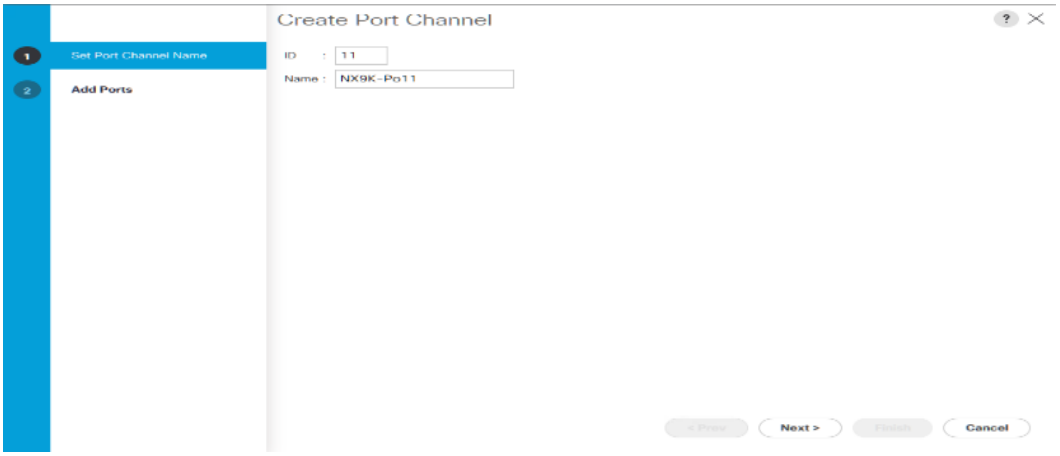

9. Select uplink ports to add as part of the port channel.

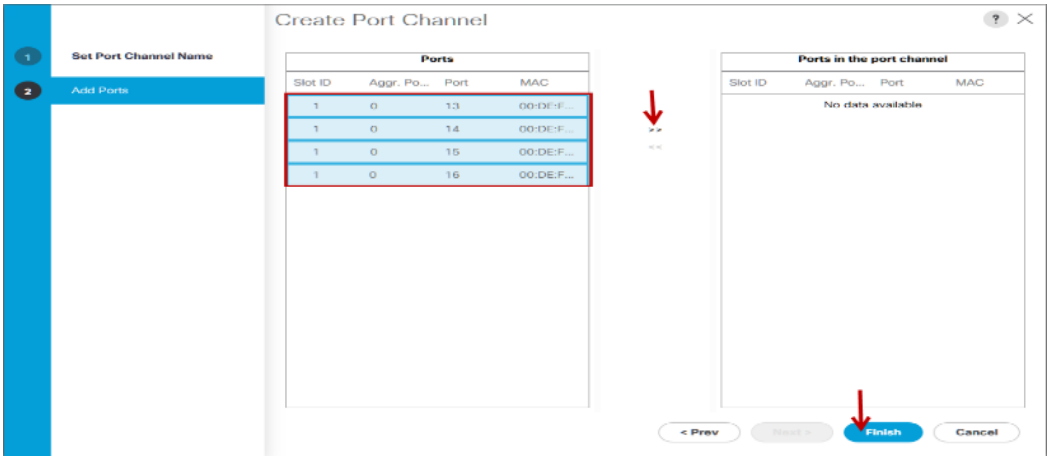

10. Click Finish.

11. Follow the previous steps to create the port channel on Fabric B, using a different port channel ID.

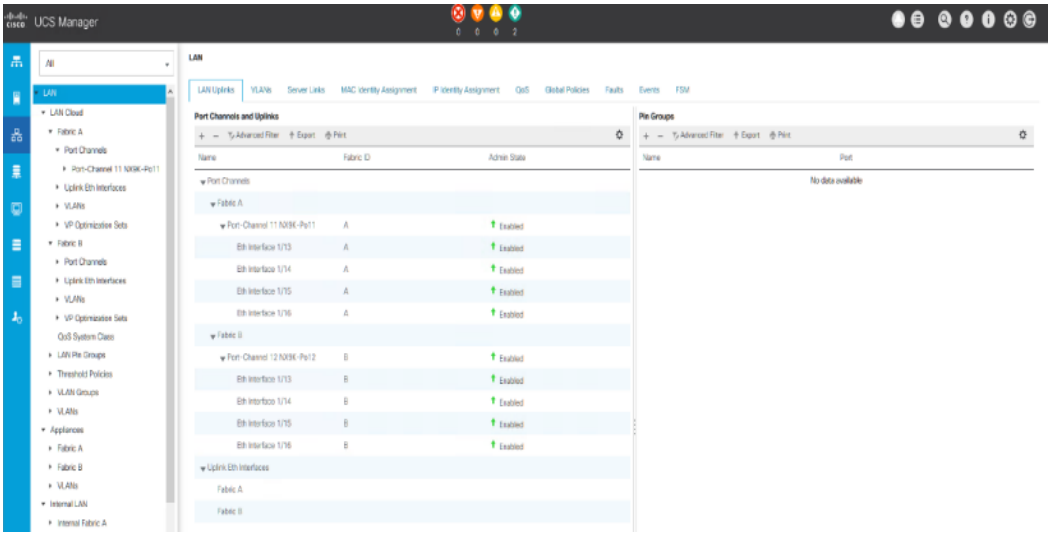

12. Configure QoS System Classes: From the LAN tab, under the Lan Cloud node, select QoS system class and configure the Platinum through Bronze system classes as shown in the following figure.

    a. Set MTU to 9216 for Platinum (Storage data) and Bronze (vMotion)

    b. Uncheck Enable Packet drop on the Platinum class

    c. Set Weight for Platinum and Gold priority class to 4 and everything else as best-effort.

    d. Enable multicast for silver class.

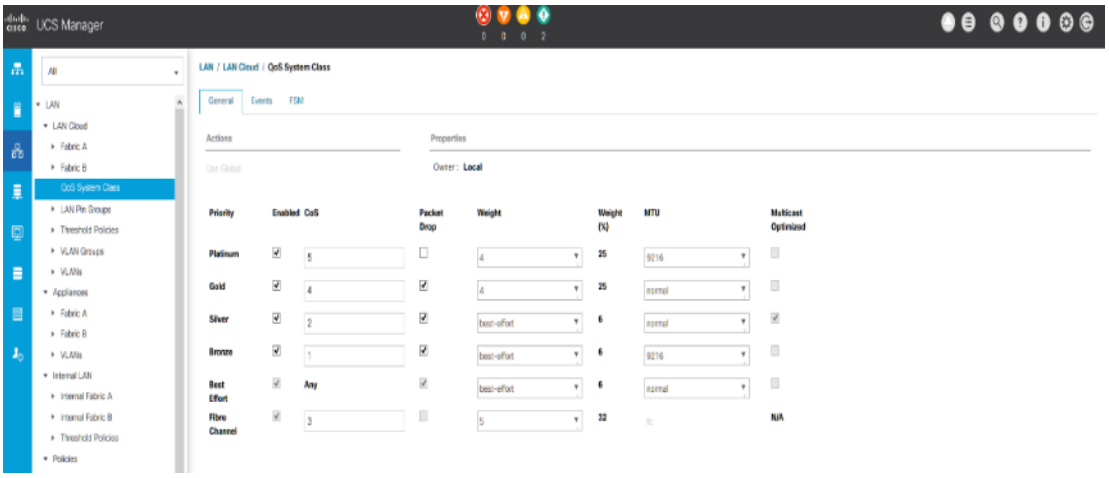

Changing QoS system class configuration on 6300 series Fabric Interconnect requires reboot of FIs.

13. Verify the UCS Manager Software Version. In the Equipment tab, select Firmware Management > Installed Firmware.

14. Check and verify both Fabric Interconnects and Cisco USC Manager are configure with Cisco UCS Manager v4.0.1a.

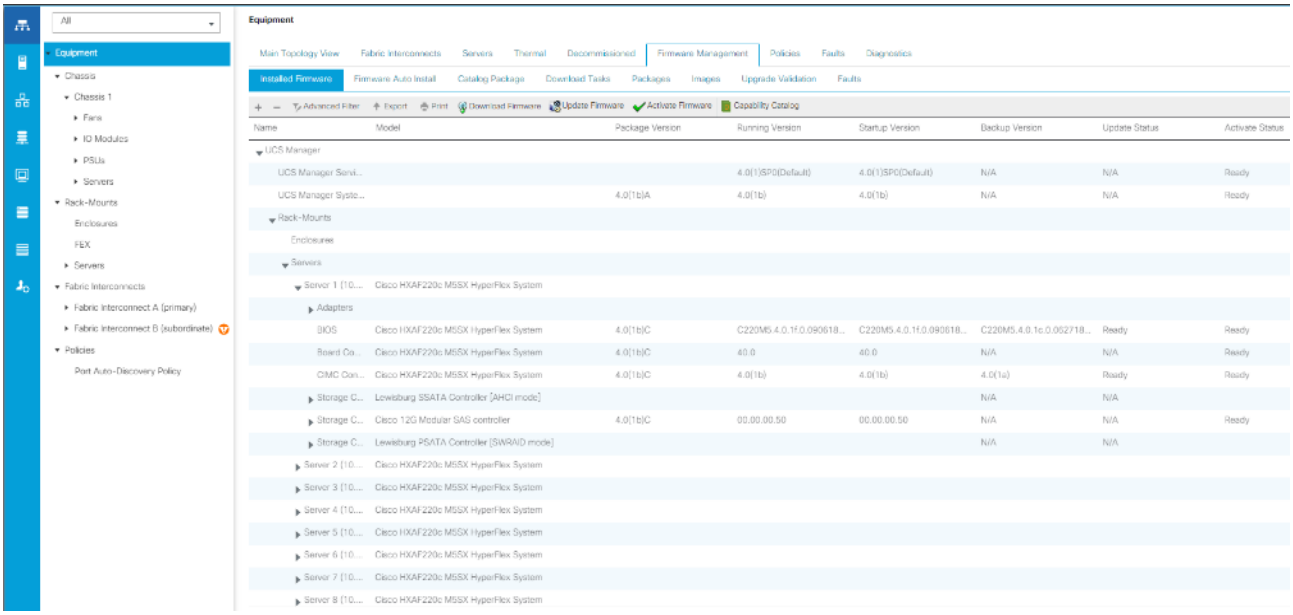

It is recommended to let the HX Installer handle upgrading the server firmware automatically as designed. This will occur once the service profiles are applied to the HX nodes during the automated deployment process.

15. Optional: If you are familiar with Cisco UCS Manager or you wish to break the install into smaller pieces, you can use the server auto firmware download to pre-stage the correct firmware on the nodes. This will speed up the association time in the HyperFlex installer at the cost of running two separate reboot operations. This method is not required or recommended if doing the install in one sitting.

Deploy Cisco HyperFlex Data Platform Installer Virtual Machine

Download the latest installer OVA from Cisco.com. Software Download Link:

https://software.cisco.com/download/home/286305544/type/286305994/release/3.5%25281a%2529

Deploy OVA to an existing host in the environment. Use either your existing vCenter Thick Client (C#) or vSphere Web Client to deploy OVA on ESXi host. This document outlines the procedure to deploy the OVA from the web client.

To deploy the OVA from the web client, complete the following steps:

1. Log into vCenter web client with your login to web browser with vCenter management IP address: https://<FQDN or IP address for VC>:9443/vcenter-client

2. Select ESXi host under hosts and cluster when HyperFlex data platform installer virtual machine to deploy.

3. Right-click ESXi host, select Deploy OVF Template.

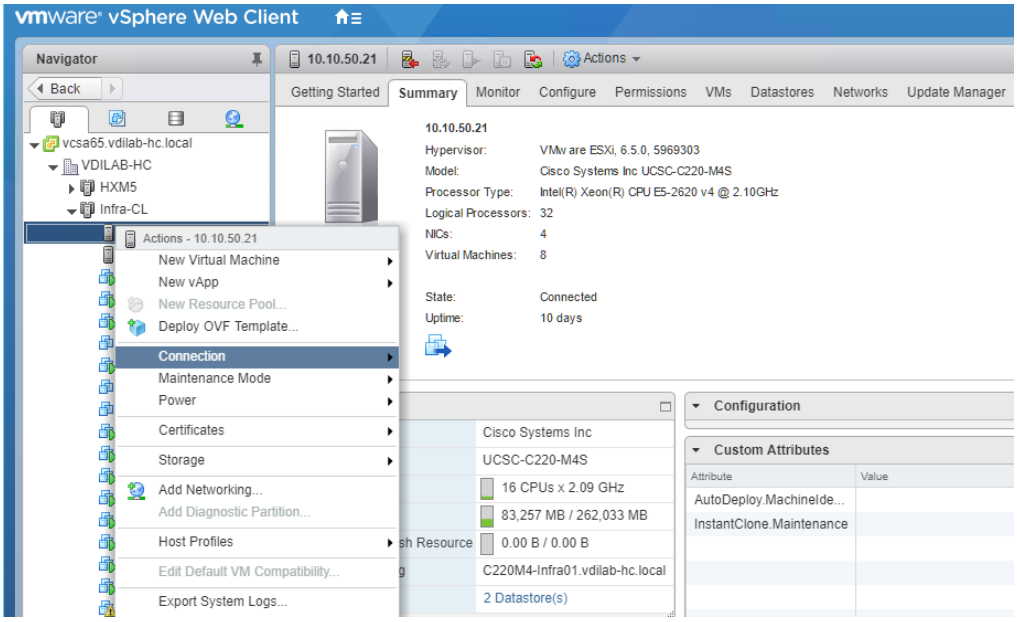

4. Follow the deployment steps to configure HyperFlex data-platform installer virtual machine deployment.

5. Select OVA file to deploy, click Next.

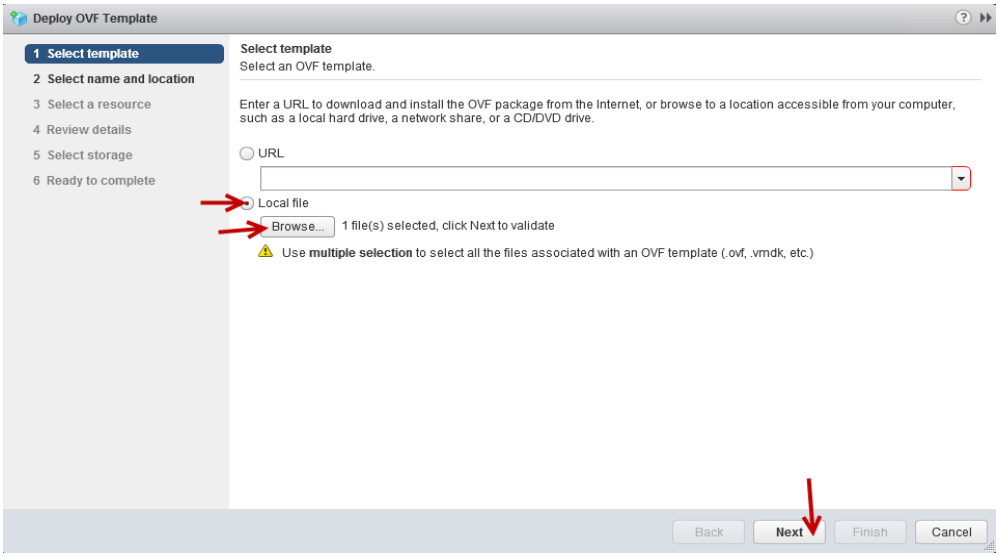

6. Review and verify the details for OVF template to deploy, click Next.

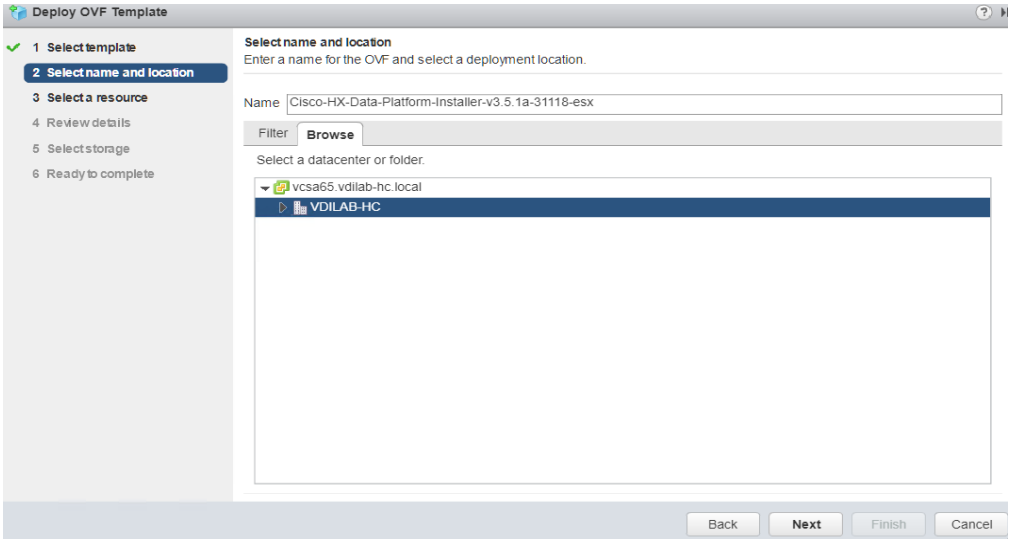

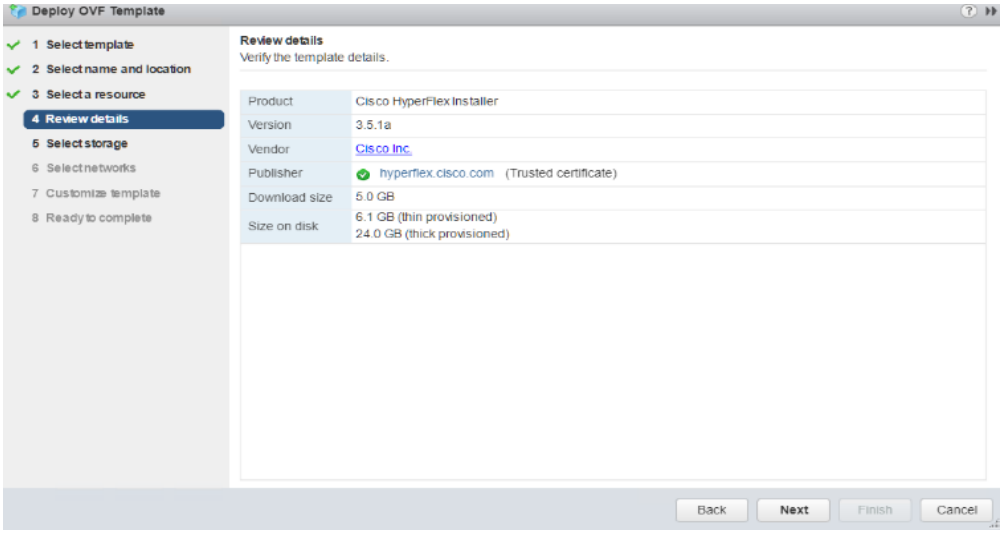

7. Enter name for OVF to template deploy, select datacenter and folder location. Click Next.

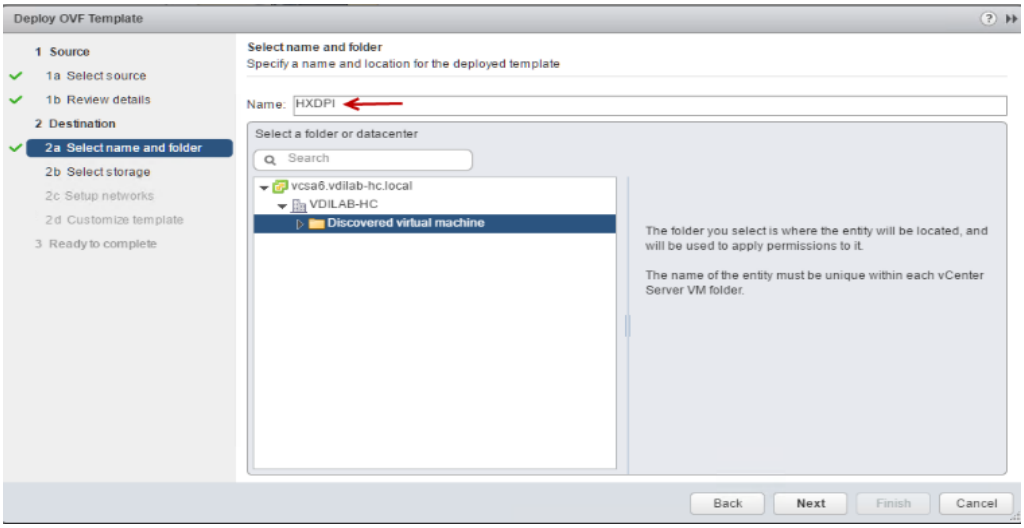

8. Select virtual disk format, virtual machine storage policy set to datastore default, select datastore for OVF deployment. Click Next.

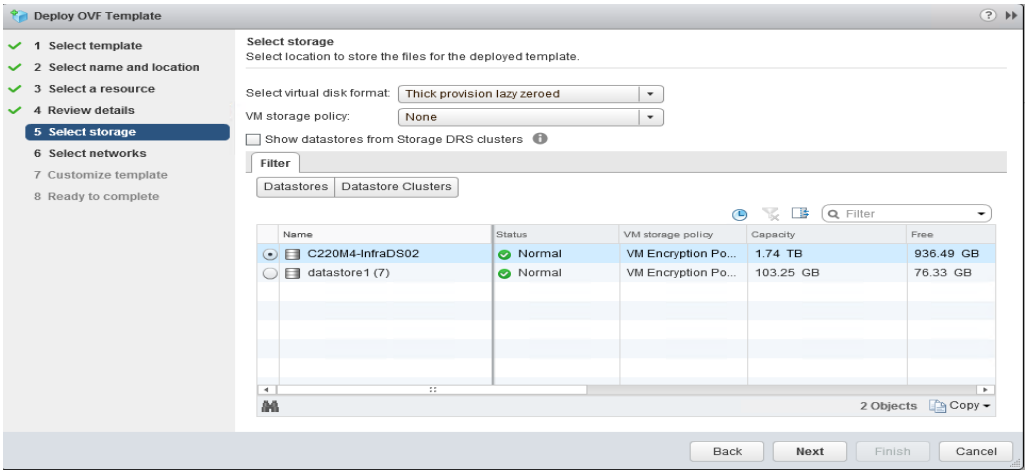

9. Select Network adapter destination port-group.

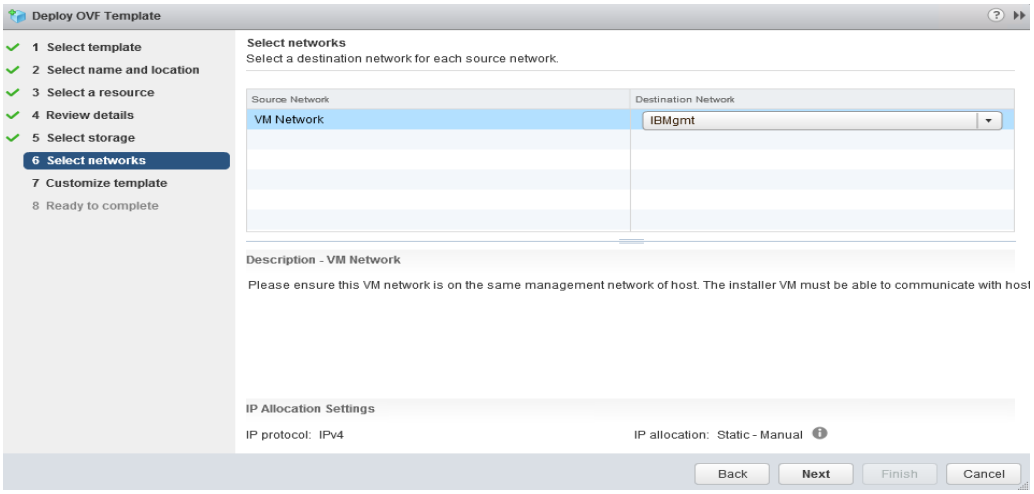

10. Fill out the parameters requested for hostname, gateway, DNS, IP address, and netmask. Alternatively, leave all blank for a DHCP assigned address.

Provide a single DNS server only. Inputting multiple DNS servers will cause queries to fail. You must connect to vCenter to deploy the OVA file and provide the IP address properties. Deploying directly from an ESXi host will not allow you to set these values correctly.

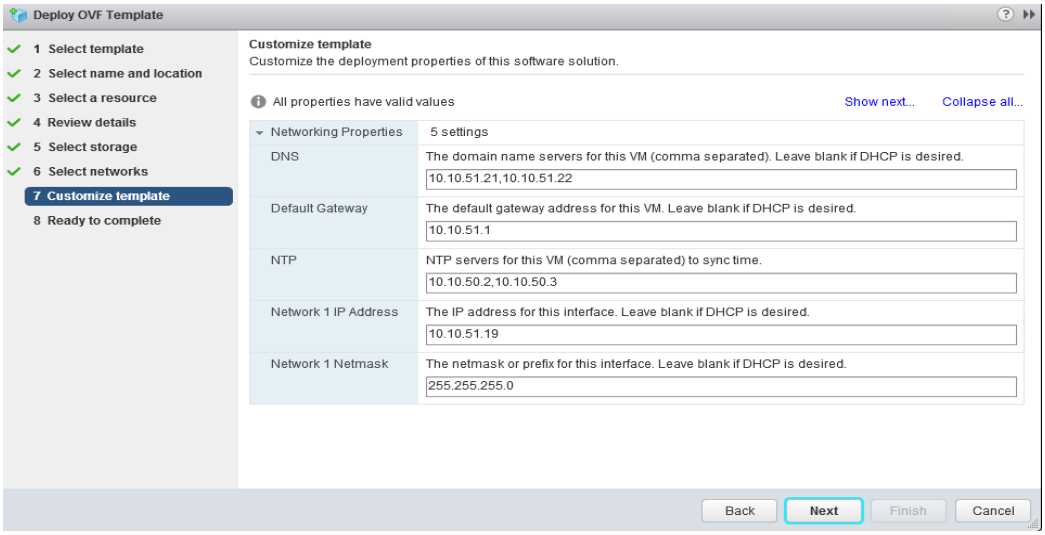

If you have internal firewall rules between these networks, please contact TAC for assistance.

If required, an additional network adapter can be added to the HyperFlex Platform Installer virtual machine after OVF deployment is completed successfully. For example, in case of a separate Inband and Out-Of-Mgmt network, see the screenshot below:

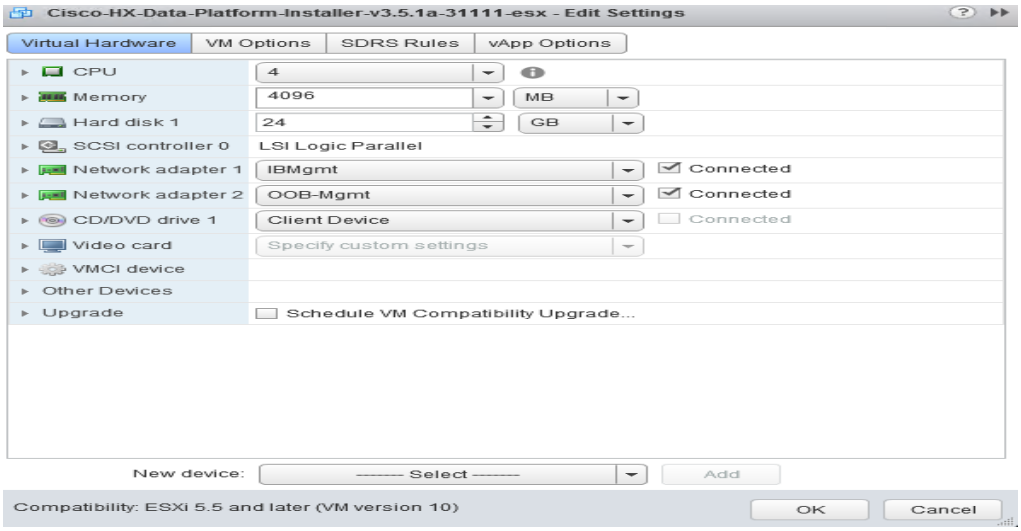

11. Review the settings selected as part of the OVF deployment, click the checkbox for Power on after deployment. Click Finish.

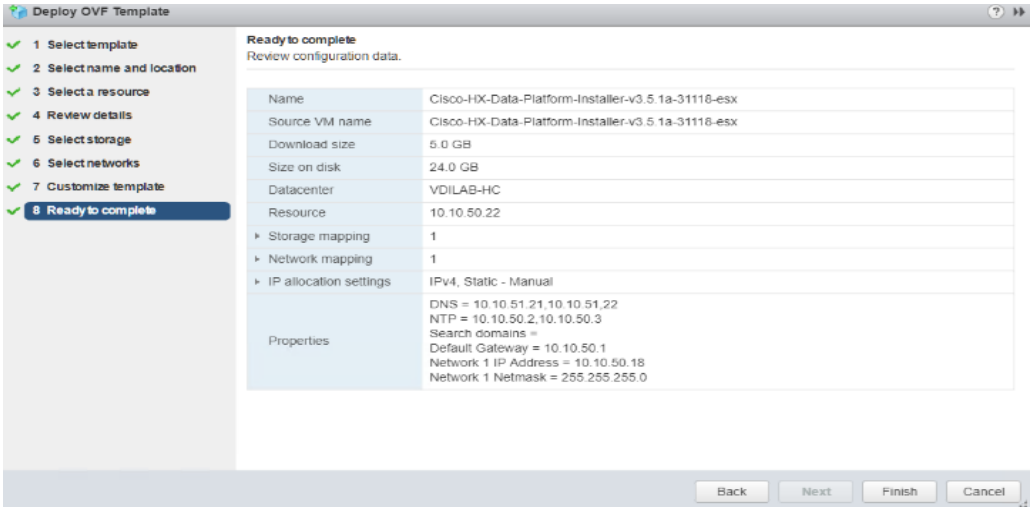

The default credentials for the HyperFlex installer virtual machine are: user name: root password: Cisco123.

Verify or Set DNS Resolution

SSH to HX installer virtual machine, verify or set DNS resolution is set on HyperFlex Installer virtual machine:

```
root@Cisco-HX-Data-Platform-Installer: # more /etc/network/eth0.interface
auto eth0
iface eth0 inet static
metric 100
address 10.10.50.18
netmask 255.255.255.0
gateway 10.10.50.1
dns-search vdilab-hc.local
dns-nameservers 10.10.51.21 10.10.51.22

root@Cisco-HX-Data-Platform-Installer:~# more /run/resolvconf/resolv.conf
# Dynamic resolv.conf(5) file for glibc resolver(3) generated by resolvconf(8)
#     DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
nameserver 10.10.51.21
nameserver 10.10.51.22
search vdilab-hc.local
```

Cisco HyperFlex Cluster Configuration

To configuring the Cisco HyperFlex Cluster, complete the following step:

1. Log into the HX Installer virtual machine through a web browser: http://<Installer_VM_IP_Address>.

Create a HyperFlex Cluster

1. Select the workflow for cluster creation to deploy a new HyperFlex cluster on eight Cisco HXAF220c-M5S nodes.

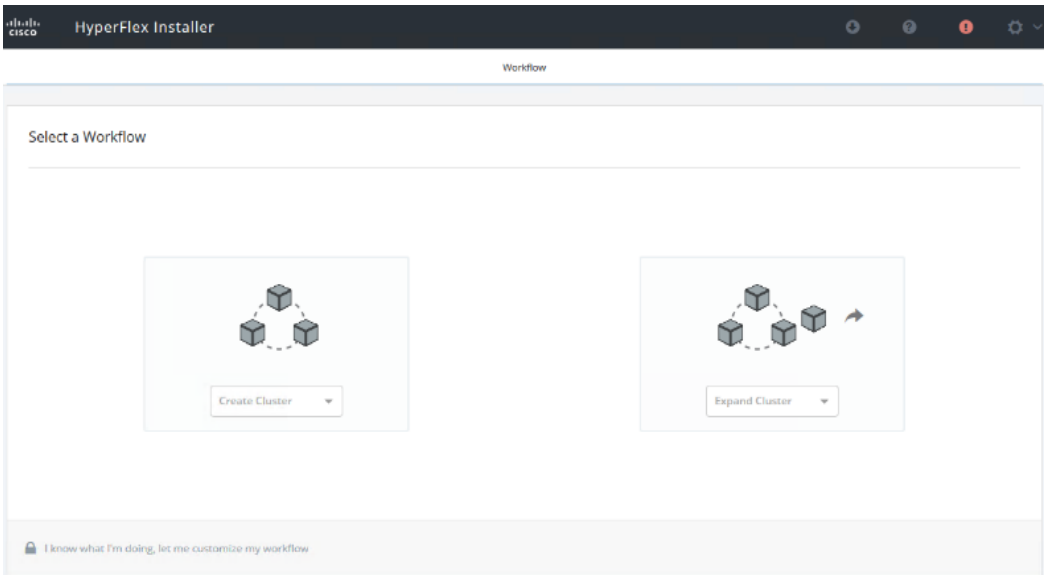

2. On the credentials page, enter the access details for Cisco UCS Manager, vCenter server, and Hypervisor. Click Continue.

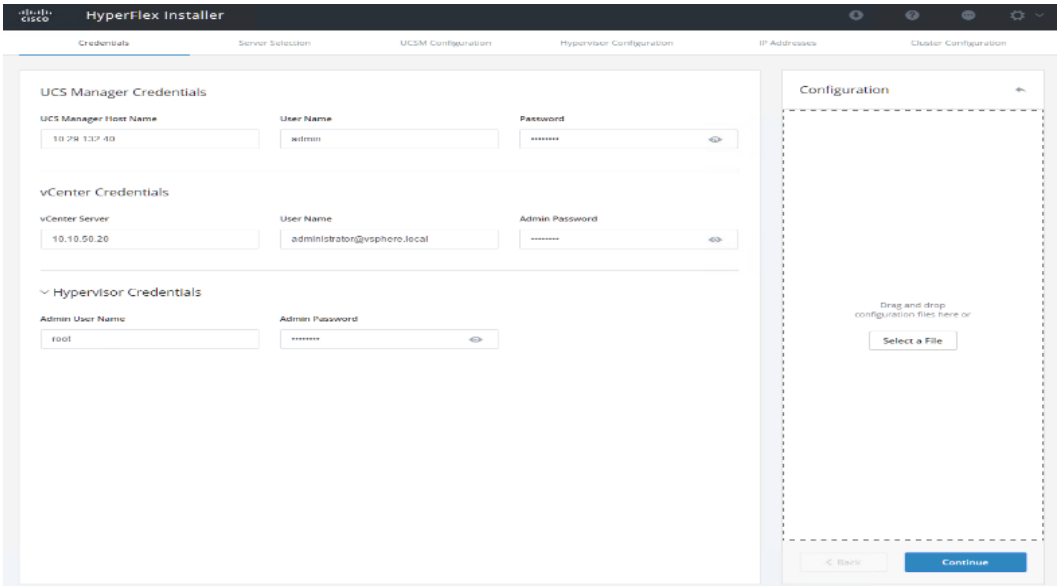

3. Select the top-most check box at the top right corner of the HyperFlex installer to select all unassociated servers. (To configure a subset of available of the HyperFlex servers, manually click the check box for individual servers.)

4. Click Continue after completing server selection.

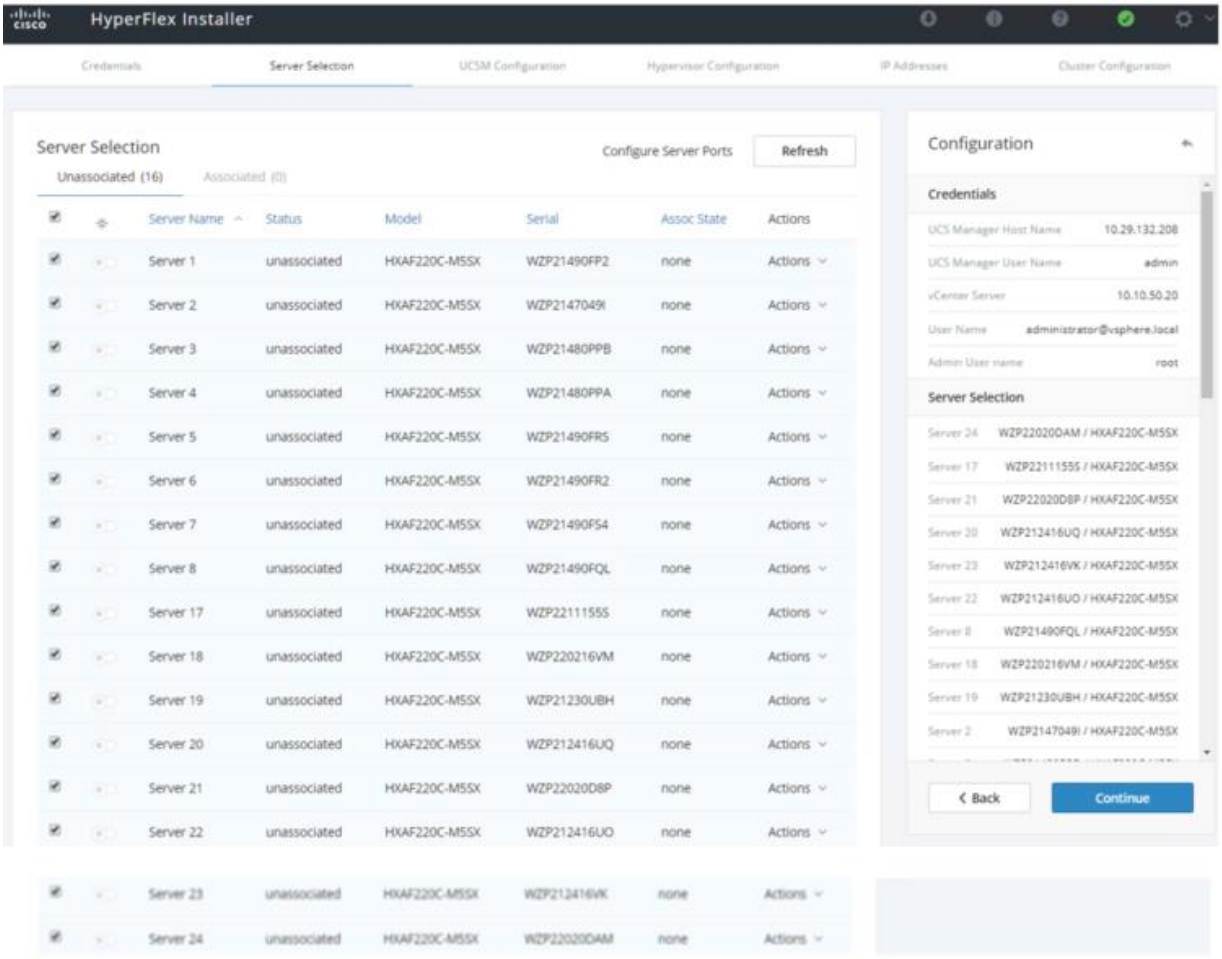

The required server ports can be configured from Installer workflow but it will extend the time to complete server discovery. Therefore, we recommend configuring the server ports and complete HX node discovery in Cisco UCS Manager as described in the Pre-requisites section above prior starting workflow for HyperFlex installer.

Configure Server Ports (Optional)

If you choose to allow the installer to configure the server ports, complete the following steps:

1. Click Configure Server Ports at the top right corner of the Server Selection window.
2. Provide the port numbers for each Fabric Interconnect in the form:

   A1/x-y,B1/x-y where A1 and B1 designate Fabric Interconnect A and B and where x=starting port number and y=ending port number on each Fabric Interconnect.
3. Click Configure.

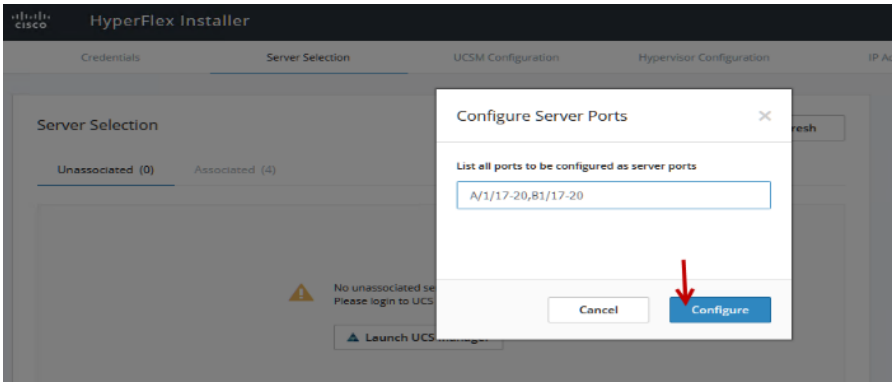

4. Enter the Details for the Cisco UCS Manager Configuration:

   a. nter VLAN ID for hx-inband-mgmt, hx-storage-data, hx-vmotion, vm-network.

   b. MAC Pool Prefix: The prefix to use for each HX MAC address pool. Please select a prefix that does not conflict with any other MAC address pool across all Cisco UCS domains.

   c. The blocks in the MAC address pool will have the following format:

      – ${prefix}:${fabric_id}${vnic_id}:{service_profile_id}

      – The first three bytes should always be “00:25:B5”.

The first three bytes should always be “00:25:B5.”

5. Enter range of IP address to create a block of IP addresses for external management and access to CIMC/KVM.

6. Cisco UCS firmware version is set to 4.0 (1b) which is the required Cisco UCS Manager release for HyperFlex v3.5(1a) installation.

7. Enter HyperFlex cluster name.

8. Enter Org name to be created in Cisco UCS Manager.

9. Click Continue.

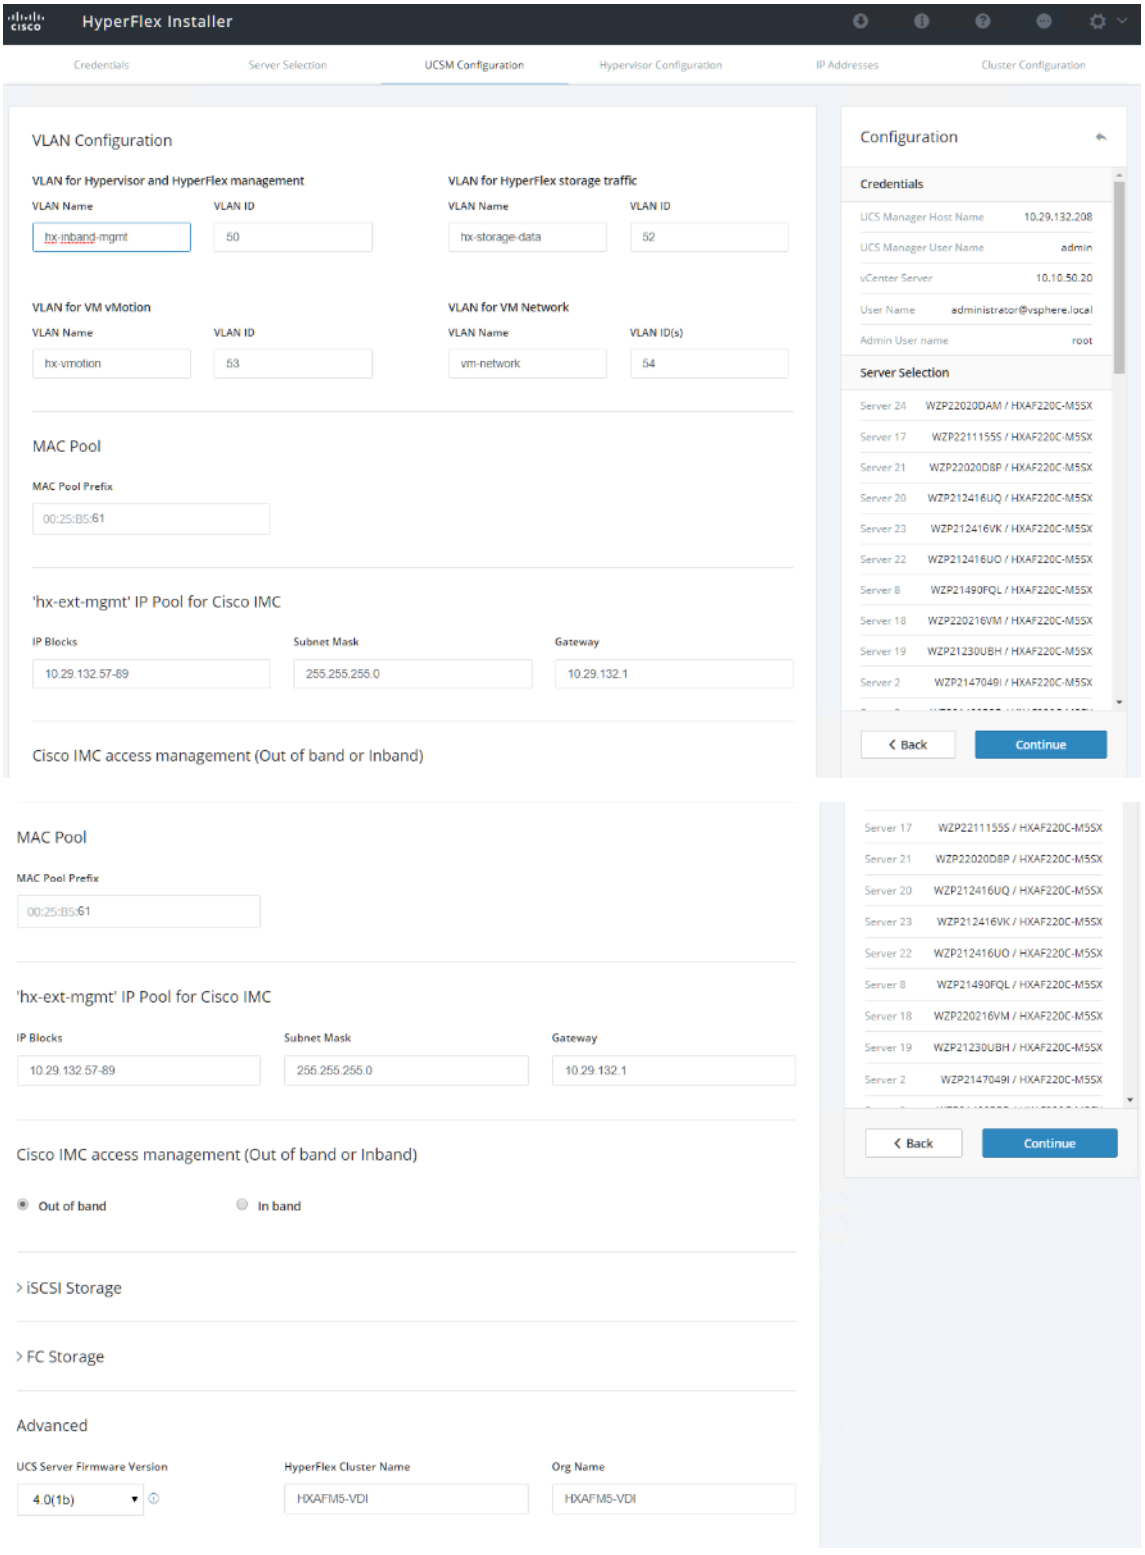

Configure Hypervisor Settings

To configure the Hypervisor settings, complete the following steps:

1. In the Configure common Hypervisor Settings section, enter:

- Subnet Mask
- Gateway
- DNS server(s)

2. In the Hypervisor Settings section:
   - Select check box Make IP Address and Hostnames Sequential if they are following in sequence.
   - Provide the starting IP Address.
   - Provide the starting Host Name or enter Static IP address and Host Names manually for each node
3. Click Continue.

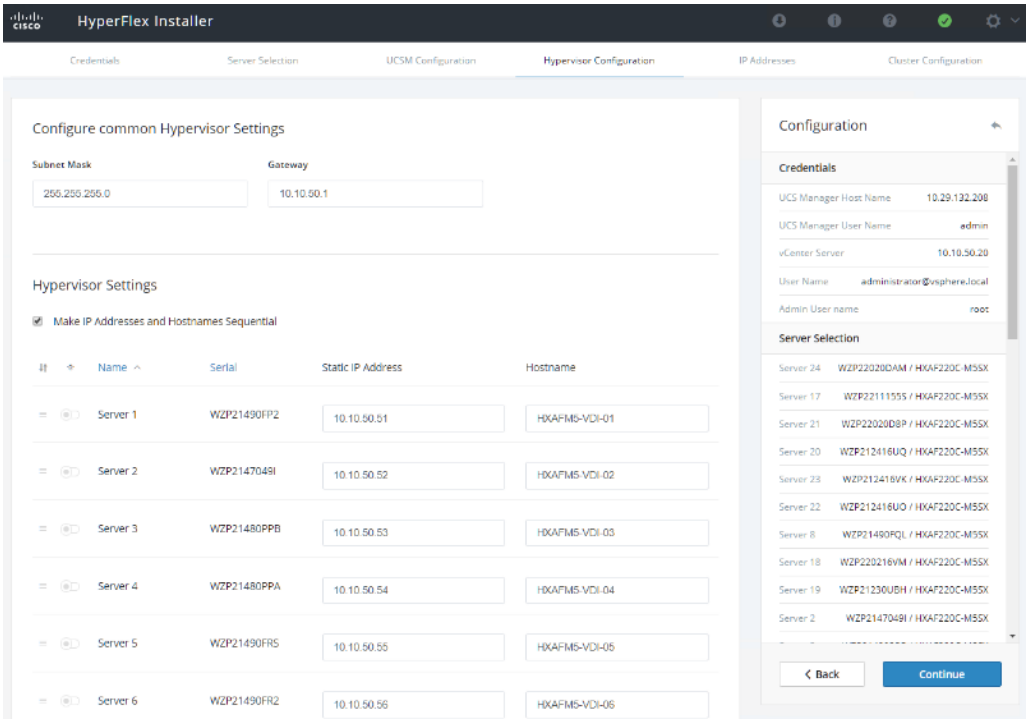

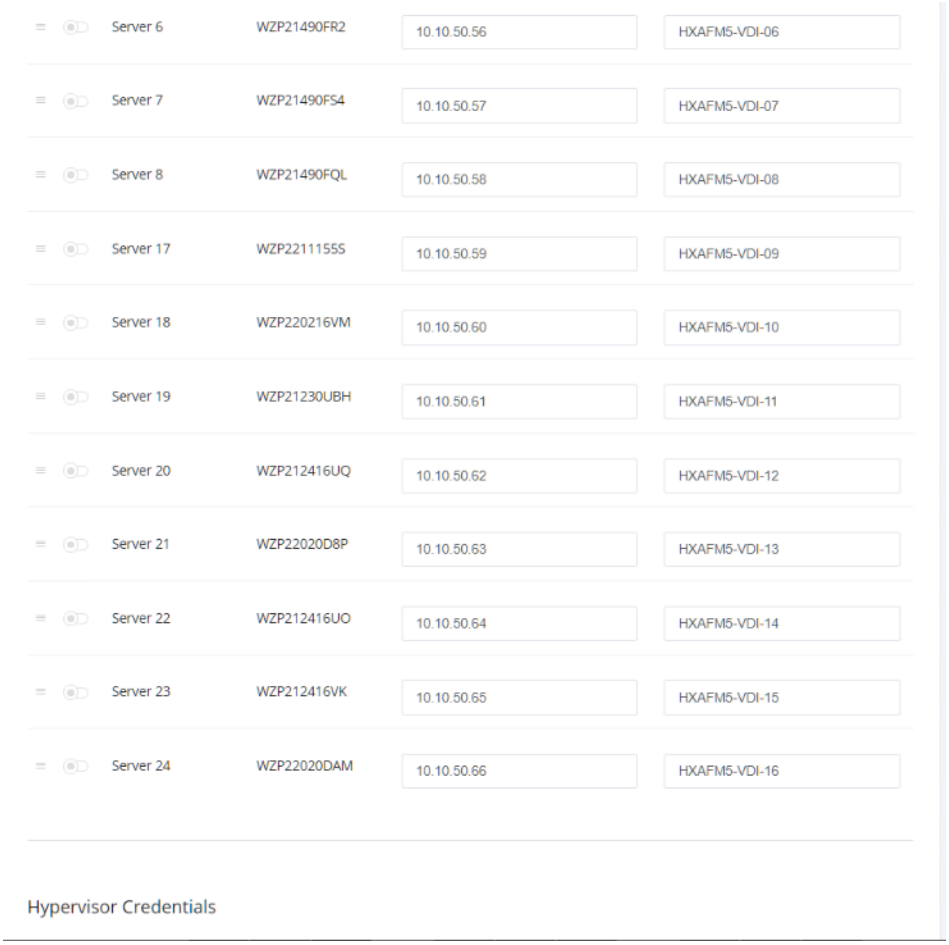

IP Addresses

To add the IP addresses, complete the following steps:

When the IP Addresses page appears, the hypervisor IP address for each node that was configured in the Hypervisor Configuration tab, appears under the Management Hypervisor column.

Three additional columns appear on this page:

- Storage Controller/Management
- Hypervisor/Data
- Storage Controller/Data

The Data network IP addresses are for vmkernel addresses for storage access by the hypervisor and storage controller virtual machine.

1. On the IP Addresses page, check the box Make IP Addresses Sequential or enter the IP address manually for each node for the following requested values:
   - Storage Controller/Management
   - Hypervisor/Data

- Storage Controller/Data

2. Enter subnet and gateway details for the Management and Data subnets configured.
3. Click Continue to proceed.

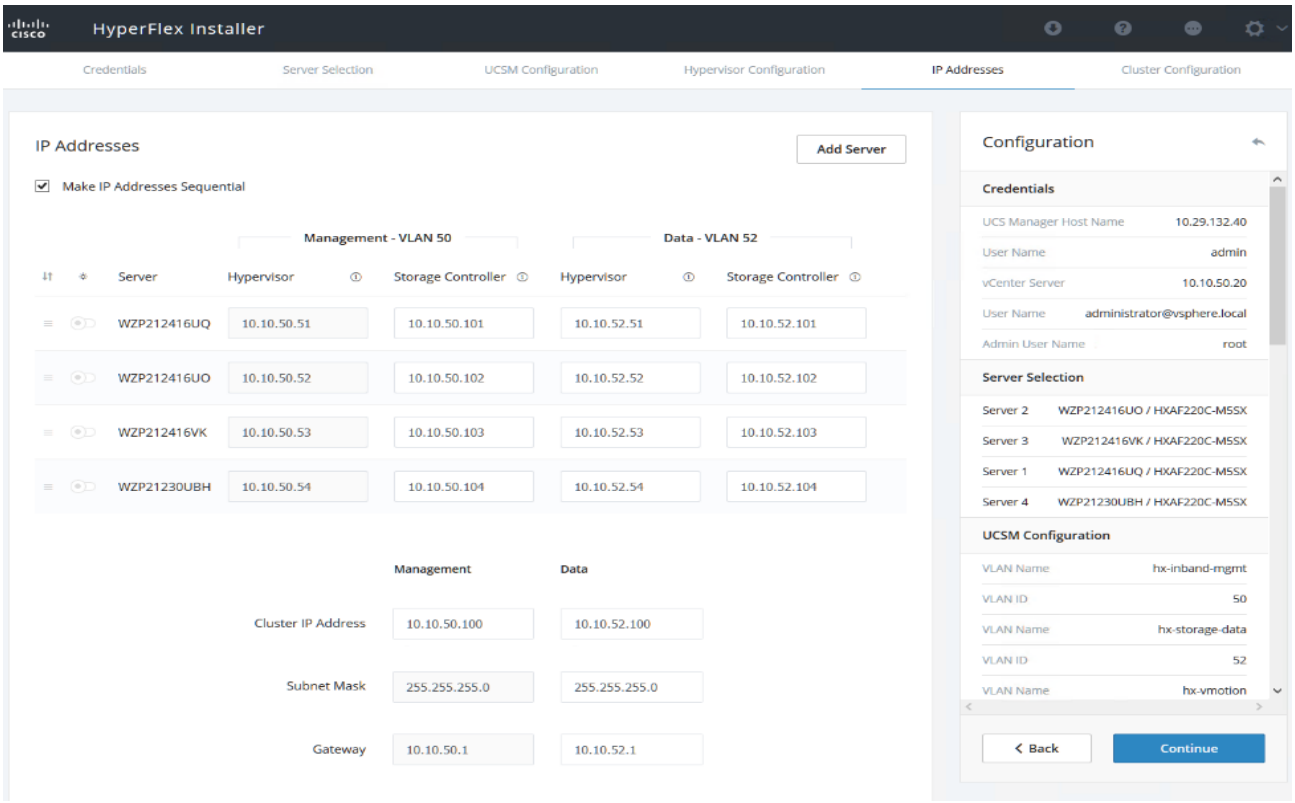

4. On the Cluster Configuration page, enter the following:
   - Cluster Name
   - Cluster management IP address
   - Cluster data IP Address
   - Set Replication Factor: 2 or 3
   - Controller virtual machine password
   - vCenter configuration
     - vCenter Datacenter name
     - vCenter Cluster name
   - System Services
     - DNS Server(s)
     - NTP Server(s)
     - Time Zone

- Auto Support
  - Click the check box for Enable Auto Support
  - Mail Server
  - Mail Sender
  - ASUP Recipient(s)
- Advanced Networking
  - Management vSwitch
  - Data vSwitch
- Advanced Configuration
  - Click the check box to Optimize for VDI only deployment
  - Enable jumbo Frames on Data Network
  - Clean up disk partitions (optional)
    - vCenter Single-Sign-On server

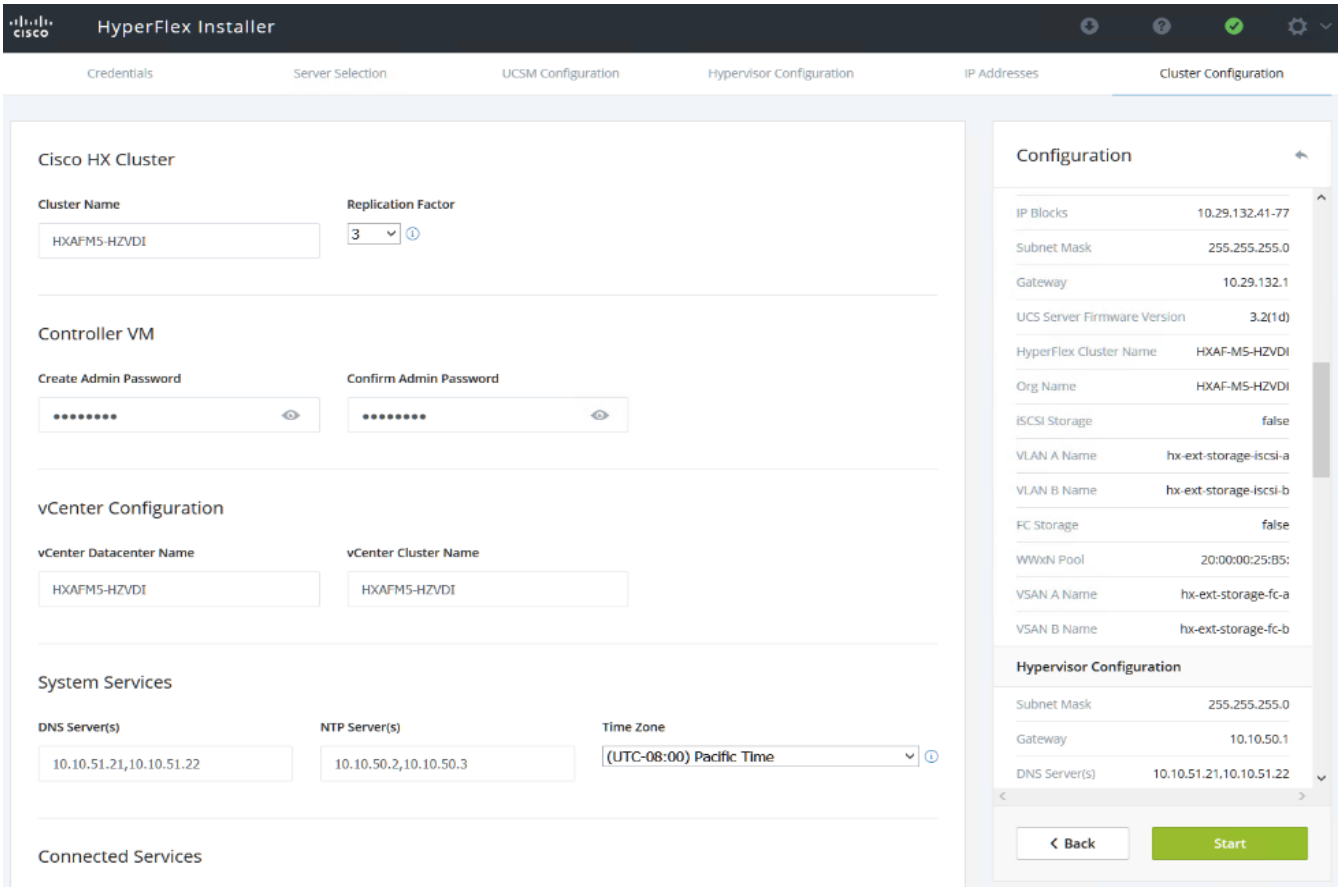

- vCenter Single-Sign-On server

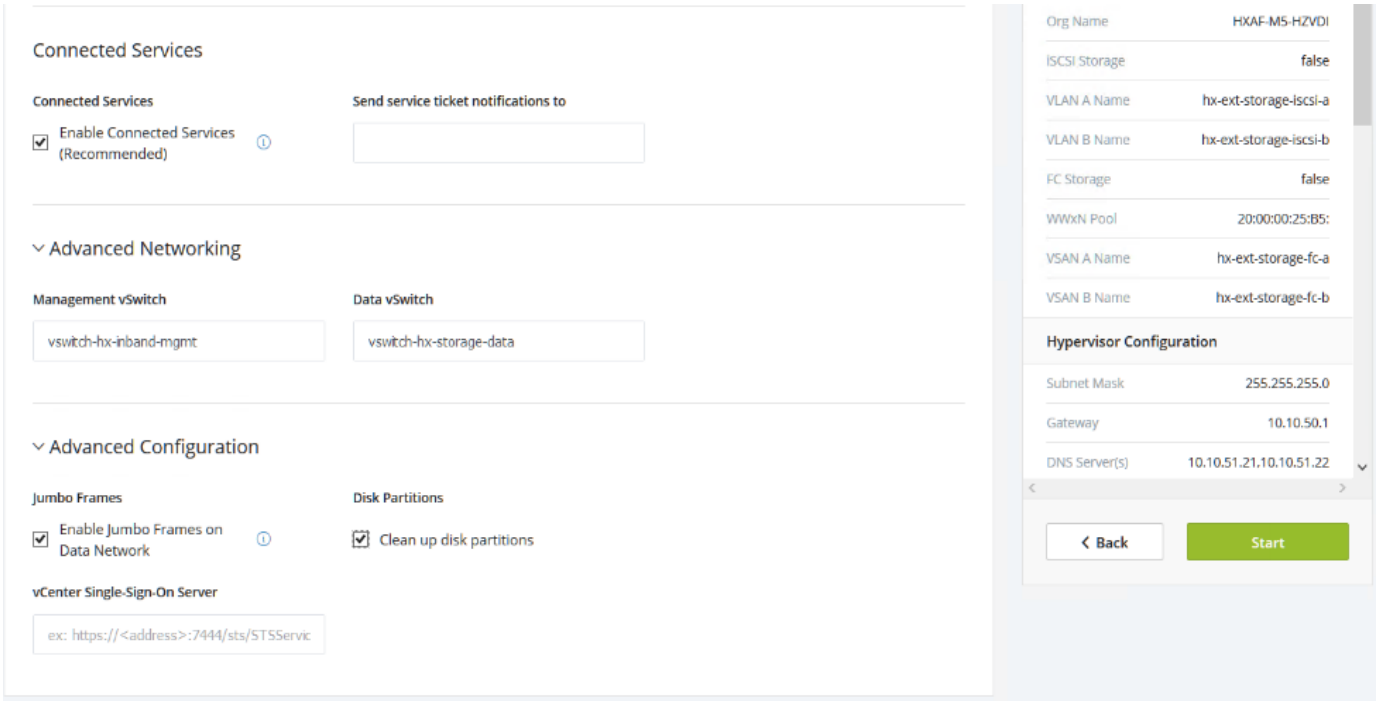

5. The configuration details can be exported to a JSON file by clicking the down arrow icon in the top right corner of the Web browser page as shown in the screenshot below.

6. Configuration details can be reviewed on Configuration page on right side section. Verify entered details for IP address entered in Credentials page, server selection for cluster deployment and creation workflow, Cisco UCS Manager configuration, Hypervisor Configuration, IP addresses.

7. Click Start after verifying details.

When the installation workflow begins, it will go through the Cisco UCS Manager validation.

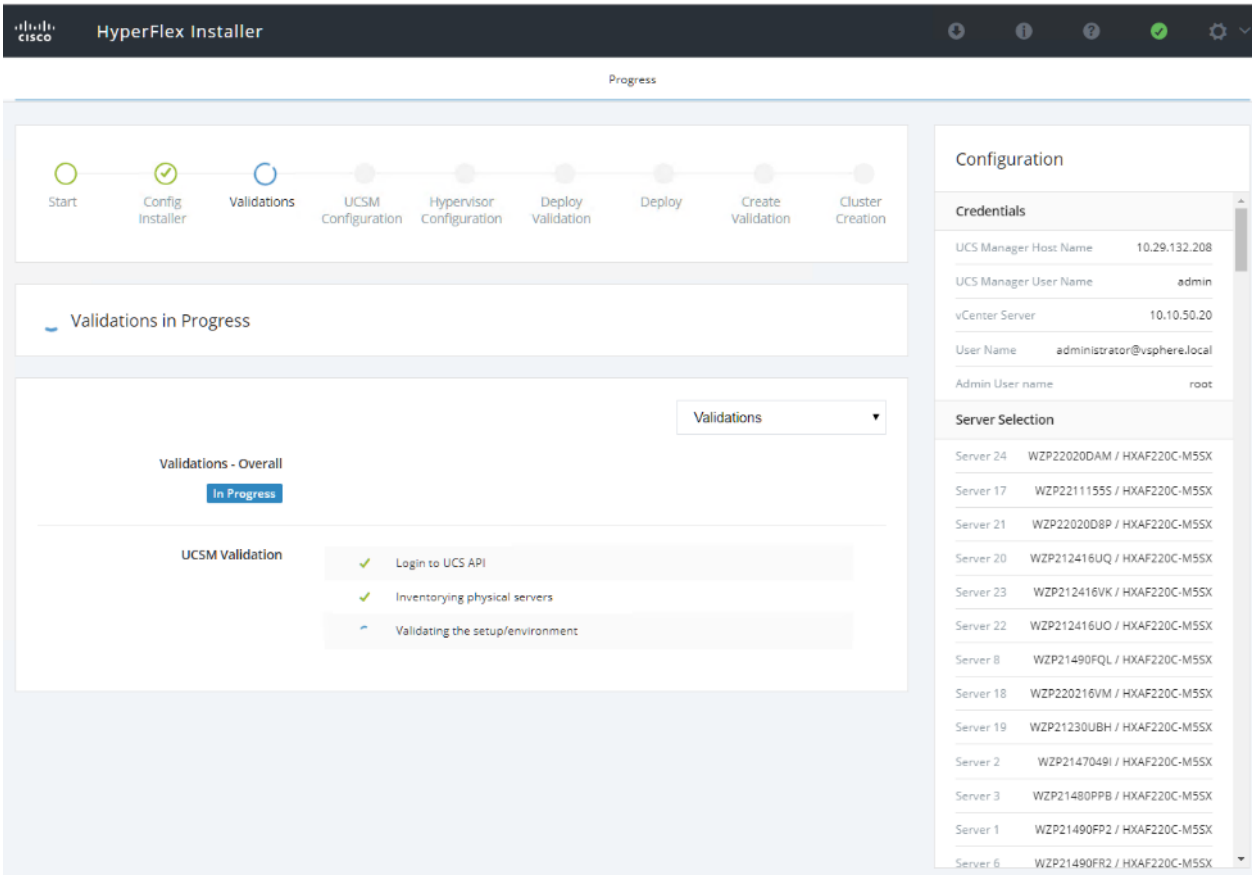
HyperFlex Installer
Progress
Start
Config Installer
Validations
UCSM Configuration
Hypervisor Configuration
Deploy Validation
Deploy
Create Validation
Cluster Creation
Validations in Progress
Validations
Validations - Overall
In Progress
UCSM Validation
Login to UCS API
Inventorying physical servers
Validating the setup/environment
Configuration
Credentials
UCS Manager Host Name 10.29.132.208
UCS Manager User Name admin
vCenter Server 10.10.50.20
User Name administrator@vsphere.local
Admin User name root
Server Selection
Server 24 WZP22020DAM / HXAF220C-M5SX
Server 17 WZP2211155S / HXAF220C-M5SX
Server 21 WZP22020D8P / HXAF220C-M5SX
Server 20 WZP212416UQ / HXAF220C-M5SX
Server 23 WZP212416VK / HXAF220C-M5SX
Server 22 WZP212416UO / HXAF220C-M5SX
Server 8 WZP21490FQL / HXAF220C-M5SX
Server 18 WZP220216VM / HXAF220C-M5SX
Server 19 WZP21230UBH / HXAF220C-M5SX
Server 2 WZP2147049I / HXAF220C-M5SX
Server 3 WZP21480PPB / HXAF220C-M5SX
Server 1 WZP21490FP2 / HXAF220C-M5SX
Server 6 WZP21490FR2 / HXAF220C-M5SX

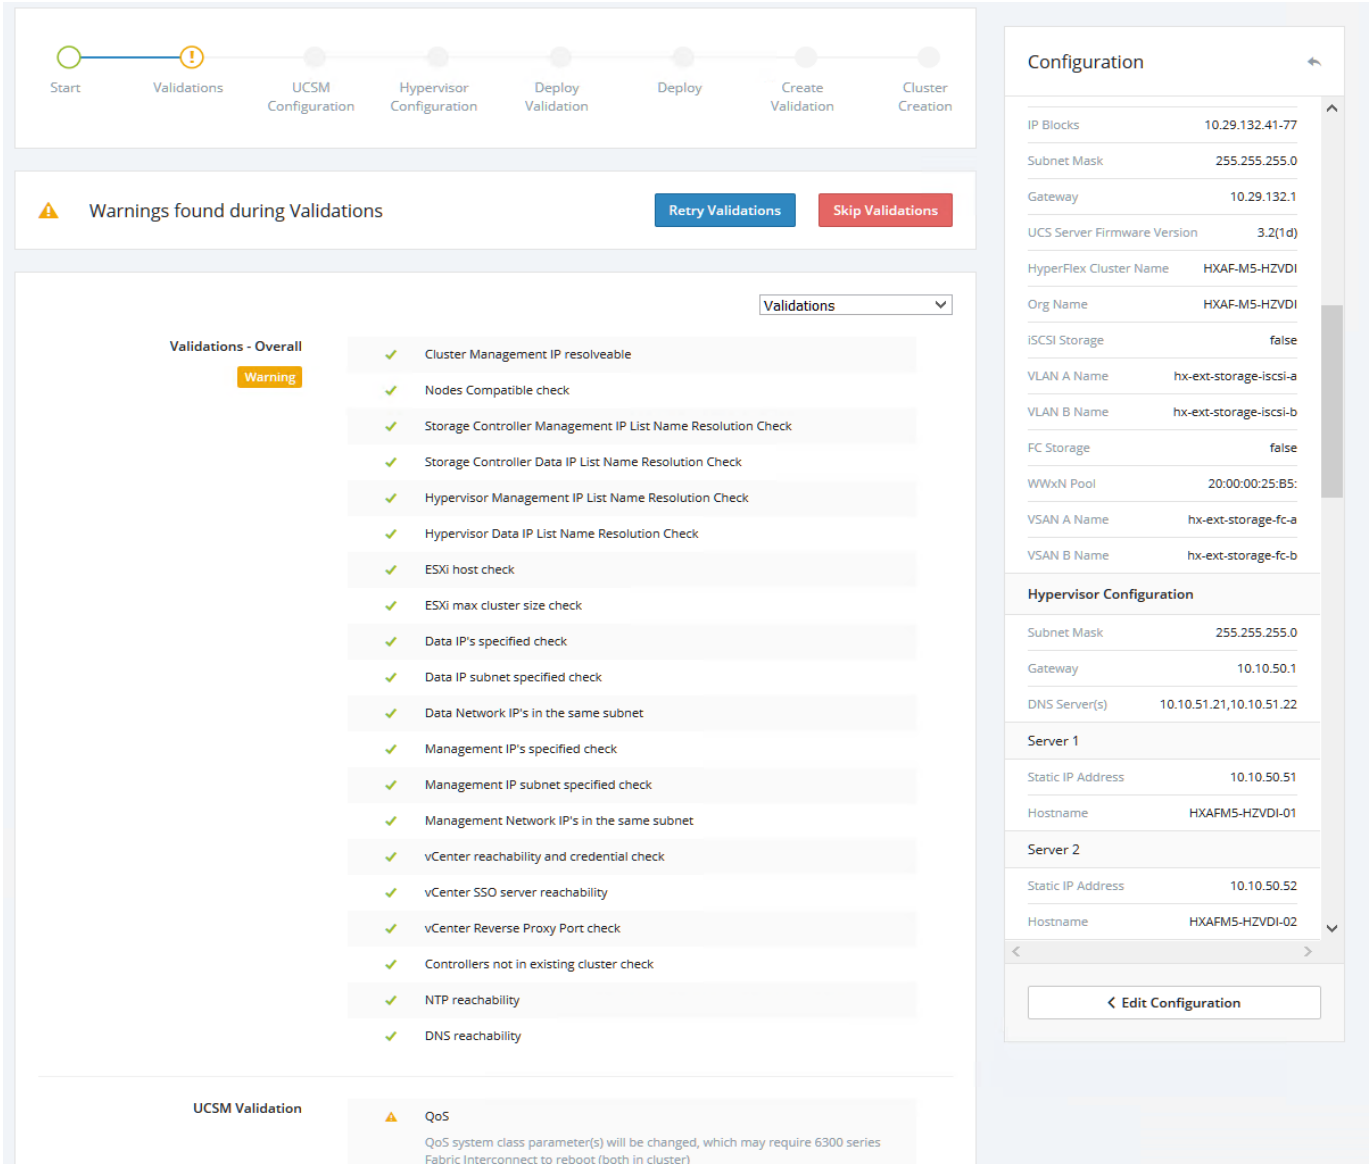

If QoS system class is not defined as per the requirement HyperFlex installer will go ahead and make required changes. There will be a warning generated accordingly in HyperFlex Installer workflow. For 6300 series Fabric Interconnect change in QoS system class requires reboot of FIs.

8. After a successful validation, the workflow continues with the Cisco UCS Manager configuration.

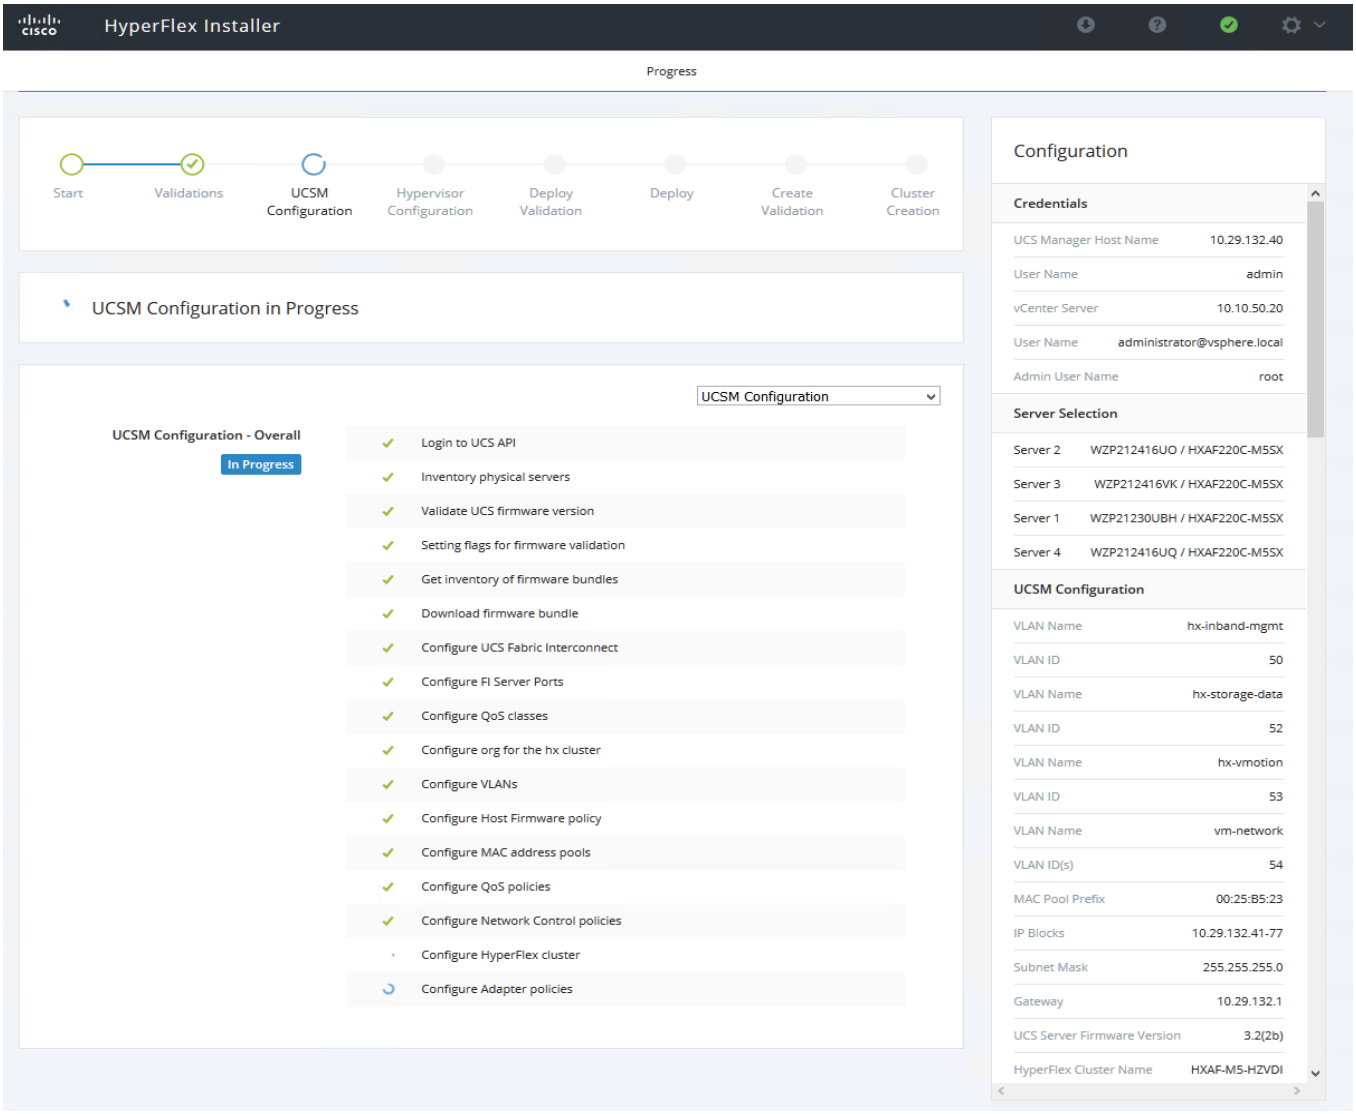

9. After a successful Cisco UCS Manager configuration, the installer proceeds with the Hypervisor configuration.

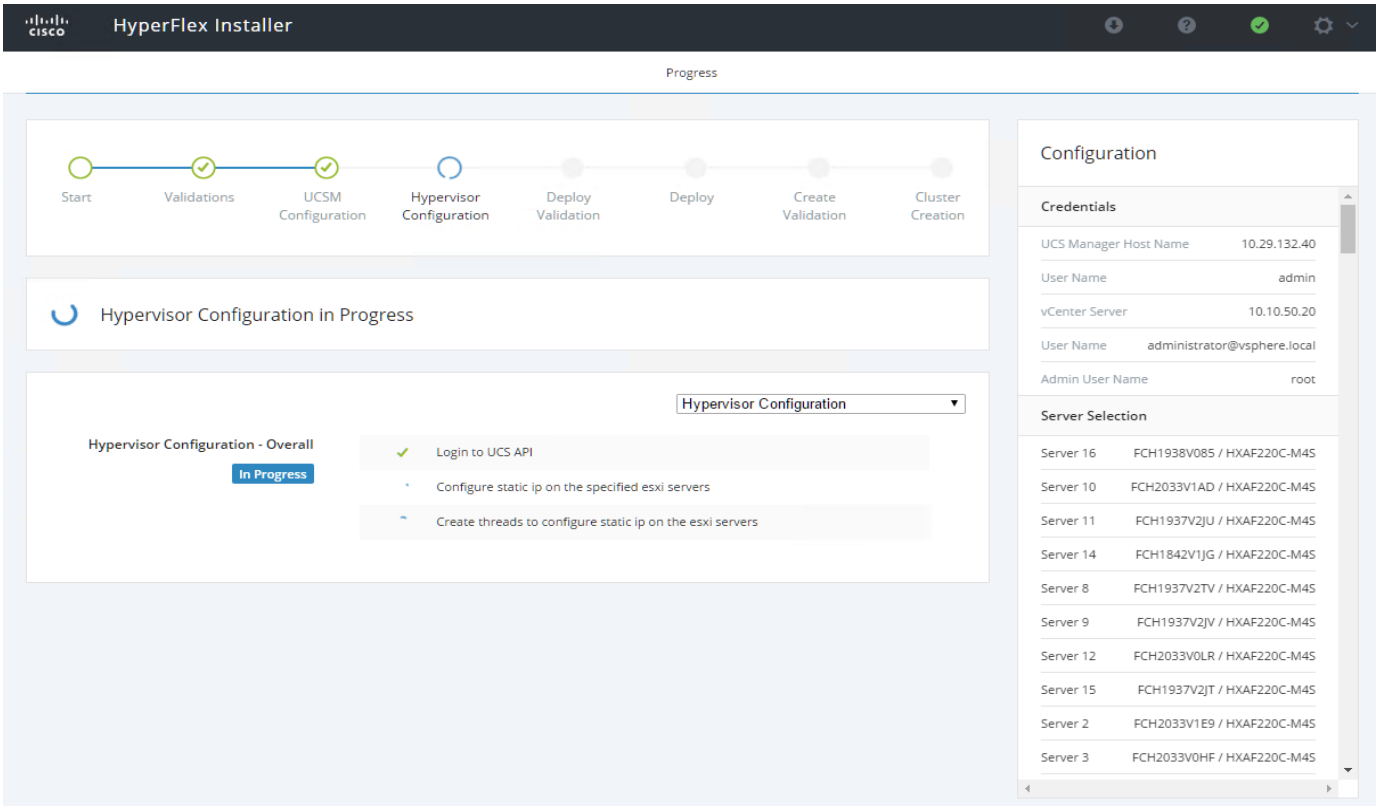

10. After a successful Hypervisor configuration, deploy validation task is performed which checks for required compoenent and accessibilty prior Deploy task is performed on Storage Controller virtual machine.

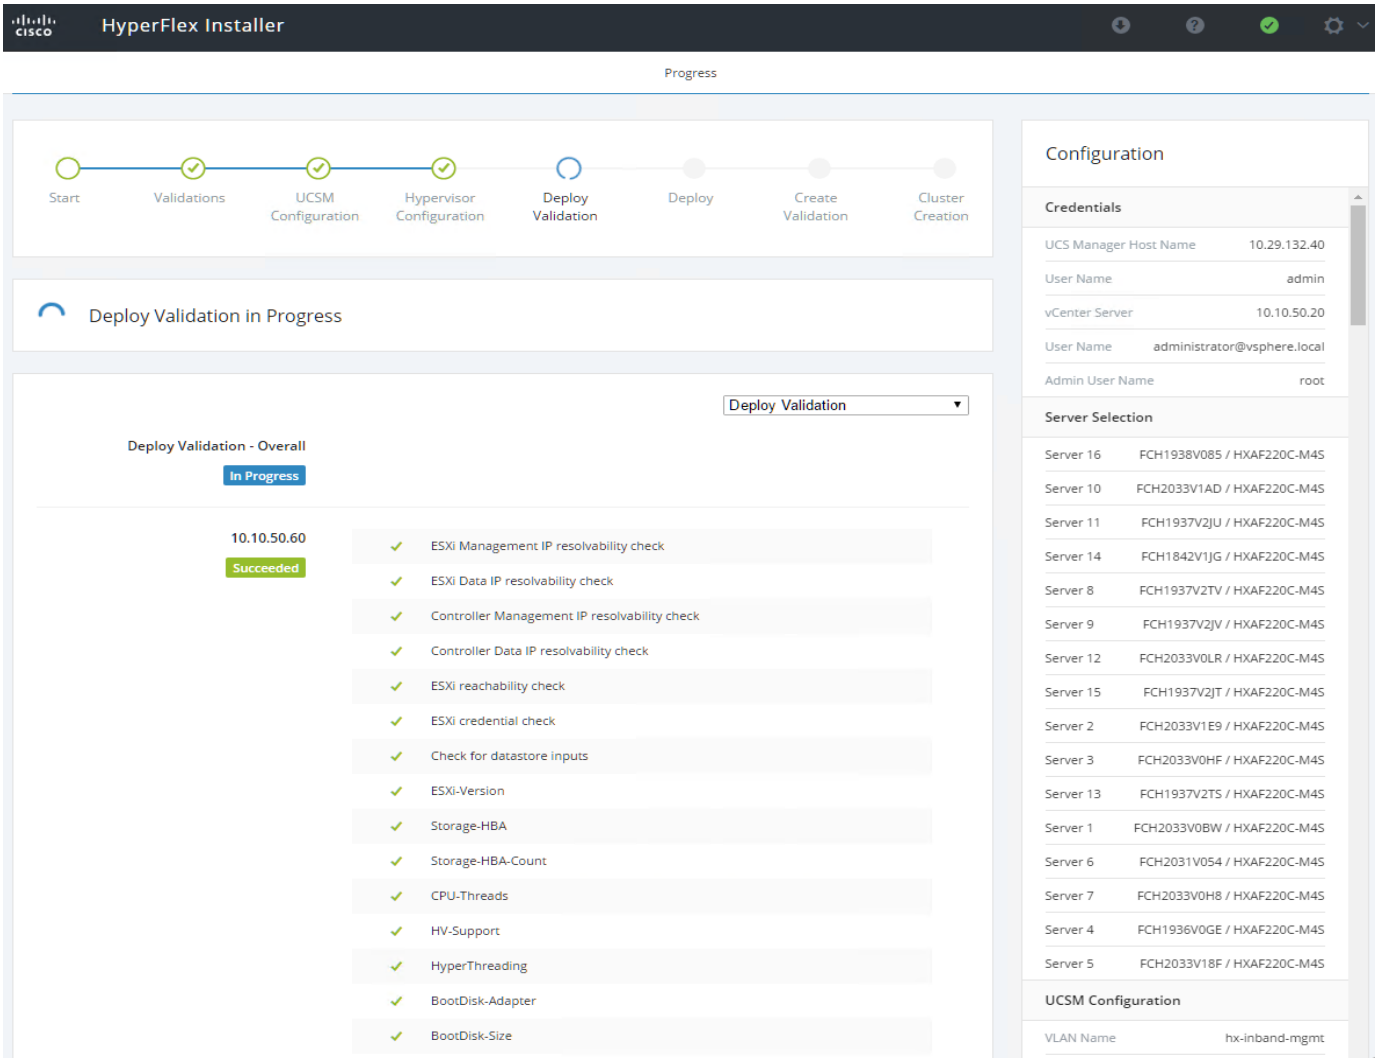

11. Installer performs deployment task after successfully validating Hypervisor configuration.

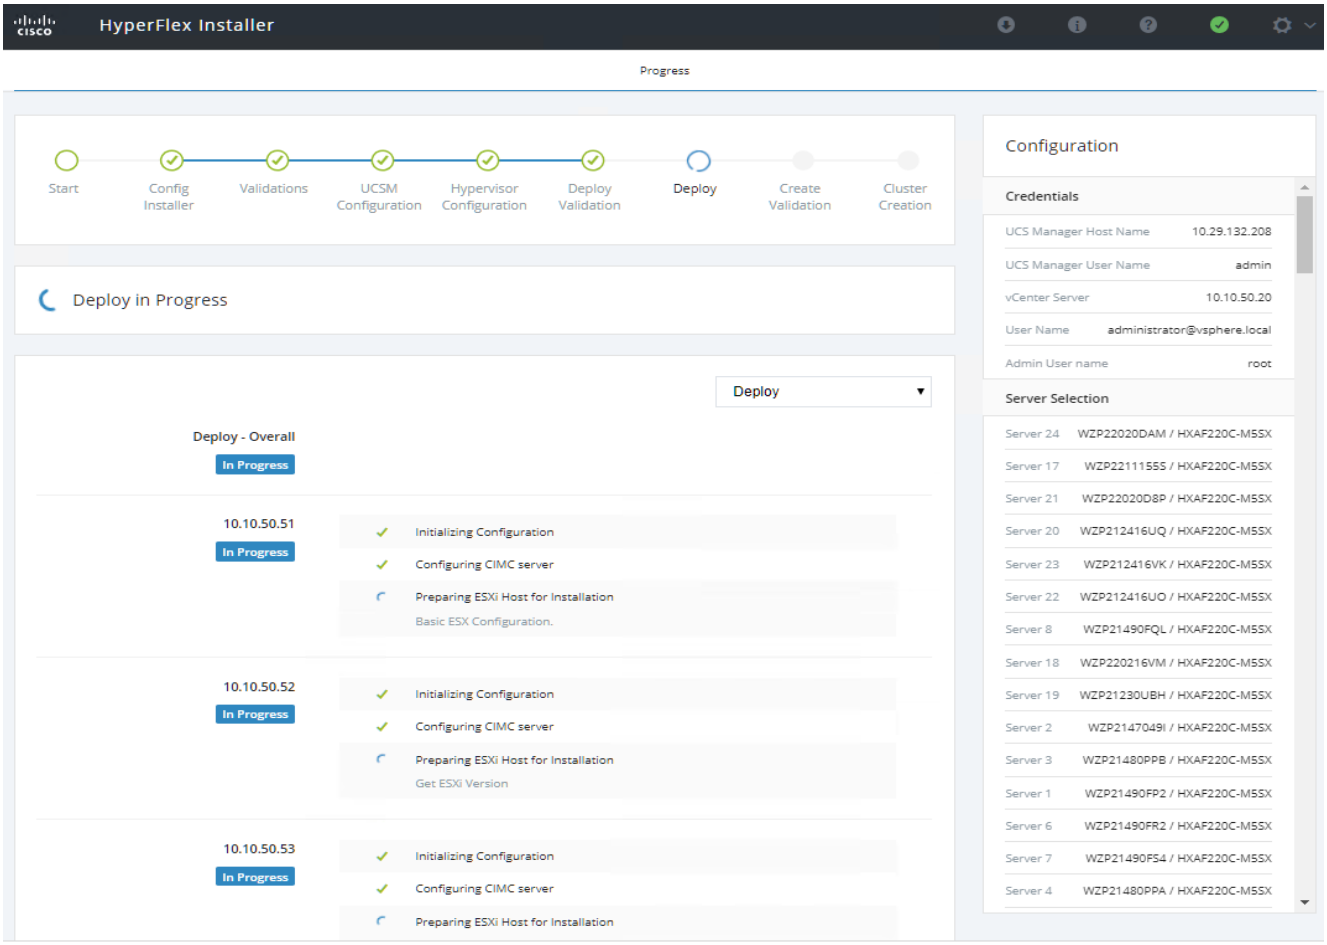

12. After a successful deployment of the ESXi hosts configuration, the Controller virtual machine software components for HyperFlex installer checks for validation prior to creating the cluster.

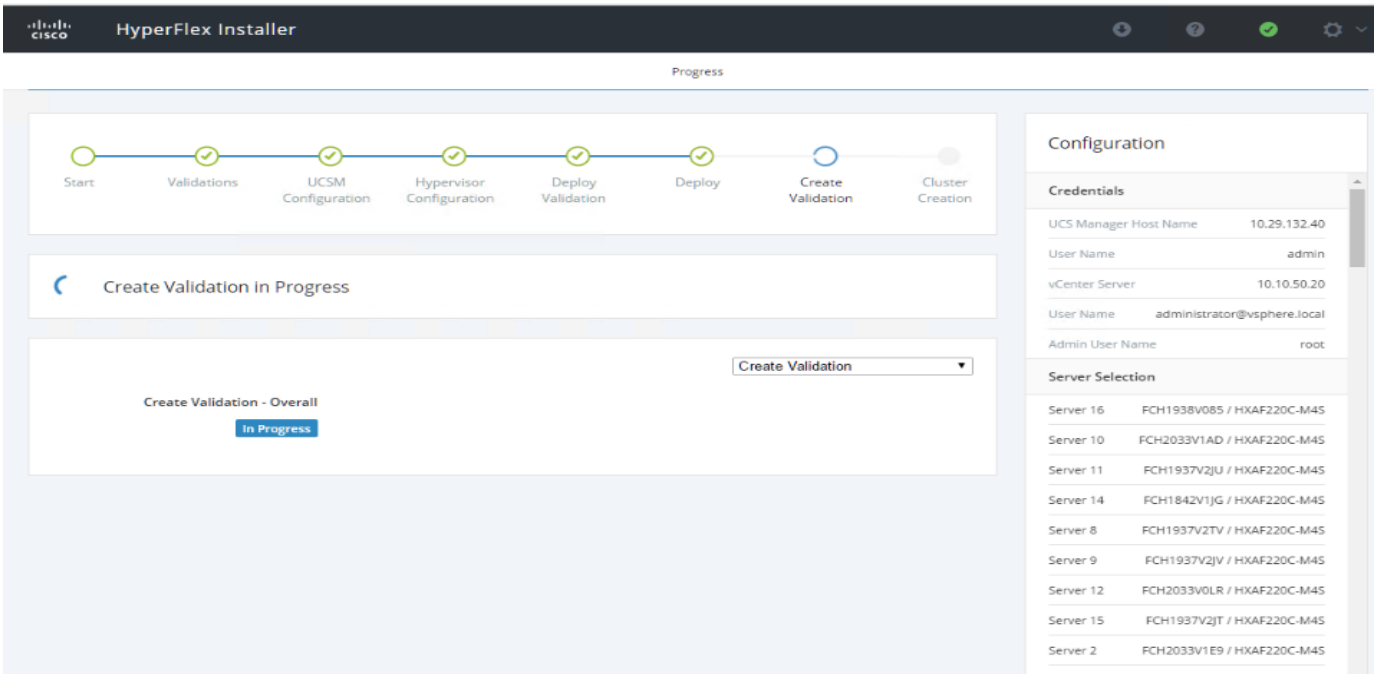

13. After a successful validation, the installer creates and starts the HyperFlex cluster service.

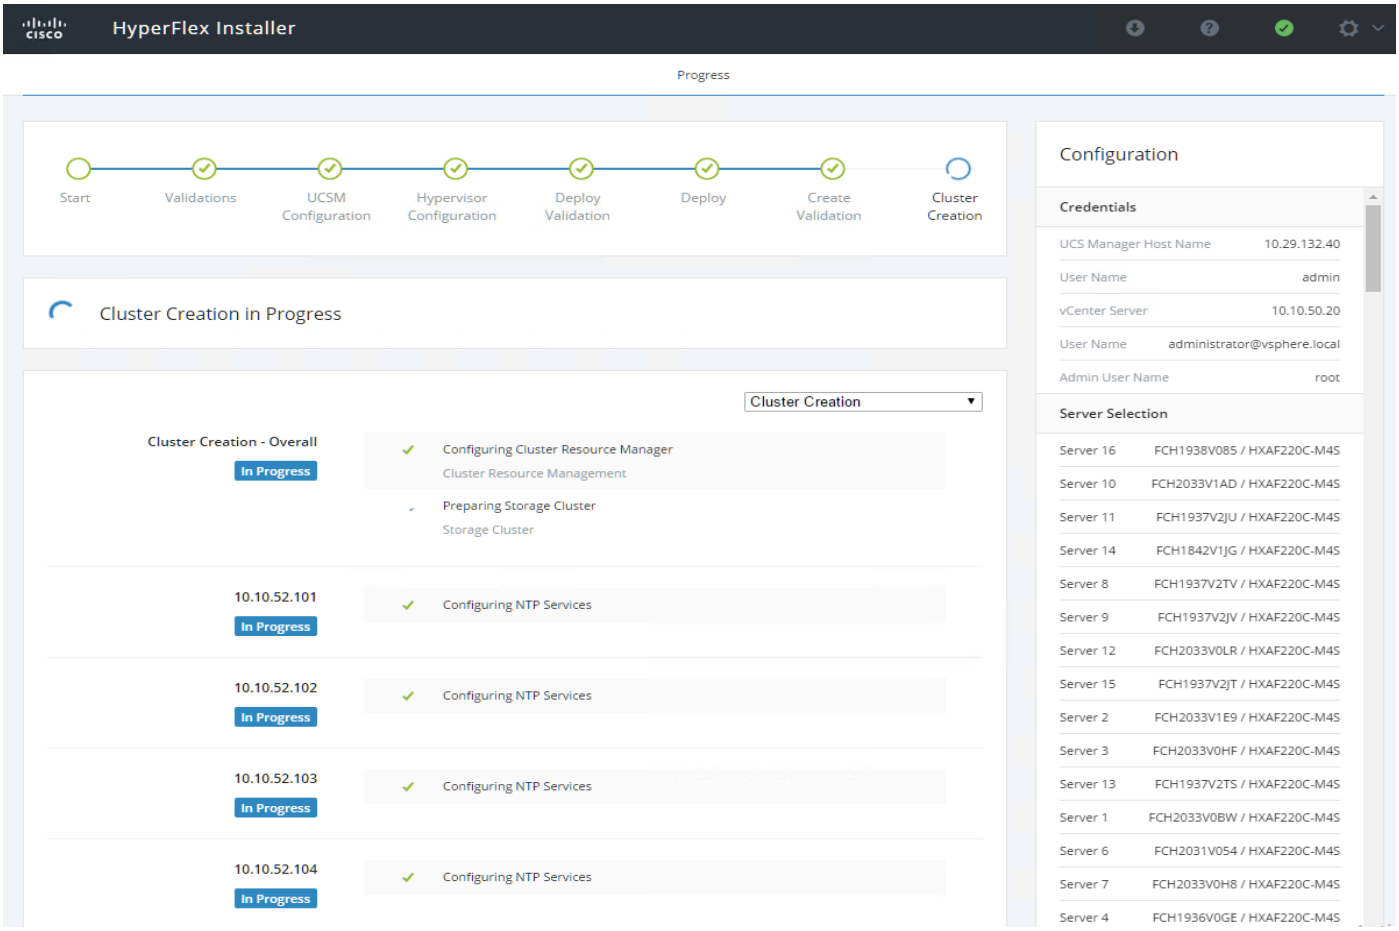

14. After a successful HyperFlex Installer virtual machine workflow completion, the installer GUI provides a summary of the cluster that has been created.

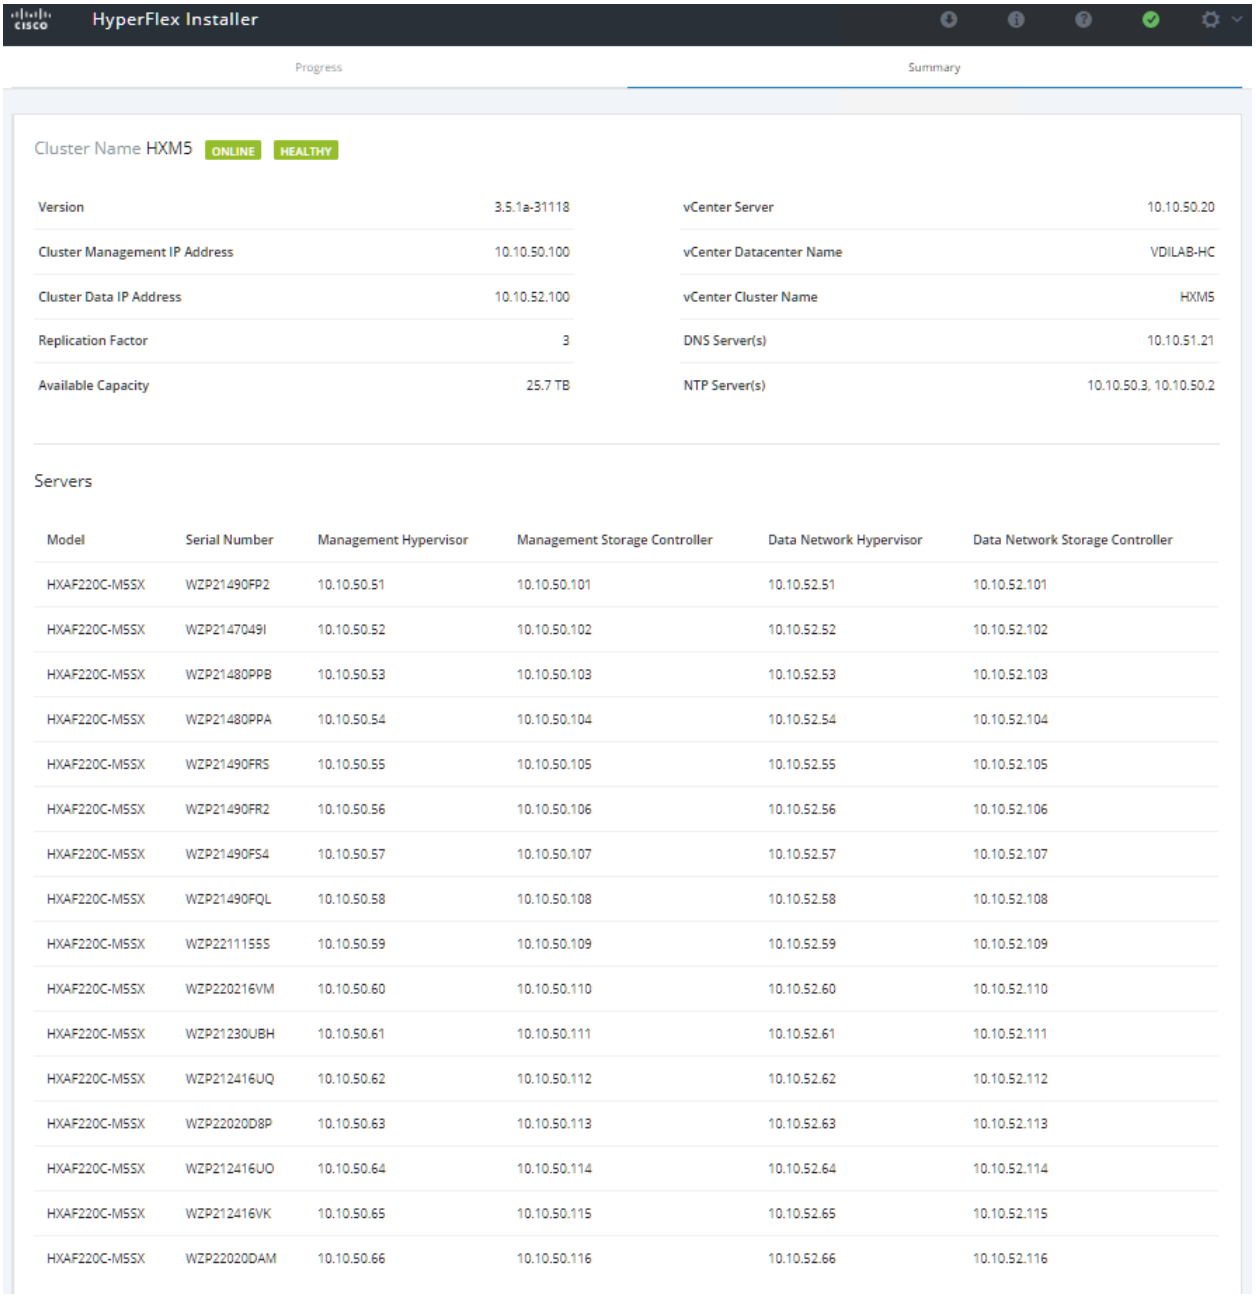

Model	Serial Number	Management Hypervisor	Management Storage Controller	Data Network Hypervisor	Data Network Storage Controller
HXAF220C-M5SX	WZP21490FP2	10.10.50.51	10.10.50.101	10.10.52.51	10.10.52.101
HXAF220C-M5SX	WZP2147049I	10.10.50.52	10.10.50.102	10.10.52.52	10.10.52.102
HXAF220C-M5SX	WZP21480PPB	10.10.50.53	10.10.50.103	10.10.52.53	10.10.52.103
HXAF220C-M5SX	WZP21480PPA	10.10.50.54	10.10.50.104	10.10.52.54	10.10.52.104
HXAF220C-M5SX	WZP21490FRS	10.10.50.55	10.10.50.105	10.10.52.55	10.10.52.105
HXAF220C-M5SX	WZP21490FR2	10.10.50.56	10.10.50.106	10.10.52.56	10.10.52.106
HXAF220C-M5SX	WZP21490FS4	10.10.50.57	10.10.50.107	10.10.52.57	10.10.52.107
HXAF220C-M5SX	WZP21490FQL	10.10.50.58	10.10.50.108	10.10.52.58	10.10.52.108
HXAF220C-M5SX	WZP22111555	10.10.50.59	10.10.50.109	10.10.52.59	10.10.52.109
HXAF220C-M5SX	WZP220216VM	10.10.50.60	10.10.50.110	10.10.52.60	10.10.52.110
HXAF220C-M5SX	WZP21230UBH	10.10.50.61	10.10.50.111	10.10.52.61	10.10.52.111
HXAF220C-M5SX	WZP212416UQ	10.10.50.62	10.10.50.112	10.10.52.62	10.10.52.112
HXAF220C-M5SX	WZP22020D8P	10.10.50.63	10.10.50.113	10.10.52.63	10.10.52.113
HXAF220C-M5SX	WZP212416UO	10.10.50.64	10.10.50.114	10.10.52.64	10.10.52.114
HXAF220C-M5SX	WZP212416VK	10.10.50.65	10.10.50.115	10.10.52.65	10.10.52.115
HXAF220C-M5SX	WZP22020DAM	10.10.50.66	10.10.50.116	10.10.52.66	10.10.52.116

Cisco HyperFlex Cluster Expansion

For this exercise, you will add the compute node workflow as part of the cluster expansion.

Prerequisites

Configure the service profile for compute-only nodes and install ESXi hypervisor.

To add the compute node workflow, complete the following steps:

1. Log into Cisco UCS Manger.
2. Under “hx-cluster” sub-organization:

a. In the existing vMedia policy “HyperFlex” add vMedia mount details to boot ESXi image from data platform installer virtual machine.

b. For Hostname/IP Address – Add IP address of data-platform installer virtual machine which can alsocommunicate with Cisco UCS Manager.

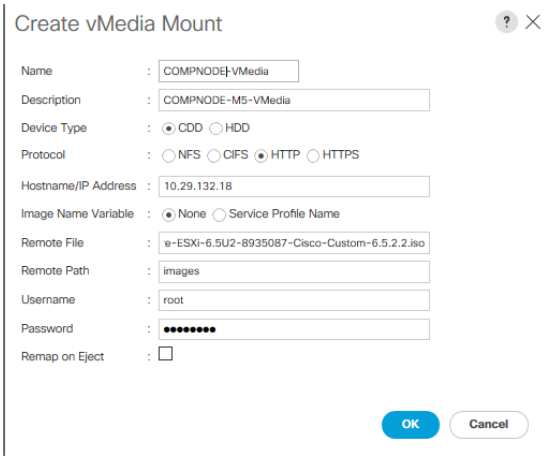

3. Change the existing service profile template to accommodate the new changes; install ESXi through vMedia policy.

4. In the existing service profile template “compute-nodes” select vMedia Policy tab.

5. Click Modify vMedia Policy.

6. From the drop-down list of vMedia Policy, select HyperFlex.

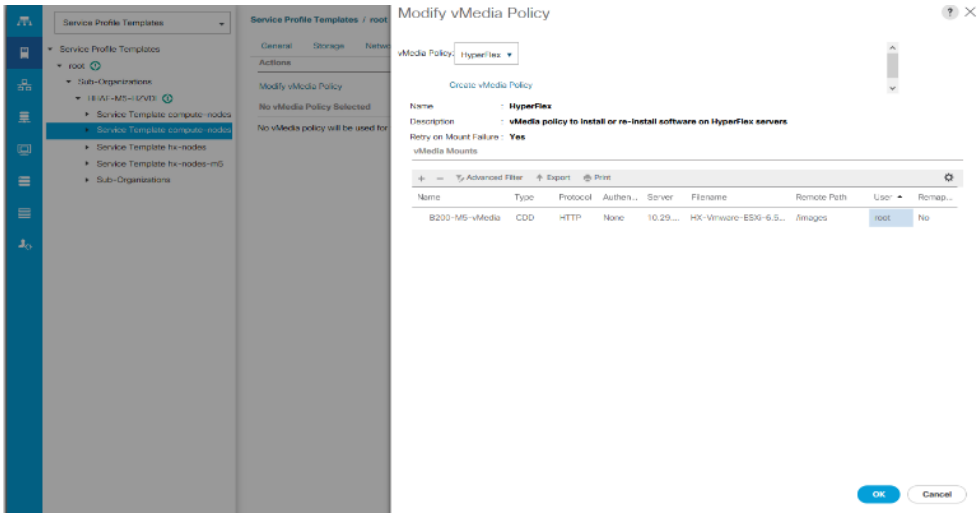

7. In the existing service profile template “compute-nodes” click Boot Order tab.

8. Click Modify Boot Policy.

9. From the drop-down list of Boot Policies, select HyperFlexInstall.

10. Save changes.

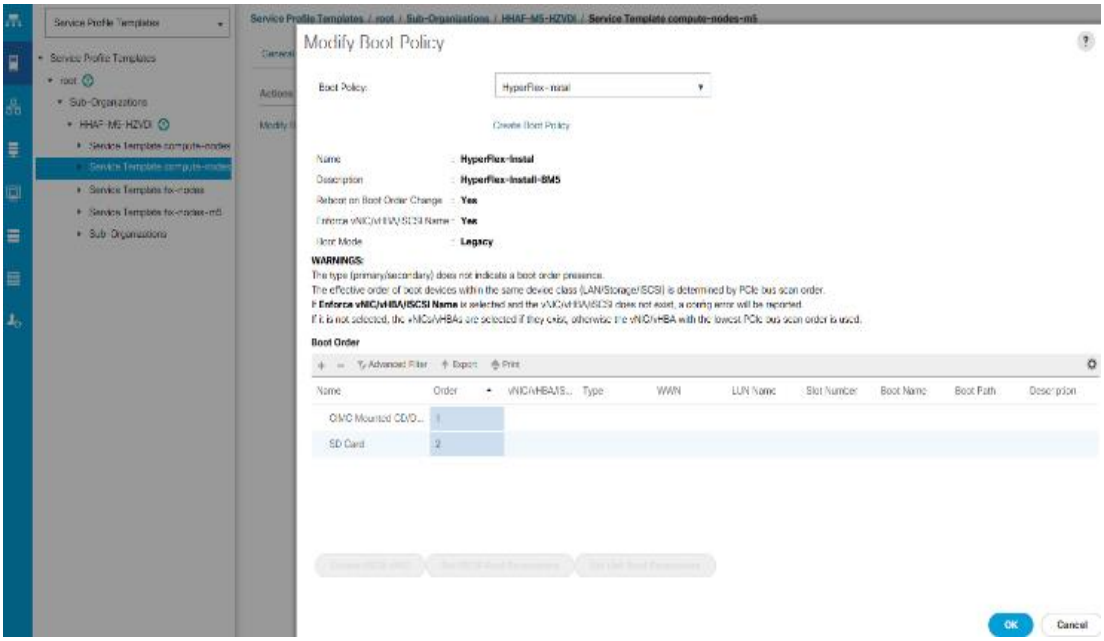

11. Create the service profile from the “compute-nodes” updating service profile template located in the HyperFlex cluster sub organization.

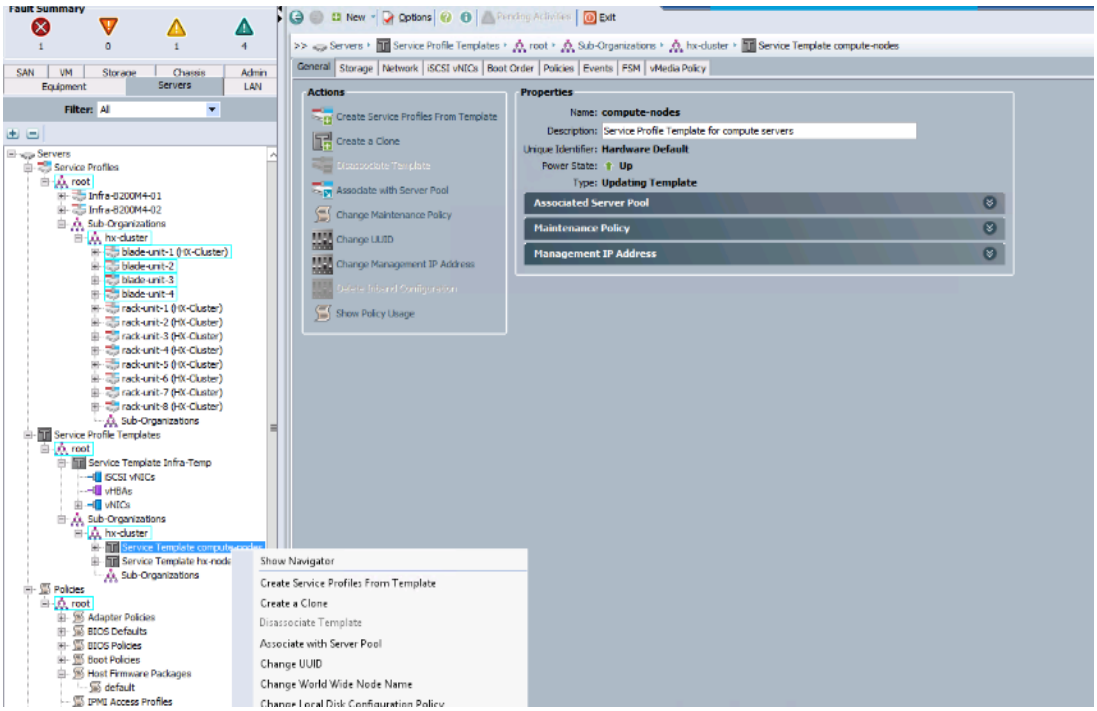

12. Add the Naming Prefix and Number of Instances to be created.

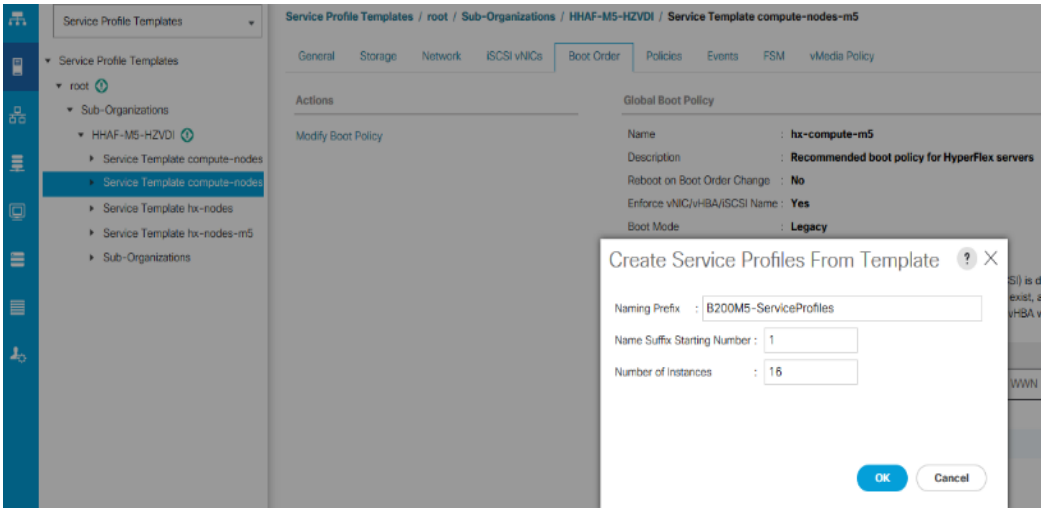

13. Click OK.

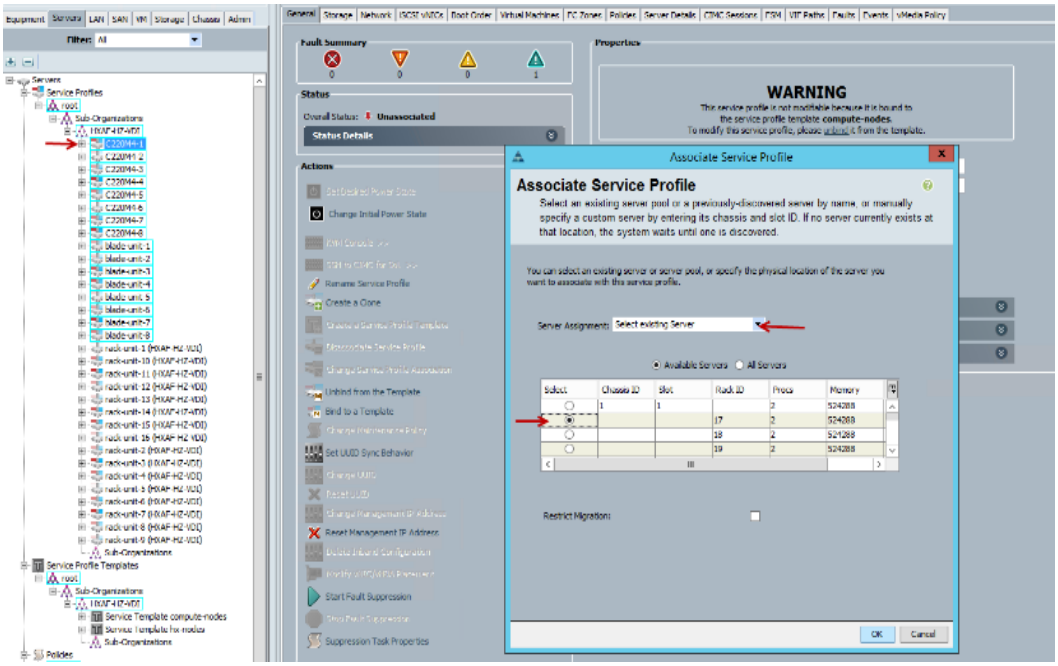

14. After the of ESXi install, assign the VLAN tag on the ESXi host; the static IP address configuration is located in the Configure Management Network section.

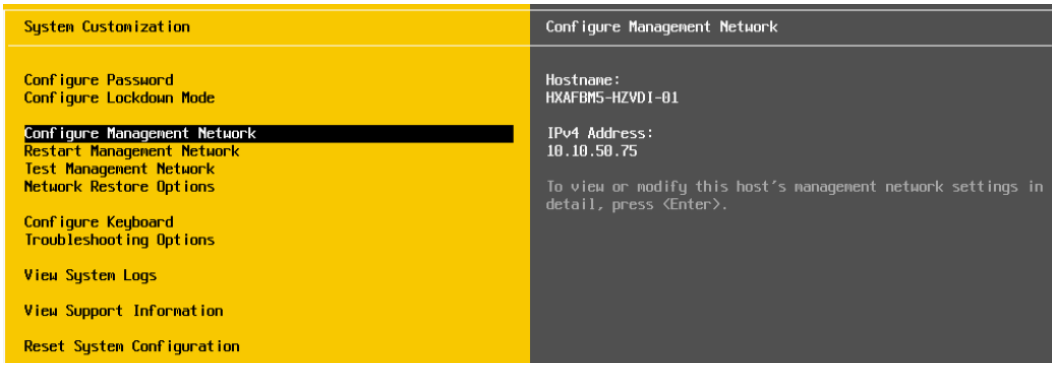

15. Log into the HyperFlex data platform installer WebUI. Click “I know what I'm doing, let me customize my workflow.”

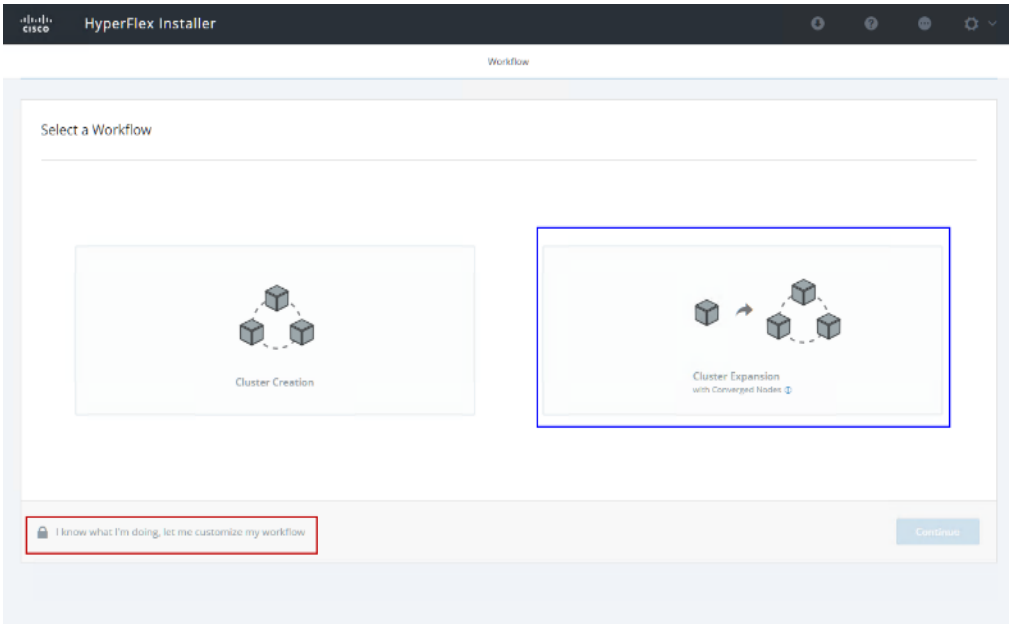

16. Select Deploy HX Software, Expand HX Cluster. Click Continue.

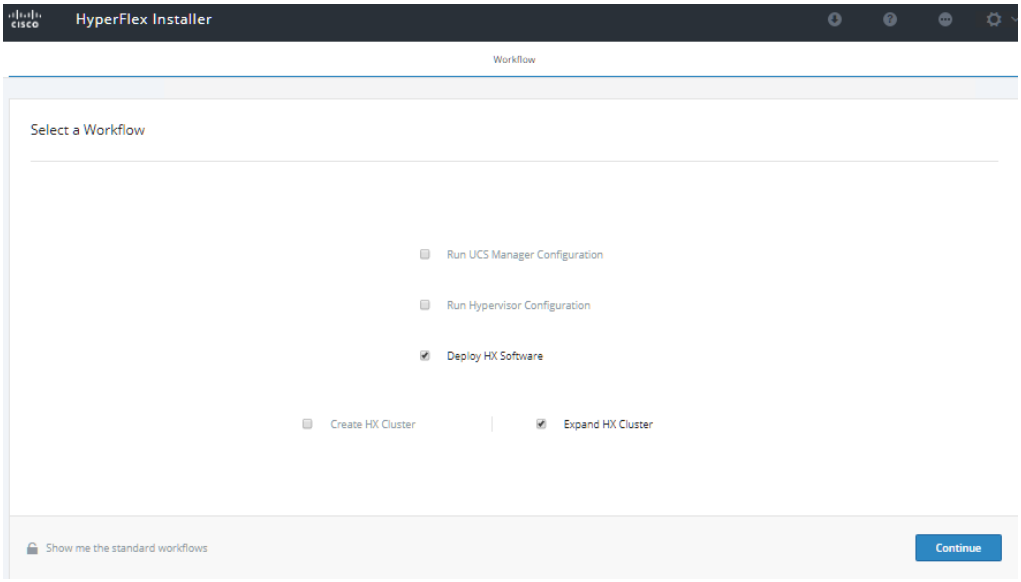

17. Enter the credentials for vCenter server, and ESXi. Click Continue.

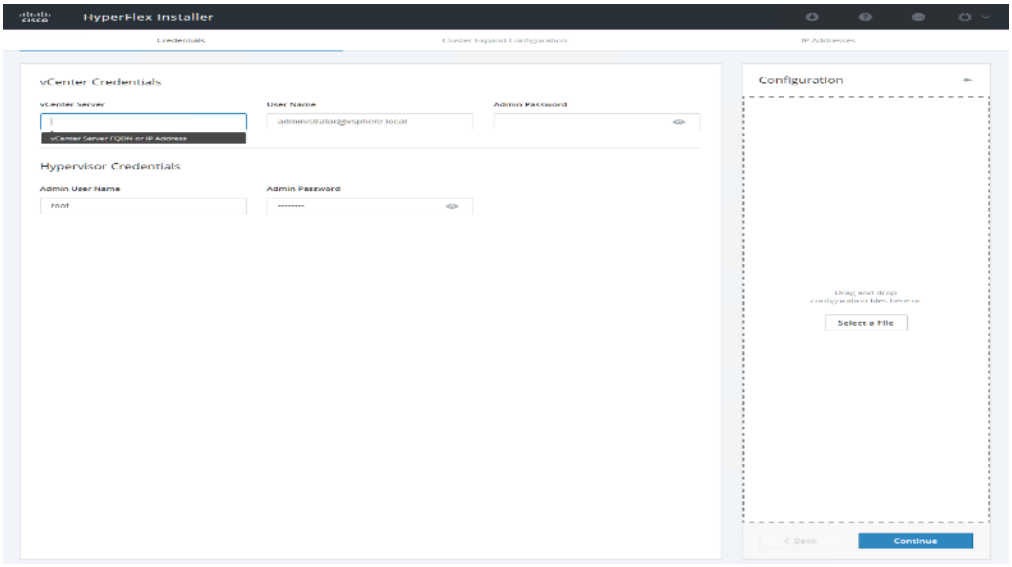

18. Select Cluster to expand, click Continue.

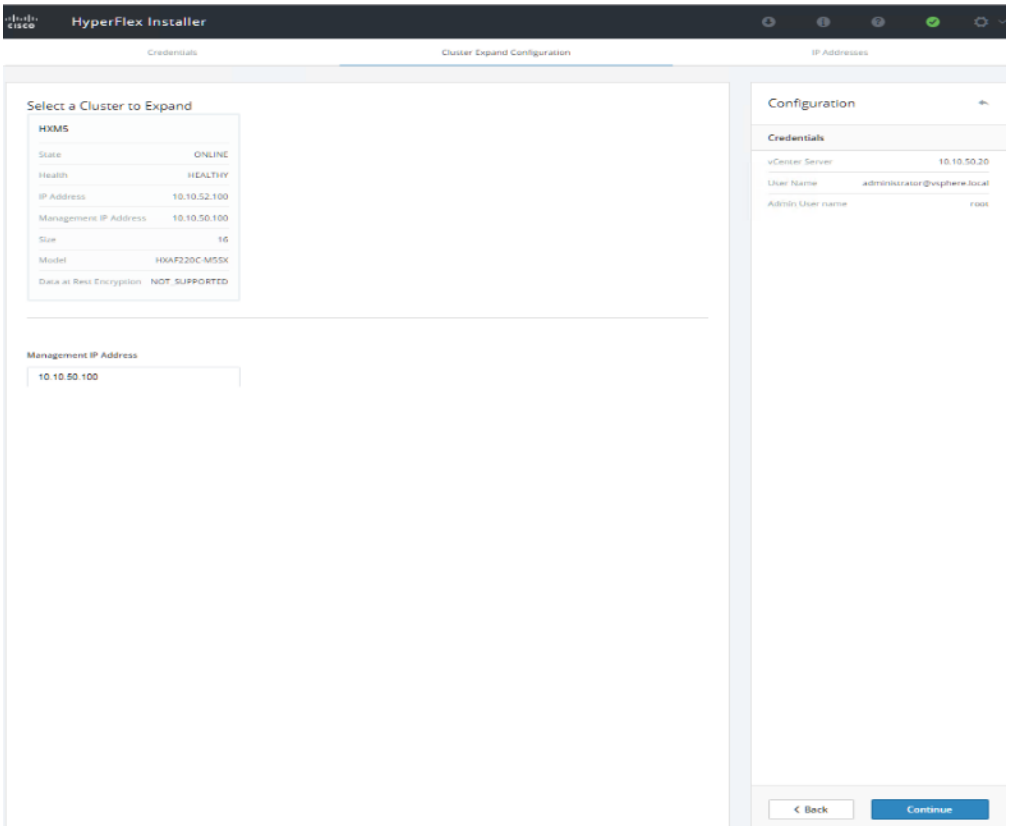

Since you are performing a compute-node only expansion, no servers report in to the Cisco UCS Manager configuration tab.

19. Click Add Compute Server tab for N number of compute-only node expansion to existing HyperFlex cluster. Provide Hypervisor Management IP address and vmkernel IP address to access storage cluster. Click Continue.

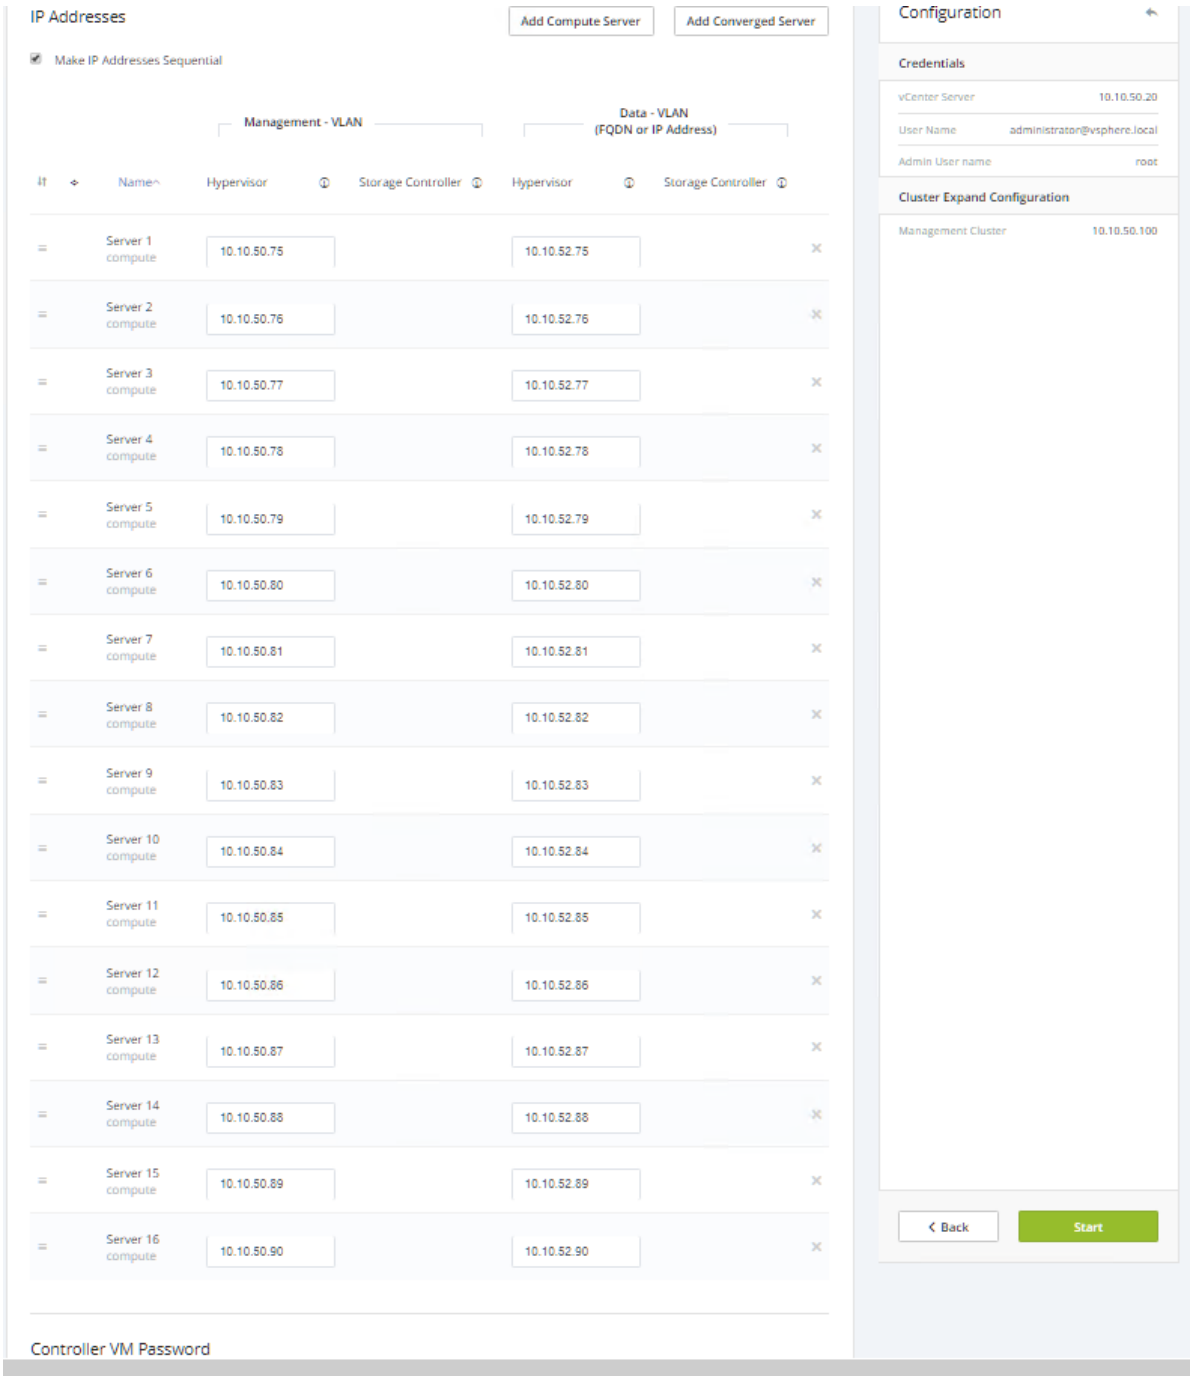

20. Cluster expansion workflow starts which performs deploy validation task first.

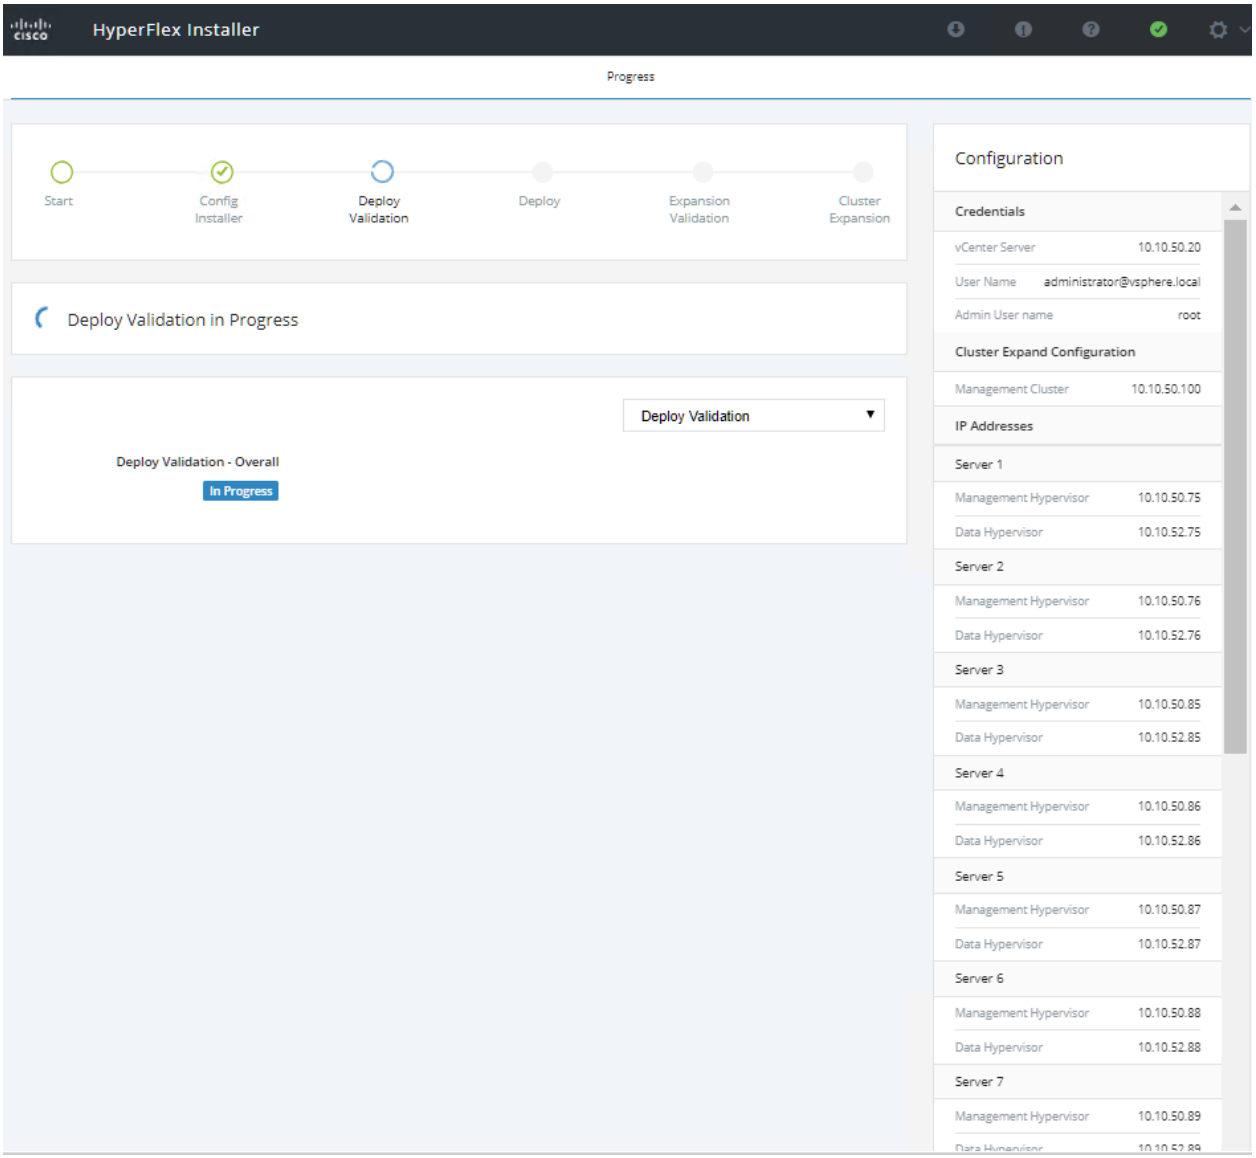

21. Performs deployment of HyperFlex controller virtual machine create and deployment task.

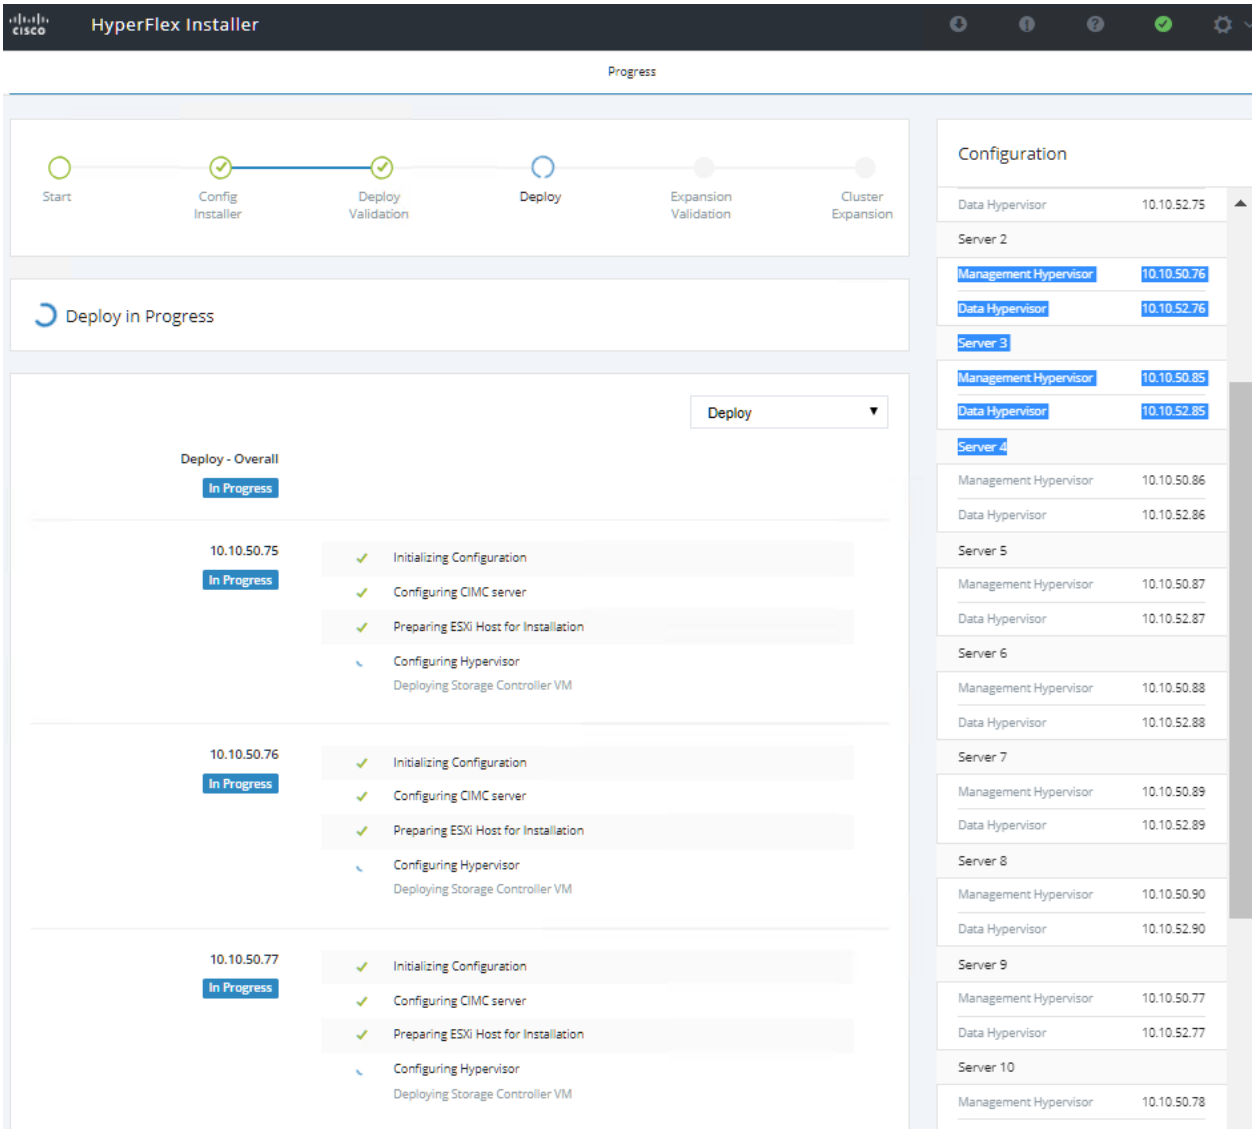

22. Performs expansion validation.

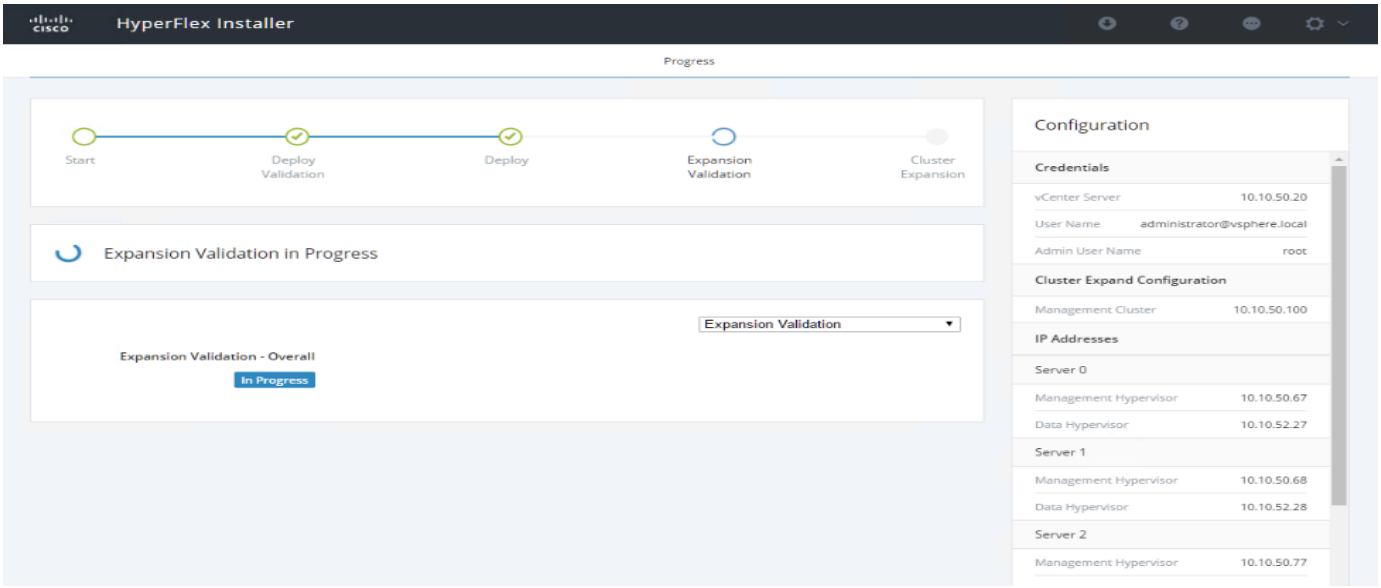
HyperFlex Installer
Progress
Start
Deploy Validation
Deploy
Expansion Validation
Cluster Expansion
Expansion Validation in Progress
Expansion Validation
Expansion Validation - Overall
In Progress
Configuration
Credentials
vCenter Server 10.10.50.20
User Name administrator@vsphere.local
Admin User Name root
Cluster Expand Configuration
Management Cluster 10.10.50.100
IP Addresses
Server 0
Management Hypervisor 10.10.50.67
Data Hypervisor 10.10.52.27
Server 1
Management Hypervisor 10.10.50.68
Data Hypervisor 10.10.52.28
Server 2
Management Hypervisor 10.10.50.77

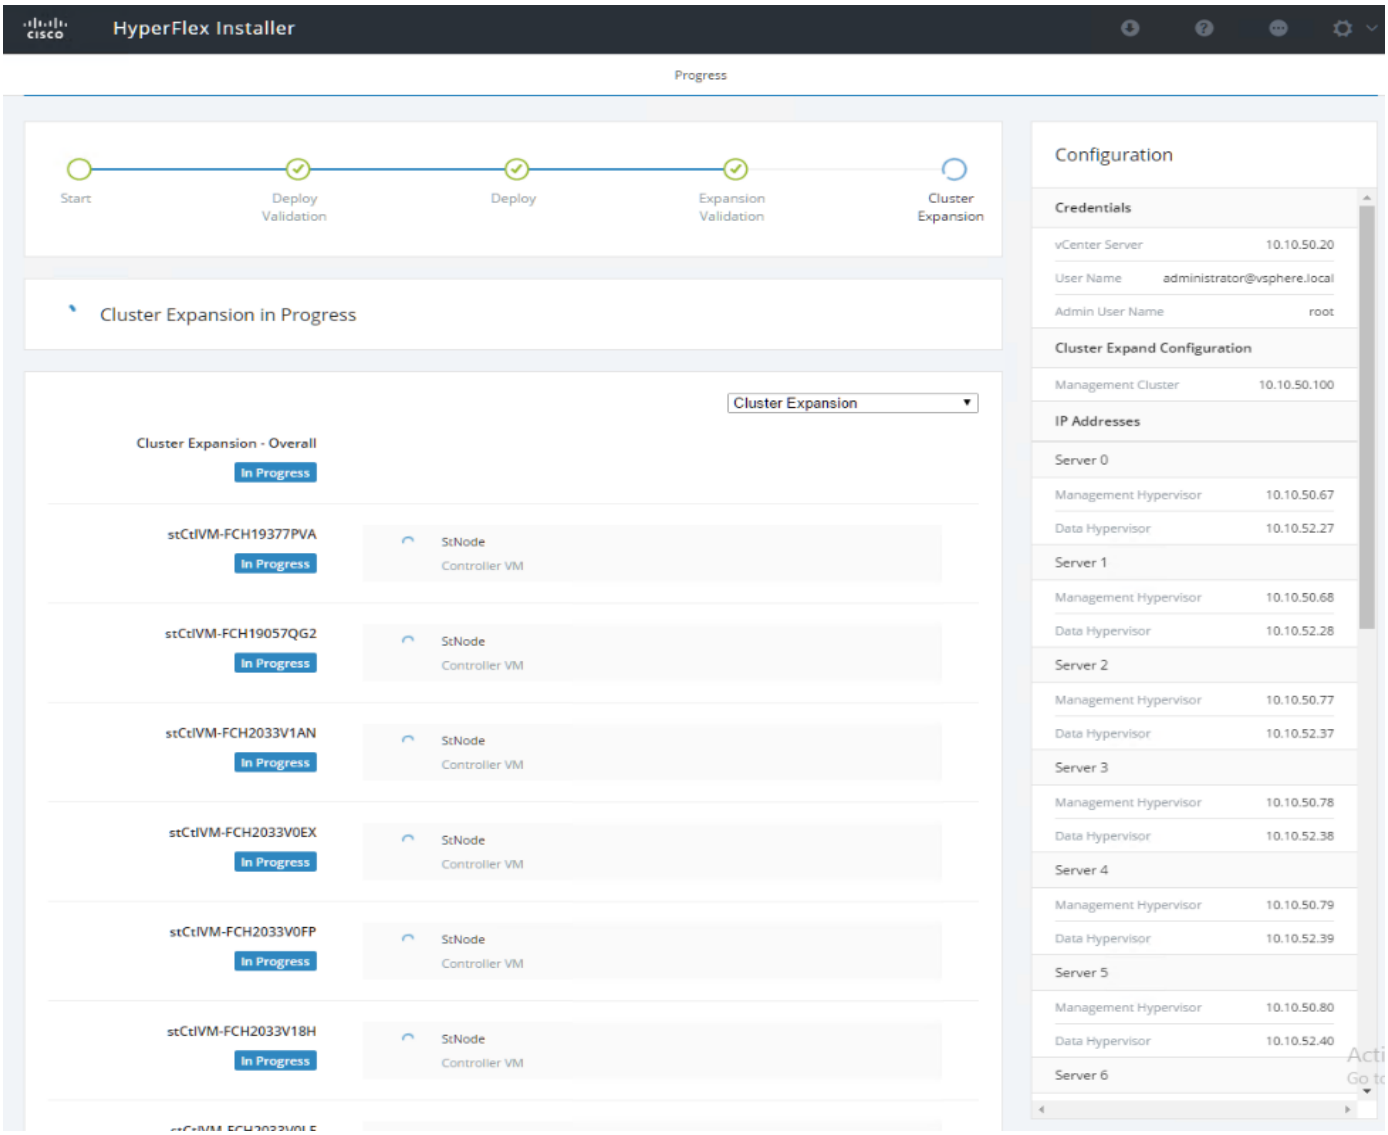
HyperFlex Installer
Progress
Start
Deploy Validation
Deploy
Expansion Validation
Cluster Expansion
Cluster Expansion in Progress
Cluster Expansion
Cluster Expansion - Overall
In Progress
stCtlVM-FCH19377PVA
In Progress
StNode
Controller VM
stCtlVM-FCH19057QG2
In Progress
StNode
Controller VM
stCtlVM-FCH2033V1AN
In Progress
StNode
Controller VM
stCtlVM-FCH2033V0EX
In Progress
StNode
Controller VM
stCtlVM-FCH2033V0FP
In Progress
StNode
Controller VM
stCtlVM-FCH2033V18H
In Progress
StNode
Controller VM
Configuration
Credentials
vCenter Server 10.10.50.20
User Name administrator@vsphere.local
Admin User Name root
Cluster Expand Configuration
Management Cluster 10.10.50.100
IP Addresses
Server 0
Management Hypervisor 10.10.50.67
Data Hypervisor 10.10.52.27
Server 1
Management Hypervisor 10.10.50.68
Data Hypervisor 10.10.52.28
Server 2
Management Hypervisor 10.10.50.77
Data Hypervisor 10.10.52.37
Server 3
Management Hypervisor 10.10.50.78
Data Hypervisor 10.10.52.38
Server 4
Management Hypervisor 10.10.50.79
Data Hypervisor 10.10.52.39
Server 5
Management Hypervisor 10.10.50.80
Data Hypervisor 10.10.52.40
Server 6

23. Summary of Expansion cluster workflow performed.

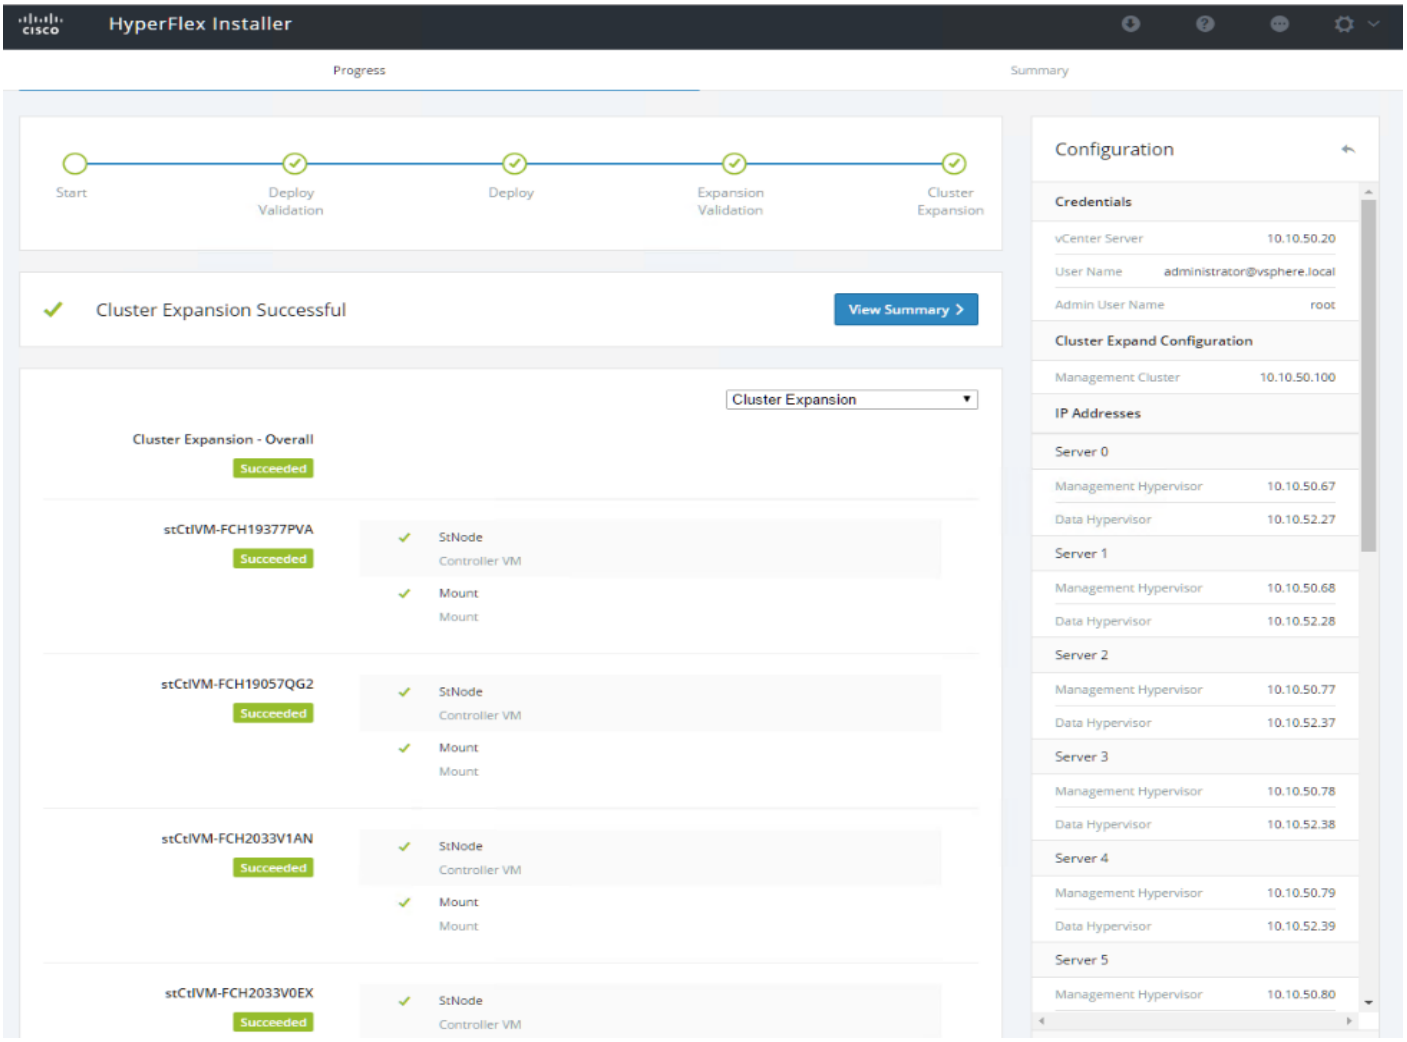

As part of the cluster creation operations, the HyperFlex Installer adds HyperFlex functionality to the vSphere vCenter identified in earlier steps. This functionality allows vCenter administrators to manage the HyperFlex cluster entirely from their vSphere Web Client.

24. Click Launch vSphere Web Client.

Cisco HyperFlex installer creates and configures a controller virtual machine on each converged or compute-only node. Naming convention used is as “stctlvm-<Serial Number for Cisco UCS Node>” shown in Figure 44.

Do **not** to change the name or any resource configuration for the controller virtual machine.

Figure 44 Cisco UCS Node Naming Convention

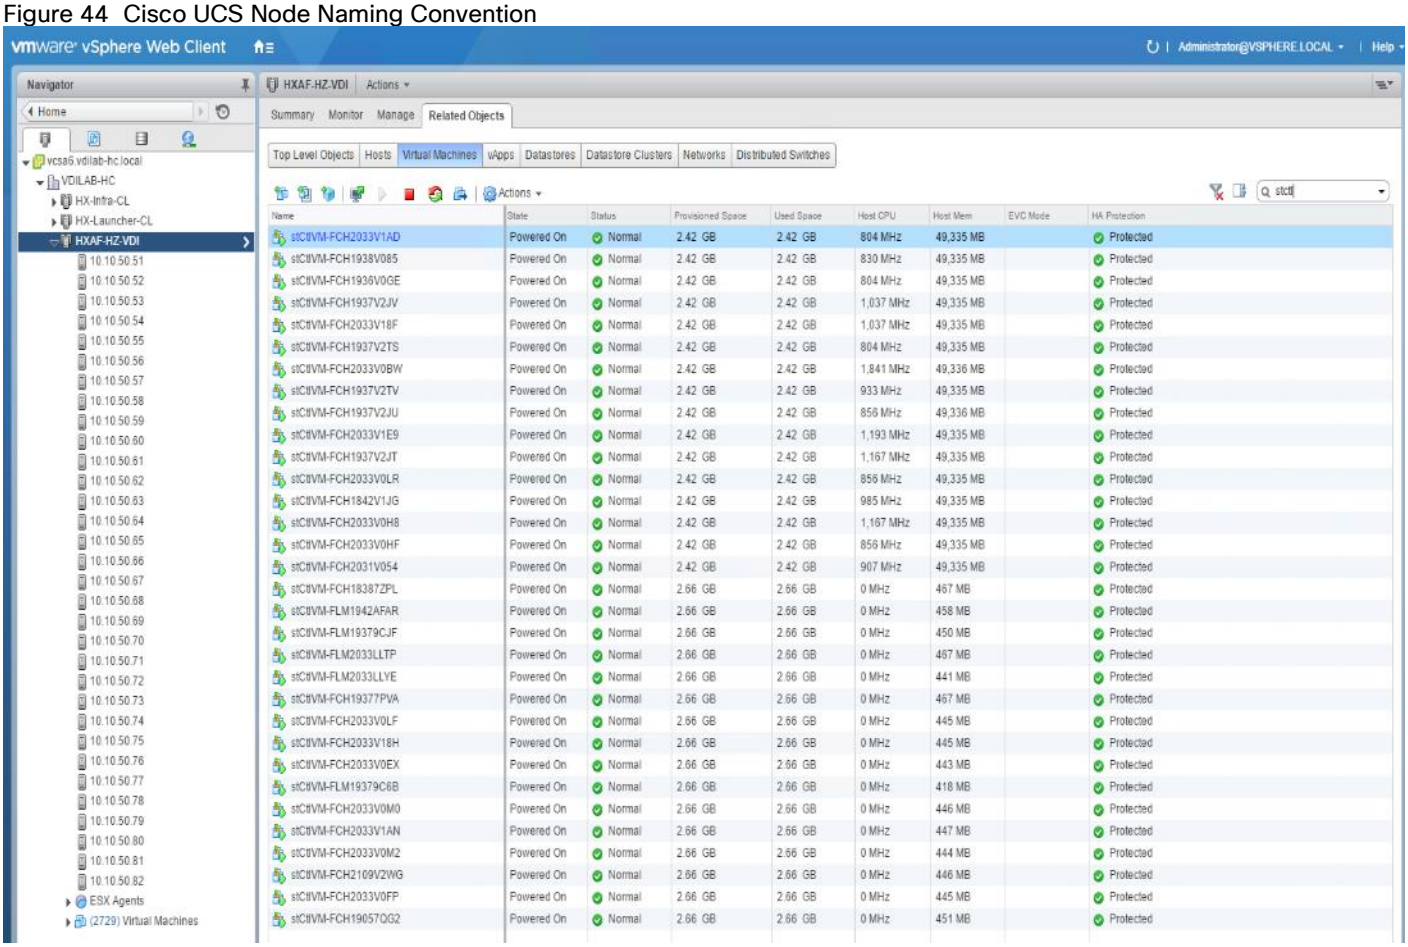

Run Cluster Post Installation Script

After a successful installation of HyperFlex cluster, run the post_install script by logging into the Data Platform Installer virtual machine through SSH, using the credentials configured earlier.

A built-in post install script automates basic final configuration tasks like enabling HA/DRS on HyperFlex cluster, configuring VMkernel for vMotion interface, creating datastore for ESXi logging, and so on, as shown in the following figures.

```
root@Cisco-HX-Data-Platform-Installer:~# post_install
Getting ESX hosts from HX cluster...
vCenter URL: 10.10.50.20
Enter vCenter username (user@domain): administrator@vsphere.local
vCenter Password:
Found datacenter VDILAB-HX
Found cluster HX-VDI-CL

Enable HA/DRS on cluster? (y/n) y

Disable SSH warning? (y/n) y

Configure ESXi logging onto HX datastore? (y/n) y
No datastores found
 Creating datastore...
 Name of datastore: HX-Logs
 Size (GB): 100
 Storing logs on datastore HX-Logs
 Creating folder [HX-Logs]/esxi_logs

Add vmotion interfaces? (y/n) y
 Netmask for vMotion: 255.255.255.0
 VLAN ID: (0-4096) 53
 vMotion IP for 10.10.50.27: 10.10.53.27
 Adding vmotion to 10.10.50.27
 Adding vmkernel to 10.10.50.27
 vMotion IP for 10.10.50.28: 10.10.53.28
 Adding vmotion to 10.10.50.28
 Adding vmkernel to 10.10.50.28
 vMotion IP for 10.10.50.29: 10.10.53.29
 Adding vmotion to 10.10.50.29
 Adding vmkernel to 10.10.50.29
 vMotion IP for 10.10.50.30: 10.10.53.30
 Adding vmotion to 10.10.50.30
 Adding vmkernel to 10.10.50.30
 vMotion IP for 10.10.50.31: 10.10.53.31
 Adding vmotion to 10.10.50.31
 Adding vmkernel to 10.10.50.31
 vMotion IP for 10.10.50.32: 10.10.53.32
 Adding vmotion to 10.10.50.32
 Adding vmkernel to 10.10.50.32
 vMotion IP for 10.10.50.33: 10.10.53.33
 Adding vmotion to 10.10.50.33
 Adding vmkernel to 10.10.50.33
 vMotion IP for 10.10.50.34: 10.10.53.34
 Adding vmotion to 10.10.50.34
 Adding vmkernel to 10.10.50.34
```

```
Add VM network VLANs? (y/n) n

Enable NTP on ESX hosts? (y/n) y
Starting ntpd service on 10.10.50.27
Starting ntpd service on 10.10.50.28
Starting ntpd service on 10.10.50.29
Starting ntpd service on 10.10.50.30
Starting ntpd service on 10.10.50.31
Starting ntpd service on 10.10.50.32
Starting ntpd service on 10.10.50.33
Starting ntpd service on 10.10.50.34

Send test email? (y/n) n

Validating cluster health and configuration...
 Found UCSM 10.29.132.11, logging with username admin.  Org is hx-vdi-org
 UCSM Password:
```

1. To run the script, first use your tool of choice to make a secure connection to the Cisco HyperFlex Data Platform installer using its IP address and port 22.

2. Authenticate with the credentials provided earlier; user name: root with password Cisco 123 if you did not change the defaults.

3. When authenticated, enter post_install at the command prompt, then press Enter.
4. Provide a valid vCenter administrator user name and password and the vCenter url IP address.
5. Type y for yes to each of the promts that follow except Add VM network VLANs? (y/n) where you can choose whether or not to send health status data through SMS to Cisco support.
6. Provide the requested user credentials, the vMotion netmask, VLAN ID and an IP address on the vMotion VLAN for each host when prompted for the vmkernel IP.
7. Sample post install input and output:

```
root@Cisco-HX-Data-Platform-Installer:root@Cisco-HX-Data-Platform-
Installer:~#post_install Getting ESX hosts from HX cluster...

vCenter URL: 10.10.50.20

Enter vCenter username (user@domain): administrator@vsphere.local

vCenter Password:

Found datacenter VDILAB-HX

Found cluster HX-VDI-CL

Enable HA/DRS on cluster? (y/n) y

Disable SSH warning? (y/n) y

Add vmotion interfaces? (y/n) y

 Netmask for vMotion: 255.255.255.0

 VLAN ID: (0-4096) 53

 vMotion IP for 10.10.50.27: 10.10.53.27

 Adding vmotion to 10.10.50.27

 Adding vmkernel to 10.10.50.27

 vMotion IP for 10.10.50.28: 10.10.53.28

 Adding vmotion to 10.10.50.28

 Adding vmkernel to 10.10.50.28

 vMotion IP for 10.10.50.29: 10.10.53.29

 Adding vmotion to 10.10.50.29

 Adding vmkernel to 10.10.50.29

 vMotion IP for 10.10.50.30: 10.10.53.30

 Adding vmotion to 10.10.50.30

 Adding vmkernel to 10.10.50.30

 vMotion IP for 10.10.50.31: 10.10.53.31
```

```
 Adding vmotion to 10.10.50.31
 Adding vmkernel to 10.10.50.31
 vMotion IP for 10.10.50.32: 10.10.53.32
 Adding vmotion to 10.10.50.32
 Adding vmkernel to 10.10.50.32
 vMotion IP for 10.10.50.33: 10.10.53.33
 Adding vmotion to 10.10.50.33
 Adding vmkernel to 10.10.50.33
 vMotion IP for 10.10.50.34: 10.10.53.34
 Adding vmotion to 10.10.50.34
 Adding vmkernel to 10.10.50.34
Add VM network VLANs? (y/n) n
Enable NTP on ESX hosts? (y/n) y
Starting ntpd service on 10.10.50.27
Starting ntpd service on 10.10.50.28
Starting ntpd service on 10.10.50.29
Starting ntpd service on 10.10.50.30
Starting ntpd service on 10.10.50.31
Starting ntpd service on 10.10.50.32
Starting ntpd service on 10.10.50.33
Starting ntpd service on 10.10.50.34
Send test email? (y/n) n
Validating cluster health and configuration...
 Found UCSM 10.29.132.11, logging with username admin.  Org is hx-vdi-org
 UCSM Password:
TChecking MTU settings
 Pinging 10.10.52.107 from vmk1
 Pinging 10.10.52.101 from vmk1
 Pinging 10.10.52.105 from vmk1
 Pinging 10.10.52.108 from vmk1
 Pinging 10.10.52.102 from vmk1
 Pinging 10.10.52.104 from vmk1
```

```
Pinging 10.10.52.106 from vmk1
Pinging 10.10.52.103 from vmk1
Setting vnic2 to active and vmnic3 to standby
Pinging 10.10.52.107 from vmk1
Pinging 10.10.52.107 with mtu 8972 from vmk1
Pinging 10.10.52.101 from vmk1
Pinging 10.10.52.101 with mtu 8972 from vmk1
Pinging 10.10.52.105 from vmk1
Pinging 10.10.52.105 with mtu 8972 from vmk1
Pinging 10.10.52.108 from vmk1
Pinging 10.10.52.108 with mtu 8972 from vmk1
Pinging 10.10.52.102 from vmk1
Pinging 10.10.52.102 with mtu 8972 from vmk1
Pinging 10.10.52.104 from vmk1
Pinging 10.10.52.104 with mtu 8972 from vmk1
Pinging 10.10.52.106 from vmk1
Pinging 10.10.52.106 with mtu 8972 from vmk1
Pinging 10.10.52.103 from vmk1
Pinging 10.10.52.103 with mtu 8972 from vmk1
Setting vmnic3 to active and vnic2 to standby
Pinging 10.10.50.33 from vmk0
Pinging 10.10.50.27 from vmk0
Pinging 10.10.50.31 from vmk0
Pinging 10.10.50.34 from vmk0
Pinging 10.10.50.28 from vmk0
Pinging 10.10.50.30 from vmk0
Pinging 10.10.50.32 from vmk0
Pinging 10.10.50.29 from vmk0
Setting vnic1 to active and vmnic0 to standby
Pinging 10.10.50.33 from vmk0
Pinging 10.10.50.27 from vmk0
Pinging 10.10.50.31 from vmk0
```

```
Pinging 10.10.50.34 from vmk0
Pinging 10.10.50.28 from vmk0
Pinging 10.10.50.30 from vmk0
Pinging 10.10.50.32 from vmk0
Pinging 10.10.50.29 from vmk0
Setting vmnic0 to active and vnic1 to standby
Pinging 10.10.53.27 from vmk2
Pinging 10.10.53.28 from vmk2
Pinging 10.10.53.29 from vmk2
Pinging 10.10.53.30 from vmk2
Pinging 10.10.53.31 from vmk2
Pinging 10.10.53.32 from vmk2
Pinging 10.10.53.33 from vmk2
Pinging 10.10.53.34 from vmk2
Setting vnic7 to active and vmnic6 to standby
Pinging 10.10.53.27 from vmk2
Pinging 10.10.53.27 with mtu 8972 from vmk2
Pinging 10.10.53.28 from vmk2
Pinging 10.10.53.28 with mtu 8972 from vmk2
Pinging 10.10.53.29 from vmk2
Pinging 10.10.53.29 with mtu 8972 from vmk2
Pinging 10.10.53.30 from vmk2
Pinging 10.10.53.30 with mtu 8972 from vmk2
Pinging 10.10.53.31 from vmk2
Pinging 10.10.53.31 with mtu 8972 from vmk2
Pinging 10.10.53.32 from vmk2
Pinging 10.10.53.32 with mtu 8972 from vmk2
Pinging 10.10.53.33 from vmk2
Pinging 10.10.53.33 with mtu 8972 from vmk2
Pinging 10.10.53.34 from vmk2
Pinging 10.10.53.34 with mtu 8972 from vmk2
Setting vmnic6 to active and vnic7 to standby
```

```
Network Summary:
 Host: 10.10.50.27
    vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
       vmnic0 - 1 - K22-HXVDI-A - active
       vmnic1 - 1 - K22-HXVDI-B - standby
         Portgroup Name  -  VLAN
         Storage Controller Management Network - 50
         Management Network - 50
    vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
       vmnic4 - 1 - K22-HXVDI-A - active
       vmnic5 - 1 - K22-HXVDI-B - active
         Portgroup Name  -  VLAN
         vm-network-54 - 54
    vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
       vmnic6 - 1 - K22-HXVDI-A - active
       vmnic7 - 1 - K22-HXVDI-B - standby
         Portgroup Name  -  VLAN
         vmotion - 53
    vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
       vmnic2 - 1 - K22-HXVDI-A - standby
       vmnic3 - 1 - K22-HXVDI-B - active
         Portgroup Name  -  VLAN
         Storage Controller Data Network - 52
         Storage Hypervisor Data Network - 52
 Host: 10.10.50.28
    vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
       vmnic0 - 1 - K22-HXVDI-A - active
       vmnic1 - 1 - K22-HXVDI-B - standby
         Portgroup Name  -  VLAN
         Storage Controller Management Network - 50
         Management Network - 50
    vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
```

```
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.29
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - K22-HXVDI-A - active
      vmnic1 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        Storage Controller Management Network - 50
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
```

```
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.30
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - K22-HXVDI-A - active
      vmnic1 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        Storage Controller Management Network - 50
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.31
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - K22-HXVDI-A - active
```

```
      vmnic1 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        Storage Controller Management Network - 50
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.32
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - K22-HXVDI-A - active
      vmnic1 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        Storage Controller Management Network - 50
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
```

```
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.33
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - K22-HXVDI-A - active
      vmnic1 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        Storage Controller Management Network - 50
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
```

```
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.34
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - K22-HXVDI-A - active
      vmnic1 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        Storage Controller Management Network - 50
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - K22-HXVDI-A - active
      vmnic5 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - K22-HXVDI-A - active
      vmnic7 - 1 - K22-HXVDI-B - standby
        Portgroup Name  -  VLAN
        vmotion - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - K22-HXVDI-A - standby
      vmnic3 - 1 - K22-HXVDI-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52
Host: 10.10.50.27
    No errors found
Host: 10.10.50.28
    No errors found
Host: 10.10.50.29
    No errors found
Host: 10.10.50.30
```

```
        No errors found
    Host: 10.10.50.31
        No errors found
    Host: 10.10.50.32
        No errors found
    Host: 10.10.50.33
        No errors found

    Host: 10.10.50.34
        No errors found
    Controller VM Clocks:
        stCtlVM-FCH1937V2JV  -  2018-10-22 05:32:09
        stCtlVM-FCH1937V2TV  -  2018-10-22 05:32:25
        stCtlVM-FCH1842V1JG  -  2018-10-22 05:32:41
        stCtlVM-FCH1936V0GE  -  2018-10-22 05:32:57
        stCtlVM-FCH1937V2JT  -  2018-10-22 05:33:14
        stCtlVM-FCH1938V085  -  2018-10-22 05:33:30
        stCtlVM-FCH1937V2TS  -  2018-10-22 05:33:46
        stCtlVM-FCH1937V2JU  -  2018-10-22 05:34:02
    Cluster:
        Version - 3.5.1a-31118
        Model - HXAF220C-M5S
        Health - HEALTHY
        Access Policy - LENIENT
        ASUP enabled - False
        SMTP Server -
root@Cisco-HX-Data-Platform-Installer:~#
```

8. Log into vSphere WebClient to create additional shared datastore.
9. Go to the Summary tab on the cluster created through the HyperFlex cluster creation workflow.
10. From Cisco HyperFlex Systems, click the cluster name.

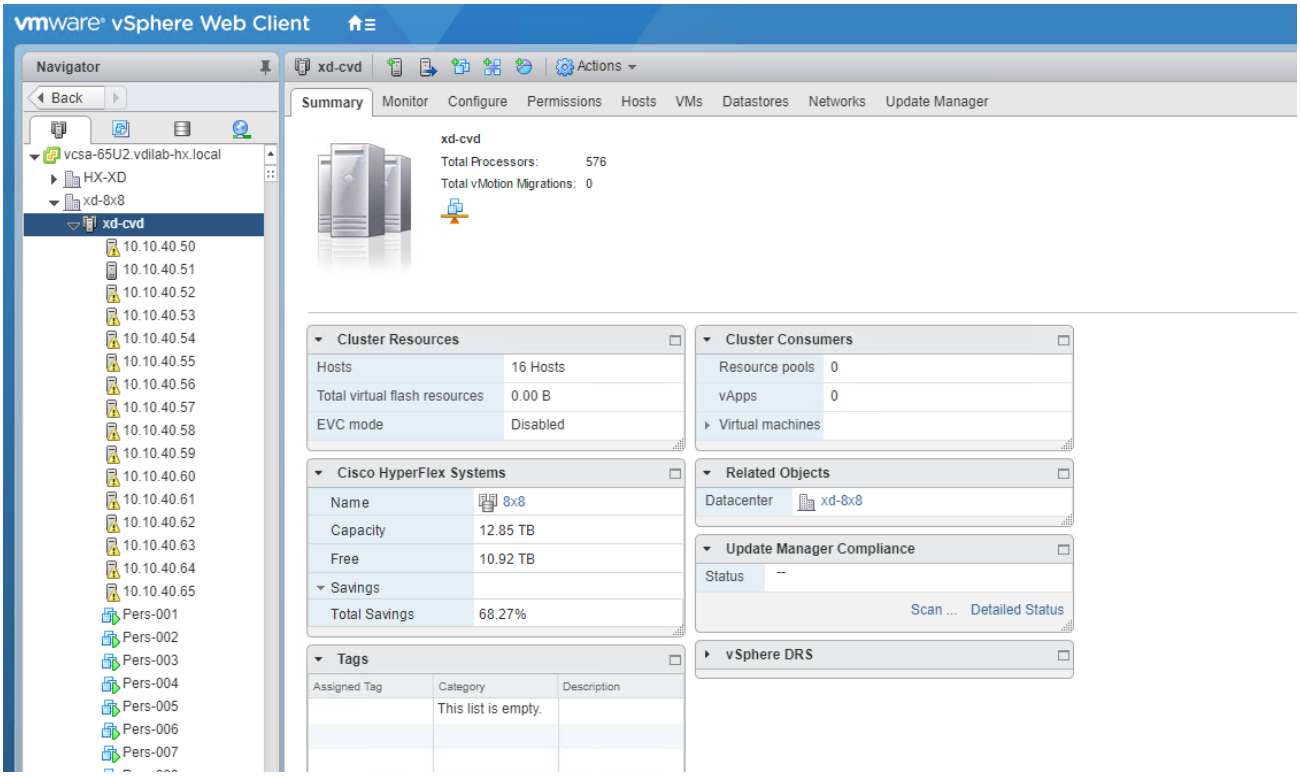

The Summary tab displays the details about the cluster status, capacity, and performance.

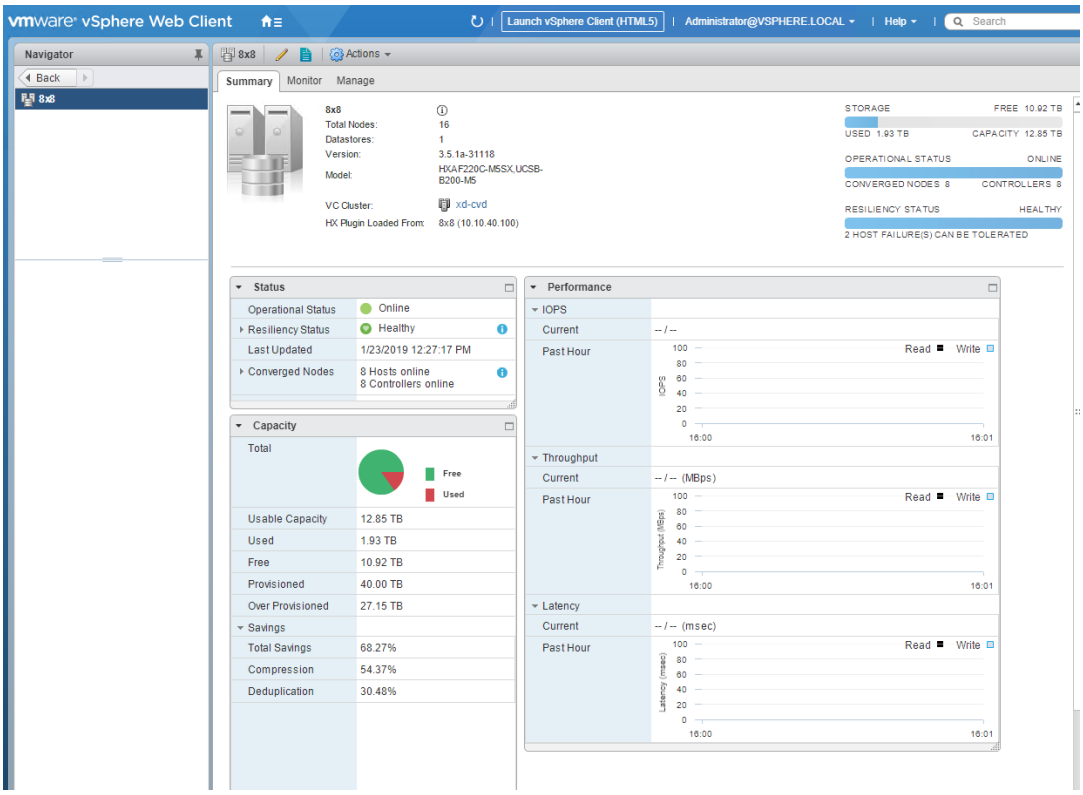

11. Click Manage, select Datastores. Click the Add datastore icon, select the datastore name and size to provision.

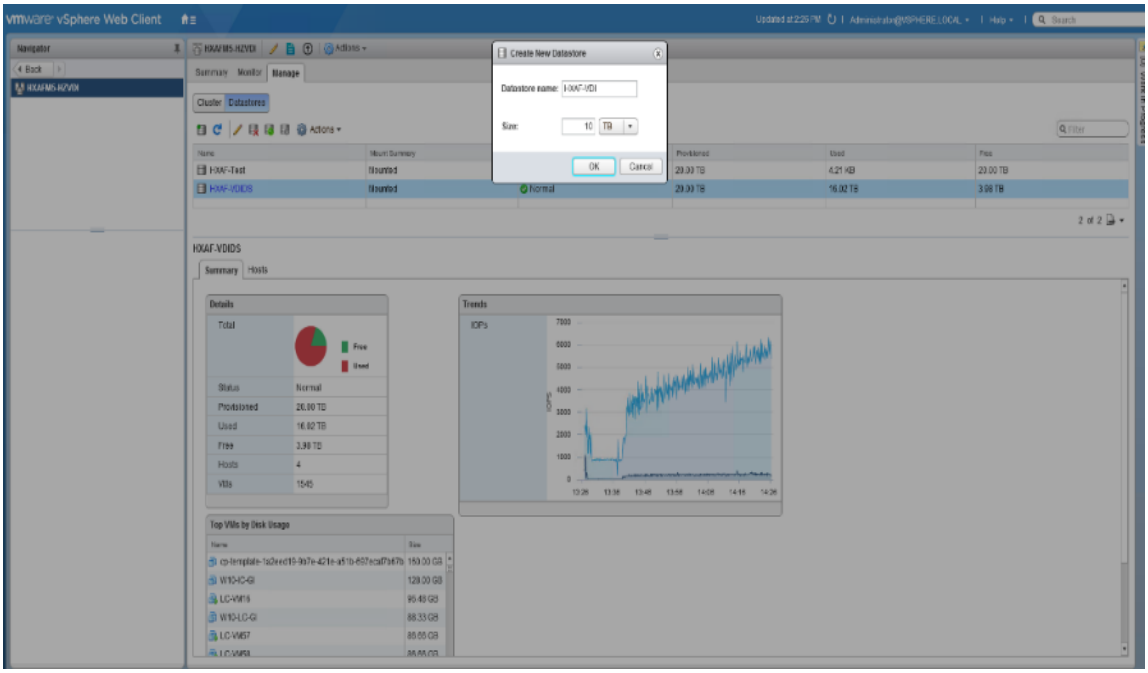

You have created a 40TB datastore for the Citrix pooled, persistent/non-persistent, and HSD server desktop performance test.

Alternatively, HyperFlex connect WebUI can be utilized as well to create a datastore. While using HyperFlex Connect UI to create a datastore there is an option to select Block size. By default datastores are created with 8K Block size using vSphere WebClient.

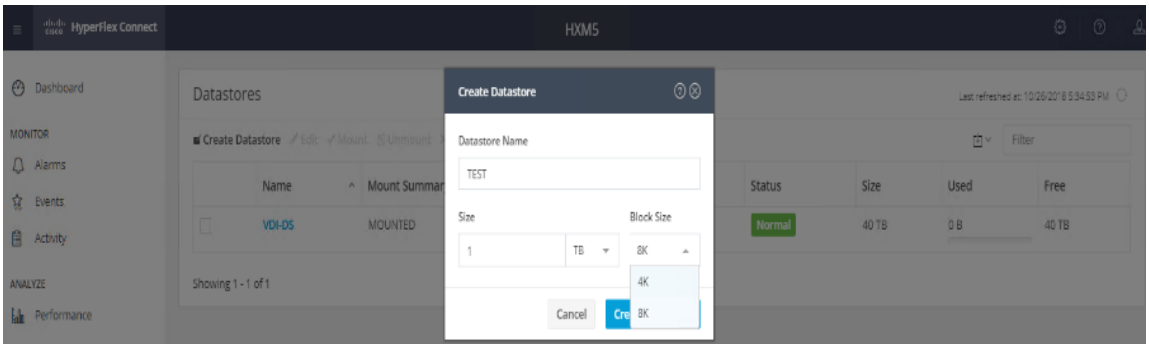

Building the Virtual Machines and Environment for Workload Testing

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

Software Infrastructure Configuration

Install and configure the infrastructure virtual machines by following the process provided in Table 12

Table 12 Test Infrastructure Virtual Machine Configuration

Configuration	Citrix Desktop Controllers Virtual Machines	Citrix Profile Servers Virtual Machines
Operating System	Microsoft Windows Server 2016	Microsoft Windows Server 2016
Virtual CPU amount	6	8
Memory amount	8 GB	8 GB
Network	VMNIC	Network
Disk-1 (OS) size and location	40 GB	Disk-1 (OS) size and location
Disk-2 size and location	–	
Configuration	Microsoft Active Directory DC's Virtual Machines	
Operating system	Microsoft Windows Server 2016	Operating system
Virtual CPU amount	4	
Memory amount	4 GB	
Network	VMNIC	
Disk size and location	40 GB	
Configuration	Microsoft SQL Server Virtual Machine	Citrix StoreFront Virtual Machine
Operating system	Microsoft Windows Server 2016	Operating system
Virtual CPU amount	4	4
Memory amount	16 GB	8 GB
Network	VMNIC	Network
Disk-1 (OS) size and location	40 GB	Disk-1 (OS) size and location
Disk-2 size and location	200 GB Infra-DS volume	Disk-2 size and location
Configuration	Citrix License Server Virtual Machine	NetScaler VPX Appliance Virtual Machine
Operating system	Microsoft Windows Server 2016	NS11.1 52.13.nc
Virtual CPU amount	4	2
Memory amount	4 GB	2 GB
Network	VMNIC	Network

Disk size and location	40 GB	20 GB

Prepare the Master Images

This section details how to create the golden (or master) images for the environment. virtual machines for the master images must first be installed with the software components needed to build the golden images. For this CVD, the images contain the basics needed to run the Login VSI workload.

To prepare the master virtual machines for the Hosted Virtual Desktops (HVDs) and Hosted Shared Desktops (HSDs), there are three major steps to complete when the base virtual machine has been created:

- Installing OS
- Installing application software
- Installing the Virtual Delivery Agents (VDAs)

The master image HVD and HSD virtual machines were configured as listed in Table 13 :

Table 13 HVD and HSD Configurations

Configuration	HVDI Virtual Machines	HSD Virtual Machines
Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2016
Virtual CPU amount	2	8
Memory amount	4.0 GB (reserved)	24 GB (reserved)
Network	VMNIC vm-network	VMNIC vm-network
Citrix PVS vDisk size and location	24 GB WriteCache Volume	40 GB WriteCache Volume
Citrix PVS write cache Disk size	6 GB	24 GB
Additional software used for testing	Microsoft Office 2016 Login VSI 4.1.32 (Knowledge Worker Workload)	Microsoft Office 2016 Login VSI 4.1.32 (Knowledge Worker Workload)

Install and Configure Citrix Desktop Delivery Controller, Citrix Licensing, and StoreFront

This section details the installation of the core components of the Citrix Virtual Apps and Desktops 1808 system. This CVD provides the process to install two Desktop Delivery Controllers to support hosted shared desktops (HSD), non-persistent virtual desktops (VDI), and persistent virtual desktops (VDI).

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

Install Citrix License Server

To install the Citrix License Server, complete the following steps:

1. To begin the installation, connect to the first Citrix License server and launch the installer from the Citrix Virtual Apps and Desktops 1808 ISO.

2. Click Start.

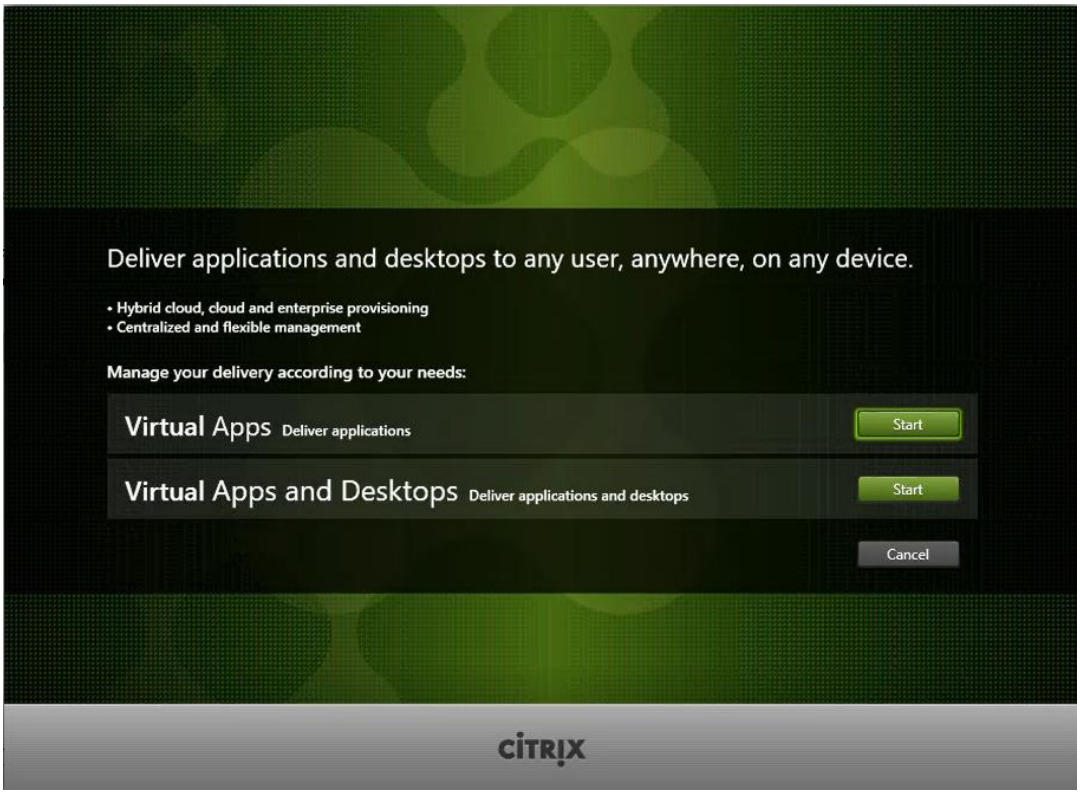

3. Click “Extend Deployment – Citrix License Server.”

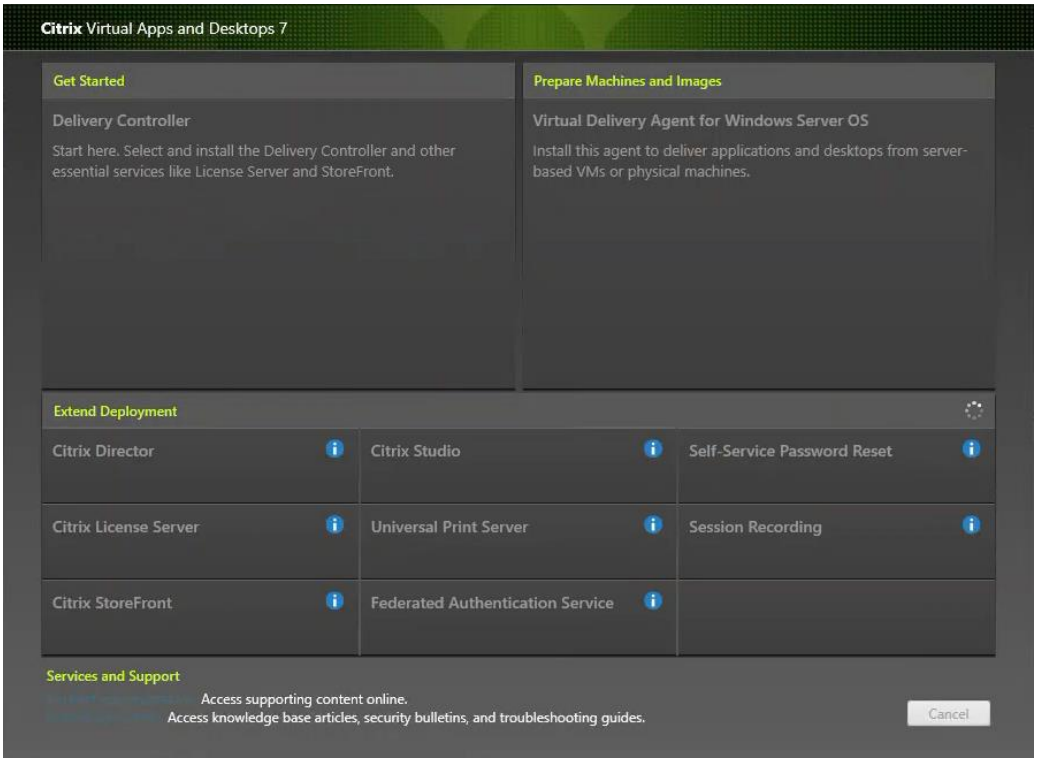

4. Read the Citrix License Agreement.

5. 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.

6. Click Next.

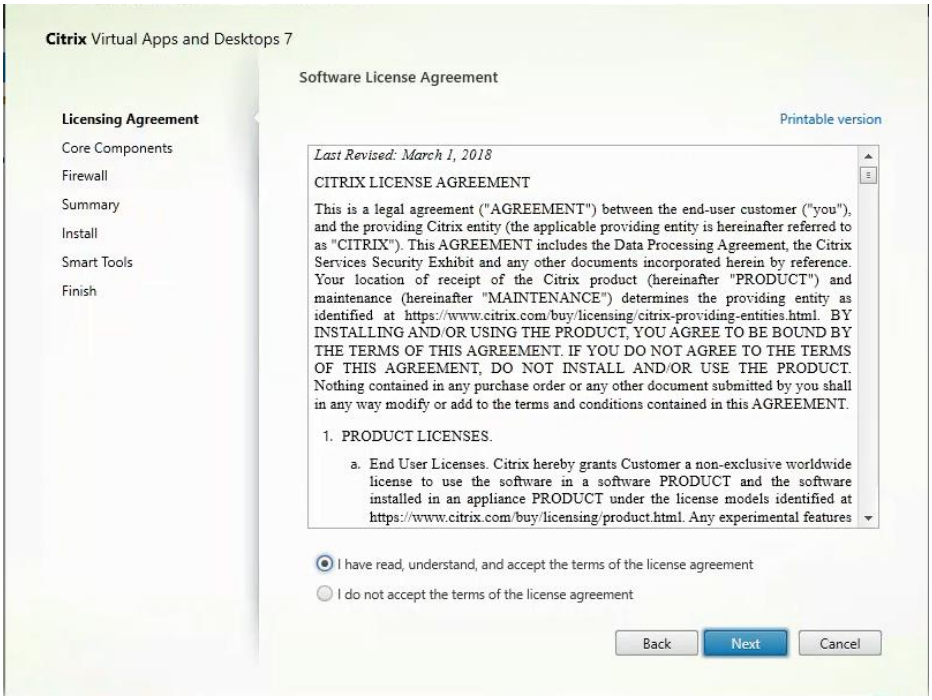

7. Click Next.

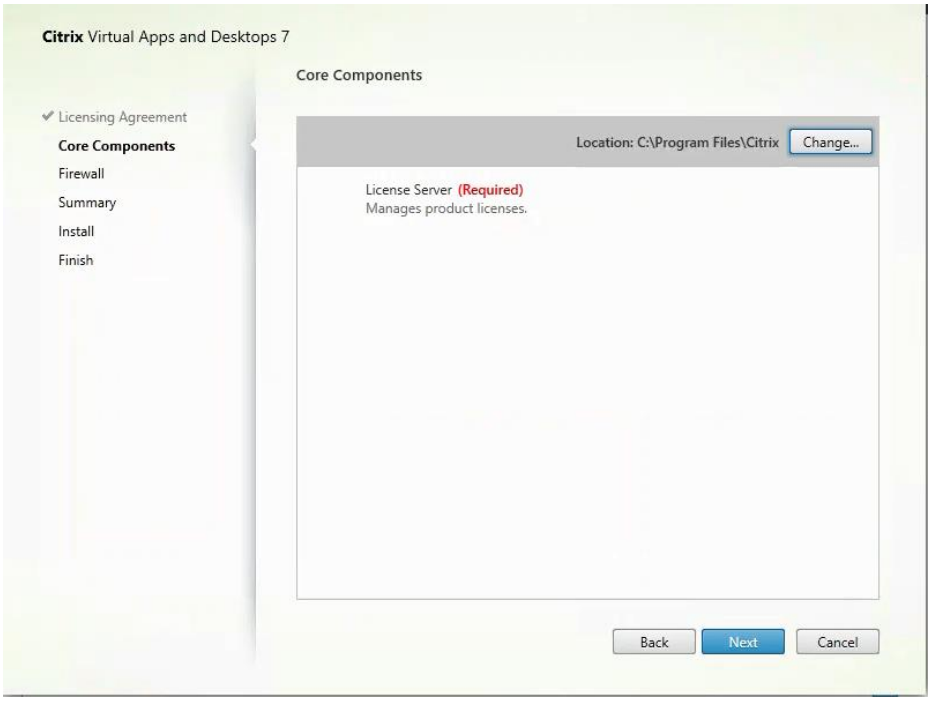

8. Select the default ports and automatically configured firewall rules.

9. Click Next.

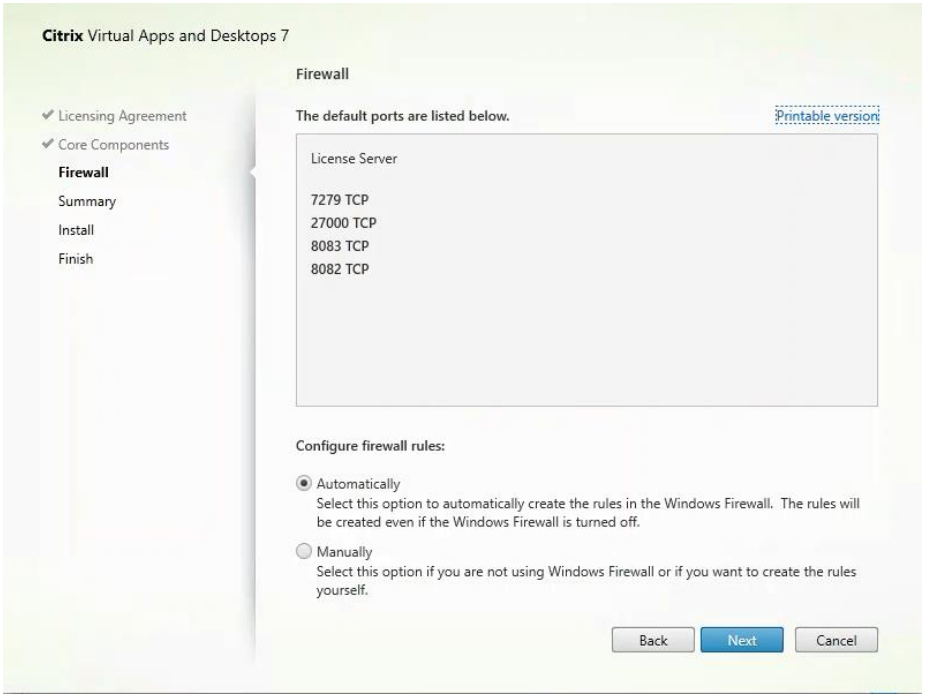

10. Click Install.

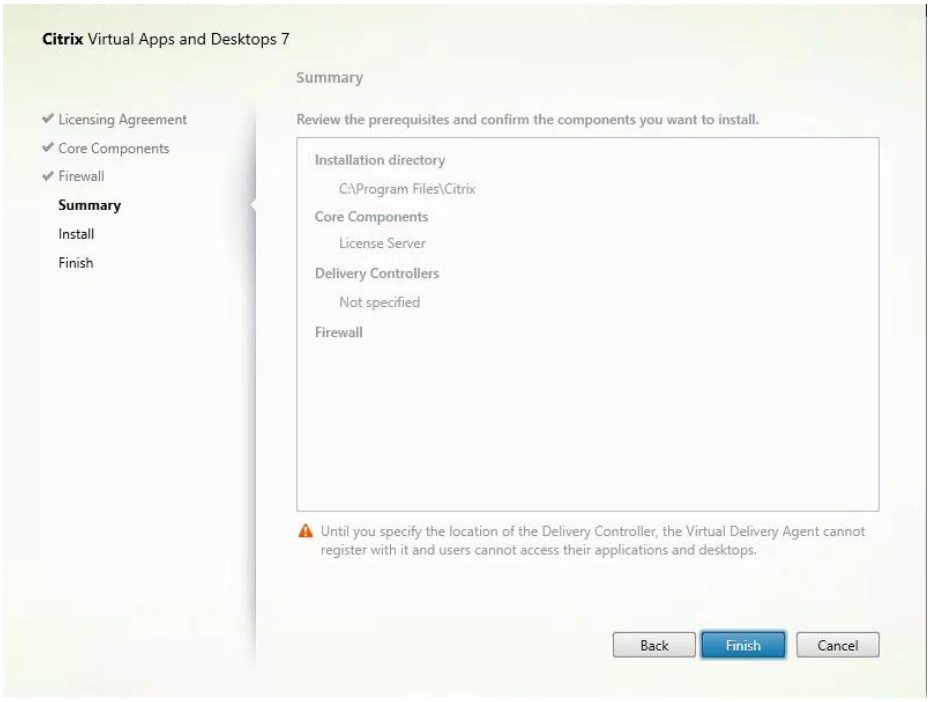

11. Click Finish to complete the installation.

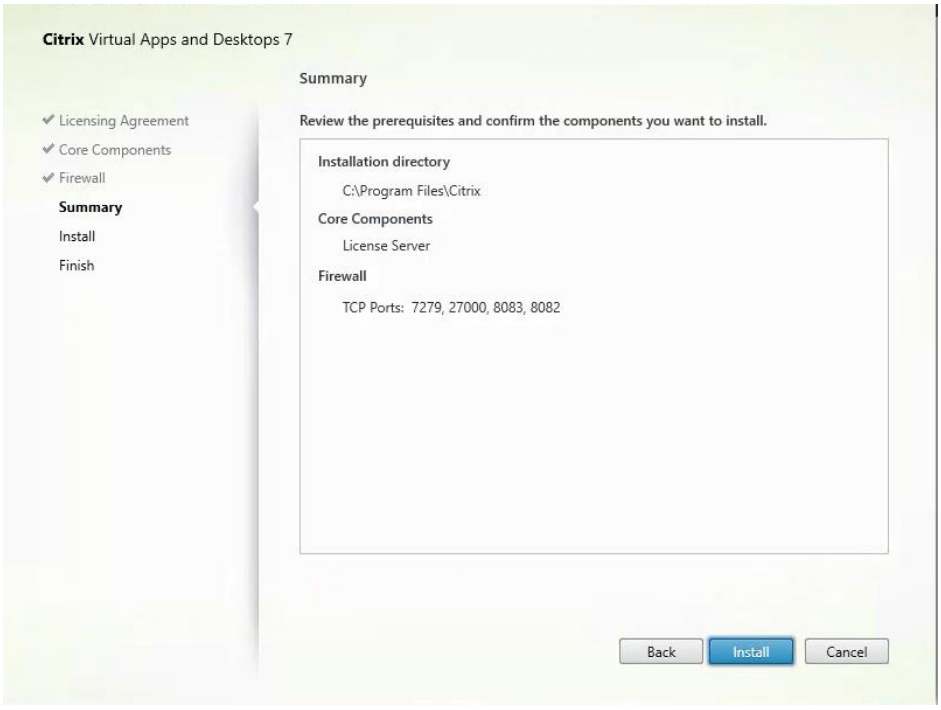

Install Citrix Licenses

To install the Citrix Licenses, complete the following steps:

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

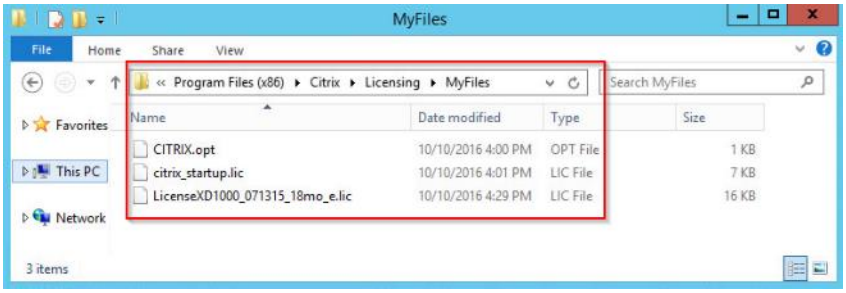

2. Restart the server or Citrix licensing services so that the licenses are activated.

3. Run the application Citrix License Administration Console.

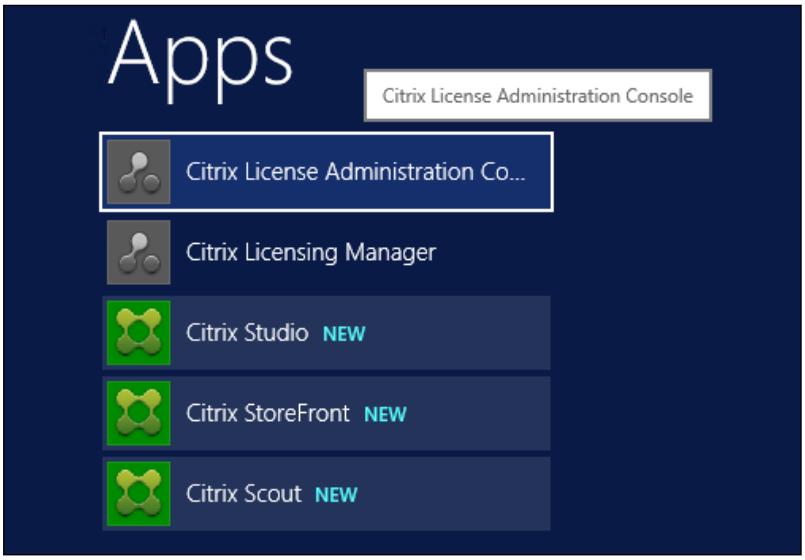

4. Confirm that the license files have been read and enabled correctly.

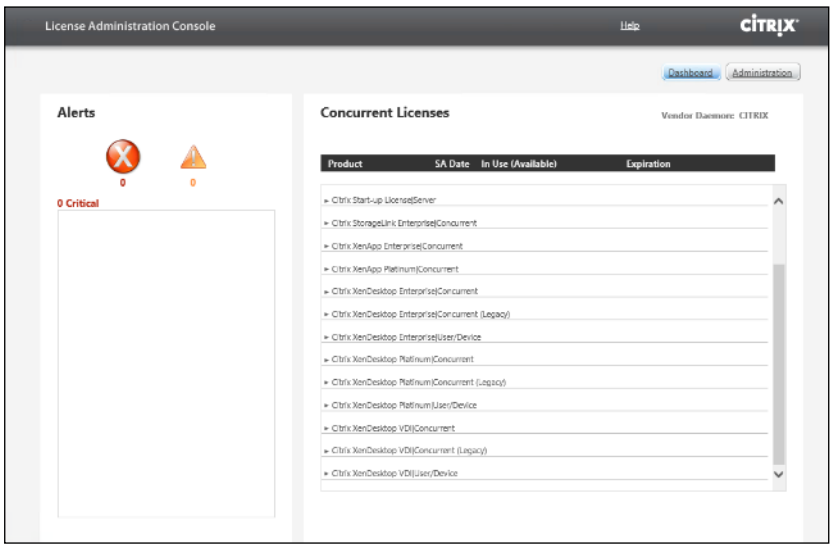

Install Citrix Desktop Broker/Studio

To install Citrix Desktop, complete the following steps:

1. Connect to the first Citrix VDI server and launch the installer from the Citrix Desktop 1808 ISO.
2. Click Start.

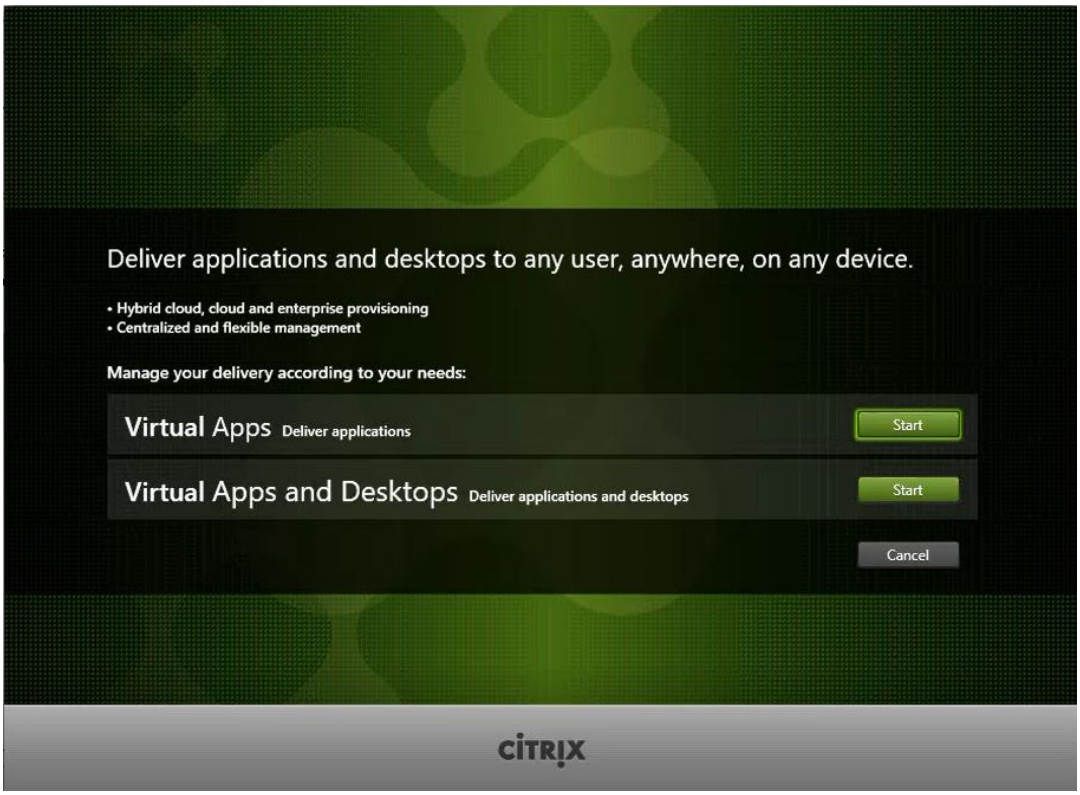

The installation wizard presents a menu with three subsections.

3. Click “Get Started – Delivery Controller.”

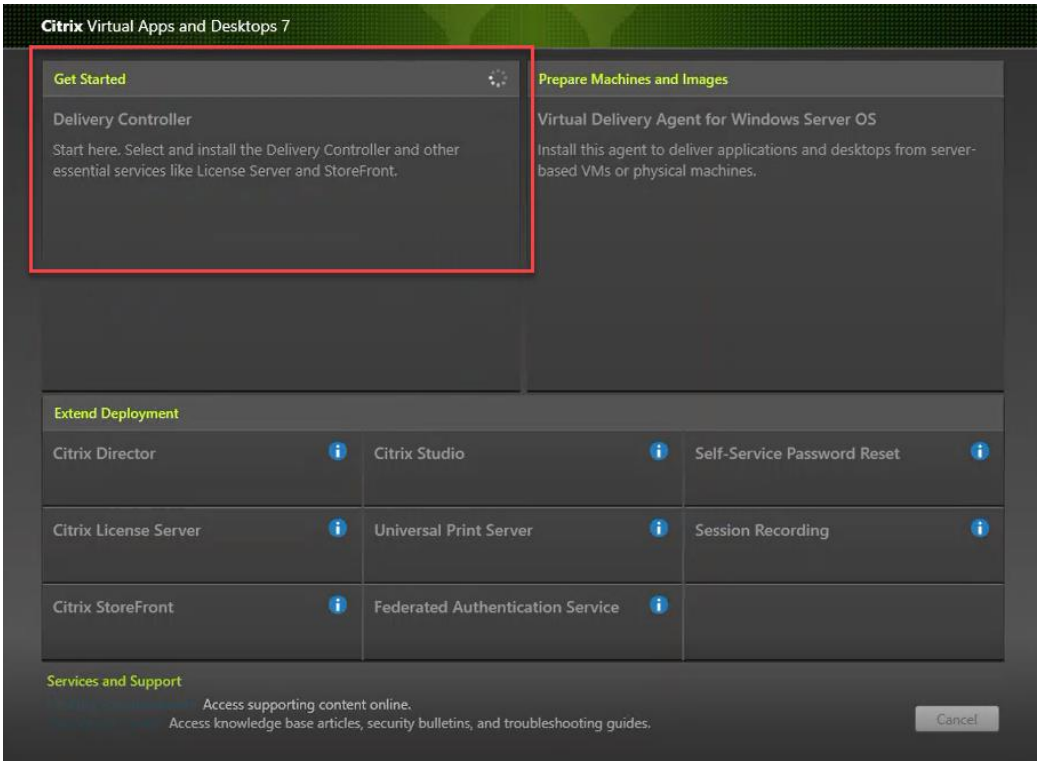

4. Read the Citrix License Agreement.

5. 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.

6. Click Next.

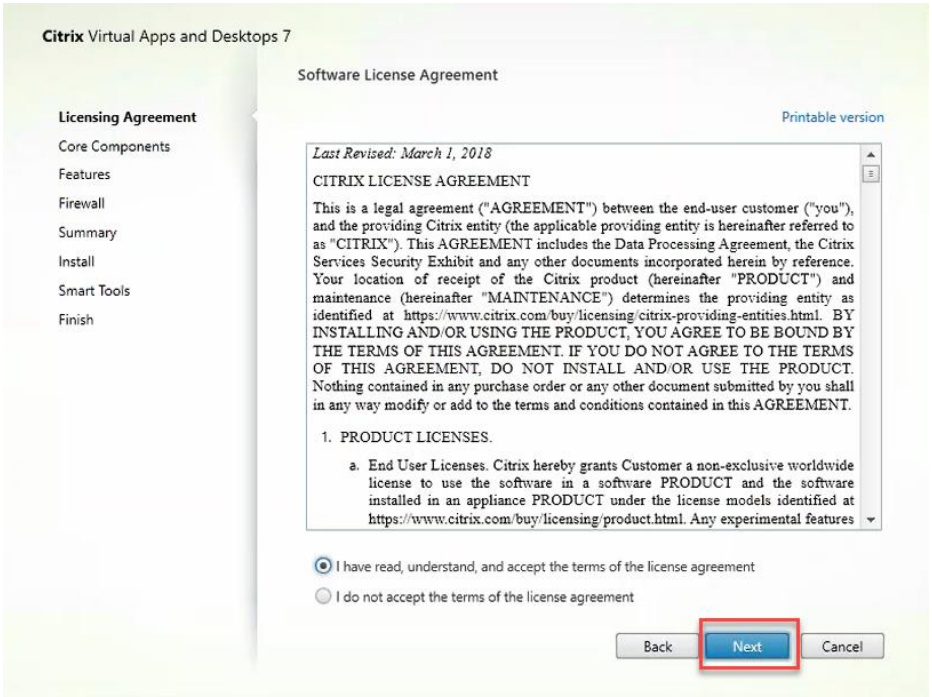

7. Select the components to be installed on the first Delivery Controller Server:
   a. Delivery Controller
   b. Studio
   c. Director
8. Click Next.

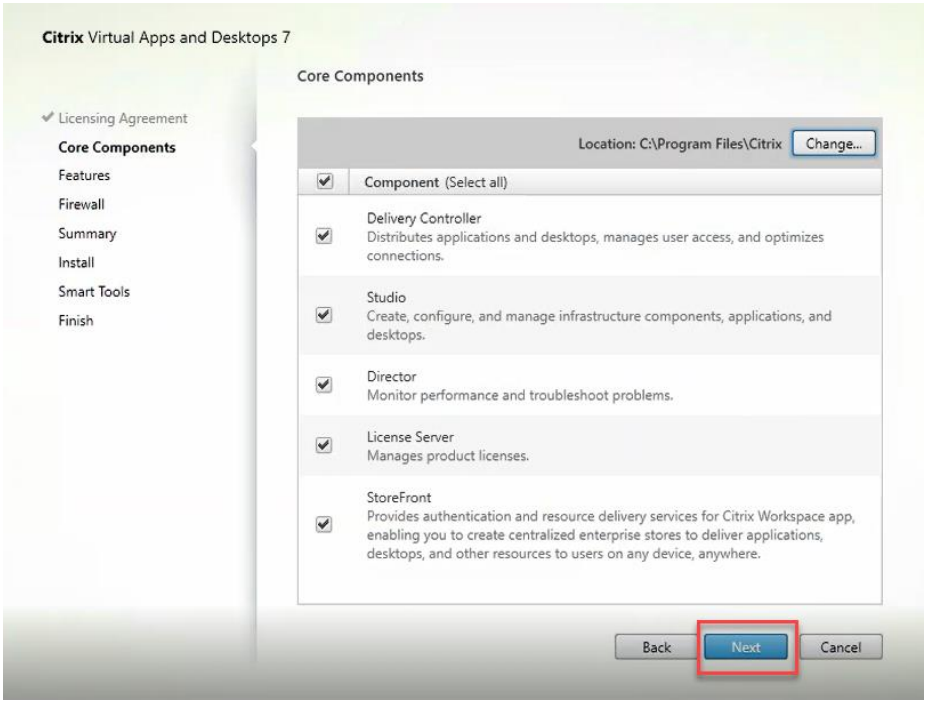

Dedicated StoreFront and License servers should be implemented for large-scale deployments.

9. Since a SQL Server will be used to Store the Database, leave "Install Microsoft SQL Server 2012 SP1 Express" unchecked.
10. Click Next.

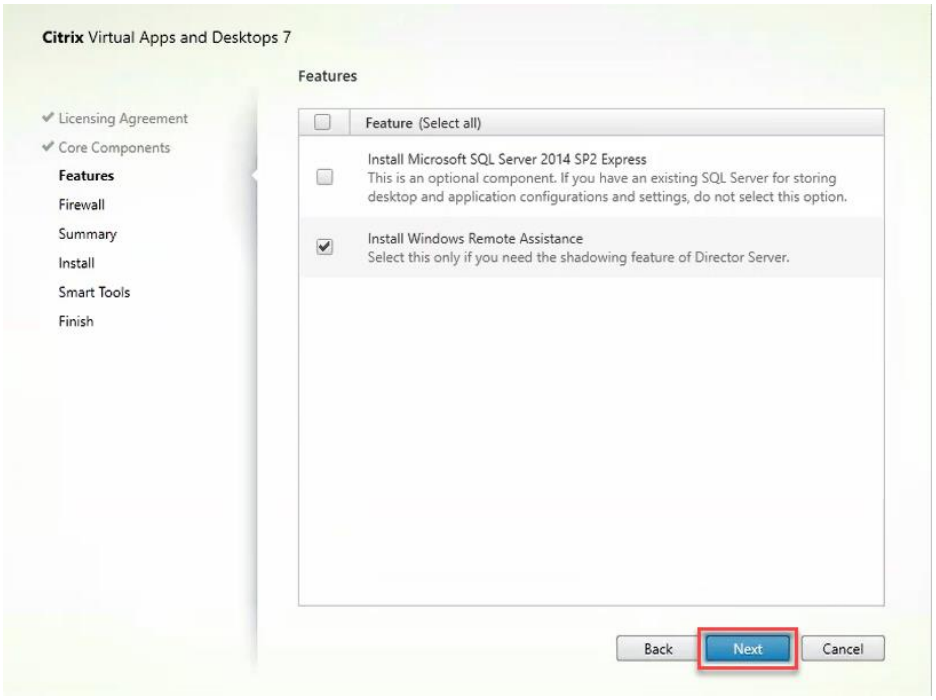

11. Select the default ports and automatically configured firewall rules.

12. Click Next.

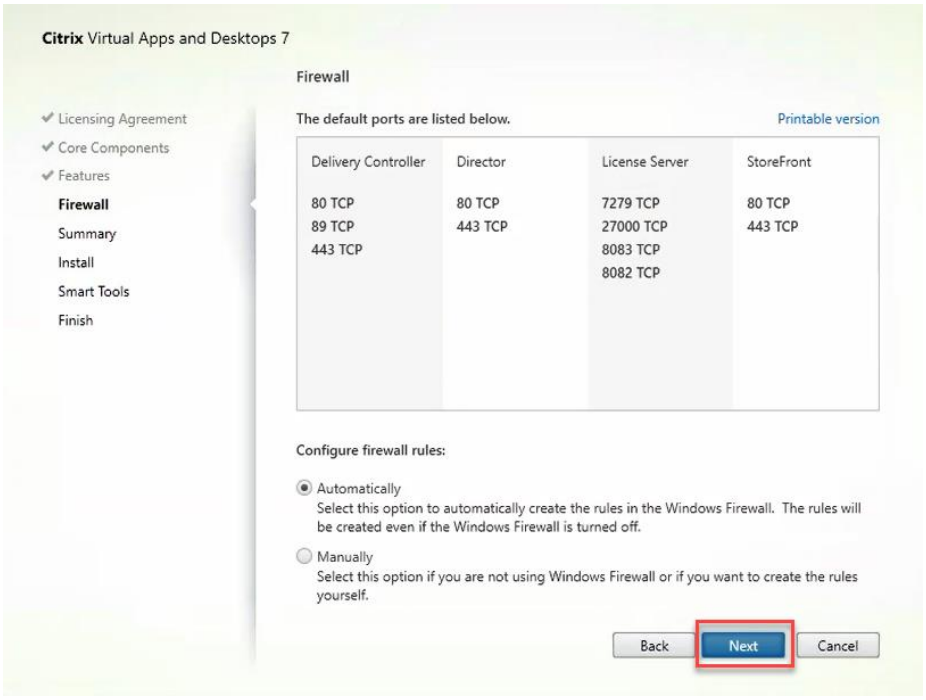

13. Click Install.

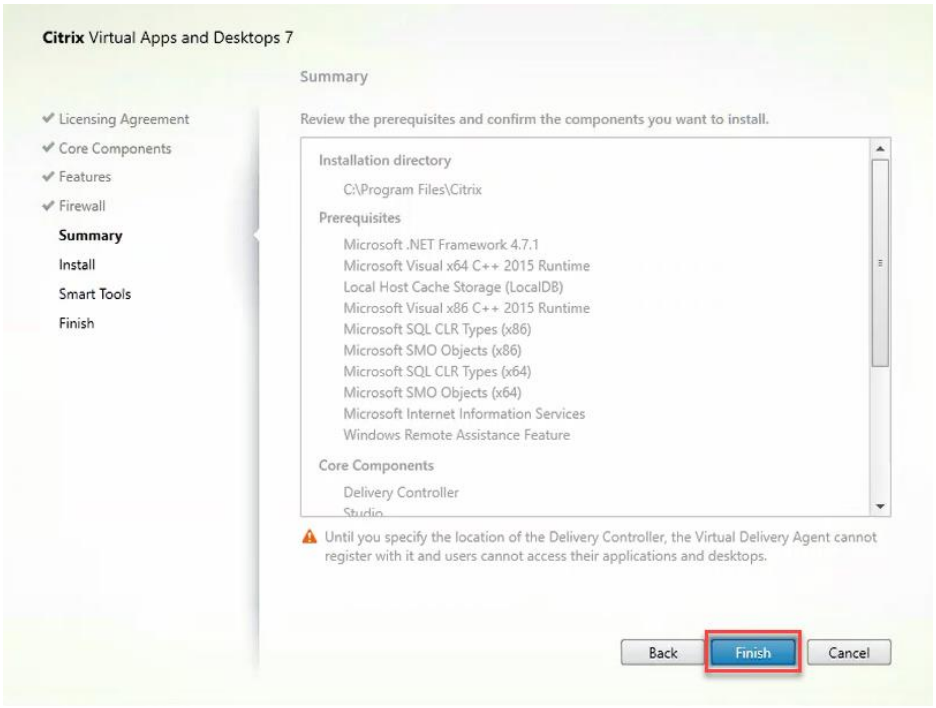

14. (Optional) Click the Call Home participation.

15. Click Next.

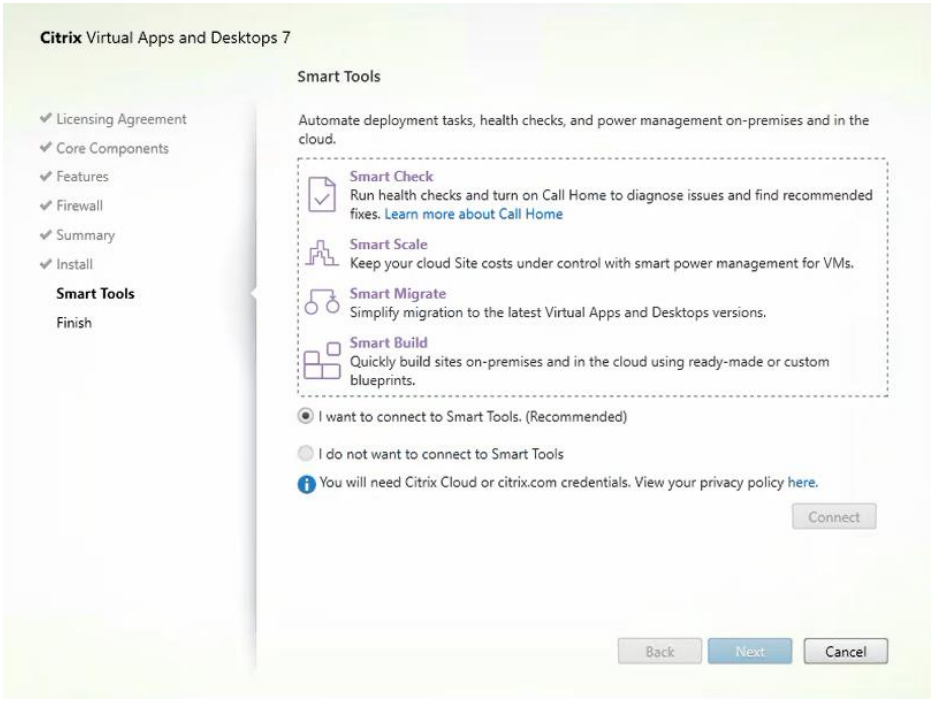

16. Click Finish to complete the installation.

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

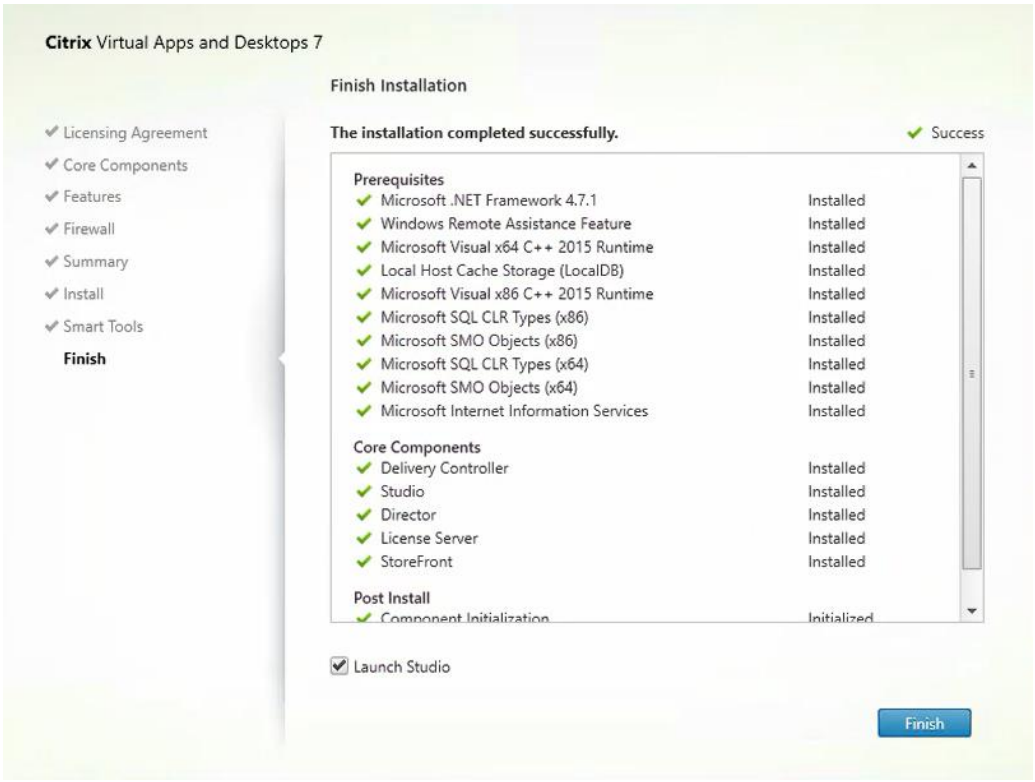

Configure the Citrix VDI Site

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

To configure Citrix VDI, complete the following steps:

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

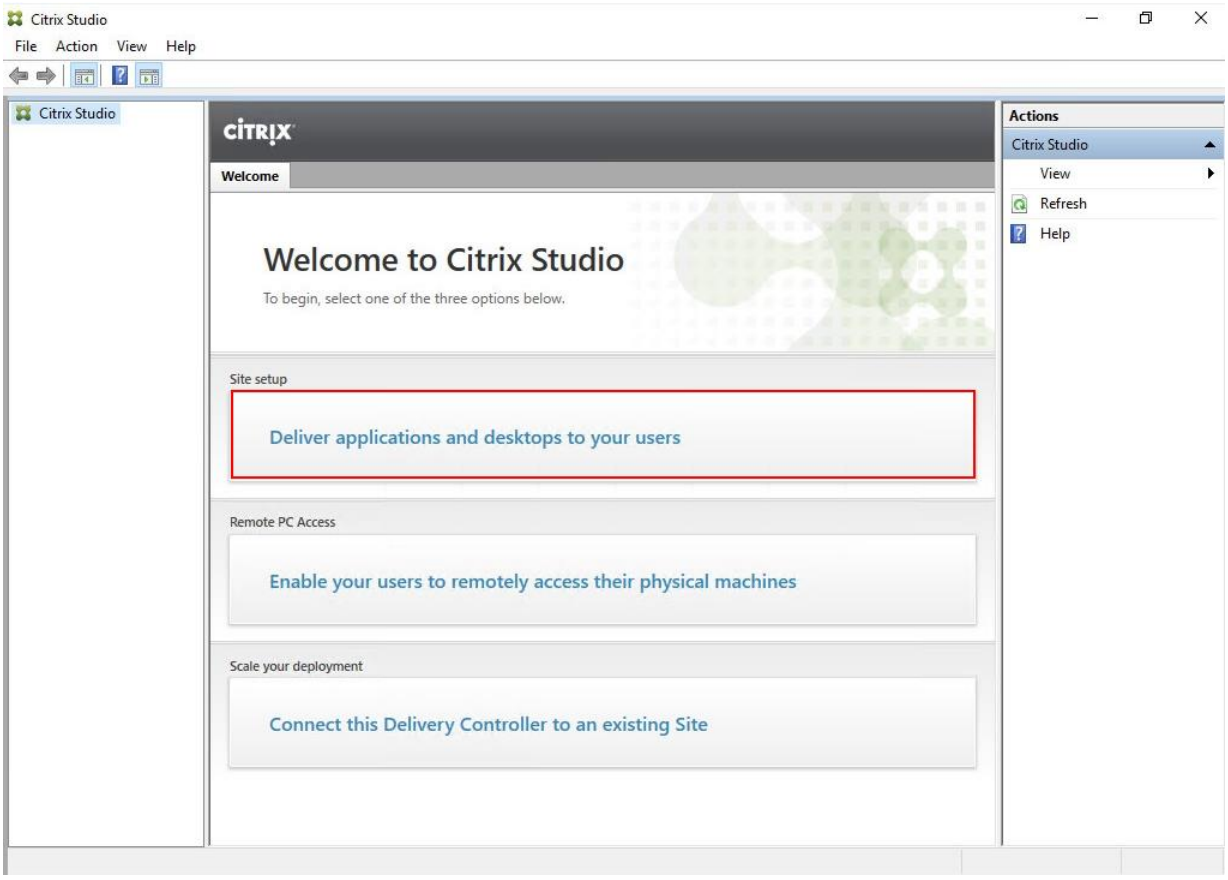

2. Select the “A fully configured, production-ready Site” radio button.
3. Enter a site name.
4. Click Next.

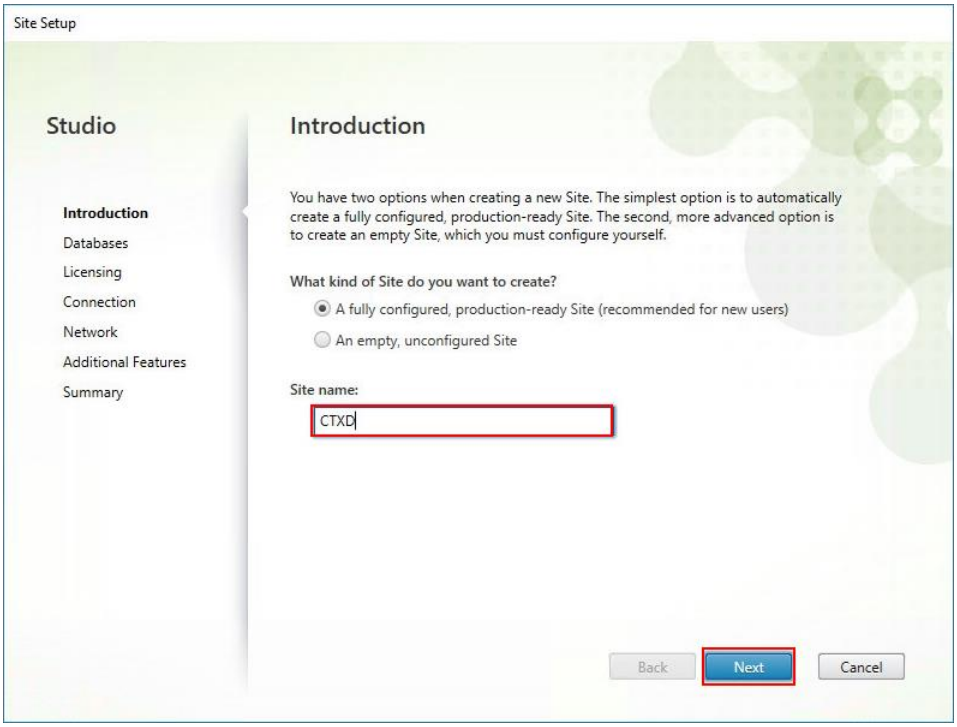

5. Provide the Database Server Locations for each data type and click Next.

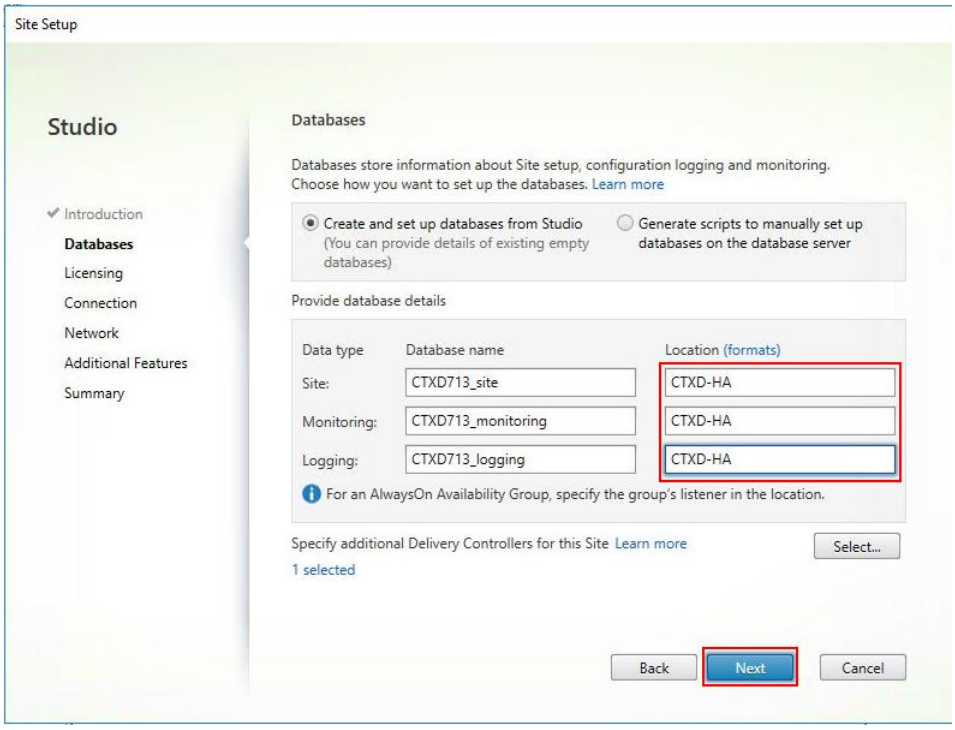

6. For an AlwaysOn Availability Group, use the group's listener DNS name.

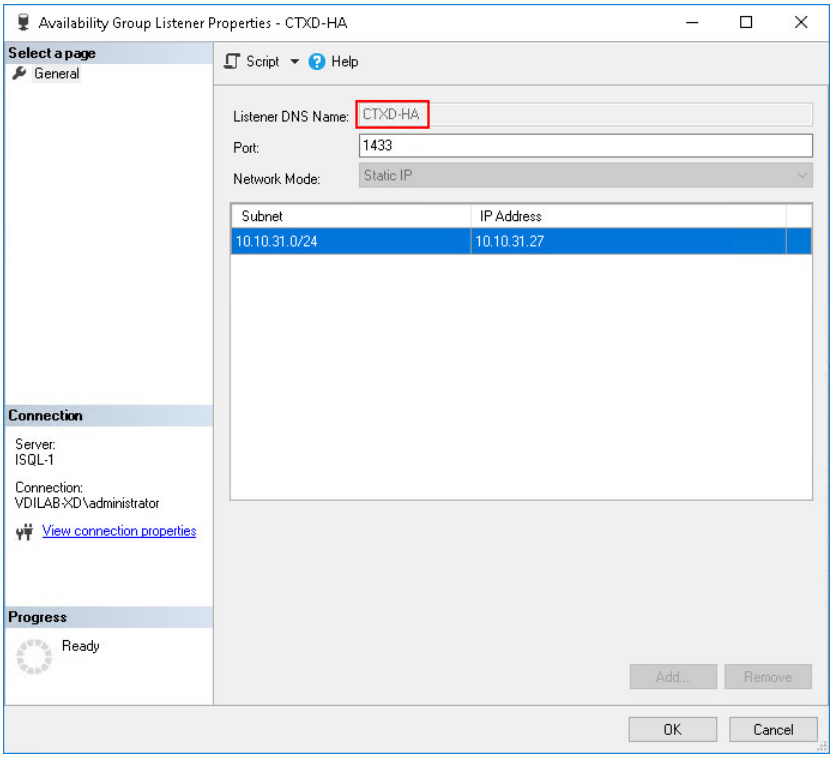

7. Provide the FQDN of the license server.

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

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

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

10. Click Next.

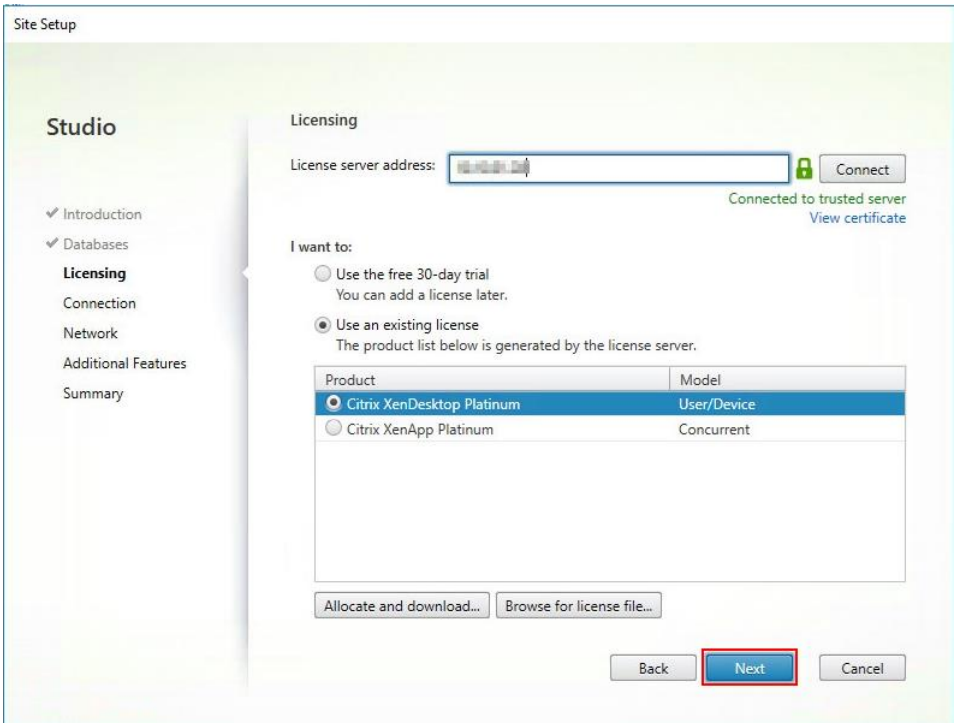

11. Select the Connection type of VMware vCenter.
12. Enter the FQDN of the vCenter server (in Server_FQDN/sdk format).
13. Enter the username (in username@domain format) for the vCenter account.
14. Provide the password for the Domain Admin account.
15. Provide a connection name.
16. Select the Other tools radio button.

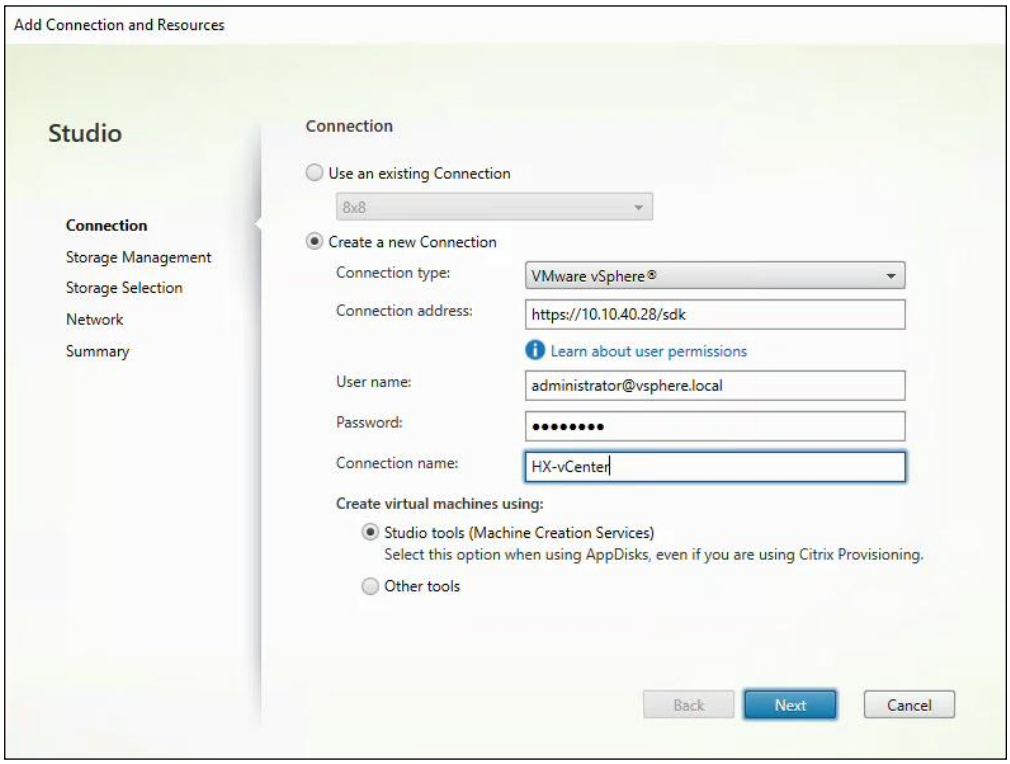

17. Click Next.
18. Select HyperFlex Cluster that will be used by this connection.
19. Check Studio Tools radio button required to support desktop provisioning task by this connection.
20. Click Next.

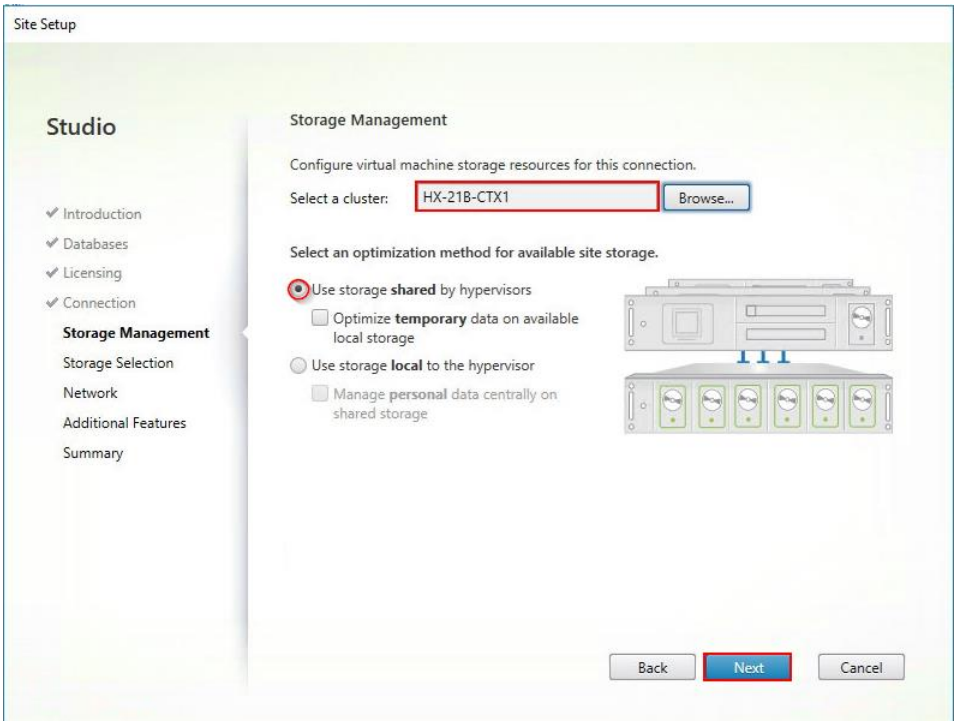

21. Make Storage selection to be used by this connection.

22. Click Next.

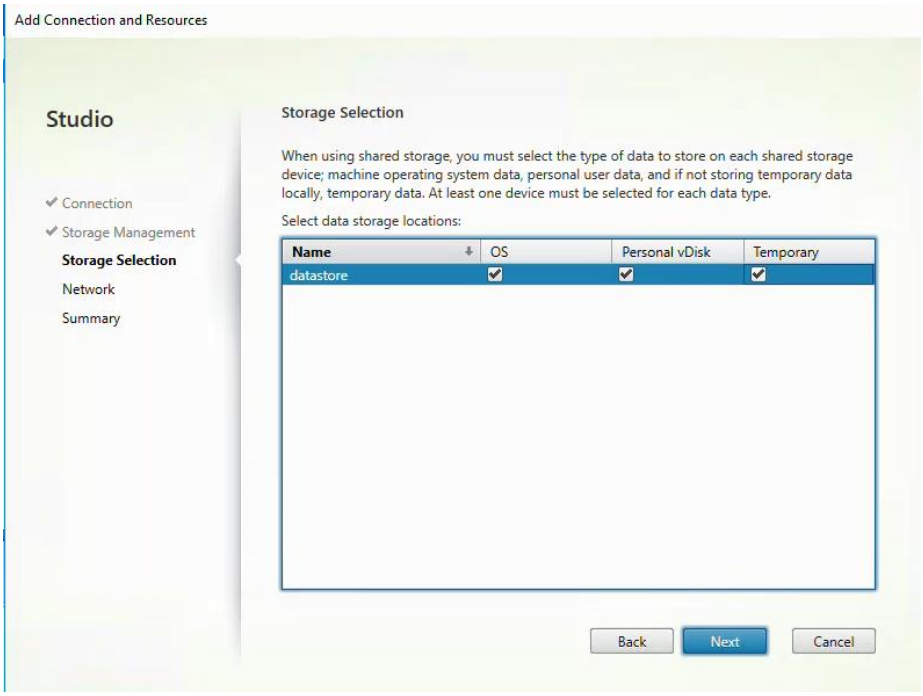

23. Make Network selection to be used by this connection.

24. Click Next.

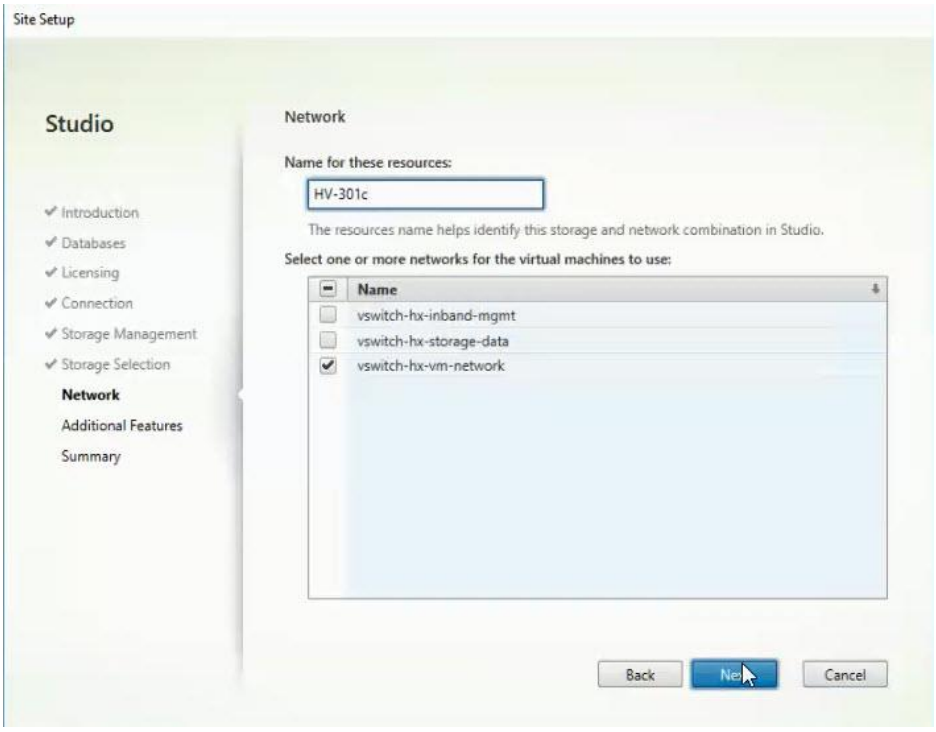

25. Select Additional features.

26. Click Next.

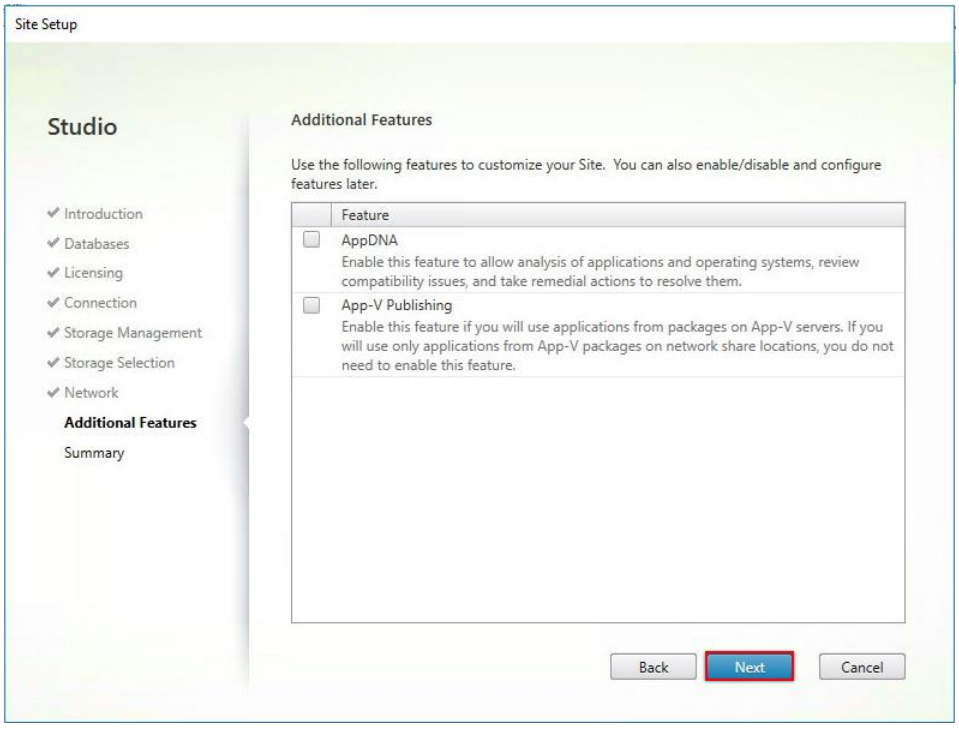

27. Review Site configuration Summary and click Finish.

Configure the Citrix VDI Site Administrators

To configure the Citrix VDI site administrators, complete the following steps:

1. Connect to the Citrix VDI server and open Citrix Studio Management console.

2. From the Configuration menu, right-click Administrator and select Create Administrator from the drop-down list.

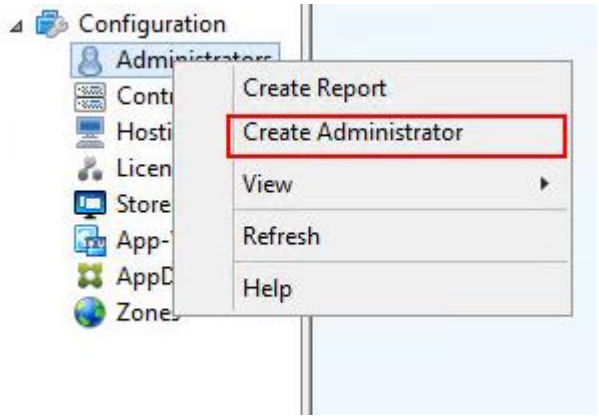

3. Select/Create appropriate scope and click Next.

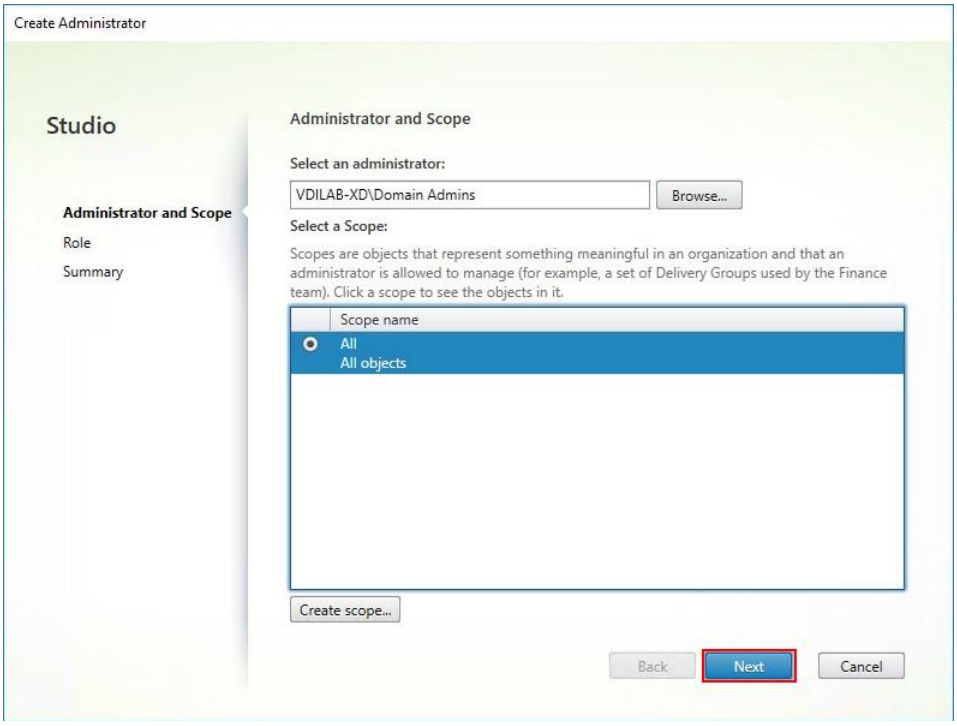

4. Choose an appropriate Role.

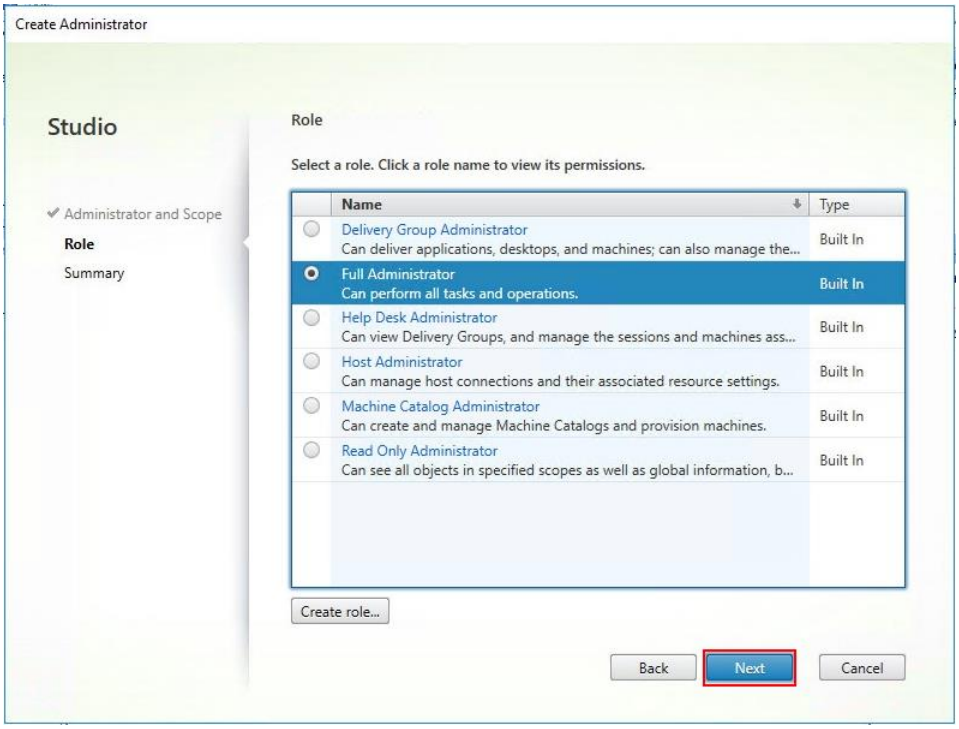

5. Review the Summary, check Enable administrator, and click Finish.

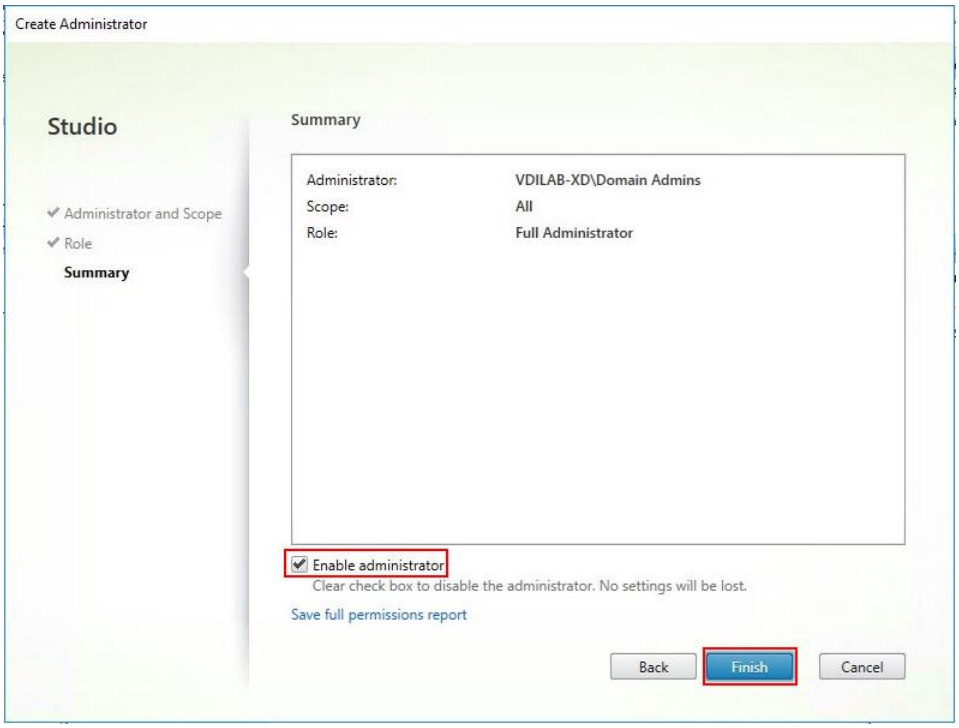

Configure Additional Desktop Controller

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 Citrix Desktop controllers, complete the following steps:

1. To begin the installation of the second Delivery Controller, connect to the second Citrix VDI server and launch the installer from the Citrix Virtual Apps and Desktops ISO.

2. Click Start.

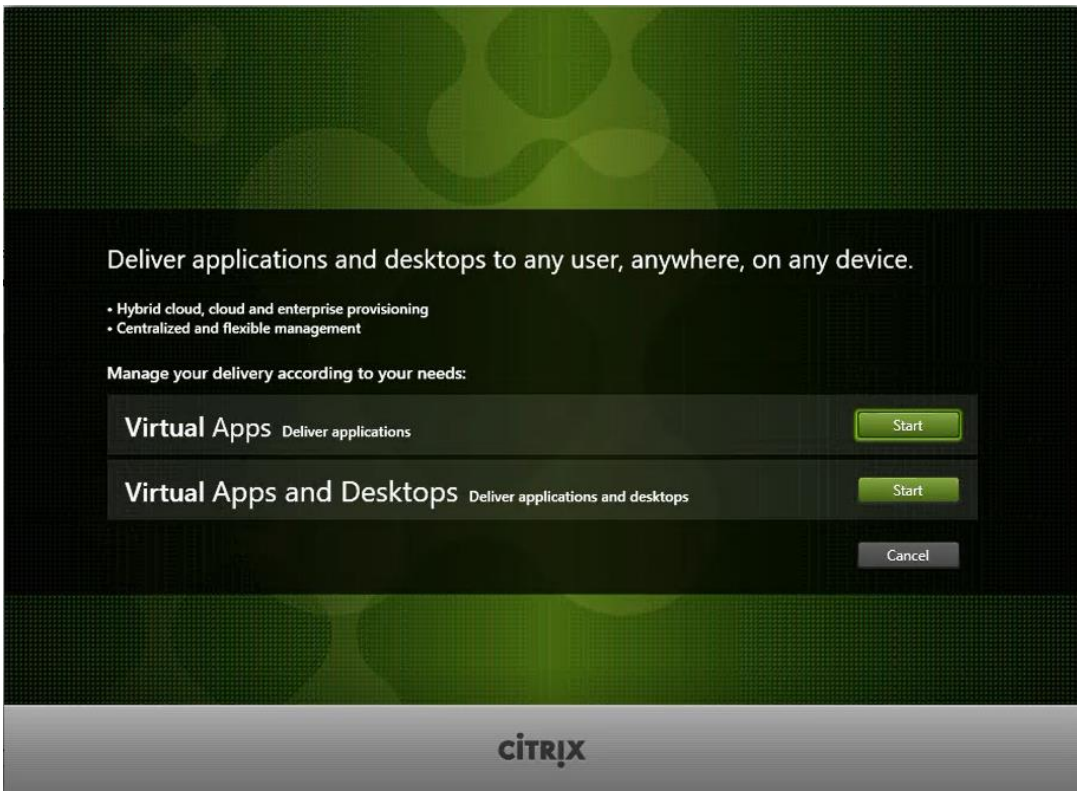

3. Click Delivery Controller.

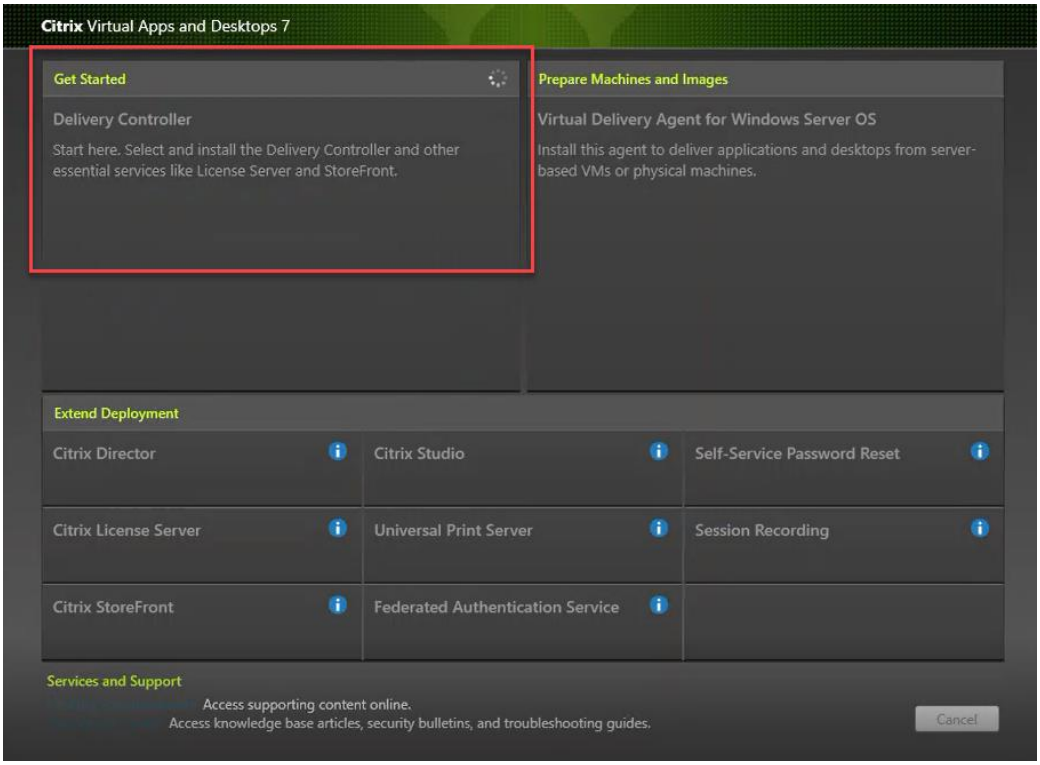

4. Repeat the same steps used to install the first Delivery Controller, including the step of importing an SSL certificate for HTTPS between the controller and Hyper-V.

5. Review the Summary configuration.

6. Click Install.

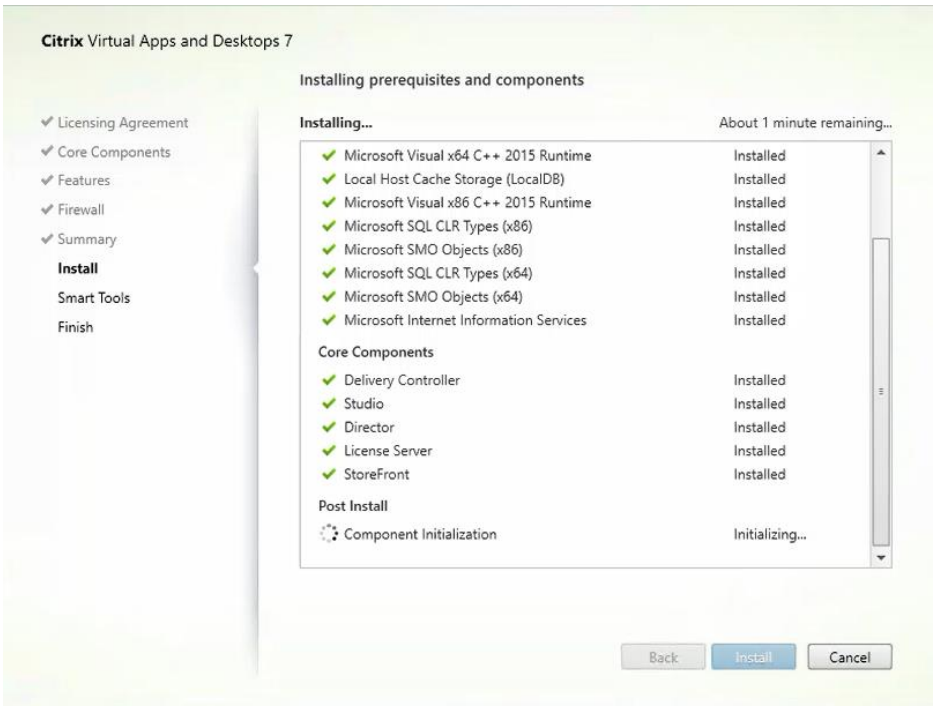

7. (Optional) Click the "I want to participate in Call Home."

8. Click Next.

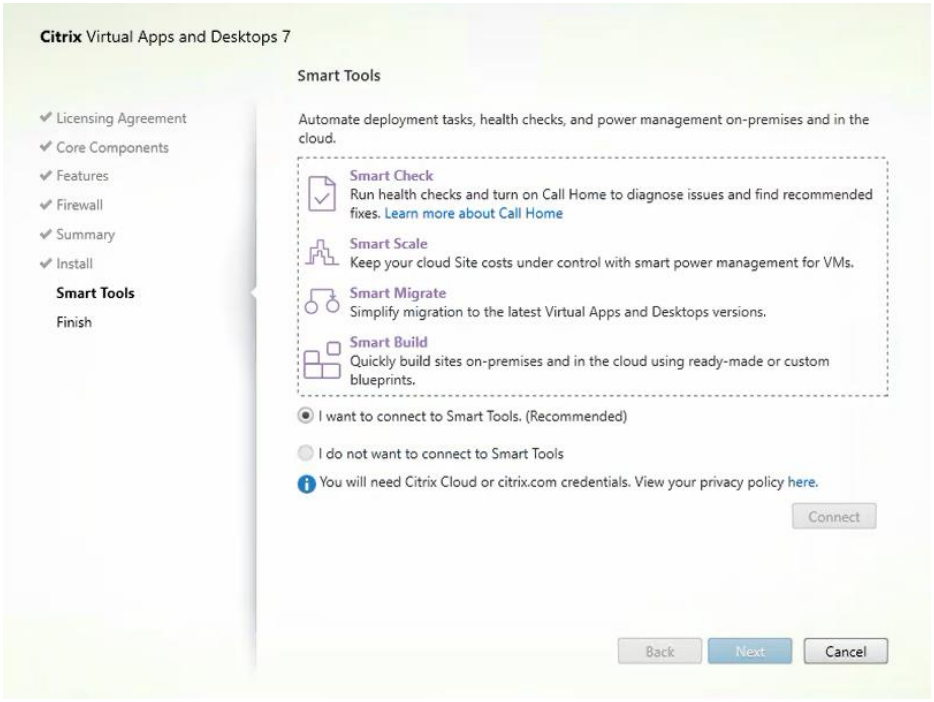

9. Verify the components installed successfully.

10. Click Finish.

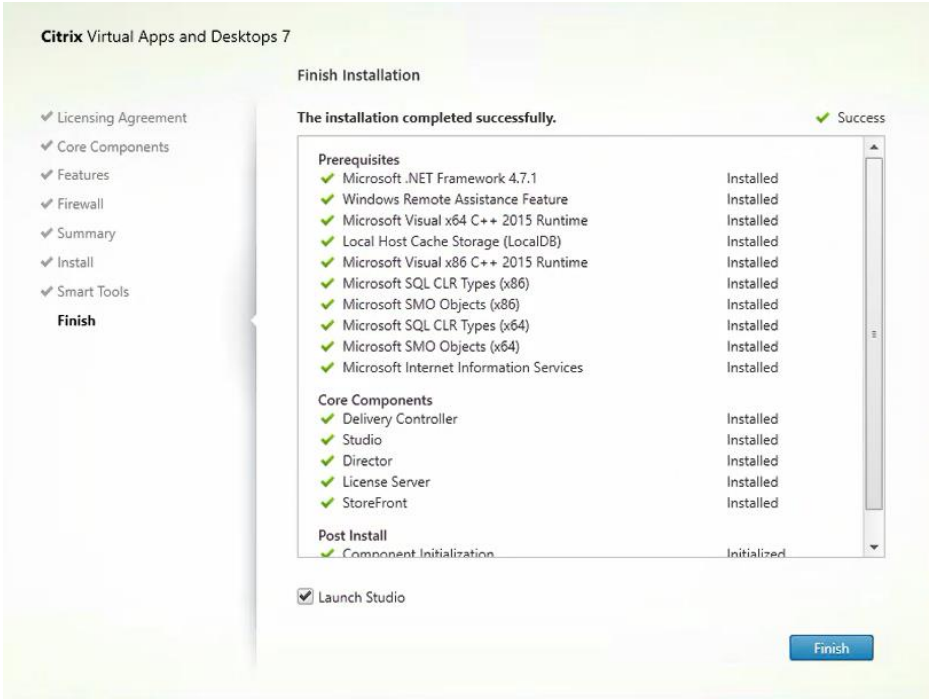

Add the Second Delivery Controller to the Citrix Desktop Site

To add the second Delivery Controller to the Citrix Desktop Site, complete the following steps:

1. In Desktop Studio click the "Connect this Delivery Controller to an existing Site" button.

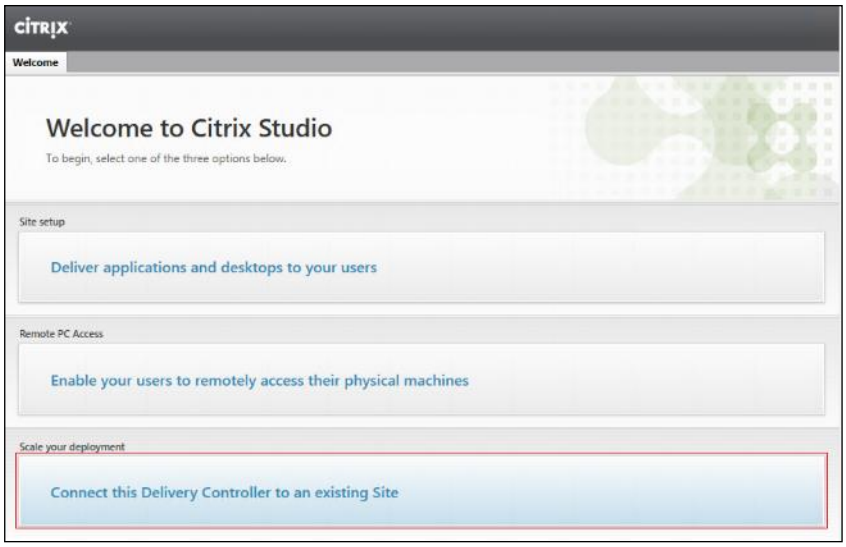

2. Enter the FQDN of the first delivery controller.

3. Click OK.

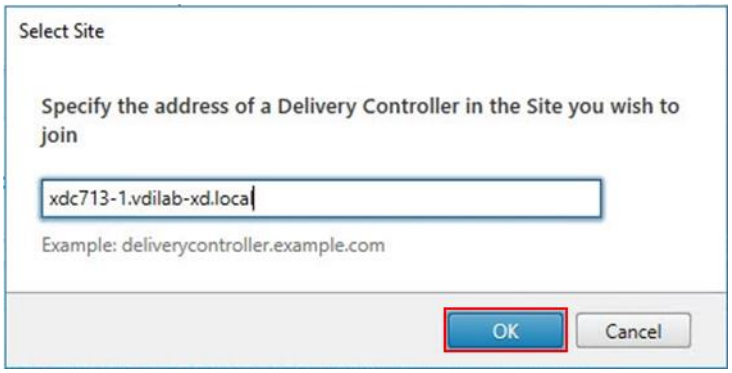

4. Click Yes to allow the database to be updated with this controller's information automatically.

5. When complete, test the site configuration and verify the Delivery Controller has been added to the list of Controllers.

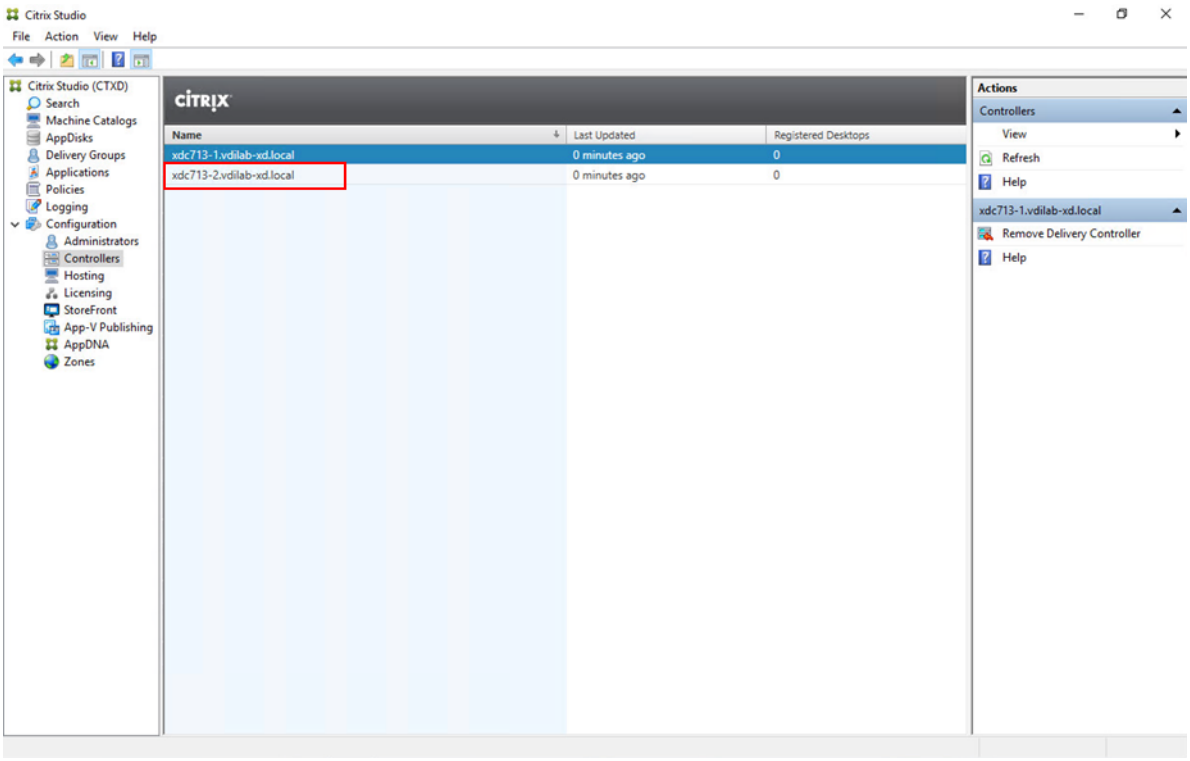

Install and Configure StoreFront

Citrix StoreFront stores aggregate desktops and applications from Citrix VDI sites, making resources readily available to users.

In this CVD, we created two StoreFront servers on dedicated virtual machines.

To install and configure StoreFront, complete the following steps:

1. To begin the installation of the StoreFront, connect to the first StoreFront server and launch the installer from the Citrix Desktop 1808 ISO.

2. Click Start.

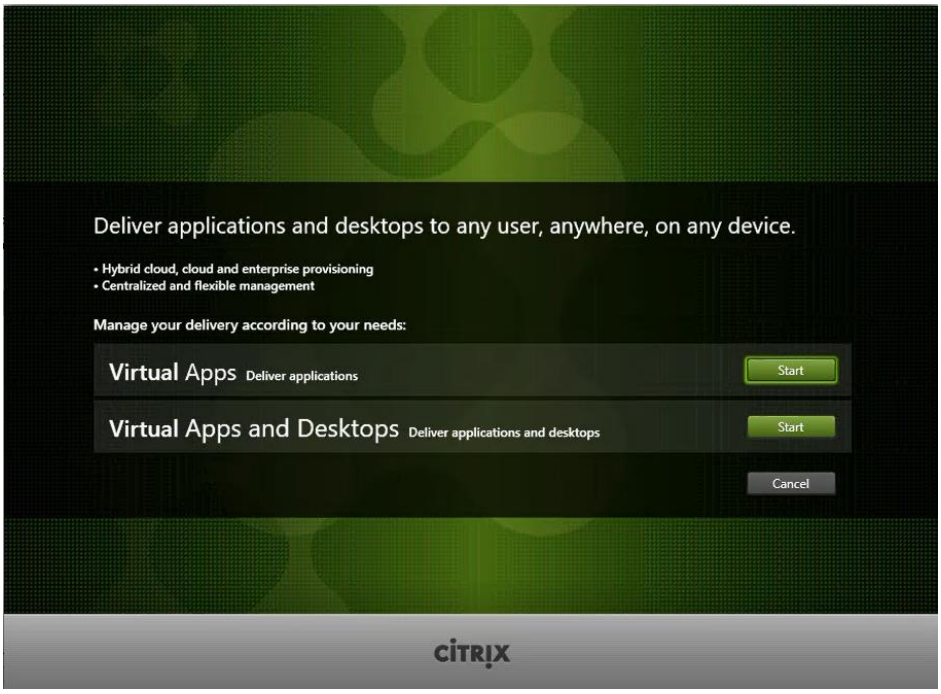

3. Click Extend Deployment Citrix StoreFront.

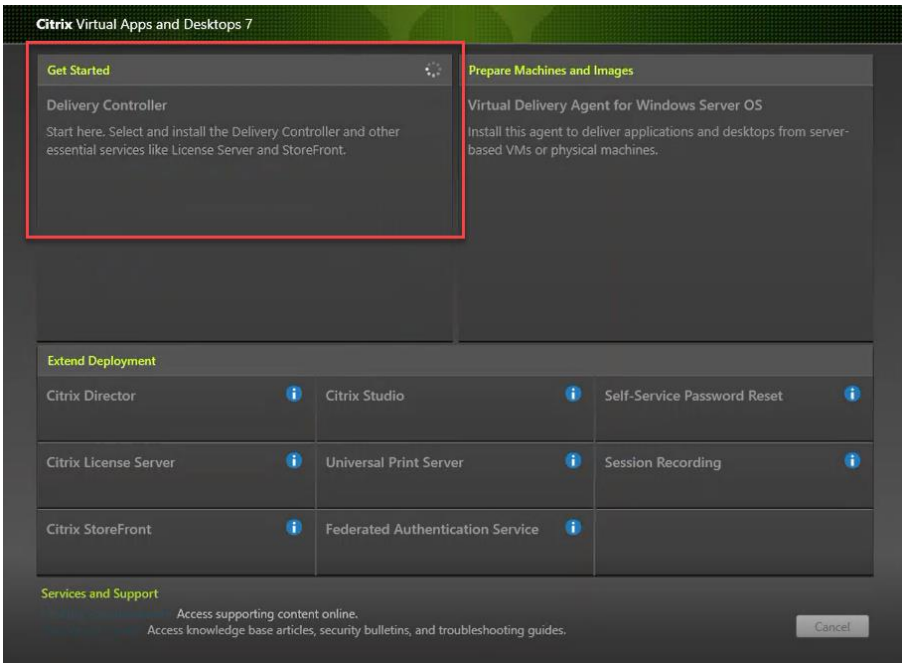

4. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

5. Click Next.

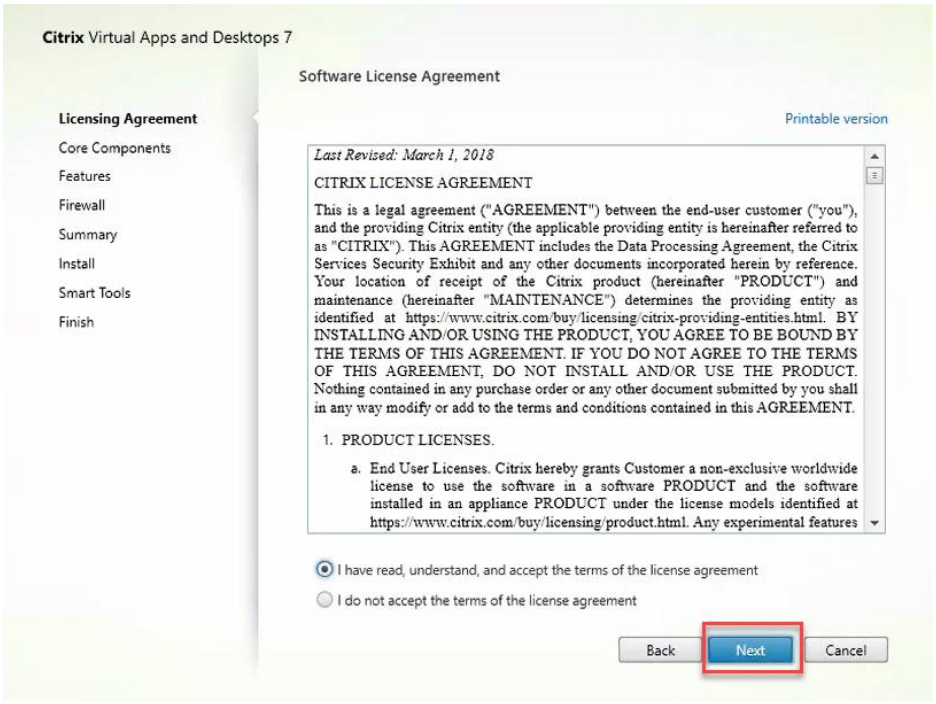

6. Select Storefront and Click Next.

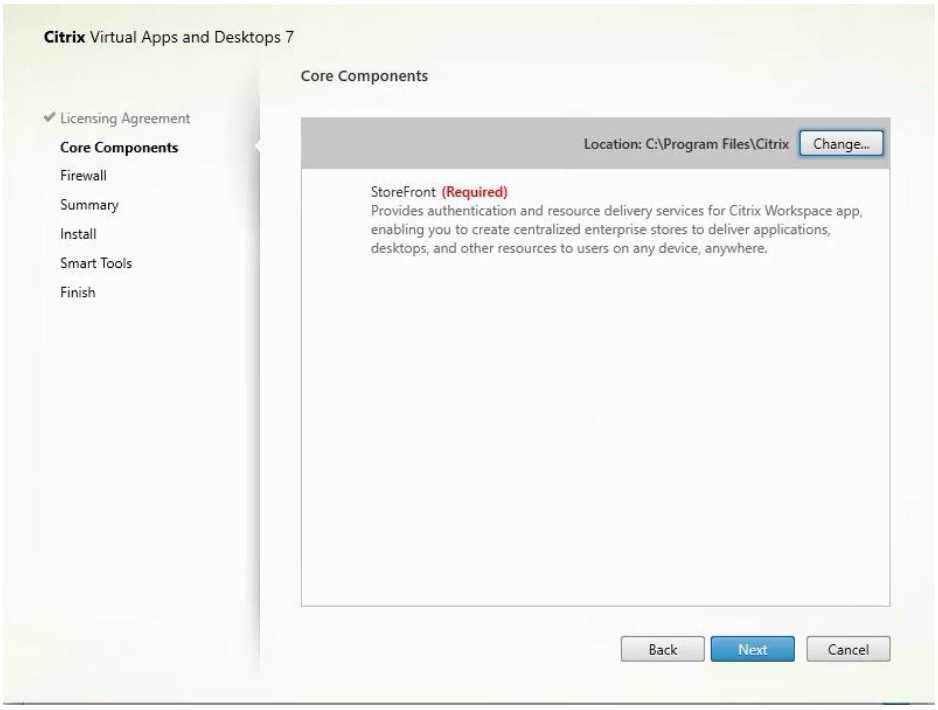

7. Select the default ports and automatically configured firewall rules.

8. Click Next.

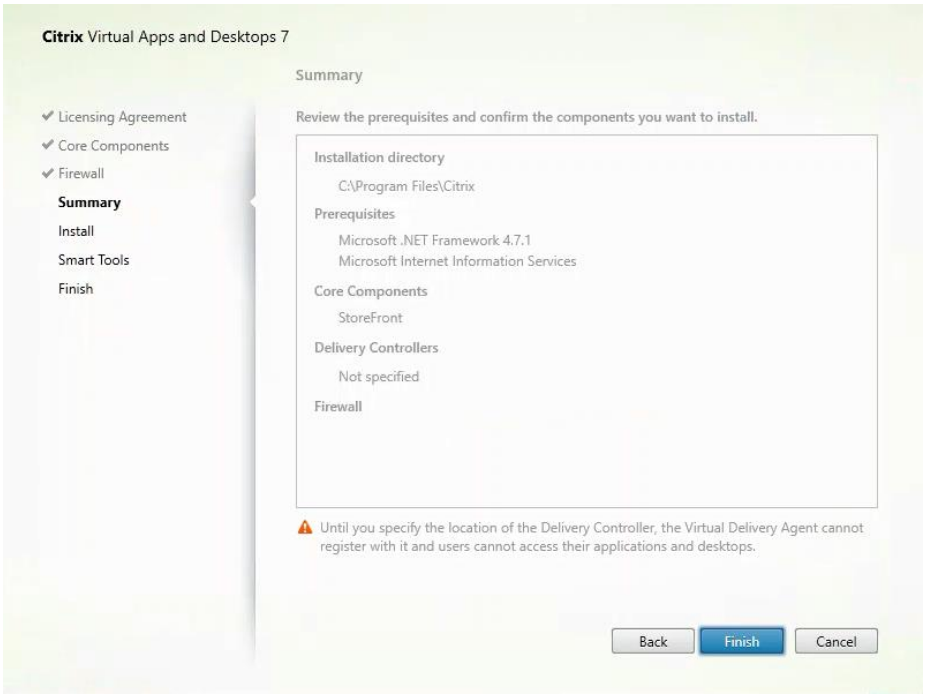

9. Click Install.

10. (Optional) Click "I want to participate in Call Home."

11. Click Next.

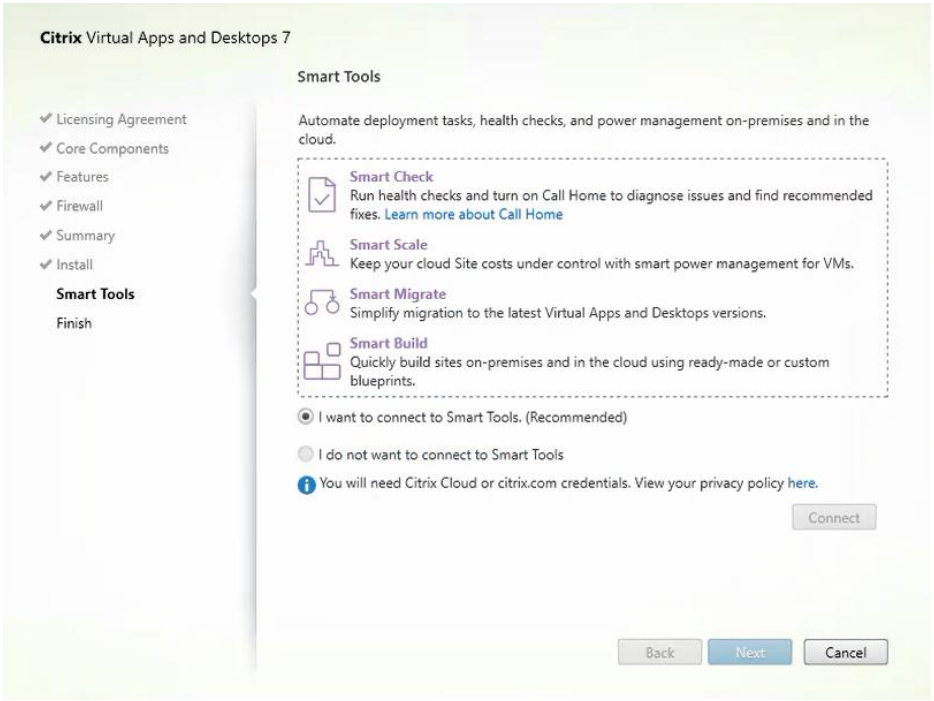

12. Check "Open the StoreFront Management Console."

13. Click Finish.

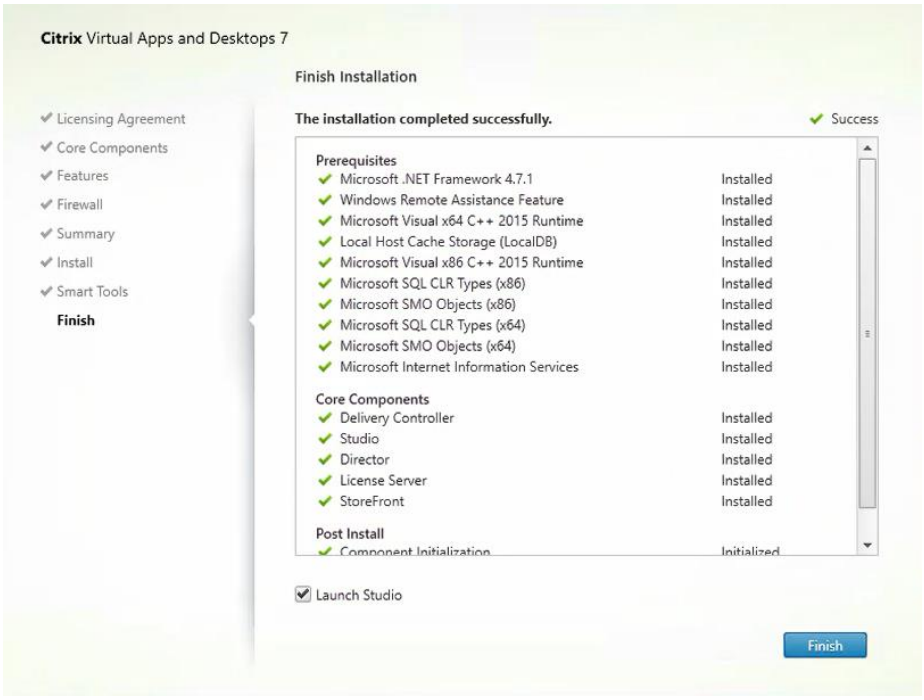

14. Click Create a new deployment.

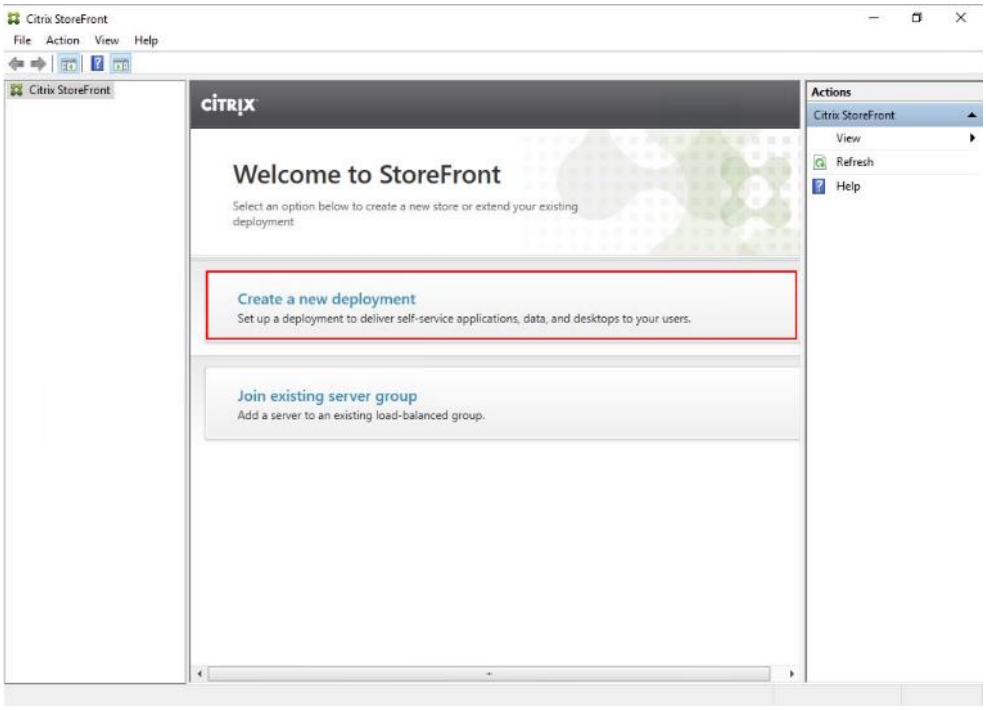

15. Specify the URL of the StoreFront server and click Next.

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

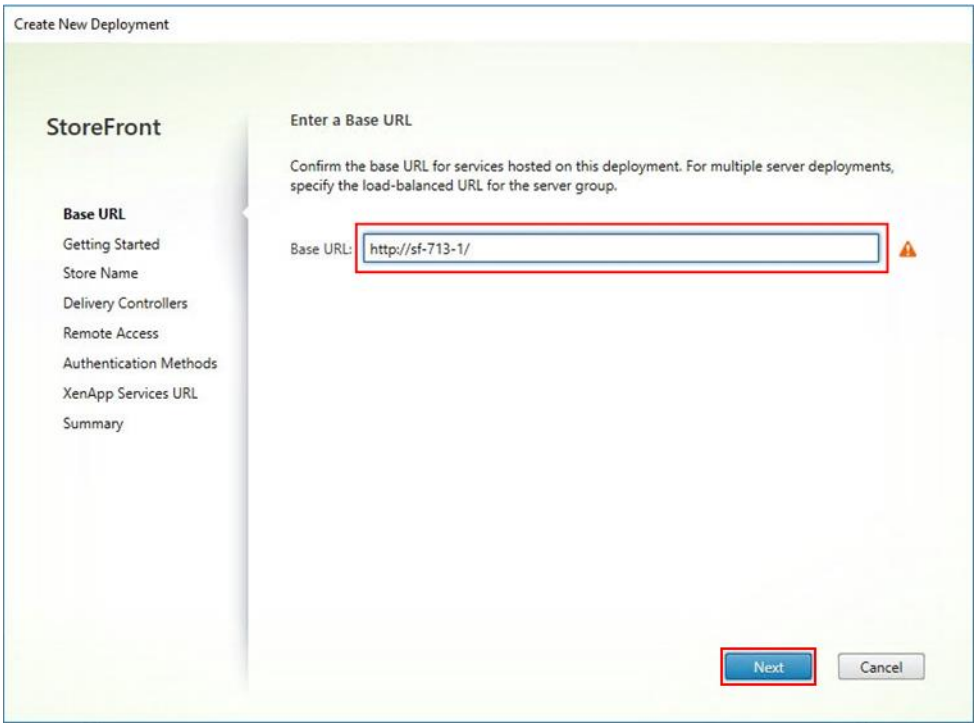

16. Click Next.

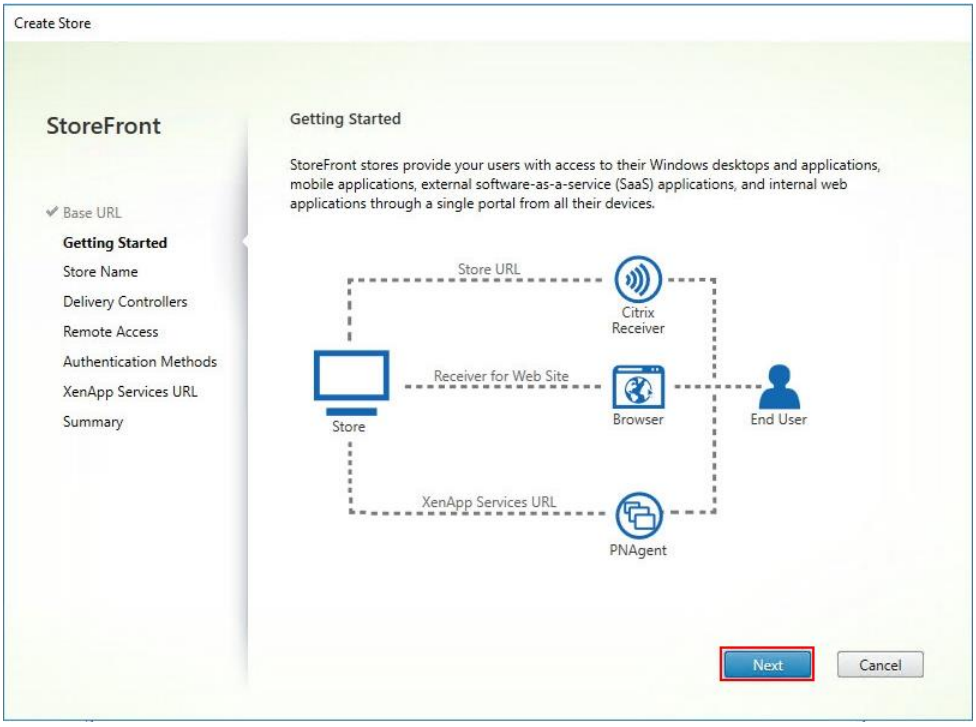

17. Specify a name for your store and click Next.

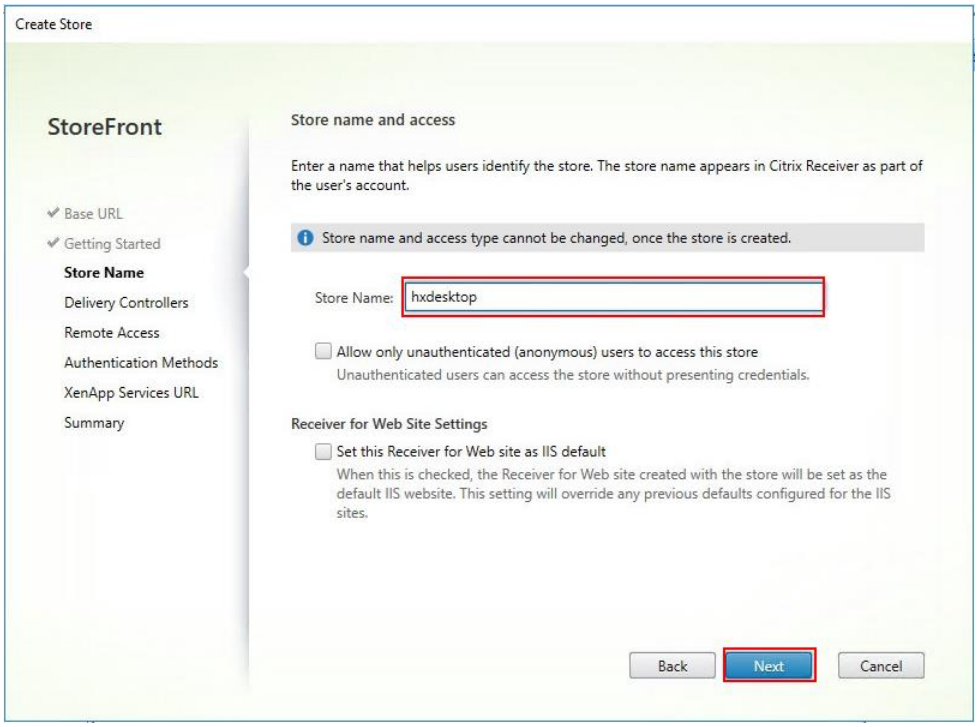

18. Add the required Delivery Controllers to the store and click Next.

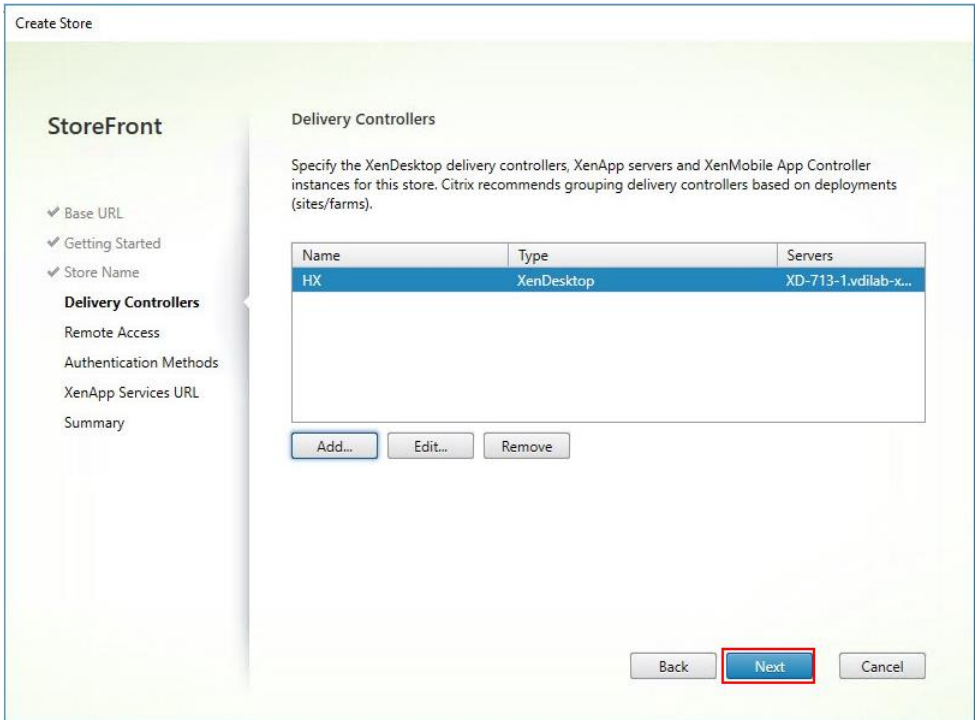

19. Specify how connecting users can access the resources, in this environment only local users on the internal network are able to access the store, and click Next.

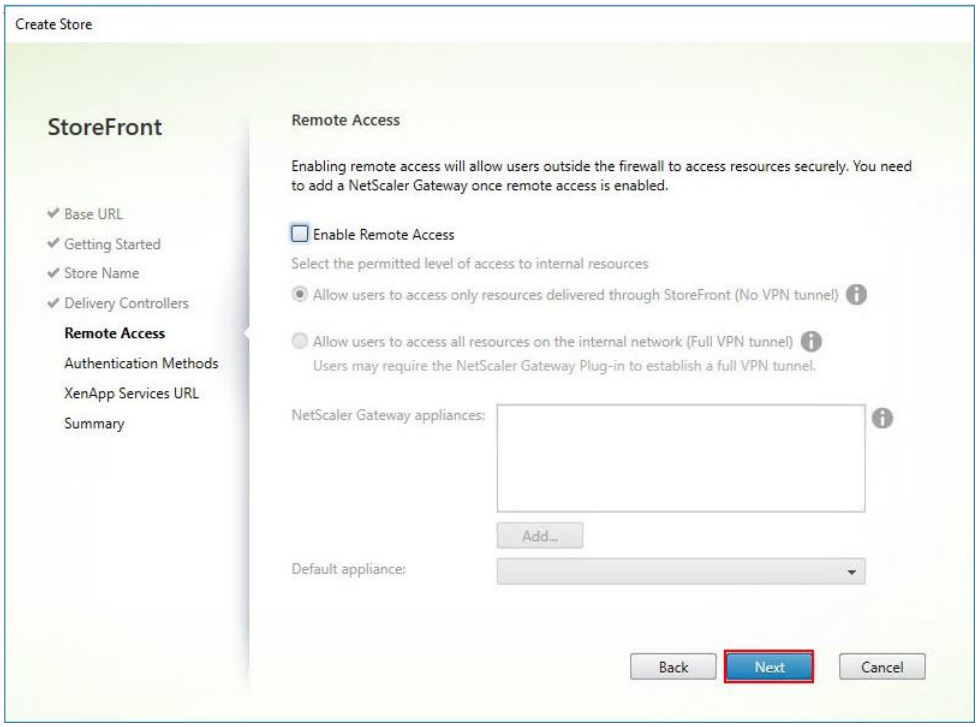

20. On the “Authentication Methods” page, select the methods your users will use to authenticate to the store and click Next. You can select from the following methods as shown below:

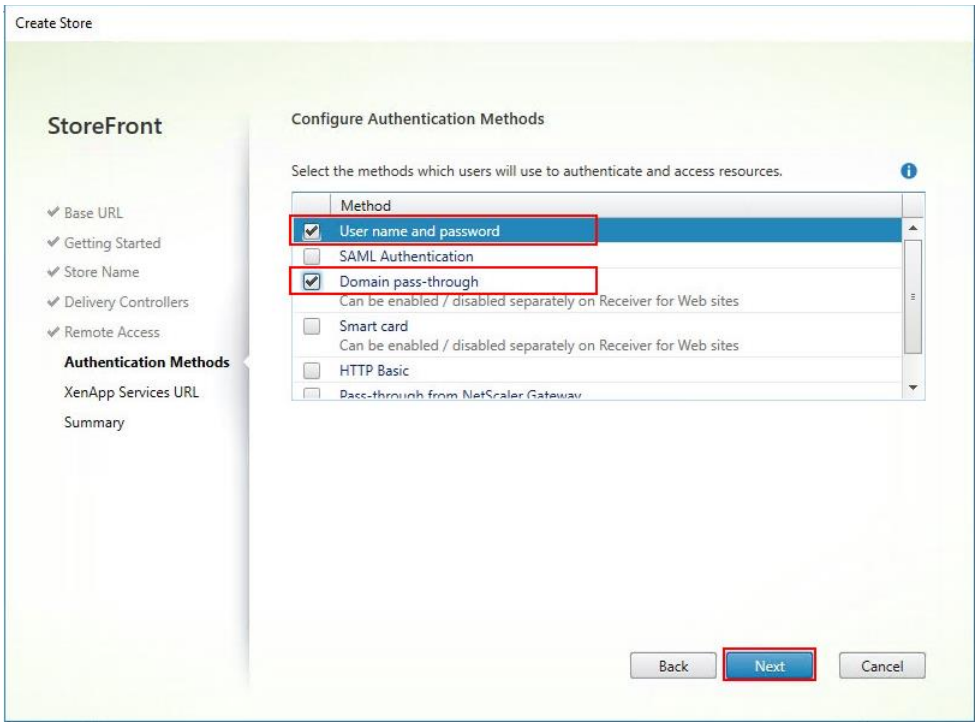

21. Username and password: Users enter their credentials and are authenticated when they access their stores.

22. Domain pass-through: Users authenticate to their domain-joined Windows computers and their credentials are used to log them on automatically when they access their stores.

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

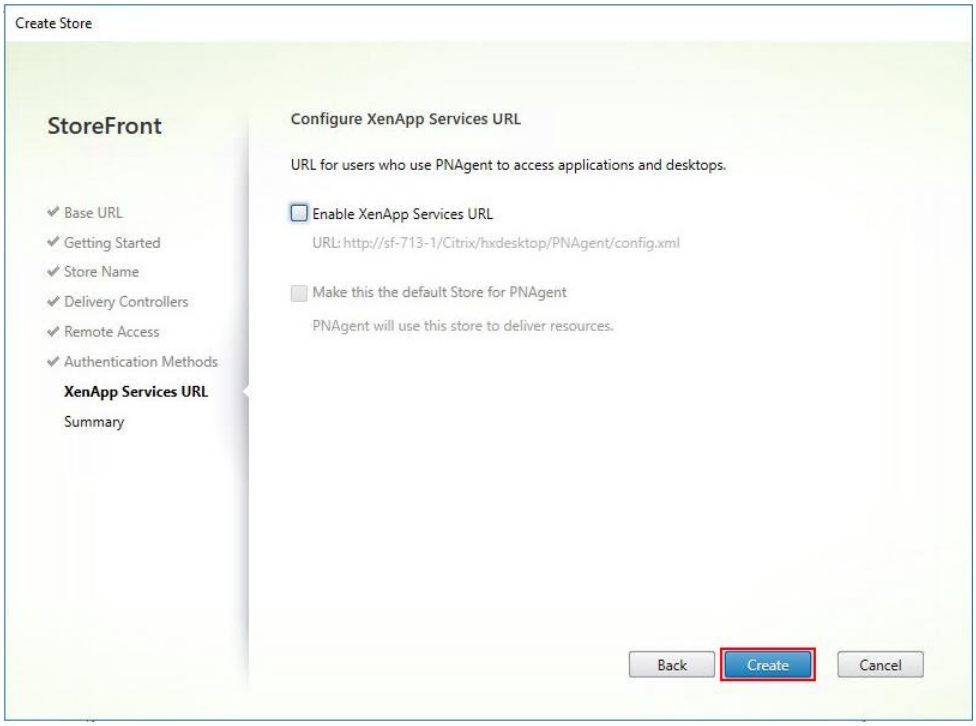

24. After creating the store click Finish.

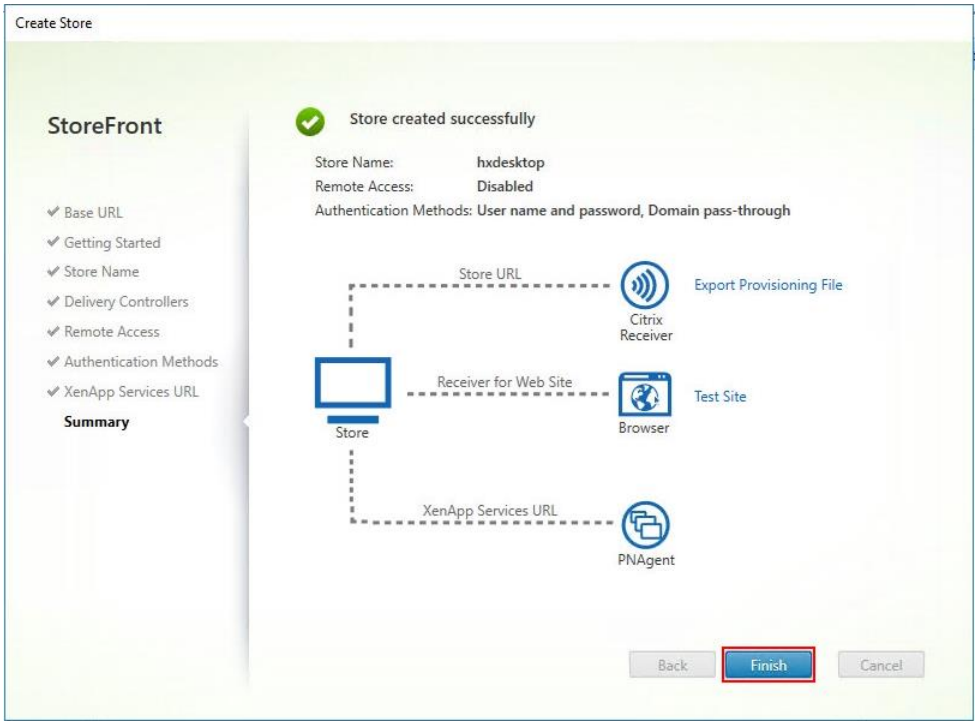

Additional StoreFront Configuration

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

To configure additional StoreFront server, complete the following steps:

1. To begin the installation of the second StoreFront, connect to the second StoreFront server and launch the installer from the Citrix Citrix VDI ISO.

2. Click Start.

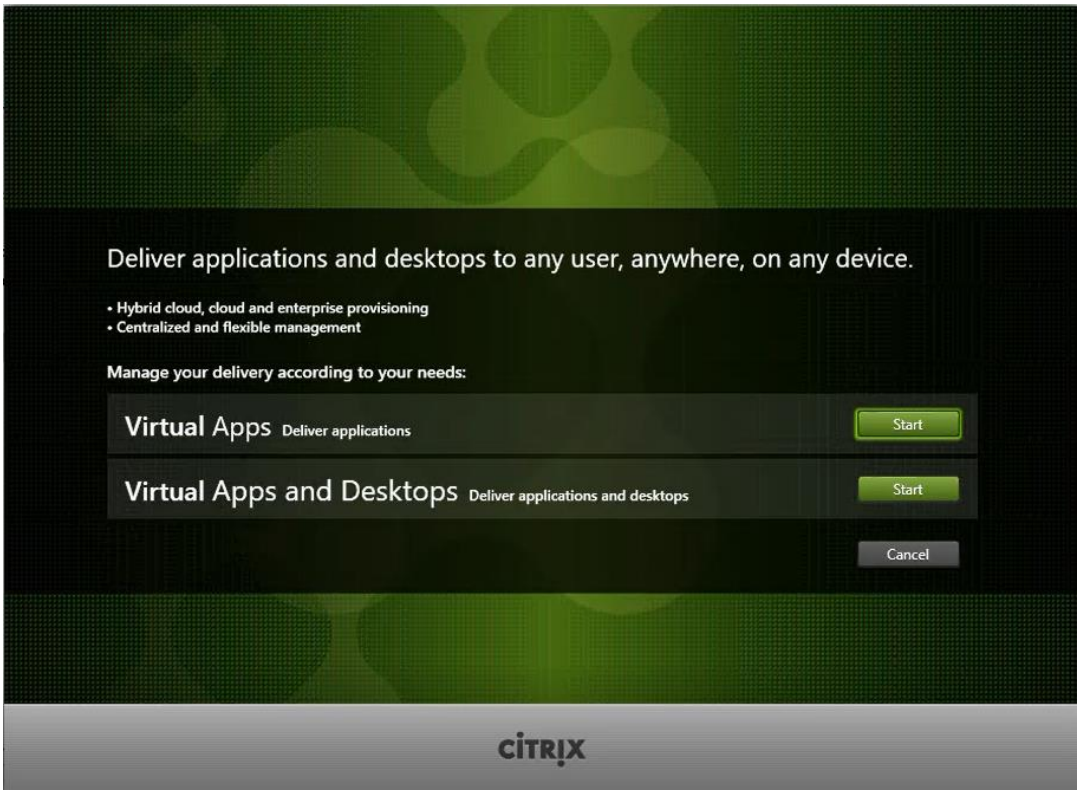

3. Click Extended Deployment Citrix StoreFront.

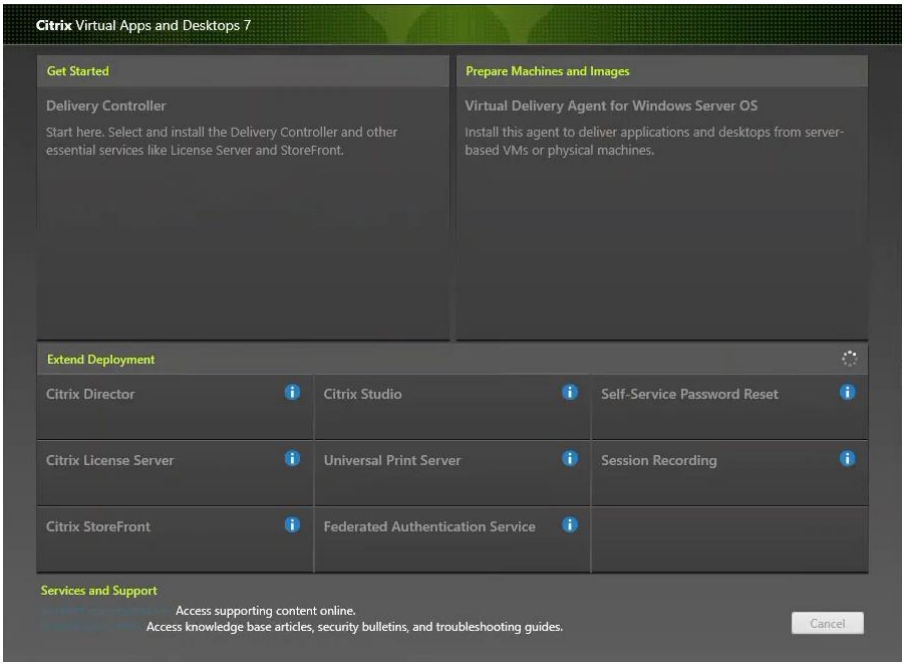

4. Repeat the same steps used to install the first StoreFront.
5. Review the Summary configuration.
6. Click Install.

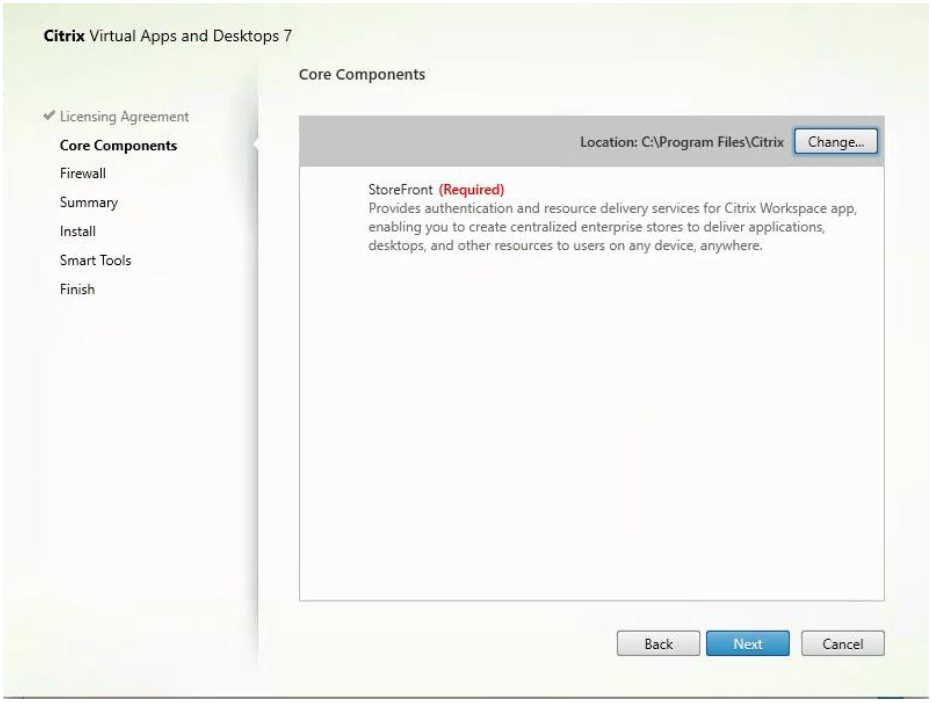

7. (Optional) Click “I want to participate in Call Home.”
8. Click Next.

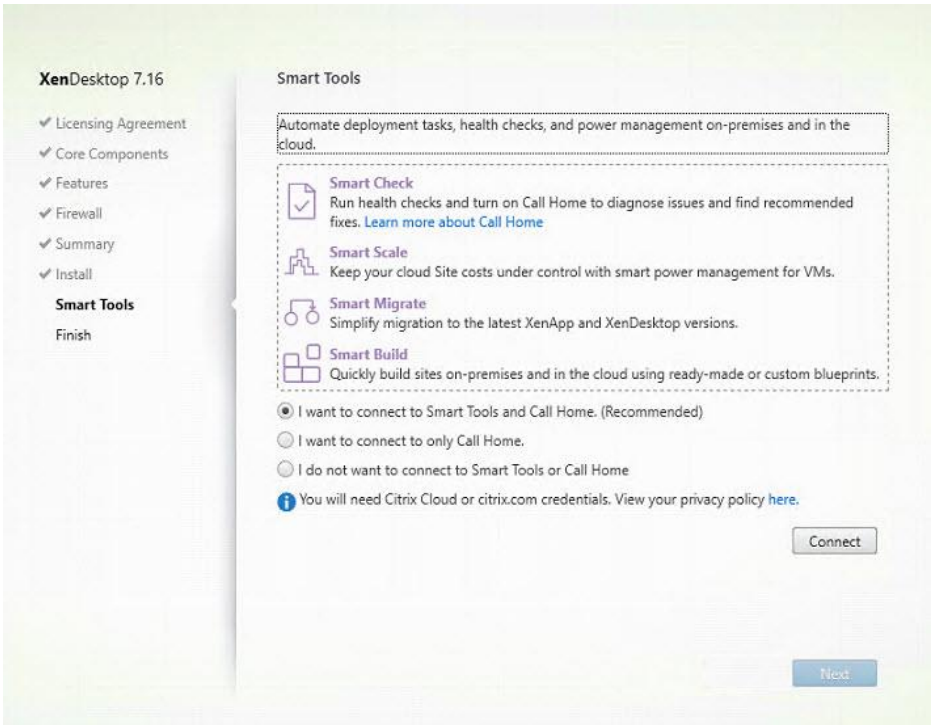

9. Check “Open the StoreFront Management Console."

10. Click Finish.

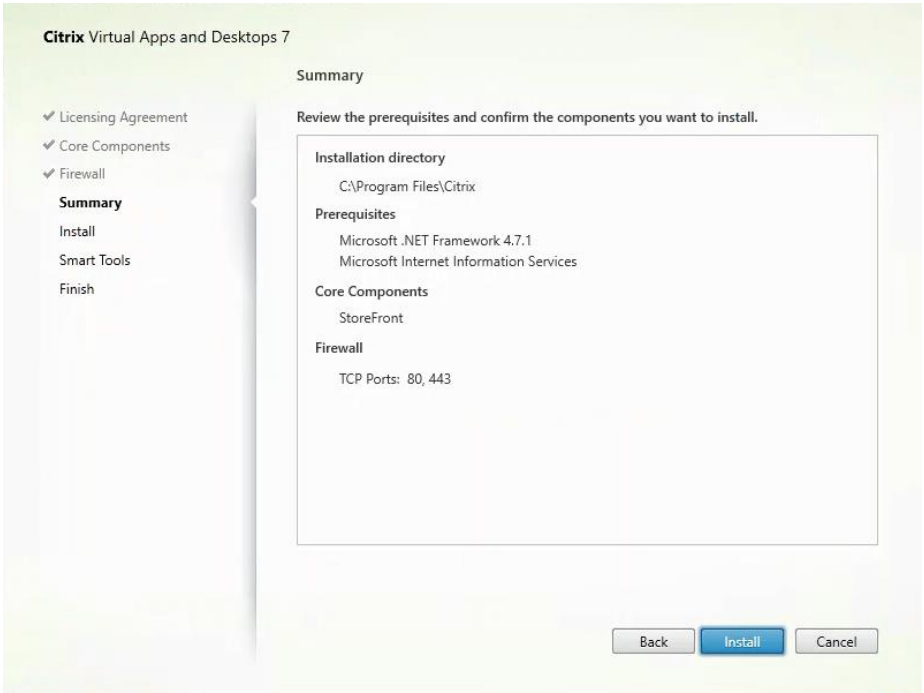

To configure the second StoreFront if used, complete the following steps:

1. From the StoreFront Console on the second server select “Join existing server group.”

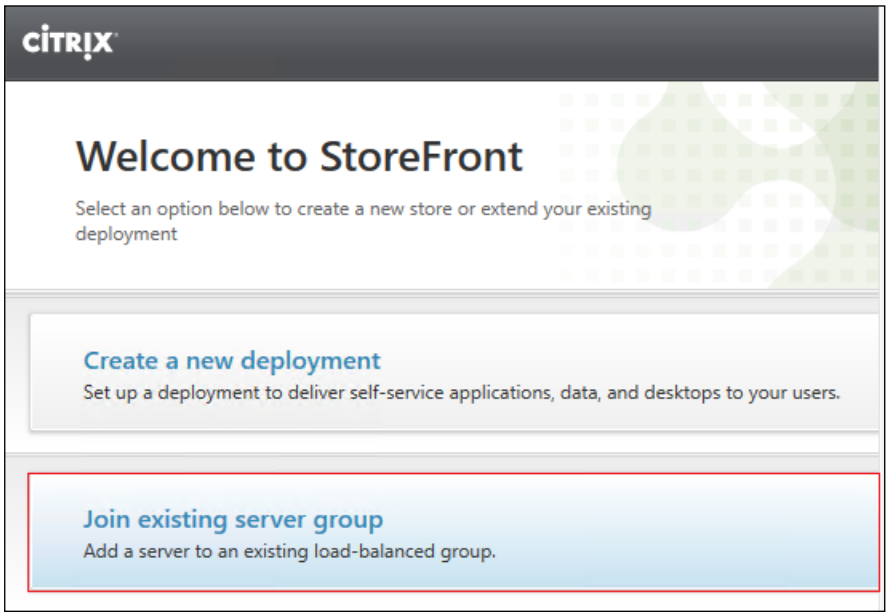

2. In the Join Server Group dialog, enter the name of the first Storefront server.

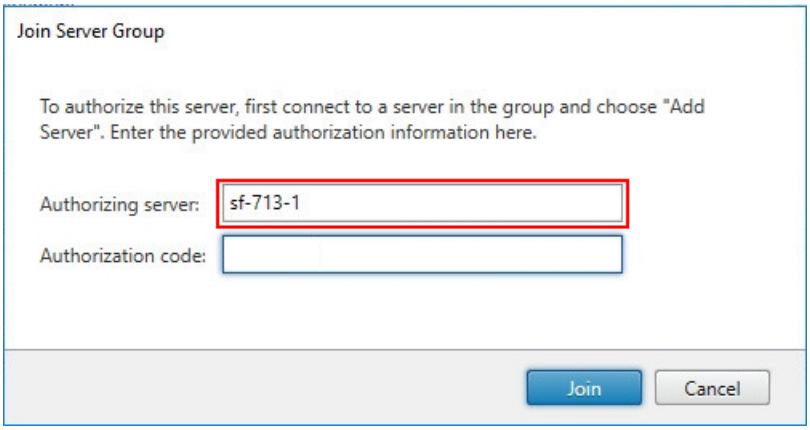

3. Before the additional StoreFront server can join the server group, you must connect to the first Storefront server, add the second server, and obtain the required authorization information.

4. Connect to the first StoreFront server.

5. Using the StoreFront menu on the left, you can scroll through the StoreFront management options.

6. Select Server Group from the menu.

7. To add the second server and generate the authorization information that allows the additional StoreFront server to join the server group, select Add Server.

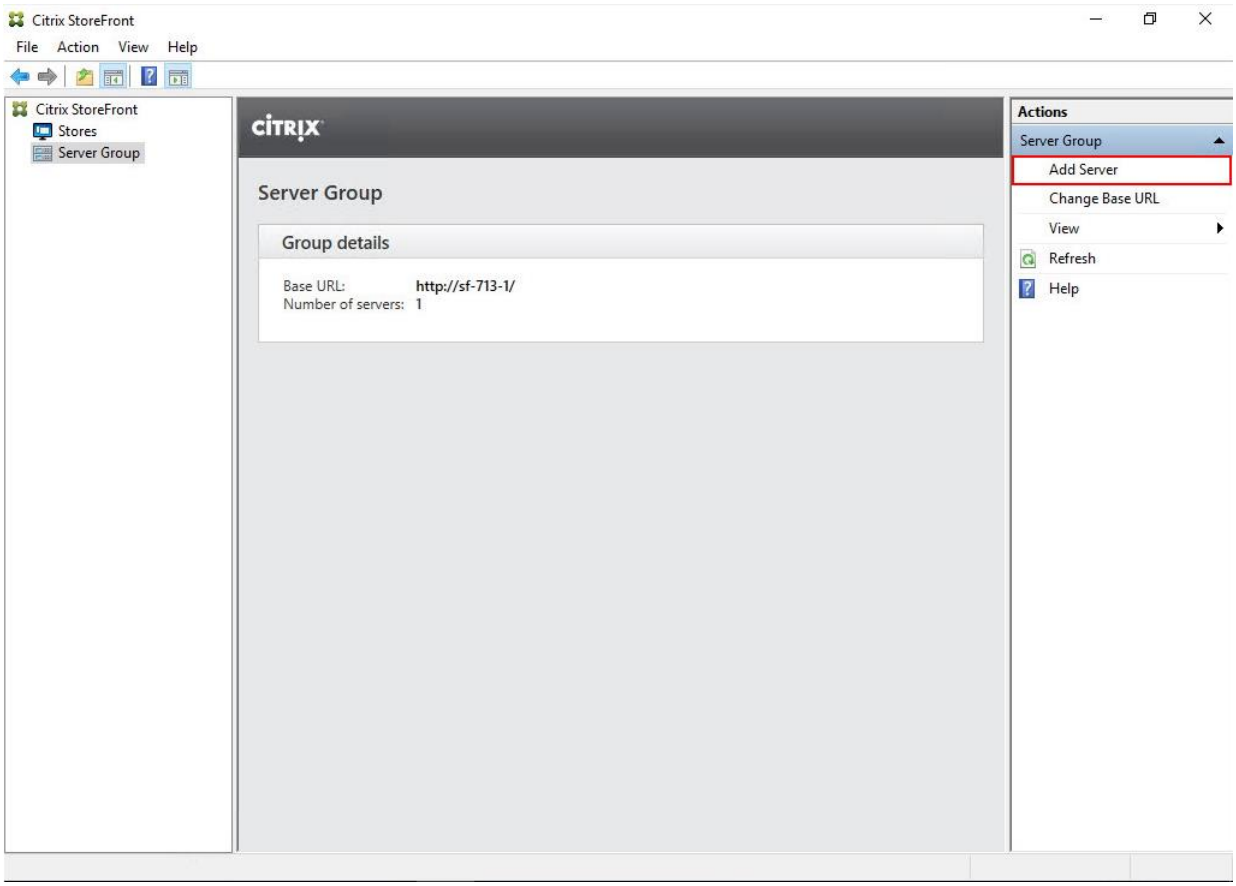

8. Copy the Authorization code from the Add Server dialog.

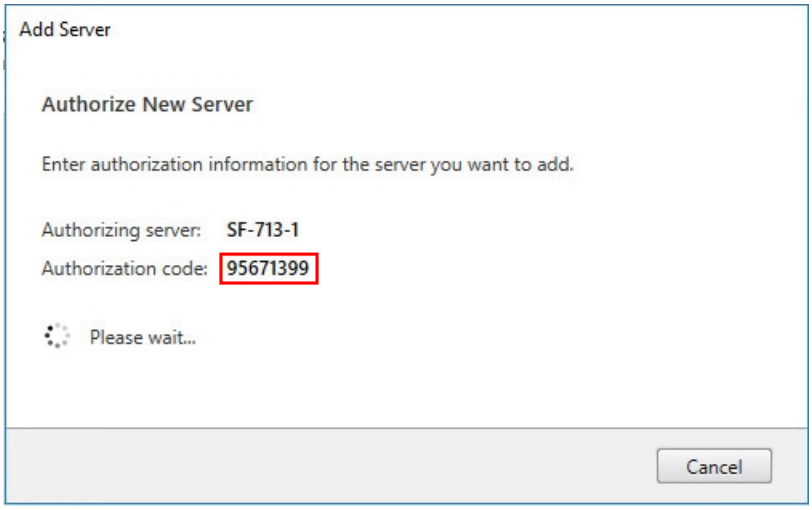

9. Connect to the second Storefront server and paste the Authorization code into the Join Server Group dialog.

10. Click Join.

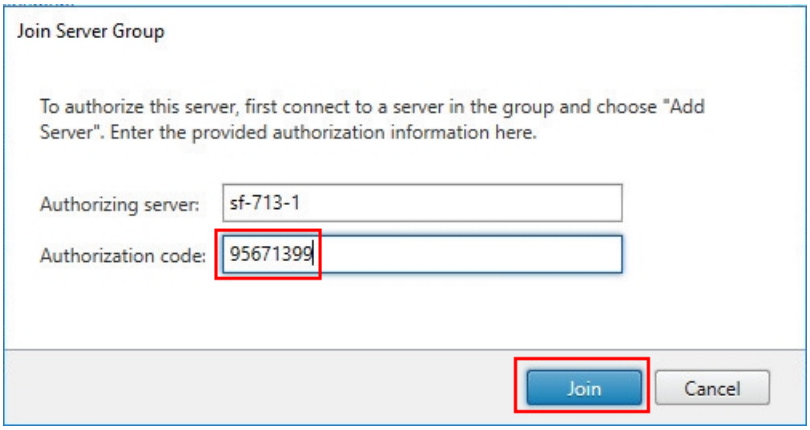

11. A message appears when the second server has joined successfully.

12. Click OK.

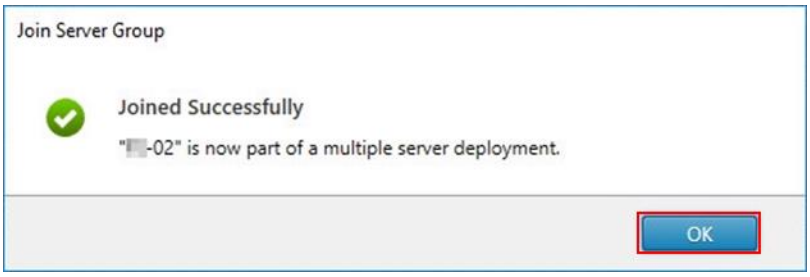

The second StoreFront is now in the Server Group.

Install the Citrix Provisioning Services Target Device Software

For non-persistent Windows 10 virtual desktops and Server 2016 XenApp virtual machines, Citrix Provisioning Services (PVS) is used for deployment. 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, complete the following steps:

The instructions below outline the installation procedure to configure a vDisk for VDI desktops. When you have completed these installation steps, repeat the procedure to configure a vDisk for RDS.

1. On the Window 10 Master Target Device, launch the PVS installer from the Provisioning Services ISO.

2. Click the Target Device Installation button.

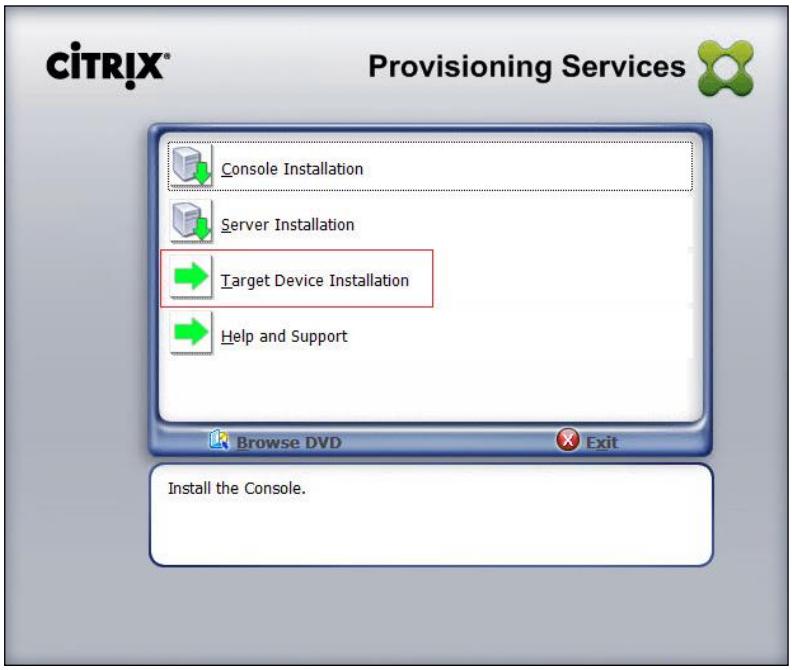

The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.

3. Click Next.

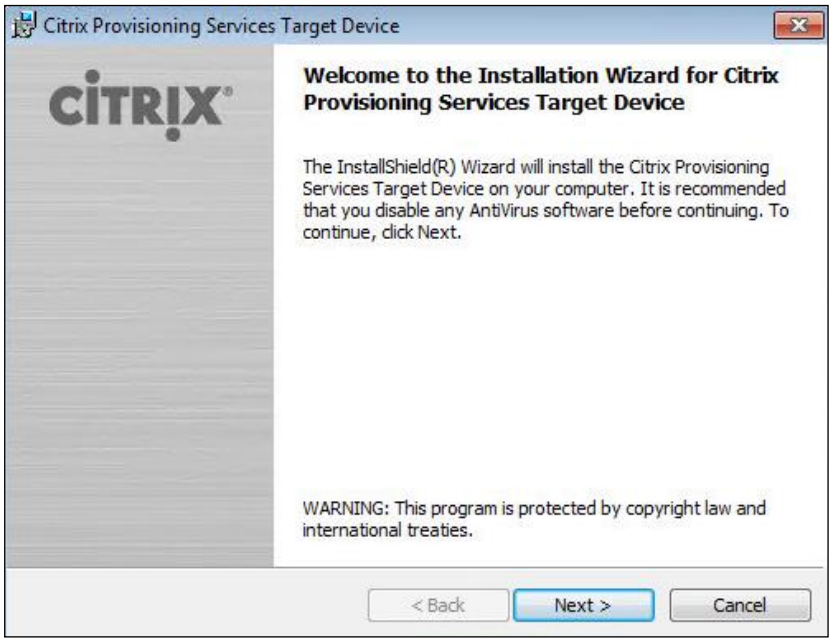

4. Confirm the installation settings and click Install.

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

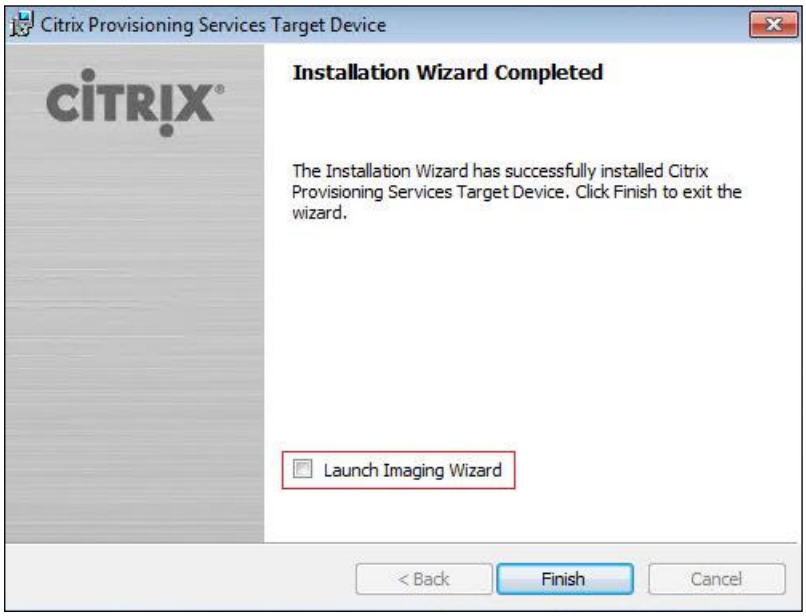

6. Reboot the machine.

Create Citrix Provisioning Services vDisks

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

The following procedure explains how to create a vDisk for VDI desktops. When you have completed these steps, repeat the procedure to build a vDisk for RDS.

1. The PVS Imaging Wizard's Welcome page appears.

2. Click Next.

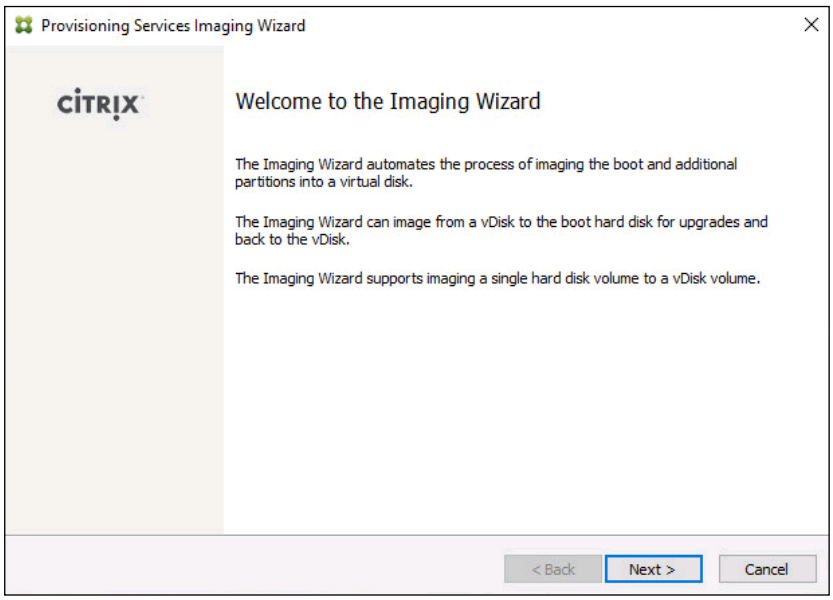

3. The Connect to Farm page appears. Enter the name or IP address of a Provisioning Server within the farm to connect to and the port to use to make that connection.

4. Use the Windows credentials (default) or enter different credentials.

5. Click Next.

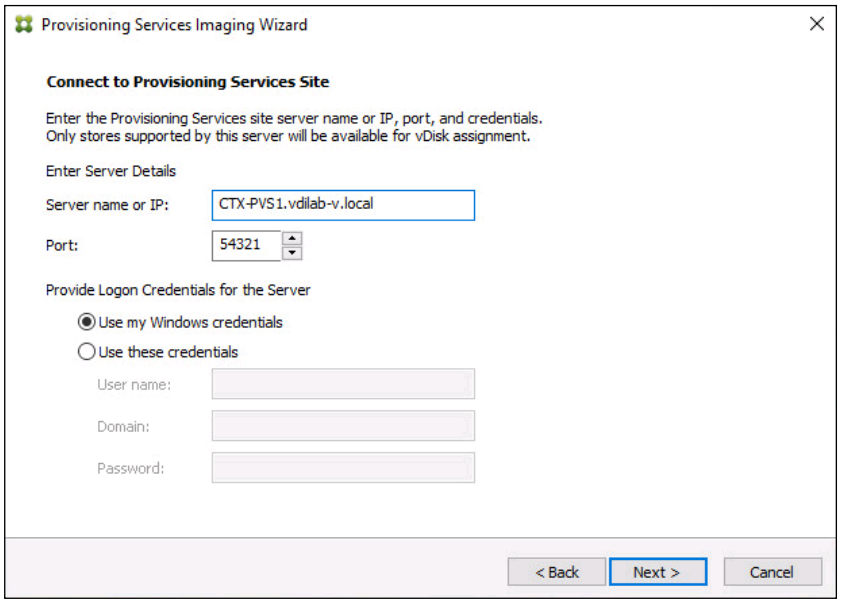

6. Select Create new vDisk.

7. Click Next.

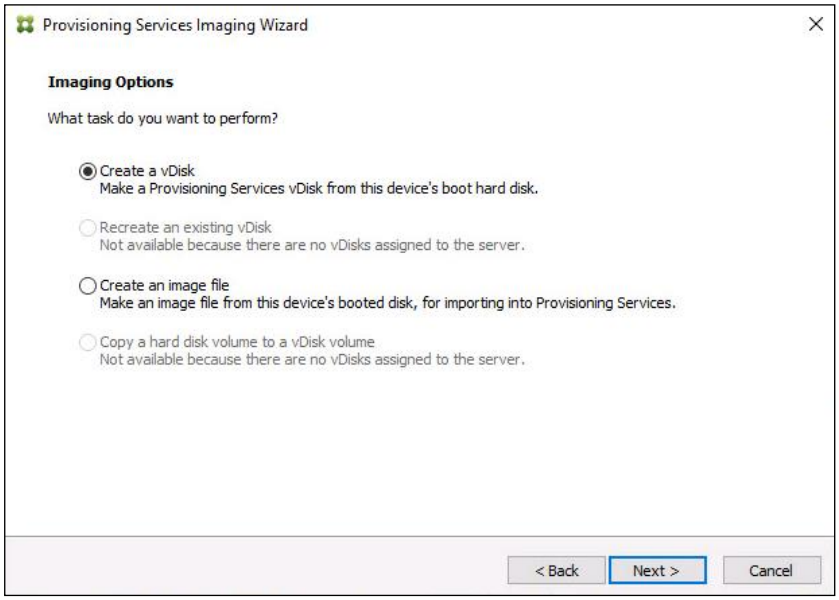

8. The Add Target Device page appears.

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.

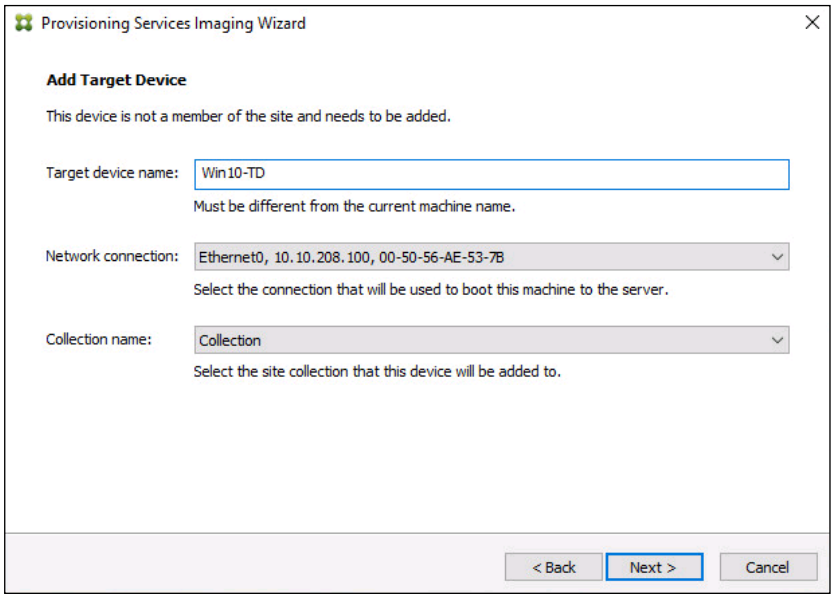

11. The New vDisk dialog displays. Enter the name of the vDisk.

12. Select the Store where the vDisk will reside. Select the vDisk type, either Fixed or Dynamic, from the drop-down menu. (This CVD used Dynamic rather than Fixed vDisks.)

13. Click Next.

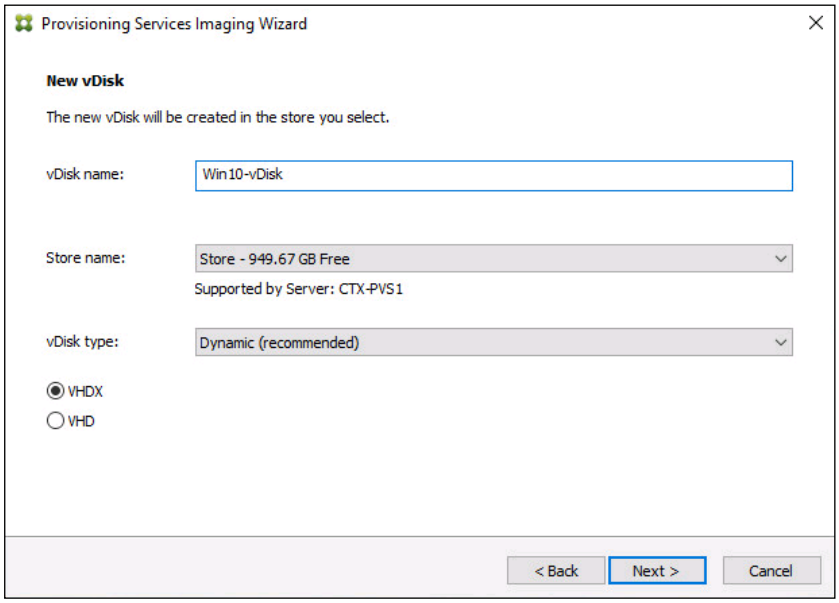

14. On the Microsoft Volume Licensing page, select the volume license option to use for target devices. For this CVD, volume licensing is not used, so the None button is selected.

15. Click Next.

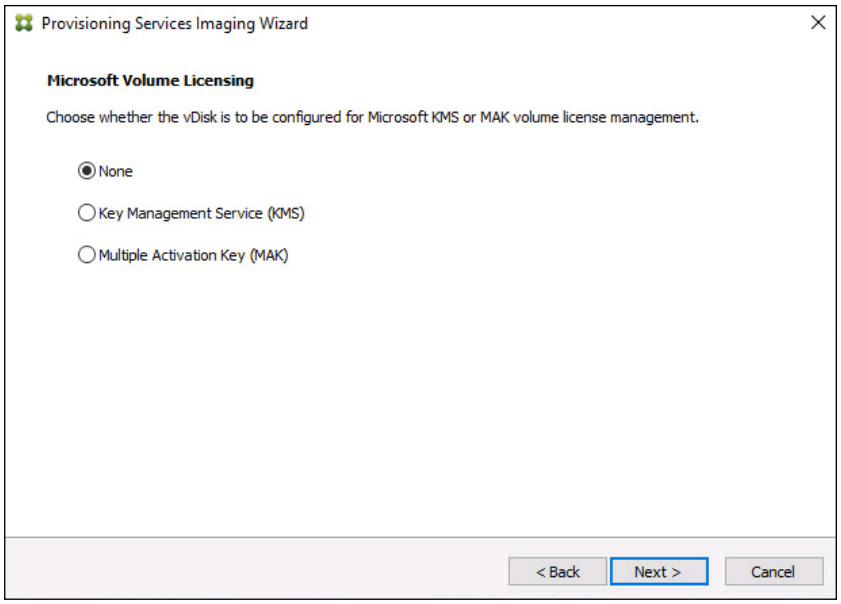

16. Select Image entire boot disk on the Configure Image Volumes page.

17. Click Next.

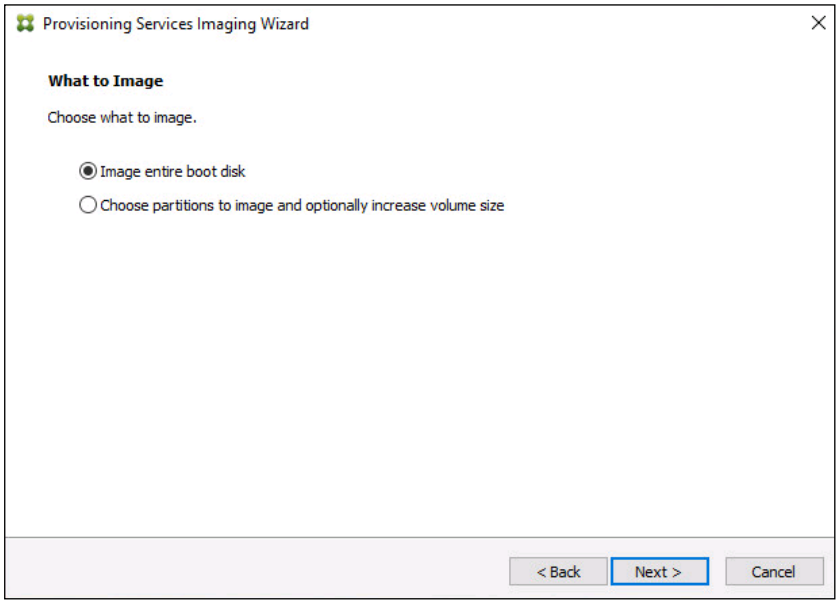

18. Select Optimize for hard disk again for Provisioning Services before imaging on the Optimize Hard Disk for Provisioning Services.

19. Click Next.

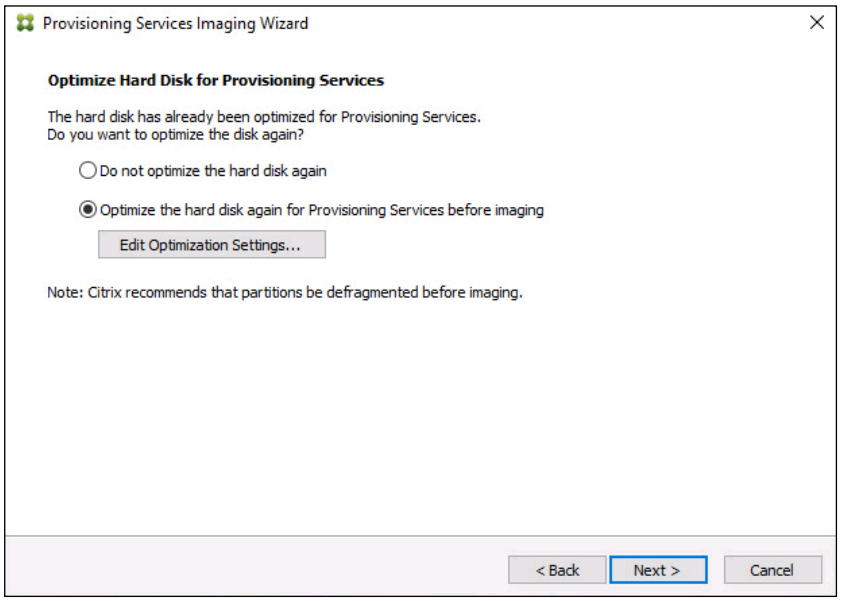

20. Select Create on the Summary page.

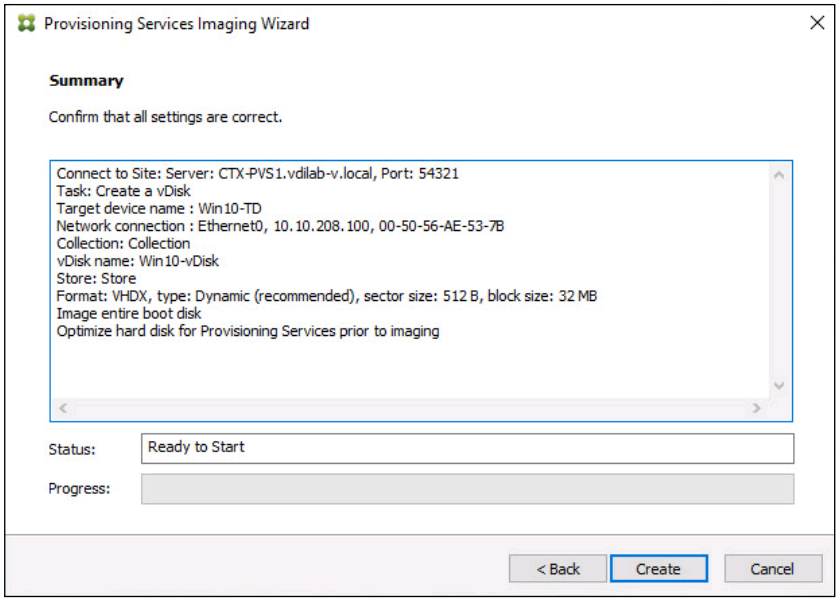

21. Review the configuration and click Continue.

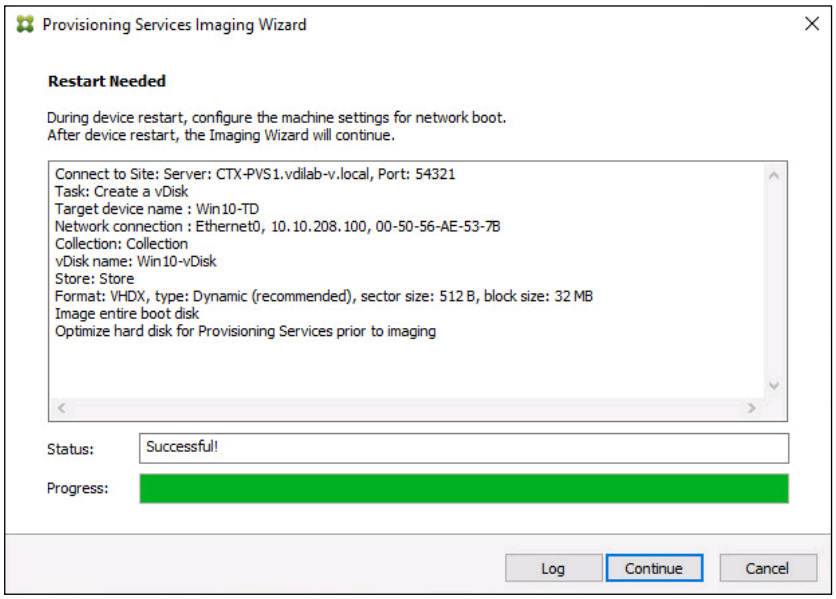

22. When prompted, click No to shut down the machine.

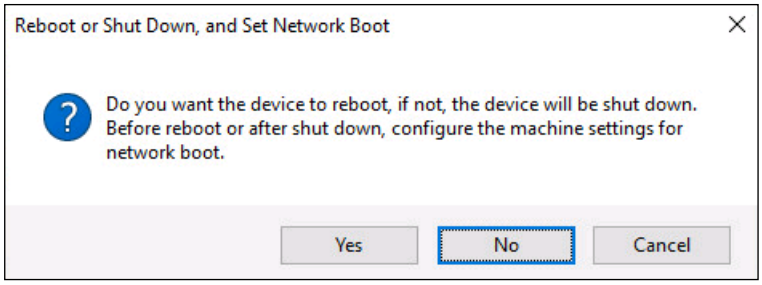

23. Edit the virtual machine settings and select VM Options > Boot Options > Force BIOS Setup under Boot Order.

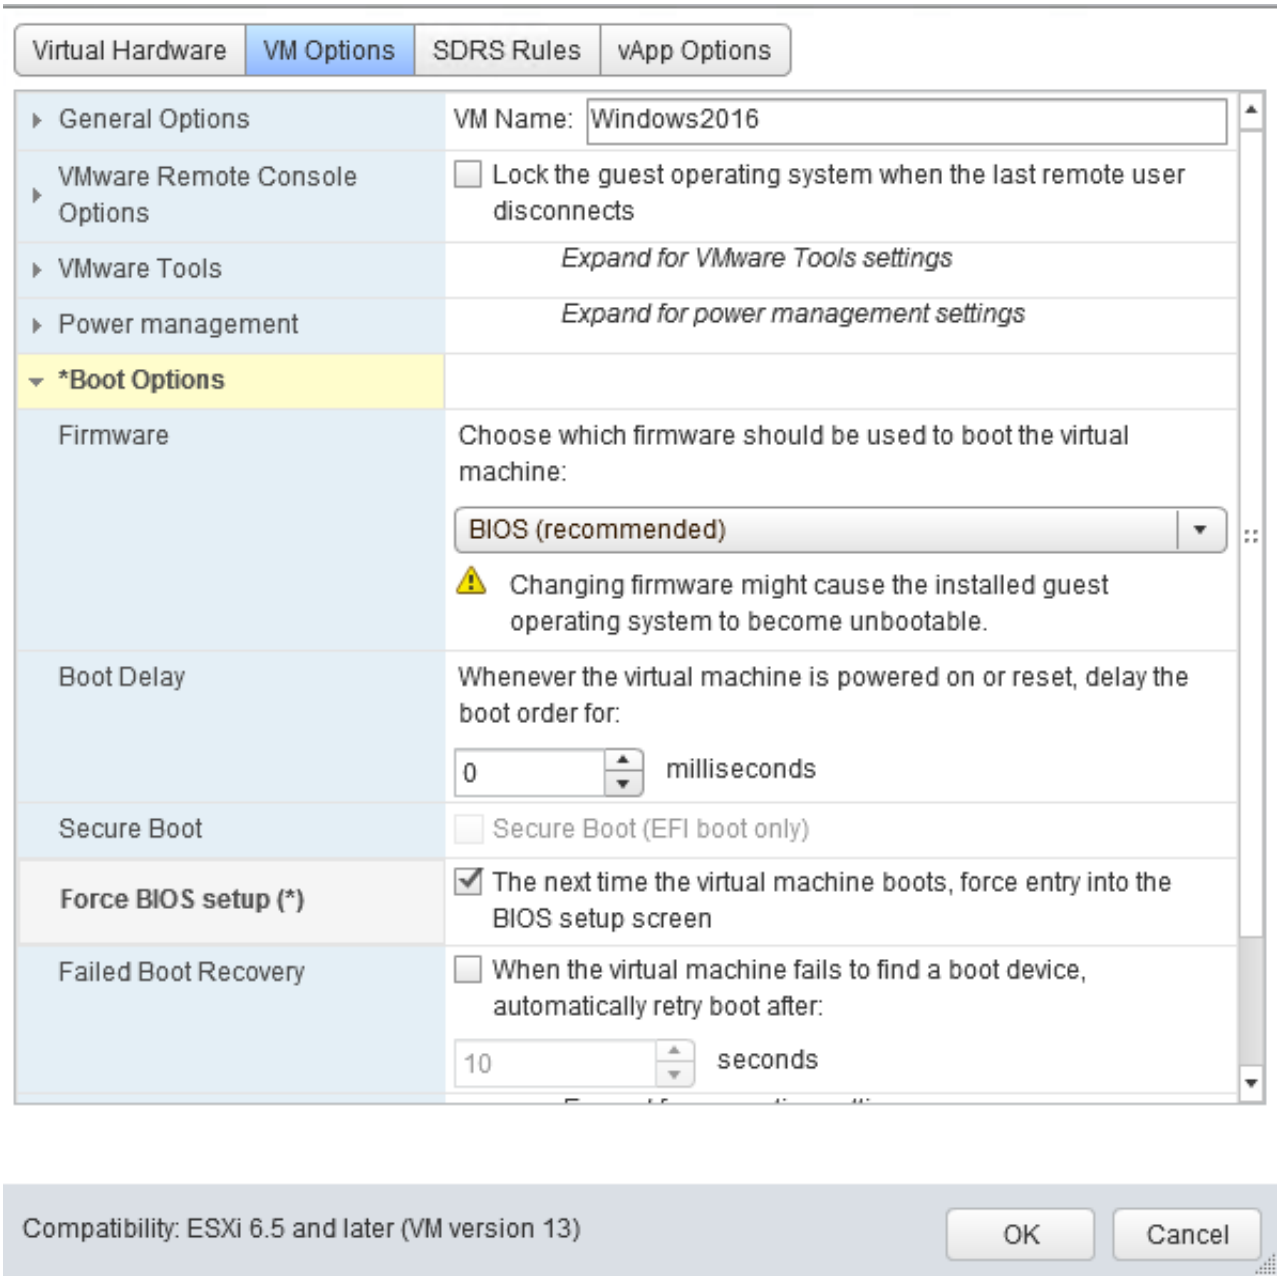

24. Select Boot and move the Network Boot to the top of the list.

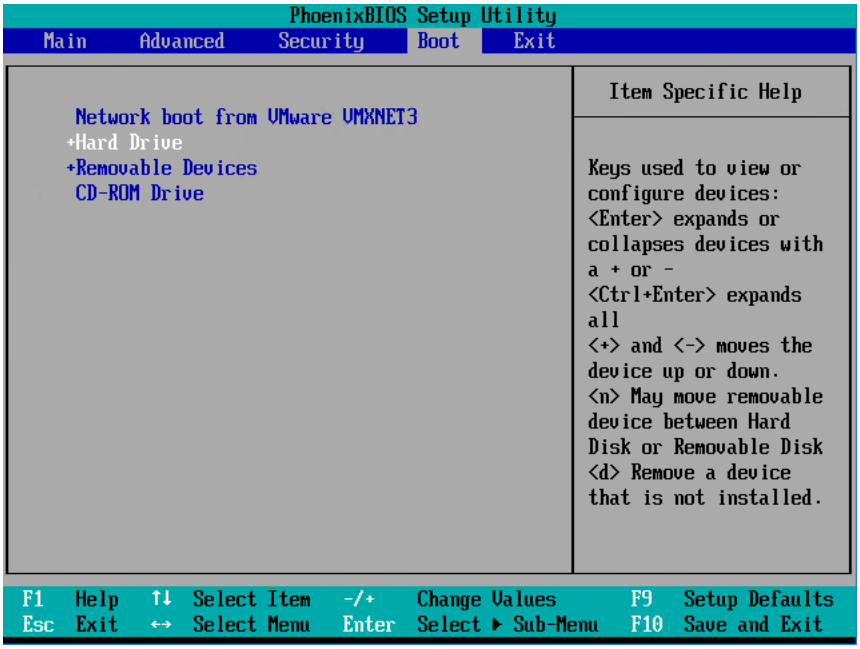

25. Restart Virtual Machine.

After restarting the virtual machine, log into the VDI or RDS master target. The PVS imaging process begins, copying the contents of the C: drive to the PVS vDisk located on the server.

26. If prompted to Restart select Restart Later.

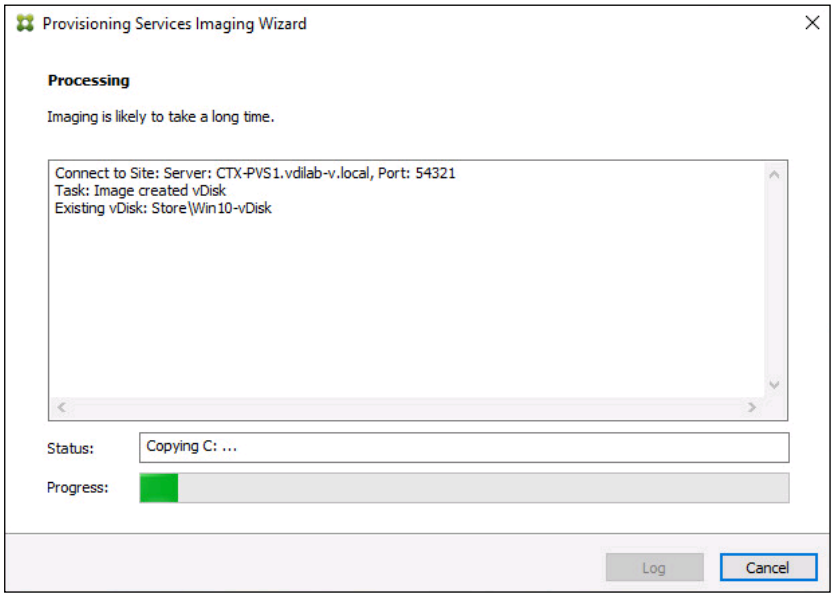

27. A message is displayed when the conversion is complete, click Done.

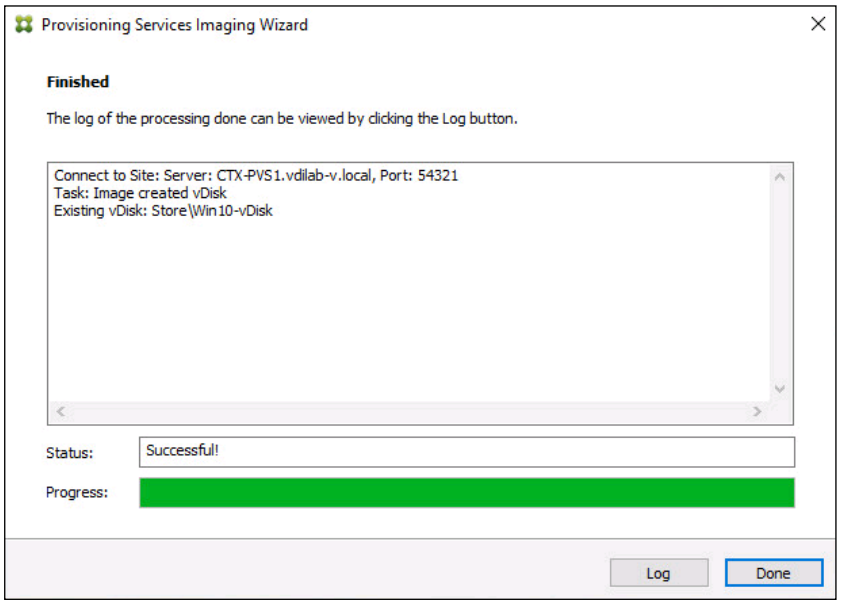

28. Shutdown the virtual machine used as the VDI or RDS master target.

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.

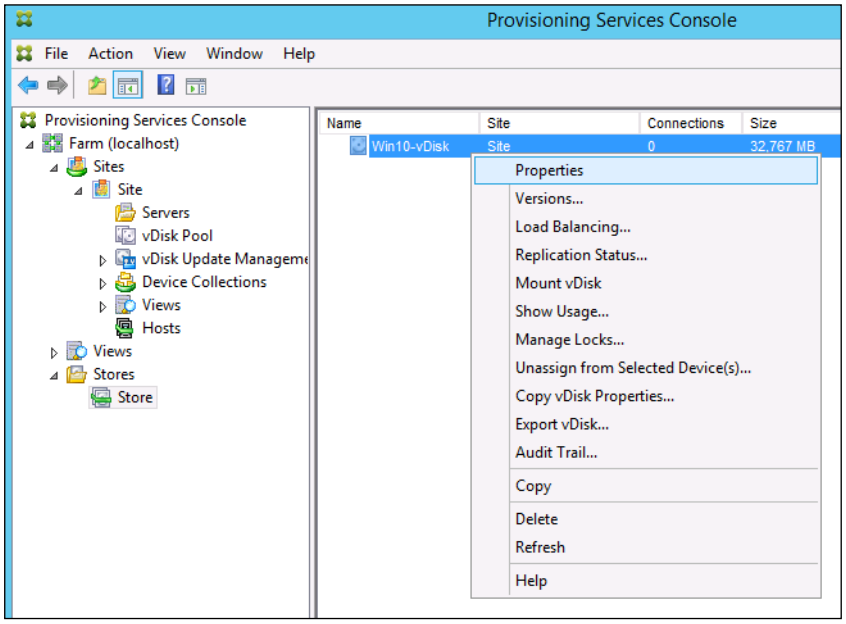

31. On the vDisk Properties dialog, change Access mode to “Standard Image (multi-device, read-only access)”.

32. Set the Cache Type to “Cache in device RAM with overflow on hard disk.”

33. Set Maximum RAM size (MBs): 256 for VDI and set 1024 MB for RDS vDisk.

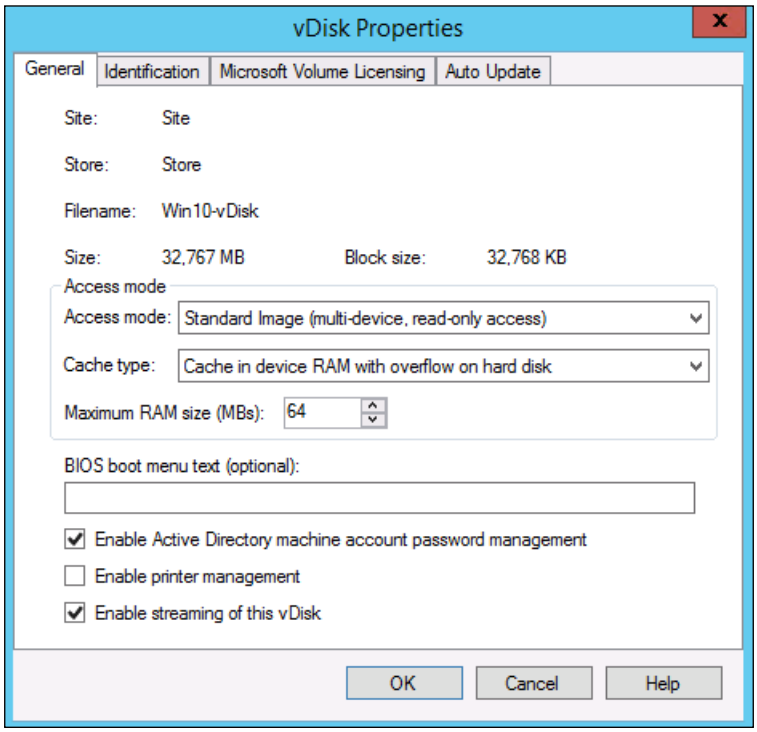

34. Click OK.

Repeat this procedure to create vDisks for both the Hosted VDI Desktops (using the Windows 10 OS image) and the Hosted Shared Desktops (using the Windows Server 2016 image).

Provision Virtual Desktop Machines

To create VDI and RDS machines, complete the following steps:

1. Select the Master Target Device virtual machine from the vCenter Client.
2. Right-click the virtual machine and select Template > Clone to template.
3. Name the cloned 'Template'.
4. Select the cluster and datastore where the first phase of provisioning will occur.

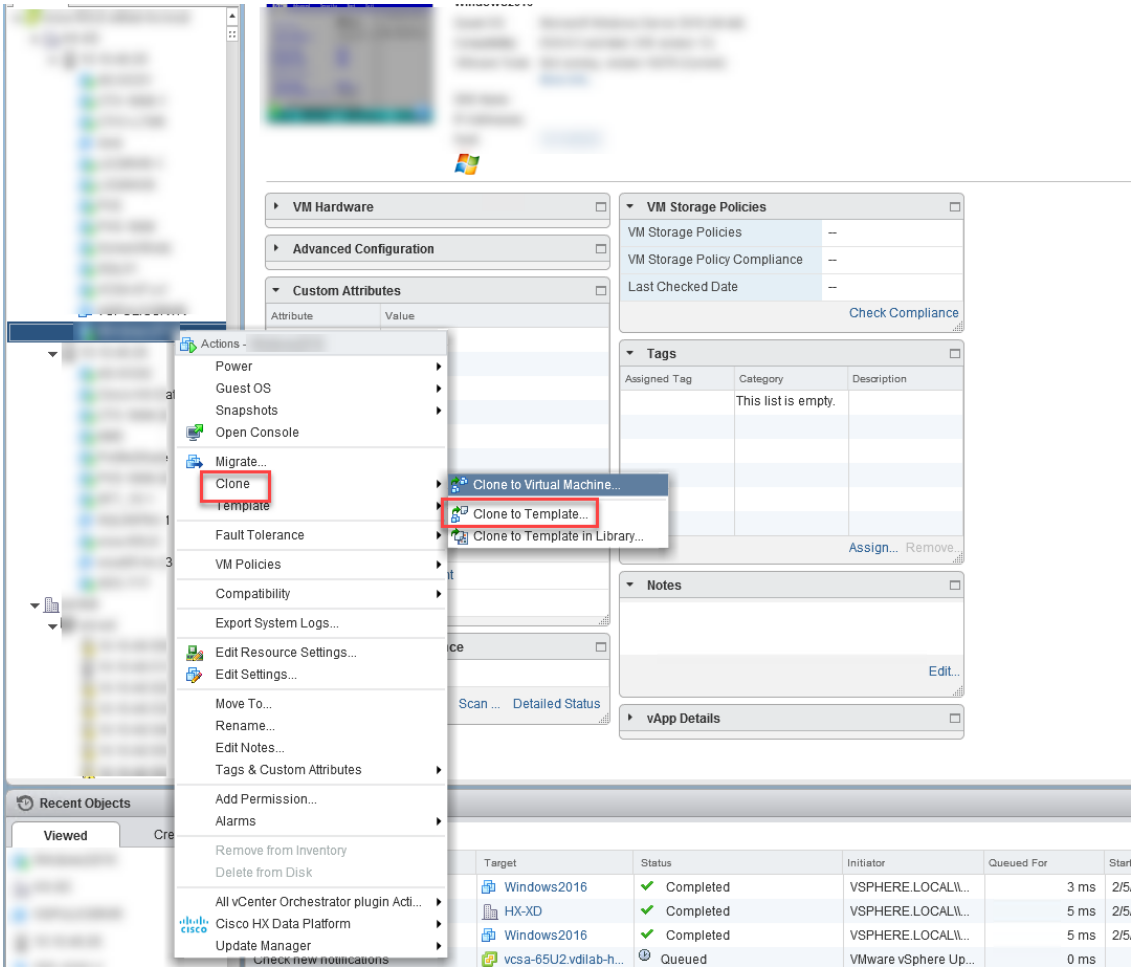

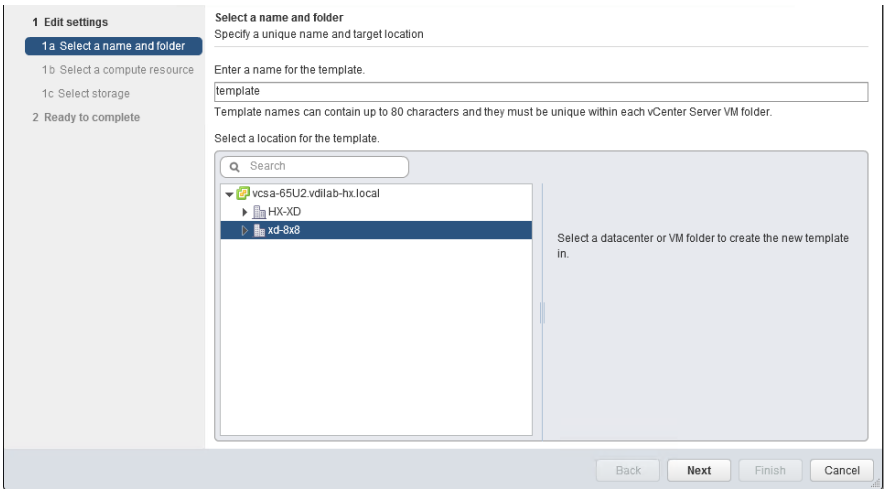

5. Click Next.
6. Click Next through the remaining screens.
7. Click Finish to create the template.
8. From Citrix Studio on the Desktop Controller, select Hosting and Add Connection and Resources.

9. Select Use an existing Connection and click Next.

10. Correspond the name of the resource with desktop machine clusters.

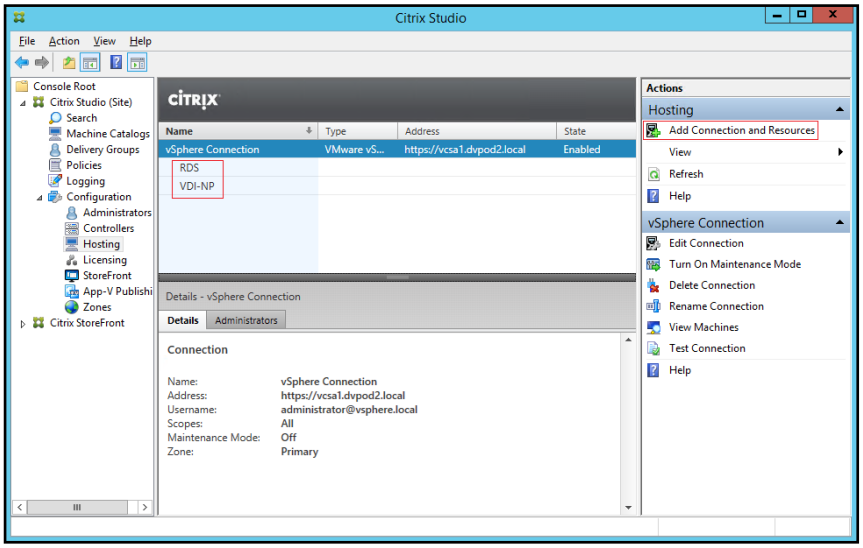

11. Browse and select the Hyper-V cluster for desktop provisioning and use the default storage method Use storage shared by hypervisors.

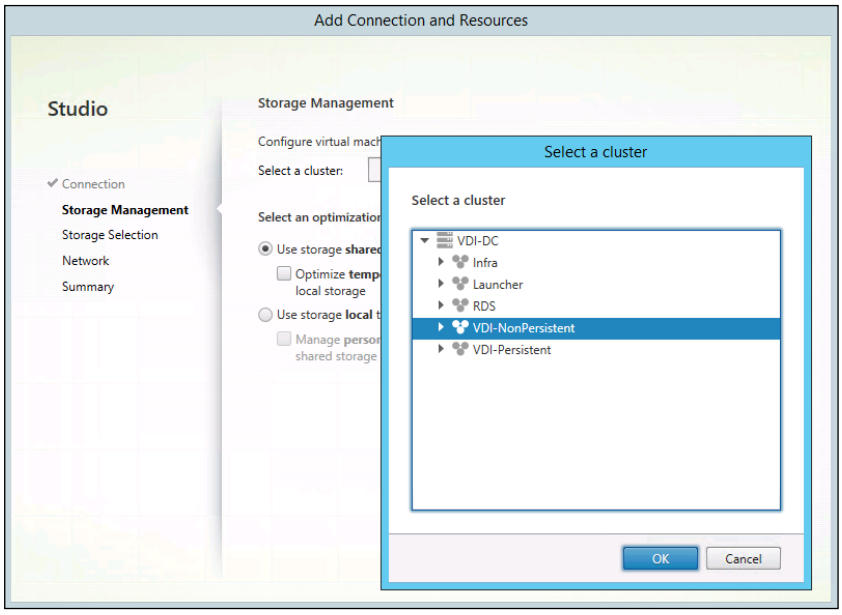

12. Select the data storage location for the corresponding resource.

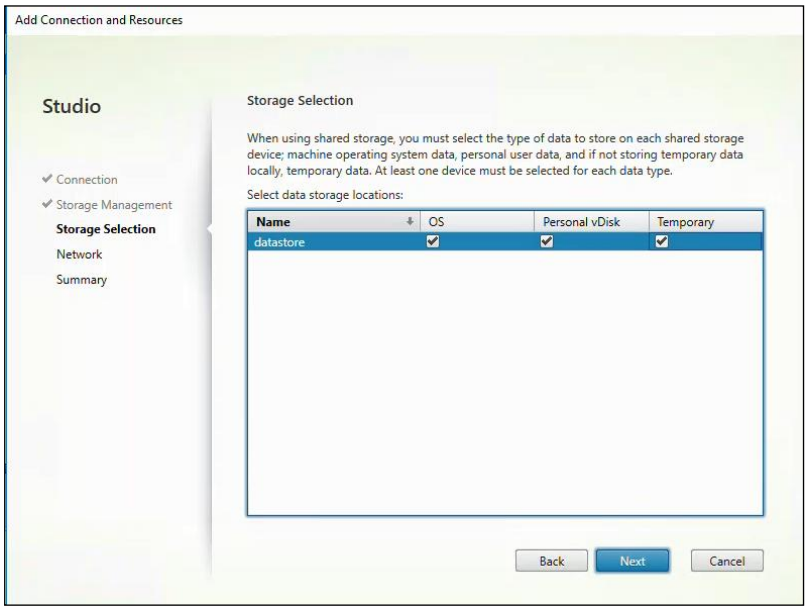

13. Select the VDI networks for the desktop machines and click Next.

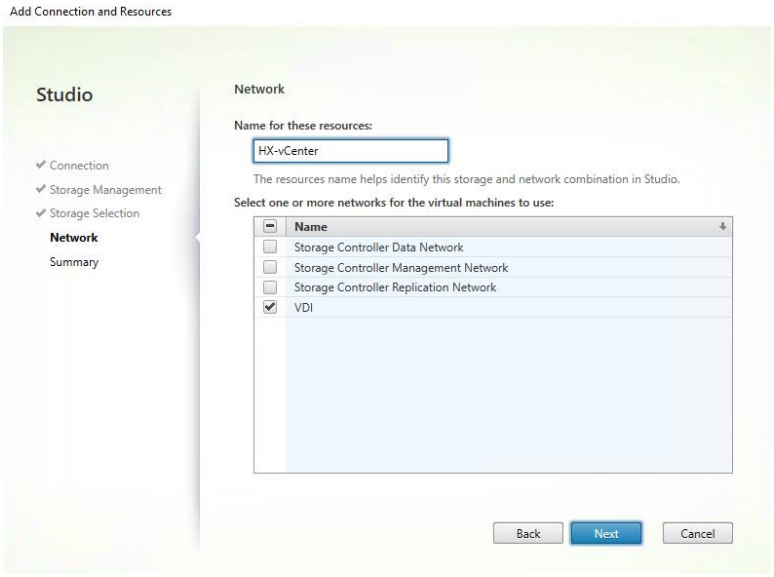

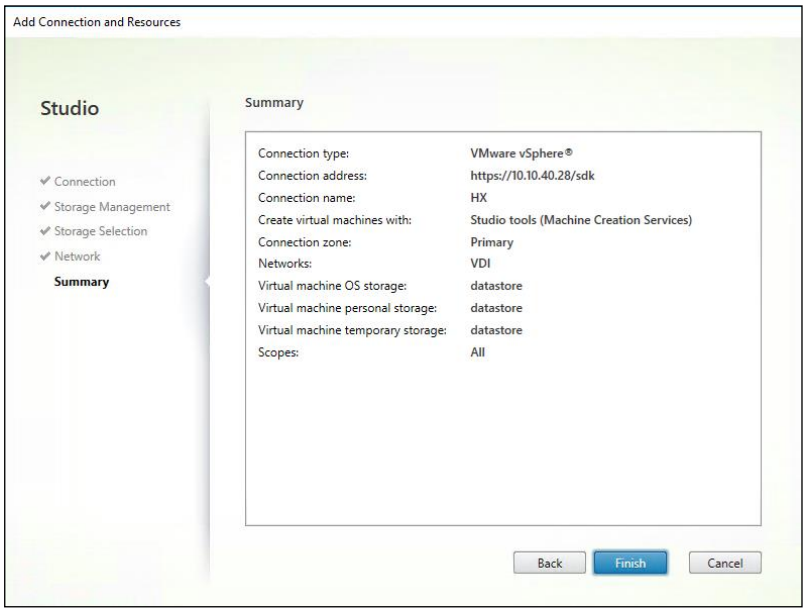

14. Click Finish.

Return to these settings to alter the datastore selection for each set of provisioned desktop machines.

Provision Desktop Machines Using the Citrix Provisioning Services Console

To provision desktop machines using the Citrix Provisioning Services Console, complete the following steps:

1. Start the Citrix Desktop Setup Wizard from the Provisioning Services Console.
2. Right-click the Site.
3. Choose Citrix VDI Setup Wizard… from the context menu.

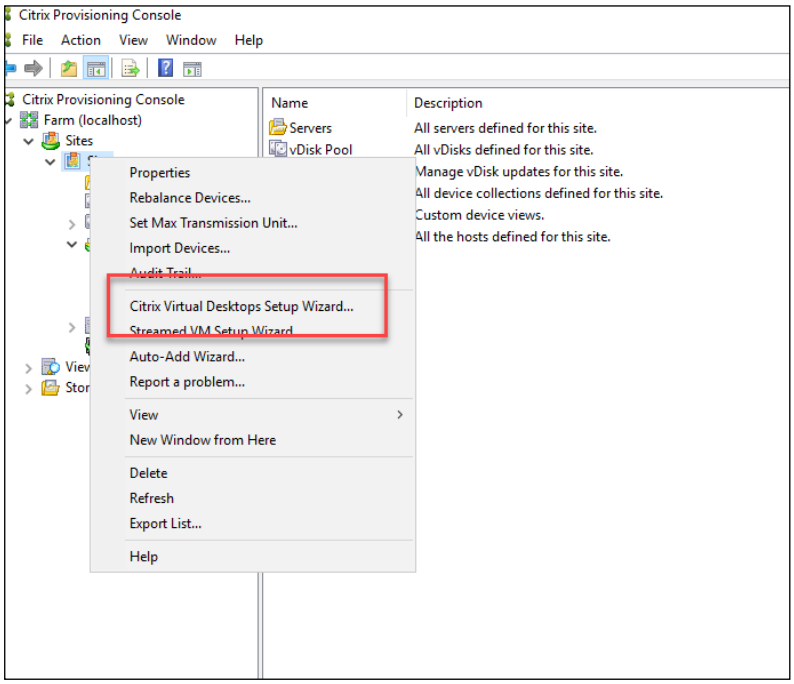

4. Click Next.

5. Enter the Citrix VDI Controller address that will be used for the wizard operations.

6. Click Next.

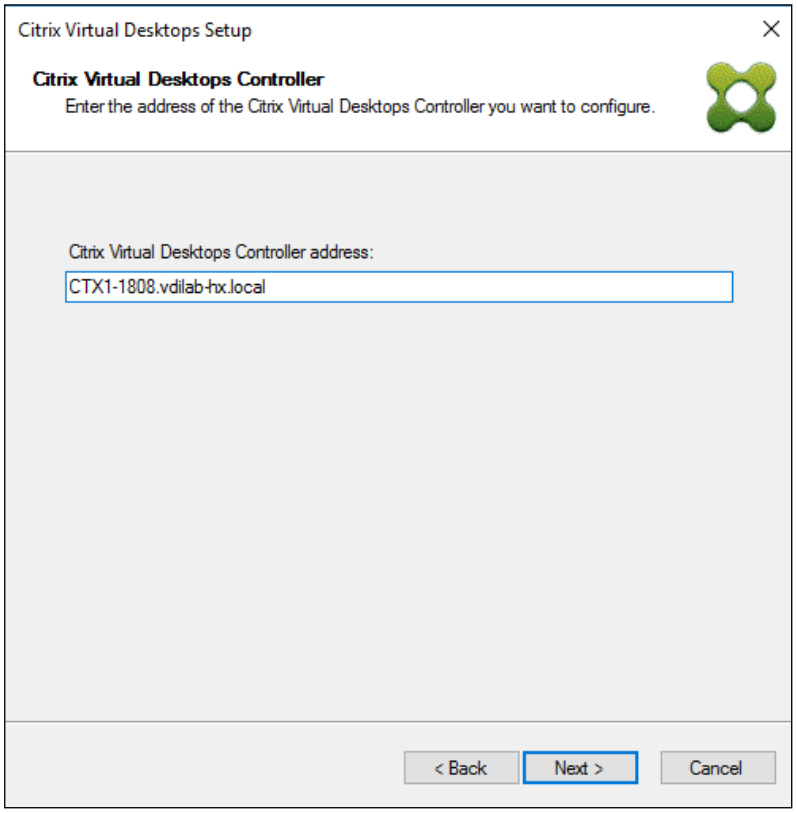

7. Select the Host Resources on which the virtual machines will be created.
8. Click Next.

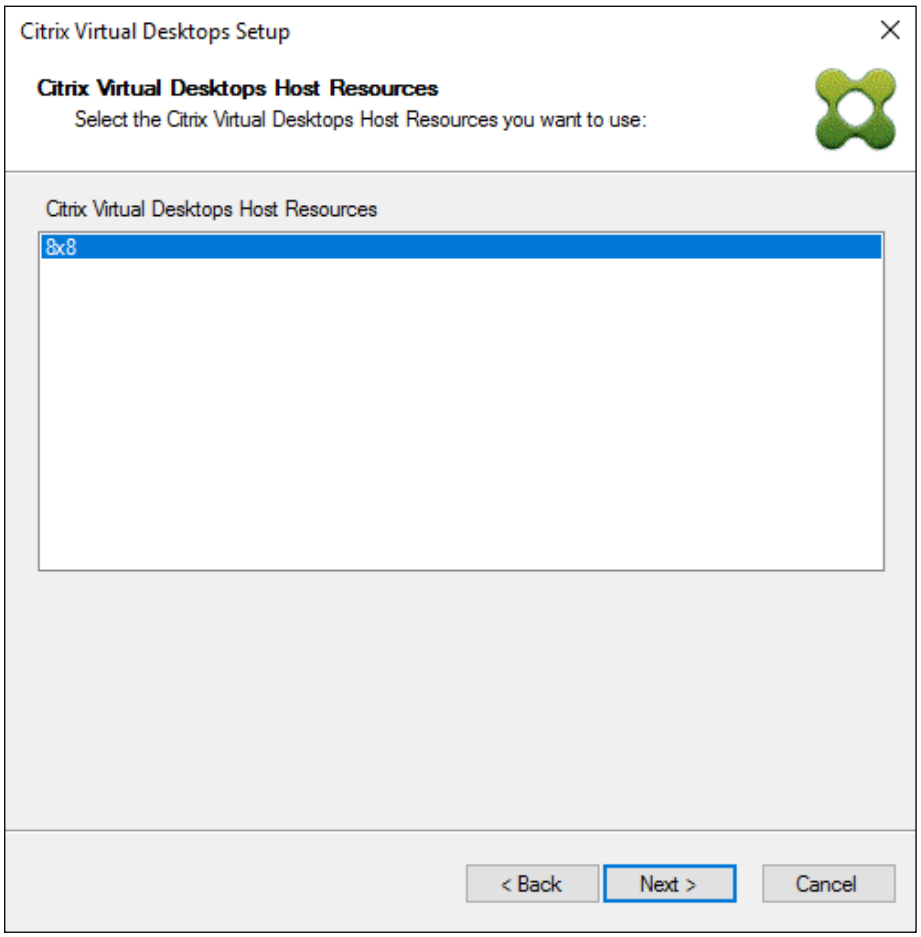

9. Provide the Host Resources Credentials (Username and Password) to the Citrix VDI controller when prompted.
10. Click OK.

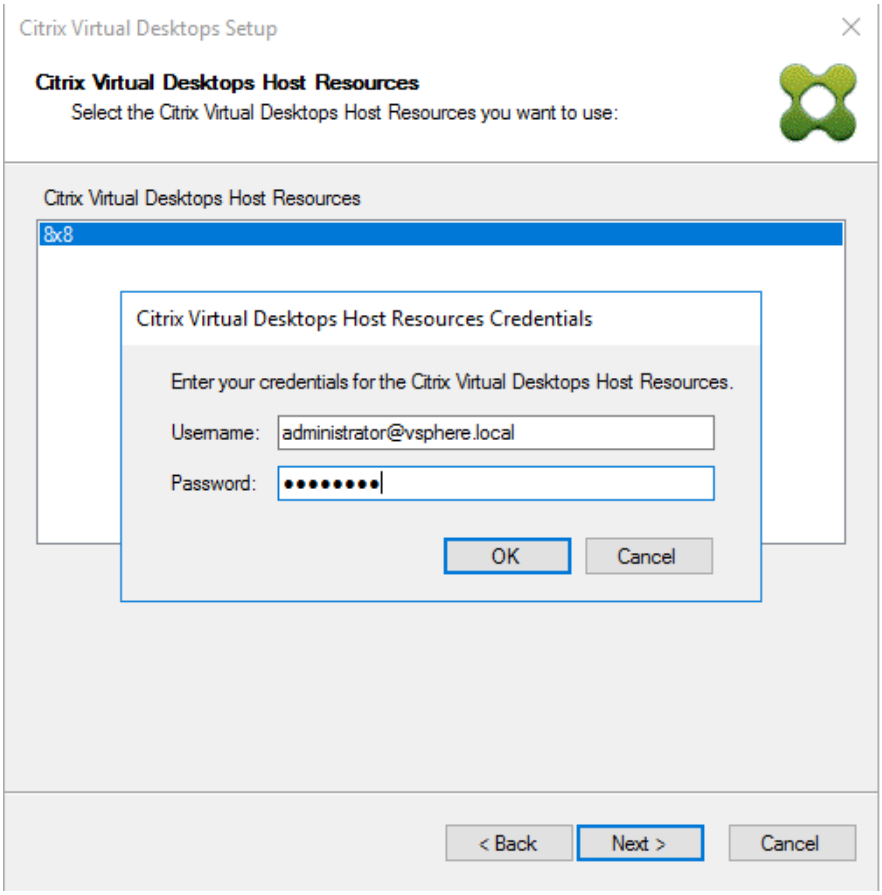

11. Select the Template created earlier.

12. Click Next.

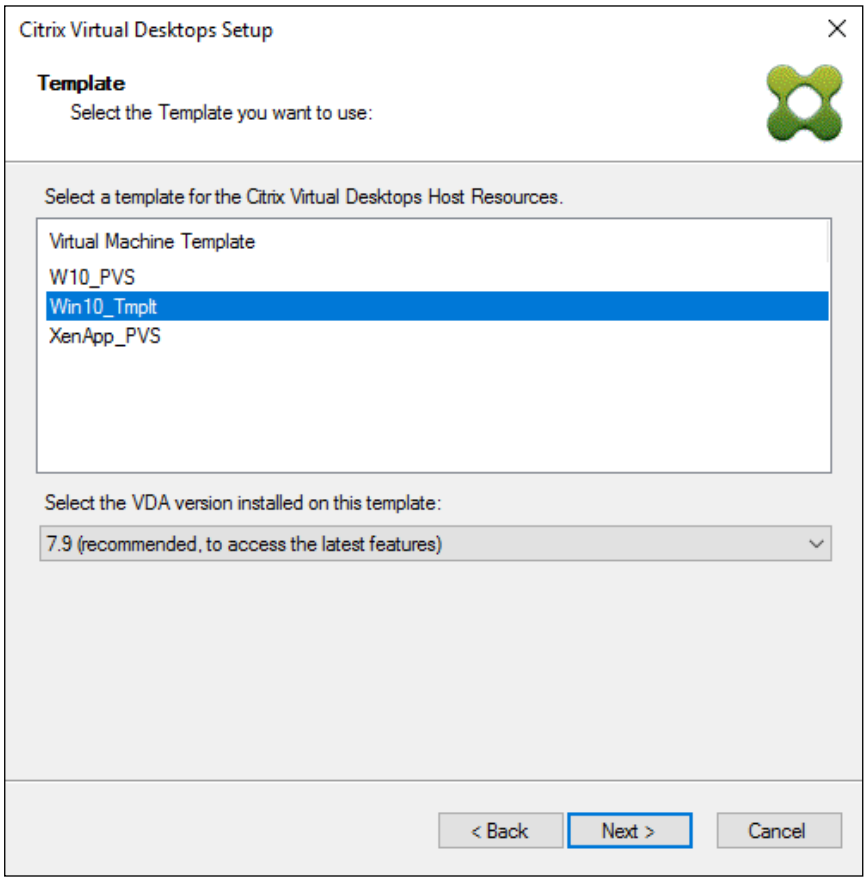

13. Select the vDisk that will be used to stream virtual machines.
14. Click Next.

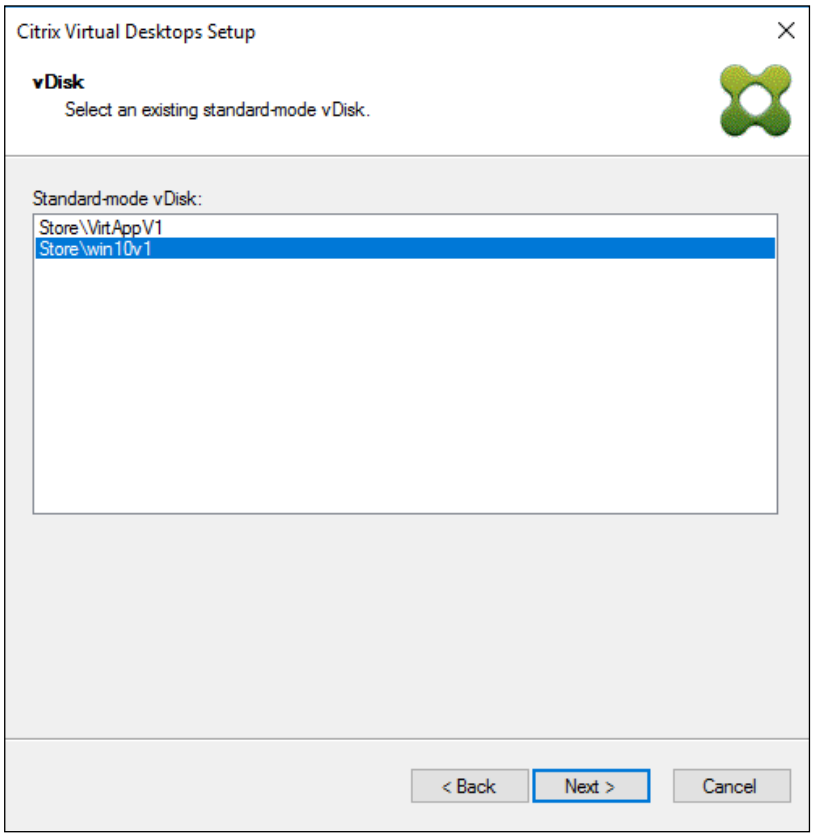

15. Select “Create a new catalog.”

The catalog name is also used as the collection name in the PVS site.

16. Click Next.

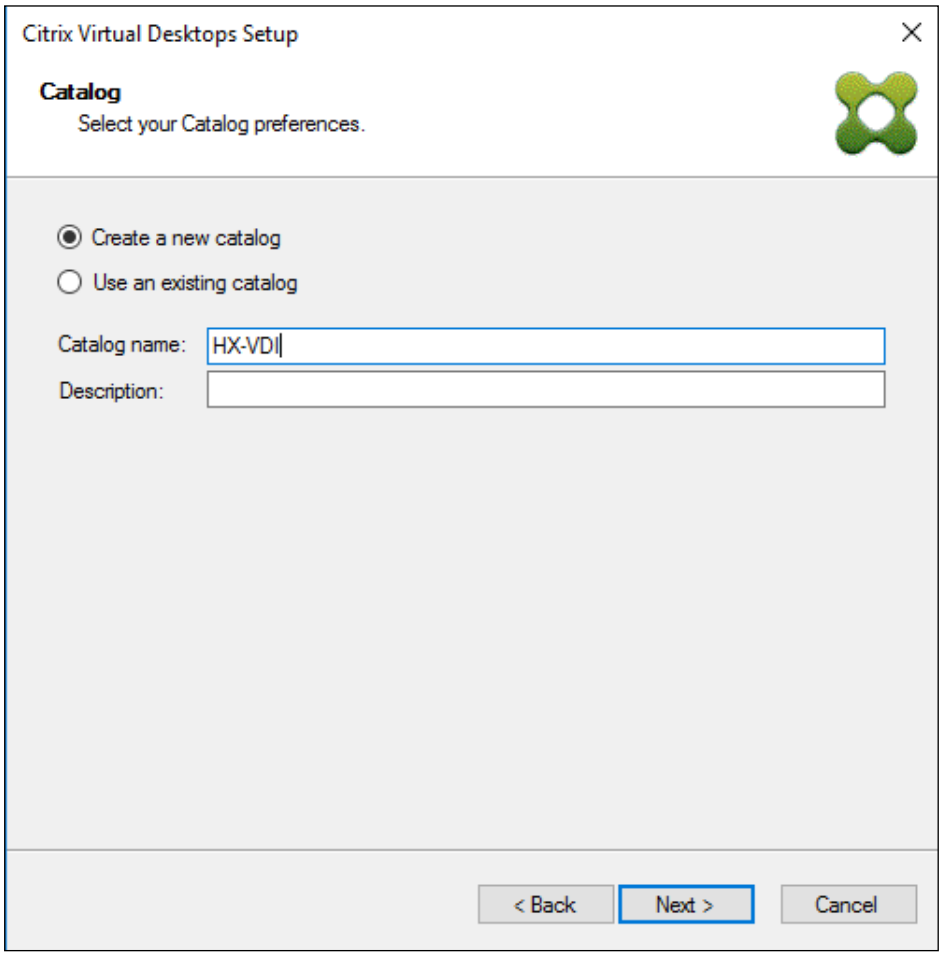

17. On the Operating System dialog, specify the operating system for the catalog. Specify Windows Desktop Operating System for VDI and Windows Server Operating System for RDS.

18. Click Next.

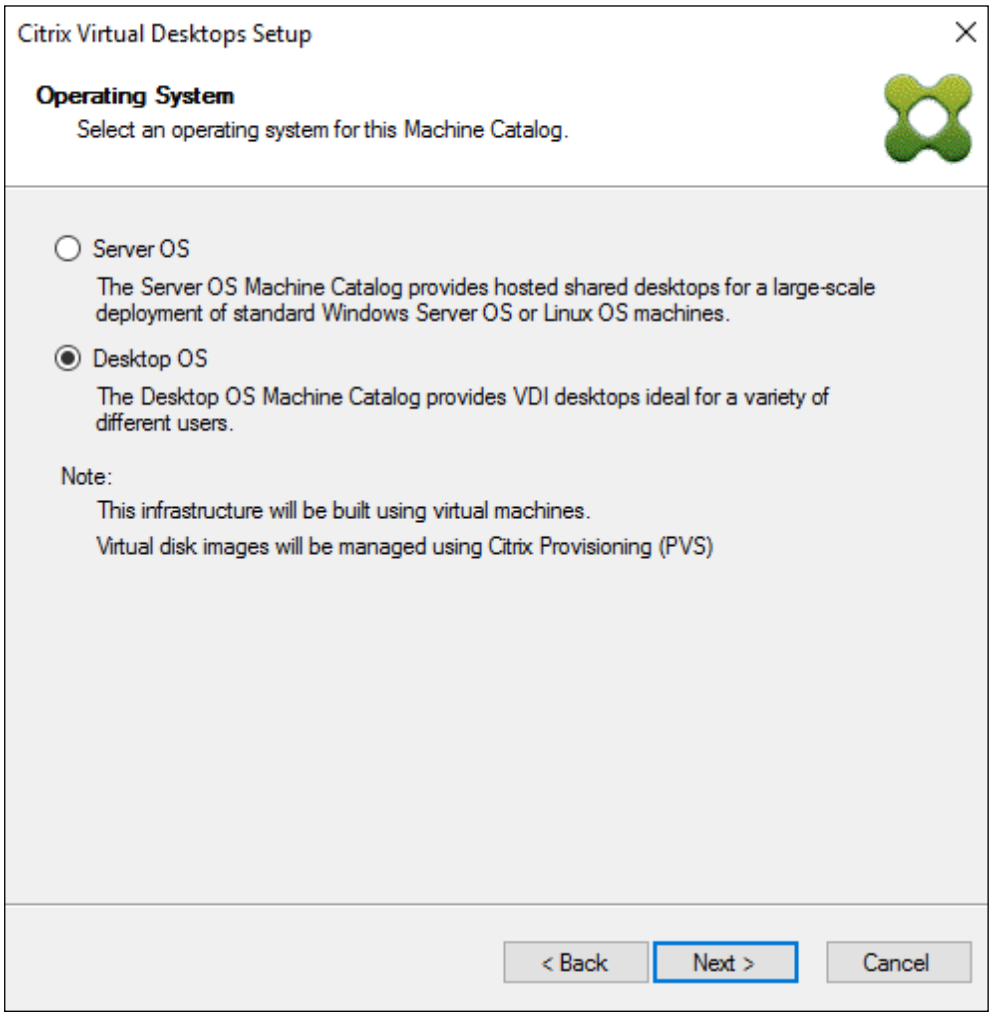

19. If you specified a Windows Desktop OS for VDIs, a User Experience dialog appears. Specify that the user will connect to "A fresh new (random) desktop each time."

20. Click Next.

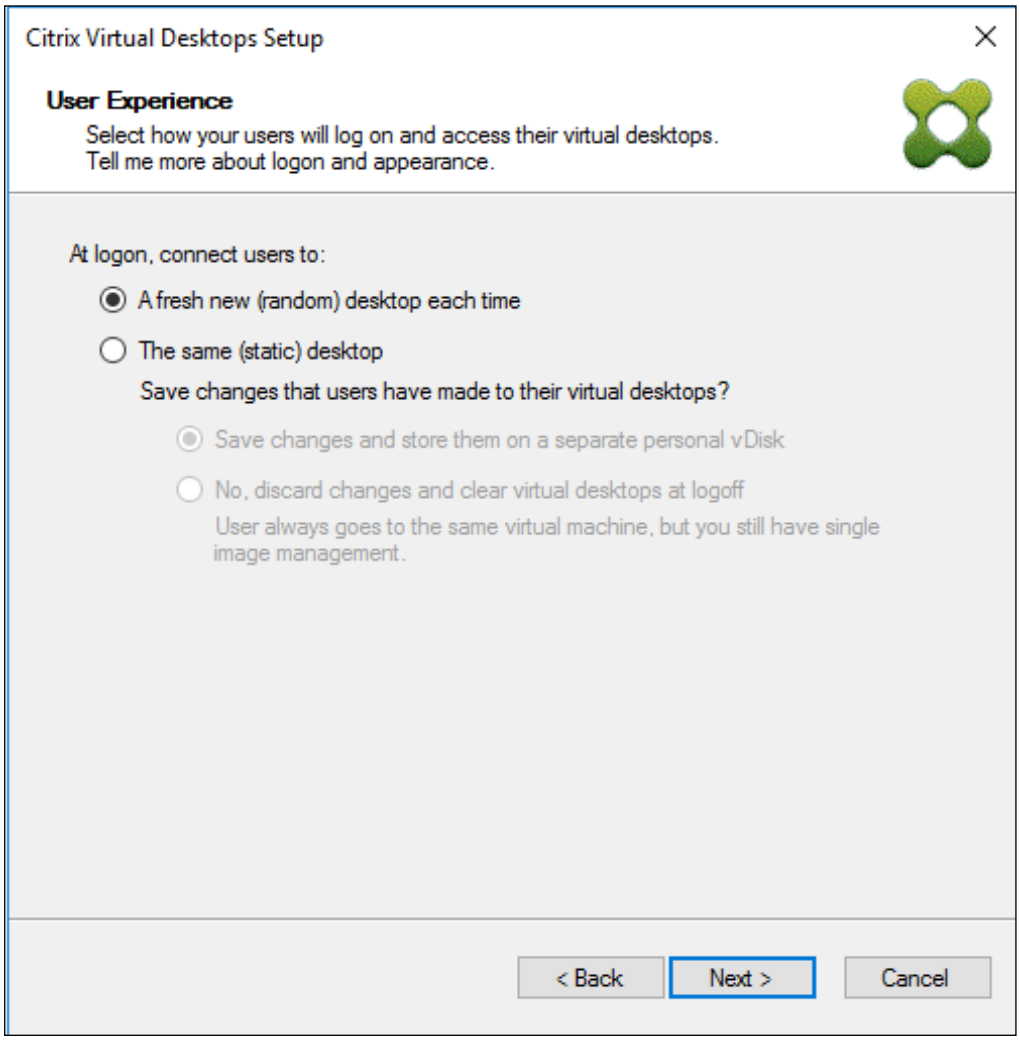

21. On the Virtual machines dialog, specify:

    a. The number of virtual machines to create.
    b. Number of vCPUs for the virtual machine (2 for VDI, 8 for RDS)
    c. The amount of memory for the virtual machine (4GB for VDI, 24GB for RDS)
    d. The write-cache disk size (10GB for VDI, 30GB for RDS)
    e. PXE boot as the Boot Mode

22. Click Next.

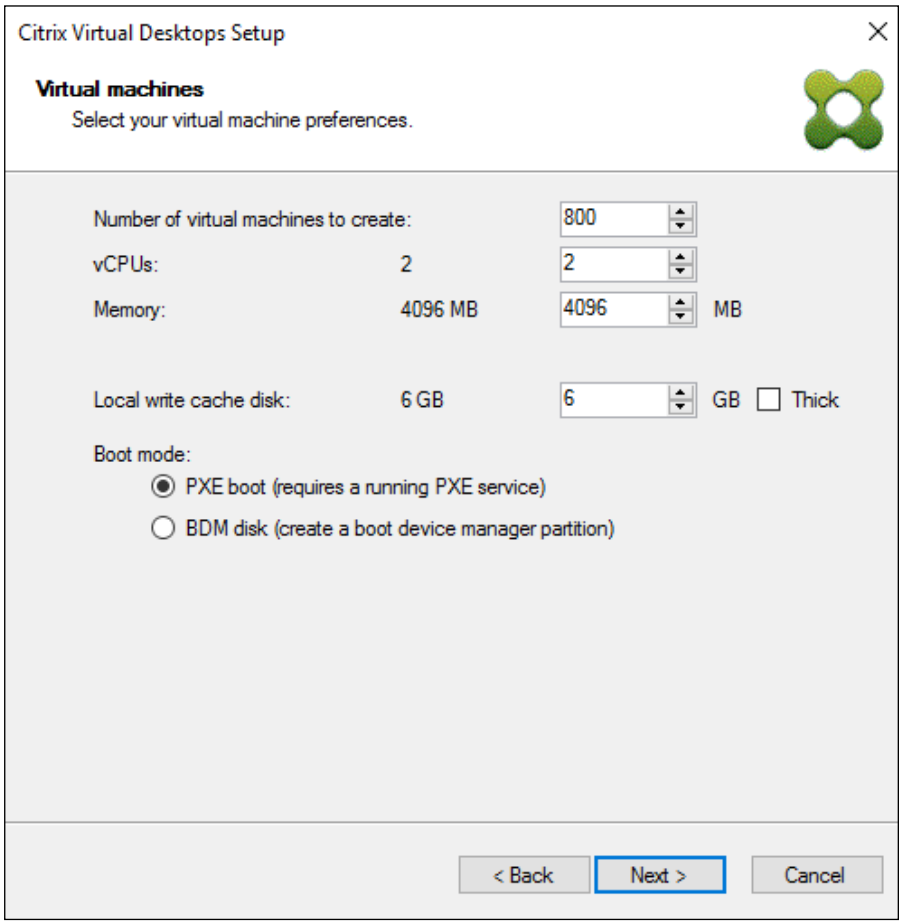

23. Select the Create new accounts radio button.

24. Click Next.

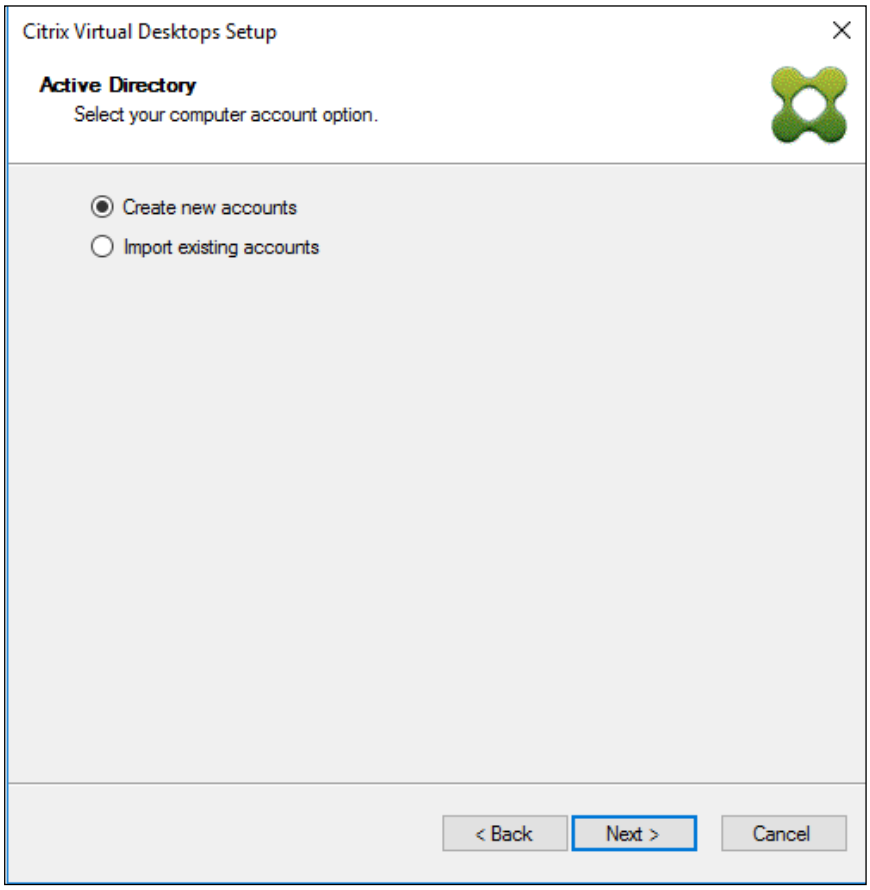

25. Specify the Active Directory Accounts and Location. This is where the wizard should create the computer accounts.

26. Provide the Account naming scheme. An example name is shown in the text box below the name scheme selection location.

27. Click Next.

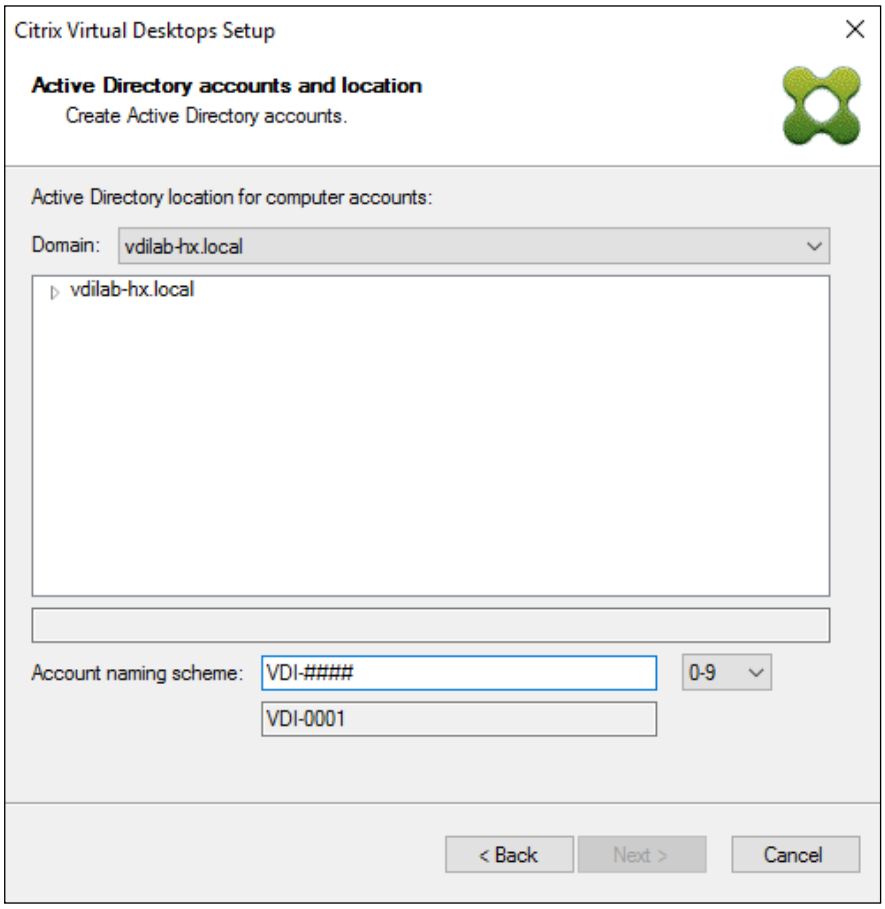

28. Click Finish to begin the virtual machine creation.

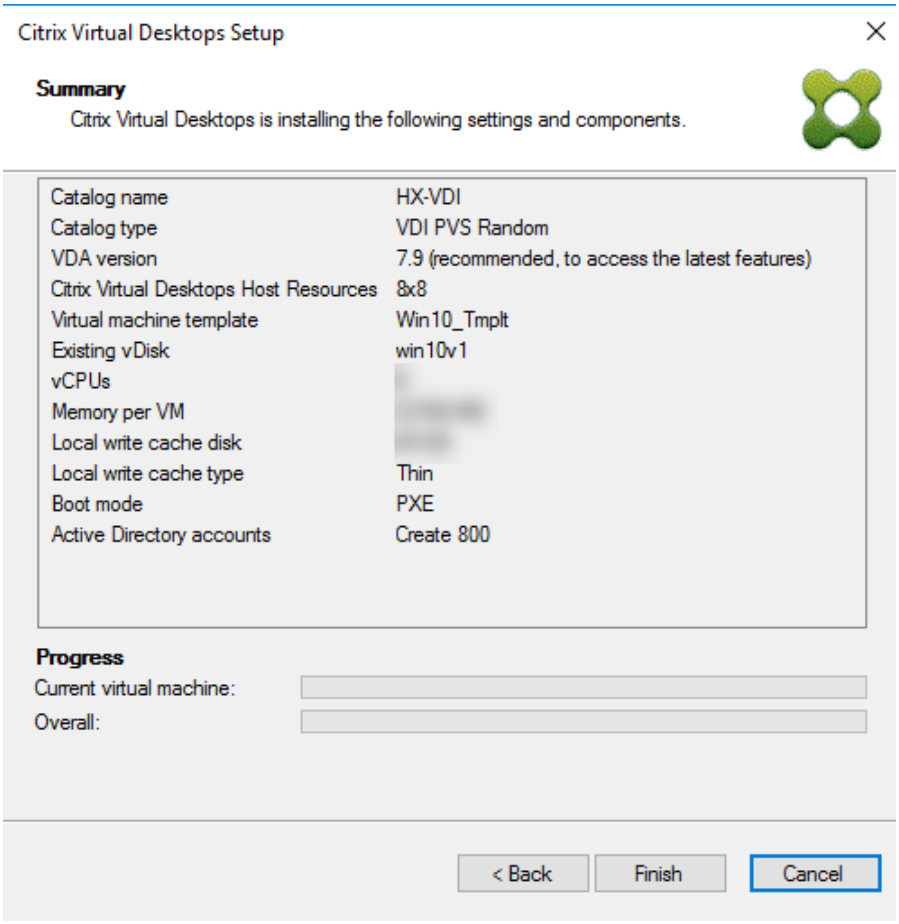

29. When the wizard is done provisioning the virtual machines, click Done.

Provisioning process takes ~10 seconds per machine.

30. Verify the desktop machines were successfully created in the following locations:

    a. PVS1 > Provisioning Services Console > Farm > Site > Device Collections > VDI-NP > CTX-VDI-001

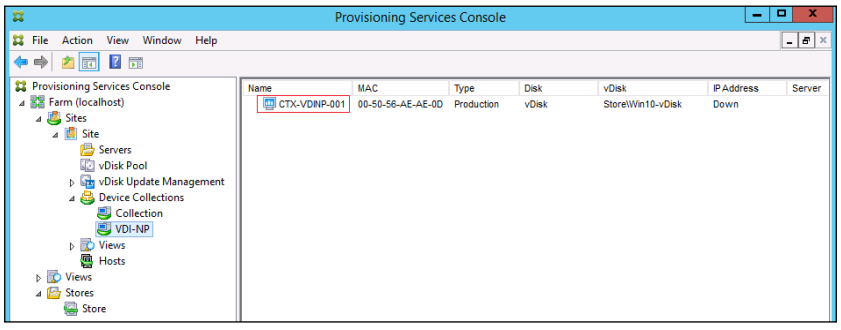

    b. CTX-XD1 > Citrix Studio > Machine Catalogs > VDI-NP

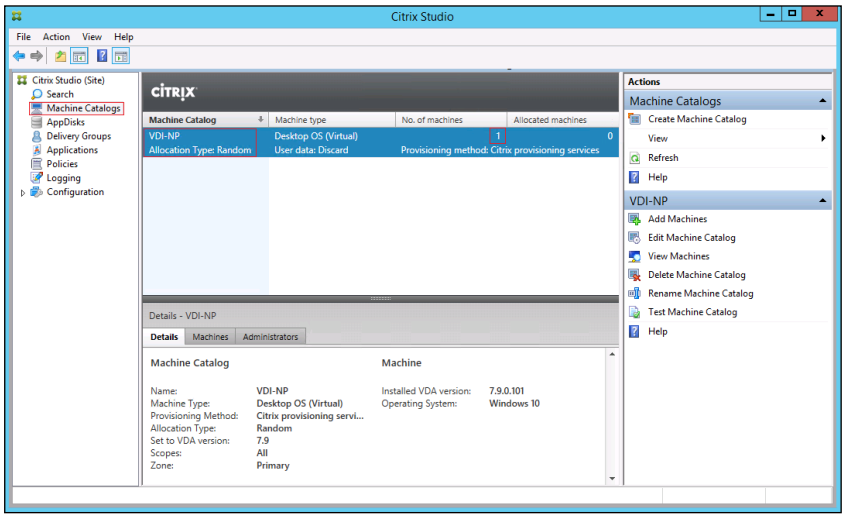

c. AD-DC1 > Active Directory Users and Computers > dvpod2.local > Computers > CTX-VDI-001

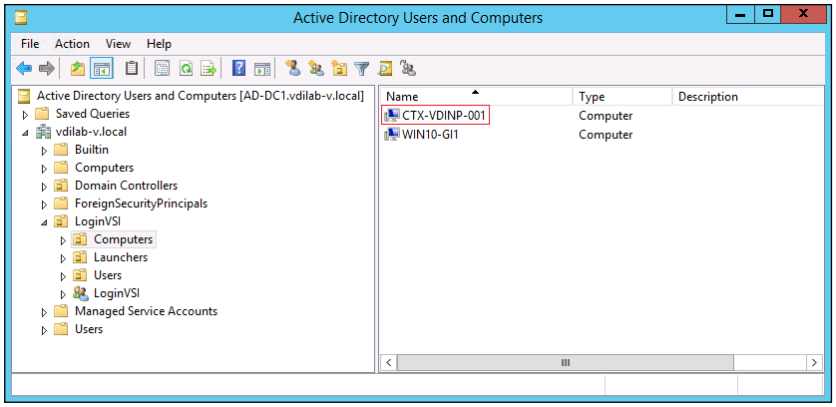

31. Log into the newly provisioned desktop machine, using the Virtual Disk Status verify the image mode is set to Ready Only and the cache type as Device Ram with overflow on local hard drive.

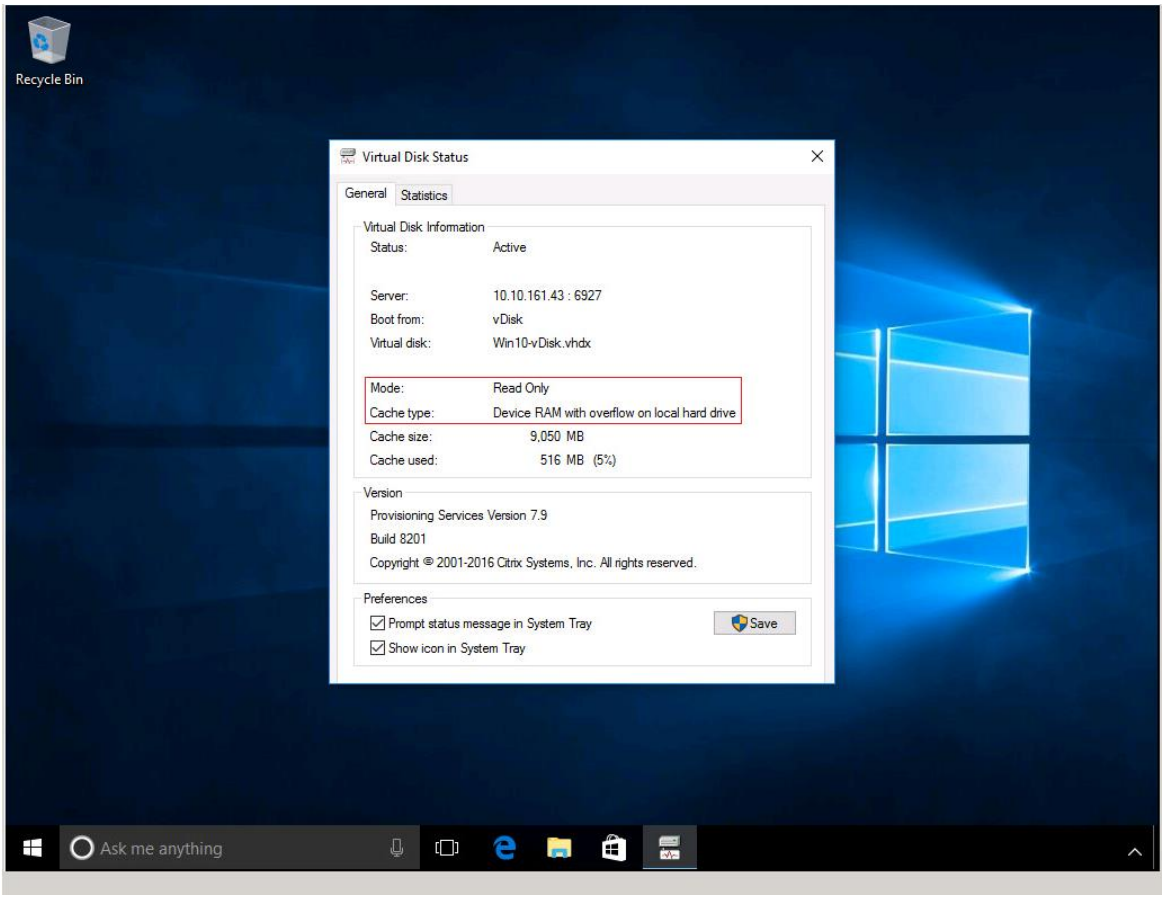

Install Citrix Virtual Apps and Desktop Virtual Desktop Agents

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 HVD and HSD environments.

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 was used for this CVD and is described in a subsequent section.

To install Citrix Desktop Virtual Desktop Agents, complete the following steps:

1. Launch the Citrix Desktop installer from the CVA Desktop 1808 ISO.

2. Click Start on the Welcome Screen.

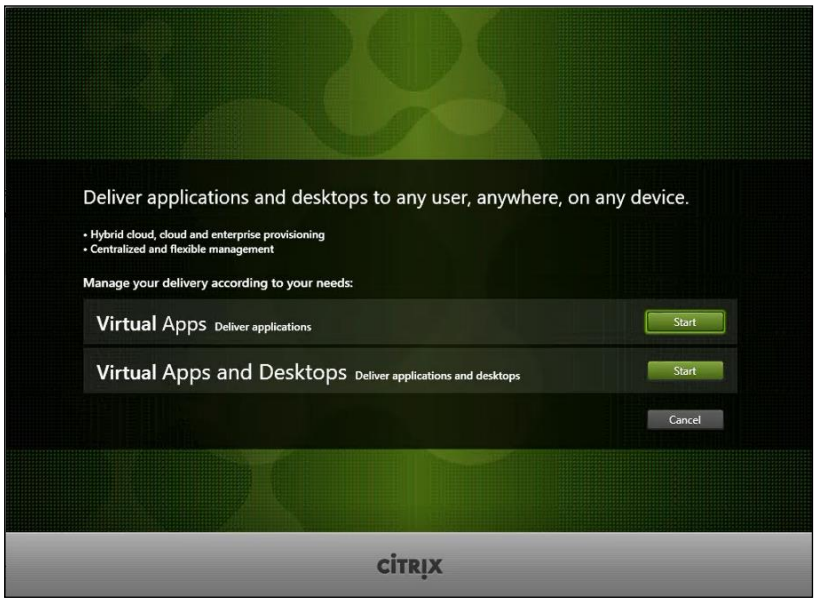

3. To install the VDA for the Hosted Virtual Desktops (VDI), select Virtual Delivery Agent for Windows Desktop OS. After the VDA is installed for Hosted Virtual Desktops, repeat the procedure to install the VDA for Hosted Shared Desktops (RDS). In this case, select Virtual Delivery Agent for Windows Server OS and follow the same basic steps.

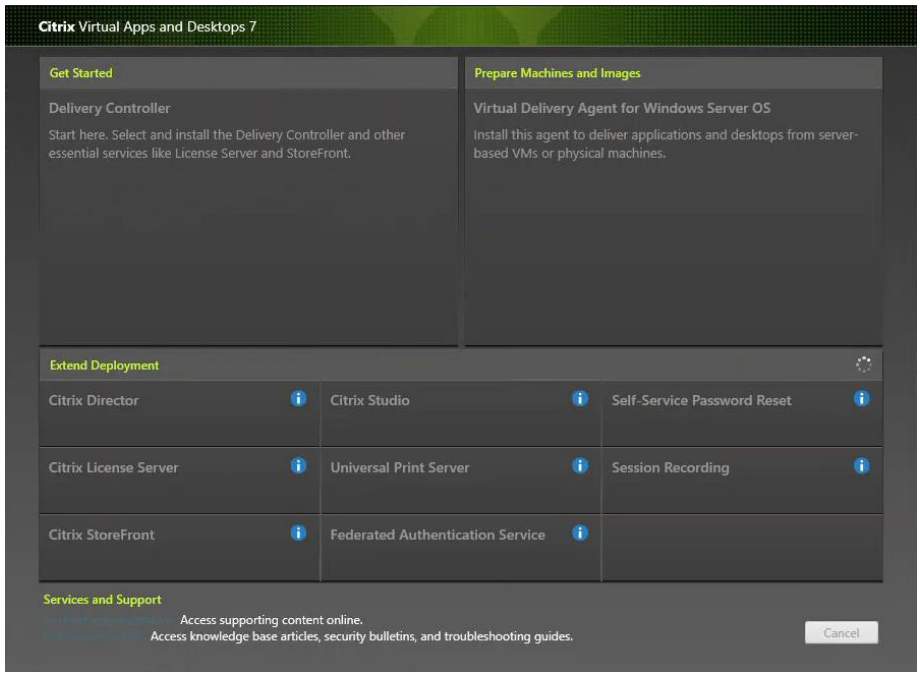

4. Select “Create a Master Image.”

5. Click Next.

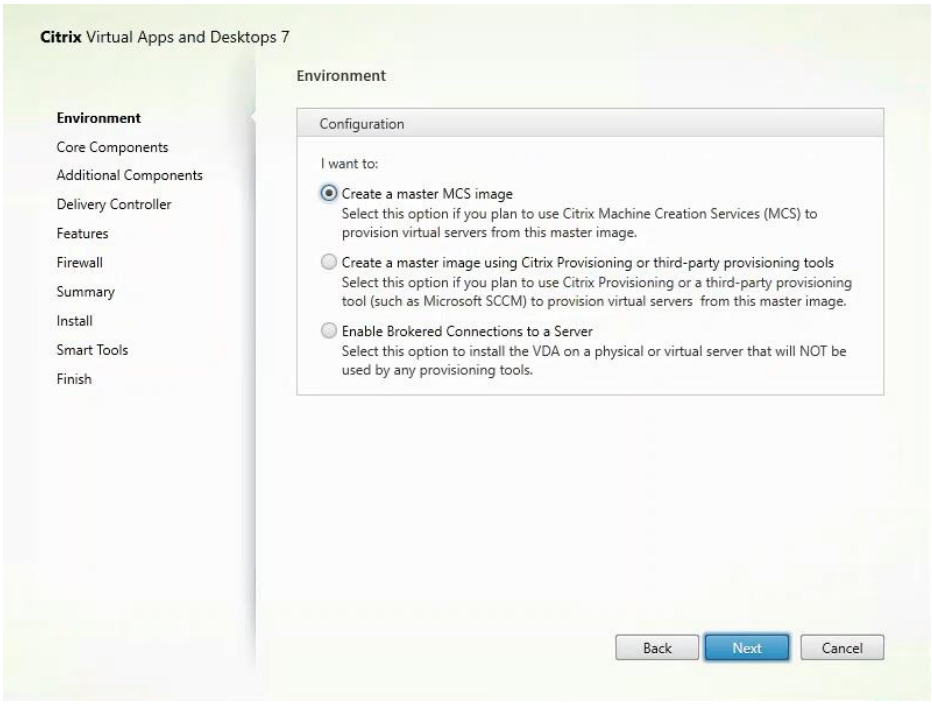

6. Optional: Select Citrix Workspace App.

7. Click Next.

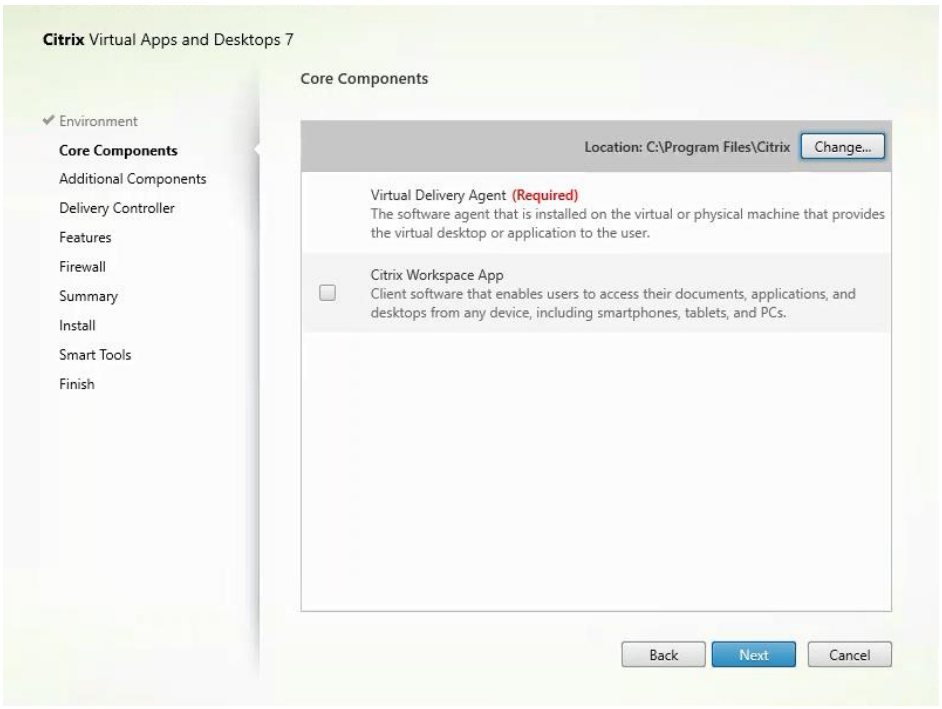

8. Click Next.

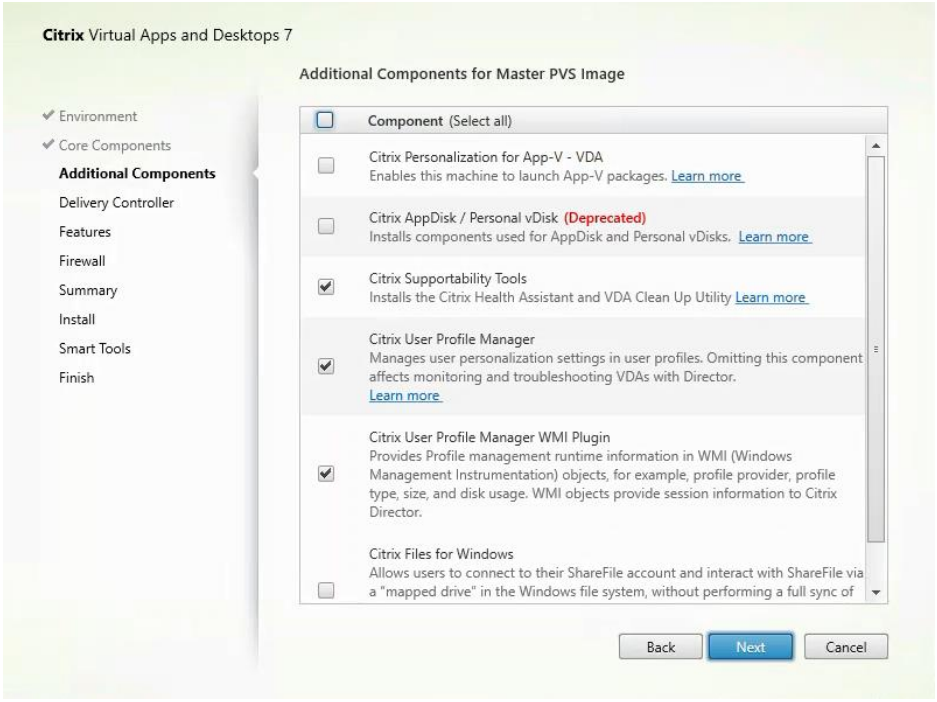

9. Select “Do it manually” and specify the FQDN of the Delivery Controllers.

10. Click Next.

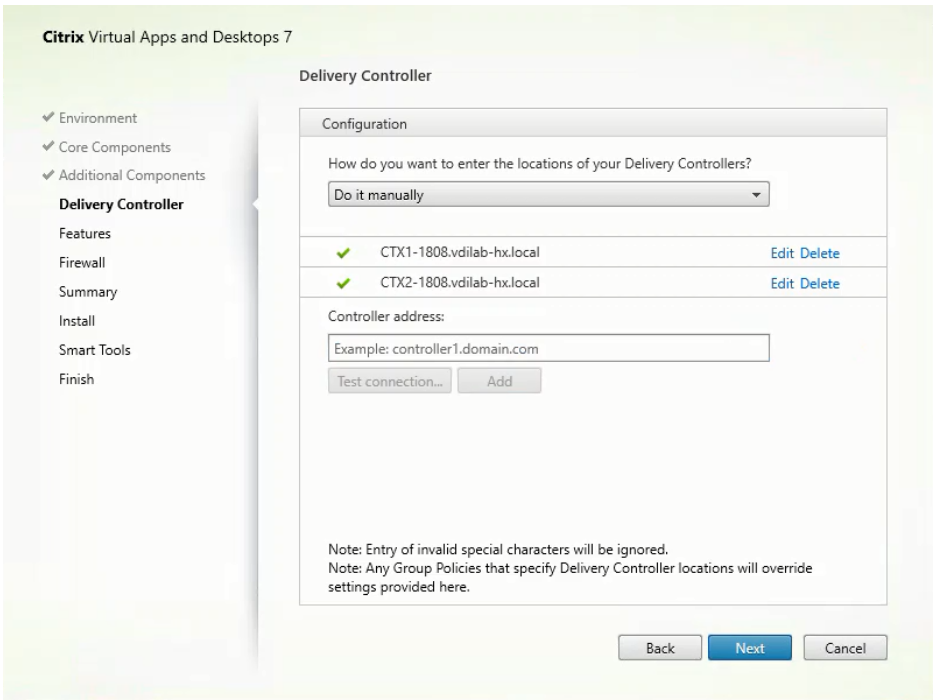

11. Accept the default features.

12. Click Next.

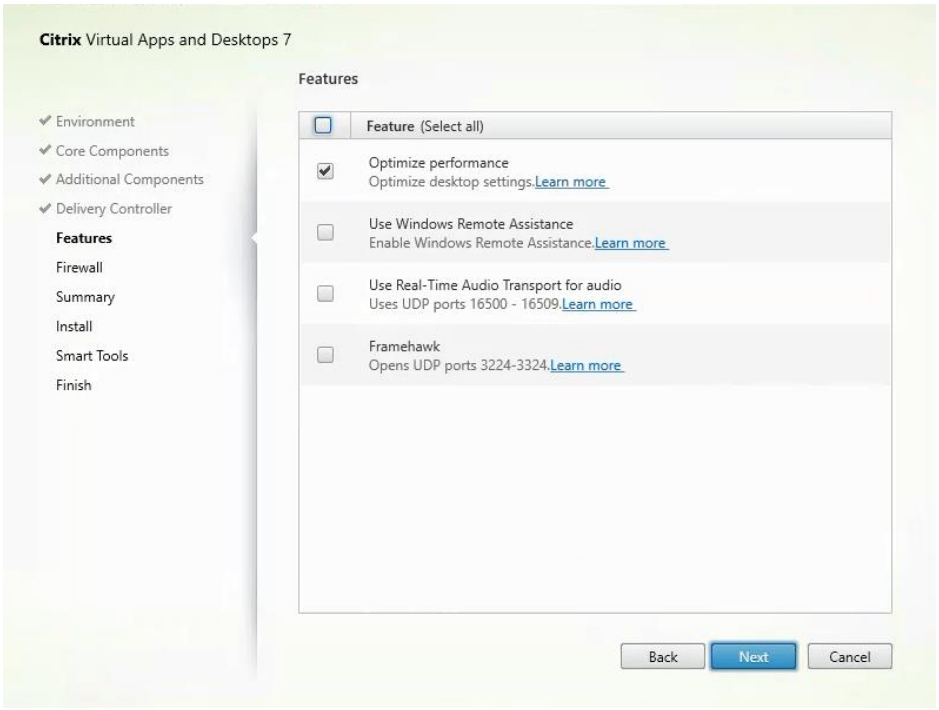

13. Allow the firewall rules to be configured automatically.

14. Click Next.

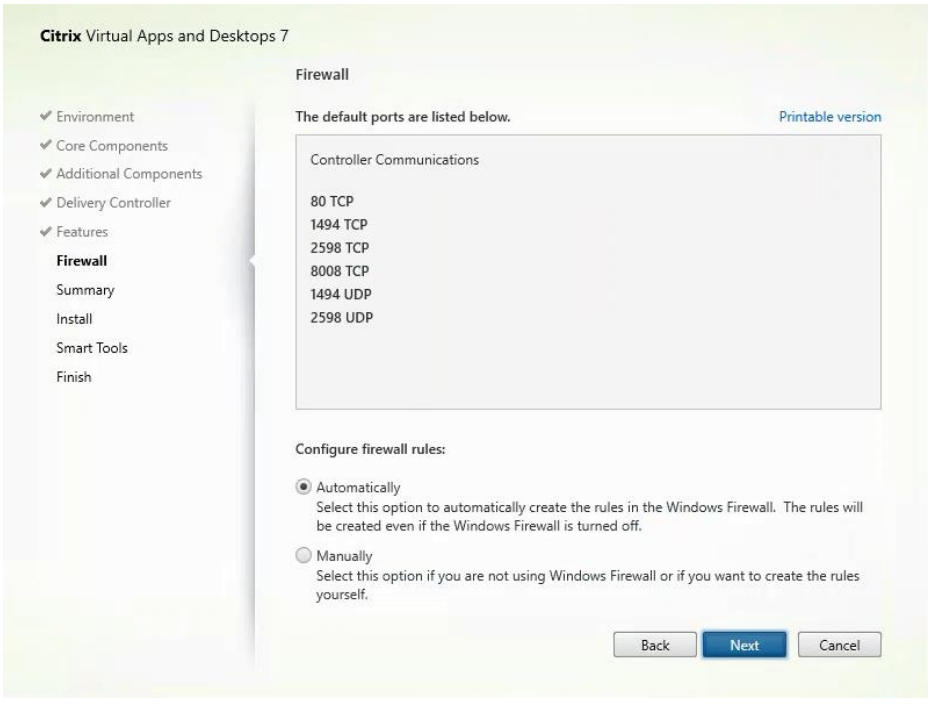

15. Verify the Summary and click Install.

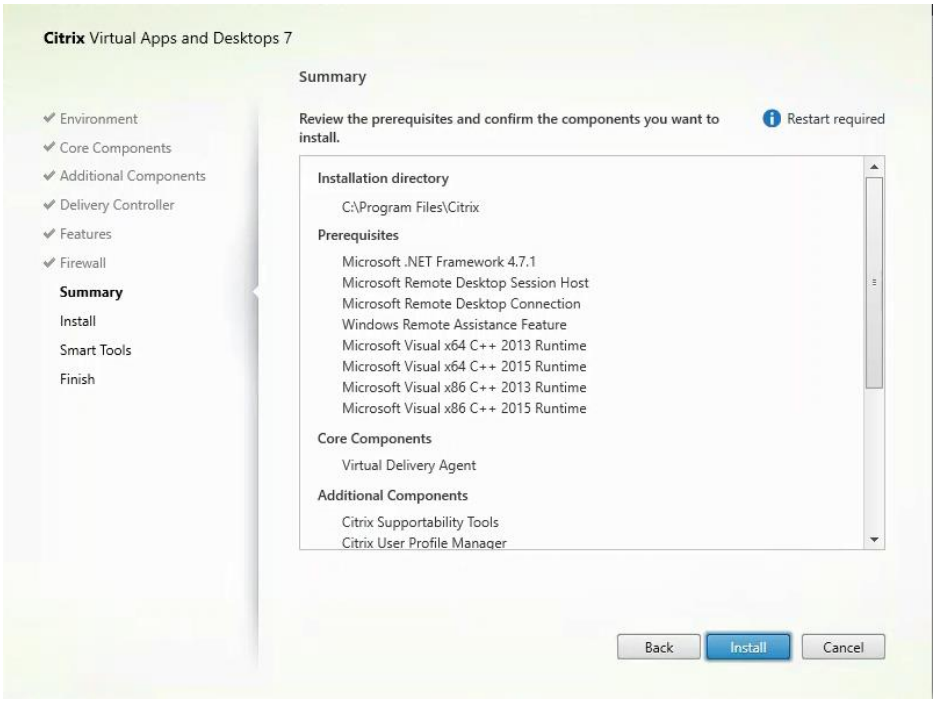

16. (Optional) Select Call Home participation.

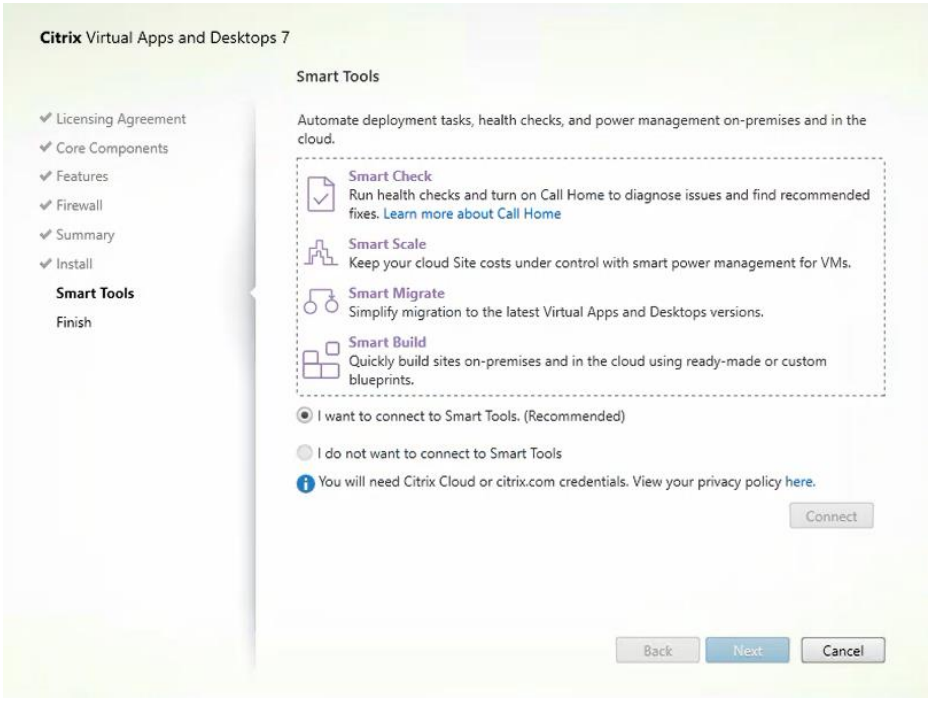

17. (Optional) check “Restart Machine.”

18. Click Finish.

19. Repeat the procedure so that VDAs are installed for both HVD (using the Windows 10 OS image) and the HSD desktops (using the Windows Server 2016 image).

20. Select an appropriate workflow for the HSD desktop.

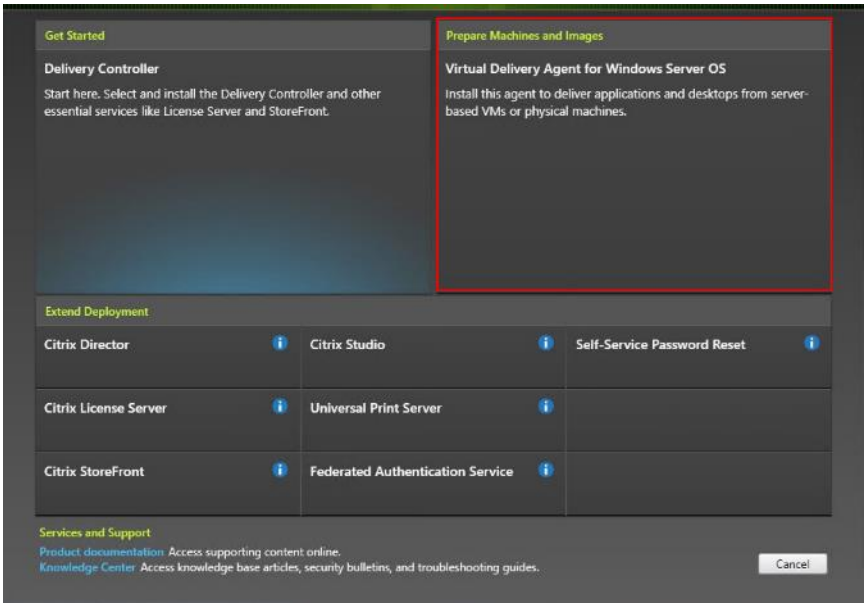

Create Delivery Groups

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, complete the following steps:

The instructions below outline the steps to create a Delivery Group for VDI desktops. When you have completed these steps, repeat the procedure to a Delivery Group for HVD desktops.

1. Connect to a Citrix Desktop server and launch Citrix Studio.

2. Choose Create Delivery Group from the drop-down list.

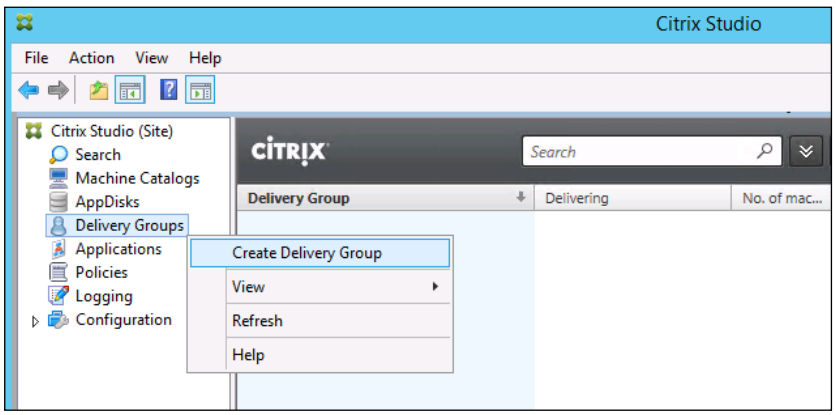

3. Click Next.

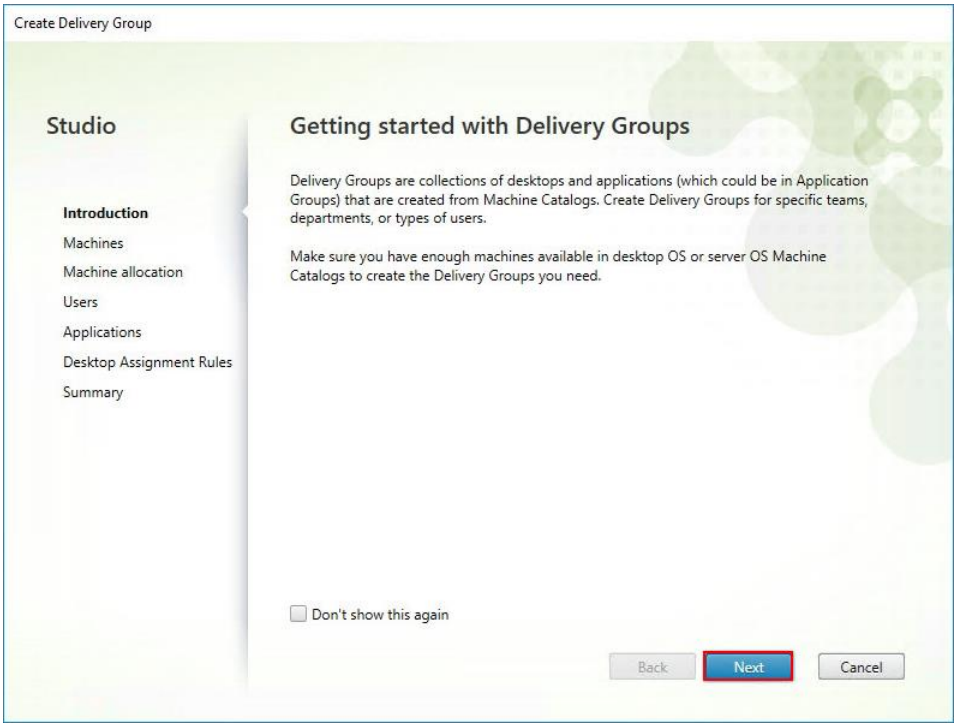

4. Select Machine catalog.

5. Provide the number of machines to be added to the delivery Group.

6. Click Next.

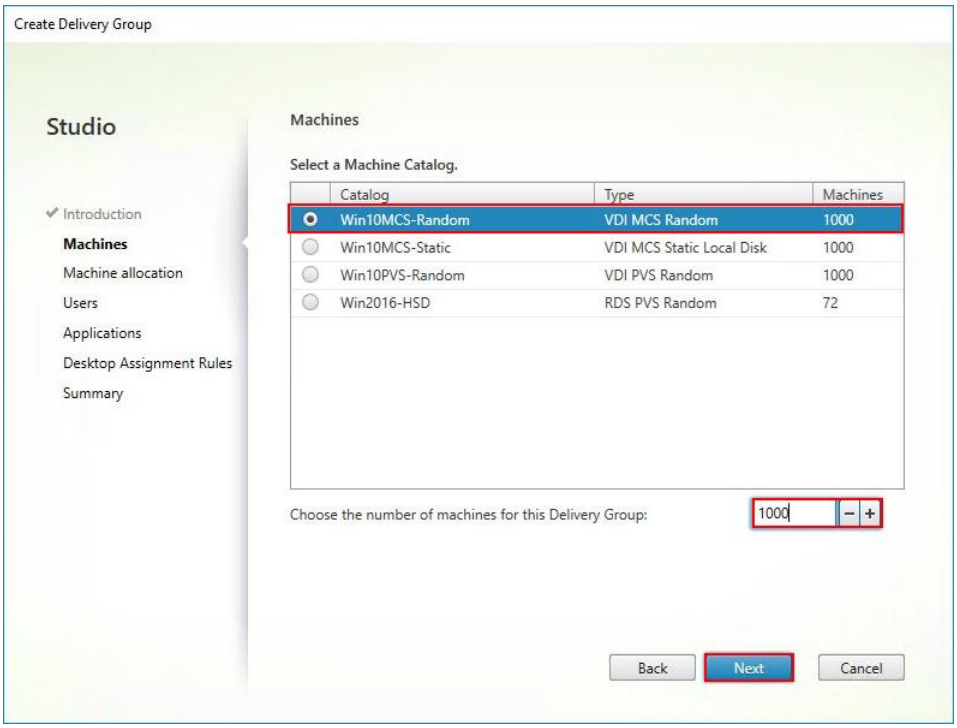

7. To make the Delivery Group accessible, you must add users, select Allow any authenticated users to use this Delivery Group.

8. Click Next.

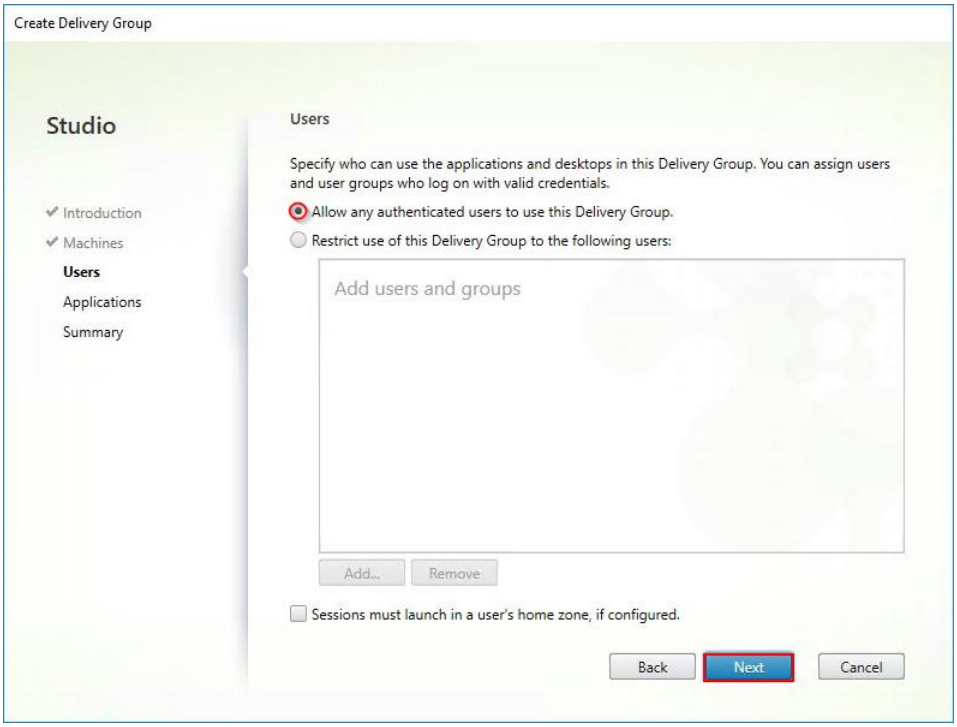

User assignment can be updated any time after Delivery group creation by accessing Delivery group properties in Desktop Studio.

9. (Optional) specify Applications catalog will deliver.

10. Click Next.

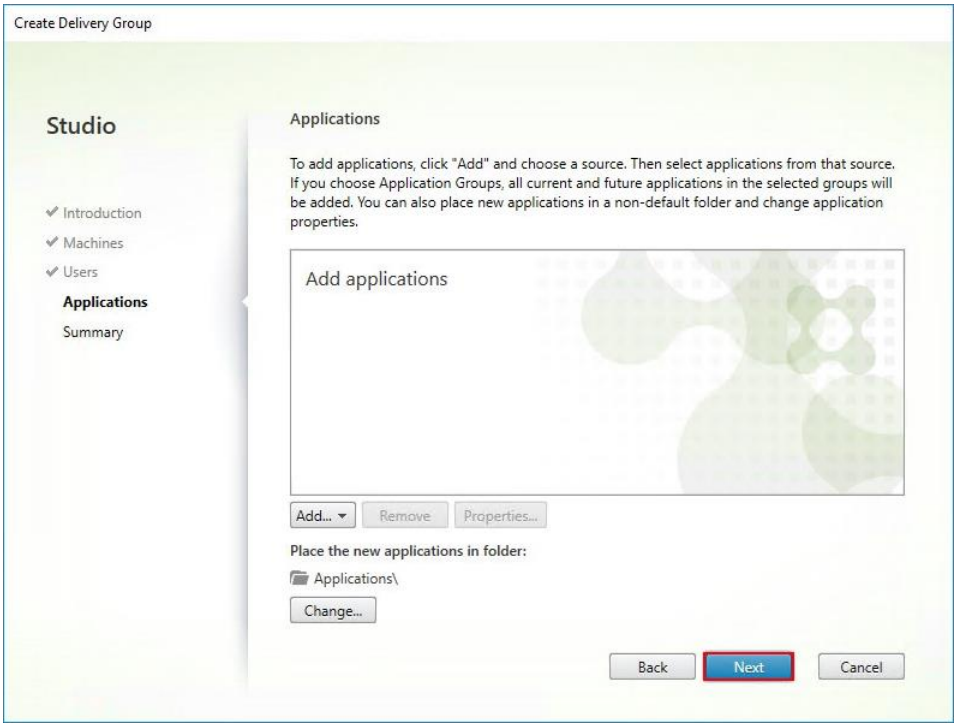

11. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Display name (for example, HVD or HSD).

12. Click Finish.

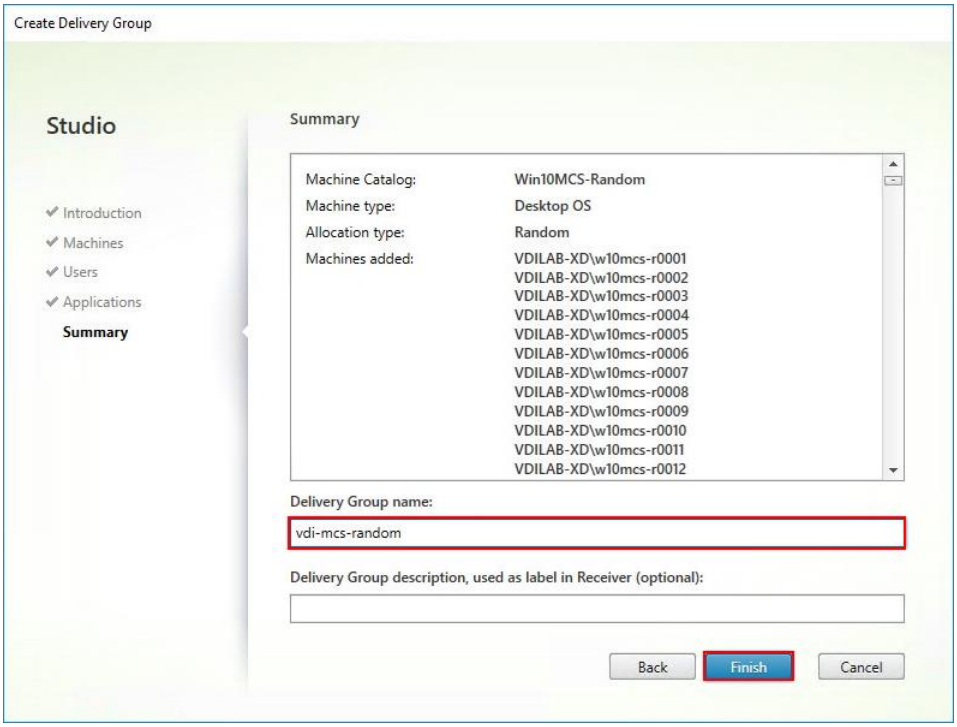

13. Citrix Studio lists the created Delivery Groups and the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab. Select Delivery Group and in Action List, select "Turn on Maintenance Mode."

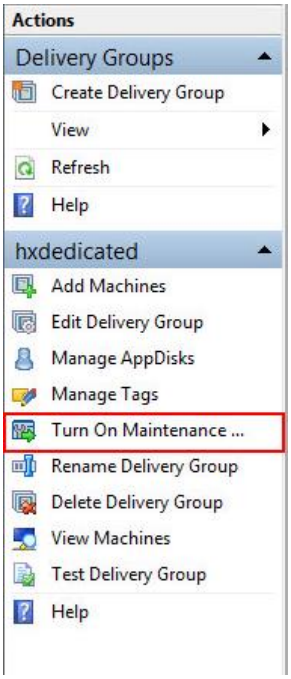

Citrix Desktop Policies and Profile Management

Policies and profiles allow the Citrix VDI environment to be easily and efficiently customized.

Configure Citrix Desktop Policies

Citrix VDI 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. (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 45 shows policies for Login VSI testing in this CVD.

Figure 45 Citrix Desktop Policy

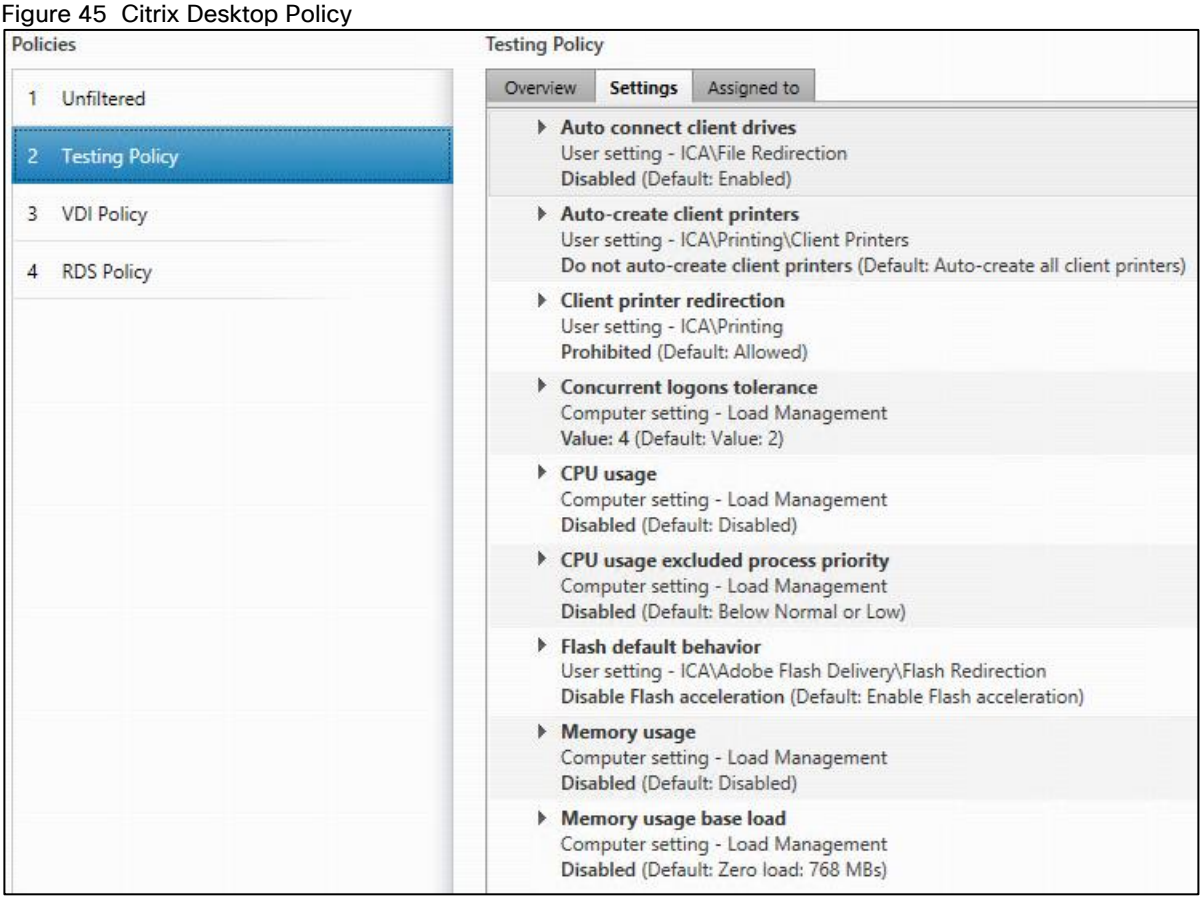

Configuring User Profile Management

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration, and provides users with fast logons and logoffs. 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. Examples of settings that can be customized are:

- Desktop settings such as wallpaper and screen saver
- Shortcuts and Start menu setting
- Internet Explorer Favorites and Home Page
- Microsoft Outlook signature
- Printers

Some user settings and data can be redirected by means of folder redirection. However, if folder redirection is not used these settings are stored within the user profile.

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature simplifies some of this decision-making for Citrix Desktop deployments. Screenshots of the User Profile Management interfaces that establish policies for this CVD's RDS and VDI users (for testing purposes) are shown below.

Basic profile management policy settings are documented here: https://docs.citrix.com/en-us/profile-management/1808.html

Figure 46 VDI User Profile Manager Policy

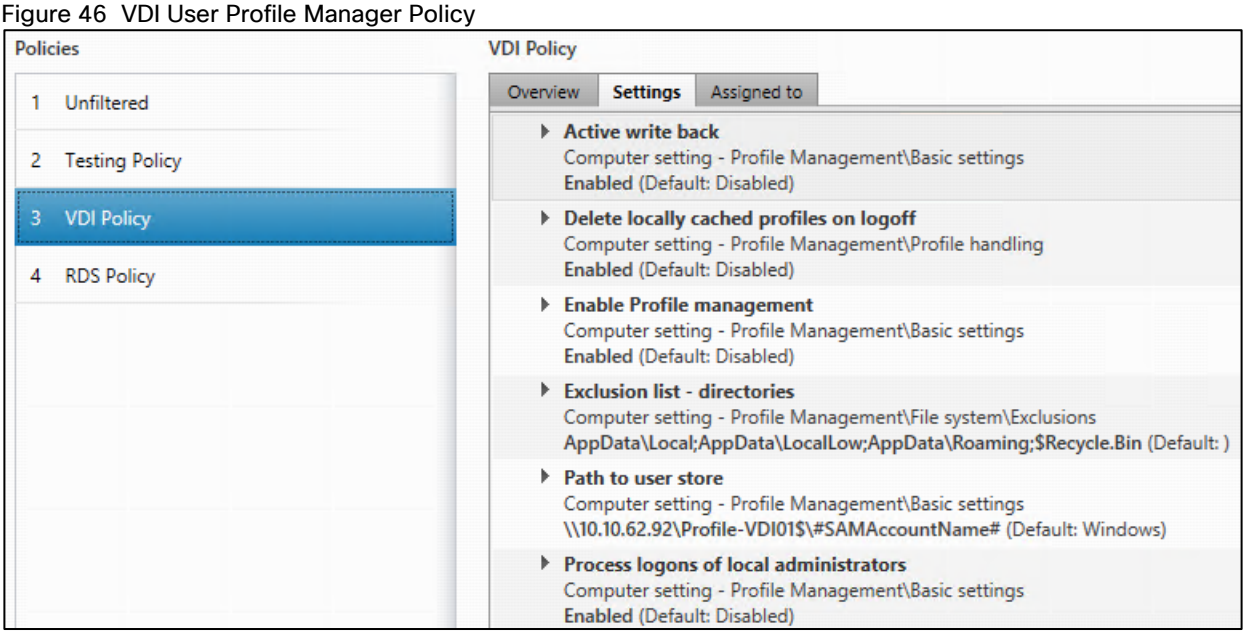

Test Setup and Configurations

In this project, we tested a single Cisco HyperFlex cluster running eight Cisco UCS HXAF220c-MS4 Rack Servers and eight Cisco UCS B200 M4 Blade Servers in a single Cisco UCS domain. This solution has been tested to illustrate linear scalability for each workload studied.

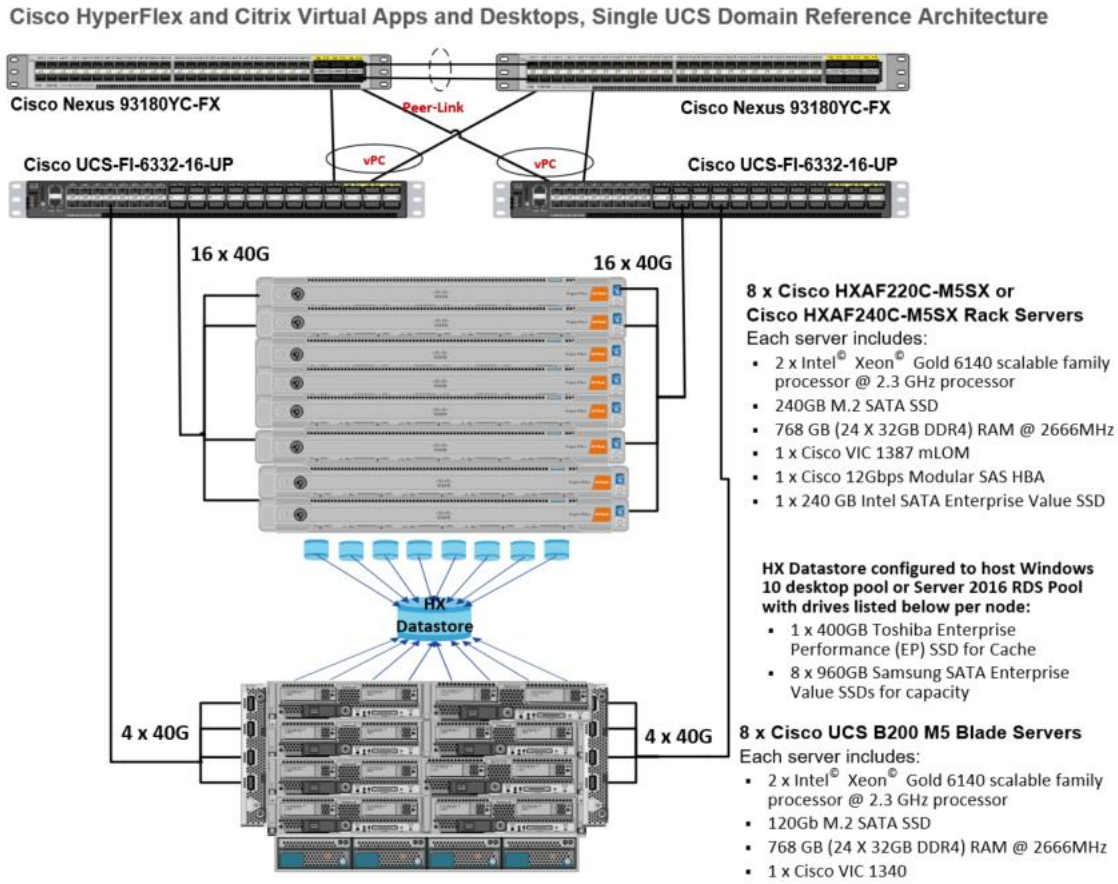

Hardware Components:

- 2 x Cisco UCS 6332 UP Fabric Interconnects
- 2 x Cisco Nexus 93180YC-FX Access Switches
- 8 x Cisco UCS HXAF220c-M5SX Rack Servers 2 Intel Xeon Gold 6140 scalable family processor at 2.3 GHz, with 768 GB of memory per server (32 GB x 24 DIMMs at 2666 MHz)
- Cisco VIC 1387 mLOM
- 12G modular SAS HBA Controller
- 240GB M.2 SATA SSD drive (Boot and HyperFlex Data Platform controller virtual machine)
- 240GB 2.5" 6G SATA SSD drive (Housekeeping)
- 400GB 2.5" 6G SAS SSD drive (Cache)

- 8 x 960GB 2.5” SATA SSD drive (Capacity)
- 1 x 32GB mSD card (Upgrades temporary cache)
- 8 x Cisco UCS B200 M5 Blade Servers (2 Intel Xeon processor 6140 CPUs at 2.3 GHz, with 768 GB of memory per server [32 GB x 24 DIMMs at 2666 MHz]).
- Cisco VIC 1340 mLOM
- 2 x 64GB SD card

Software components:

- Cisco UCS firmware 4.0(1b)
- Cisco HyperFlex data platform 3.5.1a
- VMware vSphere 6.5
- Citrix Virtual Apps and Desktops 1808
- Citrix StoreFront
- Windows File Server for User Profiles
- Microsoft SQL Server 2016
- Microsoft Windows 10
- Microsoft Windows Server 2016
- Microsoft Office 2016
- Login VSI 4.1.32.1

Testing Methodology and Success Criteria

All validation testing has been 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 HSD Servers 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 have been conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from http://www.loginvsi.com.

Testing Procedure

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

Pre-Test Setup for Testing

All virtual machines and HSD Servers have been shut down utilizing the Citrix Studio Administrator Console.

All Launchers virtual machines used for testing were 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.

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, we require all sessions started, whether 60 single server users or 2000 full-scale test users to become active within two minutes after the last session is launched.

In addition, Cisco requires that the Login VSI Benchmark method 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. Complete the following steps:

1. Time 0:00:00 Start PerfMon Logging on the following systems:
   - Infrastructure and VDI Host Blades used in test run
   - All Infrastructure virtual machines used in test run (AD, SQL, View Connection brokers, image mgmt., and so on.)
2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
3. Time 0:05: Boot RDS Machines using Citrix Studio Administrator Console.
4. Time 0:06 First machines boot.
5. Time 0:35 Single Server or Scale target number of RDS Servers registered on Citrix.

No more than 60 Minutes of rest time is allowed after the last desktop is registered and available on Citrix Studio Administrator Console dashboard. Typically a 20-30 minute rest period for Windows 10 desktops and 10 minutes for RDS virtual machines is sufficient.

6. Time 1:35 Start Login VSI 4.1.32.1 Knowledge Worker Benchmark Mode Test, setting auto-logoff time at 900 seconds, with Single Server or Scale target number of desktop virtual machines utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).
7. Time 2:23 Single Server or Scale target number of desktop virtual machines desktops launched (48 minute benchmark launch rate).
8. Time 2:25 All launched sessions must become active.

All sessions launched must become active for a valid test run within this window.

9. Time 2:40 Login VSI Test Ends (based on Auto Logoff 900 Second period designated above).
10. Time 2:55 All active sessions logged off.

11. All sessions launched and active must be logged off for a valid test run. The Citrix Studio Administrator Dashboard must show that all desktops have been returned to the registered/available state as evidence of this condition being met.

12. Time 2:57 All logging terminated; Test complete.

13. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows 7 machines.

14. Time 3:30 Reboot all hypervisors.

15. Time 3:45 Ready for new test sequence.

Success Criteria

Our "pass" criteria for this testing is as follows: Cisco will run tests at a session count levels that effectively utilize the server capacity measured by CPU, memory, storage and network utilization. We use Login VSI version 4.1.32.1 to launch Knowledge Worker workload sessions. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The Citrix Studio Server Dashboard will 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. Cisco's tolerance for Stuck Sessions is 0.5 percent (half of one percent.) If the Stuck Session count exceeds that value, we identify it as a test failure condition.

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

The purpose of this testing is to provide the data needed to validate Citrix VDI Hosted Shared Desktop with Citrix PVS provisioning using Microsoft Windows Server 2016 sessions on Cisco UCS HXAF220C-M5SX.

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 system performance and linear scalability.

VSImax 4.1.x Description

The philosophy behind Login VSI is different to conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put the faster the execution time or the bigger the throughput, the faster the system is according to the benchmark.

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for SBC or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer and Adobe PDF reader. By gradually increasing the numbers 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 solutions, 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, etc. 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 47 Sample of a VSI Max Response Time Graph, Representing a Normal Test

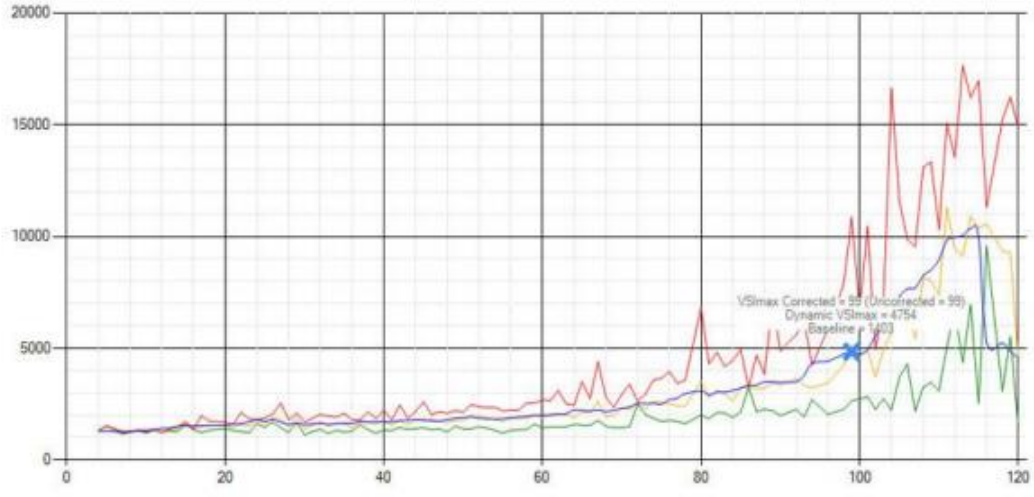

Figure 48 Sample of a VSI Test Response Time Graph with a Clear Performance Issue

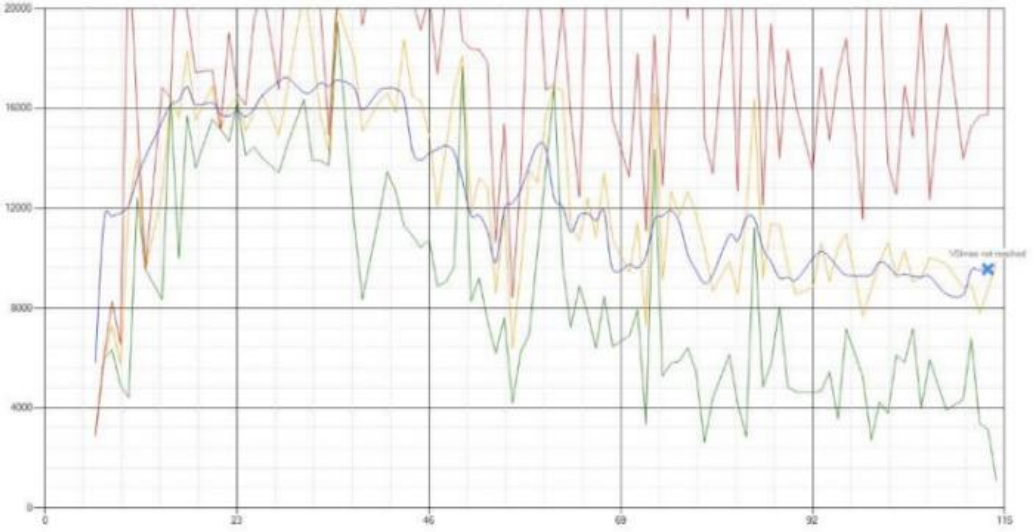

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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. This response time of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSImax models, this weighting much better represent system performance. All actions have very similar weight in the VSImax total. The following weighting of the response times are applied.

The following actions are part of the VSImax v4.1 calculation and are weighted as follows (US notation):

- Notepad File Open (NFO): 0.75
- Notepad Start Load (NSLD): 0.2
- Zip High Compression (ZHC): 0.125
- Zip Low Compression (ZLC): 0.2
- CPU: 0.75

This weighting is applied on the baseline and normal Login VSI response times.

With the introduction of Login VSI 4.1 we also created a new method to calculate the base phase of an environment. With the new workloads (Taskworker, Powerworker, etc.) enabling 'base phase' for a more reliable baseline has become obsolete. The calculation is explained below. In total 15 lowest VSI response time samples are taken from the entire test, the lowest 2 samples are removed and the 13 remaining samples are averaged. The result is the Baseline. The calculation is as follows:

- Take the lowest 15 samples of the complete test
- From those 15 samples remove the lowest 2
- Average the 13 results that are left is the baseline

The VSImax average response time in Login VSI 4.1.x is calculated on the amount of active users that are logged on the system.

Always a 5 Login VSI response time samples are averaged + 40 percent of the amount of “active” sessions. For example, if the active sessions is 60, then latest 5 + 24 (=40 percent of 60) = 31 response time measurement are used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5 percent and bottom 5 percent of the VSI response time samples are removed from the average calculation, with a minimum of 1 top and 1 bottom sample. As a result, with 60 active users, the last 31 VSI response time sample are taken. From those 31 samples the top 2 samples are removed and lowest 2 results are removed (5 percent of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

VSImax v4.1.x is reached when the VSIbase + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSImax response time can grow 2 - 3x the baseline average. In end-user computing, a 3x increase in response time in comparison to the baseline is typically regarded as the maximum performance degradation to be considered acceptable.

In VSImax v4.1.x this latency threshold is fixed to 1000ms, this allows better and fairer comparisons between two different systems, especially when they have different baseline results. Ultimately, in VSImax v4.1.x, the performance of the system is not decided by the total average response time, but by the latency is has under load. For all systems, this is now 1000ms (weighted).

The threshold for the total response time is: average weighted baseline response time + 1000ms.

When the system has a weighted baseline response time average of 1500ms, the maximum average response time may not be greater than 2500ms (1500+1000). If the average baseline is 3000 the maximum average response time may not be greater than 4000ms (3000+1000).

When the threshold is not exceeded by the average VSI response time during the test, VSImax is not hit and the amount of sessions ran successfully. This approach is fundamentally different in comparison to previous VSImax methods, as it was always required to saturate the system beyond VSImax threshold.

Lastly, VSImax v4.1.x is now always reported with the average baseline VSI response time result. For example: “The VSImax v4.1 was 125 with a baseline of 1526ms”. This helps considerably in the comparison of systems and gives a more complete understanding of the system. The baseline performance helps to understand the best performance the system can give to an individual user. VSImax indicates what the total user capacity is for the system. These two are not automatically connected and related:

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

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

With Login VSI 4.1.x a new VSImax method is introduced: VSImax v4.1. This methodology gives much better insight in system performance and scales to extremely large systems.

Test Results

Boot Storms

A key performance metric for desktop virtualization environments is the ability to boot the virtual machines quickly and efficiently to minimize user wait time for their desktop.

As part of Cisco's virtual desktop test protocol, we shut down each virtual machine at the conclusion of a benchmark test. When we run a new test, we cold boot all 2000 desktops and measure the time it takes for the 2000th virtual machine to register as available in the Citrix Administrator console.

The Cisco HyperFlex HXAF220cM5SX and Cisco UCS B200 M5 cluster running Data Platform version 3.5(1a) software can accomplish this task in 15 minutes as shown in the following charts:

Figure 49 2000 Citrix Server 2016 HSD Sessions with Office 2016 Virtual Desktops, Windows 10 Office 2016 Virtual Desktops Boot and Register as Available in Less Than 30 Minutes

Recommended Maximum Workload and Configuration Guidelines

Eight Node Cisco HXAF220c-M5S Rack Server and Eight Node Cisco UCS B200 M5 HyperFlex Cluster

For Citrix Virtual Apps and Desktops, RDS Hosted Shared Desktop and Hosted Virtual Desktop use case, the recommended maximum workload was determined based on both Login VSI Knowledge Worker workload end user experience measures and HXAF220c-M5S and Cisco UCS B200 M5 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 insure 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.

Callouts have been added throughout the data charts to indicate each phase of testing.

Test Phase	Description
Boot	Start all RDS and/or VDI virtual machines at the same time.
Login	The Login VSI phase of 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.
Logoff	Sessions finish executing the Login VSI workload and logoff.

The recommended maximum workload for a Cisco HyperFlex cluster configured on Cisco HXAF220c-M5S, and Cisco UCS B200 M5 nodes with 6140 Gold processors and 768GB of RAM for Windows Server 2016 Hosted Sessions and persistent/non-persistent Hosted Virtual Desktop users is 2000 sessions with Office 2016 virtual desktops respectively.

2000 User Full-Scale Testing on Sixteen-Node Cisco HyperFlex Cluster

This section details the key performance metrics that were captured on the Cisco UCS HyperFlex storage cluster configured with sixteen HXAF220c-M5S converged node and sixteen compute-only node (Eight Cisco UCS HXAF220C-M5 and Eight Cisco UCS B200 M5) running HSD virtual machines and VDI non -persistent/persistent performance monitoring during the full-scale testing.

The full-scale testing with 2000 users comprised of 800 RDS Hosted Server Sessions, 800 VDI Non-Persistent using Citrix PVS, and 400 VDI persistent full clone virtual machines virtual machines using Citrix MCS.

Test result highlights include:

- 0.708 second baseline response time
- 0.935 second average response time with 2000 desktop sessions running
- Average CPU utilization of 60 percent during steady state
- Average of 400 GB of RAM used out of 768 GB available
- 200 Mbps peak network utilization per host.
- Average Read Latency 1.1ms/Max Read Latency 0.75ms

- Average Write Latency 2.16ms/Max Write Latency 3.09ms
- 10000 peak I/O operations per second (IOPS) per cluster at steady state
- 600MBps peak throughput per cluster at steady state
- 29 percent Deduplication savings
- 59 percent Compression savings
- Total of 71 percent storage space savings

Figure 50 LoginVSI Analyzer Chart for 2000 Users Test

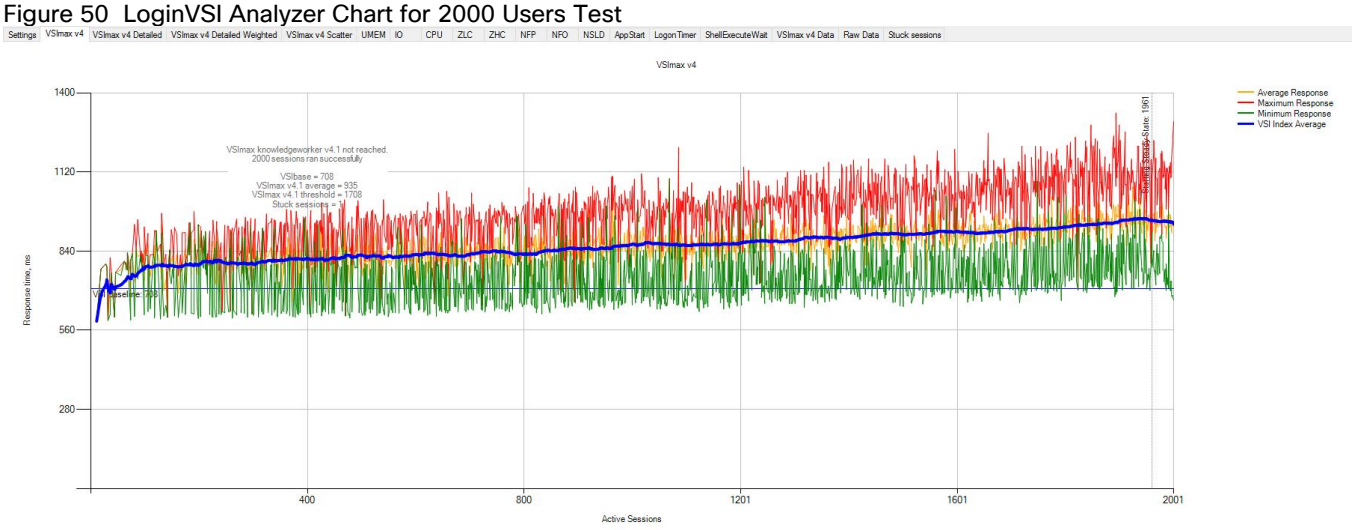

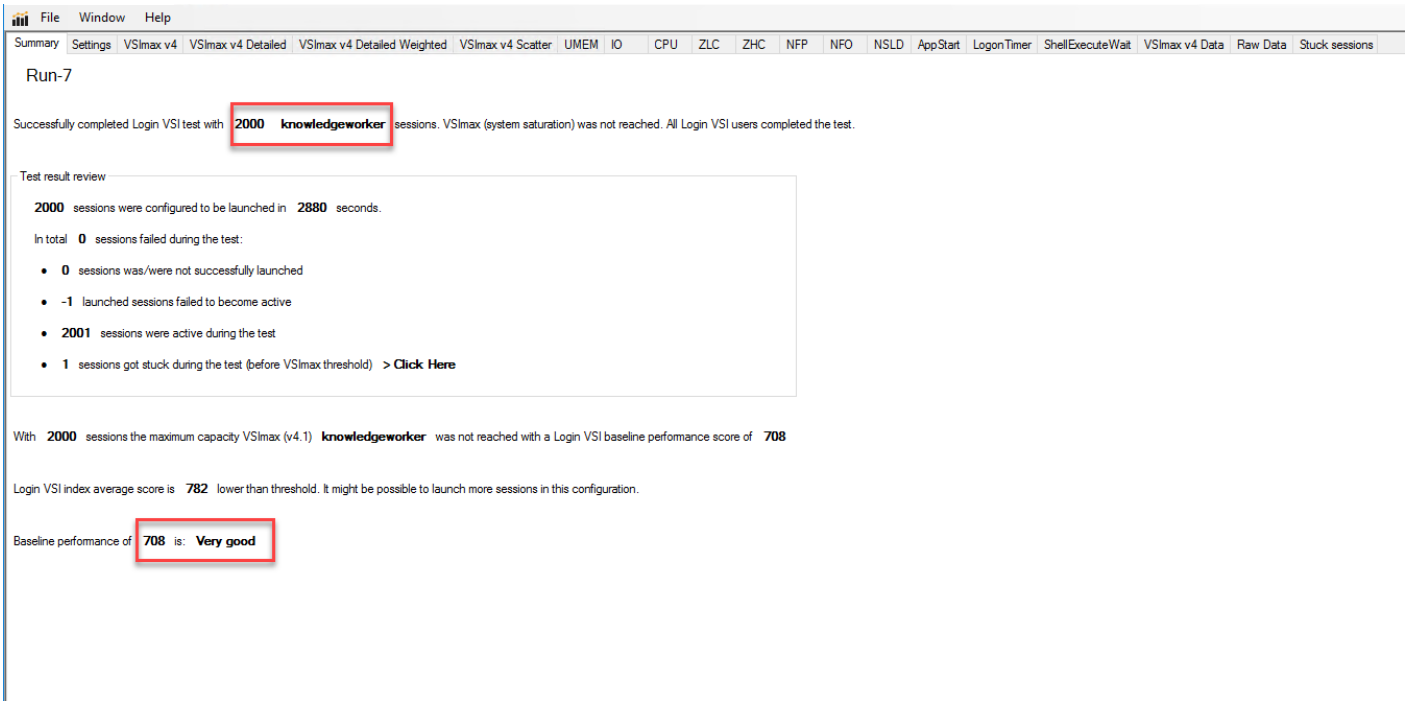

Figure 51 LoginVSI Analyzer Chart for Four Consecutive Test Running 2000 Knowledge Workload on 16 Node HyperFlex Cluster

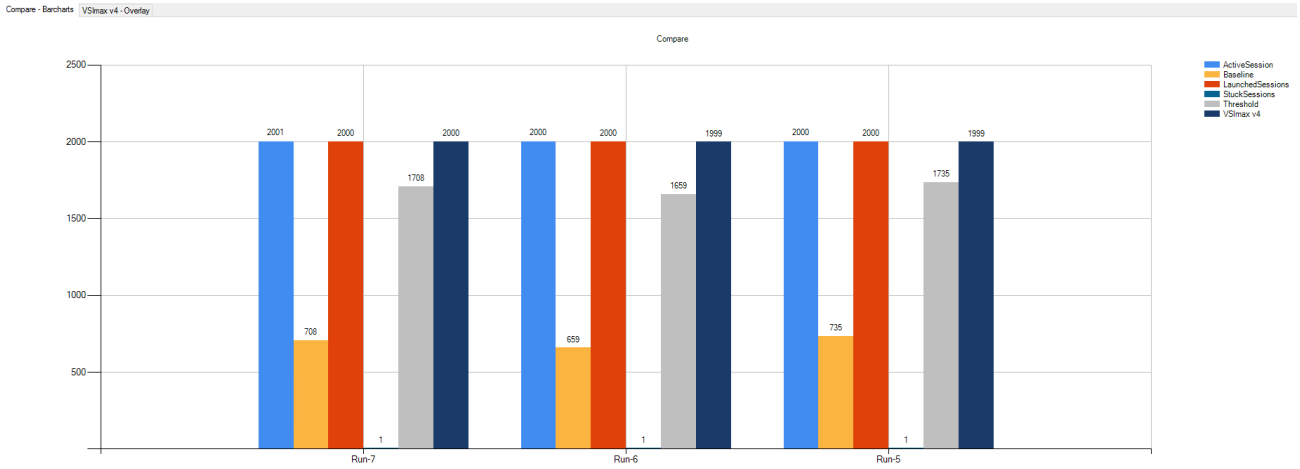

Figure 52 ESXi host CPU Core Utilization Running 2000 User Test on 8 Converged Nodes

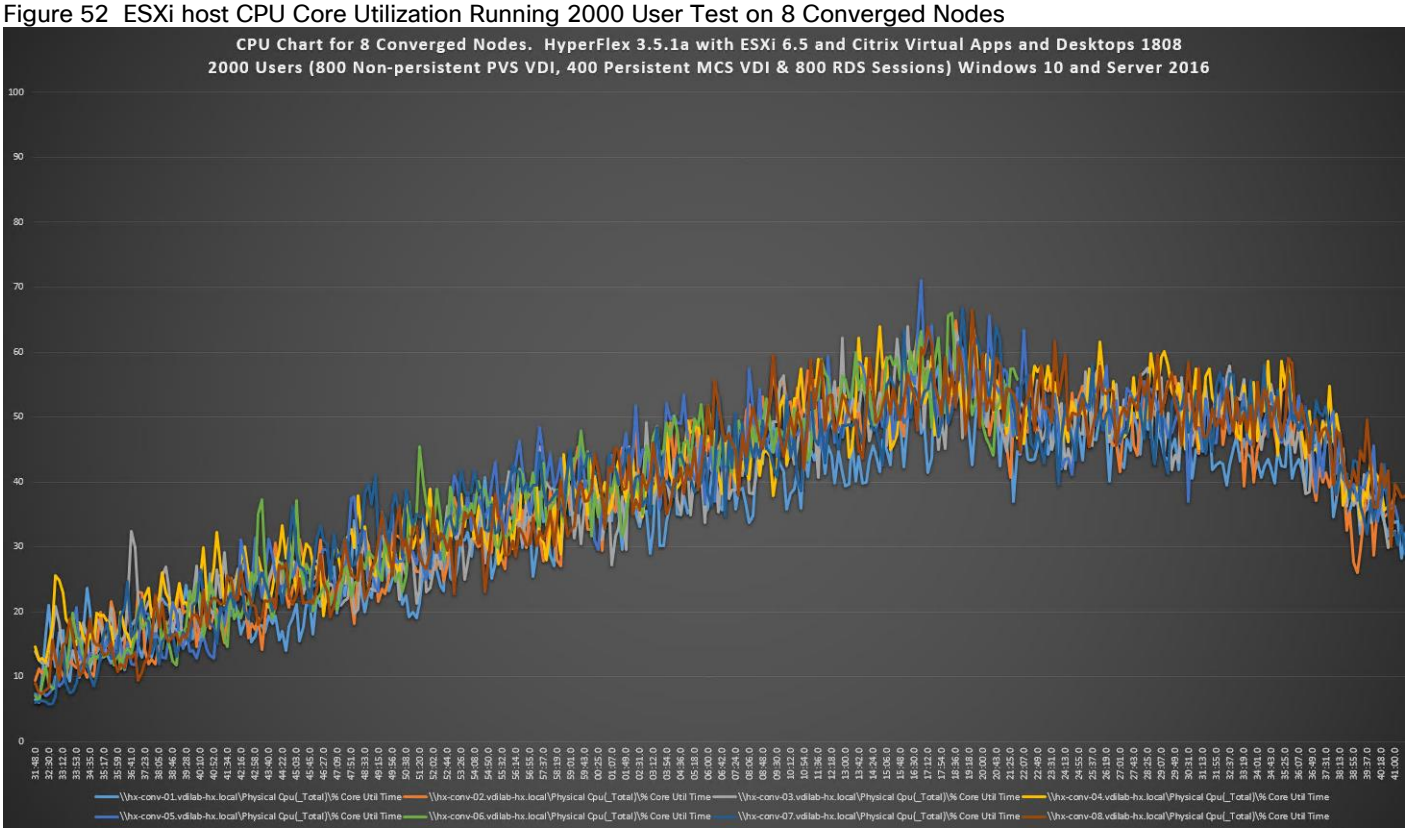

Figure 53 ESXi host CPU Core Utilization Running 2000 User Test on 8 Compute-Only Nodes

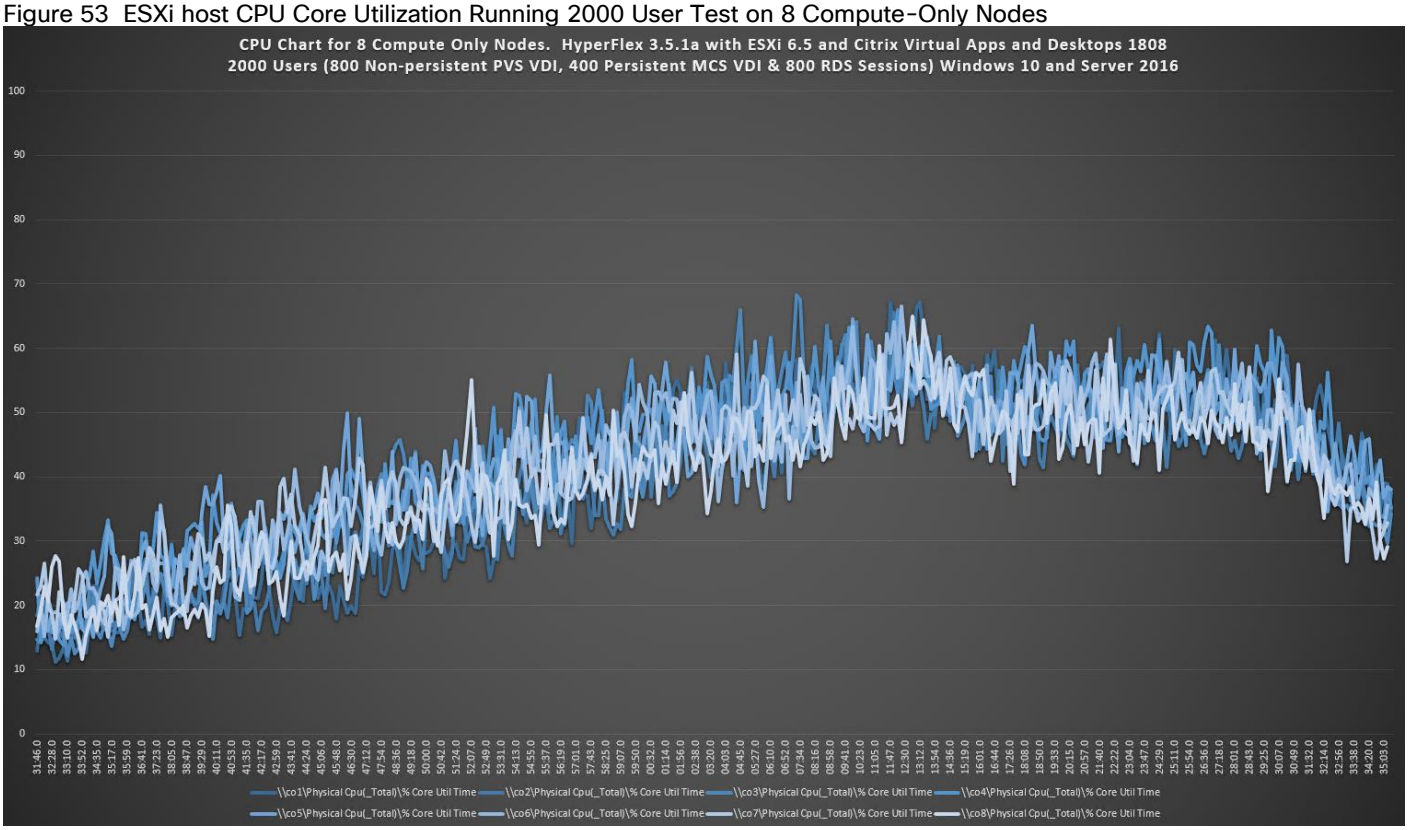

Figure 54 Sample ESXi Host Memory Usage in Mbytes running 2000 User Test on 16 Node

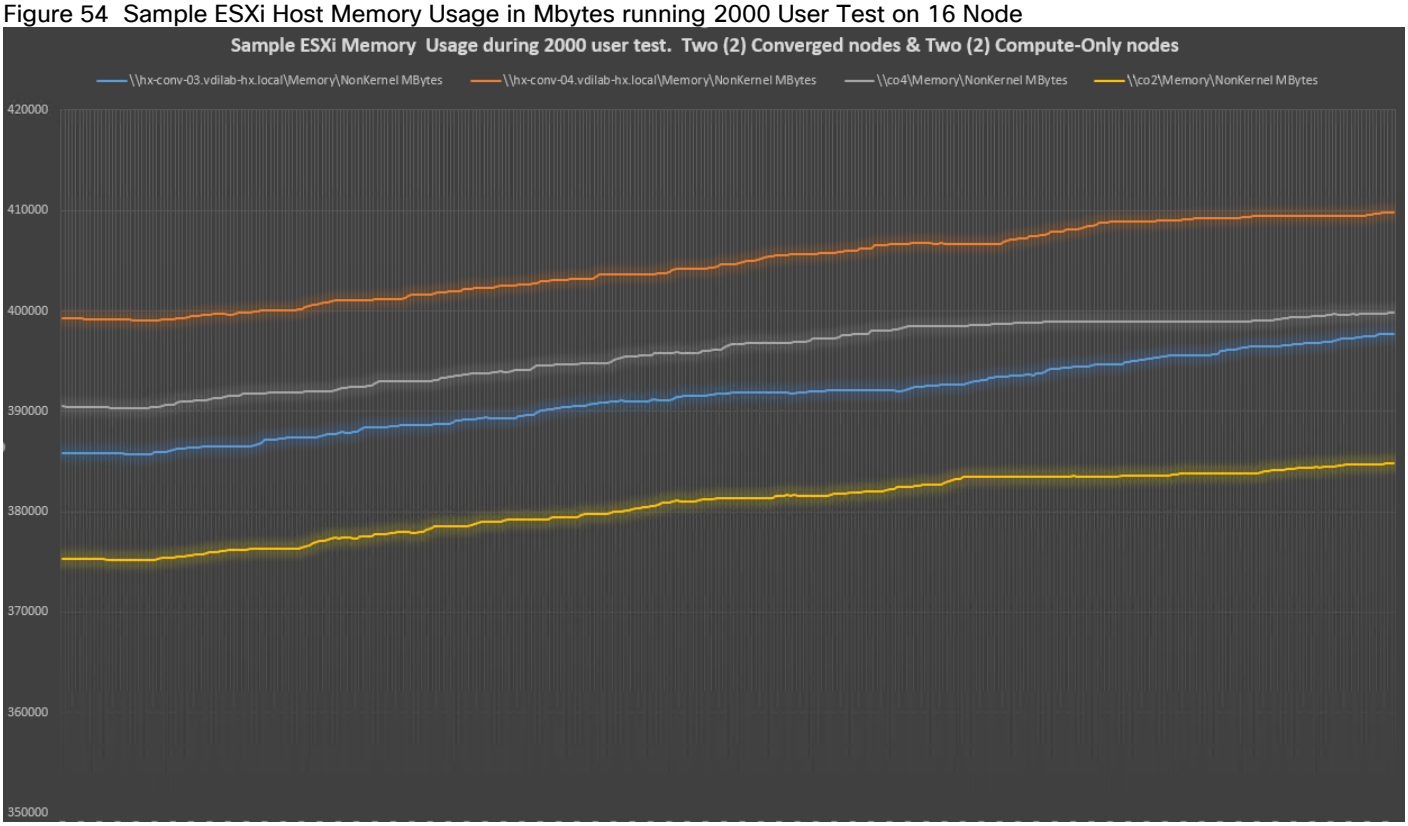

Figure 55 Sample ESXi Host Network Adapter (VMNICs) Mbits Received/ Transmitted Per Sec Running 2000 User Test on 16 Node

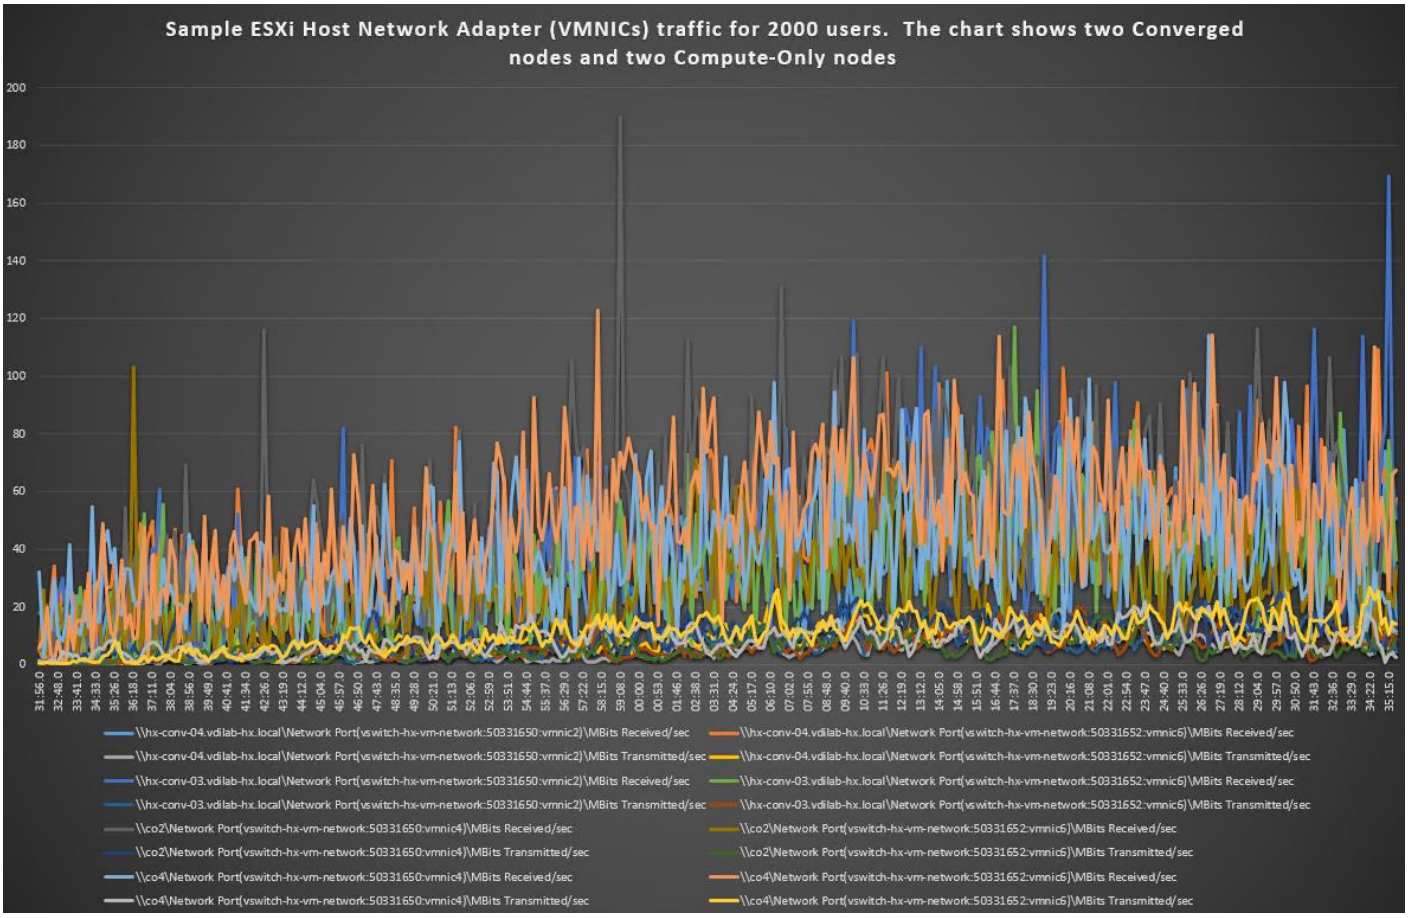

Figure 56 HyperFlex Cluster WebUI Performance Chart for Knowledge Worker Workload Running 2000 User Test on 16 Node

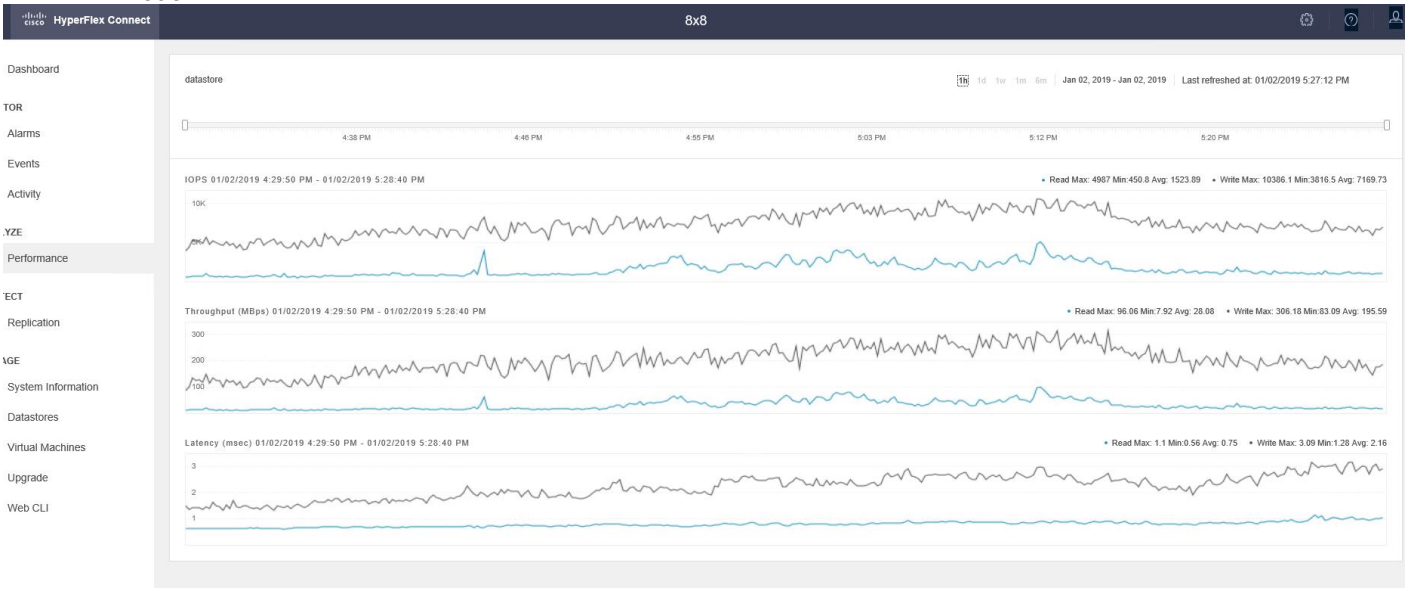

Figure 57 vCenter WebUI Reporting HyperFlex Cluster De-duplication and Compression Savings for 2000 User Sessions Supported on Windows Server 2016 Based Hosted Shared Sessions and Windows 10 Based Hosted Virtual Desktops Deployed on 16 Node HyperFlex Cluster

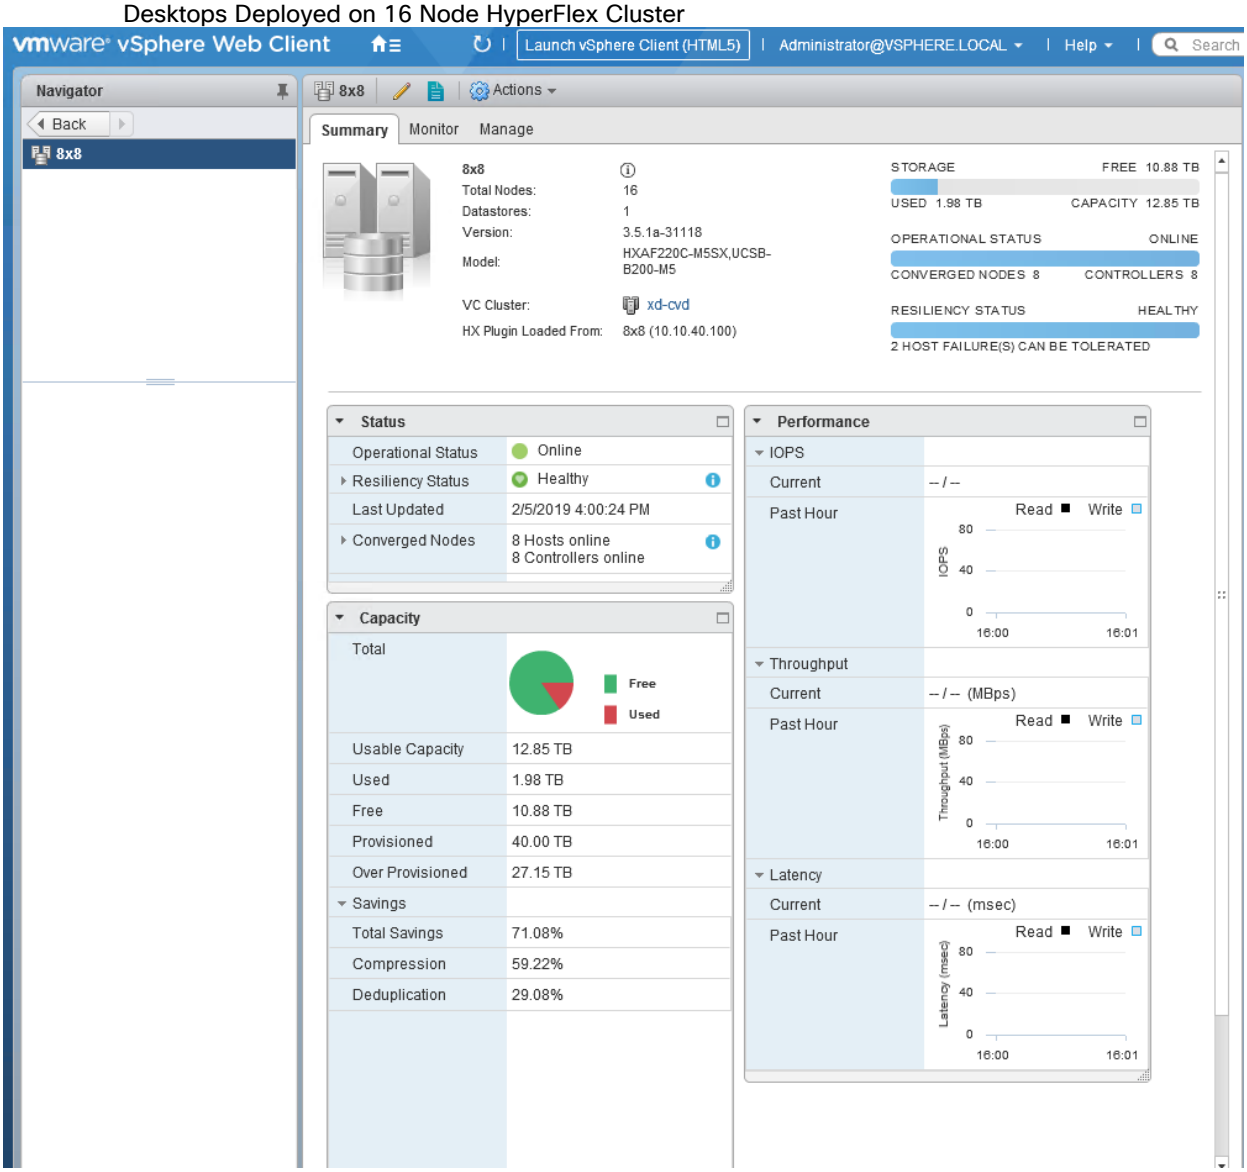

Summary

This Cisco HyperFlex solution addresses the urgent needs of IT by delivering a platform that is cost effective and simple to deploy and manage. The architecture and approach used provides for a flexible and high-performance system with a familiar and consistent management model from Cisco. In addition, the solution offers numerous enterprise-class data management features to deliver the next-generation hyperconverged system.

Only Cisco offers the flexibility to add compute only nodes to a true hyperconverged cluster for compute intensive workloads like desktop virtualization. This translates to lower cost for the customer, since no hyperconvergence licensing is required for those nodes.

Delivering responsive, resilient, high performance Citrix Virtual Apps and Desktops provisioned Microsoft Windows 10 Virtual Machines and Microsoft Windows Server 2016 for hosted Apps or desktops has many advantages for desktop virtualization administrators.

The sixteen node tested system can be expanded to 64 nodes (32 hyper converged plus 32 compute only nodes in a single UCS rack solution) for an expected user capacity of 8800 knowledge worker users.

The solution if fully capable of supporting graphics accelerated workloads. Each Cisco HyperFlex HXAF240c M5 node and each Cisco UCS C240 M5 server can support up to two NVIDIA M10 or P40 cards or up to six NVIDIA P4 cards. The Cisco UCS B200 M5 server supports up to two NVIDIA P6 cards for high density, high performance graphics workload support. See the Cisco Graphics White Paper for our fifth generation servers with NVIDIA GPUs and software for details on how to integrate this capability with VMware & Citrix.

Virtual desktop end-user experience, as measured by the Login VSI tool in benchmark mode, is outstanding with Intel Xeon Scalable Family processors and Cisco 2666Mhz memory. In fact, we set a new industry standard in performance for Desktop Virtualization on a hyperconverged platform.

About the Authors

Jeff Nichols, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Jeff Nichols is a Cisco Unified Computing System architect, focusing on Virtual Desktop and Application solutions with extensive experience with Microsoft ESX/Hyper-V, Citrix VDI, XenApp and Microsoft Remote Desktop Services. He has expert product knowledge in application, desktop and server virtualization across all three major hypervisor platforms and supporting infrastructures including but not limited to Windows Active Directory and Group Policies, User Profiles, DNS, DHCP and major storage platforms.

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:

- Mike Brennan, Product Manager, Virtual Client Computing and Graphics Solutions, Cisco Systems, Inc.

Appendix A – Cisco Nexus 93180 Switch Configuration

Switch A Configuration

```
!Command: show running-config.
!Time: Fri Nov 13 17:17:40 2018
version 7.0(3)I7(2)
switchname XXXXXXXXXXX
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
vdc XXXXXXXXXX id 1
limit-resource vlan minimum 16 maximum 4094
limit-resource vrf minimum 2 maximum 4096
limit-resource port-channel minimum 0 maximum 511
limit-resource u4route-mem minimum 248 maximum 248
limit-resource u6route-mem minimum 96 maximum 96
limit-resource m4route-mem minimum 58 maximum 58
limit-resource m6route-mem minimum 8 maximum 8
feature telnet
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock protocol ntp vdc 1
no password strength-check
username admin password 5 $1$MSJwTJtn$Bo0IrVnESUVxLcbRHg86j1  role network-admin
```

```
ip domain-lookup
no service unsupported-transceiver
class-map type qos match-all class-fcoe
policy-map type qos jumbo
class class-default
set qos-group 0
copp profile strict
snmp-server user admin network-admin auth md5 0x71d6a9cf1ea007cd3166e91a6f3807e5
priv 0x71d6a9cf1ea007cd3166e91a6f3807e5 localizedkey
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 10.10.50.2
ntp peer 10.10.50.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8
vlan 1,50-56
vlan 50
name InBand-Mgmt-C1
vlan 51
name Infra-Mgmt-C1
vlan 52
name StorageIP-C1
vlan 53
name vMotion-C1
vlan 54
name VM-Data-C1
vlan 55
name Launcher-C1
```

```
service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
ip route 0.0.0.0/0 10.29.132.1
vpc domain 50
role priority 1000
peer-keepalive destination 10.29.132.5 source 10.29.132.4
interface Vlan1
no shutdown
ip address 10.29.132.2/24
interface Vlan50
no shutdown
ip address 10.10.50.2/24
hsrp version 2
hsrp 50
preempt
priority 110
ip 10.10.50.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface Vlan51
no shutdown
ip address 10.10.51.2/24
hsrp version 2
hsrp 51
preempt
priority 110
ip 10.10.51.1
interface Vlan52
no shutdown
```

```
ip address 10.10.52.2/24
hsrp version 2
hsrp 52
preempt
priority 110
ip 10.10.52.1
interface Vlan53
no shutdown
ip address 10.10.53.2/24
hsrp version 2
hsrp 53
preempt
priority 110
ip 10.10.53.1
interface Vlan54
no shutdown
ip address 10.54.0.2/20
hsrp version 2
hsrp 54
preempt
priority 110
ip 10.54.0.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface port-channel10
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type network
service-policy type qos input jumbo
vpc peer-link
interface port-channel11
```

```
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 11
interface port-channel12
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 12
interface port-channel13
description FI-Uplink-K13
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 13
interface port-channel14
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 14
interface port-channel49
```

```
description FI-Uplink-K23
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
vpc 49
interface port-channel50
description FI-Uplink-K23
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
vpc 50
interface Ethernet1/1
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/2
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/3
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/4
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/5
```

```
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 11 mode active
interface Ethernet1/6
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 11 mode active
interface Ethernet1/7
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 12 mode active
interface Ethernet1/8
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 12 mode active
interface Ethernet1/9
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 13 mode active
interface Ethernet1/10
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 13 mode active
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
```

```
interface Ethernet1/14
interface Ethernet1/15
description HX-Infra01
switchport mode trunk
switchport trunk allowed vlan 1,50-55
interface Ethernet1/16
description HX-Infra02
switchport mode trunk
switchport trunk allowed vlan 1,50-55
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
switchport mode trunk
switchport trunk allowed vlan 1,50-54
spanning-tree port type edge trunk
interface Ethernet1/28
switchport mode trunk
switchport trunk allowed vlan 1,50-54
spanning-tree port type edge trunk
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
interface Ethernet1/33
```

```
interface Ethernet1/34
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
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 49 mode active
interface Ethernet1/50
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 50 mode active
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
vrf member management
ip address 10.29.132.4/24
```

```
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin
```

Switch B Configuration

```
!Command: show running-config
!Time: Fri Nov 13 17:18:36 2018
version 7.0(3)I7(2)
switchname XXXXXXXXXX
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
vdc XXXXXXXXXX id 1
limit-resource vlan minimum 16 maximum 4094
limit-resource vrf minimum 2 maximum 4096
limit-resource port-channel minimum 0 maximum 511
limit-resource u4route-mem minimum 248 maximum 248
limit-resource u6route-mem minimum 96 maximum 96
limit-resource m4route-mem minimum 58 maximum 58
limit-resource m6route-mem minimum 8 maximum 8
feature telnet
cfs eth distribute
feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
```

```
clock protocol ntp vdc 1
no password strength-check
username admin password 5 $1$jEwHqUvM$gpOec2hramkyX09KD3/Dn.  role network-admin
ip domain-lookup
no service unsupported-transceiver
class-map type qos match-all class-fcoe
policy-map type qos jumbo
class class-default
set qos-group 0
copp profile strict
snmp-server user admin network-admin auth md5 0x9046c100ce1f4ecdd74ef2f92c4e83f9
priv 0x9046c100ce1f4ecdd74ef2f92c4e83f9 localizedkey
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 peer 10.10.50.2
ntp server 10.10.50.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8
vlan 1,50-54
vlan 50
name InBand-Mgmt-C1
vlan 51
name Infra-Mgmt-C1
vlan 52
name StorageIP-C1
vlan 53
name vMotion-C1
vlan 54
```

```
name VM-Data-C1
service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
ip route 0.0.0.0/0 10.29.132.1
vpc domain 50
role priority 2000
peer-keepalive destination 10.29.132.4 source 10.29.132.5
interface Vlan1
no shutdown
ip address 10.29.132.3/24
interface Vlan50
no shutdown
ip address 10.10.50.3/24
hsrp version 2
hsrp 50
preempt
priority 110
ip 10.10.50.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface Vlan51
no shutdown
ip address 10.10.51.3/24
hsrp version 2
hsrp 51
preempt
priority 110
ip 10.10.51.1
interface Vlan52
```

```
no shutdown
ip address 10.10.52.3/24
hsrp version 2
hsrp 52
preempt
priority 110
ip 10.10.52.1
interface Vlan53
no shutdown
ip address 10.10.53.3/24
hsrp version 2
hsrp 53
preempt
priority 110
ip 10.10.53.1
interface Vlan54
no shutdown
ip address 10.54.0.3/20
hsrp version 2
hsrp 54
preempt
priority 110
ip 10.54.0.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface Vlan55
no shutdown
ip address 10.10.55.3/24
hsrp version 2
hsrp 55
preempt
priority 110
```

```
ip 10.55.0.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface port-channel10
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type network
service-policy type qos input jumbo
vpc peer-link
interface port-channel11
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 11
interface port-channel12
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 12
interface port-channel13
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
```

```
vpc 13
interface port-channel14
description FI-Uplink-k22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
vpc 14
interface Ethernet1/1
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/2
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/3
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/4
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/5
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 11 mode active
interface Ethernet1/6
switchport mode trunk
switchport trunk allowed vlan 1,50-55
```

```
mtu 9216
channel-group 11 mode active
interface Ethernet1/7
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 12 mode active
interface Ethernet1/8
switchport mode trunk
switchport trunk allowed vlan 1,50-54
mtu 9216
channel-group 12 mode active
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
description HX-Infra01
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/16
description HX-Infra02
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
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
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/26
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/27
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/28
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/29
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/30
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/31
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
```

```
interface Ethernet1/32
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/33
interface Ethernet1/34
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
switchport access vlan 50
interface Ethernet1/49
interface Ethernet1/50
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
vrf member management
ip address 10.29.132.5/24
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```

```
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin
no system default switchport shutdown
```