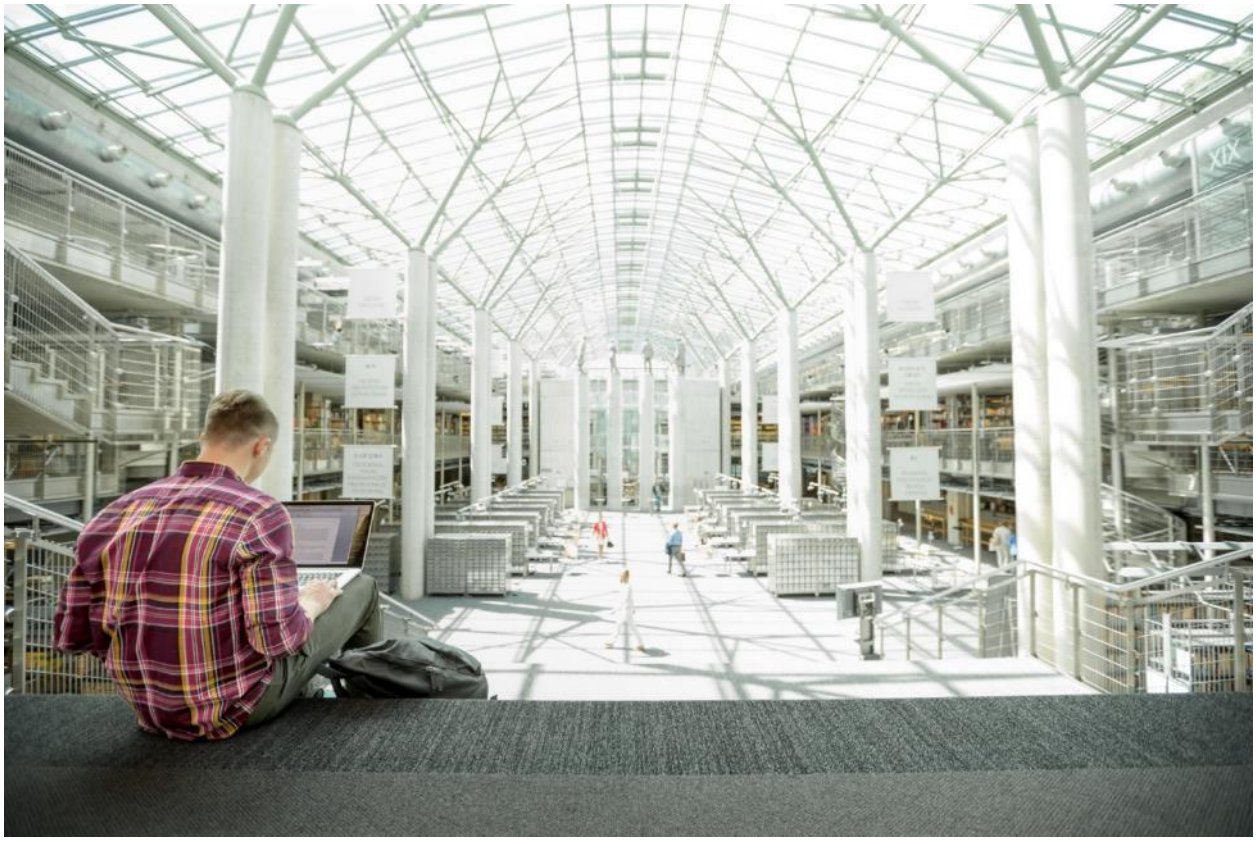

# HyperFlex M5 All-Flash Hyperconverged System with Hyper-V 2019 and Citrix Virtual Apps and Desktops

## Deployment Guide for Cisco HyperFlex with Virtual Desktop Infrastructure for Citrix Virtual Apps and Desktops using Cisco HyperFlex Data Platform v4.0.1b and Microsoft Hyper-V 2019

Published: December 18, 2019

Updated: February 10, 2020

# About the Cisco Validated Design Program

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

http://www.cisco.com/go/designzone.

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

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

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

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

© 2020 Cisco Systems, Inc. All rights reserved.

# Table of Contents


Executive Summary ..... 8
Solution Overview ..... 9
Introduction ..... 9
Audience ..... 9
Purpose of this Document ..... 9
Summary ..... 9
Solution Summary ..... 11
Cisco Desktop Virtualization Solutions: Data Center ..... 12
The Evolving Workplace ..... 12
Cisco Desktop Virtualization Focus ..... 13
Use Cases ..... 15
Physical Topology ..... 15
Fabric Interconnects ..... 16
HX-Series Rack-Mount Servers ..... 17
Cisco UCS B-Series Blade Servers ..... 17
Logical Network Design ..... 18
Technology Overview ..... 21
Cisco Unified Computing System ..... 21
Cisco Unified Computing System Components ..... 21
Supported Versions and System Requirements ..... 22
Hardware and Software Interoperability ..... 22
Software Requirements for Microsoft Hyper-V ..... 22
Cisco UCS Fabric Interconnect ..... 23
Cisco HyperFlex HX-Series Nodes ..... 24
Cisco HyperFlex Converged Data Platform Software ..... 28
Cisco HyperFlex Connect HTML5 Management Web Page ..... 29
Cisco Intersight Management Web Page ..... 29
Cisco Nexus 93108YCPX Switches ..... 37
Architectural Flexibility ..... 37
Feature-Rich ..... 37
Real-Time Visibility and Telemetry ..... 38
Highly Available and Efficient Design ..... 38
Simplified Operations ..... 38
Investment Protection ..... 38
Microsoft Hyper-V 2019 ..... 39
Microsoft System Center 2019 ..... 39
Citrix Virtual Apps and Desktops™ 1909 ..... 39
Zones ..... 41

Improved Database Flow and Configuration ..... 41
Application Limits ..... 42
Multiple Notifications before Machine Updates or Scheduled Restarts ..... 42
API Support for Managing Session Roaming ..... 42
API Support for Provisioning Virtual Machines from Hypervisor Templates ..... 42
Support for New and Additional Platforms ..... 42
Citrix Provisioning Services 1909 ..... 43
Benefits for Citrix Virtual Apps and Other Server Farm Administrators ..... 44
Benefits for Desktop Administrators ..... 44
Citrix Provisioning Services Solution ..... 44
Citrix Provisioning Services Infrastructure ..... 45
Architecture and Design of Citrix Virtual Desktops on Cisco Unified Computing System and Cisco HyperFlex Storage Design Fundamentals ..... 45
Understanding Applications and Data ..... 47
Project Planning and Solution Sizing Sample Questions ..... 47
Citrix Virtual Desktops Design Fundamentals ..... 48
Machine Catalogs ..... 48
Delivery Groups ..... 48
Example Virtual Desktops Deployments ..... 49
Distributed Components Configuration ..... 49
Multiple Site Configuration ..... 50
Citrix Cloud Services ..... 51
Designing a Virtual Desktops Environment for a Mixed Workload ..... 51
Deployment Hardware and Software ..... 53
Physical Components ..... 53
Software Components ..... 54
Licensing ..... 55
Considerations ..... 55
Version Control ..... 55
Microsoft Windows Active Directory ..... 56
Scale ..... 56
Capacity ..... 56
Physical Topology ..... 58
Topology Overview ..... 58
Fabric Interconnects ..... 59
HX-Series Rack-Mount Servers ..... 60
Logical Architecture ..... 61
Design Elements ..... 61
Network Design ..... 62
VLANs ..... 63

Jumbo Frames....63
Cisco UCS Design....64
Cisco UCS Organization....64
Cisco UCS LAN Policies....64
Cisco UCS Servers Policies....72
Cisco UCS Service Profile Templates....77
Solution Configuration....82
Installation....83
Prerequisites....83
IP Addressing....83
DHCP versus Static IP....87
Configure the Active Directory for Constrained Delegation....87
Prepopulate AD DNS with Records....88
NTP....90
VLANs....90
Network Uplinks....91
Usernames and Passwords....92
Physical Installation....93
Solution Cabling....93
Cisco Unified Computing System Configuration....96
Cisco UCS Fabric Interconnect A....96
Cisco UCS Fabric Interconnect B....97
Cisco UCS Manager....98
Cisco UCS Firmware....99
NTP....99
Uplink Ports....99
Uplink Port Channels....100
Server Ports....101
Server Discovery....103
Deploying HX Data Platform Installer on Hyper-V Infrastructure....103
Assign a Static IP Address to the HX Data Platform Installer Virtual Machine....108
HyperFlex Installer Web Page....109
HyperFlex Installation....111
Post Installation Tasks....129
Create Datastores....129
Constrained Delegation (Optional)....130
Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces....133
Rename the Cluster Network in Windows Failover Cluster - Optional....136
Configure the Windows Failover Cluster Network Roles....136
Configure the Windows Failover Cluster Network for Live Migration....137

Create Folders on the HX Datastore ... 138
Configure the Default Folder to Store Virtual Machine Files on Hyper-V ... 139
Validate the Windows Failover Cluster Configuration ... 139
Configure Quorum in Windows Server Failover Cluster ... 140
Initial Tasks and Testing ... 141
Ready Clones ... 141
Auto-Support and Notifications ... 143
Smart Licensing ... 145
HyperFlex Cluster Expansion ... 147
Expansion with Converged Nodes ... 147
Expansion with Compute-Only Nodes ... 159
Management ... 184
HyperFlex Connect ... 184
Dashboard ... 185
Monitor ... 185
Analyze ... 186
Manage ... 187
Configure Remote Management Stations for Managing VMs in HX Cluster ... 187
Microsoft Hyper-V Manager ... 189
Change the Default Location to Store the Virtual Machine Files using Hyper-V Manager ... 190
Create Virtual Machines using Hyper-V Manager ... 192
Windows Failover Cluster Manager ... 194
Post HyperFlex Cluster Installation for Hyper-V 2019 ... 199
Microsoft System Center Virtual Machine Manager 2019 ... 199
Create Run-As Account for Managing the Hyper-V Cluster ... 199
Manage Servers and Clusters ... 199
Networking ... 202
Storage ... 205
Build the Virtual Machines and Environment for Workload Testing ... 207
Software Infrastructure Configuration ... 207
Prepare the Master Images ... 208
Install and Configure Citrix Desktop Delivery Controller, Citrix Licensing, and StoreFront ... 209
Install Citrix License Server ... 209
Install Citrix Licenses ... 212
Install Citrix Desktop Broker/Studio ... 214
Configure the Citrix VDI Site ... 219
Configure the Citrix VDI Site Administrators ... 227
Configure Additional Desktop Controller ... 228
Add the Second Delivery Controller to the Citrix Desktop Site ... 232
Install and Configure StoreFront ... 233

Additional StoreFront Configuration ... 242
Install the Citrix Provisioning Services Target Device Software ... 247
Create Citrix Provisioning Services vDisks ... 249
Provision Virtual Desktop Machines ... 258
Provision Desktop Machines from Citrix Provisioning Services Console ... 263
Install Citrix Virtual Apps and Desktop Virtual Desktop Agents ... 274
Create Delivery Groups ... 280
Citrix Virtual Desktops Policies and Profile Management ... 284
Configure Citrix Virtual Desktops Policies ... 284
Configure User Profile Management ... 285
Test Setup and Configurations ... 287
Test Methodology and Success Criteria ... 288
Test Procedure ... 289
Pre-Test Setup for Testing ... 289
Test Run Protocol ... 289
Success Criteria ... 290
Test Results ... 295
Boot Storms ... 295
Recommended Maximum Workload and Configuration Guidelines ... 295
Sixteen Node Cisco HXAF220c-M5S Rack Server, B200 M5 Compute Nodes & HyperFlex All-Flash Cluster ... 295
Sixteen Node Cisco HXAF220c-M5S Rack-Mount Server, Cisco UCS B200 M5 Compute Nodes and HyperFlex All-Flash Cluster ... 296
Hyper-V Performance Counters ... 298
Hyper-V Performance Counters ... 303
Summary ... 308
About the Author ... 309
Acknowledgements ... 309
Appendix A – Cisco Nexus 93108YC Switch Configuration ... 310
Switch A Configuration ... 310
Switch B Configuration ... 321
Appendix B - Install Microsoft Windows Server 2016 ... 332
Install Microsoft Windows Server 2016 ... 332
Configure Cisco UCS Manager using HyperFlex Installer ... 332
Configure Cisco UCS vMedia and Boot Policies ... 332
Install Microsoft Windows Server 2019 OS ... 340
Undo vMedia and Boot Policy Changes ... 343

# 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 efficiently power your applications and your business

This document provides an architectural reference and design guide for up to 5000 RDS and 4000 VDI session workload on a 32-node (16x Cisco HyperFlex HXAF220C-M5SX server and 16x Cisco B200 M5 Compute only nodes) Cisco HyperFlex system. We provide deployment guidance and performance data for Citrix Virtual Desktops 1909 virtual desktops running Microsoft Windows 10 with Office 2016. The solution is a pre-integrated, best-practice data center architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches and Cisco HyperFlex Data Platform software version 4.0.1b.

The solution payload is 100 percent virtualized on Cisco HyperFlex HXAF220C-M5SX hyperconverged nodes and Cisco UCS B200 M5 Compute-Only Nodes booting through on-board M.2 SATA SSD drive running Microsoft Hyper-V 2019 hypervisor and the Cisco HyperFlex Data Platform storage controller virtual machine. The virtual desktops are configured with Virtual Desktops 1909, which incorporates both traditional persistent and non-persistent virtual Windows 7/8/10 desktops, hosted applications and remote desktop service (RDS) Microsoft Server 2008 R2, Server 2012 R2, Server 2016 or Server 2019 based desktops. The solution provides unparalleled scale and management simplicity. Citrix Virtual Desktops Provisioning Services or Machine Creation Services Windows 10 desktops, full clone desktops or Virtual Apps server-based desktops can be provisioned on an eight node Cisco HyperFlex cluster. Where applicable, this document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

The solution boots 1550 virtual desktops machines in 30 minutes or less, making sure that users will not experience delays in accessing their virtual workspace on HyperFlex.

Our past Cisco Validated Design studies with HyperFlex show linear scalability out to the cluster size limits of 8 HyperFlex hyperconverged nodes with Cisco UCS B200 M5, Cisco UCS C220 M5, or Cisco UCS C240 M5 servers. You can expect that our new HyperFlex all flash system running HX Data Platform 3.0.1 on Cisco HXAF220 M5 or Cisco HXAF240 M5 nodes will scale up to 2500 knowledge worker users per cluster with N+1 server fault tolerance.

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. 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 the fifth generation servers with NVIDIA GPUs and software for details on how to integrate this capability with Citrix Virtual Desktops.

The solution provides outstanding virtual desktop end user experience as measured by the Login VSI 4.1.32 Knowledge Worker workload running in benchmark mode. Index average end-user response times for all tested delivery methods is under one second, representing the best performance in the industry.

# Solution Overview

## Introduction

The 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 HyperFlex hyperconverged 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 audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

## Purpose of this Document

This document provides a step-by-step design, configuration, and implementation guide for the Cisco Validated Design for a Cisco HyperFlex All-Flash system running two different Citrix Virtual Desktops/Virtual Apps workloads with Cisco UCS 6300 series Fabric Interconnects and Cisco Nexus 9300 series switches.

## Summary

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

**Figure 1 Cisco HyperFlex System Overview**

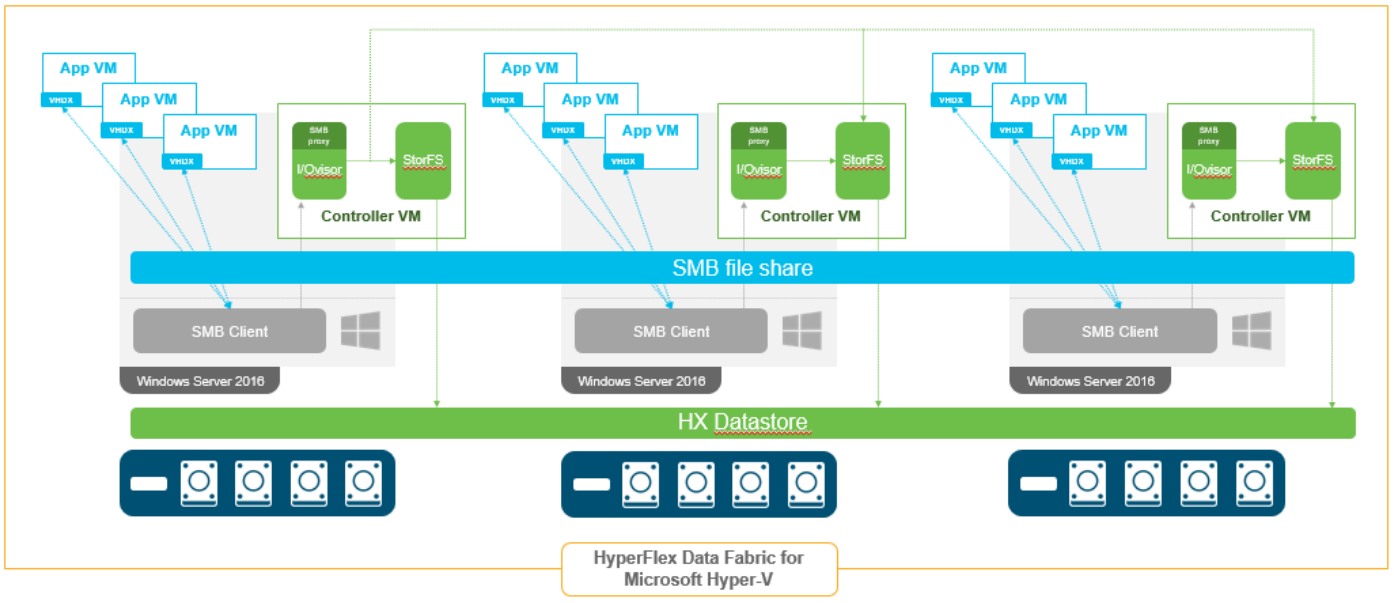

The following are the components of a Cisco HyperFlex system using Microsoft Hyper-V as the hypervisor:

- One pair of Cisco UCS Fabric Interconnects, choose from models:
  - Cisco UCS 6332 Fabric Interconnect
  - Cisco UCS 6332-16UP Fabric Interconnect
- Eight Cisco HyperFlex HX-Series Rack-Mount Servers, choose from the following models:
  - Cisco HyperFlex HXAF220c-M5SX All-Flash Rack-Mount Servers
  - Eight Cisco UCS B200 M5 Servers for Compute-Only nodes:
    - Cisco HyperFlex Data Platform Software 4.0.1b
    - Microsoft Windows Server 2019 Hyper-V Hypervisor
    - Microsoft Windows Active Directory and DNS services, RSAT tools (end-user supplied)
    - SCVMM – Needed for Virtual Desktops (end-user supplied)
    - Citrix Virtual Apps & Desktops 1909
    - Citrix Provisioning Services 1909

# Solution Summary

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both Citrix Virtual Desktops Microsoft Windows 10 virtual desktops and Citrix Virtual Apps server desktop sessions based on Microsoft Server 2019. The mixed workload solution includes Cisco HyperFlex hardware and Data Platform software, Cisco Nexus® switches, the Cisco Unified Computing System (Cisco UCS®), Citrix Virtual Desktops and Microsoft Hyper-V software in a single package. The design is efficient such that the networking, computing, and storage components occupy a 12-rack unit 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 solution 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. and Microsoft 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 20-core 2.3 GHz Intel Xeon (Gold 6230) Scalable Family processors with 768GB of 2666Mhz memory with Citrix Virtual Desktops support more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel Xeon Gold 6230 20-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 4000 user hosted virtual desktop and 5000 user hosted shared desktop environment booted and registered with the Virtual Desktops Studio in under 5 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.
- Stress-tested to the limits during simulated login storms. All 5000 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 data center. The rack space required to support the initial 5000 user system is 32 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects. Incremental Citrix Virtual Desktops users can be added to the Cisco HyperFlex cluster up to the cluster scale limits, currently 16 hyper converged and 16 compute only nodes, by adding one or more nodes.
- 100 percent virtualized: This CVD presents a validated design that is 100 percent virtualized on Microsoft Hyper-V 2019. All the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, Citrix Virtual Desktops components, Virtual Desktops VDI desktops

and Virtual Apps 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(1c) 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 v4.0.1c 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.
- Optimized for performance and scale. For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the 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.

# Cisco Desktop Virtualization Solutions: Data Center

## The Evolving Workplace

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

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

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

**Figure 2 Cisco Data Center Partner Collaboration**

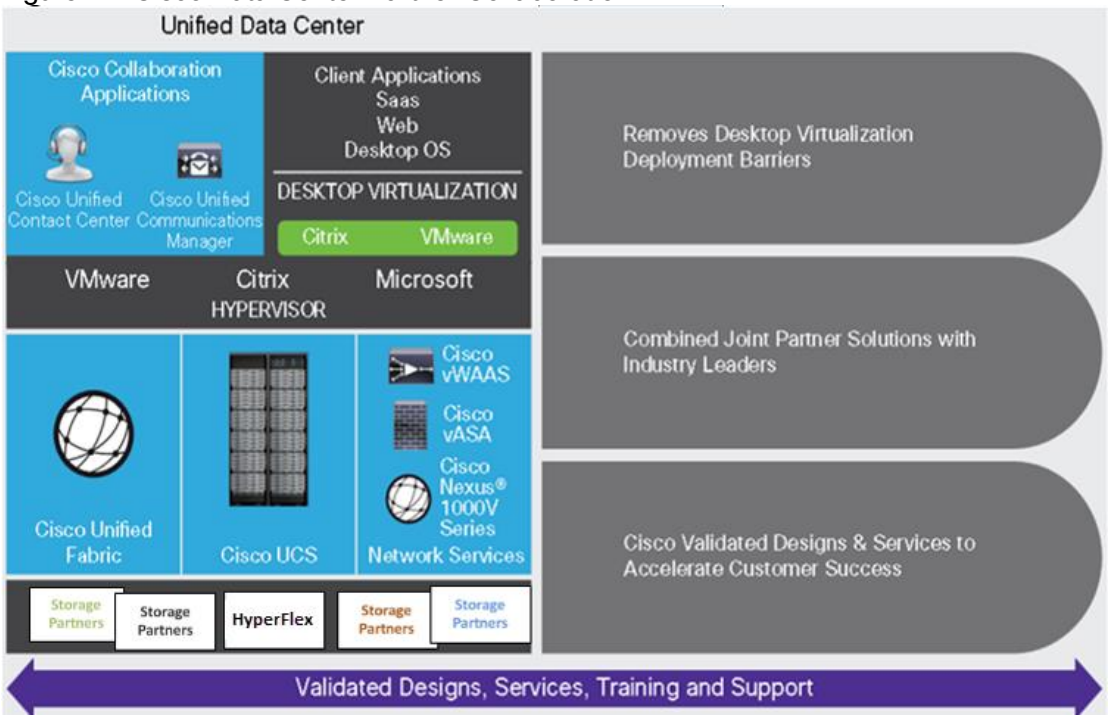


Some of the key drivers for desktop virtualization are increased data security, the ability to expand and contract capacity and reduced TCO through increased control and reduced management costs.

## Cisco Desktop Virtualization Focus

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

### Simplified

Cisco UCS and 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 and 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 Microsoft have developed integrated, validated architectures, including predefined hyper-converged architecture infrastructure packages such as HyperFlex. Cisco Desktop Virtualization Solutions have been tested with Microsoft Hyper-V.

## 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 Microsoft Live Migration, 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 1.5 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 Citrix Virtual Desktops, Cisco HyperFlex solutions have demonstrated scalability and performance, with up to 5000 hosted virtual desktops and hosted shared desktops up and running in 5 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 using 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

# Physical Topology

The Cisco HyperFlex system is composed of a pair of Cisco UCS 6200/6300 series Fabric Interconnects, along with up to 16 HXAF-Series rack mount servers per cluster Up to 8 separate HX clusters can be installed under a single pair of Fabric Interconnects. 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 data center network at the time of installation.

For this study, we uplinked the Cisco 6332-16UP Fabric Interconnects to Cisco Nexus 93108YCPX switches.

Figure 3 Cisco HyperFlex Standard Topology

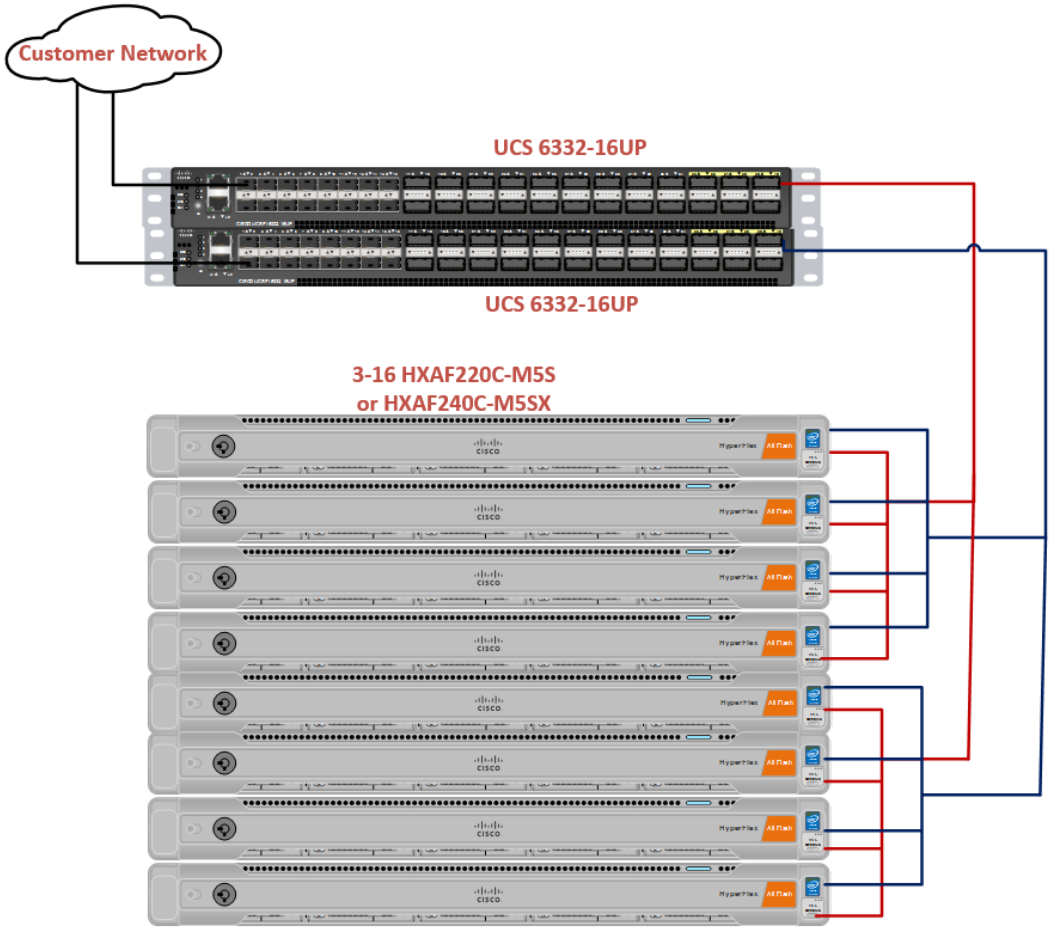


## Fabric Interconnects

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 also 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.

## HX-Series Rack-Mount Servers

The HX-Series converged 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 using a single cable for both management traffic and data traffic. Both the HXAF220C-M5SX and HXAF240C-M5SX 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 4).

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

Figure 4 HX-Series Server Connectivity

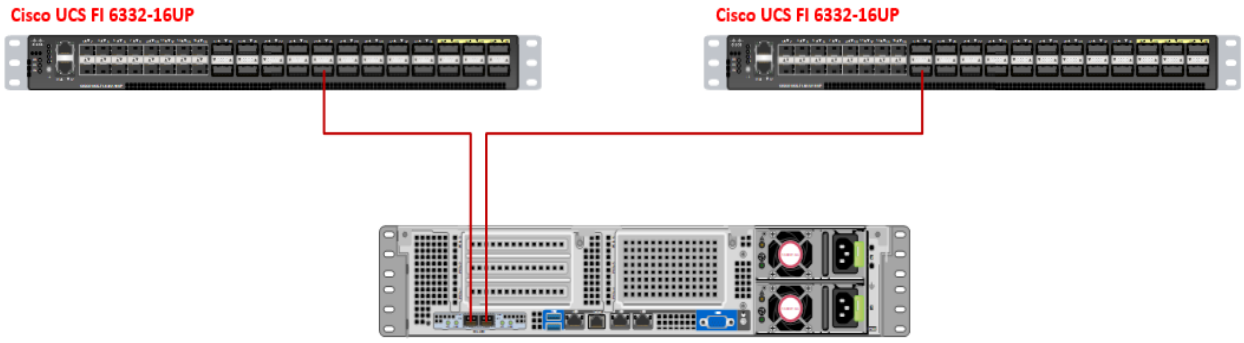


## Cisco UCS B-Series Blade Servers

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 5). All other cabling configurations are invalid, and can lead to errors, discovery failures, and loss of redundant connectivity.

Figure 5 Cisco UCS 5108 Chassis Connectivity

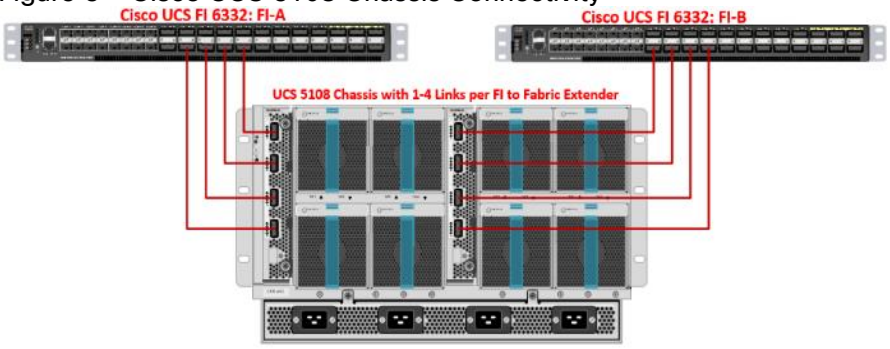

## Logical Network Design

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.
  - Hyper-V 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 via 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, Hyper-V 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 vmnic interface used for storage traffic for each Hyper-V host in the HX cluster.
  - Storage Controller virtual machine storage interfaces.
  - A roaming HX cluster storage interface.
- Live Migration Zone: This zone comprises the connections used by the Hyper-V hosts to enable LiveMigration 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 6 illustrates the logical network design.

Figure 6 Logical Network Design

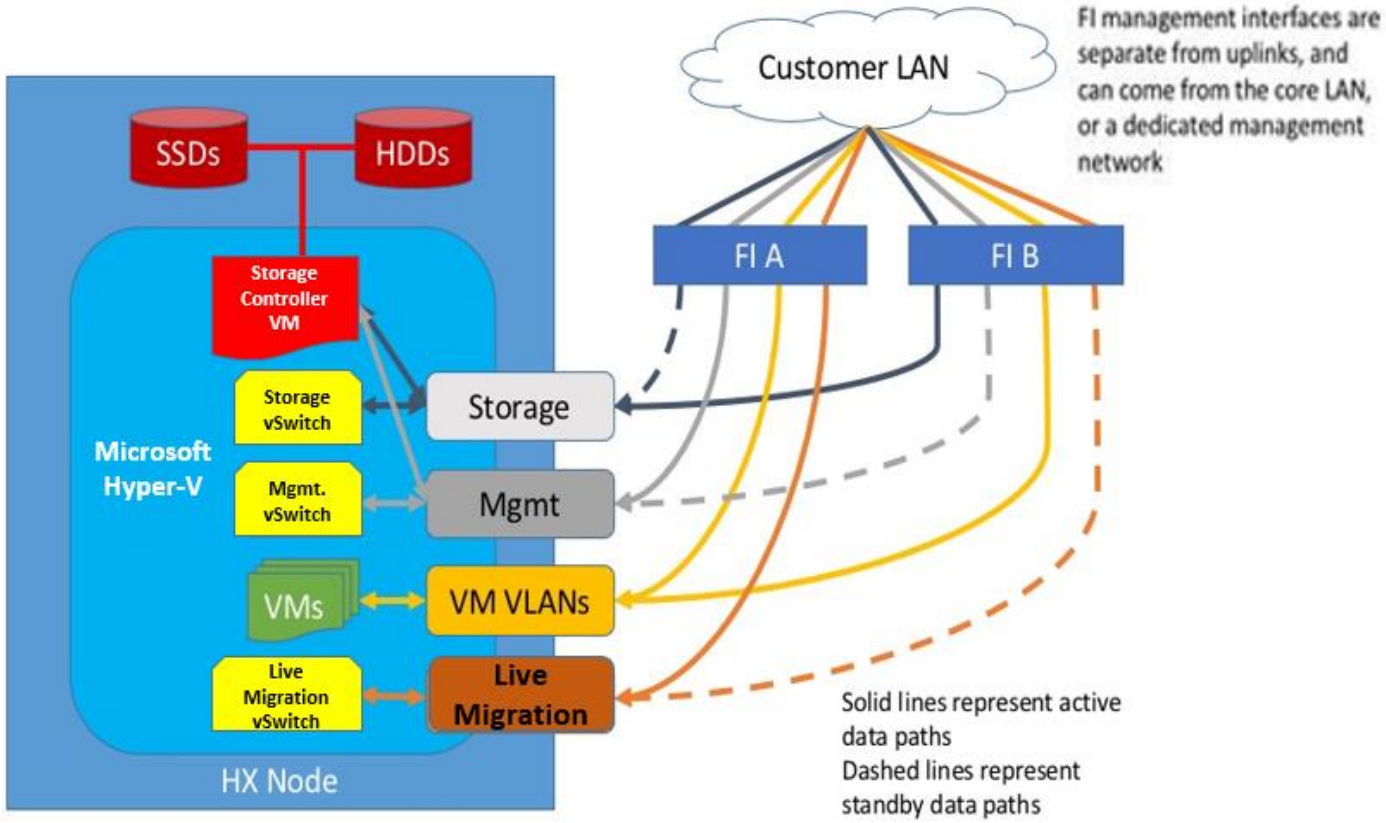


The reference hardware configuration includes:

- Two Cisco Nexus 93108YCPX switches
- Two Cisco UCS 6332-16UP fabric interconnects
- Eight Cisco HX-series Rack server running HyperFlex data platform version 4.0.1b.
- Eight Cisco B200 M5 Blade Server running HyperFlex data platform version 4.0.1b

For desktop virtualization, the deployment includes Citrix Virtual Desktops running on Microsoft Hyper-V. The design is intended to provide a large-scale building block for persistent/non-persistent desktops with following density per four-node configuration:

- 4000 Citrix Virtual Desktops Windows 10 non-persistent virtual desktops using PVS
- 5000 Citrix Virtual Apps Windows 2019 Server Desktops using PVS

**All the Windows 10 virtual desktops were provisioned with 4GB of memory for this study. Typically, persistent desktop users may desire more memory. If more than 4GB memory is needed, the second memory channel on the Cisco HXAF220c-M5SX HX-Series Rack Server should be populated.**

Data provided here will allow customers to run VDI desktops to suit their environment. For example, 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 3. These procedures explain everything from physical cabling to network, compute and storage device configurations.

# Technology Overview

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 the Cisco Unified Computing System™ (Cisco UCS) and Cisco HyperFlex through an intuitive GUI, a command-line interface (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 the nodes.
- 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. Our latest advancement offers a cloud-based management system called Cisco Intersight.

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 7 Cisco HyperFlex System Overview**

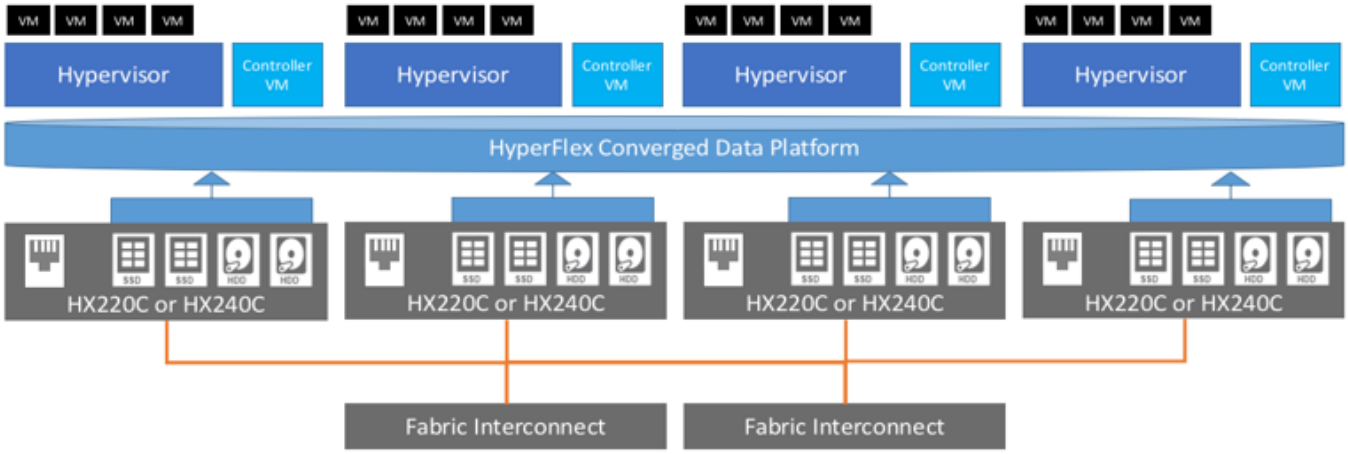


## Supported Versions and System Requirements

Cisco HX Data Platform requires specific software and hardware versions, and networking settings for successful installation.

For a complete list of requirements, see: Cisco HyperFlex Systems Installation Guide on Microsoft Hyper-V

### Hardware and Software Interoperability

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

### Software Requirements for Microsoft Hyper-V

The software requirements include verification that you are using compatible versions of Cisco HyperFlex Systems (HX) components and Microsoft Server components.

### HyperFlex Software Versions

The HX components—Cisco HX Data Platform Installer, Cisco HX Data Platform, and Cisco UCS firmware are installed on different servers. Verify that each component on each server used with and within an HX Storage Cluster are compatible.

- Verify that the preconfigured HX servers have the same version of Cisco UCS server firmware installed. If the Cisco UCS Fabric Interconnects (FI) firmware versions are different, see the Cisco HyperFlex Systems Upgrade Guide for steps to align the firmware versions.
- M5: For NEW hybrid or All Flash (Cisco HyperFlex HX240c M5 or HX220c M5) deployments, verify that Cisco UCS Manager 3.2(3b) or later is installed.

## 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 8 Cisco UCS 6332 Fabric Interconnect**

Figure 9 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 (Cisco UCS) 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 (Cisco ACI™). 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 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 Cisco UCS 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 eight nodes can be configured in one HX cluster. For detailed information, see the Cisco HyperFlex HXAF220c-M5S Spec Sheet.

Figure 10 Cisco UCS HXAF220c-M5SX Rack Server Front View

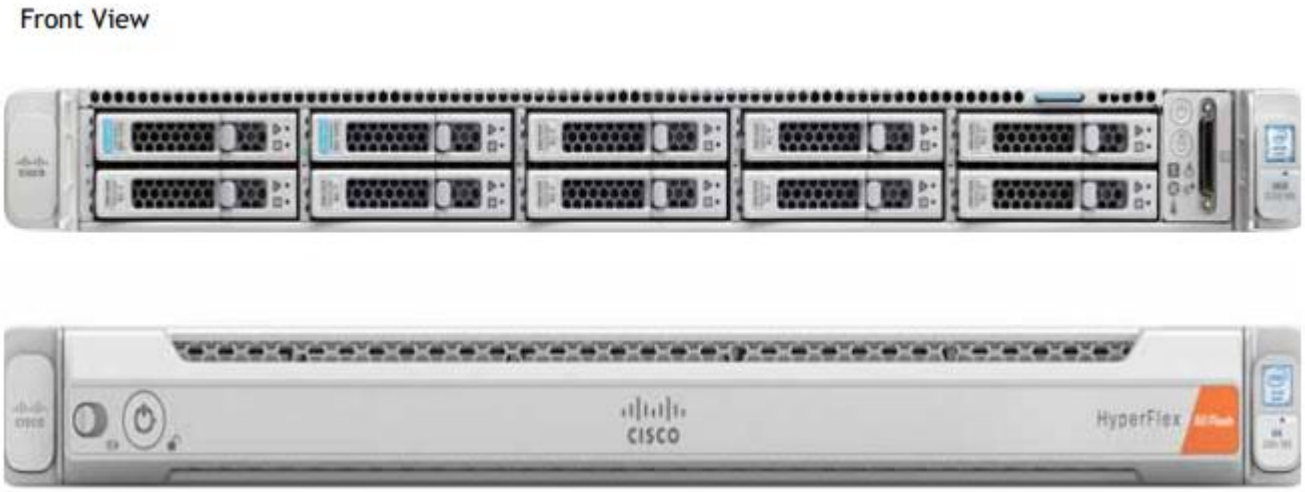

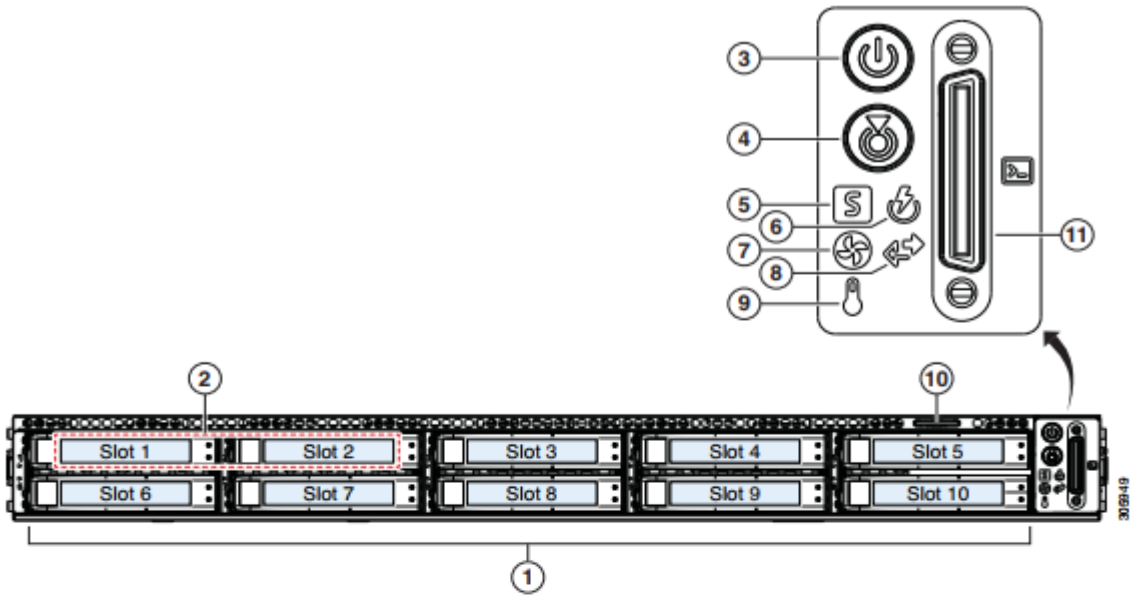


| | | | |
|---|---|---|---|
| **1** | Drive Slots<br>Slot 01 (For System/Log drive)<br>• 1 x SATA SSD<br>Slot 02 (For Cache drive)<br>• 1 x NVMe SSD **OR**<br>• 1 x SAS SSD **OR**<br>• 1 x SED SAS SSD<br>Slot 03 through 10 (For Capacity drives)<br>• Upto 8 x SATA SSD **OR**<br>• Upto 8 x SED SATA SSD **OR**<br>• upto 8 x SED SAS SSD | **7** | Fan status LED |
| **2** | N/A | **8** | Network link activity LED |
| **3** | Power button/Power status LED | **9** | Temperature status LED |
| **4** | Unit identification button/LED | **10** | Pull-out asset tag |
| **5** | System status LED | **11** | KVM connector (used with KVM cable that provides two USB 2.0, one VGA, and one serial connector) |
| **6** | Power supply status LED | **—** | — |

Figure 11 Cisco UCS HXAF220c-M5SX Rack Server Rear View

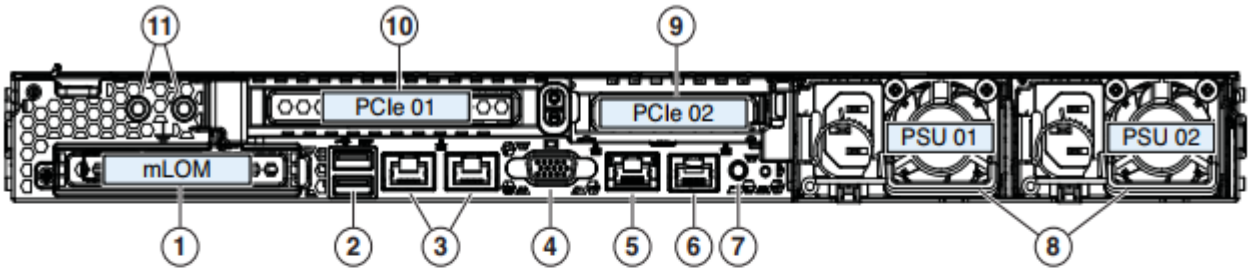


| | | | |
|---|---|---|---|
| 1 | Modular LAN-on-motherboard (mLOM) card bay (x16) | 7 | Rear unit identification button/LED |
| 2 | USB 3.0 ports (two) | 8 | Power supplies (two, redundant as 1+1) |
| 3 | Dual 1/10-Gb Ethernet ports (LAN1 and LAN2). LAN1 is left connector and LAN2 is right connector | 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
- 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 12). 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 12 Cisco VIC 1387 mLOM Card**

## Cisco HyperFlex Converged 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 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.

## 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>.

## Cisco Intersight Management Web Page

Cisco Intersight (https://intersight.com), previously known as Starship, is the latest visionary cloud-based management tool, designed to provide a centralized off-site management, monitoring and reporting tool for all your Cisco UCS based solutions. In the initial release of Cisco Intersight, monitoring and reporting is enabled against Cisco HyperFlex clusters. The Cisco Intersight website and framework can be upgraded with new and enhanced features independently of the products that are managed, meaning that many new features and capabilities can come with no downtime or upgrades required by the end users. Future releases of Cisco HyperFlex will enable further functionality along with these upgrades to the Cisco Intersight framework. This unique combination of embedded and online technologies will result in a complete cloud-based management solution that can care for Cisco HyperFlex throughout the entire lifecycle, from deployment through retirement.

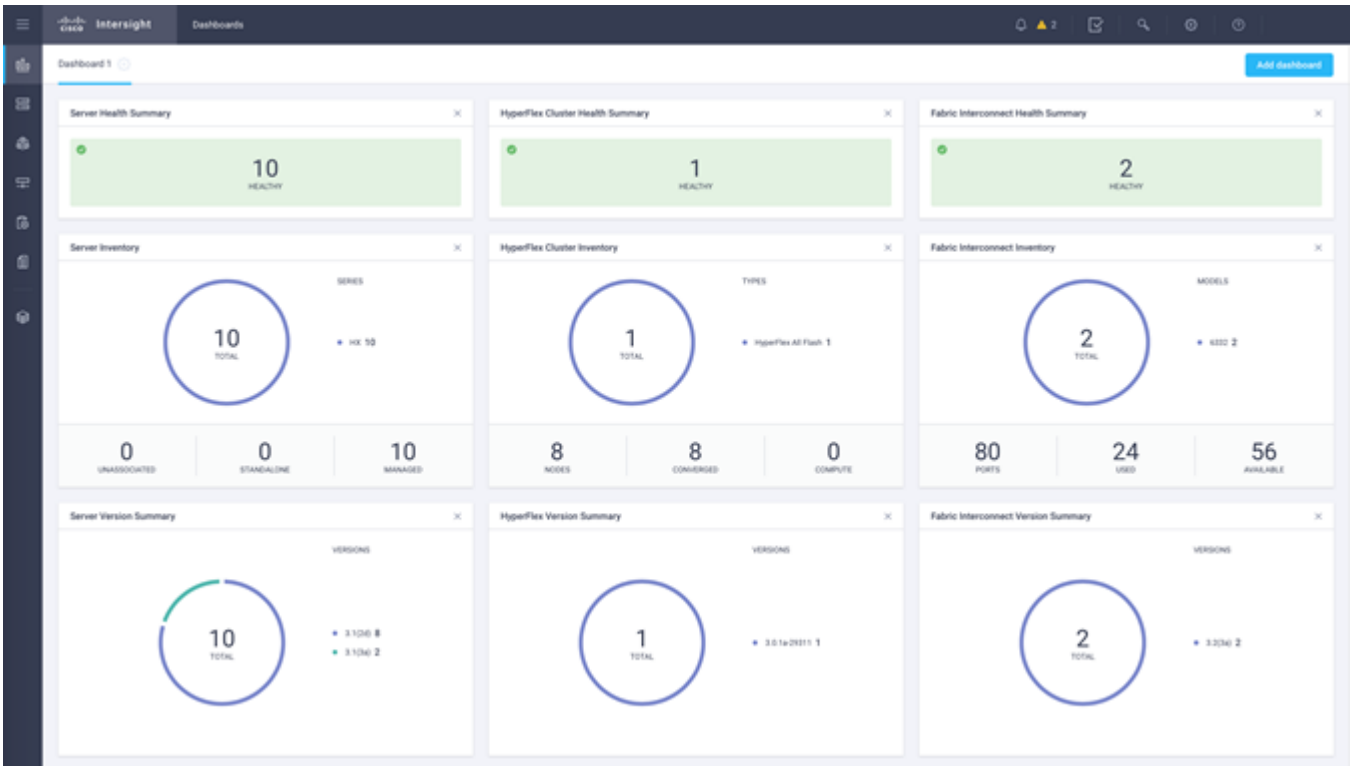

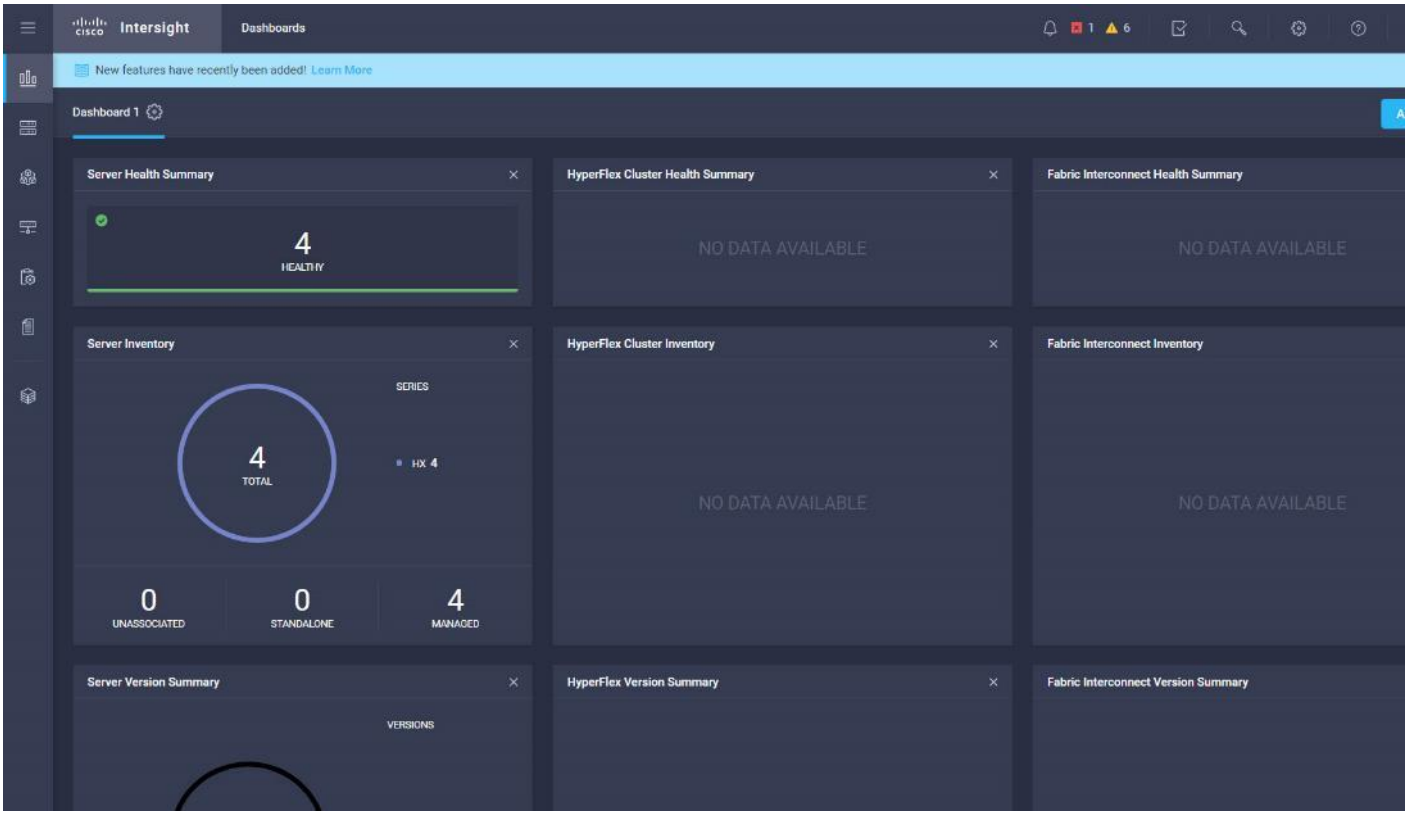
Intersight
Dashboards
New features have recently been added! Learn More
Dashboard 1
Server Health Summary
4
HEALTHY
HyperFlex Cluster Health Summary
NO DATA AVAILABLE
Fabric Interconnect Health Summary
NO DATA AVAILABLE
Server Inventory
SERIES
4
TOTAL
HX 4
0
UNASSOCIATED
0
STANDALONE
4
MANAGED
HyperFlex Cluster Inventory
NO DATA AVAILABLE
Fabric Interconnect Inventory
NO DATA AVAILABLE
Server Version Summary
VERSIONS
HyperFlex Version Summary
Fabric Interconnect Version Summary

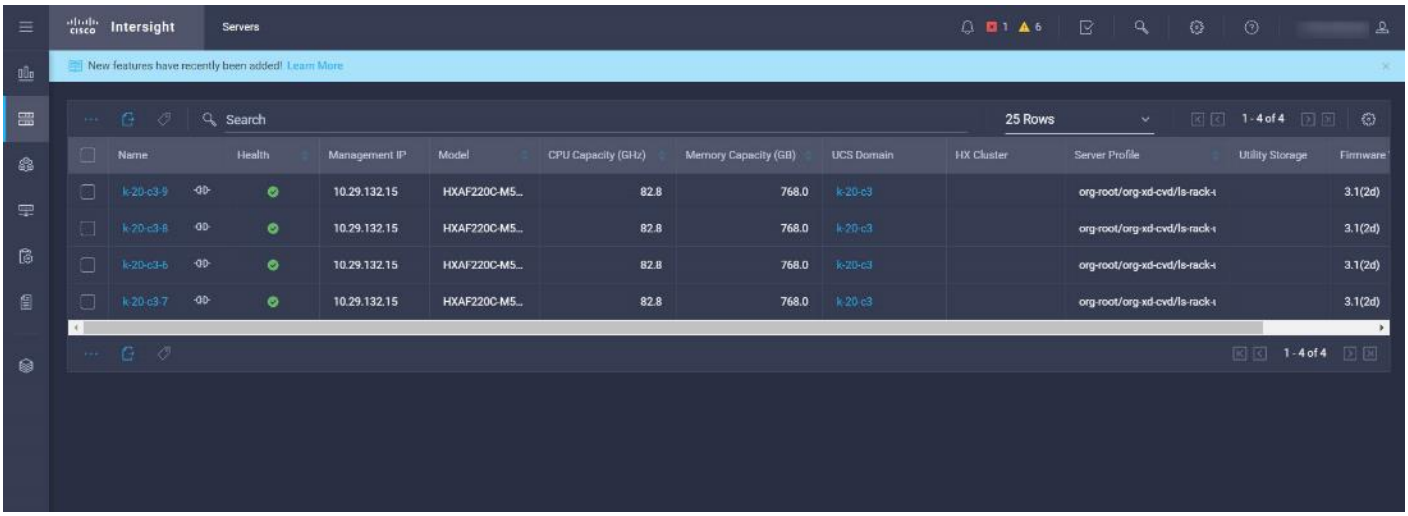
Intersight
Servers
New features have recently been added! Learn More
Search
25 Rows
1 - 4 of 4
Name
Health
Management IP
Model
CPU Capacity (GHz)
Memory Capacity (GB)
UCS Domain
HX Cluster
Server Profile
Utility Storage
Firmware
10.29.132.15
HXAF220C-M5...
82.8
768.0
3.1(2d)

Figure 13 HyperFlex Connect GUI

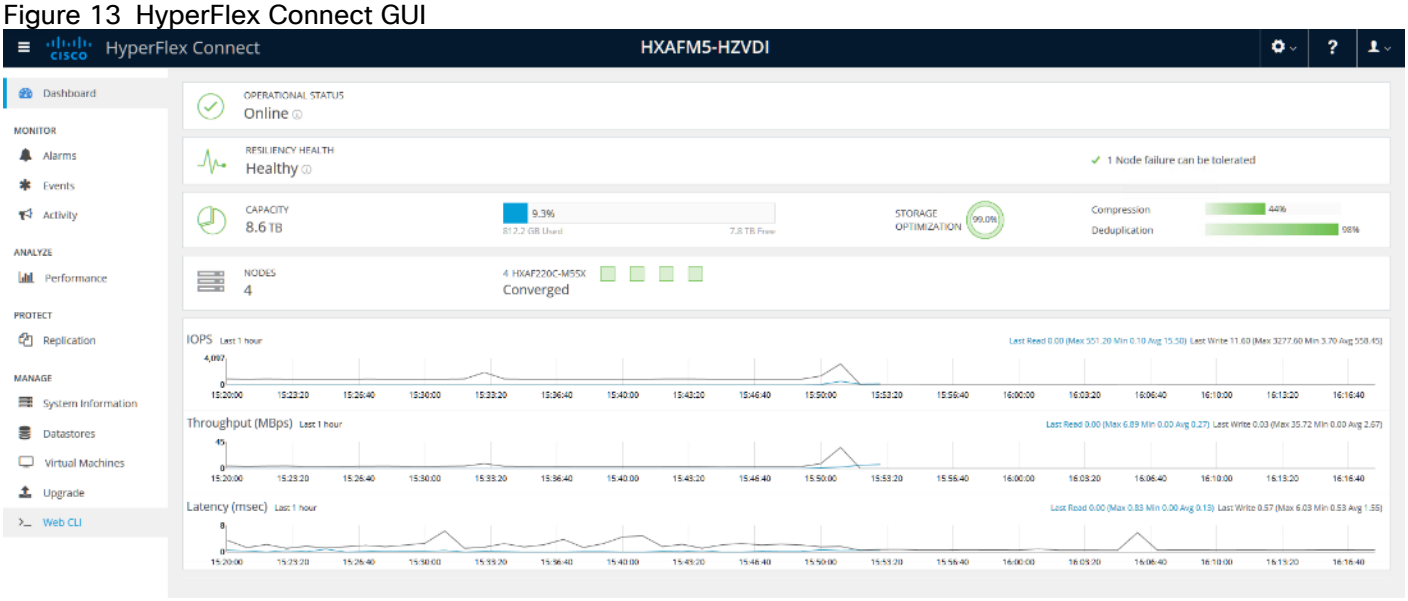


## Replication Factor

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. 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.

When moving a virtual machine to a new location using tools such as Hyper-V Cluster Optimization, 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 14).

Figure 14 Data Write Operation Flow Through the Cisco HyperFlex HX Data Platform

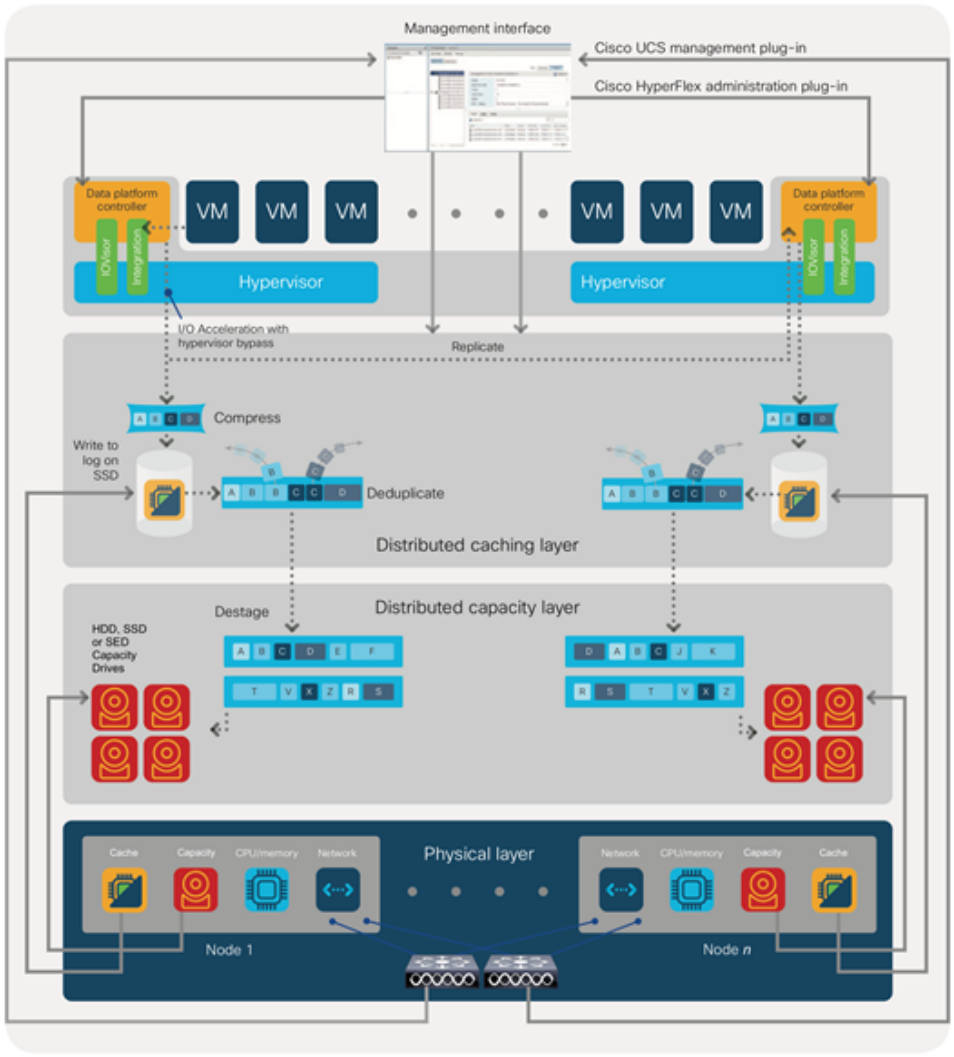


Hot data sets, data that are frequently or recently read from the capacity tier, are cached in memory. All-Flash configurations, however, does not use an SSD read cache since there is no performance benefit of such a cache; the persistent data copy already resides on high-performance SSDs. In these configurations, 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 15).

**Figure 15 Cisco HyperFlex HX Data Platform Optimizes Data Storage with No Performance Impact**

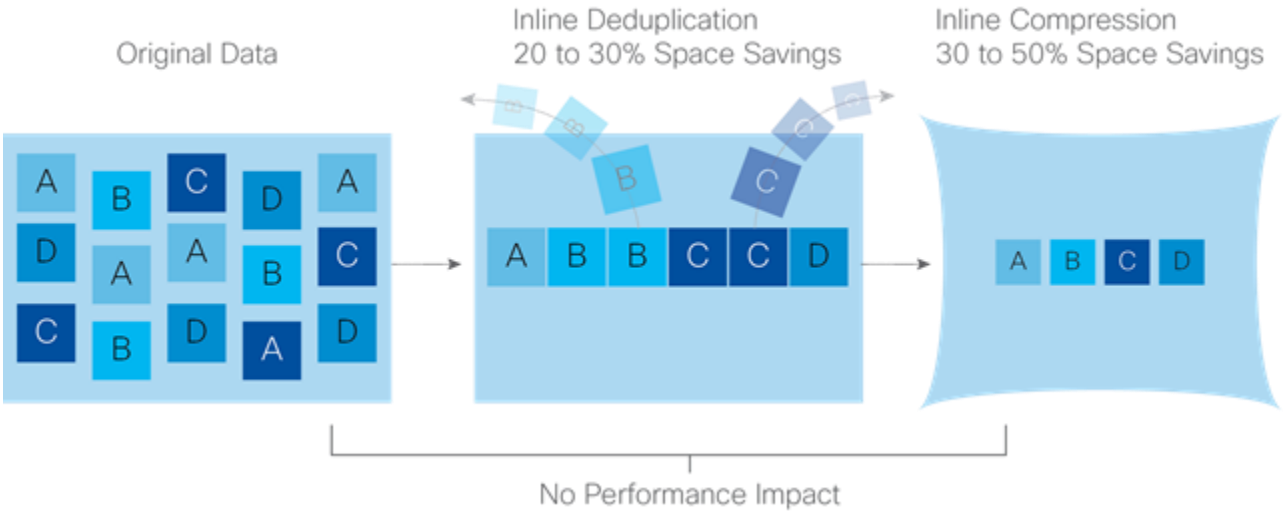


## Inline Compression

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 16). 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 16 Cisco HyperFlex HX Data Platform Optimizes Data Storage with No Performance Impact

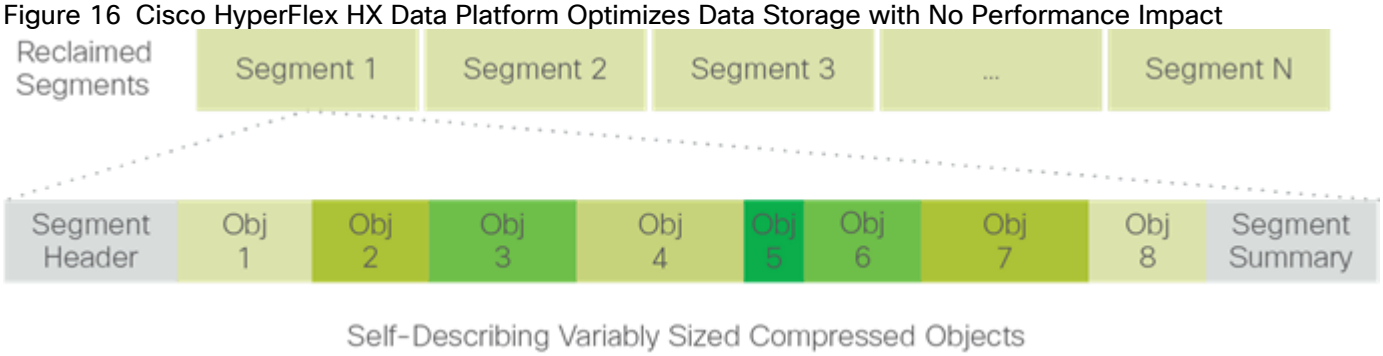


## Encryption

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 10GbE which 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 virtual machine 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 93108YCPX Switches

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

### Architectural Flexibility

- Includes top-of-rack, fabric extender aggregation, or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
- Includes leaf node support for Cisco ACI architecture
- Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support

### Feature-Rich

- Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
- ACI-ready infrastructure helps users take advantage of automated policy-based systems management
- Virtual extensible LAN (VXLAN) routing provides network services
- Rich traffic flow telemetry with line-rate data collection
- Real-time buffer utilization per port and per queue, for monitoring traffic micro-bursts and application traffic patterns

## Real-Time Visibility and Telemetry

- Cisco Tetration Analytics Platform support with built-in hardware sensors for rich traffic flow telemetry and line-rate data collection
- Cisco Nexus Data Broker support for network traffic monitoring and analysis
- Real-time buffer utilization per port and per queue, for monitor traffic micro-bursts and application traffic patterns

## Highly Available and Efficient Design

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

## Simplified Operations

- Pre-boot execution environment (PXE) and Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
- Automate and configure switches with DevOps tools like Puppet, Chef, and Ansible
- An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
- Python scripting gives programmatic access to the switch command-line interface (CLI)
- Includes hot and cold patching, and online diagnostics

## Investment Protection

- A Cisco 40-Gb bidirectional transceiver allows for reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet
- Support for 10-Gb and 25-Gb access connectivity and 40-Gb and 100-Gb uplinks facilitate data centers migrating switching infrastructure to faster speeds
- 1.44 Tbps of bandwidth in a 1 RU form factor
- 48 fixed 1/10-Gbps SFP+ ports
- 6 fixed 40-Gbps QSFP+ for uplink connectivity that can be turned into 10 Gb ports through a QSFP to SFP or SFP+ Adapter (QSA)
- Latency of 1 to 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies
- Hot swappable 2+1 redundant fan tray

Figure 17 Cisco Nexus 93108YC Switch

## Microsoft Hyper-V 2019

Hyper-V is Microsoft's hardware virtualization product. It lets you create and run a software version of a computer, called a virtual machine. Each virtual machine acts like a complete computer, running an operating system and programs. When you need computing resources, virtual machines give you more flexibility, help save time and money, and are a more efficient way to use hardware than just running one operating system on physical hardware.

Hyper-V runs each virtual machine in its own isolated space, which means you can run more than one virtual machine on the same hardware at the same time. You might want to do this to avoid problems such as a crash affecting the other workloads, or to give different people, groups or services access to different systems.

## Microsoft System Center 2019

This document does not cover the steps to install Microsoft System Center Operations Manager (SCOM) and Virtual Machine Manager (SCVMM). Follow the Microsoft guidelines to install SCOM and SCVMM 2016:

- SCOM: https://docs.microsoft.com/en-us/system-center/scom/deploy-overview
- SCVMM: https://docs.microsoft.com/en-us/system-center/vmm/install-console

## Citrix Virtual Apps and Desktops™ 1909

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 Desktops 1909, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

The Virtual Desktops 1909 release offers these benefits:

- Comprehensive virtual desktop delivery for any use case. The Virtual Desktops 1909 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 Desktops 1909 leverages common policies and cohesive tools to govern both infrastructure resources and user access.
- Simplified support and choice of BYO (Bring Your Own) devices. Virtual Desktops 1909 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. Virtual Desktops 1909 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 data center.
- 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 Virtual Desktops 1909.
- Improved high-definition user experience. Virtual Desktops 1909 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 Virtual Desktops 1909 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 Virtual Machines 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 18 Logical Architecture of Citrix Virtual Desktops

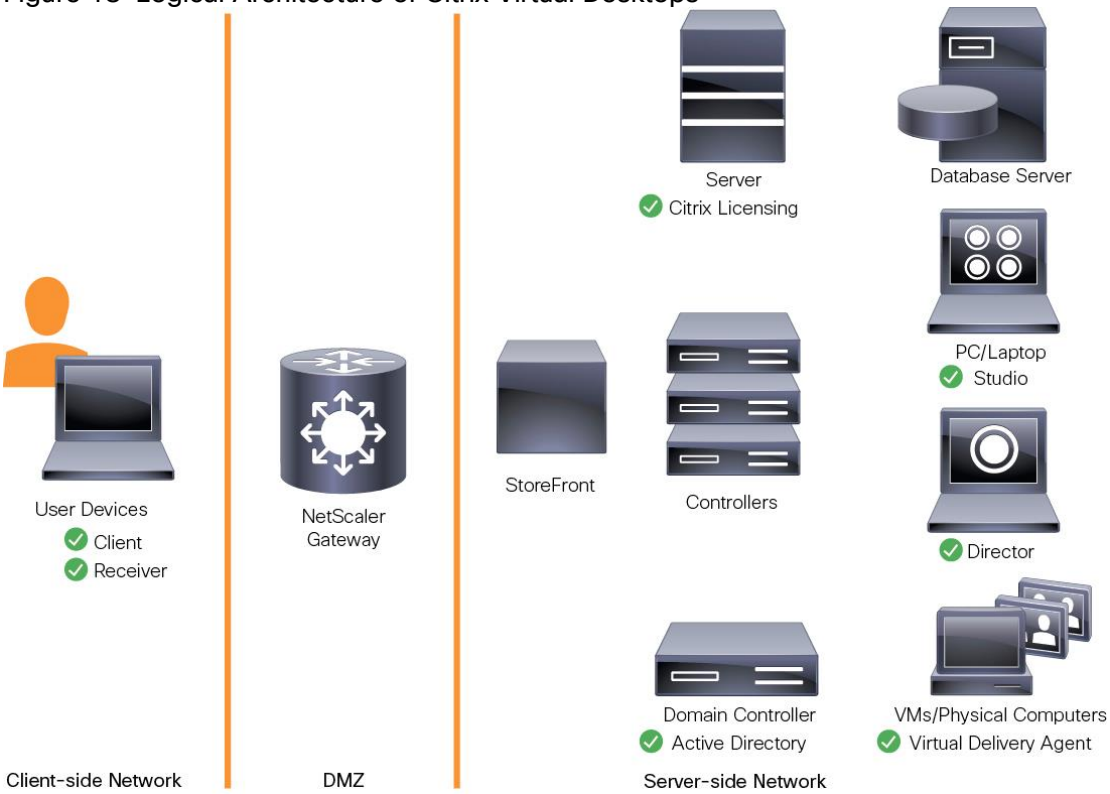


## Citrix Provisioning Services 1909

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.

## 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 19 provides a high-level view of a basic PVS infrastructure and shows how PVS components might appear within that implementation.

**Figure 19 Logical Architecture of Citrix Provisioning Services**

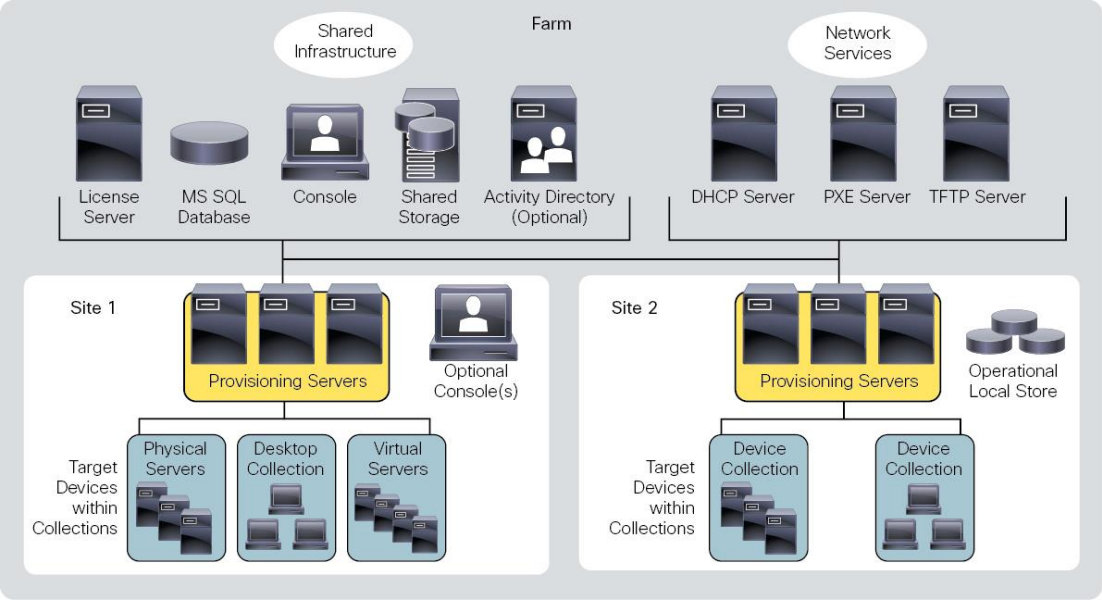


The following new features are available with Provisioning Services 1909:

- 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 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?
- What will Citrix Virtual Apps for Remote Desktop Server Hosted Sessions use?
- 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 Desktops 1909 integrates Hosted Shared and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, and manageable solution for delivering Windows applications and desktops as a service.

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

### Machine Catalogs

Collections of identical Virtual Machines or physical computers are managed as a single entity called a Machine Catalog. In this CVD, virtual machine provisioning relies on Citrix Provisioning Services 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 20 illustrates how users access desktops and applications through machine catalogs and delivery groups.

Figure 20 Access Desktops and Applications through Machine Catalogs and Delivery Groups

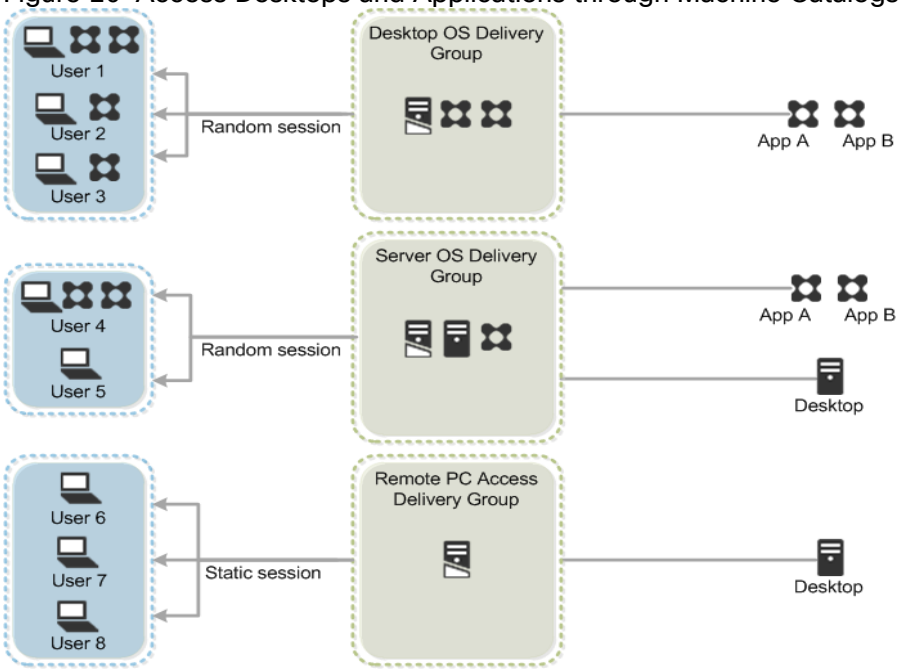


## Example Virtual Desktops Deployments

Two examples of typical Virtual Desktops deployments are:

- A distributed components configuration
- A multiple site configuration

Since Virtual Apps and Virtual Desktops 1909 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 21 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 the distributed components configuration (Figure 21). Two Cisco C220 rack servers host the required infrastructure services (AD, DNS, DHCP, Profile, SQL, Citrix Virtual Desktops management, and StoreFront servers).

Figure 21 Example of a Distributed Components Configuration

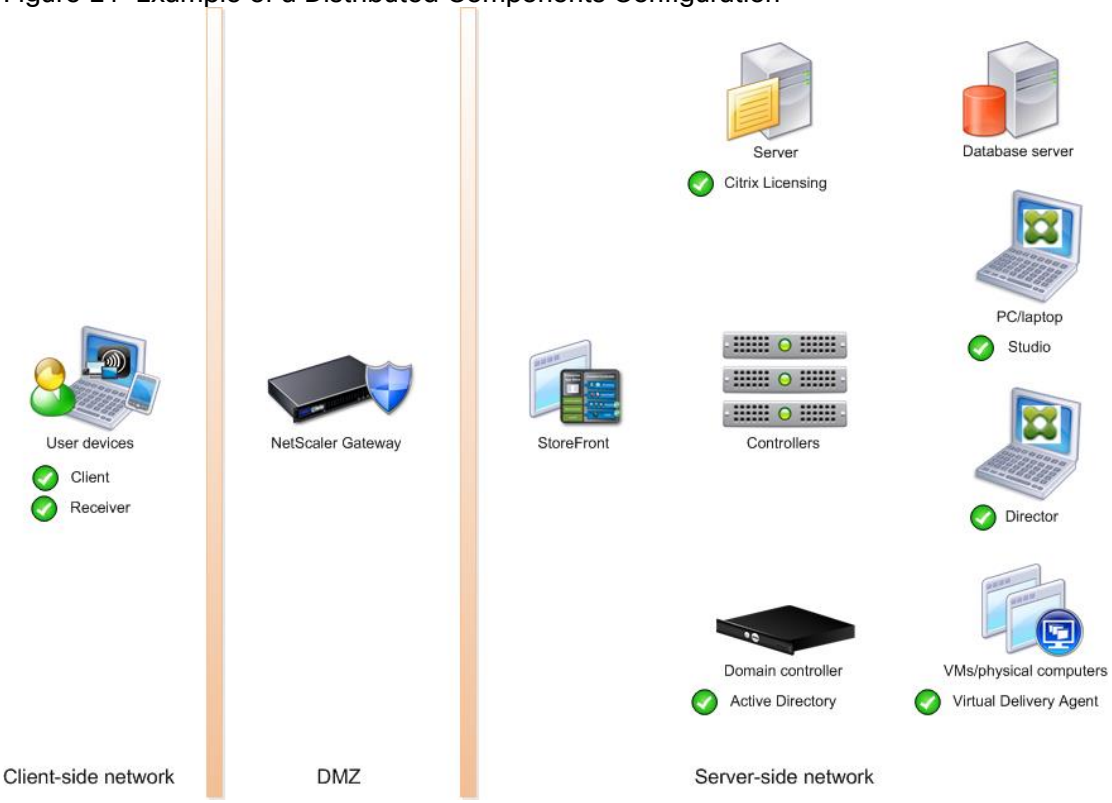


## Multiple Site Configuration

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 22 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 22 Multiple Sites

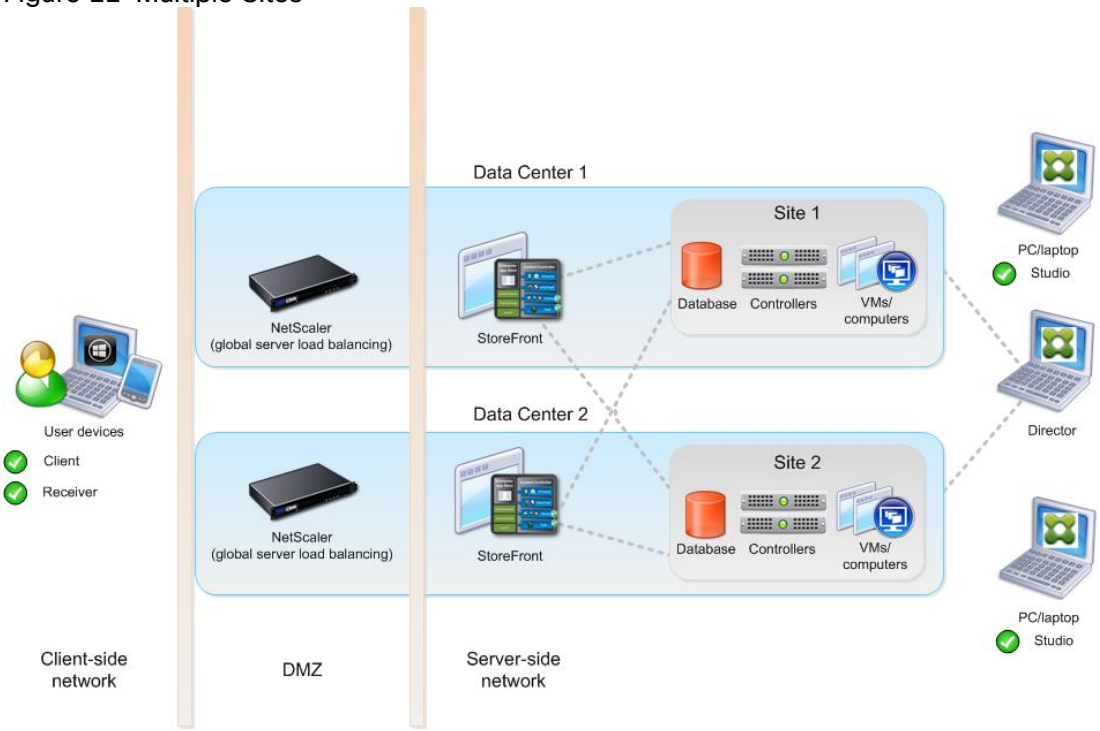


You can use StoreFront to aggregate resources from multiple sites to provide users with a single point of access with NetScaler. A separate Studio console is required to manage each site; sites cannot be managed as a single entity. You can use Director to support users across sites.

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 via a single-management plane to simplify administration

## Designing a Virtual Desktops Environment for a Mixed Workload

With Citrix Virtual Desktops 1909, 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 many users, while providing a secure, high-definition user experience.<br><br>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.<br><br>Application types: Any application. |
| Desktop OS machines | You want: A client-based application delivery solution that is secure, provides centralized management, and supports many users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition.<br><br>Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications.<br><br>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.<br><br>Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users. |
| Remote PC Access | You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the data center.<br><br>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.<br><br>Host: The same as Desktop OS machines.<br><br>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

## Physical Components

**Table 1 HyperFlex System Components**

| Component | Hardware Required |
|---|---|
| Fabric Interconnects | Two Cisco UCS 6332 Fabric Interconnects, or<br>Two Cisco UCS 6332-16UP Fabric Interconnects |
| Servers | Eight Cisco HyperFlex HXAF220c-M5SX All-Flash Rack Servers<br>Eight Cisco UCS B200 M5 Blade Servers for Compute-Only Nodes |

For complete server specifications and more information, please refer to the links below:

Compare Models:

http://www.cisco.com/c/en/us/products/hyperconverged-infrastructure/hyperflex-hx-series/index.html#compare-models

HXAF220c-M5SX Spec Sheet:

https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hxaf-220c-m5-specsheet.pdf

Table 2 lists the hardware component options for the HXAF220c-M5SX server model.

**Table 2 HXAF220c-M5SX Server Options**

| HXAF220c-M5SX options | | Hardware Required |
|---|---|---|
| Processors | | A pair of Intel Xeon Processor Scalable Family CPUs (6230 Gold) |
| Memory | | 768GB of total memory using 64 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<br>One 400 GB 2.5 Inch Enterprise Performance 12G SAS SSD<br>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 |
| Network | | Cisco UCS VIC1387 VIC MLOM |

| HXAF220c-M5SX options | Hardware Required |
|---|---|
| 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+ |

## Software Components

Table 3 lists the software components and the versions required for the Cisco HyperFlex system for Microsoft Hyper-V.

**Table 3 Software Components**

| Component | Software Required |
|---|---|
| Hypervisor | Hyper-V - Microsoft Windows Server 2019 Data Center<br><br>Note: Microsoft Windows Server with Hyper-V will NOT be installed in Cisco Factory. Customers need to bring their own Windows Server ISO image that needs to be installed at deployment site |
| Active Directory | A Windows 2012 or later domain and forest functionality level with AD integrated DNS server. |
| Management Server | Windows 10 or Windows Server 2016 with PowerShell and RSAT tools installed.<br><br>System Center VMM 2019<br><br>Windows Admin Center (Optional) |
| Cisco HyperFlex Data Platform | Cisco HyperFlex HX Data Platform Installer for Microsoft Hyper-V 4.0(1b) - Cisco-HX-Data-Platform-Installer-v4.0.1b-29681-hyperv.vhdx.zip |
| Microsoft Windows Server 2019 System Preparation Script | Cisco HyperFlex Data Platform System Preparation Script for Microsoft Windows Server 2019 with Cisco Drivers - HXInstall-HyperV-DatacenterCore-v4.0.1b-29681.img, or<br><br>Cisco HyperFlex Data Platform System Preparation Script for Microsoft Windows Server 2019 Desktop Experience with Cisco Drivers - HXInstall-HyperV-DatacenterDE-v4.0.1b-29681.img |
| Ready Clone PowerShell Script | Cisco HyperFlex Data Platform Hyper-V ReadyClone PowerShell Script<br><br>HxClone-HyperV-v4.0.1b-29681.ps1 |

| Component | Software Required |
|---|---|
| Cisco UCS Firmware | Cisco UCS Infrastructure software, Cisco UCS B-Series and C-Series bundles, revision 3.2(3g) or later. |

## Licensing

Cisco HyperFlex systems must be properly licensed using Cisco Smart Licensing, which is a cloud-based software licensing management solution used to automate many manual, time consuming and error prone licensing tasks. Cisco HyperFlex 2.5 and later communicate with the Cisco Smart Software Manager (CSSM) online service via a Cisco Smart Account, to check out or assign available licenses from the account to the Cisco HyperFlex cluster resources. Communications can be direct via the internet, they can be configured to communicate via a proxy server, or they can communicate with an internal Cisco Smart Software Manager satellite server, which caches and periodically synchronizes licensing data. In a small number of highly secure environments, systems can be provisioned with a Permanent License Reservation (PLR) which does not need to communicate with CSSM. Contact your Cisco sales representative or partner to discuss if your security requirements will necessitate use of these permanent licenses. New HyperFlex cluster installations will operate for 90 days without licensing as an evaluation period, thereafter the system will generate alarms and operate in a non-compliant mode. Systems without compliant licensing will not be entitled to technical support.

For more information about the Cisco Smart Software Manager satellite server, see: https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager-satellite.html

Beginning with Cisco HyperFlex 3.0, licensing of the system requires one license per node - Standard license.

Table 4 lists the licensing editions and the features available with each type of license.

**Table 4 HyperFlex System License Editions**

| HyperFlex Licensing Edition | Standard |
|---|---|
| Features Available | 8 Converged Nodes standard cluster with Fabric Interconnects (Compute-only nodes not supported)<br>All Cisco UCS M5 with SFF server models<br>Replication Factor 3<br>10 GbE or 40 GbE Ethernet |

# Considerations

## Version Control

The software revisions listed in Table 3 are the only valid and supported configuration at the time of the publishing of this validated design. Special care must be taken not to alter the revision of the hypervisor, vCenter server, Cisco HX platform software, or the Cisco UCS firmware without first consulting the appropriate release notes and compatibility matrixes to make sure that the system is not being modified into an unsupported configuration.

## Microsoft Windows Active Directory

The Microsoft Windows Active Directory 2012 or later is required due to the requirement Cisco HyperFlex for Microsoft Hyper-V. The Active Directory with integrated DNS server must be installed and operational prior to the installation of the Cisco HyperFlex HX Data Platform software.

**This document does not explain the installation and configuration of Microsoft Windows Active Directory and DNS server.**

## Scale

Cisco HyperFlex for Microsoft Hyper-V standard clusters currently scale from a minimum of 3 to a maximum of 8 Cisco HX-series converged nodes with small form factor (SFF) disks per cluster. A converged node is a member of the cluster which provides storage resources to the HX Distributed Filesystem. Once the maximum size of a single cluster has been reached, the environment can be “scaled out” by adding additional HX model servers to the Cisco UCS domain, installing an additional HyperFlex cluster on them, and controlling them via the same management host with PowerShell and RSAT tools installed.

**At the time of the publication of this document, Cisco HyperFlex for Microsoft Hyper-V does not support the following: Adding compute-only nodes to a cluster or expanding an existing cluster and Cisco UCS M4 server models and LFF disks are not supported**

Table 5 lists the minimum and maximum scale for various installations of the Cisco HyperFlex system with Microsoft Hyper-V.

**Table 5 HyperFlex Cluster Scale**

| Cluster Type | Minimum Converged Nodes Required | Maximum Converged Nodes Allowed | Maximum Compute-only Nodes Allowed |
|---|---|---|---|
| Standard with SFF disks | 3 | 16 | 16 |

## Capacity

Overall usable cluster capacity is based on a number of factors. The number of nodes in the cluster, the number and size of the capacity layer disks, and the replication factor of the HyperFlex HX Data Platform, all affect the cluster capacity.

Disk drive manufacturers have adopted a size reporting methodology using calculation by powers of 10, also known as decimal prefix. As an example, a 120 GB disk is listed with a minimum of 120 x 10^9 bytes of usable addressable capacity, or 120 billion bytes. However, many operating systems and filesystems report their space based on standard computer binary exponentiation, or calculation by powers of 2, also called binary prefix. In this example, 2^10 or 1024 bytes make up a kilobyte, 2^10 kilobytes make up a megabyte, 2^10 megabytes make up a gigabyte, and 2^10 gigabytes make up a terabyte. As the values increase, the disparity between the two systems of measurement and notation get worse, at the terabyte level, the deviation between a decimal prefix value and a binary prefix value is nearly 10 percent.

The International System of Units (SI) defines values and decimal prefix by powers of 10 as follows:

**Table 6 SI Unit Values (Decimal Prefix)**

| Value | Symbol | Name |
|---|---|---|
| 1000 bytes | kB | Kilobyte |
| 1000 kB | MB | Megabyte |
| 1000 MB | GB | Gigabyte |
| 1000 GB | TB | Terabyte |

The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC) defines values and binary prefix by powers of 2 in ISO/IEC 80000-13:2008 Clause 4 listed in Table 7

**Table 7 IEC Unit Values (binary prefix)**

| Value | Symbol | Name |
|---|---|---|
| 1024 bytes | KiB | Kibibyte |
| 1024 KiB | MiB | Mebibyte |
| 1024 MiB | GiB | Gibibyte |
| 1024 GiB | TiB | Tebibyte |

For the purpose of this document, the decimal prefix numbers are used only for raw disk capacity as listed by the respective manufacturers. For all calculations where raw or usable capacities are shown from the perspective of the HyperFlex software, filesystems or operating systems, the binary prefix numbers are used. This is done primarily to show a consistent set of values as seen by the end user from within the HyperFlex Connect GUI when viewing cluster capacity, allocation and consumption, and also within most operating systems.

Table 8 lists a set of HyperFlex HX Data Platform cluster usable capacity values, using binary prefix, for an array of cluster configurations. These values provide an example of the capacity calculations, for determining the appropriate size of HX cluster to initially purchase, and how much capacity can be gained by adding capacity disks. The calculations for these values are listed in Appendix A: Cluster Capacity Calculations.

**Table 8 Cluster Usable Capacities**

| HX-Series Server Model | Node Quantity | Capacity Disk Size (each) | Capacity Disk Quantity (per node) | Cluster Usable Capacity at RF=2 | Cluster Usable Capacity at RF=3 |
|---|---|---|---|---|---|
| HXAF220c-M5SX | 8 | 3.8 TB | 8 | 102.8 TiB | 68.6 TiB |
| | | 960 GB | 8 | 25.7 TiB | 17.1 TiB |
| HXAF240c- | 8 | 3.8 TB | 6 | 77.1 TiB | 51.4 TiB |

| HX-Series Server Model | Node Quantity | Capacity Disk Size (each) | Capacity Disk Quantity (per node) | Cluster Usable Capacity at RF=2 | Cluster Usable Capacity at RF=3 |
|---|---|---|---|---|---|
| M5SX (Not used in this solution) | | | 15 | 192.8 TiB | 128.5 TiB |
| | | | 23 | 295.7 TiB | 197.1 TiB |
| | | 960 GB | 6 | 19.3 TiB | 12.9 TiB |
| | | | 15 | 48.2 TiB | 32.1 TiB |
| | | | 23 | 73.9 TiB | 49.3 TiB |

**Calculations are based on the number of nodes, the number of capacity disks per node, and the size of the capacity disks. Table 8 is not a comprehensive list of all capacities and models available.**

# Physical Topology

## Topology Overview

The Cisco HyperFlex for Microsoft Hyper-V system is composed of a pair of Cisco UCS Fabric Interconnects along with up to 8 HX-Series rack-mount servers as converged nodes per cluster. Up to eight separate HX clusters can be installed under a single pair of Fabric Interconnects. The two Fabric Interconnects both connect to every HX-Series rack-mount server. Upstream network connections, also referred to as “northbound” network connections are made from the Fabric Interconnects to the customer data center network at the time of installation.

Figure 23 HyperFlex Standard Cluster Topology

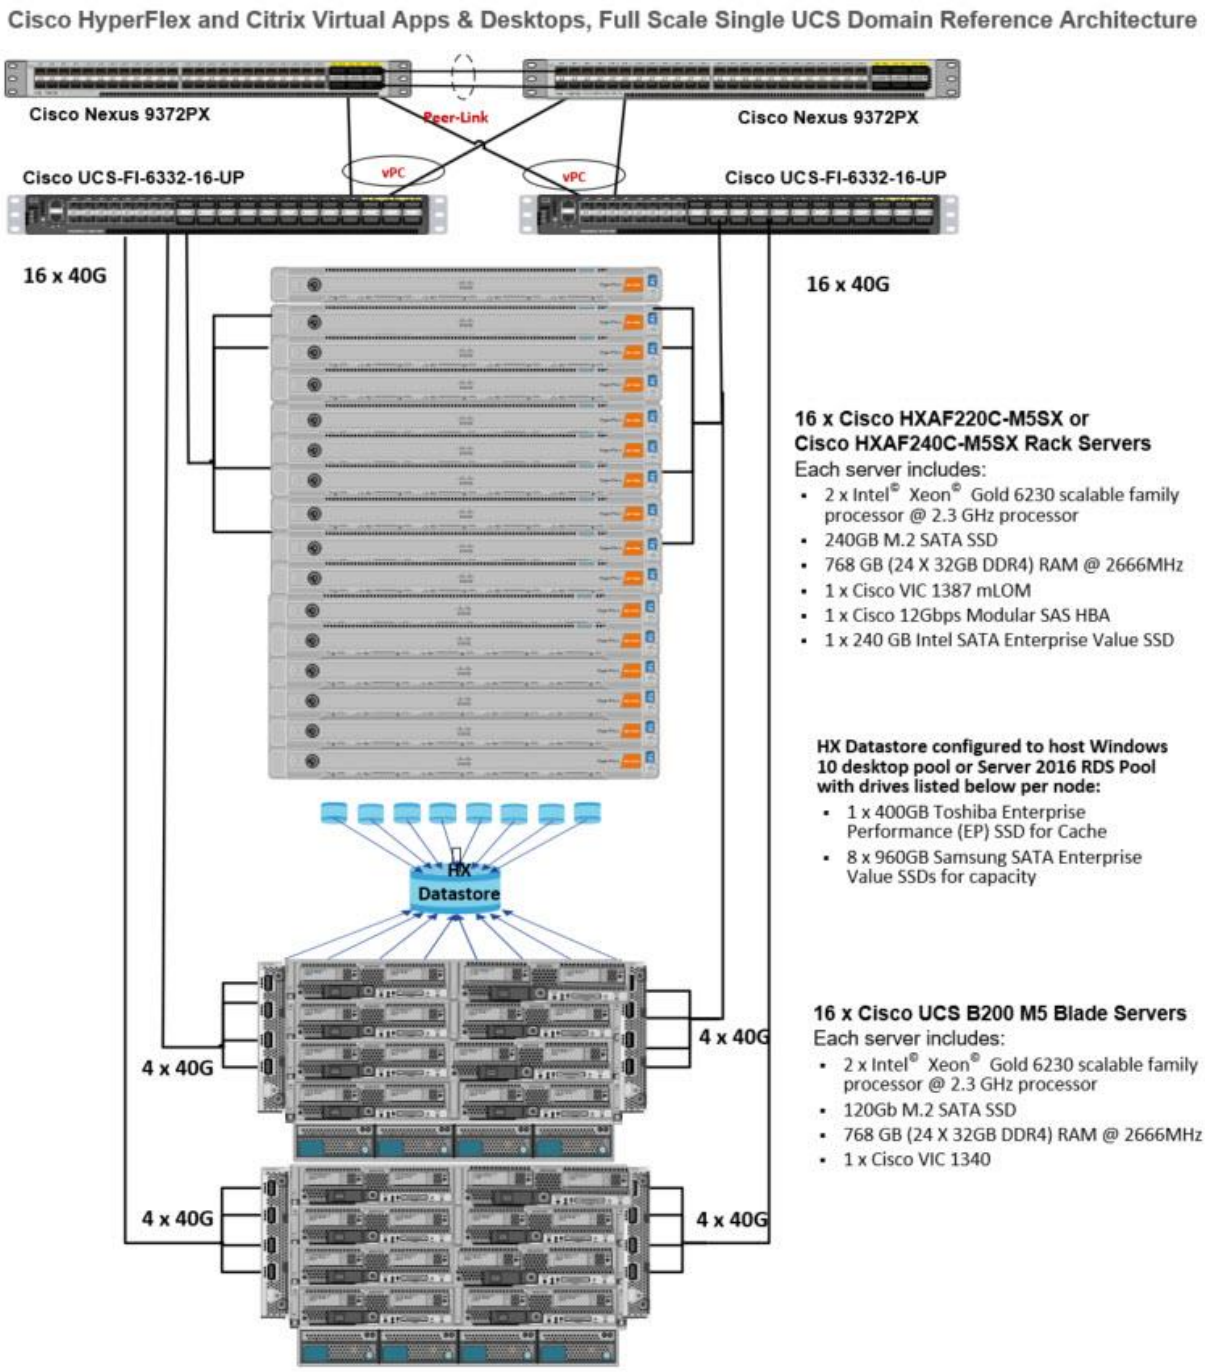


## Fabric Interconnects

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 also 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. This port is 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 port 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 port 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. This port is 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 Rack-Mount Servers

The HX-Series converged 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 using a single cable for both management traffic and data traffic. Cisco HyperFlex M5 generation servers can be configured only with the Cisco VIC 1387 card. The standard and redundant connection practice is to connect port 1 of the VIC card (the right-hand port) to a port on FI A, and port 2 of the VIC card (the left-hand port) to a port on FI B (Figure 24). The HyperFlex installer checks for this configuration, and that all servers' cabling matches. Failure to follow this cabling practice can lead to errors, discovery failures, and loss of redundant connectivity.

Various combinations of physical connectivity between the Cisco HX-series servers and the Fabric Interconnects are possible, but only specific combinations are supported. For example, use of the Cisco QSA module to convert a 40 GbE QSFP+ port into a 10 GbE SFP+ port is allowed for M5 generation servers in order to configure M5 generation servers along with model 6248 or 6296 Fabric Interconnects. Table 9 lists the possible connections, and which of these methods is supported.

**Table 9 Supported Physical Connectivity**

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

Figure 24 HX-Series Server Connectivity

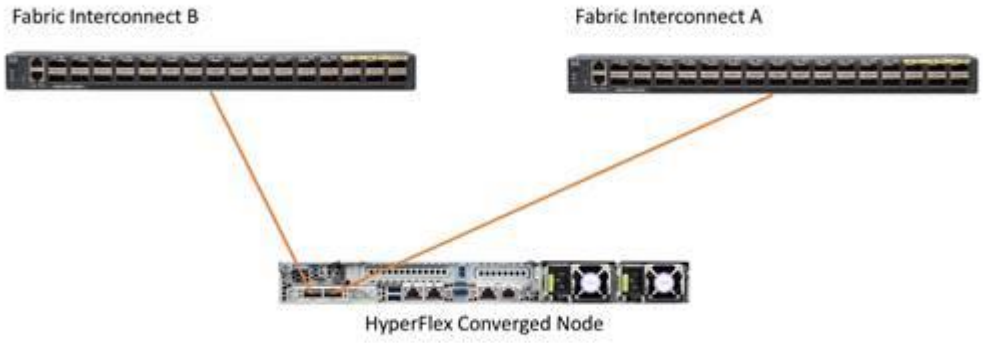


## Logical Architecture

The logical architecture of this solution is designed to support up to 2500 Hosted Virtual Microsoft Windows 10 Desktops users or 1600 RDS users within an eight node Cisco UCS HXAF220c- HyperFlex cluster, which provides physical redundancy for each workload type.

Figure 25 Logical Architecture Design

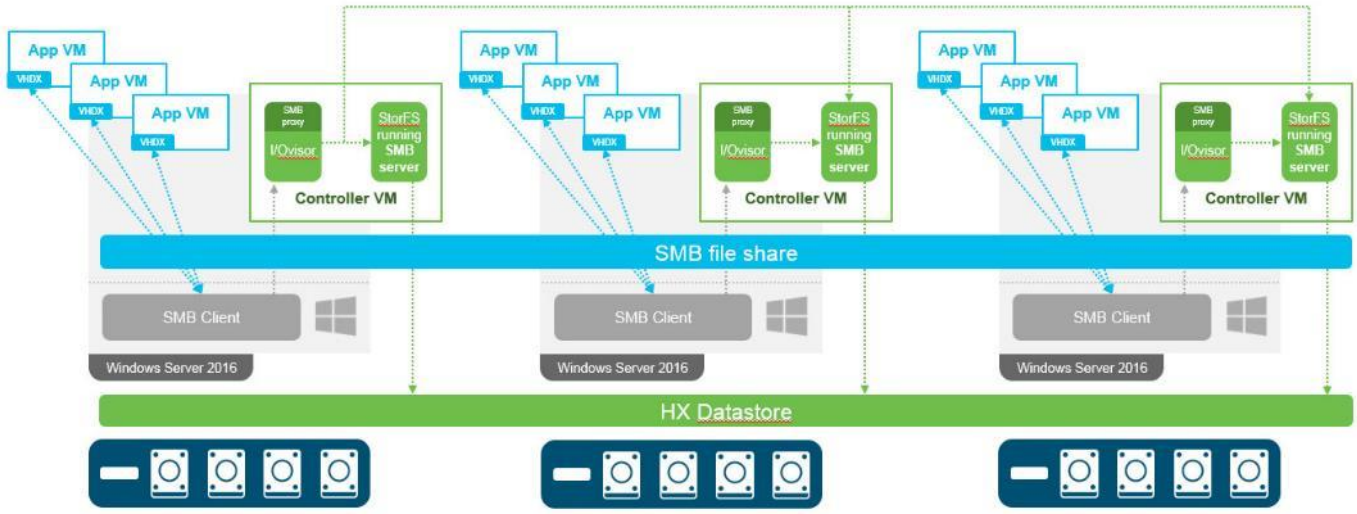


Table 3 lists the software revisions for this solution.

**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 10 through Table 34 lists the information you need to configure your environment.**

# Design Elements

Installing the HyperFlex system is primarily done through a deployable HyperFlex installer virtual machine, available for download at cisco.com as an OVA file. The installer virtual machine performs most of the Cisco UCS configuration work, it can be leveraged to simplify the installation of Windows Server 2016 on the HyperFlex hosts,

and also performs significant portions of the configuration. Finally, the installer virtual machine is used to install the HyperFlex HX Data Platform software and create the HyperFlex cluster. Because this simplified installation method has been developed by Cisco, this CVD will not give detailed manual steps for the configuration of all the elements that are handled by the installer. Instead, the elements configured will be described and documented in this section, and the subsequent sections will guide you through the manual steps needed for installation, and how to utilize the HyperFlex Installer for the remaining configuration steps.

## Network Design

### Cisco UCS Uplink Connectivity

Cisco UCS network uplinks connect "northbound" from the pair of Cisco UCS Fabric Interconnects to the LAN in the customer data center. All Cisco UCS uplinks operate as trunks, carrying multiple 802.1Q VLAN IDs across the uplinks. The default Cisco UCS behavior is to assume that all VLAN IDs defined in the Cisco UCS configuration are eligible to be trunked across all available uplinks.

Cisco UCS Fabric Interconnects appear on the network as a collection of endpoints versus another network switch. Internally, the Fabric Interconnects do not participate in spanning-tree protocol (STP) domains, and the Fabric Interconnects cannot form a network loop, as they are not connected to each other with a layer 2 Ethernet link. All link up/down decisions via STP will be made by the upstream root bridges.

Uplinks need to be connected and active from both Fabric Interconnects. For redundancy, multiple uplinks can be used on each FI, either as 802.3ad Link Aggregation Control Protocol (LACP) port-channels or using individual links. For the best level of performance and redundancy, uplinks can be made as LACP port-channels to multiple upstream Cisco switches using the virtual port channel (vPC) feature. Using vPC uplinks allows all uplinks to be active passing data, plus protects against any individual link failure, and the failure of an upstream switch. Other uplink configurations can be redundant, but spanning-tree protocol loop avoidance may disable links if vPC is not available.

All uplink connectivity methods must allow for traffic to pass from one Fabric Interconnect to the other, or from fabric A to fabric B. There are scenarios where cable, port or link failures would require traffic that normally does not leave the Cisco UCS domain, to now be forced over the Cisco UCS uplinks. Additionally, this traffic flow pattern can be seen briefly during maintenance procedures, such as updating firmware on the Fabric Interconnects, which requires them to be rebooted. The following section detail the uplink connectivity option used for this solution.

### vPC to Multiple Switches

This recommended connection design relies on using Cisco switches that have the virtual port channel feature, such as Catalyst 6000 series switches running VSS, Cisco Nexus 5000 series, and Cisco Nexus 9000 series switches. Logically the two vPC enabled switches appear as one, and therefore spanning-tree protocol will not block any links. This configuration allows for all links to be active, achieving maximum bandwidth potential, and multiple redundancy at each level.

Figure 26 Connectivity with vPC

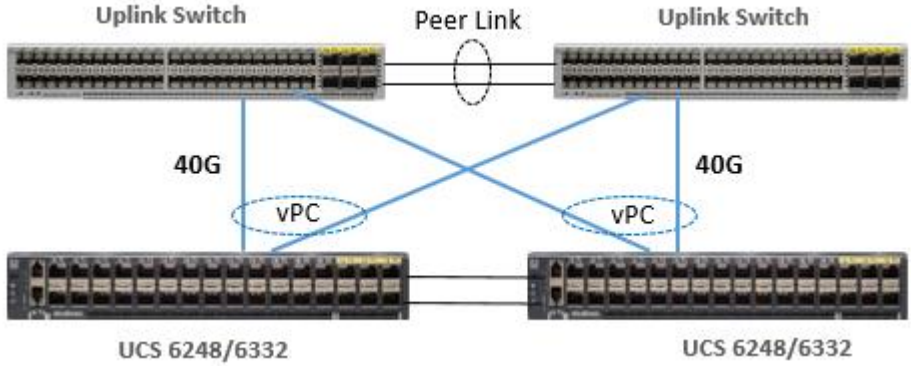


## VLANs

The VLAN configuration recommended for the environment includes a total of seven VLANs as outlined in Table 10 .

Table 10 VLANs Configured in this Study

| VLAN Name | VLAN ID | VLAN Purpose |
|---|---|---|
| Default | 1 | Native VLAN |
| Hx-in-Band-Mgmt | 30 | VLAN for in-band management interfaces |
| Infra-Mgmt | 32 | VLAN for Virtual Infrastructure |
| Hx-storage-data | 101 | VLAN for HyperFlex Storage |
| Hx-livemigration | 33 | VLAN for Hyper-V Live Migration |
| Vm-network | 34 | 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 data centers. 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 in Table 10 .**

## 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.

# Cisco UCS Design

This section describes the elements within Cisco UCS Manager that are configured 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, external management IP pools and more. Where the elements can be manually set during the installation, those items will be noted in << >> brackets.

## Cisco UCS Organization

During the HyperFlex installation a new Cisco UCS Sub-Organization is created. You must specify a unique Sub-Organization name for each cluster during the installation, for example “hx1hybrid”, or “hx2sed”. The sub-organization is created underneath the root level of the Cisco UCS hierarchy, and is used to contain all policies, pools, templates and service profiles used by HyperFlex, which prevents problems from overlapping settings across policies and pools. This arrangement also allows for organizational control using Role-Based Access Control (RBAC) and administrative locales within Cisco UCS Manager later 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 27 Cisco UCS HyperFlex Sub-Organization**

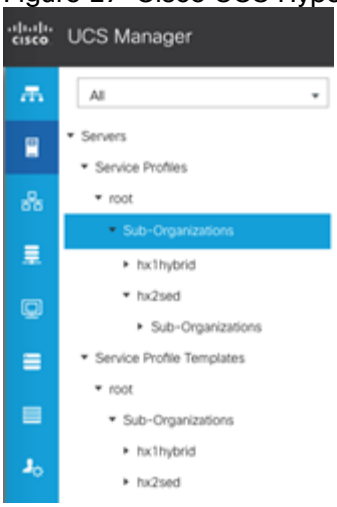


## Cisco UCS LAN Policies

### QoS System Classes

Specific Cisco UCS Quality of Service (QoS) system classes are defined for a Cisco HyperFlex system. These classes define Class of Service (CoS) values that can be used by the uplink switches north of the Cisco UCS domain, plus which classes are active, along with whether packet drop is allowed, the relative weight of the different classes when there is contention, the maximum transmission unit (MTU) size, and if there is multicast optimization applied. QoS system classes are defined for the entire Cisco UCS domain, the classes that are enabled can later be used in QoS policies, which are then assigned to Cisco UCS vNICs. Table 11 and Figure 28 list the QoS System Class settings configured for HyperFlex.

**Table 11 QoS System Classes**

| Priority | Enabled | CoS | Packet Drop | Weight | MTU | Multicast Optimized |
|---|---|---|---|---|---|---|

| Priority | Enabled | CoS | Packet Drop | Weight | MTU | Multicast Optimized |
|---|---|---|---|---|---|---|
| Platinum | Yes | 5 | No | 4 | 9216 | No |
| Gold | Yes | 4 | Yes | 4 | Normal | No |
| Silver | Yes | 2 | Yes | Best-effort | Normal | Yes |
| Bronze | Yes | 1 | Yes | Best-effort | 9216 | No |
| Best Effort | Yes | Any | Yes | Best-effort | Normal | No |
| Fibre Channel | Yes | 3 | No | 5 | FC | N/A |

Figure 28 QoS System Classes

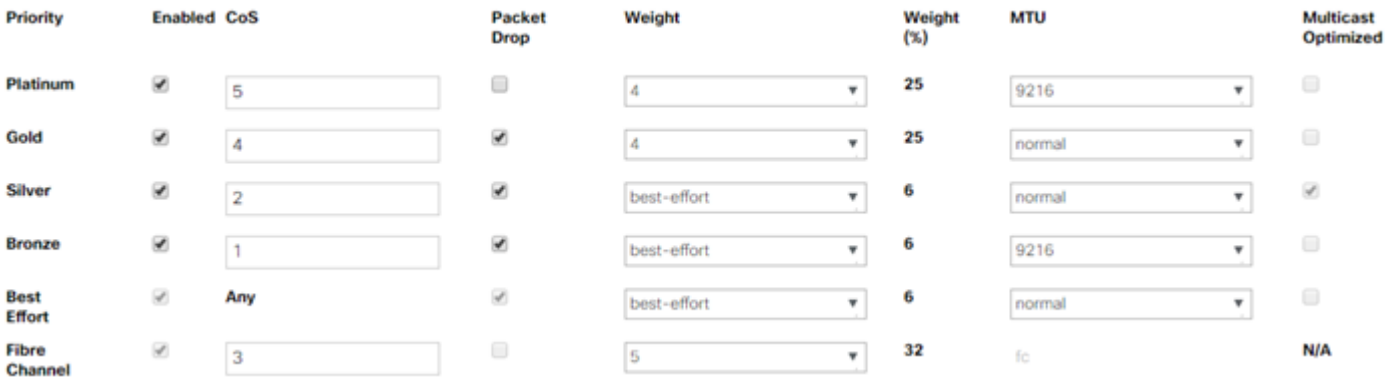


Changing the QoS system classes on a Cisco UCS 6332 or 6332-16UP model Fabric Interconnect requires both FIs to reboot in order to take effect.

## QoS Policies

In order to apply the settings defined in the Cisco UCS QoS System Classes, specific QoS Policies must be created, and then assigned to the vNICs, or vNIC templates used in Cisco UCS Service Profiles. Table 12 details the QoS Policies configured for HyperFlex, and their default assignment to the vNIC templates created:

Table 12 HyperFlex QoS Policies

| Policy | Priority | Burst | Rate | Host Control | Used by vNIC Template |
|---|---|---|---|---|---|
| Platinum | Platinum | 10240 | Line-rate | None | storage-data-a<br>storage-data-b |
| Gold | Gold | 10240 | Line-rate | None | vm-network-a<br>vm-network-b |

| Policy | Priority | Burst | Rate | Host Control | Used by vNIC Template |
|---|---|---|---|---|---|
| Silver | Silver | 10240 | Line-rate | None | hv-mgmt-a<br>hv-mgmt-b |
| Bronze | Bronze | 10240 | Line-rate | None | hv-livemigrate-a<br>hv- livemigrate -b |
| Best Effort | Best Effort | 10240 | Line-rate | None | N/A |

## Multicast Policy

A Cisco UCS Multicast Policy is configured by the HyperFlex installer, which is referenced by the VLANs that are created. The policy allows for future flexibility if a specific multicast policy needs to be created and applied to other VLANs that may be used by non-HyperFlex workloads in the Cisco UCS domain. Table 13 and Figure 29 details the Multicast Policy configured for HyperFlex:

**Table 13 Multicast Policy**

| Name | IGMP Snooping State | IGMP Snooping Queries State |
|---|---|---|
| HyperFlex | Enabled | Disabled |

**Figure 29 Multicast Policy**

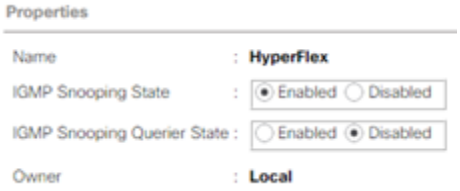


## VLANs

VLANs are created by the HyperFlex installer to support a base HyperFlex system, with a VLAN for live migrate, and a single or multiple VLANs defined for guest virtual machine traffic. Names and IDs for the VLANs are defined in the Cisco UCS configuration page of the HyperFlex installer web interface. The VLANs listed in Cisco UCS must already be present on the upstream network, and the Cisco UCS FIs do not participate in VLAN Trunk Protocol (VTP). Table 14 and Figure 30 list the VLANs configured for HyperFlex.

**Table 14 Cisco UCS VLANs**

| Name | ID | Type | Transport | Native | VLAN Sharing | Multicast Policy |
|---|---|---|---|---|---|---|
| <<hx-inband-mgmt>> | <user_defined> | LAN | Ether | No | None | HyperFlex |
| <<hx-storage-data>> | <user_defined> | LAN | Ether | No | None | HyperFlex |
| <<vm-network>> | <user_defined> | LAN | Ether | No | None | HyperFlex |

| Name | ID | Type | Transport | Native | VLAN Sharing | Multicast Policy |
|---|---|---|---|---|---|---|
| <<hx-livemigrate>> | <user_defined> | LAN | Ether | No | None | HyperFlex |

Figure 30 Cisco UCS VLANs

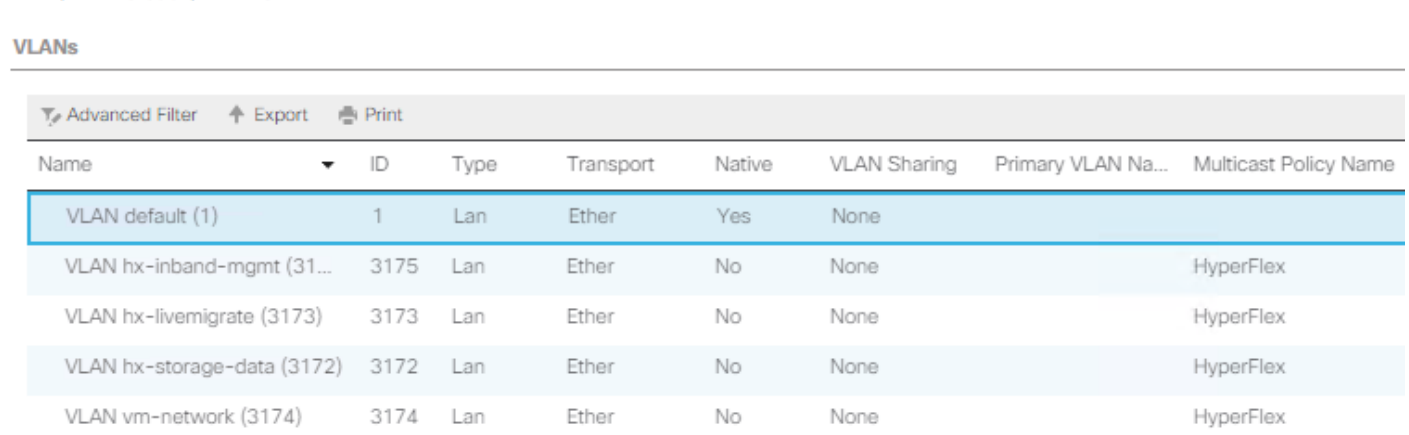


## Management IP Address Pool

A Cisco UCS Management IP Address Pool must be populated with a block of IP addresses. These IP addresses are assigned to the Cisco Integrated Management Controller (CIMC) interface of the rack mount and blade servers that are managed in the Cisco UCS domain. The IP addresses are the communication endpoints for various functions, such as remote KVM, virtual media, Serial over LAN (SoL), and Intelligent Platform Management Interface (IPMI) for each rack mount or blade server. Therefore, a minimum of one IP address per physical server in the domain must be provided. The IP addresses are an "out-of-band" address, meaning that the communication pathway uses the Fabric Interconnects' mgmt0 ports, which answer ARP requests for the management addresses. Because of this arrangement, the IP addresses in this pool must be in the same IP subnet as the IP addresses assigned to the Fabric Interconnects' mgmt0 ports. A new IP pool, named "hx-ext-mgmt" is created in the HyperFlex sub-organization, and populated with a block of IP addresses, a subnet mask, and a default gateway by the HyperFlex installer.

Figure 31 Management IP Address Pool

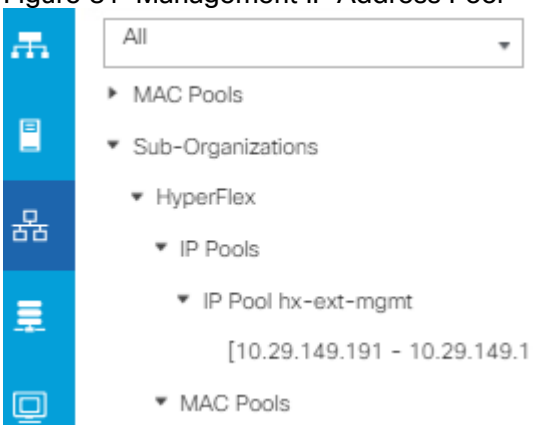


## MAC Address Pools

One of the core benefits of the Cisco UCS and Virtual Interface Card (VIC) technology is the assignment of the personality of the card via Cisco UCS Service Profiles. The number of virtual NIC (vNIC) interfaces, their VLAN

associations, MAC addresses, QoS policies and more are all applied dynamically as part of the service profile association process. Media Access Control (MAC) addresses use 6 bytes of data as a unique address to identify the interface on the layer 2 network. All devices are assigned a unique MAC address, which is ultimately used for all data transmission and reception. The Cisco UCS and VIC technology picks a MAC address from a pool of addresses and assigns it to each vNIC defined in the service profile when that service profile is created.

Best practices mandate that MAC addresses used for Cisco UCS domains use 00:25:B5 as the first three bytes, which is one of the Organizationally Unique Identifiers (OUI) registered to Cisco Systems, Inc. The fourth byte (for example, 00:25:B5:xx) is specified during the HyperFlex installation. The fifth byte is set automatically by the HyperFlex installer, to correlate to the Cisco UCS fabric and the vNIC placement order. Finally, the last byte is incremented according to the number of MAC addresses created in the pool. To avoid overlaps, when you define the values in the HyperFlex installer you must ensure that the first four bytes of the MAC address pools are unique for each HyperFlex cluster installed in the same layer 2 network, and also different from MAC address pools in other Cisco UCS domains which may exist.

Table 15 list the MAC Address Pools configured for HyperFlex and their default assignment to the vNIC templates created.

**Table 15 MAC Address Pools**

| Name | Block Start | Size | Assignment Order | Used by vNIC Template |
|---|---|---|---|---|
| hv-mgmt-a | 00:25:B5:<xx>:A1:01 | 100 | Sequential | hv-mgmt-a |
| hv-mgmt-b | 00:25:B5:<xx>:B2:01 | 100 | Sequential | hv-mgmt-b |
| hv-livemigrate-a | 00:25:B5:<xx>:A7:01 | 100 | Sequential | hv-livemigrate-a |
| hv-livemigrate-b | 00:25:B5:<xx>:B8:01 | 100 | Sequential | hv-livemigrate-b |
| storage-data-a | 00:25:B5:<xx>:A3:01 | 100 | Sequential | storage-data-a |
| storage-data-b | 00:25:B5:<xx>:B4:01 | 100 | Sequential | storage-data-b |
| vm-network-a | 00:25:B5:<xx>:A5:01 | 100 | Sequential | vm-network-a |
| vm-network-b | 00:25:B5:<xx>:B6:01 | 100 | Sequential | vm-network-b |

Figure 32 MAC Address Pools

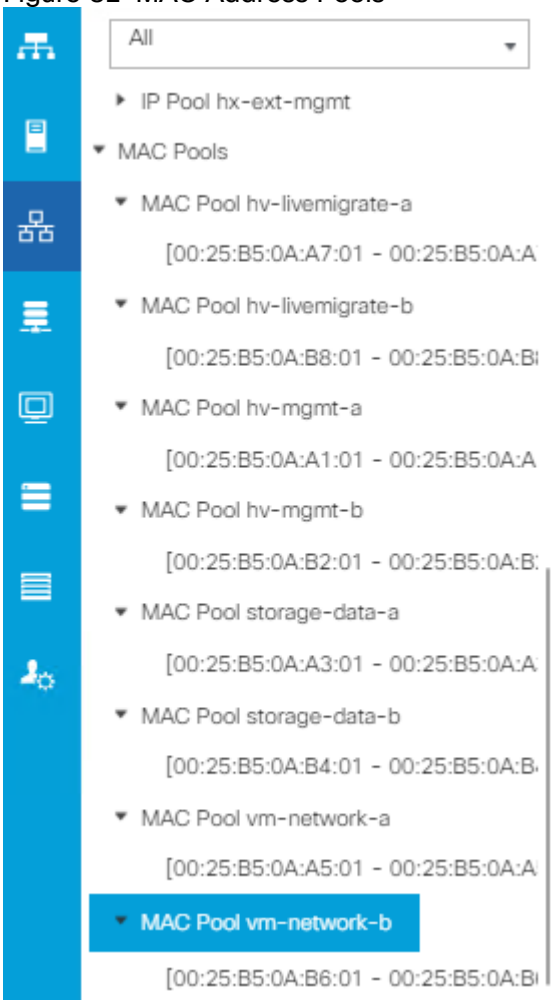


## Network Control Policies

Cisco UCS Network Control Policies control various aspects of the behavior of vNICs defined in the Cisco UCS Service Profiles. Settings controlled include enablement of Cisco Discovery Protocol (CDP), MAC address registration, MAC address forging, and the action taken on the vNIC status if the Cisco UCS network uplinks are failed. Two policies are configured by the HyperFlex Installer, HyperFlex-infra is applied to the “infrastructure” vNIC interfaces of the HyperFlex system, and HyperFlex-vm, which is only applied to the vNIC interfaces carrying guest virtual machine traffic. This allows for more flexibility, even though the policies are currently configured with the same settings. Table 16 details the Network Control Policies configured for HyperFlex, and their default assignment to the vNIC templates created:

Table 16 Network Control Policy

| Name | CDP | MAC Register Mode | Action on Uplink Fail | MAC Security | Used by vNIC Template |
|---|---|---|---|---|---|
| HyperFlex-infra | Enabled | Only Native VLAN | Link-down | Forged: Allow | hv-mgmt-a<br>hv-mgmt-b<br>hv-livemigrate-a |

| Name | CDP | MAC Register Mode | Action on Uplink Fail | MAC Security | Used by vNIC Template |
|---|---|---|---|---|---|
| | | | | | hv-livemigrate-b<br>storage-data-a<br>storage-data-b |
| HyperFlex-vm | Enabled | Only Native VLAN | Link-down | Forged: Allow | vm-network-a<br>vm-network-b |

Figure 33 Network Control Policy

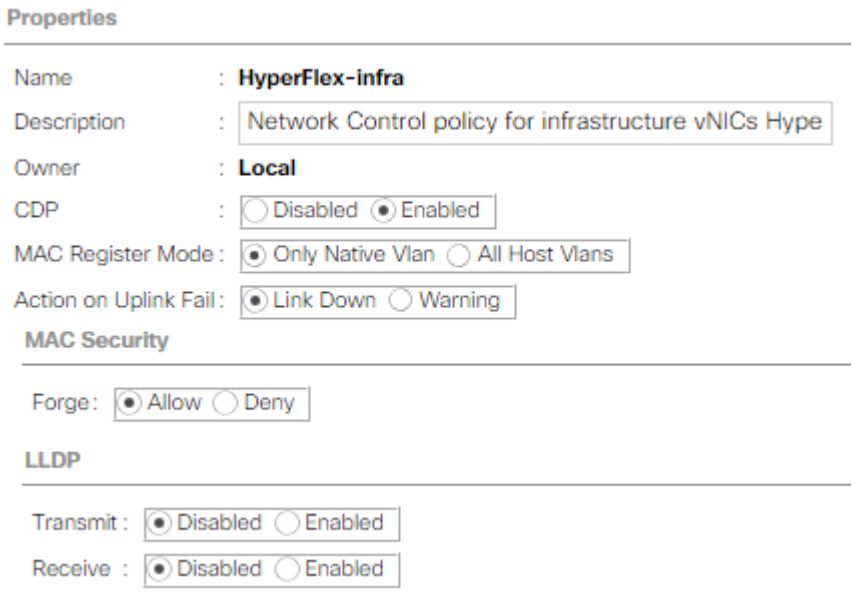


## vNIC Templates

Cisco UCS Manager has a feature to configure vNIC templates, which can be used to simplify and speed up configuration efforts. VNIC templates are referenced in service profiles and LAN connectivity policies, versus configuring the same vNICs individually in each service profile, or service profile template. VNIC templates contain all the configuration elements that make up a vNIC, including VLAN assignment, MAC address pool selection, fabric A or B assignment, fabric failover, MTU, QoS policy, Network Control Policy, and more. Templates are created as either initial templates or updating templates. Updating templates retain a link between the parent template and the child object, therefore when changes are made to the template, the changes are propagated to all remaining linked child objects. An additional feature named “vNIC Redundancy” allows vNICs to be configured in pairs, so that the settings of one vNIC template, designated as a primary template, will automatically be applied to a configured secondary template. For all HyperFlex vNIC templates, the “A” side vNIC template is configured as a primary template, and the related “B” side vNIC template is a secondary. In each case, the only configuration difference between the two templates is which fabric they are configured to connect through. The following tables list the initial settings in each of the vNIC templates created by the HyperFlex installer:

Table 17 vNIC Template hv-mgmt-a/b

| vNIC Template | hv-mgmt-a | hv-mgmt-b | storage-data-a | storage-data-b | hv-livemigrate-a | hv-livemigrate-b | vm-network-a | vm-network-b |
|---|---|---|---|---|---|---|---|---|

| Name: | | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Setting | Value | Value | Value | Value | Value | Value | Value | Value |
| Fabric ID | A | B | A | B | A | B | A | B |
| Fabric Failover | Disabled | Disabled | Disabled | Disabled | Disabled | Disabled | Disabled | Disabled |
| Target | Adapter | Adapter | Adapter | Adapter | Adapter | Adapter | Adapter | Adapter |
| Type | Updating Template | Updating Template | Updating Template | Updating Template | Updating Template | Updating Template | Updating Template | Updating Template |
| MTU | 1500 | 1500 | 9000 | 9000 | 9000 | 9000 | 1500 | 1500 |
| MAC Pool | hv-mgmt-a | hv-mgmt-b | storage-data-a | storage-data-b | hv-livemigrate-a | hv-livemigrate-b | vm-network-a | vm-network-b |
| QoS Policy | silver | silver | platinum | platinum | bronze | bronze | gold | gold |
| Network Control Policy | HyperFlex-infra | HyperFlex-infra | HyperFlex-infra | HyperFlex-infra | HyperFlex-infra | HyperFlex-infra | HyperFlex-vm | HyperFlex-vm |
| VLANs | <<hx-inband-mgmt>> | <<hx-inband-mgmt>> | <<hx-storage-data>> | <<hx-storage-data>> | <<hx-livemigrate>> | <<hx-livemigrate>> | <<vm-network>> | <<vm-network>> |
| Native VLAN | No | No | No | No | No | No | No | No |

## LAN Connectivity Policies

Cisco UCS Manager has a feature for LAN Connectivity Policies, which aggregates all the vNICs or vNIC templates desired for a service profile configuration into a single policy definition. This simplifies configuration efforts by defining a collection of vNICs or vNIC templates once and using that policy in the service profiles or service profile templates. The HyperFlex installer configures a LAN Connectivity Policy named HyperFlex, which contains all the vNIC templates defined in the previous section, along with an Adapter Policy named HyperFlex, also configured by the HyperFlex installer. Table 18 lists the LAN Connectivity Policy configured for HyperFlex.

**Table 18 LAN Connectivity Policy**

| Policy Name | Use vNIC Template | vNIC Name | vNIC Template Used | Adapter Policy |
|---|---|---|---|---|
| HyperFlex | Yes | hv-mgmt-a | hv-mgmt-a | HyperFlex |
| | | hv-mgmt-b | hv-mgmt-b | |

| Policy Name | Use vNIC Template | vNIC Name | vNIC Template Used | Adapter Policy |
|---|---|---|---|---|
| | | hv-livemigrate-a | hv-livemigrate-a | |
| | | hv-livemigrate-b | hv-livemigrate-b | |
| | | storage-data-a | storage-data-a | |
| | | storage-data-b | storage-data-b | |
| | | vm-network-a | vm-network-a | |
| | | vm-network-b | vm-network-b | |

## Cisco UCS Servers Policies

### Adapter Policies

Cisco UCS Adapter Policies are used to configure various settings of the Converged Network Adapter (CNA) installed in the Cisco UCS blade or rack-mount servers. Various advanced hardware features can be enabled or disabled depending on the software or operating system being used. The following figures detail the Adapter Policy named “HyperFlex” configured for HyperFlex.

**Figure 34 Cisco UCS Adapter Policy Resources**

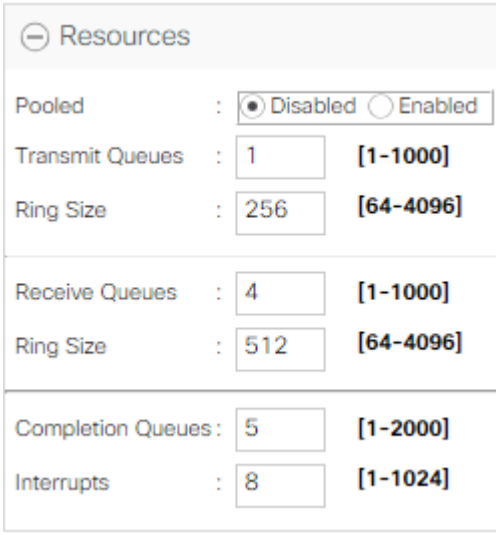

Figure 35 Cisco UCS Adapter Policy Options

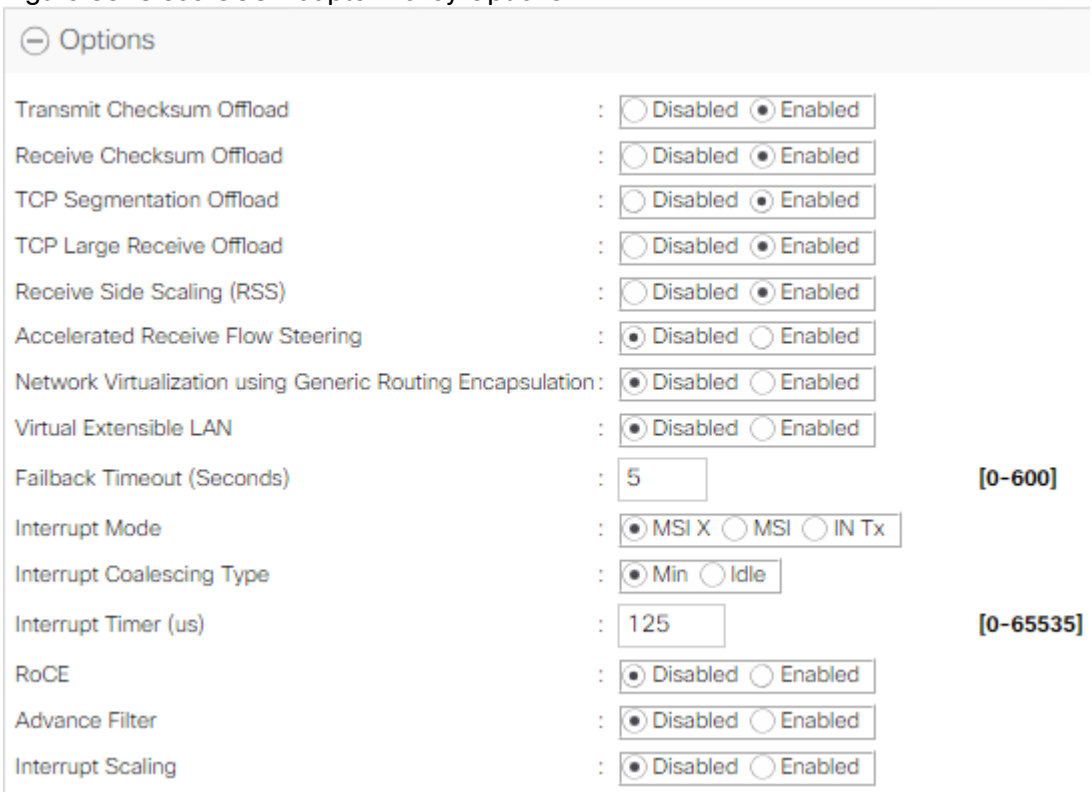


## BIOS Policies

Cisco HX-Series M5 generation servers no longer use pre-defined BIOS setting defaults derived from Cisco UCS Manager, instead the servers have default BIOS tokens set from the factory. The current default token settings can be viewed at the following website: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/Reference-Docs/Server-BIOS-Tokens/3-2/b_UCS_BIOS_Tokens.html

A BIOS policy, named “HyperFlex-m5” is created by the HyperFlex installer to modify the setting of M5 generation servers. The modified settings are as follows:

- System altitude is set to “Auto”
- CPU performance is set to “HPC”
- Processor C1E state is set to “Disabled”
- Power Technology is set to “Performance”
- Energy Performance is set to “Performance”
- Serial Port A is enabled
- Console Redirection is set to Serial Port A

## Boot Policies

Cisco UCS Boot Policies define the boot devices used by rack-mount servers, and the order that they are attempted to boot from. Cisco HX-Series M5 generation rack-mount servers have their Hyper-V hypervisors installed to an internal M.2 SSD boot drive, therefore they require a unique boot policy defining that the servers should boot from that location. The HyperFlex installer configures a boot policy named “HyperFlex-m5” specifying

boot from the M.2 SSDs, referred to as “Embedded Disk” which is used by the HyperFlex M5 converged nodes, and should not be modified.

Figure 36 details the HyperFlex Boot Policies for Cisco UCS HX-Series M5 generation rack-mount servers.

Figure 36 Cisco UCS M5 Boot Policy

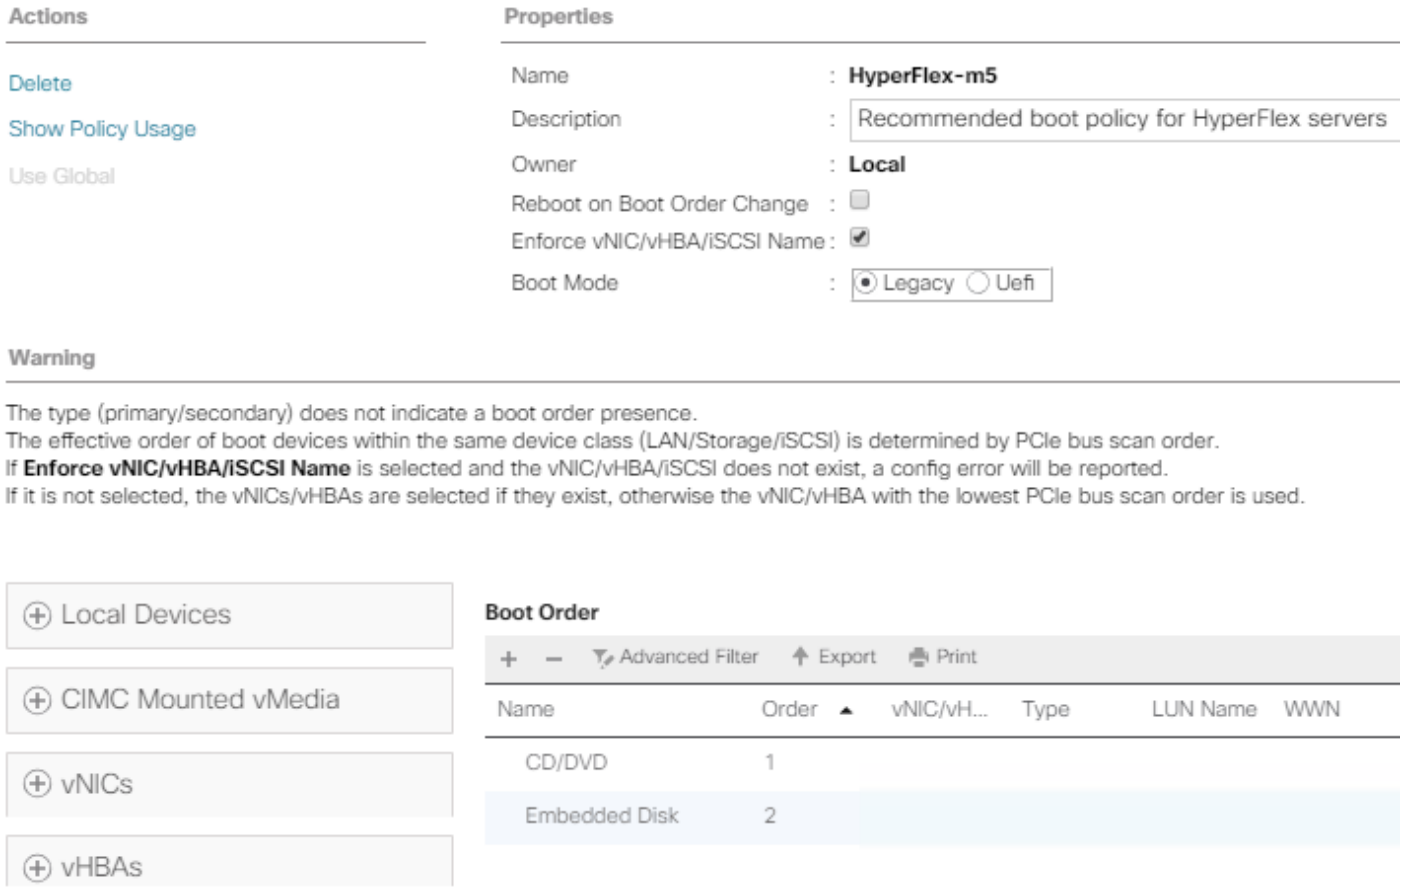


## Host Firmware Packages

Cisco UCS Host Firmware Packages represent one of the most powerful features of the Cisco UCS platform; the ability to control the firmware revision of all the managed blades and rack-mount servers via a policy specified in the service profile. Host Firmware Packages are defined and referenced in the service profiles. Once a service profile is associated to a server, the firmware of all the components defined in the Host Firmware Package are automatically upgraded or downgraded to match the package. The HyperFlex installer creates a Host Firmware Packages named “HyperFlex-m5” which use the simple package definition method, applying firmware revisions to all components that matches a specific Cisco UCS firmware bundle, versus defining the firmware revisions part by part. Figure 37 shows the Host Firmware Package configured by the HyperFlex installer for Cisco HX-Series M5 generation rack-mount servers.

Figure 37 Cisco UCS M5 Host Firmware Package

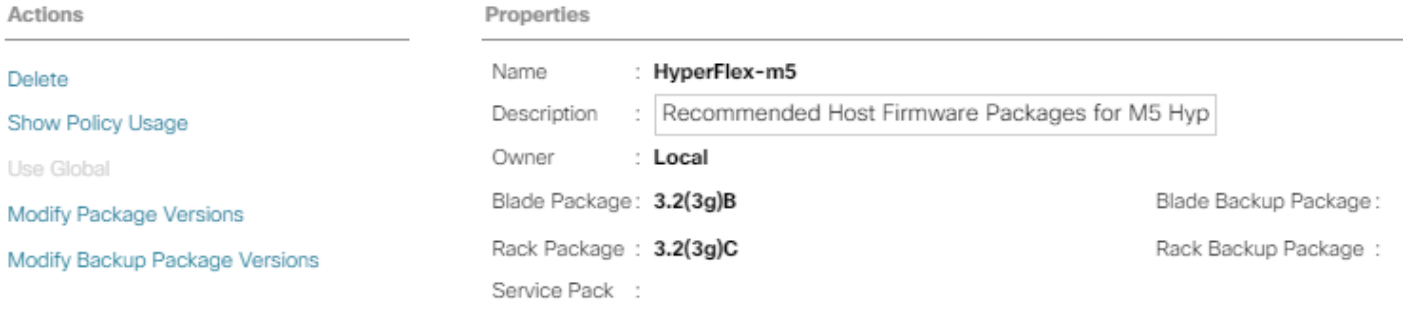


## Local Disk Configuration Policies

Cisco UCS Local Disk Configuration Policies are used to define the configuration of disks installed locally within each blade or rack-mount server, most often to configure Redundant Array of Independent/Inexpensive Disks (RAID levels) when multiple disks are present for data protection. Since HX-Series converged nodes providing storage resources do not require RAID, the HyperFlex installer creates a Local Disk Configuration Policies, named "HyperFlex-m5" which allows any local disk configuration. The policy named "HyperFlex-m5" is used by the service profile template named "hx-nodes-m5", which is for the HyperFlex M5 generation converged servers, and should not be modified.

Figure 38 details the Local Disk Configuration Policies configured by the HyperFlex installer.

Figure 38 Cisco UCS M5 Local Disk Configuration Policy

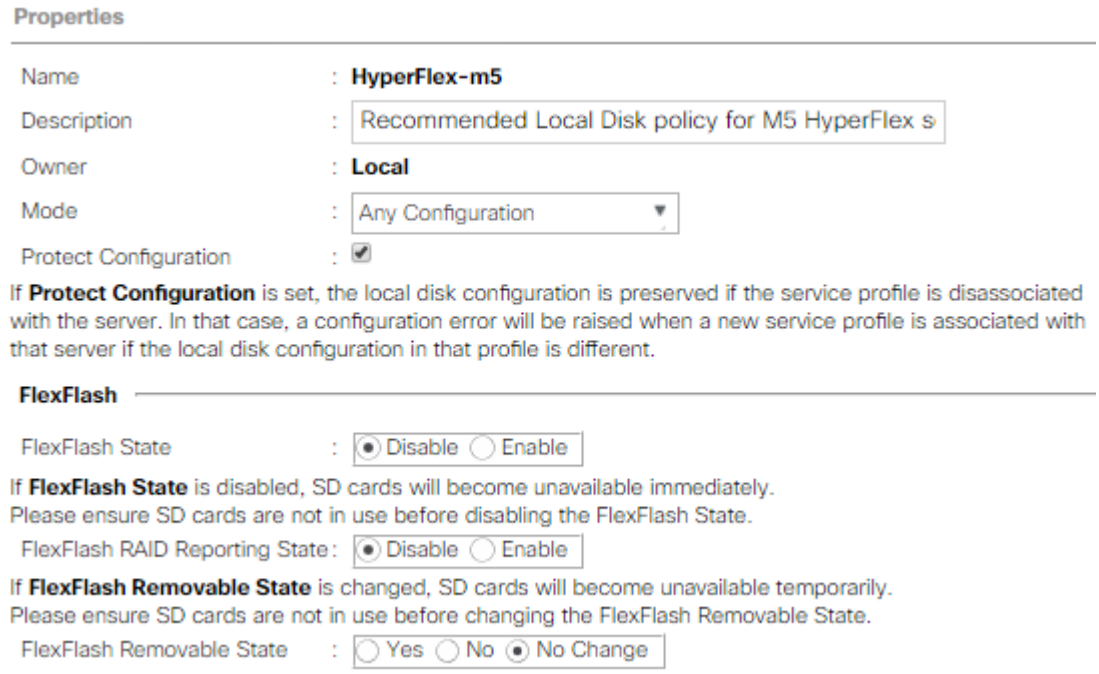


## Maintenance Policies

Cisco UCS Maintenance Policies define the behavior of the attached blades and rack-mount servers when changes are made to the associated service profiles. The default Cisco UCS Maintenance Policy setting is "Immediate" meaning that any change to a service profile that requires a reboot of the physical server will result in an immediate reboot of that server. The Cisco best practice is to use a Maintenance Policy set to "user-ack", which requires a secondary acknowledgement by a user with the appropriate rights within Cisco UCS Manager,

before the server is rebooted to apply the changes. The HyperFlex installer creates a Maintenance Policy named "HyperFlex" with the setting changed to "user-ack". In addition, the On Next Boot setting is enabled, which will automatically apply changes the next time the server is rebooted, without any secondary acknowledgement. Figure 39 details the Maintenance Policy configured by the HyperFlex installer:

**Figure 39 Cisco UCS Maintenance Policy**

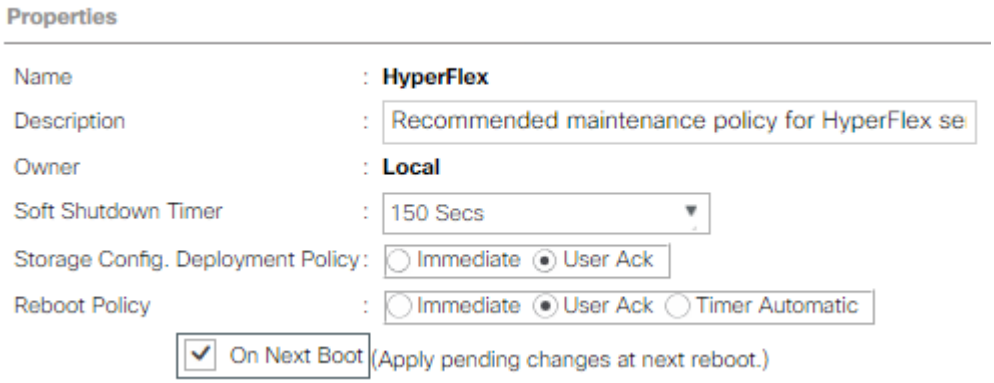


## Power Control Policies

Cisco UCS Power Control Policies allow administrators to set priority values for power application to servers in environments where power supply may be limited, during times when the servers demand more power than is available. The HyperFlex installer creates a Power Control Policy named "HyperFlex" with all power capping disabled, and fans allowed to run at full speed when necessary. Figure 40 details the Power Control Policy configured by the HyperFlex installer.

**Figure 40 Cisco UCS Power Control Policy**

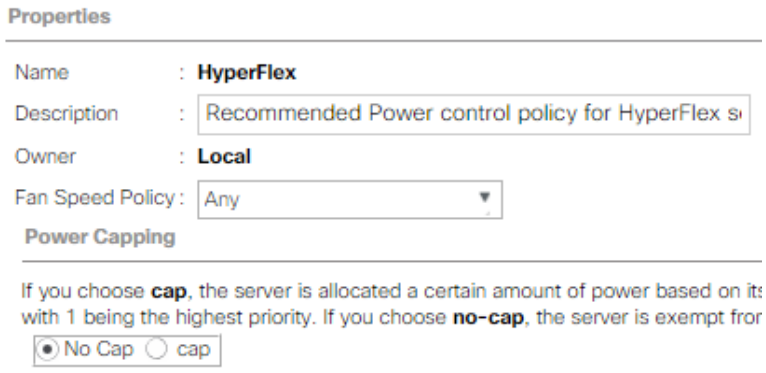


## Scrub Policies

Cisco UCS Scrub Policies are used to scrub, or erase data from local disks, BIOS settings and FlexFlash SD cards. If the policy settings are enabled, the information is wiped when the service profile using the policy is disassociated from the server. The HyperFlex installer creates a Scrub Policy named "HyperFlex" which has all settings disabled, therefore all data on local disks, SD cards and BIOS settings will be preserved if a service profile is disassociated. Figure 41 details the Scrub Policy configured by the HyperFlex installer.

Figure 41 Cisco UCS Scrub Policy

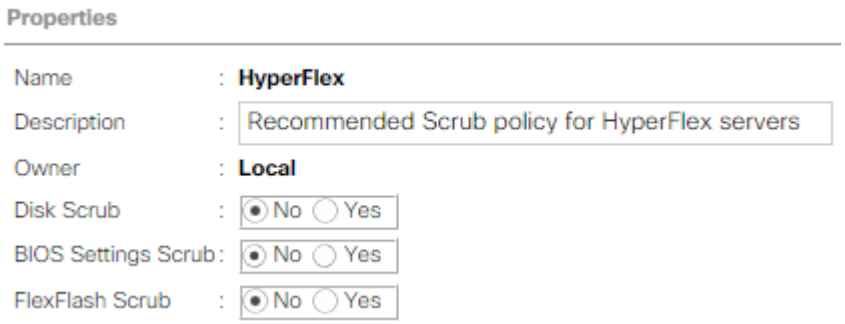


### vMedia Policies

Cisco UCS Virtual Media (vMedia) Policies automate the connection of virtual media files to the remote KVM session of the Cisco UCS blades and rack-mount servers. Using a vMedia policy can speed up installation time by automatically attaching an installation ISO file to the server, without having to manually launch the remote KVM console and connect them one-by-one. The HyperFlex installer creates a vMedia Policy named “HyperFlex” for future use, with no media locations defined.

## Cisco UCS Service Profile Templates

Cisco UCS Manager has a feature to configure service profile templates, which can be used to simplify and speed up configuration efforts when the same configuration needs to be applied to multiple servers. Service profile templates are used to spawn multiple service profile copies to associate with a group of servers, versus configuring the same service profile manually each time it is needed. Service profile templates contain all the configuration elements that make up a service profile, including vNICs, vHBAs, local disk configurations, boot policies, host firmware packages, BIOS policies and more. Templates are created as either initial templates or updating templates. Updating templates retain a link between the parent template and the child object, therefore when changes are made to the template, the changes are propagated to all remaining linked child objects. The HyperFlex installer creates a service profile templates, named “hx-nodes-m5”. The following tables list the service profile template configured by the HyperFlex installer.

Table 19 Cisco UCS Service Profile Template Settings and Values

| Service Profile Template Name: | hx-nodes-m5 |
|---|---|
| Setting | Value |
| UUID Pool | Hardware Default |
| Associated Server Pool | None |
| Maintenance Policy | HyperFlex |
| Management IP Address Policy | hx-ext-mgmt |
| Local Disk Configuration Policy | HyperFlex-m5 |
| LAN Connectivity Policy | HyperFlex |

| Service Profile Template Name: | hx-nodes-m5 |
|---|---|
| Setting | Value |
| Boot Policy | HyperFlex-m5 |
| BIOS Policy | HyperFlex-m5 |
| Firmware Policy | HyperFlex-m5 |
| Power Control Policy | HyperFlex |
| Scrub Policy | HyperFlex |
| Serial over LAN Policy | HyperFlex |
| vMedia Policy | Not defined |

## Virtual Networking Design

The Cisco HyperFlex system has a pre-defined virtual network design at the Hyper-V 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 Hyper-V 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 virtual machines 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 Hyper-V/Hyper-V
- 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 via 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 Hyper-V/Hyper-V
- 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 Hyper-V/Hyper-V
- live-migration: 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 Hyper-V/Hyper-V

The following tables and figures provide more details about the Hyper-V virtual networking design as built by the HyperFlex installer:

Table 20 Table Hyper-V Host Virtual Switch Configuration

| Virtual Switch | Port Groups | Active vmnic(s) | Passive vmnic(s) | VLAN IDs | Jumbo |
|---|---|---|---|---|---|
| vswitch-hx-inband-mgmt | Management Network<br><br>Storage Controller Management Network | vmnic0 | vmnic1 | hx-inband-mgmt | no |
| vswitch-hx-storage-data | Storage Controller Data Network<br><br>Storage Hypervisor Data Network | vmnic3 | vmnic2 | hx-storage-data | yes |
| vswitch-hx-vm-network | none | vmnic4,vmnic5 | none | vm-network | no |
| Live-migration | none | vmnic6 | vmnic7 | 33 | yes |

Figure 42 SCVMM Network Design

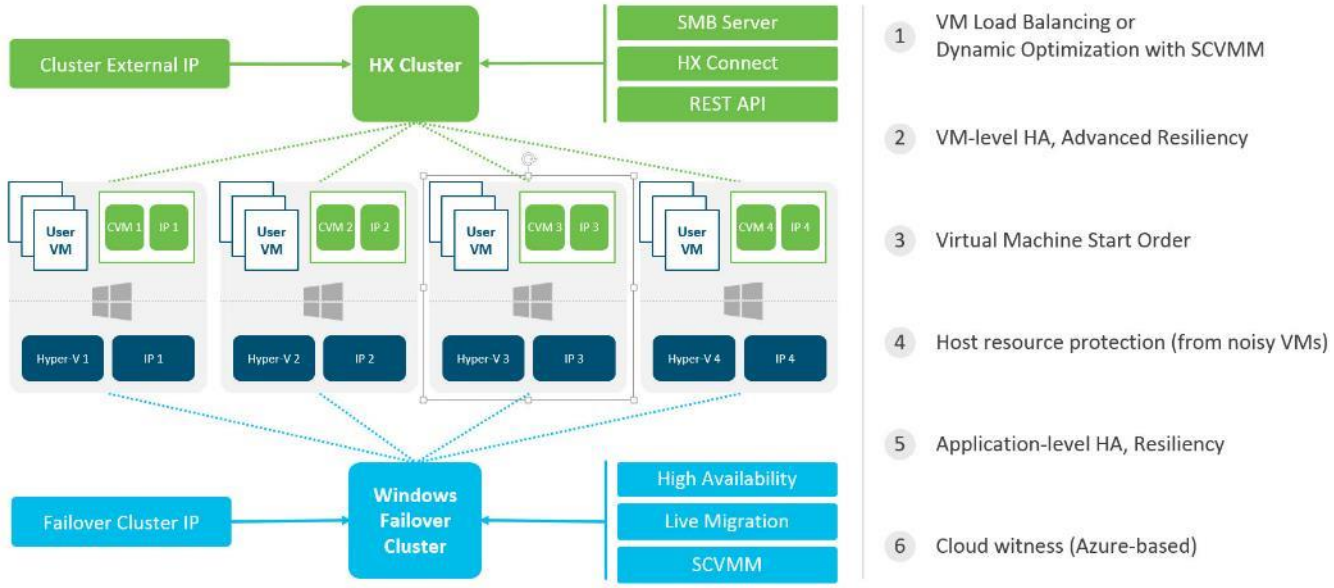


## Storage Platform Controller Virtual Machines

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 virtual machines are deployed as a Hyper-V agent, which is similar in concept to that of a Linux or Windows service. Hyper-V agents are tied to a specific host, they start and stop along with the Hyper-V hypervisor, and the system is not considered to be online and ready until both the hypervisor and the agents have started. Each Hyper-V hypervisor host has a single Hyper-V agent deployed, which is the controller virtual machine for that node, and it

cannot be moved or migrated to another host. The collective Hyper-V agents are managed via a Hyper-V agency in the Hyper-V 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 Hyper-V agents to the agency, therefore the Hyper-V hypervisors nor SCVMM 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 via a plugin installed to the SCVMM server or appliance managing the Hyper-V cluster. The plugin communicates directly with the controller virtual machines to display the information requested.

### 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 a Hyper-V 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 via SR-IOV. The configuration details of the models are described in the following subsections.

### Cisco HyperFlex Datastores

The new HyperFlex cluster has no default datastores configured for virtual machine storage, therefore the datastores must be created using the HyperFlex Connect GUI. Overall space consumption in the HyperFlex clustered filesystem is optimized by the default deduplication and compression features.

**Figure 43 Datastore Example**

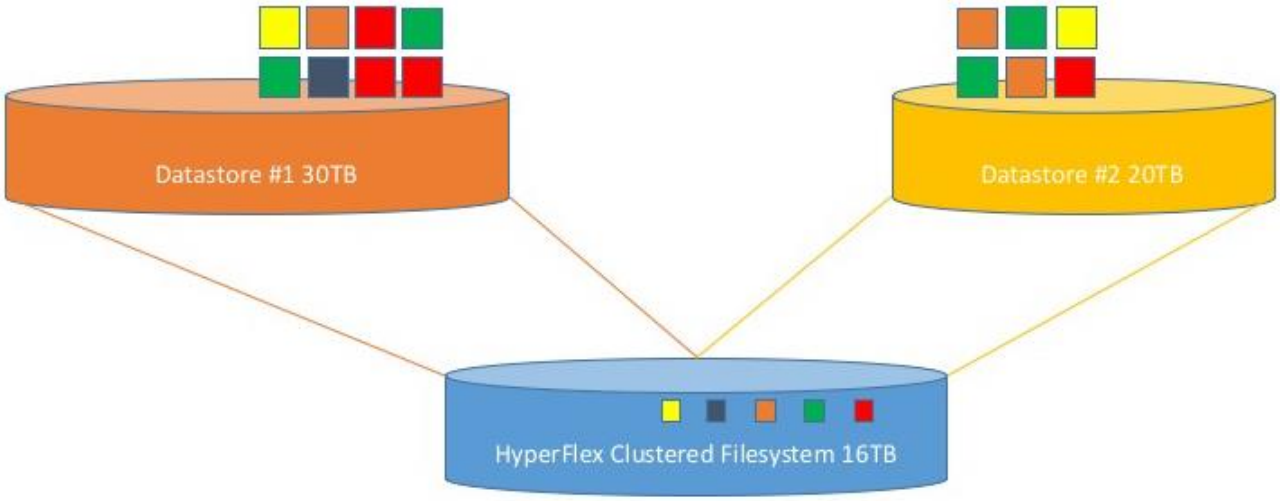


### CPU Resource Reservations

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 Hyper-V hypervisor host are being heavily consumed by the guest virtual machines. Table 21 lists the CPU resource reservation of the storage controller virtual machines.

Table 21 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 Hyper-V hypervisor host are being heavily consumed by the guest virtual machines.

Table 22 lists the memory resource reservation of the storage controller virtual machines.

Table 22 Controller Virtual Machine Memory Reservations

| Server Model | Amount of Guest Memory | Reserve All Guest Memory |
|---|---|---|
| HX220c-M5<br>HXAF220c-M5 | 48 GB | Yes |
| HX240c-M5SX<br>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 44 illustrates the configuration topology for this solution.

Figure 44 Configuration Topology for Scalable Citrix Virtual Desktops 1909 Workload with HyperFlex

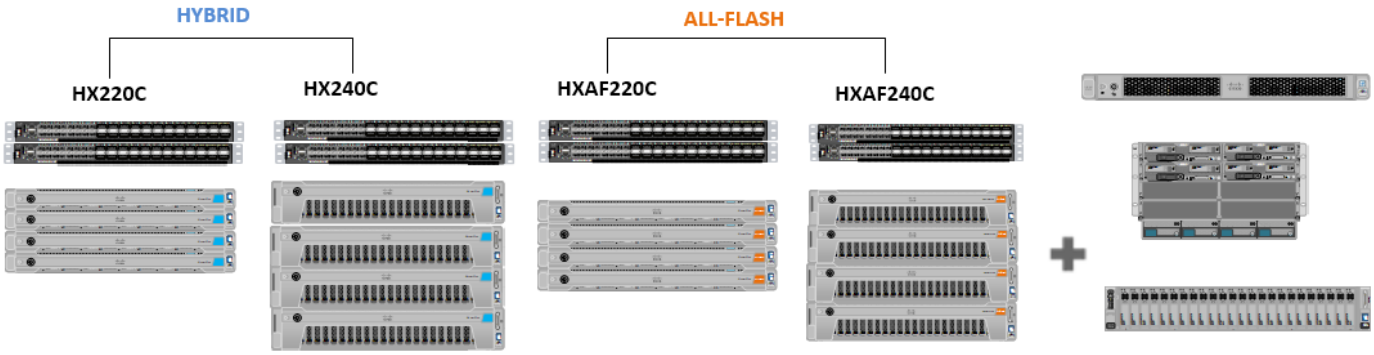

# Installation

Installing the Cisco HyperFlex system on Hyper-V is primarily done via a deployable HyperFlex installer virtual machine available for download at cisco.com as a vhdx file. The installer virtual machine performs the Cisco UCS configuration work, the configuration of Hyper-V on the HyperFlex hosts, the installation of the HyperFlex HX Data Platform software and creation of the HyperFlex cluster. Because this simplified installation method has been developed by Cisco, this CVD will not provide detailed manual steps for the configuration of all the elements that are handled by the installer. The following sections will guide you through the prerequisites and manual steps needed prior to using the HyperFlex installer, how to utilize the HyperFlex Installer, and how to perform the remaining post-installation tasks.

## Prerequisites

Prior to beginning the installation activities, it is important to gather the following information:

### IP Addressing

To install the HX Data Platform, an OVF installer appliance must be deployed on a separate virtualization host, which is not a member of the HyperFlex cluster. The HyperFlex installer requires one IP address on the management network and the HX installer appliance IP address must be able to communicate with Cisco UCS Manager, Hyper-V management IP addresses on the HX hosts, Windows Active Directory and DNS server and any management server IP addresses from where the Windows failover cluster will be managed.

Additional IP addresses for the Cisco HyperFlex system need to be allocated from the appropriate subnets and VLANs to be used. IP addresses that are used by the system fall into the following groups:

- Cisco UCS Manager: These addresses are used and assigned by Cisco UCS Manager. Three IP addresses are used by Cisco UCS Manager; one address is assigned to each Cisco UCS Fabric Interconnect, and the third IP address is a roaming address for management of the Cisco UCS cluster. In addition, at least one IP address per Cisco UCS blade or HX-series rack-mount server is required for the hx-ext-mgmt IP address pool, which are assigned to the CIMC interface of the physical servers. Since these management addresses are assigned from a pool, they need to be provided in a contiguous block of addresses. These addresses must all be in the same subnet.
- HyperFlex and Hyper-V Management: These addresses are used to manage the Hyper-V hypervisor hosts, and the HyperFlex Storage Platform Controller virtual machines. Two IP addresses per node in the HyperFlex cluster are required from the same subnet, an additional IP address is required for roaming HyperFlex cluster management interface and another additional IP address is required for the Windows failover cluster. These addresses can be assigned from the same subnet at the Cisco UCS Manager addresses, or they may be separate.
- HyperFlex Storage: These addresses are used by the HyperFlex Storage Platform Controller virtual machines, and the Hyper-V hypervisor hosts, for sending and receiving data to/from the HX Distributed Data Platform Filesystem. Two IP addresses per node in the HyperFlex cluster are required from the same subnet, and a single additional IP address is needed as the roaming HyperFlex cluster storage interface. It is recommended to provision a subnet that is not used in the network for other purposes, and it is also possible to use non-routable IP address ranges for these interfaces. Finally, if the Cisco UCS domain is going to contain multiple HyperFlex clusters, it is recommended to use a different subnet and VLAN ID for the HyperFlex storage traffic for each cluster. This is a safer method, guaranteeing that storage traffic from multiple clusters cannot intermix.

- Live Migration: These IP addresses are used by the Hyper-V hypervisor hosts on interfaces to enable live migration capabilities. One or more IP addresses per node in the HyperFlex cluster are required from the same subnet.

The following tables provide space to input the required IP addresses for the installation of an 8-node standard HyperFlex cluster by listing the addresses required, plus an example IP configuration.

**Table cells shaded in black do not require an IP address.**

Table 23 HyperFlex Standard Cluster IP Addressing

| Address Group: | UCS Management | HyperFlex and Hyper-V Management | | HyperFlex Storage | | Live Migration |
|---|---|---|---|---|---|---|
| VLAN ID: | | | | | | |
| Subnet: | | | | | | |
| Subnet Mask: | | | | | | |
| Gateway: | | | | | | |
| **Device** | **UCS Management Addresses** | **Hyper-V Management Interfaces** | **Storage Controller Virtual Machine Management Interfaces** | **Hyper-V Hypervisor Storage Interfaces** | **Storage Controller Virtual Machine Storage Interfaces** | **Live Migration Interfaces** |
| Fabric Interconnect A | | | | | | |
| Fabric Interconnect B | | | | | | |
| UCS Manager | | | | | | |
| HyperFlex Cluster | | | | | | |
| Windows Failover Cluster | | | | | | |
| HyperFlex Node #1 | | | | | | |
| HyperFlex Node #2 | | | | | | |
| HyperFlex Node #3 | | | | | | |
| HyperFlex Node #4 | | | | | | |
| HyperFlex Node #5 | | | | | | |
| HyperFlex Node #6 | | | | | | |
| HyperFlex Node #7 | | | | | | |
| HyperFlex Node #8 | | | | | | |

HyperFlex extended clusters are also addressed similarly to a standard cluster, however the compute-only nodes do not require any IP addresses for the Storage Controller virtual machines, as shown in Table 24 .

Table 24 HyperFlex Extended Cluster IP Addressing

| Address Group: | UCS Management | HyperFlex and Hyper-V Management | | HyperFlex Storage | | Live Migration |
|---|---|---|---|---|---|---|
| VLAN ID: | | | | | | |
| Subnet: | | | | | | |
| Subnet Mask: | | | | | | |
| Gateway: | | | | | | |
| Device | UCS Management Addresses | Hyper-V Management Interfaces | Storage Controller Virtual Machine Management Interfaces | Hyper-V Hypervisor Storage Interfaces | Storage Controller Virtual Machine Storage Interfaces | Live Migration Interfaces |
| Fabric Interconnect A | | | | | | |
| Fabric Interconnect B | | | | | | |
| UCS Manager | | | | | | |
| HyperFlex Cluster | | | | | | |
| Windows Failover Cluster | | | | | | |
| HyperFlex Node #1 | | | | | | |
| HyperFlex Node #2 | | | | | | |
| HyperFlex Node #3 | | | | | | |
| HyperFlex Node #4 | | | | | | |
| HyperFlex Node #5 | | | | | | |
| HyperFlex Node #6 | | | | | | |
| HyperFlex Node #7 | | | | | | |
| HyperFlex Node #8 | | | | | | |
| Compute Node #1 | | | | | | |
| Compute Node #2 | | | | | | |
| Compute Node #3 | | | | | | |
| Compute Node #4 | | | | | | |
| Compute Node #5 | | | | | | |
| Compute Node #6 | | | | | | |
| Compute Node #7 | | | | | | |
| Compute Node #8 | | | | | | |

Table 25 HyperFlex Standard Cluster Example IP Addressing for a 4-Node Cluster

| Address Group: | UCS Management | HyperFlex and Hyper-V Management | | HyperFlex Storage | | Live Migration |
|---|---|---|---|---|---|---|
| VLAN ID: | 130 | 30 | | 101 | | 33 |
| Subnet: | 10.65.121.0 | 10.104.252.0 | | 192.168.11.0 | | 192.168.73.0 |
| Subnet Mask: | 255.255.255.0 | 255.255.255.0 | | 255.255.255.0 | | 255.255.255.0 |
| Gateway: | 10.65.121.1 | 10.104.252.1 | | | | |
| Device | UCS Management Addresses | UCS Management Addresses | Hyper-V Management Interfaces | Hyper-V Hypervisor Storage Interfaces | Storage Controller Virtual Machine Storage Interfaces | Live Migration Interfaces |
| Fabric Interconnect A | 10.65.121.241 | | | | | |
| Fabric Interconnect B | 10.65.121.242 | | | | | |
| UCS Manager | 10.65.121.240 | | | | | |
| HyperFlex Cluster | | | 10.104.252.35 | | 192.168.11.35 | |
| Windows Failover Cluster | | 10.104.252.36 | | | | |
| HyperFlex Node #1 | 10.104.252.11 | 10.104.252.19 | 10.104.252.27 | 192.168.11.11 | 192.168.11.19 | 192.168.73.11 |
| HyperFlex Node #2 | 10.104.252.12 | 10.104.252.20 | 10.104.252.28 | 192.168.11.12 | 192.168.11.20 | 192.168.73.12 |
| HyperFlex Node #3 | 10.104.252.13 | 10.104.252.21 | 10.104.252.29 | 192.168.11.13 | 192.168.11.21 | 192.168.73.13 |
| HyperFlex Node #4 | 10.104.252.14 | 10.104.252.22 | 10.104.252.30 | 192.168.11.14 | 192.168.11.22 | 192.168.73.14 |
| HyperFlex Node #5 | 10.104.252.15 | 10.104.252.23 | 10.104.252.31 | 192.168.11.15 | 192.168.11.23 | 192.168.73.15 |
| HyperFlex Node #6 | 10.104.252.16 | 10.104.252.24 | 10.104.252.32 | 192.168.11.16 | 192.168.11.24 | 192.168.73.16 |
| HyperFlex Node #7 | 10.104.252.17 | 10.104.252.25 | 10.104.252.33 | 192.168.11.17 | 192.168.11.25 | 192.168.73.17 |
| HyperFlex Node #8 | 10.104.252.18 | 10.104.252.26 | 10.104.252.34 | 192.168.11.18 | 192.168.11.26 | 192.168.73.18 |

IP addresses for Cisco UCS Management, including HyperFlex and Hyper-V Management, can come from the same subnet, or can be separate subnets, if the HyperFlex installer can reach them both.

## DHCP versus Static IP

By default, the HX installation will assign a static IP address to the management interface of the Hyper-V servers. Using Dynamic Host Configuration Protocol (DHCP) for automatic IP address assignment in not recommended.

## Configure the Active Directory for Constrained Delegation

A Windows 2008 R2 and above forest/domain level Active Directory (AD) is required for the successful installation and operation of Cisco HyperFlex system with Hyper-V.

The steps in this topic must be completed to enable constrained delegation. Constrained delegation is used to join computers to the Active Directory. You provide constrained delegation information through the HX Data Platform Installer. Constrained delegation uses a service account, which is created manually. This service account is used to then log in to Active Directory, join the computers, and perform the authentications from the HyperFlex Storage Controller virtual machine.

The Active Directory computer accounts applied to every node in the HyperFlex cluster include:

- Hyper-V host
- HyperFlex Storage Controller virtual machine
- Hyper-V host Cluster namespace
- Server Message Block (SMB) Share namespace for the HyperFlex cluster

To configure constrained delegation, follow these steps:

1. Create an hxadmin domain user account as HX service account.
2. Create an Organization Unit (OU) in Active Directory (AD), for example, HyperFlex:
    a. Use the Active Directory Users and Computers management tool to create the OU. Select View > Advanced Features to enable advance features. Select the OU that you created. For example, HyperFlex > Properties > Attribute Editor.
    b. Find the distinguished name attribute in the OU, and record the information as this will be required in the Constrained Delegation wizard of the HX Data Platform Installer wizard. The values will look like this: OU=HyperFlex,DC=contoso,DC=com.
    c. Use the Get-ADOrganizationalUnit cmdlet to get an organizational unit (OU) object or to perform a search to get multiple OUs.

        *Get-ADOrganizationalUnit -Filter 'Name -like "Hyp*" ' | Format-Table Name, DistinguishedName*

**Figure 45 PowerShell Get-ADOrganizationalUnit**

```
PS C:\Users\Administrator> Get-ADOrganizationalUnit -Filter 'Name -like "Hyp*"' | Format-Table Name, DistinguishedName

Name      DistinguishedName
----      -----------------
HyperFlex OU=HyperFlex,DC=HXHVDOM,DC=LOCAL
```

3. Use Active Directory Users and Computers management tool to grant full permissions for the hxadmin user for the newly created OU. Make sure that Advanced features are enabled. If not, go back to step 2.
    a. Select the OU that you created. For example, HyperFlex > Properties > Security > Advance.
    b. Click Change Owner and choose your hxadmin user.

c. Click Add in the Advanced view.

d. Select the principal and choose the hxadmin user. Choose Full Control and click OK.

Figure 46 Change Ownership of the AD OU

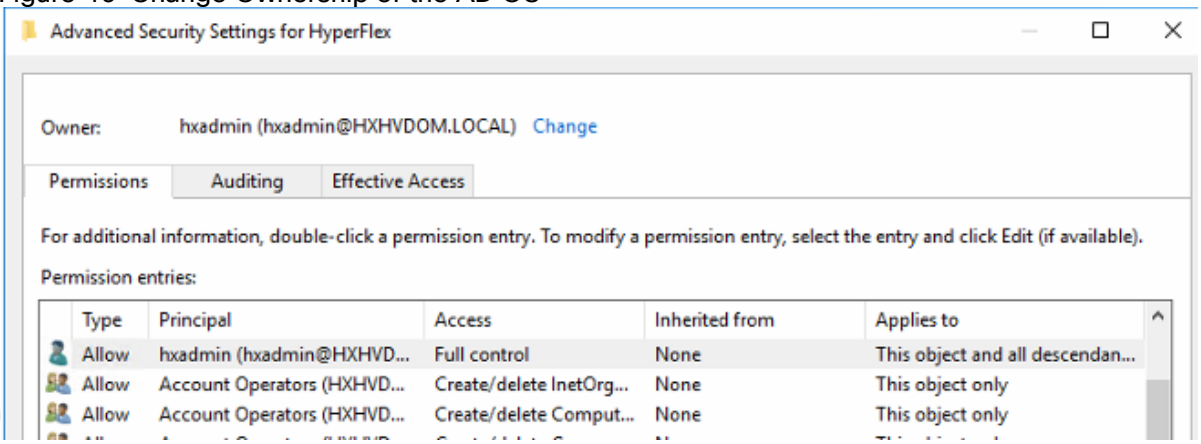


## Prepopulate AD DNS with Records

The AD integrated DNS server is also required to resolve Fully Qualified Domain Names (FQDN).

To create DNS records, follow these steps:

1. Create a record and reverse the PTR records for the listed devices to avoid installation failures:
   - For each Hyper-V hosts' management and storage interfaces
   - For each Storage Controller Nodes' management and storage interfaces
   - HX Cluster CIP
   - SMB Access point
   - Windows Failover Cluster IP

**Standalone and non-Windows DNS servers are not supported.**

The following tables provide a place to input the required DNS information for the installation and lists the information required and provide an example configuration.

Table 26 DNS Server Information

| Item | Value |
|---|---|
| DNS Server #1 | |
| AD DNS Domain | |
| UCS Domain Name | |

| Item | Value |
|---|---|
| HX Server #1 Name | |
| HX Server #2 Name | |
| HX Server #3 Name | |
| HX Server #4 Name | |
| HX Server #5 Name | |
| HX Server #6 Name | |
| HX Server #7 Name | |
| HX Server #8 Name | |

**Table 27 DNS Server Example Information**

| Item | Value |
|---|---|
| DNS Server #1 | 10.29.149.222 |
| DNS Server #2 | |
| AD DNS Domain | hxhvdom.local |
| SMTP Server Name | outbound. hxhvdom.local |
| UCS Domain Name | HXHV-FI |
| HX Server #1 Name | hxhv1. hxhvdom.local |
| HX Server #2 Name | hxhv2. hxhvdom.local |
| HX Server #3 Name | hxhv3. hxhvdom.local |
| HX Server #4 Name | Hxhv4. hxhvdom.local |
| HX Server #5 Name | Hxhv5. hxhvdom.local |
| HX Server #6 Name | Hxhv6. hxhvdom.local |

| Item | Value |
| --- | --- |
| HX Server #7 Name | Hxhv7. hxhvdom.local |
| HX Server #8 Name | Hxhv8. hxhvdom.local |

## NTP

Consistent time clock synchronization is required across the components of the HyperFlex system, provided by reliable NTP servers, accessible for querying in the Cisco UCS Management network group, and the HyperFlex and Hyper-V Management group. NTP is used by Cisco UCS Manager, the Hyper-V hypervisor hosts, and the HyperFlex Storage Platform Controller virtual machines. For HyperFlex System with Hyper-V, AD Domain Controller IP or domain name is required to be used as reliable NTP source for consistent time clock synchronization.

## VLANs

Prior to the installation, the required VLAN IDs need to be documented, and created in the upstream network if necessary. At a minimum, there are 4 VLANs that need to be trunked to the Cisco UCS Fabric Interconnects that comprise the HyperFlex system; a VLAN for the HyperFlex and Hyper-V Management group, a VLAN for the HyperFlex Storage group, a VLAN for the Live Migration group, and at least one VLAN for the guest virtual machine traffic. The VLAN IDs must be supplied during the HyperFlex Cisco UCS configuration step, and the VLAN names can optionally be customized.

The following tables provide a place to input the required VLAN information and provide an example configuration.

**Table 28 VLAN Information**

| Name | ID |
| --- | --- |
| <<hx-inband-mgmt>> | |
| <<hx-storage-data>> | |
| <<hx-vm-network>> | |
| <<hx-livemigrate>> | |

**Table 29 VLAN Example Information**

| Name | ID |
| --- | --- |
| hx-inband-mgmt | 30 |
| hx-storage-data | 101 |
| vm-network | 34 |

| Name | ID |
|---|---|
| hx-livemigrate | 33 |

## Network Uplinks

The Cisco UCS uplink connectivity design needs to be finalized prior to beginning the installation. One of the early manual tasks to be completed is to configure the Cisco UCS network uplinks and verify their operation, prior to beginning the HyperFlex installation steps. Refer to the network uplink design possibilities in the Network Design section.

The following tables provide a place to input the required network uplink information for the installation and provide an example configuration:

**Table 30 Network Uplink Configuration**

| Fabric Interconnect Port | | Port Channel | Port Channel Type | Port Channel ID | Port Channel Name |
|---|---|---|---|---|---|
| A | | ☐ Yes ☐ No | ☐ LACP ☐ vPC | | |
| | | ☐ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |
| B | | ☐ Yes ☐ No | ☐ LACP ☐ vPC | | |
| | | ☐ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |

**Table 31 Network Uplink Example Configuration**

| Fabric Interconnect Port | | Port Channel | Port Channel Type | Port Channel ID | Port Channel Name |
|---|---|---|---|---|---|
| A | 1/25 | ☒ Yes ☐ No | ☐ LACP ☒ vPC | 10 | vpc-10 |
| | 1/26 | ☒ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |

| Fabric Interconnect Port | | Port Channel | Port Channel Type | Port Channel ID | Port Channel Name |
|---|---|---|---|---|---|
| B | 1/25 | ☒ Yes ☐ No | ☐ LACP ☒ vPC | 20 | vpc-20 |
| | 1/26 | ☒ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |
| | | ☐ Yes ☐ No | | | |

## Usernames and Passwords

Several usernames and passwords need to be defined or known as part of the HyperFlex installation process. The following tables provide a place to input the required username and password information and provide an example configuration.

**Table 32 Usernames and Passwords**

| Account | Username | Password |
|---|---|---|
| HX Installer Administrator | root | <<hx_install_root_pw>> |
| UCS Administrator | admin | <<ucs_admin_pw>> |
| Hyper-V Local Administrator | root | <<hyperv_local_pw>> |
| HyperFlex Administrator | root | <<hx_admin_pw>> |
| AD Domain Admin or Service Account | <<administrator>> | <<ad_admin_pw>> |

**Table 33 Example Usernames and Passwords**

| Account | Username | Password |
|---|---|---|
| HX Installer Administrator | root | Cisco123 |
| UCS Administrator | admin | Cisco123 |
| Hyper-V Local Administrator | root | Cisco123 |
| HyperFlex Administrator | root | Cisco123!! |
| AD Domain Admin or Service Account | administrator@domain.local | !QAZ2wsx |

## Physical Installation

Install the Fabric Interconnects, the HX-Series rack-mount servers according to their corresponding hardware installation guides listed below:

- Cisco UCS 6200 Series Fabric Interconnect Hardware Installation Guide
- Cisco UCS 6300 Series Fabric Interconnect Hardware Installation Guide
- Cisco HX220c M5 HyperFlex Node Installation Guide (Hybrid and All-Flash Models)
- Cisco HX240c M5 HyperFlex Node (Hybrid and All-Flash Models) Installation Guide

## 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 list 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 47 shows a cabling diagram for a Citrix Virtual Desktops configuration using the Cisco Nexus 9000 and Cisco UCS Fabric Interconnect.

**Table 34 Cisco Nexus 93108YC-Cabling Information**

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---|---|---|---|---|
| Cisco Nexus 93108YC A | Eth1/1 | 10GbE | Cisco Nexus 93108YC B | Eth1/1 |
| | Eth1/2 | 10GbE | Cisco Nexus 93108YC B | Eth1/2 |
| | Eth1/3 | 10GbE | Cisco UCS fabric interconnect A | Eth1/13 |
| | Eth1/4 | 10GbE | Cisco UCS fabric interconnect A | Eth1/14 |
| | Eth1/5 | 10GbE | Cisco UCS fabric interconnect B | Eth1/13 |
| | Eth1/6 | 10GbE | Cisco UCS fabric interconnect B | Eth1/14 |
| | Eth1/25 | 10GbE | Infra-host-01 | Port01 |

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---|---|---|---|---|
| | Eth1/26 | 10GbE | Infra-host-02 | Port01 |
| | Eth1/27 | 10GbE | Launcher-host-01 | Port01 |
| | Eth1/28 | 10GbE | Launcher-host-02 | Port01 |
| | Eth1/29 | 10GbE | Launcher-host-03 | Port01 |
| | Eth1/30 | 10GbE | Launcher-host-04 | Port01 |
| | MGMT0 | GbE | GbE management switch | Any |

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

**Table 35 Cisco Nexus 93108YC-B Cabling Information**

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---|---|---|---|---|
| Cisco Nexus 93108YC B | Eth1/1 | 10GbE | Cisco Nexus 93108YC A | Eth1/1 |
| | Eth1/2 | 10GbE | Cisco Nexus 93108YC A | Eth1/2 |
| | Eth1/3 | 10GbE | Cisco UCS fabric interconnect A | Eth1/15 |
| | Eth1/4 | 10GbE | Cisco UCS fabric interconnect A | Eth1/16 |
| | Eth1/5 | 10GbE | Cisco UCS fabric interconnect B | Eth1/15 |
| | Eth1/6 | 40GbE | Cisco UCS fabric interconnect B | Eth1/16 |
| | Eth1/25 | 10GbE | Infra-host-01 | Port02 |
| | Eth1/26 | 10GbE | Infra-host-02 | Port02 |
| | Eth1/27 | 10GbE | Launcher-host-01 | Port02 |
| | Eth1/28 | 10GbE | Launcher-host-02 | Port02 |
| | Eth1/29 | 10GbE | Launcher-host-03 | Port02 |
| | Eth1/30 | 10GbE | Launcher-host-04 | Port02 |
| | MGMT0 | GbE | GbE management switch | Any |

Table 36 Cisco UCS Fabric Interconnect A Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---|---|---|---|---|
| Cisco UCS fabric interconnect A | Eth1/13 | 10GbE | Cisco Nexus 93108YC A | Eth1/3 |
| | Eth1/14 | 10GbE | Cisco Nexus 93108YC A | Eth1/4 |
| | Eth1/15 | 10GbE | Cisco Nexus 93108YC B | Eth1/5 |
| | Eth1/16 | 10 GbE | Cisco Nexus 93108YC B | Eth 1/6 |
| | MGMT0 | GbE | GbE management switch | Any |
| | L1 | GbE | Cisco UCS fabric interconnect B | L1 |
| | L2 | GbE | Cisco UCS fabric interconnect B | L2 |

Table 37 Cisco UCS Fabric Interconnect B Cabling Information

| Local Device | Local Port | Connection | Remote Device | Remote Port |
|---|---|---|---|---|
| Cisco UCS fabric interconnect B | Eth1/13 | 10GbE | Cisco Nexus 93108YC B | Eth1/3 |
| | Eth1/14 | 10GbE | Cisco Nexus 93108YC B | Eth1/4 |
| | Eth1/15 | 10GbE | Cisco Nexus 93108YC A | Eth1/5 |
| | Eth1/16 | 10GbE | Cisco Nexus 93108YC A | Eth 1/6 |
| | MGMT0 | GbE | GbE management switch | Any |
| | L1 | GbE | Cisco UCS fabric interconnect A | L1 |
| | L2 | GbE | Cisco UCS fabric interconnect A | L2 |

Figure 47 Cable Connectivity Between Cisco Nexus 93108YC A and B to Cisco UCS 6248 Fabric A and B

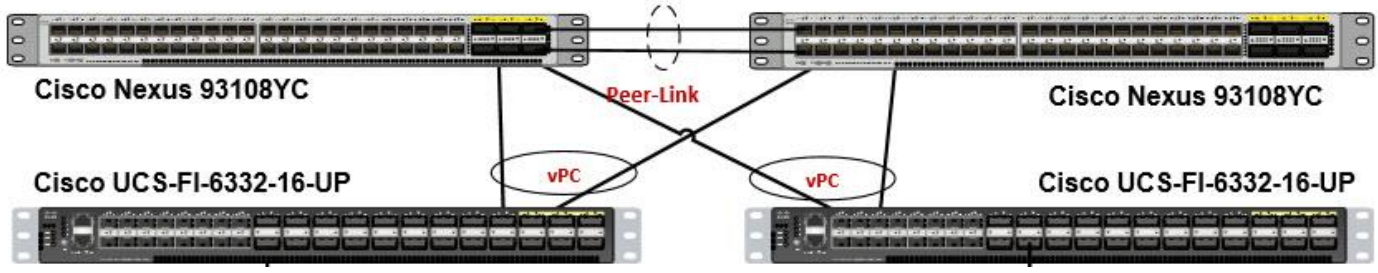

# Cisco Unified Computing System Configuration

This section describes the steps to initialize and configure the Cisco UCS Fabric Interconnects and how to prepare them for the HyperFlex installation.

## Cisco UCS Fabric Interconnect A

To configure Fabric Interconnect A, follow these steps:

1. Make sure the Fabric Interconnect cabling is properly connected, including the L1 and L2 cluster links, and power the Fabric Interconnects on by inserting the power cords.

2. Connect to the console port on the first Fabric Interconnect, which will be designated as the A fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.

3. Start your terminal emulator software.

4. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, and 1 stop bit.

5. Open the connection just created. You may have to press ENTER to see the first prompt.

6. Configure the first Fabric Interconnect, using the following example as a guideline:

```
         ---- Basic System Configuration Dialog ----
  This setup utility will guide you through the basic configuration of
  the system. Only minimal configuration including IP connectivity to
  the Fabric interconnect and its clustering mode is performed through these steps.


  Type Ctrl-C at any time to abort configuration and reboot system.
  To back track or make modifications to already entered values, complete input till
end of section and answer no when prompted to apply configuration.
  Enter the configuration method. (console/gui) ? console
  Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
  You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
  Enforce strong password? (y/n) [y]: y
  Enter the password for "admin":
  Confirm the password for "admin":
  Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no)
[n]: yes
  Enter the switch fabric (A/B) []: A
  Enter the system name:  HXHV-FI-A
  Physical Switch Mgmt0 IP address : 10.29.149.203
  Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
  IPv4 address of the default gateway : 10.29.149.1
  Cluster IPv4 address : 10.29.149.205
```

```
  Configure the DNS Server IP address? (yes/no) [n]: yes
    DNS IP address : 10.29.149.222
  Configure the default domain name? (yes/no) [n]: yes
    Default domain name : hxhvdom.local
  Join centralized management environment (UCS Central)? (yes/no) [n]: no
  Following configurations will be applied:
    Switch Fabric=A
    System Name=HXHV-FI-A
    Enforced Strong Password=no
    Physical Switch Mgmt0 IP Address=10.29.149.203
    Physical Switch Mgmt0 IP Netmask=255.255.255.0
    Default Gateway=10.29.149.1
    Ipv6 value=0
    DNS Server=10.29.149.222
    Domain Name=hx.lab.cisco.com
    Cluster Enabled=yes
    Cluster IP Address=10.29.149.205
    NOTE: Cluster IP will be configured only after both Fabric Interconnects are ini-
tialized
 Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
  Applying configuration. Please wait.
  Configuration file - Ok
```

## Cisco UCS Fabric Interconnect B

To configure Fabric Interconnect B, follow these steps:

1. Connect to the console port on the first Fabric Interconnect, which will be designated as the B fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.

2. Start your terminal emulator software.

3. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, and 1 stop bit.

4. Open the connection just created. You may have to press ENTER to see the first prompt.

5. Configure the second Fabric Interconnect, using the following example as a guideline:

```
         ---- Basic System Configuration Dialog ----
  This setup utility will guide you through the basic configuration of
  the system. Only minimal configuration including IP connectivity to
  the Fabric interconnect and its clustering mode is performed through these steps.
  Type Ctrl-C at any time to abort configuration and reboot system.
  To back track or make modifications to already entered values,
```

```
  complete input till end of section and answer no when prompted
  to apply configuration.

  Enter the configuration method. (console/gui) ? console
  Installer has detected the presence of a peer Fabric interconnect. This Fabric inter-
connect will be added to the cluster. Continue (y/n) ? y
  Enter the admin password of the peer Fabric interconnect:
    Connecting to peer Fabric interconnect... done
    Retrieving config from peer Fabric interconnect... done
    Peer Fabric interconnect Mgmt0 IPv4 Address: 10.29.149.204
    Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
    Cluster IPv4 address          : 10.29.149.205
    Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0
IPv4 Address
  Physical Switch Mgmt0 IP address : 10.29.149.204
  Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
  Applying configuration. Please wait.
  Configuration file - Ok
```

## Cisco UCS Manager

Log into the Cisco UCS Manager environment and follow these steps:

1. Open a web browser and navigate to the Cisco UCS Manager Cluster IP address, for example https://10.29.149.205

**Figure 48 Cisco UCS Manager**

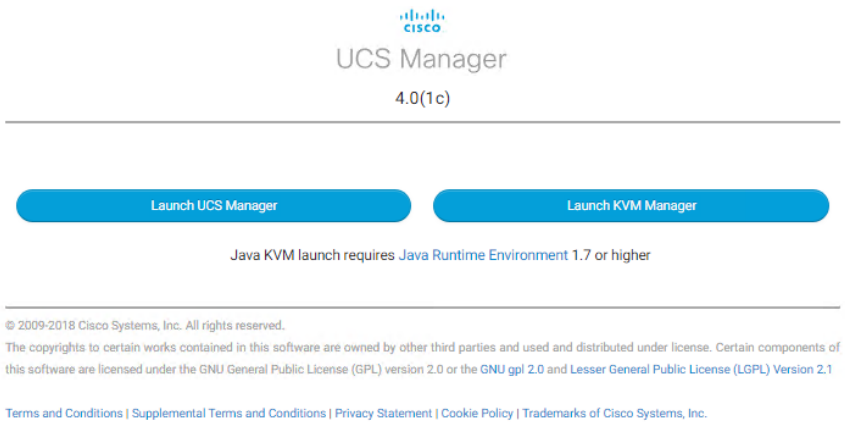


2. Click the “Launch UCS Manager” HTML link to open the Cisco UCS Manager web client.

3. At the login prompt, enter “admin” as the username, and enter the administrative password that was set during the initial console configuration.

4. Click No when prompted to enable Cisco Smart Call Home. This feature can be enabled at a later time.

## Cisco UCS Firmware

Your Cisco UCS firmware version should be correct as shipped from the factory, as documented in the Software Components section. This document is based on Cisco UCS infrastructure, B-series bundle, and C-Series bundle software versions 4.0(1c). If the firmware version of the Fabric Interconnects is older than this version, the firmware must be upgraded to match the requirements prior to completing any further steps.

To upgrade the Cisco UCS Manager version, the Fabric Interconnect firmware, and the server bundles, refer to these instructions:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Firmware-Mgmt/4-0/b_UCSM_GUI_Firmware_Management_Guide_4-0.html

## NTP

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

1. In Cisco UCS Manager, click the Admin button.
2. In the navigation pane, select All > Time Zone Management, and click the carat next to Time Zone Management to expand it.
3. Click Timezone.
4. In the Properties pane, select the appropriate time zone in the Time Zone menu.
5. Click Add NTP Server.
6. Enter the NTP server IP address and click OK.
7. Click OK.
8. Click Save Changes and then click OK.

## Uplink Ports

The Ethernet ports of a Cisco UCS Fabric Interconnect are all capable of performing several functions, such as network uplinks or server ports, and more. By default, all ports are unconfigured, and their function must be defined by the administrator. To define the specified ports to be used as network uplinks to the upstream network, follow these steps:

1. In Cisco UCS Manager, click the Equipment button.
2. Select Fabric Interconnects > Fabric Interconnect A > Fixed Module or Expansion Module as appropriate > Ethernet Ports.
3. Select the ports that are to be uplink ports, right-click them, and click Configure as Uplink Port.
4. Click Yes to confirm the configuration and click OK.
5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module as appropriate > Ethernet Ports.

6. Select the ports that are to be uplink ports, right-click them, and click Configure as Uplink Port.
7. Click Yes to confirm the configuration and click OK.
8. Verify all the necessary ports are now configured as uplink ports, where their role is listed as “Network.”

**Figure 49 Uplinks Ports**

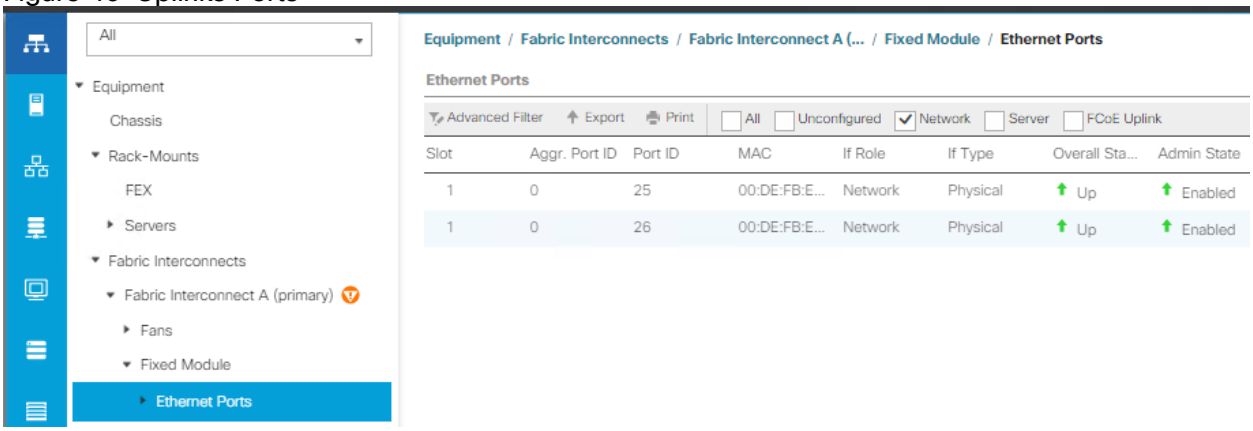


## Uplink Port Channels

If the Cisco UCS uplinks from one Fabric Interconnect are to be combined into a port channel or vPC, you must separately configure the port channels, which will use the previously configured uplink ports. To configure the necessary port channels in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the LAN button.
2. Under LAN > LAN Cloud, click the carat to expand the Fabric A tree.
3. Right-click Port Channels underneath Fabric A, then click Create Port Channel.
4. Enter the port channel ID number as the unique ID of the port channel (this does not have to match the port-channel ID on the upstream switch).
5. Enter the name of the port channel.
6. Click Next.
7. Click each port from Fabric Interconnect A that will participate in the port channel and click the >> button to add them to the port channel.
8. Click Finish.
9. Click OK.
10. Under LAN > LAN Cloud, click the carat to expand the Fabric B tree.
11. Right-click Port Channels underneath Fabric B, then click Create Port Channel.
12. Enter the port channel ID number as the unique ID of the port channel (this does not have to match the port-channel ID on the upstream switch).

13. Enter the name of the port channel.
14. Click Next.
15. Click each port from Fabric Interconnect B that will participate in the port channel and click the >> button to add them to the port channel.
16. Click Finish.
17. Click OK.
18. Verify the necessary port channels have been created. It can take a few minutes for the newly formed port channels to converge and come online.

**Figure 50 Uplink Port Channels**

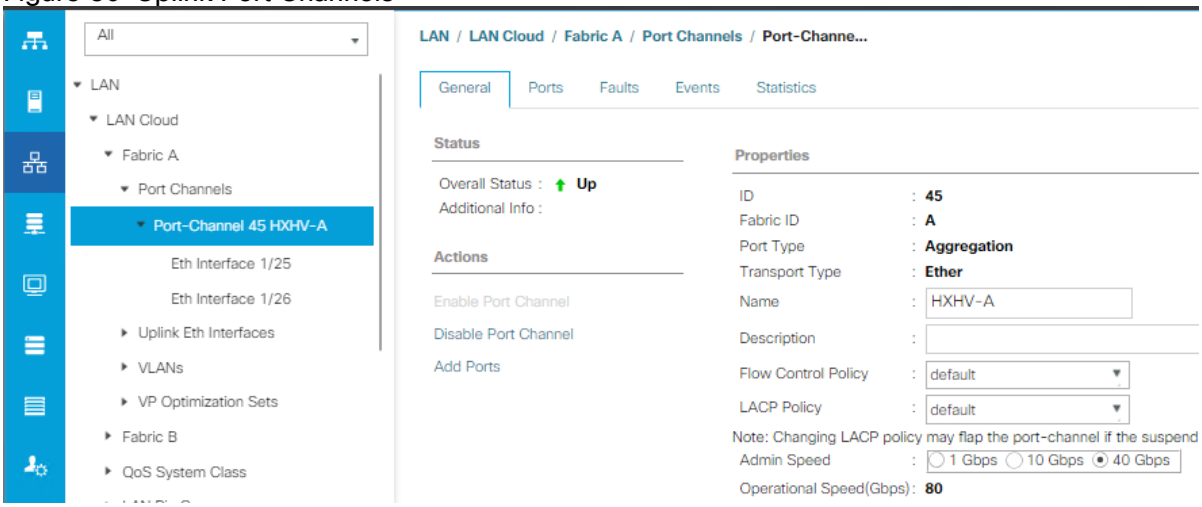


## Server Ports

The Ethernet ports of a Cisco UCS Fabric Interconnect connected to the rack-mount servers must be defined as server ports. When a server port is activated, the connected server or chassis will begin the discovery process shortly afterwards. Rack-mount servers are automatically numbered in Cisco UCS Manager in the order which they are first discovered. For this reason, it is important to configure the server ports sequentially in the order you wish the physical servers and/or chassis to appear within Cisco UCS Manager. For example, if you installed your servers in a cabinet or rack with server #1 on the bottom, counting as you go higher in the cabinet or rack, then you need to enable the server ports to the bottom-most server first, and enable them one-by-one as you move upward. You must wait until the server appears in the Equipment tab of Cisco UCS Manager before configuring the ports for the next server.

### Auto Configuration

A new feature in Cisco UCS Manager 3.1(3a) and later is Server Port Auto-Discovery, which automates the configuration of ports on the Fabric Interconnects as server ports when a Cisco UCS rack-mount server is connected to them. The firmware on the rack-mount servers must already be at version 3.1(3a) or later for this feature to function properly. Enabling this policy eliminates the manual steps of configuring each server port, however it does configure the servers in a somewhat random order. For example, the rack-mount server at the bottom of the stack, which you may refer to as server #1, and you may have plugged into port 1 of both Fabric Interconnects, could be discovered as server 2, or server 5, and so on. In order to have fine control of the rack-

mount server or chassis numbering and order, the manual configuration steps listed in the next section must be followed.

To configure automatic server port definition and discovery, follow these steps:

1. In Cisco UCS Manager, click the Equipment button.
2. In the navigation tree, under Policies, click Port Auto-Discovery Policy.
3. In the properties pane, set Auto Configure Server Port option to Enabled.
4. Click Save Changes.
5. Click OK.
6. Wait for a brief period, until the rack-mount servers appear in the Equipment tab underneath Equipment > Rack Mounts > Servers, or the chassis appears underneath Equipment > Chassis.

**Figure 51 Port Auto-Discovery Policy**

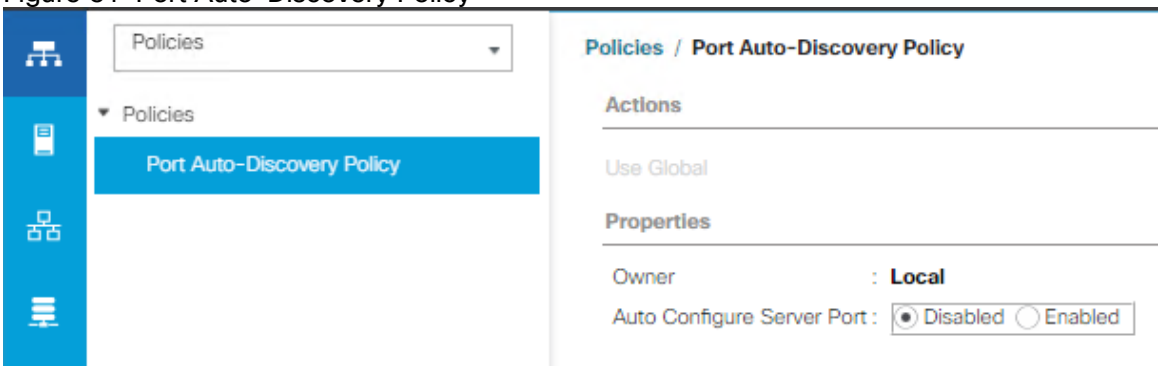


## Manual Configuration

To manually define the specified ports to be used as server ports and have control over the numbering of the servers, follow these steps:

1. In Cisco UCS Manager, click the Equipment button.
2. Select Fabric Interconnects > Fabric Interconnect A > Fixed Module or Expansion Module as appropriate > Ethernet Ports.
3. Select the first port that is to be a server port, right-click it, and click Configure as Server Port.
4. Click Yes to confirm the configuration and click OK.
5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module as appropriate > Ethernet Ports.
6. Select the matching port as chosen for Fabric Interconnect A that is to be a server port, right-click it, and click Configure as Server Port.
7. Click Yes to confirm the configuration and click OK.

8. Wait for a brief period, until the rack-mount server appears in the Equipment tab underneath Equipment > Rack Mounts > Servers, or the chassis appears underneath Equipment > Chassis.

9. Repeat steps 1-8 for each pair of server ports, until all rack-mount servers and chassis appear in the order desired in the Equipment tab.

### Server Discovery

As previously described, when the server ports of the Fabric Interconnects are configured and active, the servers connected to those ports will begin a discovery process. During discovery, the servers' internal hardware inventories are collected, along with their current firmware revisions. Before continuing with the HyperFlex installation processes, which will create the service profiles and associate them with the servers, wait for all the servers to finish their discovery process and to show as unassociated servers that are powered off, with no errors.

## Deploying HX Data Platform Installer on Hyper-V Infrastructure

To deploy HX Data Platform Installer using Microsoft Hyper-V Manager to create a HX Data Platform Installer virtual machine, follow these steps:

1. Locate and download the HX Data Platform Installer.vhdx zipped file (for example, Cisco-HX-Data-Platform-Installer-v4.0.1b-33133-hyperv.vhdx.zip) from the Cisco Software Downloads site.

2. Extract the zipped folder to your local computer and copy the .vhdx file to the Hyper-V host where you want to host the HX Data Platform Installer. For example,

   \\hyp-v-host01\....\HX-Installer\Cisco-HX-Data-Platform-Installer-v4.0.1b-33133-hyperv.vhdx.zip

3. In Hyper-V Manager, navigate to one of the Hyper-V servers.

4. Select the Hyper-V server, and right-click and select New > Create a virtual machine. The Hyper-V Manager New Virtual Machine Wizard displays.

**Figure 52 Hyper-V Manager – New Virtual Machine**

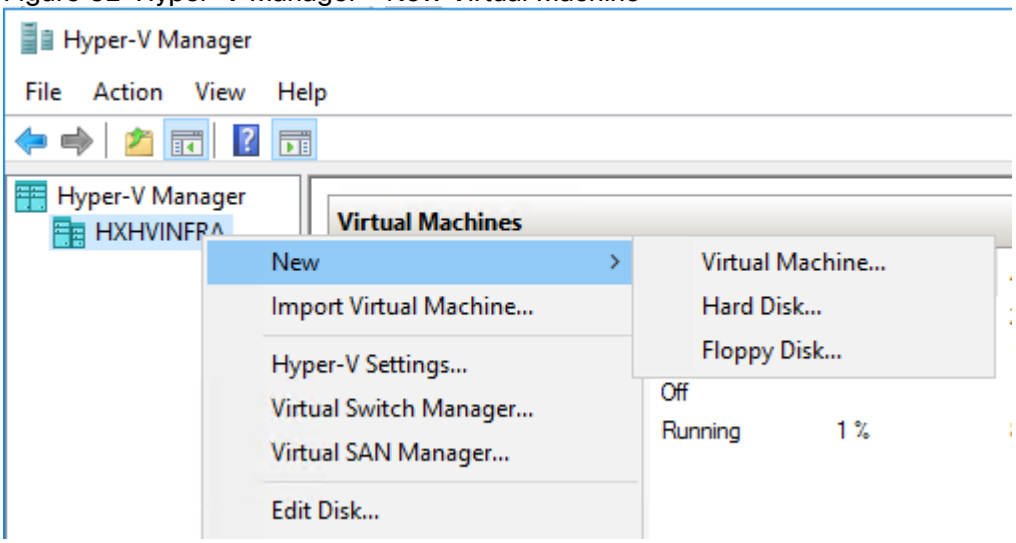


5. In the Before you Begin page, click Next.

6. In the Specify Name and Location page, enter a name and location for the virtual machine where the virtual machine configuration files will be stored. Click Next.

Figure 53 Hyper-V Manager – Specify Virtual Machine Name

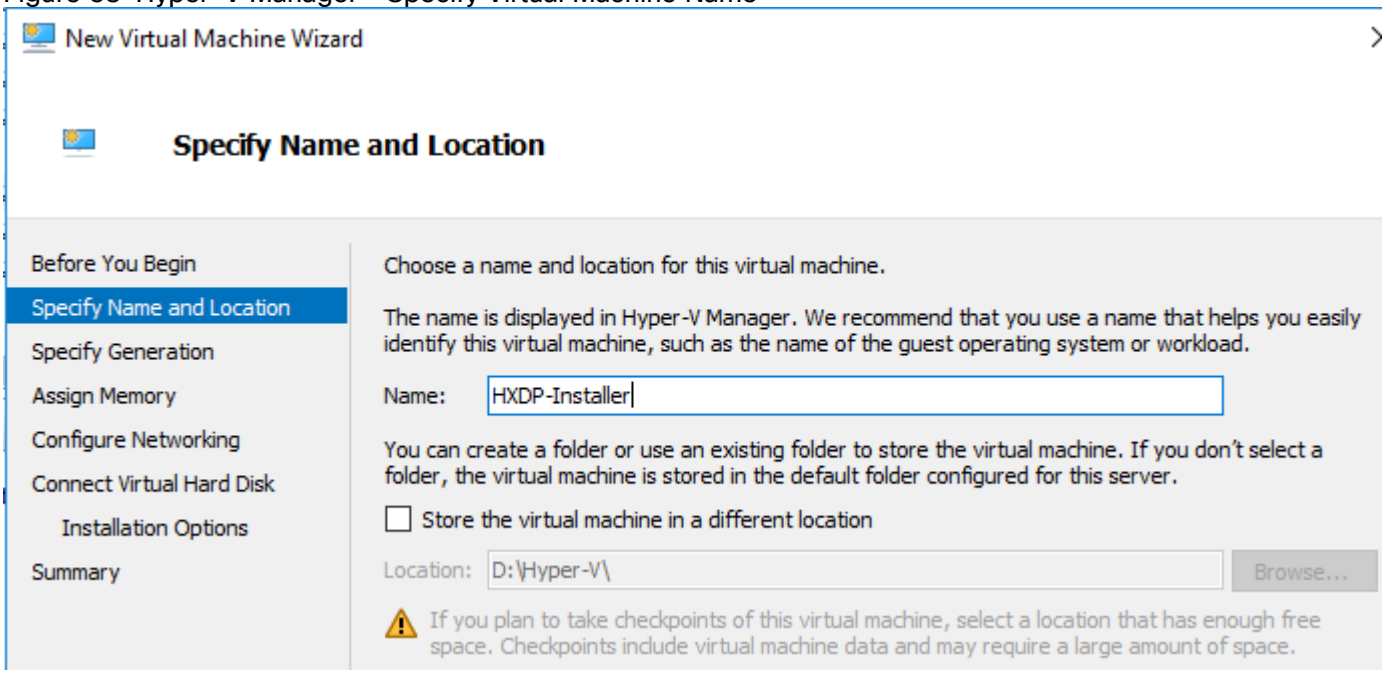


As a best practice, store the virtual machine together with the `.vhdx` file.

7. In the Specify Generation page, select Generation 1. Click Next.

Figure 54 Hyper-V Manger – Specify Virtual Machine Generation

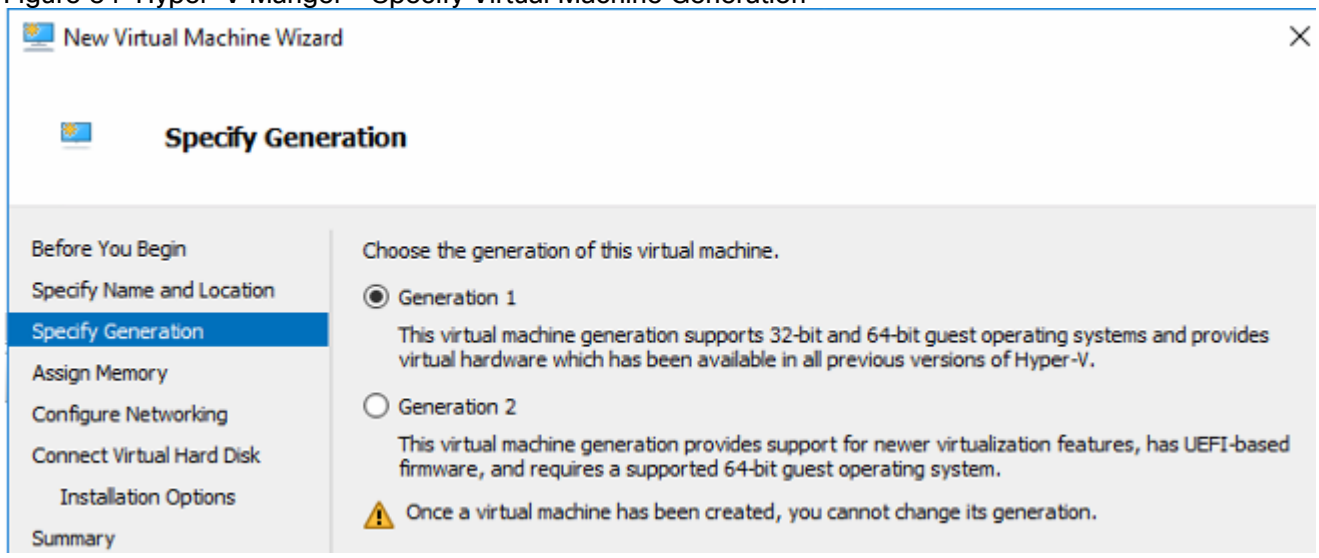


If you select Generation 2, the virtual machine may not boot.

8. In the Assign Memory page, set the startup memory value to 4096 MB. Click Next.

Figure 55 Hyper-V Manager – Assign Virtual Machine Memory

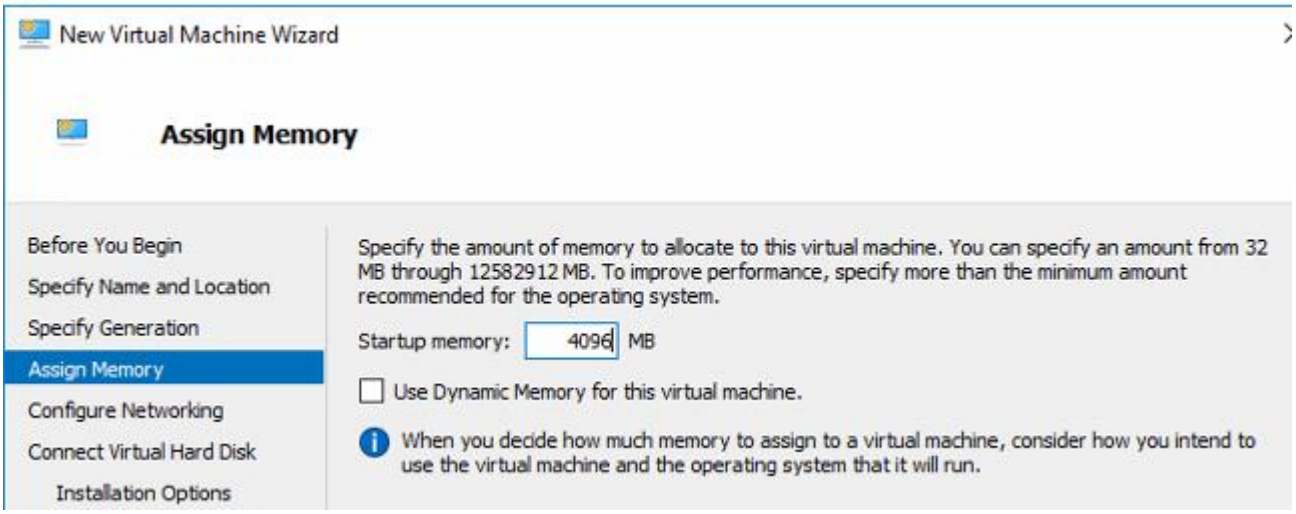


9. In the Configure Networking page, select a network connection for the virtual machine to use from a list of existing virtual switches. Click Next.

Figure 56 Hyper-V Manager – Configure Virtual Machine Networking

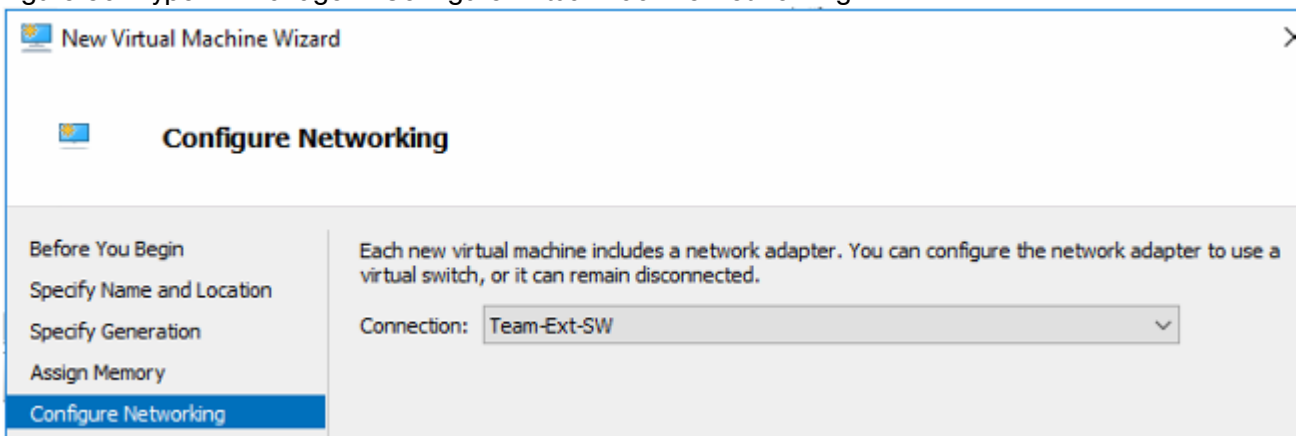


10. In the Connect Virtual Hard Disk page, select Use an existing virtual hard disk, and browse to the folder on your Hyper-V host that contains the `.vhdx` file. Click Next.

Figure 57 Hyper-V Manager – Connect Virtual Hard Disk

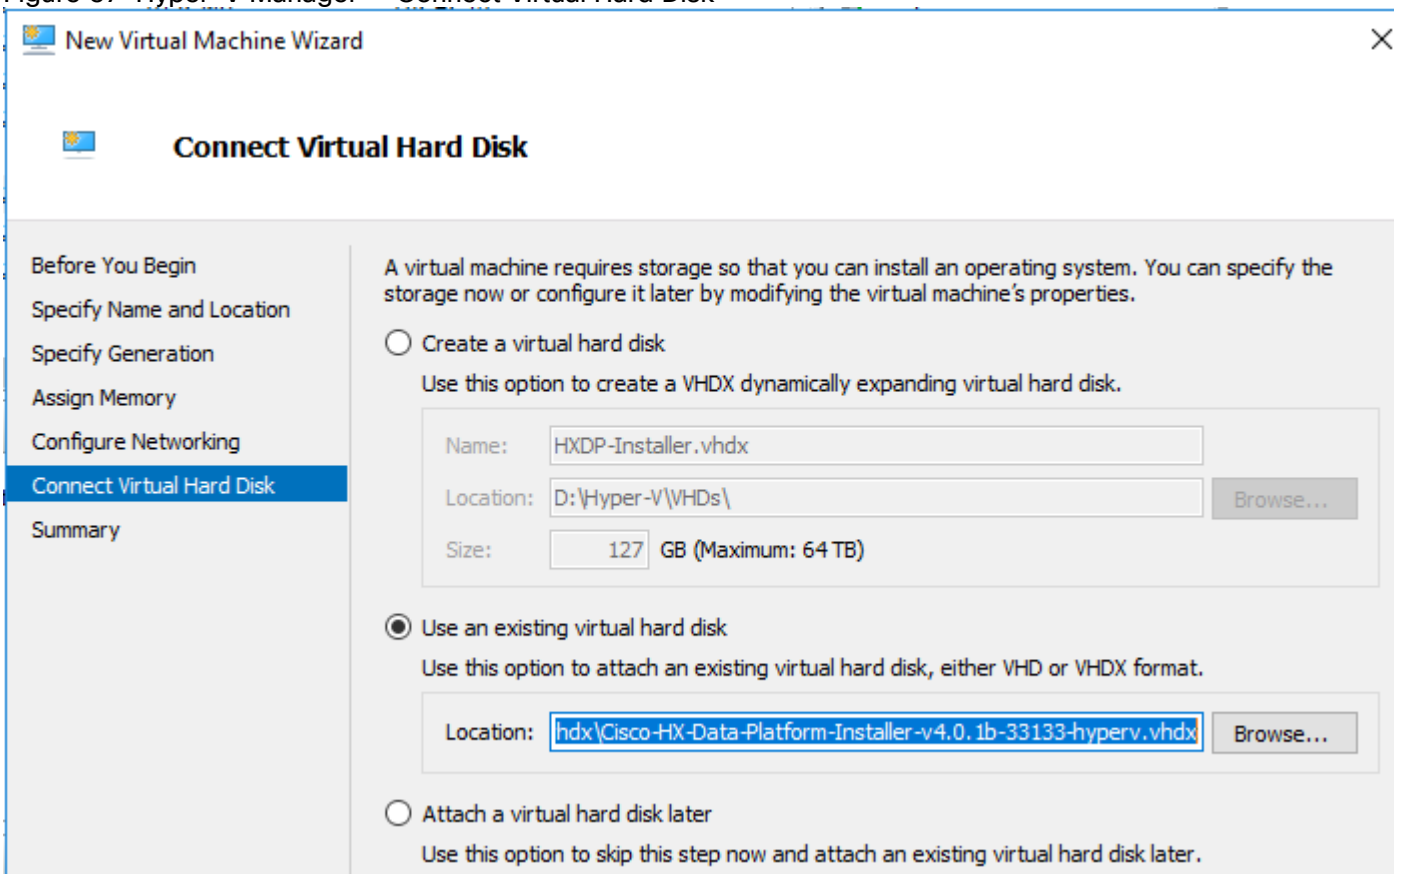


11. In the Summary page, verify that the list of options displayed are correct. Click Finish.

12. After the virtual machine is created, edit settings and assign four virtual CPUs as shown below.

Figure 58 Hyper-V Manager - Assign Virtual CPUs

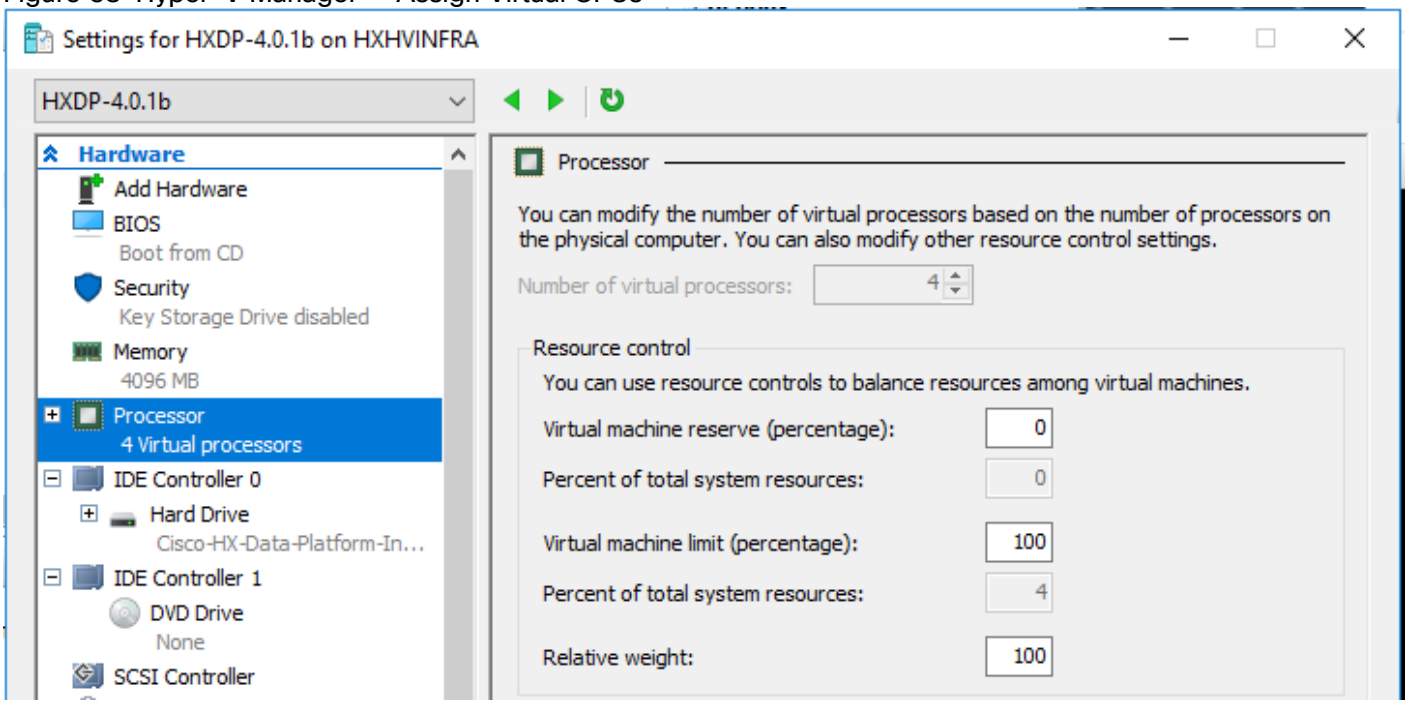


13. Review the final configuration summary and click Finish.
14. Right-click the virtual machine and choose Connect.
15. Choose Action > Start (Ctrl+S).
16. When the virtual machine is booted, login as 'root' with default password. The default password is 'Cisco123' (without quotes)
17. After logging in with default password, you are prompted to change the default password for root.

Figure 59 HXDP Installer

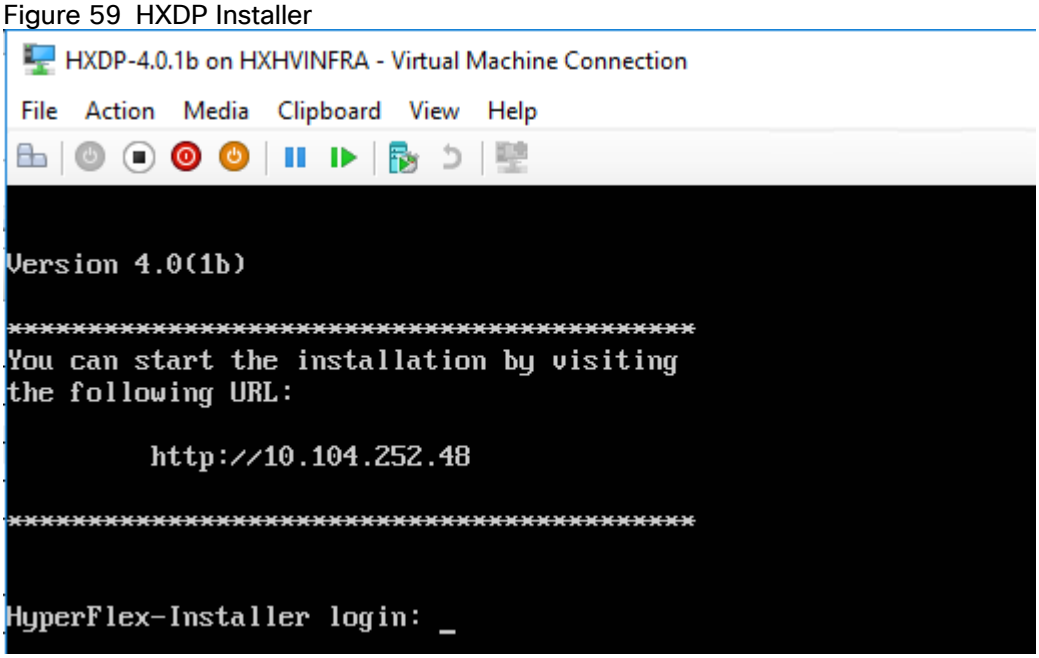


## Assign a Static IP Address to the HX Data Platform Installer Virtual Machine

During a default installation of the virtual machine, the HXDP Installer will try and automatically obtain an IP address using DHCP. To ensure that you have the same IP address on every boot, you can assign a static IP address on the virtual machine.

To configure your network interface (/etc/network/interfaces) with a static IP address. Make sure you change the relevant settings to suit your network and follow these steps:

1. Log into your Installer machine via the Hyper-V Console.
2. Type 'Ifdown eth0' to shoutdown the interface.
3. Edit the '/etc/network/eth0.interface' file and add the following lines to the file:

```
auto eth0
iface eth0 inet static
metric 100
address 10.104.252.48
netmask 255.255.255.0
gateway 10.104.252.1
dns-nameservers 10.104.252.48
dns-search hxhvdom2.local
```

4. Press 'ESC' to exit the insert mode and type '.wq' to save and quit the VI.
5. Type 'Ifup eth0' to bring up the interface.
6. Reboot the virtual machine for changes to take effect.
7. Verify the settings as shown in the following figures.

Figure 60 Show Interfaces

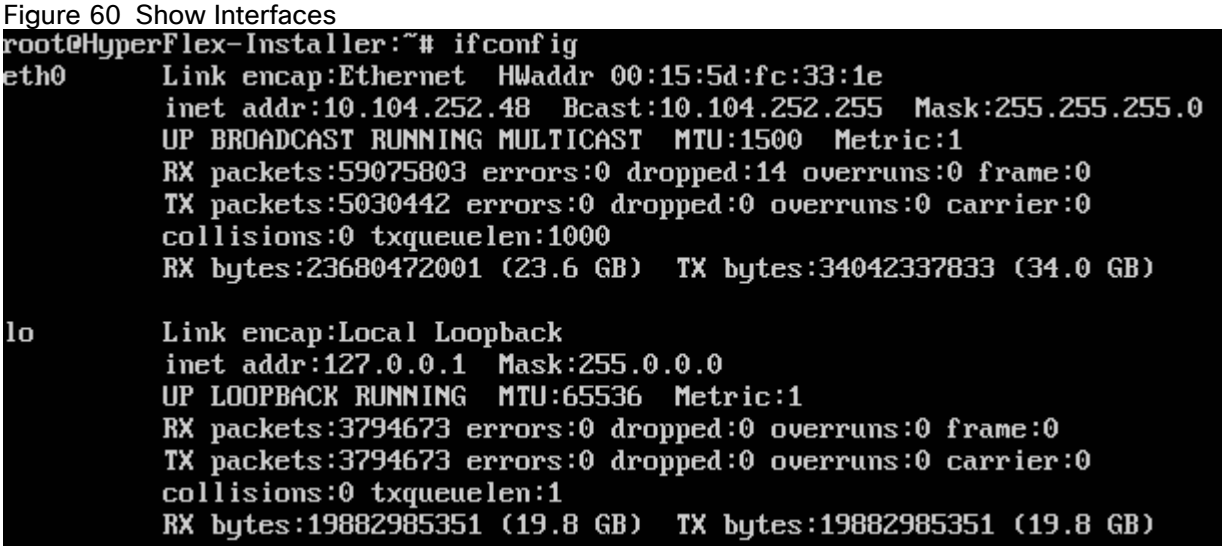


Figure 61 Routing Table

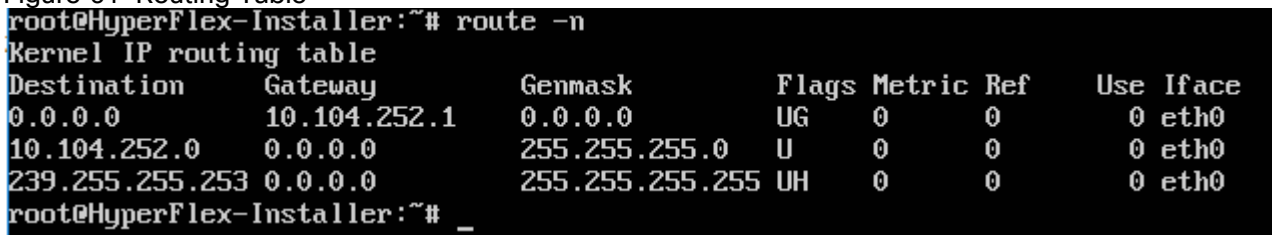


Figure 62 /etc/resolv.conf file for Nameserver and Search Domain

## HyperFlex Installer Web Page

The HyperFlex installer is accessed via a webpage using your local computer and a web browser. If the HyperFlex installer was deployed with a static IP address, then the IP address of the website is already known.

If DHCP was used, open the local console of the installer virtual machine. In the console, you will see an interface like the example below, showing the IP address that was leased:

Figure 63 HyperFlex Installer Virtual Machine IP Address

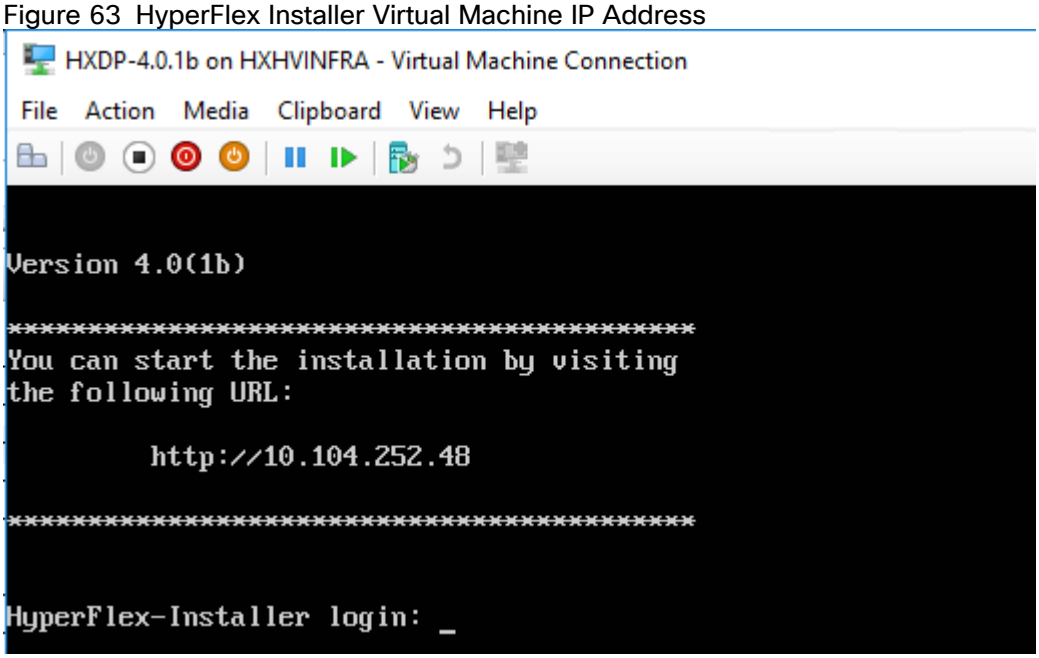


To access the HyperFlex installer webpage, follow these steps:

1. Open a web browser on the local computer and navigate to the IP address of the installer virtual machine. For example, open http://10.104.252.48
2. Click accept or continue to bypass any SSL certificate errors.
3. At the login screen, enter the username: root
4. At the login screen, enter the password
5. Verify the version of the installer in the lower right-hand corner of the Welcome page is the correct version.
6. Check the box for “I accept the terms and conditions” and click Login.

Figure 64 HyperFlex Connect

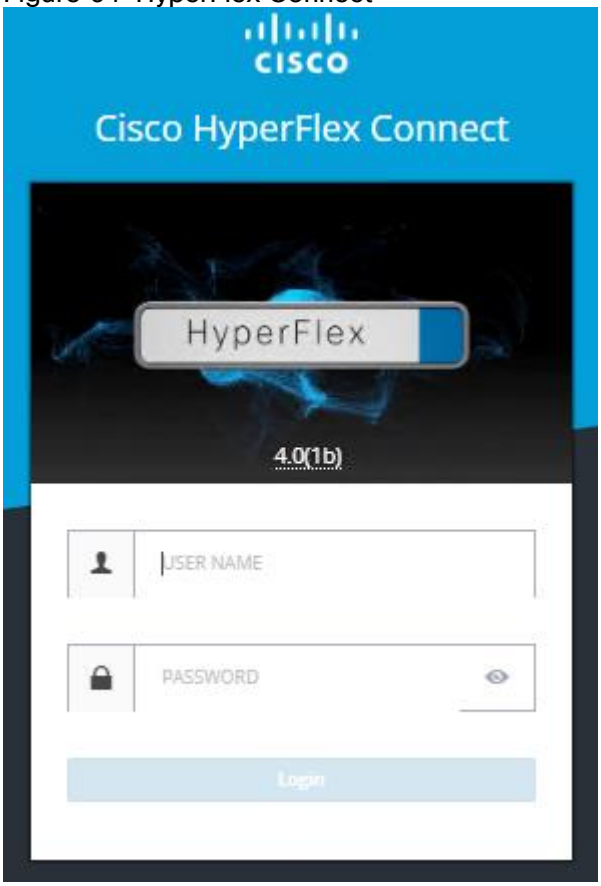


## HyperFlex Installation

The HyperFlex installer will guide you through the process of setting up your cluster. The Windows OS is not factory installed and requires the customer to provide media for the installation. It will configure Cisco UCS policies, templates, service profiles, settings and install Windows Server 2016/2019, as well as assigning IP addresses to the HX servers after the OS installation. The installer will deploy the HyperFlex controller virtual machines and software on the nodes, add the nodes to the Windows failover cluster, then finally create the HyperFlex cluster and distributed filesystem. All these processes can be completed through a single workflow from the HyperFlex Installer webpage.

To install and configure a HyperFlex cluster, follow these steps:

1. On the HyperFlex installer webpage click “Cluster Creation with HyperFlex (FI).”

Figure 65 HyperFlex Installer - Workflow

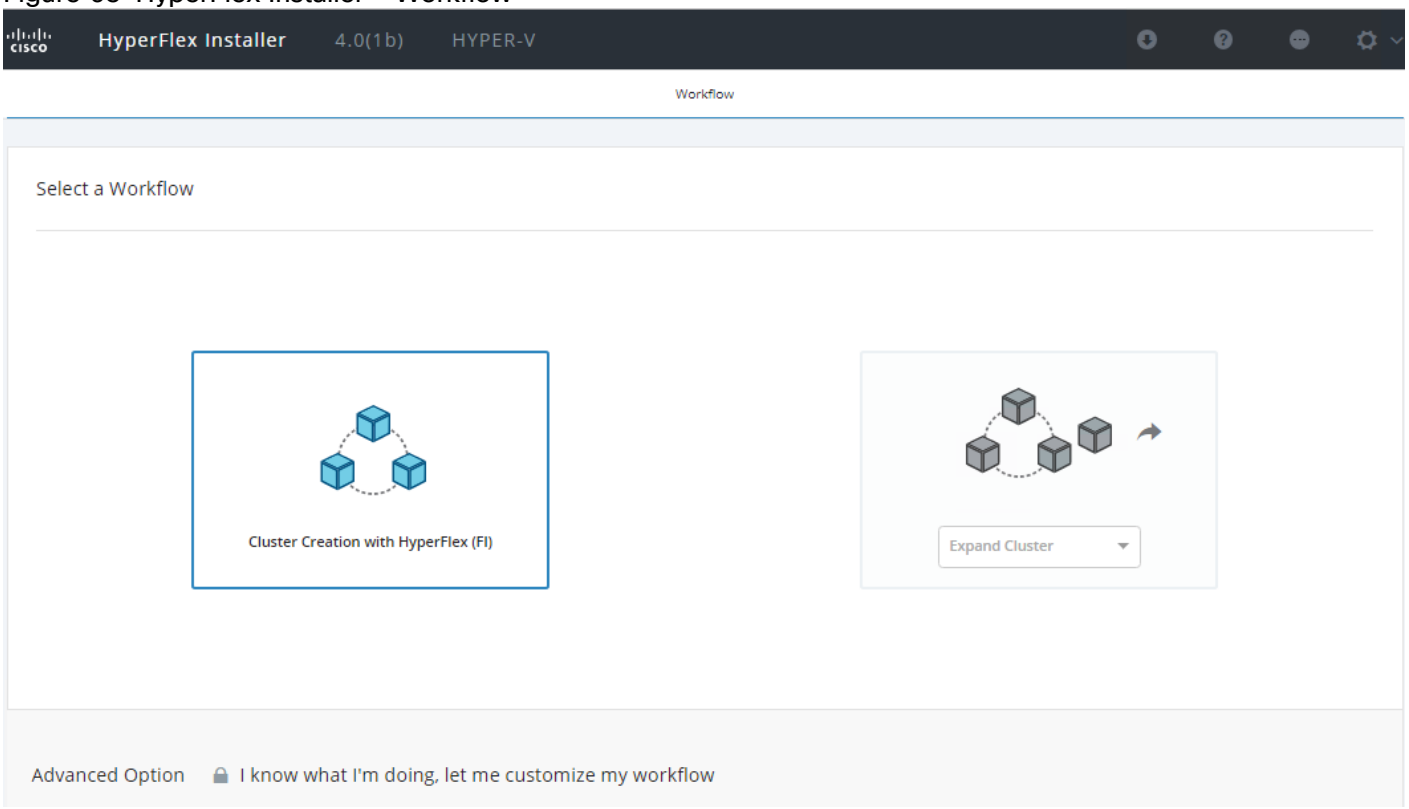


2. On the Credentials page:

   a. Under the "UCS Manager Credentials", enter the Cisco UCS Manager Host Name or IP Address, UCS Manager User Name and Password

   b. Under the "Domain Information", enter the AD Domain Name, DNS Server IP Address, HX Service Account user name and password. It is recommended to select "Configure Constrained Delegation now" and select "Use HX Service Account" if HX service account is member of AD Domain Admin group, else provide Domain Admin credentials (which is a one-time requirement). To configure Constrained Delegation later, refer the appendix section of this document.

   c. And, under "Advanced Attributes (optional)", enter the Domain Controller IP address and the distinguished name of Organization Unit (OU). Providing distinguished name of the OU is required, if you want the Computer objects created to be placed under a specific OU instead of the default built-in "Computers" OU.

3. Optionally, you can import a JSON file that has the configuration information, except for the appropriate passwords.

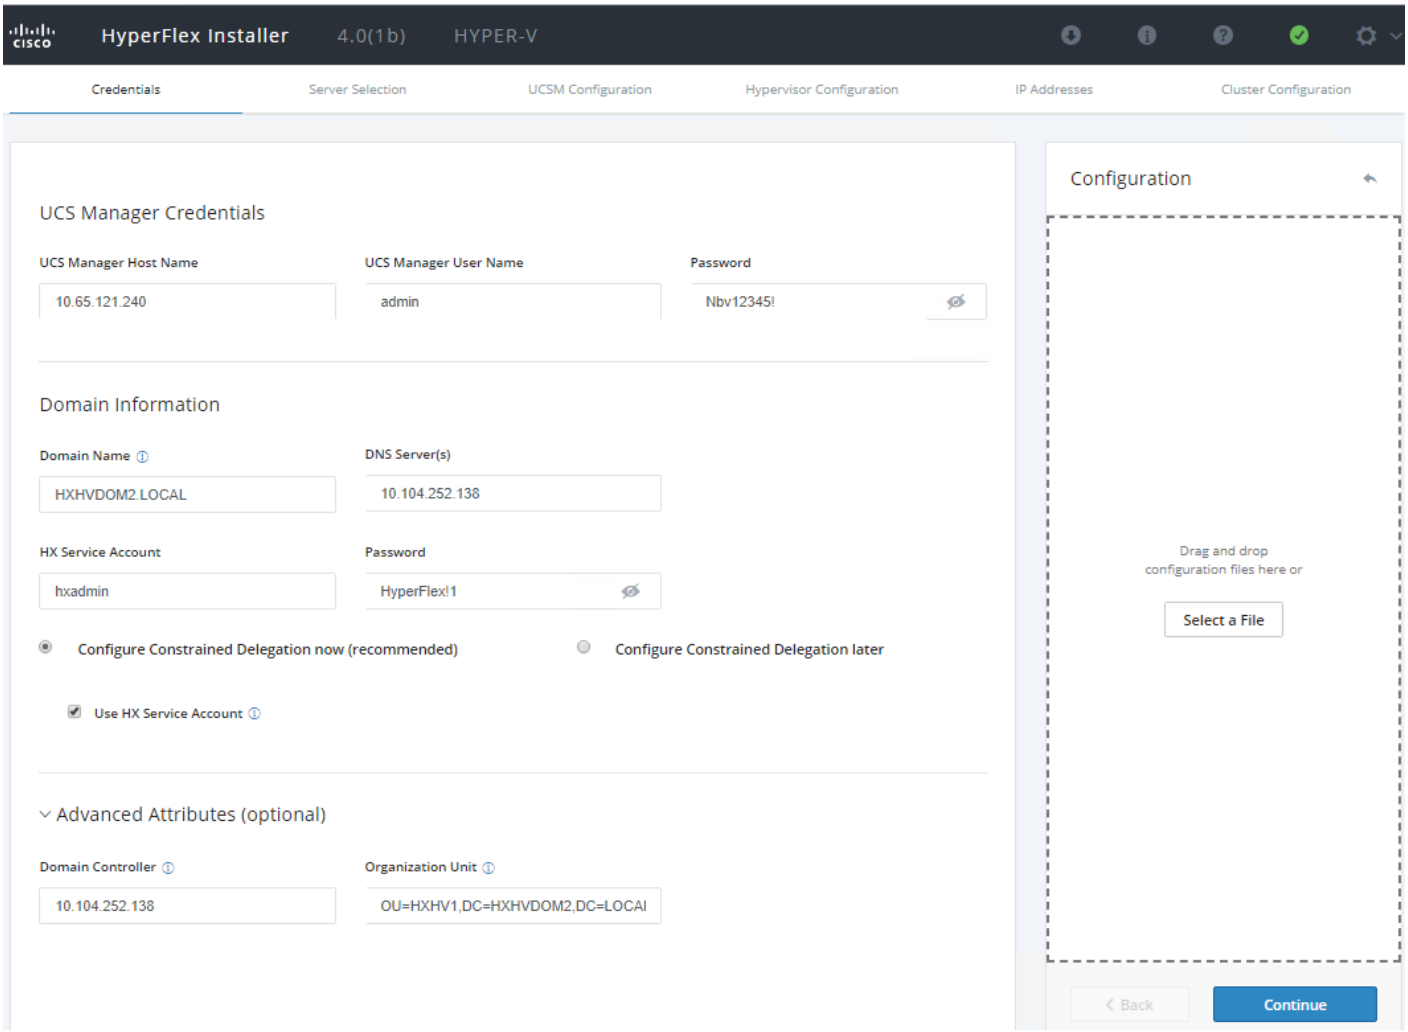


4. Click Continue.

5. In the Server Selection page:

   a. Select the Unassociated HX server models that are to be used in the new HX cluster and click Continue. We have selected three nodes to demonstrate deployment of 3-node hx cluster.

   b. If the Fabric Interconnect server ports were not enabled in the earlier step, you have the option to enable them here to begin the discovery process by clicking the Configure Server Ports link.

**HyperFlex for Hyper-V only supports M5 Servers for converged nodes.**

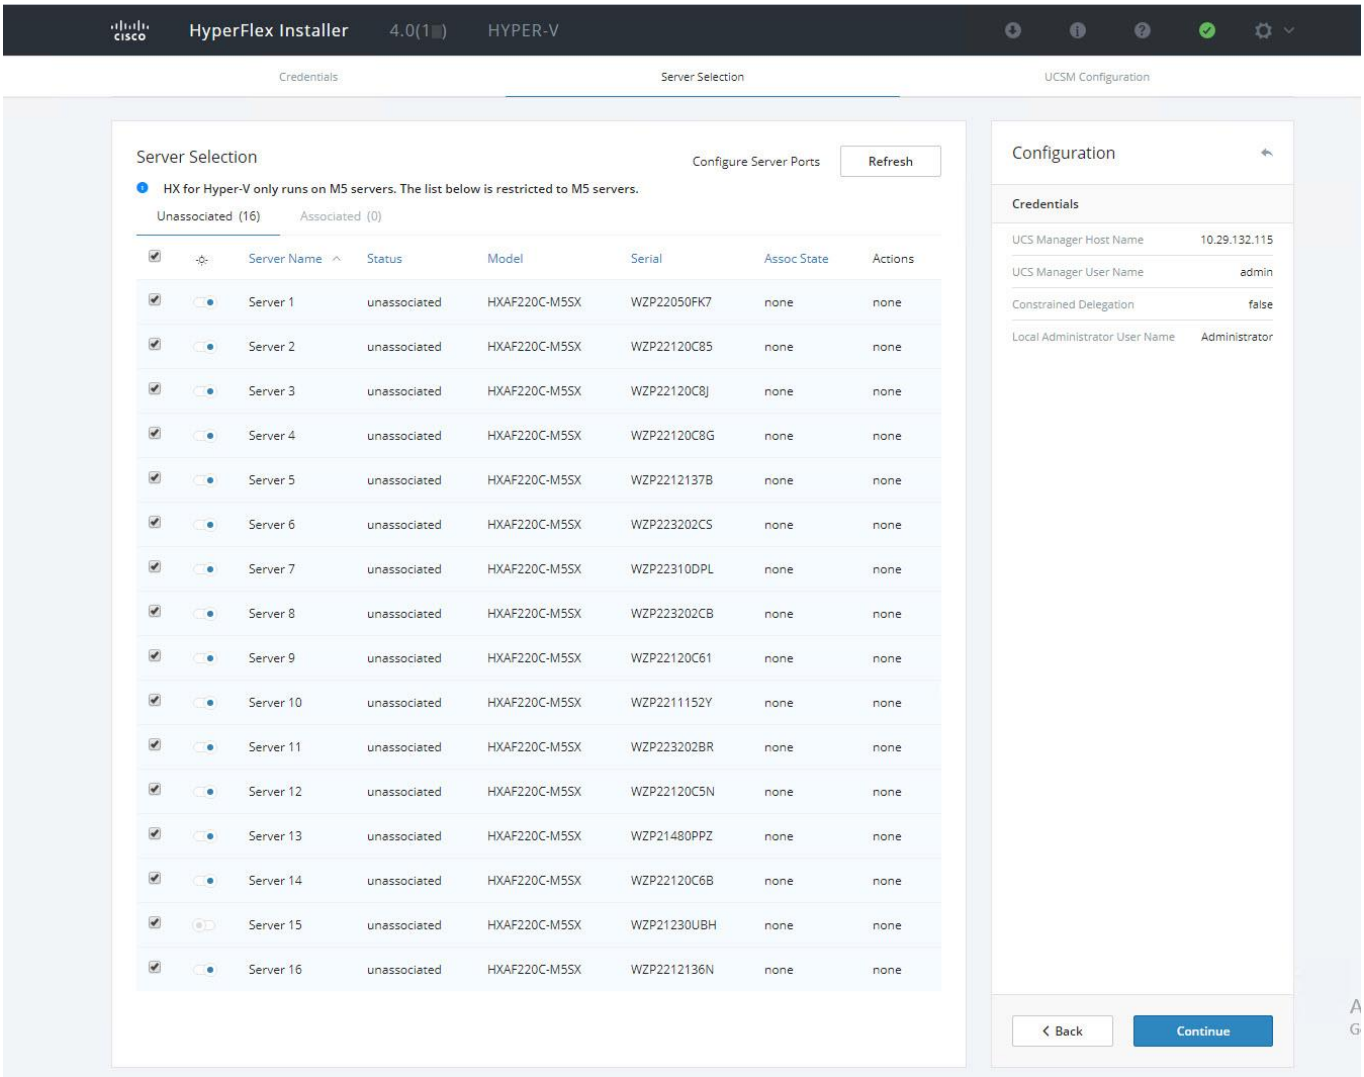


**Using the option to enable the server ports within the HX Installer will not allow you to finely control the server number order, as would be possible when performing this step manually before installing the HyperFlex cluster. To have control of the server number order, perform the steps outlined earlier for manually configuring the server ports. The server discovery can take several minutes to complete, and it will be necessary to periodically click the Refresh button to see the unassociated servers appear once discovery is completed.**

6. On the UCSM Configuration page:

   a. VLAN Configuration – HyperFlex needs to have at least 4 VLANs to function; each VLAN needs to be on different IP subnets and extended from the fabric interconnects to the connecting uplink switches, to make sure that traffic can flow from Primary Fabric Interconnect (Fabric A) to Subordinate Fabric Interconnect (Fabric B).

   b. Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed here.

**Do not use VLAN 1, since it is not a best practice and can cause issues with disjoint layer 2.**

vm-network can have multiple VLAN IDs added as a comma separated list (as shown in the below screenshot).

Renaming the 4 core networks is not supported.

c. MAC Pool – Enter the MAC Pool prefix, only enter the 4th byte value, for example: 00:25:B5:0A.

d. 'hx' IP Pool for Cisco IMC – Enter the IP address range, subnet mask and gateway to be used by the CIMC interfaces of the servers in this HX cluster.

e. Cisco IMC access Management (Out of band or inband) – Select the recommended 'in band' option for faster installation of hypervisor OS on all the hx nodes.

f. The Out-Of-Band network needs to be on the same subnet as the Cisco UCS Manager. You can add multiple blocks of addresses as a comma separated line.

g. VLAN for Inband Cisco IMC Connectivity – Enter a VLAN name and ID.

h. Advanced – If multiple firmware packages exist on the Fabric Interconnect, choose the version to be installed on the servers that will comprise this cluster. Note that for HXDP 4.0(1b) release with M5 generation servers running on 2nd Generation Intel® Xeon® Scalable Processors, the supported and recommended version of UCS FI firmware's is 4.0(4d).

i. Enter a unique Org name for the HyperFlex Cluster.

The Cisco UCS B and C packages must exist on the Fabric interconnect otherwise the installation will fail. If the right version is not available in the drop-down list, then upload it to Cisco UCS Manager before continuing.

j. iSCSI/FC Storage (optional) – iSCSI Storage and FC Storage are used for adding external storage to the HyperFlex cluster. Not defined for this setup.

**Important!** When deploying a second or any additional clusters, you must put them into a different sub-org, use a different MAC Pool prefix, a unique pool of IP addresses for the CIMC interfaces, and you should also create new VLAN names for the additional clusters. Even if reusing the same VLAN ID, it is prudent to create a new VLAN name to avoid conflicts. For example, for a second cluster change the VLAN names, use a unique MAC Pool prefix, IP address pool, Cluster Name and Org Name to not overwrite the original cluster information.

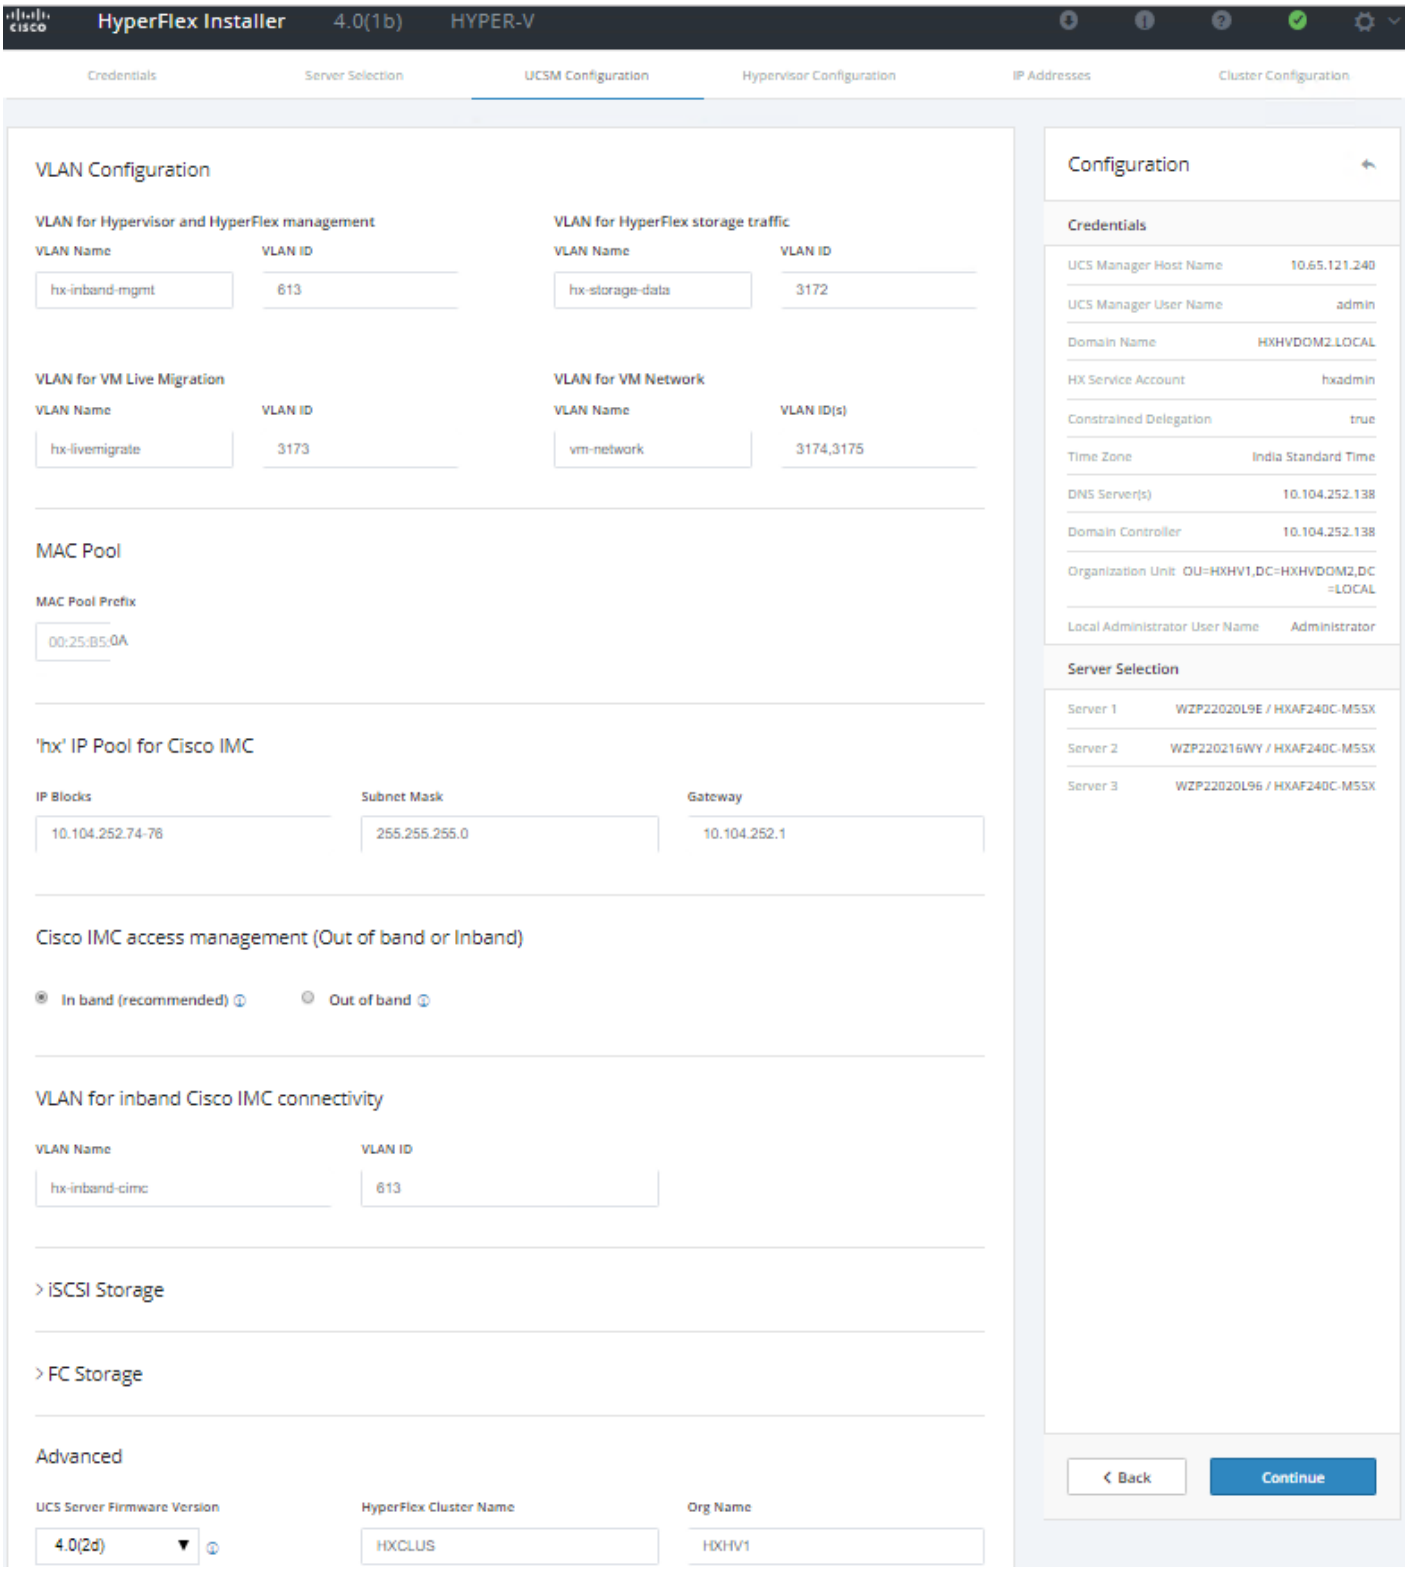


7. Click Continue.

8. On the Hypervisor Configuration page:

a. Bare Metal Configuration – If Windows Server 2016/2019 is not installed on the nodes, select “Install Hypervisor (Hyper-V)”, drag or click browse to upload OS media file in the box and select the radio button to choose OS you want to install.

b. Configure Common Hypervisor Settings – Enter the subnet mask, gateway, DNS Server IP Address.

c. Hypervisor Settings – Enter IP addresses and hostnames for the Hypervisors that were created in the pre-installation section phase. The IP addresses will be assigned via Serial over Lan (SoL) through Cisco UCS Manager to the Hyper-V host systems as their management IP addresses.

d. Primary DNS Suffix – Add any additional DNS suffixes.

9. Click Continue.

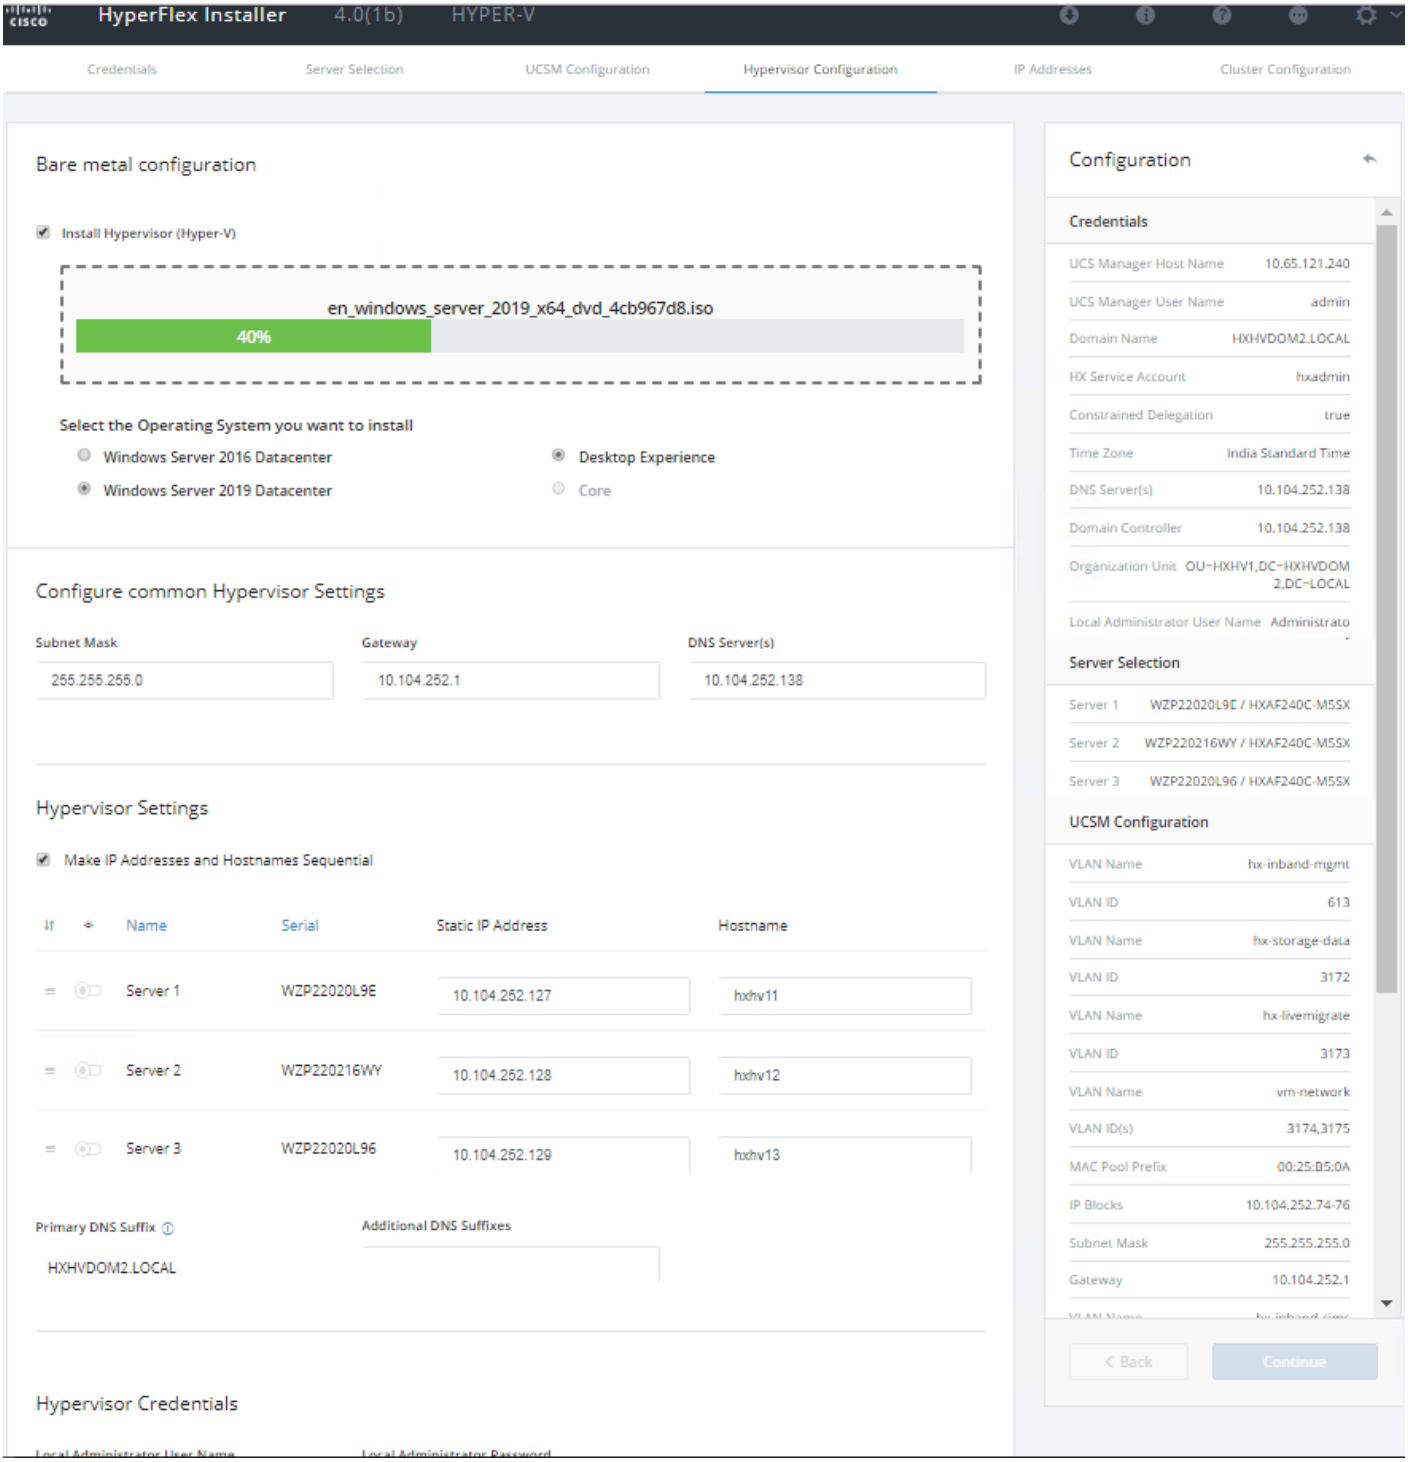


10. On the IP Addresses page:

    a. Assign the hostnames for the Storage Controllers Management that were created in the pre-installation phase.

**If you leave the checkbox, Make IP Addresses and Hostnames Sequential as checked, then the installer will automatically fill the rest of the servers sequentially.**

b. Assign the additional IP addresses for the Management and Data networks as well as the cluster IP addresses, then click Continue.

**A default gateway is not required for the data network, as those interfaces normally will not communicate with any other hosts or networks, and the subnet can be non-routable.**

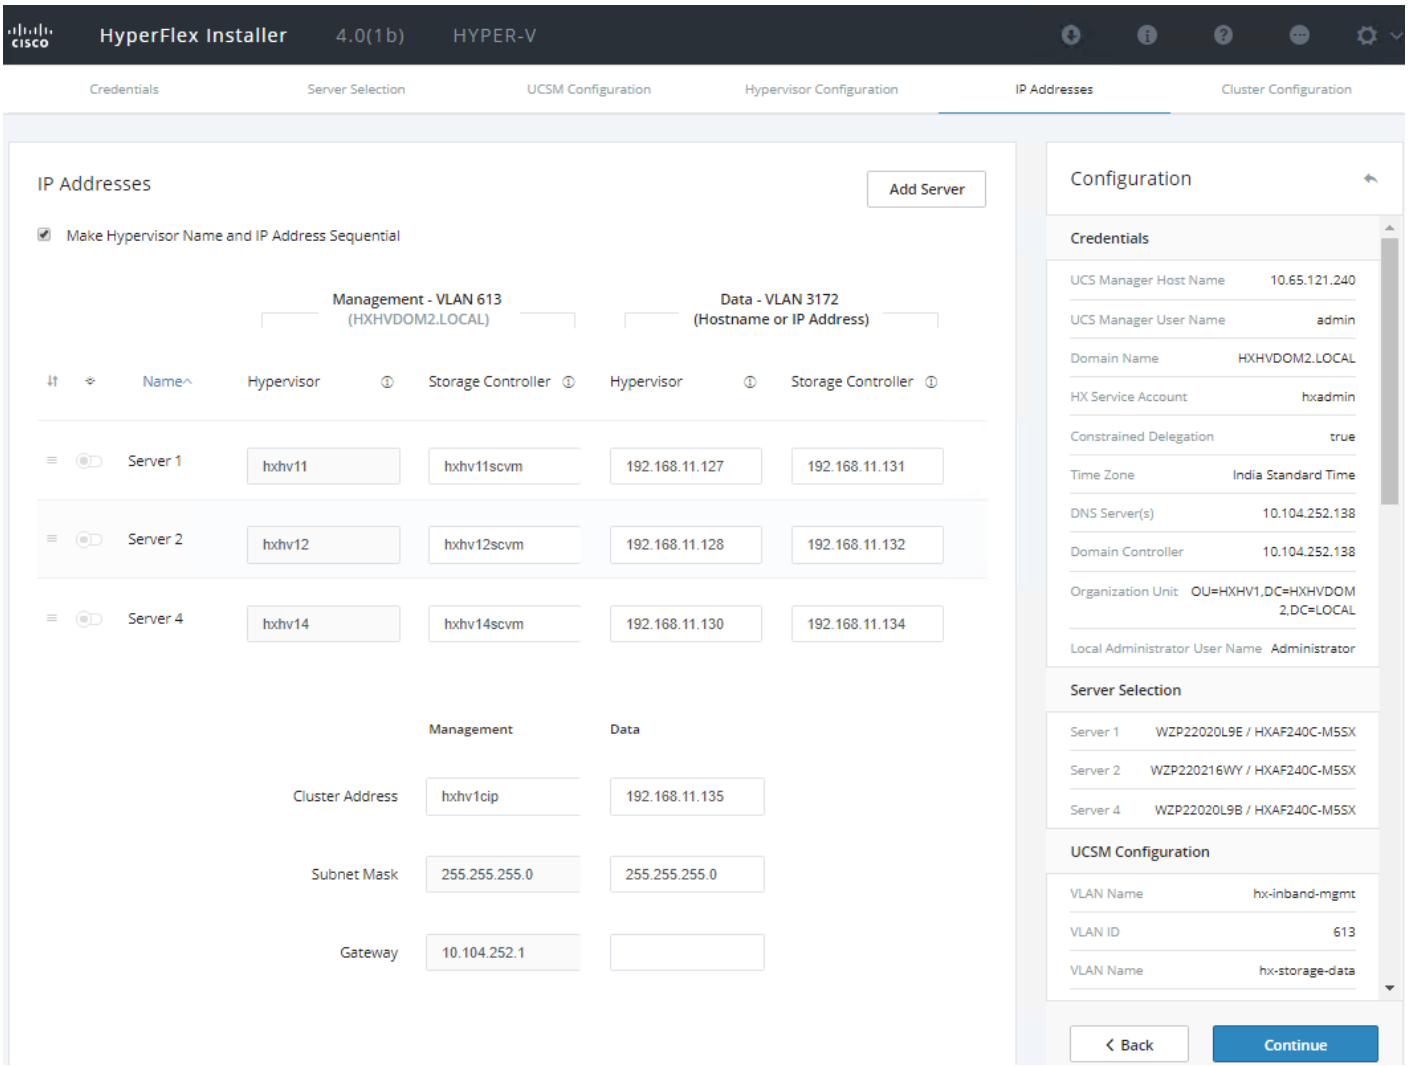


11. On the Cluster Configuration page:

a. Cisco HX Cluster - Enter the Cluster Name (SMB Access Point). Select a Replication Factor from the drop-down menu and enter a Windows Failover Cluster Name that was created in the pre-installation phase.

b. Controller virtual machine - Enter the Password that will be assigned to the Controller virtual machines.

c. System Services - Enter the AD DNS server IP address, and make sure to use the Active Directory domain name for NTP Server for Controller virtual machines to synchronize time with the Active Directory .

d. Time Zone – Select a time zone from the drop-down li.st

e. Auto Support - Enable Connected Services in order to enable management via Cisco Intersight and enter the email address to receive service ticket alerts, then scroll down.

f. Advanced Configuration – Jumbo Frames should be enabled to ensure the best performance, unless the upstream network is not capable of being configured to transmit jumbo frames. It is not necessary to select Clean up disk partitions for a new cluster installation, but an installation using previously used converged nodes should have the option checked.

12. Click Start.

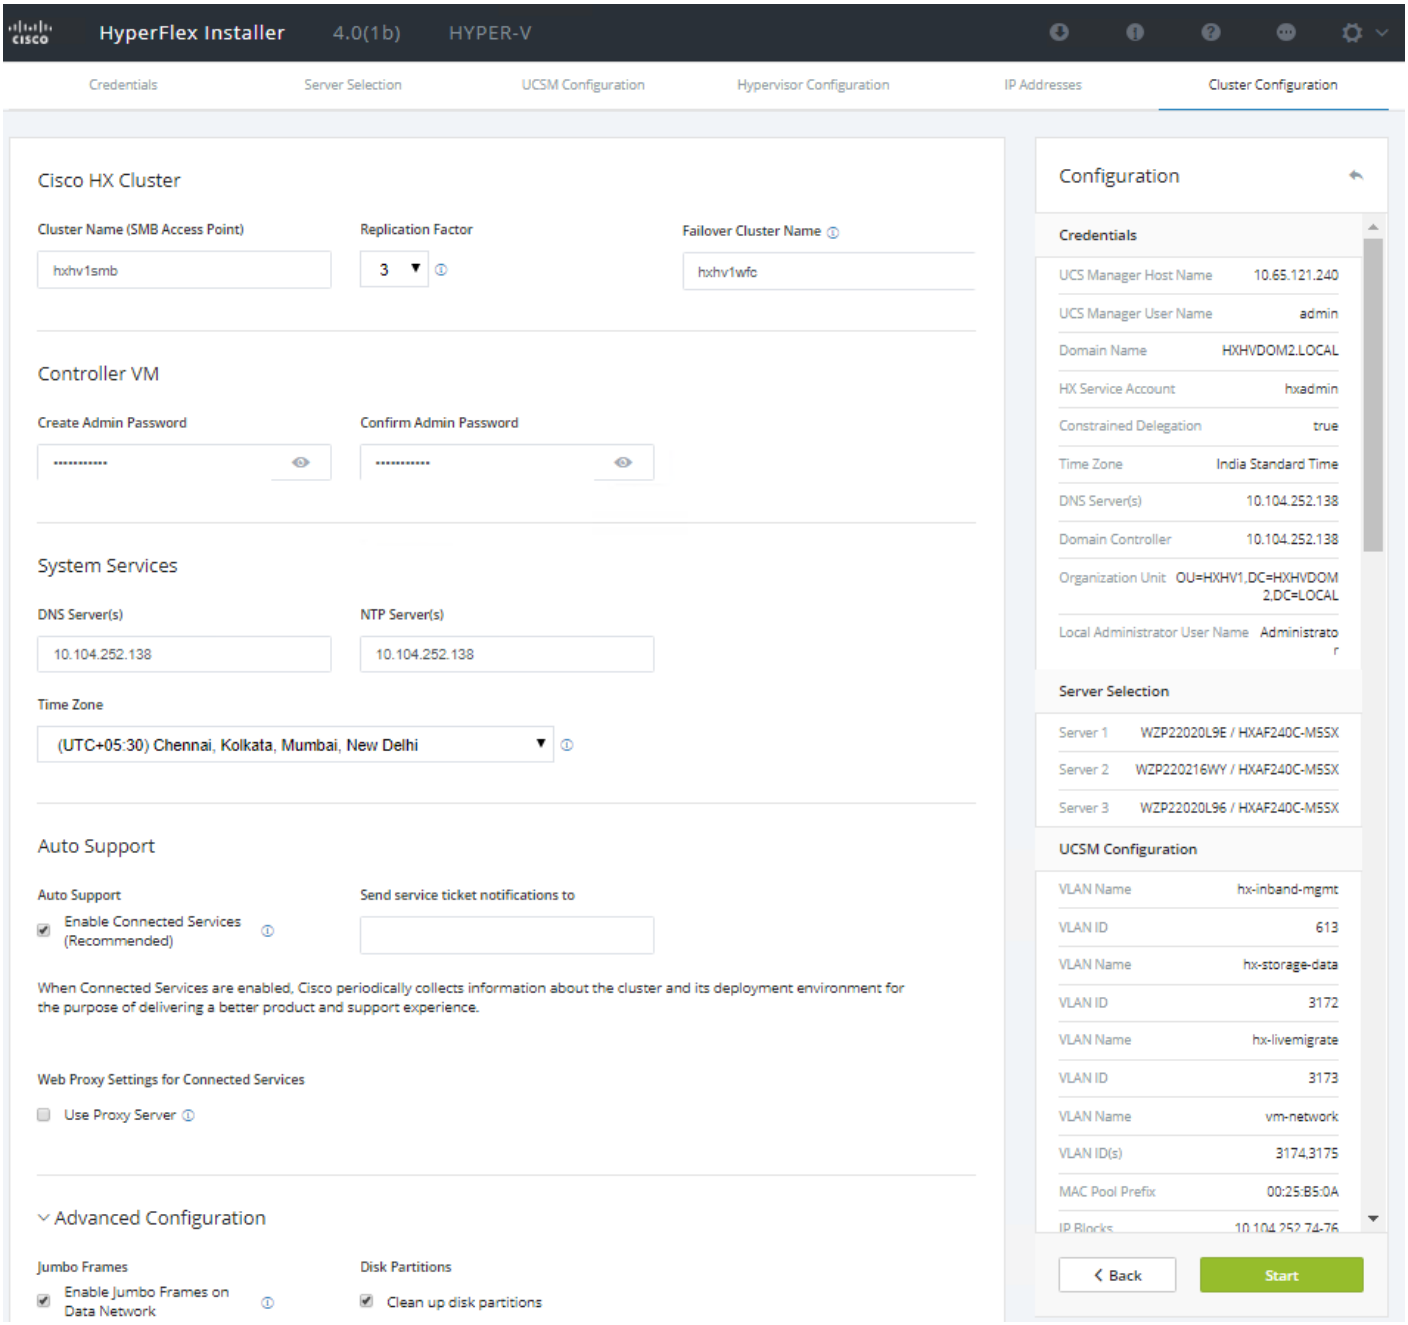


13. Validation of the configuration will now start. If there are warnings, you can review them and click “Skip Validation” if the warnings are acceptable. If there are no warnings, the installer will automatically continue to the configuration process.

The initial validation always fails when using new Cisco UCS 6332 or 6332-16UP model Fabric Interconnects. This is because changes to the QoS system classes require these models to reboot. If the validation is skipped, the HyperFlex installer will continue the installation and automatically reboot both Fabric Interconnects sequentially. If this is an initial setup of these Fabric Interconnects, and no other systems are running on them yet, then it is safe to proceed. However, if these Fabric Interconnects are already in use for other workloads, then caution must be taken to ensure that the sequential reboots of both Fabric Interconnects will not interrupt those workloads, and that the QoS changes will not cause traffic drops. Contact Cisco TAC for assistance if this situation applies.

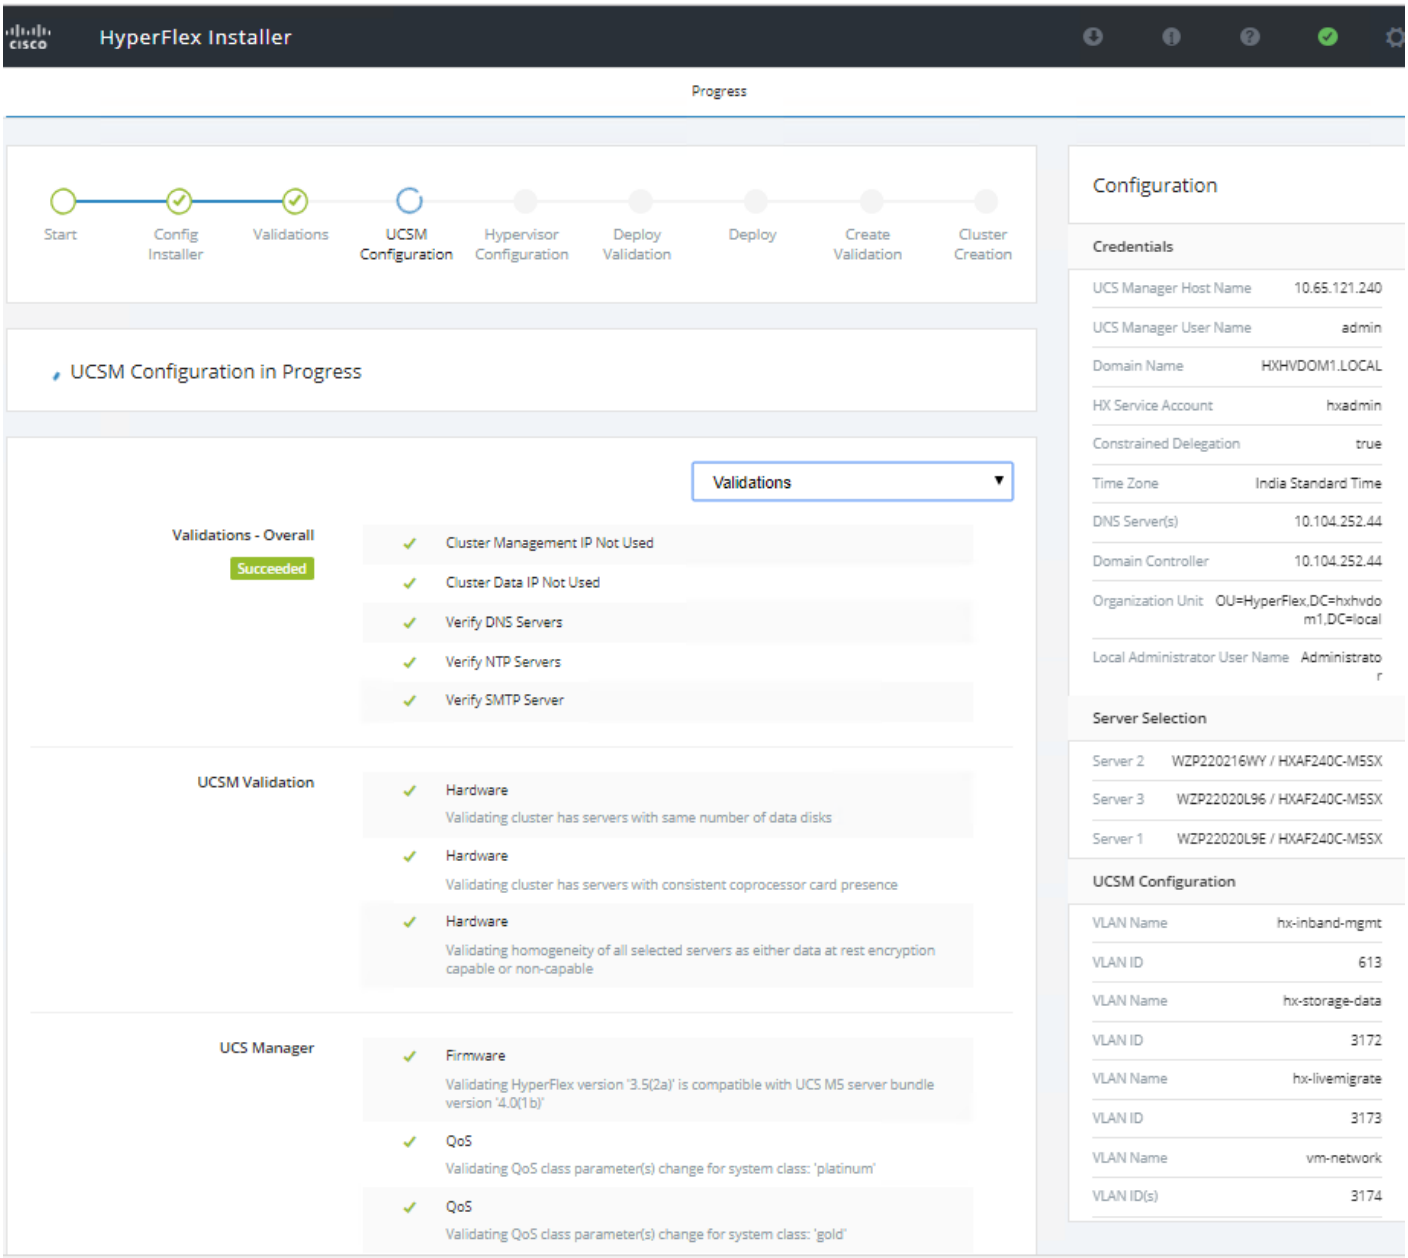

14. After the pre-installation validations, the HX installer will proceed to complete the deployment and perform all the steps listed at the top of the screen along with their status. The process can also be monitored in Cisco UCS Manager while the profiles and cluster are created.

The figure below shows the Cisco UCS Manager service profile association in progress:

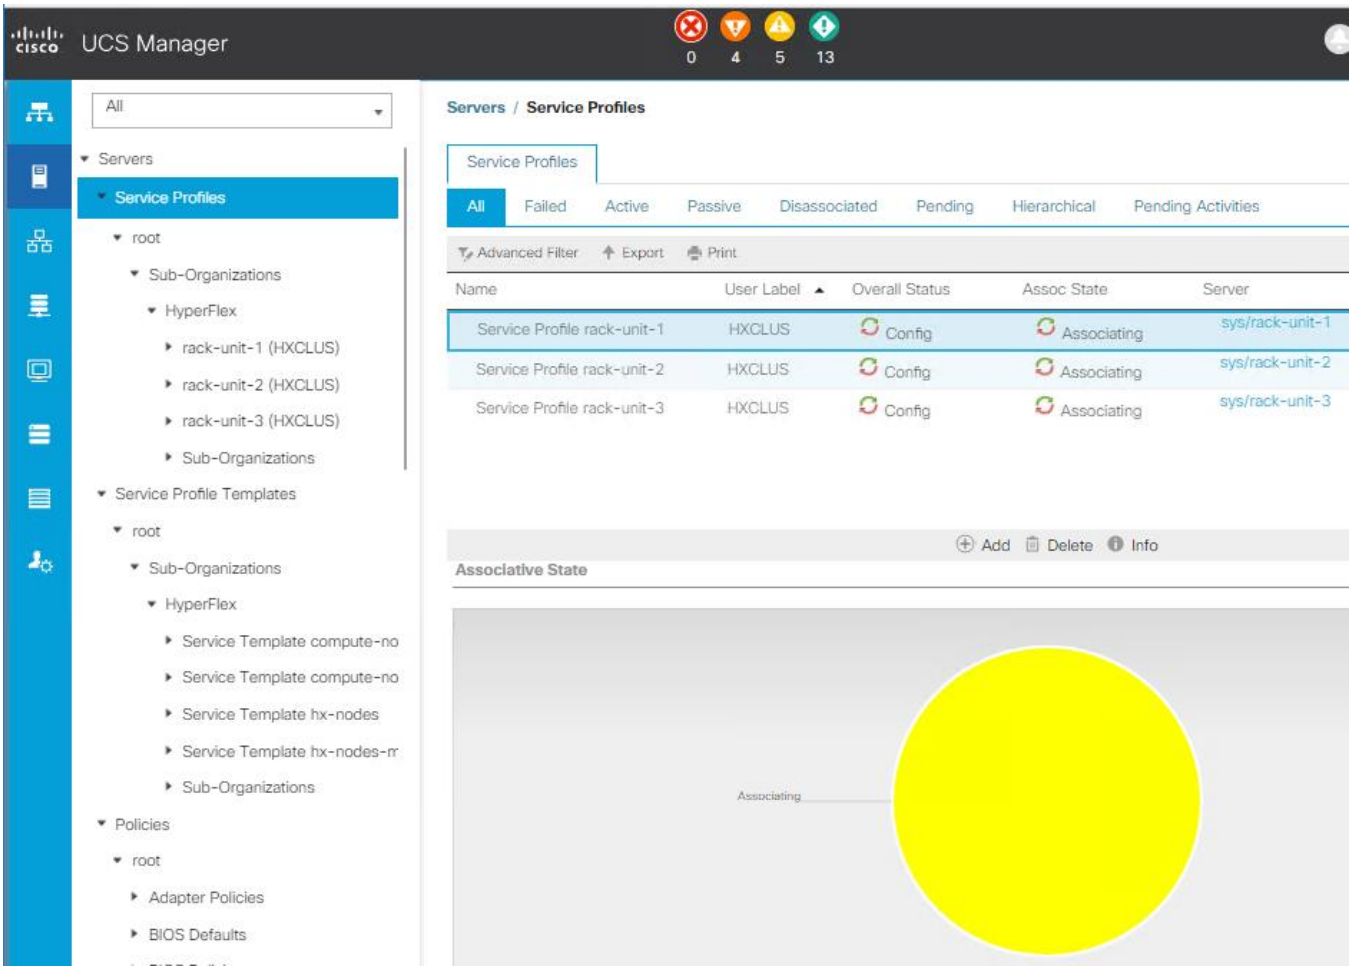


The figure below shows the UCSM > Equipment > Inventory > CIMC with two images in mounted state during the Hypervisor configuration stage in HyperFlex Installer. One is Windows Server 2016 ISO image for OS installation and the second one (latest.img) image file for preparing the system for hx installation after the OS installation is complete.

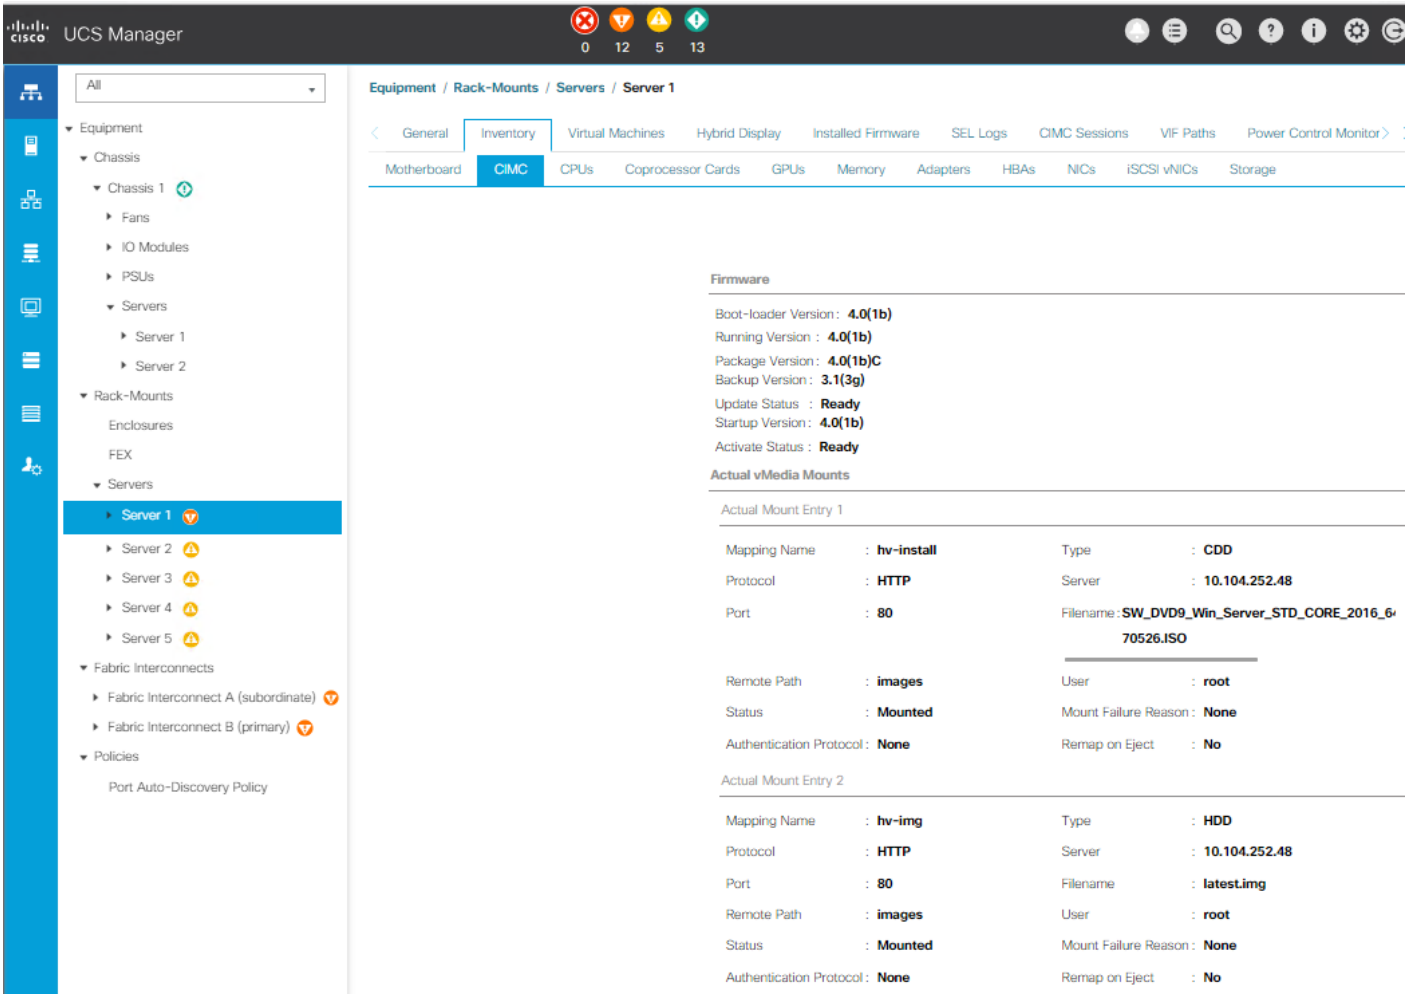


The figure below shows the HyperFlex Installer > Hypervisor Configuration in progress:

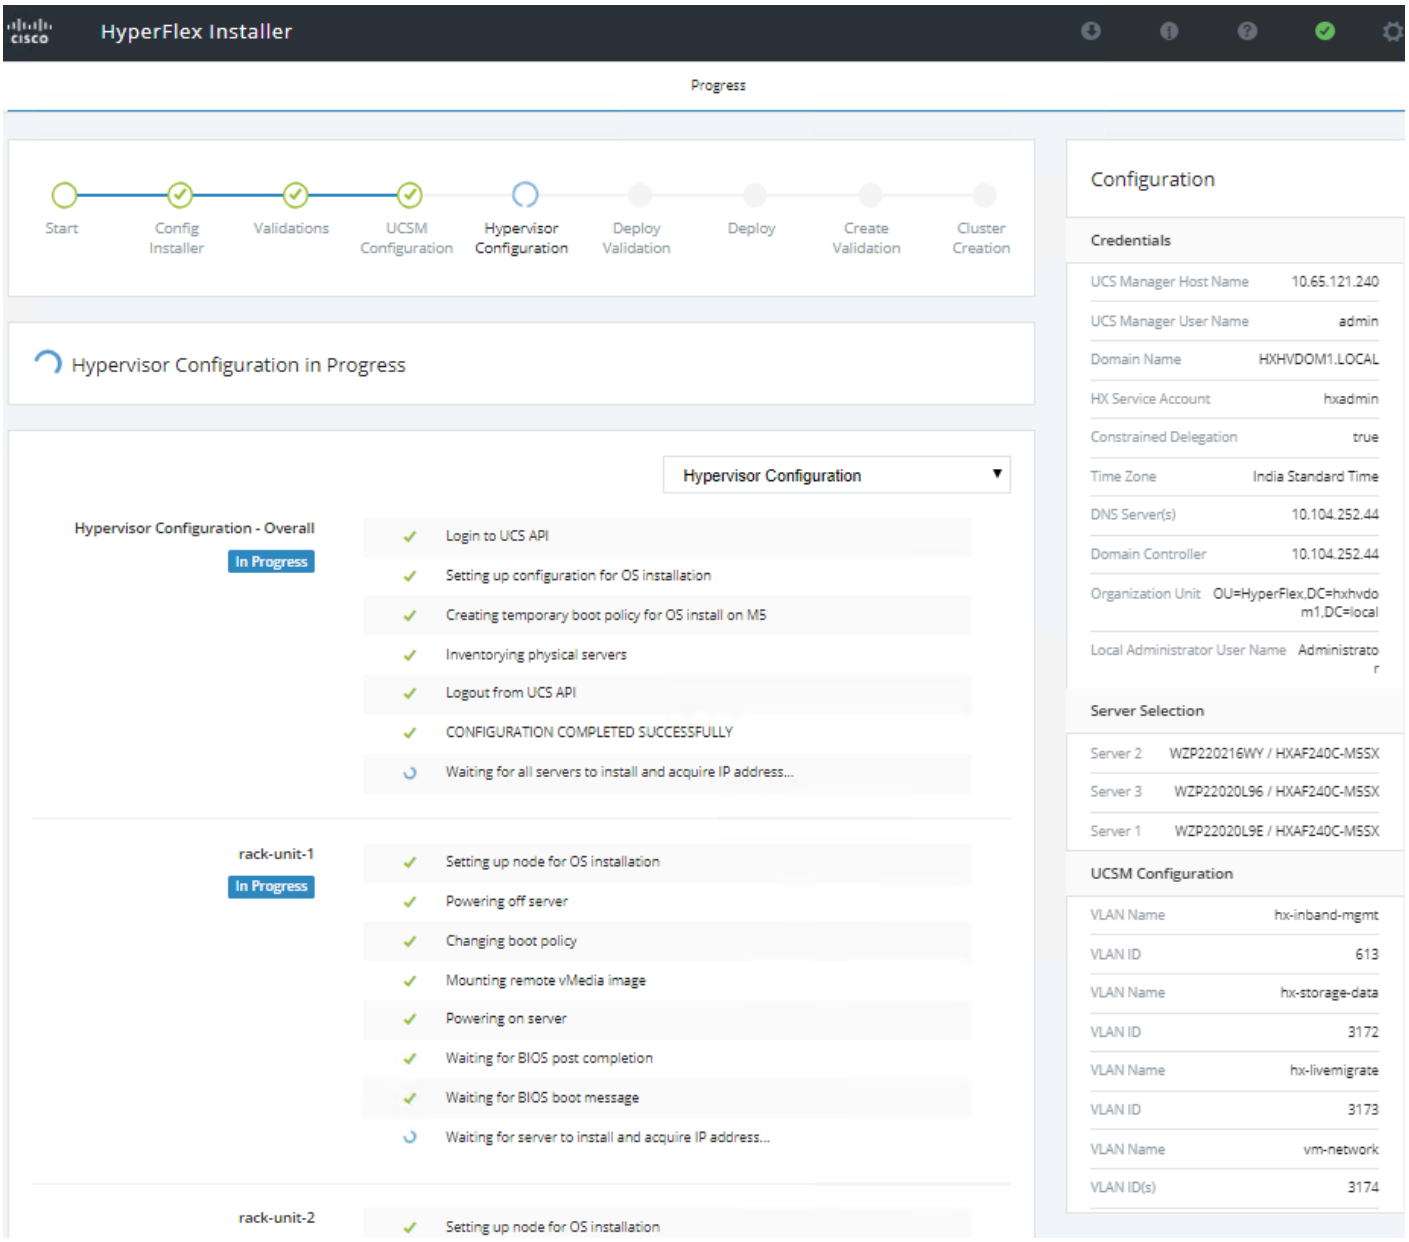


The figure below shows the OS installation in progress in the background during the Hypervisor configuration stage of HyperFlex Installer:

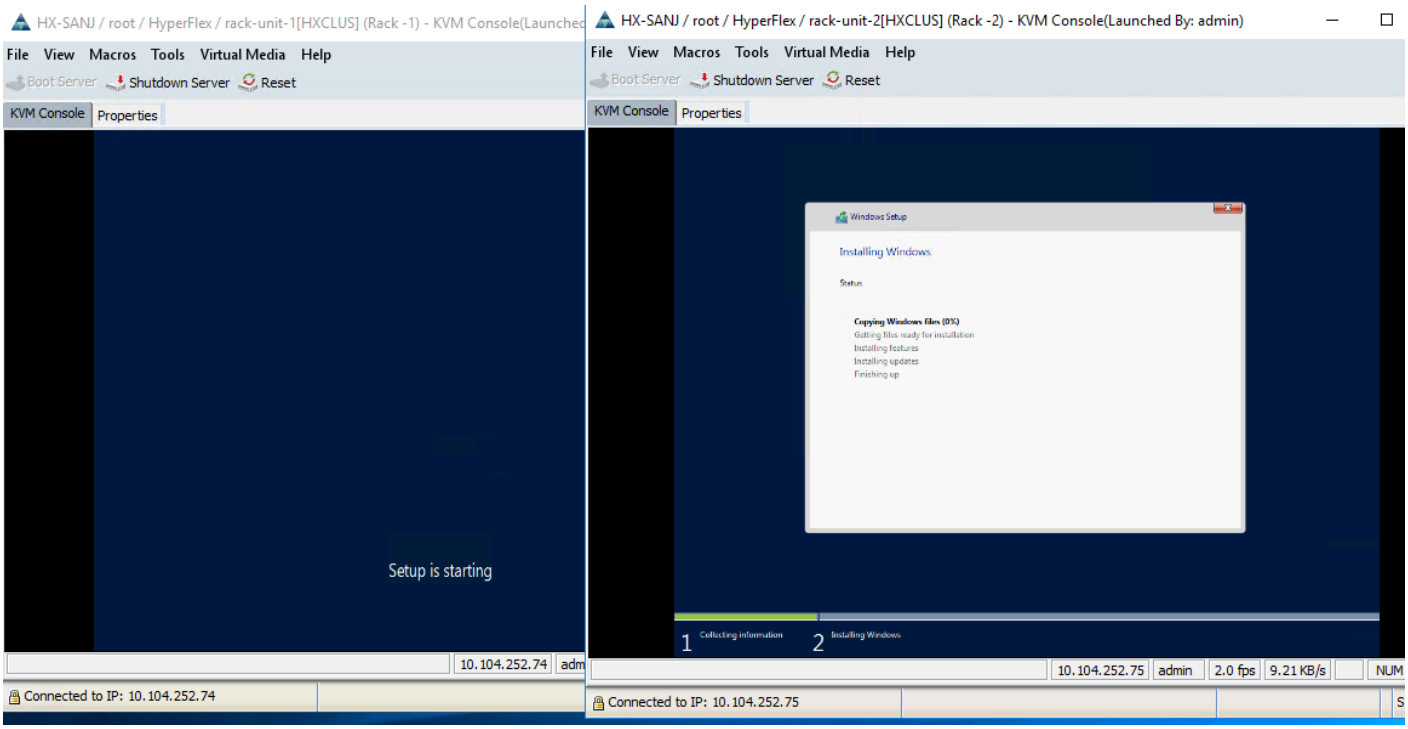


The figure below shows the HyperFlex Installer > Deploy in progress:

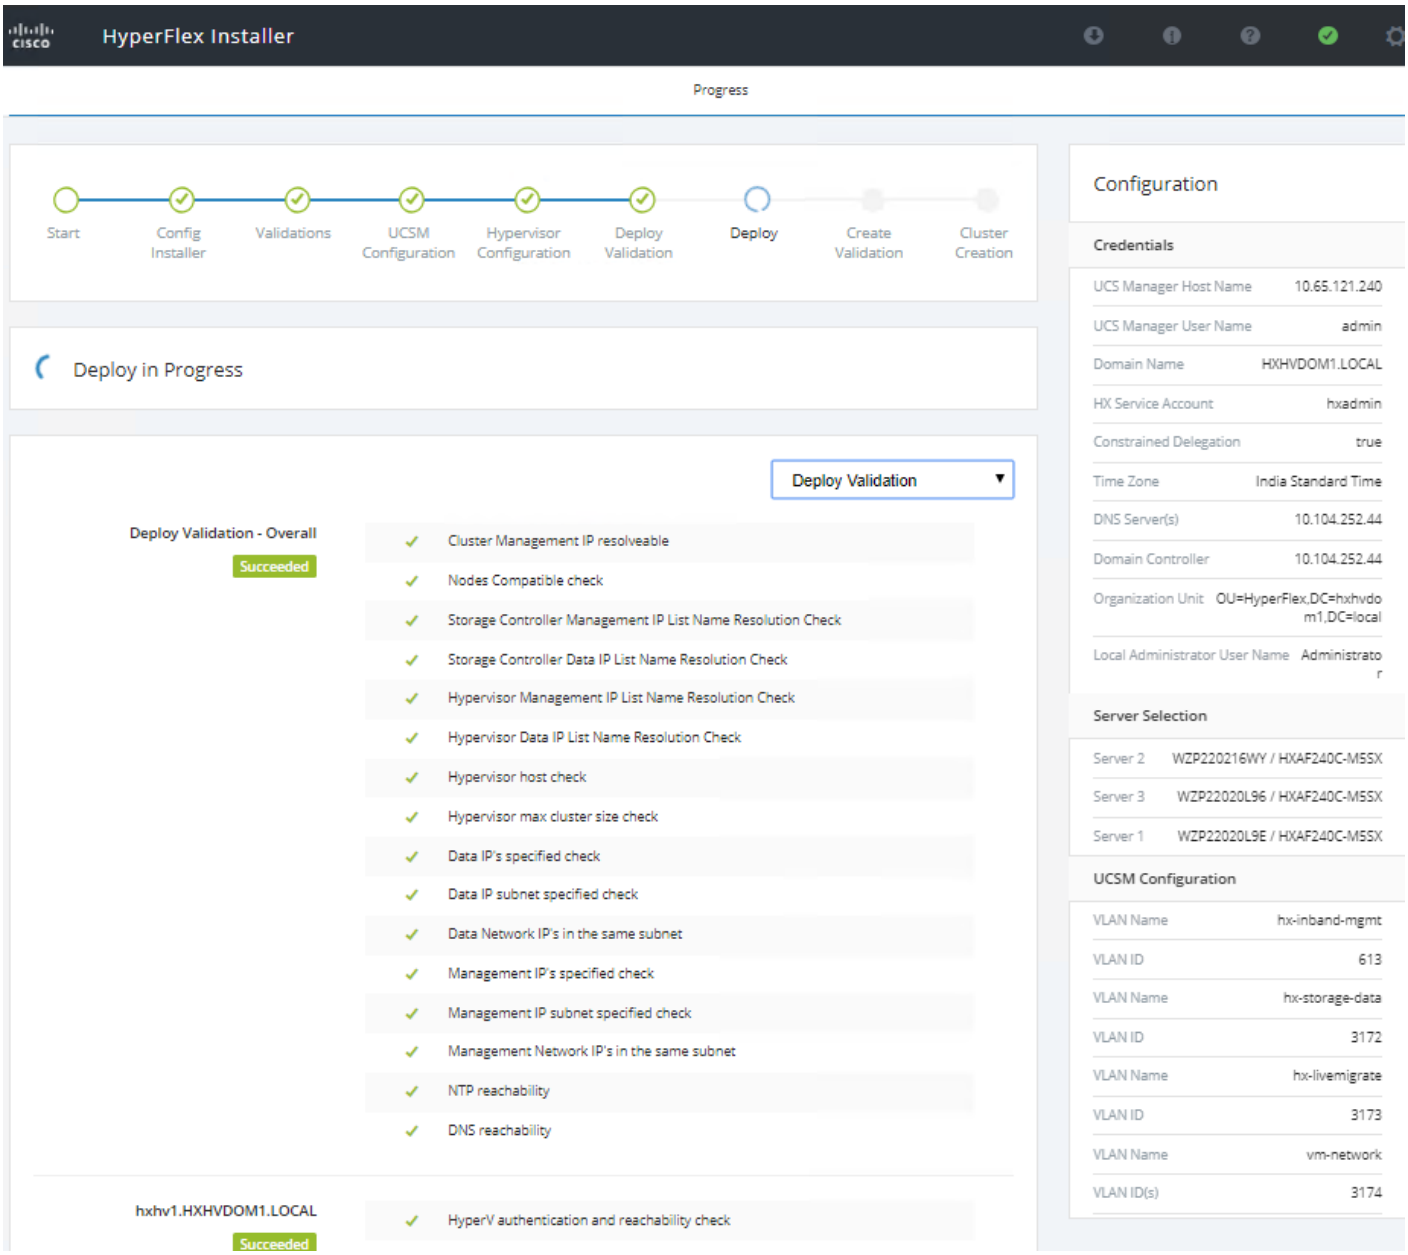


15. Review the Summary screen after the installation completes by selecting Summary.

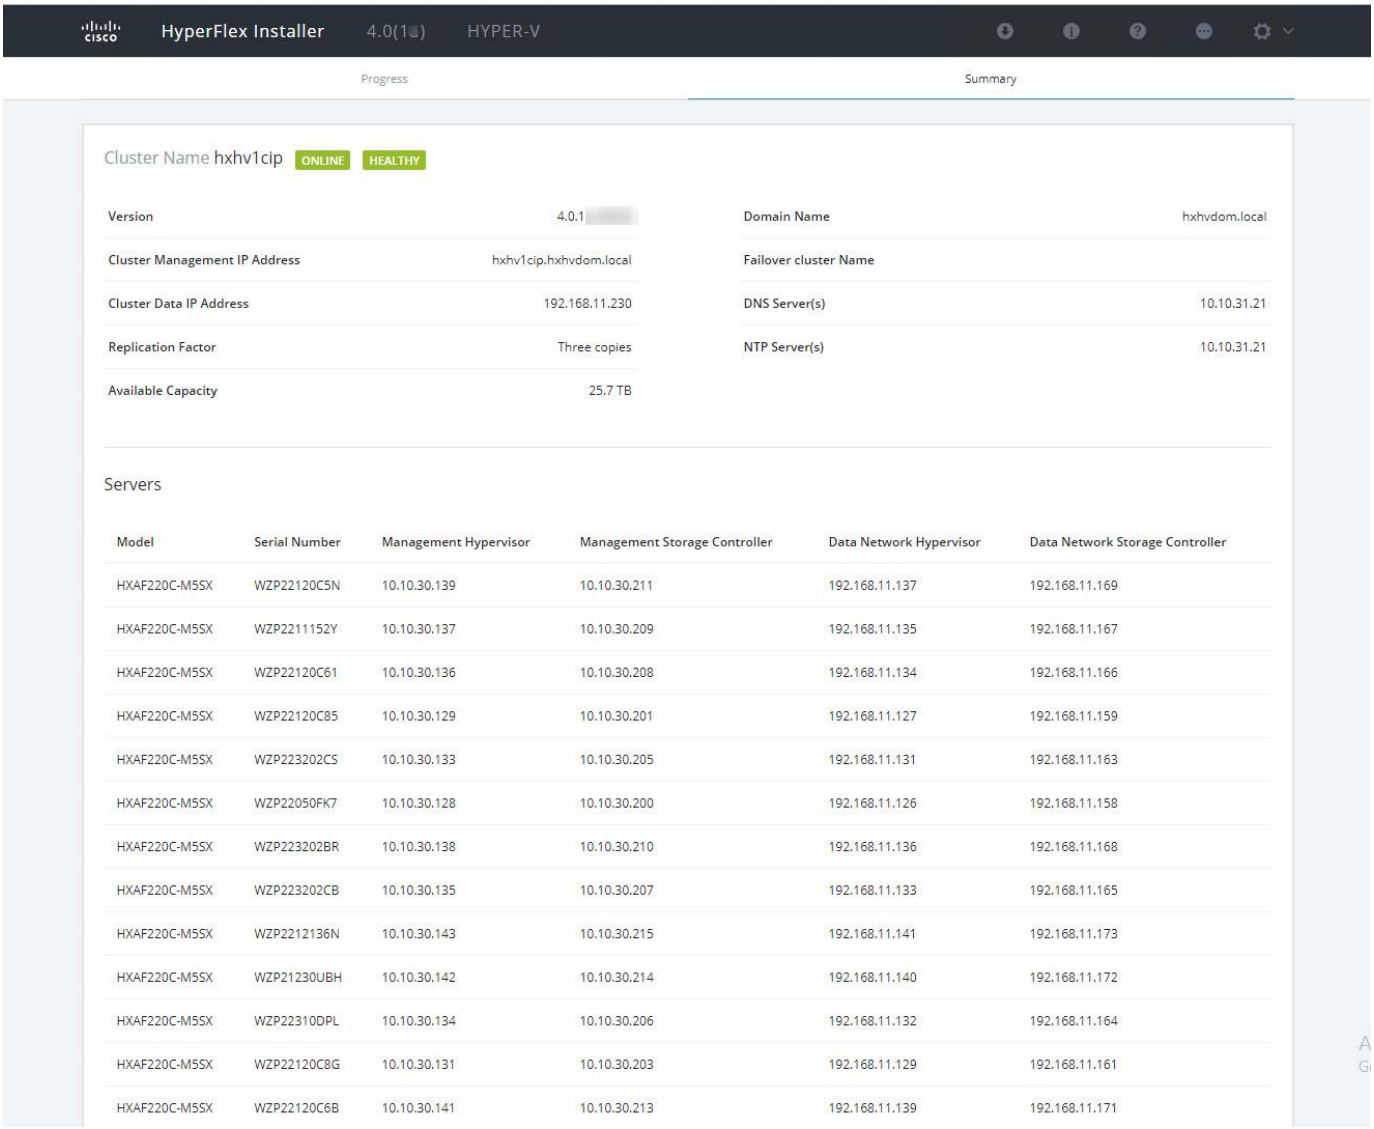


16. After the install completes, you may export the cluster configuration by clicking on the downward arrow icon in the top right of the screen. Click OK to save the configuration to a JSON file. This file can be imported to save time if you need to rebuild the same cluster in the future and be kept as a record of the configuration options and settings used during the installation.

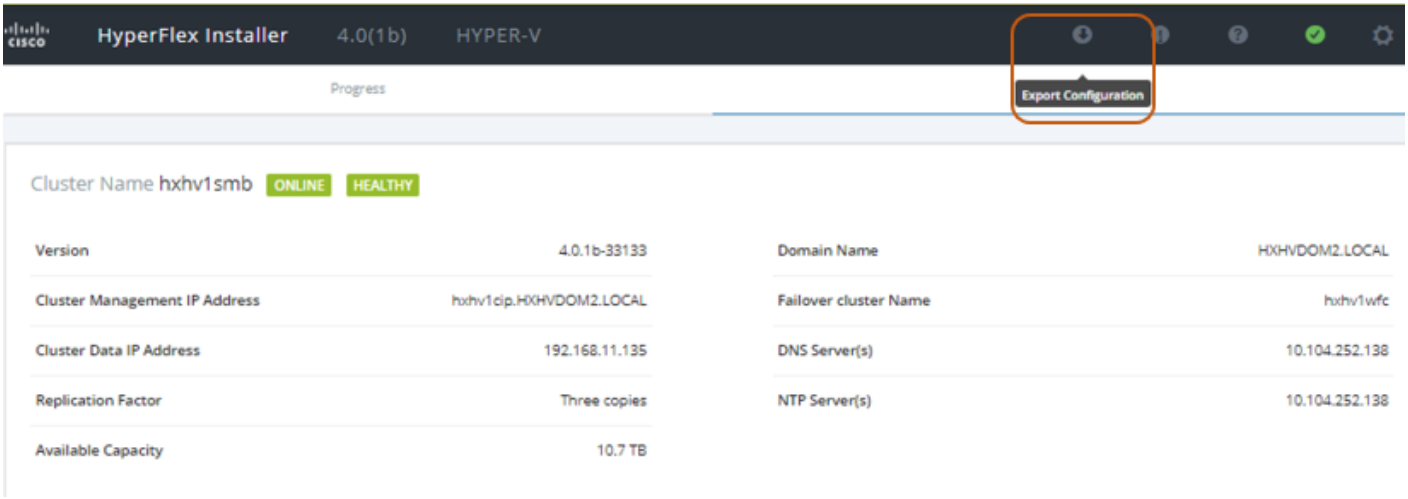


17. After the installation completes, you can click Launch HyperFlex Connect to immediately log into the HTML5 management GUI.

You can go to the cluster expansion immediately after this standard cluster deployment:

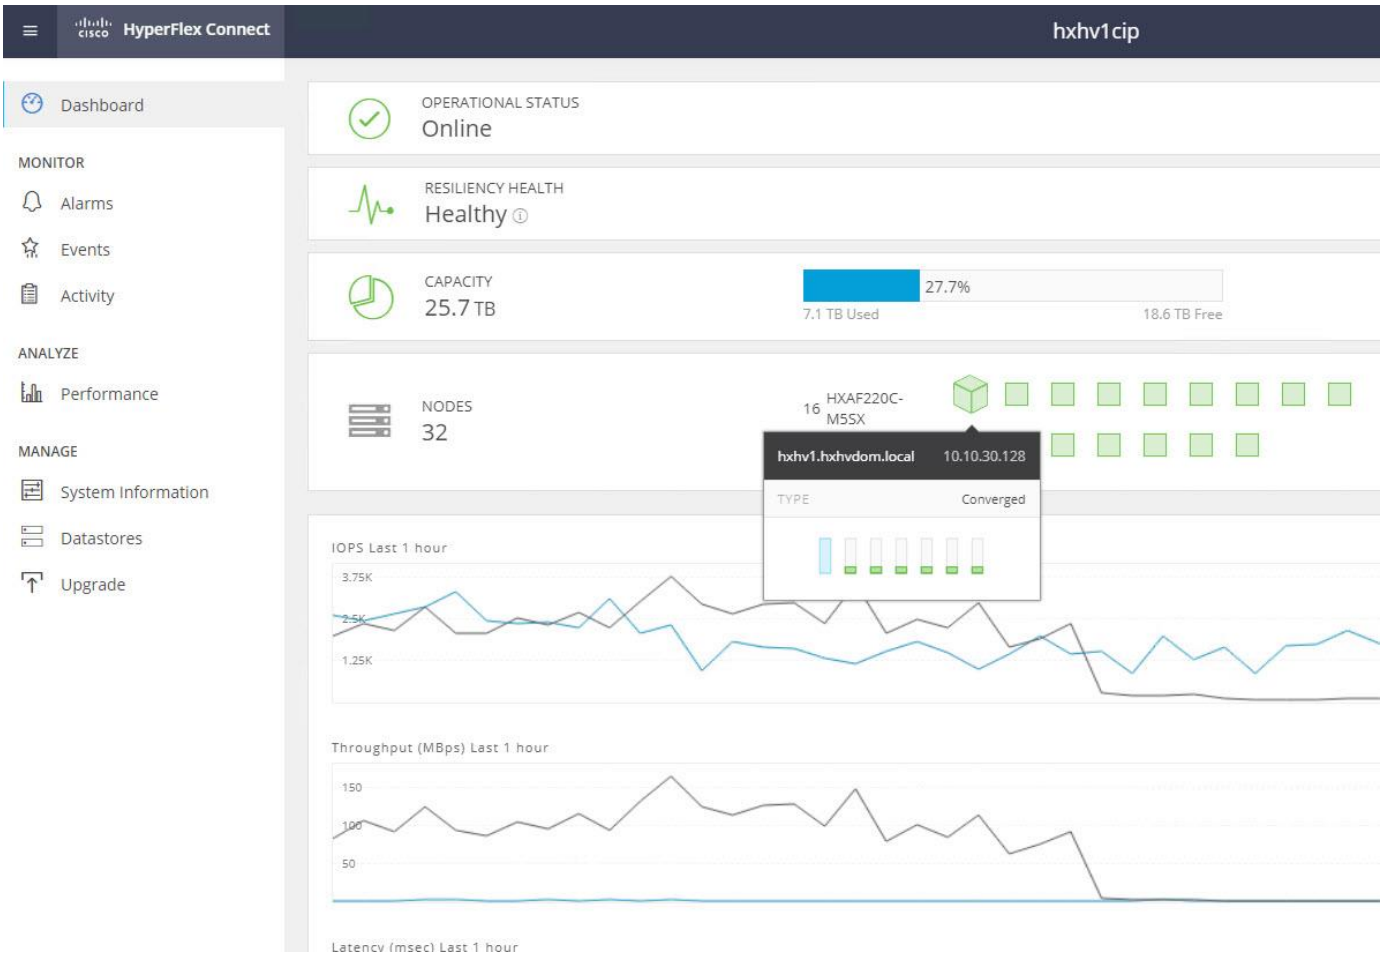

# Post Installation Tasks

## Create Datastores

Create a datastore for storing the virtual machines. This task can be completed by using the HyperFlex Connect HTML management webpage. The Datastores created using HX Connect creates a SMB share which the HyperFlex Hyper-V nodes can use it to store virtual machine files. To configure a new datastore through the HyperFlex Connect webpage, follow these steps:

**Cisco recommends an 8K block size for best performance and as few datastores as possible for ease of management.**

1. Use a web browser to open the HX cluster IP management URL.
2. Enter the credentials.
3. Click Login.
4. Click Datastores in the left pane and click Create Datastore.

**Figure 66 HX Connect - Datastores**

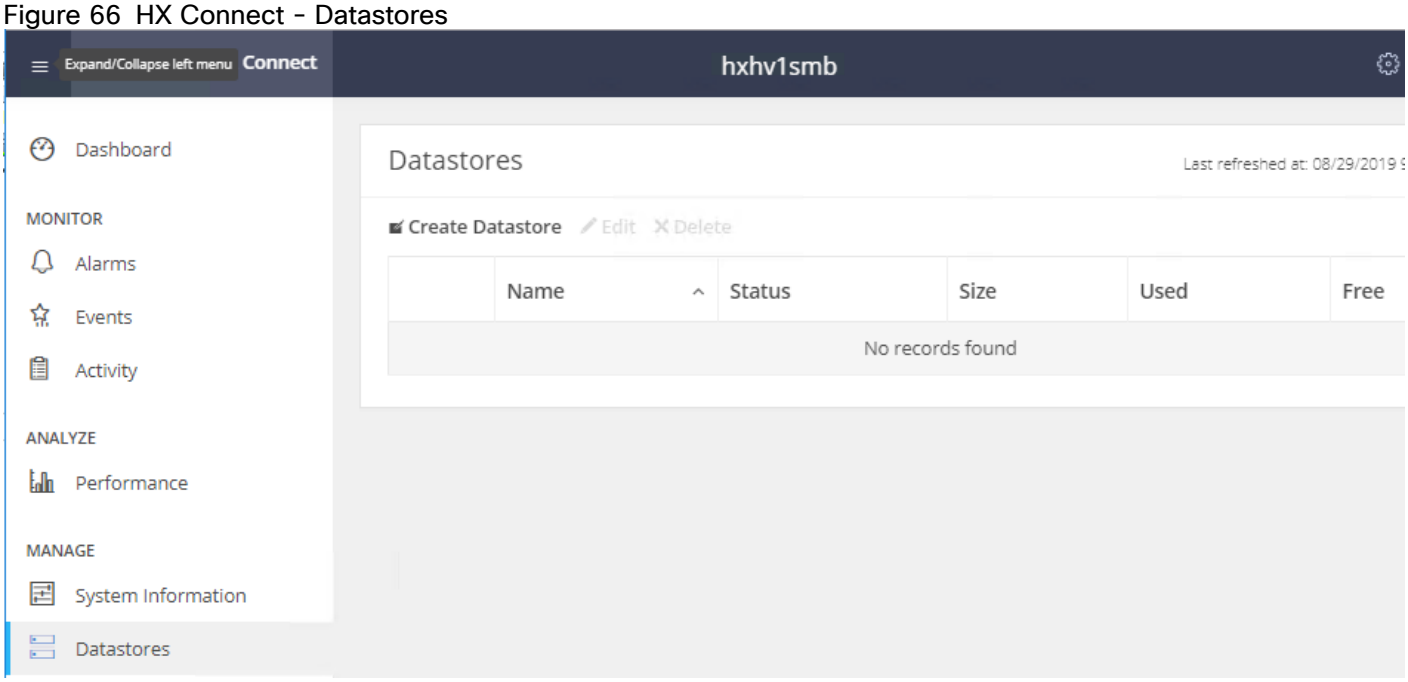


5. In the popup, enter the Datastore Name and size. For most applications, leave the Block Size at the default of 8K. Only dedicated Virtual Desktop Infrastructure (VDI) environments should choose the 4K Block Size option.

Figure 67 HX Connect – Create Datastore

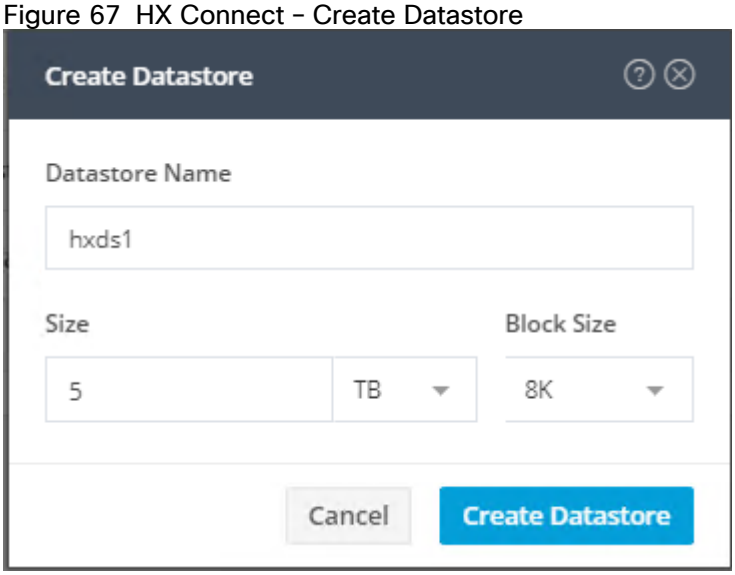


6. Click Create Datastore.

Figure 68 HX Connect – Datastore Status

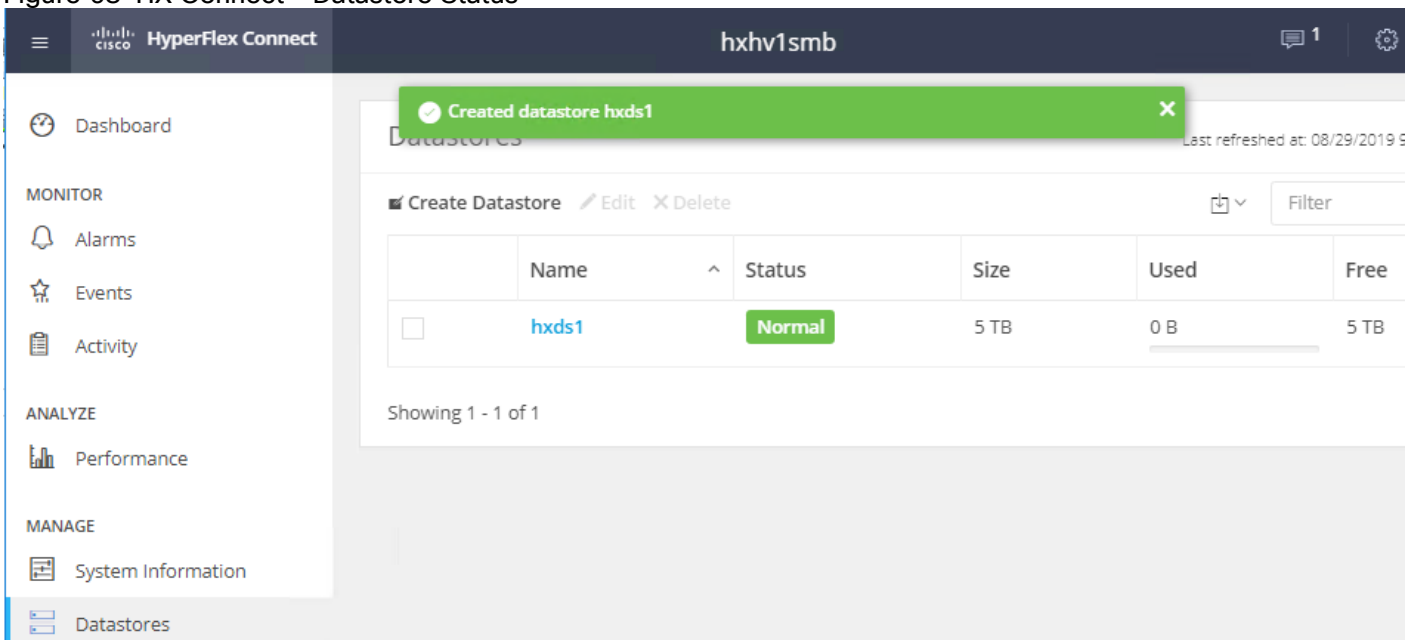


## Constrained Delegation (Optional)

Windows provides a safer form of delegation that could be used by services. When it is configured, constrained delegation restricts the services to which the specified server can act on the behalf of a user. In other words. Constrained Delegation gives granular control over impersonation. When the remote management requests are made to the Hyper-V hosts, it needs to make those requests to the storage on behalf of the caller. This is allowed if that host is trusted for delegation for the CIFS service principal of HX Storage.

Constrained Delegation requires that the option for the security setting User Account Control: Behavior of the elevation prompt for Administrators in Admin Approval Mode is set to Elevate without Prompting. This will prevent the global AD policy from overriding policy on HX OU.

**This step must be performed only if Constrained Delegation was not configured during initial installation. It is recommended that you perform this procedure using the HX Installer and not as part of post-installation.**

To configure constrained delegation using the domain administrator, follow these steps on each Hyper-V host in the HX Cluster and on management hosts (with RSAT tools from where you want to remotely perform administrator tasks):

1. Open the Active Directory Users and Computers snap-in. (From Server Manager, select the server if it is not selected, click Tools >> Active Directory Users and Computers).

2. From the navigation pane in Active Directory Users and Computers, select the domain and double-click the Computers folder.

3. From the Computers folder, right-click the computer account of the source server and then click Properties.

Figure 69 Active Directory Users and Computers

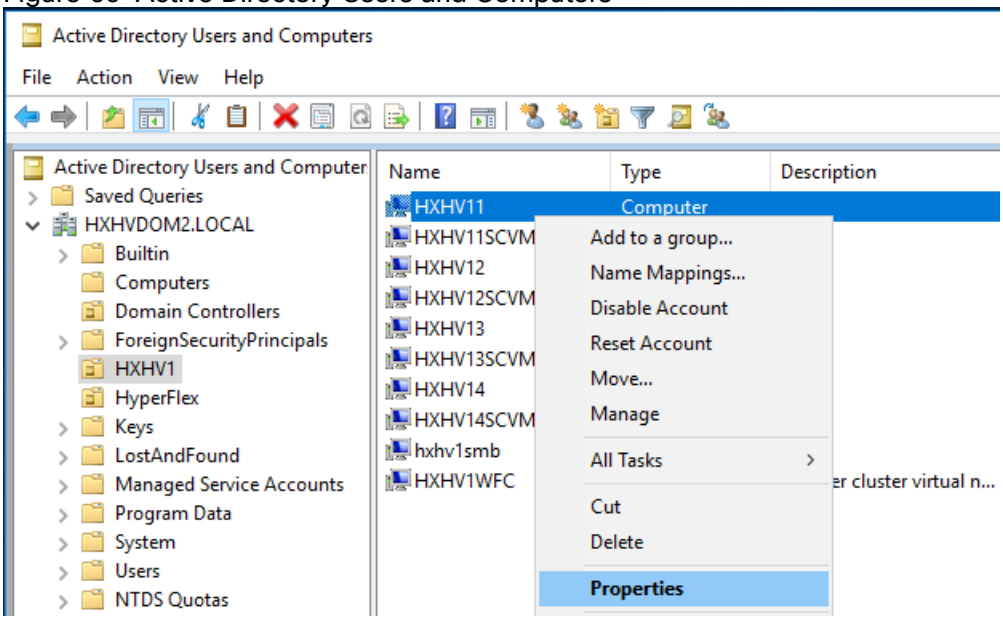


4. From the Properties tab, click the Delegation tab.

5. On the delegation tab, select Trust this computer for delegation to the specified services only and then select Use any authentication protocol.

Figure 70 Active Directory Users and Computers – Server Properties

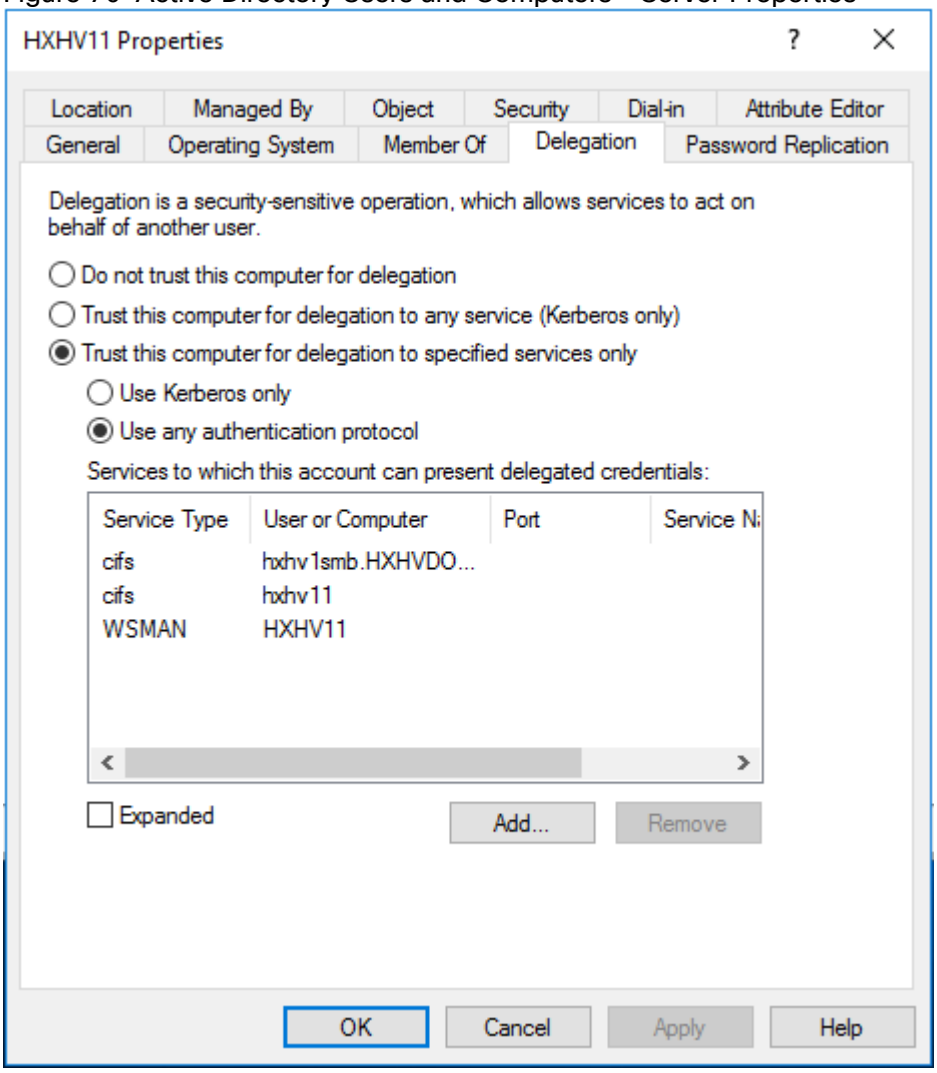


6. Click Add.
7. From Add Services, click Users or Computers.
8. From Select Users or Computers, type the name of the destination server.
9. Click Check Names to verify it and then click OK.
10. From Add Services, in the list of available services, do the following and then click OK:
    a. To move virtual machine storage, select cifs. This is required if you want to move the storage along with the virtual machine, as well as if you want to move only a virtual machine's storage. If the server is configured to use SMB storage for Hyper-V, this should already be selected.
    b. To move virtual machines, select Microsoft Virtual System Migration Service.

Figure 71 Active Directory Users and Computers – Constrained Delegation

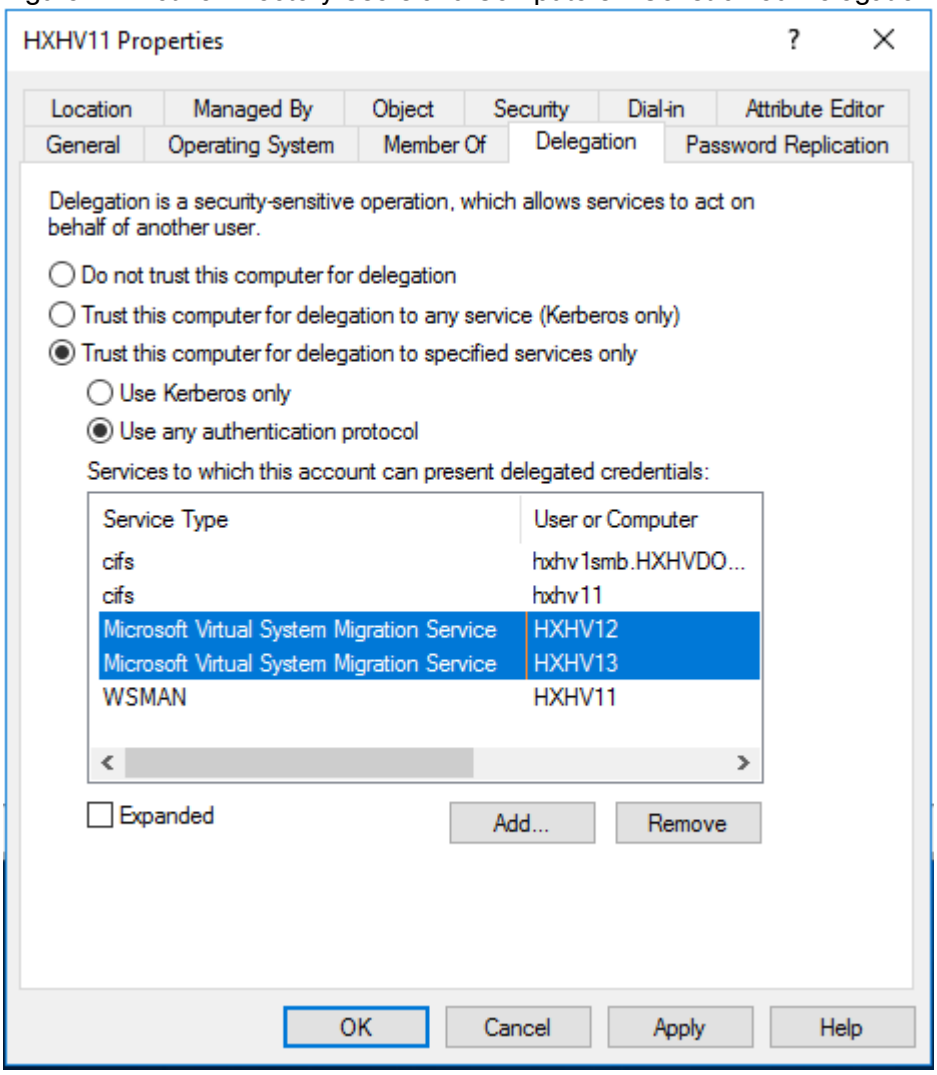


11. On the Delegation tab of the Properties dialog box, verify that the services you selected in the previous step are listed as the services to which the destination computer can present delegated credentials. Click OK.

12. From the Computers folder, select the computer account of the destination server and repeat the process. In the Select Users or Computers dialog box, be sure to specify the name of the source server.

## Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces

To assign a static IP address to Live Migration and Network Interfaces, log into each Hyper-V node and execute the following commands in PowerShell, follow these steps:

1. Use the following PowerShell command (from a remote management station) to check if there is vSwitch created for Live Migration network on Hyper-V hosts by the HX installer:

```
Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Get-VMSwitch -
name * | select PSComputerName,Name} | Format-Table -AutoSize
```

Figure 72 PowerShell - Get VMSwitch

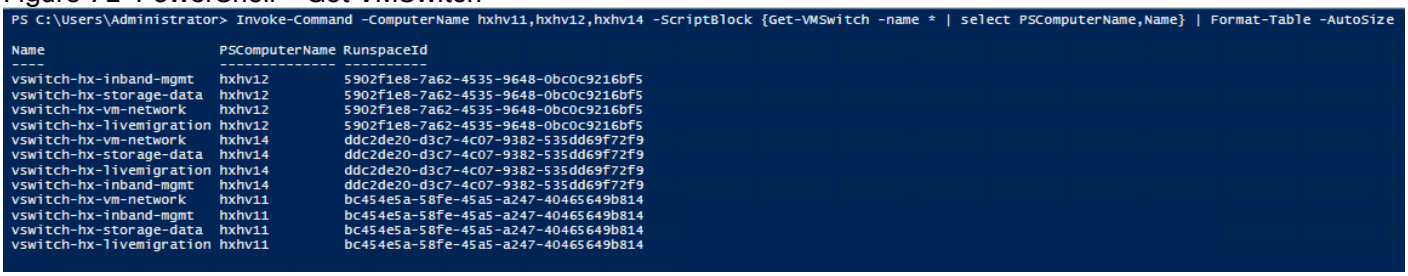


2. Optional - remove the vSwitch named 'vswitch-hx-livemigration' using the following PowerShell command:

```
Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Remove-VMSwitch -
Name vswitch-hx-livemigration}
```

Figure 73 PowerShell - Remove VMSwitch

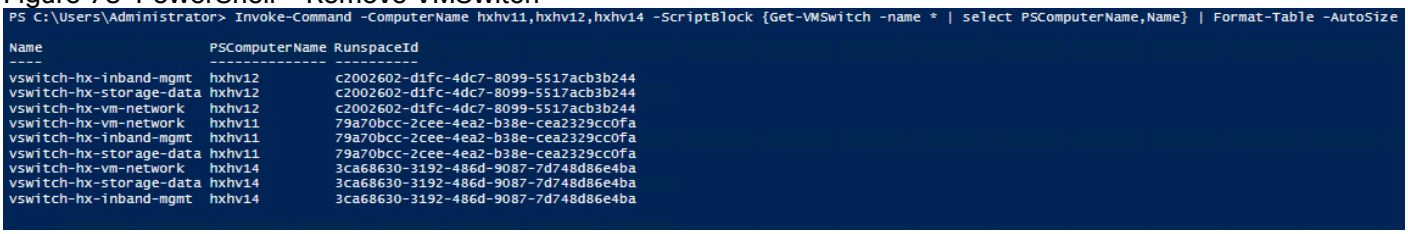


3. Assign a static IP address to the teamed interface named "team-hx-livemigration" or "vswitch-hx-livemigration" using the following PowerShell command:

```
Invoke-Command -ComputerName hxhv11 -ScriptBlock {New-NetIPAddress -ifAlias
"team-hx-livemigration" -IPAddress 192.168.73.127 -PrefixLength 24}
```

Figure 74 PowerShell – Assign Static IP

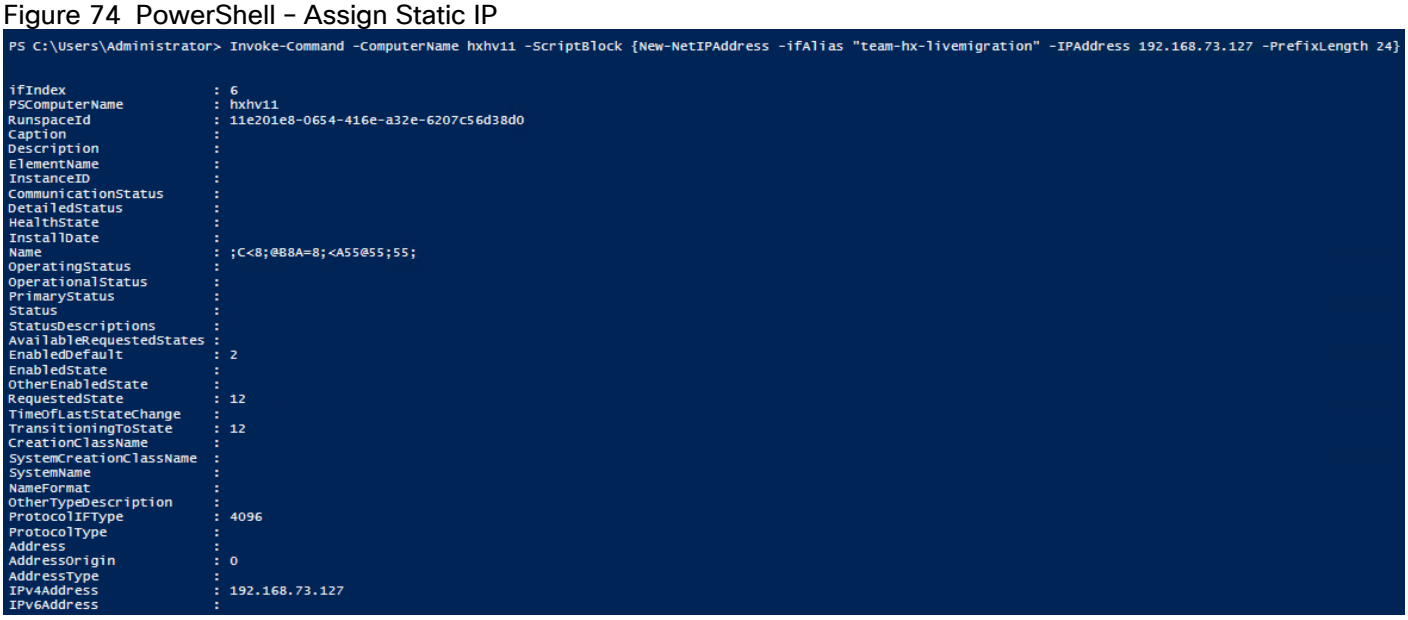


4. This step is optional. If there is a requirement for the Hyper-V host also to communicate on virtual machine network, then assign a static IP address to "team-hx-livemigration" using the following PowerShell command:

```
Invoke-Command -ComputerName hxhv11 -ScriptBlock {New-NetIPAddress -ifAlias
"vswitch-hx-vm-network" -IPAddress 172.18.1.127 -PrefixLength 16}
```

5. Repeat steps 3 and 4 to assign static IP addresses to the live migration and VM networks on all the Hyper-V hosts

6. Verify the jumbo frame settings on the live migration network adapters using the below command:

```
Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Get-
NetAdapterAdvancedProperty -name hv-livemigrate* | Where-Object {$_.DisplayName -
Match "Jumbo*"}}
```

Figure 75 PowerShell – Verify Jumbo Frame Settings

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Get-NetAdapterAdvancedProperty -name hv-livemigrate* | Where-Object {$_.DisplayName -Match "Jumbo*"}}

Name                  DisplayName           DisplayValue          RegistryKeyword RegistryValue   PSComputerName
----                  -----------           ------------          --------------- -------------   --------------
hv-livemigrate-a      Jumbo Packet          Bytes 1514            *JumboPacket    {1514}          hxhv11
hv-livemigrate-b      Jumbo Packet          Bytes 1514            *JumboPacket    {1514}          hxhv11
hv-livemigrate-b      Jumbo Packet          Bytes 1514            *JumboPacket    {1514}          hxhv12
hv-livemigrate-a      Jumbo Packet          Bytes 1514            *JumboPacket    {1514}          hxhv12
hv-livemigrate-b      Jumbo Packet          Bytes 1514            *JumboPacket    {1514}          hxhv14
hv-livemigrate-a      Jumbo Packet          Bytes 1514            *JumboPacket    {1514}          hxhv14
```

7. Configure the jumbo frame settings on the live migration network adapters using the below command to set the mtu size to 9014 Bytes.

```
Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Set-
NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword
"*JumboPacket" -RegistryValue 9014}
```

Figure 76 PowerShell – Configure Jumbo Frame Settings

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Set-NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword "*JumboPacket" -RegistryValue 9014}
```

8. Verify and validate the jumbo frame setting using the below ping command:

```
PS C:\Users\administrator.HXHVDOM2> ping -l 8972 -f 192.168.73.128

Pinging 192.168.73.128 with 8972 bytes of data:
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128

Ping statistics for 192.168.73.128:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Users\administrator.HXHVDOM2> ping -l 8972 -f 192.168.73.130

Pinging 192.168.73.130 with 8972 bytes of data:
Reply from 192.168.73.130: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.130: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.130: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.130: bytes=8972 time<1ms TTL=128

Ping statistics for 192.168.73.130:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Users\administrator.HXHVDOM2> _
```

## Rename the Cluster Network in Windows Failover Cluster - Optional

To rename the default cluster network names assigned during cluster creation to more meaningful names, run the following PowerShell commands from any one HyperFlex Hyper-V host:

1. Run the "*Get-ClusterNetwork*" from one of the Hyper-V nodes as shown below to view information about the cluster network.

**Figure 77 PowerShell – Get Cluster Network**

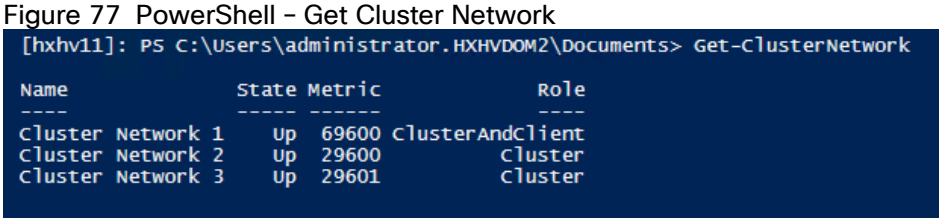


2. Run the below PowerShell command to rename the cluster networks:

```
Get-ClusterNetwork | Where-Object {$_.Address -eq "10.29.149.0"}).Name = "hx-inband-mgmt"

(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.11.0"}).Name = "hx-storage-data"

(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.73.0"}).Name = "LiveMigration"

(Get-ClusterNetwork | Where-Object {$_.Address -eq "172.18.0.0"}).Name = "vm-network"
```

**Figure 78 PowerShell – Rename the Cluster Network**

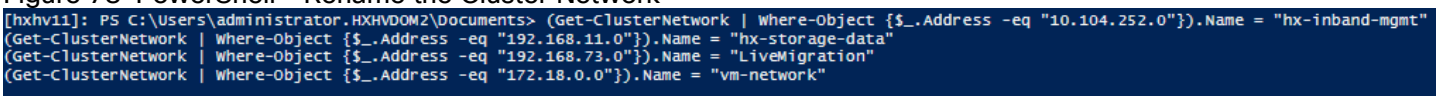


3. Verify now the cluster network using the "*Get-ClusterNetwork*" command

**Figure 79 PowerShell – Verify Cluster Network**

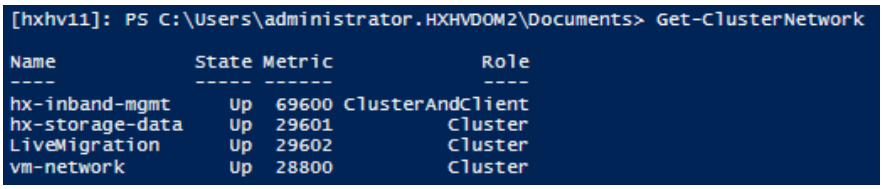


## Configure the Windows Failover Cluster Network Roles

Cluster networks are automatically configured during the cluster creation. To manually configure the cluster network roles based on their type of function, run the following PowerShell commands on any one HyperFlex Hyper-V host:

1. Execute the following PowerShell commands to configure the cluster networks roles:

```
(Get-ClusterNetwork -Name "hx-inband-mgmt").Role = 3

(Get-ClusterNetwork -Name "hx-storage-data").Role = 0
```

```
(Get-ClusterNetwork -Name "LiveMigration").Role = 1

(Get-ClusterNetwork -Name "vm-network").Role = 0
```

Figure 80 PowerShell – Configure Cluster Network Roles

```
[hxhv11]: PS C:\Users\administrator.HXHVDOM2\Documents> (Get-ClusterNetwork -Name "hx-inband-mgmt").Role = 3
(Get-ClusterNetwork -Name "hx-storage-data").Role = 0
(Get-ClusterNetwork -Name "LiveMigration").Role = 1
(Get-ClusterNetwork -Name "vm-network").Role = 0
```

Role = 0 to disable cluster communication

Role = 1 to enable only cluster communication

Role = 3 to enable both cluster & client communication

Figure 81 Failover Cluster Manager – Networks

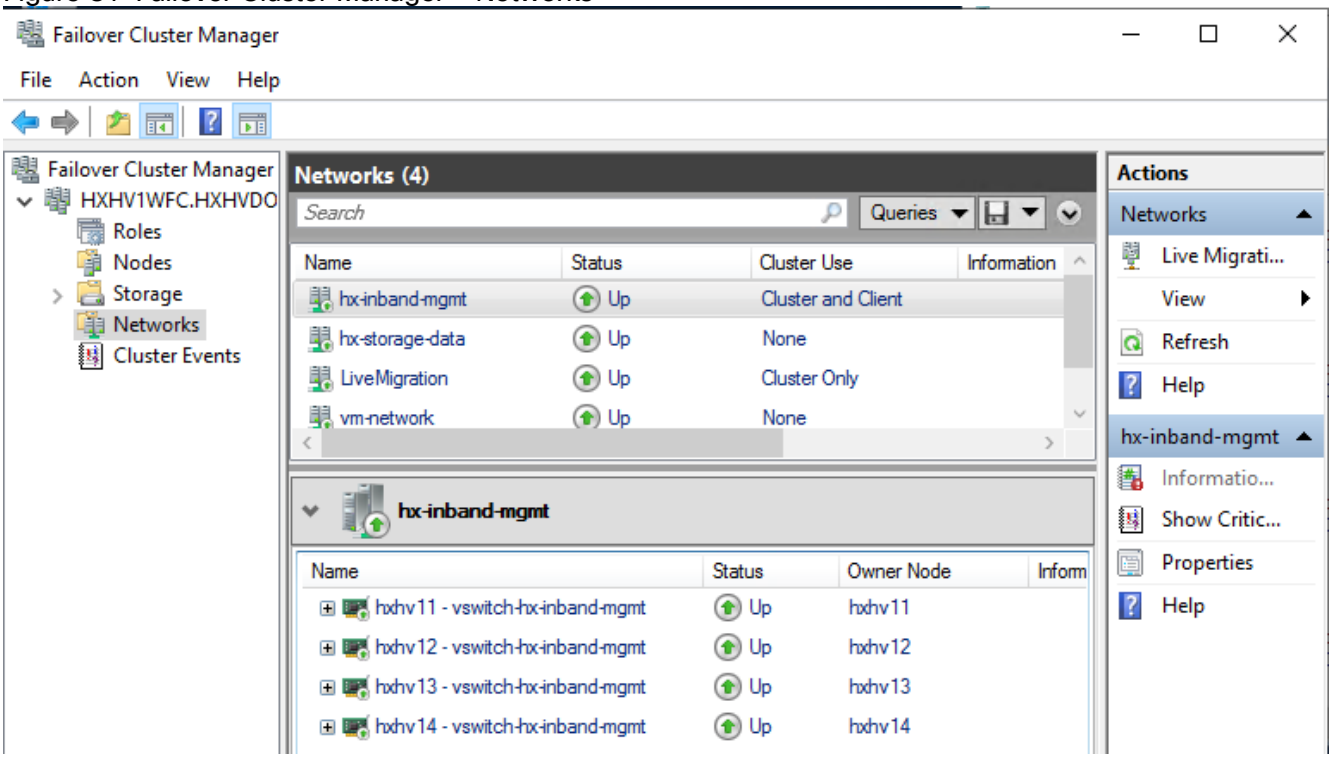


## Configure the Windows Failover Cluster Network for Live Migration

To make sure that you are using the appropriate cluster network for Live Migration traffic configure the Live Migration settings by following these steps:

1. Run the PowerShell command shown below to configure the cluster network for live migration traffic:

```
Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -Name
MigrationExcludeNetworks -Value ([String]::Join(";",(Get-ClusterNetwork | Where-
Object {$_.Name -ne "LiveMigration"}).ID))
```

Figure 82 PowerShell – Configure Live Migration Network

```
[hxhv11]: PS C:\Users\administrator.HXHVDOM2\Documents> Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -Name MigrationExcludeNetworks -Value ([String]::Join(";",(Get-ClusterNetwork | Where-Object {$_.Name -ne "LiveMigration"}).ID))
```

Figure 83 Failover Cluster Manager – Live Migration Settings

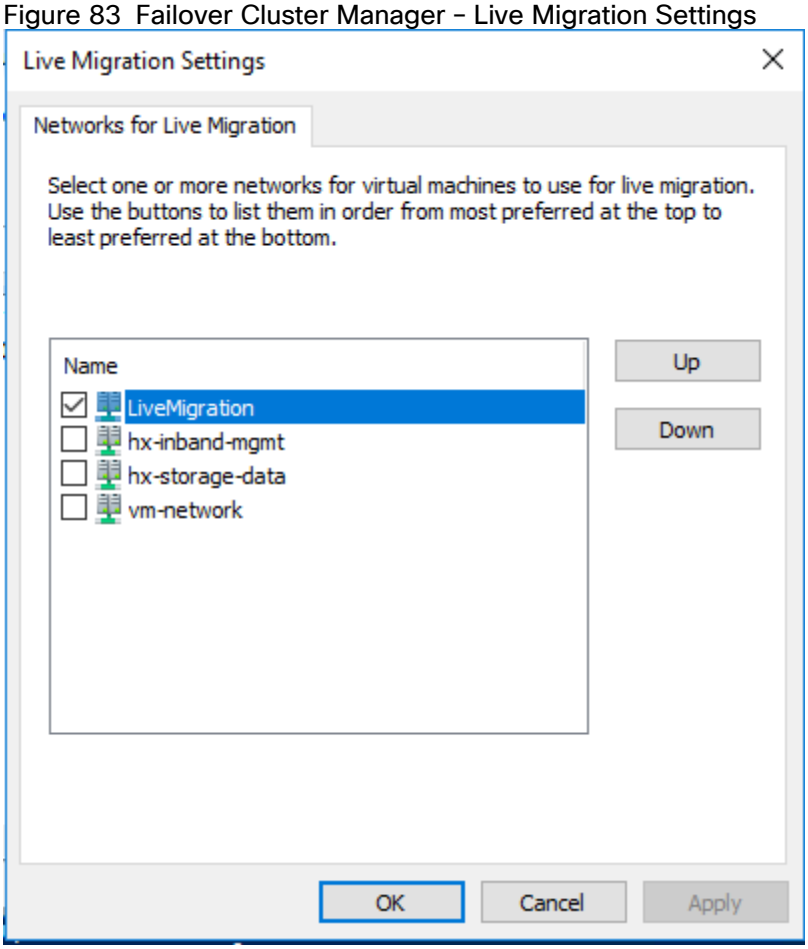


## Create Folders on the HX Datastore

To create folders on the newly created HX Datastore, follow these steps:

1. To create a folder, log in to a HyperFlex Hyper-V node and run the following command:

   ```
   mkdir \\hxhv1smb.hxhvdom.local\hxds1\<folder Name>
   ```

Figure 84 Create a Folder on SMB Share

```
PS C:\Users\administrator.HXHVDOM2> dir \\hxhv1smb.hxhvdom2.local\hxds1\
PS C:\Users\administrator.HXHVDOM2> mkdir \\hxhv1smb.hxhvdom2.local\hxds1\Hyper-V\VHDs


    Directory: \\hxhv1smb.hxhvdom2.local\hxds1\Hyper-V


Mode                 LastWriteTime         Length Name
----                 -------------         ------ ----
d-----         9/3/2019   3:33 PM                VHDs


PS C:\Users\administrator.HXHVDOM2> _
```

2. Create folders for different purposes and requirements.

## Configure the Default Folder to Store Virtual Machine Files on Hyper-V

By default, Hyper-V stores virtual machine files at the following specified locations:

- "C:\ProgramData\Microsoft\Windows\Hyper-V" for virtual machine configuration files
- "C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks" for virtual hard drives

To store the virtual machine files on the newly created highly available HX Datastore as the default folder. Complete the following step on each HyperFlex Hyper-V hosts:

1. Run the following PowerShell command from a remote management station to set/change the Hyper-V default store location for virtual hard disk and virtual machine configuration files on all the Hyper-V hosts and verify the settings as shown below:

```
$hosts = "hxhv11","hxhv12","hxhv13"

Invoke-Command -ComputerName $hosts -ScriptBlock {SET-VMHOST -virtualharddiskpath
"\\hxhv1smb.hxhvdom2.local\hxds1\Hyper-V\VHDs" -virtualmachinepath
"\\hxhv2smb.hxhvdom2.local\HXDS1\Hyper-V\"}
```

Figure 85 PowerShell – Configure Virtual Machine Files Store Location

```
PS C:\Users\Administrator> $hosts = "hxhv11","hxhv12","hxhv14"

PS C:\Users\Administrator> Invoke-Command -ComputerName $hosts -ScriptBlock {SET-VMHOST -virtualharddiskpath "\\hxhv1smb.hxhvdom2.local\hxds1\Hyper-V\VHDs" -virtualmachinepath "\\hxhv1smb.hxhvdom2.local\hxds1\Hyper-V"}

PS C:\Users\Administrator>
```

Figure 86 PowerShell – Configure Virtual Machine Files Store Location

```
PS C:\Users\Administrator> Invoke-Command -ComputerName $hosts -ScriptBlock {Get-VMHost | Select-Object -Property  ComputerName, VirtualHardDiskPath, VirtualMachinePath}


ComputerName        : HXHV12
VirtualHardDiskPath : \\hxhv1smb.hxhvdom2.local\HXDS1\Hyper-V\VHDs
VirtualMachinePath  : \\hxhv1smb.hxhvdom2.local\HXDS1\Hyper-V
PSComputerName      : hxhv12
RunspaceId          : 53d20d3a-8907-48d0-ac63-1e97b5d80f1e

ComputerName        : HXHV11
VirtualHardDiskPath : \\hxhv1smb.hxhvdom2.local\HXDS1\Hyper-V\VHDs
VirtualMachinePath  : \\hxhv1smb.hxhvdom2.local\HXDS1\Hyper-V
PSComputerName      : hxhv11
RunspaceId          : 2077babb-fe03-4ab2-98e5-f0cc6ff0cfa3

ComputerName        : HXHV14
VirtualHardDiskPath : \\hxhv1smb.hxhvdom2.local\HXDS1\Hyper-V\VHDs
VirtualMachinePath  : \\hxhv1smb.hxhvdom2.local\HXDS1\Hyper-V
PSComputerName      : hxhv14
RunspaceId          : 6eef39c1-ee2c-43be-bdf6-b7e726503dfc
```

Figure 87 Hyper-V Settings – VHD Store Location

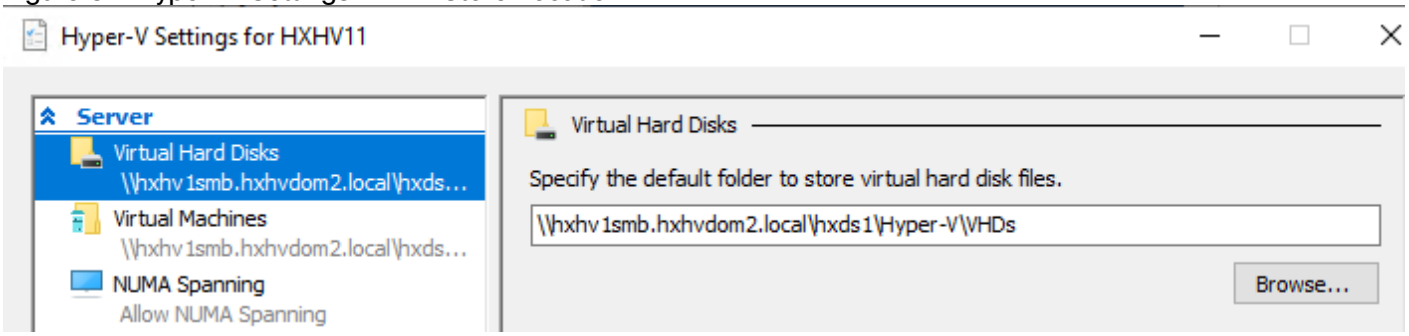


## Validate the Windows Failover Cluster Configuration

It is a good practice to validate the Windows failover cluster by running the Validate a Configuration Wizard from the Failover Cluster Manager, or the Test-Cluster Windows PowerShell cmdlet and fix any errors or warnings reported in the results page. Figure 88 shows the command to run the cluster validation.

Figure 88 Failover Cluster – Validate

```
PS C:\Users\administrator.HXHVDOM2> Test-Cluster
WARNING: Cluster Configuration - Validate Quorum Configuration: The test reported some warnings..
WARNING: Cluster Configuration - Validate Resource Status: The test reported some warnings..
WARNING: System Configuration - Validate Software Update Levels: The test reported some warnings..
WARNING:
Test Result:
HadUnselectedTests, ClusterConditionallyApproved
Testing has completed for the tests you selected. You should review the warnings in the Report.  A cluster solution is supported by Microsoft only
if you run all cluster validation tests, and all tests succeed (with or without warnings).
Test report file path: C:\Users\administrator.HXHVDOM2\AppData\Local\Temp\2\Validation Report 2019.09.05 At 17.28.48.htm

Mode                 LastWriteTime         Length Name
----                 -------------         ------ ----
-a----         9/5/2019   5:32 PM        1398579 Validation Report 2019.09.05 At 17.28.48.htm
```

## Configure Quorum in Windows Server Failover Cluster

The quorum is automatically configured during the creation of a new cluster based on the number of nodes and the availability of shared storage. However, as a best practice, run the cluster validation tool as shown in the above section and review the quorum configuration and fix any warnings related to quorum configuration.

Review the information about quorum resources using the “Get-ClusterQuorum” PowerShell cmdlet or from the Summary page of failover cluster manager as shown below:

Figure 89 PowerShell – Get Cluster Quorum

```
PS C:\Users\Administrator> Get-ClusterQuorum -Cluster hxhv1wfc.hxhvdom2.local

Cluster              QuorumResource
-------              --------------
HXHV1WFC
```

Figure 90 Failover Cluster Manager

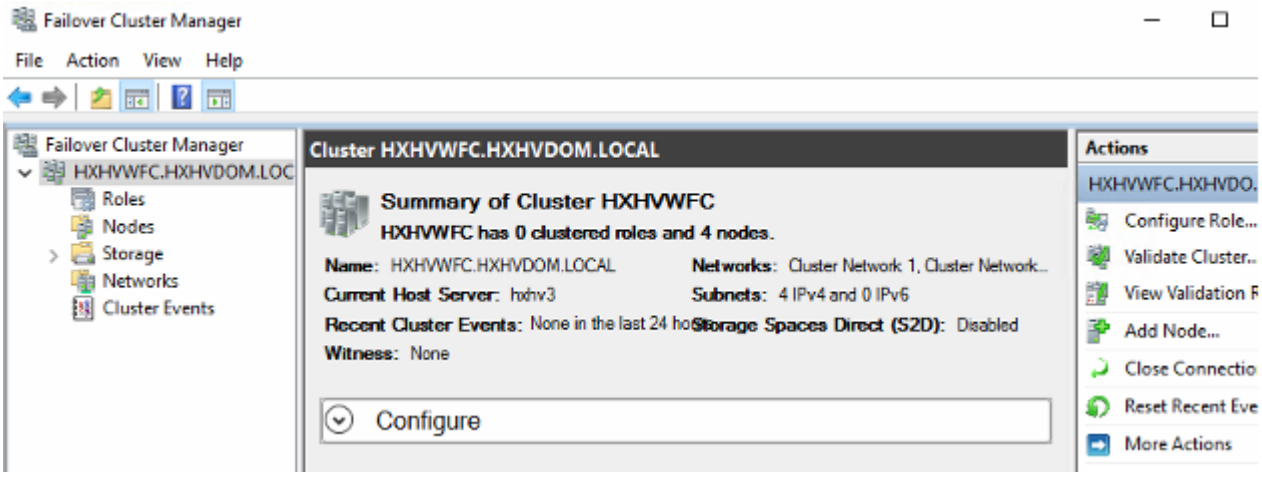


Run the following PowerShell command to configure the cluster quorum by placing the witness on a file share residing on the HX Datastore:

```
Set-ClusterQuorum -NodeAndFileShareMajority "\\fileserver\fsw"
```

Figure 91 PowerShell – Configure Cluster Quorum

```
PS C:\Users\Administrator> Set-ClusterQuorum -Cluster hxhv1wfc.hxhvdom2.local -NodeAndFileShareMajority \\hxhv2dc\Cluster_FSW

Cluster              QuorumResource
-------              --------------
HXHV1WFC             File Share Witness
```

Figure 92 Failover Cluster Manager – Quorum Configuration

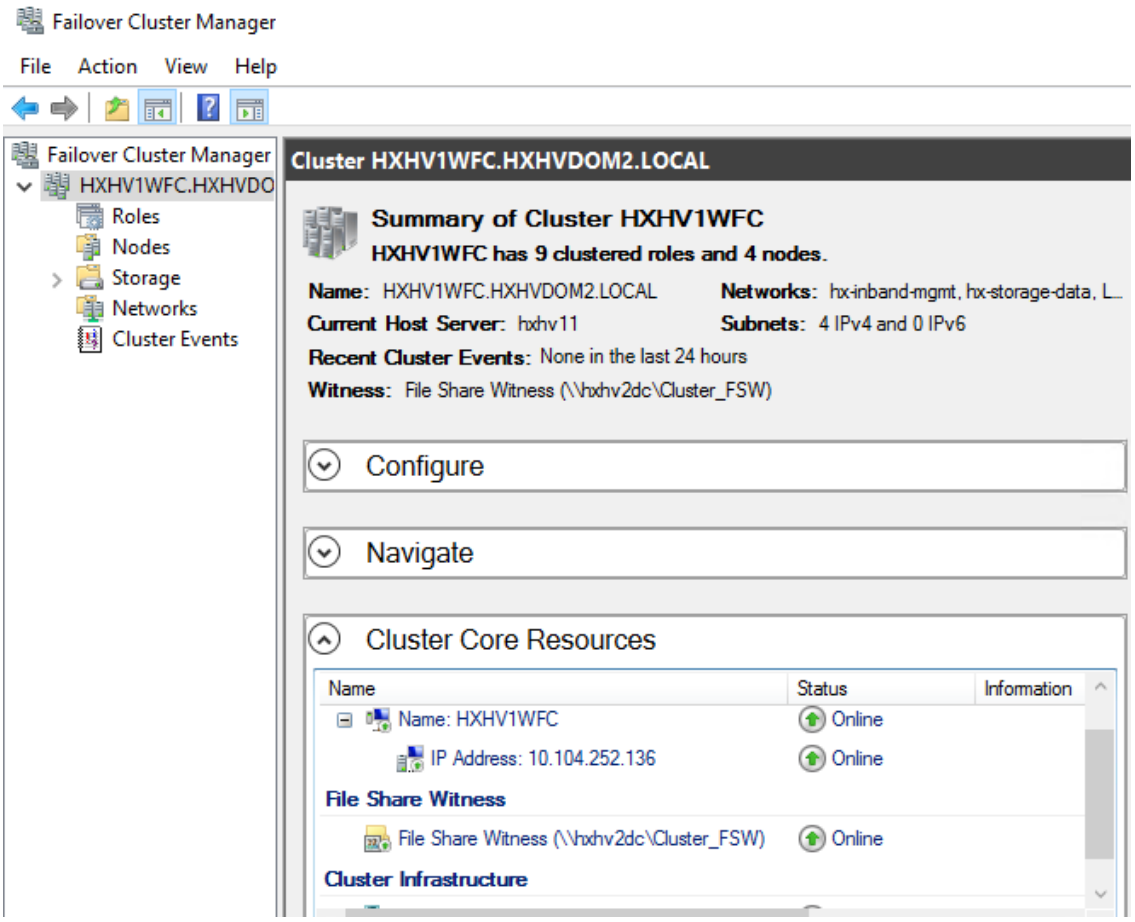


The file server can run on any virtual/physical machine if it is not hosted on the same cluster that uses the file share witness.

## Initial Tasks and Testing

In order to perform initial testing and learn about the features in the HyperFlex cluster, create a test virtual machine stored on your new HX datastore in order to take a snapshot and perform a cloning operation.

### Ready Clones

HXDP ReadyClones in a Hyper-V environment are created using a PowerShell Script that is available for download from Cisco CCO web site. Basically, two or three steps take place in the background when ReadyClone VMs are created. In the first step, the original VM is exported to a temporary folder and in the second step the saved VM is imported to a new location and registered. Later the VM is added to the cluster if that option is chosen. After the successful creation ReadyClone VMs, the exported temp folder will be deleted automatically

To create ReadyClones from a running VM, follow these steps:

1. Download the Cisco HyperFlex Data Platform Hyper-V ReadyClone PowerShell Script from the below location:

   https://software.cisco.com/download/home/286305544/type/286305994/release/4.0(1b)

2. Log into a HX Hyper-V node or a remote management station with RSAT tools

3. Open PowerShell and run the downloaded script as shown below.

   ```
   HxClone-HyperV-v4.0.1b-33133.ps1 -VmName <VM Name> -ClonePrefix <Prefix> -
   CloneCount <number> -AddToCluster <$false/$true>
   ```

Figure 93 PowerShell – Create Ready Clones

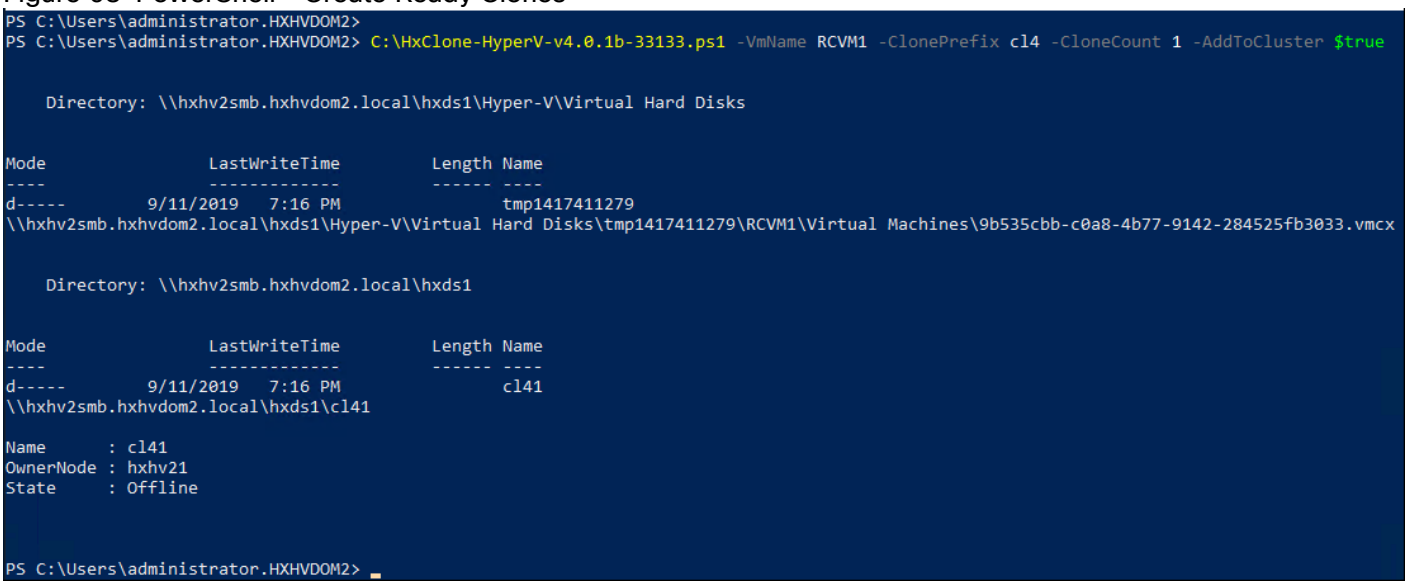


4. From the Failover Cluster manager or Hyper-V manager, select the ReadyClone VM and turn ON since it will be in the saved state.

Figure 94 Failover Cluster Manager – Roles View

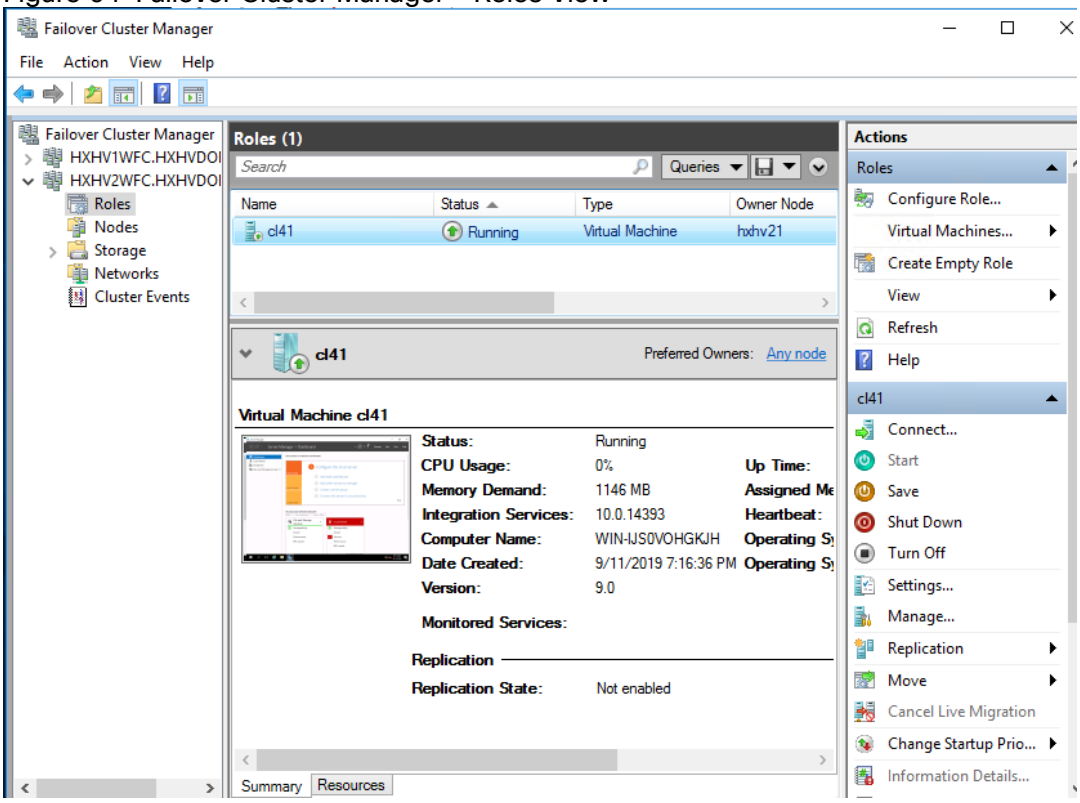


Table 38 lists the information for the script parameters.

Table 38 Script Parameters

| PowerShell Script Parameters | Information |
|---|---|
| VmName | Enter the Name of the running VM used for creating ReadyClones |
| ClonePrefix | Enter a prefix for the guest virtual machine name.<br><br>This prefix is added to the name of each ReadyClone created. |
| CloneCount | Enter a value to create the number of ReadyClones |
| AddToCluster | $false – creates standalone VMs (only visible in Hyper-V Manager)<br><br>$true – creates a highly available clustered ReadyClone VMs (visible in Failover Cluster Manager and Hyper-V Manager as well) |

## Auto-Support and Notifications

Auto-Support should be enabled for all clusters during the initial HyperFlex installation. Auto-Support enables Call Home to automatically send support information to Cisco TAC, and notifications of tickets to the email address

specified. If the settings need to be modified, they can be changed in the HyperFlex Connect HTML management webpage.

A list of events that automatically open a support ticket with Cisco TAC are as follows:

- Cluster Capacity Changed
- Cluster Unhealthy
- Cluster Health Critical
- Cluster Read Only
- Cluster Shutdown
- Space Warning
- Space Alert
- Space Critical
- Disk Blacklisted
- Infrastructure Component Critical
- Storage Timeout

To change Auto-Support settings, follow these steps:

1. From the HyperFlex Connect webpage, click the gear shaped icon in the upper right-hand corner, and click Auto-Support Settings.
2. Enable or disable Auto-Support as needed.
3. Enter the email address to receive alerts when Auto-Support events are generated.
4. Enable or disable Remote Support as needed. Remote support allows Cisco TAC to connect to the HX cluster and accelerate troubleshooting efforts.
5. Enter in the information for a web proxy if needed.
6. Click OK.

Figure 95 HyperFlex Connect – Auto Support Settings

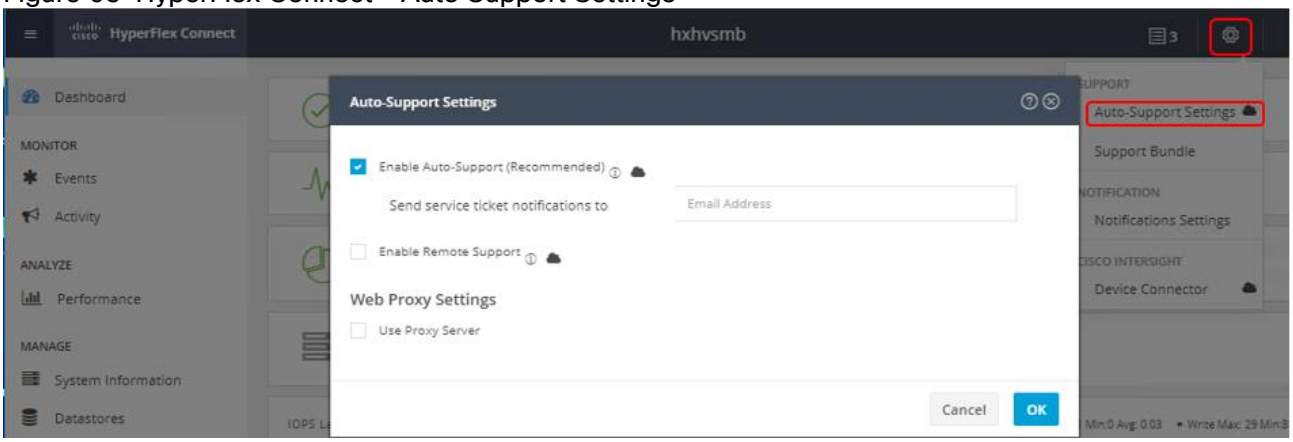


**Email notifications that come directly from the HyperFlex cluster can also be enabled.**

To enable direct email notifications, follow these steps:

1. From the HyperFlex Connect webpage, click the gear shaped icon in the upper right-hand corner, and click Notifications Settings.
2. Enter the DNS name or IP address of the outgoing email server or relay, an email address the notifications will come from, and the recipients.
3. Click OK.

Figure 96 HyperFlex Connect – Notification Settings

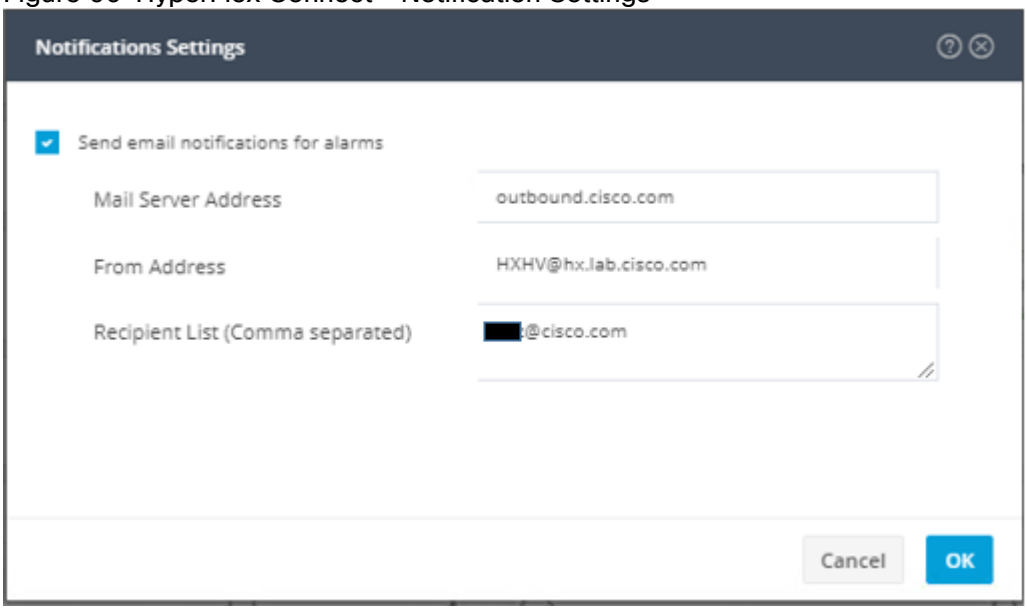


## Smart Licensing

HyperFlex 2.5 and later utilizes Cisco Smart Licensing, which communicates with a Cisco Smart Account to validate and check out HyperFlex licenses to the nodes, from the pool of available licenses in the account. At the beginning, Smart Licensing is enabled but the HX storage cluster is unregistered and in a 90-day

evaluation period or EVAL MODE. For the HX storage cluster to start reporting license consumption, it must be registered with the Cisco Smart Software Manager (SSM) through a valid Cisco Smart Account. Before beginning, verify that you have a Cisco Smart account, and valid HyperFlex licenses are available to be checked out by your HX cluster.

To create a Smart Account, go to Create Smart Accounts.

To activate and configure smart licensing, follow these steps:

1. Log into a controller virtual machine. Confirm that your HX storage cluster is in Smart Licensing mode by entering the following:

   ```
   stcli license show status
   ```

Figure 97 Storage Controller Virtual Machine – View License

```
root@hxhv11scvm:~# stcli license show status

Smart Licensing is ENABLED

Registration:
  Status: UNREGISTERED
  Export-Controlled Functionality: Not Allowed

License Authorization:
  Status: EVAL MODE
  Evaluation Period Remaining: 64 days, 2 hr, 23 min, 45 sec
  Last Communication Attempt: NONE

License Conversion:
 Automatic Conversion Enabled: true
 Status: NOT STARTED

Export Authorization Key:
  Last request status:
  Features Authorized:
   None
  Last return status:
  Return Keys in process:
   None

Utility:
  Status: DISABLED

Transport:
  Type: TransportCallHome
root@hxhv11scvm:~#
```

**Feedback will show Smart Licensing is ENABLED, Status: UNREGISTERED, and the amount of time left during the 90-day evaluation period (in days, hours, minutes, and seconds).**

2. Navigate to Cisco Software Central (https://software.cisco.com/) and log in to your Smart Account.

3. From Cisco Smart Software Manager, generate a registration token.

4. In the License pane, click Smart Software Licensing to open Cisco Smart Software Manager.

5. Click Inventory.
6. From the virtual account where you want to register your HX storage cluster, click General, and then click New Token.
7. In the Create Registration Token dialog box, add a short Description for the token, enter the number of days you want the token to be active and available to use on other products, and check Allow export-controlled functionality on the products registered with this token.
8. Click Create Token.
9. From the New ID Token row, click the Actions drop-down list, and click Copy.
10. Log into a controller virtual machine.
11. Register your HX storage cluster, where idtoken-string is the New ID Token from Cisco Smart Software Manager.

    ```
    # stcli license register --idtoken idtoken-string
    ```

12. Confirm that your HX storage cluster is registered.

    ```
    # stcli license show summary
    ```

## HyperFlex Cluster Expansion

The process to expand a HyperFlex cluster can be used to grow an existing HyperFlex cluster with additional converged storage nodes, or to expand an existing cluster with additional compute-only nodes to create an extended cluster.

### Expansion with Converged Nodes

The HX installer has a wizard for Cluster Expansion with Converged Nodes. This procedure is very similar to the initial HyperFlex cluster setup. The following process assumes a new Cisco HX node has been ordered, therefore it is pre-configured from the factory with the proper hardware and firmware installed.

#### Prerequisites

Refer to sections IP Addressing, Prepopulate AD DNS with Records, and Solution Cabling before continuing with the following steps.

To add converged storage nodes to an existing HyperFlex cluster, follow these steps:

1. On the HyperFlex installer webpage click the drop-down list for Expand Cluster, then click Converged Node.

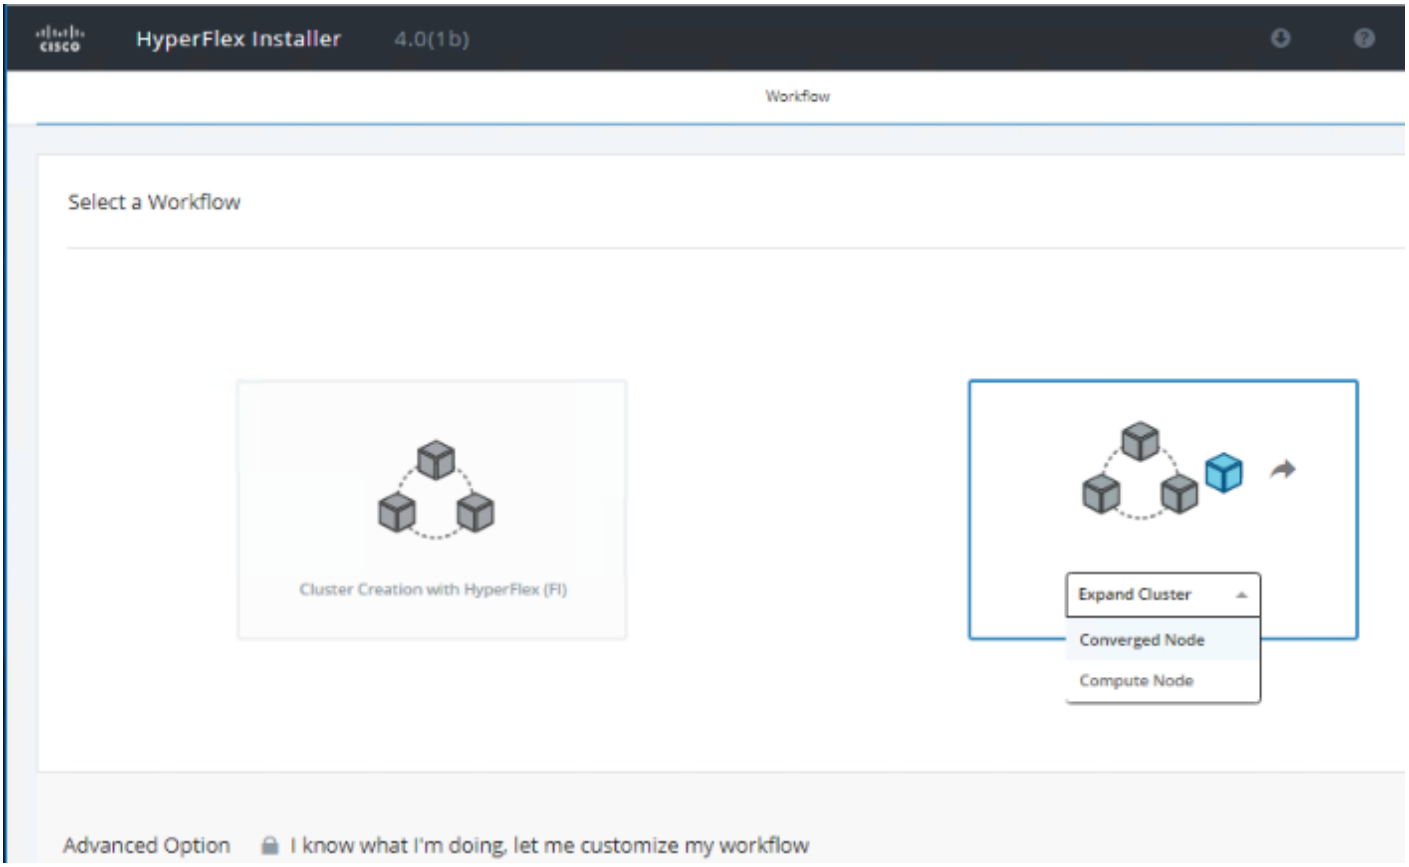


2. On the Cluster Page:

   a. Enter the HX Cluster Management IP address, Cluster Admin User name and password. You can select the option to see the passwords in clear text.

   b. Optionally, you can import a JSON file that has the configuration information, except for the appropriate passwords.

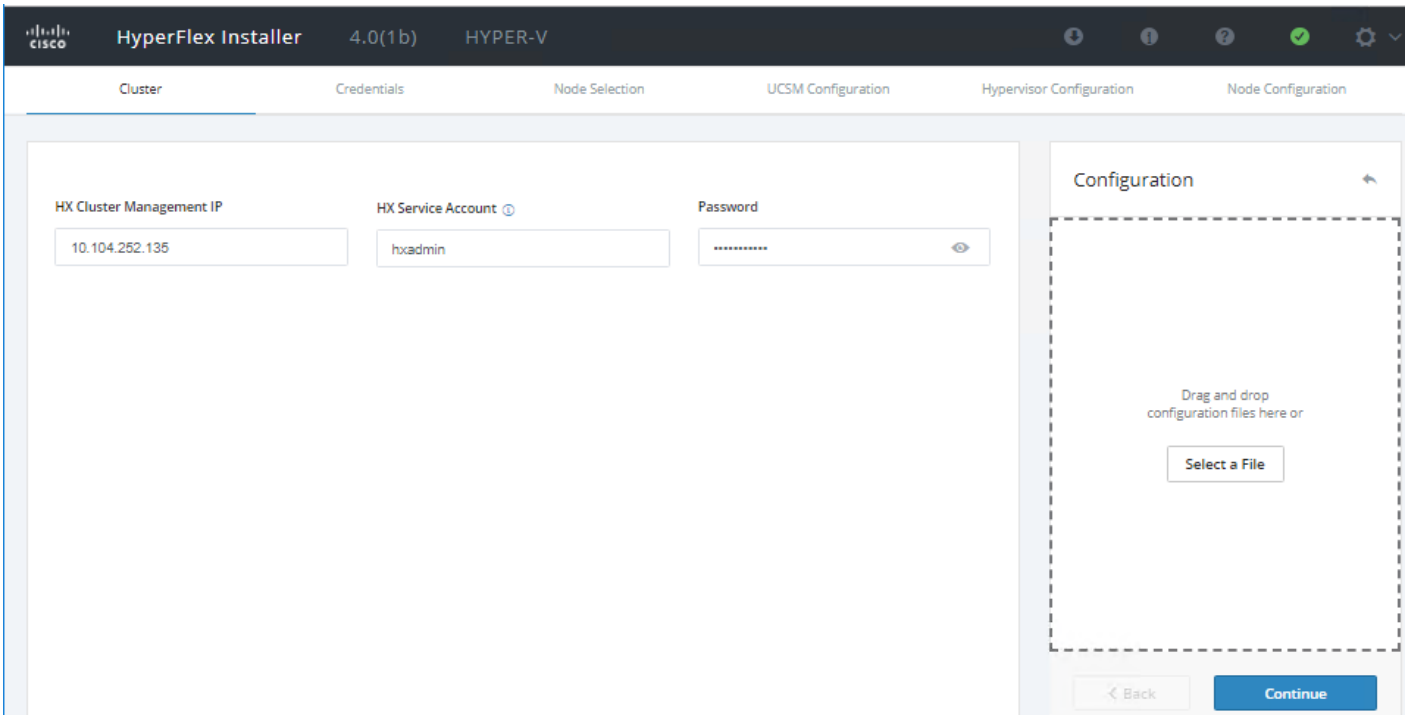


3. Click Continue.
4. On the Credentials Page:
    a. Cisco UCS Manager Credentials – Enter the Cisco UCS Manager DNS hostname or IP address, the admin usernames, and the passwords.
    b. Domain Information – Enter the HX Service Account user name and password. It is recommended to select “Constrained Delegation” here to avoid configuring it manually after the installation completion. Then select “Use HX Service Account”, if HX service account is member of AD Domain Admin group, else provide Domain Admin credentials (which is a one-time requirement)

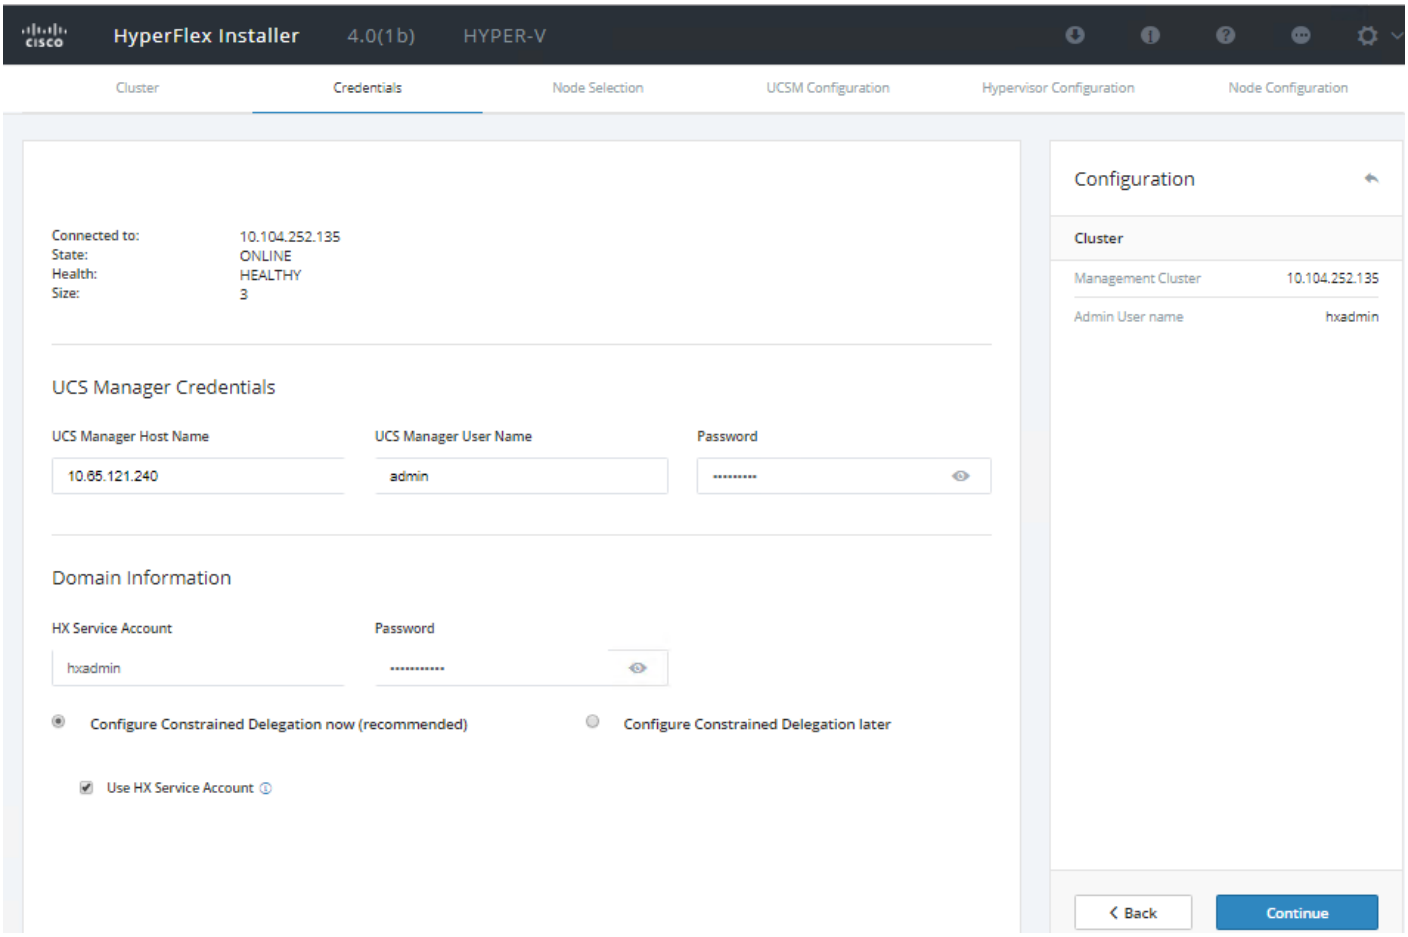


5. On the Node Selection page:

   a. Select the unassociated HX servers you want to add to the existing HX cluster.

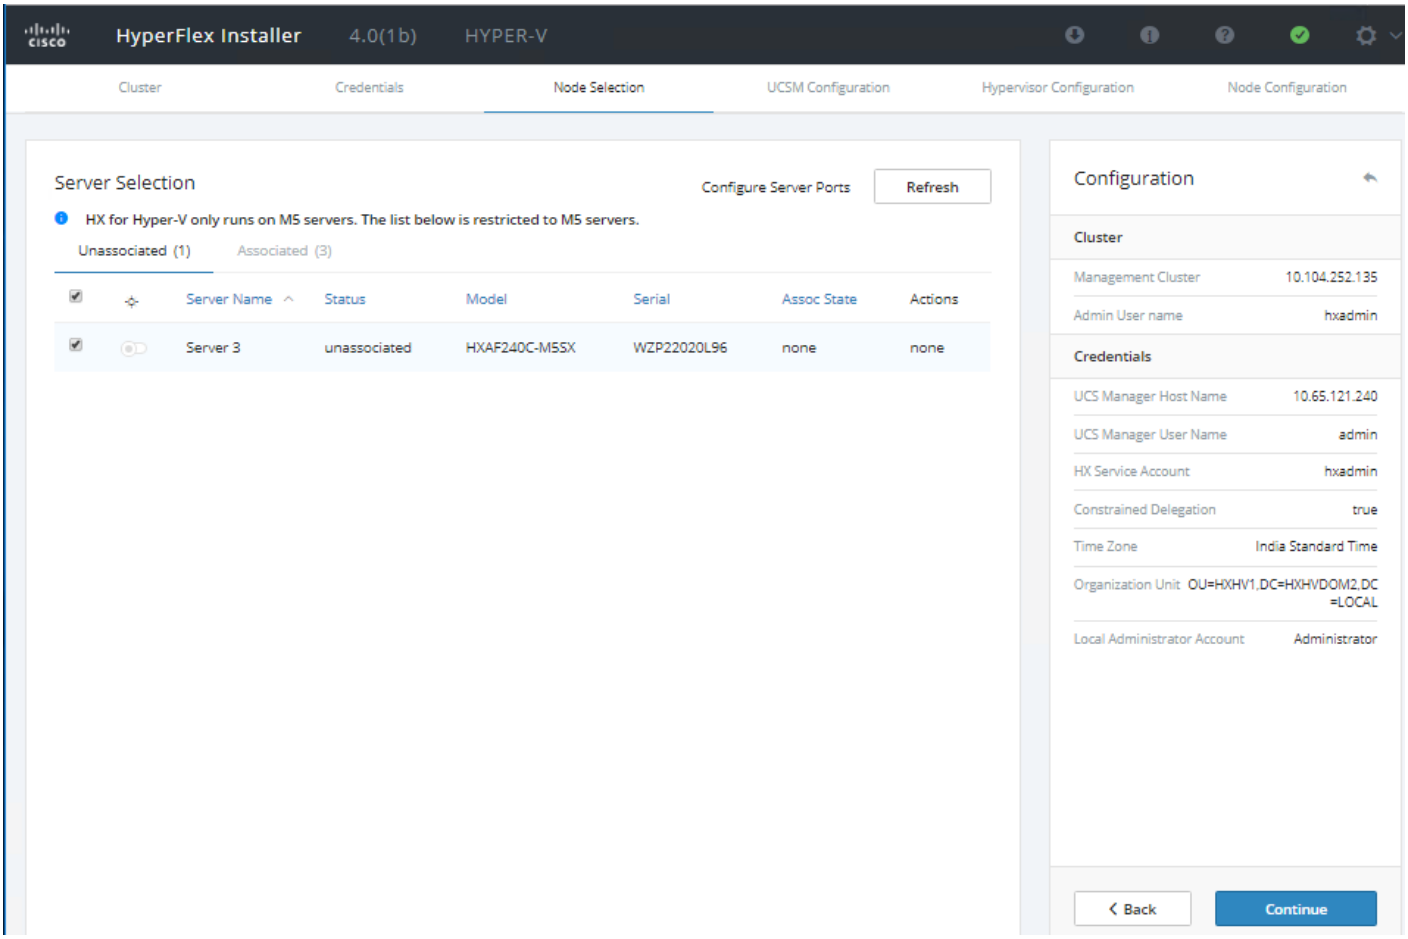


6. Click Continue

7. On the Cisco UCSM Configuration page:

   a. VLAN Configuration - Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed here.

   b. MAC Pool - Enter the MAC Pool prefix, only enter the 4th byte value, for example: 00:25:B5:0A.

   c. 'hx' IP Pool for Cisco IMC - Enter the IP address range, subnet mask and gateway to be used by the CIMC interfaces of the servers in this HX cluster.

   d. Cisco IMC access Management (Out-of-band or inband) – Select the recommended 'in band' option for faster installation of hypervisor OS on all the hx nodes.

   e. The Out-Of-Band network needs to be on the same subnet as the Cisco UCS Manager. You can add multiple blocks of addresses as a comma separated line

   f. VLAN for Inband Cisco IMC Connectivity – Enter a VLAN name and ID

   g. iSCSI/FC Storage (optional) – iSCSI Storage and FC Storage are used for adding external storage to the HyperFlex cluster. Not defined for this setup.

   h. Advanced - If multiple firmware packages exist on the Fabric Interconnect, choose the version to be installed on the servers that will comprise this cluster. Enter a unique Org name for the HyperFlex Cluster.

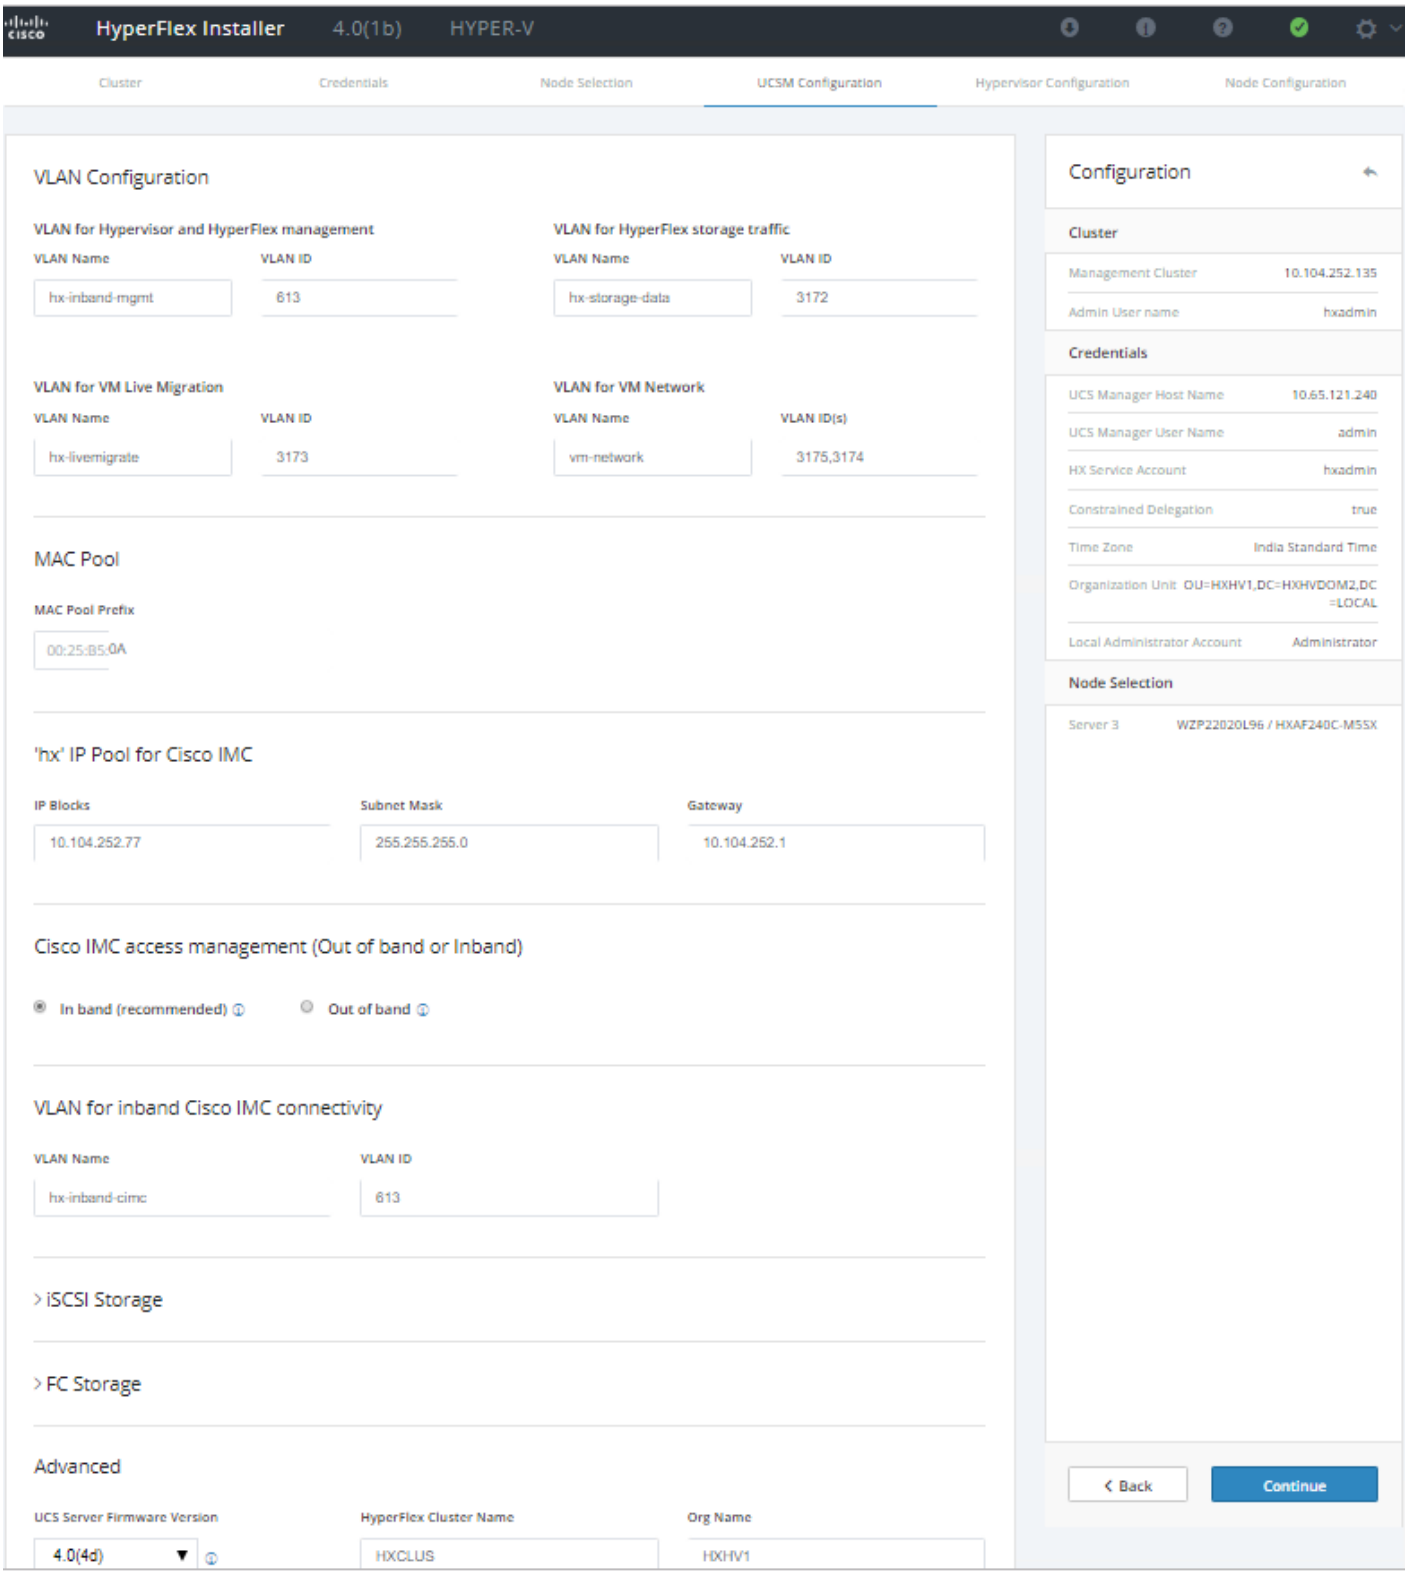


8. Click Continue.

9. On the Hypervisor Configuration page:

a. Bare Metal Configuration - If Windows Server 2016 is not installed on the nodes, select "Install Hypervisor (Hyper-V)", drag or click Browse to upload OS media file in the box and select the radio button to choose OS you want to install.

b. Hypervisor Settings - Enter IP addresses and hostnames for the Hypervisors that were created in the pre-installation section phase. The IP addresses will be assigned through Serial over Lan (SoL) through Cisco UCS Manager to the Hyper-V host systems as their management IP addresses.

c. Hypervisor Credentials – this field is pre-populated and can't be edited.

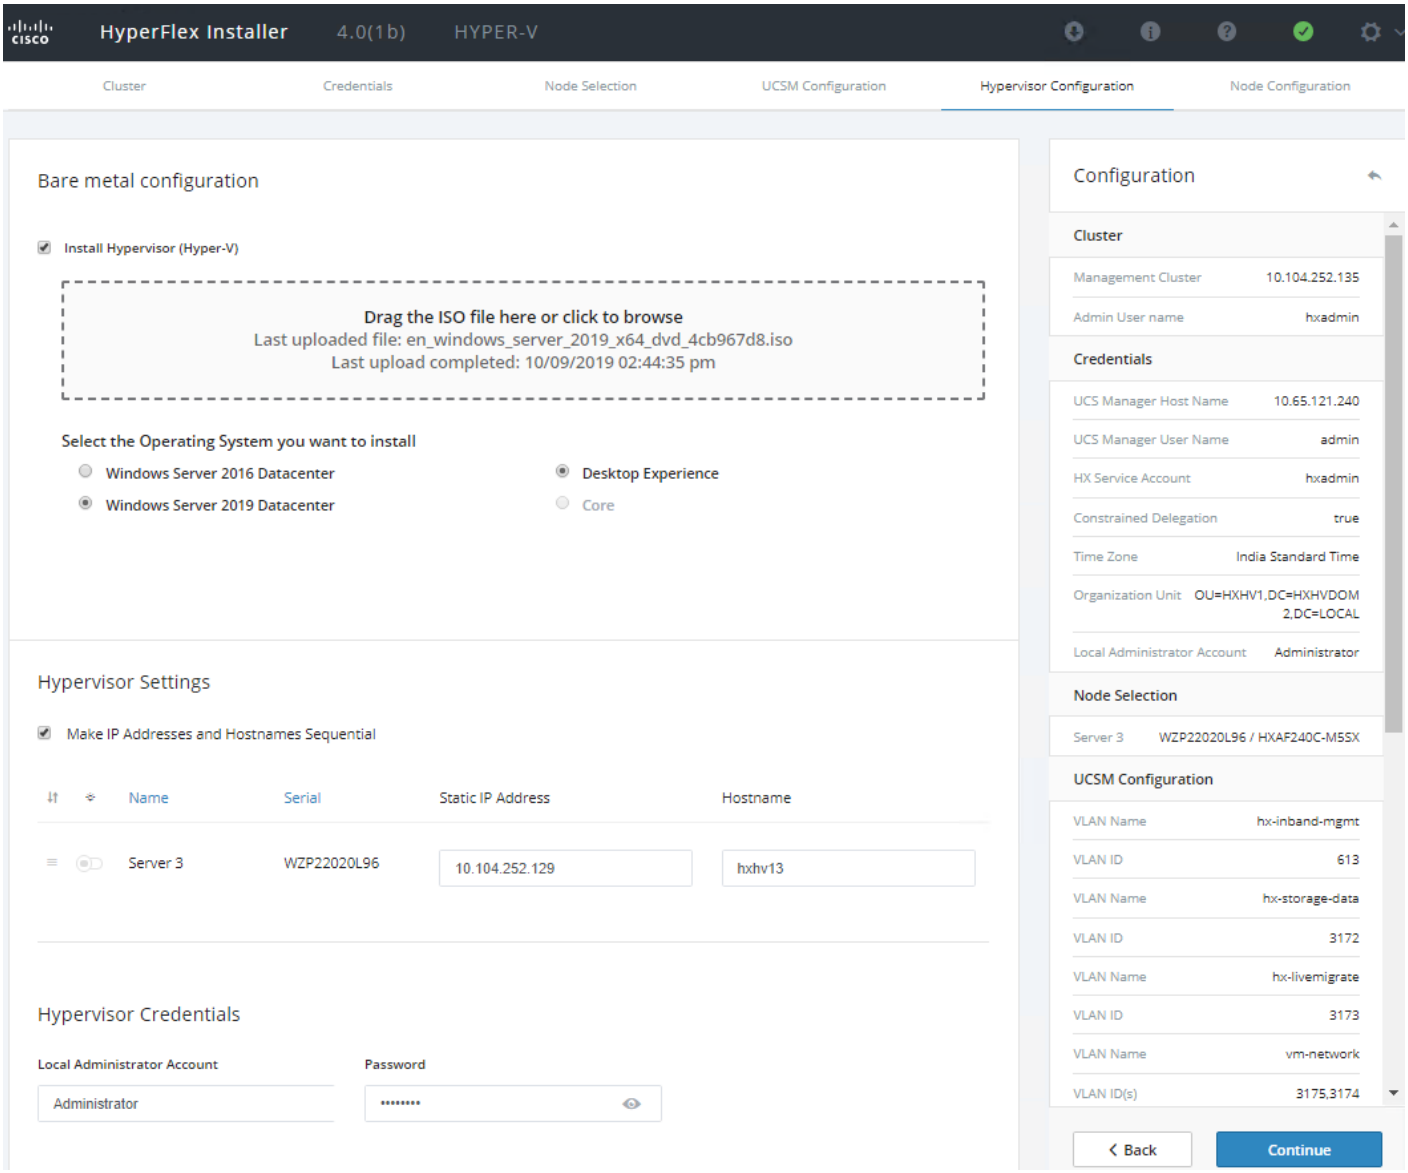


10. Click Continue.

11. On the Node Configuration Page:

a. Hypervisor Settings – Here enter the Subnet Mask, Gateway and DNS Server IP addresses.

b. Failover Cluster Name – Enter the Windows failover cluster name.

c. IP Addresses – If you leave the checkbox, Make IP Addresses and Hostnames Sequential as checked, then the installer will automatically fill the rest of the servers sequentially.

d. Assign the additional IP addresses for the Management and Data networks as well as the cluster IP addresses

e. (Optional) At this step you can manually add more servers for expansion if these servers already have service profiles associated and the hypervisor is ready, by clicking on Add Converged Server and then entering the IP addresses for the storage controller management and data networks.

f. Advanced Configuration – Select Clean up disk partitions in case if the server has been used previously

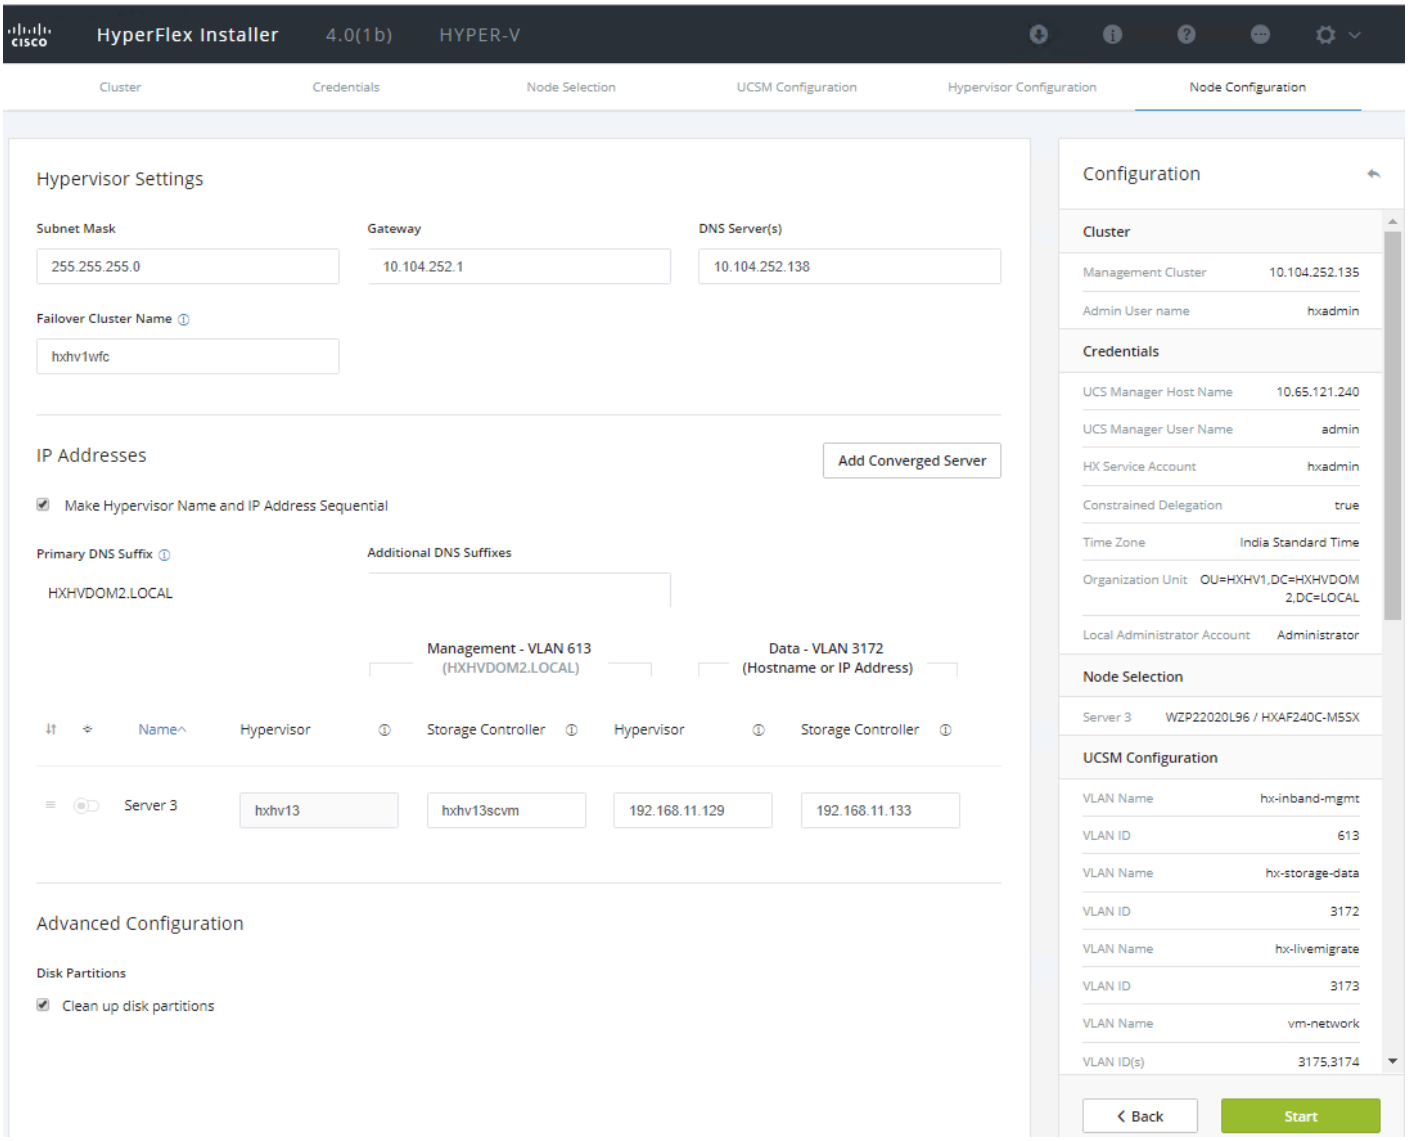


12. Click Start.

13. Validation of the configuration will now start. If there are warnings, you can review and click “Skip Validation” if the warnings are acceptable (for example, you might get the warning from Cisco UCS Manger validation that the guest VLAN is already assigned). If there are no warnings, the validation will automatically continue to the configuration process.

The HX installer will now proceed to complete the deployment and perform all the and list their status.

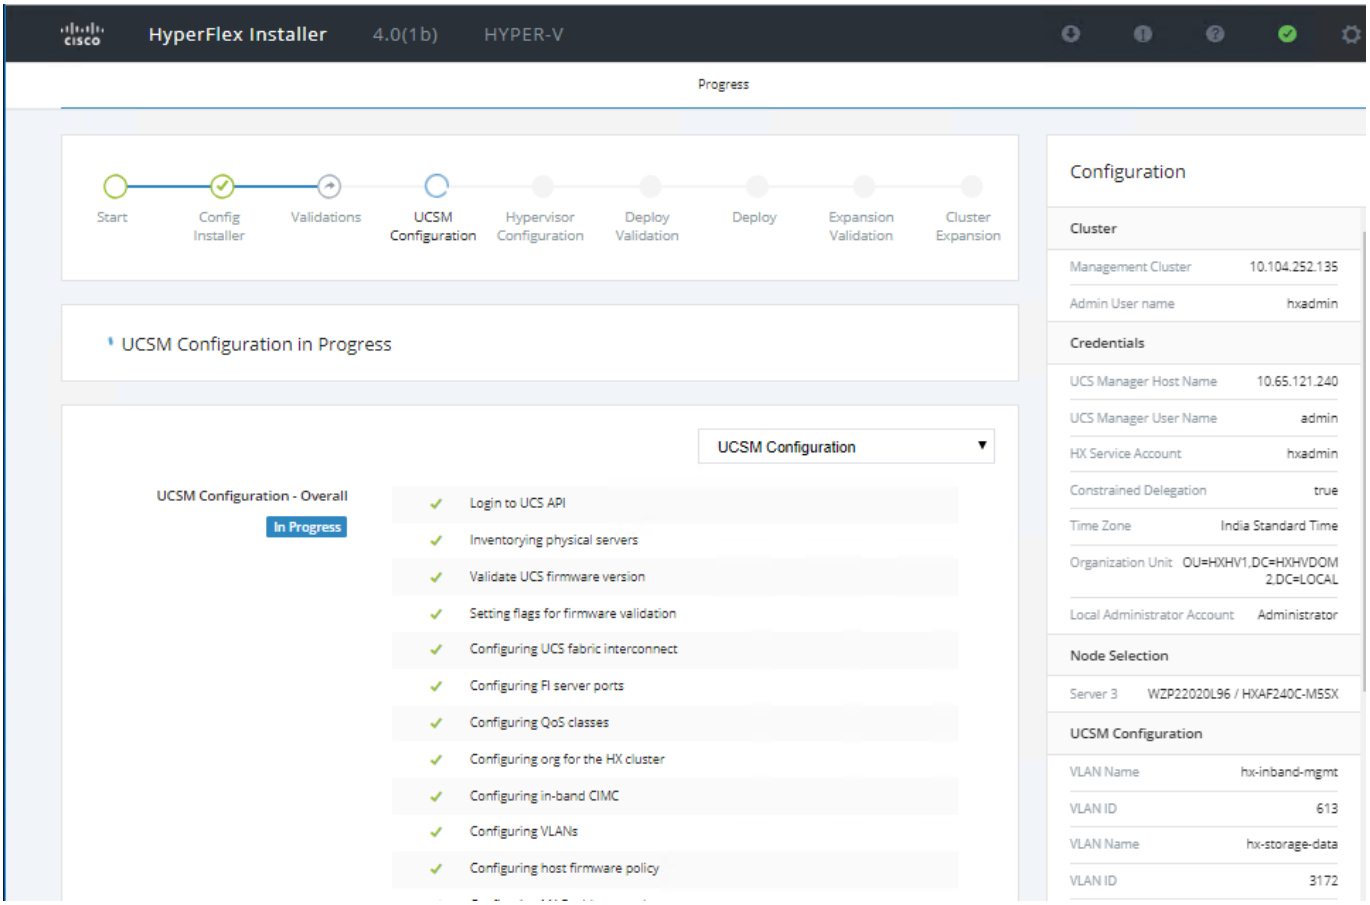


14. You can review the summary screen after the installation completes by selecting Summary.

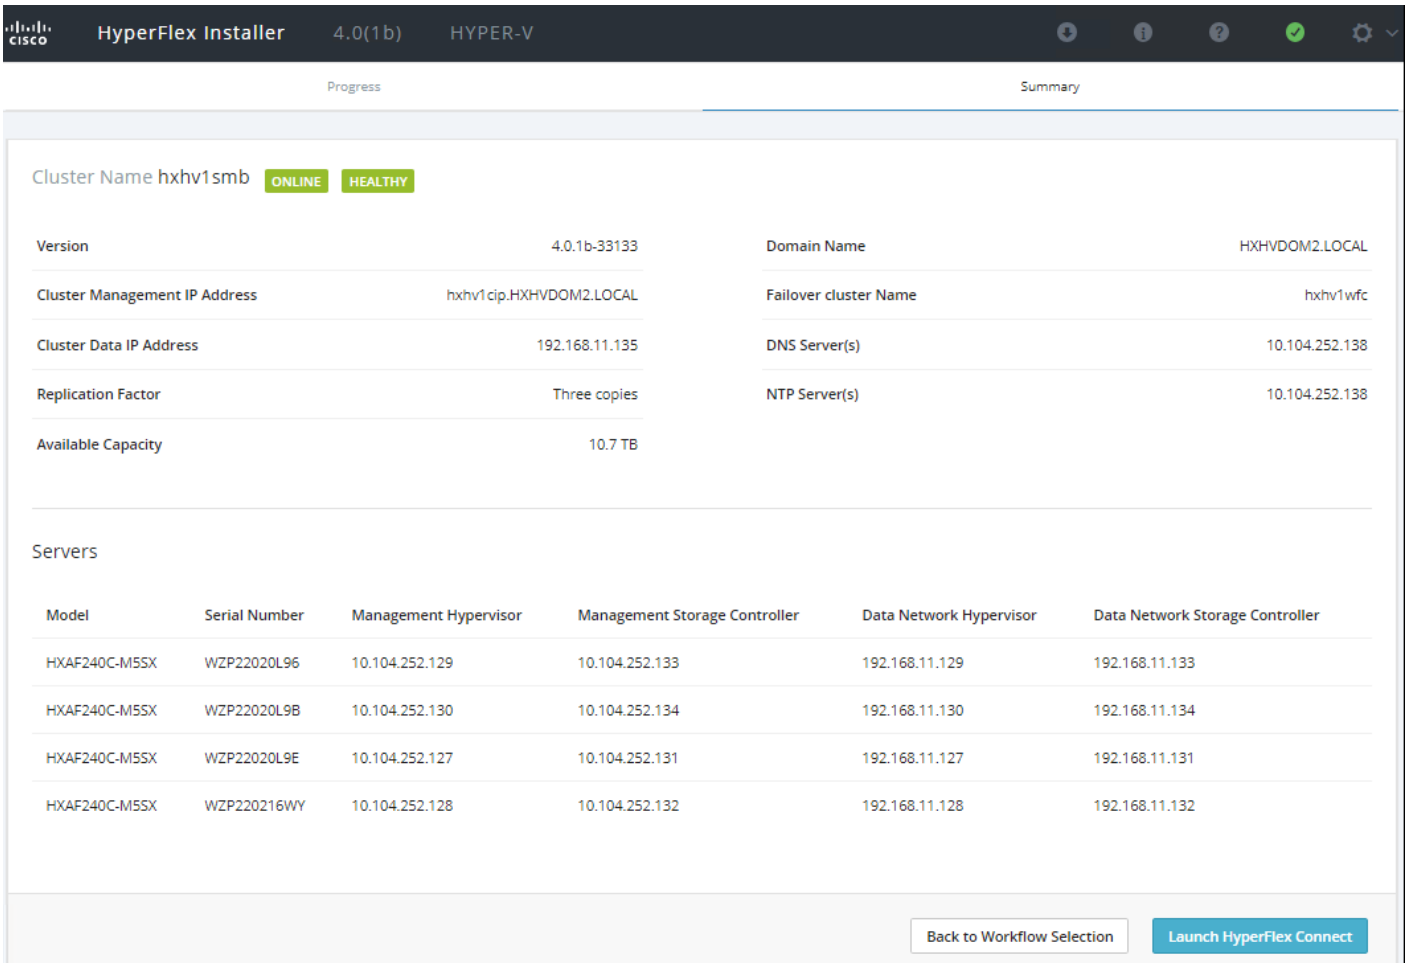


15. After the installation has completed, the new converged node is added to the cluster, and its storage, CPU, and RAM resources are immediately available. Launch HyperFlex Connect to verify the expansion of the cluster using converged node in the Dashboard, Activity and System Information sections.

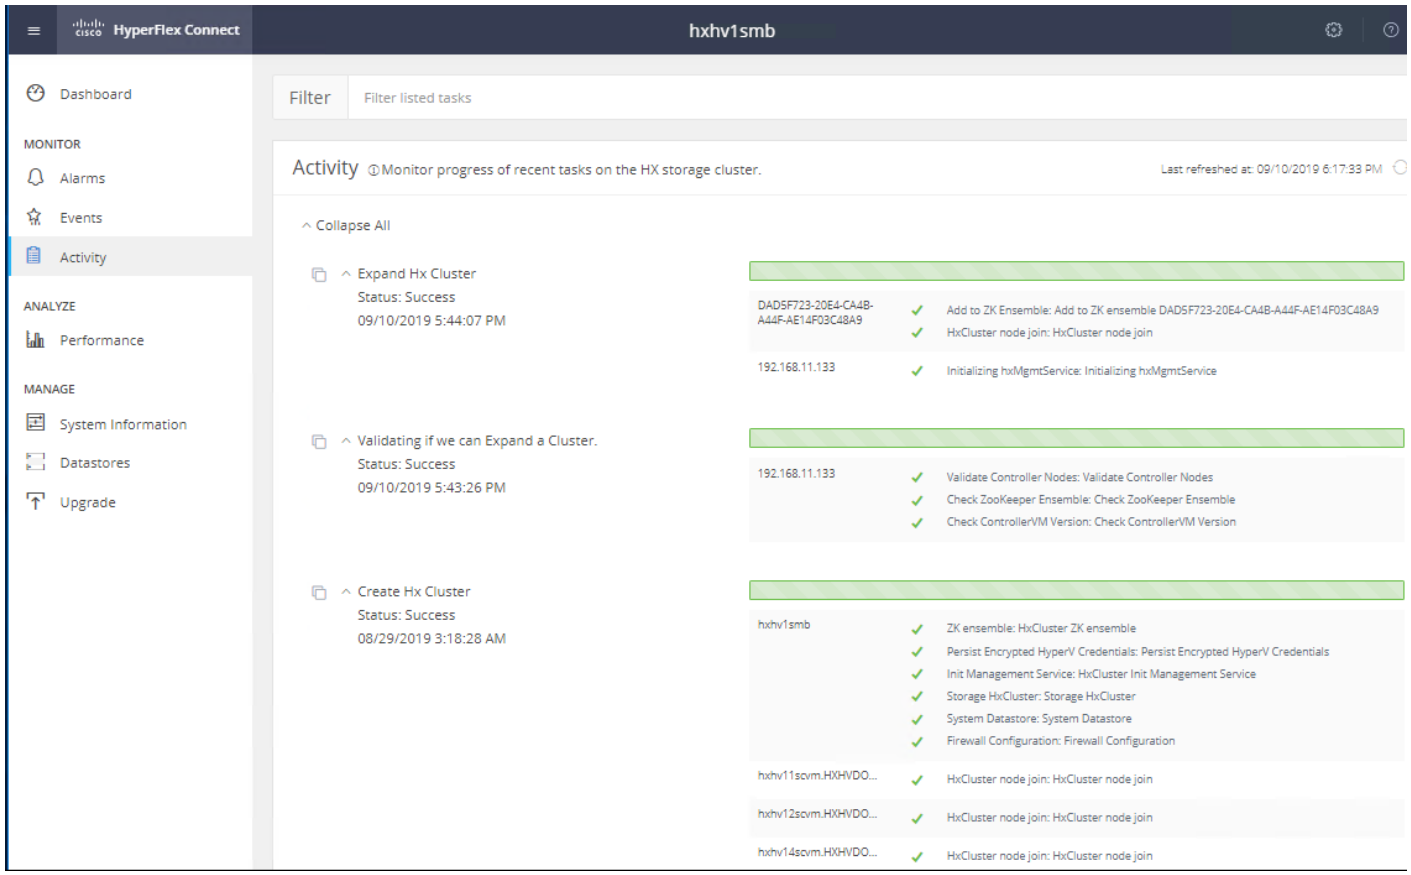
HyperFlex Connect
hxhv1smb
Dashboard
MONITOR
Alarms
Events
Activity
ANALYZE
Performance
MANAGE
System Information
Datastores
Upgrade
Filter
Filter listed tasks
Activity
Monitor progress of recent tasks on the HX storage cluster.
Last refreshed at: 09/10/2019 6:17:33 PM
Collapse All
Expand Hx Cluster
Status: Success
09/10/2019 5:44:07 PM
DAD5F723-20E4-CA4B-A44F-AE14F03C48A9
Add to ZK Ensemble: Add to ZK ensemble DAD5F723-20E4-CA4B-A44F-AE14F03C48A9
HxCluster node join: HxCluster node join
192.168.11.133
Initializing hxMgmtService: Initializing hxMgmtService
Validating if we can Expand a Cluster.
Status: Success
09/10/2019 5:43:26 PM
192.168.11.133
Validate Controller Nodes: Validate Controller Nodes
Check ZooKeeper Ensemble: Check ZooKeeper Ensemble
Check ControllerVM Version: Check ControllerVM Version
Create Hx Cluster
Status: Success
08/29/2019 3:18:28 AM
hxhv1smb
ZK ensemble: HxCluster ZK ensemble
Persist Encrypted HyperV Credentials: Persist Encrypted HyperV Credentials
Init Management Service: HxCluster Init Management Service
Storage HxCluster: Storage HxCluster
System Datastore: System Datastore
Firewall Configuration: Firewall Configuration
hxhv11scvm.HXHVDO...
HxCluster node join: HxCluster node join
hxhv12scvm.HXHVDO...
HxCluster node join: HxCluster node join
hxhv14scvm.HXHVDO...
HxCluster node join: HxCluster node join

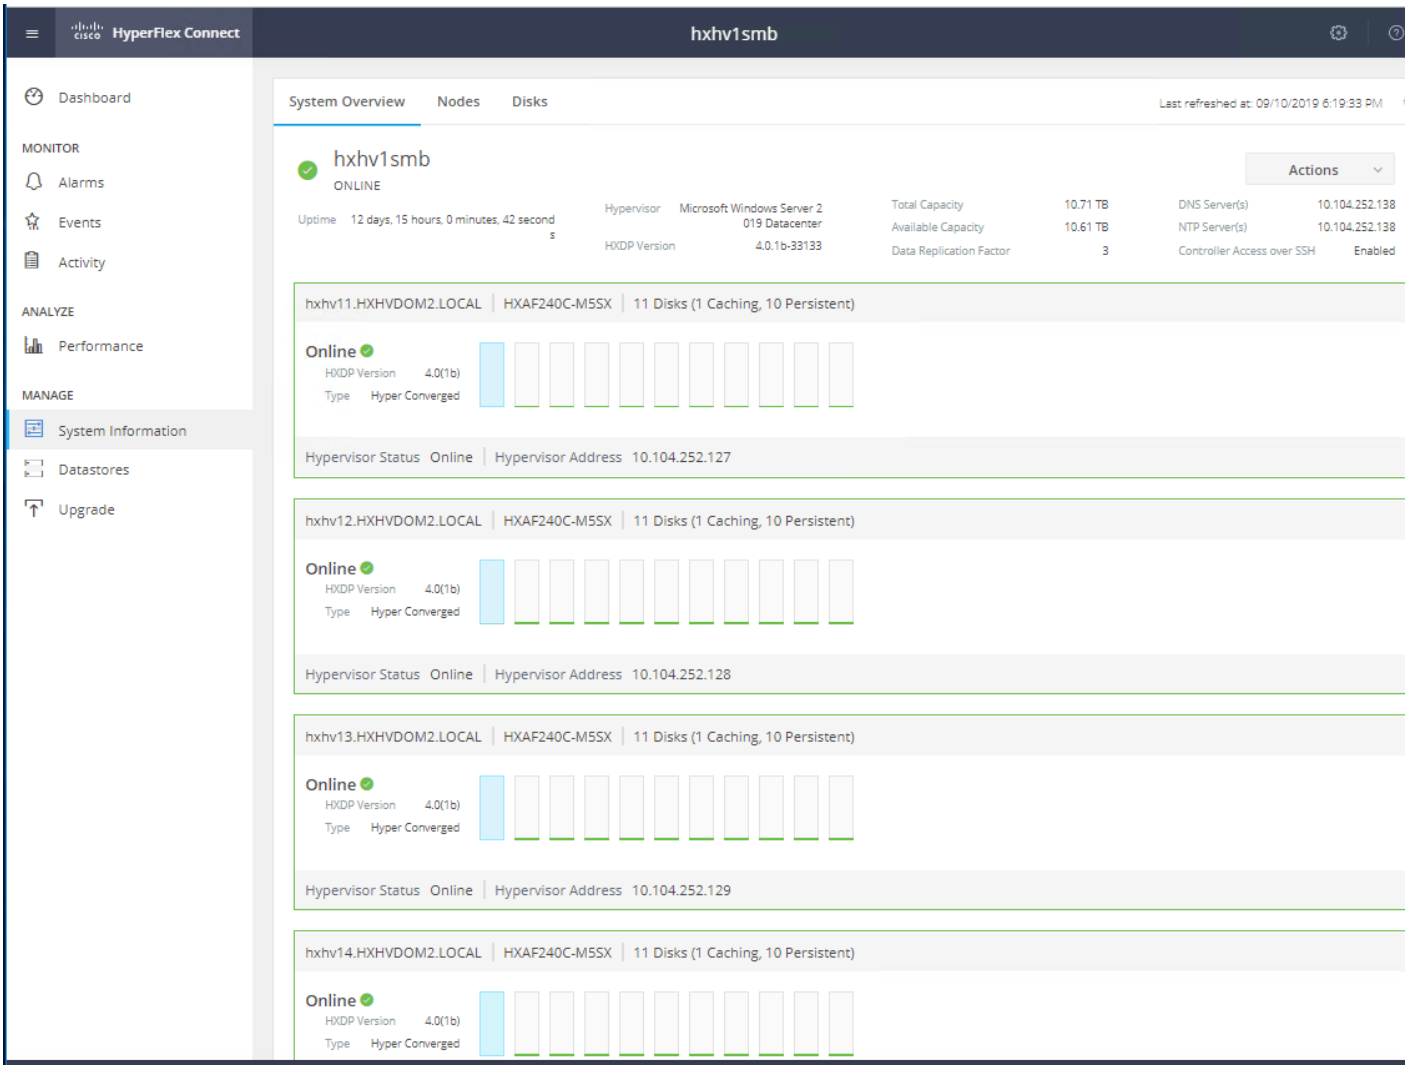


16. Launch Failover Cluster Manager to verify the expansion of the Windows cluster also.

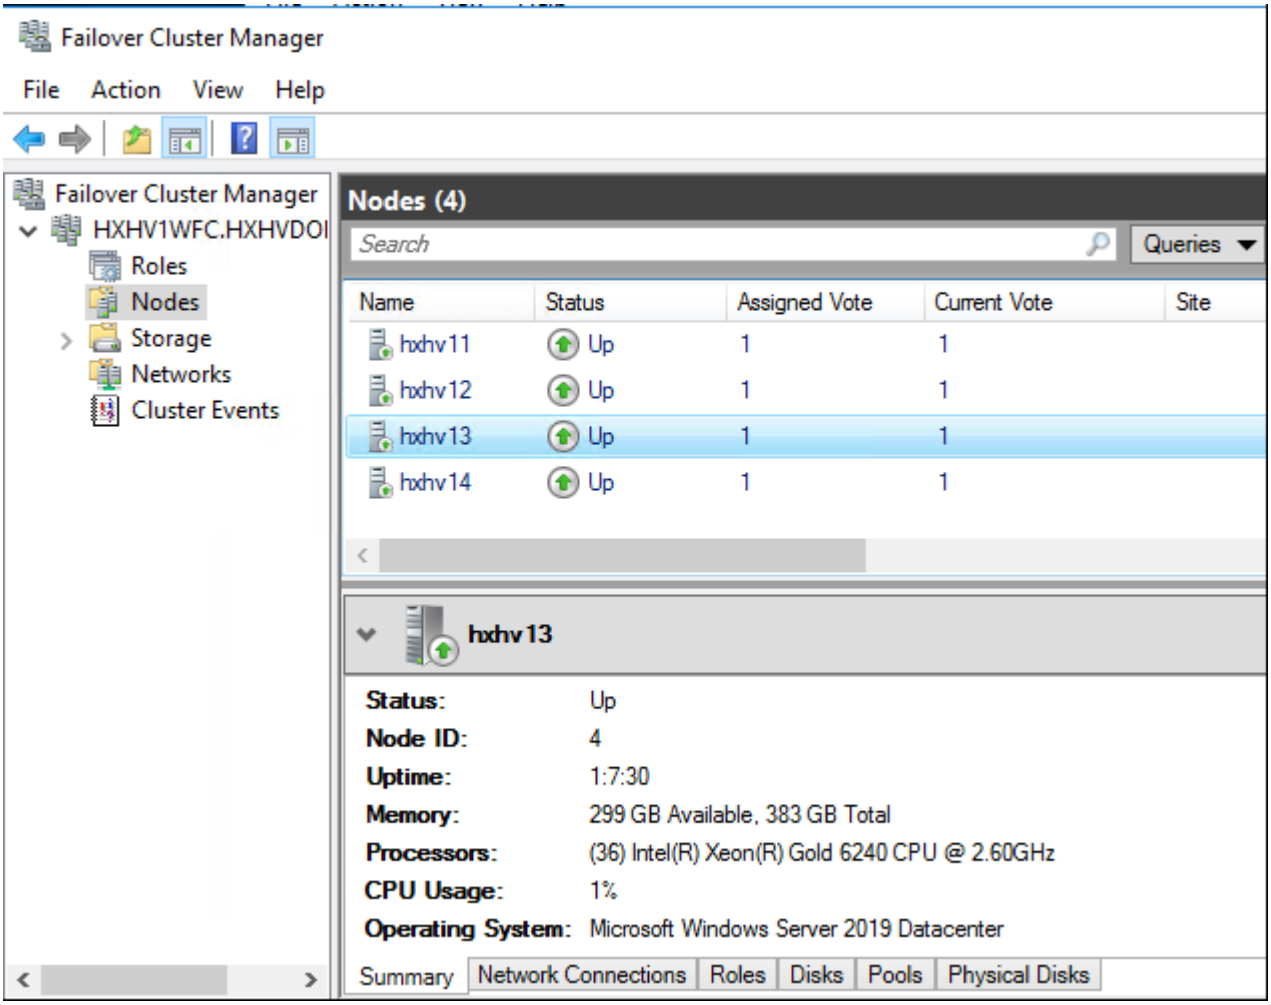


17. Refer to section Post Installation Tasks > Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces.

## Expansion with Compute-Only Nodes

The following technical guidelines must be followed when adding compute-only nodes to a Cisco HyperFlex cluster:

- The number of compute-only nodes cannot exceed the number of HyperFlex converged nodes within a single HyperFlex cluster..
- The Cisco UCS infrastructure firmware revision, which provides the firmware for Cisco UCS Manager and the Fabric Interconnects, must be maintained at the minimum version required for the HyperFlex converged nodes, or higher, at all times.
- The version of Microsoft Hyper-V installed on the compute-only nodes must be compatible with the Cisco HyperFlex version in use, and it must match the version installed on the HyperFlex converged nodes.
- While the CPU models and memory capacities between the compute-only nodes and the HyperFlex converged nodes do not have to match, configuring the nodes to have similar capacities is recommended.

- Care must be taken that the addition of the compute-only nodes will not significantly impact the HyperFlex cluster by creating additional load, or by consuming too much space. Pay close attention to the space consumption and performance requirements of any net-new virtual machines that will run on the additional compute-only nodes and note the current cluster performance and space utilization. If no new virtual machines will be created, then the current cluster performance will not be impacted.
- Mixing different models of compute-only nodes is allowed within the same cluster. Example: using Cisco UCS B220 M4, B200 M5 and C220 M5 servers as compute-only nodes is allowed.
- Connectivity between compute-only nodes and the HyperFlex cluster must be within the same Cisco UCS domain and must be 10 GbE or 40 GbE. Connecting compute-only nodes from a different Cisco UCS domain is not allowed and connecting standalone rack-mount servers from outside of the Cisco UCS domain is not allowed.
- Blade servers installed in the Cisco UCS 5108 Blade Chassis can connect through 10 GbE or 40 GbE chassis links, using the Cisco UCS 2204XP, 2208XP, or 2304 model Fabric Extenders. The Fabric Extenders, also called I/O Modules (IOMs), are typically installed in pairs, and connect the 5108 chassis to the Fabric Interconnects, which provide all the networking and management for the blades. Care must be taken not to oversubscribe and saturate the chassis links.
- Mixing CPU generations will require configuring Processor Compatibility in order to allow Live Migration to work between the compute-only nodes and the converged nodes. Enabling Processor Compatibility typically requires all virtual machines to be powered off. If it is known ahead of time that Processor Compatibility will be needed, then it is easier to enable it on the Hyper-V Manager prior to installing HyperFlex as shown below:

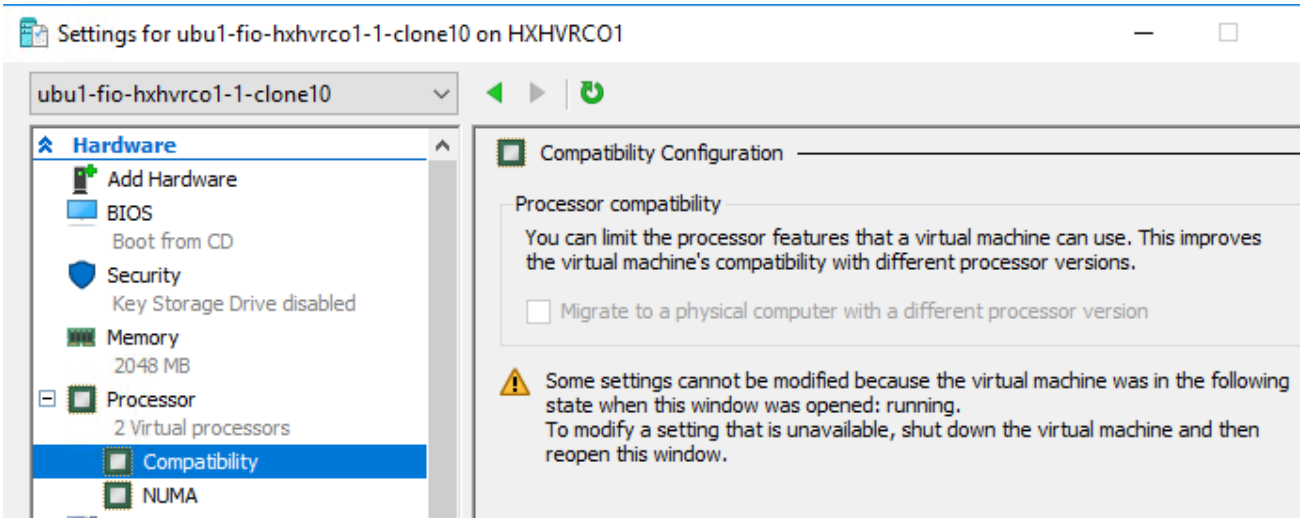


- Compute-only nodes can be configured to boot from SAN or local disks. No other internal storage should be present in a compute-only node. Manual configuration of the boot policy will be necessary if booting from any device other than M.2 drive.

## Prerequisites

Refer to sections IP Addressing, Prepopulate AD DNS with Records, and Solution Cabling before continuing with the following steps. The HX installer has a wizard for Cluster Expansion with converged nodes and compute-only nodes, however the compute-only node process requires some additional manual steps to install the Windows Server 2016 on the nodes. To expand an existing HyperFlex cluster with compute-only nodes, follow these steps:

1. On the HyperFlex installer webpage click the drop-down list for Expand Cluster, then click Compute Node.

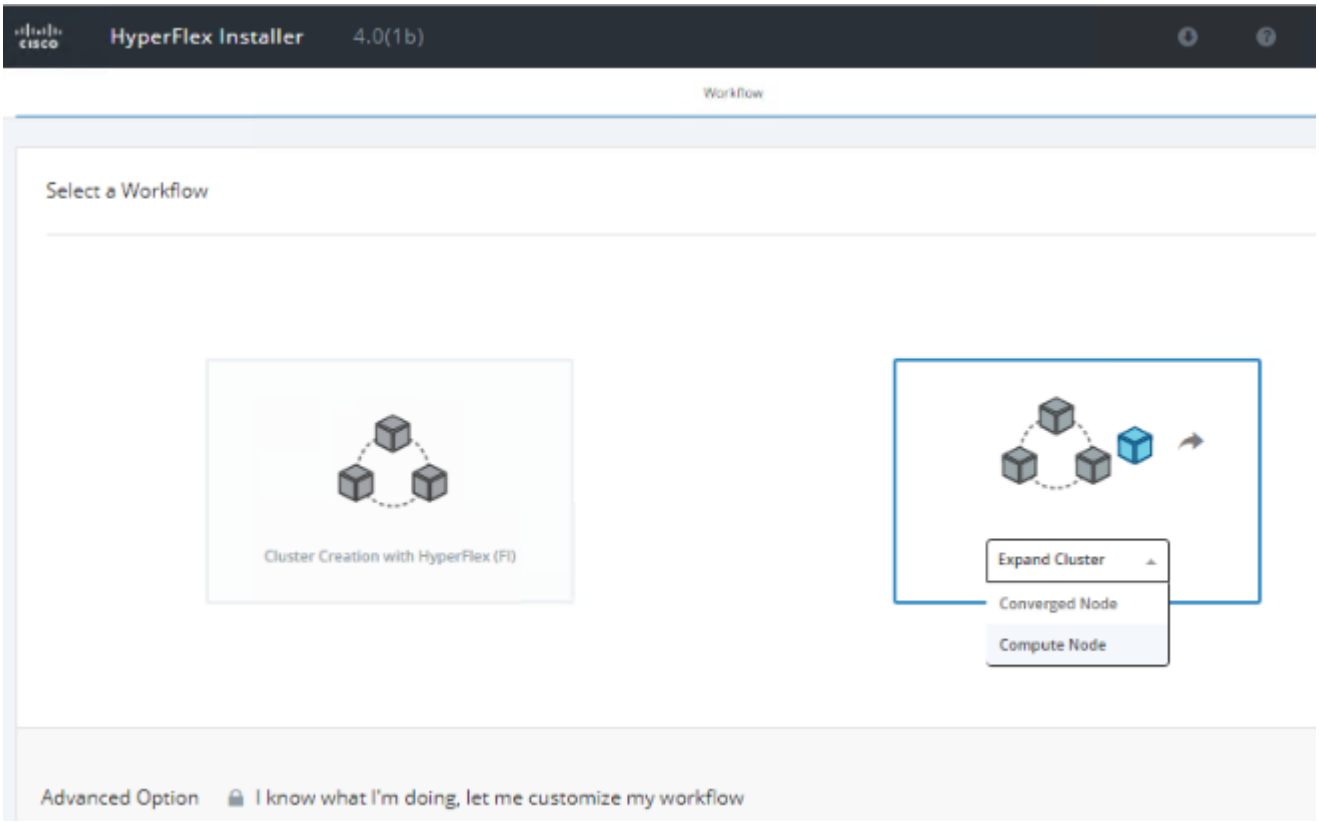


2. Select “Run UCS Manager Configuration” as shown below:

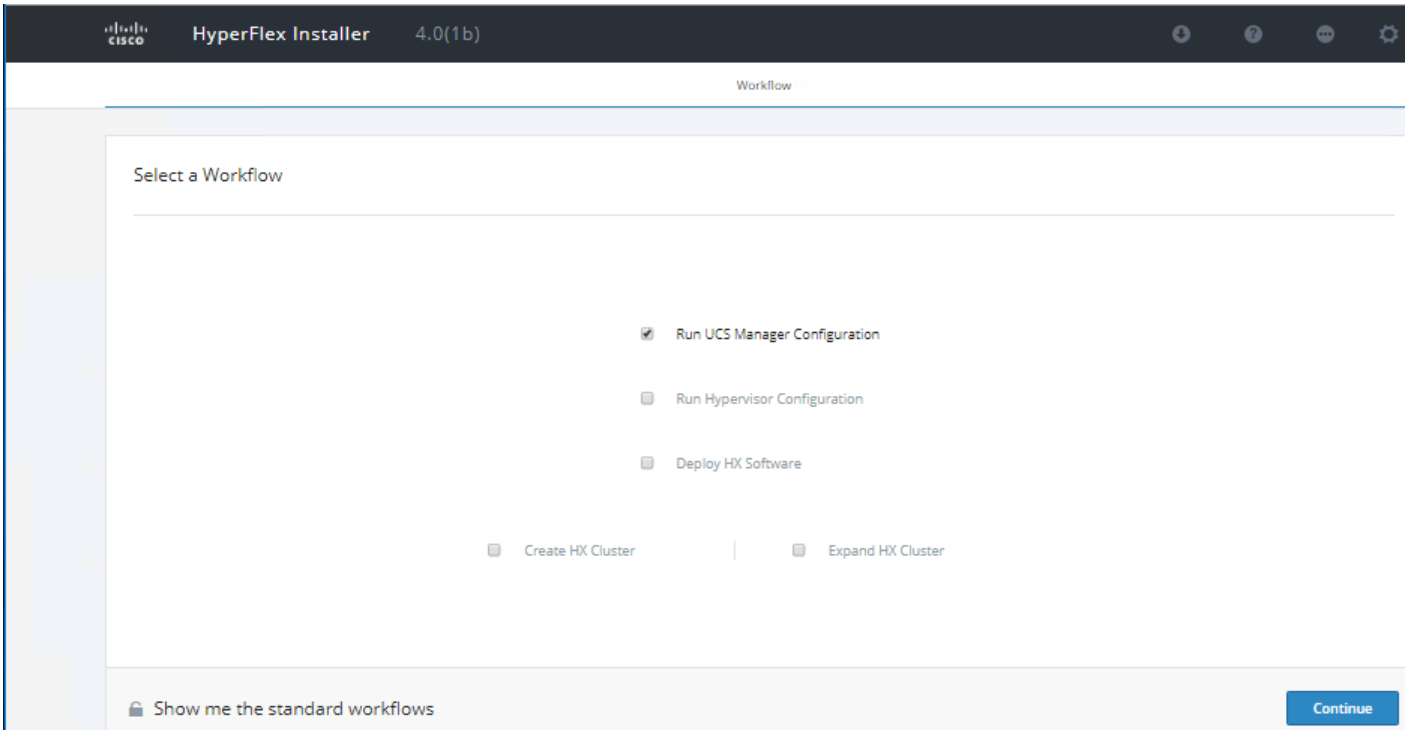


3. Click Continue and click Confirm and Proceed on the warning pop-up message.

4. On the Credentials Page:

   – Cisco UCS Manager Credentials – Enter the Cisco UCS Manager DNS hostname or IP address, the admin usernames, and the passwords.

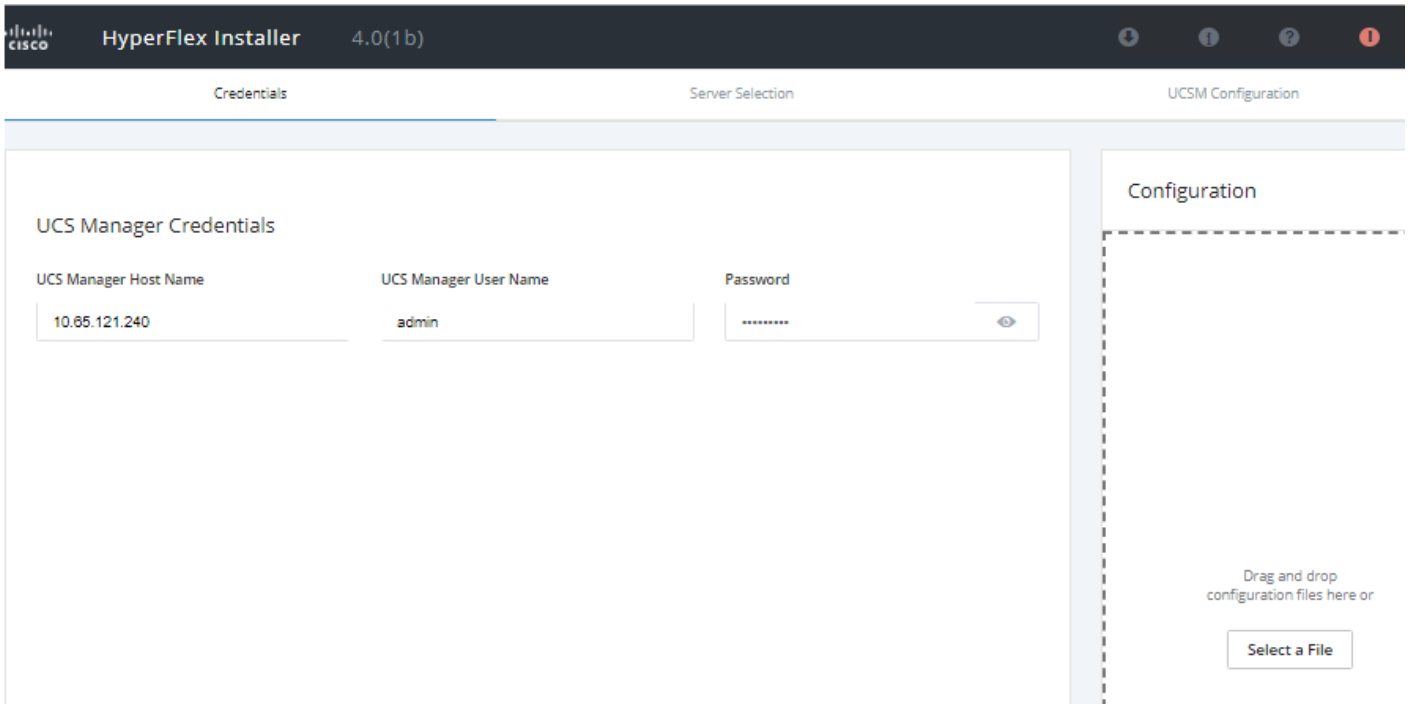


5. On the Server Selection page:

   – Select the unassociated server/s you want to add to the existing HX cluster. In this case UCSC-C220-M5SX is selected.

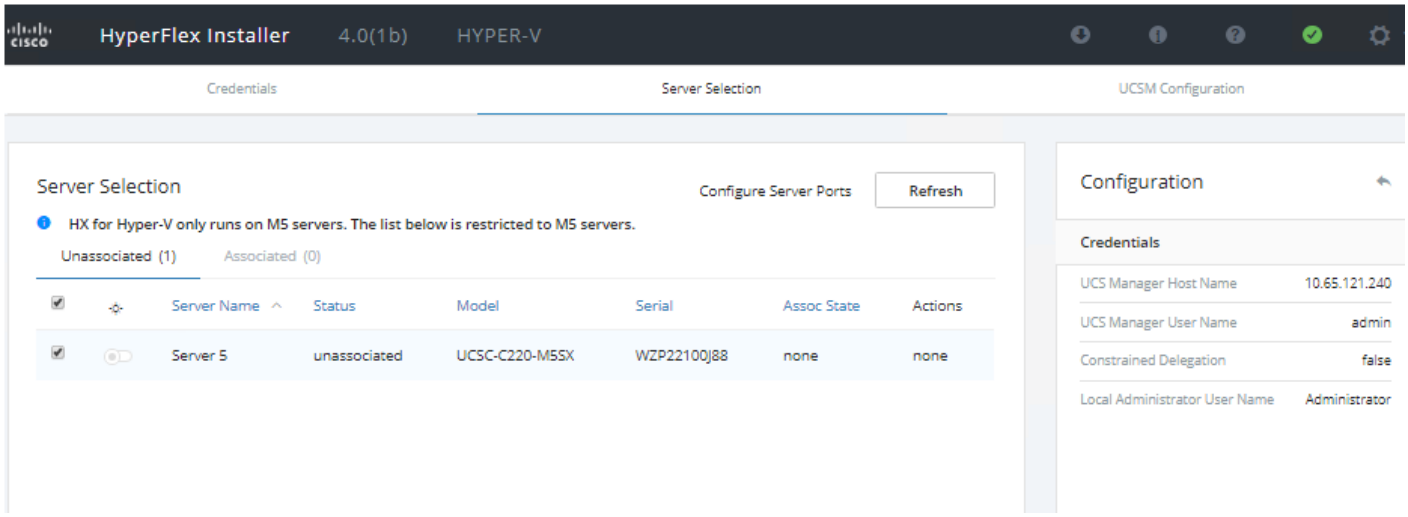


6. Click Continue.

7. On the Cisco UCSM Configuration page:

a. VLAN Configuration – Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed here.

b. MAC Pool – Enter the MAC Pool prefix, only enter the 4th byte value, for example: 00:25:B5:0A.

c. 'hx' IP Pool for Cisco IMC – Enter the IP address range, subnet mask and gateway to be used by the CIMC interfaces of the servers in this HX cluster.

d. Cisco IMC access Management (Out of band or inband) – Select the recommended 'in band' option for faster installation of hypervisor OS on all the hx nodes.

e. The Out-Of-Band network needs to be on the same subnet as the Cisco UCS Manager. You can add multiple blocks of addresses as a comma separated line

f. VLAN for Inband Cisco IMC Connectivity – Enter a VLAN name and ID

g. iSCSI/FC Storage (optional) – iSCSI Storage and FC Storage are used for adding external storage to the HyperFlex cluster. Not defined for this setup.

h. Advanced – If multiple firmware packages exist on the Fabric Interconnect, choose the version to be installed on the servers that will comprise this cluster. Enter a unique Org name for the HyperFlex Cluster.

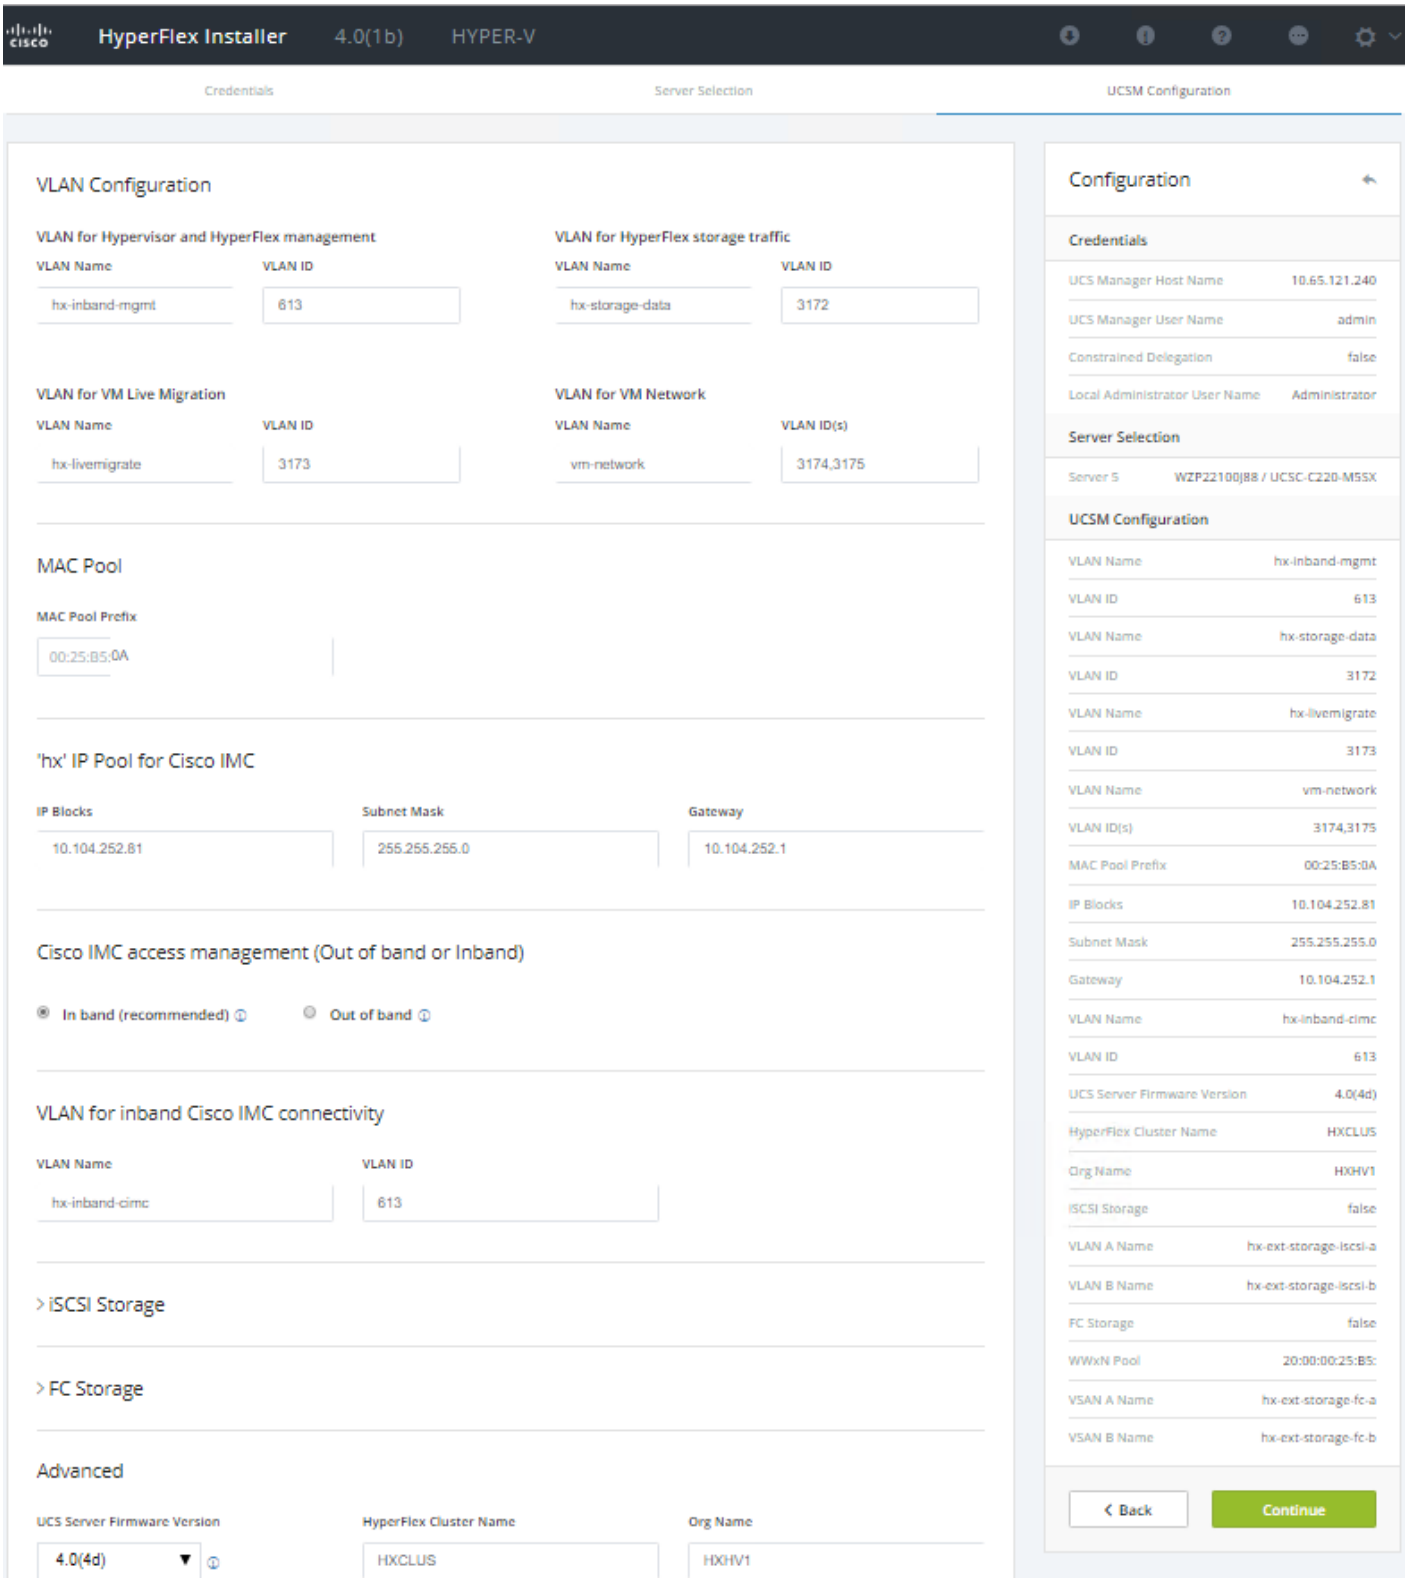


8. Click Start.

9. Validation of the configuration will now start. After validation, the installer will create the compute-only node service profiles and associate them with the selected servers.

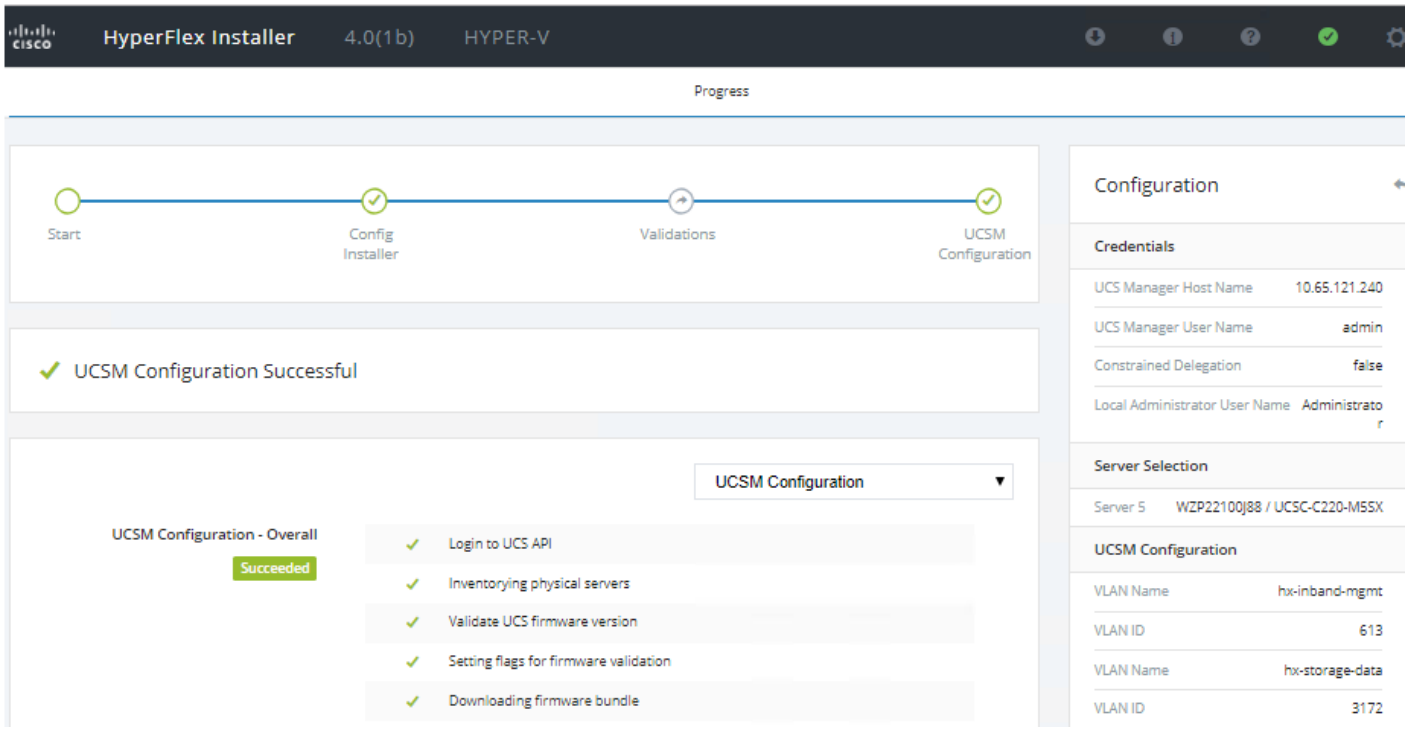


10. If the Hyper-V hypervisor has not been previously installed on the compute-only nodes, refer to section HyperFlex Installation.

11. If the hypervisor is already installed, then continue with the next steps in this section.

12. You might need to “start over” since the previous workflow was finished. Click the gear icon in the top right corner and select Start Over.

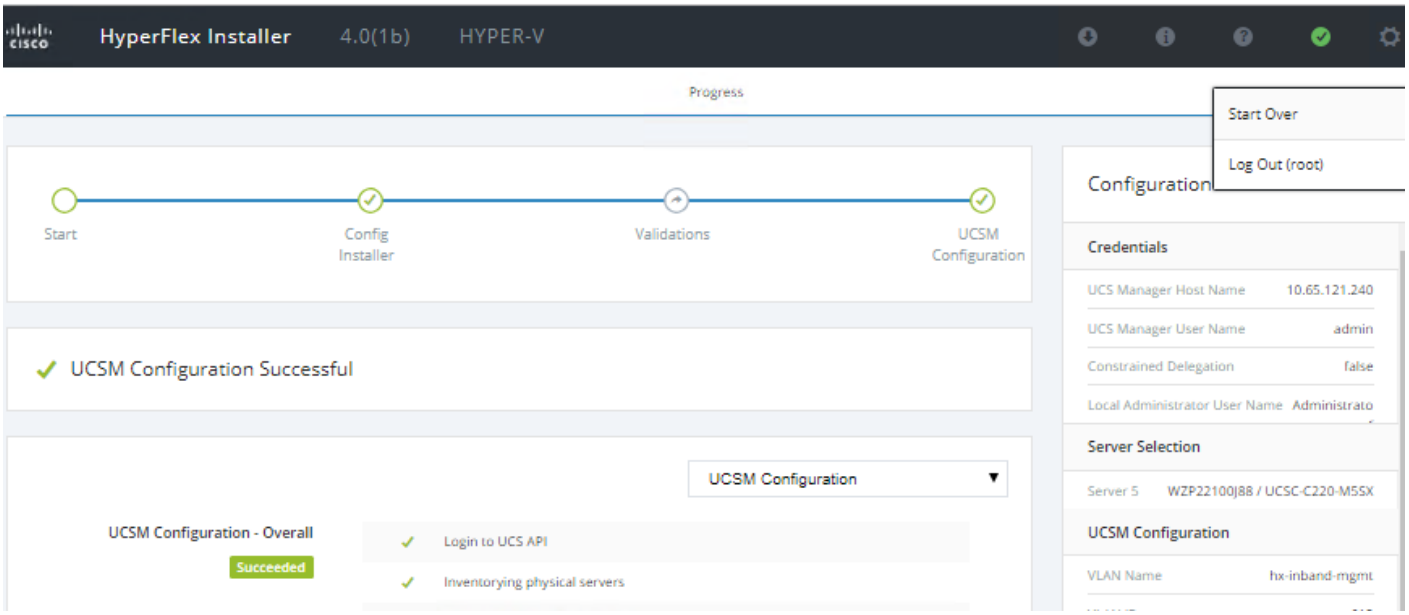


13. On the HyperFlex installer webpage click the drop-down list for Expand Cluster, then click Compute Node.

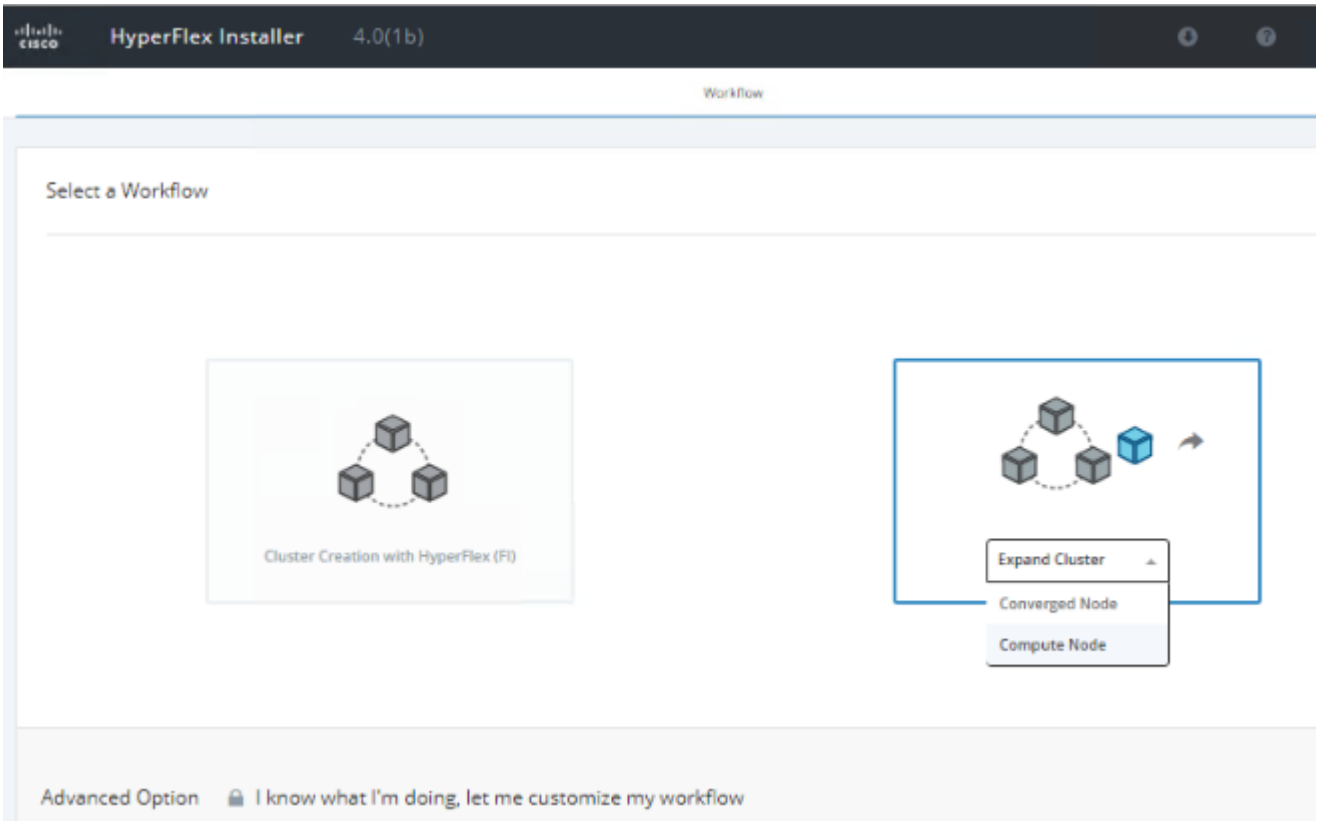


14. On the Select a Workflow page, select “Is OS installed on the node”. With this selection, Run Hypervisor Configuration, Deploy HX Software and Expand Cluster are also automatically selected.

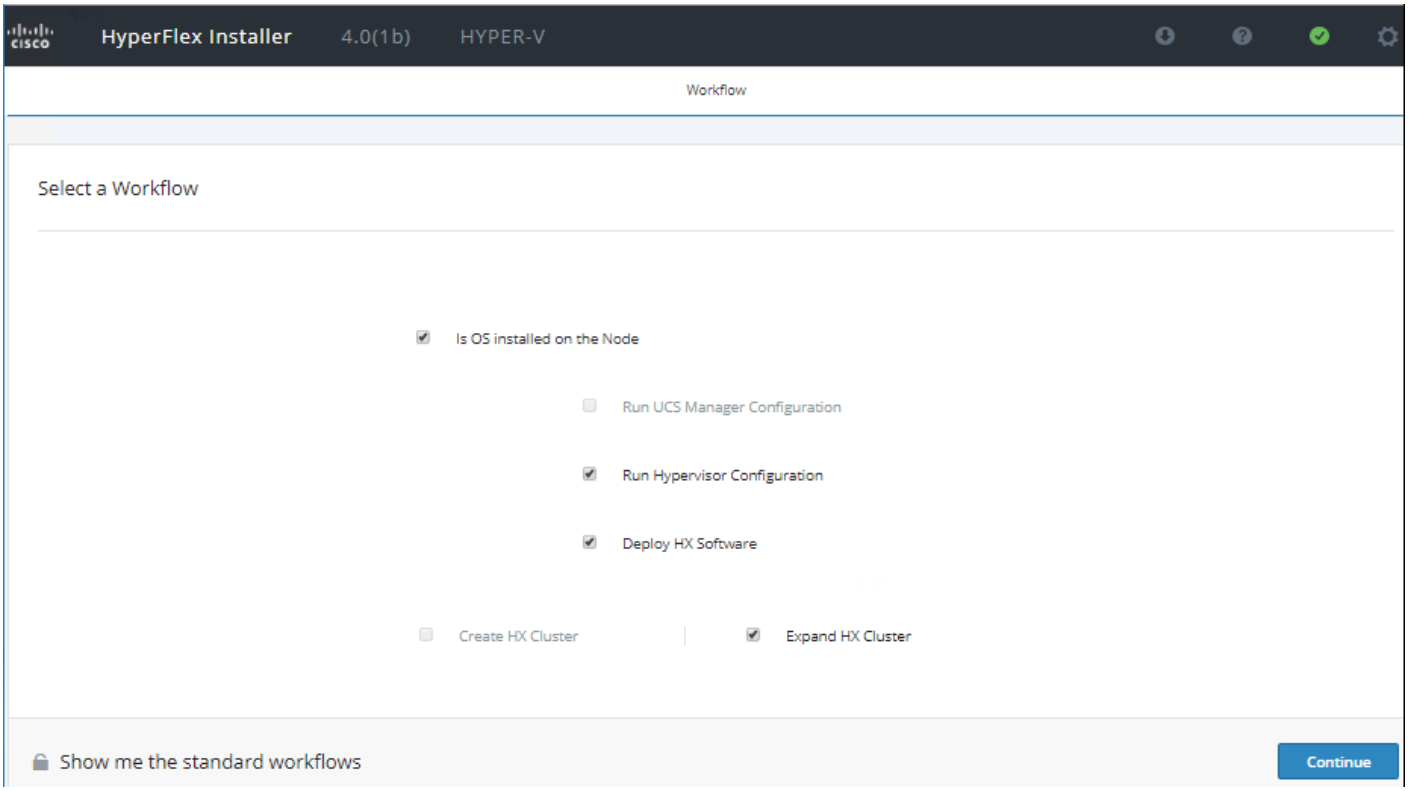

15. Click Continue and click Confirm and Proceed on the warning pop-up message.

16. On the Cluster Page:

    a. Enter the HX Cluster Management IP address, Cluster Admin User name and password. You can select the option to see the passwords in clear text.

    b. Optionally, you can import a JSON file that has the configuration information, except for the appropriate passwords.

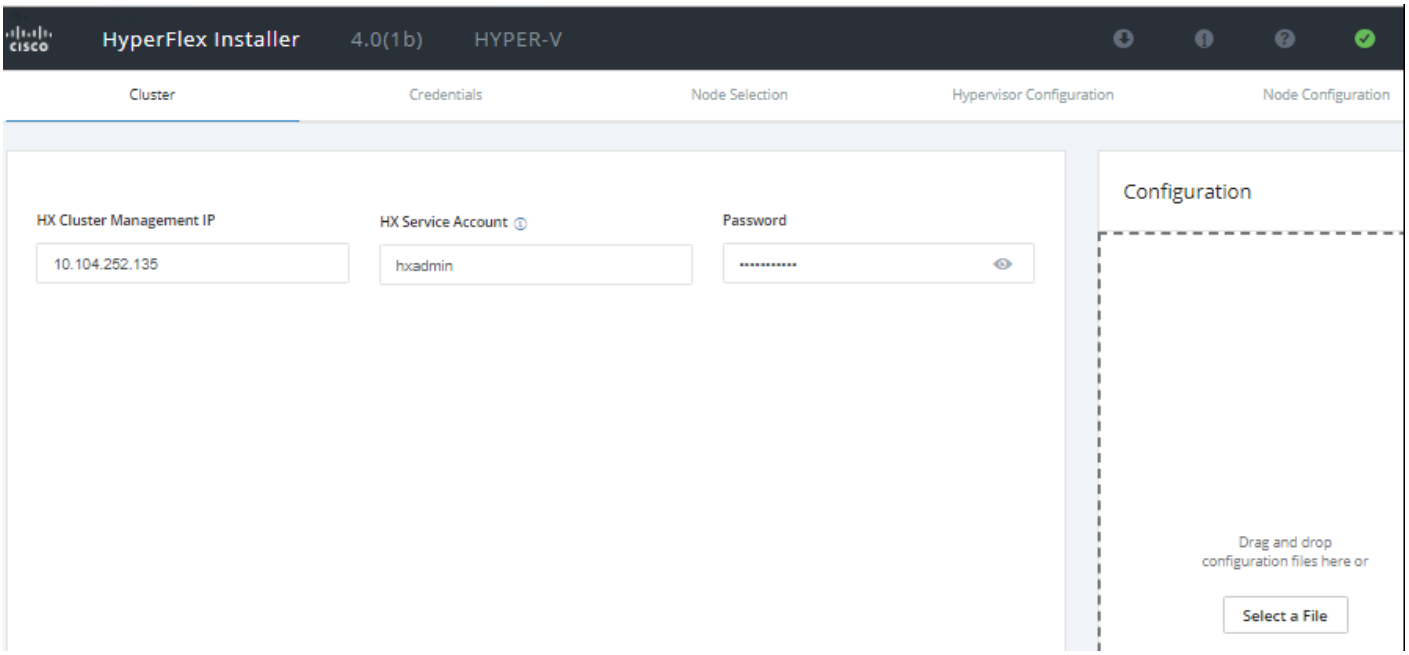


17. Click Continue.

18. On the Credentials Page:

    a. Cisco UCS Manager Credentials – Enter the Cisco UCS Manager DNS hostname or IP address, the admin usernames, and the passwords.

    b. Domain Information – Enter the HX Service Account user name and password. It is recommended to select “Constrained Delegation” here to avoid configuring it manually after the installation completion. Then select “Use HX Service Account”, if HX service account is member of AD Domain Admin group, else provide Domain Admin credentials (which is a one-time requirement).

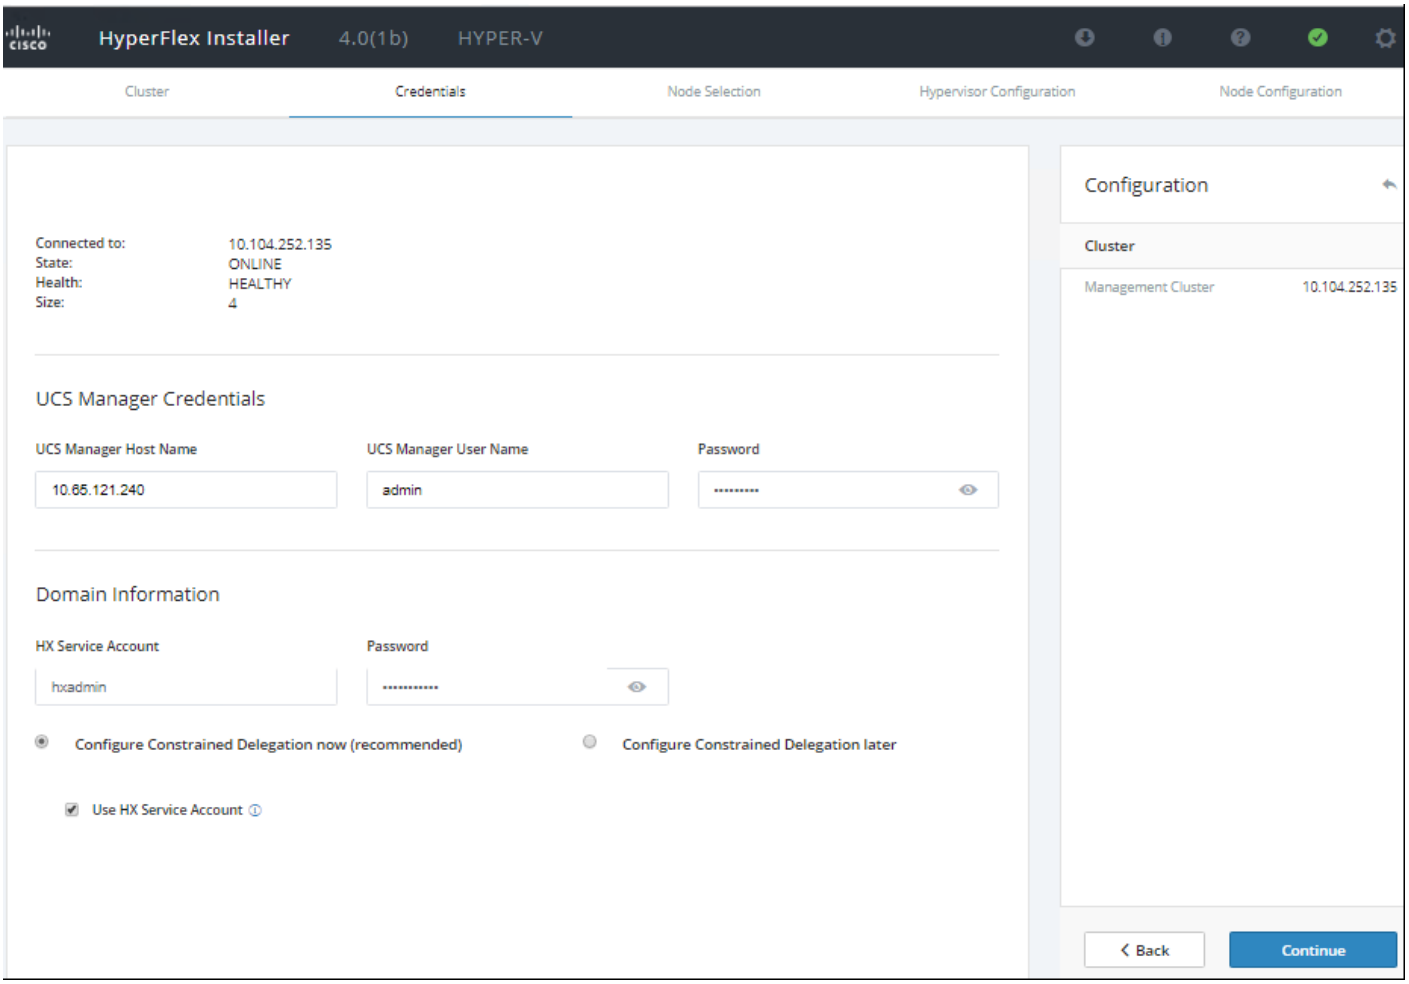


19. On the Node Selection page:

    – Select the unassociated HX servers you want to add to the existing HX cluster.

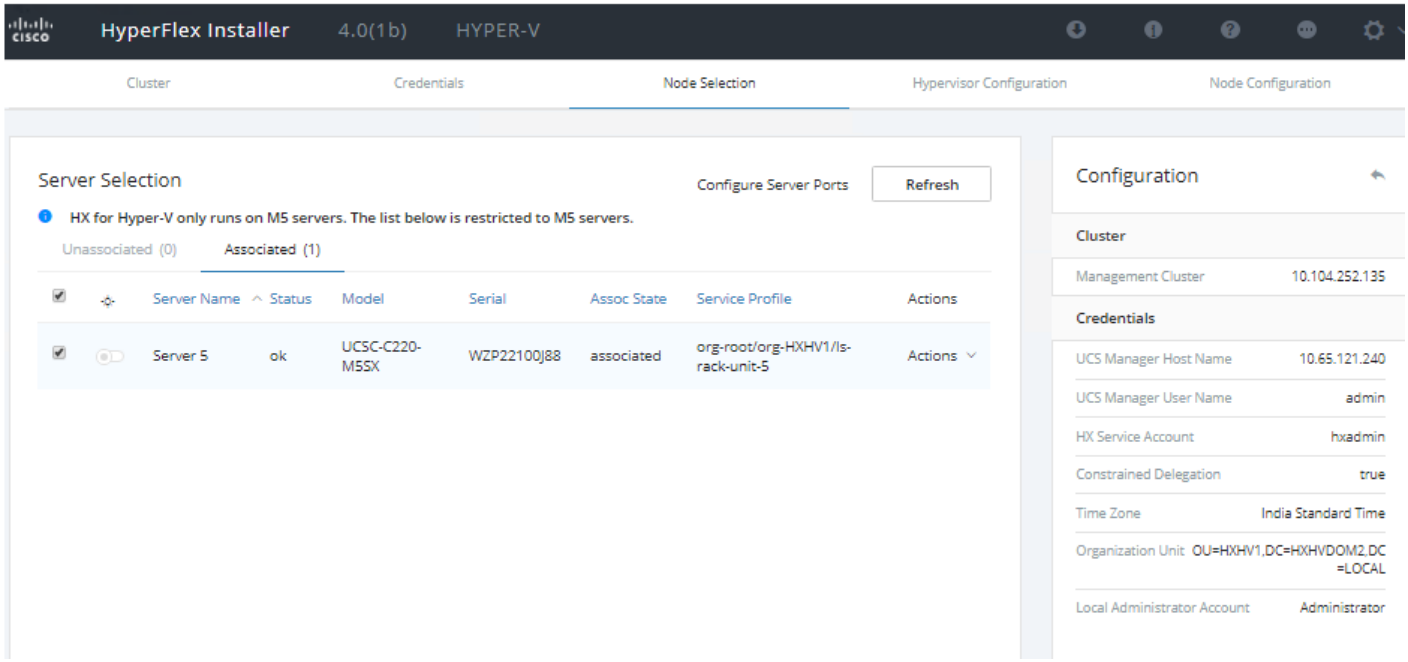


20. Click Continue.

21. On the Hypervisor Configuration page:

    a. VLAN Configuration - Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed here.

    b. Hypervisor Settings - Enter IP addresses and hostnames for the Hypervisors that were created in the pre-installation section phase. The IP addresses will be assigned via Serial over Lan (SoL) through Cisco UCS Manager to the Hyper-V host systems as their management IP addresses.

    c. Hypervisor Credentials – this field is pre-populated and can't be edited.

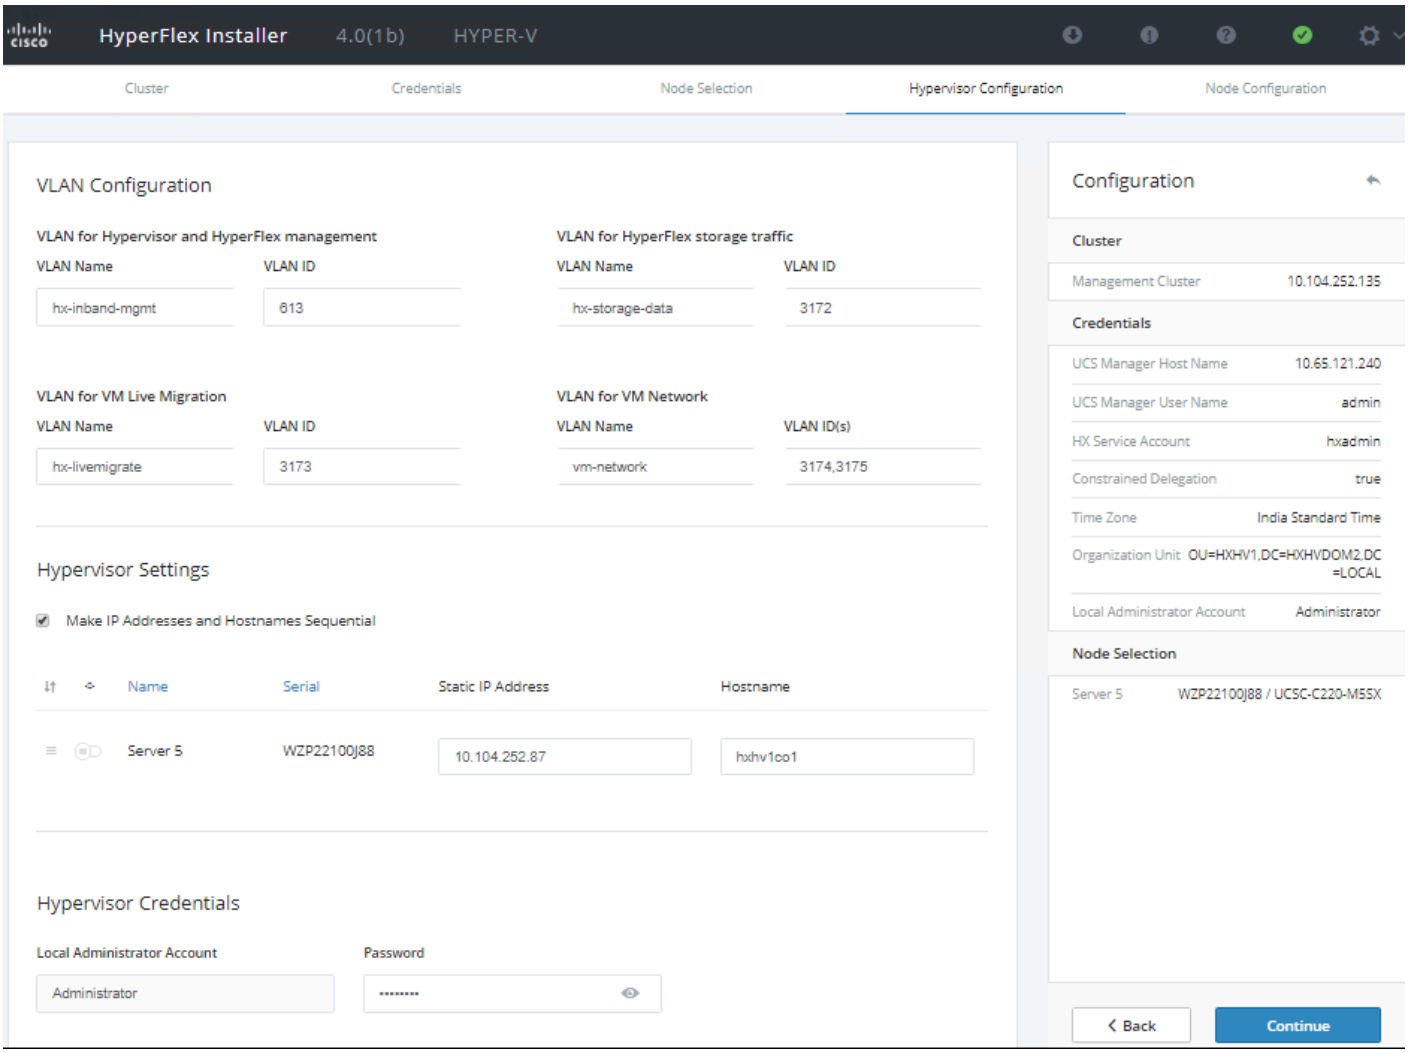


22. Click Continue.

23. On the Node Configuration Page:

- Hypervisor Settings –Enter the Subnet Mask, Gateway and DNS Server IP addresses.
- Failover Cluster Name – Enter the Windows failover cluster name.
- IP Addresses - If you leave the checkbox, Make IP Addresses and Hostnames Sequential as checked, then the installer will automatically fill the rest of the servers sequentially.
- Assign the IP address for the Hypervisor Management and Data networks.
- (Optional) At this step you can manually add more servers for expansion if these servers already have service profiles associated and the hypervisor is ready, by clicking Add Compute Server and then entering the IP addresses for the storage controller management and data networks.

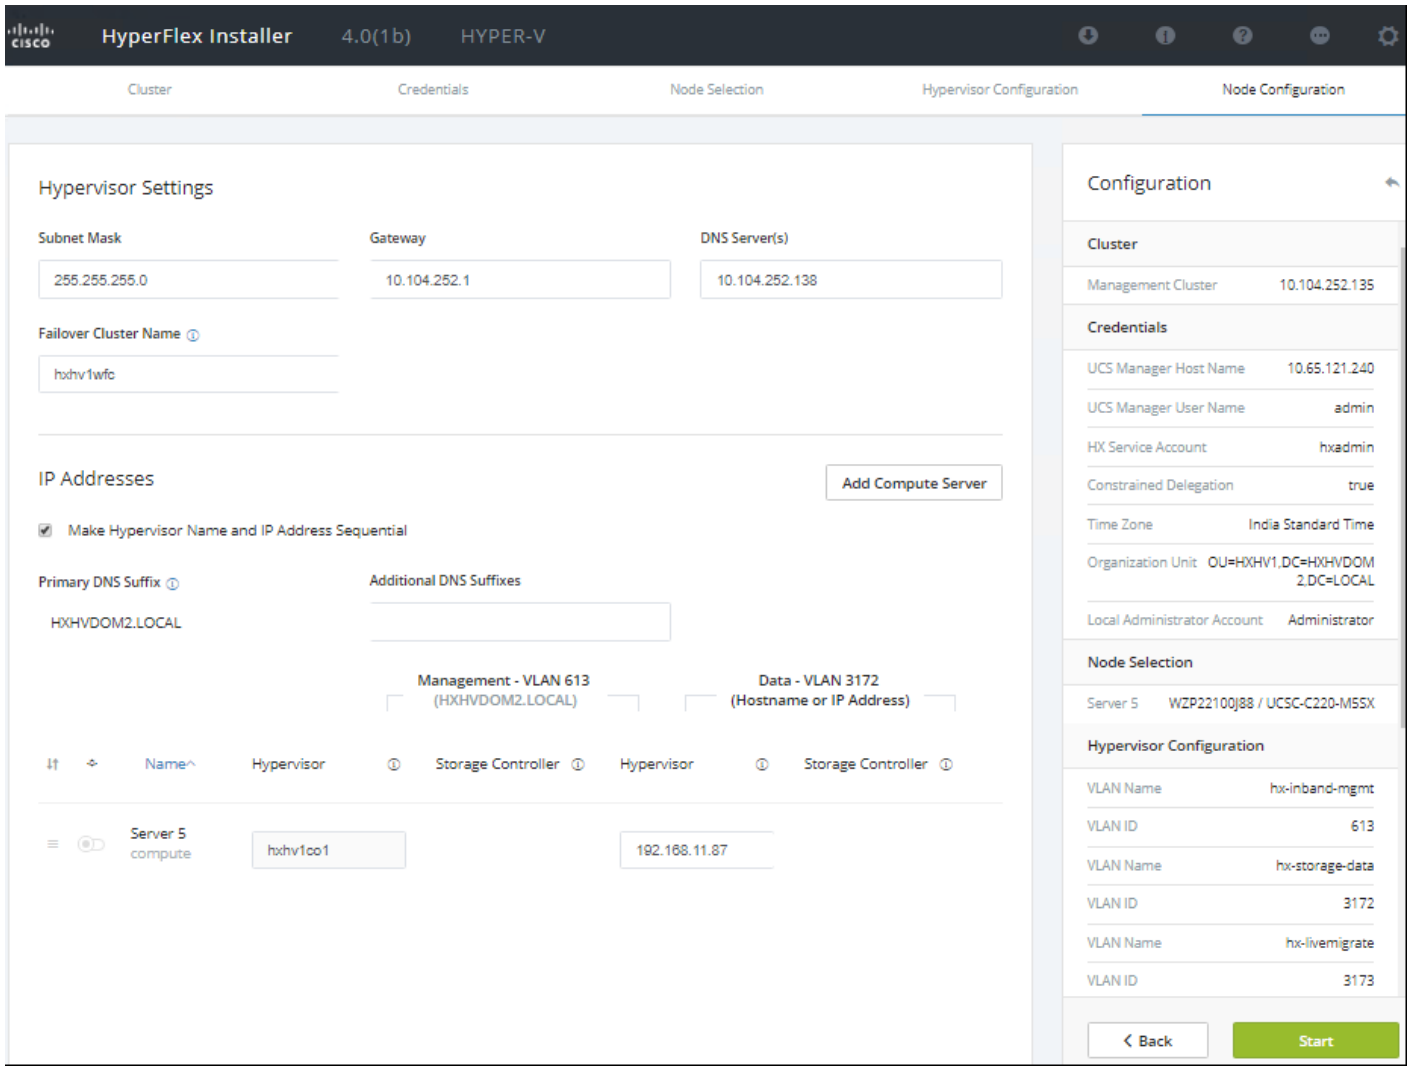


24. Click Start.

25. Click Continue on the warning pop-up message.

26. Validation of the configuration will now start. If there are warnings, you can review and click “Skip Validation” if the warnings are acceptable (such as, you might get the warning from Cisco UCS Manger validation that the guest VLAN is already assigned). If there are no warnings, the validation will automatically continue to the configuration process.

27. The HX installer will now proceed to complete the deployment and perform all the steps listed at the top of the screen along with their status.

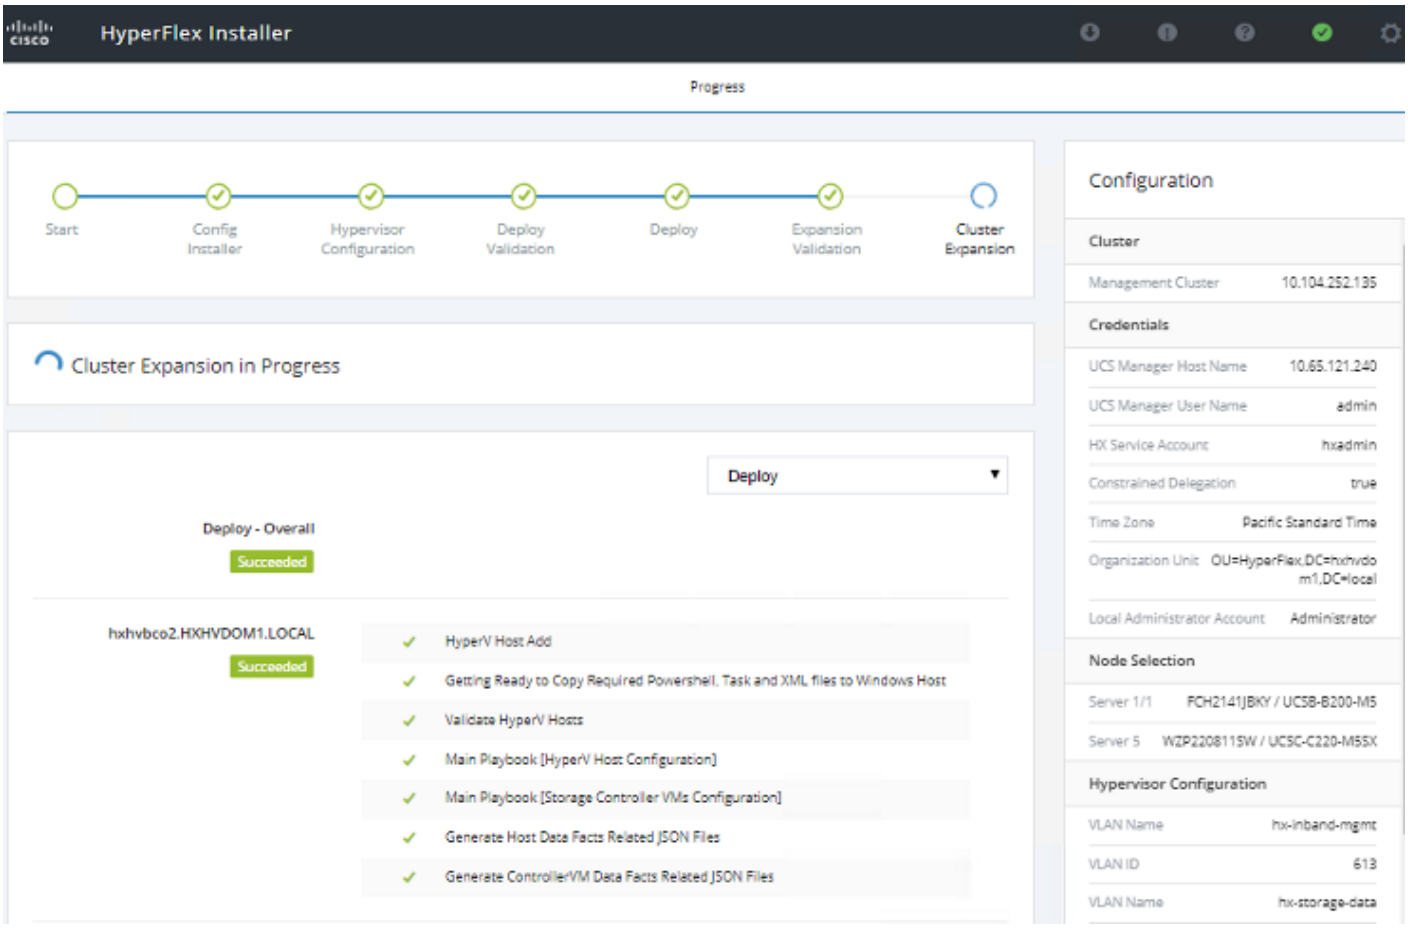


28. You can review the summary screen after the install completes by selecting Summary on the top right of the window.

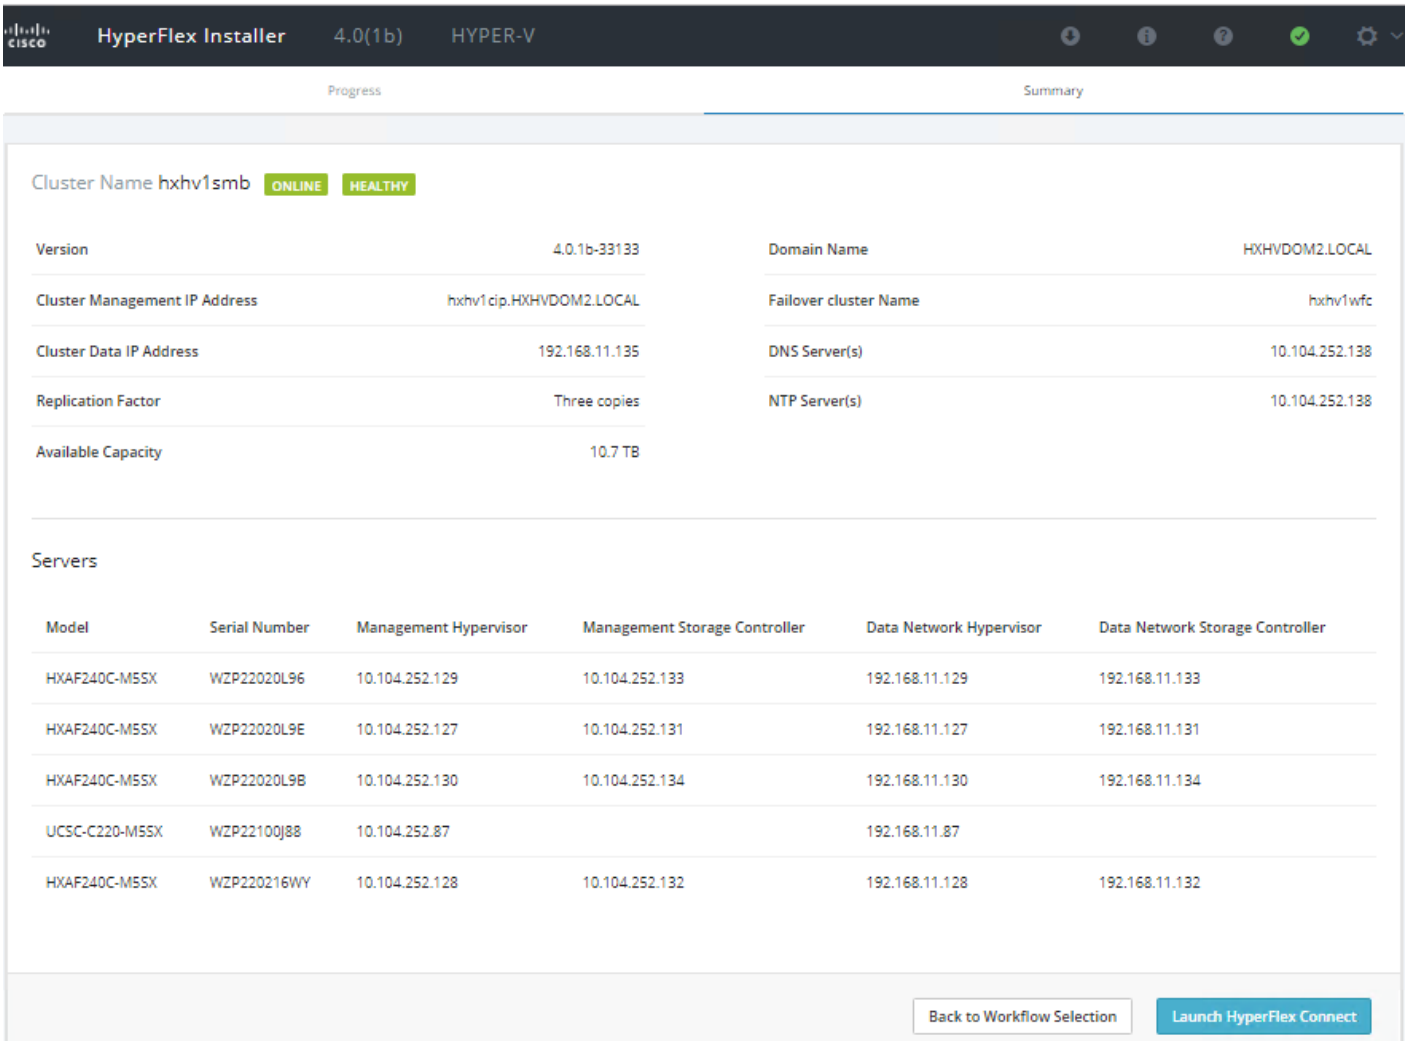


29. After the installation has completed, the new converged node is added to the cluster, and its CPU, and RAM resources are immediately available. Launch HyperFlex Connect to verify the expansion of the cluster using converged node in the Dashboard, Activity and System Information sections as shown below.

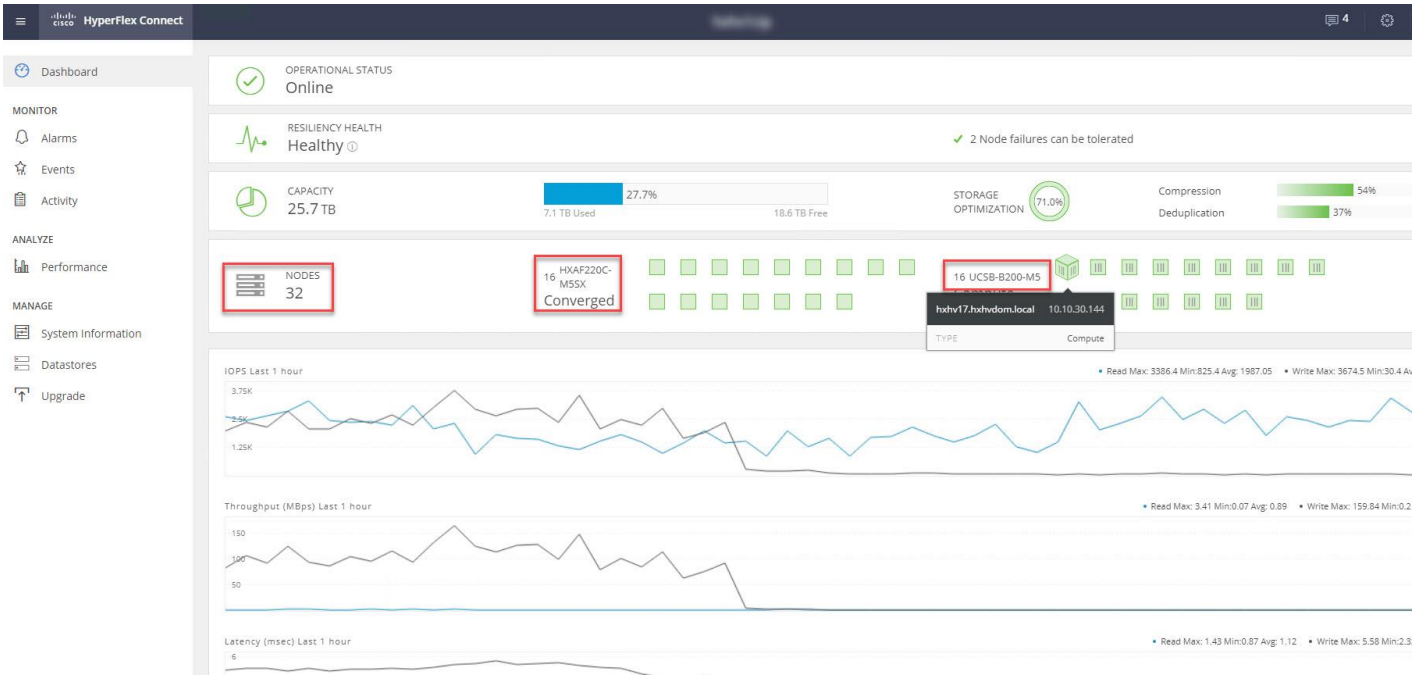


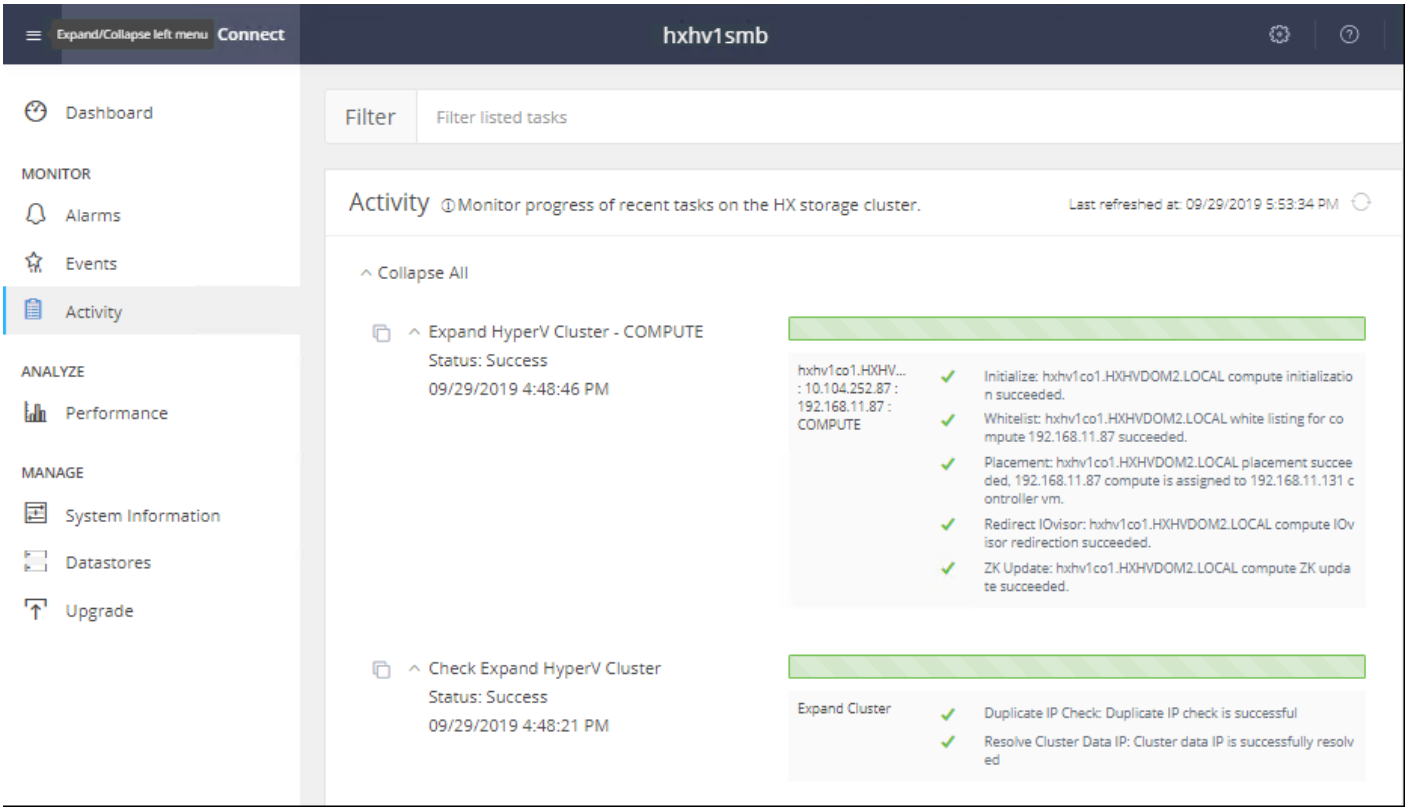


There is no controller virtual machine deployed on the compute-only node. Instead, the SMB client on it is redirected to an existing controller vm on a converged node.

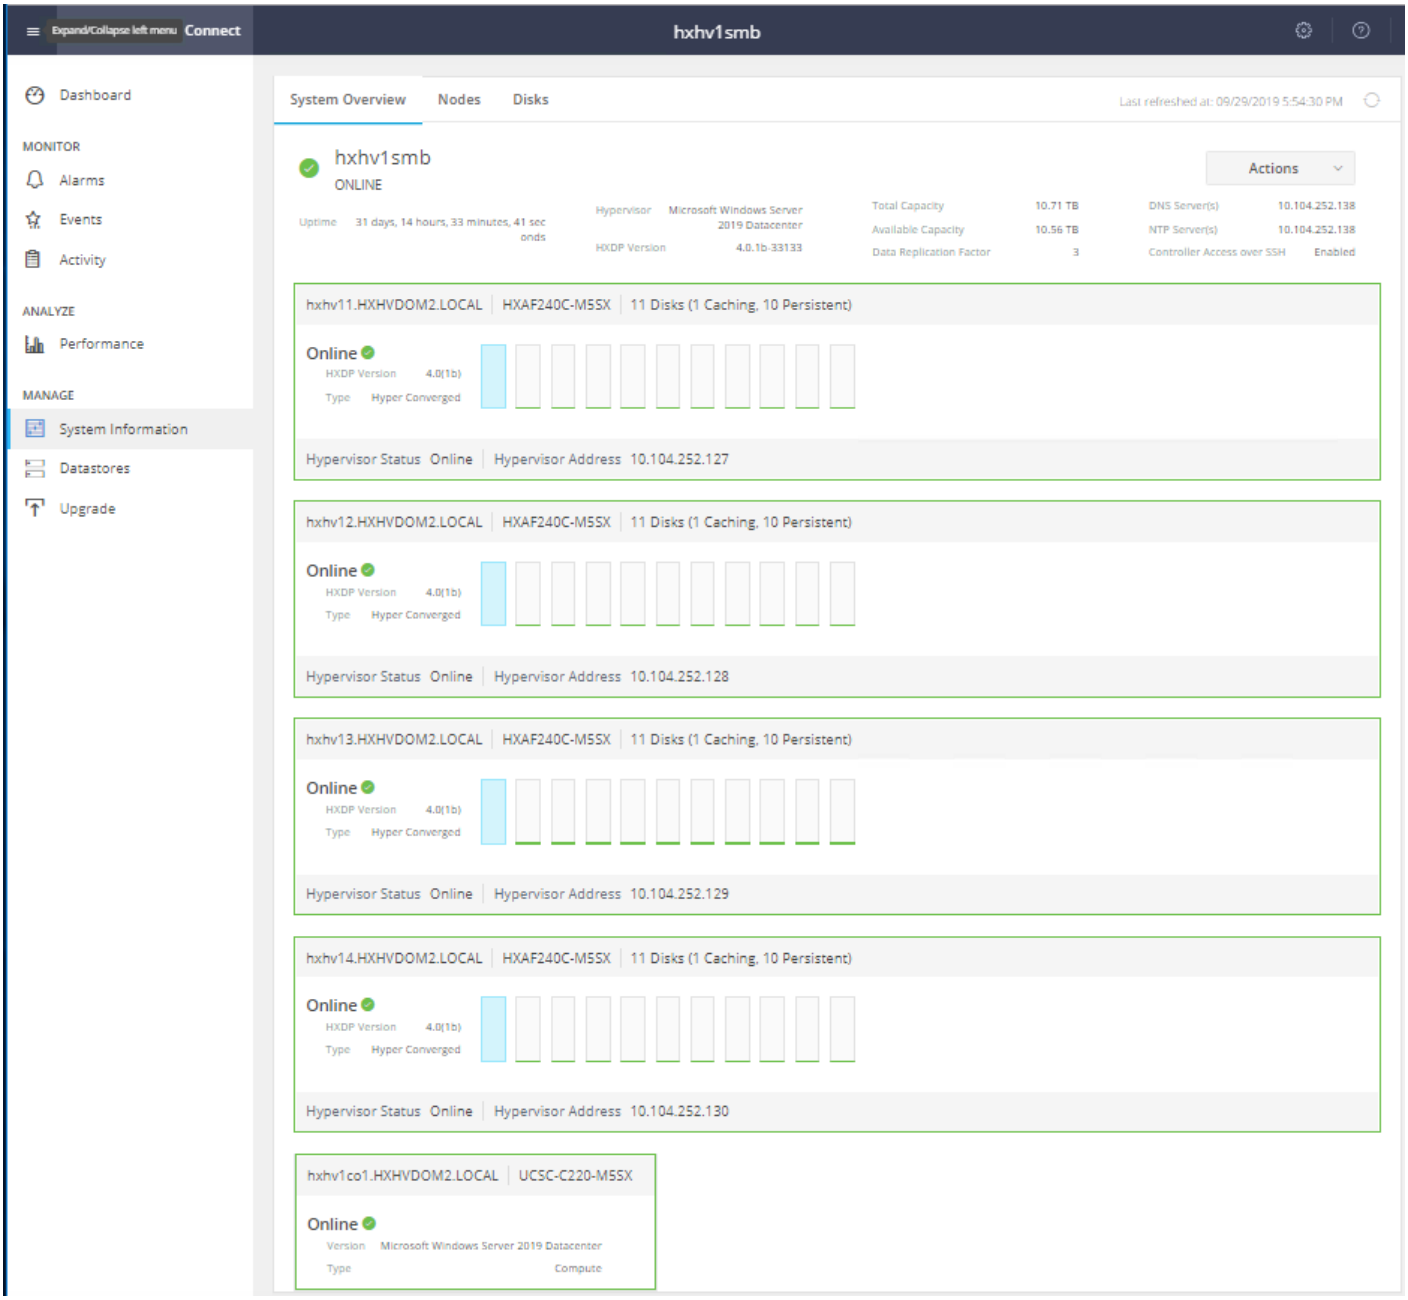
Expand/Collapse left menu
Connect
hxhv1smb
Dashboard
MONITOR
Alarms
Events
Activity
ANALYZE
Performance
MANAGE
System Information
Datastores
Upgrade
System Overview
Nodes
Disks
Last refreshed at: 09/29/2019 5:54:30 PM
hxhv1smb
ONLINE
Actions
Uptime 31 days, 14 hours, 33 minutes, 41 seconds
Hypervisor Microsoft Windows Server 2019 Datacenter
HXDP Version 4.0.1b-33133
Total Capacity 10.71 TB
Available Capacity 10.56 TB
Data Replication Factor 3
DNS Server(s) 10.104.252.138
NTP Server(s) 10.104.252.138
Controller Access over SSH Enabled
hxhv11.HXHVDOM2.LOCAL | HXAF240C-M5SX | 11 Disks (1 Caching, 10 Persistent)
Online
HXDP Version 4.0(1b)
Type Hyper Converged
Hypervisor Status Online | Hypervisor Address 10.104.252.127
hxhv12.HXHVDOM2.LOCAL | HXAF240C-M5SX | 11 Disks (1 Caching, 10 Persistent)
Online
HXDP Version 4.0(1b)
Type Hyper Converged
Hypervisor Status Online | Hypervisor Address 10.104.252.128
hxhv13.HXHVDOM2.LOCAL | HXAF240C-M5SX | 11 Disks (1 Caching, 10 Persistent)
Online
HXDP Version 4.0(1b)
Type Hyper Converged
Hypervisor Status Online | Hypervisor Address 10.104.252.129
hxhv14.HXHVDOM2.LOCAL | HXAF240C-M5SX | 11 Disks (1 Caching, 10 Persistent)
Online
HXDP Version 4.0(1b)
Type Hyper Converged
Hypervisor Status Online | Hypervisor Address 10.104.252.130
hxhv1co1.HXHVDOM2.LOCAL | UCSC-C220-M5SX
Online
Version Microsoft Windows Server 2019 Datacenter
Type Compute

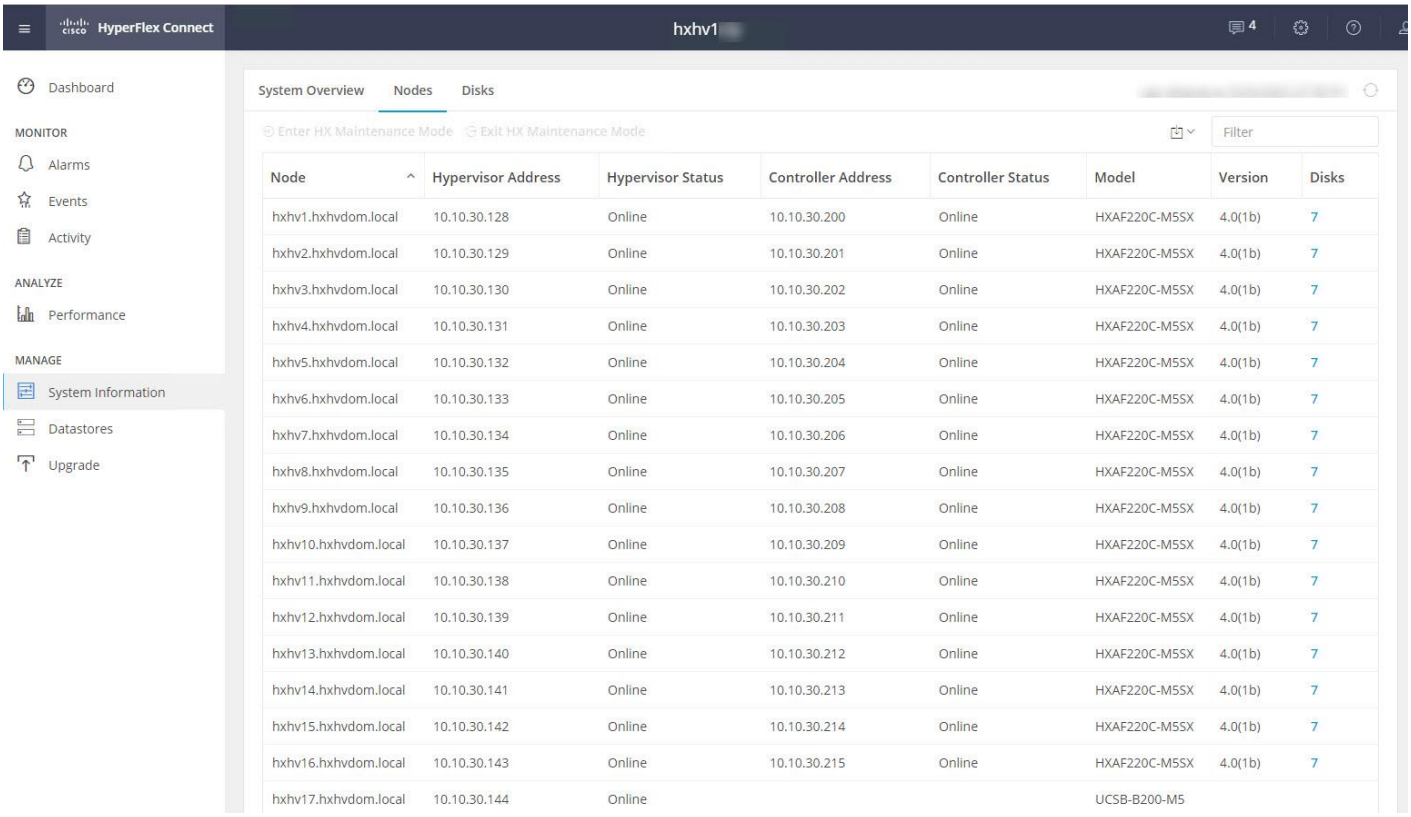
HyperFlex Connect
hxhv1
4
Dashboard
MONITOR
Alarms
Events
Activity
ANALYZE
Performance
MANAGE
System Information
Datastores
Upgrade
System Overview
Nodes
Disks
Enter HX Maintenance Mode
Exit HX Maintenance Mode
Filter
Node
Hypervisor Address
Hypervisor Status
Controller Address
Controller Status
Model
Version
Disks
hxhv1.hxhvdom.local 10.10.30.128 Online 10.10.30.200 Online HXAF220C-M5SX 4.0(1b) 7
hxhv2.hxhvdom.local 10.10.30.129 Online 10.10.30.201 Online HXAF220C-M5SX 4.0(1b) 7
hxhv3.hxhvdom.local 10.10.30.130 Online 10.10.30.202 Online HXAF220C-M5SX 4.0(1b) 7
hxhv4.hxhvdom.local 10.10.30.131 Online 10.10.30.203 Online HXAF220C-M5SX 4.0(1b) 7
hxhv5.hxhvdom.local 10.10.30.132 Online 10.10.30.204 Online HXAF220C-M5SX 4.0(1b) 7
hxhv6.hxhvdom.local 10.10.30.133 Online 10.10.30.205 Online HXAF220C-M5SX 4.0(1b) 7
hxhv7.hxhvdom.local 10.10.30.134 Online 10.10.30.206 Online HXAF220C-M5SX 4.0(1b) 7
hxhv8.hxhvdom.local 10.10.30.135 Online 10.10.30.207 Online HXAF220C-M5SX 4.0(1b) 7
hxhv9.hxhvdom.local 10.10.30.136 Online 10.10.30.208 Online HXAF220C-M5SX 4.0(1b) 7
hxhv10.hxhvdom.local 10.10.30.137 Online 10.10.30.209 Online HXAF220C-M5SX 4.0(1b) 7
hxhv11.hxhvdom.local 10.10.30.138 Online 10.10.30.210 Online HXAF220C-M5SX 4.0(1b) 7
hxhv12.hxhvdom.local 10.10.30.139 Online 10.10.30.211 Online HXAF220C-M5SX 4.0(1b) 7
hxhv13.hxhvdom.local 10.10.30.140 Online 10.10.30.212 Online HXAF220C-M5SX 4.0(1b) 7
hxhv14.hxhvdom.local 10.10.30.141 Online 10.10.30.213 Online HXAF220C-M5SX 4.0(1b) 7
hxhv15.hxhvdom.local 10.10.30.142 Online 10.10.30.214 Online HXAF220C-M5SX 4.0(1b) 7
hxhv16.hxhvdom.local 10.10.30.143 Online 10.10.30.215 Online HXAF220C-M5SX 4.0(1b) 7
hxhv17.hxhvdom.local 10.10.30.144 Online UCSB-B200-M5
MANAGE
System Information
Datastores
Upgrade
hxhv17.hxhvdom.local 10.10.30.144 Online UCSB-B200-M5
hxhv18.hxhvdom.local 10.10.30.145 Online UCSB-B200-M5
hxhv19.hxhvdom.local 10.10.30.146 Online UCSB-B200-M5
hxhv20.hxhvdom.local 10.10.30.147 Online UCSB-B200-M5
hxhv21.hxhvdom.local 10.10.30.148 Online UCSB-B200-M5
hxhv22.hxhvdom.local 10.10.30.149 Online UCSB-B200-M5
hxhv23.hxhvdom.local 10.10.30.150 Online UCSB-B200-M5
hxhv24.hxhvdom.local 10.10.30.151 Online UCSB-B200-M5
hxhv25.hxhvdom.local 10.10.30.152 Online UCSB-B200-M5
Showing 1 - 25 of 32
1

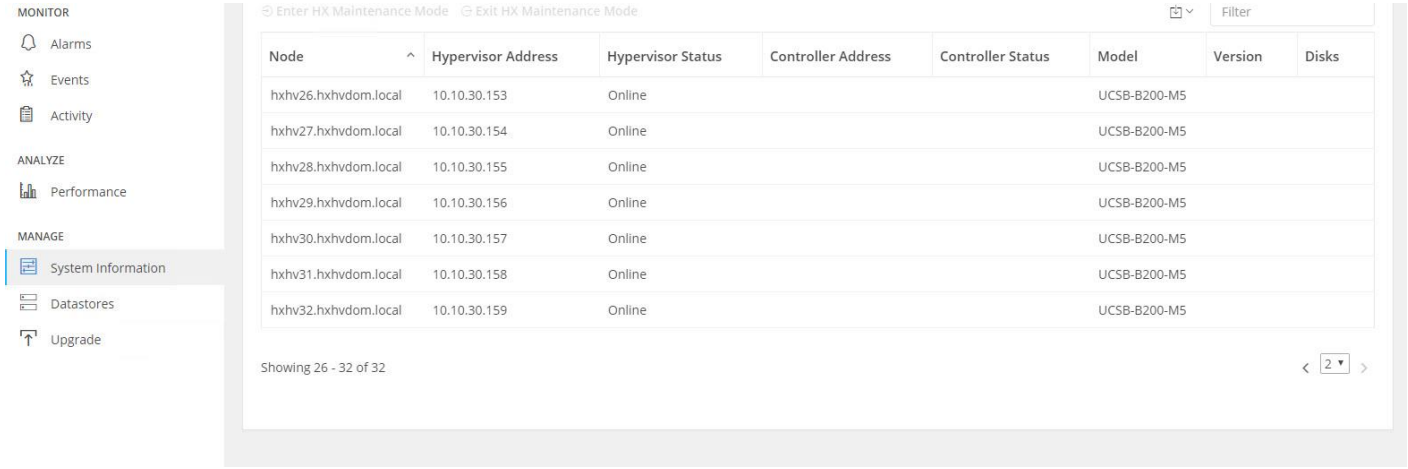


30. Log into the Windows failover cluster manager to verify the compute nodes are now part of the Windows Failover cluster.

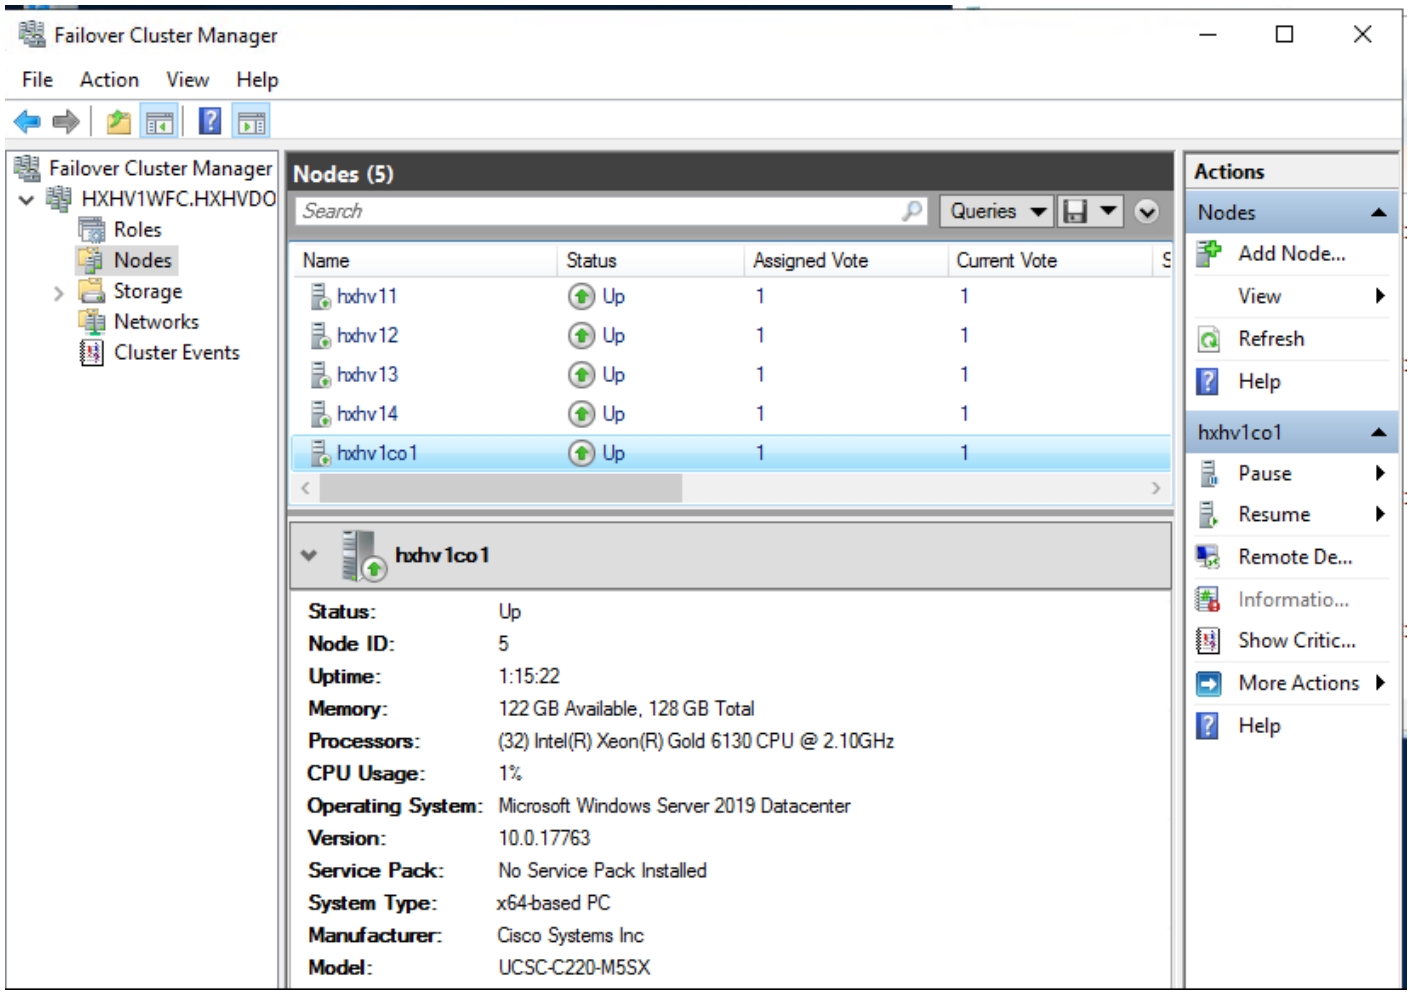


31. Refer to section Post Installation Tasks> Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces.

32. After assigning IP addresses, verify the jumbo frame settings on the live migration network of newly expanded compute node using the below PowerShell cmdlet

    a. From local host, use the below command

    ```
    Get-NetAdapterAdvancedProperty -name hv-livemigrate* | Where-Object
    {$_.DisplayName -Match "Jumbo*"}
    ```

    b. From a remote management host, use the below command

    ```
    Invoke-Command -ComputerName hxhv1co1 -ScriptBlock {Get-
    NetAdapterAdvancedProperty -name hv-livemigrate* | Where-Object {$_.DisplayName -
    Match "Jumbo*"}}
    ```

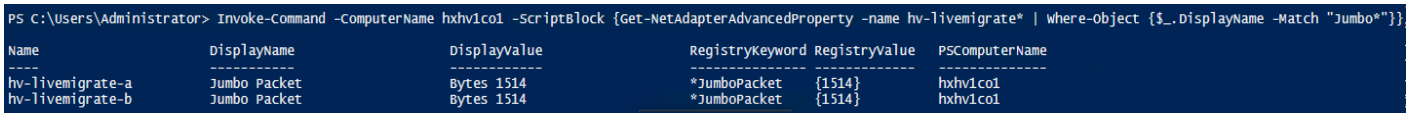


33. For better performance of live migration of VMs, configure the mtu size to 9014 as shown below.

    a. From local host, use the below command:

    ```
    Set-NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword
    "*JumboPacket" -RegistryValue 9014
    ```

    b. From a remote management host, use the below command:

    ```
    Invoke-Command -ComputerName hxhv1co1 -ScriptBlock {Set-
    NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword
    "*JumboPacket" -RegistryValue 9014}
    ```

34. On the newly added Compute node, open Server Manager > Local Server and click on Enabled next to the NIC Teaming.

35. On the NIC Teaming window, right-click “team-hx-livemigration” and click properties under the “Adapter and Interfaces” section as shown below.

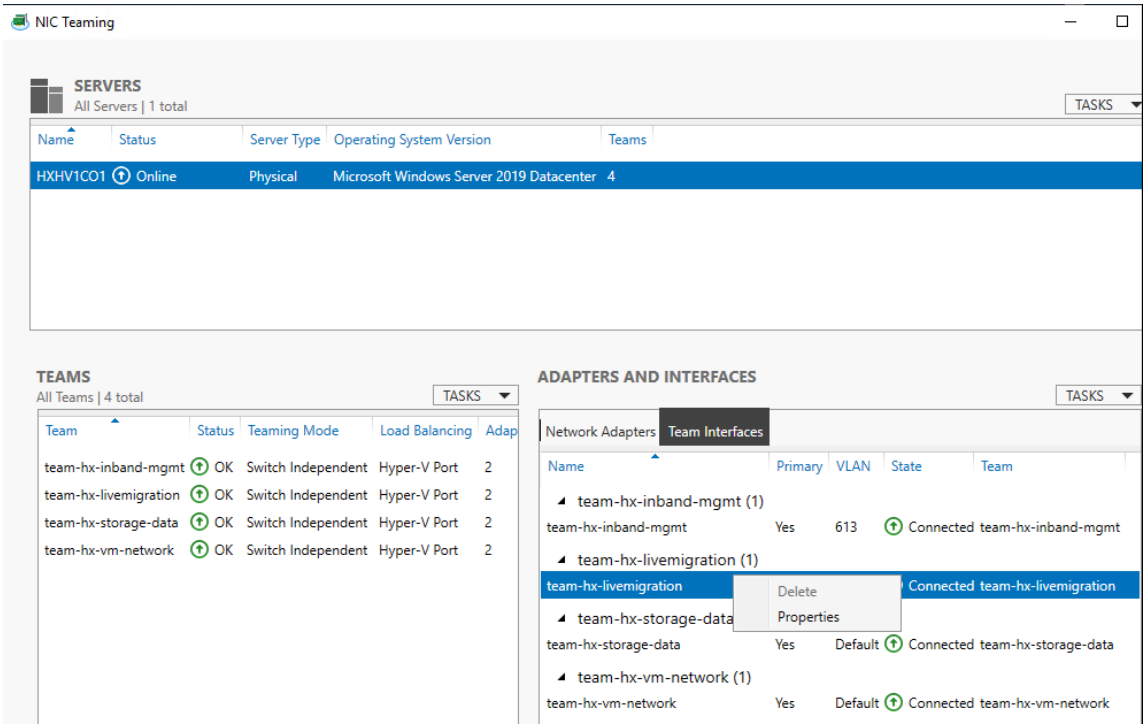


36. Select “Specific VLAN” and enter a VLAN ID used for live migration network in the NIC Teaming Adapters Properties window. Also make sure the name under General information is exactly “team-hx-livemigration” as shown in the screenshot below.

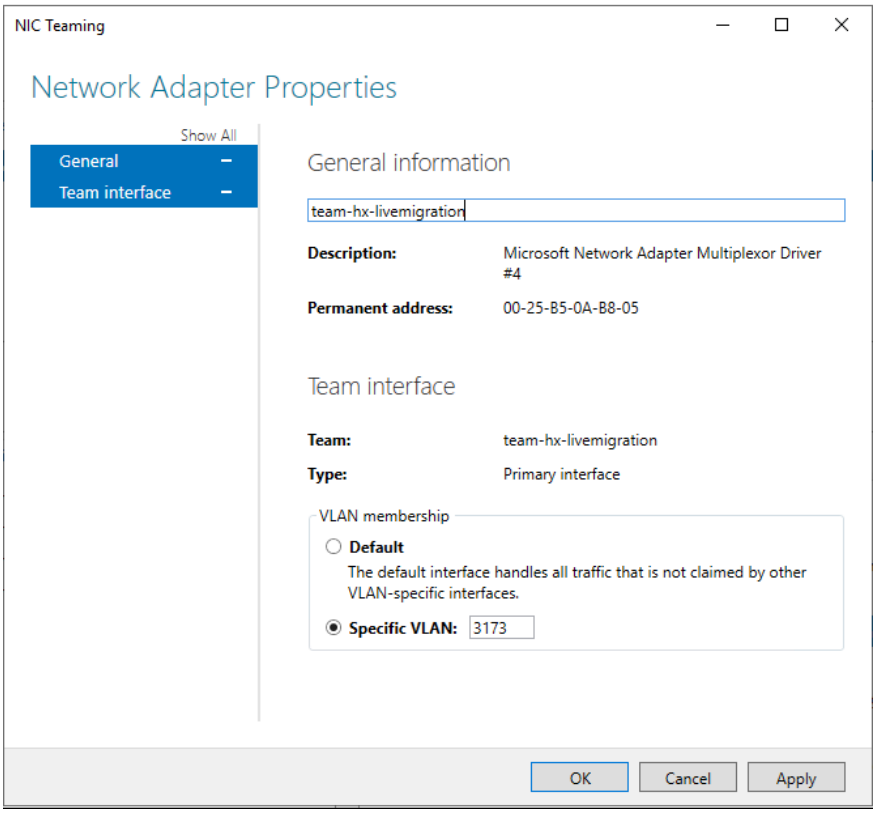

37. Verify the configuration in the main NIC Teaming window as shown below.

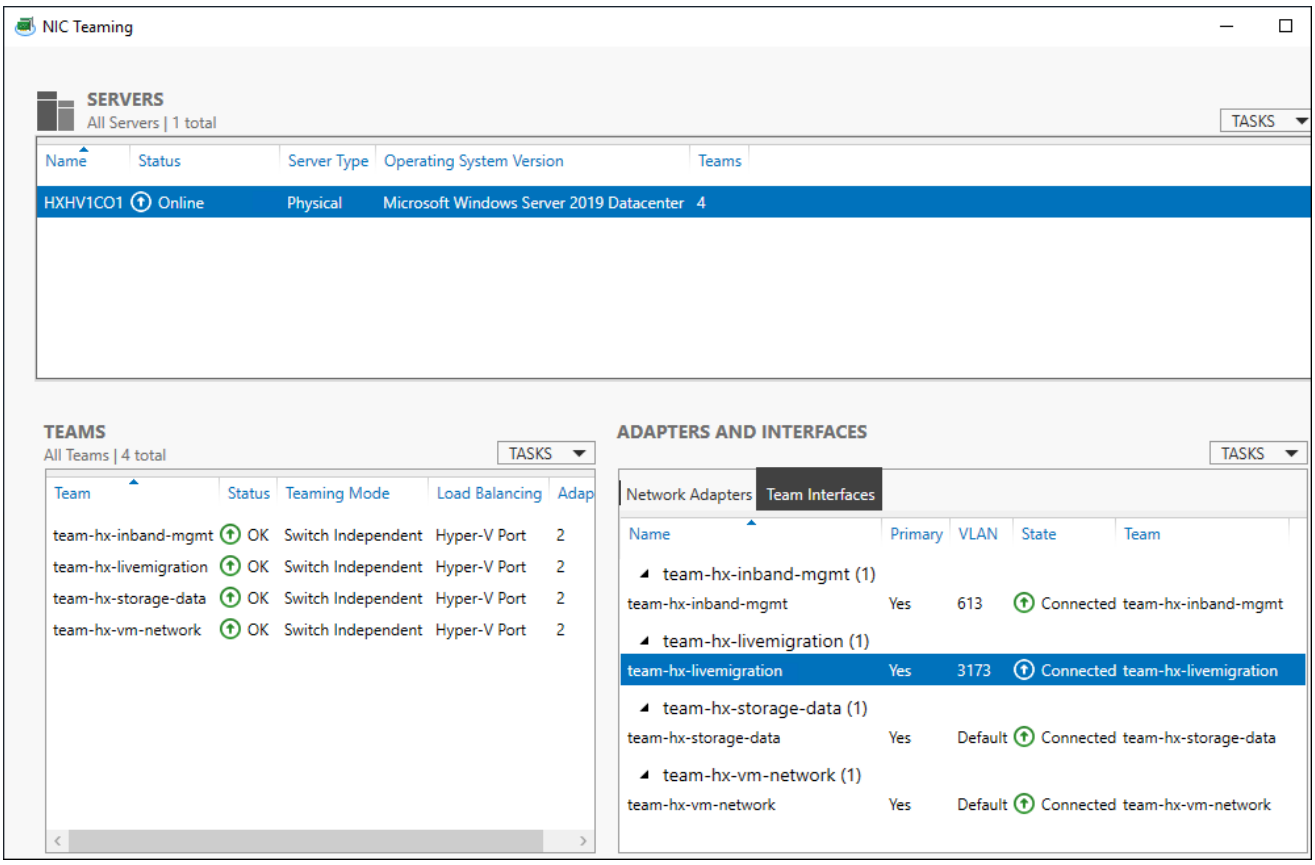


38. Test/validate by pinging the LM network interface of another hx node from the newly added compute node if jumbo frame is working as shown below:

```
ping -l 8972 -f 192.168.73.128
```

```
PS C:\Users\Administrator.HXHVDOM2> ping -l 8972 -f 192.168.73.128

Pinging 192.168.73.128 with 8972 bytes of data:
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
Reply from 192.168.73.128: bytes=8972 time<1ms TTL=128
```

39. Open Hyper-V Manger on the newly added compute node, select the server and in the Actions pane, click Hyper-V Settings > Live Migrations.

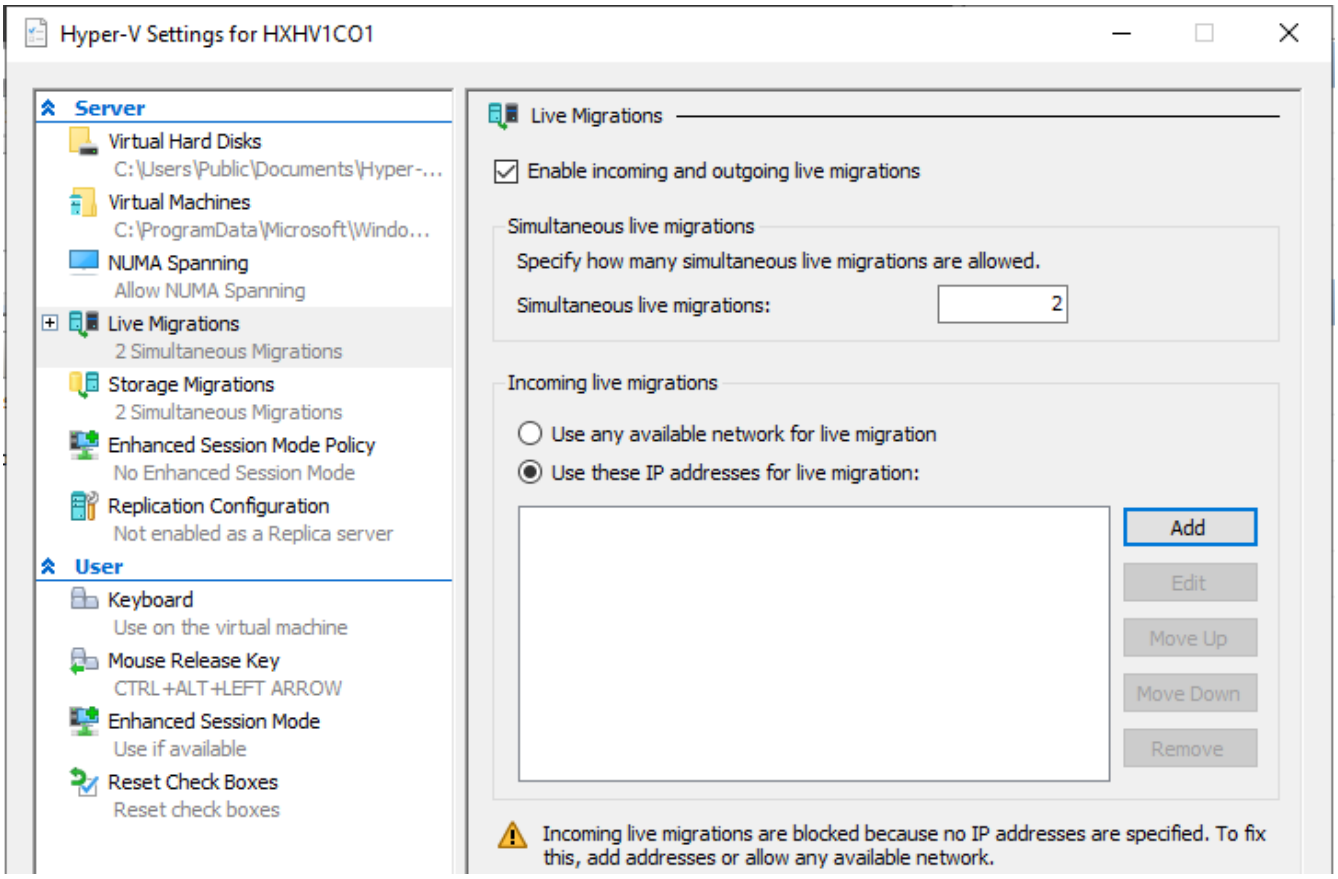


40. Under Incoming live migrations, select “Use these IP addresses for live migration:” and click Add to add the IP address assigned to the live migration network interface on this host as shown below and click OK.

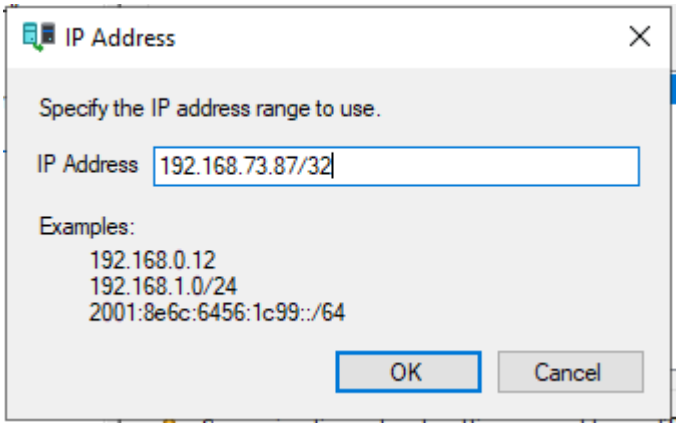


41. The following screenshot shows the IP address added under “use these IP addresses for live migration:” in the previous steps.

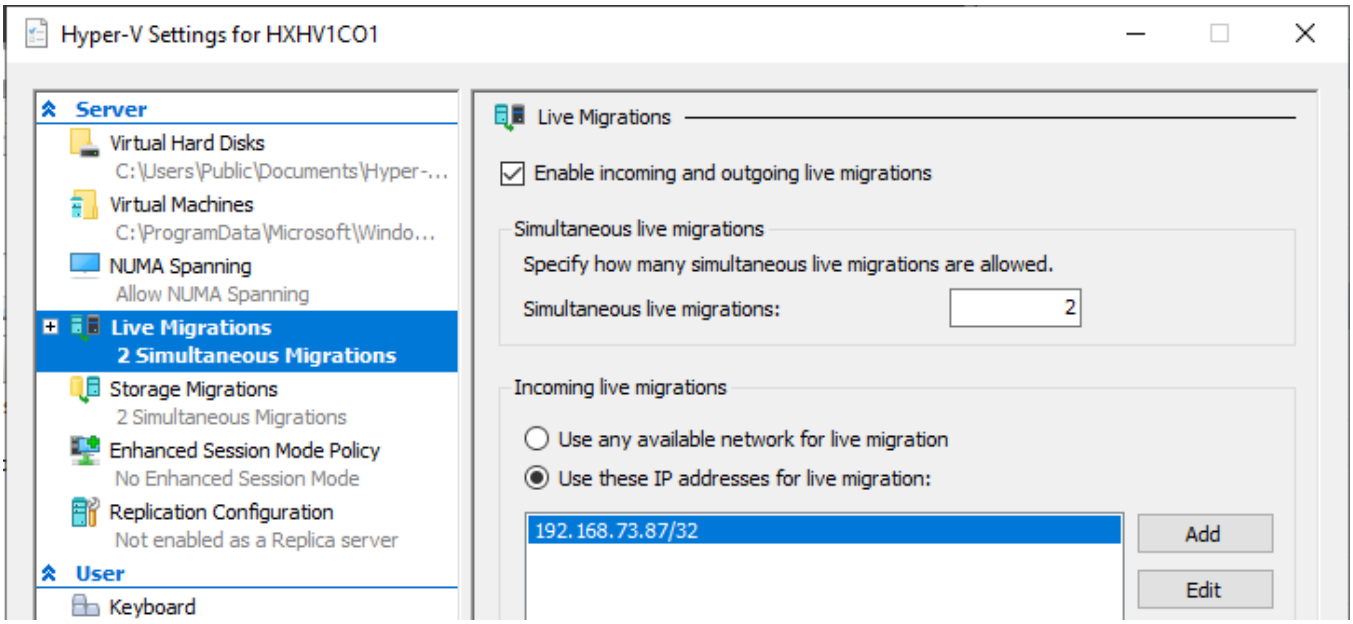


42. Verify the configuration by live migrating a VM from a converged node to a compute node.

```
PS C:\Users\administrator.HXHVDOM2> Get-VM

Name     State   CPUUsage(%) MemoryAssigned(M) Uptime            Status             Version
----     -----   ----------- ----------------- ------            ------             -------
DemoVM   Running 0           4096              00:00:35.9470000  Operating normally 9.0
StCtlVM  Running 0           73728             6.09:23:45.3590000 Operating normally 9.0
VM01     Running 0           2048              00:33:16.8980000  Operating normally 9.0
VM05     Running 0           8192              6.09:17:27.6200000 Operating normally 9.0
VM06     Running 0           8192              6.09:17:14.1760000 Operating normally 9.0


PS C:\Users\administrator.HXHVDOM2> Measure-Command { Move-VM -Name DemoVM -DestinationHost HXHV1CO1 }


Days              : 0
Hours             : 0
Minutes           : 0
Seconds           : 3
Milliseconds      : 795
Ticks             : 37955036
TotalDays         : 4.39294398148148E-05
TotalHours        : 0.00105430655555556
TotalMinutes      : 0.0632583933333333
TotalSeconds      : 3.7955036
TotalMilliseconds : 3795.5036
```

43. Make sure the below VM setting is selected before live migration if there exists different processor version between source and target hosts.

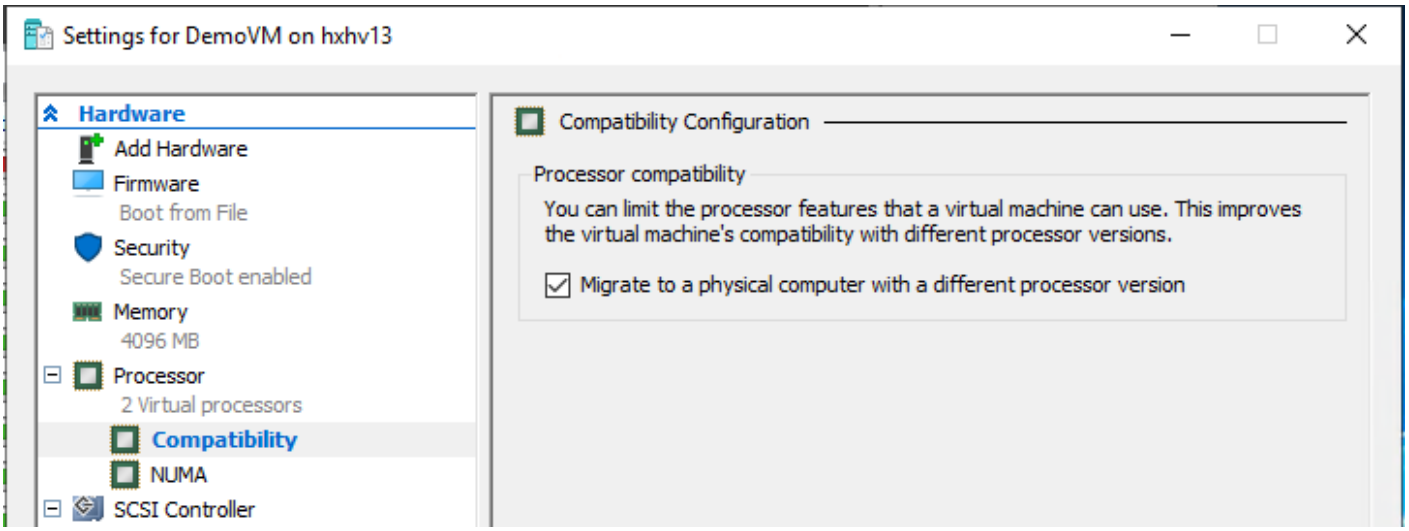
Settings for DemoVM on hxhv13
Hardware
Add Hardware
Firmware
Boot from File
Security
Secure Boot enabled
Memory
4096 MB
Processor
2 Virtual processors
Compatibility
NUMA
SCSI Controller
Compatibility Configuration
Processor compatibility
You can limit the processor features that a virtual machine can use. This improves the virtual machine's compatibility with different processor versions.
Migrate to a physical computer with a different processor version

# Management

## HyperFlex Connect

Cisco HyperFlex Connect provides robust, secure, and simple management in an intuitive user interface. It lets you manage and monitor your clusters anywhere, anytime, and delivers metrics to support your entire HyperFlex management lifecycle. HyperFlex Connect is an HTML5 web-based GUI tool which runs on all the HX nodes and is accessible through the cluster management IP address.

To manage the HyperFlex cluster using HyperFlex Connect, follow these steps:

1. Using a web browser, open the HyperFlex cluster's management IP address through HTTPS, for example, https://10.29.149.230
2. Enter a local credential, such as local/root, and the password.
3. Click Login.
4. The Dashboard view displays after a successful login.

**Figure 98 Cisco HyperFlex Connect – Login Page**

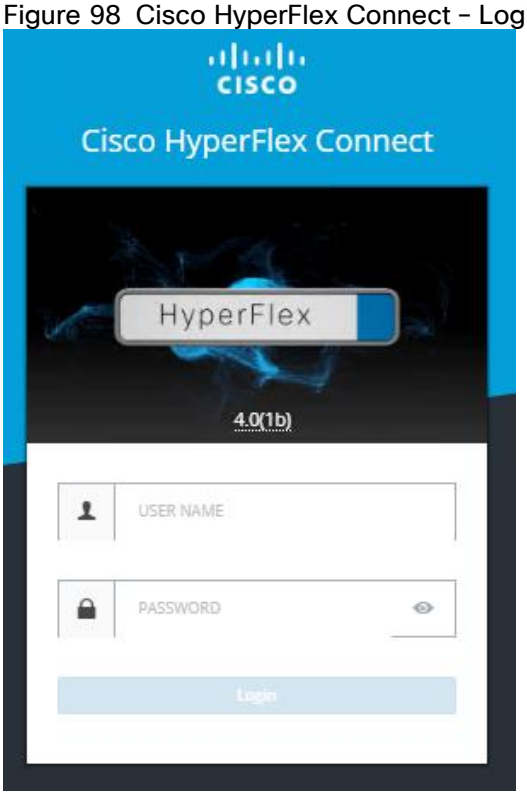

Figure 99 Cisco HyperFlex Connect – Dashboard

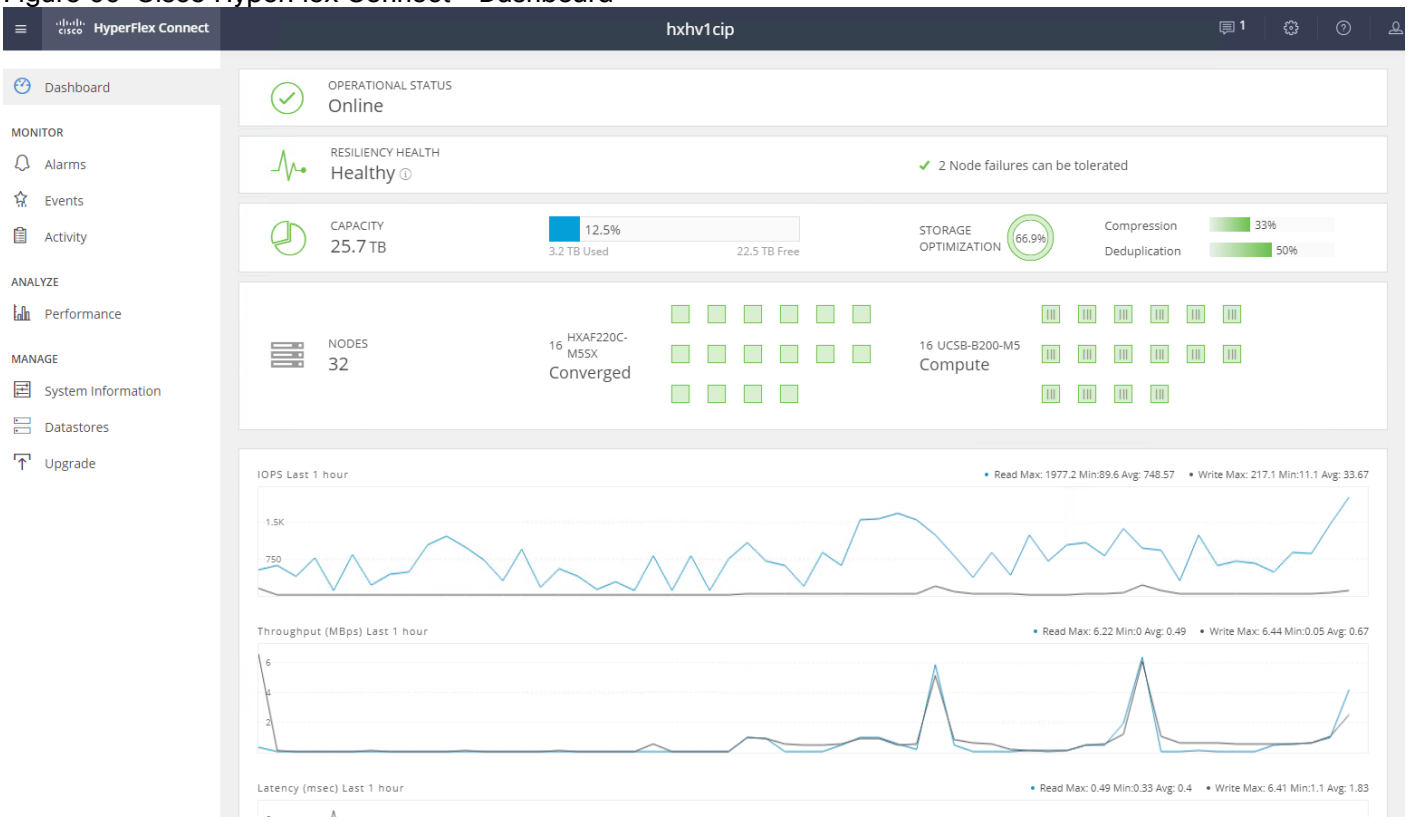


## Dashboard

From the Dashboard view, the following elements are presented:

- Cluster operational status, overall cluster health, and the cluster's current node failure tolerance.
- Cluster storage capacity used and free space, compression and deduplication savings, and overall cluster storage optimization statistics.
- Cluster size and individual node health.
- Cluster IOPs, storage throughput, and latency for the past 1 hour.

## Monitor

HyperFlex Connect provides for additional monitoring capabilities, including:

- Event Log: The cluster event log can be viewed, specific events can be filtered for, and the log can be exported.
- Activity Log: Recent job activity, such as ReadyClones can be viewed and the status can be monitored.

Figure 100 Cisco HyperFlex Connect – Events

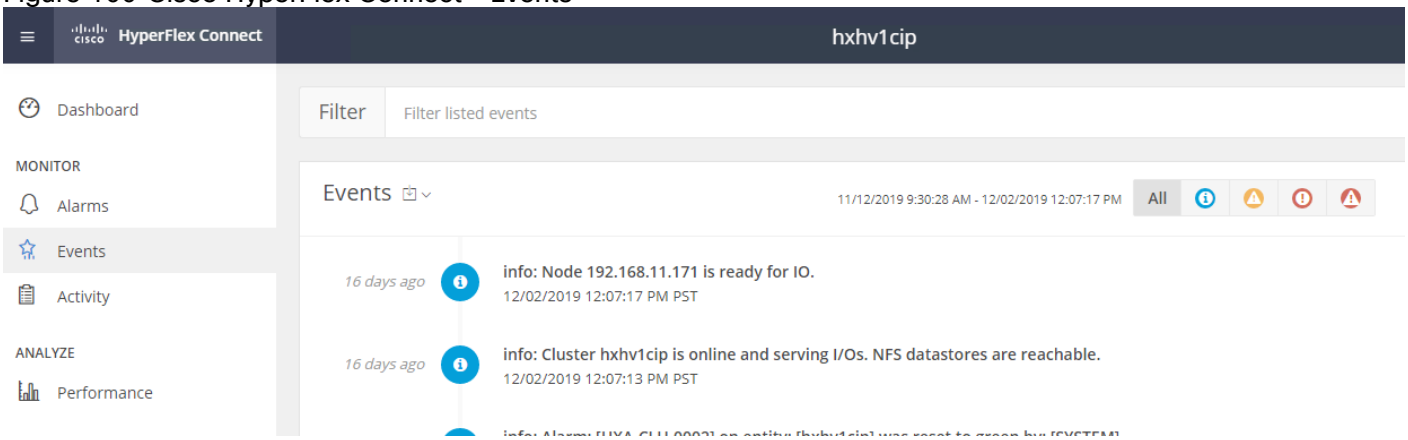


## Analyze

The historical and current performance of the HyperFlex cluster can be analyzed through the built-in performance charts. The default view shows read and write IOPs, bandwidth, and latency over the past hour (1) for the entire cluster. Views can be customized to see individual nodes or datastores, and change the timeframe shown in the charts.

Figure 101 Cisco HyperFlex Connect – Performance

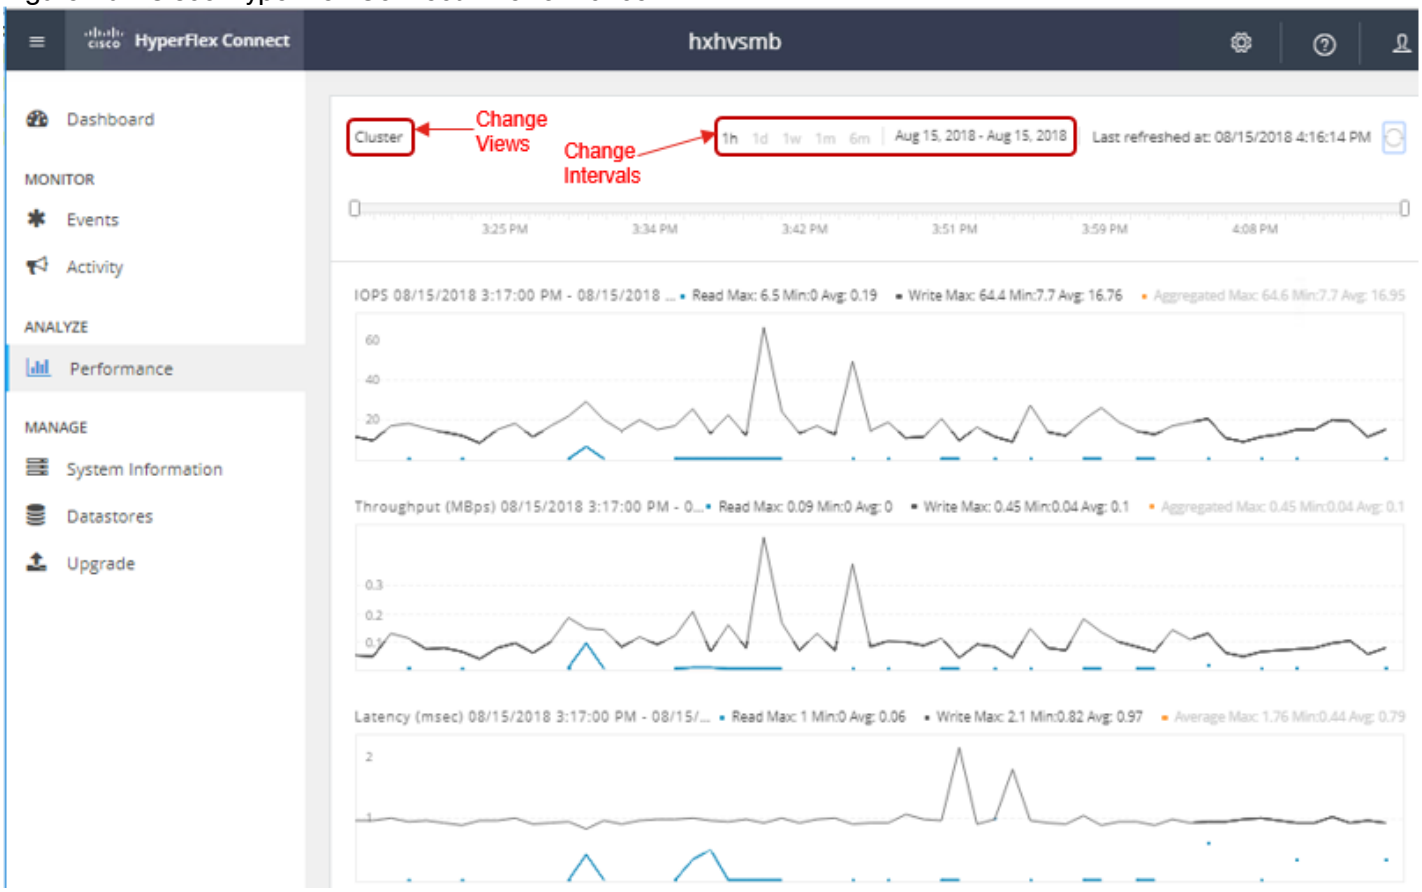

## Manage

HyperFlex Connect presents several views and elements for managing the HyperFlex cluster:

- System Information: Presents a detailed view of the cluster configuration, software revisions, hosts, disks, and cluster uptime. Support bundles can be generated to be shared with Cisco TAC when technical support is needed. Views of the individual nodes and the individual disks are available. In these views, nodes can be placed into HX Maintenance Mode, and disks can be securely erased, as described later in this document.
- Datastores: Presents the datastores present in the cluster, and allows for datastores to be created, mounted, unmounted, edited or deleted, as described earlier in this document as part of the cluster setup.
- Upgrade: Upgrades to the HXDP software and Cisco UCS firmware can be initiated from this view.

**Figure 102 Cisco HyperFlex Connect – System Information**

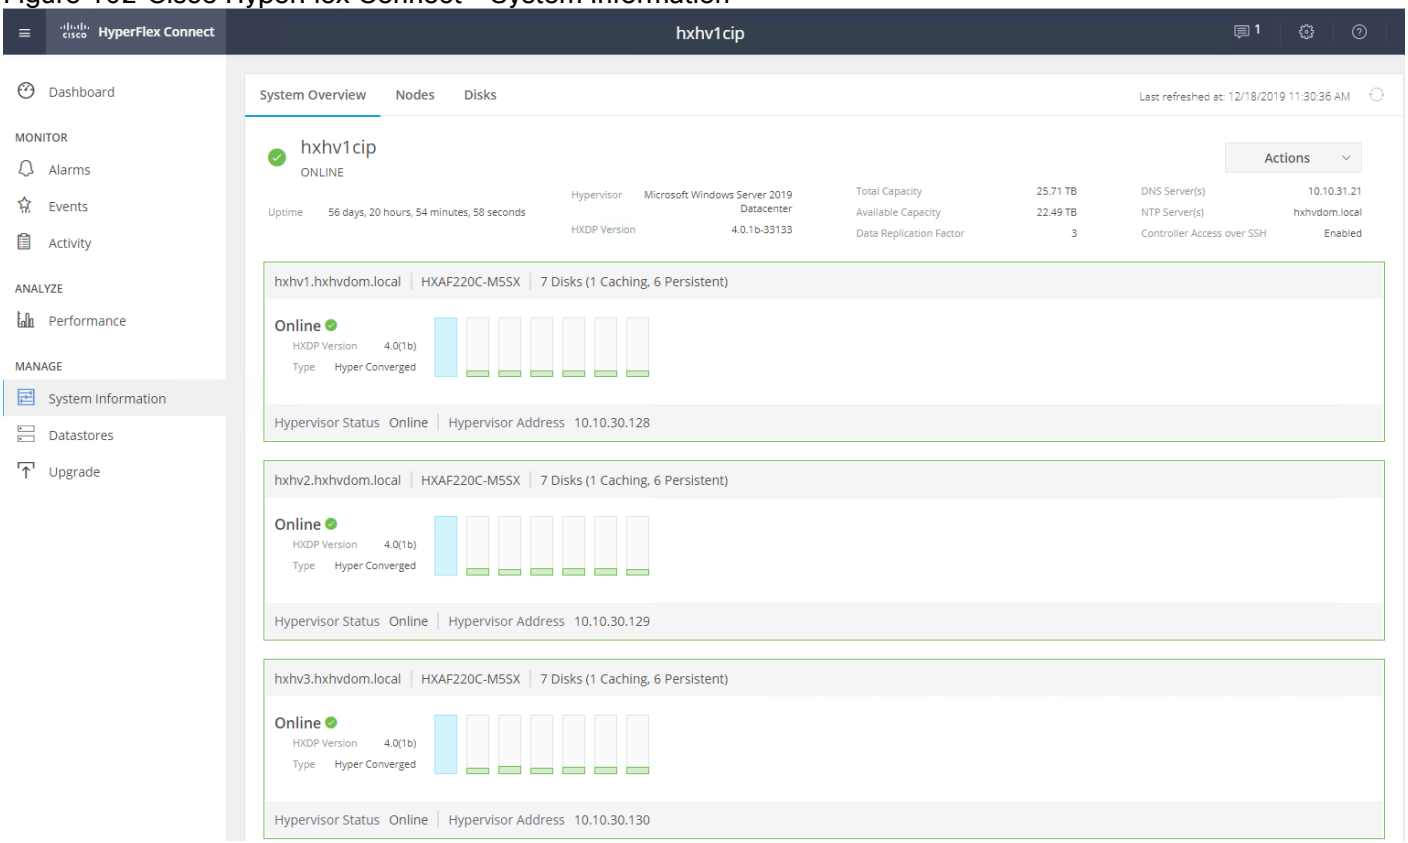


# Configure Remote Management Stations for Managing VMs in HX Cluster

By default, only the Hyper-V nodes in the HX Cluster have access to the SMB share exposed by the HX datastore. Hence, these nodes require no additional configuration for deploying and managing virtual machines in HX cluster.

For deploying and managing virtual machines in HX cluster from a remote management station (running Windows Remote Server Administration Tools - RSAT, System Center VMM, Citrix PVS, Citrix CVAD, etc.) requires opening of port 445 to provide access to the SMB share exposed by HX Datastore.

To allow access to SMB share on remote management stations, follow these steps:

1. SSH into one of the HX Storage Controller VMs and find out the ensemble members in the cluster by executing this command:

   ```
   cat /etc/springpath/storfs.cfg | grep crmZKEnsemble
   ```

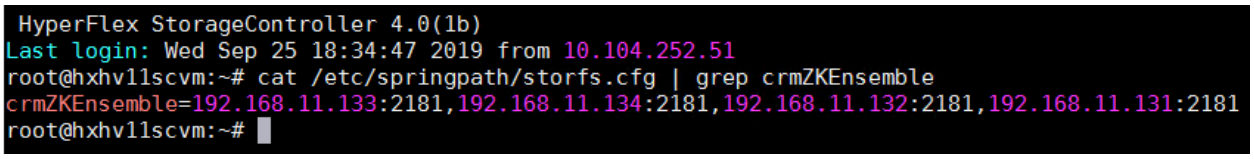


2. To allow access to the SMB share on remote management station, execute the below command without any parameters as shown in the below screenshot

   ```
   python /opt/springpath/storfs-hyperv/FixScvmmAccess.py
   ```

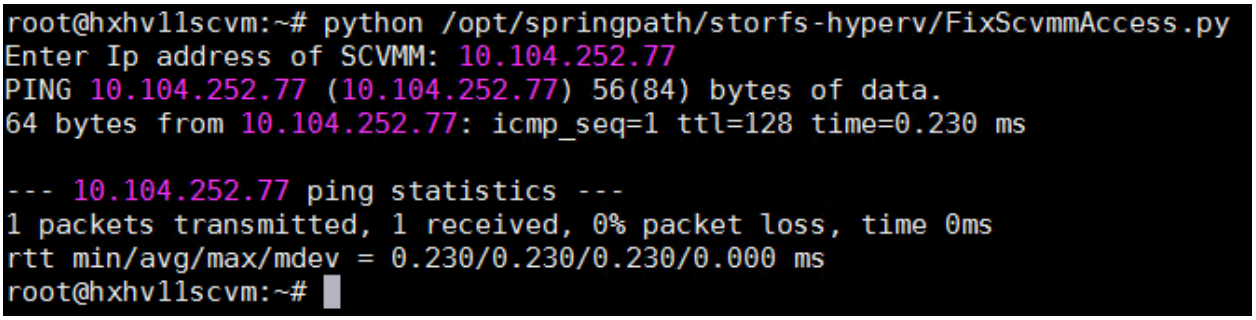


3. From the remote management station execute the below PowerShell script to verify access to the SMB share:

   ```
   Test-Path \\FQDN HX SMB Access Point>
   ```

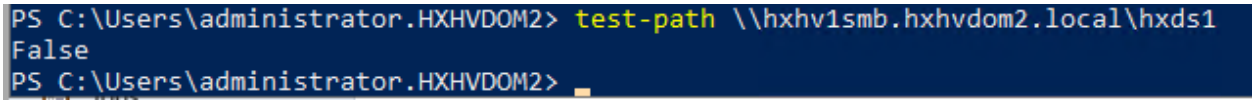


4. If the output of the above command is “false”, edit/modify the hosts file on the remote management station by adding an entry to map the FQDN of HX SMB Access point with its IP address for name resolution as shown in below:

```
PS C:\Users\administrator.HXHVDOM2> Get-Content C:\Windows\System32\drivers\etc\hosts
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.97     rhino.acme.com          # source server
#       38.25.63.10     x.acme.com              # x client host

# localhost name resolution is handled within DNS itself.
#       127.0.0.1       localhost
#       ::1             localhost

10.104.252.135          hxhv1smb.hxhvdom2.local
PS C:\Users\administrator.HXHVDOM2>
```

5. Verify the above configuration using the “test-path” PowerShell command and/or by accessing the SMB share using file explorer as shown in the following screenshots.

```
PS C:\Users\administrator.HXHVDOM2> test-path \\hxhv1smb.hxhvdom2.local\hxds1
True
PS C:\Users\administrator.HXHVDOM2>
```

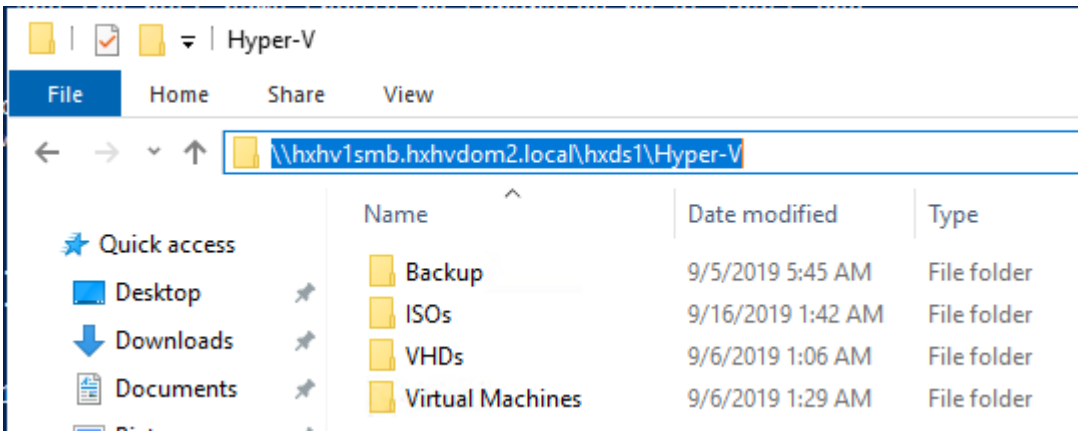


## Microsoft Hyper-V Manager

Hyper-V hosts in a Cisco HyperFlex Systems can be managed both remotely and locally using the Hyper-V Manager. It's installed when you install the Hyper-V Management Tools, which you can do either through a full Hyper-V installation or a tools-only installation on a remote Windows 10 or 2016 Server.

To manage HyperFlex Hyper-V nodes from a remote management host, follow these steps:

1. Install the RSAT tools for Hyper-V using the following PowerShell command:

   ```
   Install-WindowsFeature rsat-hyper-v-tools
   ```

2. Configure the remote management host with constrained delegation as described in the above “Post Installation Task > Constrained Delegation” section.

3. If the remote management host with RSAT tools is outside the HX cluster AD domain, pointing to the same DNS server and host file mapping between HyperFlex SMB namespace and Cluster management IP (CIP) may be required for successful name resolution.

4. For example, add the following entry to the hosts file on the (remote) machine running Hyper-V manager/Failover Cluster Manager/SCVMM Console. Host file in windows is located here - C:\Windows\System32\drivers\etc\hosts

   ```
   cluster_mgmt_ip      \\smb_namespace_name\datastore_name

         10.29.149.235        \\hxcluster.company.com\ds1
   ```

Basic management functions like changing the default store location for virtual machine files and creating virtual machines are described in the following sections. For more information about managing Hyper-V using Hyper-V manager refer to the Microsoft website.

## Change the Default Location to Store the Virtual Machine Files using Hyper-V Manager

To change the default location, follow these steps:

1. Open the Server Manager dashboard and click Tools. Click Hyper-V Manager. The Hyper-V Manager console appears.

**Figure 103 Hyper-V Manager – Connect to Server**

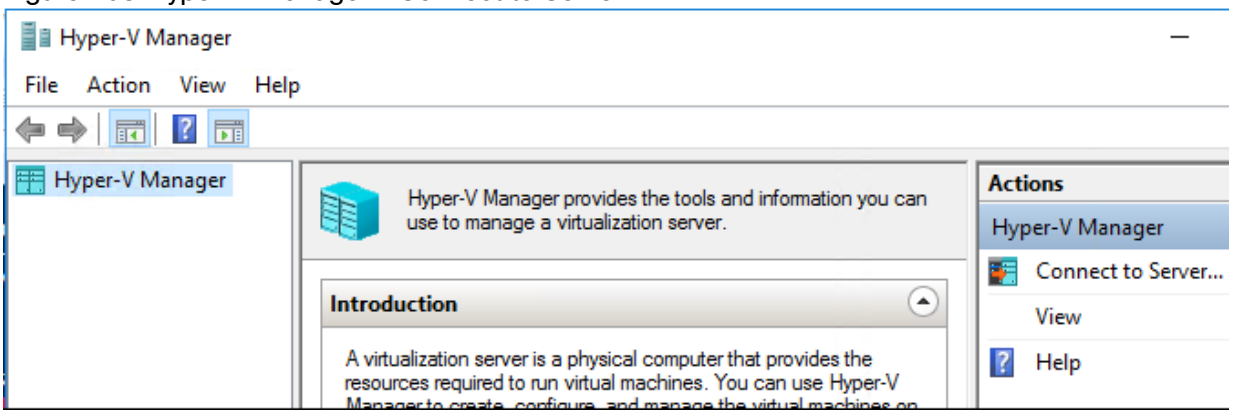


2. In the left pane, select Hyper-V Manager and click Connect to Server....

Figure 104 Hyper-V Manager – Select Computer

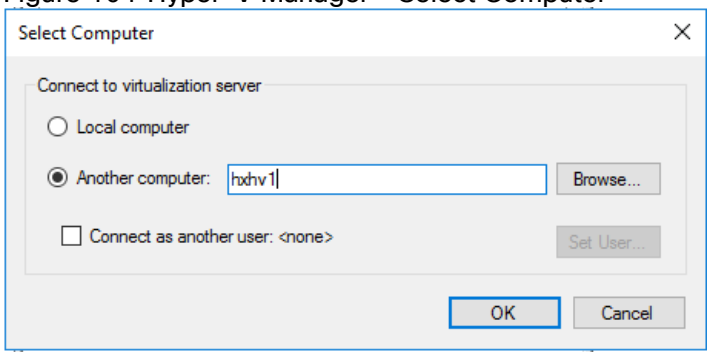


3. In the Select Computer dialog box, select Another computer and type in the name of the Hyper-V node (for example, HXHV1) that belongs to the Hyper-V cluster. Click OK.

4. Repeat steps 1-3 for each Hyper-V node in the HyperFlex cluster.

Figure 105 Hyper-V Manager – View

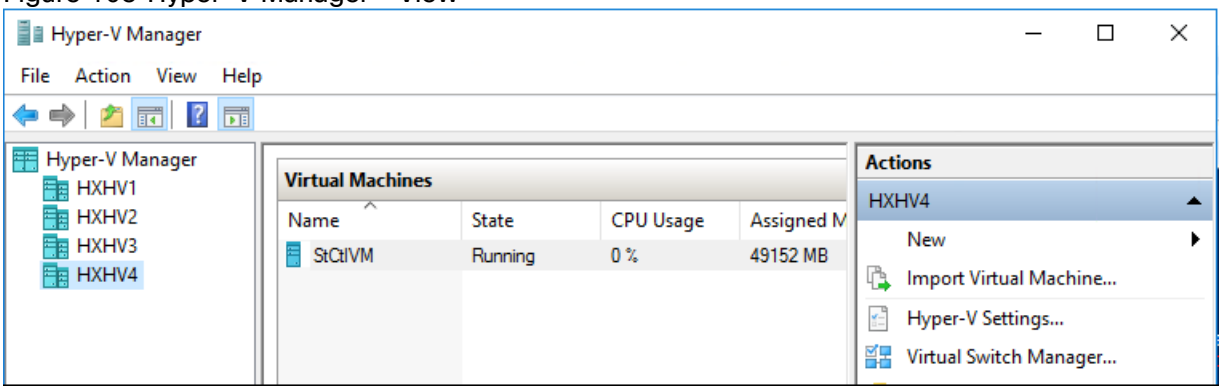


For a fresh installation, the storage controller virtual machine (StCtlVM) in the only virtual machine that appears in Virtual Machines pane in the Hyper-V Manager console. Virtual machines appear in the list under this pane as they are added in each node.

5. Select a Hyper-V server and click the Hyper-V settings and change the default folder location to store the virtual hard disk and virtual machine files as shown below.

Figure 106 Hyper-V Manager – Virtual Machine VHD Location

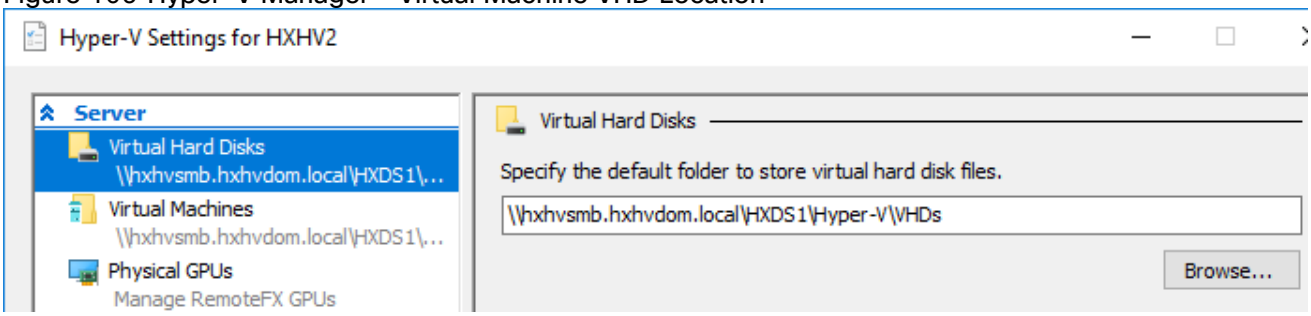

Figure 107 Hyper-V Manager – Virtual Machine Configuration Files Location

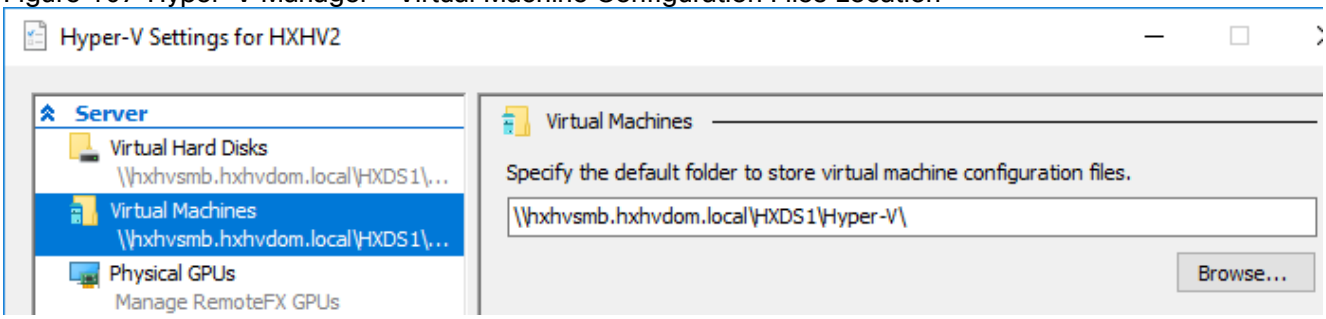


## Create Virtual Machines using Hyper-V Manager

To create Virtual machines using the Hyper-V manager, follow these steps:

1. Open Hyper-V Manager from the Server Manager > Tools.

2. Select a Hyper-V server, and right-click and select New > Create a virtual machine. The Hyper-V Manager New Virtual Machine wizard displays.

Figure 108 Hyper-V Manager – Create New Virtual Machine

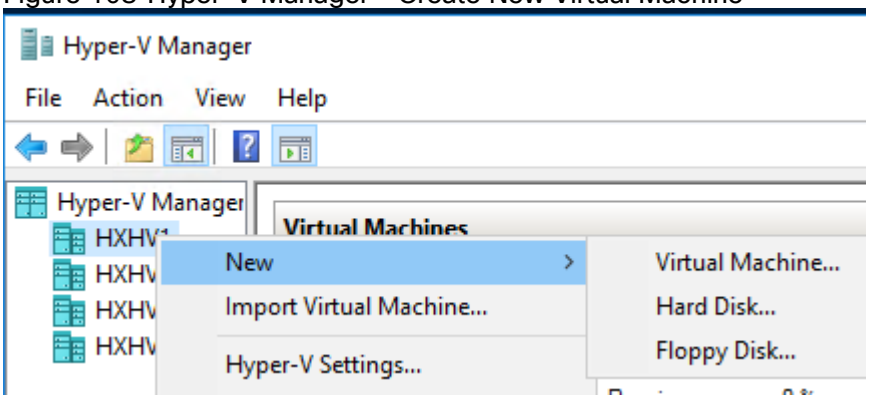


3. In the Before you Begin page, click Next.

4. In the Specify Name and Location page, enter a name for the virtual machine configuration file. The location for the virtual machine click Next.

5. In the Specify Generation page, choose either Generation 1 or Generation 2.

6. In the Assign Memory page, set the start memory value 2048 MB. Click Next.

Figure 109 Hyper-V Manager – Assign New Virtual Machine Memory

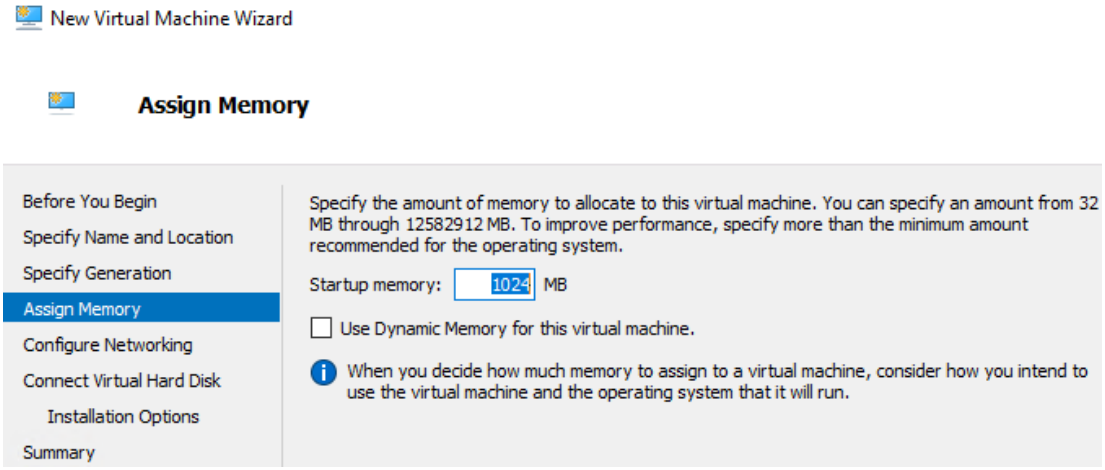


7. In the Configure Networking page, select a network connection for the virtual machine to use from a list of existing virtual switches.

Figure 110 Hyper-V Manager – Configure Networking

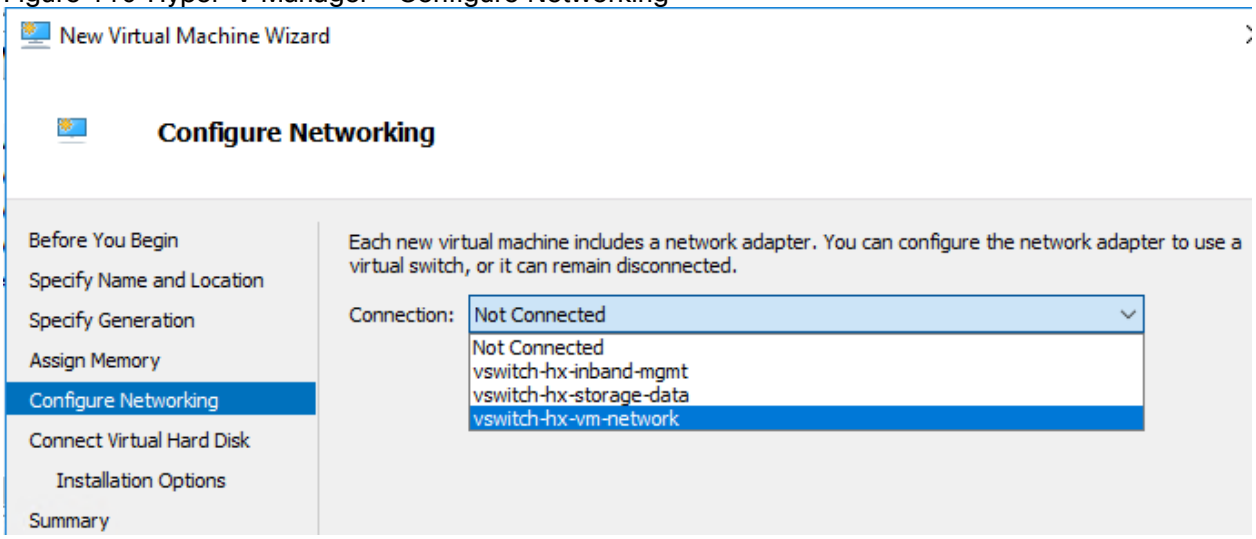


8. In the Connect Virtual Hard Disk page, select Create a Virtual Hard Disk page, and enter the name, location and size for the virtual hard disk. Click Next.

Figure 111 Hyper-V Manager – Connect VHD to New Virtual Machine

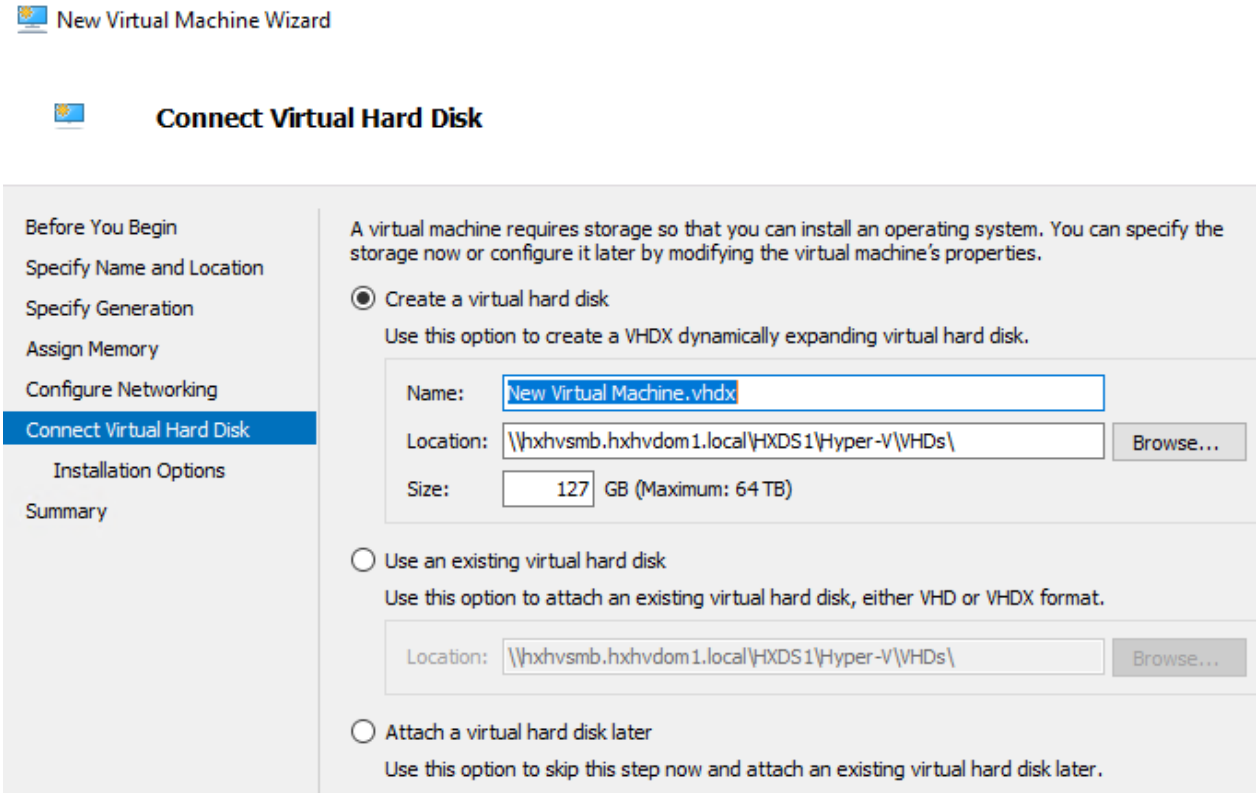


9. In the Installation Options, you can leave the default option Install an operating system later selected. Click Next.

10. In the Summary page, verify that the list of options displayed are correct. Click Finish.

11. In Hyper-V Manager, right-click the virtual machine to perform various operations like Connect, Edit Settings, Start/Stop, and so on.

---

**By default, virtual machines created using Hyper-V Manager are not highly available and they become unavailable if the host server goes down for some reason. However, you can convert these standalone (non-highly available) virtual machines into clustered virtual machine roles using failover cluster manager for high availability.**

---

## Windows Failover Cluster Manager

Cisco HyperFlex installer creates the Hyper-V Failover Cluster, and this can be managed both remotely and locally using the Failover Cluster Manager. It is installed when you install the Failover Clustering Tools, which you can do either through a full Failover Cluster installation or a tools-only installation on a remote Windows 10 or 2016 Server.

---

**By default, virtual machines created using Failover Cluster Manager are highly available and are treated as clustered virtual machine roles. You can also convert a standalone virtual machine into a clustered role using failover cluster manager. By creating clustered virtual machines, you can consolidate multiple servers on one physical server without causing that server to become a single point of failure. Instead, if**

that server, or cluster node, fails or requires scheduled maintenance, then another node begins to run the virtual machines instead through a process known as failover.

To manage the Hyper-V Failover Cluster in HyperFlex System from a remote management host, follow these steps:

1. Install the RSAT tools for Failover Cluster using the below PowerShell command:

   ```
   Install-WindowsFeature RSAT-Clustering-MGMT
   ```

2. Configure the remote management host with constrained delegation as described in section Constrained Delegation (Optional).

3. If the remote management host with RSAT tools is outside the HX cluster AD domain, pointing to the same DNS server and host file mapping between HyperFlex SMB namespace and Cluster management IP (CIP) may be required for successful name resolution.

4. For example, add the following entry to the hosts file on the (remote) machine running Hyper-V manager/Failover Cluster Manager/SCVMM Console. Host file in windows is located here - C:\Windows\System32\drivers\etc\hosts

   ```
   cluster_mgmt_ip      \\smb_namespace_name\datastore_name

   10.29.149.235        \\hxcluster.company.com\ds1
   ```

To create the clustered virtual machine role using the Failover Cluster Manager, follow these steps:

For more information about managing Hyper-V Cluster using Failover Cluster Manager refer to the Microsoft website.

1. In the Failover Cluster Manager console, under the Actions pane, click Connect to Cluster...

Figure 112 Failover Cluster Manager – Connect to Cluster

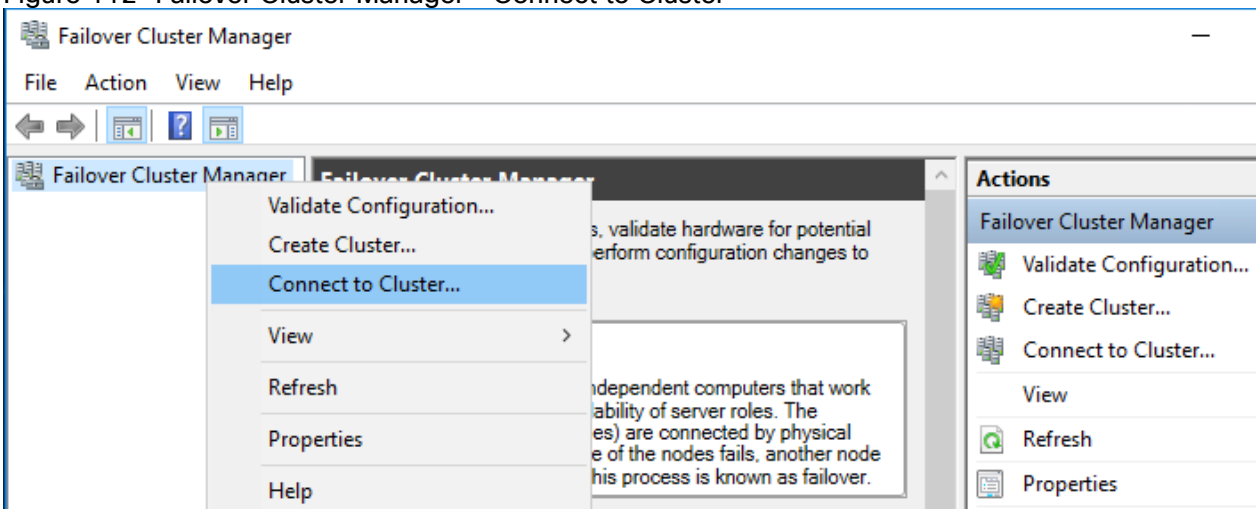


2. In the Select Cluster dialog box, click Browse to navigate to the Hyper-V HX cluster. Click OK.

Figure 113 Failover Cluster Manager – View

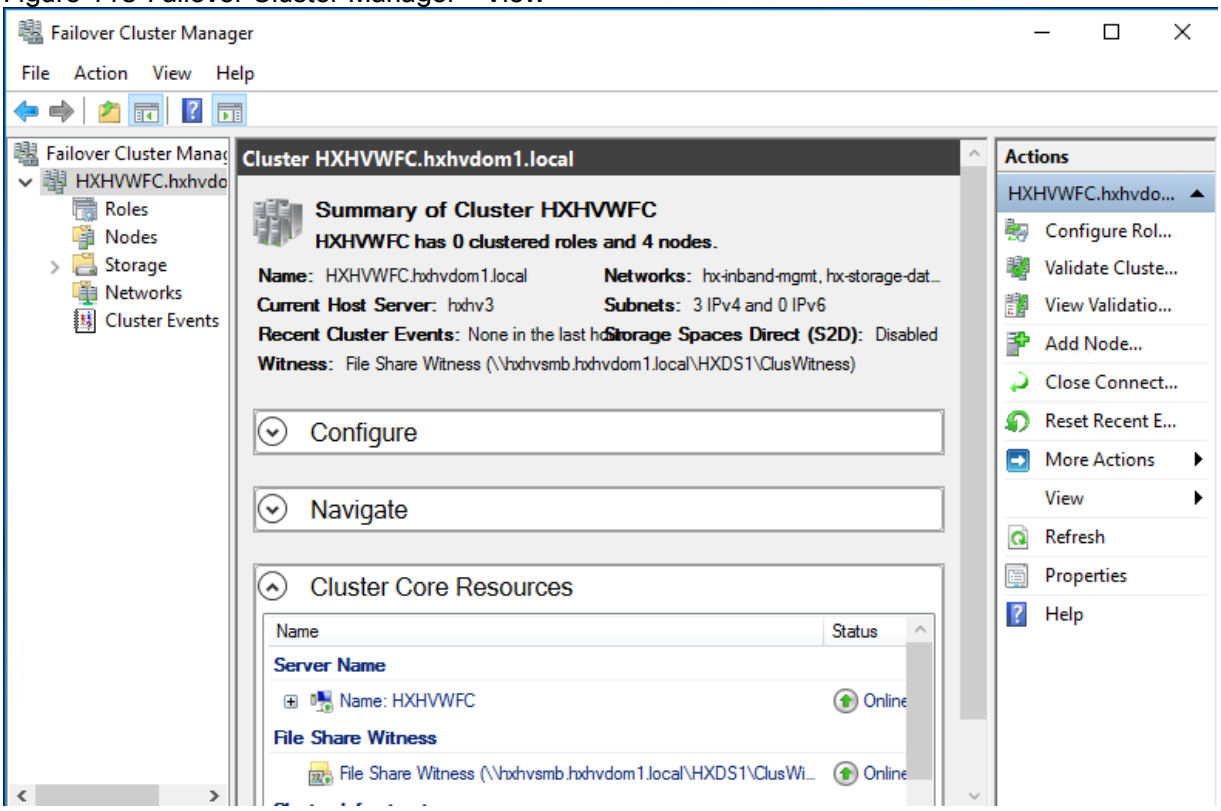


3. In the left pane, click Roles > Virtual Machines... > New Virtual Machines....

Figure 114 Failover Cluster Manager – Create a Role (Virtual Machine)

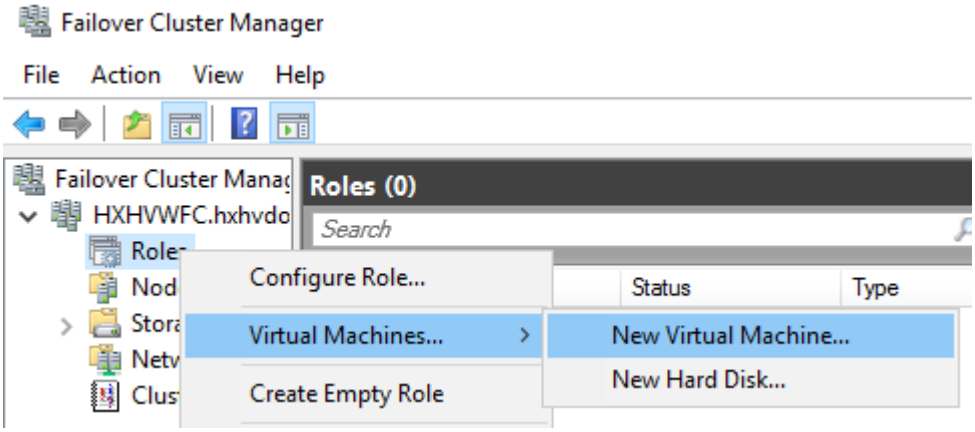


4. In the New Virtual Machine dialog box, search and select the Hyper-V node where you wish to create new virtual machines. Click OK. The New Virtual Machine wizard appears.

Figure 115 Failover Cluster Manager – New Virtual Machine Target Host

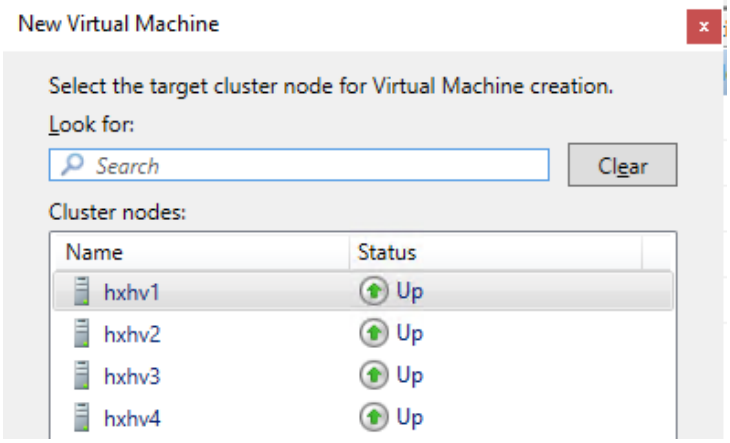


5. In the Before You Begin page, click Next.
6. In the Specify Name and Location page, choose a name for the virtual machine, and specify the location or drive where the virtual machine will be stored. Click Next.
7. In the Specify Generation page, select the generation of virtual machine you want to use (Generation 1 or Generation 2) and click Next.
8. In the Assign Memory page, enter the amount of memory that you want for the virtual machine. Click Next.
9. In the Connect Virtual Hard Disk page, enter the name, location and hard drive size. Click Next.
10. In the Installation Options page, select the install location for the OS. Click Next.
11. In the Summary page, review the options selected and click Finish.

Figure 116 Failover Cluster Manager – New Virtual Machine Summary

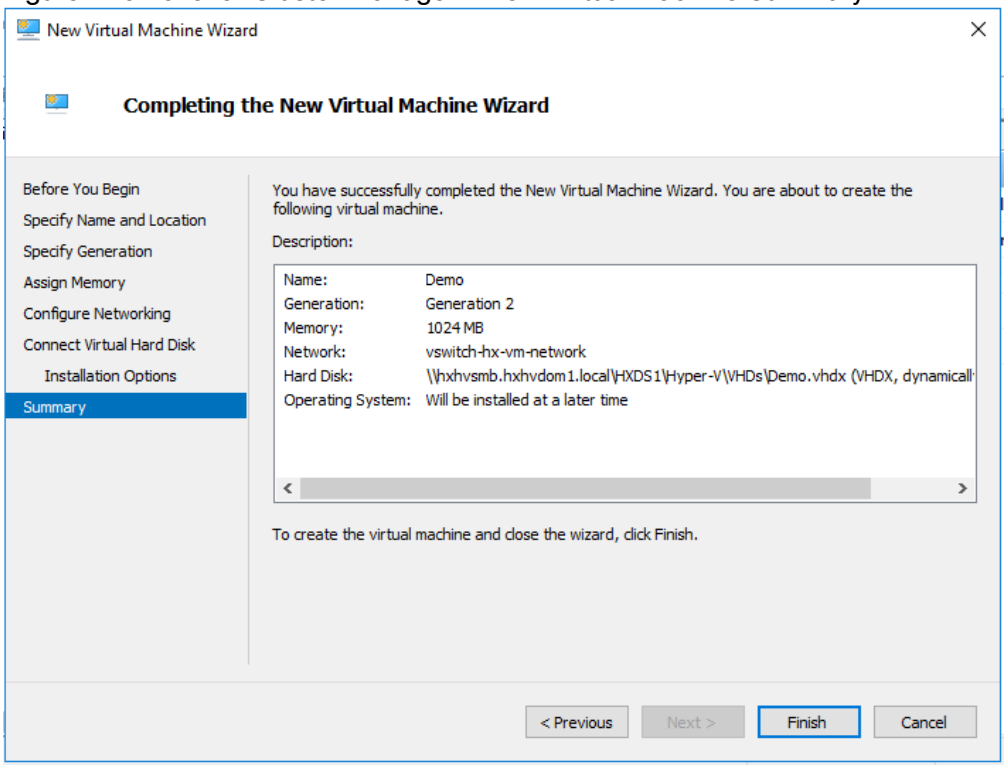

# Post HyperFlex Cluster Installation for Hyper-V 2019

The HyperFlex installer configures all components of the environment from the Cisco UCS policies to the Hyper-V networking. In order to manage the environment from SCVMM for the purpose of VDI, you must perform some post-installation steps to allow Citrix Virtual Desktops to use the environment.

This section explains the steps to prepare the environment for VDI use.

## Microsoft System Center Virtual Machine Manager 2019

The Hyper-V Cluster created by the HyperFlex Installer can also be managed using the Microsoft System Center Virtual Manager. At the time of the publishing this document, there is no HX plug-in or SMI-S integration with the HX Storage. However, the Hyper-V Cluster can still be managed using the SCVMM without these features.

### Create Run-As Account for Managing the Hyper-V Cluster

A Run As account is a container for a set of stored credentials. In VMM a Run As account can be provided for any process that requires credentials. Administrators and Delegated Administrators can create Run As accounts. For this deployment, a Run As account should be created for adding Hyper-V hosts and clusters.

To create a Run As account, follow these steps:

1. Click Settings and in Create click Create Run As Account.
2. In Create Run As Account specify name and optional description to identify the credentials in VMM.
3. In User name and Password specify the credentials. The credentials can be a valid Active Directory user or group account, or local credentials.
4. Clear Validate domain credentials if it is not required and click OK to create the Run As account.

### Manage Servers and Clusters

To add the HyperFlex Hyper-V Cluster to the SCVMM, follow these steps:

1. Open the SCVMM administrator Console and click Fabric > Servers > All Hosts.
2. Right-click All Hosts and Create a folder for the HyperFlex Hyper-V Cluster.
3. Right-click the newly created folder and click Add Hyper-V Hosts and Clusters.

Figure 117 SCVMM – Create a Host Group

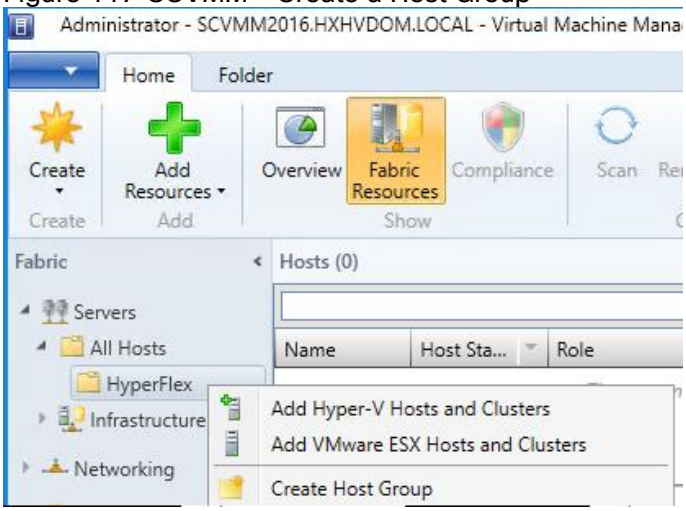


4. In the Credentials section, select Use an existing Run As account and select the account created in the previous section.

5. In the Discovery scope, enter the FQDN of HyperFlex Hyper-V Cluster as shown below.

Figure 118 SCVMM – Discovery Scope

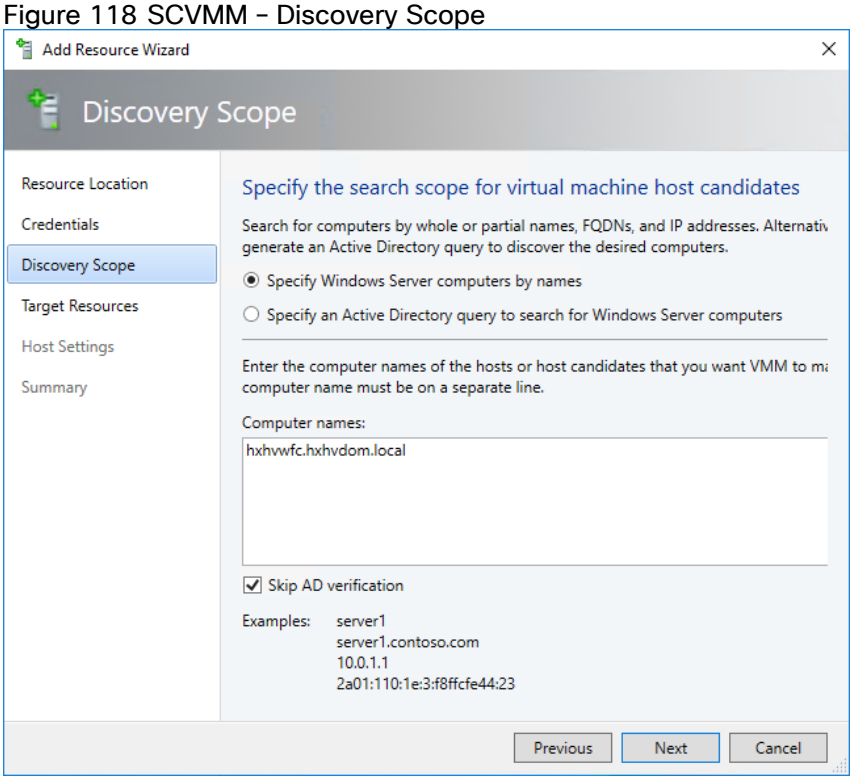


6. In the Target Resources page, select all the discovered hosts and click Next.

Figure 119 SCVMM – Target Resources

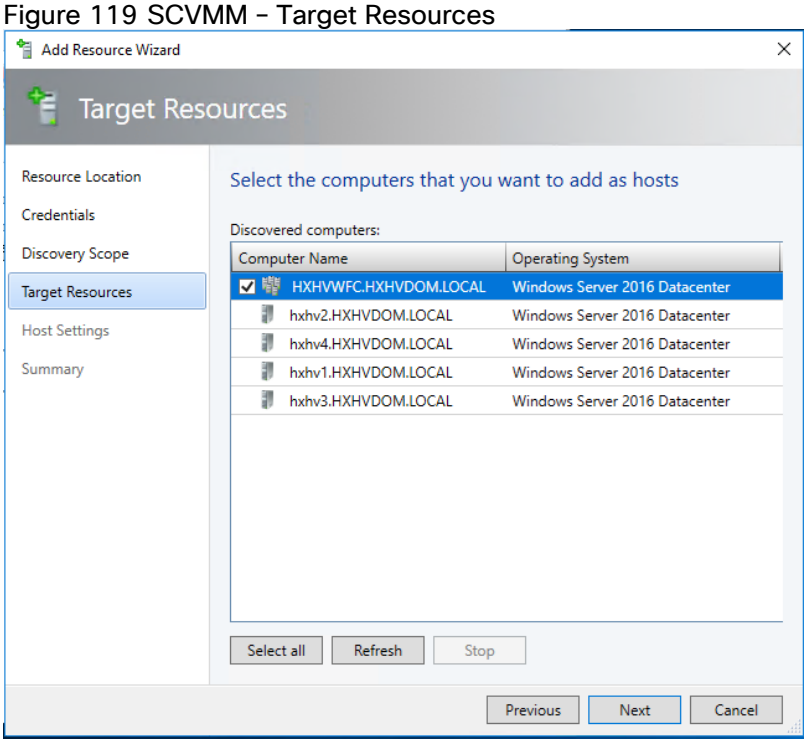


7. In the Host Settings page, select the appropriate Host group and click Next.

Figure 120 SCVMM – Host Settings

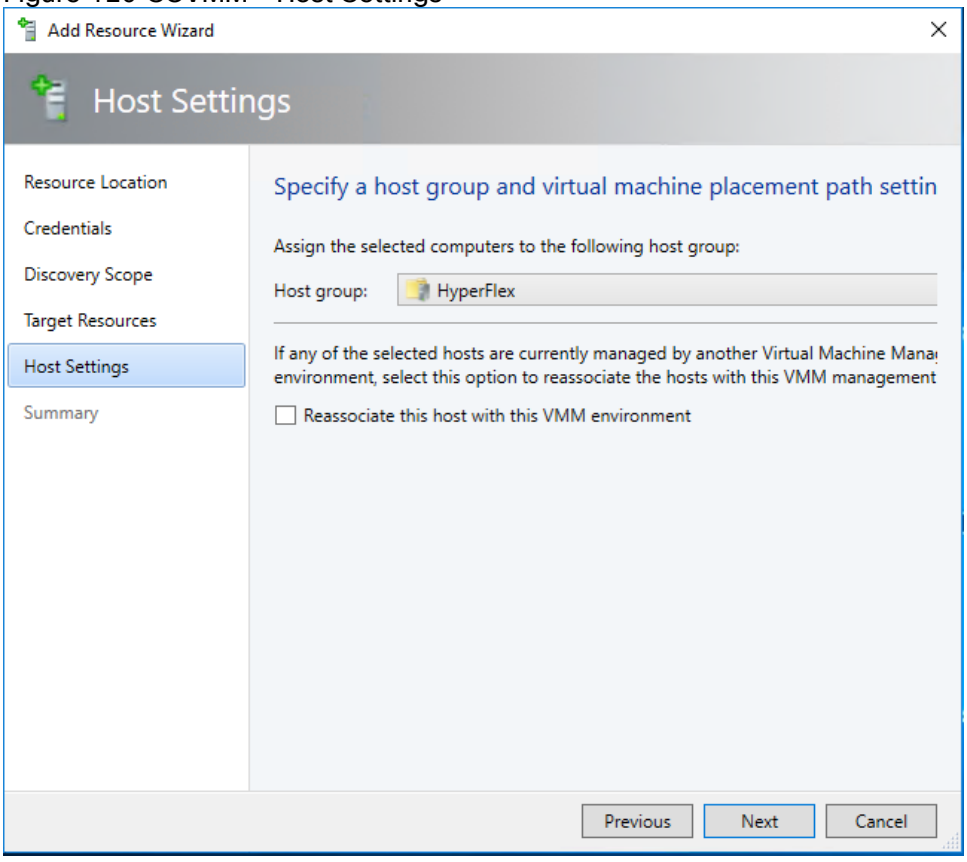

8. In the summary page, confirm the settings and click Finish.

Figure 121 SCVMM – Server View

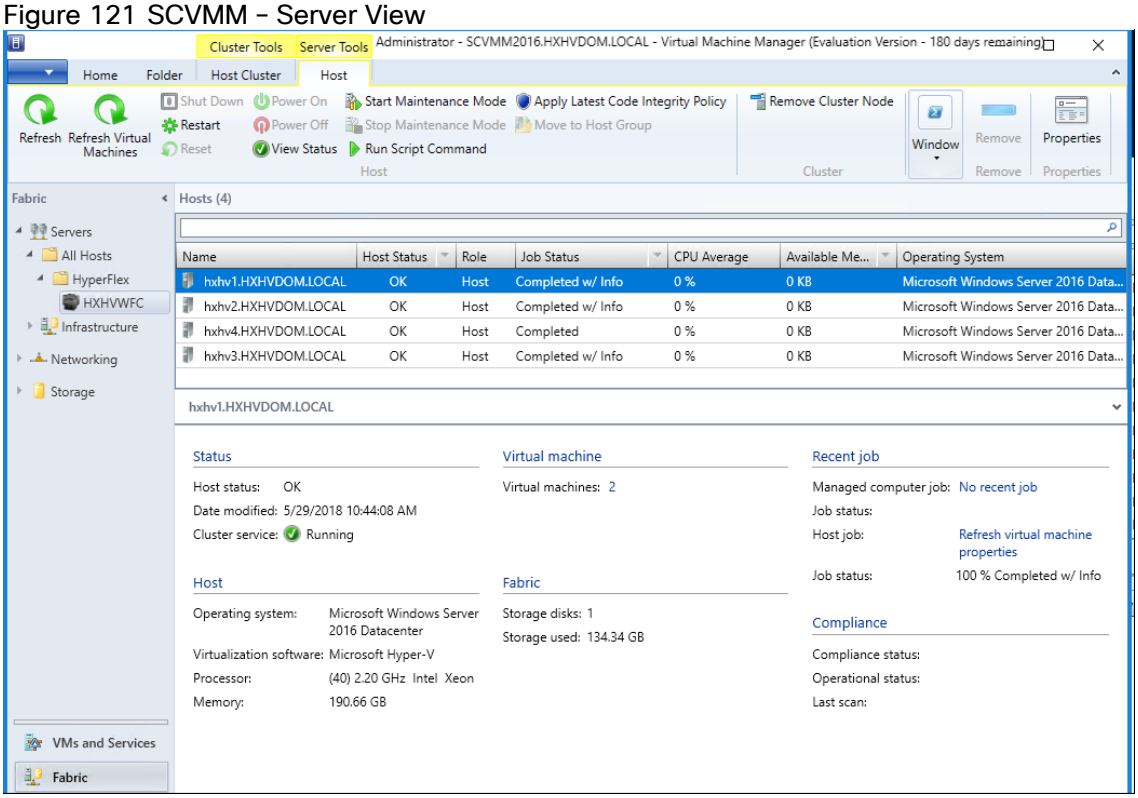


## Networking

Network switches and interfaces are created by the HyperFlex installer. A network team is created for the Management, Storage, Virtual Machine Network and Live Migration networks specified during the installer.

**When the teams are created, the Network Sites must be created and added to the logical networks created by the installer.**

To create the Network Sites, follow these steps:

1. Under Fabric > Networking > Logical Networks, find the Logical Network created by the installer.

**In this example, the Logical Network created by the installer is 'vswitch-hx-vm-network'.**

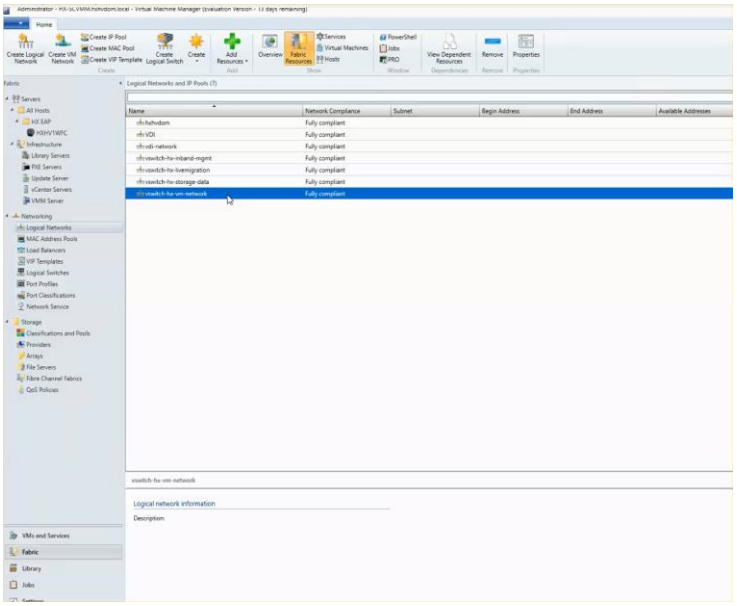

2. Right-click and select 'Properties' of the logical network.

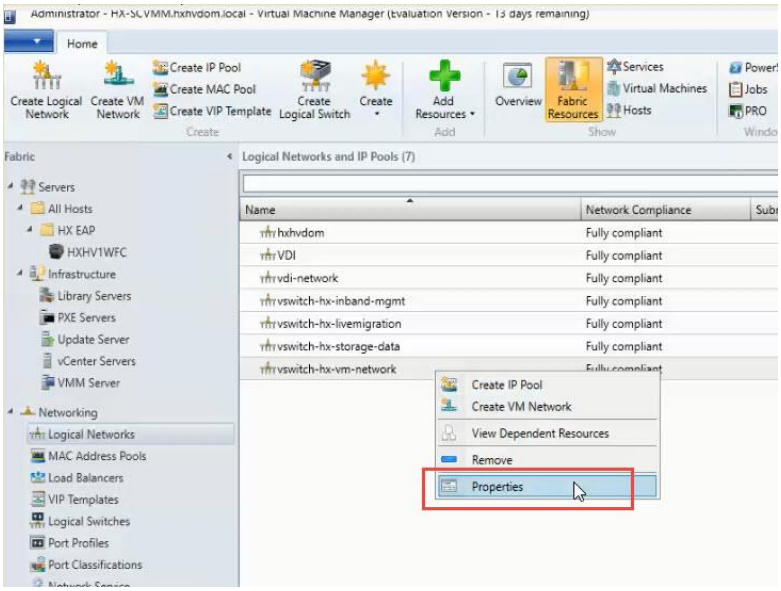


3. Under Network Site, add a site so a VLAN can be specified on the Logical Network.

> **In this example, VLAN 34 was used for the virtual machine network and VLAN 33 was used for Live Migration.**

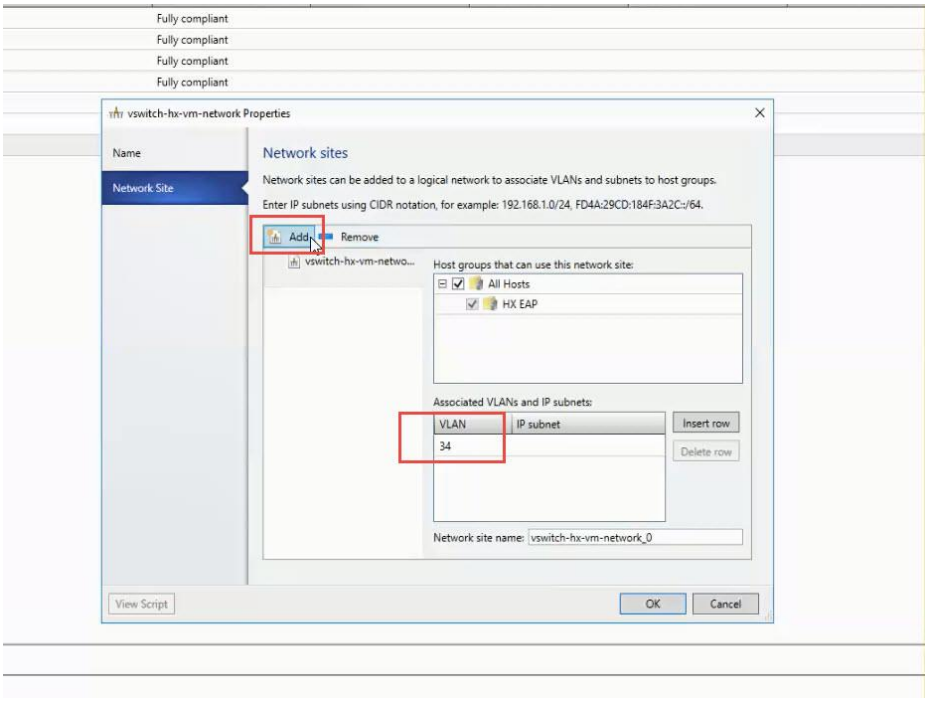


4. When the network site is created, make sure each host in the cluster has the proper VLAN checked in its properties. This can be found under the properties of each host, under Hardware -> and scroll to the 'team-hx-vm-network'.

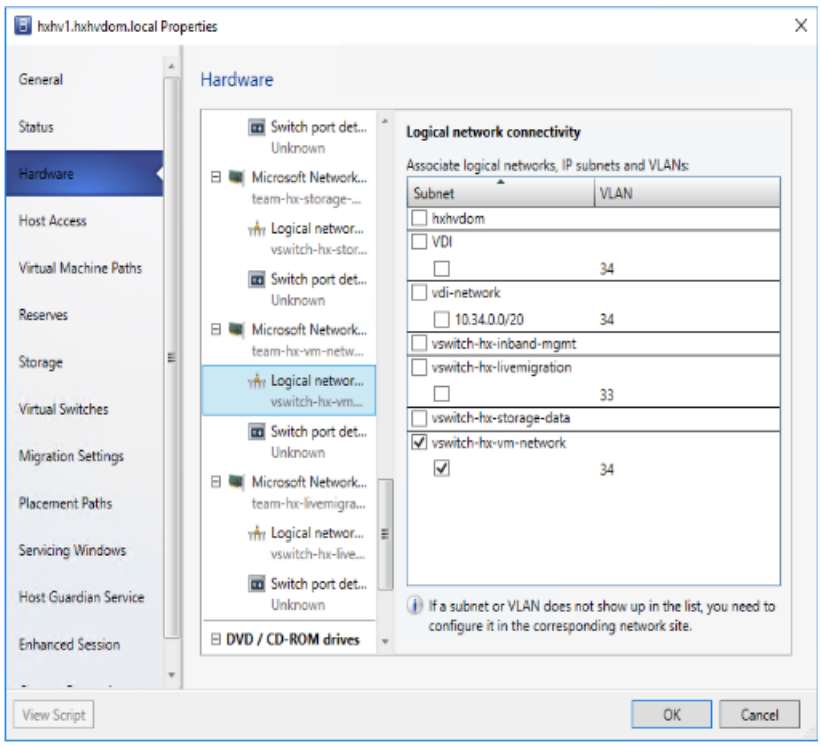

## Storage

The data stores created in the HyperFlex Connect interface, creates an SMB share to use to place Virtual Machines. The naming convention is '\\hxClustername\DatastoreName'.

To add the HX Datastores to the HX cluster, follow these steps:

1. Right-click the Cluster 'HXHV1WFC', select Properties and click 'File Share Storage'.

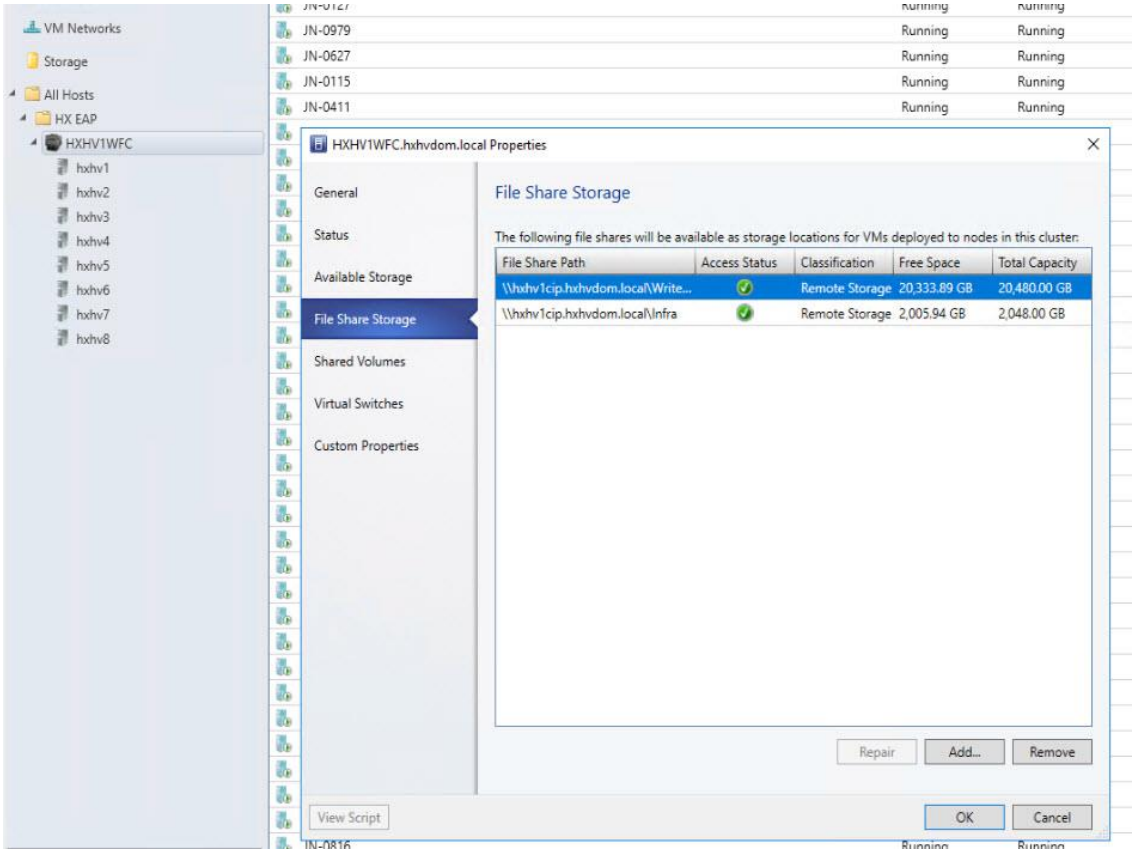


2. Click Add to specify the UNC path for the datastore.

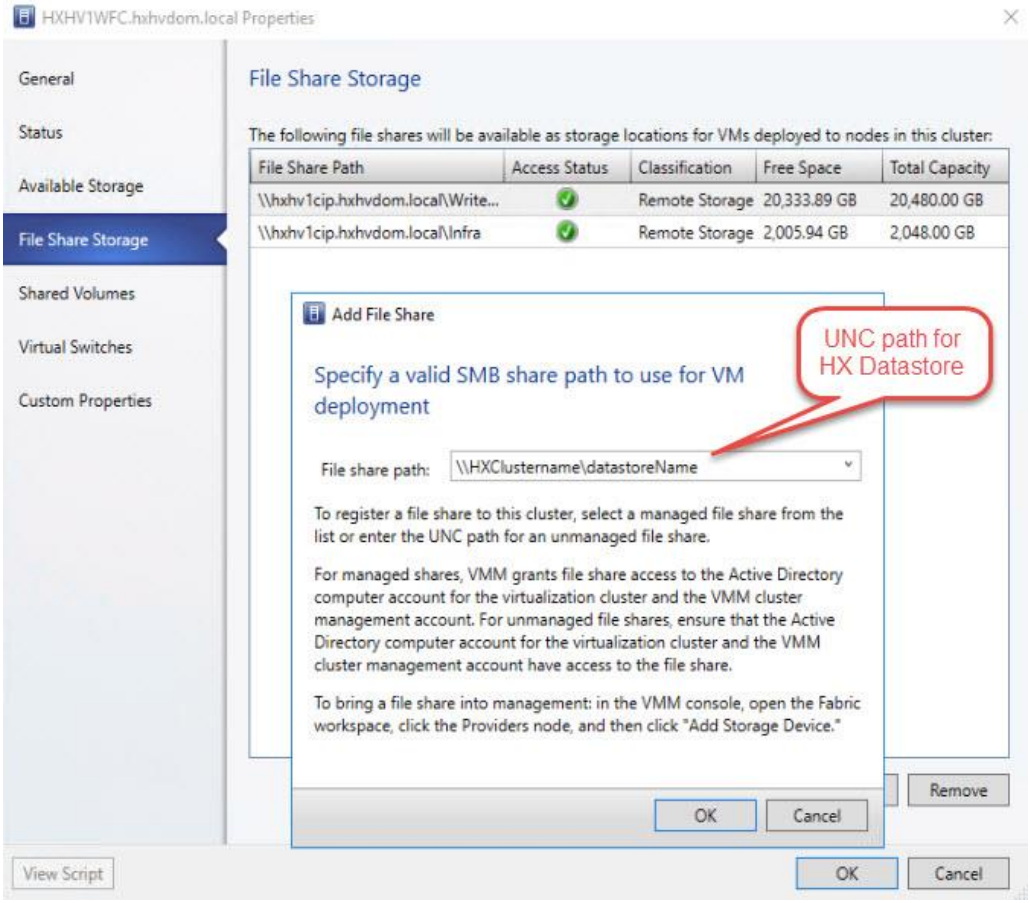
HXHV1WFC.hxhvdom.local Properties
General
Status
Available Storage
File Share Storage
Shared Volumes
Virtual Switches
Custom Properties
File Share Storage
The following file shares will be available as storage locations for VMs deployed to nodes in this cluster:
File Share Path
Access Status
Classification
Free Space
Total Capacity
\\hxhv1cip.hxhvdom.local\Write...
Remote Storage
20,333.89 GB
20,480.00 GB
\\hxhv1cip.hxhvdom.local\Infra
Remote Storage
2,005.94 GB
2,048.00 GB
Add File Share
Specify a valid SMB share path to use for VM deployment
UNC path for HX Datastore
File share path:
\\HXClustername\datastoreName
To register a file share to this cluster, select a managed file share from the list or enter the UNC path for an unmanaged file share.
For managed shares, VMM grants file share access to the Active Directory computer account for the virtualization cluster and the VMM cluster management account. For unmanaged file shares, ensure that the Active Directory computer account for the virtualization cluster and the VMM cluster management account have access to the file share.
To bring a file share into management: in the VMM console, open the Fabric workspace, click the Providers node, and then click "Add Storage Device."
OK
Cancel
Remove
View Script
OK
Cancel

3. Click OK.

# Build the Virtual Machines and Environment for Workload Testing

## Software Infrastructure Configuration

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

Install and configure the infrastructure virtual machines by following the process provided in Table 39

**Table 39 Test Infrastructure Virtual Machine Configuration**

| Configuration | Citrix Virtual Desktops 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 40

**Table 40 HVD and HSD Configurations**

| Configuration | HVDI<br>Virtual Machines | HSD<br>Virtual Machines |
|---|---|---|
| Operating system | Microsoft Windows 10 64-bit | Microsoft Windows Server 2019 |
| Virtual CPU amount | 2 | 8 |
| Memory amount | 4.0 GB (reserved) | 24 GB (reserved) |
| Network | VMNIC<br><br>vm-network | VMNIC<br><br>vm-network |
| Citrix PVS vDisk size and location | 24 GB | 40 GB |
| Citrix PVS write cache<br><br>Disk size | 6 GB | 24 GB |
| Additional software used for testing | Microsoft Office 2016<br><br>Login VSI 4.1.32 (Knowledge Worker Workload) | Microsoft Office 2016<br><br>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 1909 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, follow these steps:

1. To begin the installation, connect to the first Citrix License server and launch the installer from the Citrix Virtual Apps and Desktops 1909 ISO.

2. Click Start.

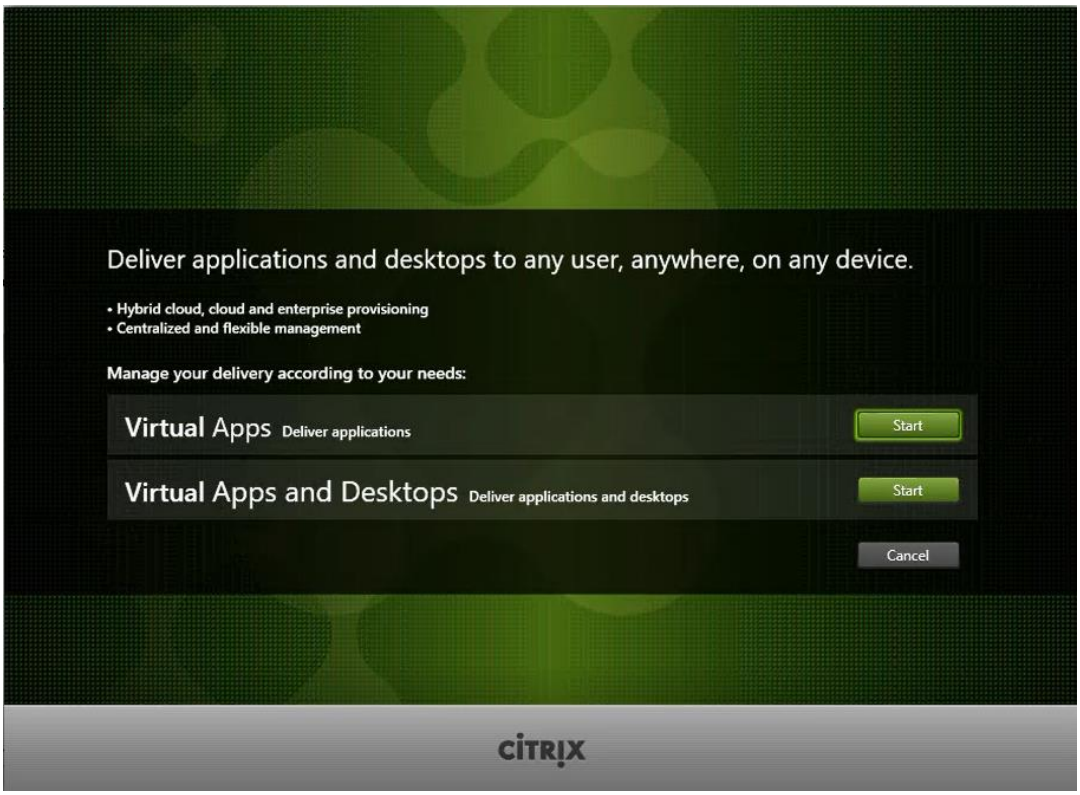


3. Click “Extend Deployment – Citrix License Server.”

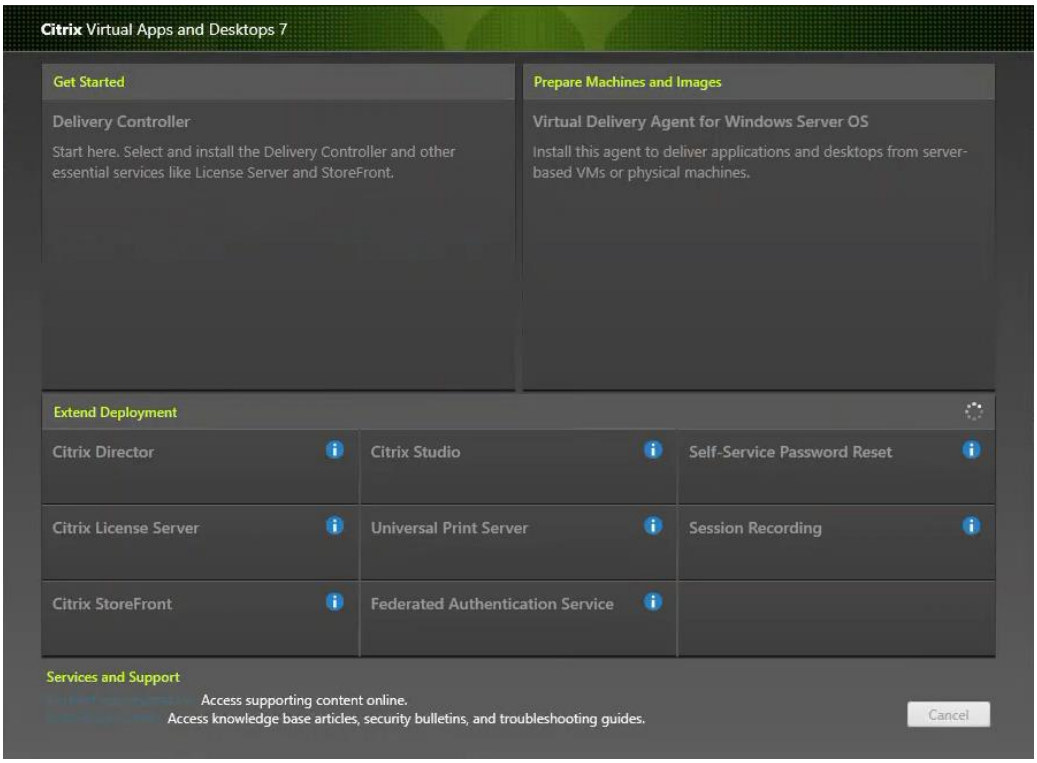


4. Read the Citrix License Agreement.

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.

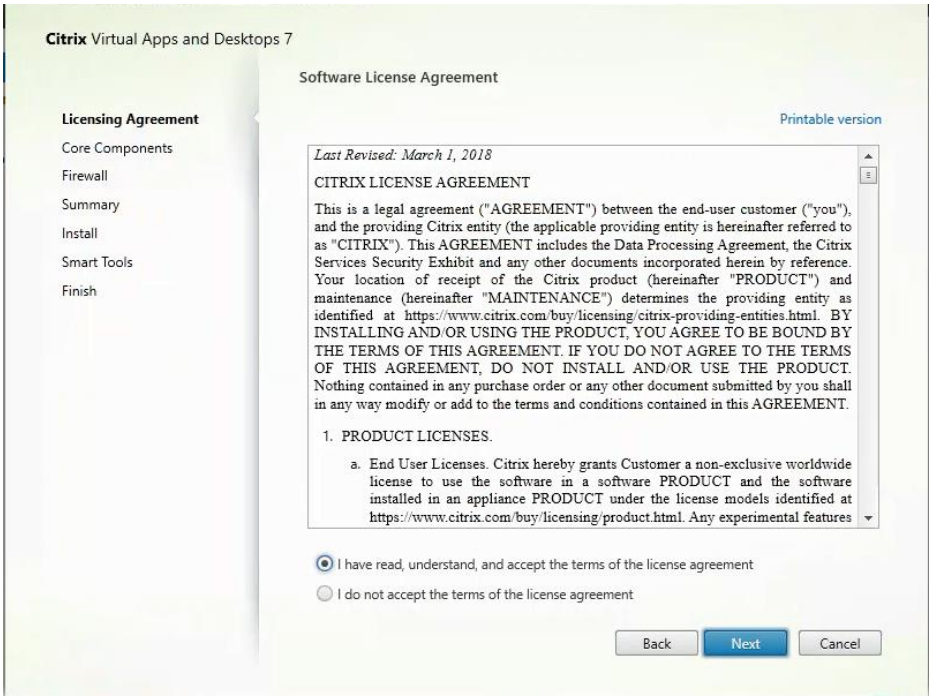

7. Click Next.

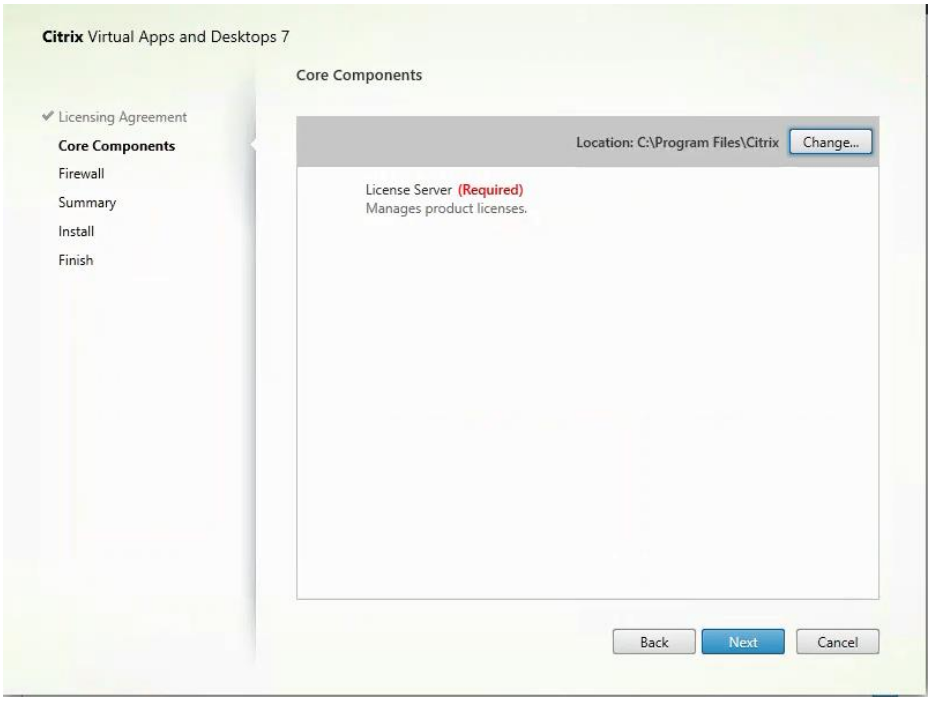


8. Select the default ports and automatically configured firewall rules.

9. Click Next.

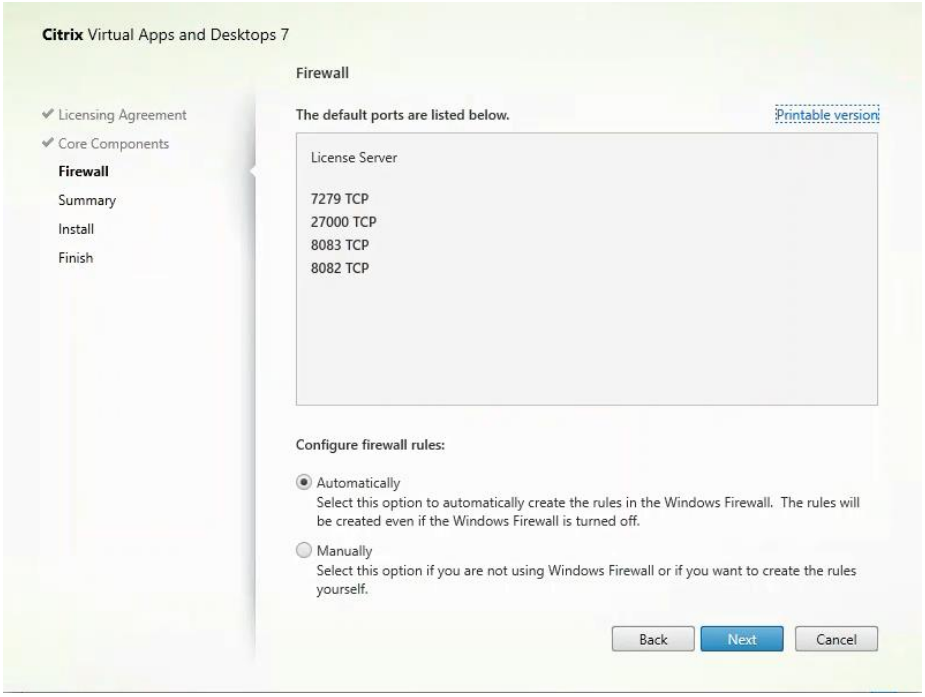


10. Click Install.

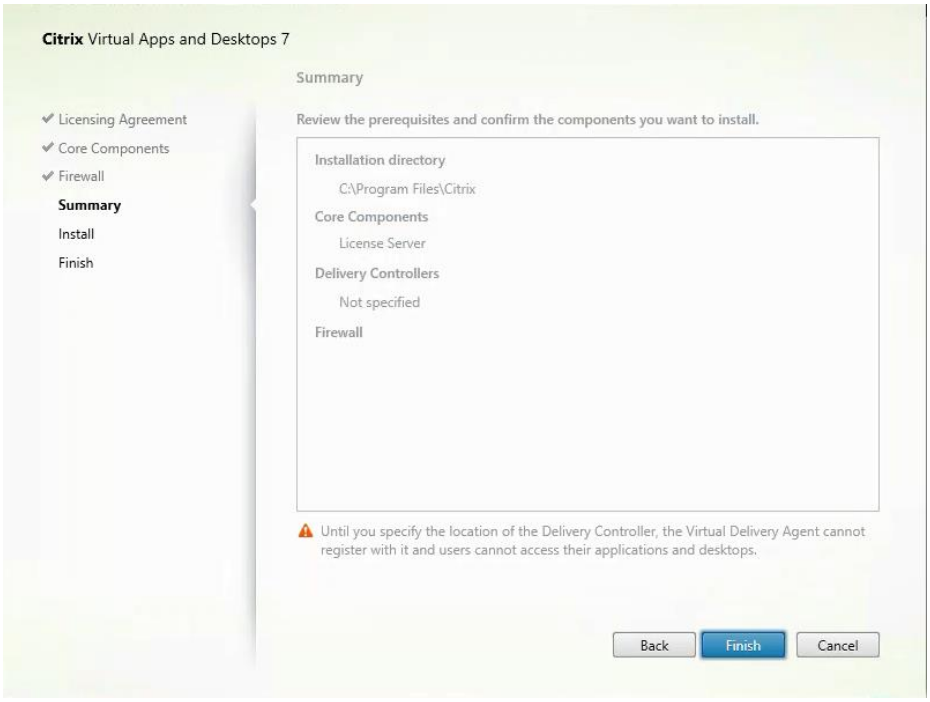


11. Click Finish to complete the installation.

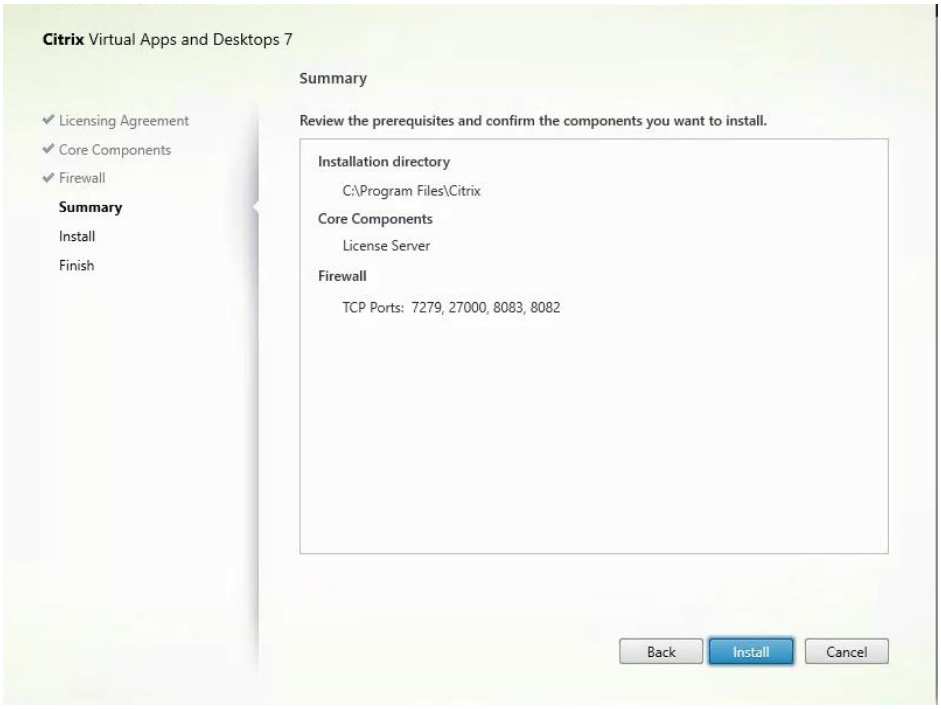


## Install Citrix Licenses

To install the Citrix Licenses, follow these steps:

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

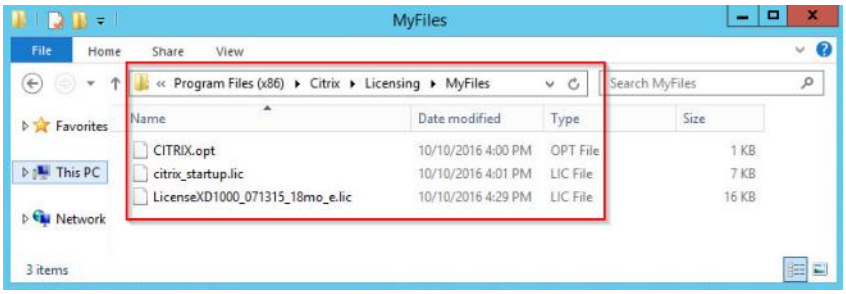


2. Restart the server or Citrix licensing services so that the licenses are activated.

3. Run the application Citrix License Administration Console.

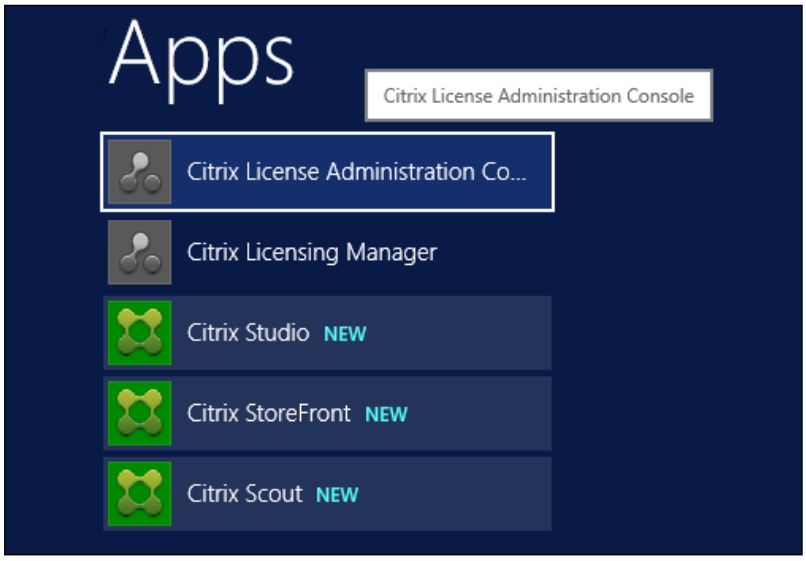


4. Confirm that the license files have been read and enabled correctly.

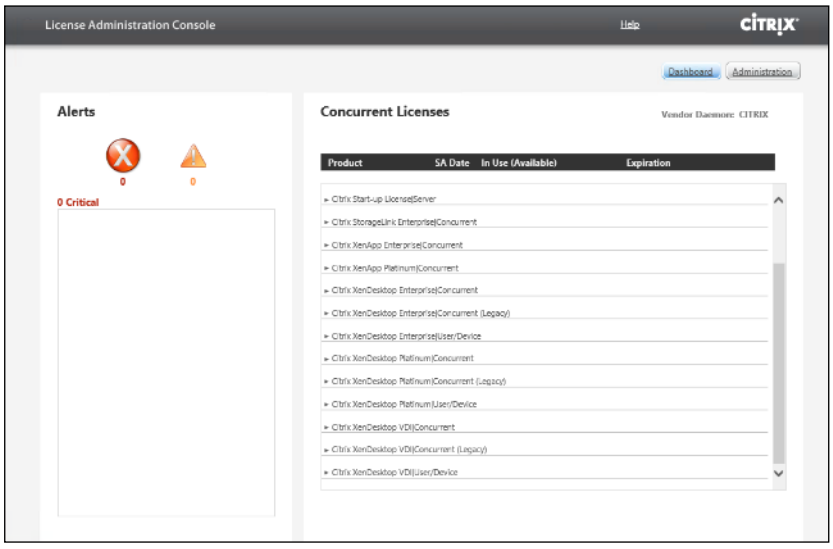

## Install Citrix Desktop Broker/Studio

To install Citrix Desktop, follow these steps:

1. Connect to the first Citrix VDI server and launch the installer from the Citrix Desktop 1909 ISO.
2. Click Start.

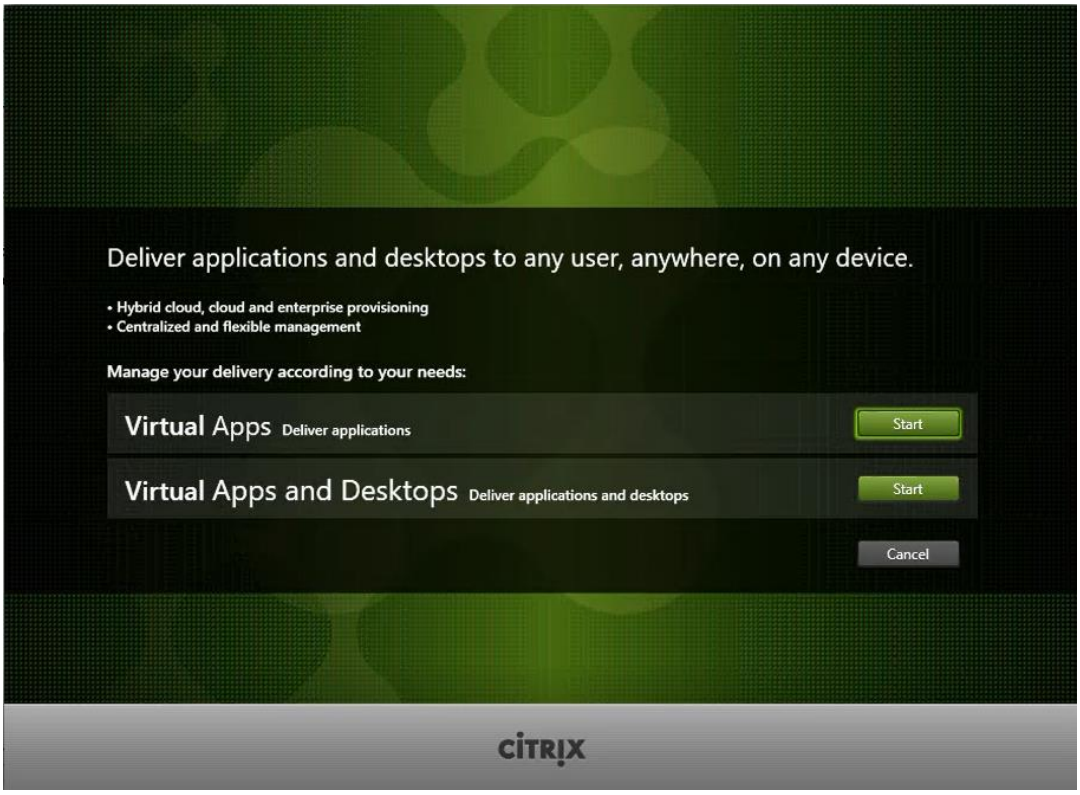


The installation wizard presents a menu with three subsections.

3. Click “Get Started – Delivery Controller.”

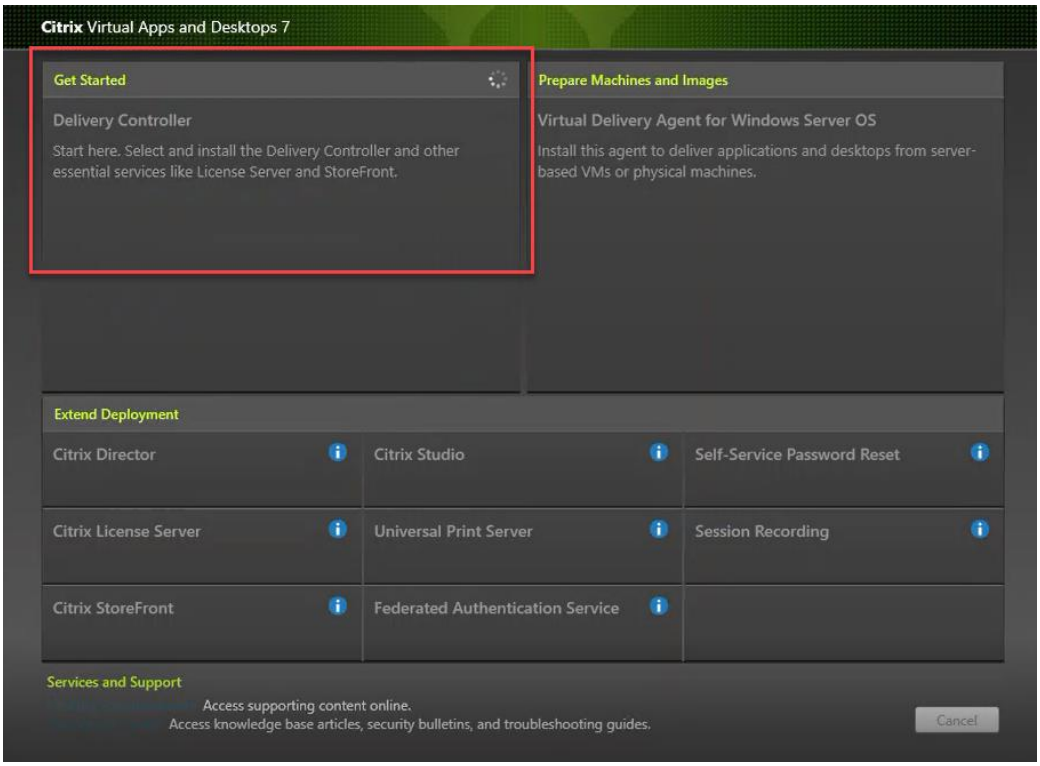


4. Read the Citrix License Agreement and if acceptable, indicate your acceptance of the license by selecting the "I have read, understand, and accept the terms of the license agreement" radio button.

5. Click Next.

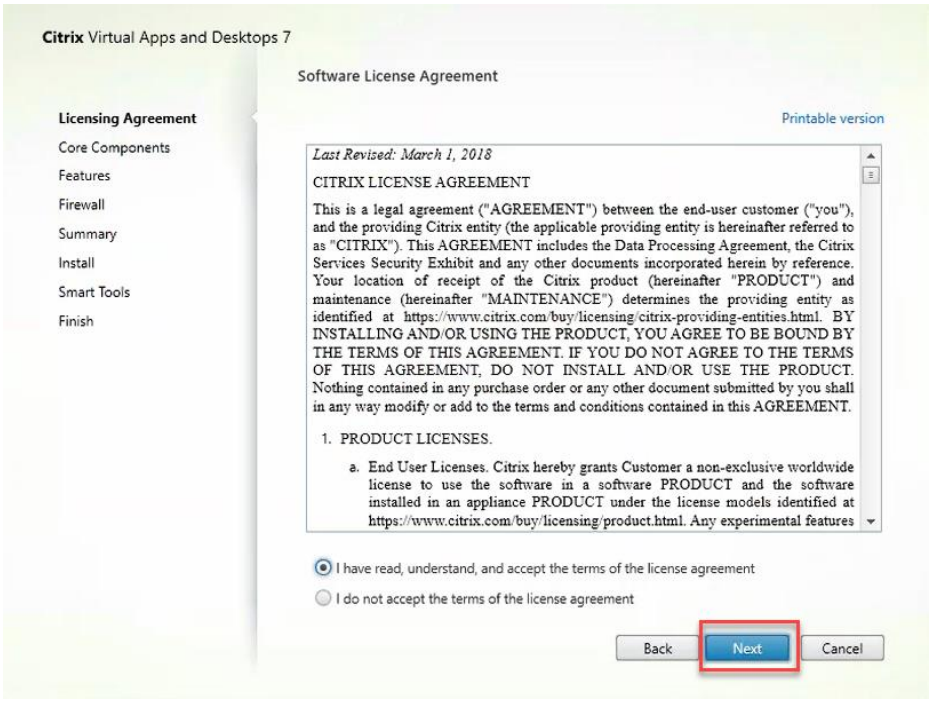


6. Select the components to be installed on the first Delivery Controller Server:

a. Delivery Controller
b. Studio
c. Director

7. Click Next.

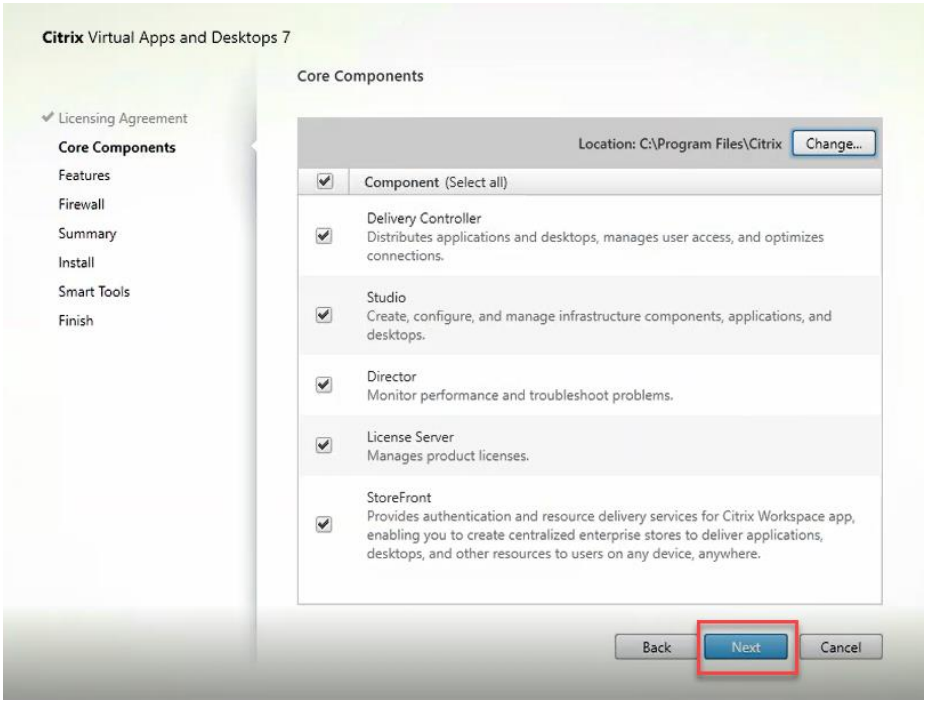

**Dedicated StoreFront and License servers should be implemented for large-scale deployments.**

8. Since a SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2012 SP1 Express” unchecked.

9. Click Next.

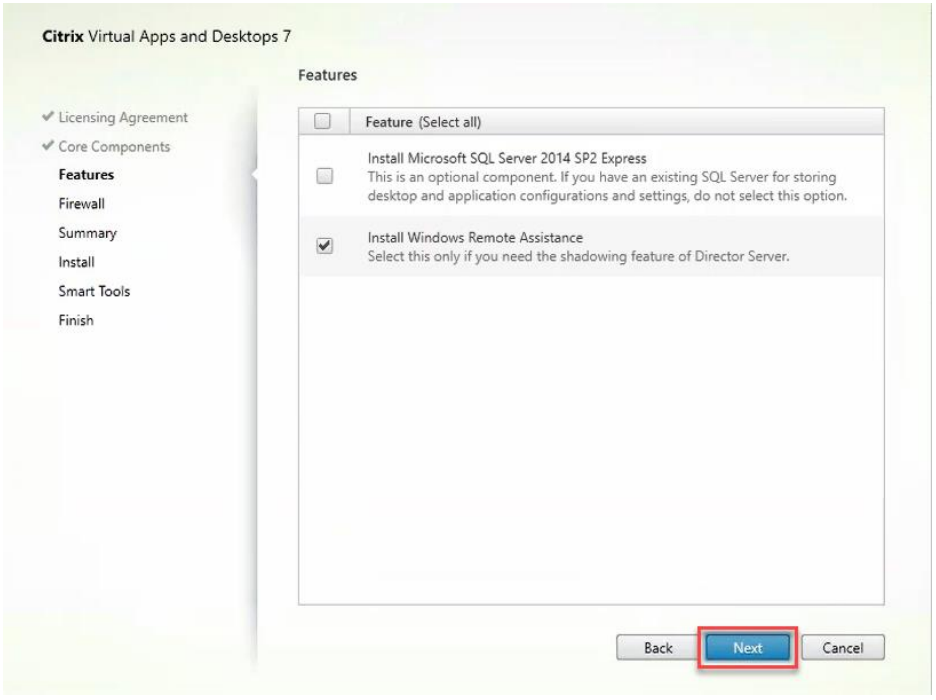


10. Select the default ports and automatically configured firewall rules.

11. Click Next.

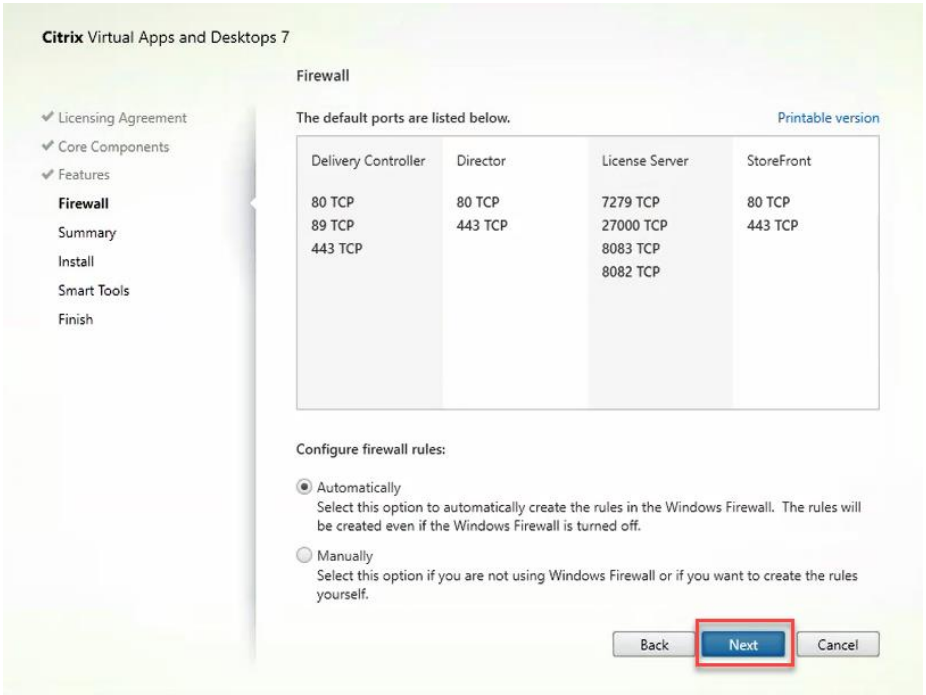


12. Click Install.

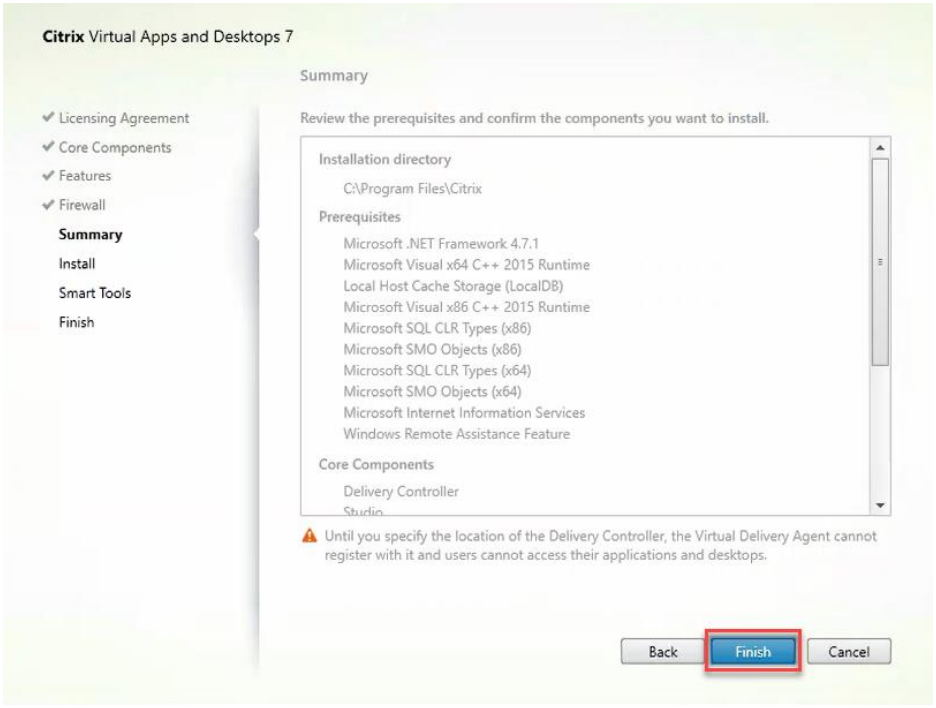


13. (Optional) Click the Call Home participation.

14. Click Next.

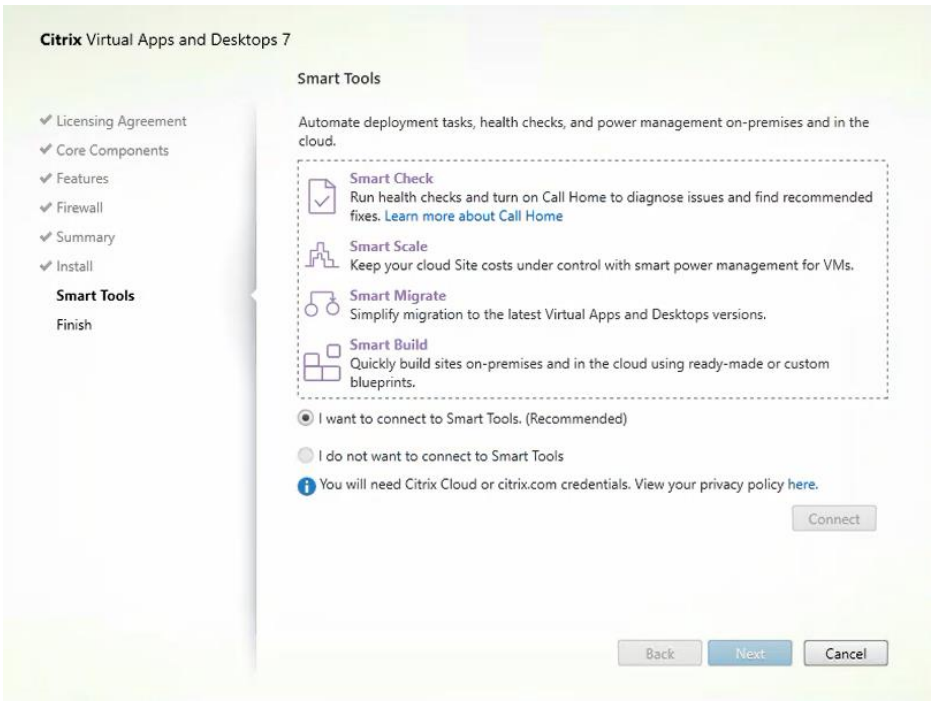


15. Click Finish to complete the installation.

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

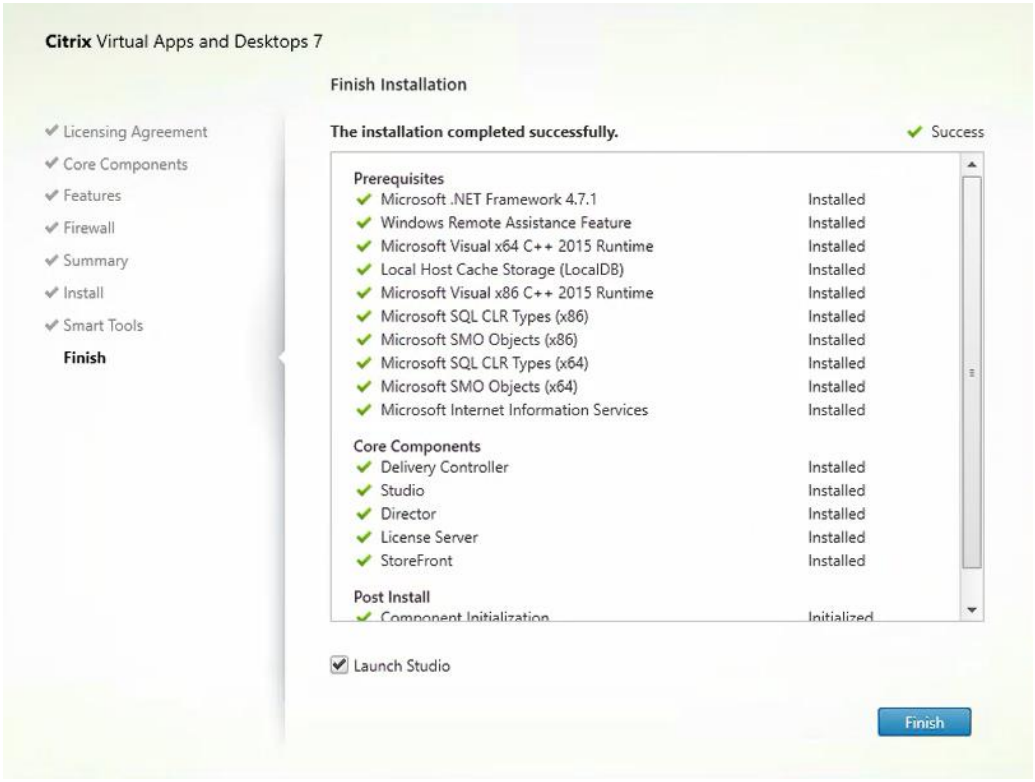


## Configure the Citrix VDI Site

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, follow these steps:

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

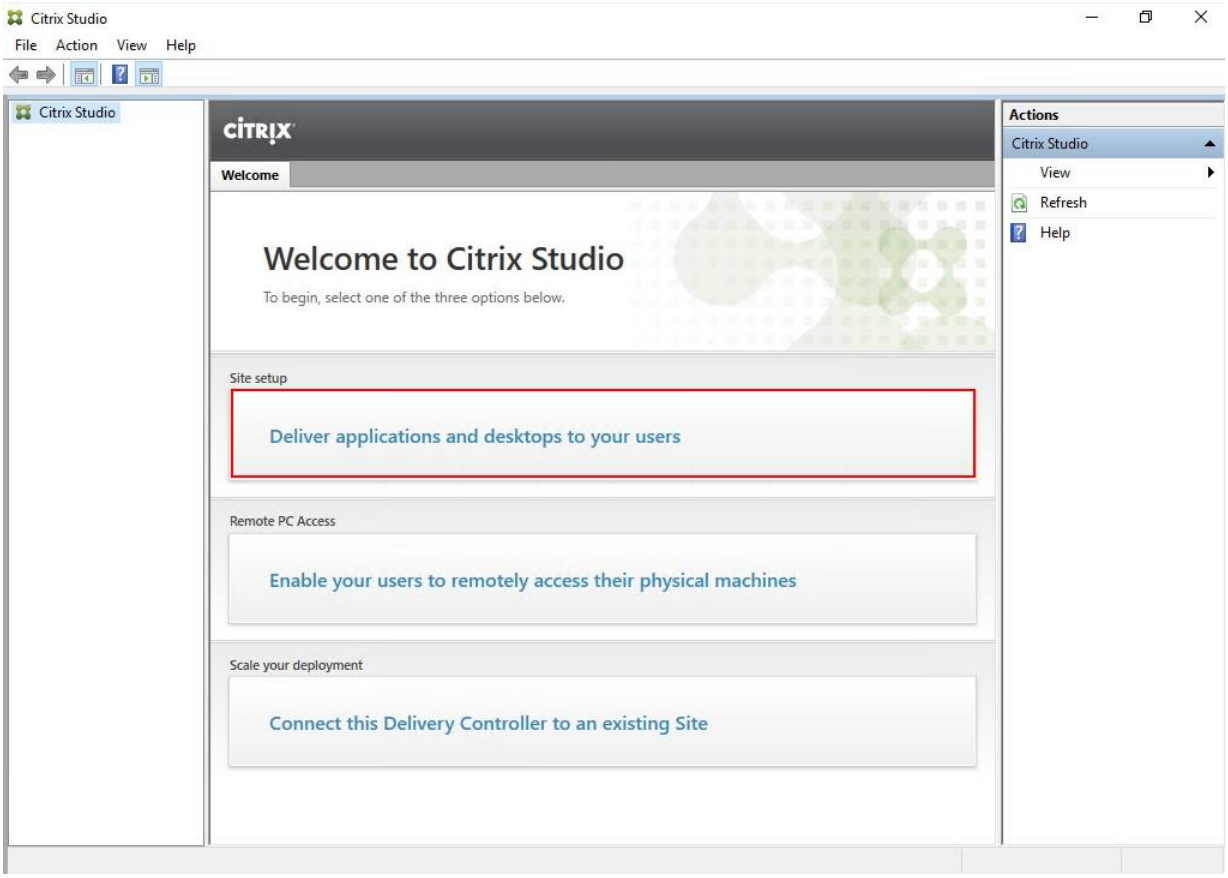


2. Select the “A fully configured, production-ready Site” radio button.
3. Enter a site name.
4. Click Next.

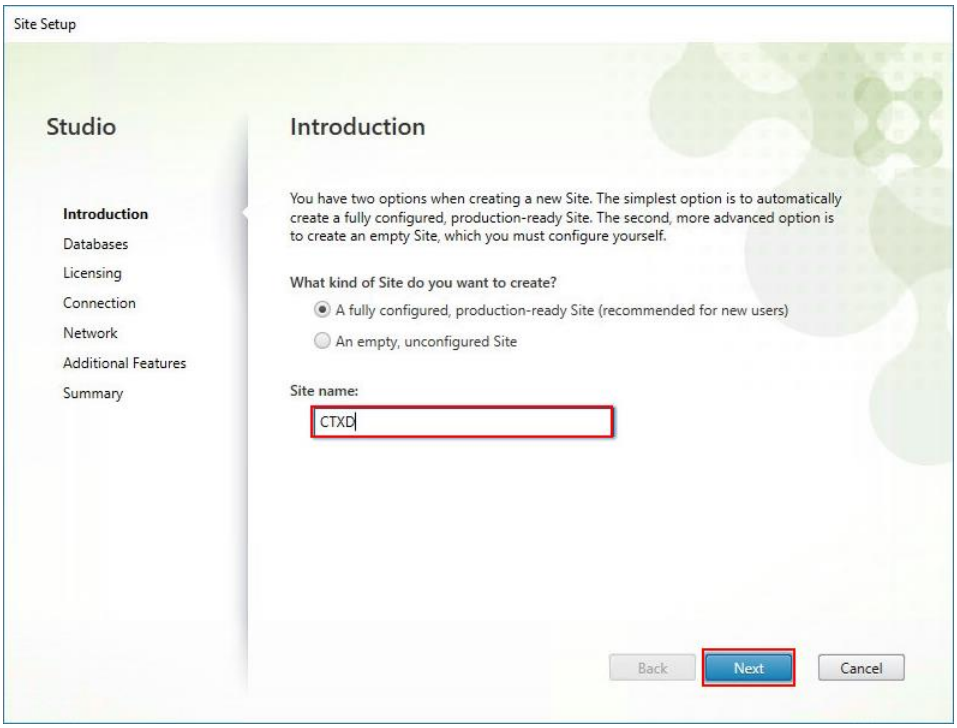


5. Provide the Database Server Locations for each data type and click Next.

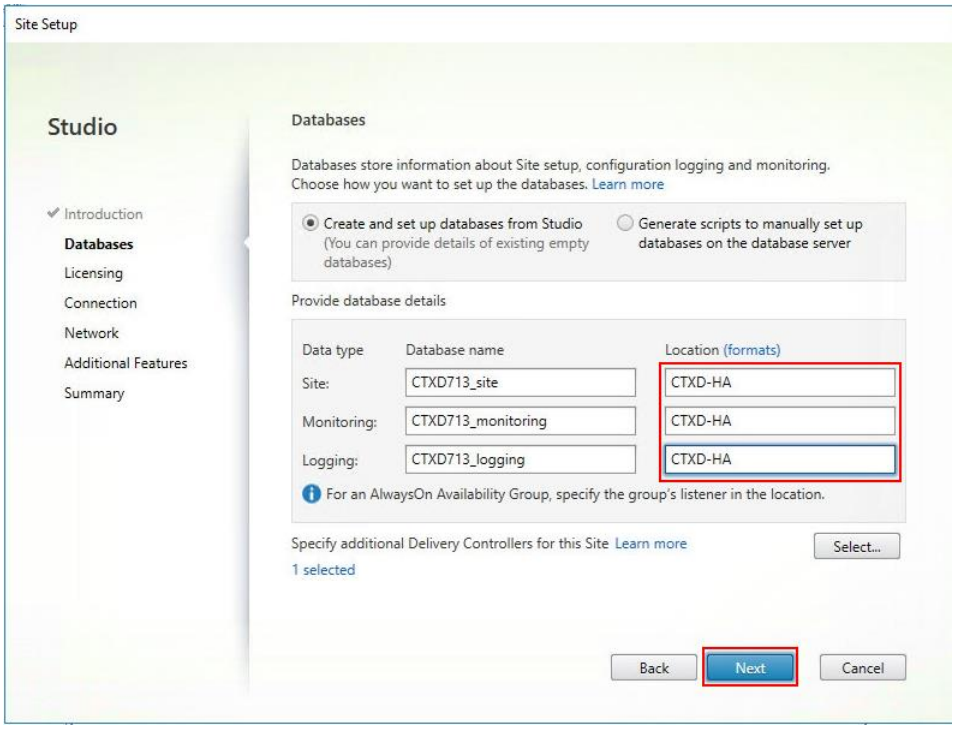


6. For an AlwaysOn Availability Group, use the group's listener DNS name.

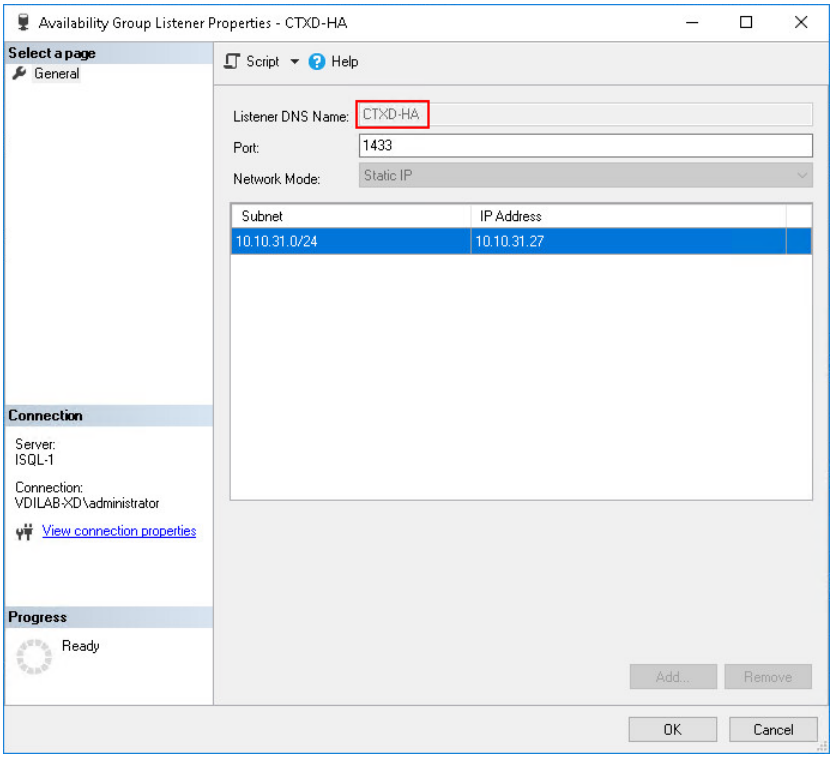


7. Provide the FQDN of the license server.
8. Click Connect to validate and retrieve any licenses from the server.

   **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.

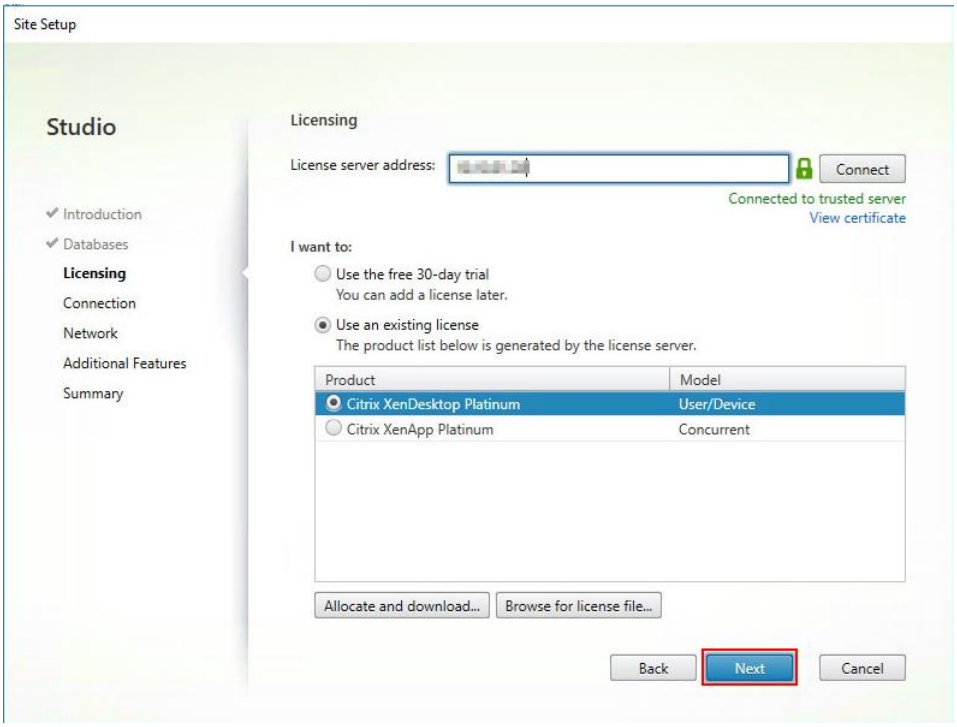


11. Select the Connection type of ‘Microsoft System Center Virtual Machine Manager’.
12. Enter the Connection Address to the SCVMM Server.
13. Enter the username (in username@domain format) for the vCenter account.
14. Provide the password for the Domain Admin account.
15. Provide a connection name.
16. Select the Studio tools radio button.

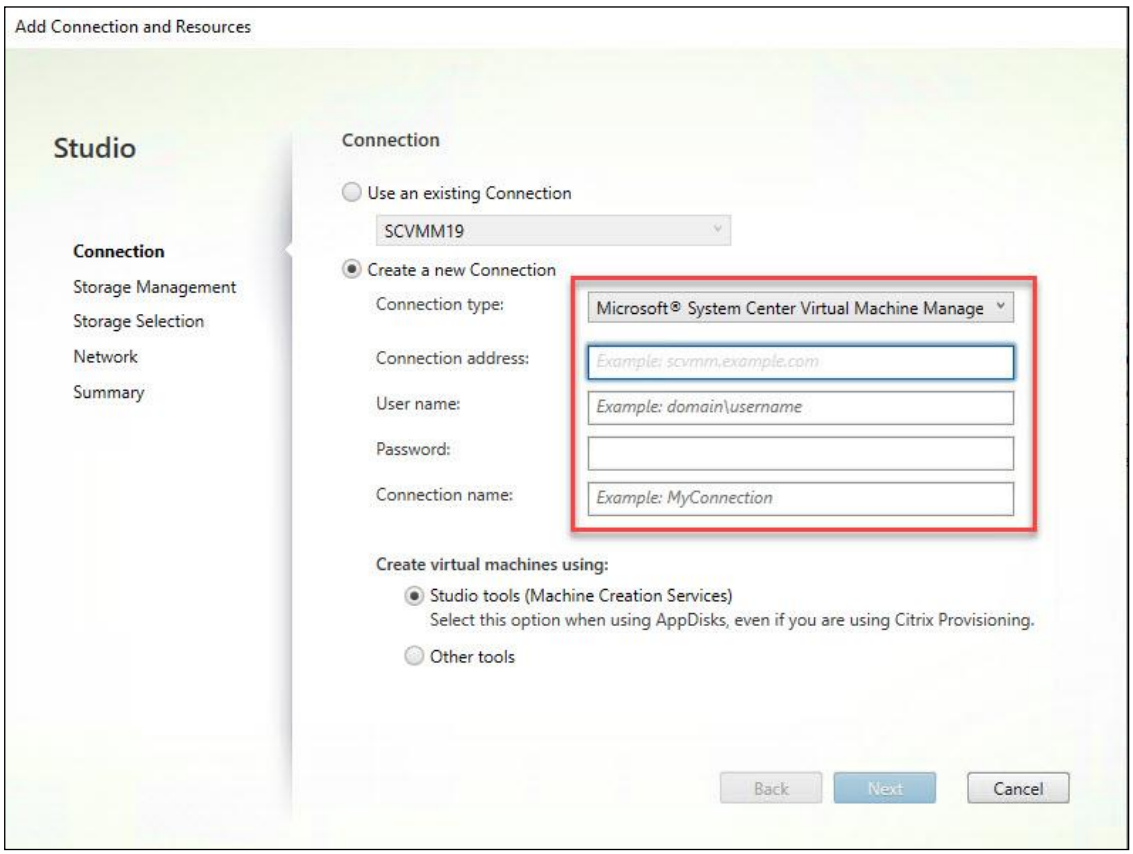


17. Click Next.
18. Select HyperFlex Cluster that will be used by this connection.
19. Check Studio Tools radio button required to support desktop provisioning task by this connection.
20. Click Next.

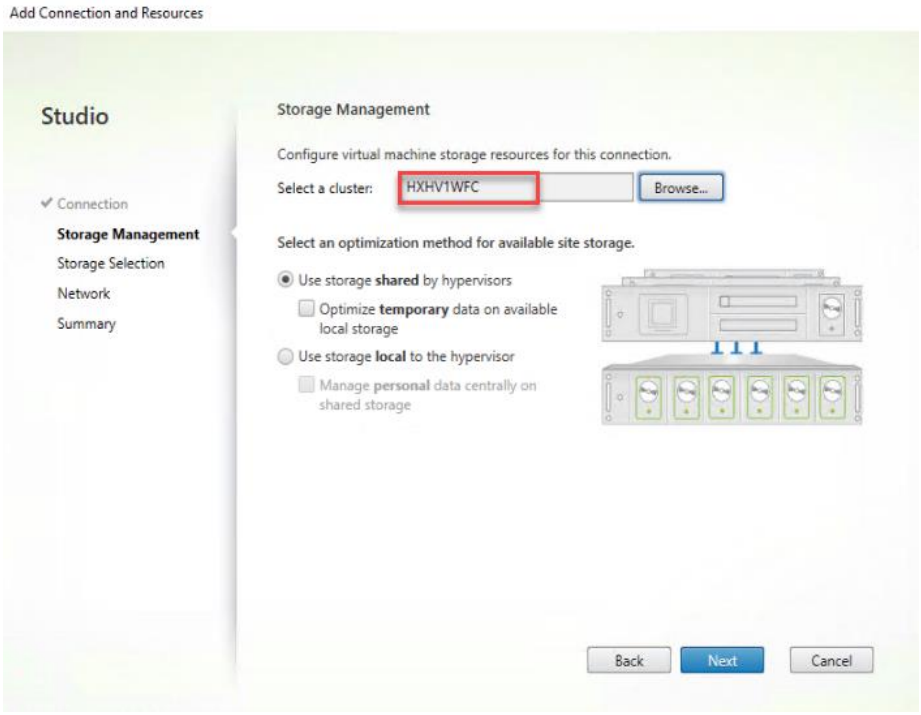


21. Make Storage selection to be used by this connection.

22. Click Next.

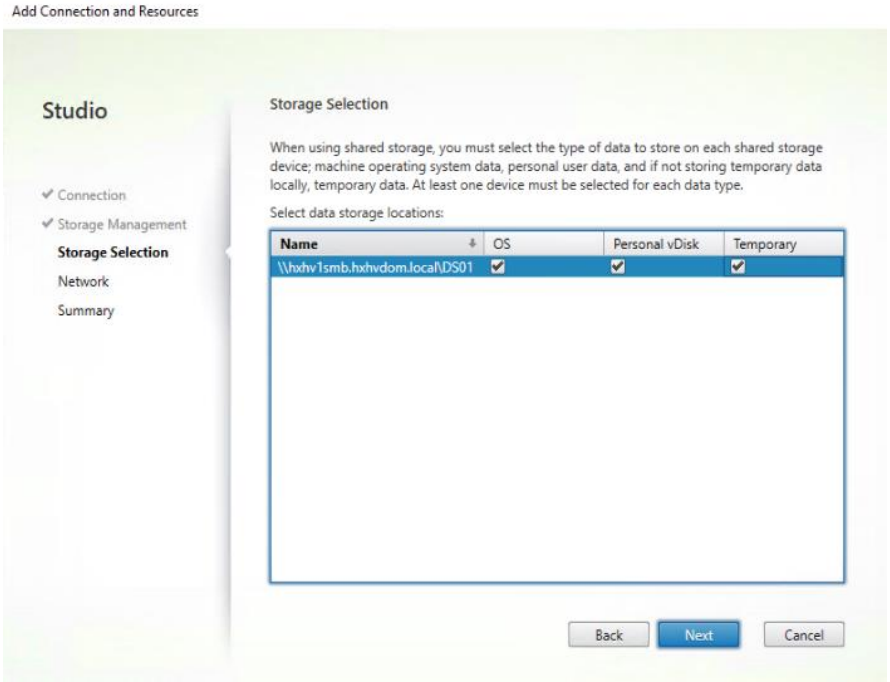


23. Make Network selection to be used by this connection.

24. Click Next.

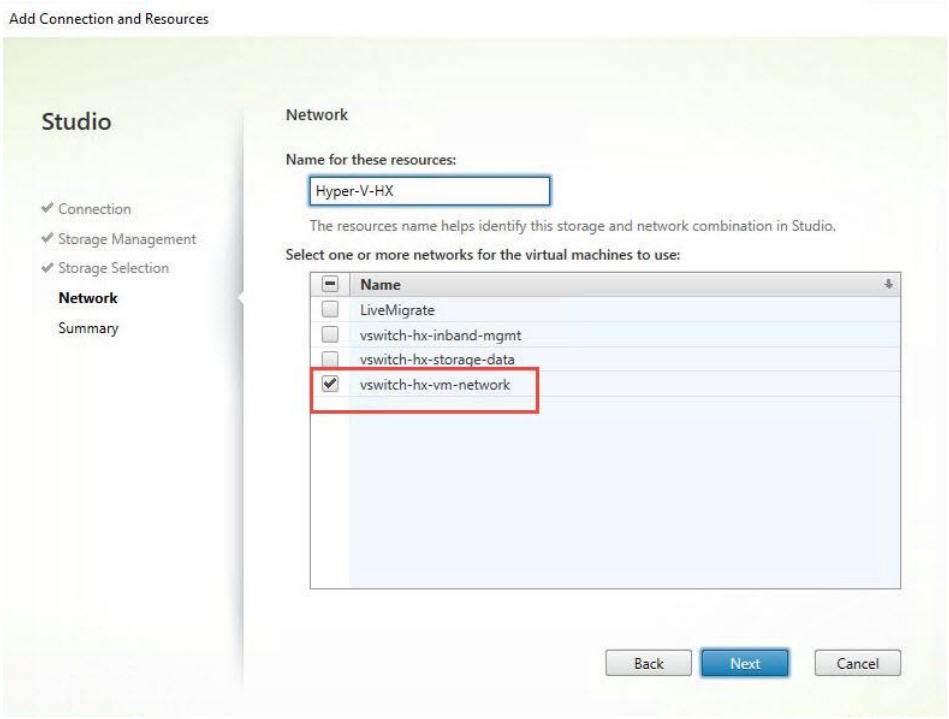


25. Select Additional features.

26. Click Next.

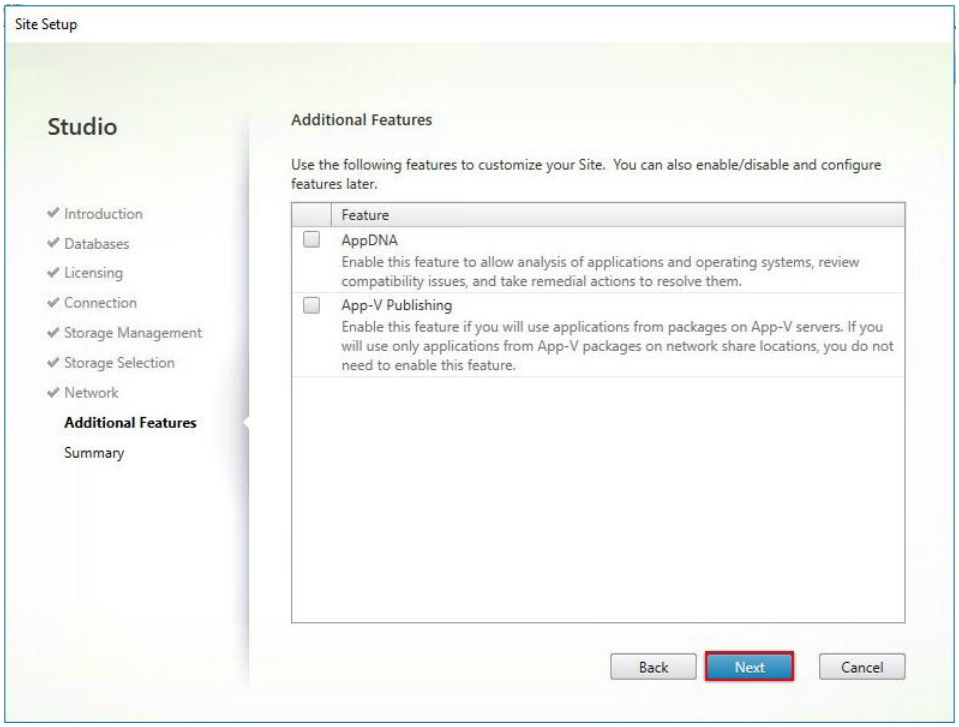


27. Review Site configuration Summary and click Finish.

## Configure the Citrix VDI Site Administrators

To configure the Citrix VDI site administrators, follow these 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.

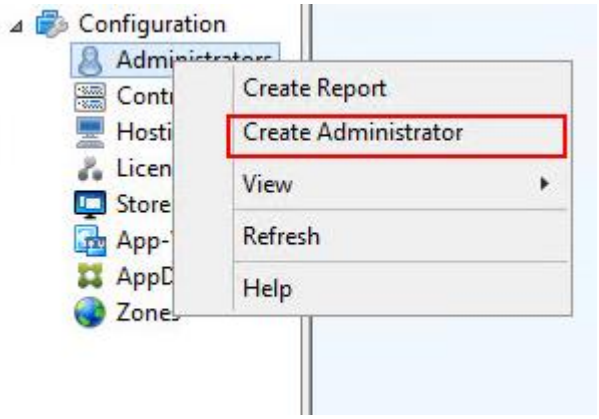


3. Select/Create appropriate scope and click Next.

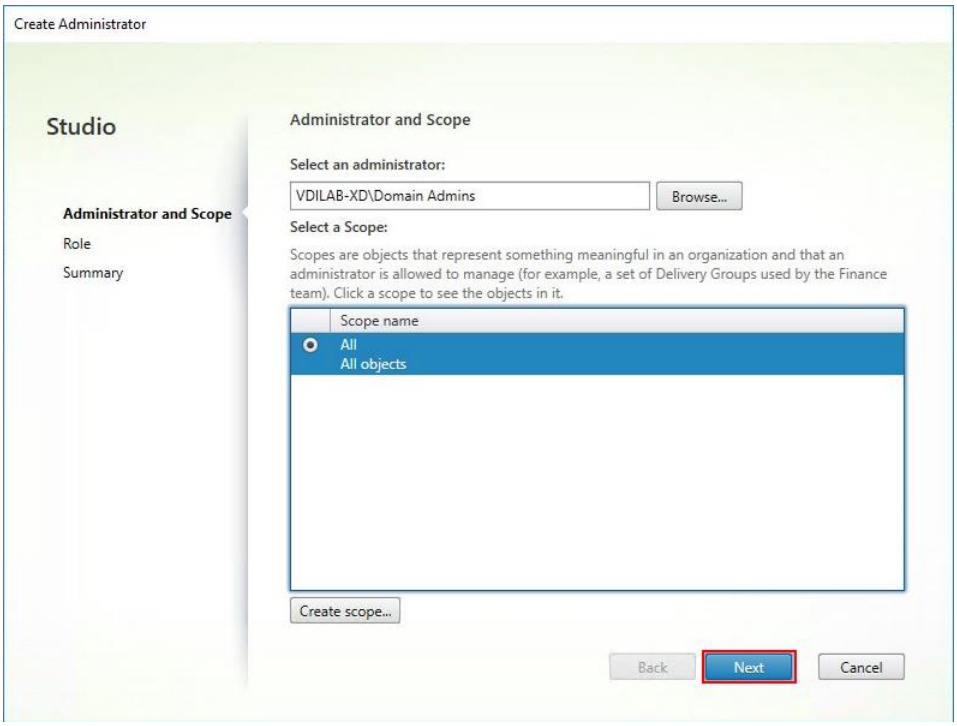


4. Choose an appropriate Role.

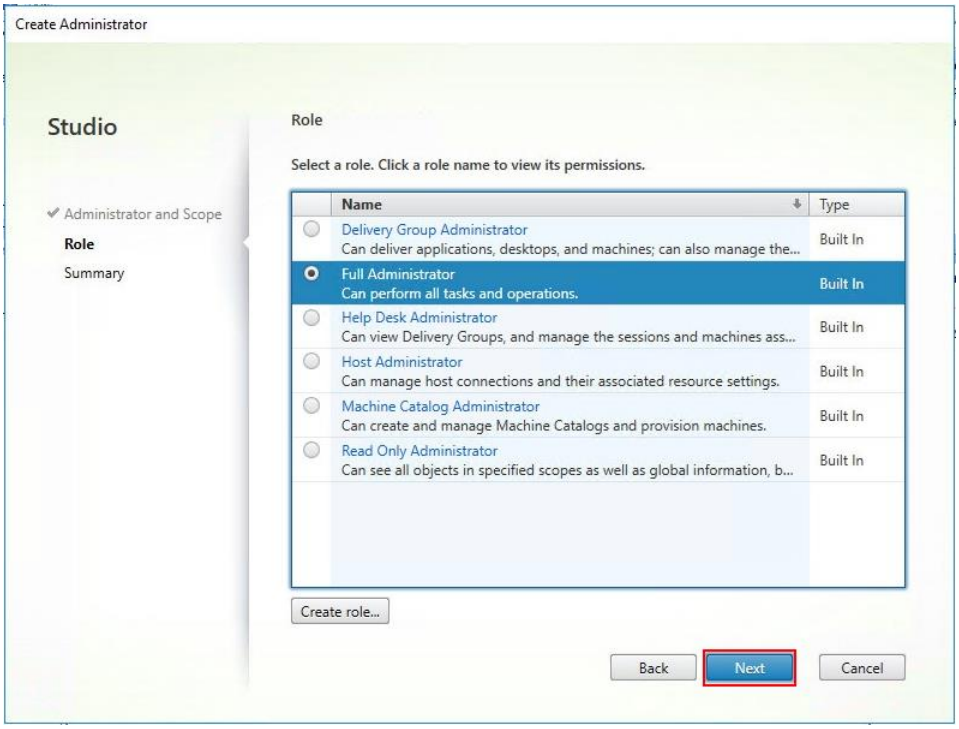


5. Review the Summary, check Enable administrator, and click Finish.

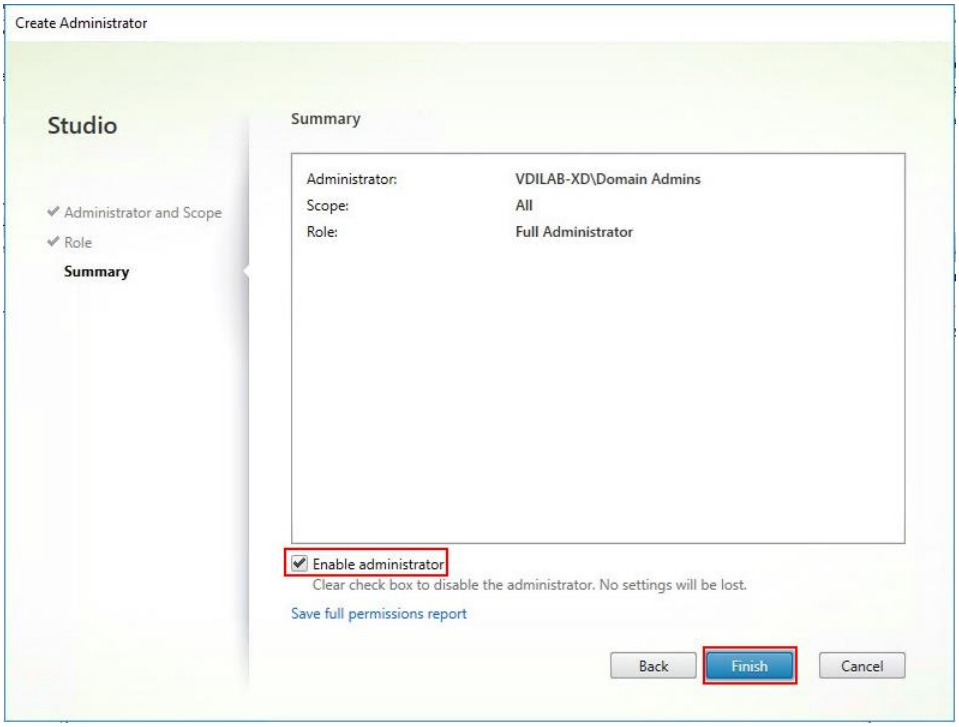


## Configure Additional Desktop Controller

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, follow these 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.

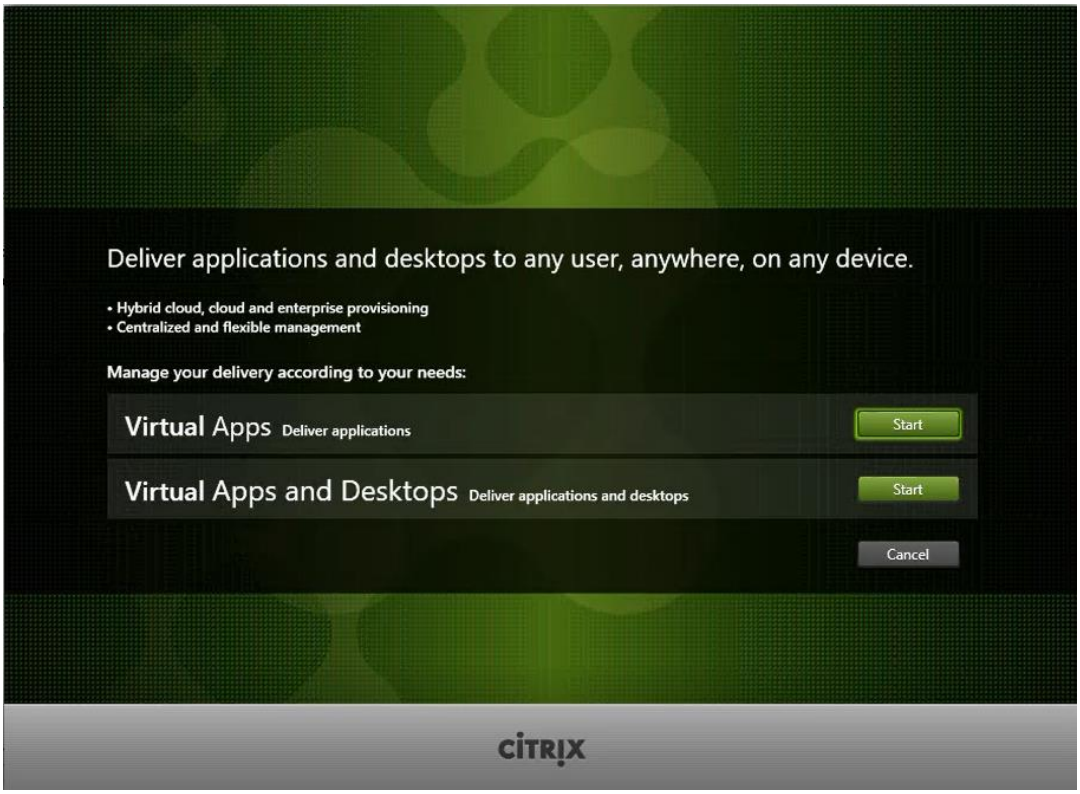


3. Click Delivery Controller.

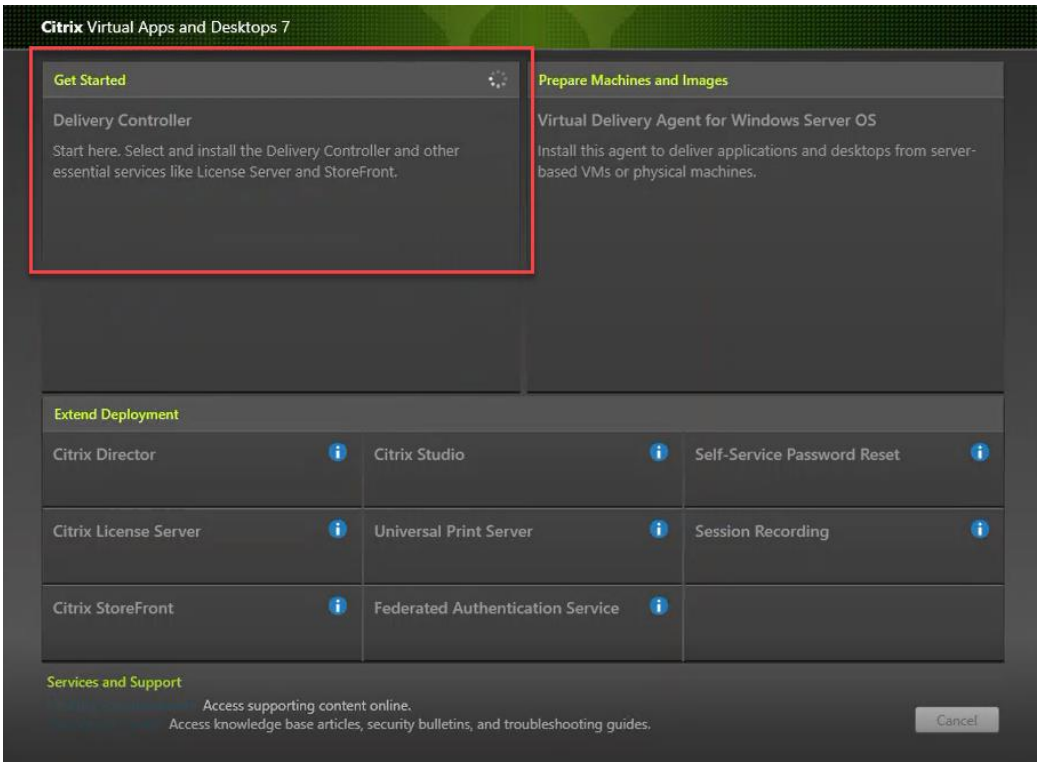


4. Repeat these steps used to install the first Delivery Controller, including the step of importing an SSL certificate for HTTPS between the controller and Hyper-V.

5. Review the Summary configuration.

6. Click Install.

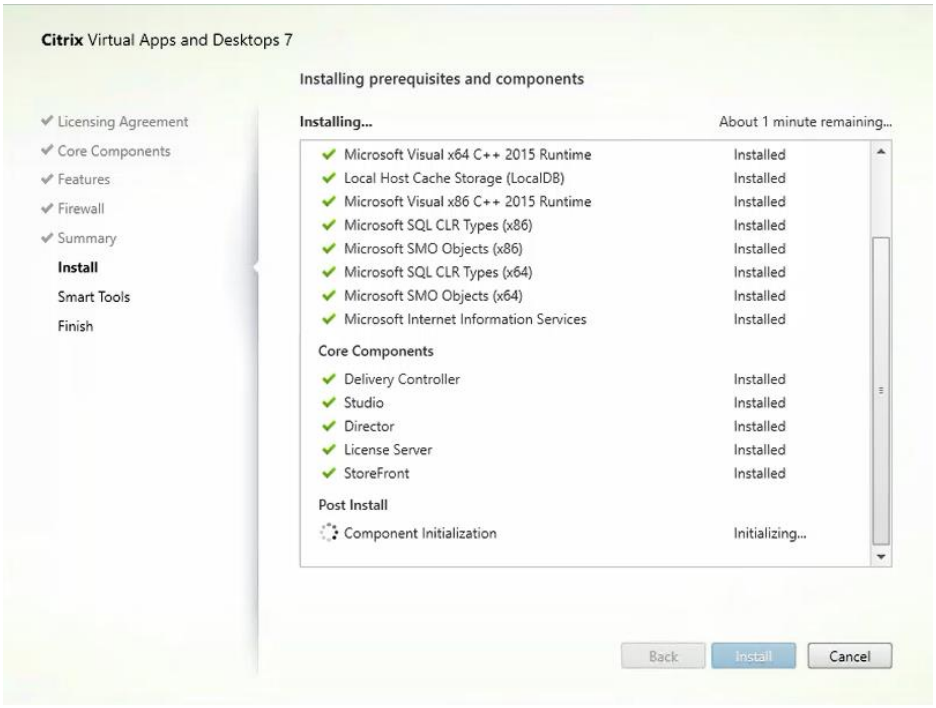

7. (Optional) Click the “I want to participate in Call Home.”

8. Click Next.

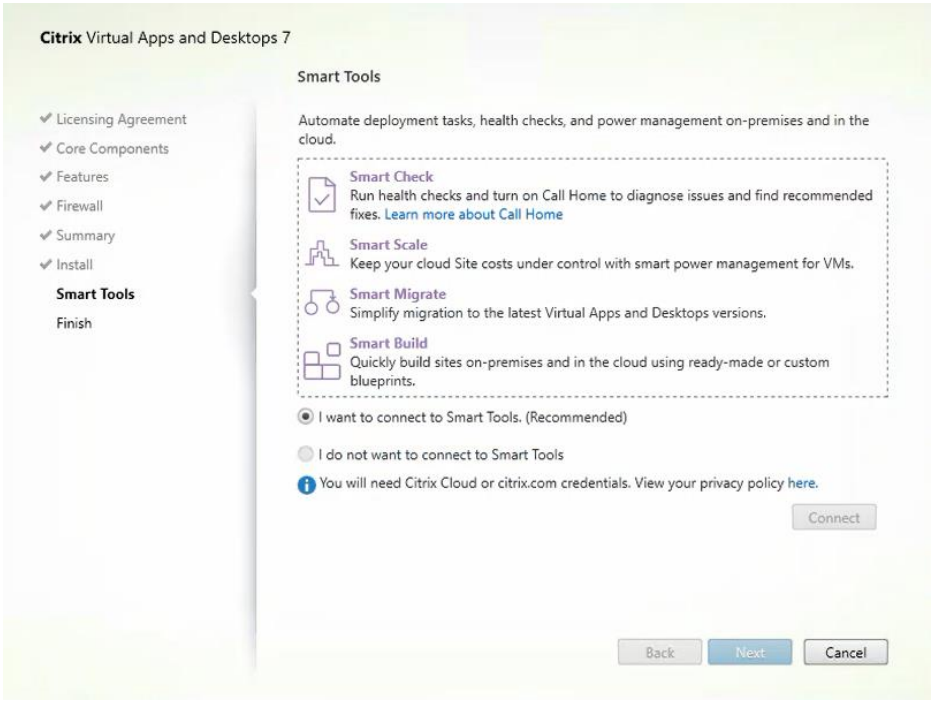


9. Verify the components installed successfully.

10. Click Finish.

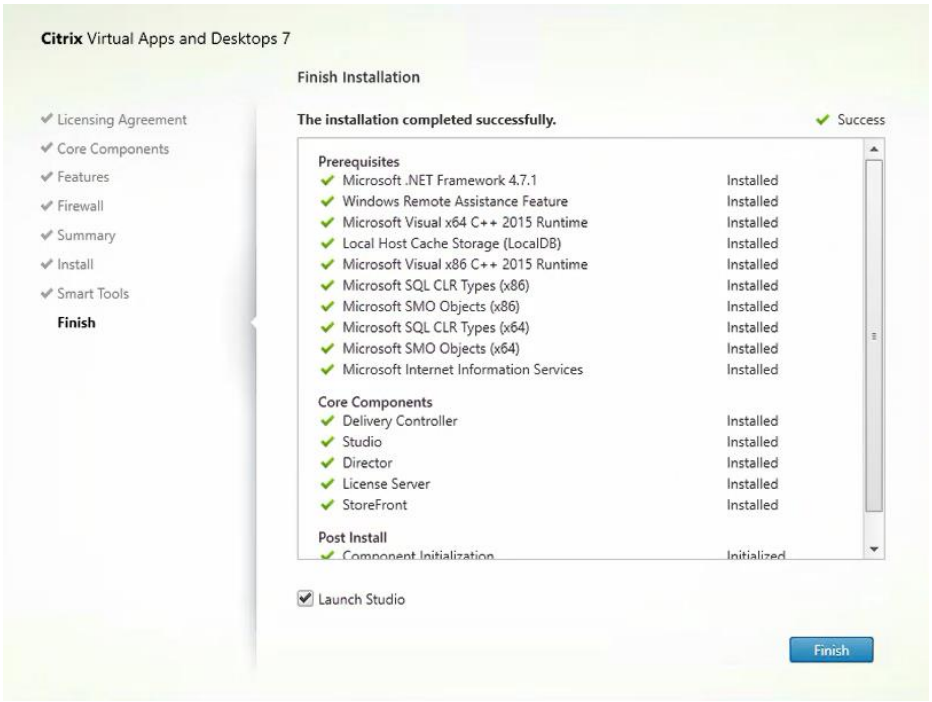

## Add the Second Delivery Controller to the Citrix Desktop Site

To add the second Delivery Controller to the Citrix Desktop Site, follow these steps:

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

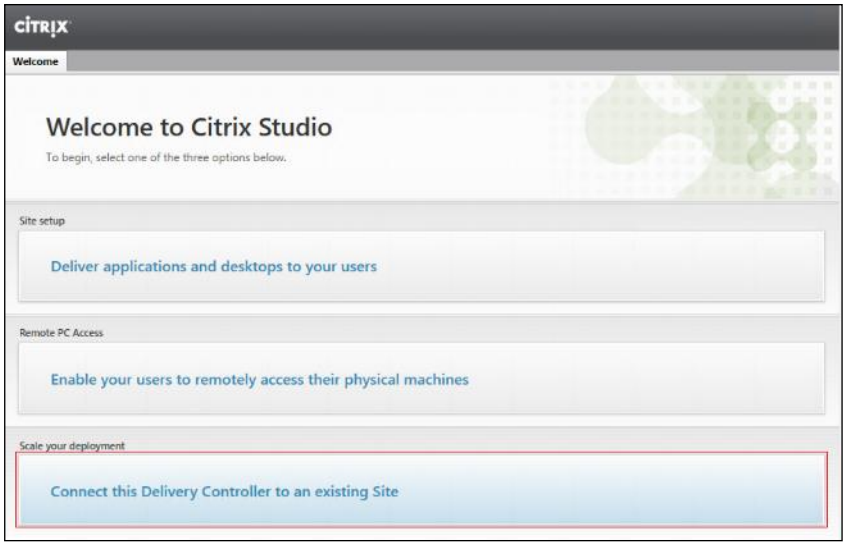


2. Enter the FQDN of the first delivery controller.

3. Click OK.

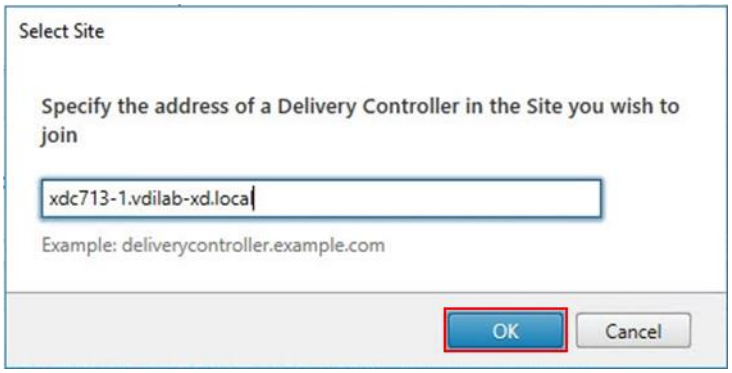


4. Click Yes to allow the database to be updated with this controller's information automatically.

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

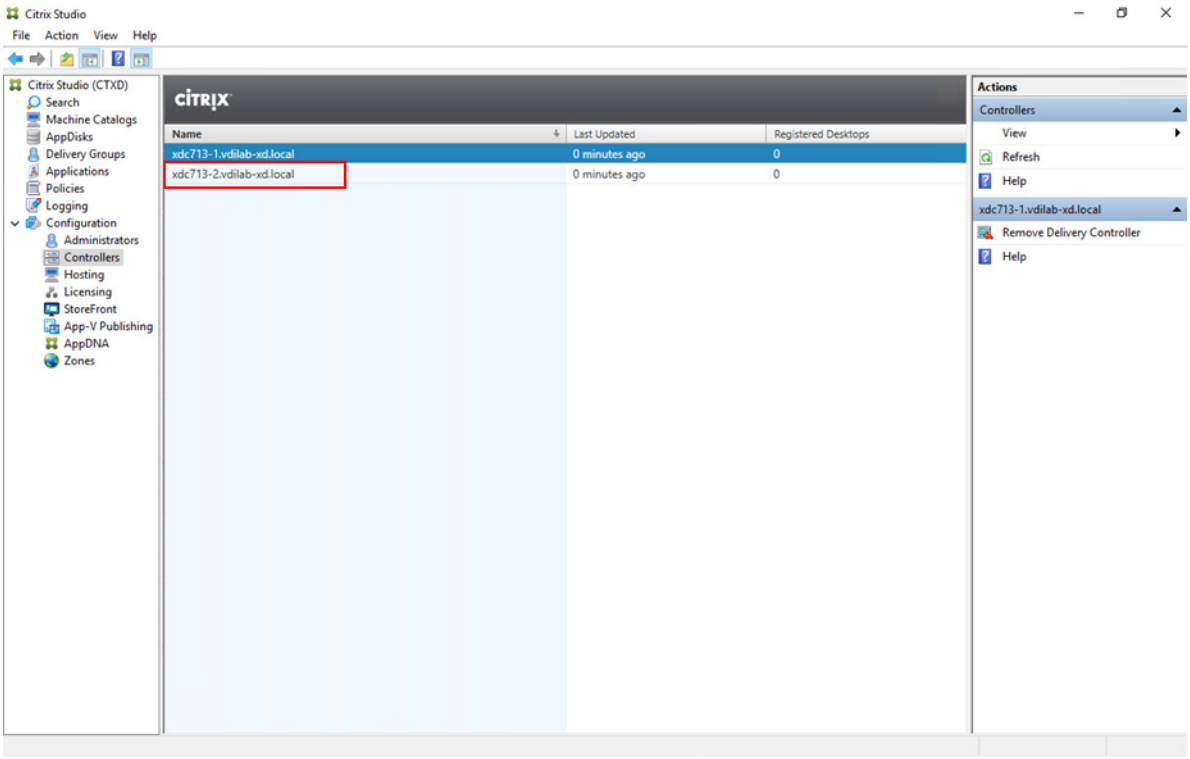


## Install and Configure StoreFront

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, follow these steps:

1. To begin the installation of the StoreFront, connect to the first StoreFront server and launch the installer from the Citrix Desktop 1909 ISO.

2. Click Start.

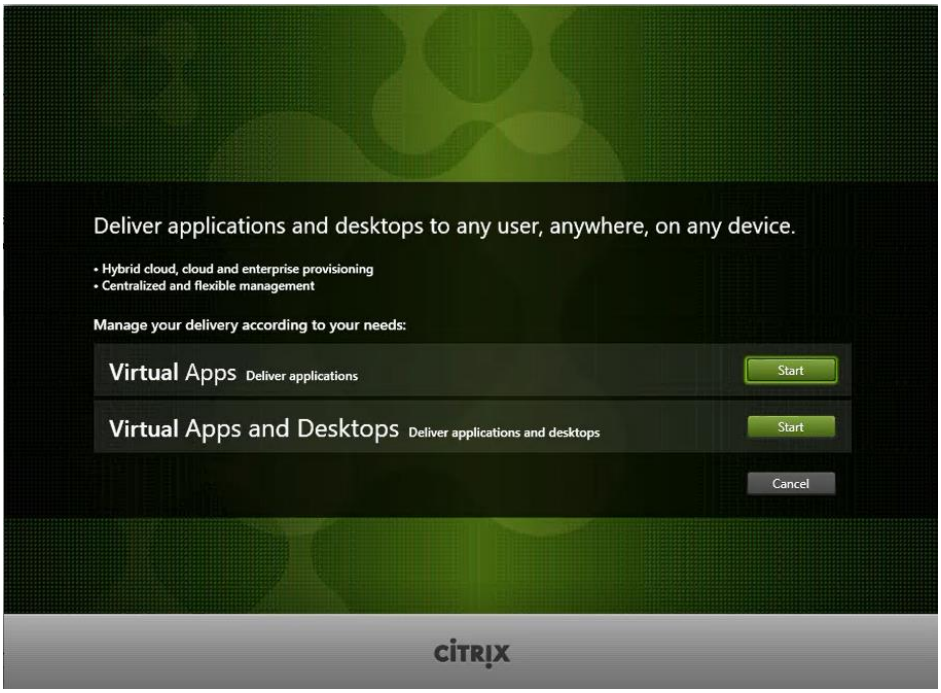


3. Click Extend Deployment Citrix StoreFront.

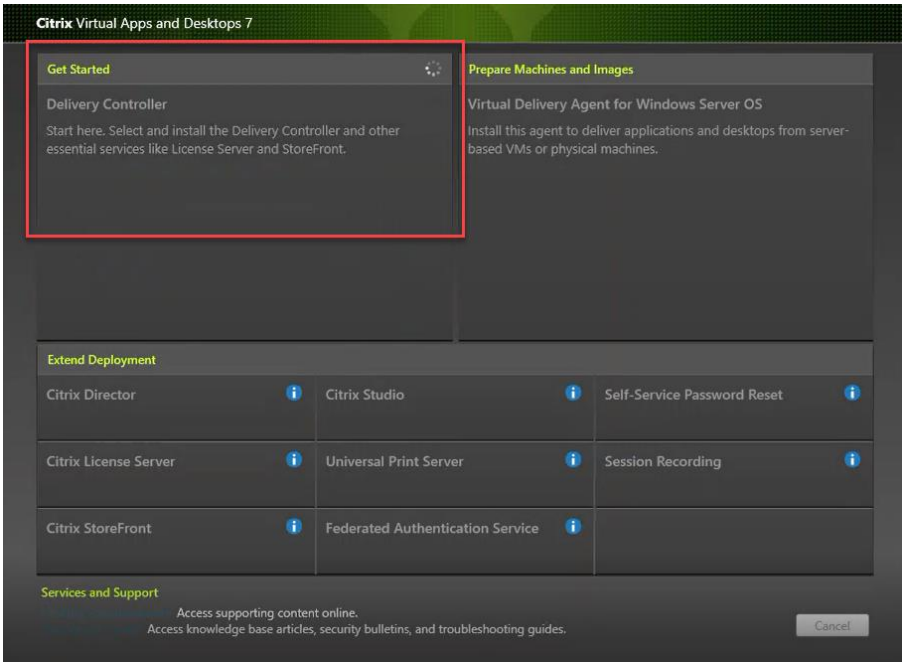


4. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

5. Click Next.

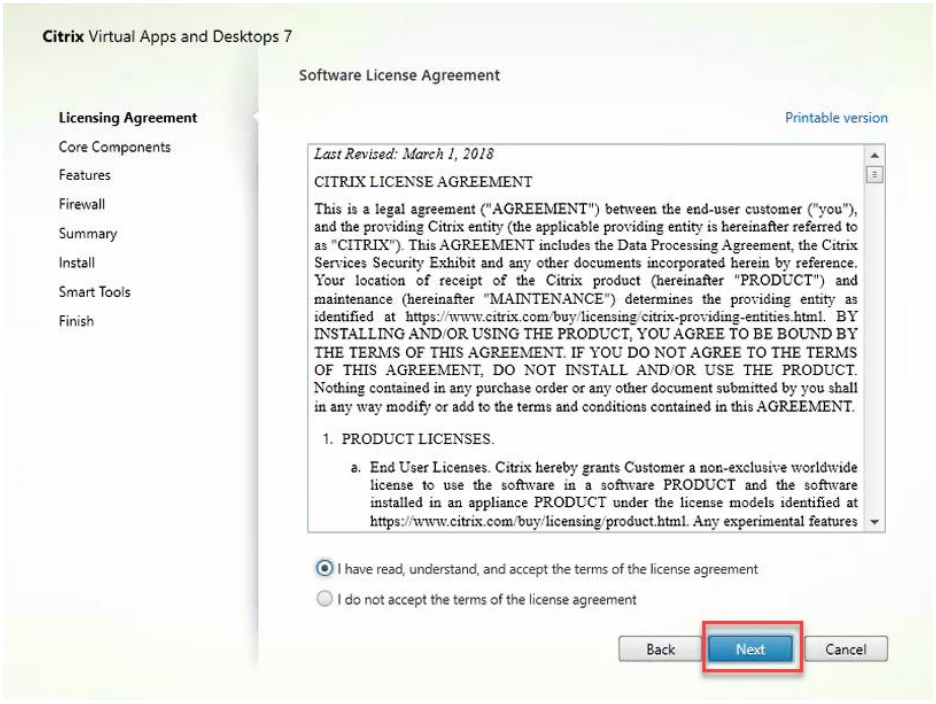


6. Select Storefront and Click Next.

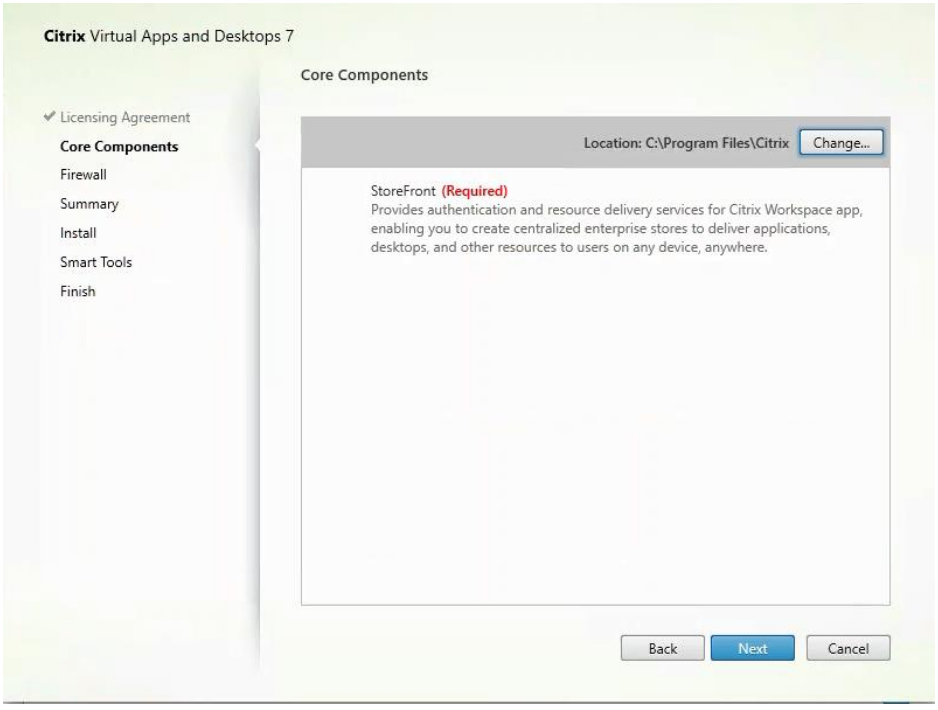


7. Select the default ports and automatically configured firewall rules.

8. Click Next.

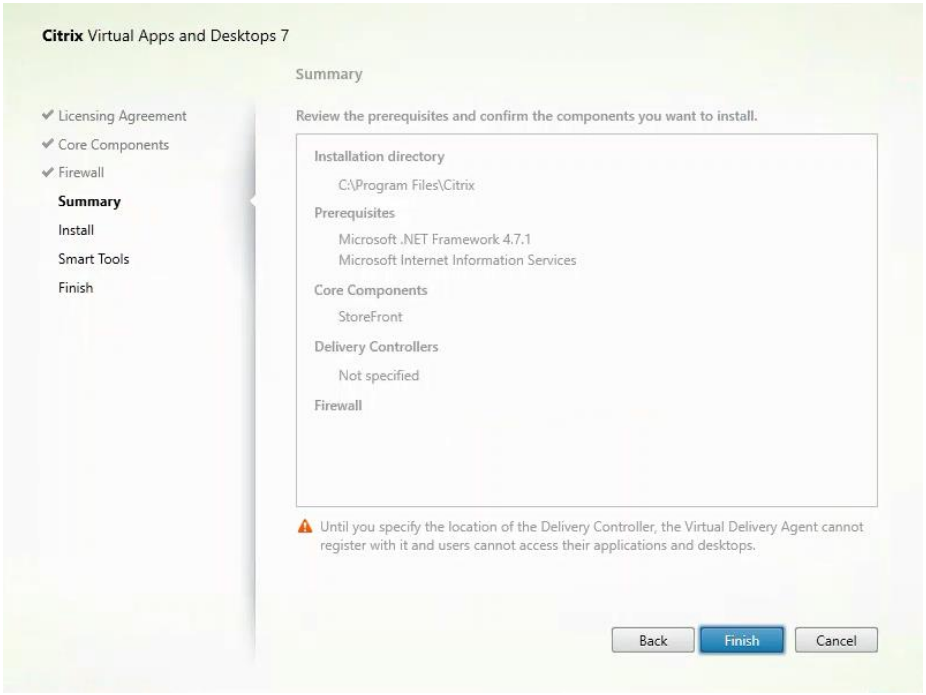


9. Click Install.

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

11. Click Next.

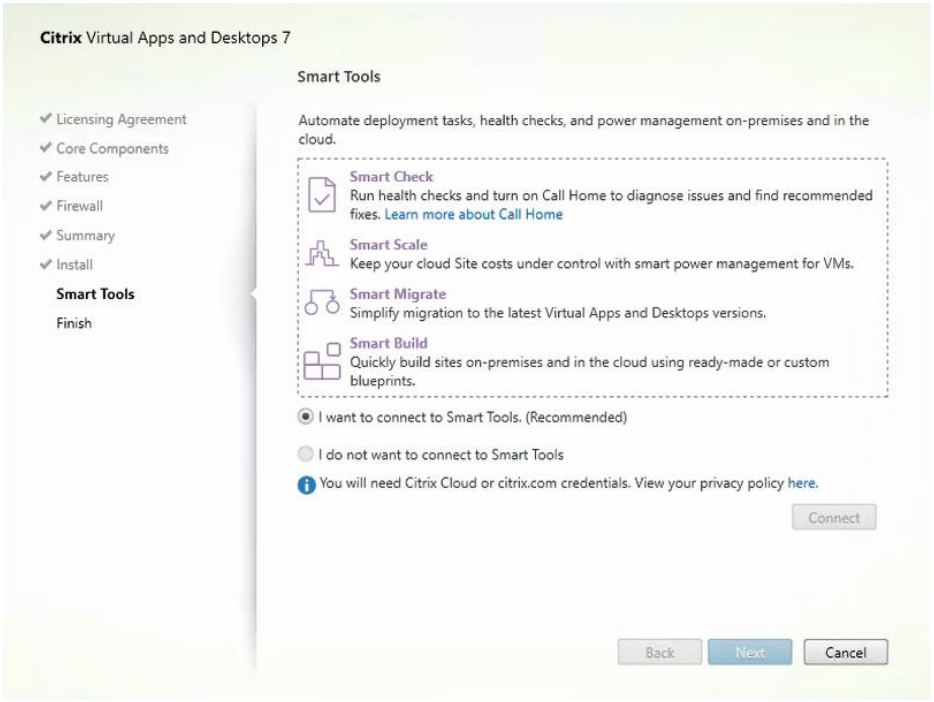


12. Check “Open the StoreFront Management Console.”

13. Click Finish.

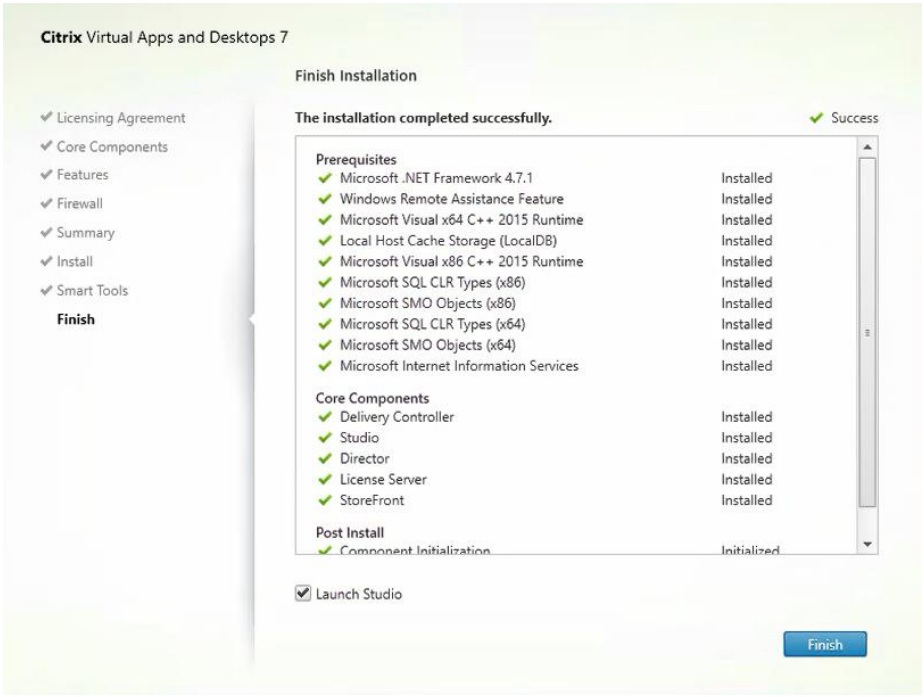


14. Click Create a new deployment.

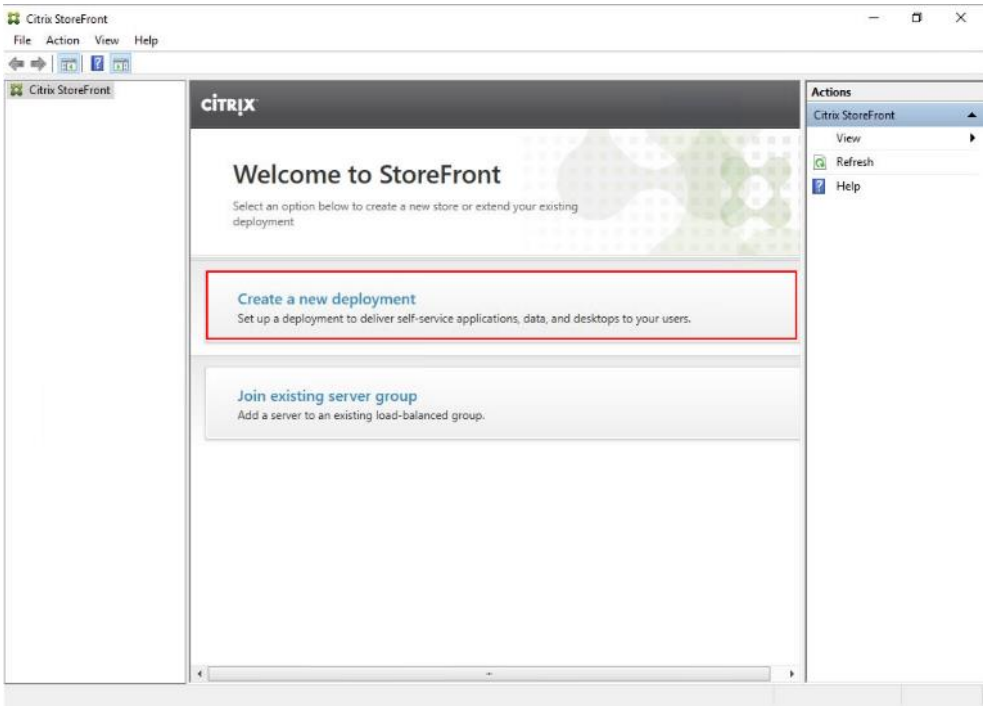


15. Specify the URL of the StoreFront server and click Next.

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

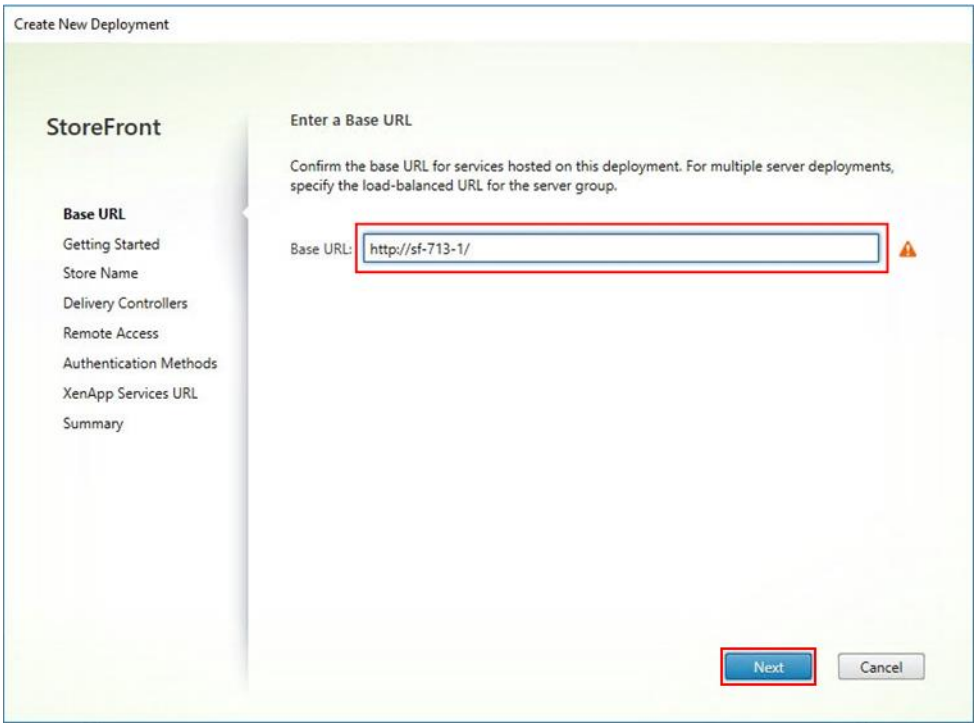


16. Click Next.

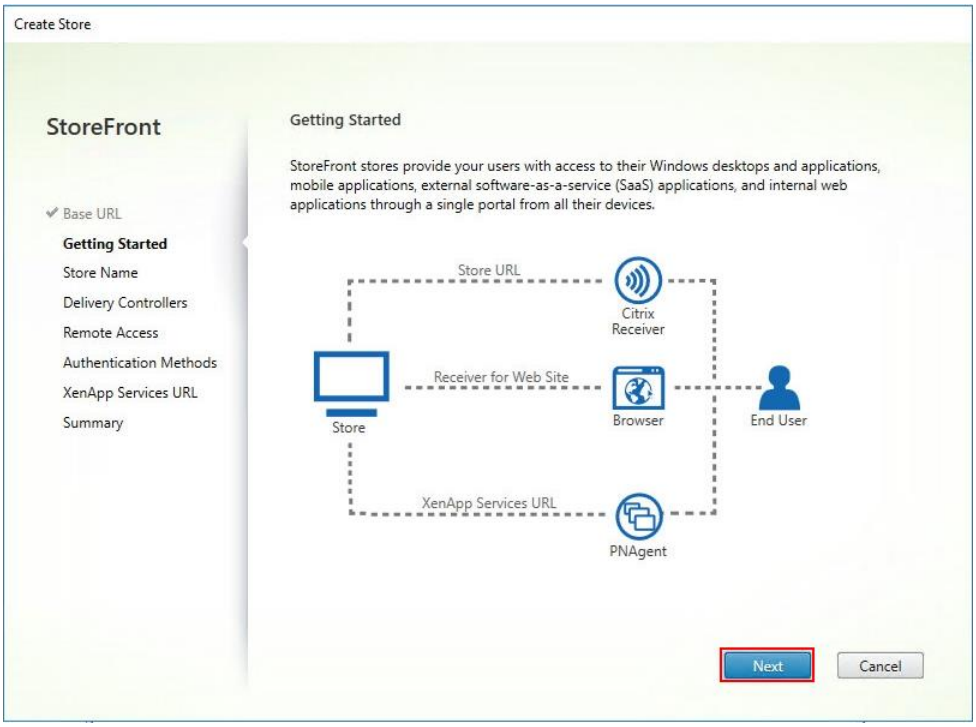


17. Specify a name for your store and click Next.

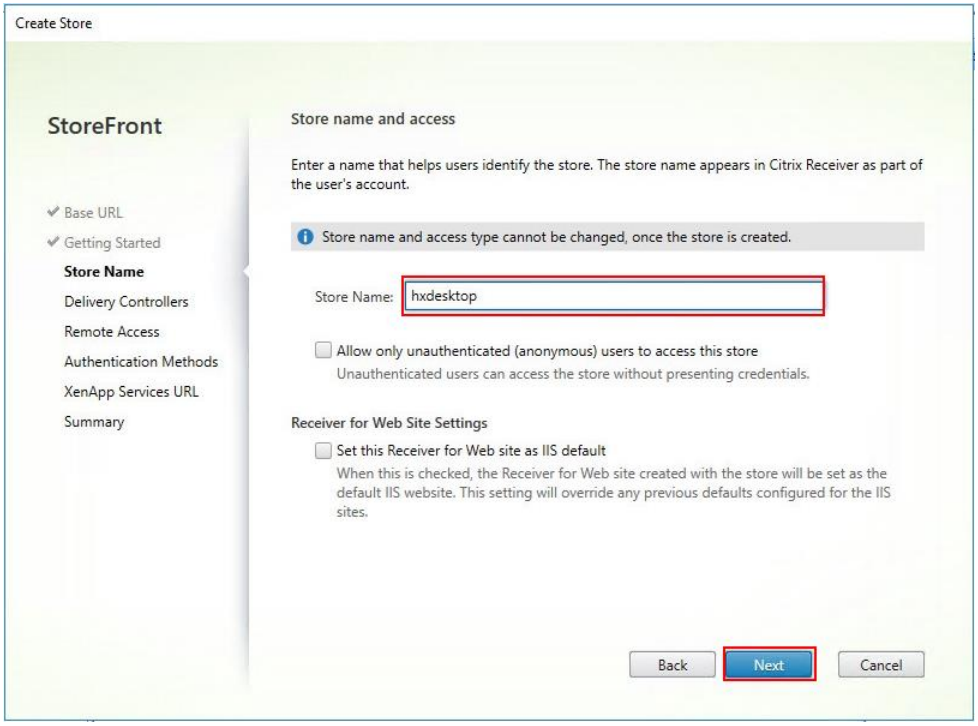


18. Add the required Delivery Controllers to the store and click Next.

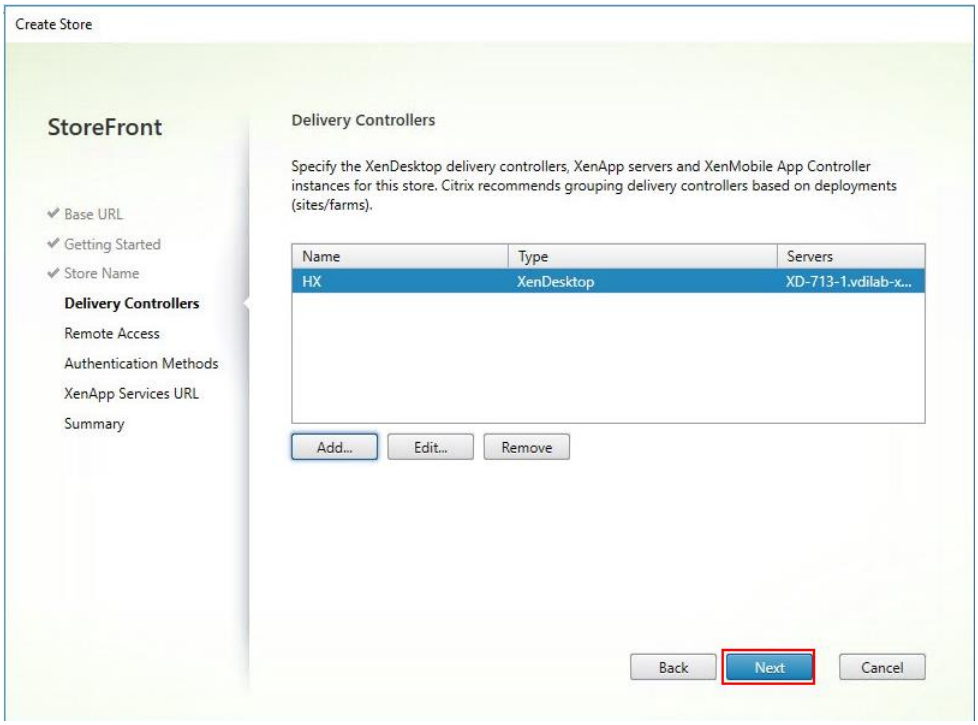


19. Specify how connecting users can access the resources, in this environment only local users on the internal network are able to access the store and click Next.

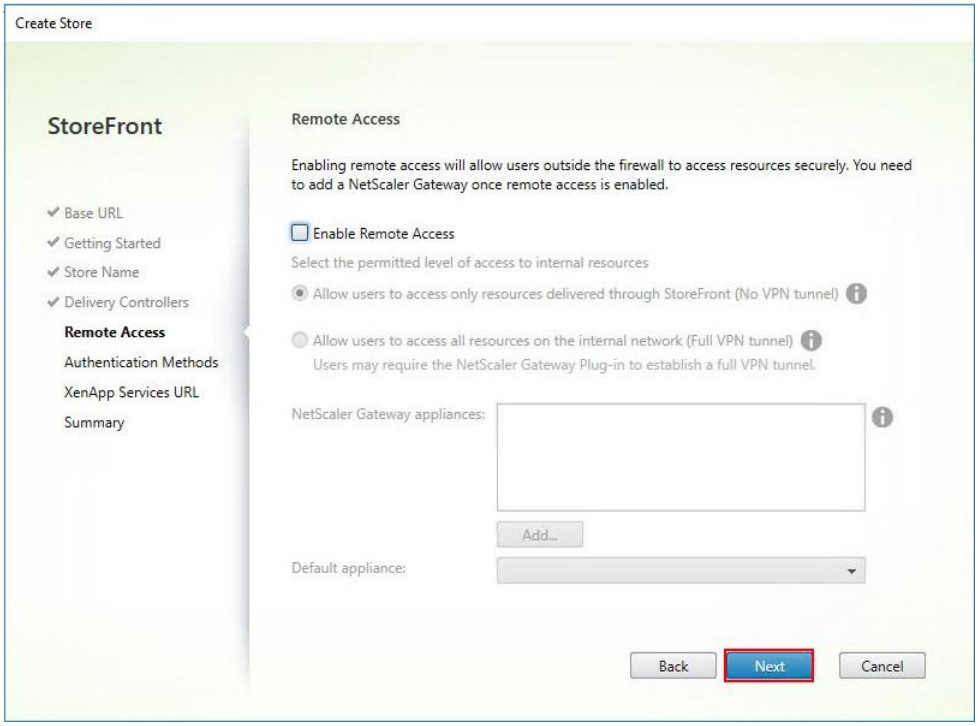


20. On the “Authentication Methods” page, select the methods your users will use to authenticate to the store and click Next. You can select from the following methods as shown below:

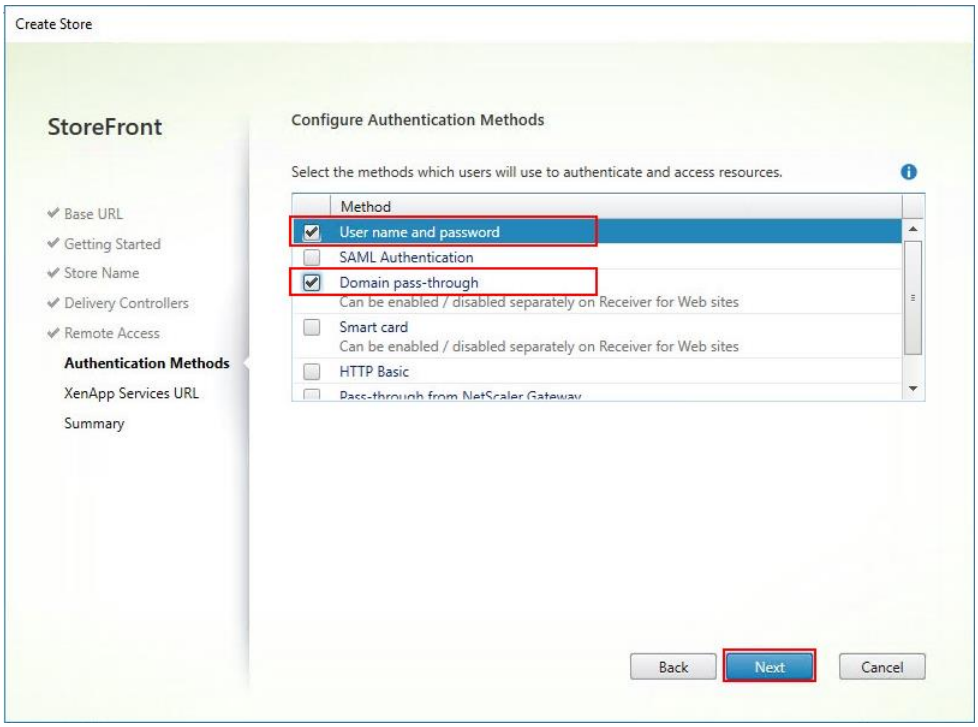


21. Username and password: Users enter their credentials and are authenticated when they access their stores.

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.

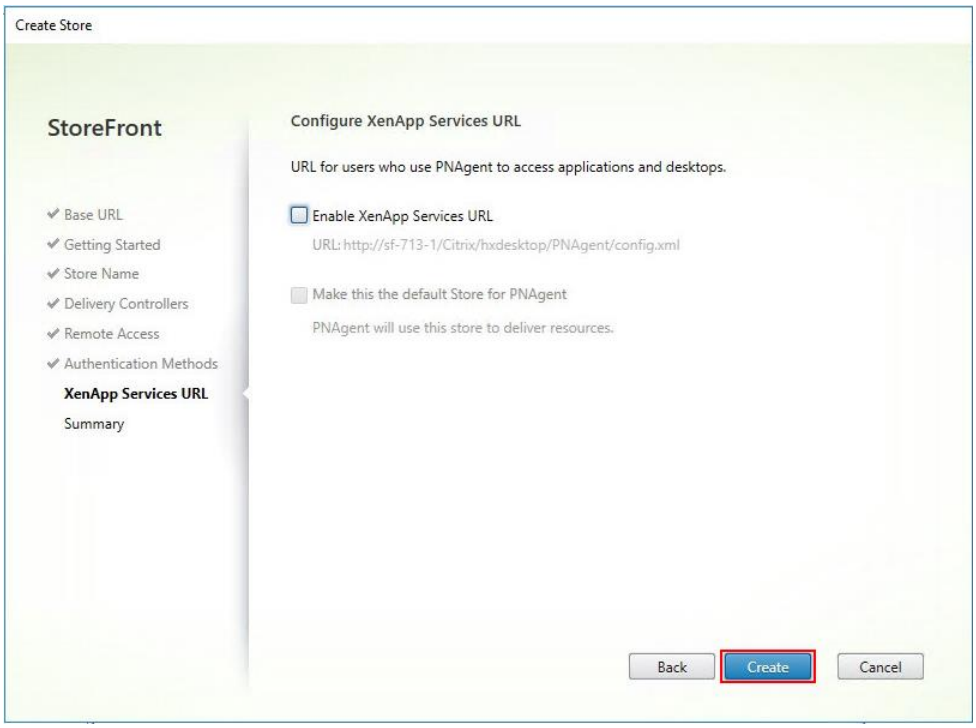


24. After creating the store click Finish.

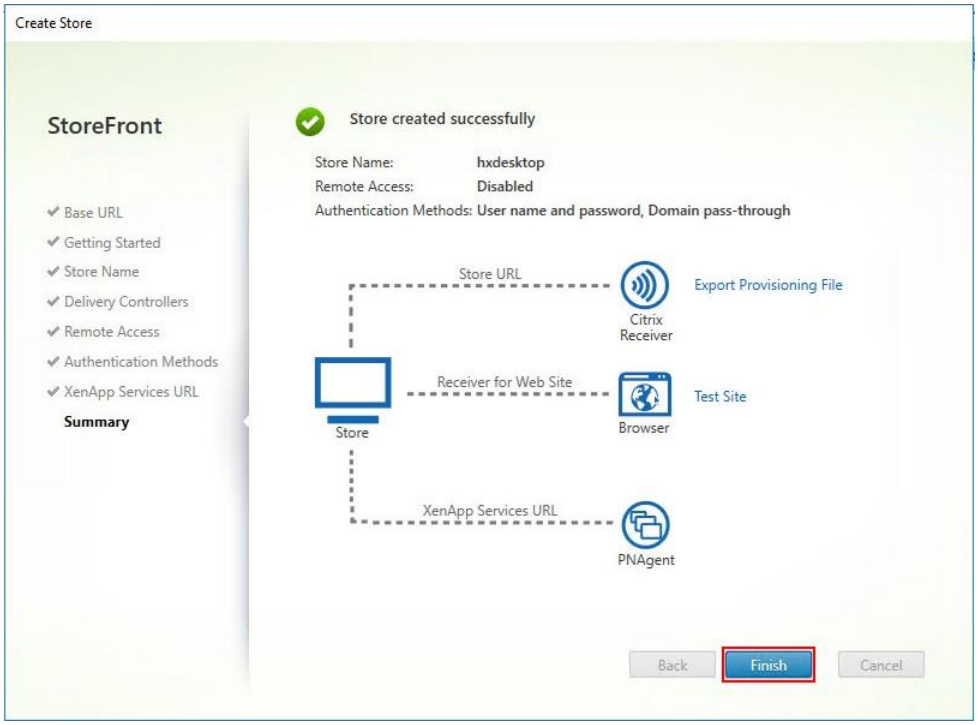

## Additional StoreFront Configuration

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

To configure additional StoreFront server, follow these steps:

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

2. Click Start.

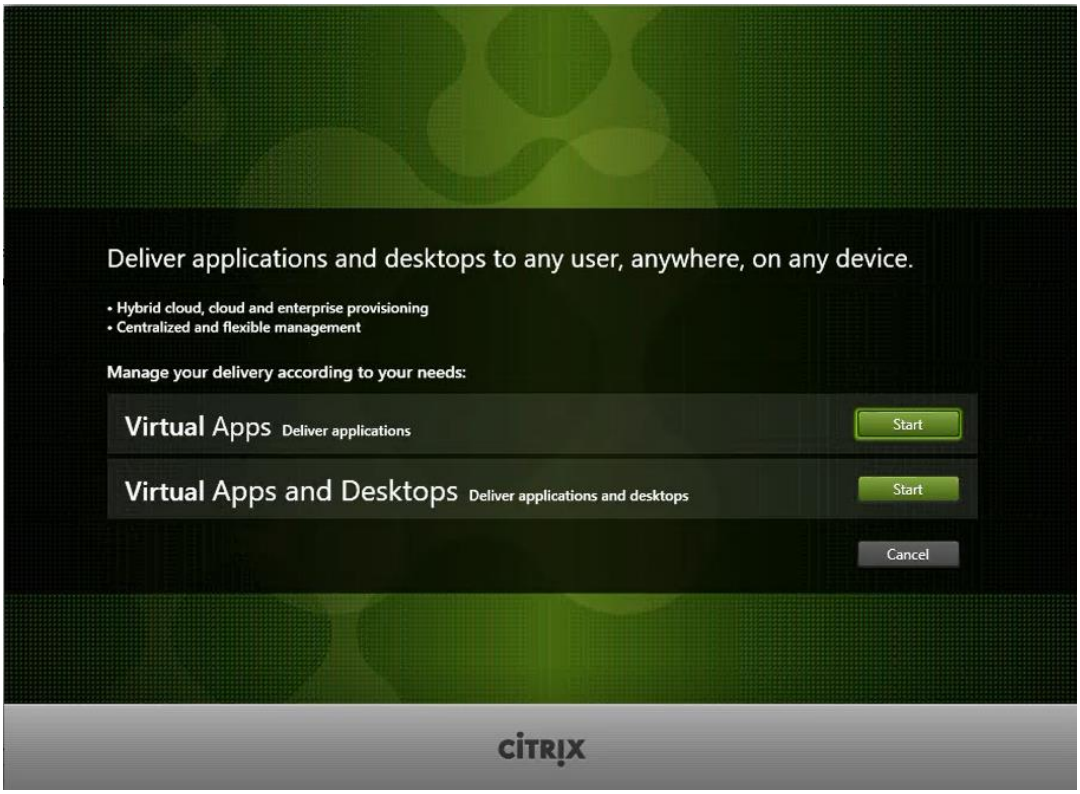


3. Click Extended Deployment Citrix StoreFront.

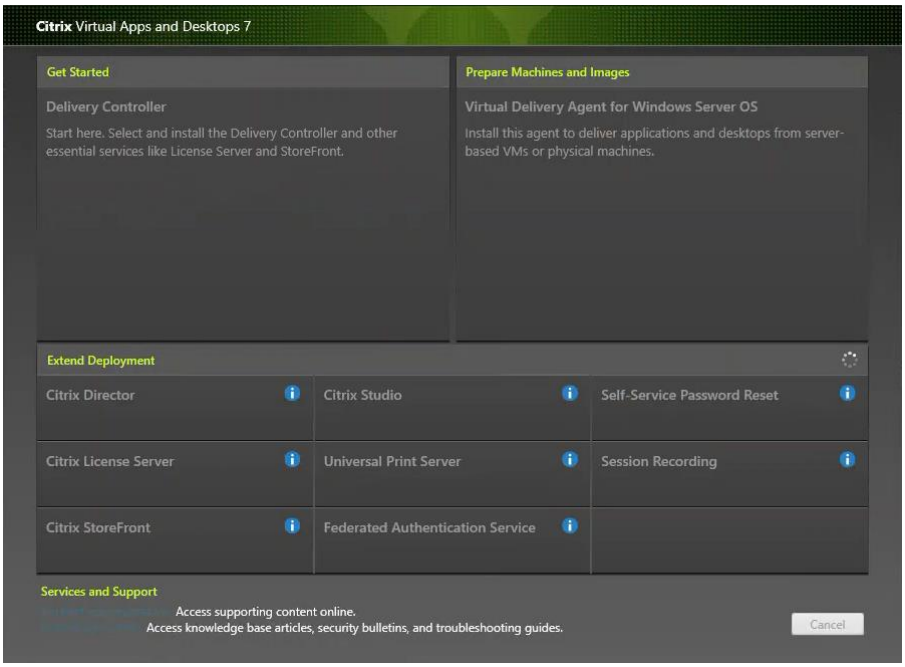


4. Repeat the same steps used to install the first StoreFront.
5. Review the Summary configuration.
6. Click Install.

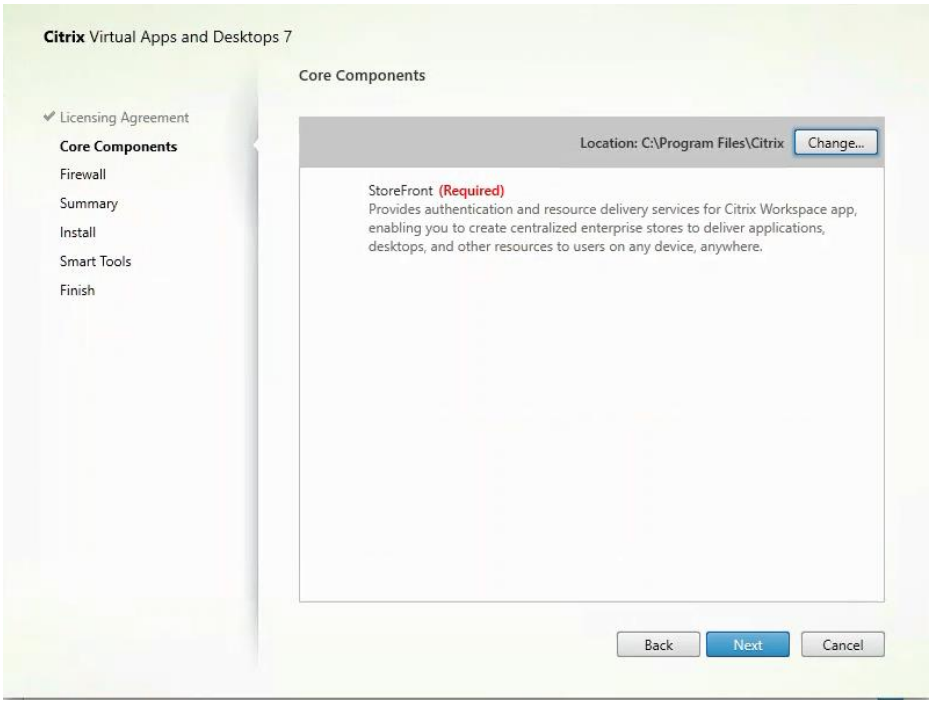


7. (Optional) Click “I want to participate in Call Home.”
8. Click Next.

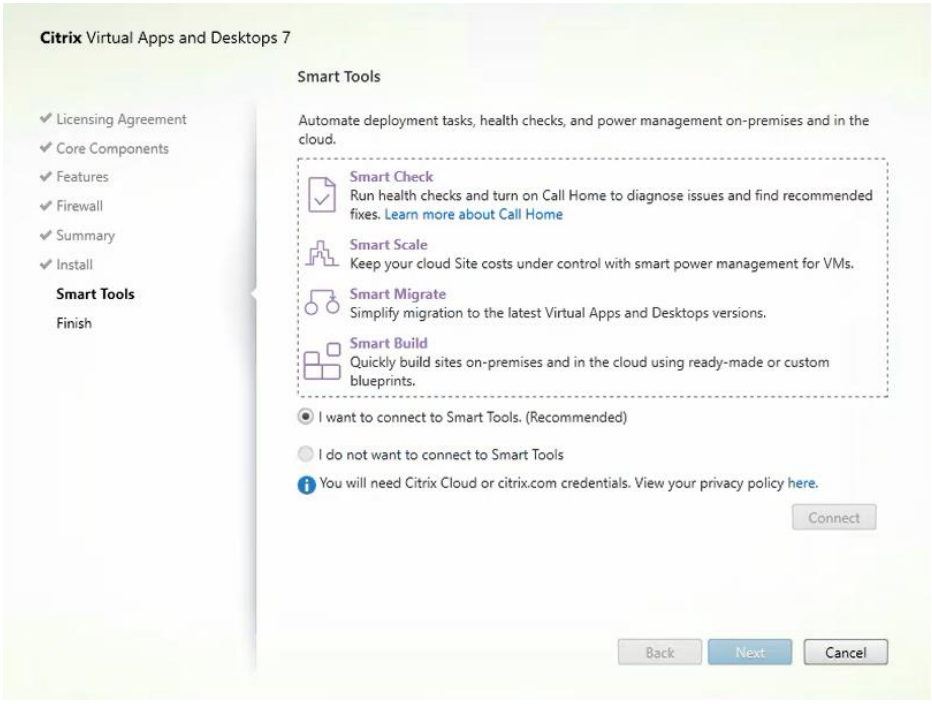


9. Check “Open the StoreFront Management Console."

10. Click Finish.

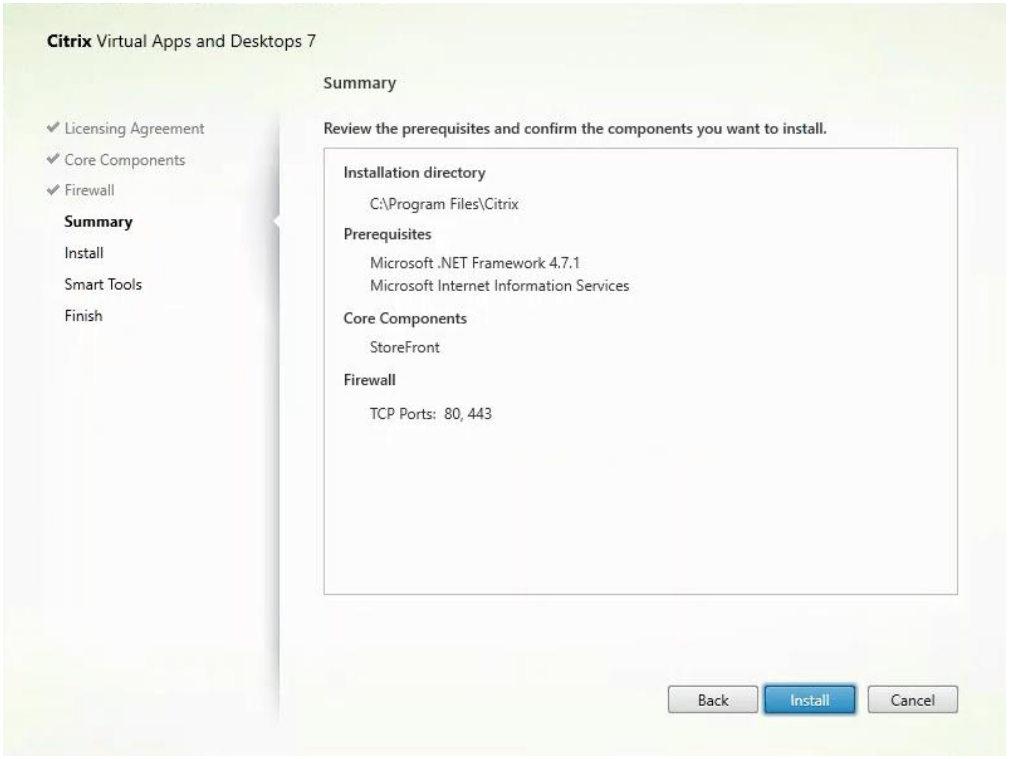


To configure the second StoreFront if used, follow these steps:

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

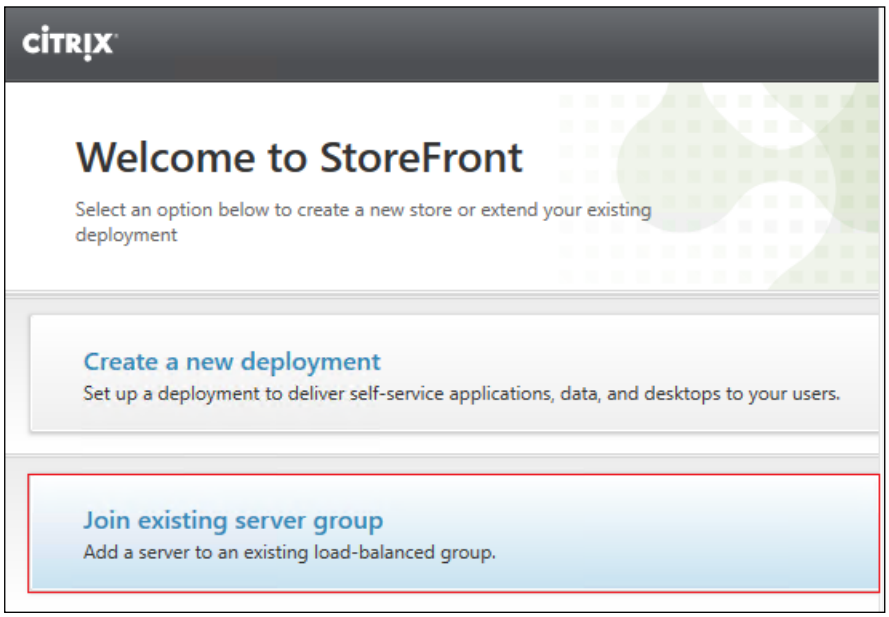


2. In the Join Server Group dialog, enter the name of the first Storefront server.

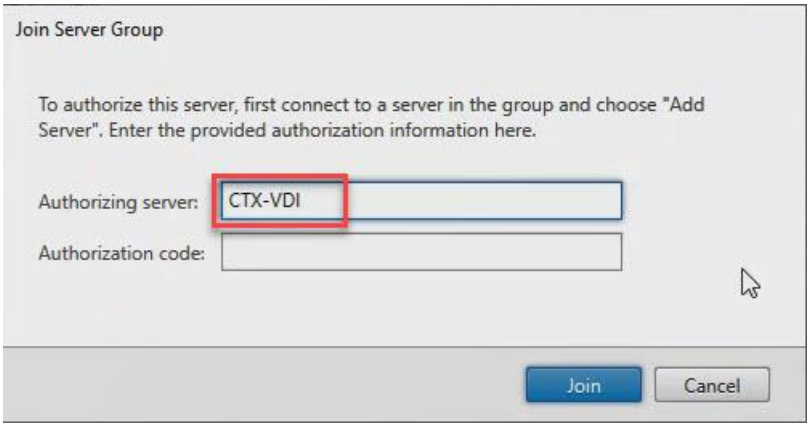


3. Before the additional StoreFront server can join the server group, you must connect to the first Storefront server, add the second server, and obtain the required authorization information.

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.

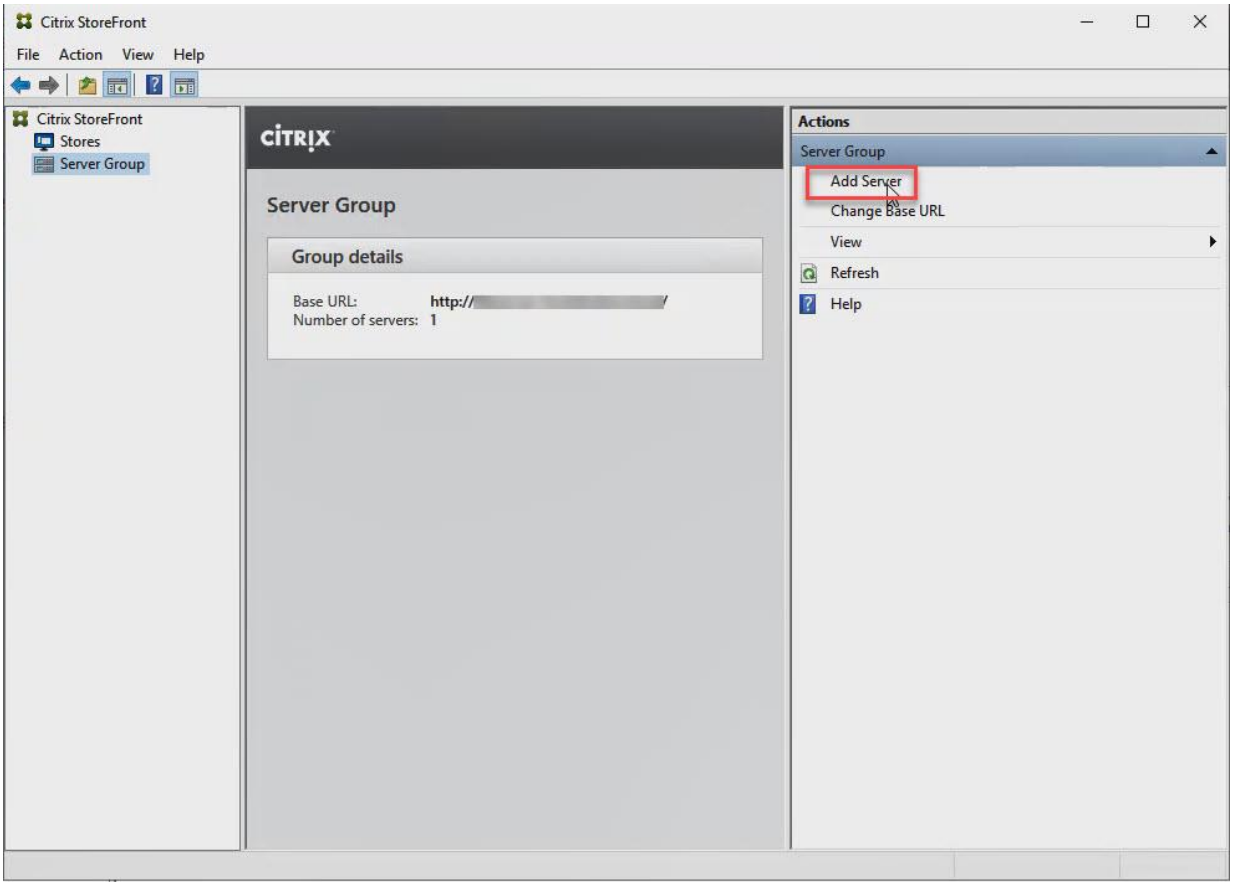


8. Copy the Authorization code from the Add Server dialog.

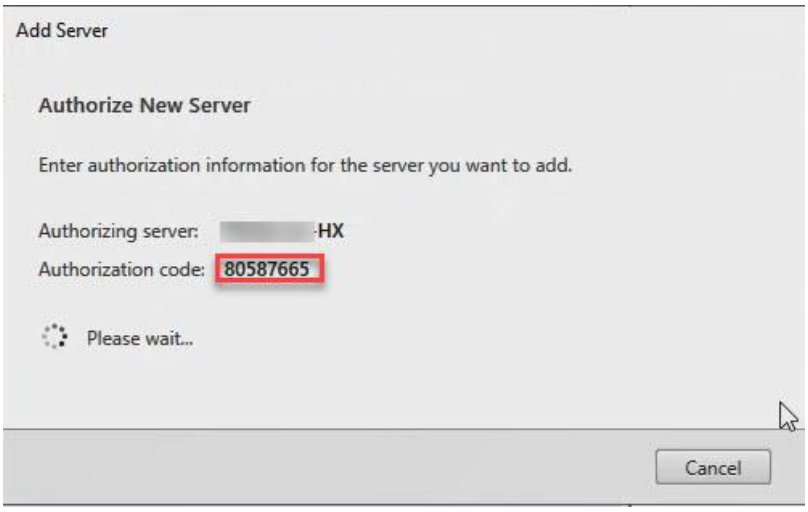


9. Connect to the second Storefront server and paste the Authorization code into the Join Server Group dialog.

10. Click Join.

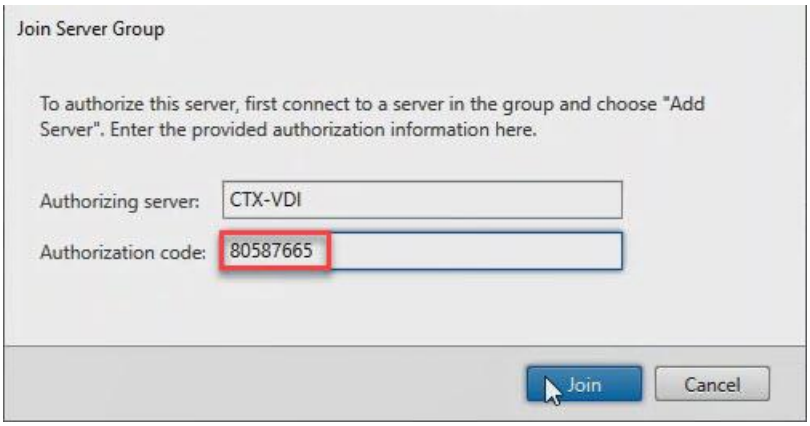


11. A message appears when the second server has joined successfully.

12. Click OK.

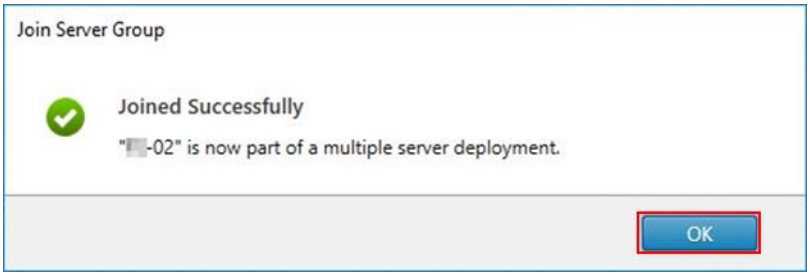


The second StoreFront is now in the Server Group.

## Install the Citrix Provisioning Services Target Device Software

For non-persistent Windows 10 virtual desktops and Server 2016 RDS 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, follow these 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.

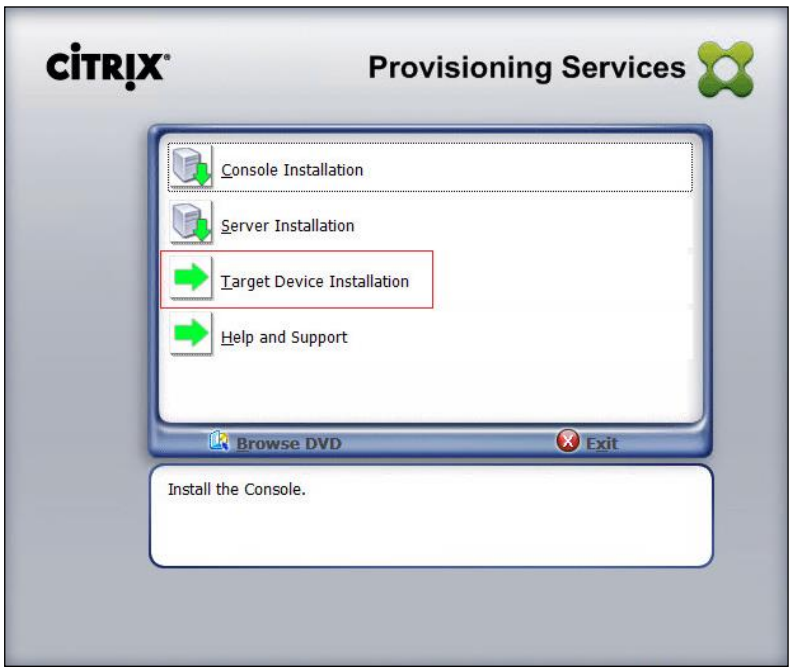


The installation wizard will check to resolve dependencies and then begin the PVS target device installation process.

3. Click Next.

4. Confirm the installation settings and click Install.

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

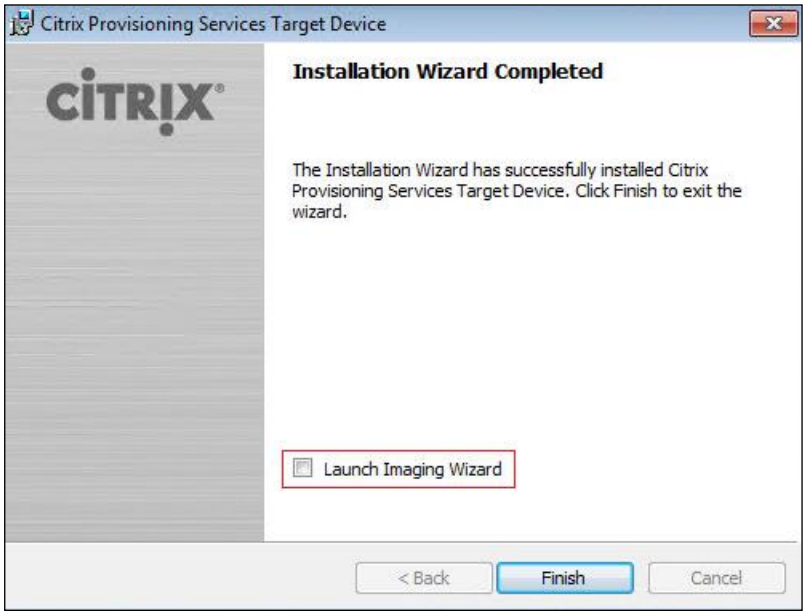


6. Reboot the machine.

## 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, follow these 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.

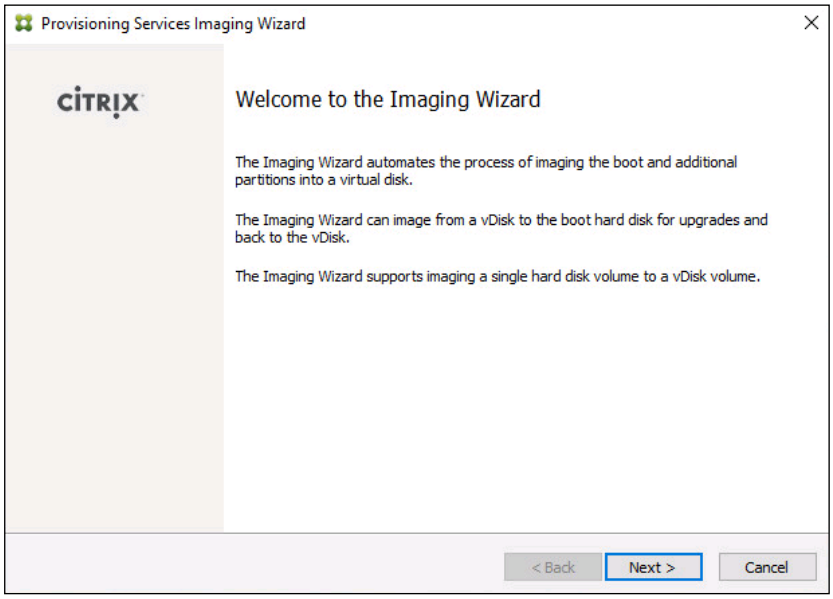

3. The Connect to Farm page appears. Enter the name or IP address of a Provisioning Server within the farm to connect to and the port to use to make that connection.

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

5. Click Next.

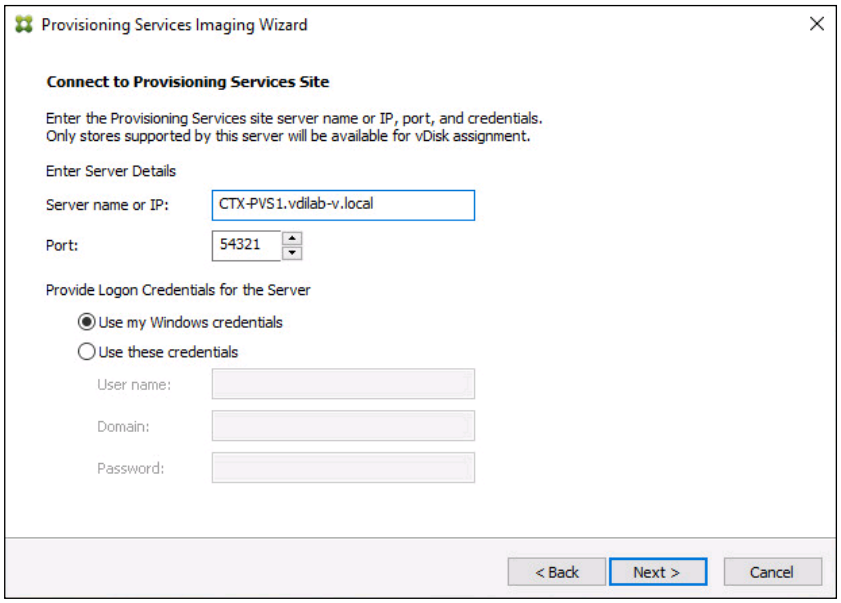


6. Select Create new vDisk.

7. Click Next.

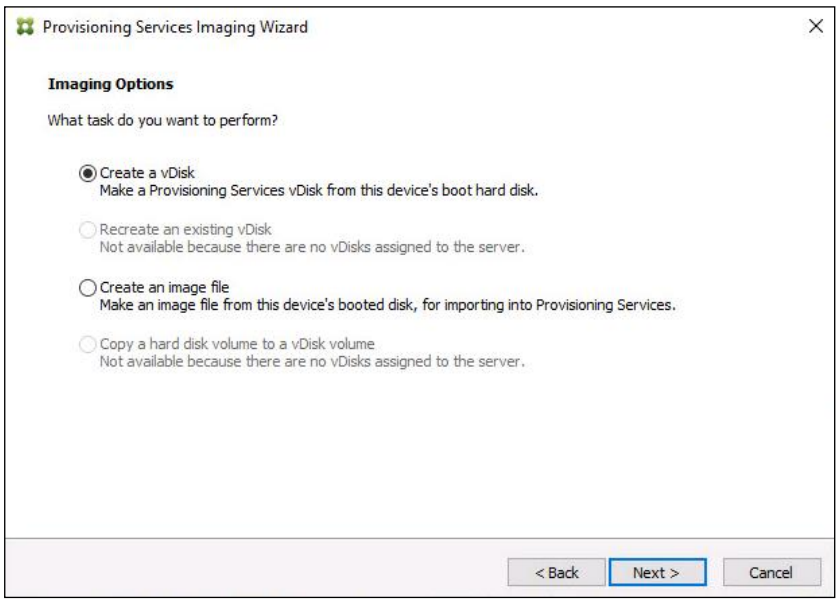


8. The Add Target Device page appears.

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.

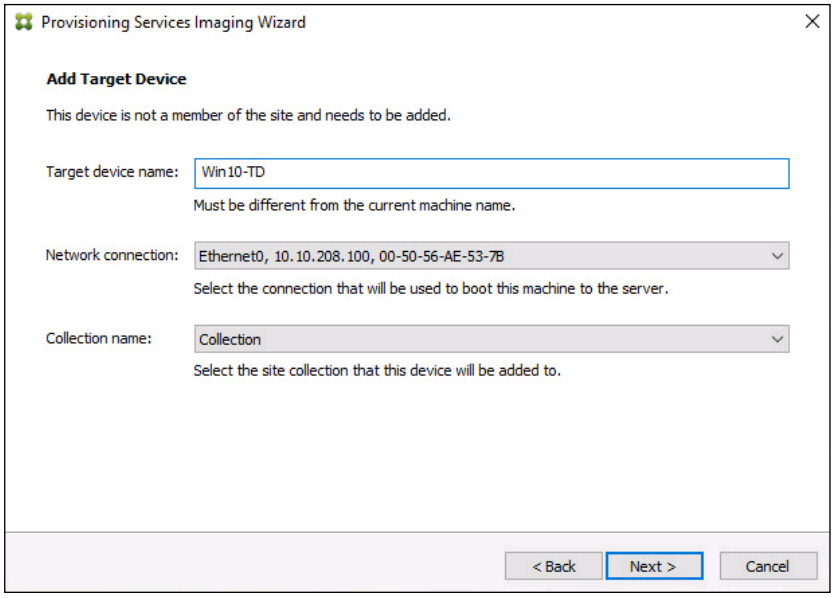


11. The New vDisk dialog displays. Enter the name of the vDisk.

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.

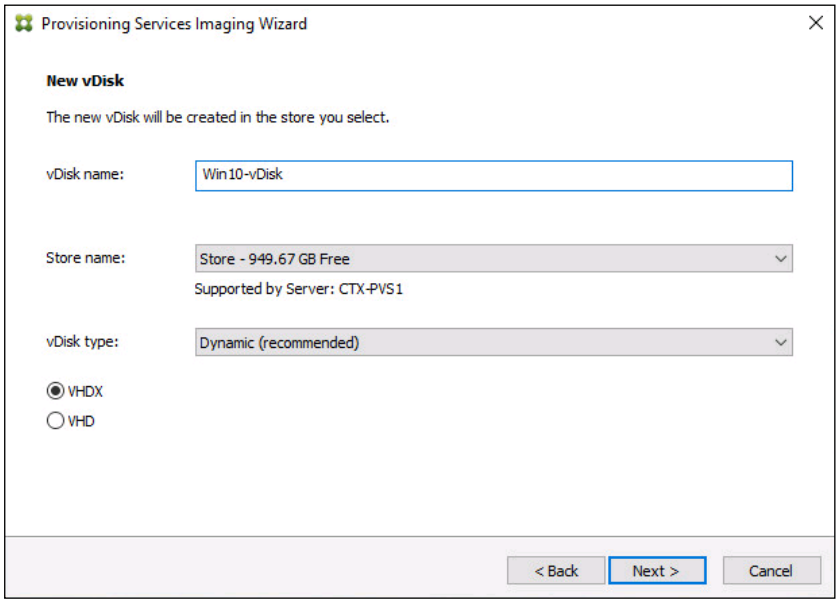

14. On the Microsoft Volume Licensing page, select the volume license option to use for target devices. For this CVD, volume licensing is not used, so the None button is selected.

15. Click Next.

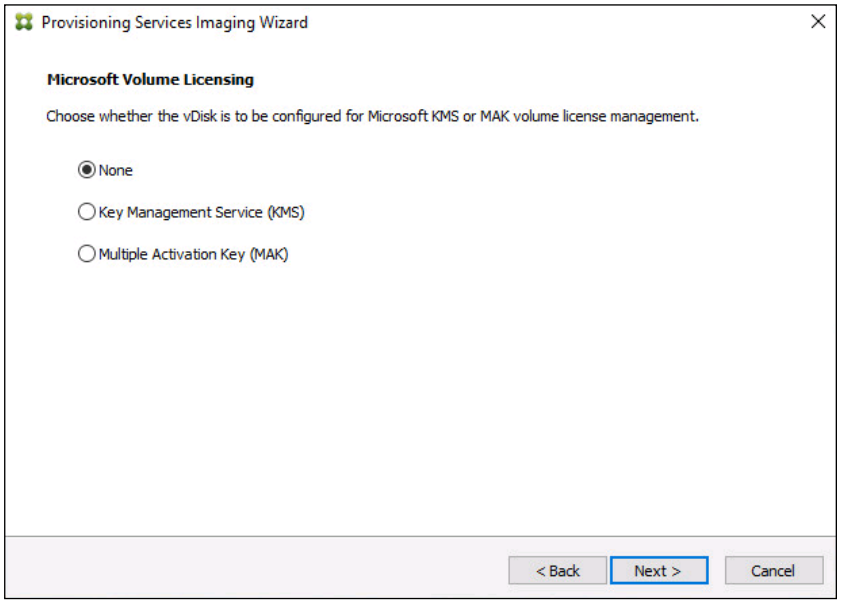


16. Select Image entire boot disk on the Configure Image Volumes page.

17. Click Next.

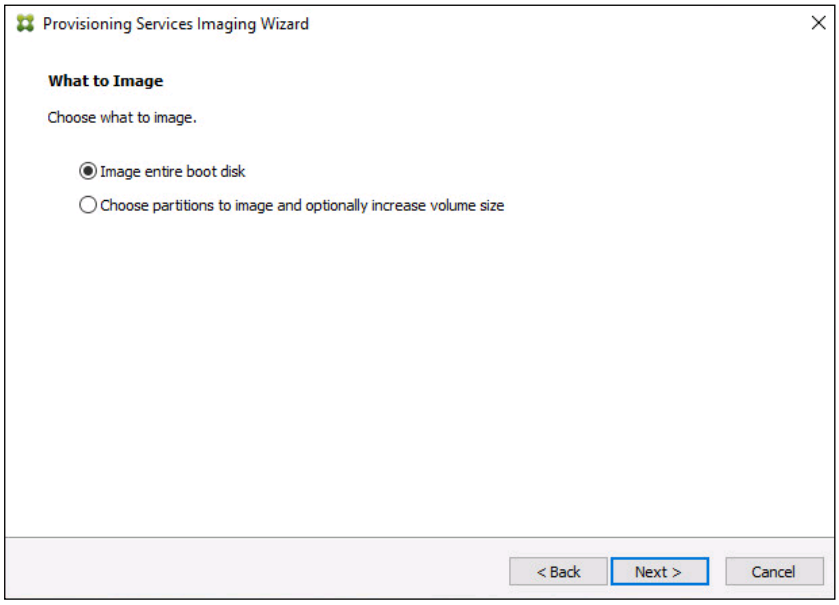


18. Select Optimize for hard disk again for Provisioning Services before imaging on the Optimize Hard Disk for Provisioning Services.

19. Click Next.

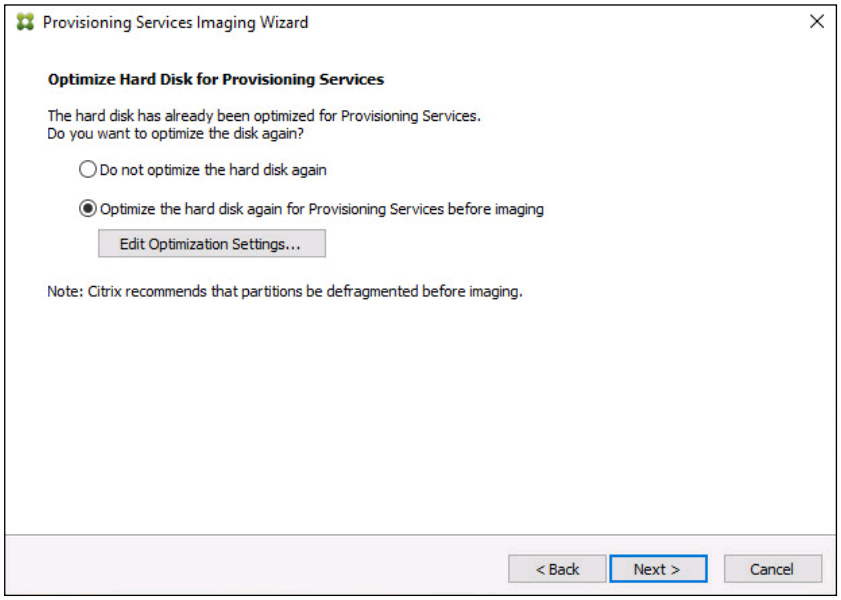


20. Select Create on the Summary page.

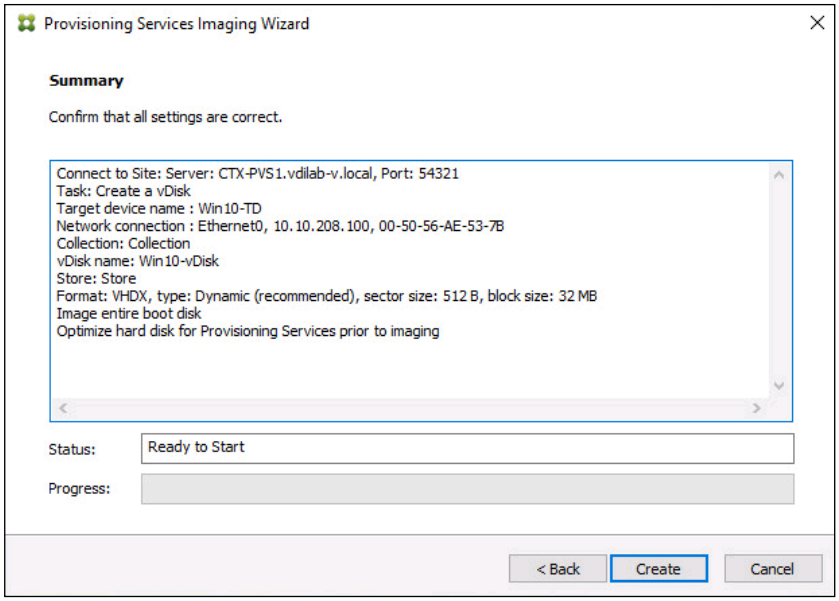


21. Review the configuration and click Continue.

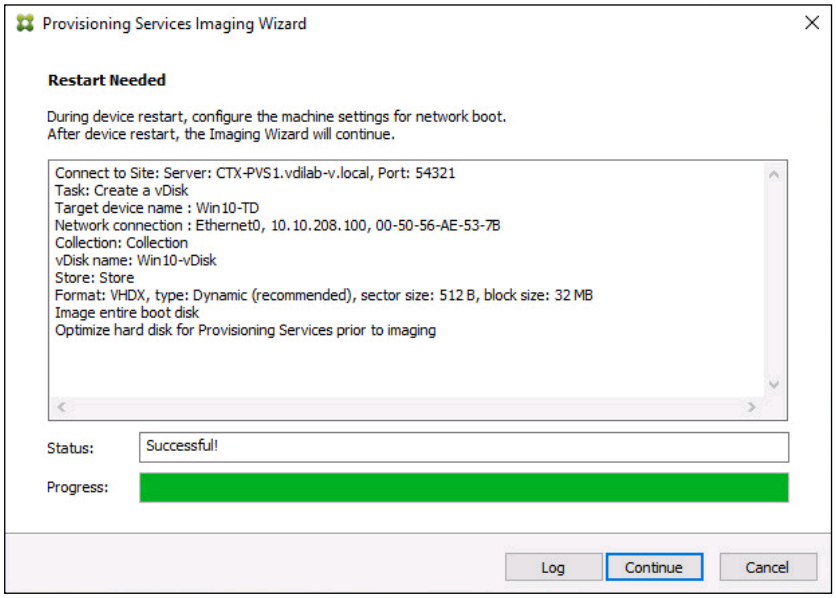


22. When prompted, click No to shut down the machine.

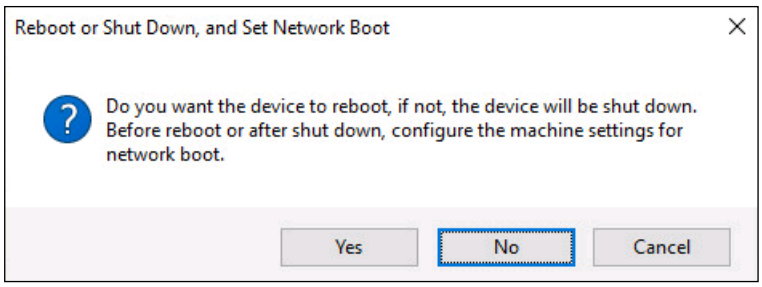


23. Edit the virtual machine settings and select Boot from Network Adapter under Boot Order.

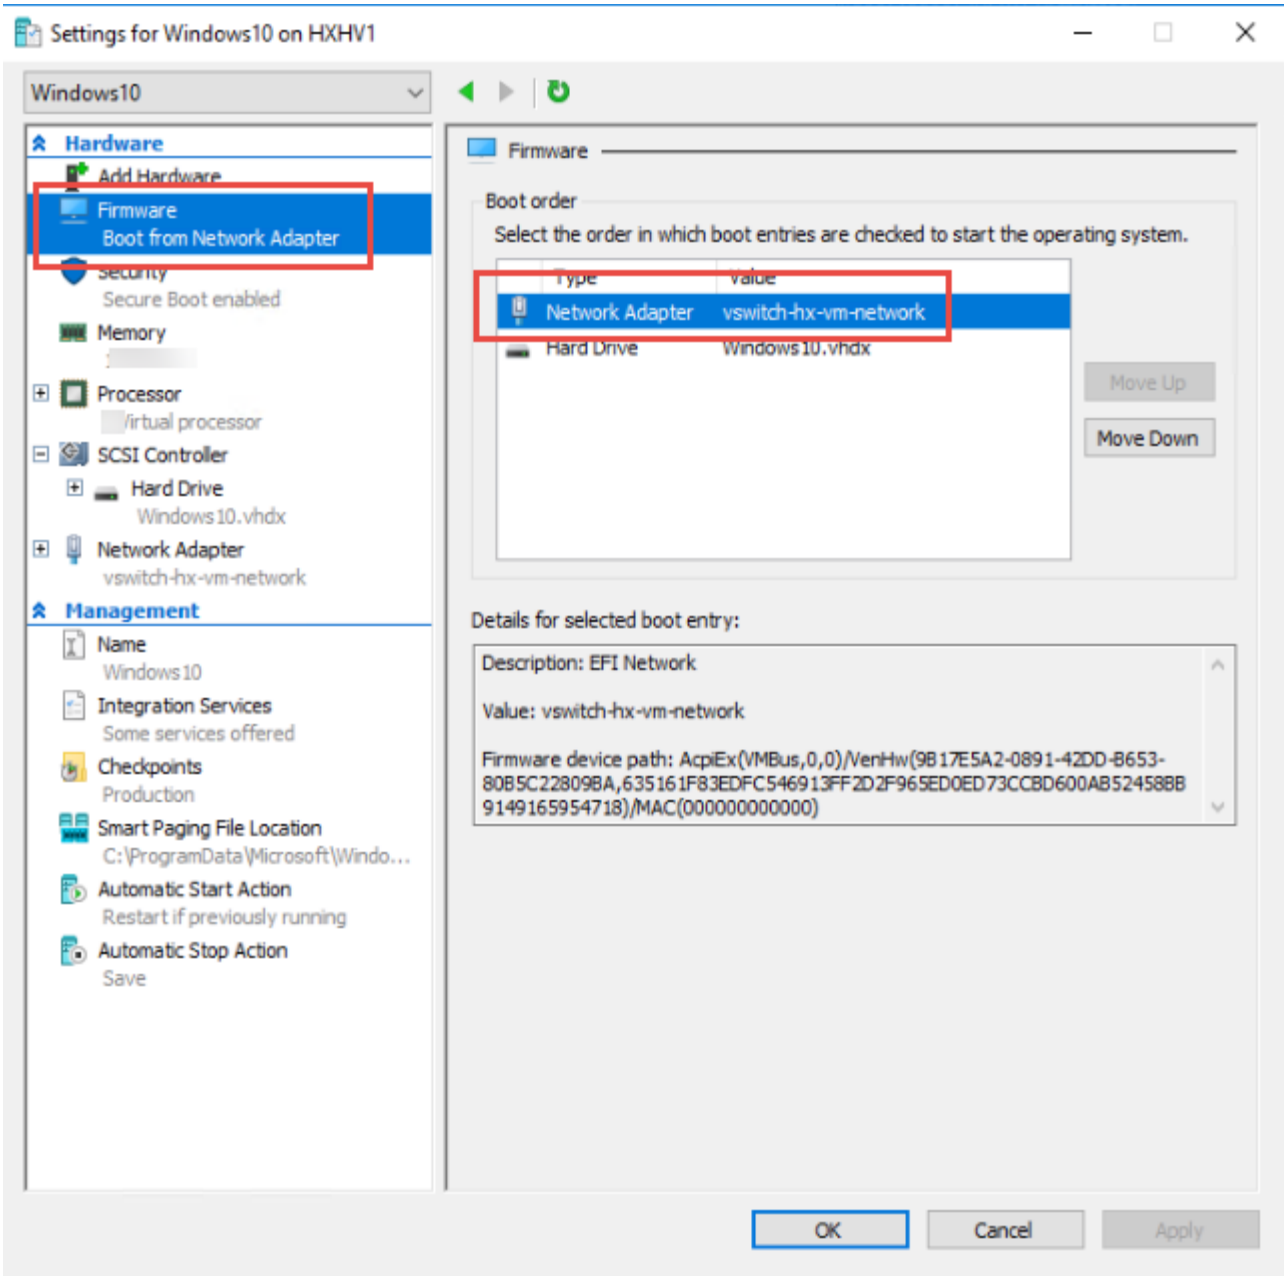


24. Restart Virtual Machine.

> **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.**

25. If prompted to Restart select Restart Later.

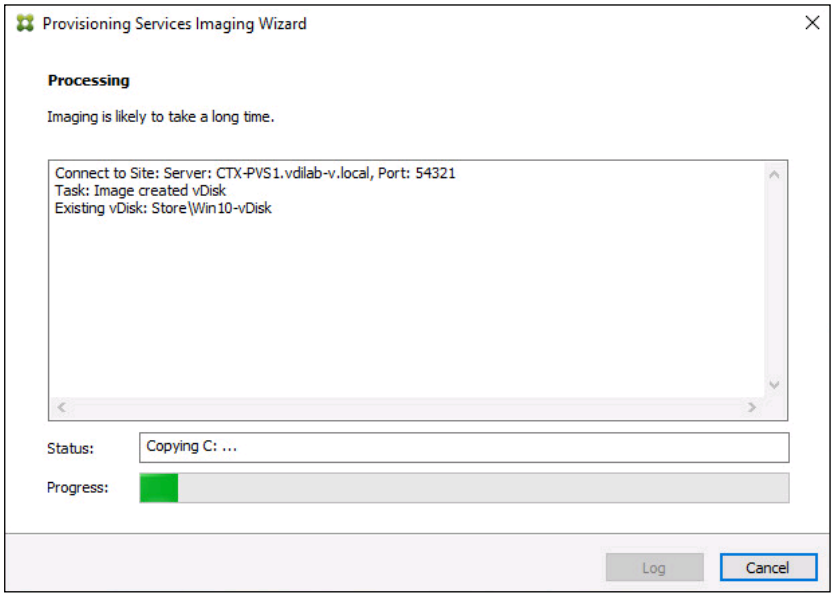


26. A message is displayed when the conversion is complete, click Done.

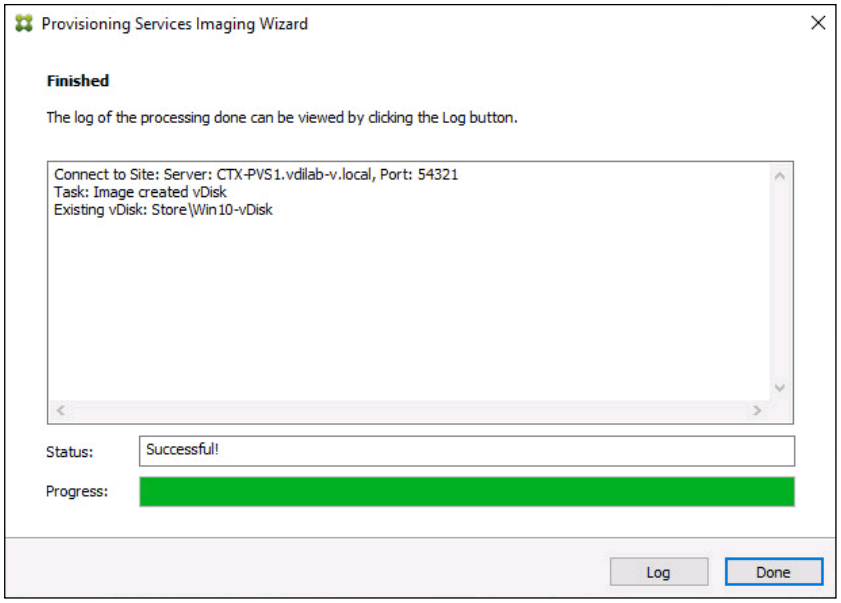


27. Shutdown the virtual machine used as the VDI or RDS master target.

28. Connect to the PVS server and validate that the vDisk image is available in the Store.

29. Right-click the newly created vDisk and select Properties.

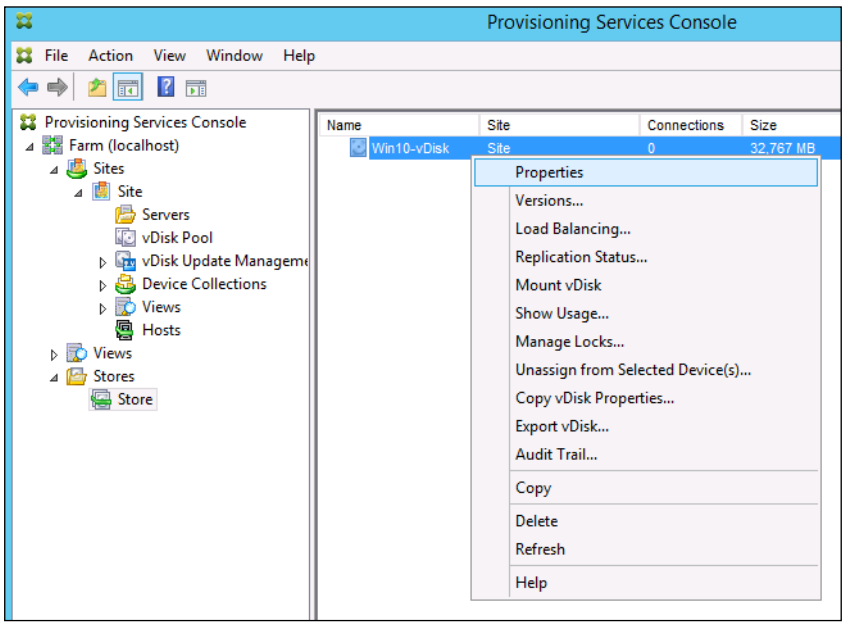


30. On the vDisk Properties dialog, change Access mode to “Standard Image (multi-device, read-only access)”.

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

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

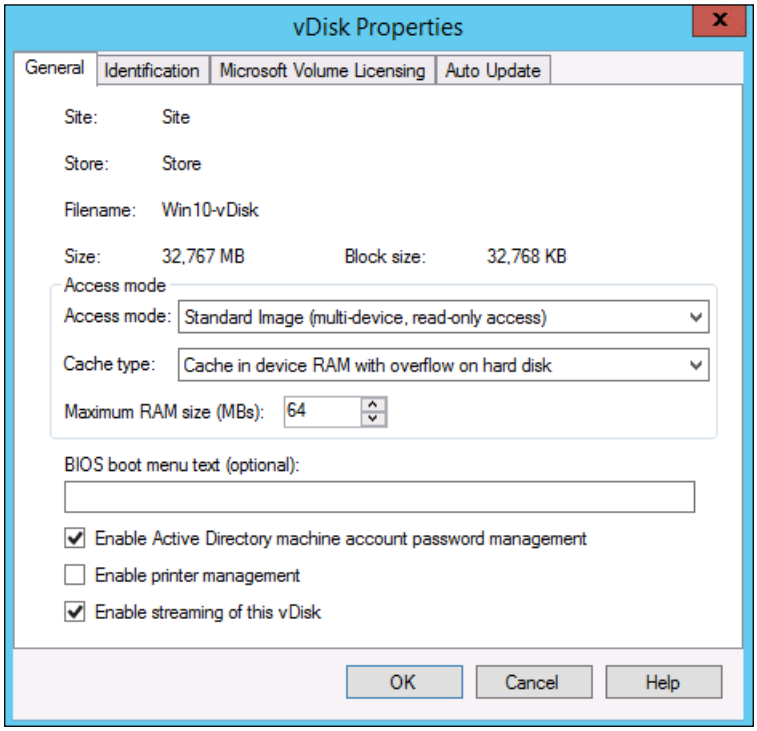


33. Click OK.

**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, follow these steps:

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

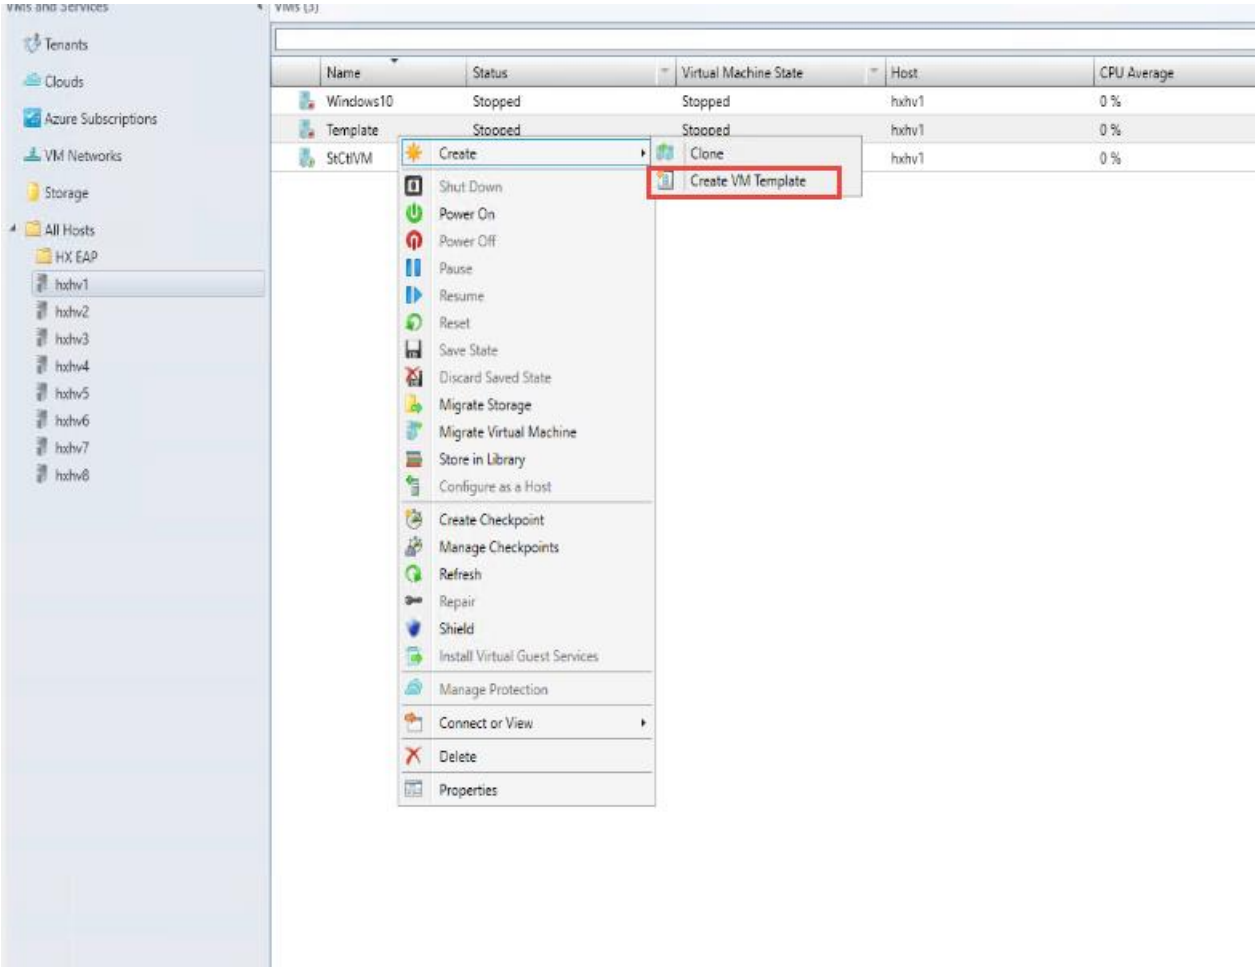


5. Click Yes to the warning that your machine will be destroyed by the Template process.

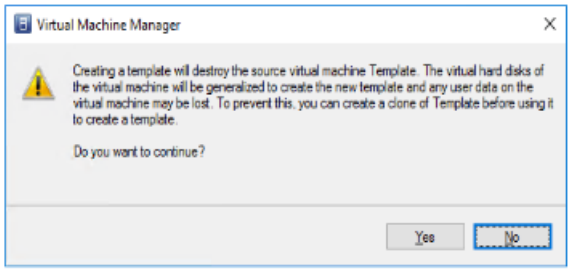
Virtual Machine Manager
Creating a template will destroy the source virtual machine Template. The virtual hard disks of the virtual machine will be generalized to create the new template and any user data on the virtual machine may be lost. To prevent this, you can create a clone of Template before using it to create a template.
Do you want to continue?
Yes
No

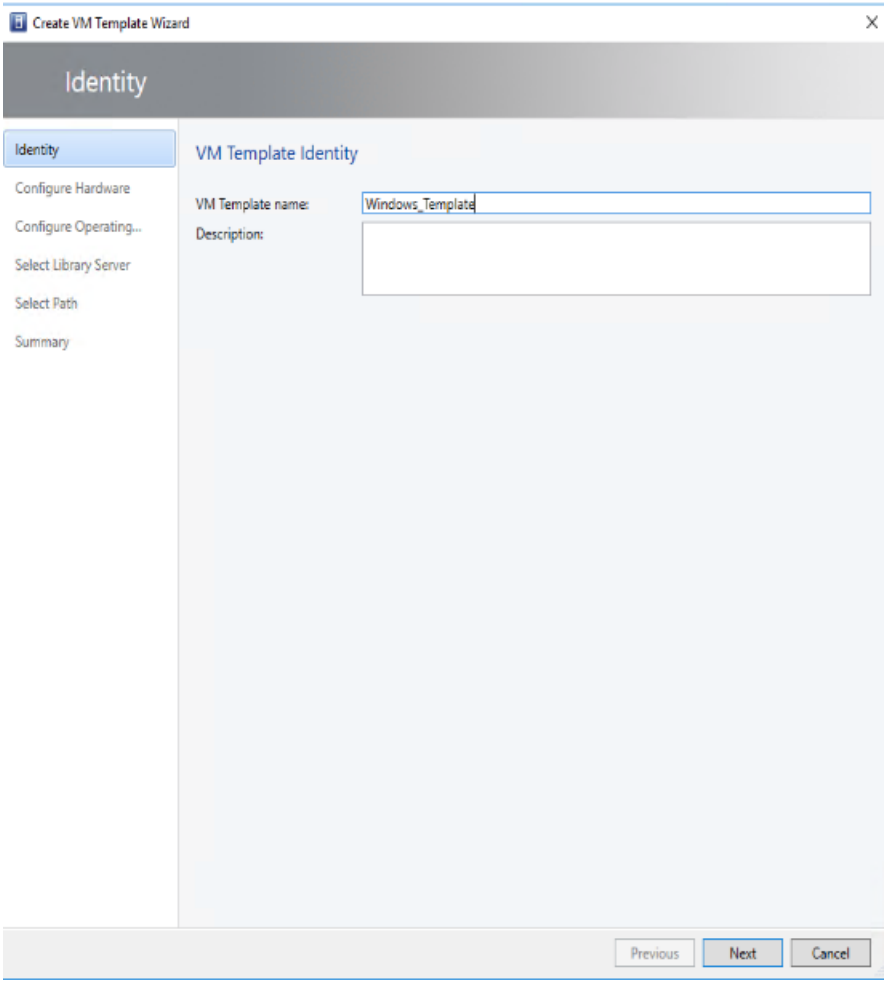
Create VM Template Wizard
Identity
Identity
Configure Hardware
Configure Operating...
Select Library Server
Select Path
Summary
VM Template Identity
VM Template name:
Windows_Template
Description:
Previous
Next
Cancel

6. Click Next.

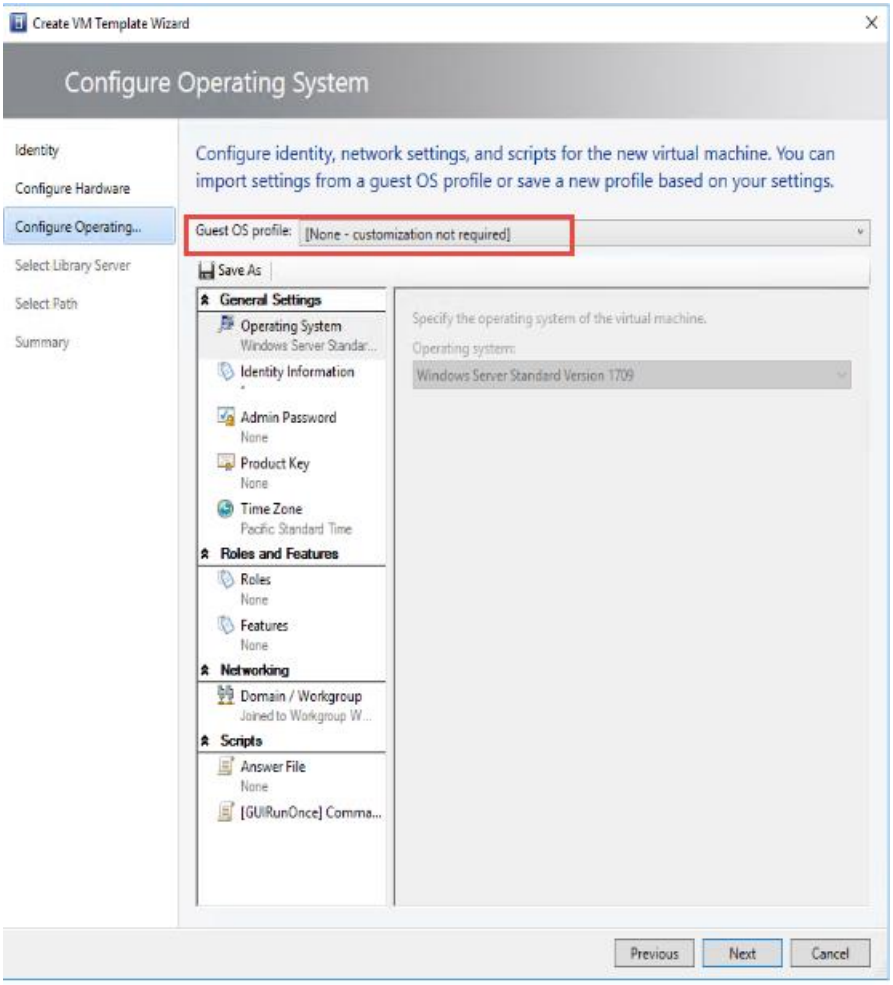


7. Under ‘Configure Operating System’, ensure that customization is disabled.
8. Click Next through the remaining screens
9. Click Finish to create the template.
10. From Citrix Studio on the Desktop Controller, select Hosting and Add Connection and Resources.
11. Select Use an existing Connection and click Next.
12. Correspond the name of the resource with desktop machine clusters.

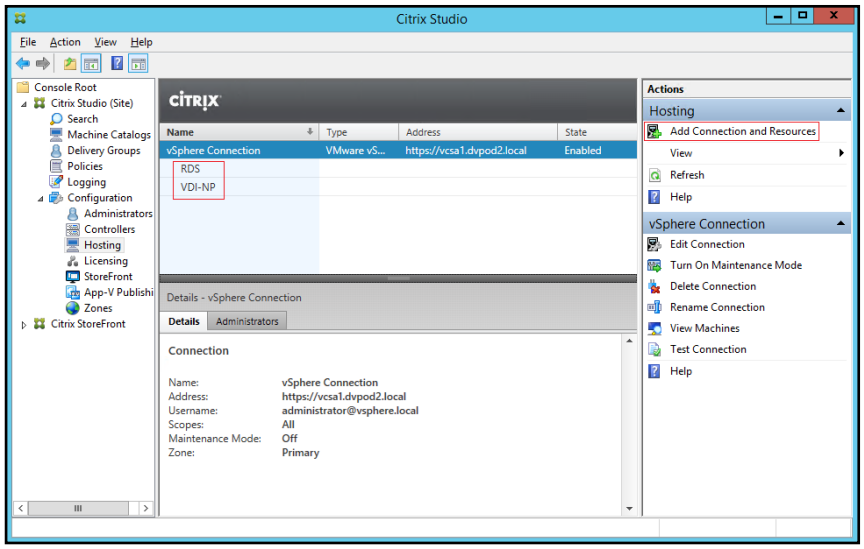


13. Browse and select the Hyper-V cluster for desktop provisioning and use the default storage method Use storage shared by hypervisors.

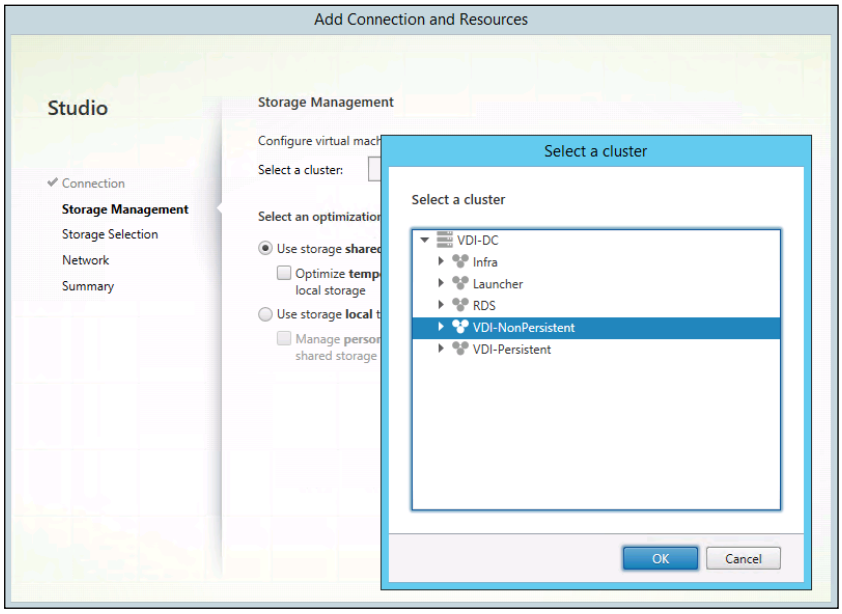


14. Select the data storage location for the corresponding resource.

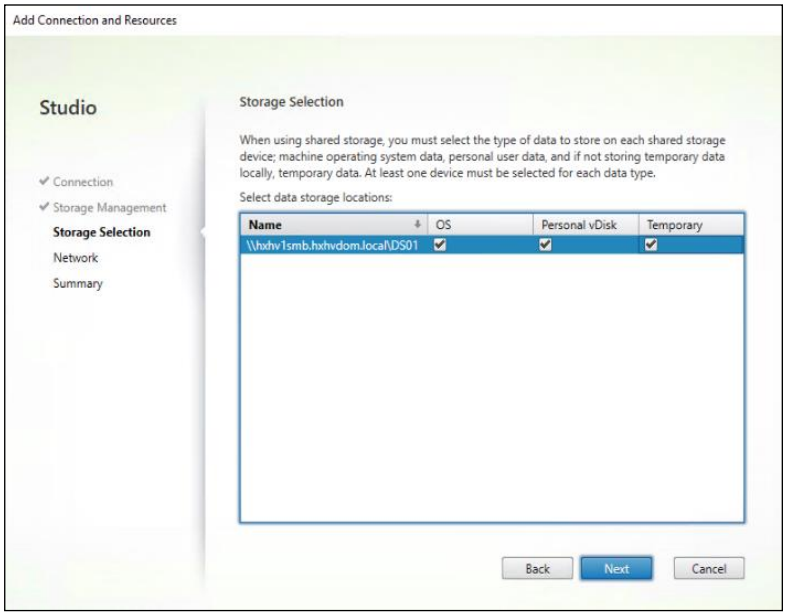


15. Select the VDI networks for the desktop machines and click Next.

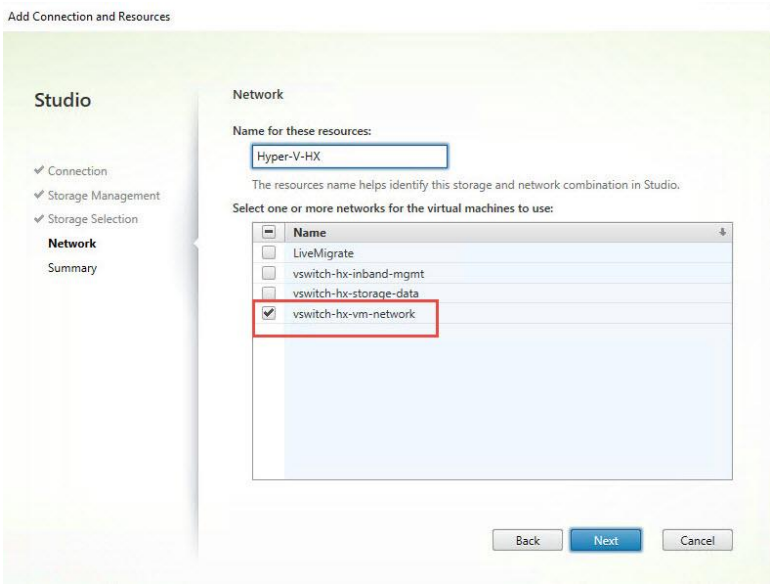

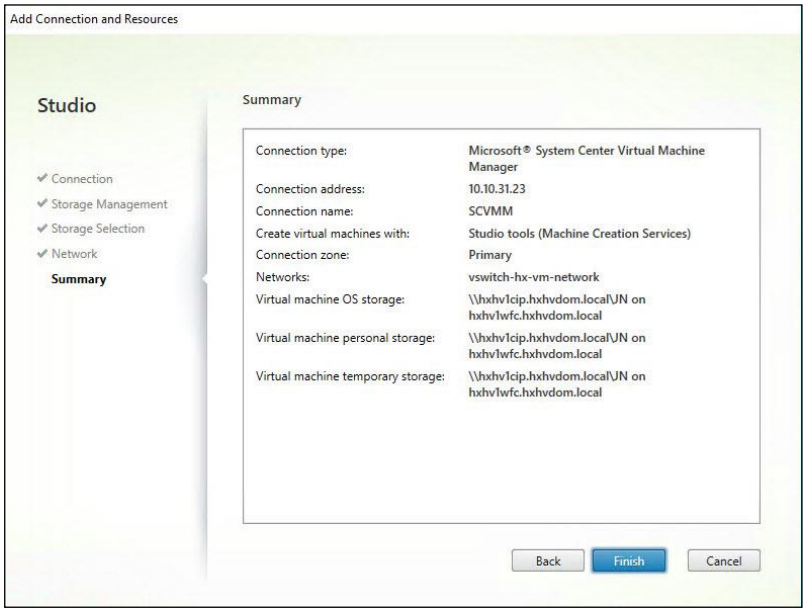


16. Click Finish.

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

## Provision Desktop Machines from Citrix Provisioning Services Console

To provision the desktop machines using the Citrix Provisioning Service Console, follow these steps:

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

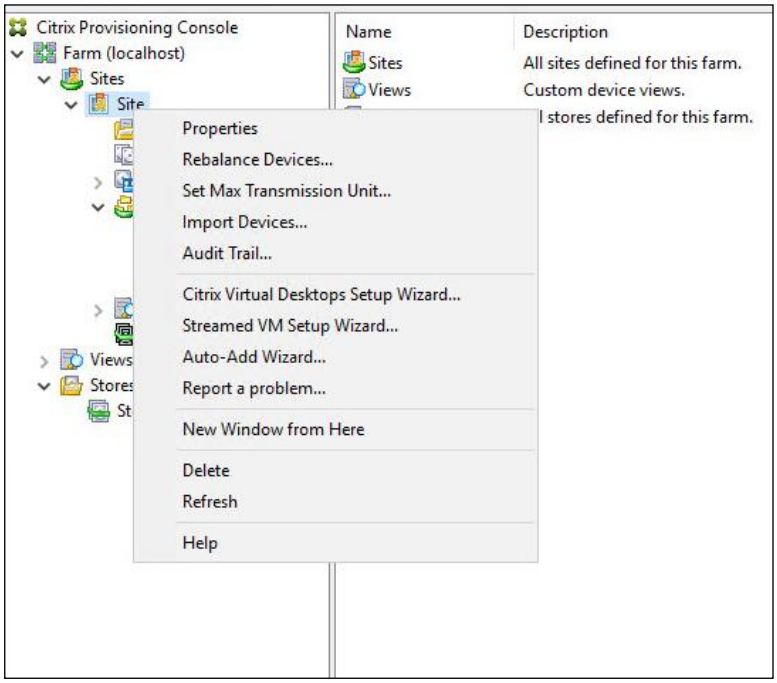


4. Click Next.

5. Enter the Virtual Desktops 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.

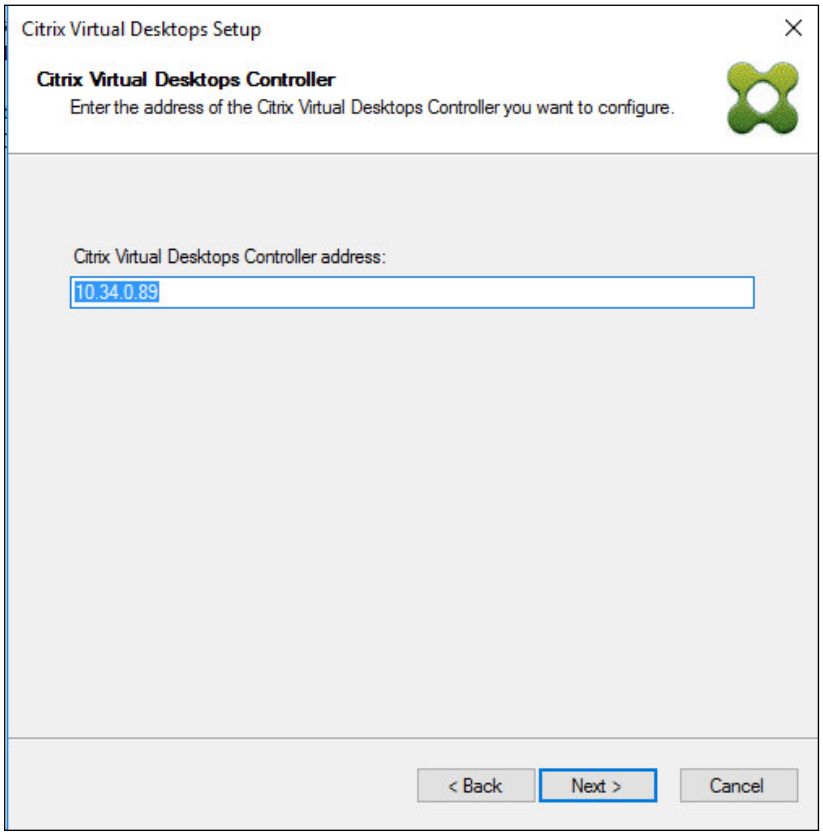


9. Provide the Host Resources Credentials (Username and Password) to the Virtual Desktops controller when prompted.
10. Click OK.

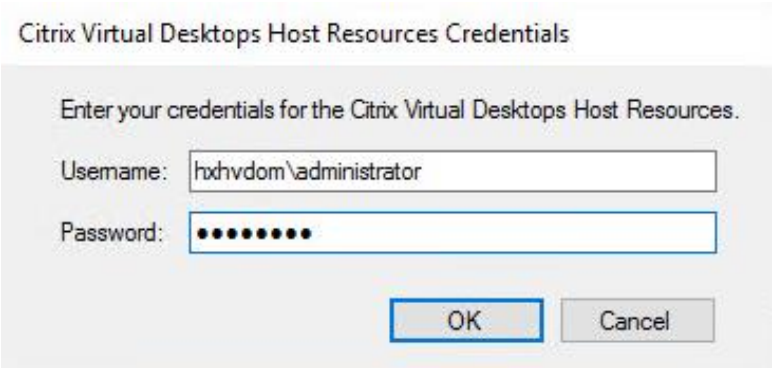


11. Select the Template created earlier.
12. Click Next.

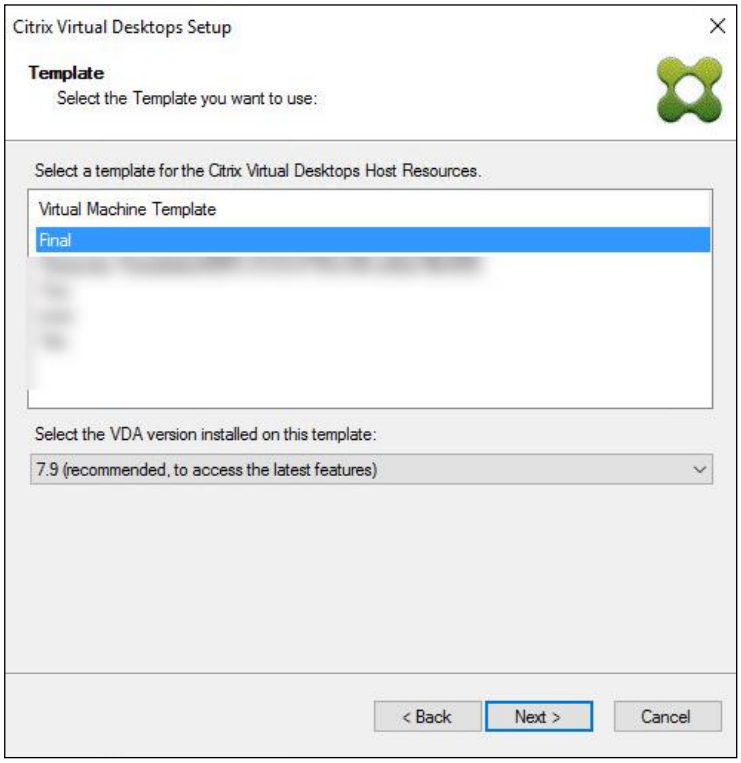


13. Select the vDisk that will be used to stream virtual machines.

14. Click Next.

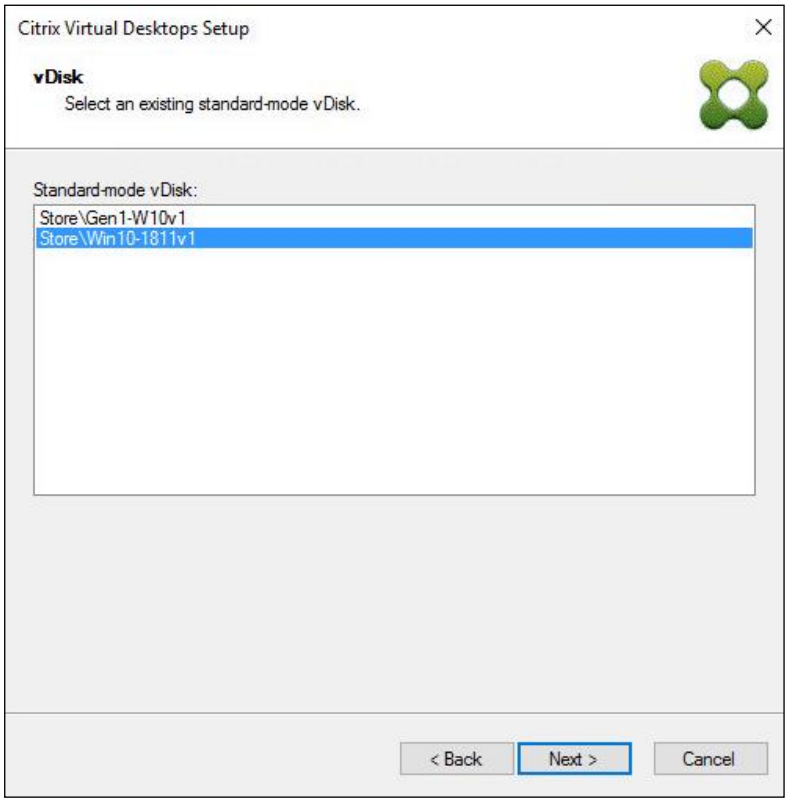

15. Select “Create a new catalog.”

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

16. Click Next.

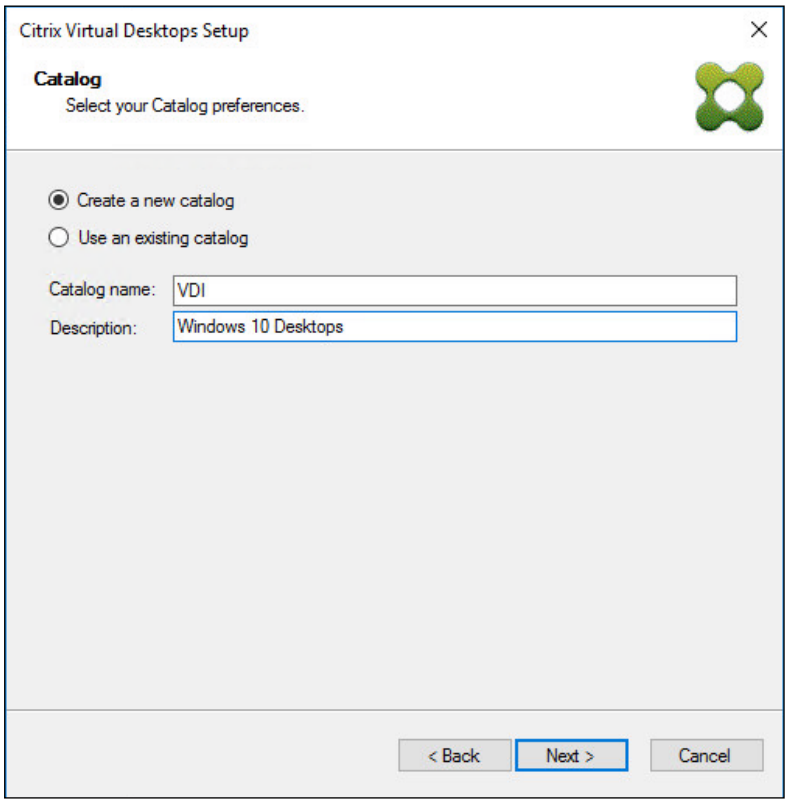


17. On the Operating System dialog, specify the operating system for the catalog. Specify Windows Desktop Operating System for VDI and Windows Server Operating System for RDS.

18. Click Next.

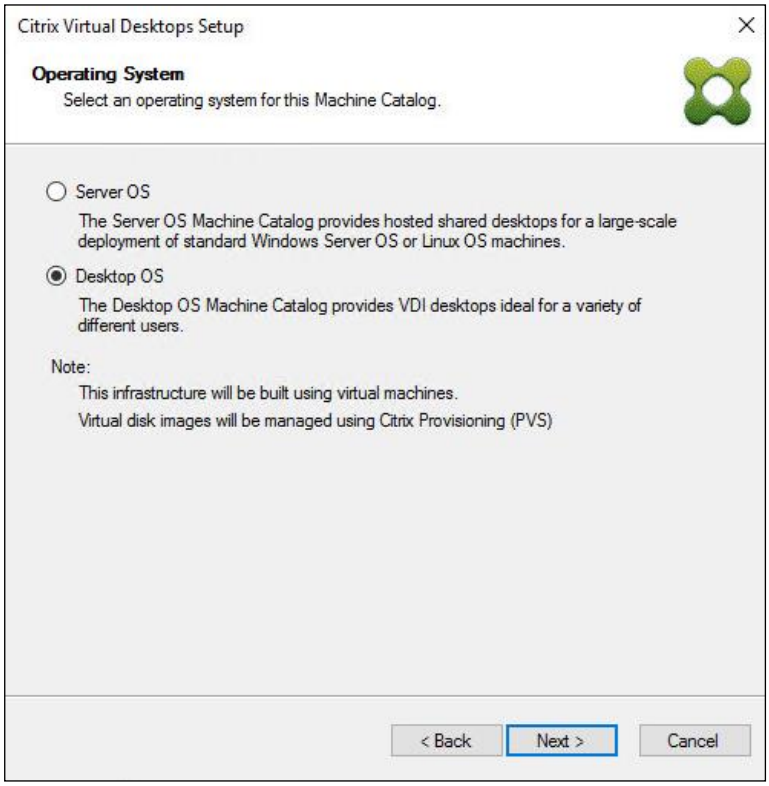


19. If you specified a Windows Desktop OS for VDIs, a User Experience dialog appears. Specify that the user will connect to “A fresh new (random) desktop each time.”

20. Click Next.

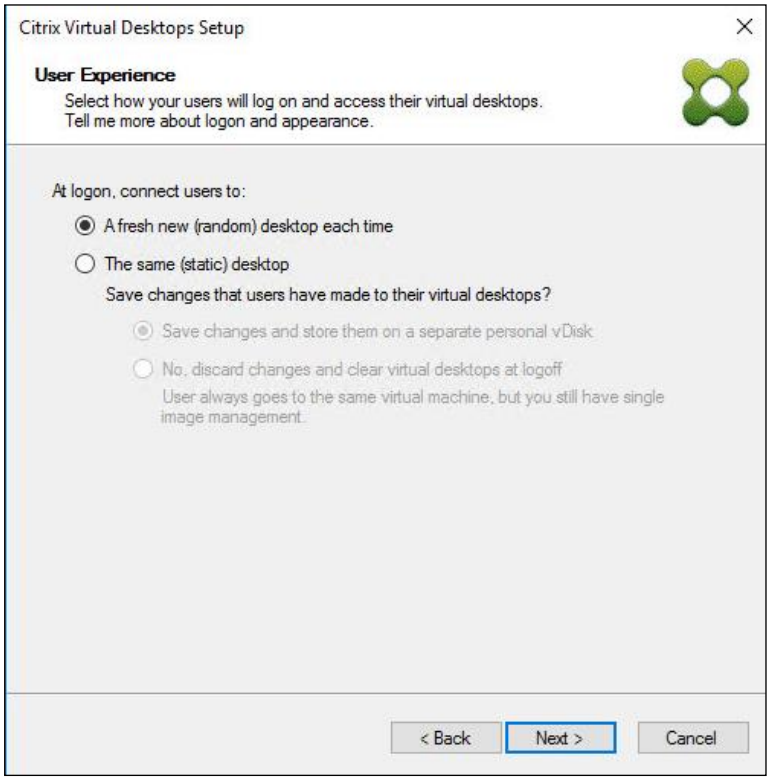


21. On the Virtual machines dialog, specify:
    a. The number of virtual machines to create.
    b. Number of vCPUs for the virtual machine (2 for VDI, 8 for RDS).
    c. The amount of memory for the virtual machine (4GB for VDI, 24GB for RDS).
    d. The write-cache disk size (10GB for VDI, 30GB for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

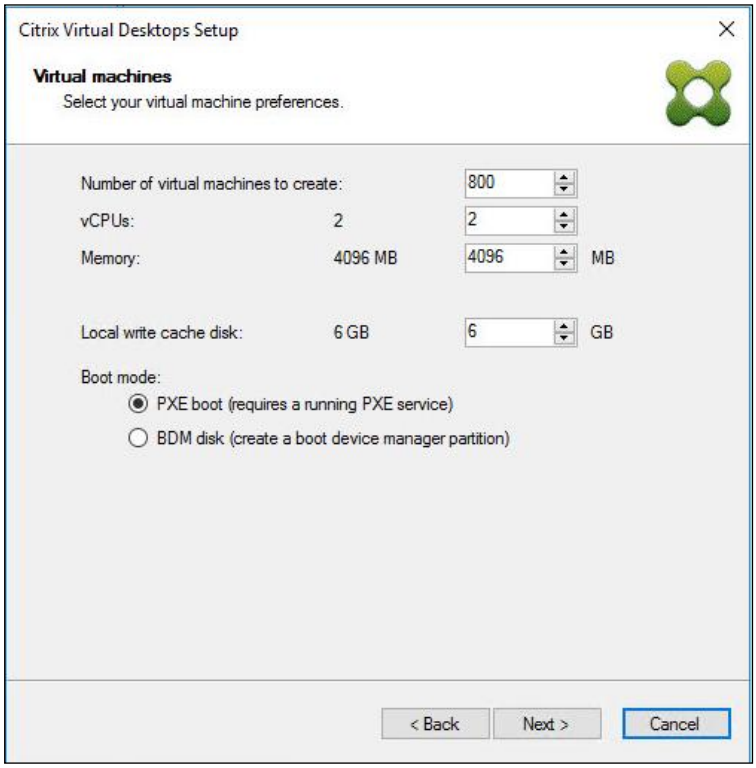


23. Select the Create new accounts radio button.

24. Click Next.

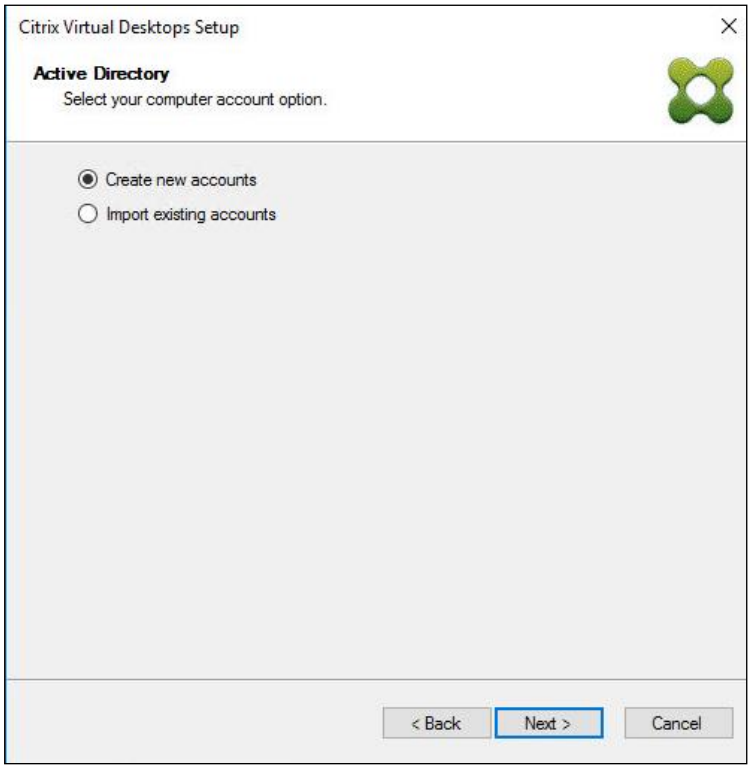

25. Specify the Active Directory Accounts and Location. This is where the wizard should create the computer accounts.

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

27. Click Next.

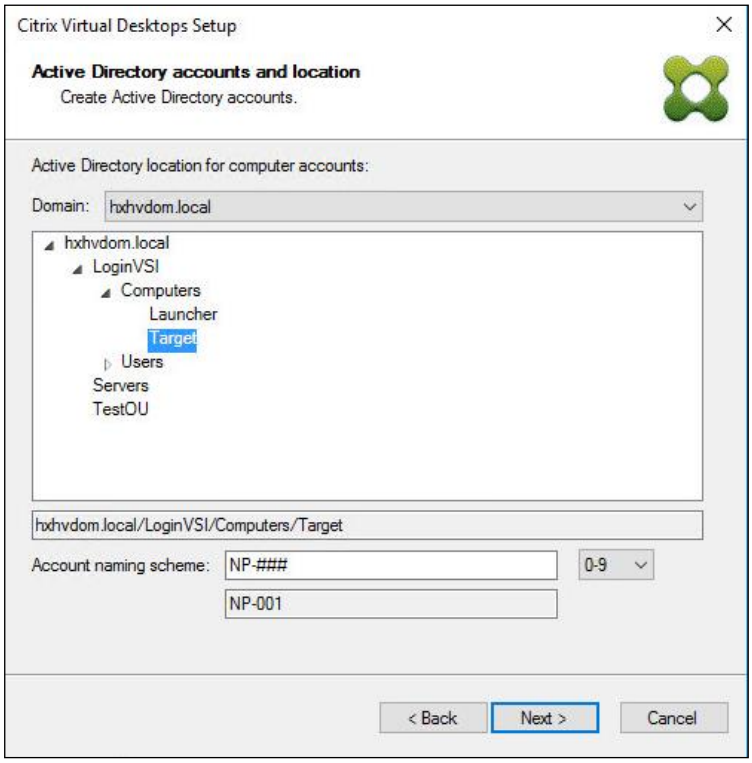


28. Click Finish to begin the virtual machine creation.

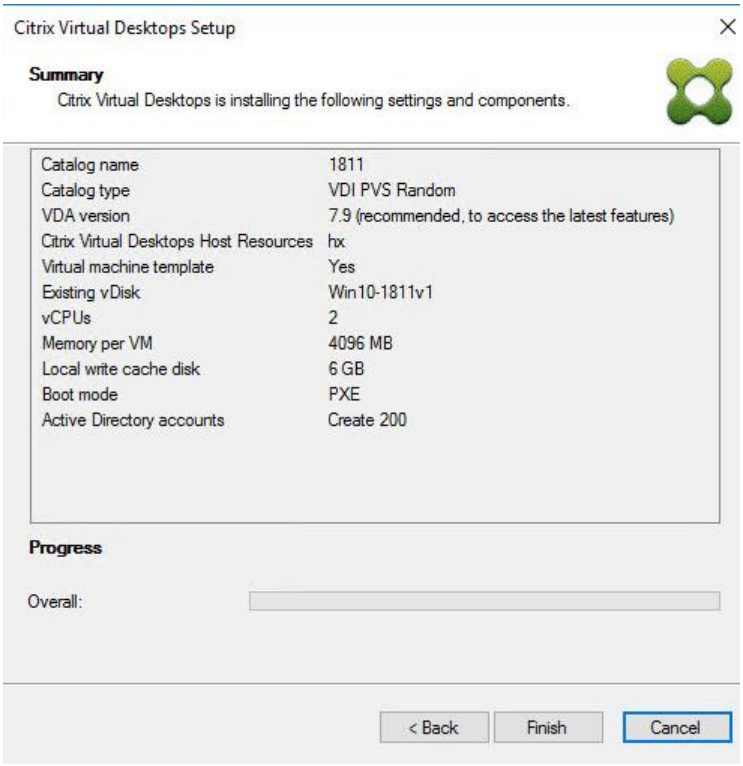


29. When the wizard is done provisioning the virtual machines, click Done.

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

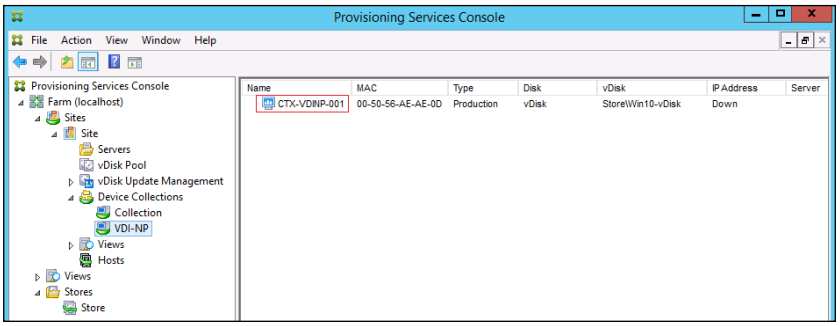


    b. CTX-XD1 > Citrix Studio > Machine Catalogs > VDI-NP

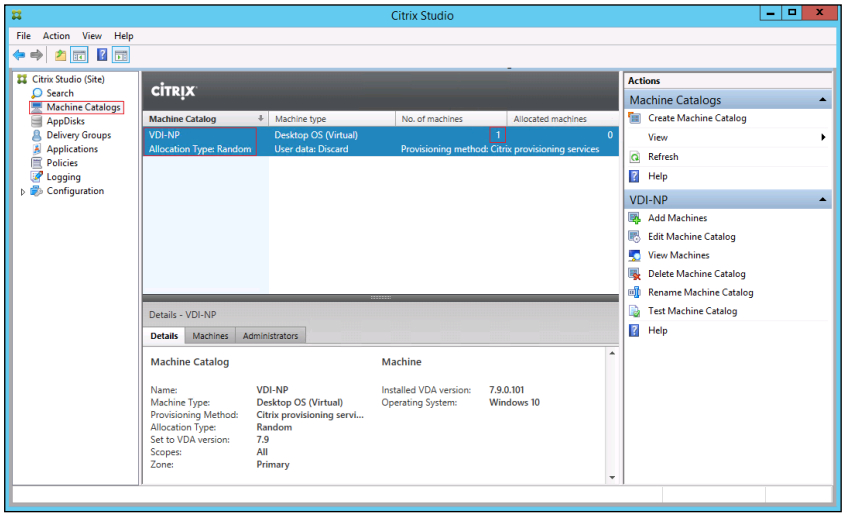


c. AD-DC1 > Active Directory Users and Computers > hxhvdom.local > Computers > CTX-VDI-001

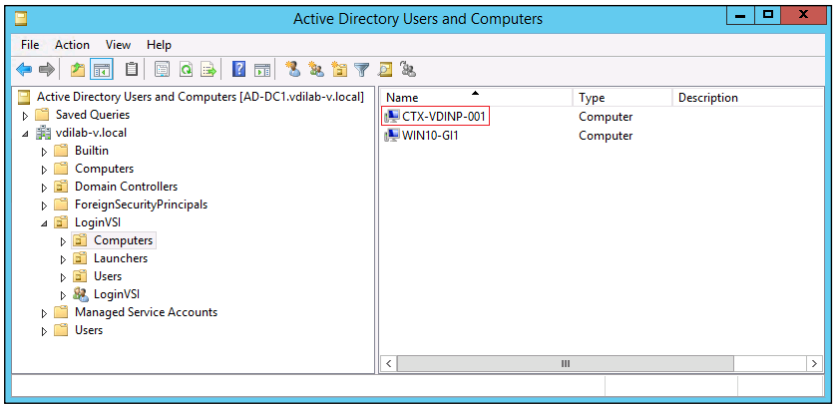


31. Log into the newly provisioned desktop machine, using the Virtual Disk Status verify the image mode is set to Ready Only and the cache type as Device Ram with overflow on local hard drive.

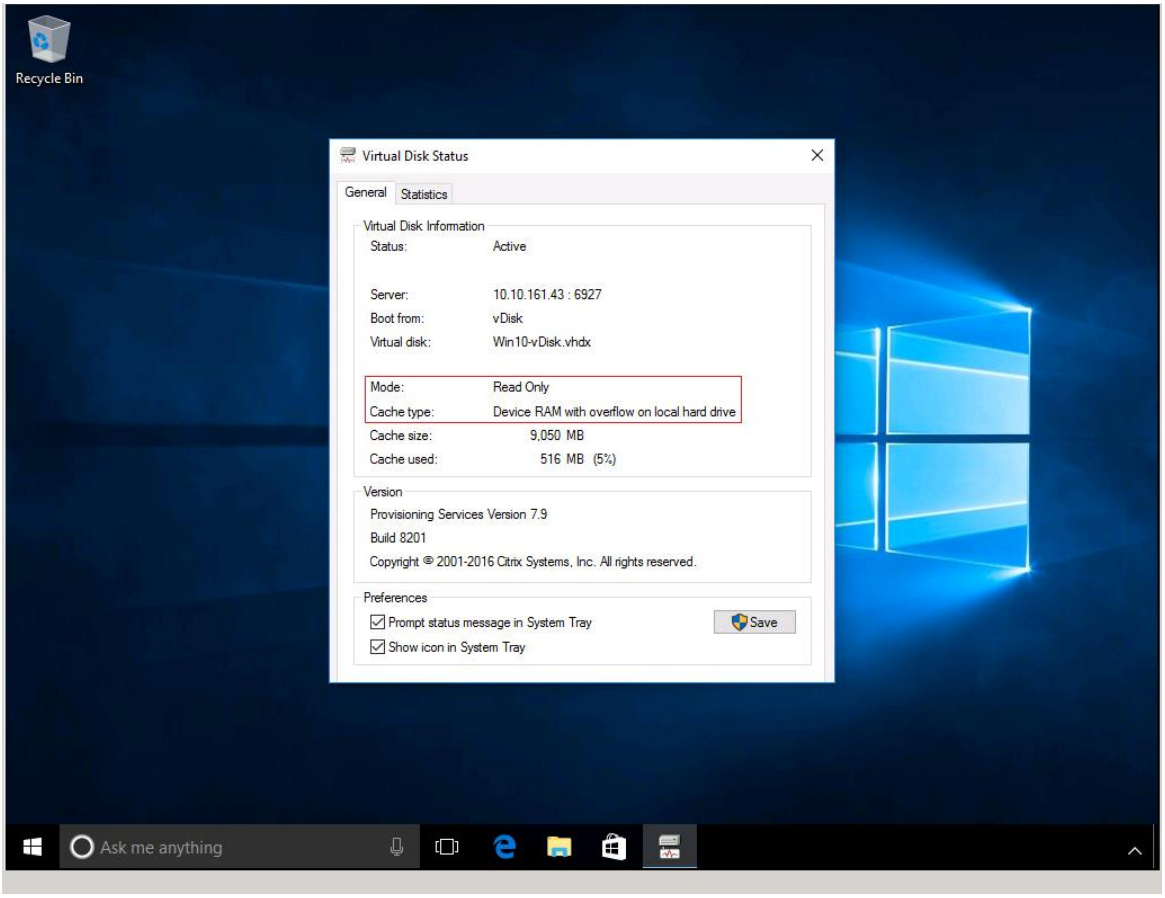


## Install Citrix Virtual Apps and Desktop Virtual Desktop Agents

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, follow these steps:

1. Launch the Citrix Desktop installer from the CVA Desktop 1909 ISO.
2. Click Start on the Welcome Screen.

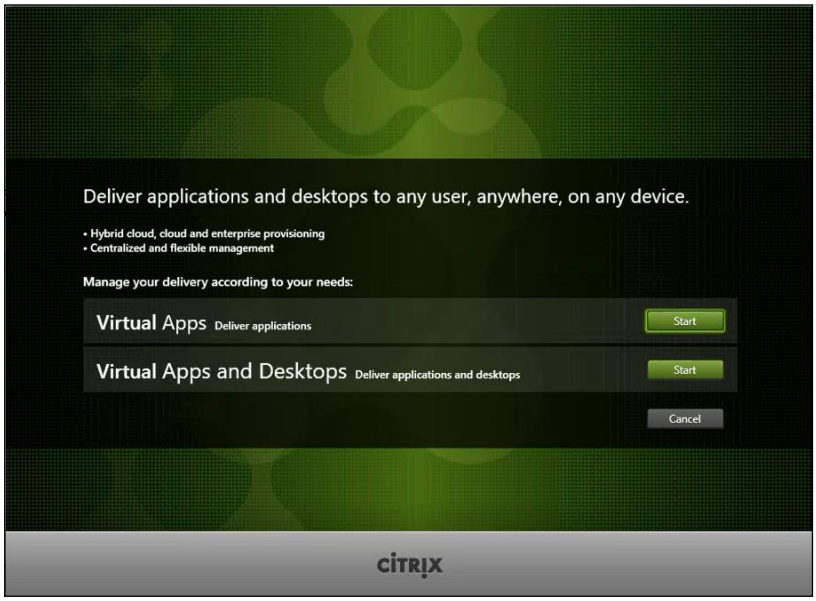


3. To install the VDA for the Hosted Virtual Desktops (VDI), select Virtual Delivery Agent for Windows Desktop OS. After the VDA is installed for Hosted Virtual Desktops, repeat the procedure to install the VDA for Hosted Shared Desktops (RDS). In this case, select Virtual Delivery Agent for Windows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

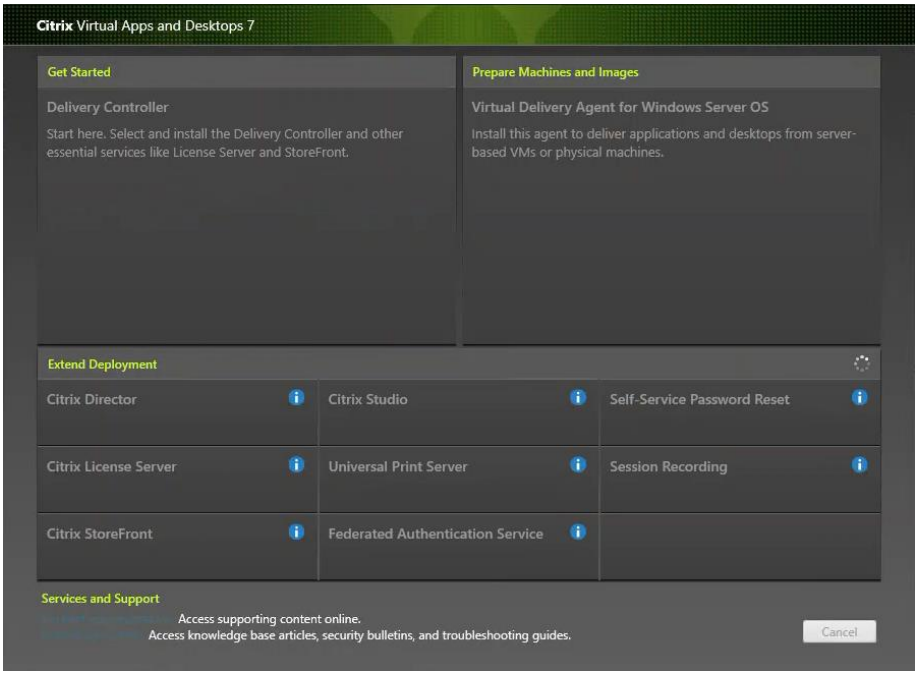


4. Select “Create a Master Image.”(Be sure to select the proper provisioning technology here)

5. Click Next.

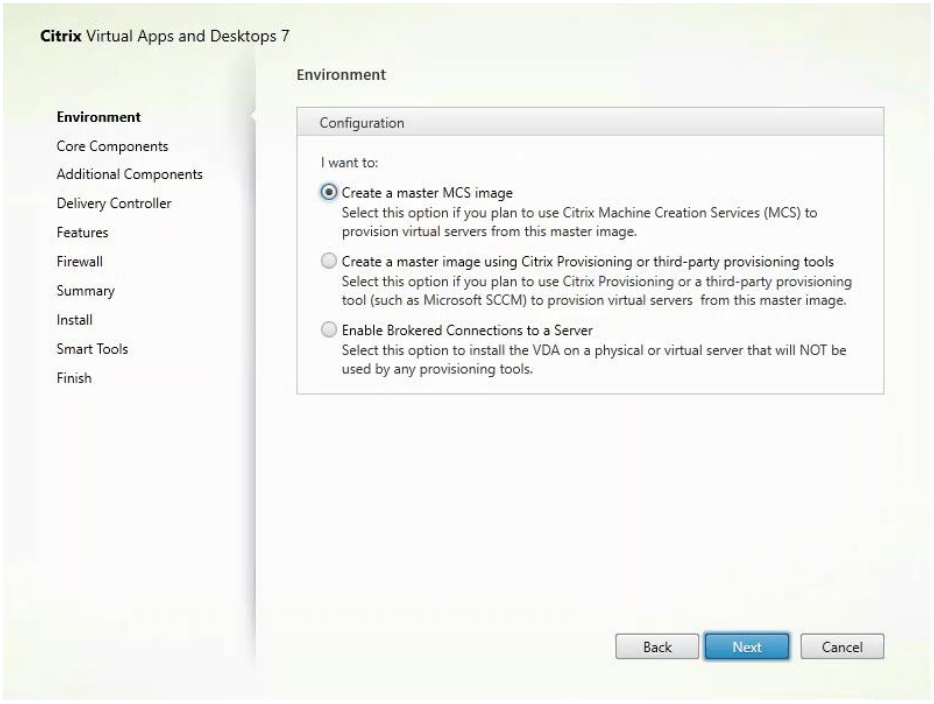


6. Optional: Select Citrix Workspace App.

7. Click Next.

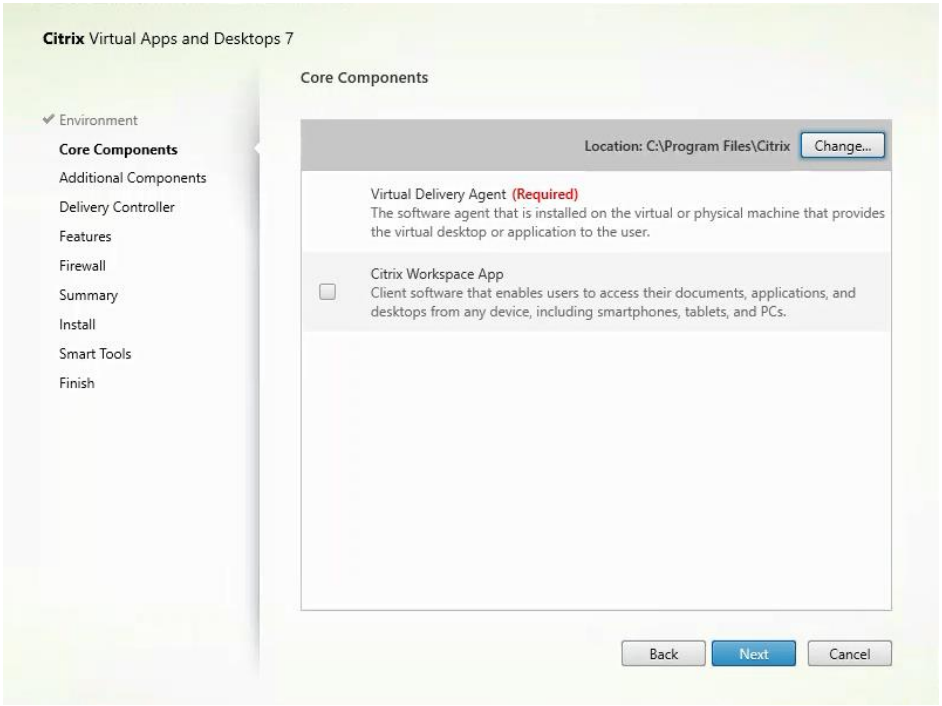


8. Click Next.

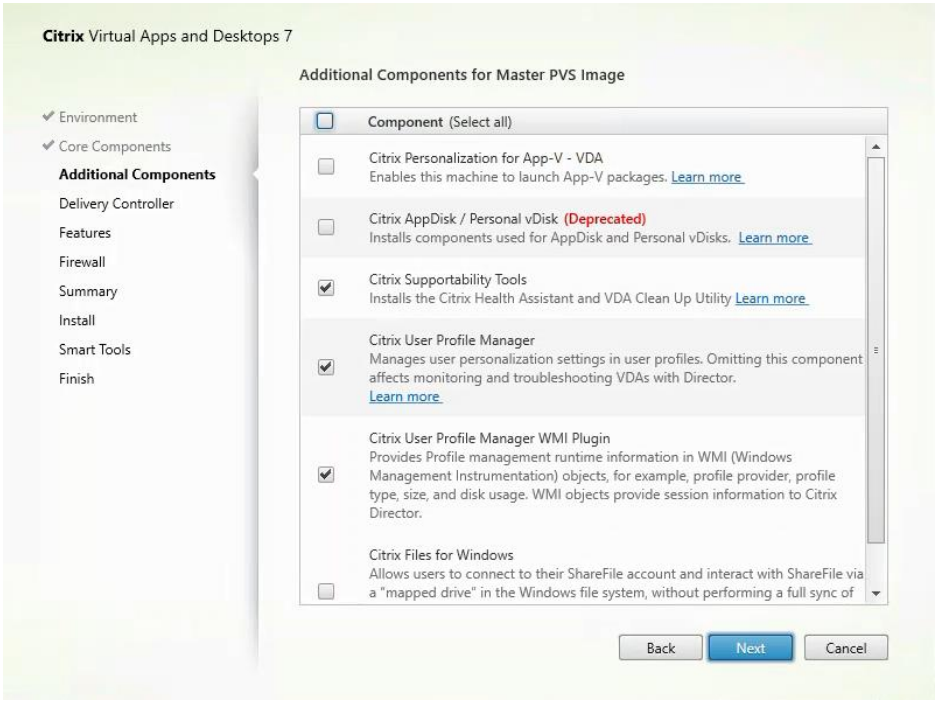


9. Select “Do it manually” and specify the FQDN of the Delivery Controllers.

10. Click Next.

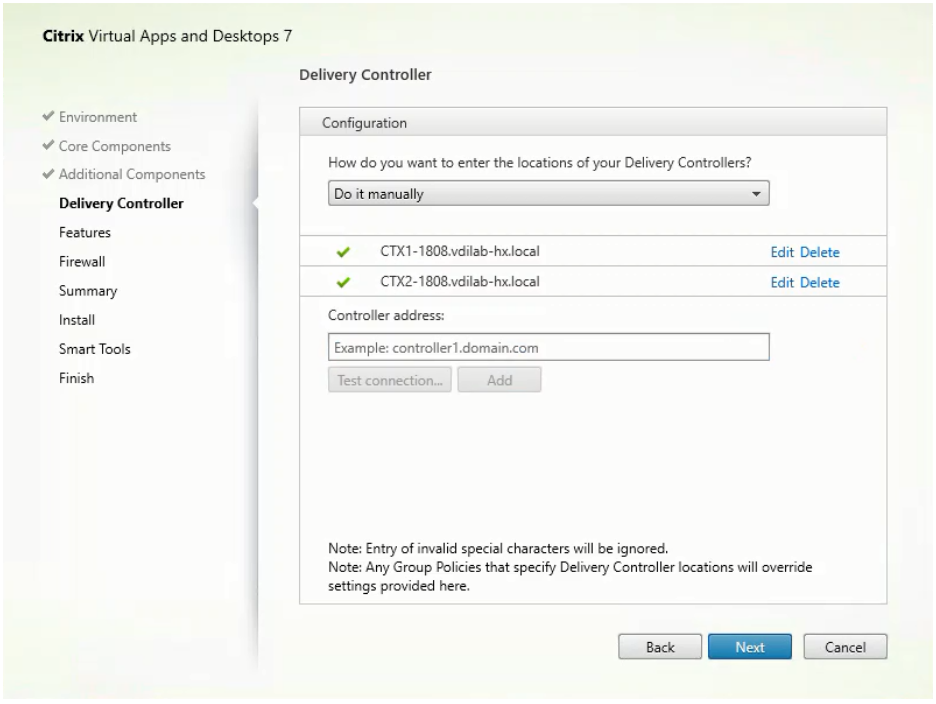


11. Accept the default features.

12. Click Next.

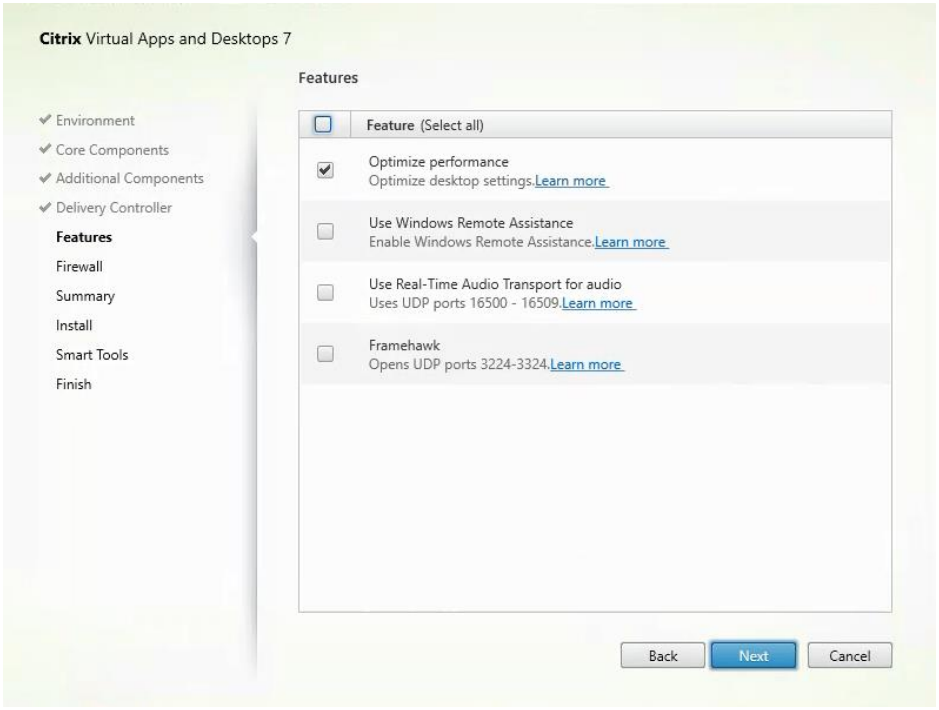


13. Allow the firewall rules to be configured automatically.

14. Click Next.

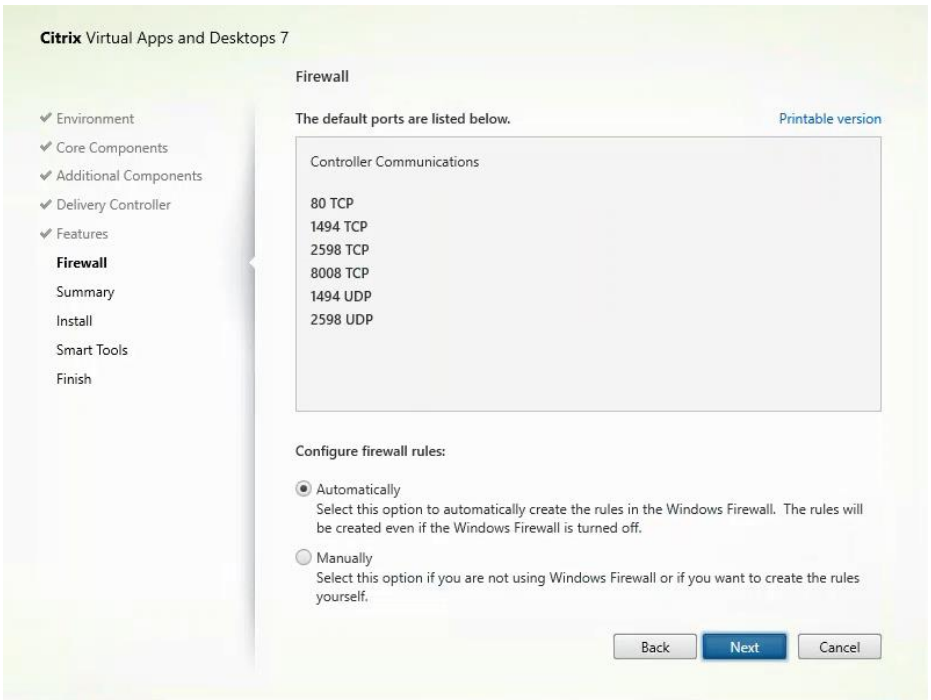


15. Verify the Summary and click Install.

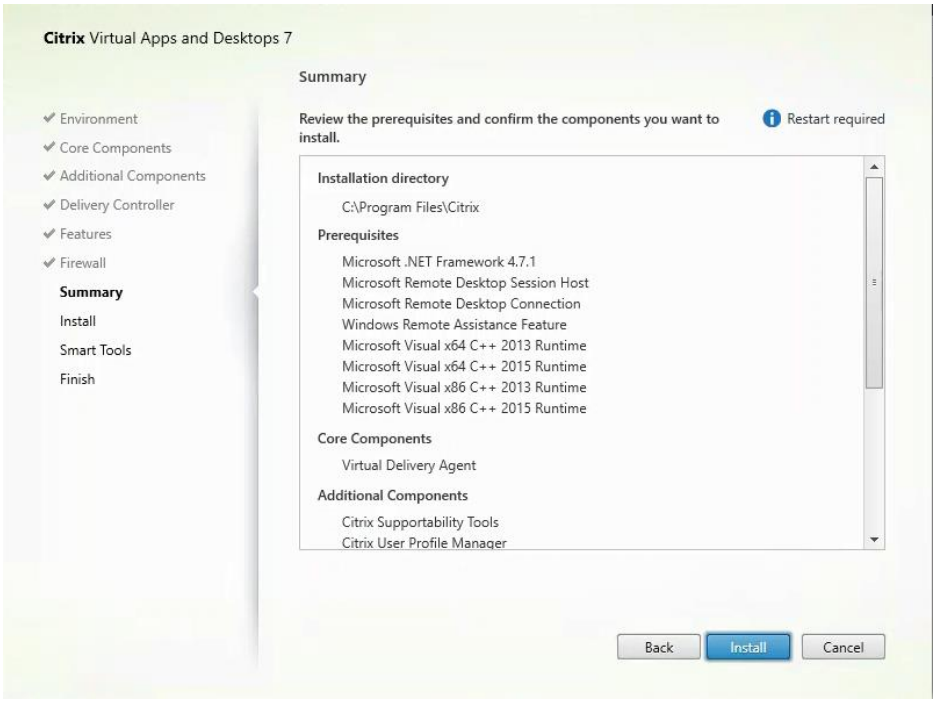


16. (Optional) Select Call Home participation.

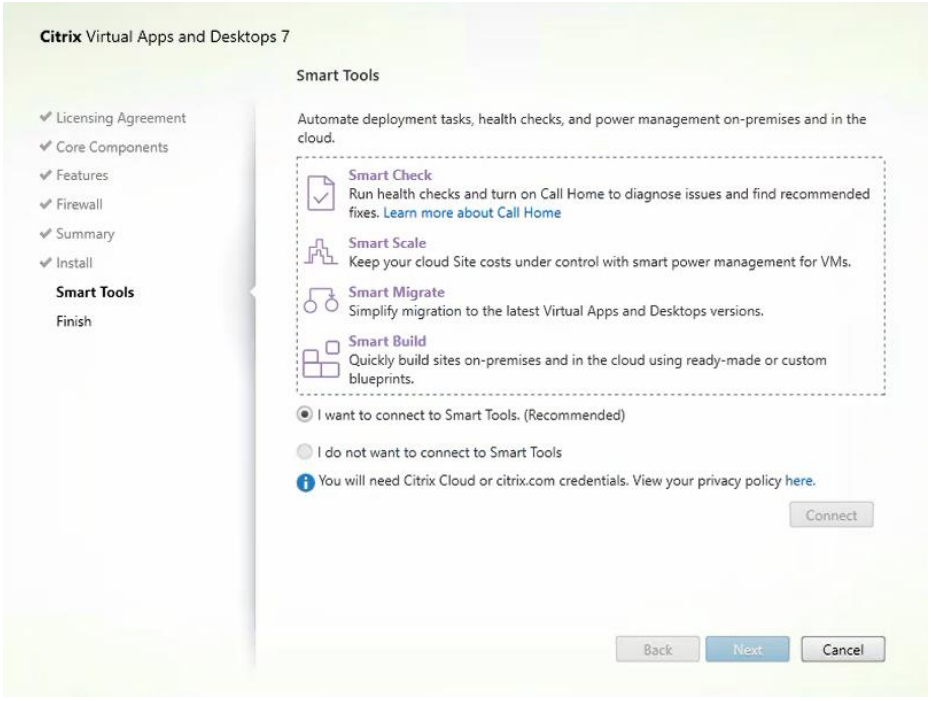


17. (Optional) check “Restart Machine.”

18. Click Finish.

19. Repeat these procedure so that VDAs are installed for both HVD (using the Windows 10 OS image) and the HSD desktops (using the Windows Server 2019 image).

20. Select an appropriate workflow for the HSD desktop.

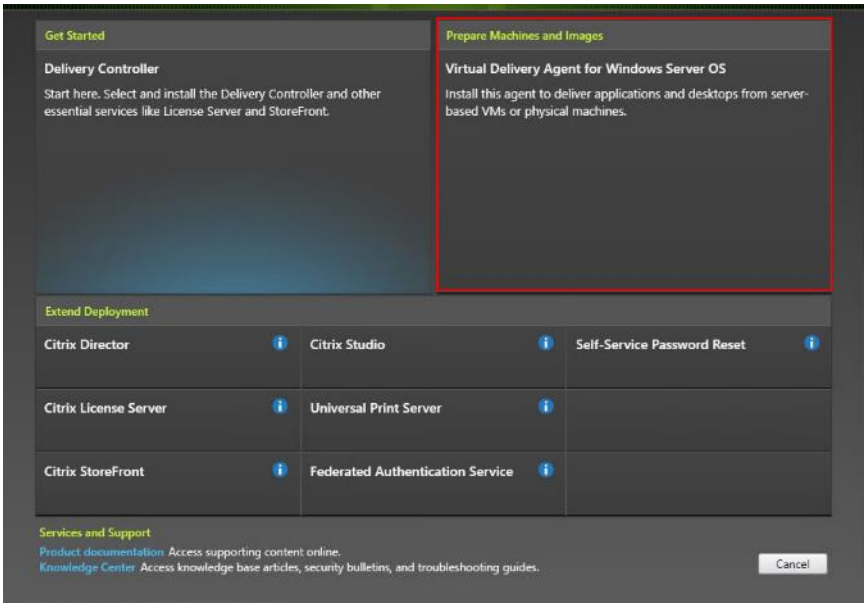


## Create Delivery Groups

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, follow these 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 Virtual Desktops server and launch Citrix Studio.

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

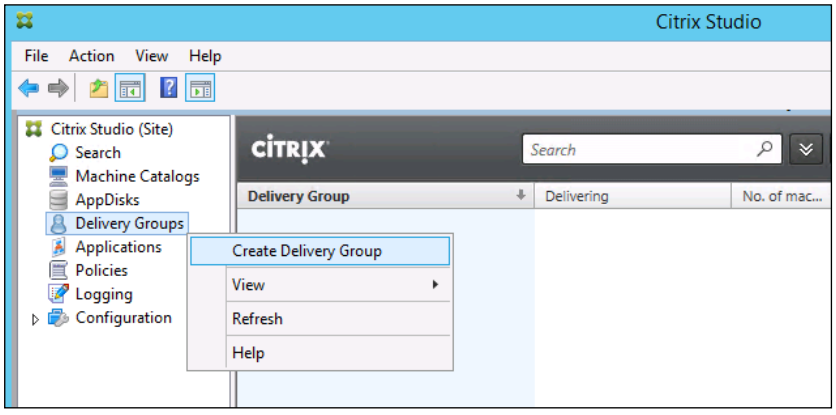


3. Click Next.

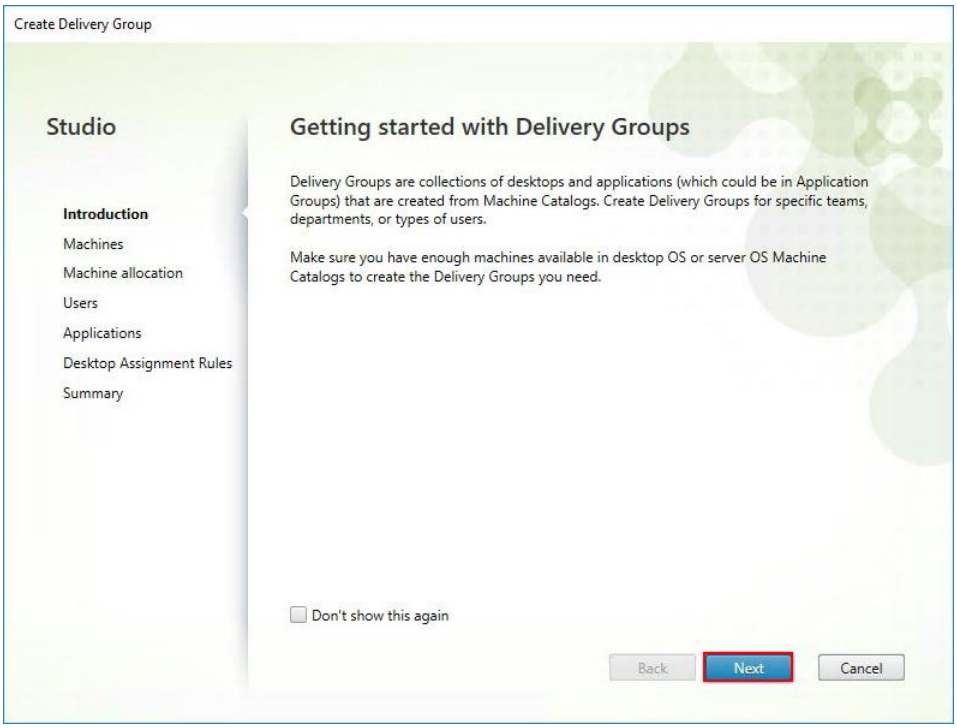


4. Select Machine catalog.
5. Provide the number of machines to be added to the delivery Group.
6. Click Next.

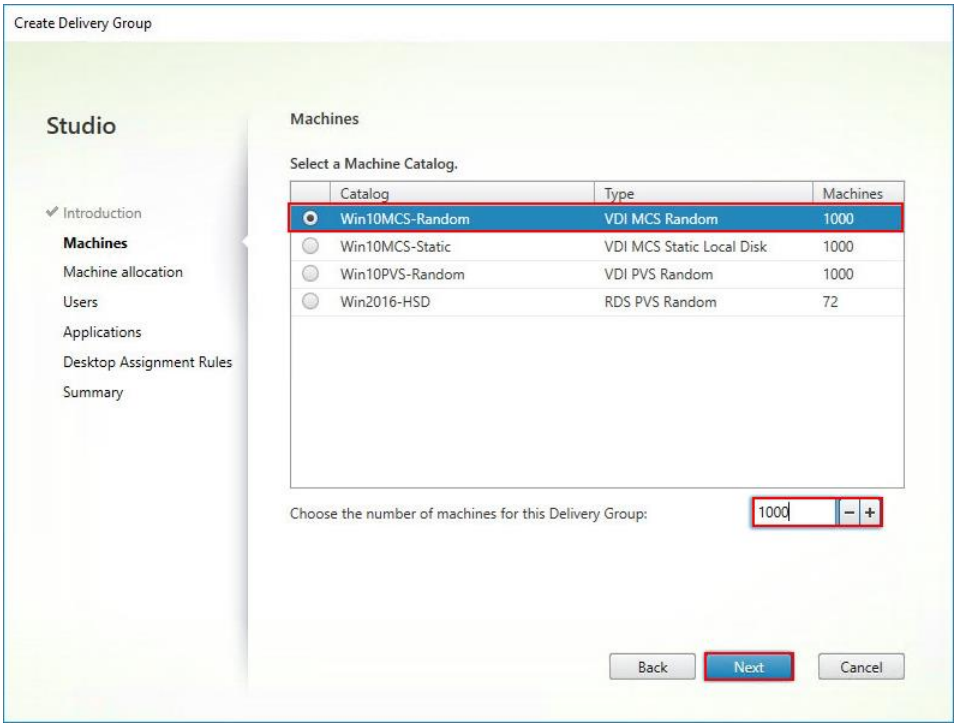

7. To make the Delivery Group accessible, you must add users, select Allow any authenticated users to use this Delivery Group.

8. Click Next.

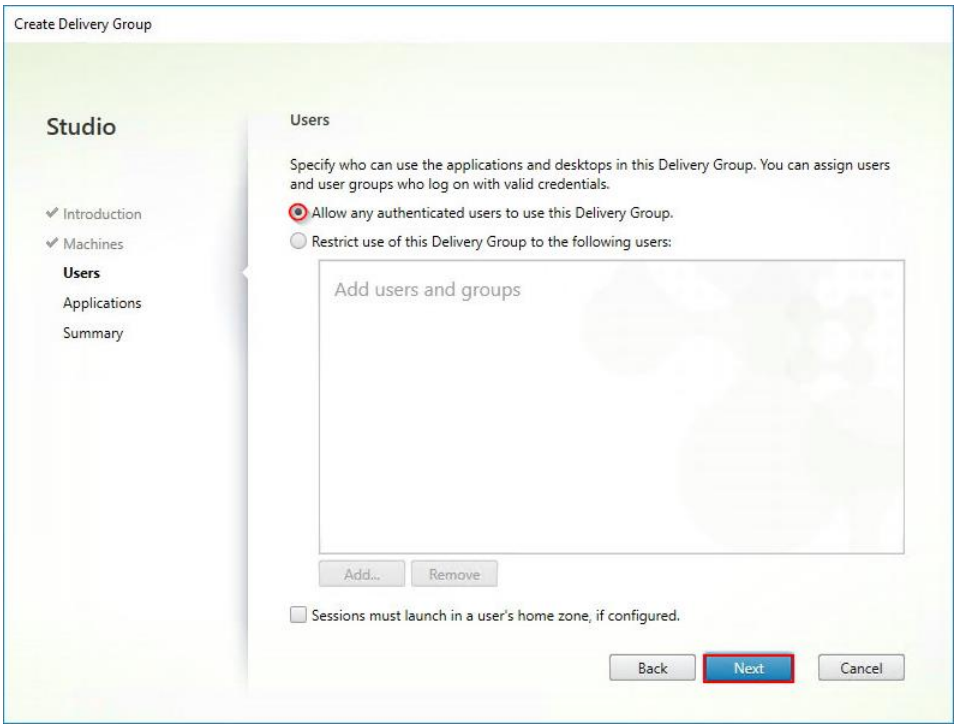

User assignment can be updated any time after Delivery group creation by accessing Delivery group properties in Desktop Studio.

9. (Optional) specify Applications catalog will deliver.

10. Click Next.

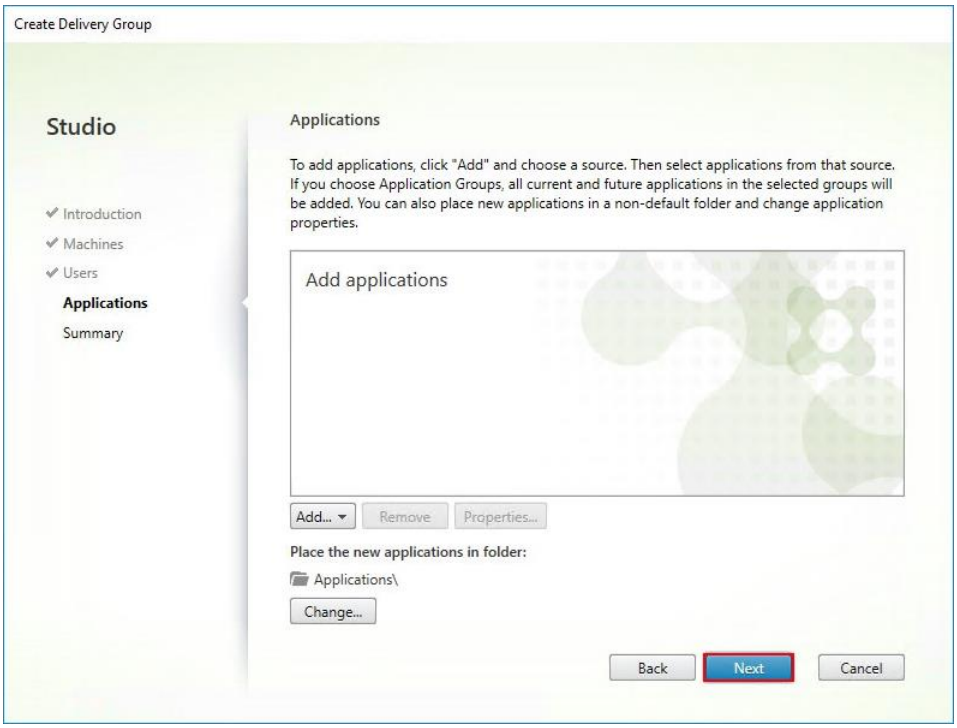


11. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Display name (for example, HVD or HSD).

12. Click Finish.

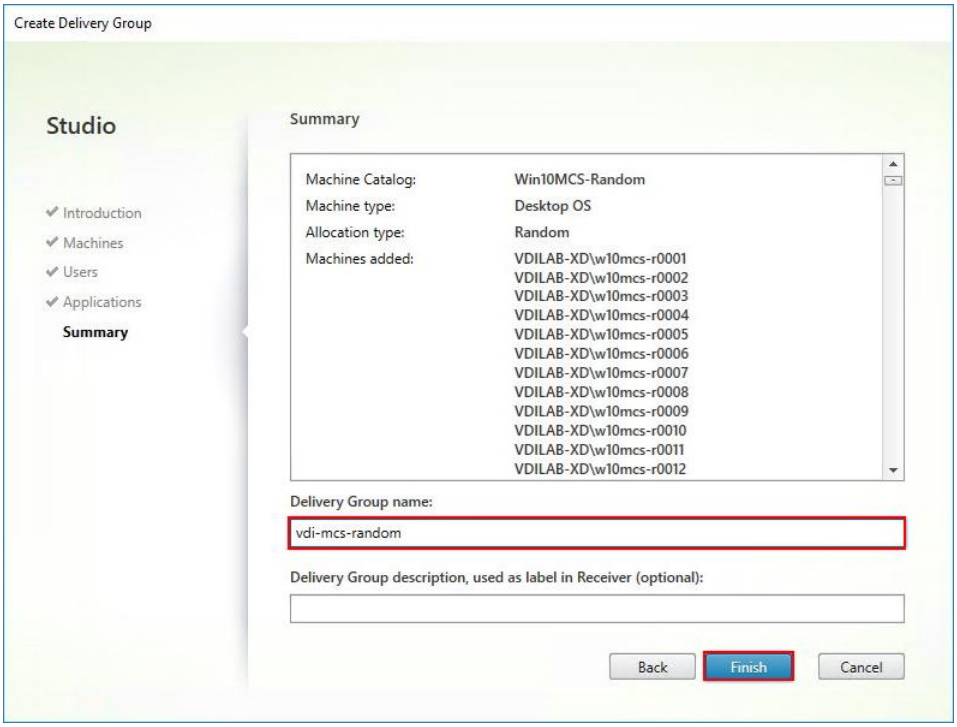

13. Citrix Studio lists the created Delivery Groups and the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab. Select Delivery Group and in Action List, select "Turn on Maintenance Mode."

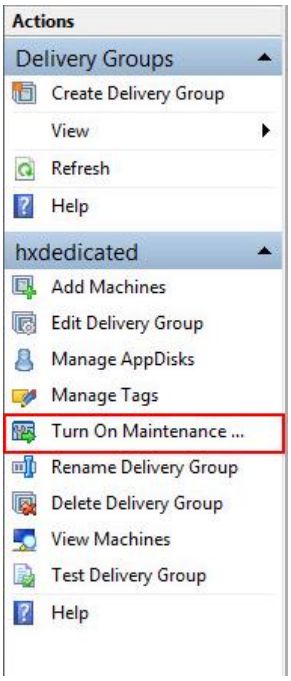


## Citrix Virtual Desktops Policies and Profile Management

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

### Configure Citrix Virtual Desktops Policies

Citrix Virtual Desktops policies control user access and session environments, and are the most efficient method of controlling connection, security, and bandwidth settings. You can create policies for specific groups of users, devices, or connection types with each policy. Policies can contain multiple settings and are typically defined through Citrix Studio. (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 122 shows policies for Login VSI testing in this CVD.

Figure 122 Virtual Desktops Policy

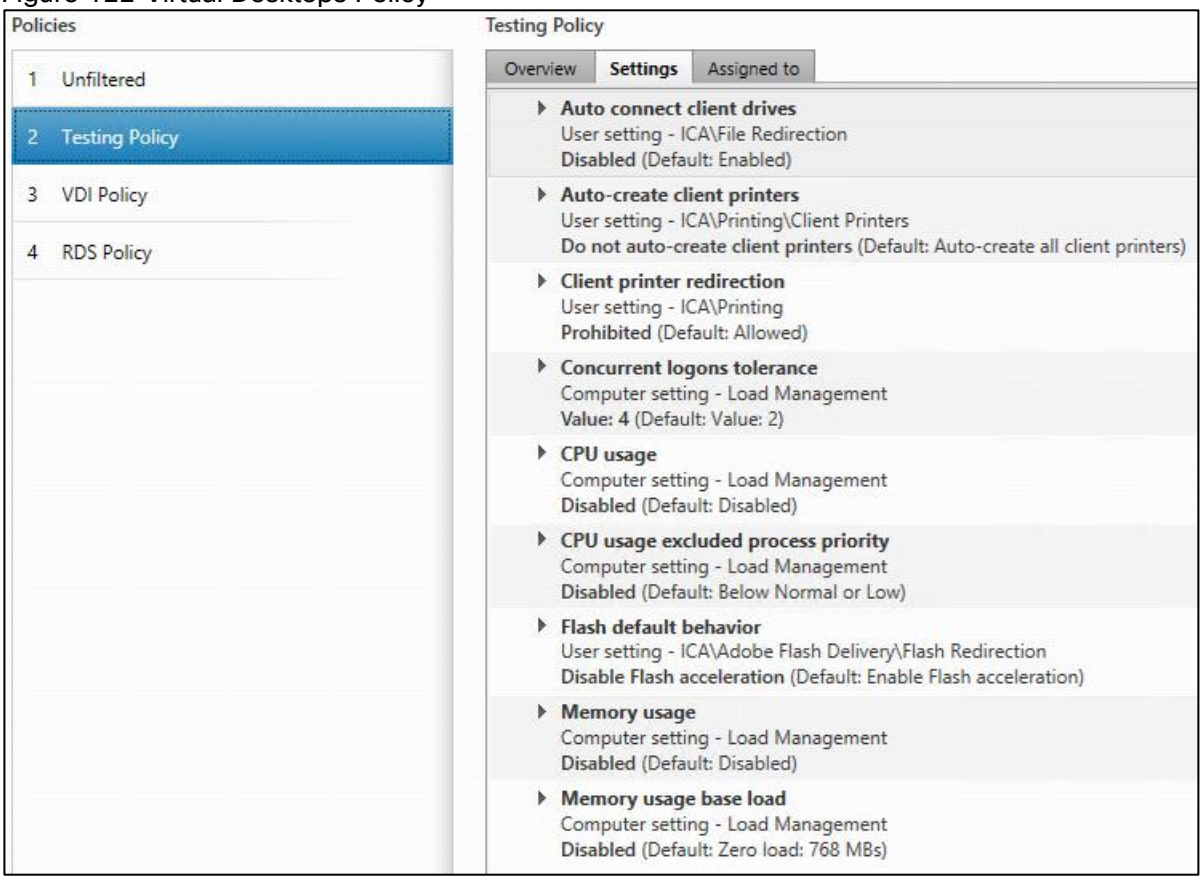


## Configure User Profile Management

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 Virtual Desktops 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/citrix-virtual-apps-desktops

Figure 123 VDI User Profile Manager Policy

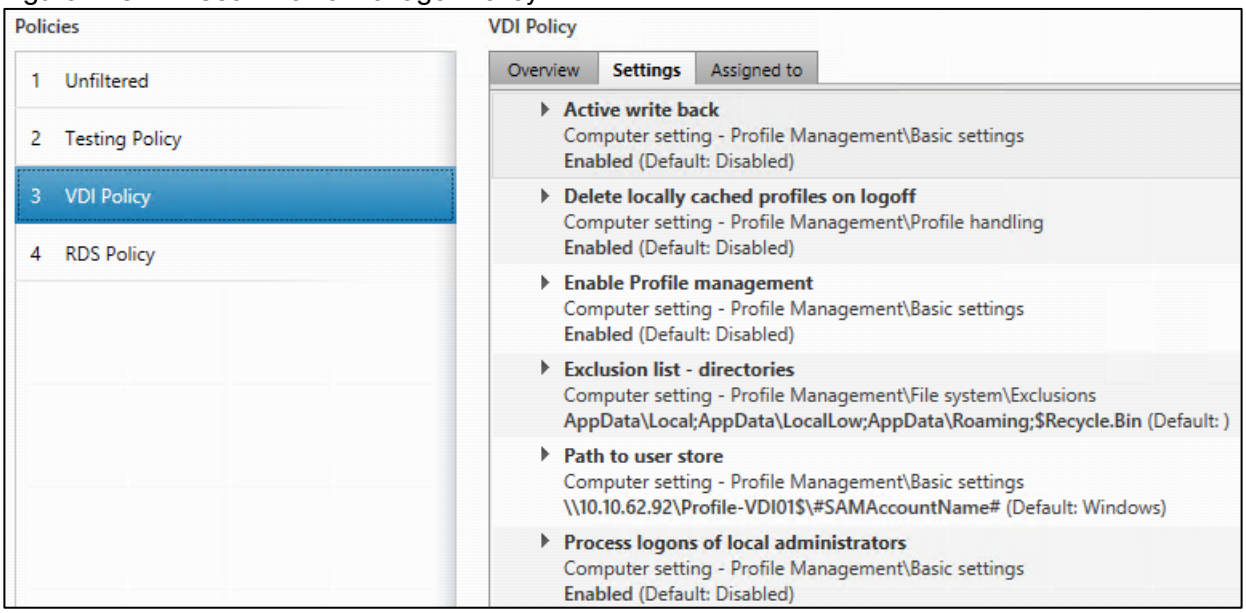

# Test Setup and Configurations

In this project, we tested a single Cisco HyperFlex cluster running four Cisco UCS HXAF220C-M5SX Rack Servers in a single Cisco UCS domain. This solution is tested to illustrate linear scalability for each workload studied.

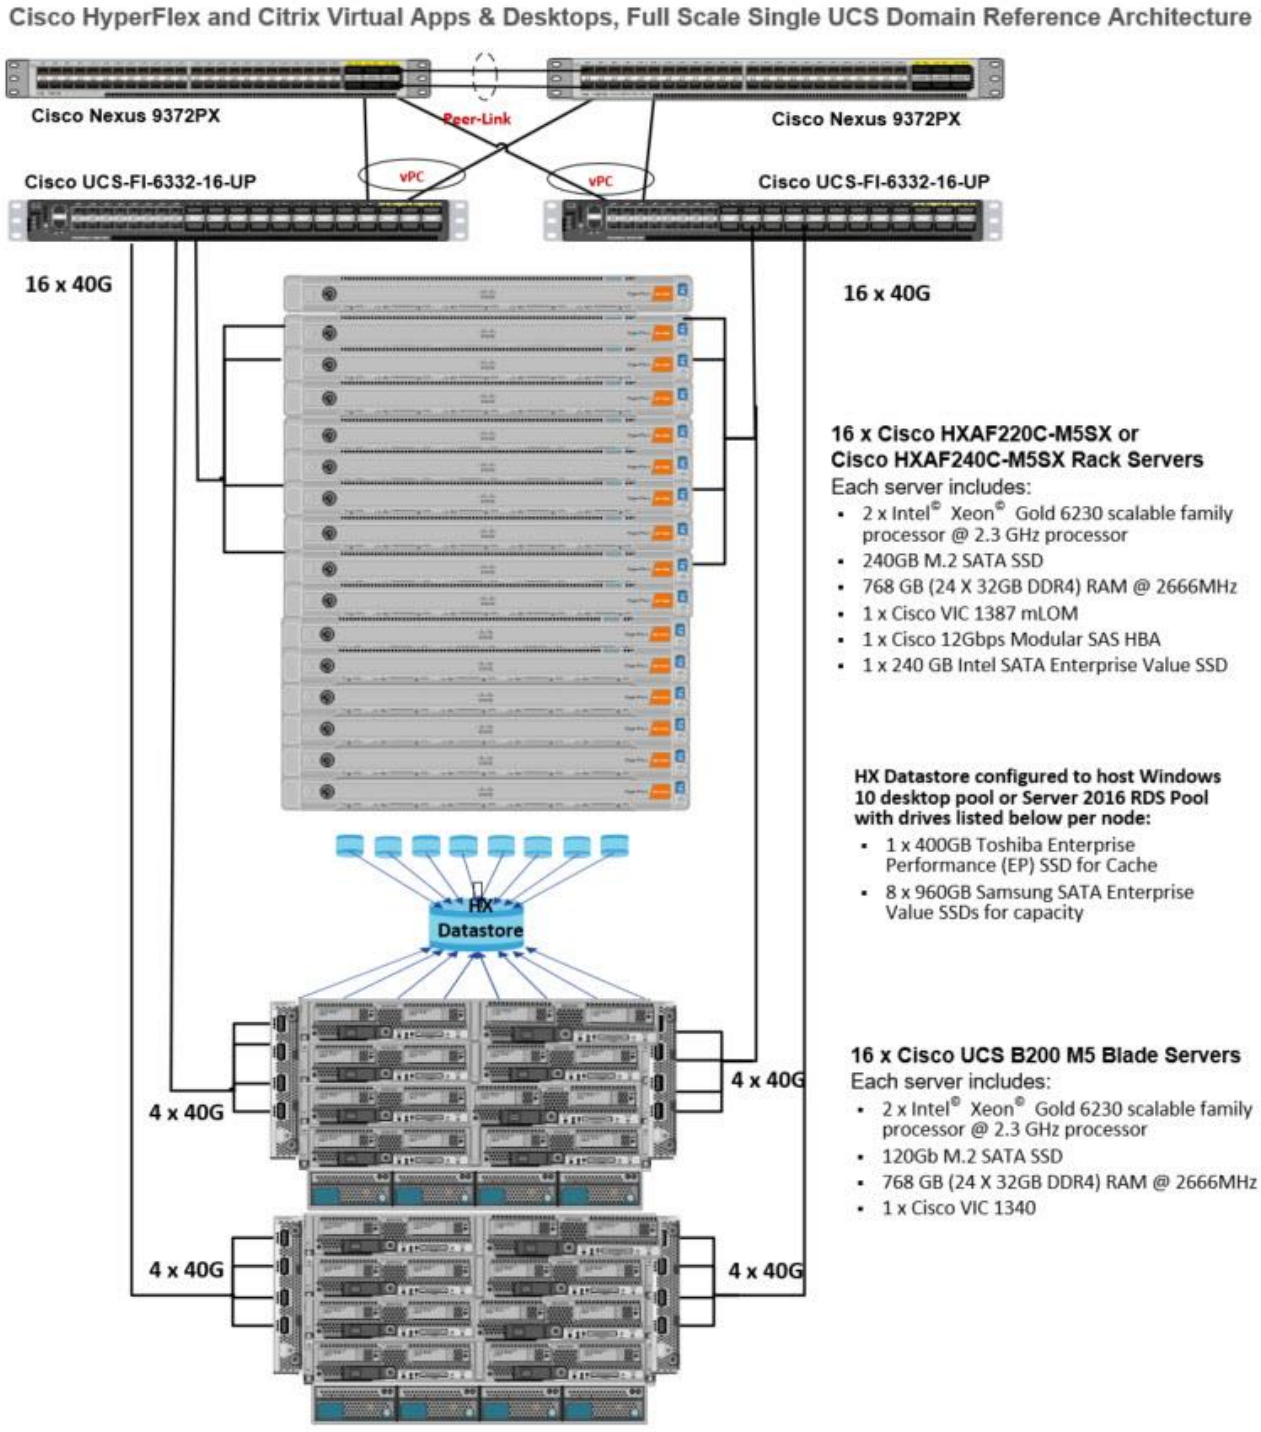

Hardware Components:

- 2 x Cisco UCS 6332-16UP Fabric Interconnects
- 2 x Cisco Nexus 93108YCPX Access Switches
- 16 x Cisco UCS HXAF220c-M5SX Rack Servers (2 Intel Xeon Gold 62306230 scalable family processor at 2.3 GHz, with 768 GB of memory per server [32 GB x 24 DIMMs at 2666 MHz])
- 16 x Cisco UCS B200 M5 Blade Servers (2 Intel Xeon Gold 62306230 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)

Software Components:

- Cisco UCS firmware 4.0(4d)
- Cisco HyperFlex Data Platform 4.0.1b
- Microsoft Hyper-V 2019
- Citrix Virtual Desktops 1909
- Citrix User Profile Management
- Microsoft SQL Server 2016
- Microsoft Windows 10
- Microsoft Windows 2019
- Microsoft Office 2016
- Login VSI 4.1.32

## Test Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the Hosted Shared Desktop Session under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users

completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

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

## Test Procedure

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

### Pre-Test Setup for Testing

All machines were shut down utilizing the Citrix Virtual Desktops 1909 Administrator Console.

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

### 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 4000 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. To run the test protocol, follow these steps:

1. Time 0:00:00 Start esxtop 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 Virtual Desktops 1909 Administrator Console.
4. Time 0:06 First machines boot.
5. Time 0:35 Single Server or Scale target number of RDS Servers registered on XD.

**No more than 60 Minutes of rest time is allowed after the last desktop is registered and available on Citrix Virtual Desktops 1909 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 4.1.5 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 Virtual Desktops 1909 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 4.1.25 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 Virtual Desktops Studio 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 Virtual Desktops 1909 Hosted Shared Desktop with Citrix Virtual Desktops 1909 Composer provisioning using Microsoft Windows Server 2016 sessions on Cisco UCS HXAF220c-M4S, Cisco UCS 220 M4 and Cisco UCS B200 M4 servers.

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 Citrix and Microsoft 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 number of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what is its true maximum user capacity.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSImax. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time, it will be clear the response times escalate at saturation point.

This VSImax is the “Virtual Session Index (VSI)”. With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

### Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system and are initiated at logon within the simulated user’s desktop session context.

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent 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 124 Sample of a VSI Max Response Time Graph, Representing a Normal Test

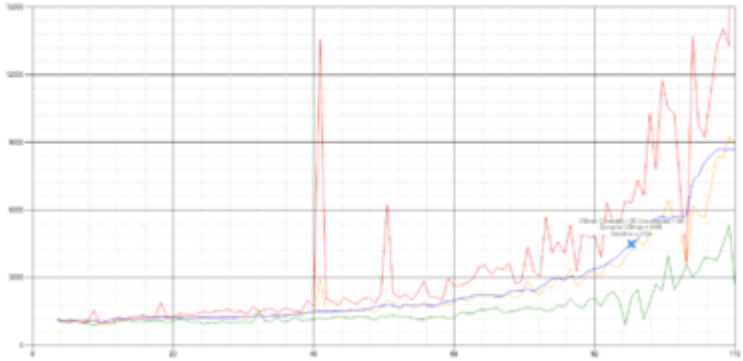

Figure 125 Sample of a VSI Test Response Time Graph with a Clear Performance Issue

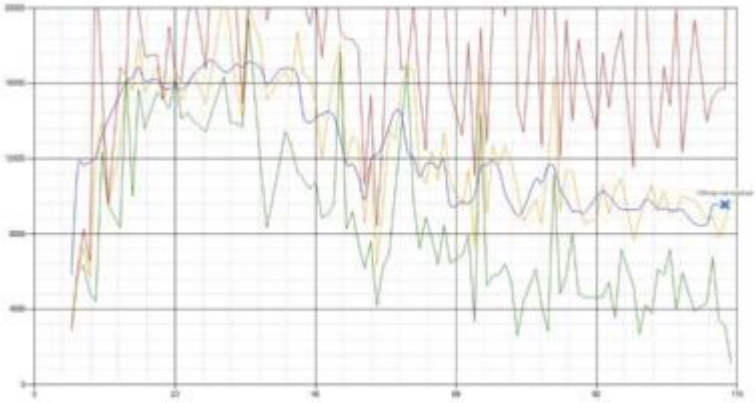

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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 represents 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 number 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 are 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 1244 desktops and measure the time it takes for the 1244th virtual machine to register as available in the Virtual Desktops Administrator console.

The Cisco HyperFlex HXAF220c-M5SX based All-Flash cluster running Data Platform version 4.0.1b software can accomplish this task in 5 minutes.

## Recommended Maximum Workload and Configuration Guidelines

### Sixteen Node Cisco HXAF220c-M5S Rack Server, B200 M5 Compute Nodes & HyperFlex All-Flash Cluster

For Citrix Virtual Apps 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-M5SX server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to ensure that end-user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.

**Memory should never be oversubscribed for Desktop Virtualization workloads.**

**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. |

| Test Phase | Description |
| --- | --- |
| Logoff | Sessions finish executing the Login VSI workload and logoff. |

## Sixteen Node Cisco HXAF220c-M5S Rack-Mount Server, Cisco UCS B200 M5 Compute Nodes and HyperFlex All-Flash Cluster

For Citrix Virtual Apps 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-M5SX server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to ensure that end-user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.

**Memory should never be oversubscribed for Desktop Virtualization workloads.**

| 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-M5SX with Intel Xeon Gold 6230 scalable family processors and 768GB of RAM for Windows 10 desktops with Office 2016 is 4000 virtual desktops.**

**The recommended maximum workload for a Cisco HyperFlex cluster configured on Cisco HXAF220c-M5SX with Intel Xeon Gold 6230 scalable family processors and 768GB of RAM for Windows Server 2019 RDS desktop sessions with Office 2016 is 5000 virtual desktops.**

### 5000 RDS Sessions and 4000 Windows 10 Citrix PVS Non-Persistent Testing 32 Node Cisco HyperFlex Cluster

Hosted Shared desktops with 5000 user sessions on 250 Windows Server 2019 virtual machines on a 32-node HyperFlex cluster.

Test results for 5000 user sessions on Citrix RDS highlights include:

- 0.707 second baseline response time
- 0.958 second average response time with 5000 desktops running
- Average CPU utilization of 35 percent during steady state
- Average of 400 GB of RAM used out of 768 GB available
- 22,000 peak I/O operations per second (IOPS) per cluster at steady state
- 650MBps peak throughput per cluster at steady state

Figure 126 Login VSI Analyzer Chart for 5000 Windows 2019 Citrix Shared Desktops

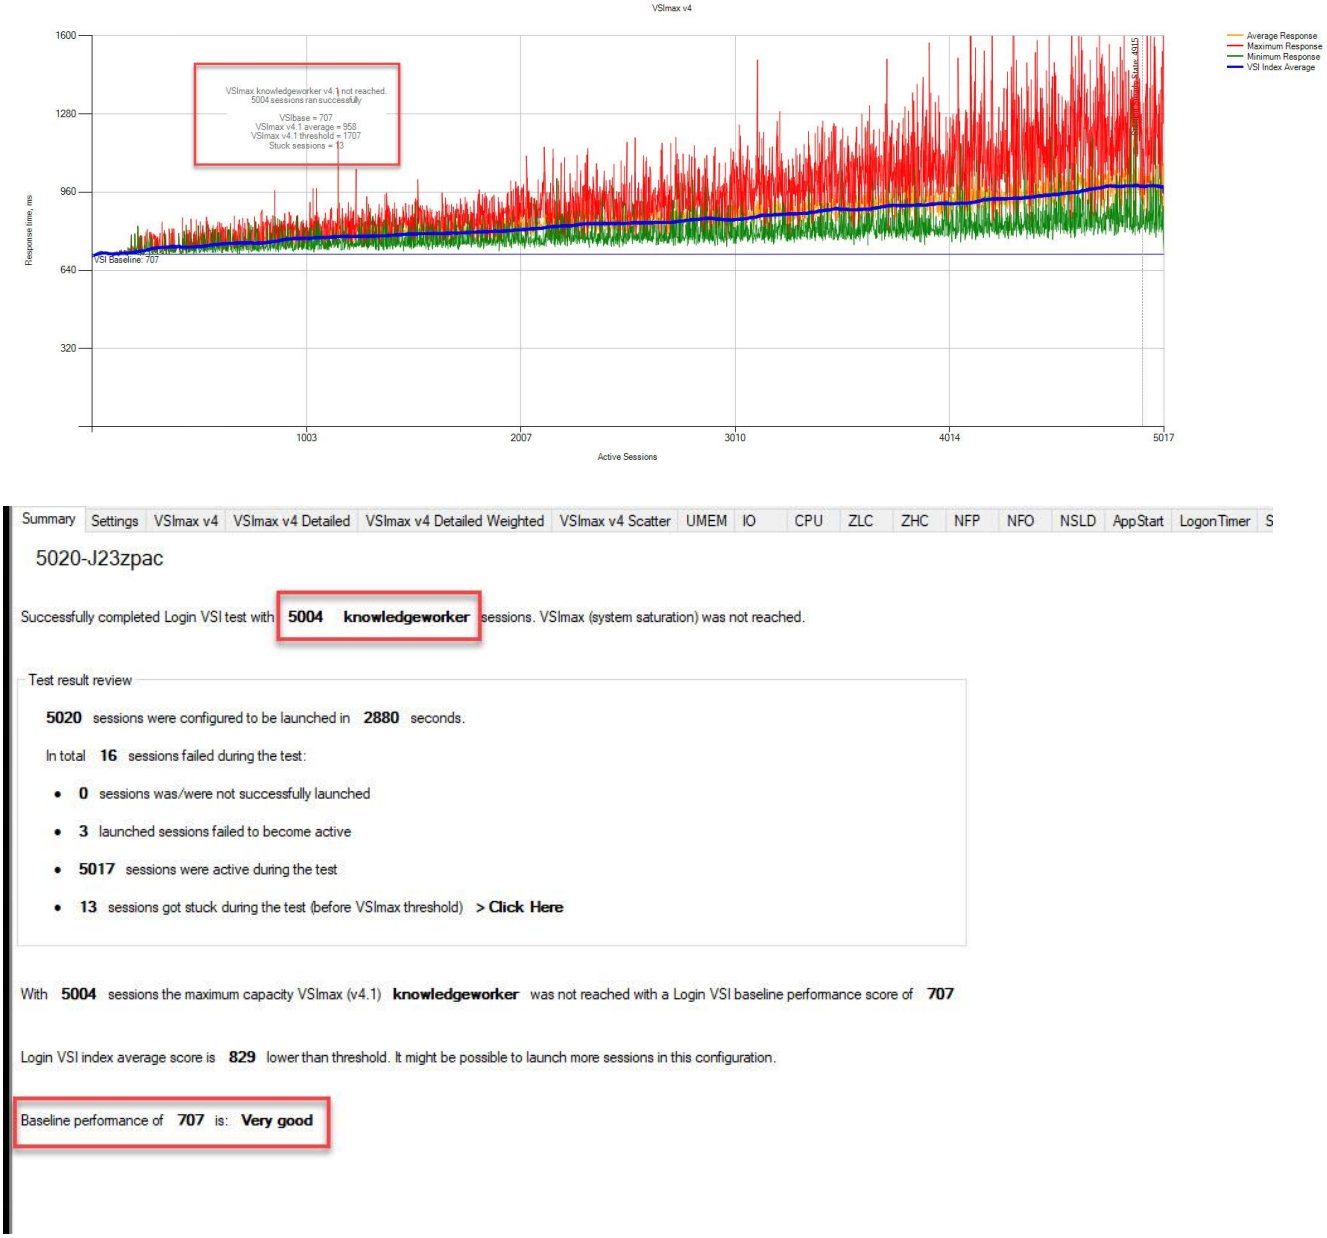

Figure 127 Three Consecutive Login VSI Analyzer Chart for 5000 Windows 2019 Citrix Shared Desktops

## Hyper-V Performance Counters

When running a Hyper-V environment for our Citrix Virtual Desktop workloads, it's important to monitor a few key performance counters to ensure the best end-user experience. We typically look for CPU utilization, memory availability, network throughput and Storage performance:

- CPU Performance: With Hyper-V our main counter is % Total Run Time for the Hypervisor's Logical Processors.
- Memory Availability: We measure the memory available in megabytes to ensure that memory is not being consumed at a high level.
- Network throughput: We measure the bytes sent and received by the VM Network vswitch on each Hyper-V host.
- Storage performance: We use HyperFlex Connect to monitor and review storage performance during VDI.

The following figures show the results of our workload testing.

Figure 128 32 Hyper-V Hosts CPU Core Utilization Running 5000 Windows Server 2019 Citrix Hosted Shared Desktops (Hyper-V Logical Processor % Total Run Time)

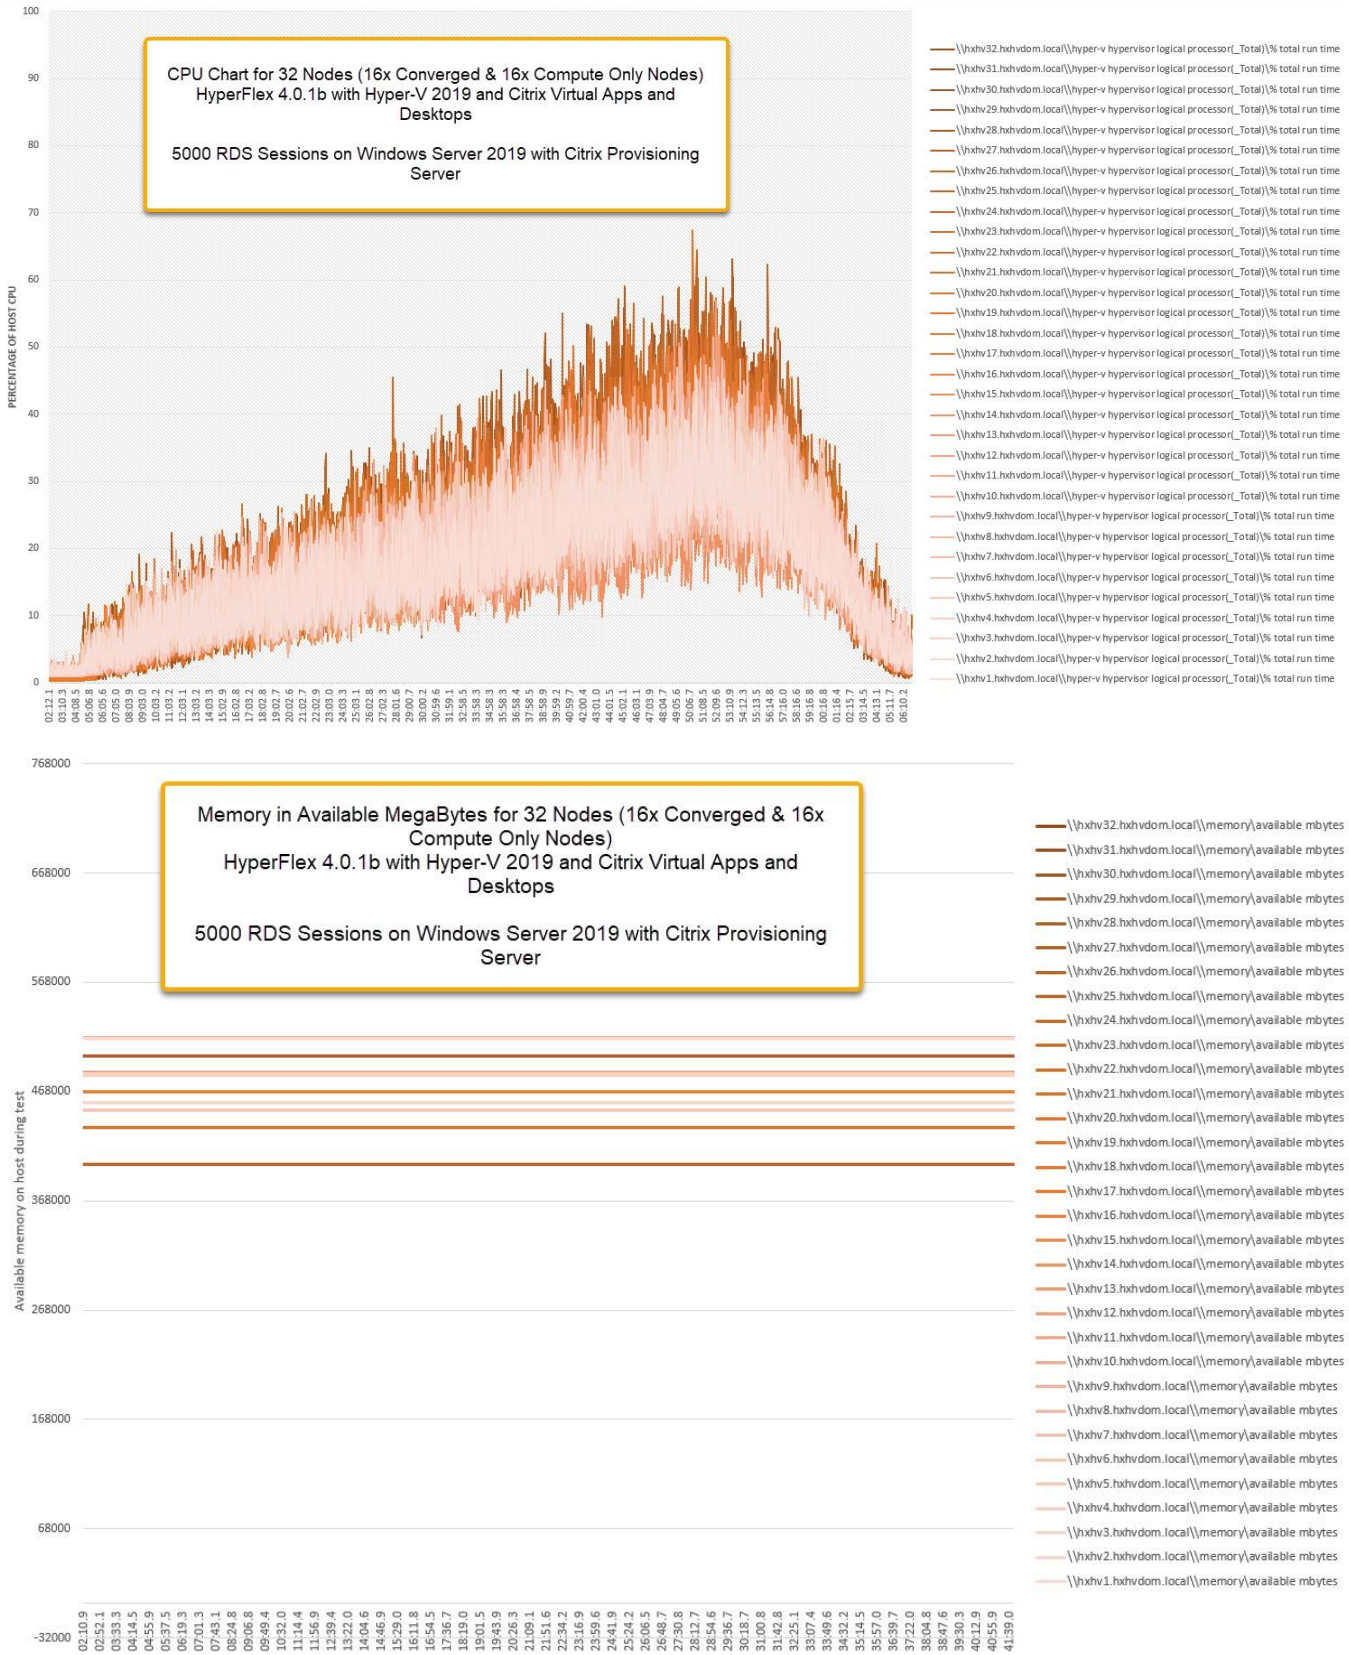

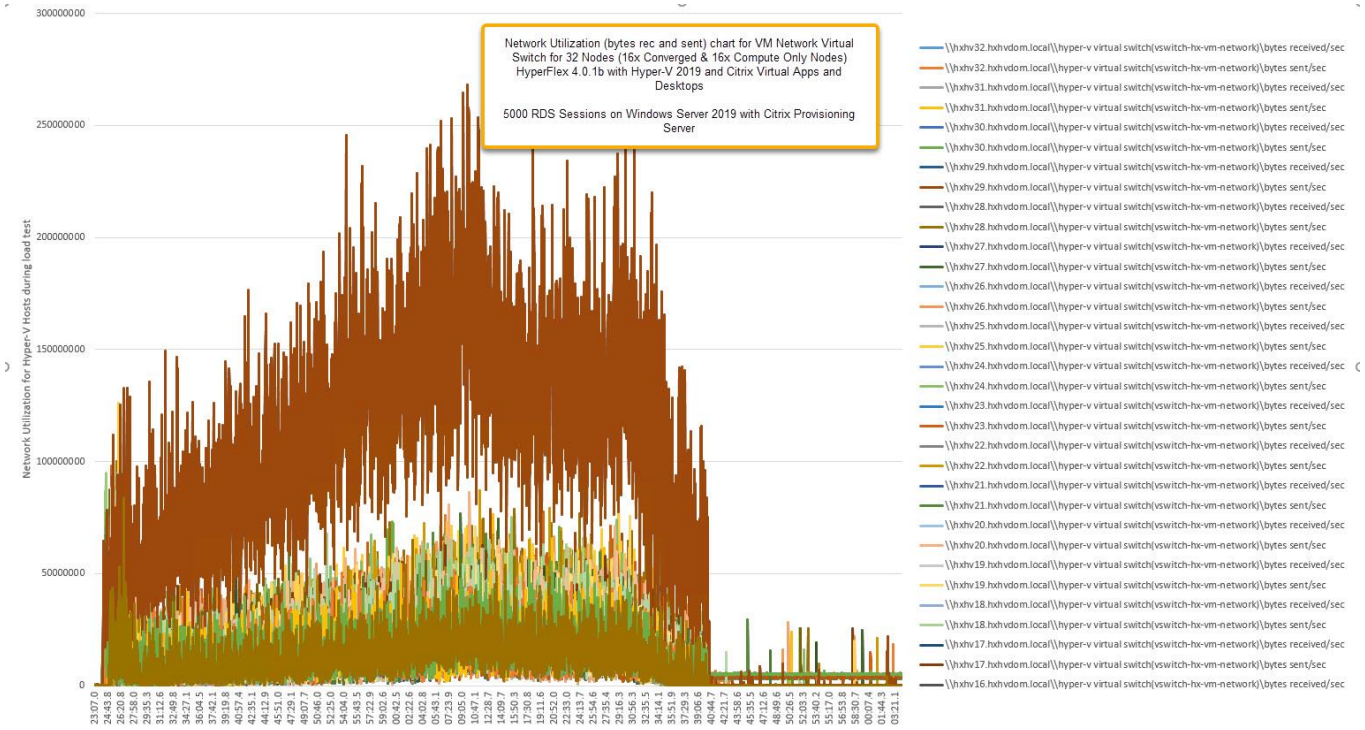
Network Utilization (bytes rec and sent) chart for VM Network Virtual Switch for 32 Nodes (16x Converged & 16x Compute Only Nodes) HyperFlex 4.0.1b with Hyper-V 2019 and Citrix Virtual Apps and Desktops
5000 RDS Sessions on Windows Server 2019 with Citrix Provisioning Server
Network Utilization for Hyper-V Hosts during load test
300000000
250000000
200000000
150000000
100000000
50000000
0
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv21.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv21.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv20.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv20.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv19.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv19.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv18.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv18.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv17.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv17.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv16.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec

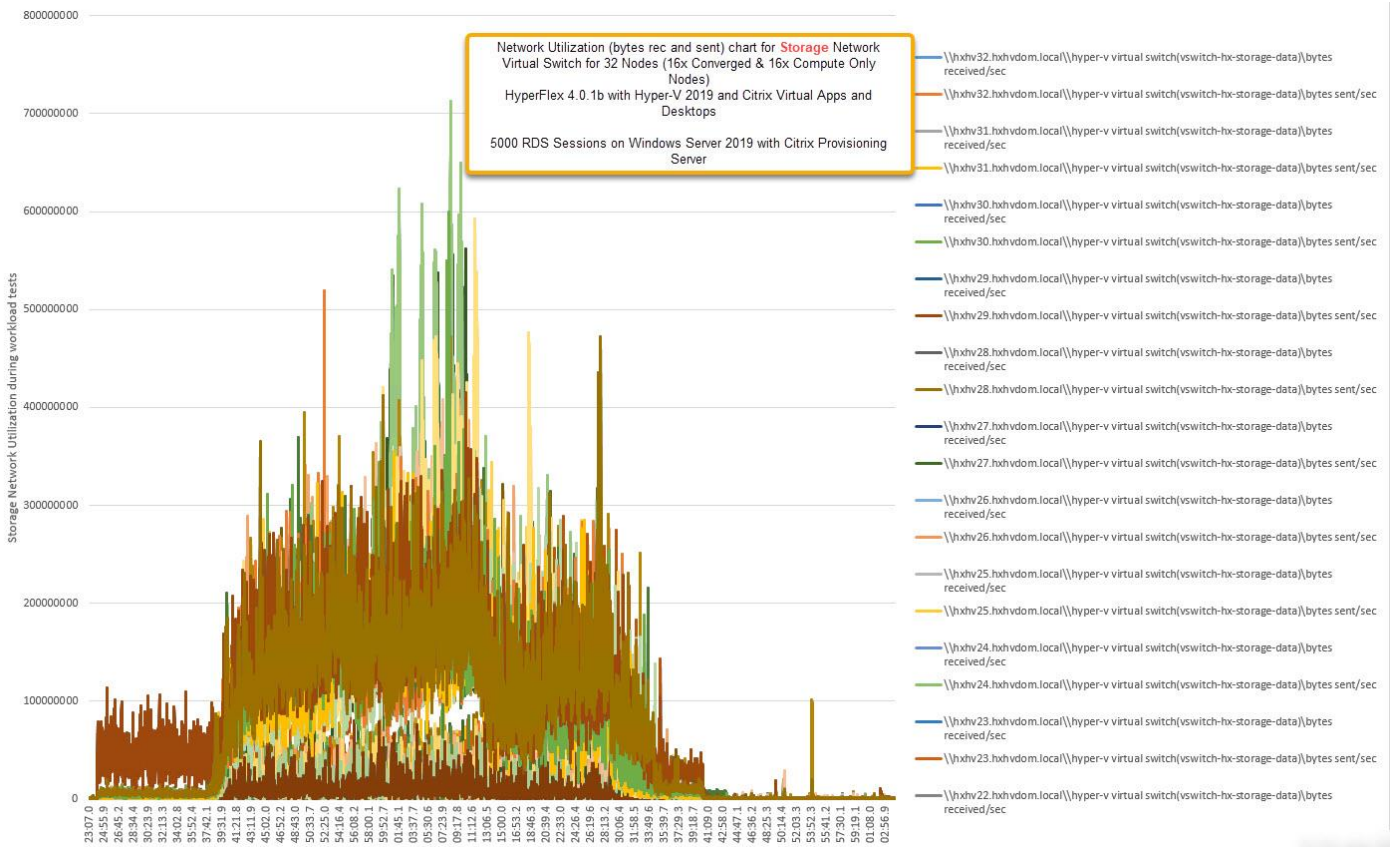
Network Utilization (bytes rec and sent) chart for Storage Network Virtual Switch for 32 Nodes (16x Converged & 16x Compute Only Nodes)
HyperFlex 4.0.1b with Hyper-V 2019 and Citrix Virtual Apps and Desktops
5000 RDS Sessions on Windows Server 2019 with Citrix Provisioning Server
Storage Network Utilization during workload tests
800000000
700000000
600000000
500000000
400000000
300000000
200000000
100000000
0
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec

Figure 129 HyperFlex Cluster UI Performance Chart for Knowledge Worker Workload Running 5000 User Test on Citrix Windows Server 2019

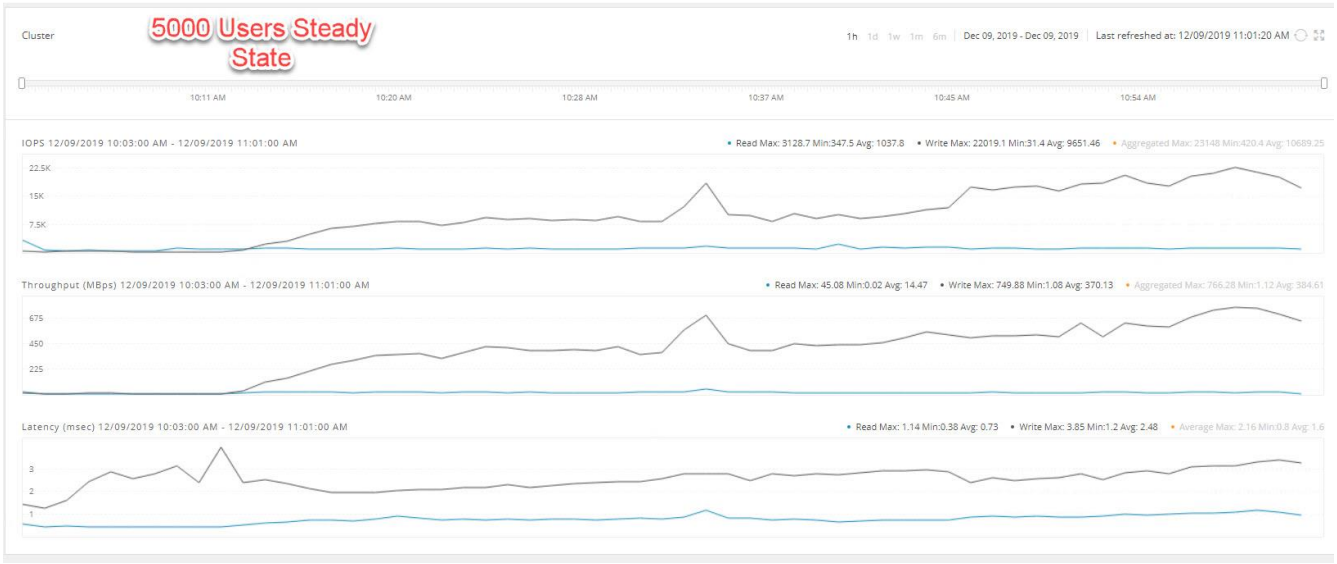


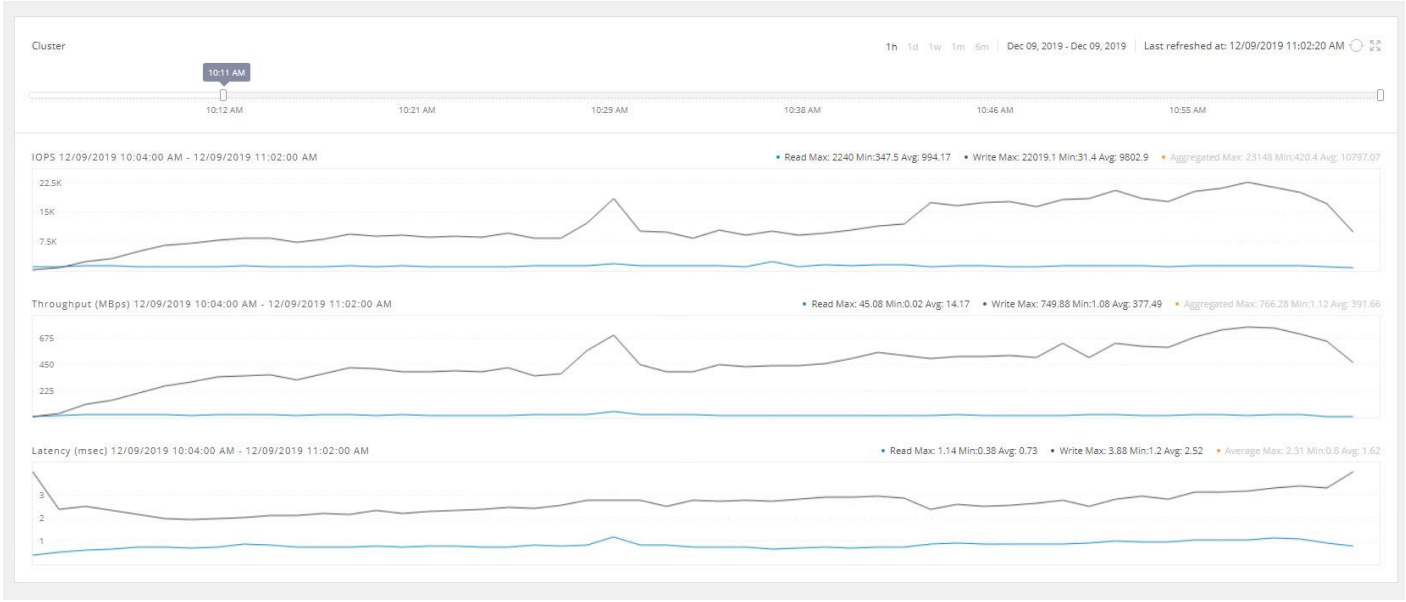


Test results for 4000 Citrix VDI Desktops highlights include:

- 0.851 second baseline response time
- 0.970 second average response time with 4000 desktops running
- Average CPU utilization of 40 percent during steady state
- Average of 350 GB of RAM used out of 768 GB available
- 4500 peak I/O operations per second (IOPS) per cluster at steady state
- 150MBps peak throughput per cluster at steady state

Figure 130 Login VSI Analyzer Chart for 4000 Windows 10 Citrix Virtual Desktops

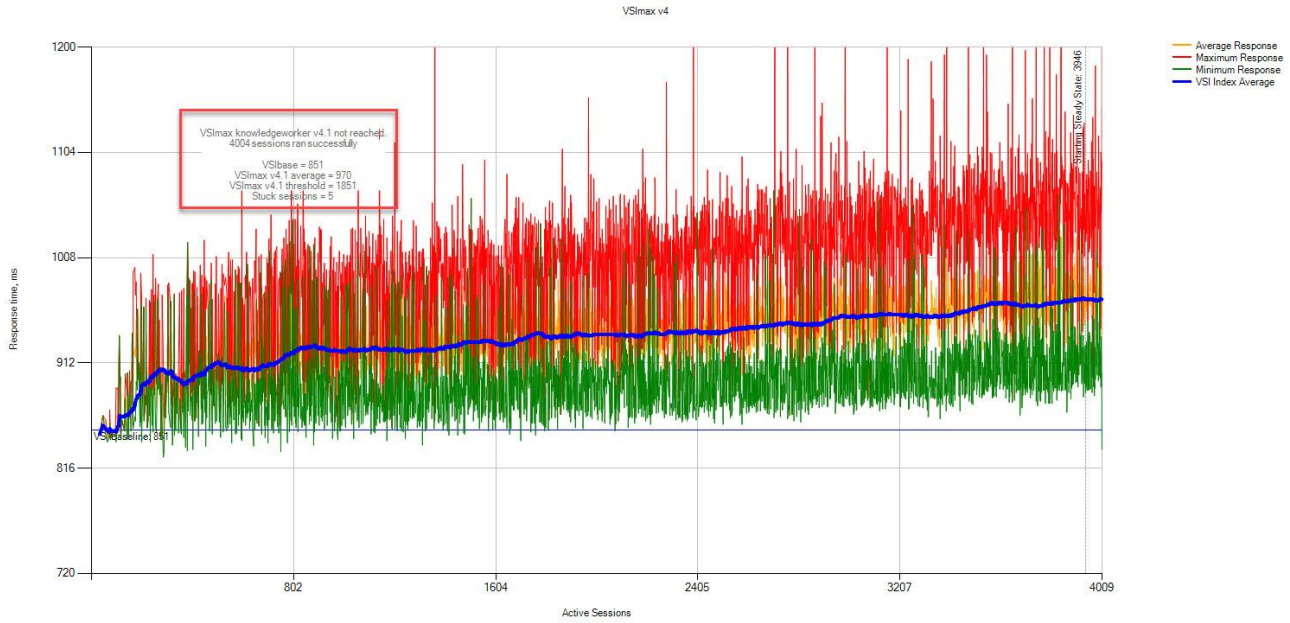


40402-8

Successfully completed Login VSI test with **4004** **knowledgeworker** sessions. VSImax (system saturation) was not reached.

Test result review

**4045** sessions were configured to be launched in **2880** seconds.

In total **41** sessions failed during the test:

- **0** sessions was/were not successfully launched
- **36** launched sessions failed to become active
- **4009** sessions were active during the test
- **5** sessions got stuck during the test (before VSImax threshold) **> Click Here**

With **4004** sessions the maximum capacity VSImax (v4.1) **knowledgeworker** was not reached with a Login VSI baseline performance score of **851**

Login VSI index average score is **896** lower than threshold. It might be possible to launch more sessions in this configuration.

Baseline performance of **851** is: **Good**

Figure 131 Three Consecutive Login VSI Analyzer Chart for 4000 Windows 10 Citrix PVS Non-persistent Virtual Desktops

## Hyper-V Performance Counters

When running a Hyper-V environment for our Citrix Virtual Desktop workloads, it's important to monitor a few key performance counters to ensure the best end-user experience. We typically look for CPU utilization, memory availability, network throughput and Storage performance:

- CPU Performance: With Hyper-V our main counter is % Total Run Time for the Hypervisor's Logical Processors.
- Memory Availability: We measure the memory available in megabytes to ensure that memory is not being consumed at a high level.
- Network throughput: We measure the bytes sent and received by the VM Network vswitch on each Hyper-V host.
- Storage performance: We use HyperFlex Connect to monitor and review storage performance during VDI.

The following figures show the results of our workload testing.

Figure 132 Thirty-Two Hyper-V Hosts CPU Core Utilization Running 4000 Windows 10 Citrix PVS Non-persistent Virtual Desktops (Hyper-V Logical Processor % Total Run Time)

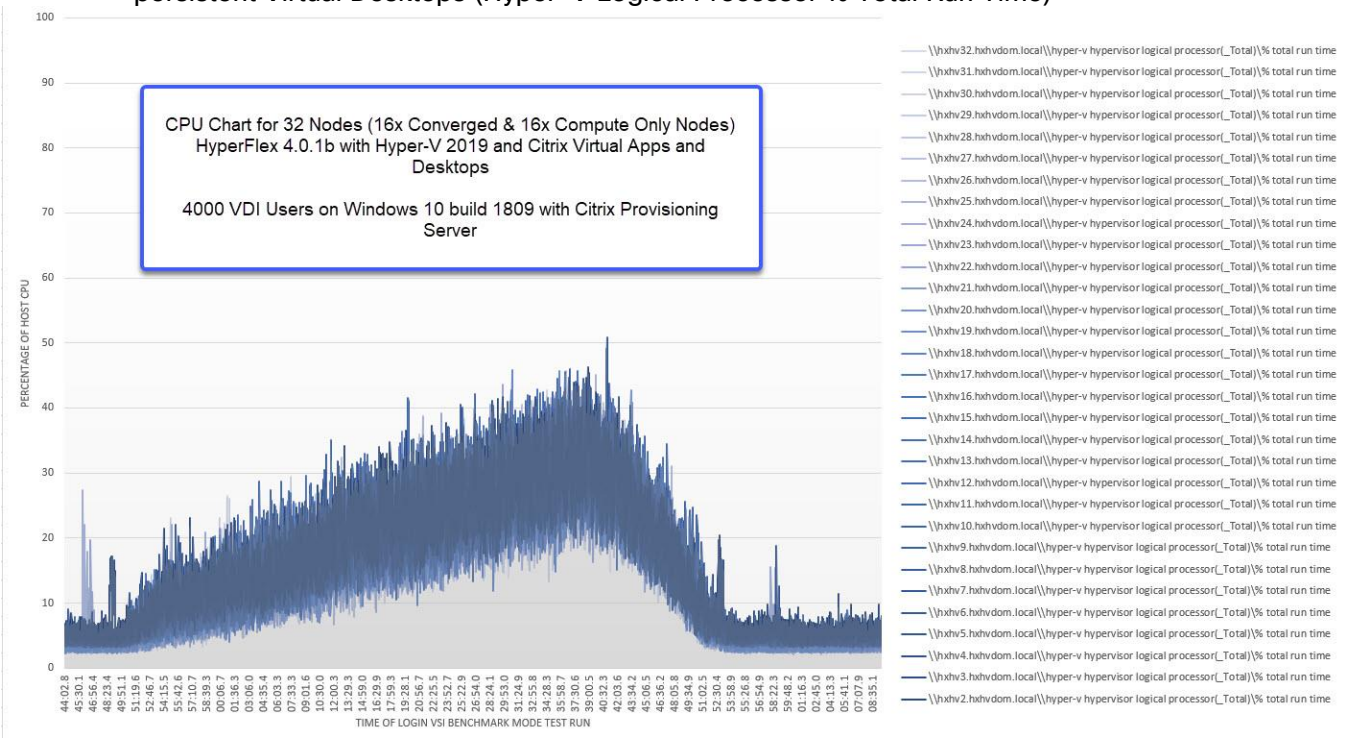

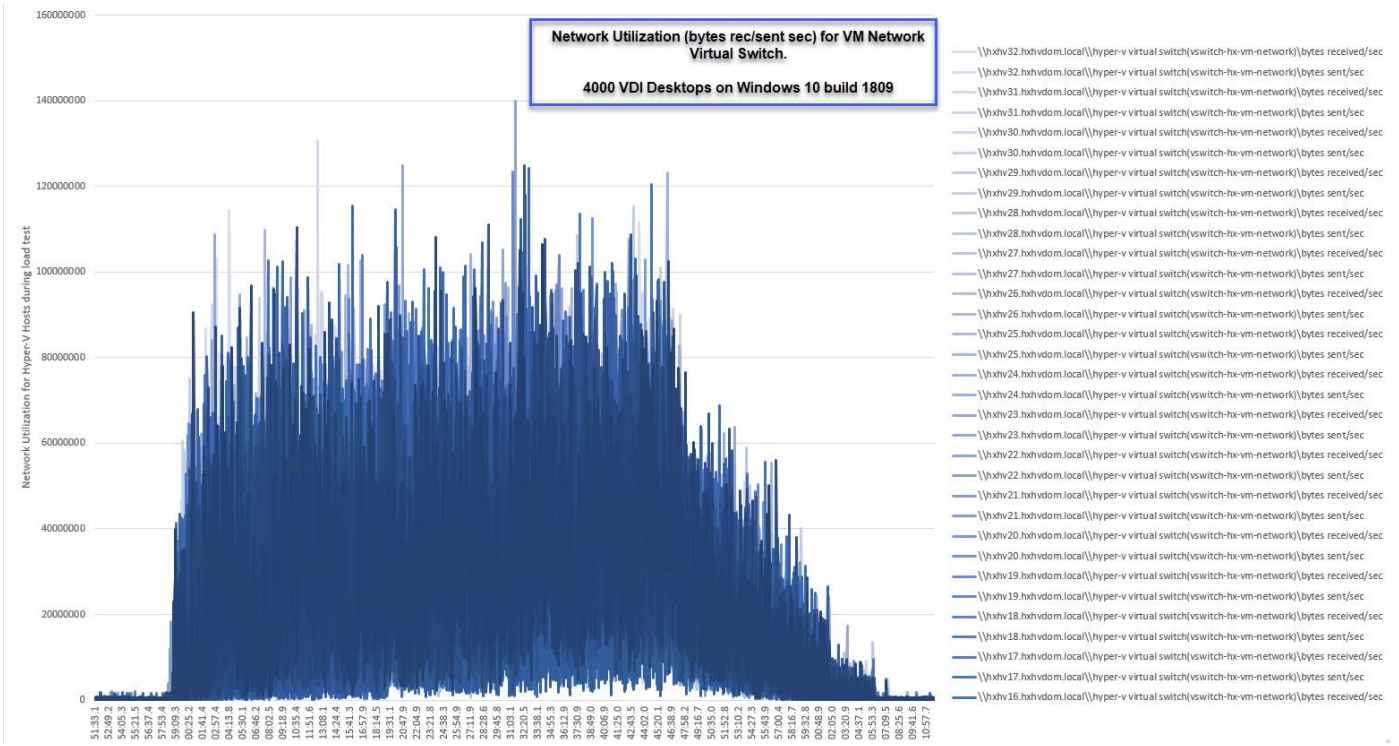
Network Utilization (bytes rec/sent sec) for VM Network Virtual Switch.
4000 VDI Desktops on Windows 10 build 1809
Network Utilization for Hyper-V Hosts during load test
160000000
140000000
120000000
100000000
80000000
60000000
40000000
20000000
0
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv21.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv21.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv20.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv20.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv19.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv19.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv18.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv18.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv17.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec
\\hxhv17.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes sent/sec
\\hxhv16.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-vm-network)\bytes received/sec

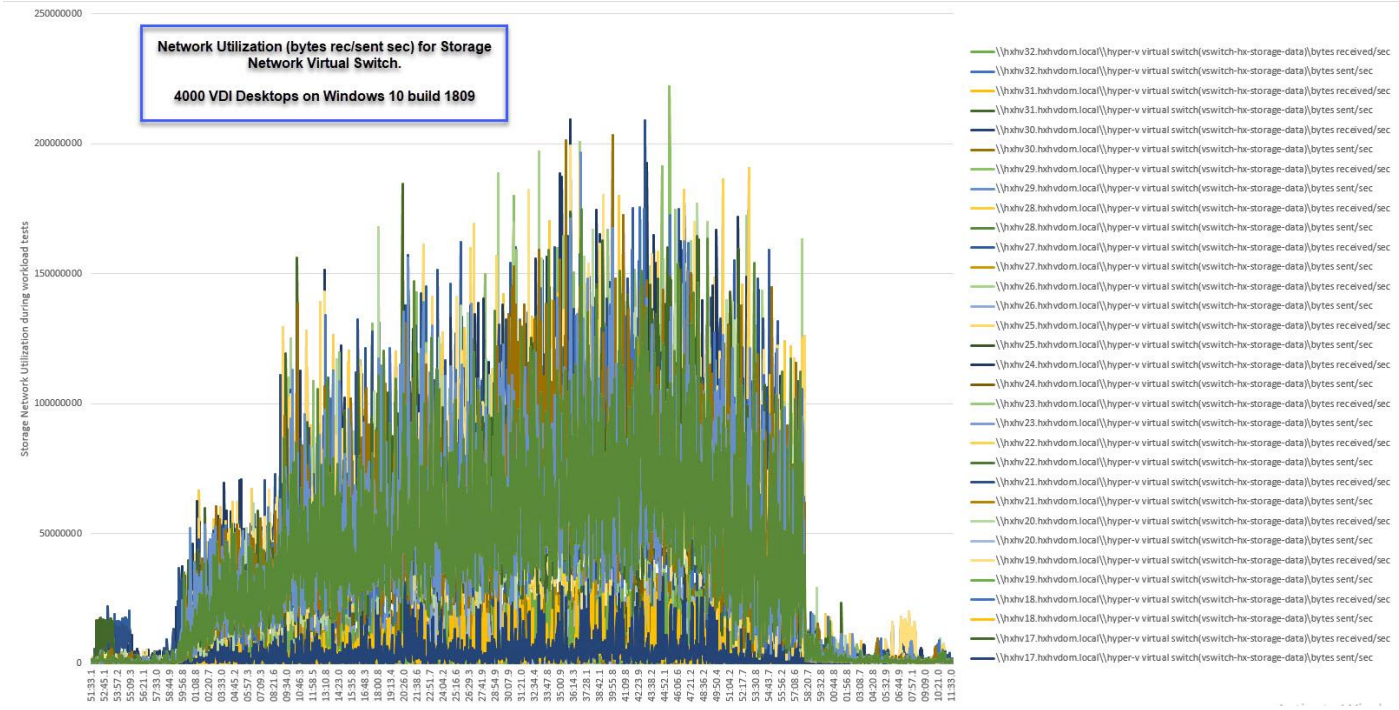
Network Utilization (bytes rec/sent sec) for Storage Network Virtual Switch.
4000 VDI Desktops on Windows 10 build 1809
Storage Network Utilization during workload tests
250000000
200000000
150000000
100000000
50000000
0
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv32.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv31.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv30.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv29.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv28.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv27.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv26.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv25.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv24.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv23.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv22.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv21.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv21.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv20.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv20.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv19.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv19.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv18.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv18.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec
\\hxhv17.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes received/sec
\\hxhv17.hxhvdom.local\\hyper-v virtual switch(vswitch-hx-storage-data)\bytes sent/sec

Figure 133 HyperFlex Cluster UI Performance Chart for Knowledge Worker Workload Running 4000 User Test on Citrix Windows 10 Virtual Desktops

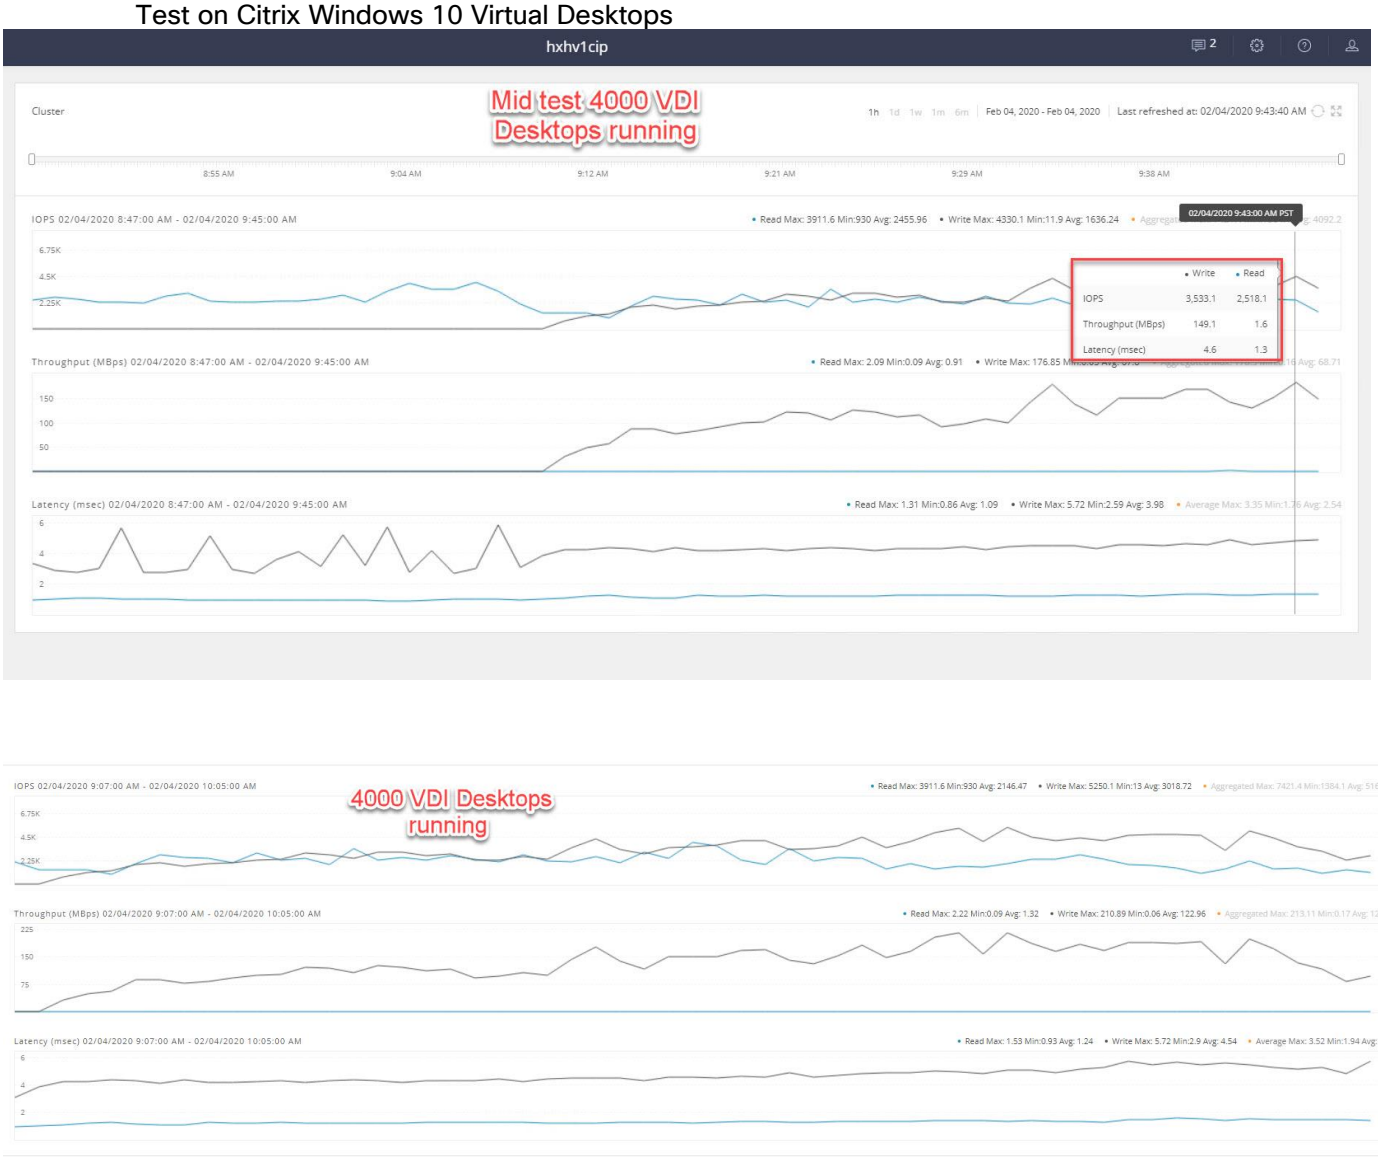

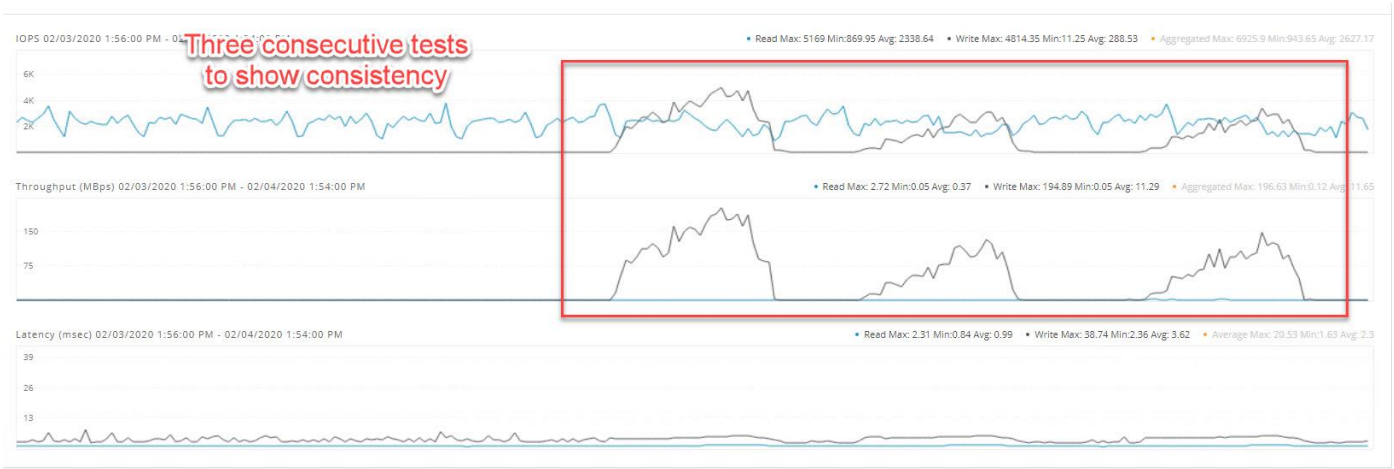
Three consecutive tests to show consistency
IOPS 02/03/2020 1:56:00 PM - 02/04/2020 1:54:00 PM
Read Max: 5169 Min:869.95 Avg: 2338.64
Write Max: 4814.35 Min:11.25 Avg: 288.53
Aggregated Max: 6925.9 Min:943.65 Avg: 2627.17
6K
4K
2K
Throughput (MBps) 02/03/2020 1:56:00 PM - 02/04/2020 1:54:00 PM
Read Max: 2.72 Min:0.05 Avg: 0.37
Write Max: 194.89 Min:0.05 Avg: 11.29
Aggregated Max: 196.63 Min:0.12 Avg: 11.65
150
75
Latency (msec) 02/03/2020 1:56:00 PM - 02/04/2020 1:54:00 PM
Read Max: 2.31 Min:0.84 Avg: 0.99
Write Max: 38.74 Min:2.36 Avg: 3.62
Average Max: 20.53 Min:1.63 Avg: 2.3
39
26
13

# Summary

This Cisco HyperFlex solution addresses 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 hyper-converged system.

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

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

The solution is 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. 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 Citrix Virtual Desktops.

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 have set a new industry standard in performance for Desktop Virtualization on a hyper-converged platform.

# About the Author

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, Virtual Desktops, Virtual Apps 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, the author would like to acknowledge the following for their contribution and expertise that resulted in developing this document:

- Mike Brennan, Product Manager, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.
- Sanjeev Naldurgkar, Technical Marketing Engineer, Cisco UCS Data Center Solutions Group, Cisco Systems, Inc.

# Appendix A – Cisco Nexus 93108YC Switch Configuration

## Switch A Configuration

```
!Command: show running-config

version 7.0(3)I2(2d)
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 1,30-34
vlan 50
  name InBand-Mgmt-C1
vlan 51
  name Infra-Mgmt-C1
vlan 52
  name StorageIP-C1
vlan 53
```

```
  name LiveMigration-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 1000
  peer-keepalive destination 10.29.132.20 source 10.29.132.19



interface Vlan1
  no shutdown
  ip address 10.29.132.2/24

vlan 1,30-34
vlan 30
  name InBand-Mgmt-C1
vlan 31
  name Infra-Mgmt-C1
vlan 32
  name StorageIP-C1
vlan 33
  name LiveMigration-C1
vlan 34
  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 10.29.132.19 source 10.29.132.20



interface Vlan1
  no shutdown
  ip address 10.29.132.3/24

interface Vlan30
  no shutdown
  ip address 10.10.30.3/24
  hsrp version 2
  hsrp 30
    preempt
    priority 110
    ip 10.10.30.1

  ip dhcp relay address 10.10.31.21
  ip dhcp relay address 10.10.31.22

interface Vlan31
  no shutdown
  ip address 10.10.31.3/24
  hsrp version 2
  hsrp 31
```

```
    preempt
    priority 110
    ip 10.10.31.1

interface Vlan32
  no shutdown
  ip address 10.10.32.3/24
  hsrp version 2
  hsrp 32
    preempt
    priority 110
    ip 10.10.32.1

interface Vlan33
  no shutdown
  ip address 10.10.33.3/24
  hsrp version 2
  hsrp 33
    preempt
    priority 110
    ip 10.10.33.1

interface Vlan34
  no shutdown
  ip address 10.34.0.3/20
  hsrp version 2
  hsrp 34
    preempt
    priority 110
    ip 10.34.0.1
  ip dhcp relay address 10.10.31.21
  ip dhcp relay address 10.10.31.22
```

```
interface port-channel10
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  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 1,30-34
  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 1,30-34
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 12

interface Ethernet1/1
  switchport mode trunk

  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active
```

```
interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/4
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/5
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/6
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/7
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  mtu 9216
  channel-group 12 mode active
```

```
interface Ethernet1/8
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  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

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21
```

```
interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  spanning-tree port type edge trunk

interface Ethernet1/26
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  spanning-tree port type edge trunk

interface Ethernet1/27
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  spanning-tree port type edge trunk

interface Ethernet1/28
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34

  spanning-tree port type edge trunk

interface Ethernet1/29
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  spanning-tree port type edge trunk
```

```
interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  spanning-tree port type edge trunk

interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  spanning-tree port type edge trunk

interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  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

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 10.29.132.19/24
```

```
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.I2.2d.bin
```

## Switch B Configuration

```
!Command: show running-config
!Time: Fri Dec 15 17:18:36 2017

version 7.0(3)I2(2d)
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 10.10.30.2
ntp server 10.10.30.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8

vlan 1,30-34
vlan 30
  name InBand-Mgmt-C1
```

```
vlan 31
  name Infra-Mgmt-C1
vlan 32
  name StorageIP-C1
vlan 33
  name LiveMigration-C1
vlan 34
  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 10.29.132.19 source 10.29.132.20


interface Vlan1
  no shutdown
  ip address 10.29.132.3/24

interface Vlan30
  no shutdown
  ip address 10.10.30.3/24
  hsrp version 2
  hsrp 30
    preempt
    priority 110
```

```
    ip 10.10.30.1

  ip dhcp relay address 10.10.31.21
  ip dhcp relay address 10.10.31.22

interface Vlan31
  no shutdown
  ip address 10.10.31.3/24
  hsrp version 2
  hsrp 31
    preempt
    priority 110
    ip 10.10.31.1

interface Vlan32
  no shutdown
  ip address 10.10.32.3/24
  hsrp version 2
  hsrp 32
    preempt
    priority 110
    ip 10.10.32.1

interface Vlan33
  no shutdown
  ip address 10.10.33.3/24
  hsrp version 2
  hsrp 33
    preempt
    priority 110
    ip 10.10.33.1
```

```
interface Vlan34
  no shutdown
  ip address 10.34.0.3/20
  hsrp version 2
  hsrp 34
    preempt
    priority 110
    ip 10.34.0.1
  ip dhcp relay address 10.10.31.21
  ip dhcp relay address 10.10.31.22

interface port-channel10
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  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 1,30-34
  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 1,30-34
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 12

interface Ethernet1/1
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/4
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  channel-group 10 mode active

interface Ethernet1/5
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  mtu 9216
  channel-group 11 mode active
```

```
interface Ethernet1/6
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  mtu 9216
  channel-group 11 mode active

interface Ethernet1/7
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  mtu 9216
  channel-group 12 mode active

interface Ethernet1/8
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  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
```

```
interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  spanning-tree port type edge trunk

interface Ethernet1/26
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  spanning-tree port type edge trunk

interface Ethernet1/27
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
```

```
  spanning-tree port type edge trunk

interface Ethernet1/28
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34

  spanning-tree port type edge trunk

interface Ethernet1/29
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  spanning-tree port type edge trunk

interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  spanning-tree port type edge trunk

interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 1, 30-34
  spanning-tree port type edge trunk

interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,30-34
  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 30

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 10.29.132.20/24
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.I2.2d.bin
```

# Appendix B – Install Microsoft Windows Server 2016

## Install Microsoft Windows Server 2016

The Windows ISO and HyperFlex Driver image files must be placed on a shared location (such as HX Installer) that is reachable from Cisco UCS Manager and the Out-of-band subnet. Use the following steps to download and host the HyperFlex Driver Image and Windows ISO in a shared location within the installer VM.

To install Windows Server 2016/2019 and apply Cisco HyperFlex driver image on all HX nodes, follow these high-level steps.

- Configure Cisco UCS Manager using HyperFlex Installer
- Configure Cisco UCS vMedia and Boot Policies
- Install Microsoft Windows Server 2019 OS
- Undo vMedia and Boot Policy Changes

### Configure Cisco UCS Manager using HyperFlex Installer

To complete this step, refer to HyperFlex Installation.

### Configure Cisco UCS vMedia and Boot Policies

By using a Cisco UCS vMedia policy, the Windows Server 2016 media ISO file and Cisco HyperFlex Driver image can be mounted to all of the HX servers automatically. The existing vMedia policy, named "HyperFlex" must be modified to mount this file, and the boot policy must be modified temporarily to boot from the remotely mounted vMedia file. Once these two tasks are completed, the servers can be rebooted, and they will automatically boot from the When mounted vMedia file, installing and configuring Windows Server 2016 on the HX nodes.

**WARNING!** While vMedia policies are very efficient for installing multiple servers, using vMedia policies as described could lead to an accidental reinstall of Windows on any existing server that is rebooted with this policy. Please be certain that the servers being rebooted while the policy is in effect are the servers you wish to reinstall. Even though the custom ISO will not continue without a secondary confirmation, extreme caution is recommended. This procedure needs to be carefully monitored and the boot policy should be changed back to original settings immediately after the intended servers are rebooted, and the Windows installation begins. Using this policy is only recommended for new installs or rebuilds. Alternatively, you can manually select the boot device using the KVM console during boot, and pressing F6, instead of making the vMedia device the default boot selection.

To configure the Cisco UCS vMedia and Boot Policies, follow these steps:

1. Connect to your HX Installer VM and browse to the folder that contains the /var/www/localhost/images/.
2. Copy the HyperFlex Driver Image "latest.img" (/opt/springpath/packages/latest.img) to the images folder ( /var/www/localhost/images/).
3. Copy *Windows Server 2016 iso* image to the HXDP installer's images folder (/var/www/localhost/images/ )

Figure 134 Upload Windows ISO and Cisco Driver and System Preparation Script

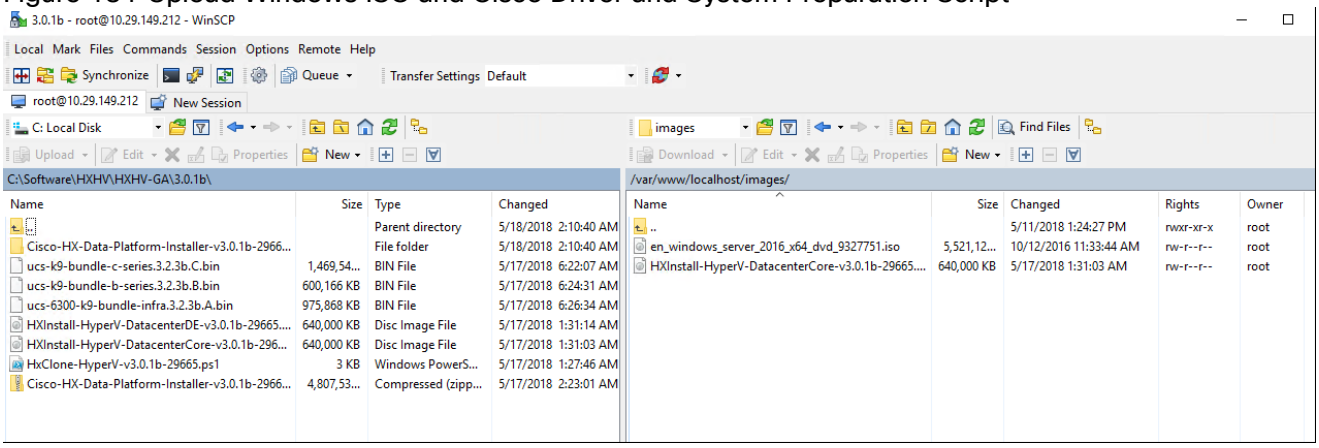


4. For Windows Server 2019, perform the following additional steps to prepare a suitable Driver image for automated OS install:

   a. Copy the HyperFlex Driver Image. For example, run the following command:

   ```
   rsync -avzP /opt/springpath/packages/latest.img
   /var/www/localhost/images/install.img
   ```

```
root@HyperFlex-Installer:~# rsync -avzP /opt/springpath/packages/latest.img /var/www/localhost/images/install.img
sending incremental file list
latest.img
    655,360,000 100%   26.52MB/s    0:00:23 (xfr#1, to-chk=0/1)

sent 583,007,126 bytes  received 35 bytes  23,796,210.65 bytes/sec
total size is 655,360,000  speedup is 1.12
root@HyperFlex-Installer:~# _
```

   b. Mount the HyperFlex Driver Image. For example, run the following command:

   ```
   mkdir -p /mnt/install-img && mount -o loop,rw
   /var/www/localhost/images/install.img /mnt/install-img
   ```

   c. Copy the answer file specific to Windows Server 2019. For example, run the following command:

   ```
   cp
   /opt/springpath/packages/FactoryUnattendXML/WindowsServer2019/Autounattend.xml.ro
   /mnt/install-img/Autounattend.xml
   ```

   d. Unmount the HyperFlex Driver Image. For example, run the following command:

   ```
   umount /mnt/install-img
   ```

```
root@HyperFlex-Installer:~# cp /opt/springpath/packages/FactoryUnattendXML/WindowsServer2019/Autounattend.xml.ro /mnt/install-img/Autounattend.x
ml
root@HyperFlex-Installer:~# umount /mnt/install-img/
root@HyperFlex-Installer:~#
```

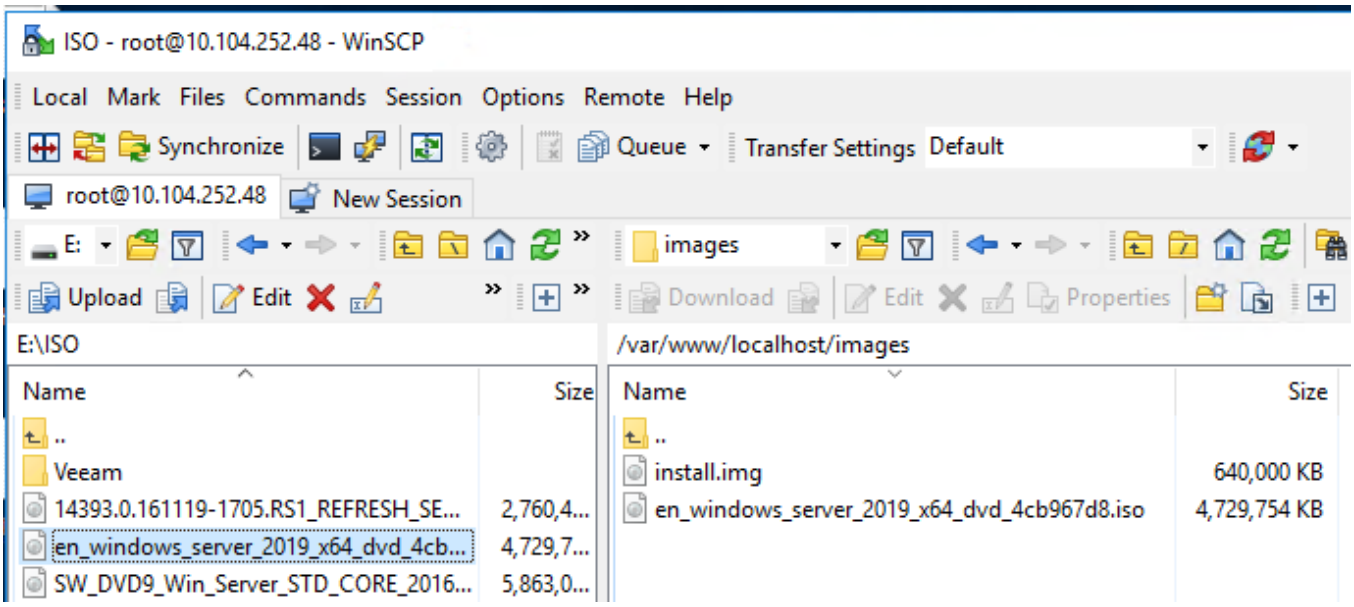


Make sure network connectivity exists between the file share and all server management IPs.

5. Configure the vMedia and Boot policies using Cisco UCS Manager to mount the above images
6. Launch Cisco UCS Manager by accessing the Cisco UCS Manager IP address in a browser of your choice.
7. Click Launch UCS Manager and log in with administrator username and the password you used at the beginning of the installation.
8. In the left navigation pane, click Servers.
9. Expand Servers > Policies > root > Sub-Organizations > hx-cluster_name>vMedia Policies to view the list of vMedia Policies.

Figure 135 Cisco UCS Manager – Create vMedia Policy

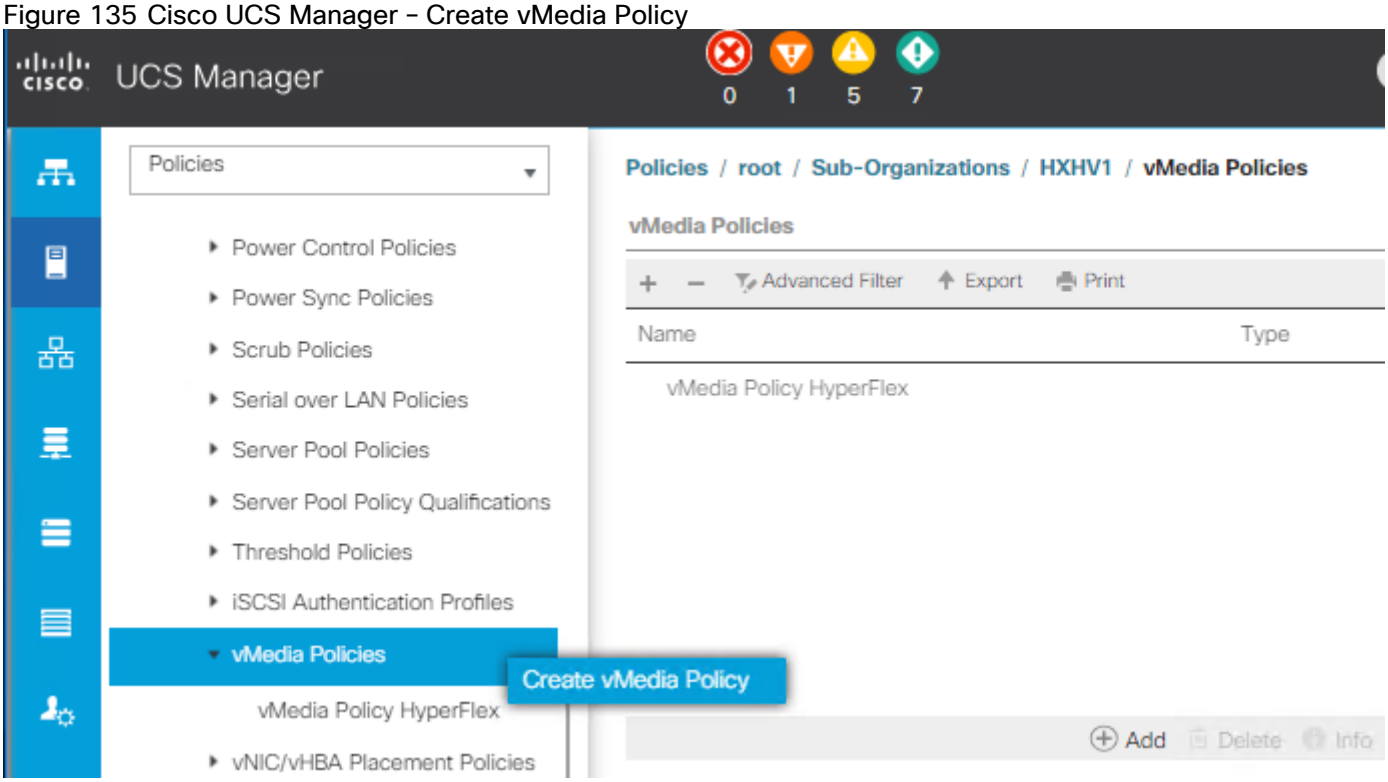


10. Double-click vMedia Policy HyperFlex.

11. In the properties for vMedia Policy HyperFlex, click Create vMedia Mount to add the mount points.

12. In the Create vMedia Mount dialog box, complete the following fields in Figure 136.

Table 41 Create vMedia Mount Details

| Field Name | Action | Example Value |
|---|---|---|
| Name | Name for the mount point. | Windows -ISO |
| Description | Can be used for more information. | |
| Device Type | Type of image that you want to mount | CDD |
| Protocol | The protocol used for accessing the share where the ISO files are located. | HTTP |
| Hostname/IP Address | IP address or FQDN of the server hosting the images. | 10.104.252.48 |
| Image Name Variable | This value is not used in HyperFlex installation. | None |
| Remote File | The filename of the ISO file that you want to mount. | |

| Field Name | Action | Example Value |
|---|---|---|
| Remote Path | The path on the remote server to where the file resides | |
| Username | If you use CIFS or NFS a username might be necessary | |
| Password | If you use CIFS or NFS a password might be necessary | |

Figure 136 Cisco UCS Manager – Create vMedia Mount CDD

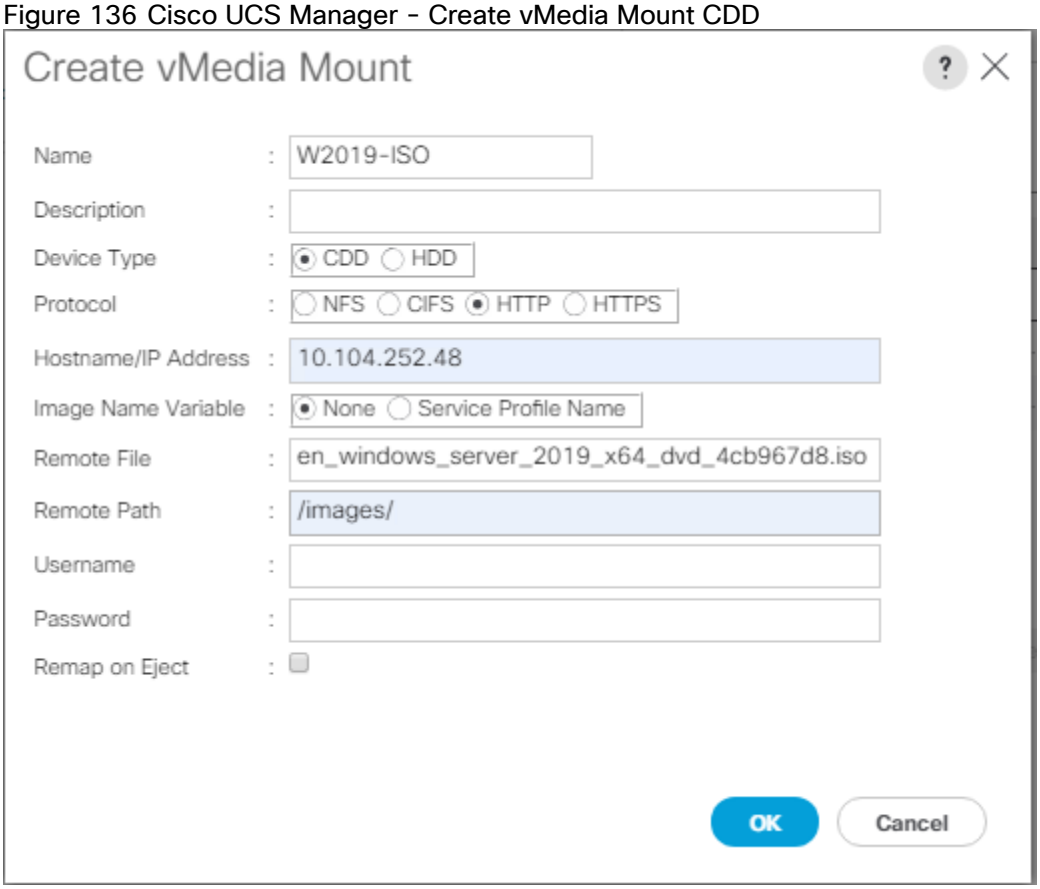


13. Click Save Changes and click OK.

14. Click OK. When you click OK, you are returned to the vMedia policy and will see the information that you submitted.

15. Repeat steps 5 and 6 but change the type to HDD and the filename to the Cisco HyperFlex driver image.

Figure 137 Cisco UCS Manager – Create vMedia Mount HDD

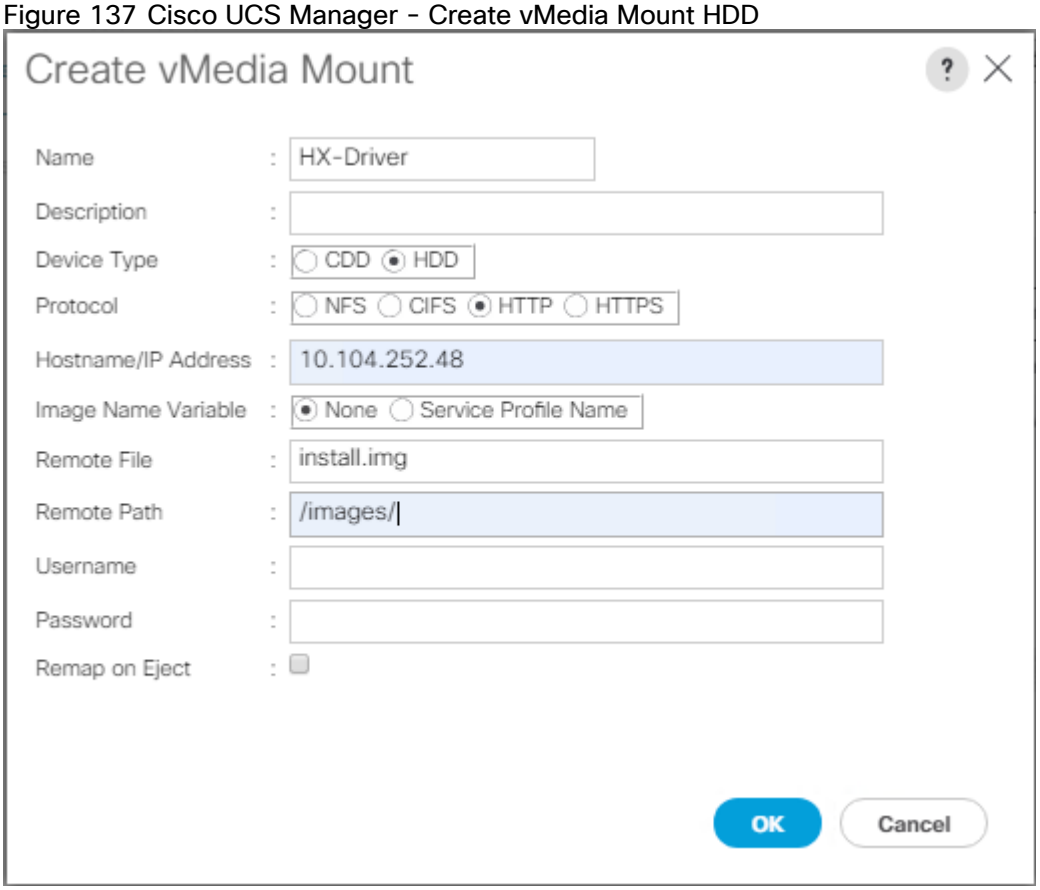


16. On completion, the following screen displays:

Figure 138 Cisco UCS Manager – Create vMedia Policy

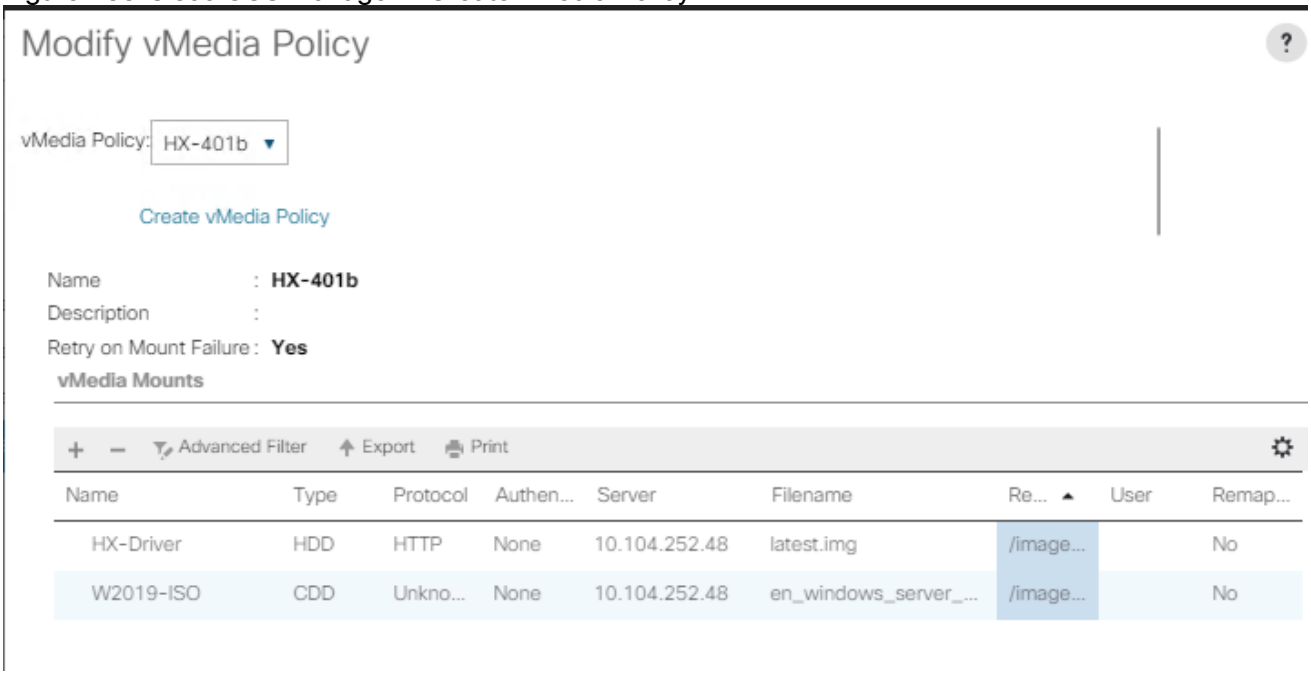

17. In the left navigation pane, select Servers > Service Profile Templates > root > Sub-Organizations > hx-cluster_name > Service Template hx-nodes_name (example: compute-nodes-m5).

Figure 139 Cisco UCS Manager – Service Template

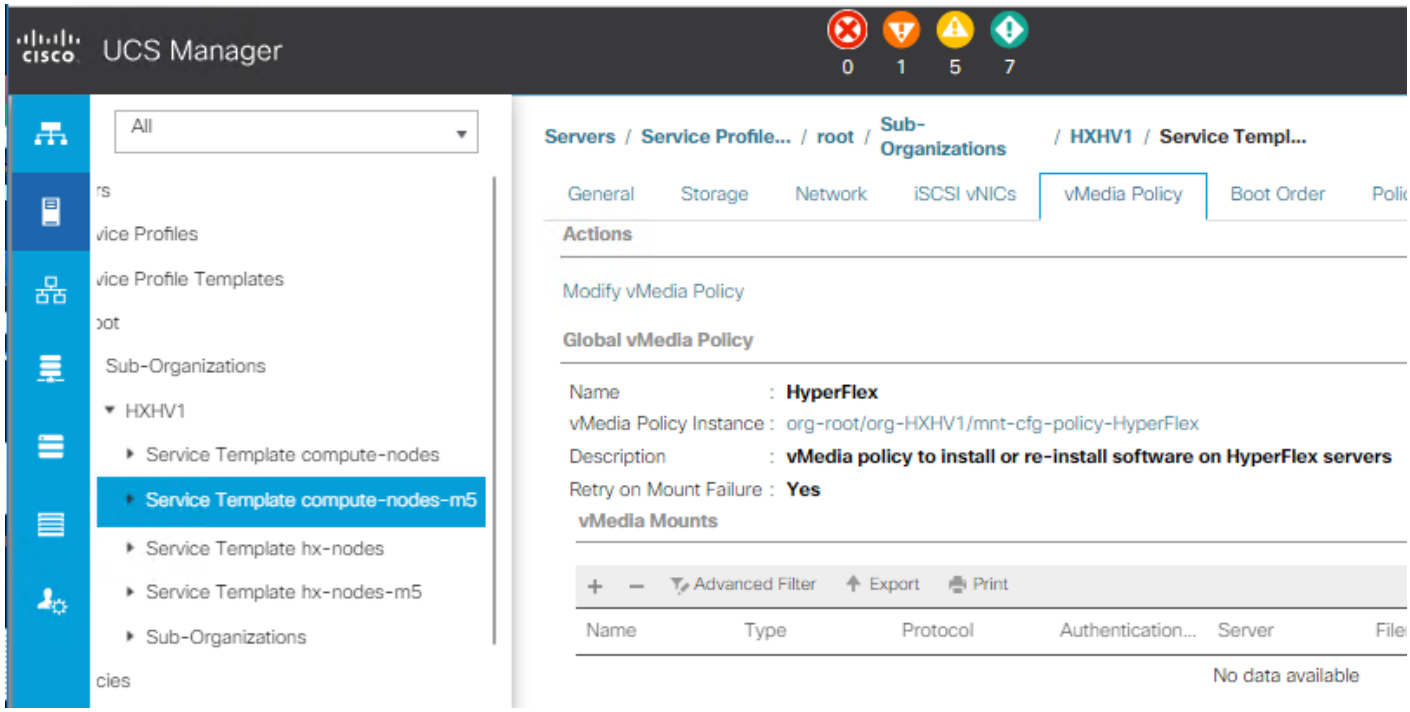


18. Choose the HyperFlex vMedia Policy from the drop-down list and click OK twice.

Figure 140 Cisco UCS Manager – Modify vMedia Policy

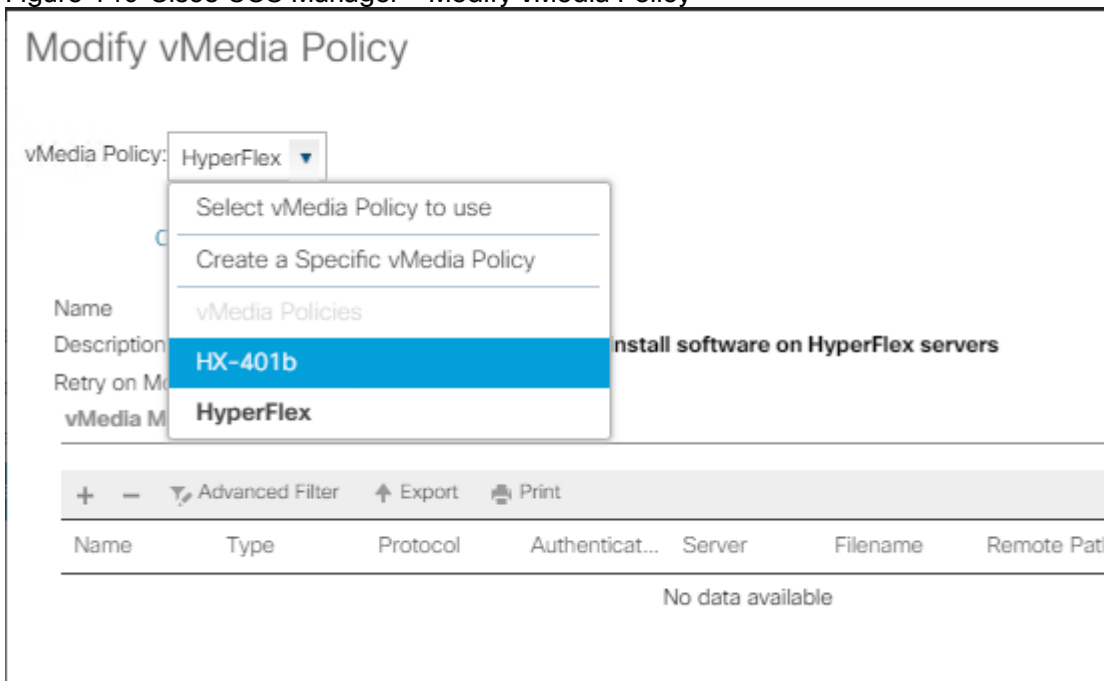

**The vMedia policy is assigned to the HyperFlex Template during the Cisco UCS Manager phase of the HyperFlex deployment.**

19. Select Servers > Policies > root > Sub-Organizations > hx-cluster_name > Boot Policies Boot Policy HyperFlex-m5.

20. In the configuration pane, click CIMC Mounted vMedia. Click Add CIMC Mounted CD/DVD to add this to the boot order.

21. Select the CIMC Mounted CD/DVD entry in the list and move it to the top of the boot order by clicking Move Up.

Figure 141 Cisco UCS Manager – Boot Order

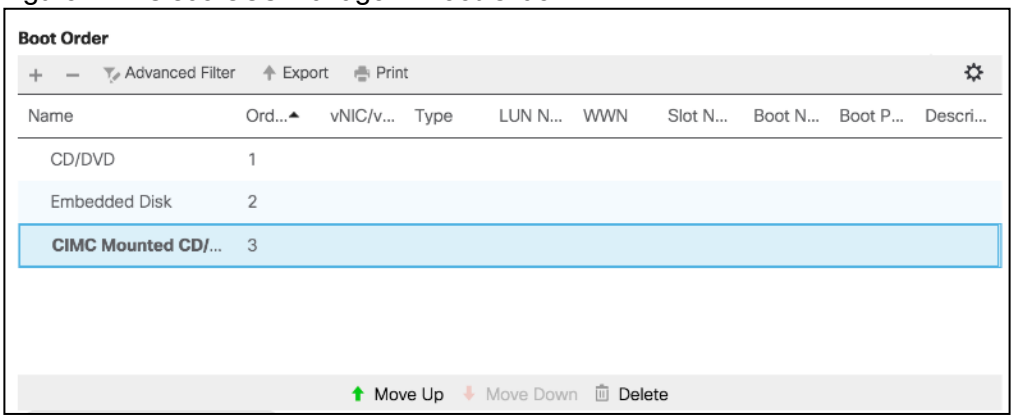


22. Click Save Changes and click OK. The boot policy is saved.

To verify the images are mounted correctly, follow these steps:

1. On the Equipment tab, select one of the servers.

2. Click Inventory > CIMC, scroll down and make sure for the mount entry #1(OS image) and mount entry #2 (Cisco HyperFlex driver image) the status is Mounted and there are no failures.

Figure 142 Cisco UCS Manager – Validate vMedia Mount

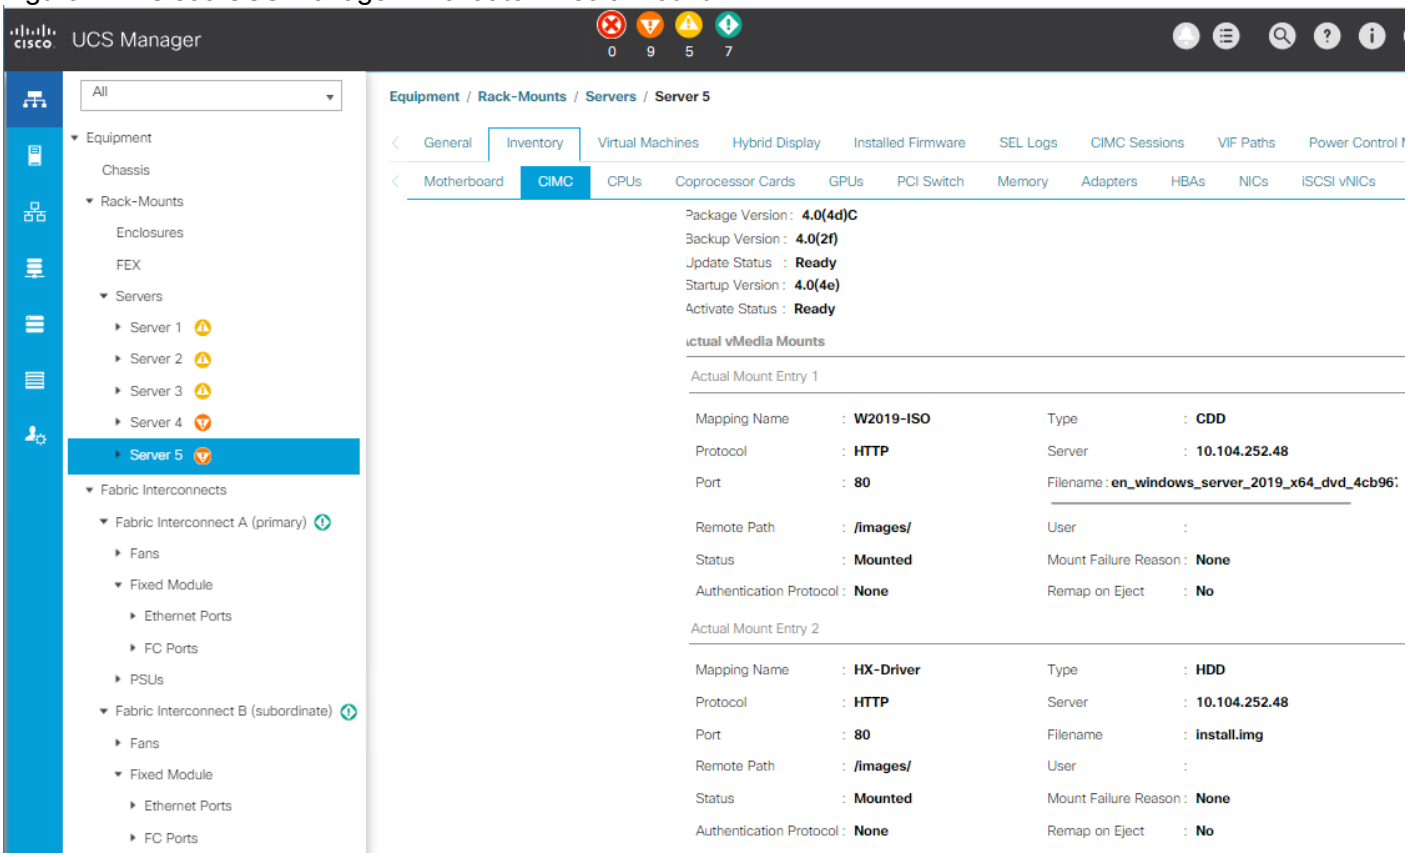


## Install Microsoft Windows Server 2019 OS

To install Microsoft Windows Server 2016 OS, follow these steps:

1. In the menu bar, click Servers and choose the first HyperFlex service profile.
2. Click the General tab and choose Actions > KVM Console.

The KVM console will try to open in a new browser. Be aware of any pop-up blockers. Allow the Pop-Ups and re-open the KVM.

Figure 143 Cisco UCS Manager – Launch KVM Console

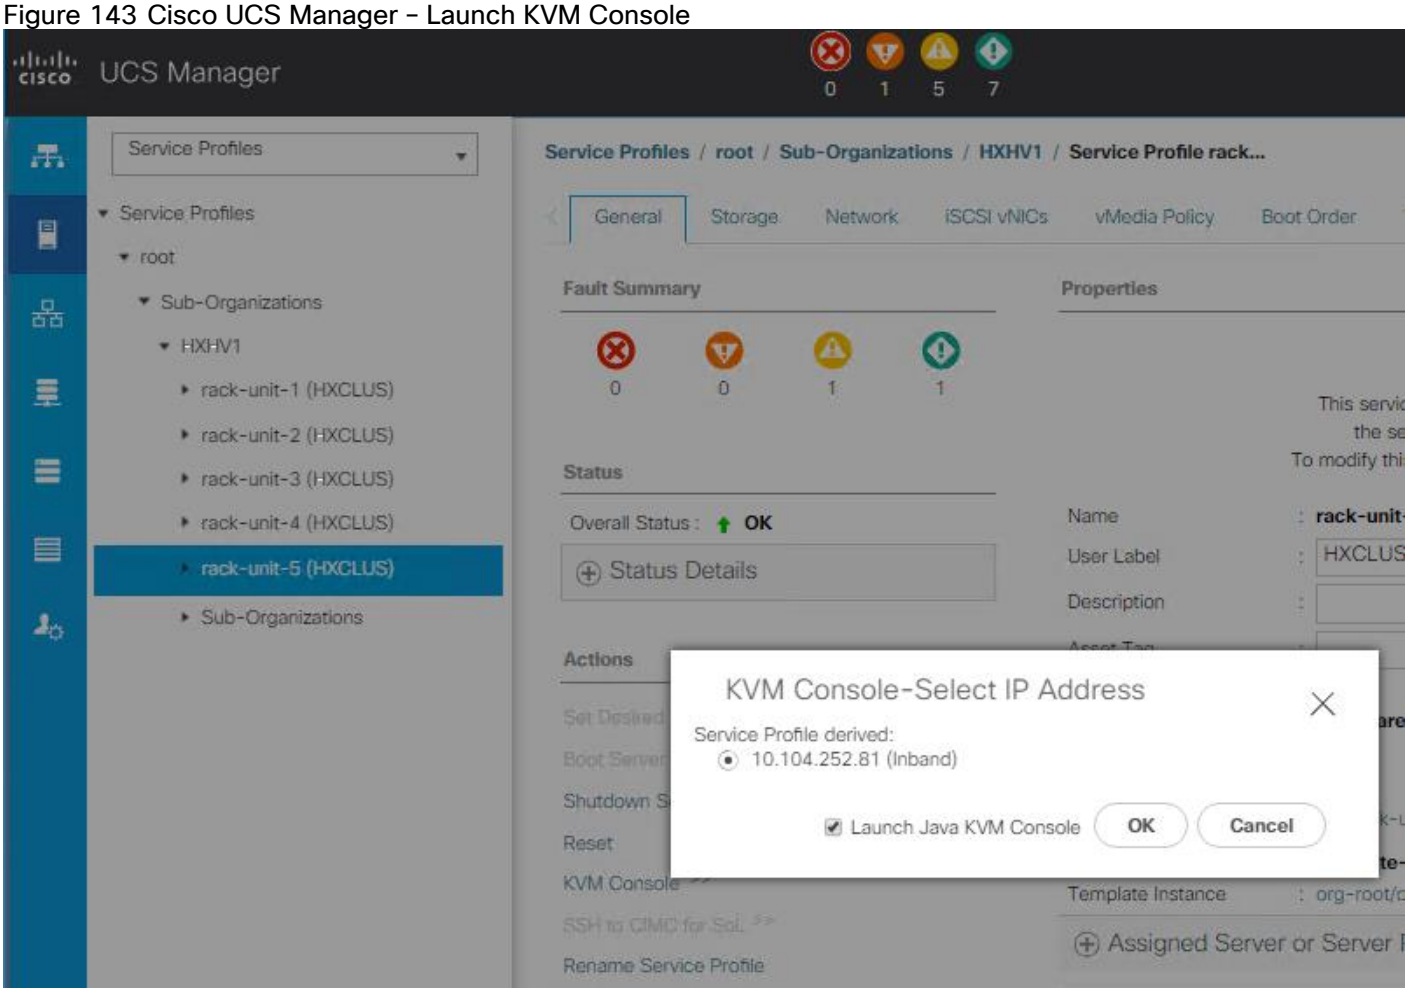


3. Reboot the server. In the KVM console choose Server Actions and click Reset.

Figure 144 Cisco UCS Manager – Server KVM Console

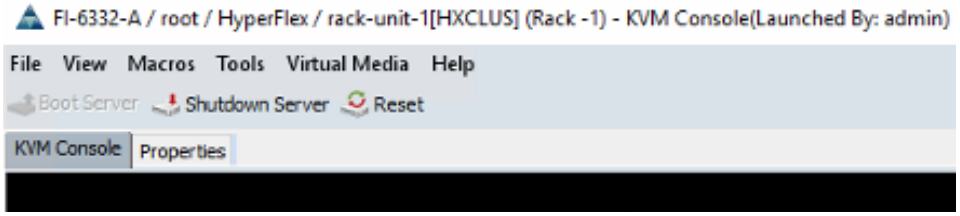


4. Choose Power Cycle.

5. When the server is rebooting, remember to press any key to start the Windows installation.

Figure 145 Cisco UCS Manager – KVM Console Server Boot

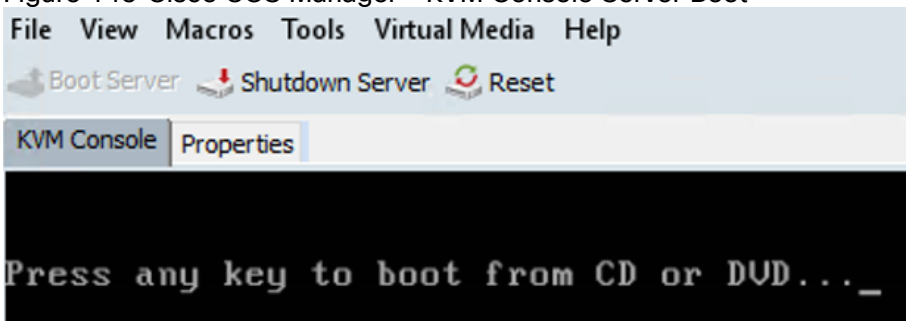


If you miss clicking any key, the server will display in the windows installation or an error page displays stating no OS is installed.

6. When the Windows installation is complete, you will see some tasks running as shown in the below and the host will reboot a few times. Allow some time for the system preparation to complete.

Figure 146 System Preparation

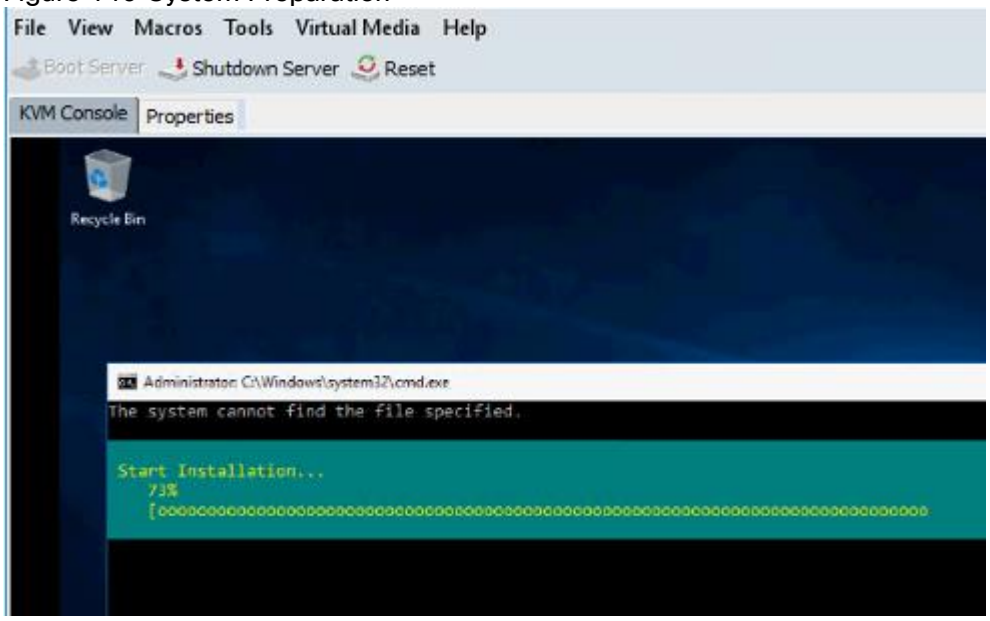


7. The installation is complete when a clean command prompt with no activity is displayed as shown below.

Figure 147 Windows Server Command Prompt

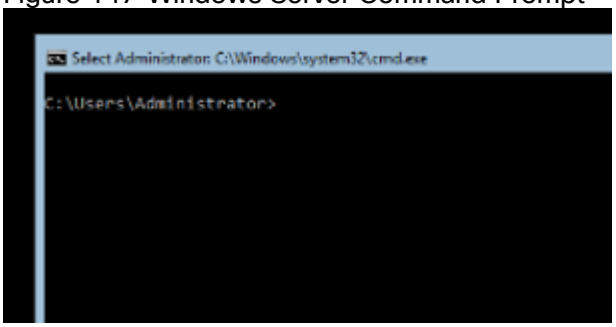

8. Repeat steps 1-7 on all the HX nodes in the cluster and verify the below task is running as shown in below. The 'HXInstallbootstraplauncherTask' in running state is an indication of successful installation Windows OS and system preparation.

Figure 148 Validate Windows Server Installation Completion

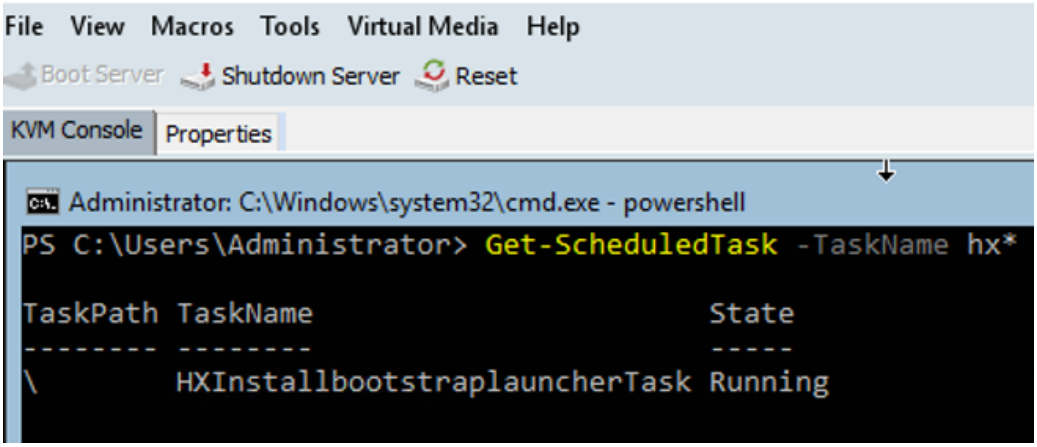


## Undo vMedia and Boot Policy Changes

To undo vMedia and boot policy changes, follow these steps:

1. When all the servers have booted from the remote vMedia file and begun their installation process, the changes to the boot policy need to be quickly undone, to prevent the servers from going into a boot loop, constantly booting from the installation ISO file. To revert the boot policy settings, follow these steps: In Cisco UCS Manager select Servers > Polices > Root > Sub-Organizations > `HX-Cluster_name` > vMedia polices.

2. Click the vMedia Policy HyperFlex. Click the mount points one at a time and delete both of them. Accept the changes.

Figure 149 Cisco UCS Manager – Undo vMedia and Boot Policy Changes

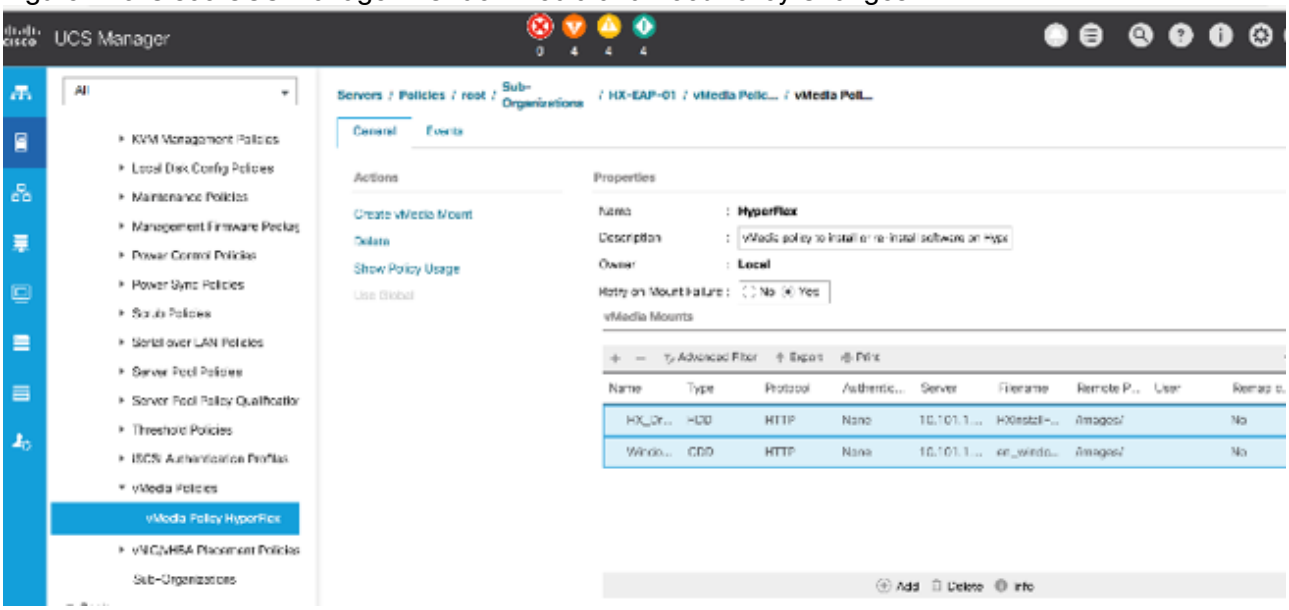

3. Go to the boot policy by selecting Servers > Polices > Root > Sub-Organizations > HX-Cluster_name > boot polices > Boot Policy HyperFlex-m5.

4. Select the CIMC mounted CD/DVD, click Delete and accept the changes.