

# Cisco HyperFlex M5 All-Flash Hyperconverged System with Hyper-V 2016 and Citrix XenDesktop

Design and Deployment of Cisco HyperFlex for Virtual  
Desktop Infrastructure with Citrix XenDesktop 7.16

Last Updated: July 11, 2018



# 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, refer 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)

© 2018 Cisco Systems, Inc. All rights reserved.

# Table of Contents

Executive Summary .....	14
Solution Overview .....	16
Introduction .....	16
Audience .....	16
Purpose of this Document.....	16
<b>What's New?</b> .....	16
Solution Summary.....	18
Cisco Desktop Virtualization Solutions: Data Center .....	19
The Evolving Workplace.....	19
Cisco Desktop Virtualization Focus .....	21
Use Cases .....	23
Physical Topology.....	25
Fabric Interconnects .....	27
HX-Series Rack Mount Servers.....	28
Cisco UCS B-Series Blade Servers.....	28
Logical Network Design .....	29
Configuration Guidelines.....	31
Solution Design.....	32
Cisco Unified Computing System.....	32
Cisco Unified Computing System Components.....	32
Enhancements for Version 3.0.1 .....	33
New Software Features .....	33
New Hardware Features .....	34
Supported Versions and System Requirements .....	35
Hardware and Software Interoperability.....	35
Software Requirements for Microsoft Hyper-V .....	35
Cisco UCS Fabric Interconnect .....	35
Cisco HyperFlex HX-Series Nodes .....	37
Cisco HyperFlex Compute Nodes .....	43
Cisco UCS B200-M5 Blade .....	43
Cisco VIC1340 Converged Network Adapter .....	45
Cisco UCS 5108 Blade Chassis .....	46
Cisco UCS 2304XP Fabric Extender .....	46
Cisco UCS C220-M5 Rack Server .....	47



Cisco UCS C240-M5 Rack Server .....	48
Cisco HyperFlex Converged Data Platform Software .....	48
Cisco HyperFlex Connect HTML5 Management Web Page .....	49
Cisco Intersight Management Web Page .....	49
Cisco Nexus 93108YCPX Switches .....	57
Architectural Flexibility .....	57
Feature-Rich .....	57
Real-Time Visibility and Telemetry .....	58
Highly Available and Efficient Design .....	58
Simplified Operations .....	58
Investment Protection .....	58
Microsoft Hyper-V 2016 .....	60
Microsoft System Center 2016 .....	60
<b>Citrix XenApp™ and XenDesktop™ 7.16</b> .....	60
Zones .....	62
Improved Database Flow and Configuration .....	62
Application Limits .....	63
Multiple Notifications before Machine Updates or Scheduled Restarts .....	63
API Support for Managing Session Roaming .....	63
API Support for Provisioning VMs from Hypervisor Templates .....	63
Support for New and Additional Platforms .....	64
Architecture and Design of Citrix XenDesktop on Cisco Unified Computing System and Cisco HyperFlex Storage Design Fundamentals .....	65
Understanding Applications and Data .....	66
Project Planning and Solution Sizing Sample Questions .....	67
Citrix XenDesktop Design Fundamentals .....	68
Machine Catalogs .....	68
Delivery Groups .....	68
Example XenDesktop Deployments .....	69
Distributed Components Configuration .....	69
Multiple Site Configuration .....	70
Citrix Cloud Services .....	71
Designing a XenDesktop Environment for a Mixed Workload .....	71
Deployment Hardware and Software .....	73
Products Deployed .....	73
Hardware Deployed .....	75

Software Deployed .....	75
Logical Architecture.....	76
VLANs .....	77
Jumbo Frames.....	78
Solution Configuration.....	82
Cisco UCS Compute Platform .....	82
Physical Infrastructure .....	82
Cisco Unified Computing System Configuration .....	85
Deploy and Configure HyperFlex Data Platform .....	86
Prerequisites .....	86
Deploying HX Data Platform Installer on Hyper-V Infrastructure .....	91
Deploy the HX Data Platform Installer OVA with a Static IP Address.....	98
Cisco UCS Manager Configuration using HX Data Platform Installer .....	100
Microsoft Windows OS and Hyper-V Installation.....	107
Configure vMedia and Boot Policies through Cisco UCS Manager .....	108
Next Steps .....	116
Hypervisor Configuration .....	123
Deploying HX Data Platform Installer and Cluster Configuration.....	128
Cluster Validation.....	135
Post HyperFlex Cluster Installation for Hyper-V 2016.....	137
Create Run As Account in VMM .....	137
Servers .....	137
Networking .....	138
Storage.....	140
Build the Virtual Machines and Environment for Workload Testing.....	143
Software Infrastructure Configuration .....	143
Prepare the Master Images.....	144
Install and Configure XenDesktop Delivery Controller, Citrix Licensing, and StoreFront.....	145
Install Citrix License Server.....	145
Install Citrix Licenses .....	149
Install the XenDesktop .....	150
Configure the XenDesktop Site.....	156
Configure the XenDesktop Site Administrators .....	164
Configure Additional XenDesktop Controller .....	166
Add the Second Delivery Controller to the XenDesktop Site .....	169
Install and Configure StoreFront.....	171

Additional StoreFront Configuration .....	180
Install XenDesktop Virtual Desktop Agents.....	186
Persistent Static Provisioned with MCS .....	192
Create Delivery Groups.....	200
Citrix XenDesktop Policies and Profile Management .....	204
Configure Citrix XenDesktop Policies.....	205
Configuring User Profile Management .....	205
Test Setup and Configurations.....	207
Testing Methodology and Success Criteria .....	208
Testing Procedure .....	208
Pre-Test Setup for Testing .....	209
Test Run Protocol .....	209
Success Criteria .....	210
Test Results.....	215
Boot Storms.....	215
Recommended Maximum Workload and Configuration Guidelines.....	215
Four Node Cisco HXAF220c-M5S Rack Server and HyperFlex All-Flash Cluster .....	215
Summary .....	220
About the Authors.....	221
Acknowledgements .....	221
Appendix A – Cisco Nexus 93108YC Switch Configuration.....	222
Switch A Configuration .....	222
Switch B Configuration .....	233



## 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 a 450 user mixed workload on a 4-node (4 Cisco HyperFlex HXAF220C-M5SX server) Cisco HyperFlex system. We provide deployment guidance and performance data for Citrix XenDesktop 7.16 virtual desktops running Microsoft Windows 10 with Office 2016 Machine Creation Services. The solution is a pre-integrated, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches and Cisco HyperFlex Data Platform software version 3.0.1a.

The solution payload is 100 percent virtualized on Cisco HyperFlex HXAF220C-M5SX hyperconverged nodes booting via on-board M.2 SATA SSD drive running Microsoft Hyper-V 2016 hypervisor and the Cisco HyperFlex Data Platform storage controller VM. The virtual desktops are configured with XenDesktop 7.16, 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 or Server 2016 based desktops. The solution provides unparalleled scale and management simplicity. Citrix XenDesktop Provisioning Services or Machine Creation Services Windows 10 desktops (450,) or full clone desktops (450) or XenApp server based desktops (600) can be provisioned on a four node Cisco HyperFlex cluster. Where applicable, this document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

The solution boots 450 virtual desktops machines in 5 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 1000 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 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 XenDesktop.

The solution provides outstanding virtual desktop end user experience as measured by the Login VSI 4.1.25 Knowledge Worker workload running in benchmark mode. Index average end-user response times for all tested delivery methods is under 1 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 hyper converged solution can be quickly deployed, thereby increasing agility and reducing costs. Cisco HyperFlex uses best of breed storage, server and network components to serve as the foundation for desktop virtualization workloads, enabling efficient architectural designs that can be quickly and confidently deployed and scaled-out.

### Audience

The 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 four different Citrix XenDesktop/XenApp workloads with Cisco UCS 6300 series Fabric Interconnects and Cisco Nexus 9300 series switches.

### What's New?

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

- Validation of Cisco Nexus 9000 with Cisco HyperFlex Support for the Cisco UCS 3.2(3) release and Cisco HyperFlex Data Platform v 3.0.1a.
- Microsoft Hyper-V 2016

Citrix XenDesktop 7.16 Persistent desktops with Citrix Machine Creation Services

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

For a complete list of hardware and software inter-dependencies, refer to the Cisco UCS Manager web page [Hardware and Software Interoperability for Cisco HyperFlex HX-Series](#).

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

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

The use cases include:

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



## Solution Summary

---

This Cisco Validated Design prescribes a defined set of hardware and software that serves as an integrated foundation for both Citrix XenDesktop Microsoft Windows 10 virtual desktops and Citrix XenApp server desktop sessions based on Microsoft Server 2016. The mixed workload solution includes Cisco HyperFlex hardware and Data Platform software, Cisco Nexus® switches, the Cisco Unified Computing System (Cisco UCS®), Citrix XenDesktop and Microsoft Hyper-V software in a single package. The design is efficient such that the networking, computing, and storage components occupy an 8-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 18-core 2.3 GHz Intel Xeon (Gold 6140) Scalable Family processors with 768GB of 2666Mhz memory with Citrix XenDesktop support more virtual desktop workloads than the previously released generation processors on the same hardware. The Intel Xeon Gold 6140 18-core scalable family processors used in this study provided a balance between increased per-server capacity and cost.
- Fault-tolerance with high availability built into the design. The various designs are based on multiple Cisco HX-Series nodes, Cisco UCS rack servers and Cisco UCS blade servers for virtual desktop and infrastructure workloads. The design provides N+1 server fault tolerance for every payload type tested.
- Stress-tested to the limits during aggressive boot scenario. The 450 user mixed hosted virtual desktop and 600 user hosted shared desktop environment booted and registered with the XenDesktop 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 450 users logged in and started running workloads up to steady state in 48-minutes without overwhelming the processors, exhausting memory or exhausting the storage subsystems, providing customers with a desktop virtualization system that can easily handle the most demanding login and startup storms.
- Ultra-condensed computing for the datacenter. The rack space required to support the initial 450 user system is 8 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects. Incremental Citrix XenDesktop 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 2016. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, Citrix XenDesktop components, XenDesktop VDI desktops and XENAPP 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 VMs 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 datacenter management: Cisco maintains industry leadership with the new Cisco UCS Manager 3.2(2) software that simplifies scaling, guarantees consistency, and eases maintenance. **Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director** insure that customer environments are consistent locally, across Cisco UCS Domains and across the globe. Cisco UCS software suite offers increasingly simplified operational and deployment management, and it continues to widen the span of control for customer organizations' **subject matter** experts in compute, storage and network.
- Cisco 40G Fabric: Our 40G unified fabric story gets additional validation on 6300 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- Cisco HyperFlex Connect (HX Connect): An all-new HTML 5 based Web UI Introduced with HyperFlex v2.5 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

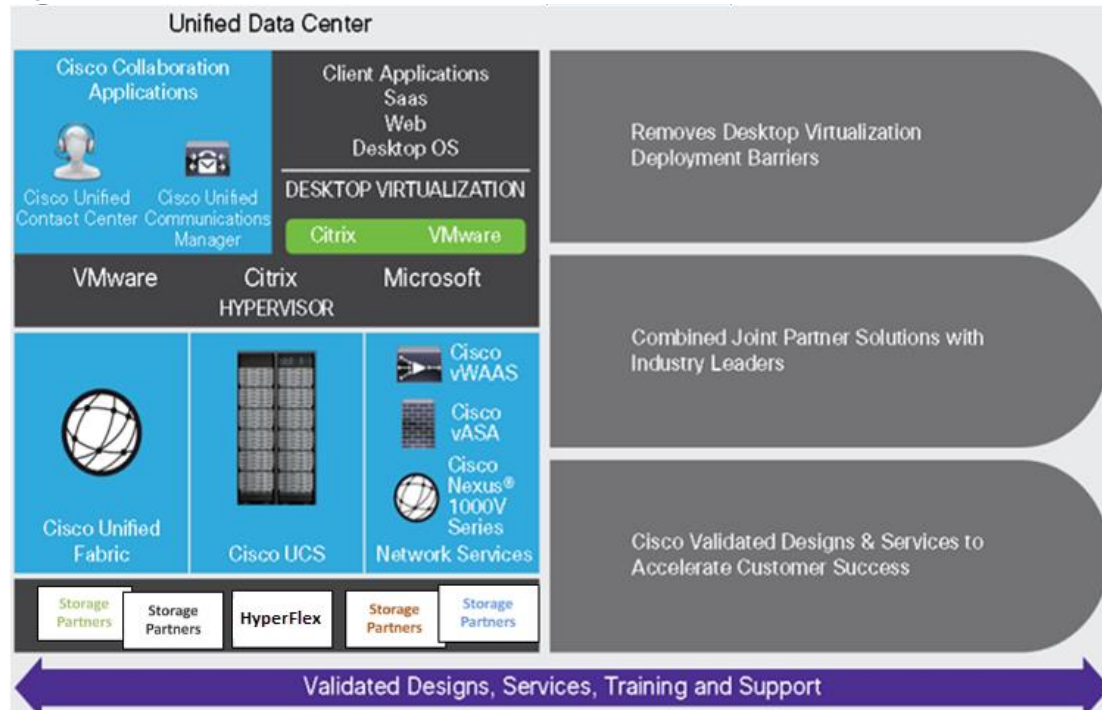
### The Evolving Workplace

**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 1).

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

Figure 1 Cisco Data Center Partner Collaboration



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

## Cisco Desktop Virtualization Focus

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

### Simplified

Cisco UCS and Cisco HyperFlex provide a radical new approach to industry-standard computing and provides the core of the data center infrastructure for desktop virtualization. Among the many features and benefits of Cisco UCS are the drastic reduction in the number of servers needed 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 XenDesktop, Cisco HyperFlex solutions have demonstrated scalability and performance, with up to 450 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 through the use of Cisco's scalable, flexible, and secure** architecture as the platform for desktop virtualization.

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

## Use Cases

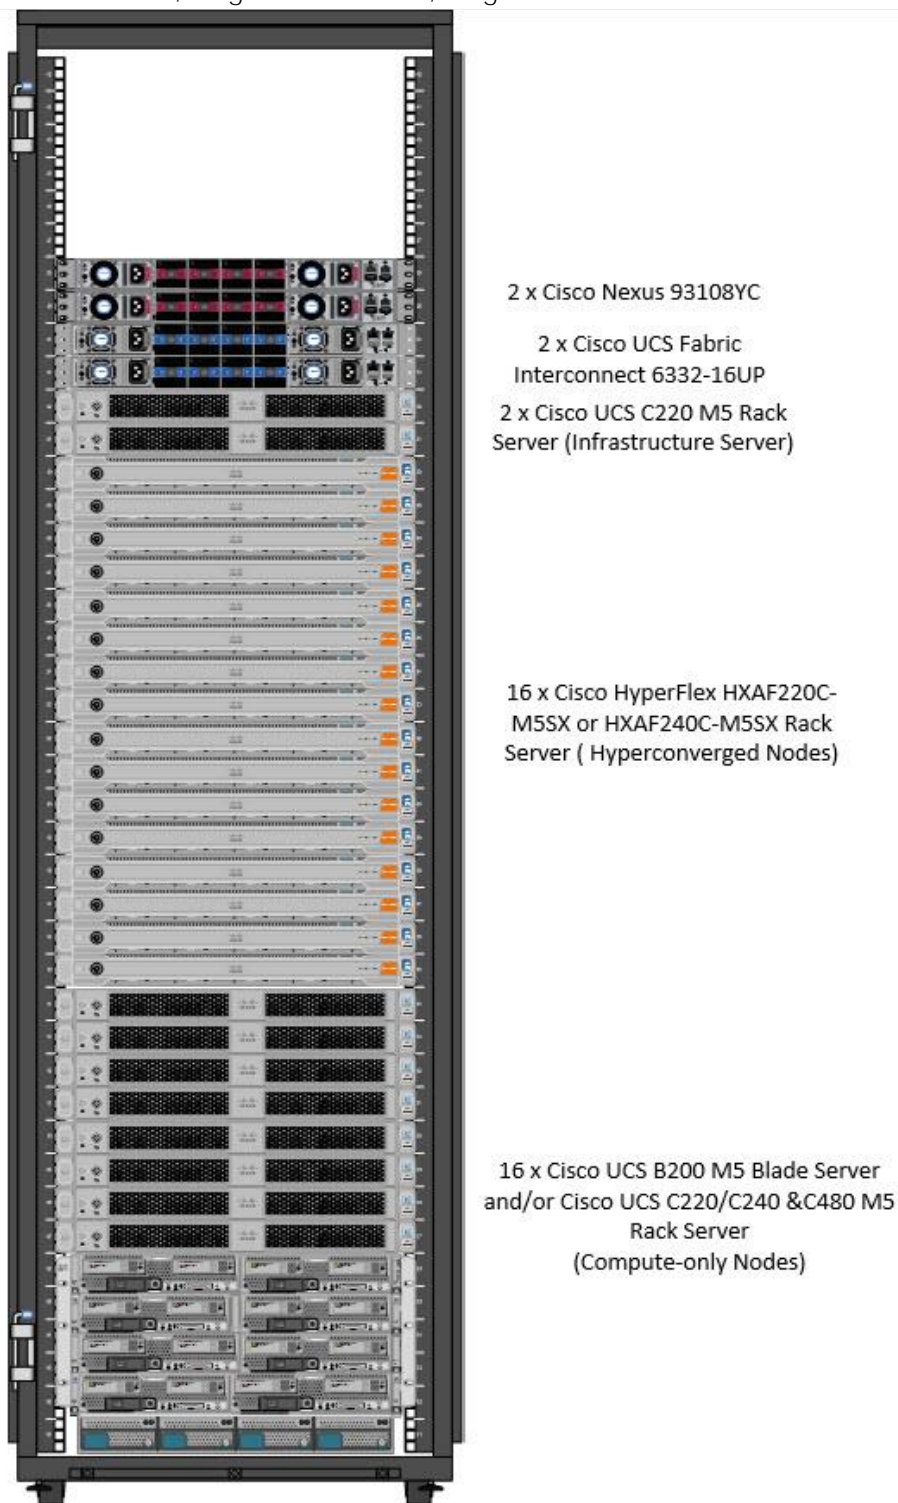
- Healthcare: Mobility between desktops and terminals, compliance, and cost
- Federal government: Teleworking initiatives, business continuance, continuity of operations (COOP), and training centers
- Financial: Retail banks reducing IT costs, insurance agents, compliance, and privacy
- Education: K-12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

Figure 2 shows the Citrix XenDesktop on Hyper-V 2016 built on Cisco Validated Design components and the network connections. The reference architecture reinforces the "wire-once" strategy, because as additional

storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.



Figure 2 Full Scale, Single UCS Domain, Single Cisco Rack Architecture



## 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. In addition, up to 16 compute only servers

can be added per cluster. Adding Cisco UCS 5108 Blade chassis allows use of Cisco UCS B200-M5 blade servers for additional compute resources in a hybrid cluster design. Cisco UCS C240 and C220 servers can also be used for additional compute resources. 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 datacenter network at the time of installation.



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

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

Figure 3 Cisco HyperFlex Standard Topology

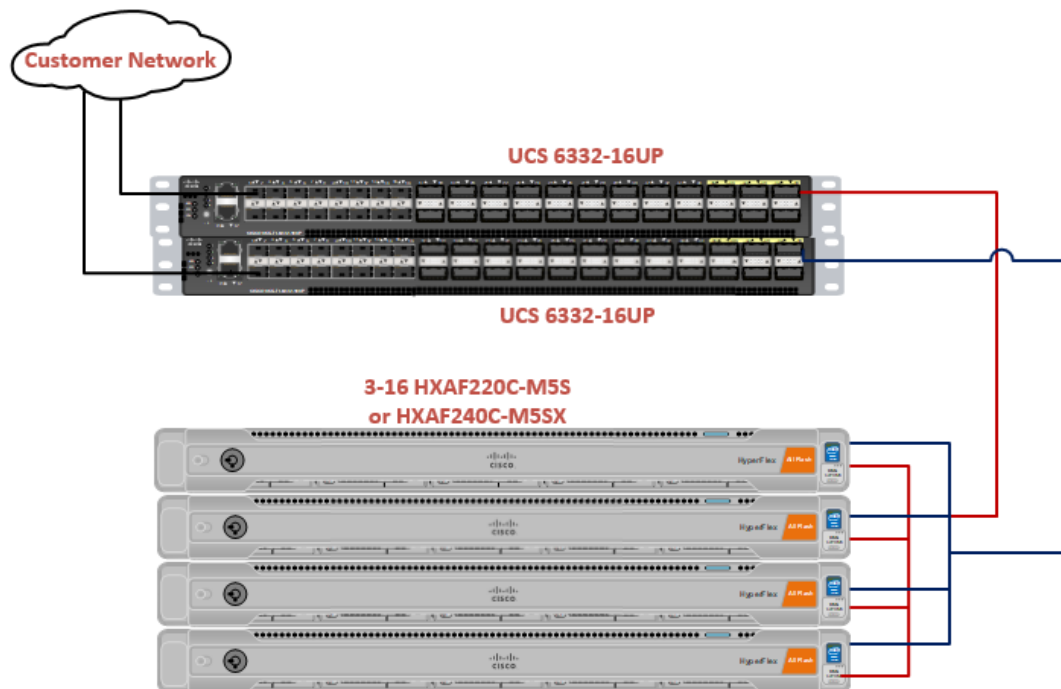
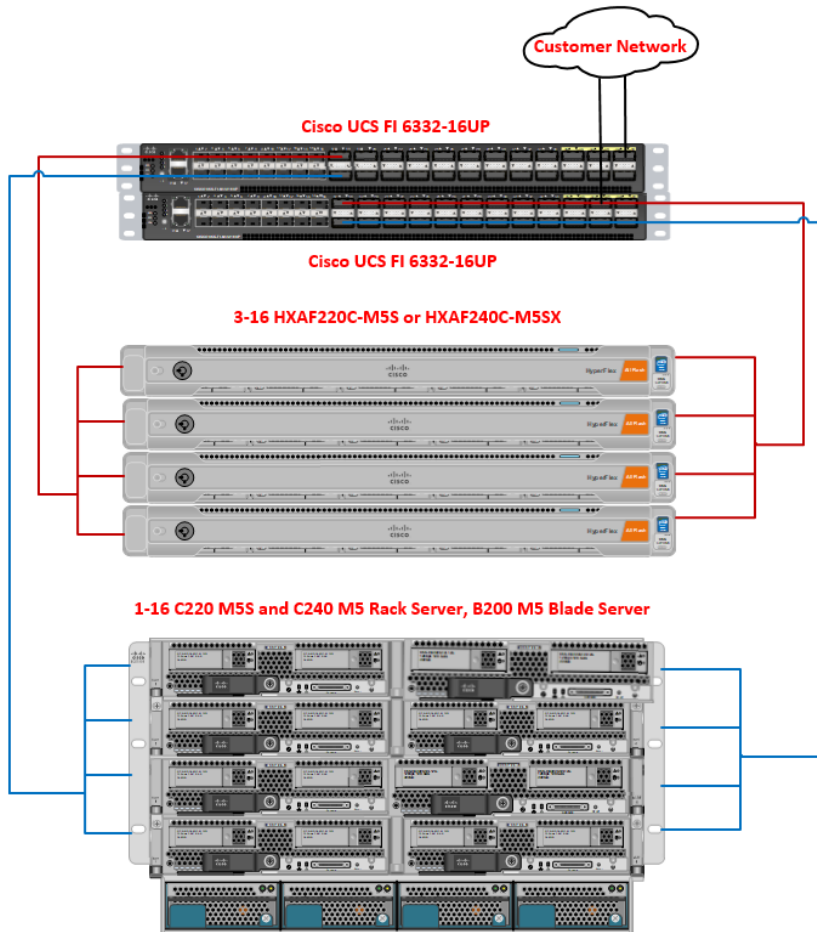


Figure 4 Cisco HyperFlex Hyperconverged plus Compute Only Node Topology



## Fabric Interconnects

Fabric Interconnects (FI) are deployed in pairs, wherein the two units operate as a management cluster, while forming two separate network fabrics, referred to as the A side and B side fabrics. Therefore, many design elements will refer to FI A or FI B, alternatively called fabric A or fabric B. Both Fabric Interconnects are active at all times, passing data on both network fabrics for a redundant and highly available configuration. Management services, including Cisco UCS Manager, are 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 via GUI and CLI tools. Also used by remote KVM, IPMI and SoL sessions to the managed servers within the domain. This is typically connected to the customer management network.
- **L1:** A cross connect port for forming the Cisco UCS management cluster. This is connected directly to the L1 port of the paired Fabric Interconnect using a standard CAT5 or CAT6 Ethernet cable with RJ45 plugs. It is not necessary to connect this to a switch or hub.

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


## HX-Series 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 5).



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

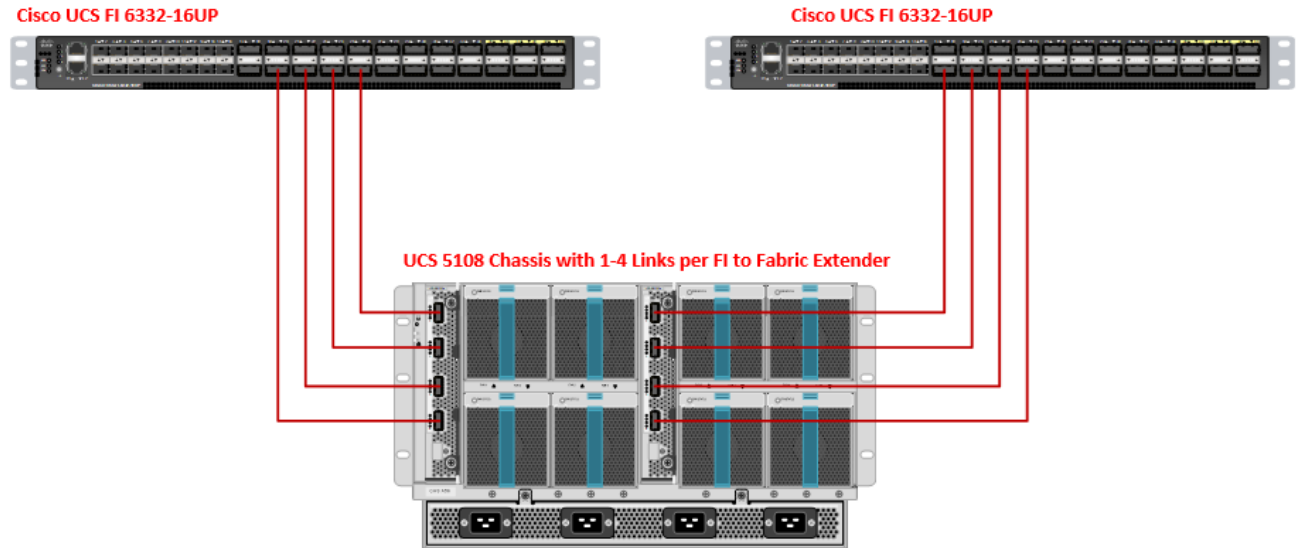
Figure 5 HX-Series Server Connectivity



## Cisco UCS B-Series Blade Servers

Hybrid HyperFlex clusters also incorporate 1-8 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 10 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 10 GbE links from the right side IOM, or IOM 2, to FI B (Figure 6). All other cabling configurations are invalid, and can lead to errors, discovery failures, and loss of redundant connectivity.

Figure 6 Cisco UCS 5108 Chassis Connectivity



## Logical Network Design

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

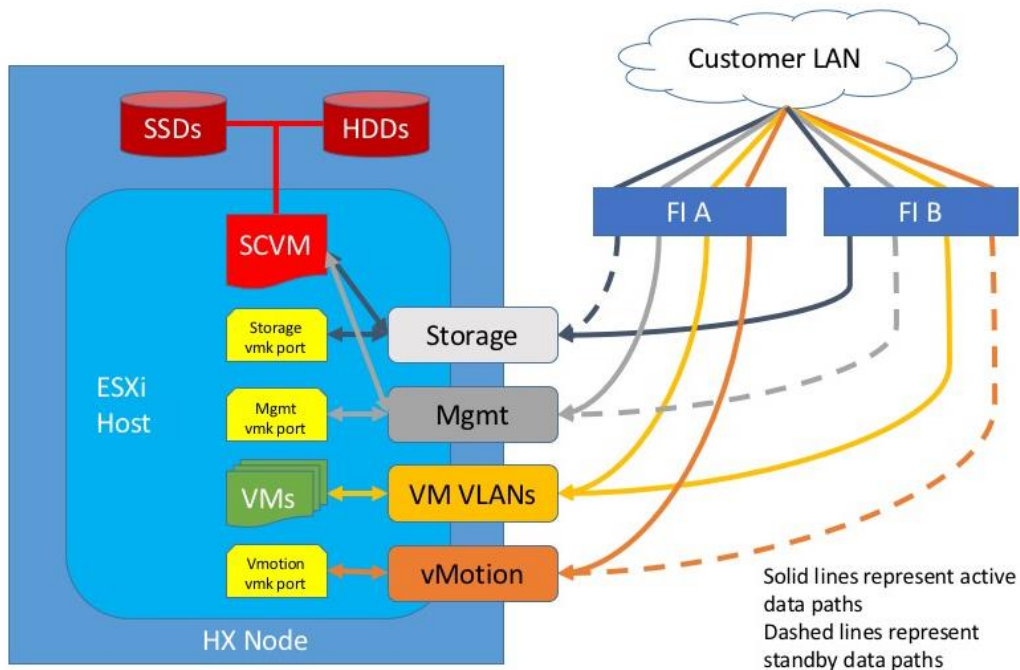
- Management Zone: This zone comprises the connections needed to manage the physical hardware, the hypervisor hosts, and the storage platform controller virtual machines (SCVM). These interfaces and IP addresses need to be available to all staff who will administer the HX system, throughout the LAN/WAN. This zone must provide access to Domain Name System (DNS) and Network Time Protocol (NTP) services, and allow Secure Shell (SSH) communication. In this zone are multiple physical and virtual components:
  - Fabric Interconnect management ports.
  - Cisco UCS external management interfaces used by the servers and blades, which answer through the FI management ports.
  - Hyper-V host management interfaces.
  - Storage Controller VM management interfaces.
  - A roaming HX cluster management interface.
- VM Zone: This zone comprises the connections needed to service network IO to the guest VMs 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 VMs 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 VMs 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 VM 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 VMs from host to host. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B, and vice-versa.

Figure 7 illustrates the logical network design.

Figure 7 Logical Network Design



The reference hardware configuration includes:

- Two Cisco Nexus 93108YCPX switches
- Two Cisco UCS 6332-16UP fabric interconnects
- Four Cisco HX-series Rack server running HyperFlex data platform version 3.0.1a.



For desktop virtualization, the deployment includes Citrix XenDesktop 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:

- 450 Citrix XenDesktop Windows 10 non-persistent virtual desktops using MCS



All of 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 2. These procedures covers everything from physical cabling to network, compute and storage device configurations.

## Configuration Guidelines

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



## Solution Design

---

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

### Cisco Unified Computing System

Cisco UCS Manager (UCSM) provides unified, embedded management of all software and hardware **components of 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 of 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](#).

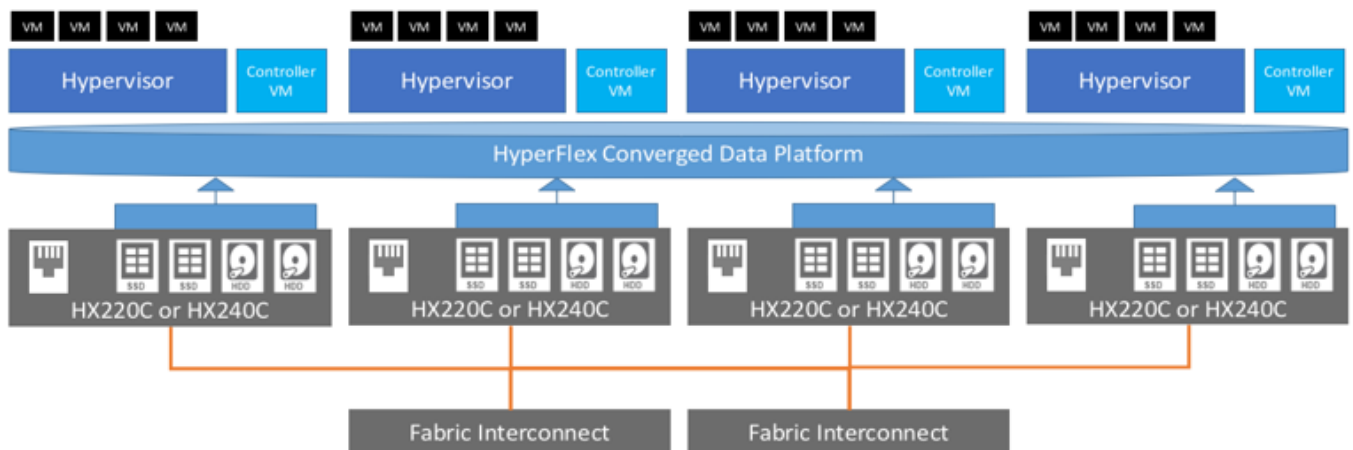
Cisco UCS and Cisco HyperFlex are designed to deliver:

- 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 VM storage, and the hypervisor software for running the virtualized servers, all within a single Cisco UCS management domain.

Figure 8 Cisco HyperFlex System Overview



## Enhancements for Version 3.0.1

### New Software Features

This release delivers key advancements in mission critical and cloud-native workload support.

- Multiple Hypervisors—Support for Microsoft Hyper-V and Microsoft Hyper-V Server 2016. For detailed information on installation requirements and procedures, see Cisco HyperFlex Systems Installation Guide for VMware, and [Cisco HyperFlex Systems Installation Guide on Microsoft Hyper-V](#).

- Stretched clusters—Ability to stretch a single HyperFlex Cluster across two datacenters enabling the highest level of availability. For more information, see [Cisco HyperFlex Systems Stretched Cluster Guide](#).
- Kubernetes FlexVolume driver—Enables the automated provisioning of persistent volumes to Kubernetes pods.
- Improved resiliency— Enabled by Logical Availability Zones (LAZ) that, when enabled, automatically partition the cluster so that it is more resilient to node and disk failure scenarios. This feature can only be enabled on HyperFlex clusters with 8 or more converged nodes. For a LAZ enabled cluster, expansion must be performed one node at a time. For more information, see [Open Caveats in Release 3.0\(1a\)](#).
- Disaster Recovery Workflow—Enhanced Disaster Recovery workflow (Planned and unplanned VM migration) using Cisco HX Connect. For more information, see [Cisco HyperFlex Systems Administration Guide, Release 3.0](#).
- REST APIs—Additional developer guidance in the form of a quick start guide for HyperFlex REST APIs on Cisco DevNet. For more information see, [Cisco HyperFlex Systems REST API Getting Started Guide](#).
- Linked mode—**HyperFlex Plugin Support for environments utilizing vCenter’s enhanced linked mode feature.**

### New Hardware Features

- Enhanced HX Capacity Scaling options
  - Large Form Factor (LFF) HX M5 240 chassis with support up to 12 drives 6TB or 8 TB drives. Note: Currently supported in HX M5 240 nodes only. For more information, see [Cisco HyperFlex HX-Series Spec Sheets](#)
  - Support for 1.8 TB SFF HDD with maximum cluster capacity of 18.06TiB for Hybrid M5 Edge nodes only. For more information, see [Cisco HyperFlex HX-Series Spec Sheets](#)
  - Enhanced Node Scaling options—Support for up to 32-nodes converged (per cluster) with 32 compute nodes. [Cisco HyperFlex HX-Series Spec Sheets](#).
- Intel Optane Support for higher drive level performance and higher endurance—HyperFlex has qualified the latest flash memory innovation, 3D XPoint. Added Intel Optane NVMe SSD HX-NVMEXP-I375 as a new caching drive option.



Supported in M5 All Flash configurations only. For more information, see [Cisco HyperFlex HX-Series Spec Sheets](#), Cisco HX220c M5 HyperFlex Node Installation Guide, and Cisco HX240c M5 HyperFlex Node Installation Guide.

---

- Support for AMD Multiuser GPU (MxGPU) hardware-based graphics virtualization on HX240c M5 nodes—AMD FirePro S7150 series GPUs are now available for HX240c M5 nodes. These graphic accelerators enable highly secure, high performance, and cost effective VDI deployments. For more information see, see [Cisco HyperFlex HX-Series Spec Sheets](#), and the [Cisco HX240c M5 HyperFlex Node Installation Guide](#). For instructions on deployment, see: [Deploying AMD GPUs](#).

- Expanded HyperFlex Edge configurations—New HyperFlex Edge ordering PIDs provide more flexibility, simplify configuration and lower costs. For more information see, see [Cisco HyperFlex HX-Series Spec Sheets](#)

## 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 respective Cisco UCS Manager release version of [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.
- For SED-based HyperFlex systems, ensure that the A (Infrastructure) and C (Rack server) bundles are at Cisco UCS Manager version 3.1(3h) or higher for M4 SED systems. Make sure that all bundles are at Cisco UCS Manager version 3.2(3b) or higher for M5 SED systems.

### 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 9 Cisco UCS 6332 Fabric Interconnect

#### Front View



#### Rear View



Figure 10 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 HXAF220c M5 servers extend the capabilities of Cisco's HyperFlex portfolio in a 1U form factor with the addition of the Intel® Xeon® Processor Scalable Family, 24 DIMM slots for 2666MHz DIMMs, up to 128GB individual DIMM capacities and up to 3.0TB of total DRAM capacities.**


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

Figure 11 Cisco UCS HXAF220c-M5SX Rack Server Front View

**Front View**








1	<b>Drive Slots</b> Slot 01 (For System/Log drive) <ul style="list-style-type: none"> <li>• 1 x SATA SSD</li> </ul> Slot 02 (For Cache drive) <ul style="list-style-type: none"> <li>• 1 x NVMe SSD OR</li> <li>• 1 x SAS SSD OR</li> <li>• 1 x SED SAS SSD</li> </ul> Slot 03 through 10 (For Capacity drives) <ul style="list-style-type: none"> <li>• Upto 8 x SATA SSD OR</li> <li>• Upto 8 x SED SATA SSD OR</li> <li>• upto 8 x SED SAS SSD</li> </ul>	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 12 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-M5S offers 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 HyperFlex HXAF240c-M5SX Nodes

This capacity optimized configuration contains a minimum of three nodes, up to twenty three SED SATA or SAS SSD drives that contribute to cluster storage, a single 240 GB SATA SSD housekeeping drive, a single 400GB SAS SSD caching drive, and M.2 SATA SSD drive that acts as the boot drives. For detailed information, see the [Cisco HyperFlex HXAF240c M5 Node Spec Sheet](#).


Figure 13 HXAF240c-M4SX Node



### Cisco HyperFlex HX220c-M4S Hybrid Node

This small footprint configuration contains a minimum of three nodes with six 1.2 terabyte (TB) SAS drives that contribute to cluster storage capacity, a 240 GB SSD housekeeping drive, a 480 GB SSD caching drive, and 240Gb SATA M.2 SSD hat act as boot drives. For detailed information, see the [Cisco HyperFlex HX220c M5 Node Spec Sheet](#).


Figure 14 HX220c-M4S Node



#### Cisco HyperFlex HX240c-M4SX Hybrid Node

This capacity optimized configuration contains a minimum of three nodes, a minimum of fifteen and up to twenty-three 1.2 TB SAS drives that contribute to cluster storage, a single 240 GB SSD housekeeping drive, a single 1.6 TB SSD caching drive, and 240Gb SATA M.2 SSD that act as the boot drives. For detailed information, see the [Cisco HyperFlex HX240c M5 Node Spec Sheet](#).

Figure 15 HX240c-M5SX Node







### 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 (0). The mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases. The card enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile.

Figure 16 Cisco VIC 1387 mLOM Card



## Cisco HyperFlex Compute Nodes

### Cisco UCS B200-M5 Blade

For workloads that require additional computing and memory resources, but not additional storage capacity, a compute-intensive hybrid cluster configuration is allowed. This configuration requires a minimum of three (up to sixteen) HyperFlex converged nodes with one to sixteen Cisco UCS B200-M5 Blade Servers for additional computing capacity. The HX-series Nodes are configured as described previously, and the Cisco

UCS B200-M5 servers are equipped with boot drives. Use of the Cisco UCS B200-M5 compute nodes also requires the Cisco UCS 5108 blade server chassis, and a pair of Cisco UCS 2300/2200 series Fabric Extenders. For detailed information, see the [Cisco UCS B200 M5 Blade Server Spec Sheet](#).

Figure 17 Cisco UCS B200 M5 Server



## Cisco VIC1340 Converged Network Adapter

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


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

Figure 18 Cisco UCS VIC 1340



Figure 19 illustrates the Cisco UCS VIC 1340 Virtual Interface Cards Deployed in the Cisco UCS B-Series B200 M4 Blade Servers.

Figure 19 Cisco UCS VIC 1340 Deployed in the Cisco UCS B200 M5



## Cisco UCS 5108 Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis. The Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+1 redundant, and grid redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS Fabric Extenders. A passive mid-plane provides up to 40 Gbps of I/O bandwidth per server slot from each Fabric Extender. The chassis is capable of supporting 40 Gigabit Ethernet standards.

Figure 20 Cisco UCS 5108 Blade Chassis Front and Rear Views



## Cisco UCS 2304XP Fabric Extender

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



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

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

Cisco UCS 2304 Fabric Extenders fit into the back of the Cisco UCS 5100 Series chassis. Each Cisco UCS 5100 Series chassis can support up to two fabric extenders, allowing increased capacity and redundancy (Figure 22).

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

Figure 21 Cisco UCS 2304XP Fabric Extender



### Cisco UCS C220-M5 Rack Server

The Cisco UCS C220 M5 Rack Server is an enterprise-class infrastructure server in a 1RU form factor. It incorporates the Intel Xeon scalable family processor product family, next-generation DDR4 memory, and 12-Gbps SAS throughput, delivering significant performance and efficiency gains. Cisco UCS C220 M5 Rack Server can be used to build a compute-intensive hybrid HX cluster, for an environment where the workloads require additional computing and memory resources but not additional storage capacity, along with the HX-series converged nodes. This configuration contains a minimum of three (up to sixteen) HX-series converged nodes with one to sixteen Cisco UCS C220-M4 Rack Servers for additional computing capacity.

Figure 22 Cisco UCS C220 M5 Rack Server



### Cisco UCS C240-M5 Rack Server

The Cisco UCS C240 M5 Rack Server is an enterprise-class 2-socket, 2-rack-unit (2RU) rack server. It incorporates the Intel Xeon scalable family processor product family, next-generation DDR4 memory, and 12-Gbps SAS throughput that offers outstanding performance and expandability for a wide range of storage and I/O-intensive infrastructure workloads. Cisco UCS C240 M5 Rack Server can be used to expand additional computing and memory resources into a compute-intensive hybrid HX cluster, along with the HX-series converged nodes. This configuration contains a minimum of three (up to sixteen) HX-series converged nodes with one to sixteen Cisco UCS C240-M4 Rack Servers for additional computing capacity.

Figure 23 Cisco UCS C240 M5 Rack Server



### 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).

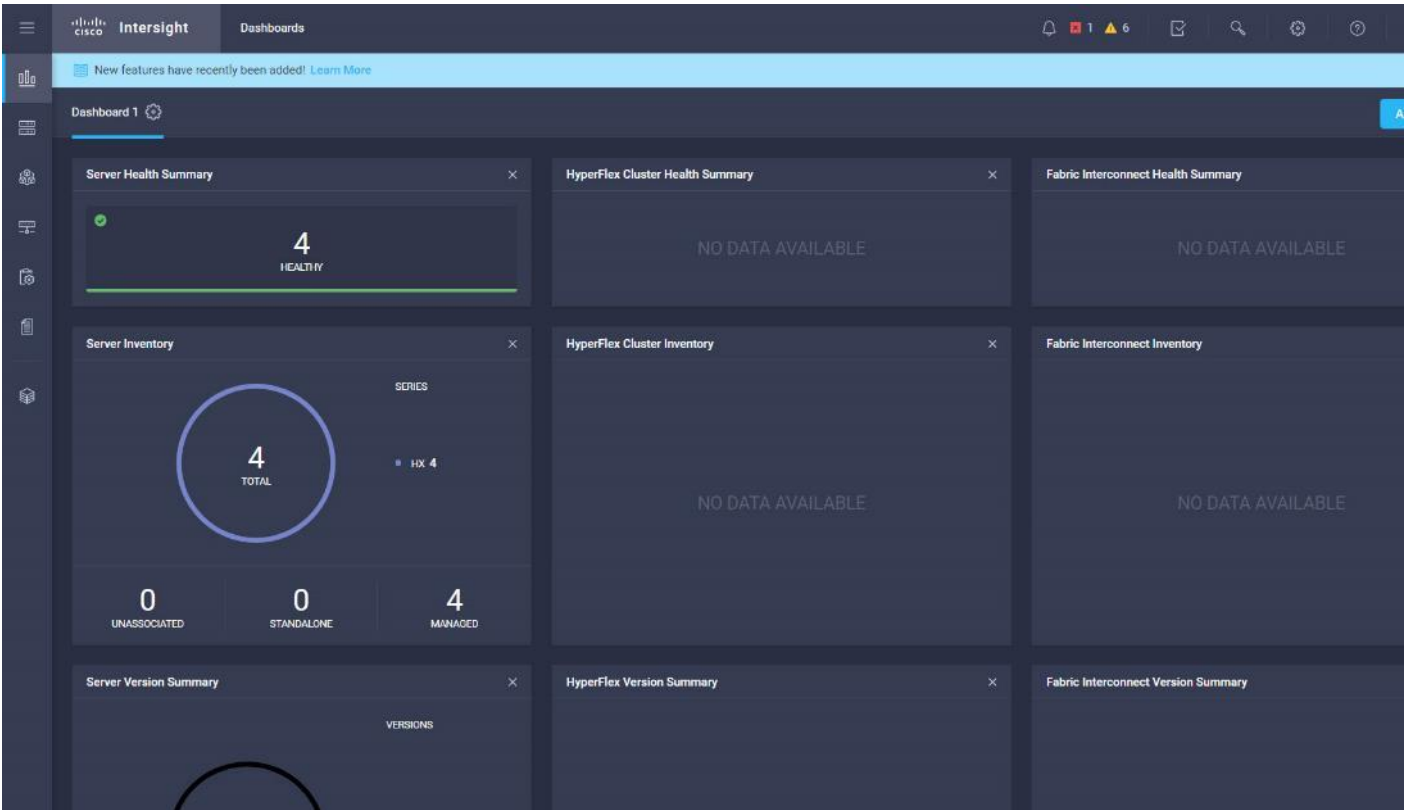
- Deduplication is always on, helping reduce storage requirements in virtualization clusters in which multiple operating system instances in client virtual machines result in large amounts of replicated data.
- Compression further reduces storage requirements, reducing costs, and the log-structured file 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

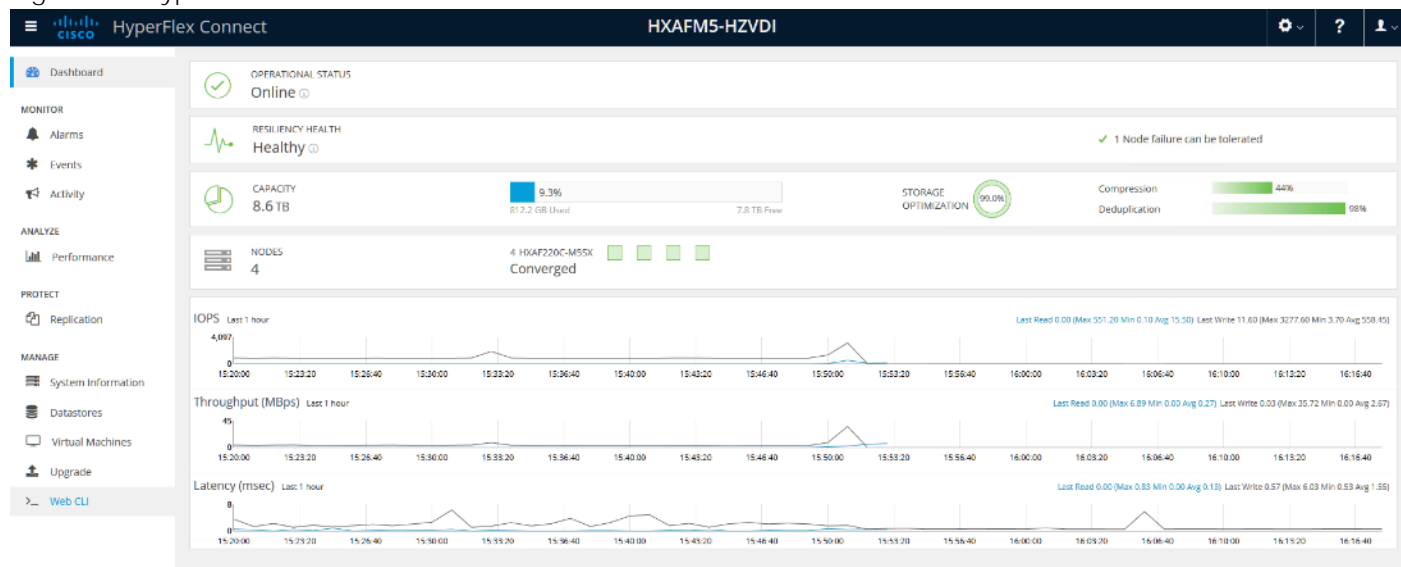
Cisco Intersight simplifies and automates IT operations management (ITOM) to make daily activities easier and more efficient. We have extended our vision of adaptive management to Cisco UCS and HyperFlex systems through the Cisco Intersight cloud-based platform. You can efficiently implement operations automation of your IT infrastructure from the data center to the edge.



The screenshot shows the 'Servers' section of the Cisco Intersight interface. It features a table with 11 columns: Name, Health, Management IP, Model, CPU Capacity (GHz), Memory Capacity (GiB), UCS Domain, HX Cluster, Server Profile, Utility Storage, and Firmware. The table displays 4 rows of data, with a '25 Rows' dropdown menu indicating more data is available. The first four rows show servers with names like 'k-20-c3-9', 'k-20-c3-8', 'k-20-c3-6', and 'k-20-c3-7', all with a 'HEALTHY' status (green checkmark) and a management IP of '10.29.132.15'. The table also includes a search bar and pagination controls at the bottom.

Name	Health	Management IP	Model	CPU Capacity (GHz)	Memory Capacity (GiB)	UCS Domain	HX Cluster	Server Profile	Utility Storage	Firmware
k-20-c3-9	HEALTHY	10.29.132.15	HXAF220C-M5...	82.8	768.0	k-20-c3		org-root/org-xd-cvd/ls-rack-1		3.1(2d)
k-20-c3-8	HEALTHY	10.29.132.15	HXAF220C-M5...	82.8	768.0	k-20-c3		org-root/org-xd-cvd/ls-rack-1		3.1(2d)
k-20-c3-6	HEALTHY	10.29.132.15	HXAF220C-M5...	82.8	768.0	k-20-c3		org-root/org-xd-cvd/ls-rack-1		3.1(2d)
k-20-c3-7	HEALTHY	10.29.132.15	HXAF220C-M5...	82.8	768.0	k-20-c3		org-root/org-xd-cvd/ls-rack-1		3.1(2d)

Figure 24 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.

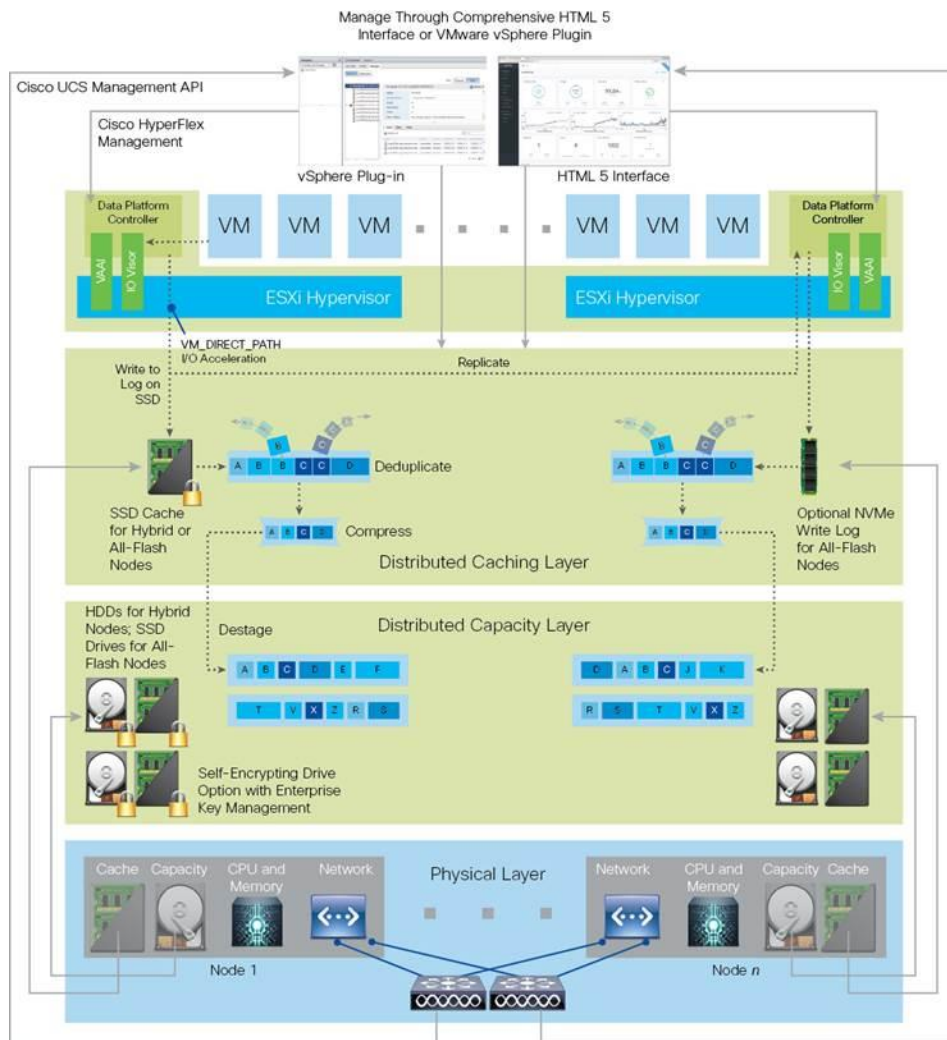
Incoming data is striped across the number of nodes required to satisfy availability requirements—usually two or three nodes. Based on policies you set, incoming write operations are acknowledged as persistent after they are replicated to the SSD drives in other nodes in the cluster. This approach reduces 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 25).

Figure 25 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

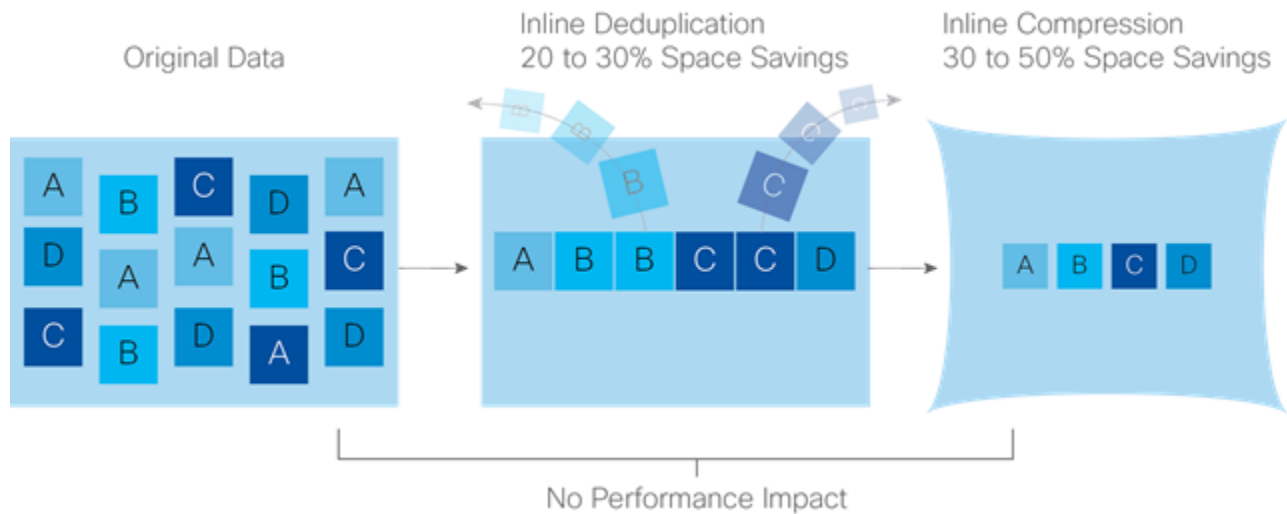
The Cisco HyperFlex HX Data Platform provides finely detailed inline deduplication and variable block inline compression that is always on for objects in the cache (SSD and memory) and capacity (SSD or HDD) layers. Unlike other solutions, which require you to turn off these features to maintain performance, 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 26).

Figure 26 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 27). 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 27 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 VM or SQL Server offline.

### Fast, Space-Efficient Clones

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

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

### Data Replication and Availability

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

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

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

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

### Data Rebalancing

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

### Online Upgrades

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

## Cisco Nexus 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 28 Cisco Nexus 93108YC Switch



## Microsoft Hyper-V 2016

Microsoft Hyper-V Server 2016 is a stand-alone product that contains only the Windows hypervisor, a Windows Server driver model, and virtualization components. It provides a simple and reliable virtualization solution to help you improve your server utilization and reduce costs.

The Windows hypervisor technology in Microsoft Hyper-V Server 2016 is the same as what's in the Hyper-V role on Windows Server 2016. So, much of the content available for the Hyper-V role on Windows Server 2016 also applies to Microsoft Hyper-V Server 2016.

## Microsoft System Center 2016

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 XenApp™ and XenDesktop™ 7.16

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 XenDesktop 7.16, IT can effectively control app and desktop provisioning while securing data assets and lowering capital and operating expenses.

The XenDesktop 7.16 release offers these benefits:

- Comprehensive virtual desktop delivery for any use case. The XenDesktop 7.16 release incorporates the full power of XenApp, delivering full desktops or just applications to users. Administrators can deploy both XenApp 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 XenDesktop 7.16 leverages common policies and cohesive tools to govern both infrastructure resources and user access.

- Simplified support and choice of BYO (Bring Your Own) devices. XenDesktop 7.16 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. XenDesktop 7.16 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 XenDesktop application and desktop workloads to private or public clouds.
- Protection of sensitive information through centralization. XenDesktop decreases the risk of corporate data loss, enabling access while securing intellectual property and centralizing applications since assets reside in the datacenter.
- Virtual Delivery Agent improvements. Universal print server and driver enhancements and support for the HDX 3D Pro graphics acceleration for Windows 10 are key additions in XenDesktop 7.16
- Improved high-definition user experience. XenDesktop 7.16 continues the evolutionary display protocol leadership with enhanced Thinwire display remoting protocol and Framehawk support for HDX 3D Pro.

Citrix XenApp and XenDesktop 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. XenApp and XenDesktop 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 XenApp delivers:

- XenApp 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 XenApp 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.
- XenApp 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 XenApp deployment.
- Citrix XenDesktop: 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 XenApp and XenDesktop: 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 XenApp and XenDesktop farms, the XenDesktop 7.16 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 XenDesktop 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 XenDesktop 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 XenDesktop editions include the features available in XenApp.

---

## Zones

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

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

For more information, see the [Zones](#) article.

## Improved Database Flow and Configuration

When you configure the databases during Site creation, you can now specify separate locations for the Site, Logging, and Monitoring databases. Later, you can specify different locations for all three databases. In



previous releases, all three databases were created at the same address, and you could not specify a different address for the Site database later.

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

For more information, see the [Databases](#) and [Controllers](#) articles.

## Application Limits

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

For more information, see the [Manage applications](#) article.

## Multiple Notifications before Machine Updates or Scheduled Restarts

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

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

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

For more information, see the [Manage Machine Catalogs](#) and [Manage machines in Delivery Groups](#) articles.

## API Support for Managing Session Roaming

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



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

---

For more information, see the [Sessions](#) article.

## API Support for Provisioning VMs from Hypervisor Templates

When using the PowerShell SDK to create or update a Machine Catalog, you can now select a template from other hypervisor connections. This is in addition to the currently-available choices of VM 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 29 Logical Architecture of Citrix XenDesktop



## Architecture and Design of Citrix XenDesktop on Cisco Unified Computing System and Cisco HyperFlex Storage Design Fundamentals

There are many reasons to consider a virtual desktop solution such as an ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Computer (BYOC) to work programs. The first step in designing a virtual desktop solution is to understand the user community and the type of tasks that are required to successfully execute their role. The following user classifications are provided:

- Knowledge Workers today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- External Contractors are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.
- Task Workers perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.
- Mobile Workers need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by

installing their own applications and storing their own data, such as photos and music, on these devices.

- Shared Workstation users are often found in state-of-the-art universities and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

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

- **Traditional PC:** A traditional PC is what typically constituted a desktop environment; physical device with a locally installed operating system.
- **Hosted Shared Desktop:** A hosted, server-based desktop is a desktop where the user interacts through a delivery protocol. With hosted, server-based desktops, a single installed instance of a server operating system, such as Microsoft Windows Server 2012, is shared by multiple users simultaneously. Each user receives a desktop "session" and works in an isolated memory space. Changes made by one user could impact the other users.
- **Hosted Virtual Desktop:** A hosted virtual desktop is a virtual desktop running either on virtualization layer (ESX) or on bare metal hardware. The user does not work with and sit in front of the desktop, but instead the user interacts through a delivery protocol.
- **Published Applications:** Published applications run entirely on the Microsoft Session Hosts and the user interacts through a delivery protocol. With published applications, a single installed instance of an application, such as Microsoft, is shared by multiple users simultaneously. Each user receives an application "session" and works in an isolated memory space.
- **Streamed Applications:** Streamed desktops and applications run entirely on the user's local client device and are sent from a server on demand. The user interacts with the application or desktop directly but the resources may only available while they are connected to the network.
- **Local Virtual Desktop:** A local virtual desktop is a desktop running entirely on the user's local device and continues to operate when disconnected from the network. In this case, the user's local device is used as a type 1 hypervisor and is synced with the data center when the device is connected to the network.

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

## Understanding Applications and Data

When the desktop user groups and sub-groups have been identified, the next task is to catalog group application and data requirements. This can be one of the most time-consuming processes in the VDI planning exercise, **but is essential for the VDI project's success. If the applications and data are not identified and co-located, performance will be negatively affected.**

The process of analyzing the variety of application and data pairs for an organization will likely be complicated by the inclusion cloud applications, like Salesforce.com. This application and data analysis is

beyond the scope of this Cisco Validated Design, but should not be omitted from the planning process. There are a variety of third party tools available to assist organizations with this crucial exercise.

### Project Planning and Solution Sizing Sample Questions

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

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

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

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

- What is the desktop OS planned? Windows 7, Windows 8, or Windows 10?
- 32-bit or 64-bit desktop OS?
- How many virtual desktops will be deployed in the pilot? In production? All Windows 7/8/10?
- How much memory per target desktop group desktop?
- Are there any rich media, Flash, or graphics-intensive workloads?
- What is the end point graphics processing capability?
- Will Citrix XenApp for Remote Desktop Server Hosted Sessions used?
- What is the hypervisor for the solution?
- What is the storage configuration in the existing environment?
- Are there sufficient IOPS available for the write-intensive VDI workload?
- Will there be storage dedicated and tuned for VDI service?
- Is there a voice component to the desktop?
- Is anti-virus a part of the image?
- Is user profile management (for example, non-roaming profile based) part of the solution?
- What is the fault tolerance, failover, disaster recovery plan?
- Are there additional desktop sub-group specific questions?

## Citrix XenDesktop 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 XenDesktop 7.16 integrates Hosted Shared and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, and manageable solution for delivering Windows applications and desktops as a service.

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

### Machine Catalogs

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Machine Catalog. In this CVD, VM provisioning relies on Citrix Provisioning Services to make sure that the machines in the catalog are consistent. In this CVD, machines in the Machine Catalog are configured to run either a Windows Server OS (for RDS hosted shared desktops) or a Windows Desktop OS (for hosted pooled VDI desktops).

### Delivery Groups

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

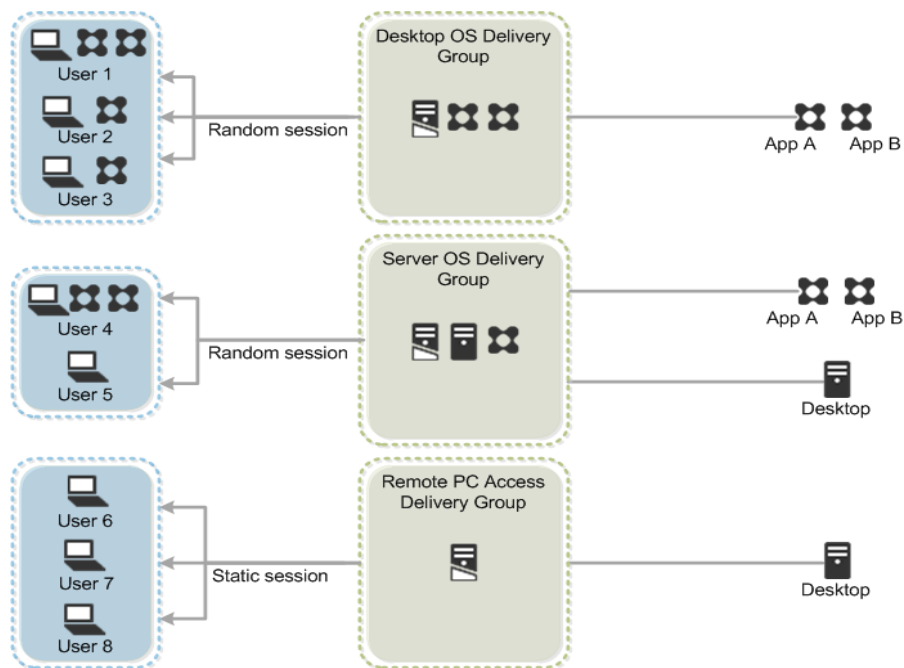
- Use machines from multiple catalogs
- Allocate a user to multiple machines
- Allocate multiple users to one machine

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

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

Figure 30 illustrates how users access desktops and applications through machine catalogs and delivery groups.

Figure 30 Access Desktops and Applications through Machine Catalogs and Delivery Groups



## Example XenDesktop Deployments

Two examples of typical XenDesktop deployments are:

- A distributed components configuration
- A multiple site configuration

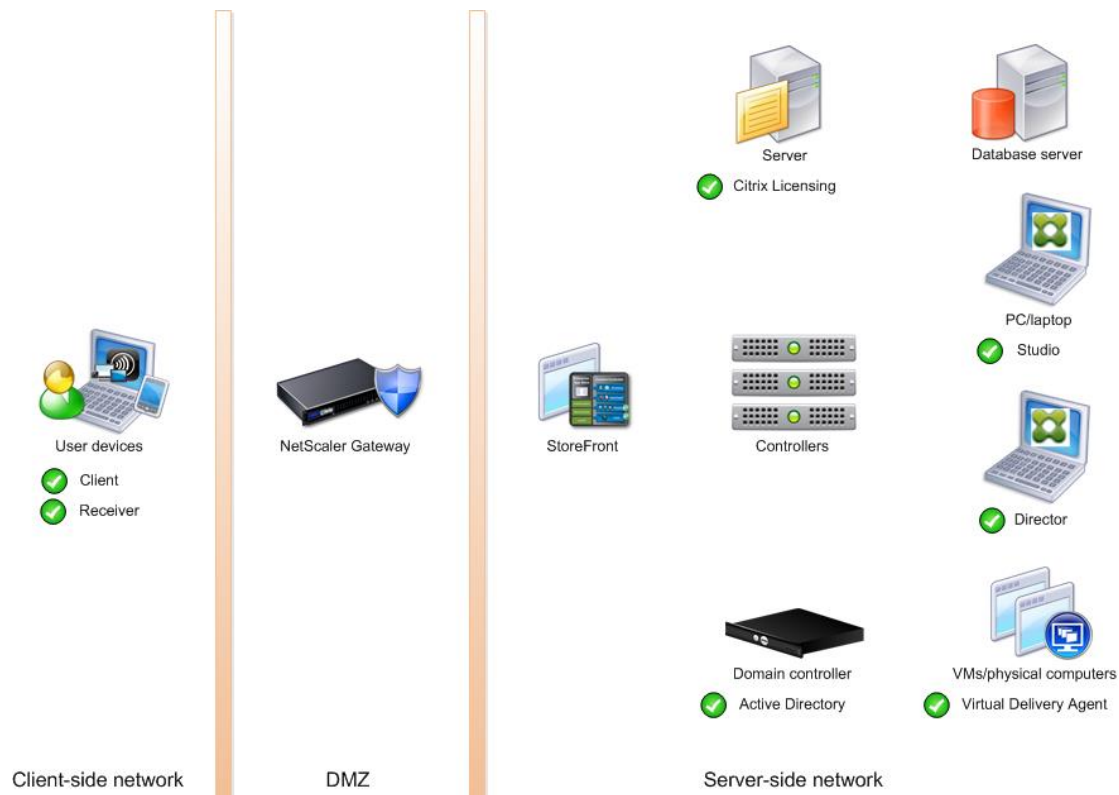
Since XenApp and XenDesktop 7.16 are based on a unified architecture, combined they can deliver a combination of Hosted Shared Desktops (HSDs, using a Server OS machine) and Hosted Virtual Desktops (HVDs, using a Desktop OS).

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

Figure 31 Example of a Distributed Components Configuration



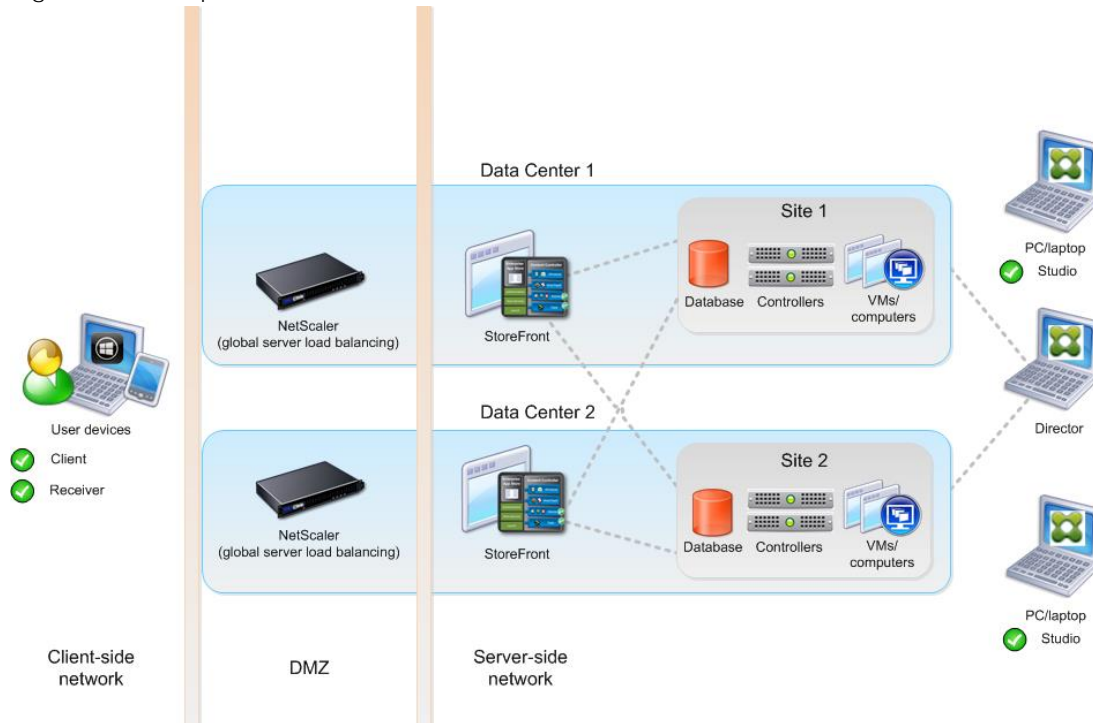
## Multiple Site Configuration

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 32 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 32 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 XenDesktop Environment for a Mixed Workload

With Citrix XenDesktop 7.16, 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.

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

## Deployment Hardware and Software

---

### Products Deployed

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

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

This validated design document details the deployment of the multiple configurations extending to 450 users for virtual desktop and hosted shared desktop workload featuring the following deployment methods:

- Citrix XenDesktop 7.16 persistent HVDs provisioned with Citrix Machine Creation Services (MCS) and using full copy on Cisco HyperFlex
- Microsoft Windows Server 2016 for User Profile Manager
- Microsoft Windows 2016 Server for Login VSI Management and data servers to simulate real world VDI workload
- Microsoft Hyper-V 2016
- Microsoft System Center Virtual Machine Manager 2016
- Windows 10 64-bit Operating Systems for VDI virtual machines
- Microsoft SQL Server 2016
- Cisco HyperFlex data platform v3.0.1a

Figure 33 Detailed Reference Architecture with Physical Hardware Cabling Configured to Enable the Solution

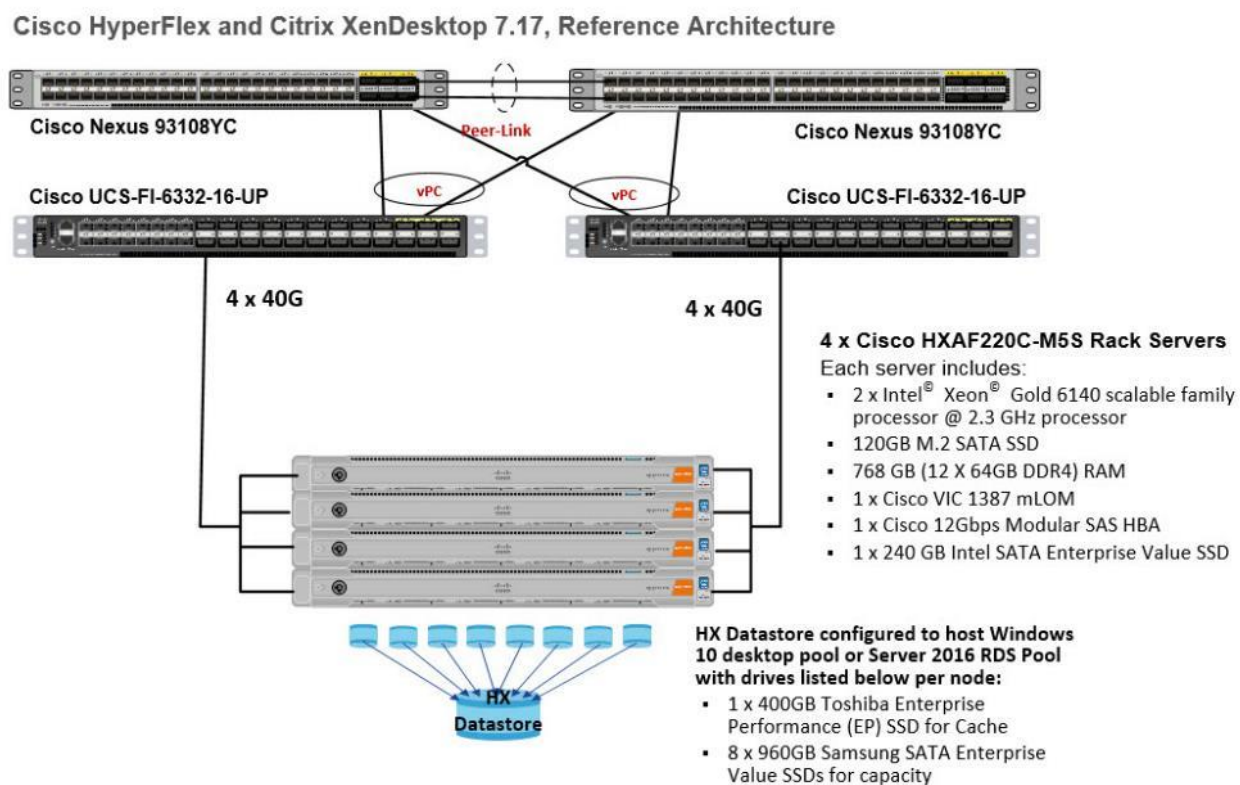
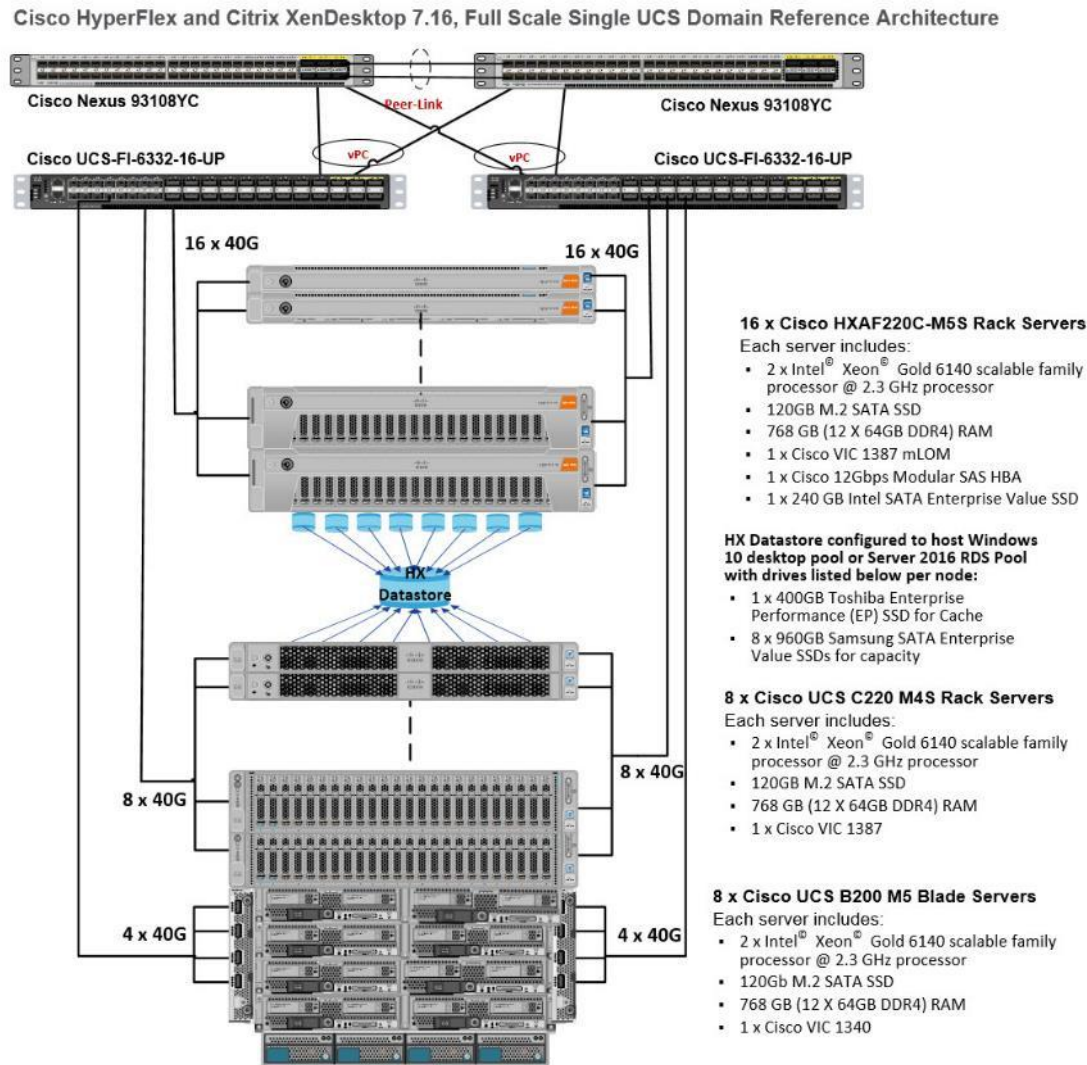


Figure 34 Detailed Reference Architecture with Physical Hardware Cabling Configured to Enable the Scale-Out Solution as Per the Current Cluster Limit



## Hardware Deployed

The solution contains the following hardware as shown in Figure 34:

- Two Cisco Nexus 93108YC Layer 2 Access Switches
- Two Cisco UCS C220 M4 Rack Servers with dual socket Intel Xeon E5-2620v4 2.1-GHz 8-core processors, 128GB RAM 2133-MHz and VIC1227 mLOM card for the hosted infrastructure with N+1 server fault tolerance. (Not show in the diagram).
- Four Cisco UCS HXAF220c-M5S Rack Servers with Intel Xeon Gold 6140 scalable family 2.3-GHz 18-core processors, 768GB RAM 2666-MHz and VIC1387 mLOM cards running Cisco HyperFlex data platform v3.0.1a for the virtual desktop workloads with N+1 server fault tolerance.

## Software Deployed

Table 1 lists the software and firmware version used in the study.

Table 1 Software and Firmware Versions

Vendor	Product	Version
Cisco	UCS Component Firmware	3.2(3d) bundle release
Cisco	UCS Manager	3.2(3d) bundle release
Cisco	UCS HXAF220c-M5S rack server	3.2(3d) bundle release
Cisco	VIC 1387	4.2(2d)
Cisco	HyperFlex Data Platform	3.0.1a-26588
Network	Cisco Nexus 9000 NX-OS	7.0(3)I2(2d)
Citrix	XenDesktop	7.16
Citrix	Provisioning Services	7.16
Citrix	User Profile Manager	
Citrix	Receiver	4.11
Microsoft	SCVMM	2016
Microsoft	Hyper-V Server	2016

## Logical Architecture

The logical architecture of this solution is designed to support up to 450 Hosted Virtual Microsoft Windows 10 Desktops users within a four node Cisco UCS HXAF220c- HyperFlex cluster, which provides physical redundancy for each workload type.

Figure 35 Logical Architecture Design

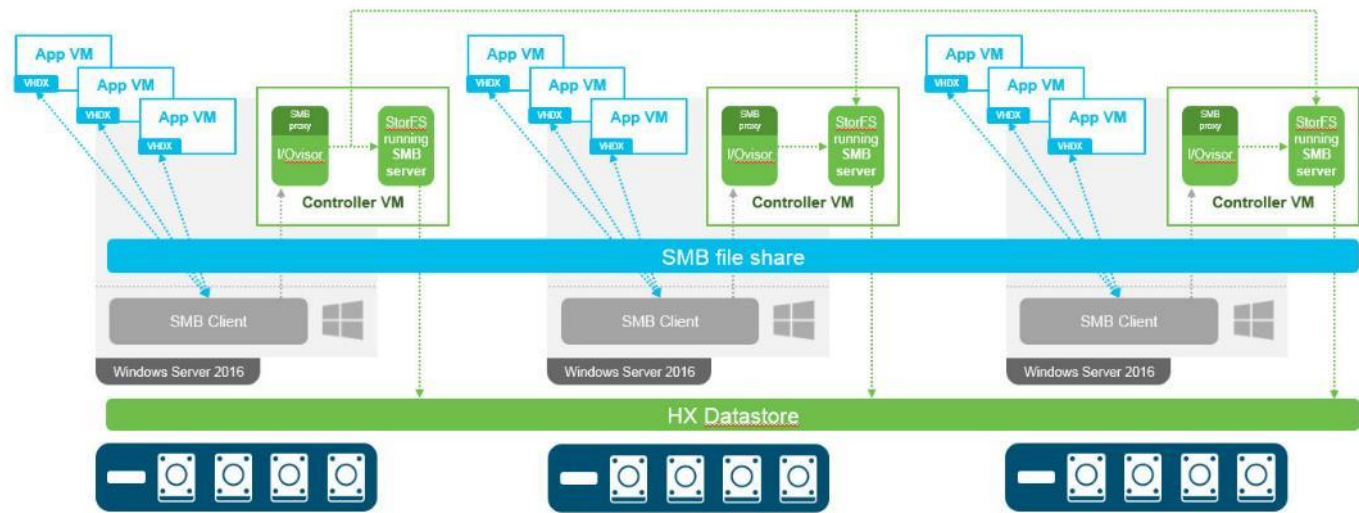



Table 1 lists the software revisions for this solution.


 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 2 through Table 6 lists the information you need to configure your environment.

VLANs

The VLAN configuration recommended for the environment includes a total of seven VLANs as outlined in Error! Reference source not found.2.

Table 2 Table 2 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 datacenters. In this scenario, the mgmt0 interfaces of the two Fabric Interconnects would be connected to that dedicated network or subnet. This is a valid configuration for HyperFlex installations with the following caveat; wherever the HyperFlex installer is deployed it must have IP connectivity to the subnet of the mgmt0 interfaces of the



Fabric Interconnects, and also have IP connectivity to the subnets used by the hx-inband-mgmt VLANs listed above.

---

## Jumbo Frames

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

## 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 VMs to connect to with their individual management interfaces. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in 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 VMs 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

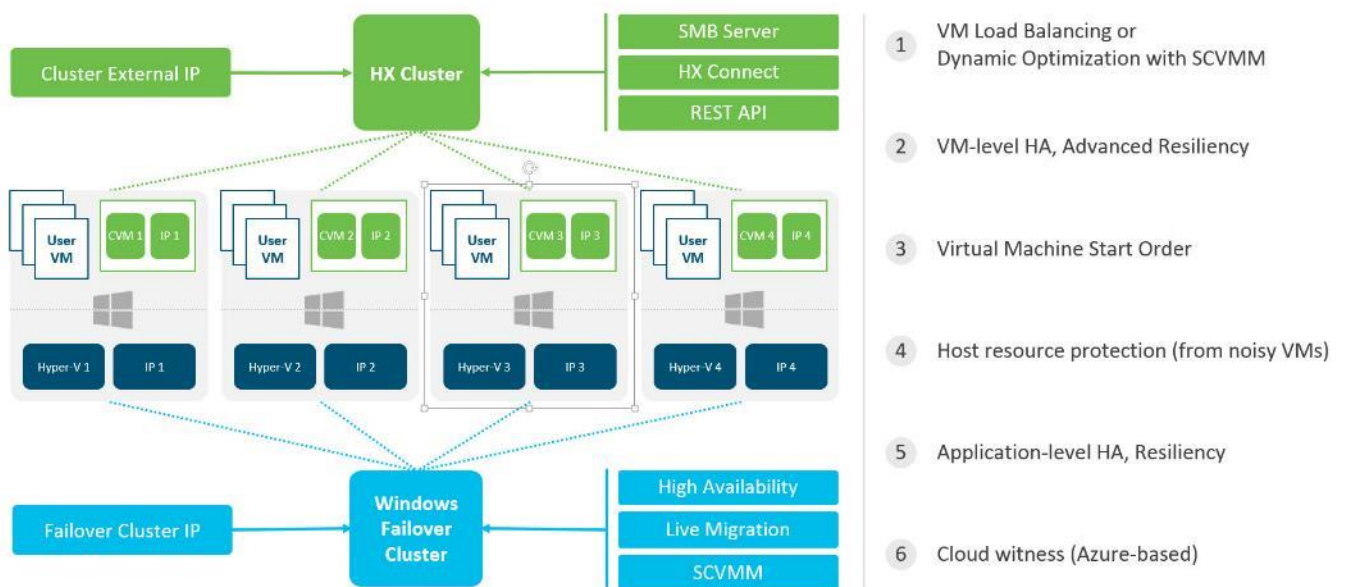
The following table and figures help give more details into the Hyper-V virtual networking design as built by the HyperFlex installer:

Table 3 Table Hyper-V Host Virtual Switch Configuration

Virtual Switch	Port Groups	Active vmnic(s)	Passive vmnic(s)	VLAN IDs	Jumbo
vswitch-hx-	Management	vmnic0	vmnic1	hx-inband-	no

Virtual Switch	Port Groups	Active vmnic(s)	Passive vmnic(s)	VLAN IDs	Jumbo
inband-mgmt	Network  Storage Controller Management Network			mgmt	
vswitch-hx-storage-data	Storage Controller Data Network  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 36 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 VMs cooperate to form and coordinate the Cisco HX Distributed Filesystem, and service all the guest VM IO requests. The controller VMs are deployed as a Hyper-V agent, which is similar in concept to that of a Linux or Windows service. Hyper-V agents are

ties 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 VM 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 VM runs custom software and services that manage and maintain the Cisco HX Distributed Filesystem. The services and processes that run within the controller VMs 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 VMs. Management and visibility into the function of the controller VMs, 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 VMs to display the information requested.


### Controller VM Locations

The physical storage location of the controller VM is similar between the Cisco HXAF220c-M5S and HXAF240c-M5SX model servers. The storage controller VM is operationally no different from any other typical virtual machines in a Hyper-V environment. The VM must have a virtual disk with the bootable root filesystem available in a location separate from the SAS HBA that the VM is controlling via SR-IOV. The configuration details of the models are as follows:

### 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 37 Datastore Example



### CPU Resource Reservations

Since the storage controller VMs provide critical functionality of the Cisco HX Distributed Data Platform, the HyperFlex installer will configure CPU resource reservations for the controller VMs. This reservation guarantees that the controller VMs 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 VMs. Table 4 details the CPU resource reservation of the storage controller VMs.

Table 4 Controller VM CPU Reservations

Number of vCPU	Shares	Reservation	Limit
8	Low	10800 MHz	unlimited

## Memory Resource Reservations

Since the storage controller VMs provide critical functionality of the Cisco HX Distributed Data Platform, the HyperFlex installer will configure memory resource reservations for the controller VMs. This reservation guarantees that the controller VMs 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 VMs.

Table 5 lists the memory resource reservation of the storage controller VMs.

Table 5 Controller VM Memory Reservations

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HX220c-M5 HXAF220c-M5	48 GB	Yes
HX240c-M5SX HXAF240c-M5SX	72 GB	Yes

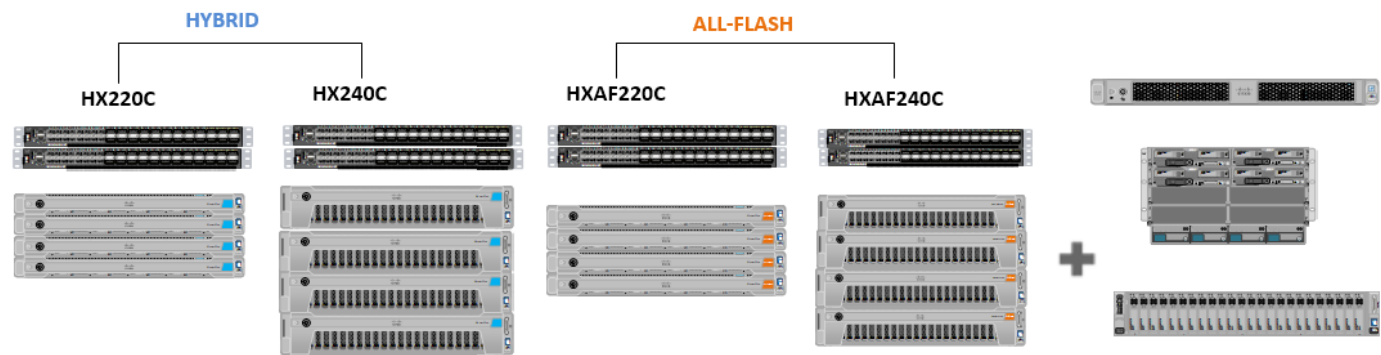


The Cisco UCS compute-only nodes have a lightweight storage controller VM; it is configured with only 1 vCPU and 512 MB of memory reservation.

# Solution Configuration

This section details the configuration and tuning that was performed on the individual components to produce a complete, validated solution. Figure 38 illustrates the configuration topology for this solution.

Figure 38 Configuration Topology for Scalable Citrix XenDesktop 7.16 Workload with HyperFlex



## Cisco UCS Compute Platform

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


### Physical Infrastructure

#### Solution Cabling

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

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

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



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

Figure 39 shows a cabling diagram for a Citrix XenDesktop configuration using the Cisco Nexus 9000 and Cisco UCS Fabric Interconnect.

Table 6 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

Local Device	Local Port	Connection	Remote Device	Remote Port
	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
	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 7 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

Local Device	Local Port	Connection	Remote Device	Remote Port
	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 8 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 9 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 39 Cable Connectivity Between Cisco Nexus 93108YC A and B to Cisco UCS 6248 Fabric A and B





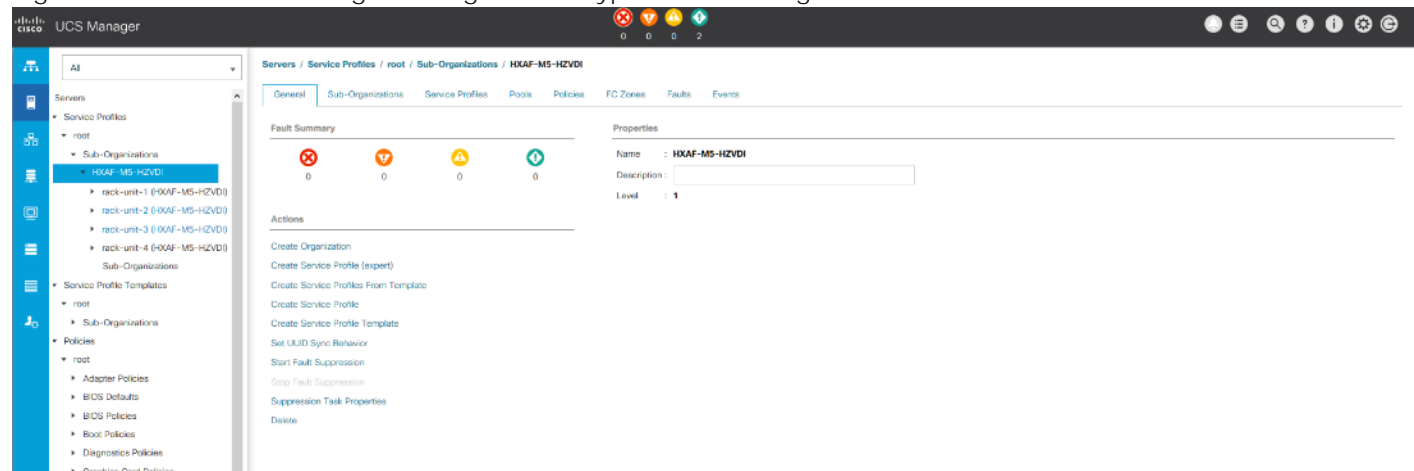
## Cisco Unified Computing System Configuration

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

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

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

Figure 40 Cisco UCS Manager Configuration: HyperFlex Sub-organization

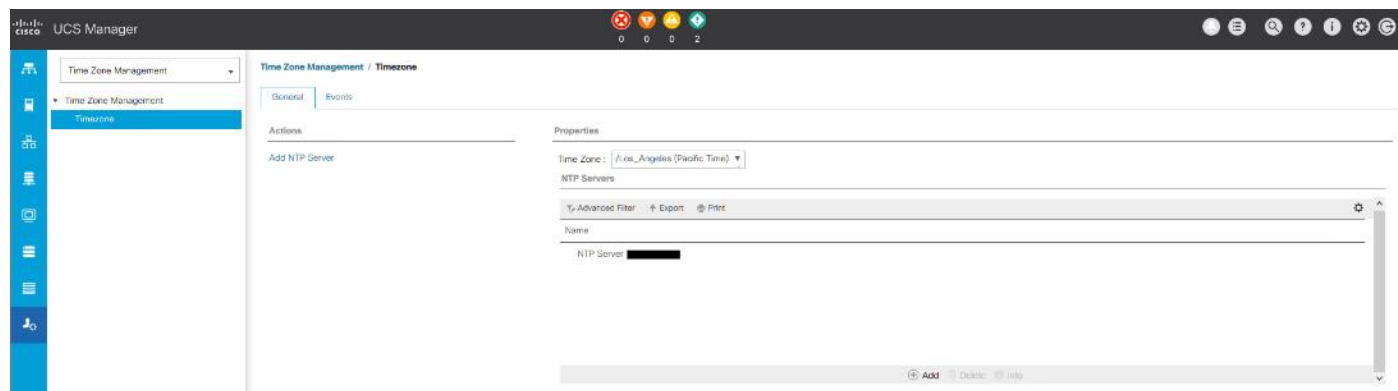


## Deploy and Configure HyperFlex Data Platform

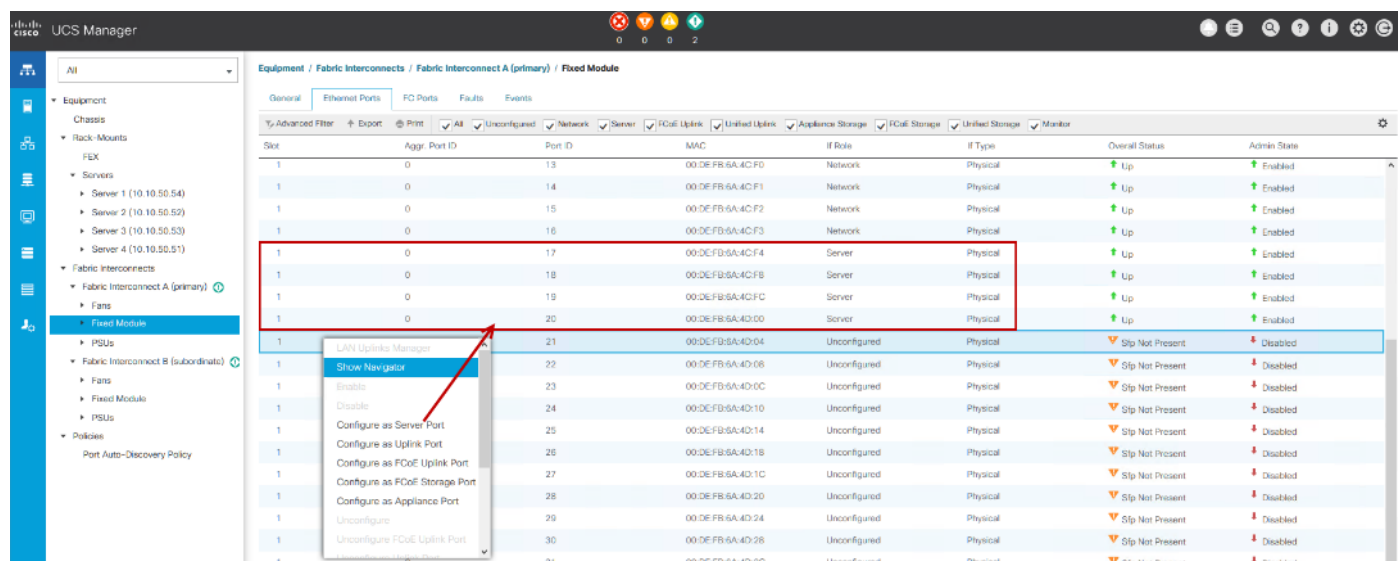
### Prerequisites

To deploy and configure the HyperFlex Data Platform, you must complete the following prerequisites:

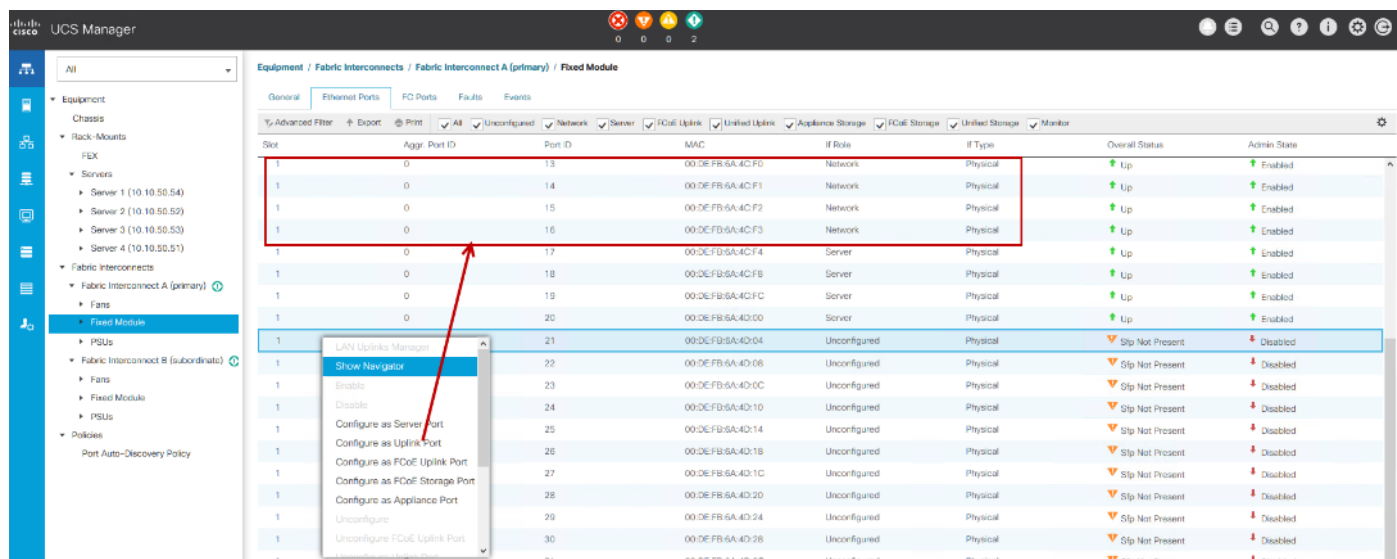
1. Set Time Zone and NTP: From the Cisco UCS Manager, from the Admin tab, Configure TimeZone and add NTP server. Save changes.



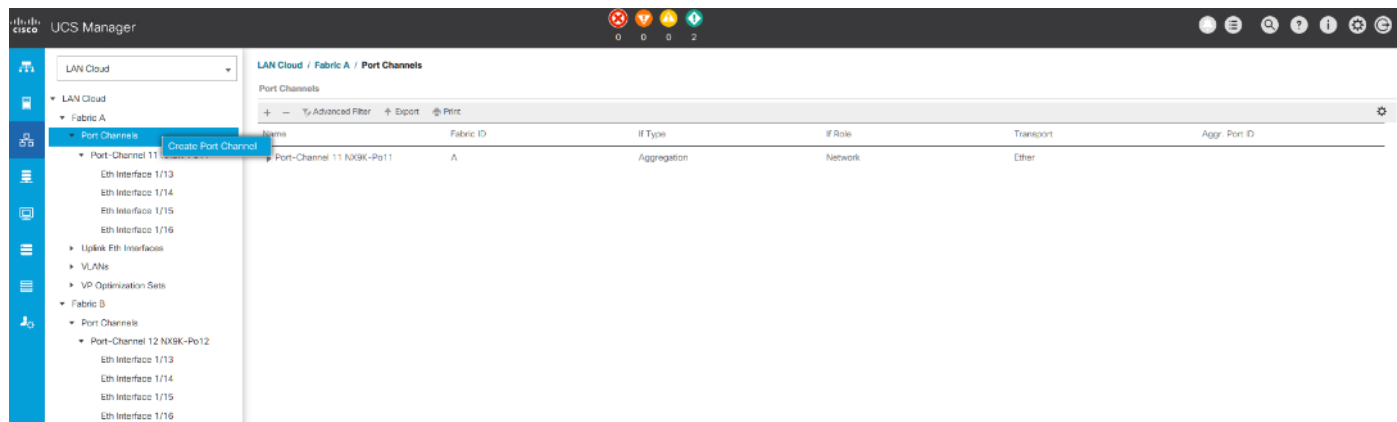
2. Configure Server Ports: Under the Equipment tab, Select Fabric A, select port to be configured as server port to manager HyperFlex rack server through Cisco UCS Manager.



3. Repeat this step to configure server port on Fabric B.
4. Configure Uplink Ports: On Fabric A, Select port to be configured as uplink port for network connectivity to north bound switch.



- Repeat this same on Fabric B.
- Create Port Channels: Under LAN tab, select expand LAN > LAN cloud > Fabric A. Right-click Port Channel.
- Select Create port-channel to connect with upstream switch as per Cisco UCS best practice. For our reference architecture, we connected a pair of Nexus 93108YCPX switches.



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

1

Set Port Channel Name

2

Add Ports

Create Port Channel

?

×

ID :

Name :

< Prev

Next >

Finish

Cancel

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

10. Click Finish.

**Create Port Channel**

**1 Set Port Channel Name**

**2 Add Ports**

Ports			
Slot ID	Aggr. Po...	Port	MAC
1	0	13	00:DE:F...
1	0	14	00:DE:F...
1	0	15	00:DE:F...
1	0	16	00:DE:F...

>>  
<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
No data available			

< Prev   Next >   **Finish**   Cancel

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

**UCS Manager**

**LAN**

**Port Channels and Uplinks**

Name	Fabric ID	Admin State
▼ Fabric A		
Port-Channel 11 NX9K-Po11	A	Enabled
Eth Interface 1/13	A	Enabled
Eth Interface 1/14	A	Enabled
Eth Interface 1/15	A	Enabled
Eth Interface 1/16	A	Enabled
▼ Fabric B		
Port-Channel 12 NX9K-Po12	B	Enabled
Eth Interface 1/13	B	Enabled
Eth Interface 1/14	B	Enabled
Eth Interface 1/15	B	Enabled
Eth Interface 1/16	B	Enabled
▼ Uplink Eth Interfaces		
Fabric A		
Fabric B		

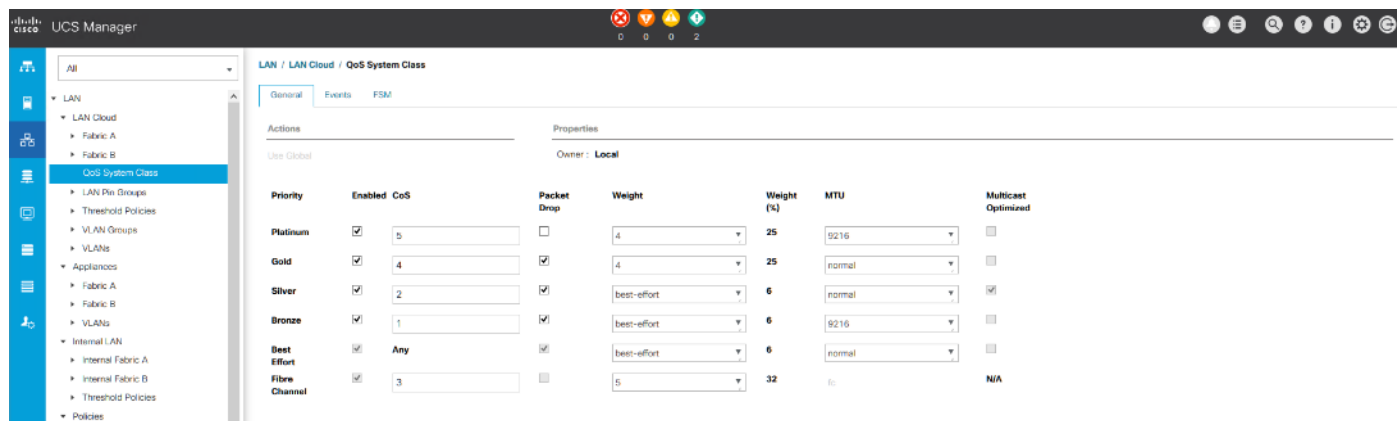
**Pin Groups**

Name	Port
No data available	

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

- Set MTU to 9216 for Platinum (Storage data) and Bronze (LiveMigration)
- Uncheck Enable Packet drop on the Platinum class

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



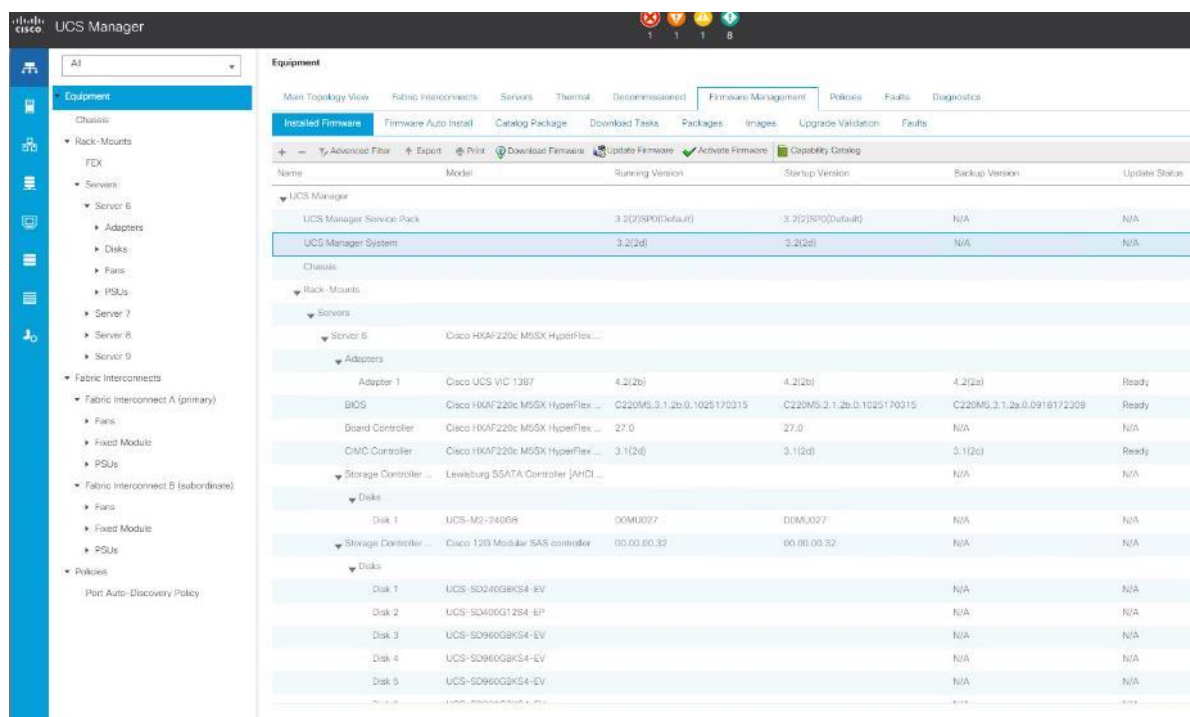
The screenshot shows the Cisco UCS Manager interface for configuring QoS System Classes. The left sidebar shows the navigation tree with 'QoS System Class' selected. The main panel displays the configuration for the 'Local' owner. The configuration table is as follows:

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input checked="" type="checkbox"/>	5	<input type="checkbox"/>	4	25	9216	<input type="checkbox"/>
Gold	<input checked="" type="checkbox"/>	4	<input checked="" type="checkbox"/>	4	25	normal	<input type="checkbox"/>
Silver	<input checked="" type="checkbox"/>	2	<input checked="" type="checkbox"/>	best-effort	6	normal	<input checked="" type="checkbox"/>
Bronze	<input checked="" type="checkbox"/>	1	<input checked="" type="checkbox"/>	best-effort	6	9216	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	best-effort	6	normal	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	32	6	N/A



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

- Verify UCS Manager Software Version: In the Equipment tab, select Firmware Management > Installed Firmware.
- Check and verify, both Fabric Interconnects and Cisco UCS Manager are configured with Cisco UCS Manager v3.2.3d.



The screenshot shows the Cisco UCS Manager interface with the 'Equipment' tab selected. The 'Firmware Management' sub-tab is active, displaying the 'Installed Firmware' table. The table lists the following components and their versions:

Name	Model	Running Version	Startup Version	Backup Version	Update Status
UCS Manager					
UCS Manager Service Pack		3.2(2)(SP0)(Default)	3.2(2)(SP0)(Default)	N/A	N/A
UCS Manager System		3.2(2d)	3.2(2d)	N/A	N/A
Chassis					
Rack-Mounts					
Servers					
Server 6	Cisco HXAF220c M5SX HyperFlex ...				
Adapters					
Adaptor 1	Cisco UCS VIC 1287	4.2(2b)	4.2(2b)	4.2(2a)	Ready
BIOS	Cisco HXAF220c M5SX HyperFlex ...	C220M5.3.1.2b.0.1020170315	C220M5.3.1.2b.0.1020170315	C220M5.3.1.2b.0.0916172309	Ready
Board Controller	Cisco HXAF220c M5SX HyperFlex ...	27.0	27.0	N/A	N/A
CIMC Controller	Cisco HXAF220c M5SX HyperFlex ...	3.1(2d)	3.1(2d)	3.1(2c)	Ready
Storage Controller	Lewieburg SSATA Controller [AHC] ...			N/A	N/A
Disks					
Disk 1	UCS-M2-240GB	DDMU027	DDMU027	N/A	N/A
Storage Controller	Cisco 12GB Modular SAS controller	00.00.00.32	00.00.00.32	N/A	N/A
Disks					
Disk 1	UCS-S0240GBK54-EV			N/A	N/A
Disk 2	UCS-S0400G1284-EV			N/A	N/A
Disk 3	UCS-S0960GBK54-EV			N/A	N/A
Disk 4	UCS-S0960GBK54-EV			N/A	N/A
Disk 5	UCS-S0960GBK54-EV			N/A	N/A



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

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

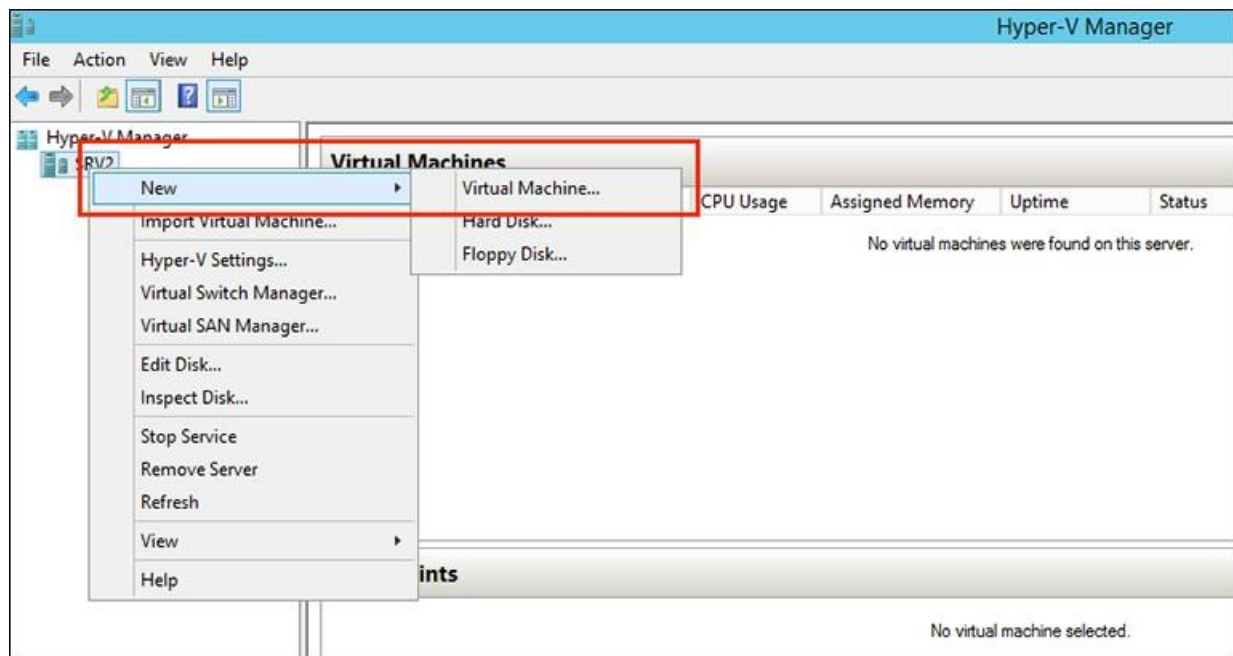
### Deploying HX Data Platform Installer on Hyper-V Infrastructure

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

- Locate and download the HX Data Platform Installer.vhdx zipped file (for example, **Cisco-HX-Data-Platform-Installer-v3.0.1a-build-hyperv.vhdx**) from the Cisco Software Downloads site.
- 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,

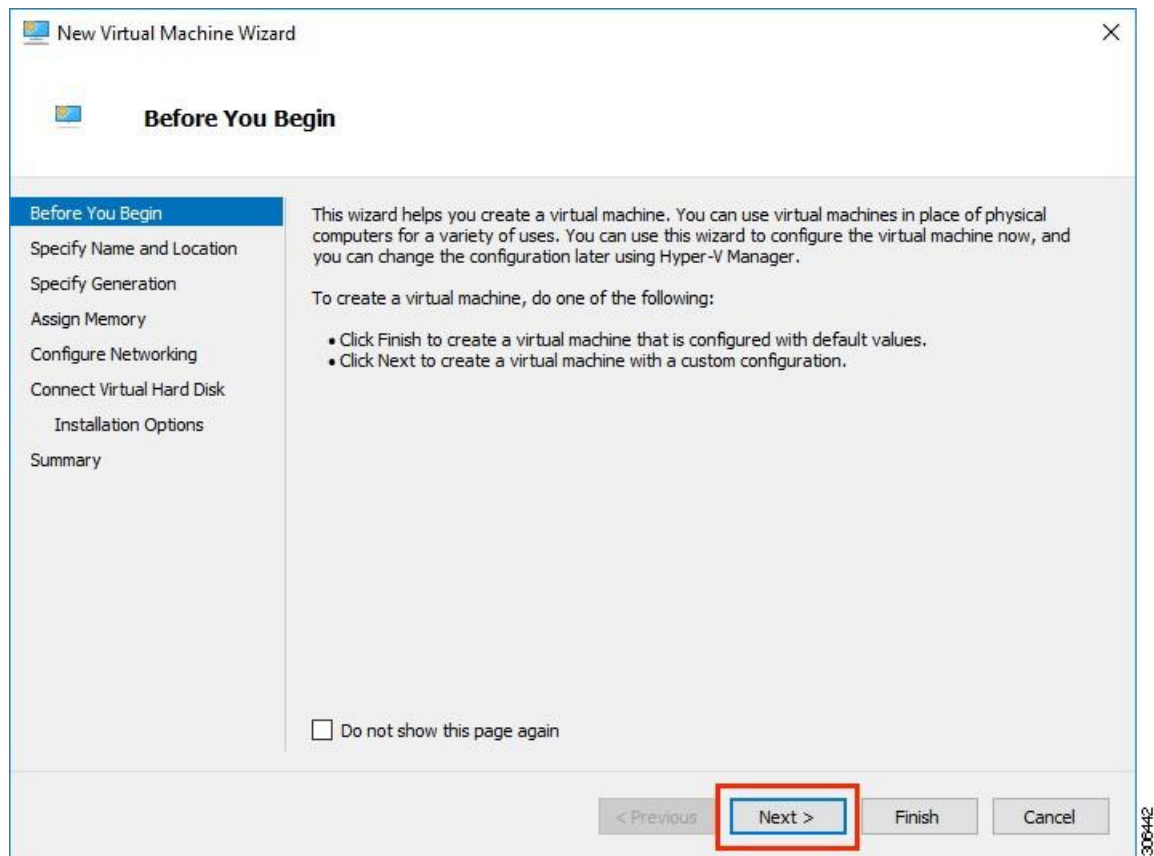
```
\\hyp-v-host01\...\HX-Installer\Cisco-HX-Data-Platform-Installer-v3.0.1a-29499-hyperv.vhdx
```

- In Hyper-V Manager, navigate to one of the Hyper-V servers.
- Select the Hyper-V server, and right click and select New > Create a virtual machine. The Hyper-V Manager New Virtual Machine Wizard displays.

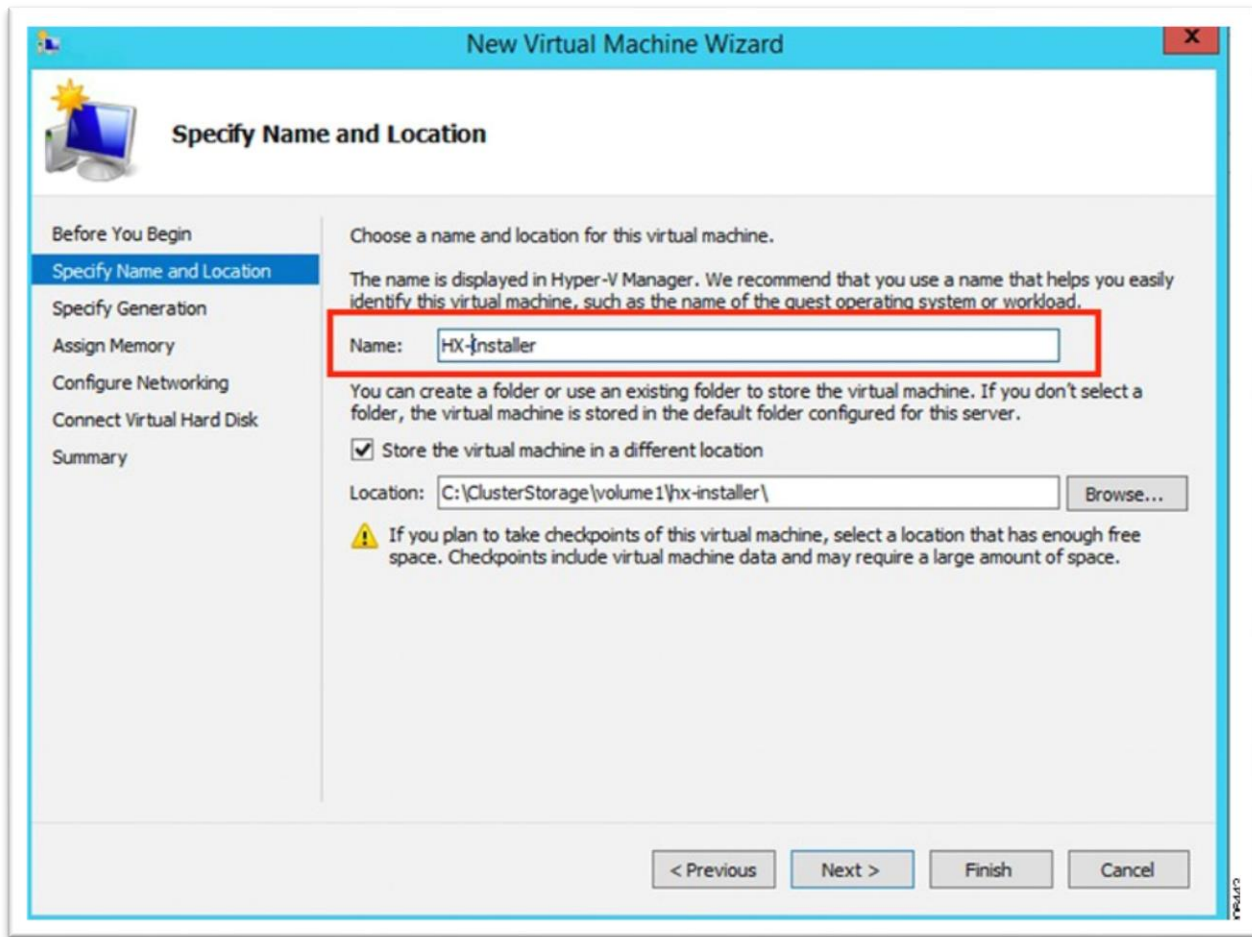


- In the Before you Begin page, click Next.





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.

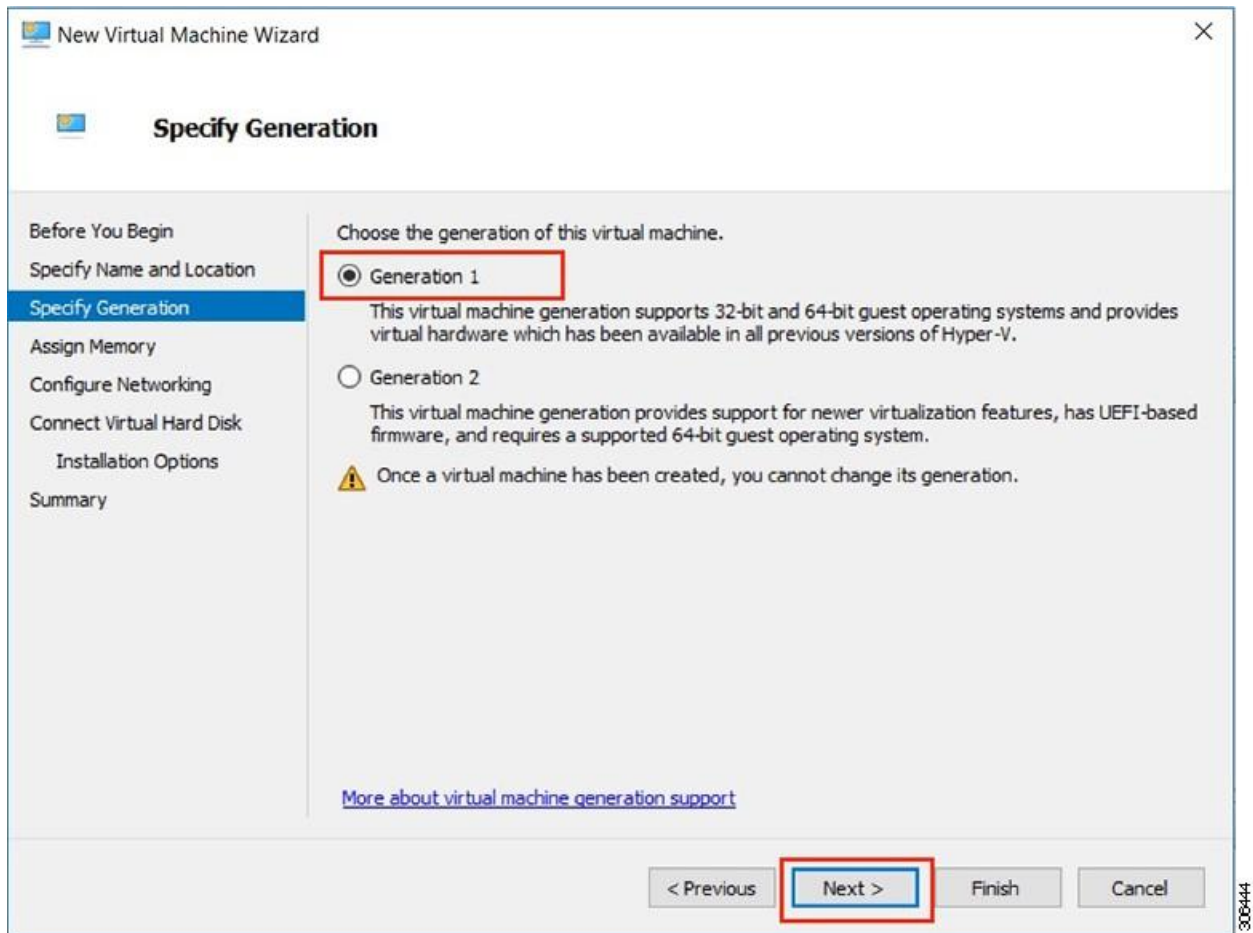


As a best practice, store the VM together with the .vhdx file.

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



If you select Generation 2, the VM may not boot.



The image shows the 'Specify Generation' step of the 'New Virtual Machine Wizard'. The wizard has a sidebar on the left with steps: 'Before You Begin', 'Specify Name and Location', 'Specify Generation' (highlighted), 'Assign Memory', 'Configure Networking', 'Connect Virtual Hard Disk', 'Installation Options', and 'Summary'. The main area is titled 'Specify Generation' and contains the instruction 'Choose the generation of this virtual machine.' Below this, there are two radio button options: 'Generation 1' (selected and highlighted with a red box) and 'Generation 2'. A description for 'Generation 1' states it supports 32-bit and 64-bit guest operating systems and provides virtual hardware available in previous versions of Hyper-V. A description for 'Generation 2' states it provides support for newer virtualization features, has UEFI-based firmware, and requires a supported 64-bit guest operating system. A warning icon and text state: 'Once a virtual machine has been created, you cannot change its generation.' At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a red box), 'Finish', and 'Cancel'. A link for 'More about virtual machine generation support' is also present. The window title is 'New Virtual Machine Wizard' and the ID '306444' is in the bottom right corner.

New Virtual Machine Wizard

**Specify Generation**

Before You Begin  
Specify Name and Location  
**Specify Generation**  
Assign Memory  
Configure Networking  
Connect Virtual Hard Disk  
Installation Options  
Summary

Choose the generation of this virtual machine.

☒ Generation 1  
This virtual machine generation supports 32-bit and 64-bit guest operating systems and provides virtual hardware which has been available in all previous versions of Hyper-V.

☐ Generation 2  
This virtual machine generation provides support for newer virtualization features, has UEFI-based firmware, and requires a supported 64-bit guest operating system.

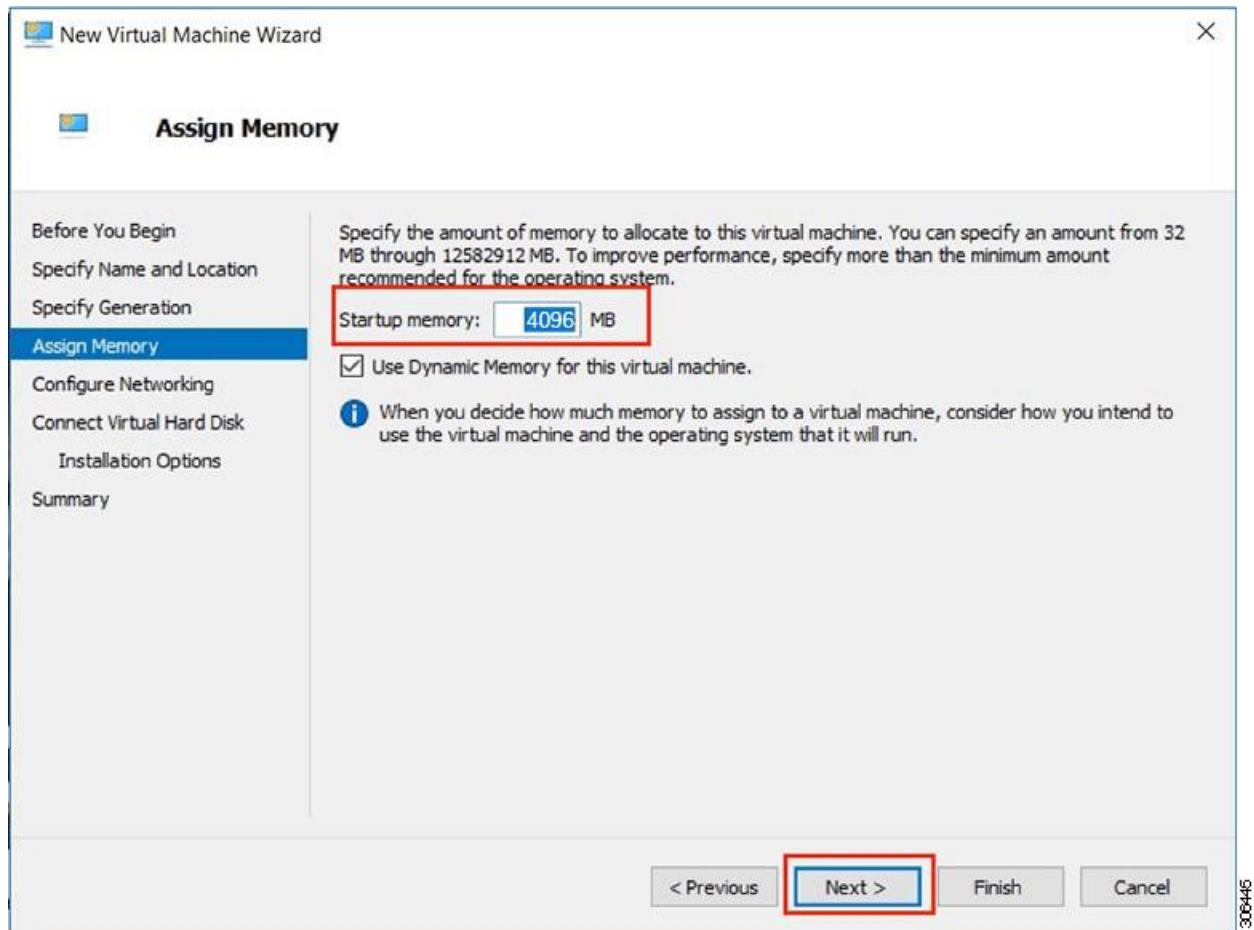
⚠ Once a virtual machine has been created, you cannot change its generation.

[More about virtual machine generation support](#)

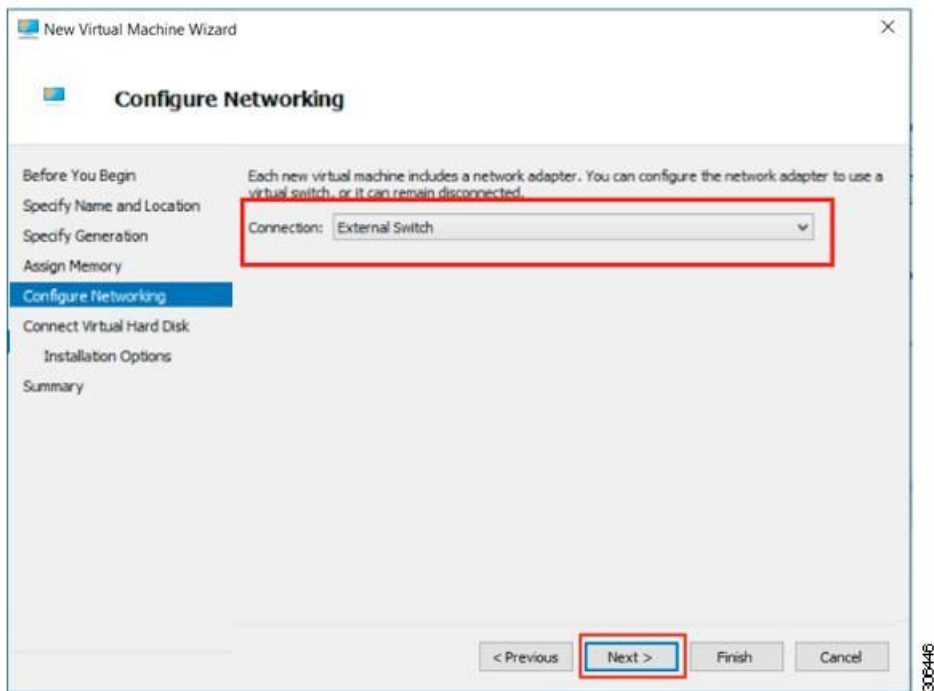
< Previous   **Next >**   Finish   Cancel

306444

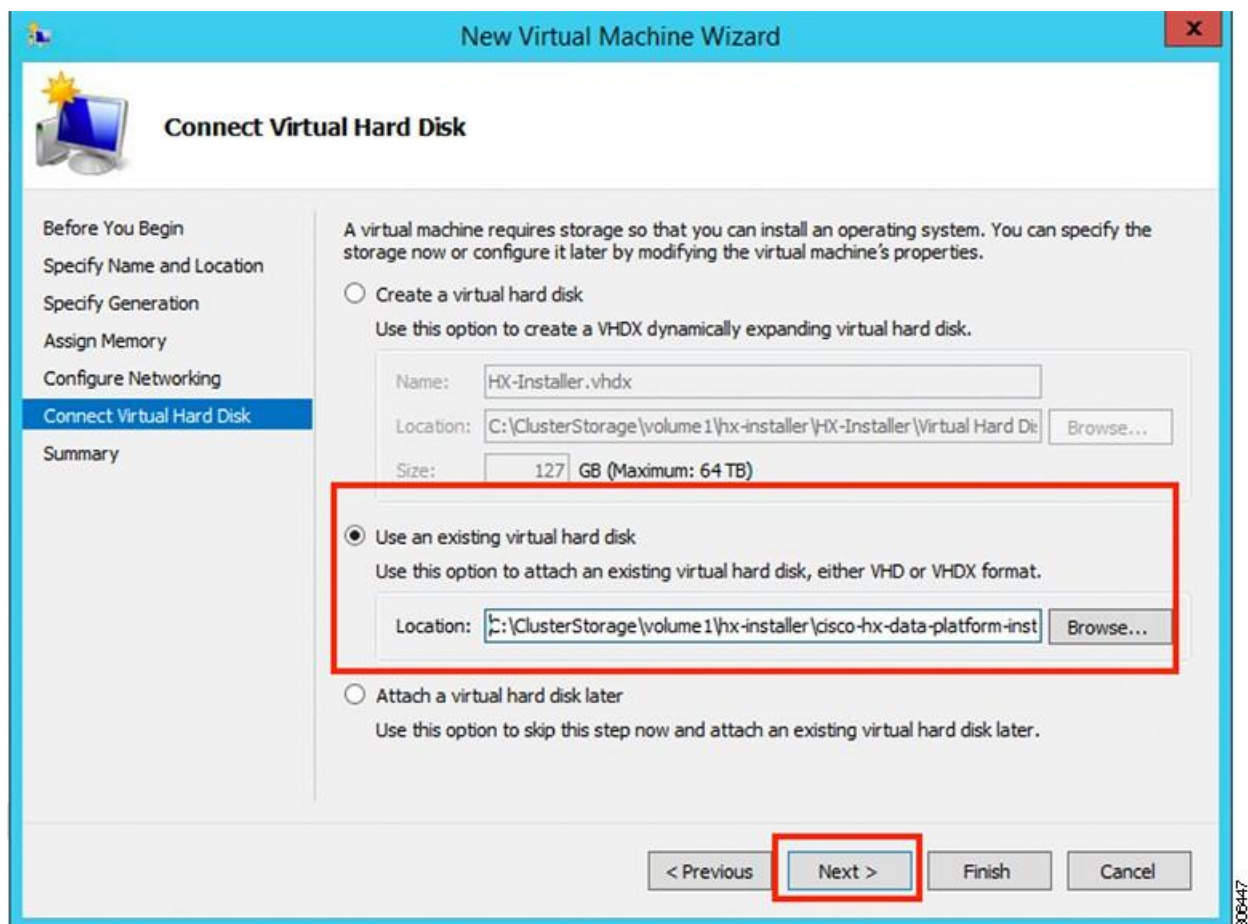
8. In the Assign Memory page, set the startup memory value to **4096 MB**. Click Next.



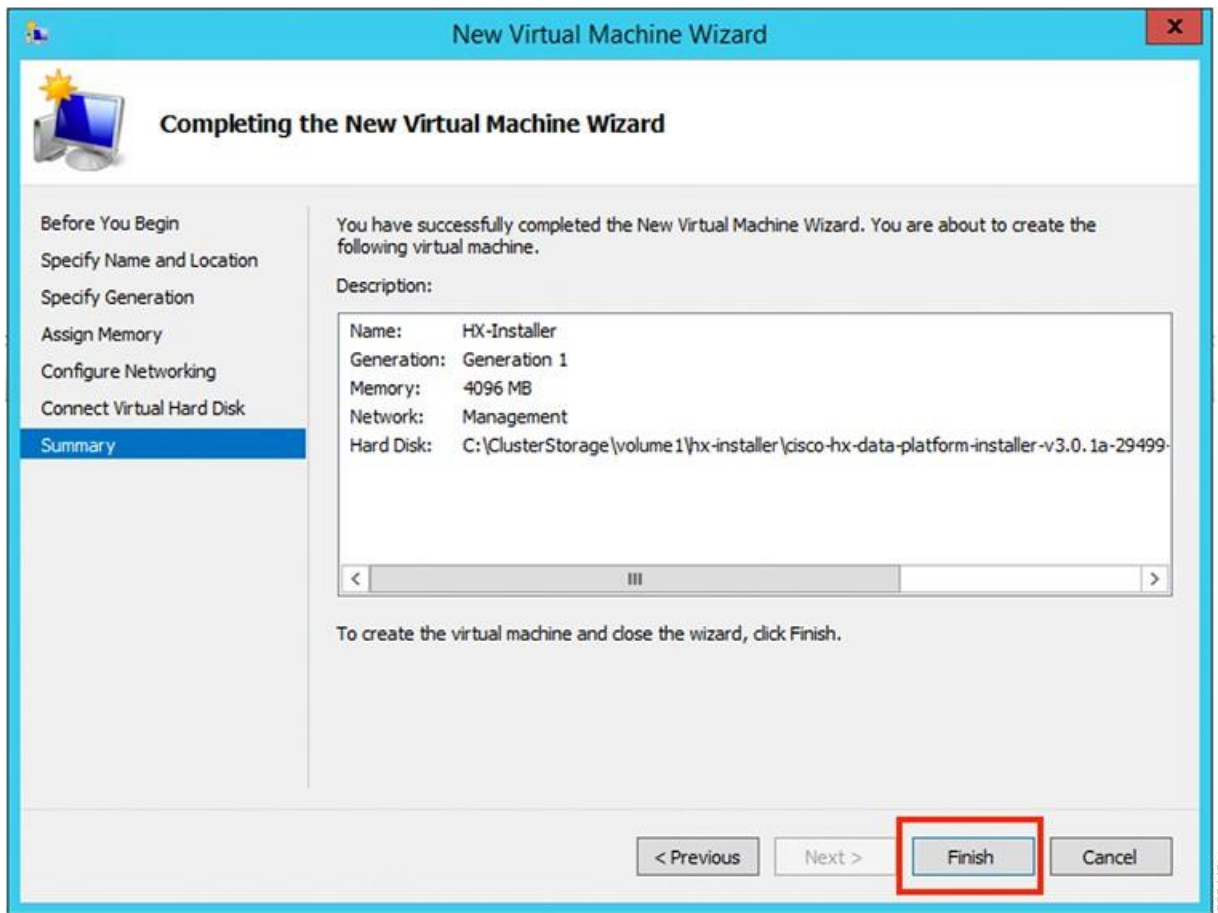
9. In the Configure Networking page, select a network connection for the virtual machine to use from a list of existing virtual switches. Click Next.



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.




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



12. After the VM is created, power it ON, and launch the GUI.

- a. Right-click the VM and choose Connect.
- b. Choose Action > Start (Ctrl+S).
- c. When the VM is booted, make a note of the URL (IP address of the VM). You will need this information in the following steps in the installation.



### Deploy the HX Data Platform Installer OVA with a Static IP Address

If your hypervisor wizard defaults to DHCP for assigning IP addresses to new VMs, deploy the Cisco HX Data Platform Installer using the following steps:

1. Log in to your Installer machine via the Hyper-V Console or putty.
2. Use ipconfig to assign an IP address to your NIC:
  - a. For example, **ifconfig eth0 x.x.x.x netmask 255.255.255.0**
3. For static IPs that will be persistent use the following configuration

```
[root@frida root]# cat /etc/sysconfig/network
```

```
NETWORKING=yes
```

```
HOSTNAME=frida.localdomain
```

```
GATEWAY=172.30.10.1
```

```
[root@frida root]#
```

```
[root@frida root]# cat /etc/sysconfig/network-scripts/ifcfg-eth0
```

```
DEVICE=eth0
```

```
BOOTPROTO=static
```

```
BROADCAST=172.30.10.255
```



```

IPADDR=172.30.10.101
NETMASK=255.255.255.0
ONBOOT=yes
[root@frida root]#

```

```

[root@frida root]# service network restart

Shutting down interface eth0: [ OK ]
Shutting down loopback interface: [ OK ]
Setting network parameters: [ OK ]
Bringing up loopback interface: [ OK ]
Bringing up interface eth0: [ OK ]
[root@frida root]#

```

b. Step 2 - Configure DNS server

```
echo "nameserver 207.62.187.54" > /etc/resolv.conf
```

4. Verify settings

a. Show interfaces and routing table

```

[root@frida root]# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:0C:29:7E:C3:EC
          inet addr:172.30.10.101  Bcast:172.30.10.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:1586 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1264 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:100
          RX bytes:181119 (176.8 Kb)  TX bytes:193645 (189.1 Kb)
          Interrupt:9 Base address:0x10a4

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:454110 errors:0 dropped:0 overruns:0 frame:0
          TX packets:454110 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:31015690 (29.5 Mb)  TX bytes:31015690 (29.5 Mb)

[root@frida root]#

```

```

[root@frida root]# route -n
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref    Use Iface
172.30.10.0      0.0.0.0         255.255.255.0   U        0      0        0 eth0
169.254.0.0      0.0.0.0         255.255.0.0     U        0      0        0 eth0
127.0.0.0        0.0.0.0         255.0.0.0       U        0      0        0 lo
0.0.0.0          172.30.10.1    0.0.0.0         UG       0      0        0 eth0
[root@frida root]#


```

### Cisco UCS Manager Configuration using HX Data Platform Installer

To configure Cisco UCS Manager using HX Installer, complete the following steps:

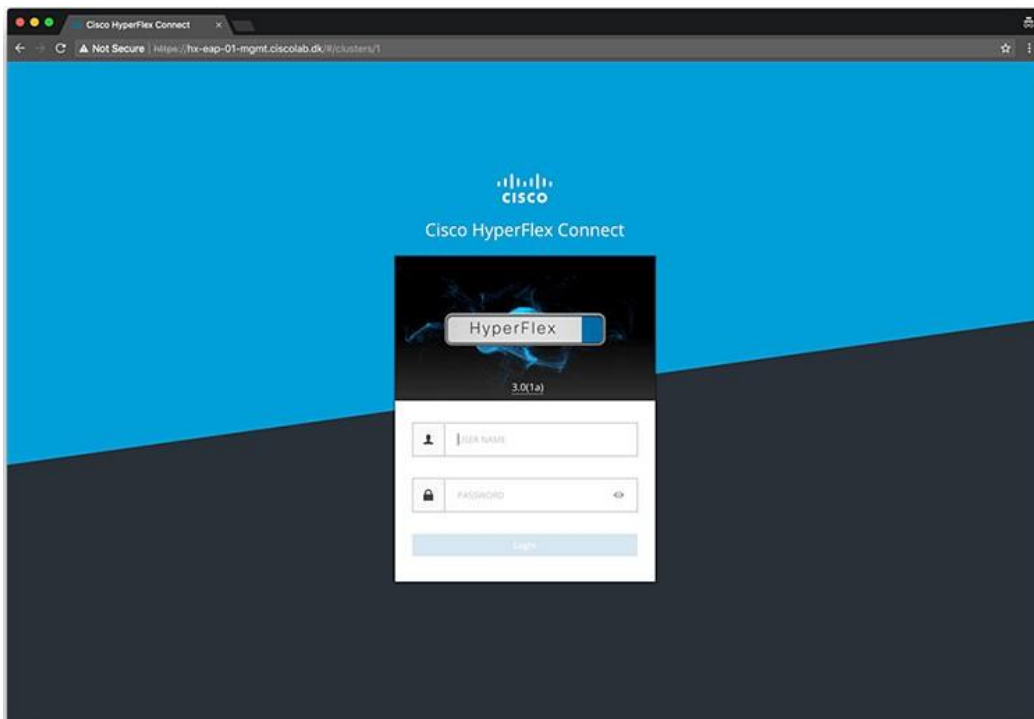
1. Log into the HX Data Platform Installer using the following steps:
  - a. In a browser, enter the URL for the VM where HX Data Platform Installer was installed. If you do not have the URL, go back to Step 13 in the earlier section on [Deploying HX Data Platform Installer](#).
  - b. Use the credentials: `username: root`, `password: Cisco123`

---

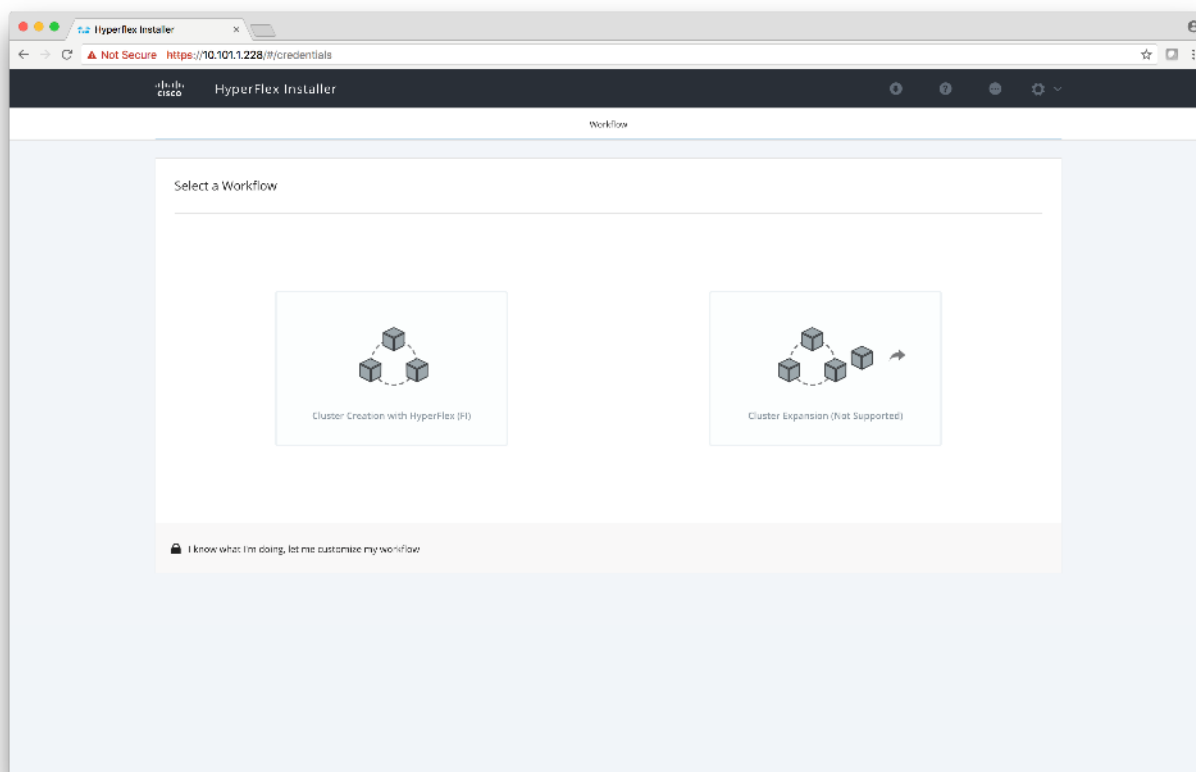
 **Important!** Systems ship with a default password of `cisco123` that must be changed during installation; you cannot continue installation unless you specify a new user supplied password.

---

2. Read the EULA. Click I accept the terms and conditions.
3. Verify the product version listed in the lower right corner is correct. This version must be 3.0(1a) or later. Click Login.



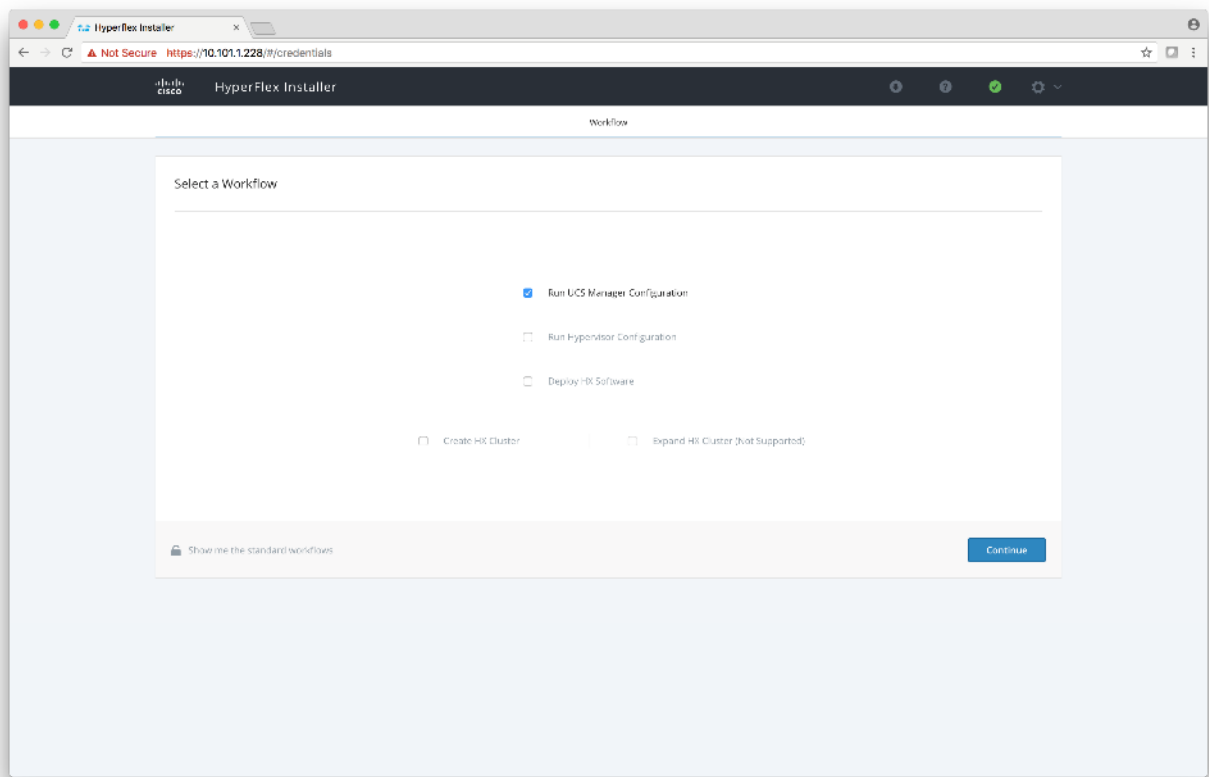
4. From the HX Data Platform Installer Workflow page, select I know what I'm doing, let me customize my workflow.




5. On the next screen, click Run UCS Manager Configuration and then click Continue.



Do not choose any other workflow options.

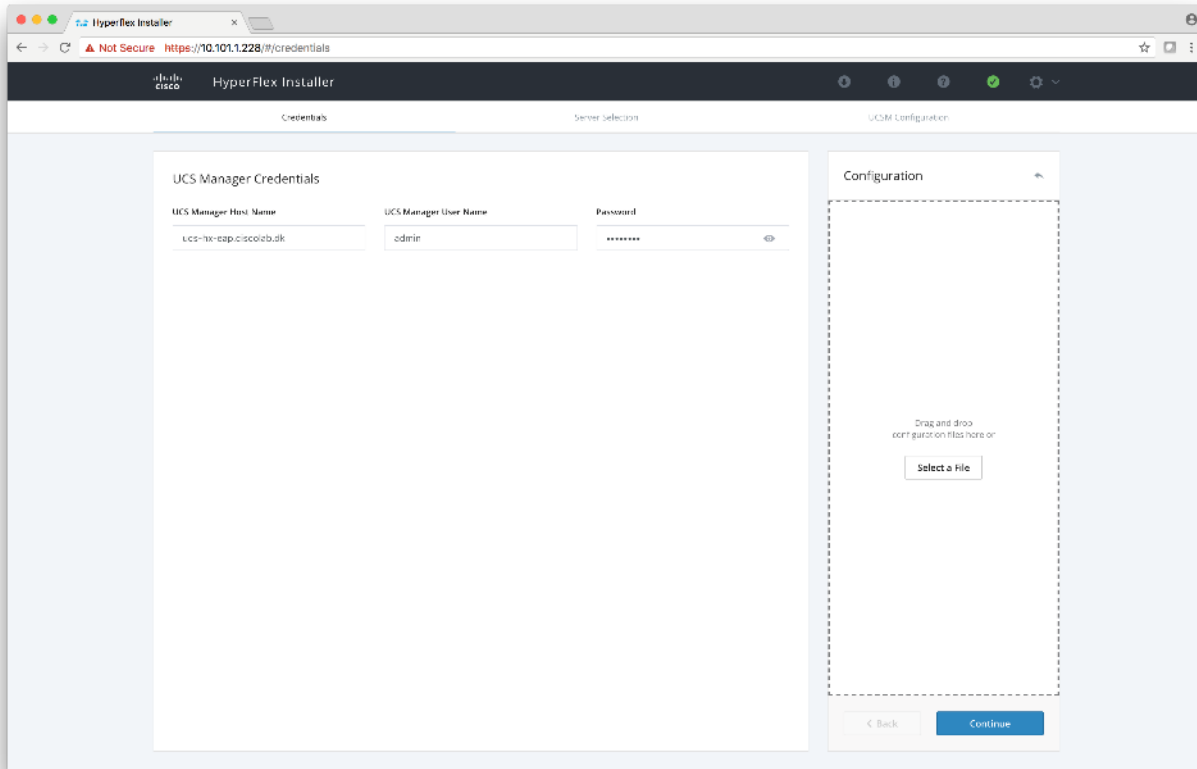


- 6. Click Confirm in the popup that displays.
- 7. Enter the UCS Manager credentials.

 The right side of the page is unused. Further in the setup process a configuration JSON is saved, so in subsequent installations the JSON file can be imported to add the data quickly.

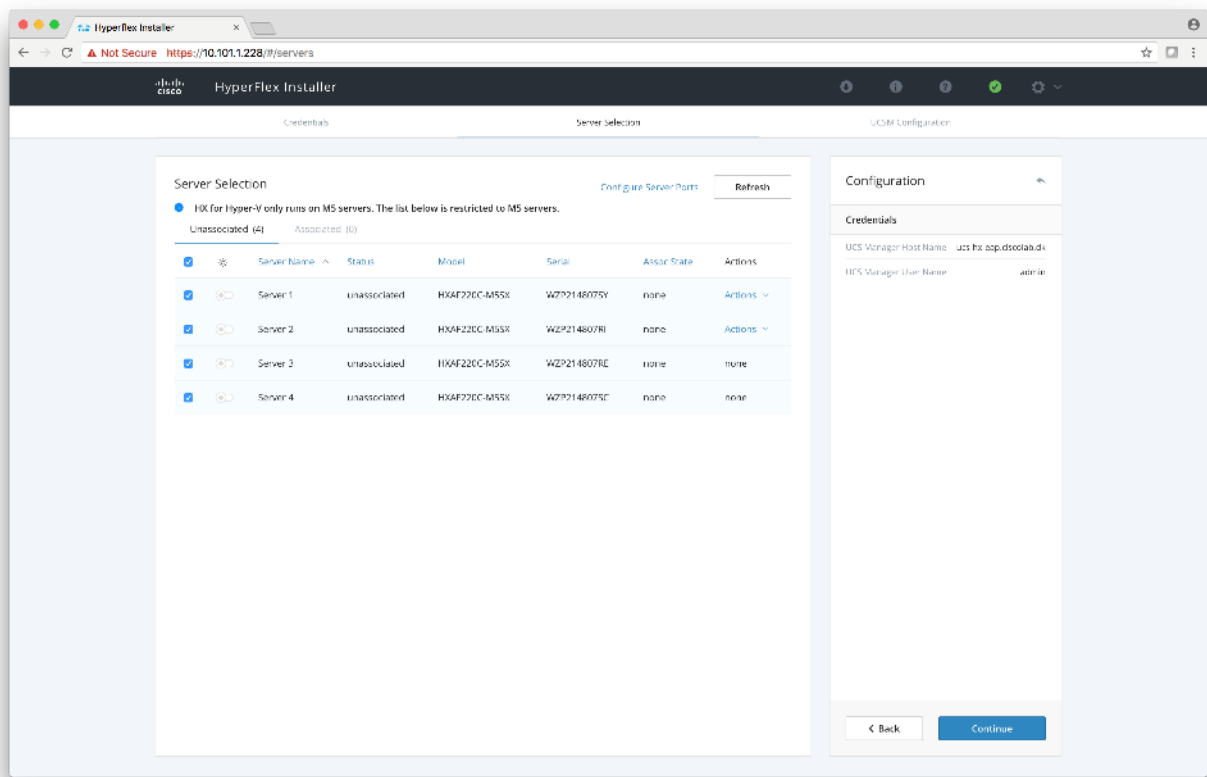
- 8. Click Confirm and Proceed to bypass the warning. Complete the following fields for Cisco UCS Manager.

Field	Description
Cisco UCS Manager Host Name	FQDN or the VIP address of Cisco UCS Manager
Cisco UCS Manager User Name and Password	Administrator user and password or an user with Cisco UCS Manager admin rights



9. Click Continue.

The installer will connect to UCSM and query for available servers. The configuration pane is populated as the installer progresses. You can at any time save the JSON file so you can re-use it for subsequent installations. This feature works on all the different workflows in the installer. After the query finishes a screen displays showing the available servers.




10. Choose all the servers that you want to install in the cluster and click Continue.


 HyperFlex for Hyper-V only supports M5 Servers.

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).

Name	Usage	ID
hx-inband-mgmt	Hyper-V and HyperFlex VM management.	30
hx-storage-data	HyperFlex storage traffic	101
hx-livemigrate	Hyper-V LiveMigration network	33
vm-network	VM guest network	34,35

 Do not use vlan 1 as it is not best practice and can cause issues with disjoint layer 2.

 vm-network can be multiple VLANs added as a comma separated list.



Renaming the 4 core networks is not supported.

The following illustration shows the various fields in the VLAN Configuration pane where you need to enter values.

11. Enter the remaining network configuration.

VLAN Configuration

VLAN for Hypervisor and HyperFlex management

VLAN Name

hx-inband-mgmt

VLAN ID

VLAN for HyperFlex storage traffic

VLAN Name

hx-storage-data

VLAN ID

VLAN for VM Live Migration

VLAN Name

hx-livemigrate

VLAN ID

VLAN for VM Network

VLAN Name

vm-network

VLAN ID(s)

Field	Description	Value
MAC pool prefix	MAC address pool for the HX cluster, to be configured in UCSM by the installer. Ensure that the mac address pool isn't used anywhere else in your layer 2 environment.	00:25:b5:xx
IP blocks	The range of IP addresses that are used for Out-Of-Band management of the hyperflex nodes	10.193.211.124-.127
Subnet Mask	The subnet mask for the Out-Of-Band network	255.255.0.0
Gateway	The gateway address for the Out-Of-Band network	10.193.0.1

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.

MAC Pool

MAC Pool Prefix

00:25:b5:

'hx-ext-mgmt' IP Pool for Out-of-band CIMC

IP Blocks

ex: 10.193.211.124-127;10.193.211.158-161

Subnet Mask

ex: 255.255.0.0

Gateway

ex: 10.193.0.1



iSCSI Storage and FC Storage are used for adding external storage to the HyperFlex cluster.



This is currently not supported for the Hyper-V Edition.

---

Advanced Settings

Use the table below to complete the fields in this section.

Table 10    Field Settings

Field	Description	Example Value
UCS Firmware Server Version	Choose the appropriate UCS Server Firmware version.	3.2(3a)
HyperFlex Cluster Name	This user defined name will be used as part of the service profile naming In UCSM for easier identification	
Org Name	The org. name is used for isolating the HX environment from the rest of the UCS platform to ensure consistency.	HX-Cluster1



The Cisco UCS C and B 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.

---



The supported version for HyperFlex Hyper-V is 3.2(3a).

---

### VLAN Configuration

**VLAN for Hypervisor and HyperFlex management**

VLAN Name:  VLAN ID:

**VLAN for HyperFlex storage traffic**

VLAN Name:  VLAN ID:

**VLAN for VM Live Migration**

VLAN Name:  VLAN ID:

**VLAN for VM Network**

VLAN Name:  VLAN ID(s):

### MAC Pool

MAC Pool Prefix:

### "hx-ext-mgmt" IP Pool for Out-of-band CIMC

IP Blocks:  Subnet Mask:  Gateway:

> iSCSI Storage

> FC Storage

### Advanced

UCS Server Firmware Version:  HyperFlex Cluster Name:  Org Name:

### Configuration

**Credentials**

UCS Manager Host Name:  UCS Manager User Name:

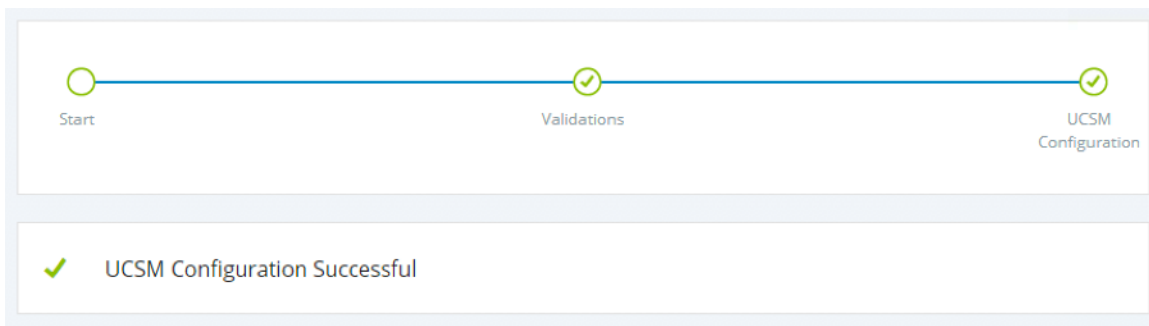
**Server Selection**

Server 2	WQF214007R1 / HXAF223C-MDSX
Server 3	WQF214007R6 / HXAF223C-MDSX
Server 5	WQF214007R1 / HXAF223C-MDSX
Server 6	WQF214007R1 / HXAF223C-MDSX

< Back Start

12. Click Start. The installer validates your input and then begins configuring Cisco UCS Manager.

13. When the HX Data Platform Installer is finished, then you are ready to proceed to next step, [Microsoft Windows OS and Hyper-V Installation, on page 26](#).



## Microsoft Windows OS and Hyper-V Installation

For this part of the installation, you need the Windows 2016 Datacenter Edition ISO and the Cisco provided [Cisco HyperFlex Driver image](#).

The two files must be placed on a share that is reachable from the Cisco UCS Manager and the Out-of-band subnet that was used in the previous installation step.

The following protocols are supported:

- NFS
- CIFS
- HTTP

If you do not have a place to serve the files from, you can use the installer to host the files. Please see the section: [How to Upload the ISO and IMG File to the Installer VM using WinSCP](#).



Make sure network connectivity exists between the file share and all server management IPs.

---

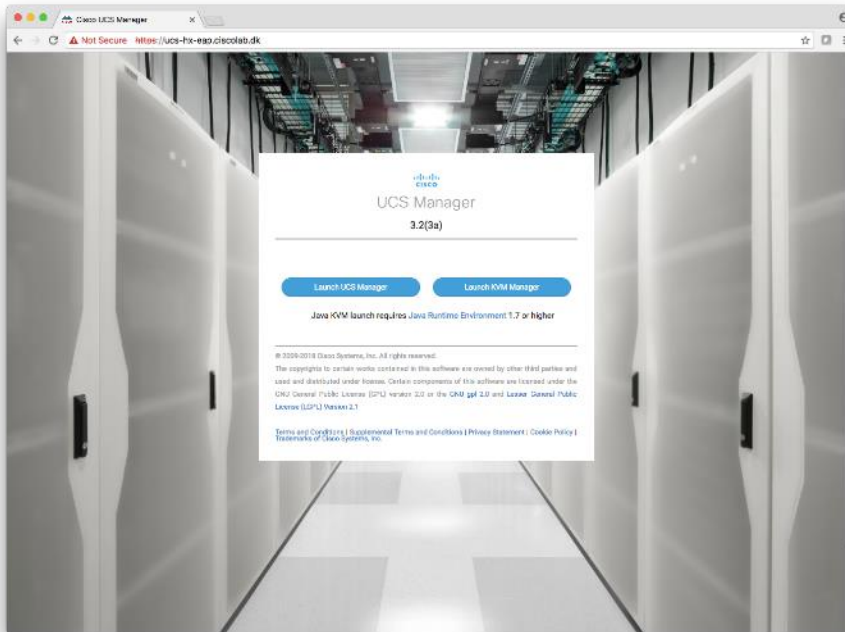
Below is a summary of steps:

1. Attach the two images to the service profiles. Follow the steps described in [Configure vMedia and Boot Policies through Cisco UCS Manager, on page 26](#).
2. Verify the images are mounted correctly.
3. Reboot the servers and verify the OS installation completes successfully.
4. Clean up the service profiles so the HX Installer can continue.

### Configure vMedia and Boot Policies through Cisco UCS Manager

To configure the Cisco UCS vMedia and Boot Policies using Cisco UCS Manager, complete the following steps:

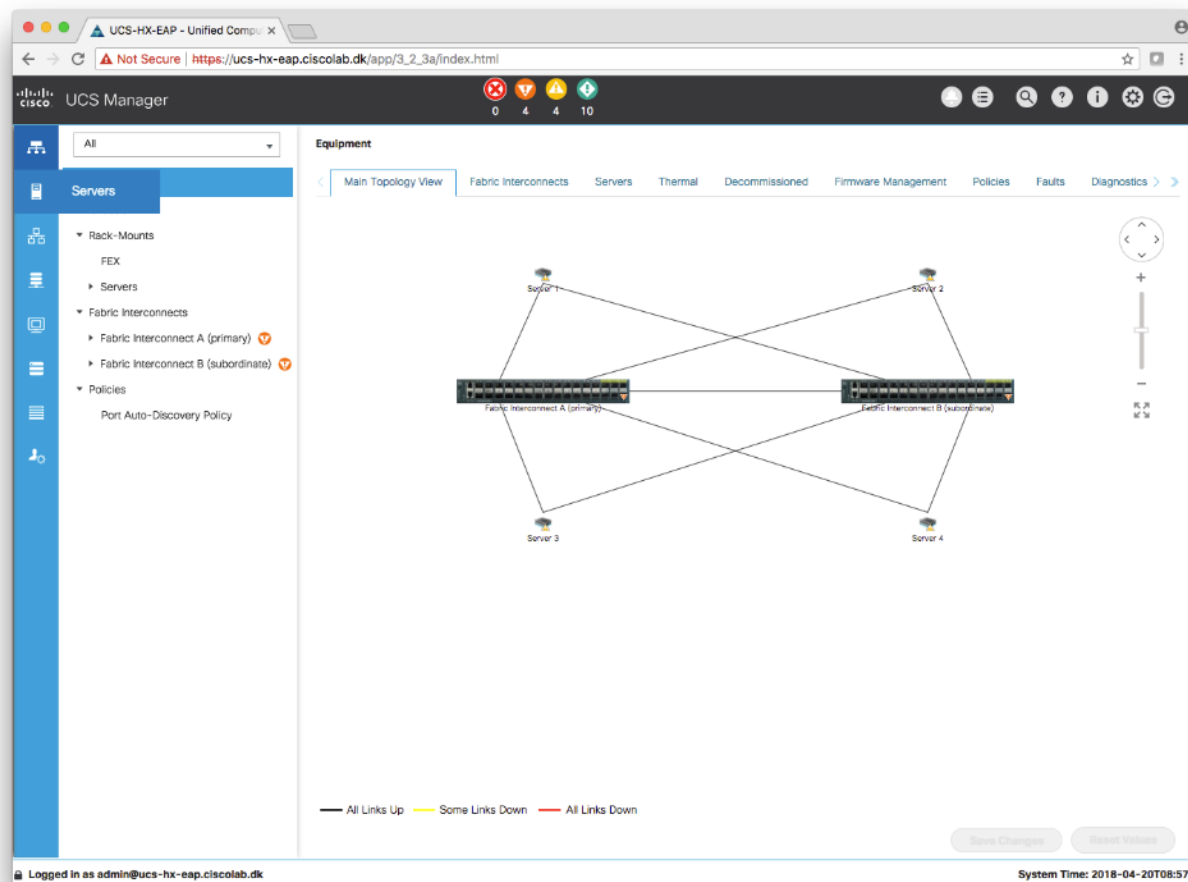
1. Launch Cisco UCS Manager by accessing the Cisco UCS Manager IP address in a browser of your choice.



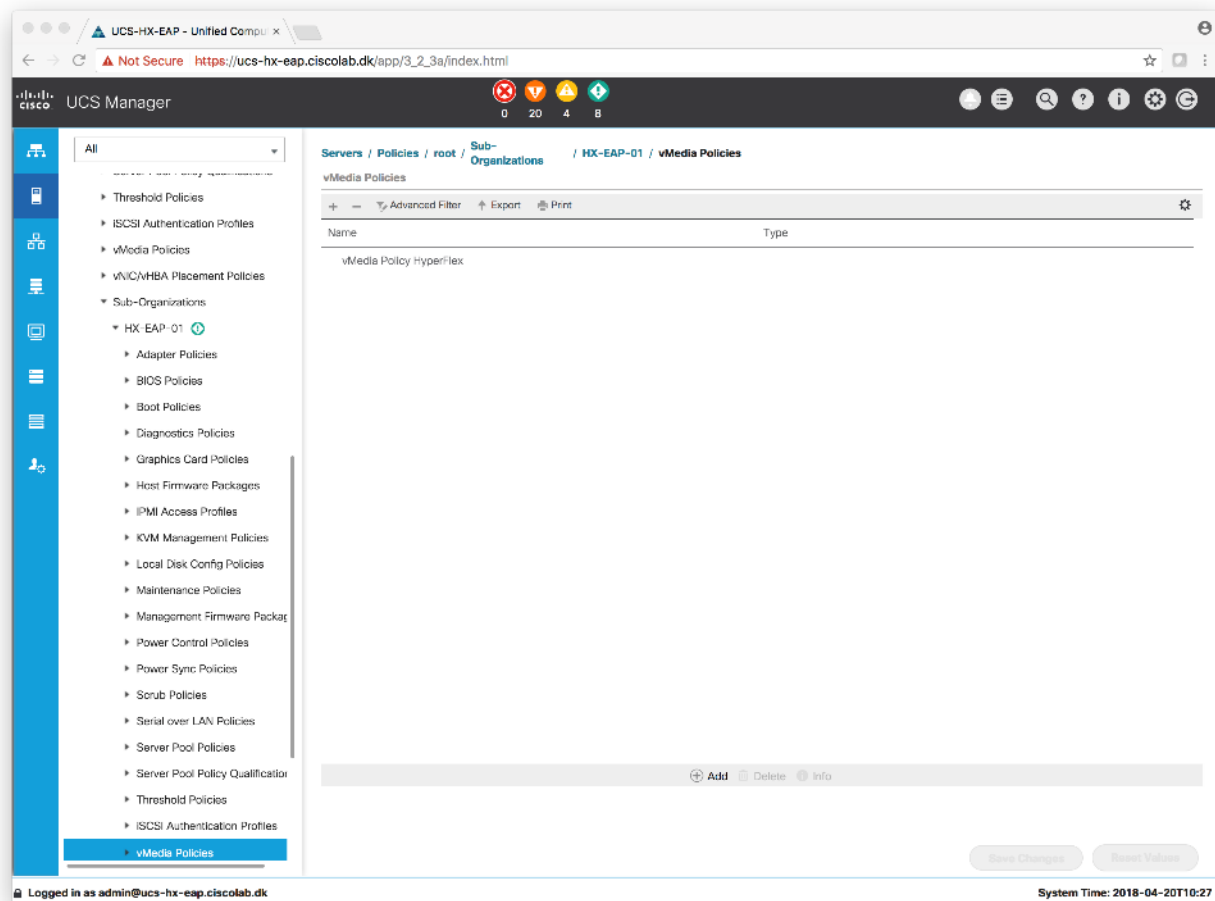
2. Click Launch UCS Manager and log in with administrator username and the password you used at the beginning of the installation.

A screenshot of the Cisco UCS Manager login page. The page has a clean, white background with the 'Cisco UCS Manager 3.2(3a)' logo at the top. Below the logo, there are two input fields: 'Username' with the text 'admin' and 'Password' which is masked with dots. A blue 'Log In' button is positioned below the password field. At the bottom of the page, there is a small information icon and a note: 'For best results use a supported browser'. Below this, there is a copyright notice for Cisco Systems, Inc. 2009-2018.

3. In the left navigation pane, click Servers.



4. Expand Servers > Policies > root > Sub-Organizations > hx-cluster\_name > vMedia Policies to view the list of vMedia Policies.



5. Double-click vMedia Policy HyperFlex.

Properties for: vMedia Policy HyperFlex

GeneralEvents

Actions

Create vMedia Mount

Delete

Show Policy Usage

Use Global

Properties

Name : HyperFlex

Description : vMedia policy to install or re-install software on Hyp

Owner : Local

Retry on Mount Failure : ☐ No ☒ Yes

vMedia Mounts

+ - Advanced Filter Export Print

Name	Type	Protocol	Authent...	Server	Filename	Remote...	User	Remap ...
No data available								

+ Add Delete Info

OK

Apply

Cancel

Help

6. In the properties for vMedia Policy HyperFlex, click Create vMedia Mount to add the mount points.
7. In the Create vMedia Mount dialog box, complete the following fields:

Create vMedia Mount

Name :

Description :

Device Type : ☐ CDD ☐ HDD

Protocol : ☐ NFS ☐ CIFS ☐ HTTP ☐ HTTPS

Hostname/IP Address :

Image Name Variable : ☒ None ☐ Service Profile Name

Remote File :

Remote Path :

Username :

Password :

OK

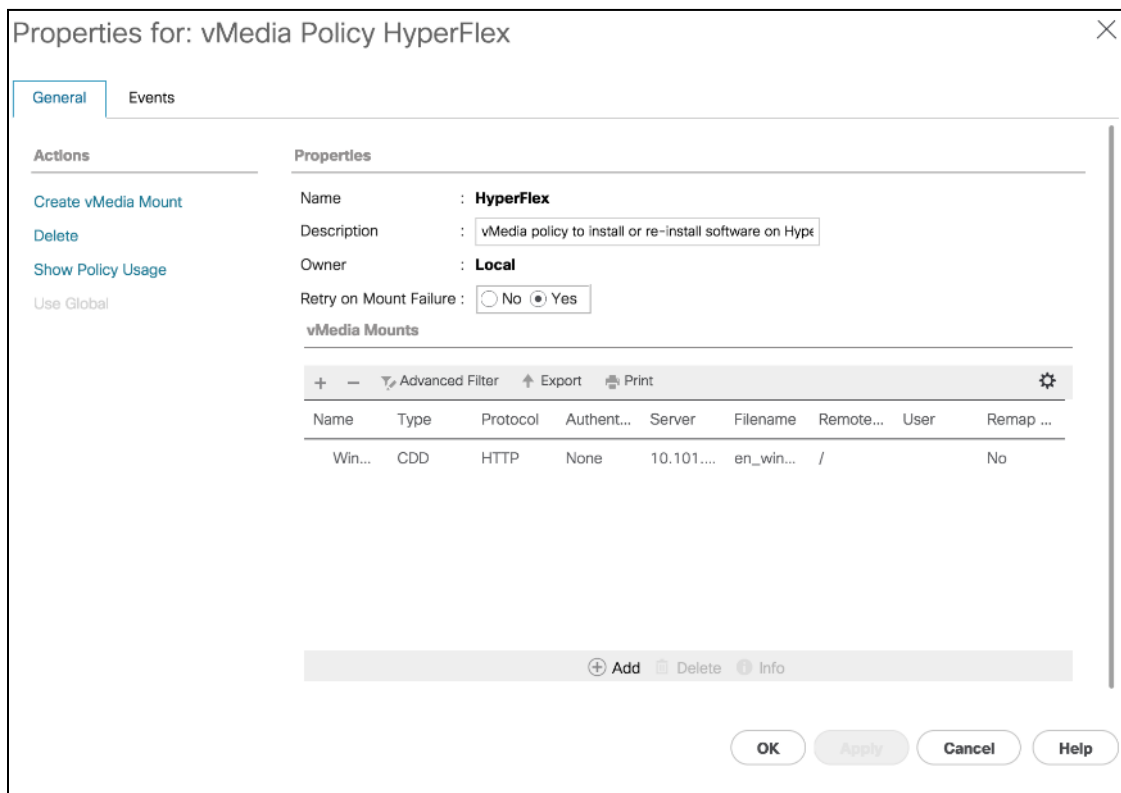
Cancel

Field Name	Action	Example Value
Name	Name for the mount point.	OS Install
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.	10.101.1.92
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.	
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	

8. Click Save Changes and click OK.
9. Click OK. When you click OK, you are returned to the vMedia policy and will see the information that you submitted.



10. Repeat steps 5 and 6 but change the type to HDD and the filename to the Cisco HyperFlex driver image.

Create vMedia Mount

Name

:

HX\_Driver

Description

:

Cisco HyperFlex driver image

Device Type

:

CDD

HDD

Protocol

:

NFS

CIFS

HTTP

HTTPS

Hostname/IP Address

:

10.101.1.92

Image Name Variable

:

None

Service Profile Name

Remote File

:

HXInstall-HyperV-v3.0.1a-29499.img

Remote Path

:

/

Username

:

Password

:

OK

Cancel

On completion, the following screen displays:

Properties for: vMedia Policy HyperFlex

GeneralEvents

Actions

Create vMedia Mount

Delete

Show Policy Usage

Use Global

Properties

Name

:

HyperFlex

Description

:

vMedia policy to install or re-install software on Hyp

Owner

:

Local

Retry on Mount Failure

:

No

Yes

vMedia Mounts

+ - Advanced Filter Export Print

Name	Type	Protocol	Authent...	Server	Filename	Remote...	User	Remap ...
HX_...	HDD	HTTP	None	10.101....	HXInsta...	/		No
Win...	CDD	HTTP	None	10.101....	en_win...	/		No

Add

Delete

Info

OK

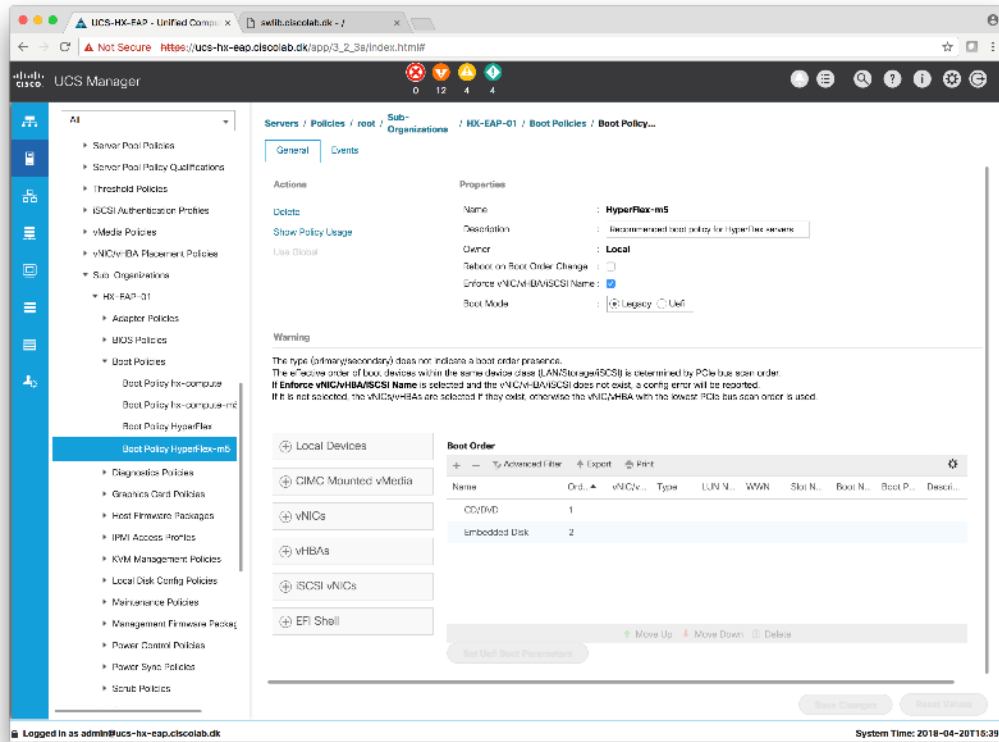
Apply

Cancel

Help

11. In the left navigation pane, select Servers > Service Profile Templates > root > Sub-Organizations > hx-cluster\_name > Service Template hx-nodes\_name (example:hx-nodes-m5)

114



12. Choose the HyperFlex vMedia Policy from the drop-down list and click OK twice.

The vMedia policy is assigned to the HyperFlex Template during the Cisco UCS Manager phase of the HyperFlex deployment.

13. Select Servers > Policies > root > Sub-Organizations > hx-cluster\_name > Boot Policies Boot Policy Hyper-Flex-m5.

14. In the configuration pane, click CIMC Mounted vMedia. Click Add CIMC Mounted CD/DVD to add this to the boot order.

15. Select the CIMC Mounted CD/DVD entry in the list and move it to the top of the boot order by pressing the Move Up button.

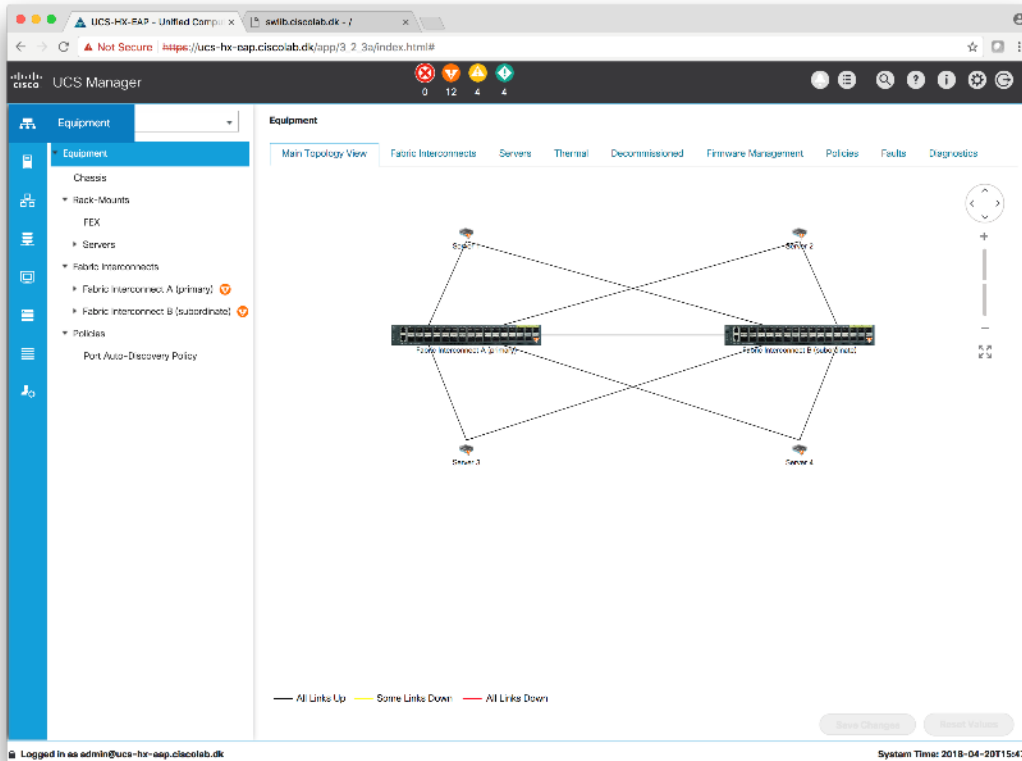
Boot Order									
<div> + - Advanced Filter Export Print </div>									
Name	Ord...	vNIC/v...	Type	LUN N...	WWN	Slot N...	Boot N...	Boot P...	Descri...
CD/DVD	1								
Embedded Disk	2								
CIMC Mounted CD/...	3								
<div> Move Up Move Down Delete </div>									

16. Click Save Changes and click OK. The boot policy is saved.

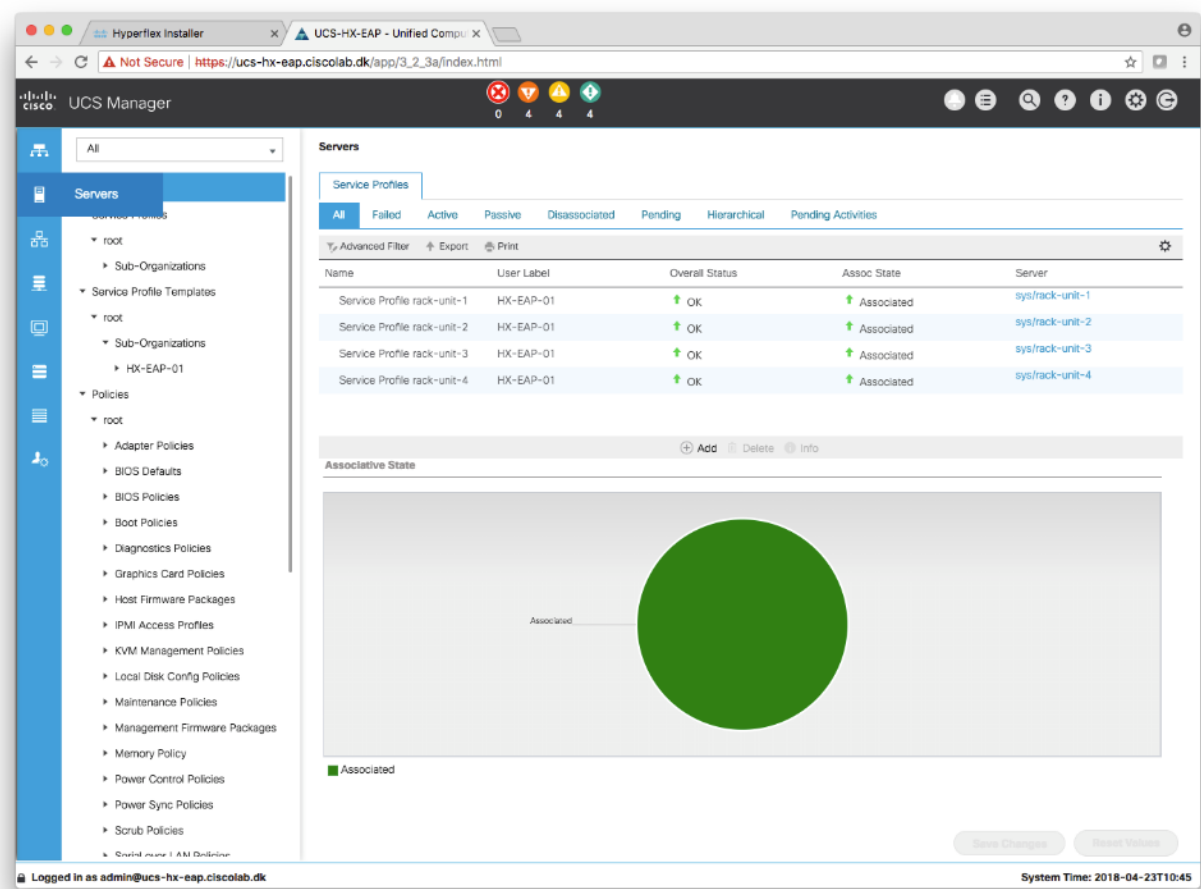
## Next Steps

To verify the images are mounted correctly, complete the following steps:

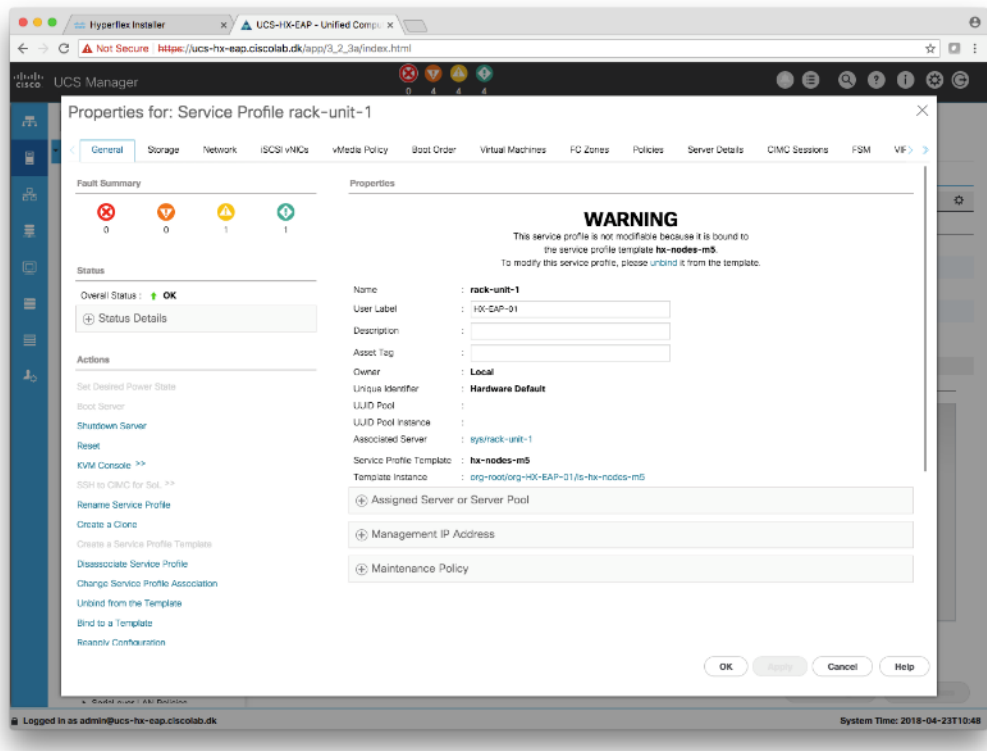
1. On the Equipment tab, select one of the servers.



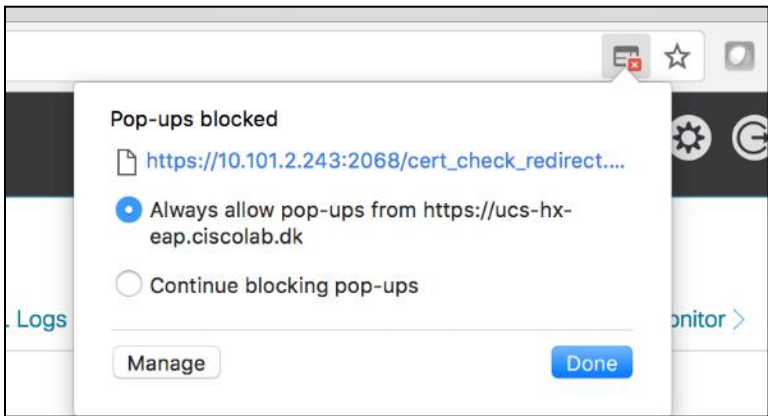
2. Click Inventory > CIMC, scroll down and make sure for the mount entry #1(OS image) and mount entry #2 (Cisco HyperFlex driver image) the status is Mounted and there are no failures.
3. In the menu bar, click Servers and choose the first HyperFlex service profile.




4. Click the General tab and choose 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.



5. Reboot the server. In the KVM console choose Server Actions and press Reset.



### 6. Choose Power Cycle.

Do you want to reset the selected servers? ✕

You are attempting to reset a server. The server can be reset by gracefully restarting the OS or via a brute force power cycle. How would you like to reset?

☒ Power Cycle  
☐ Gracefully restart OS

If Graceful OS Restart is not supported by the OS or it does not happen within a reasonable amount of time, the system will perform a power cycle.

The UCS system might be in the process of performing some tasks on this server. Would you like this operation to wait until the completion of outstanding activities?

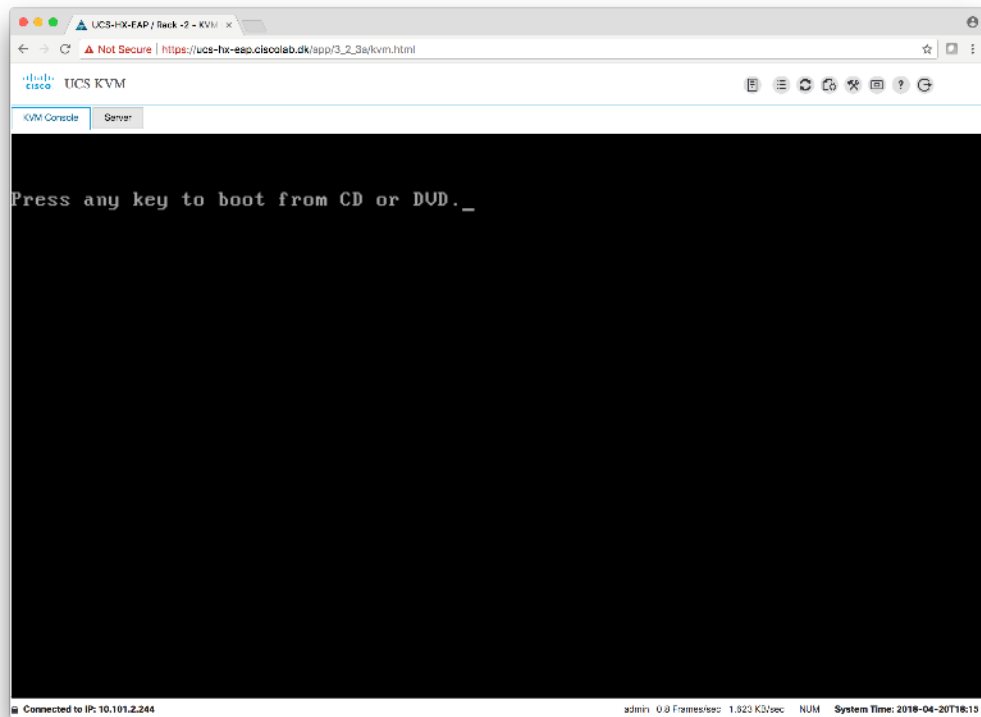
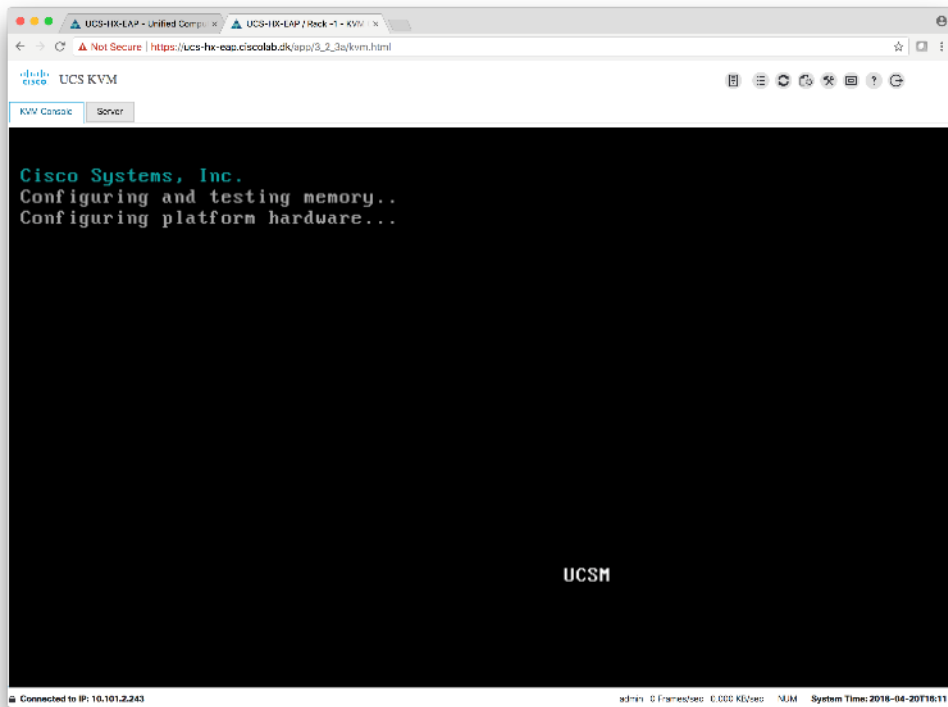
☐ Wait for completion of outstanding UCS tasks on this server.

OK

Cancel

The server should reboot when the server has finished. Remember to press any key to start the Windows installer.

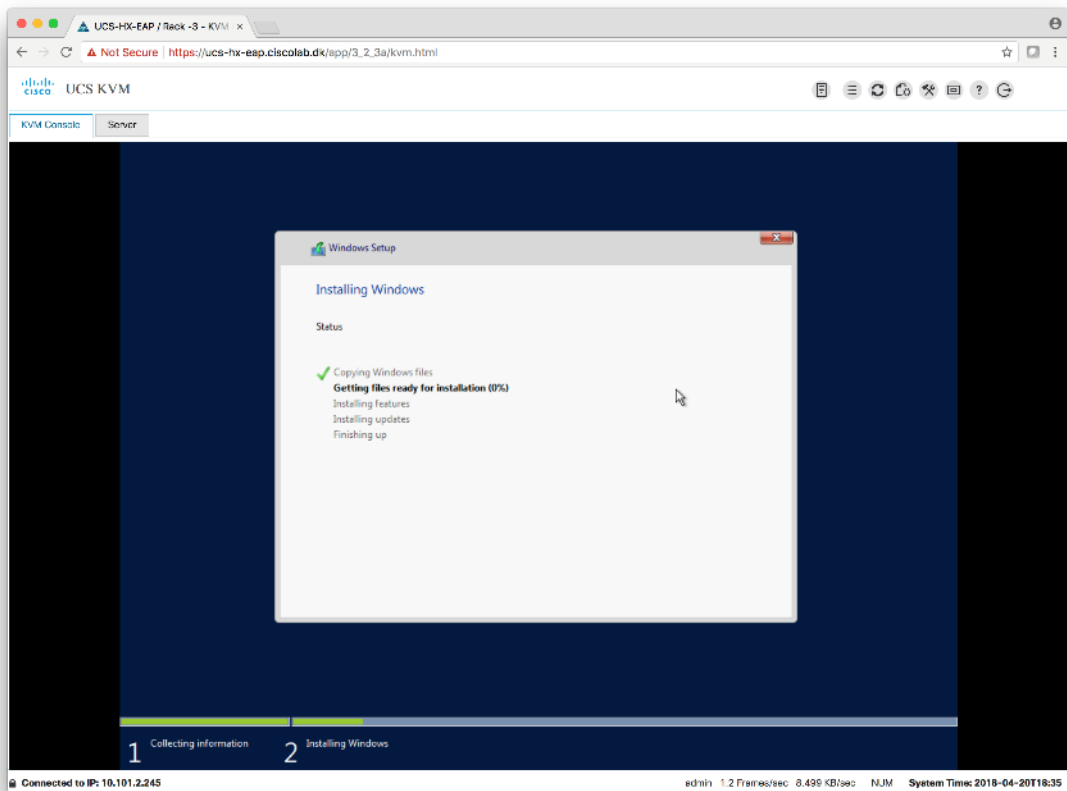




If you miss pressing any key, the server will display in the windows installation or an error page displays stating no OS is installed.

## Solution Configuration

When the server starts booting the Windows image, it starts the windows installation and continues the installation automatically.




During the installation the server will reboot automatically a few times.



Do not stop the installation and allow the process to continue.

The installation is complete when a clean command prompt displays as shown below:

## Solution Configuration



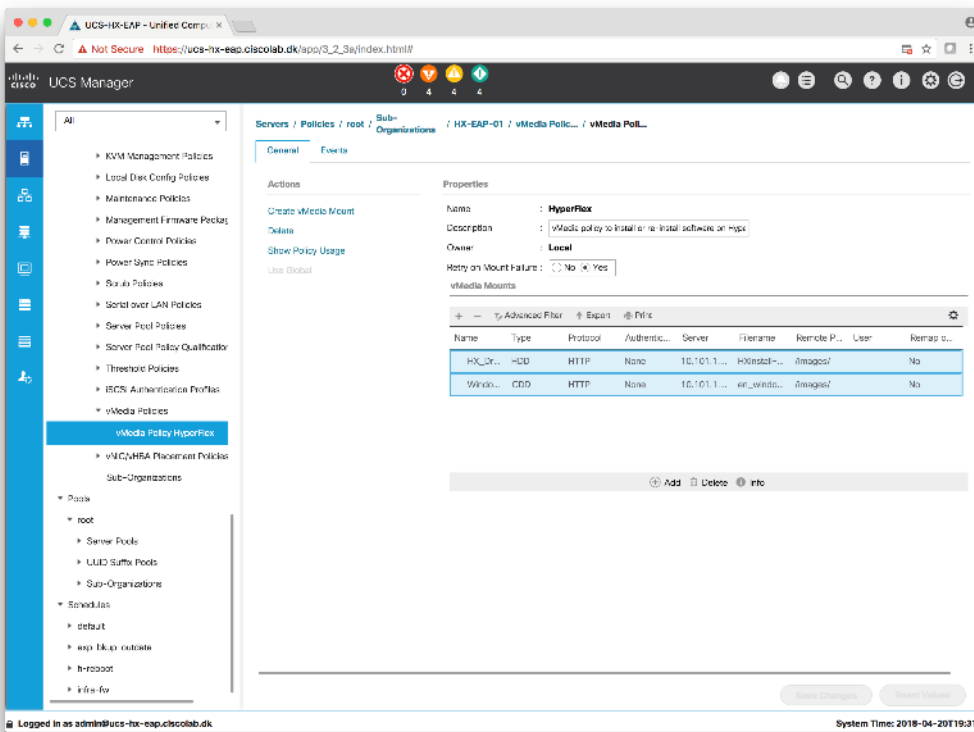
7. Repeat these steps on all the nodes in the cluster.

## Clean Up

When the installation is complete on all nodes, you must clean up the vMedia policy and the boot policy so they return to their previous state before continuing.

1. In Cisco UCS Manager select Servers > Policies > Root > Sub-Organizations > HX-Cluster\_name > vMedia policies
2. Click the vMedia Policy HyperFlex. Click the mount points one at a time and delete both of them. Accept the changes.

## Solution Configuration



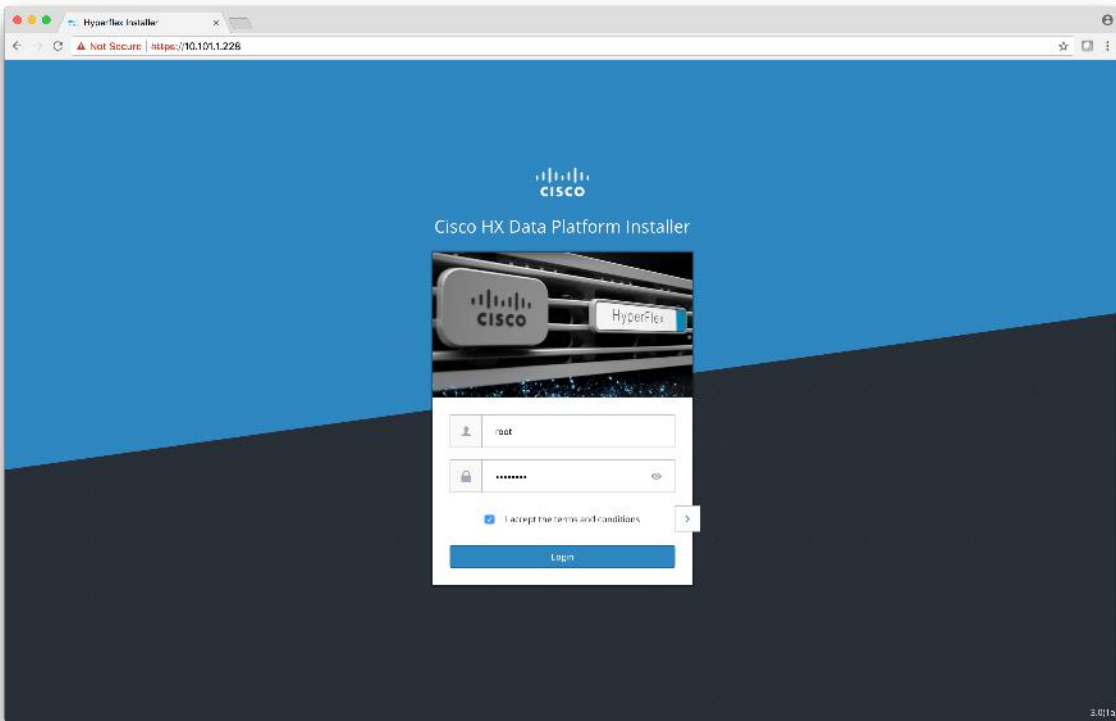
3. Go to the boot policy by selecting Servers > Policies > Root > Sub-Organizations > HX-Cluster\_name > boot policies > Boot Policy HyperFlex-m5
4. Select the CIMC mounted CD/DVD, click Delete and accept the changes.

## Hypervisor Configuration

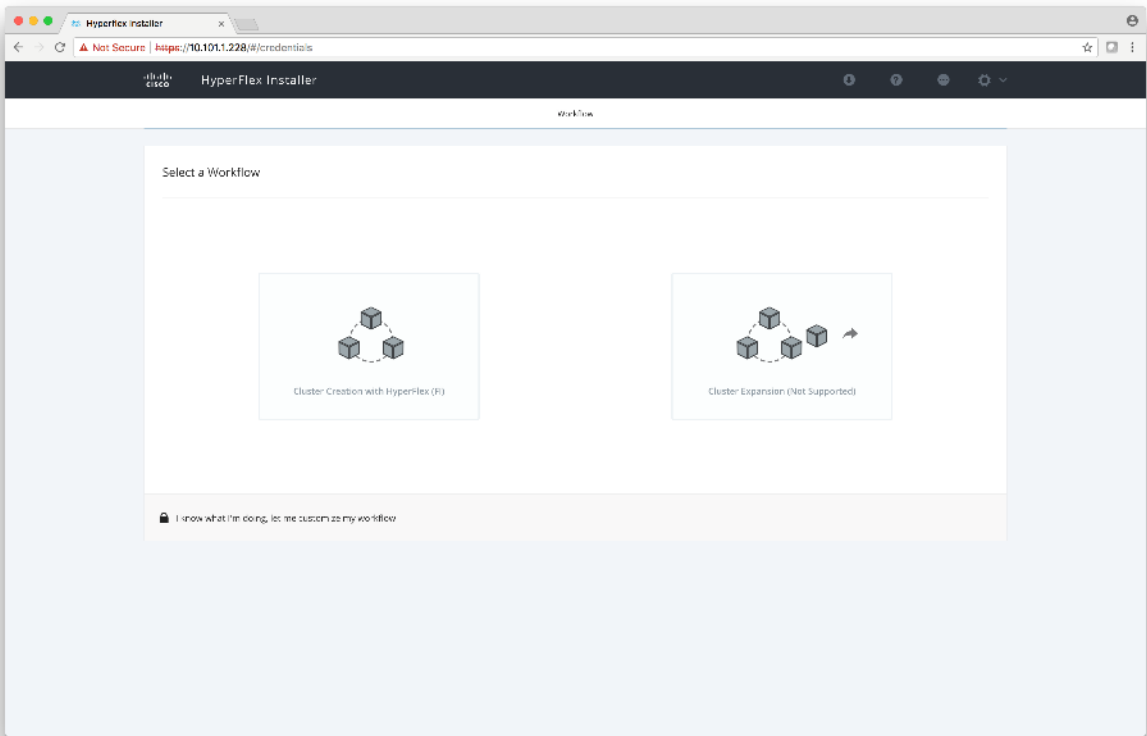
After the OS is installed, you need to configure the hypervisor before installing the HX Software and configure the cluster.

To configure the Hypervisor, complete the following steps:

1. Open the HX Data Platform Installer and log in.




2. 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.
3. In the main menu, select I know what I'm doing, let me customize my workflow. In the Warning dialog box, click Confirm and Proceed.



- 4. Select Run Hypervisor Configuration and click Continue.
- 5. Complete the information for the UCS Manager, Domain Information, and Hypervisor Credentials.

Field	Description	Default Value
UCS Manager Host Name	FQDN or the VIP address of the UCSM	
UCS Manager User Name	Admin user or an user with UCSM admin rights	Admin
Password	Password for the UCS Manager User Name	
Domain Name	Active Directory domain name that the HyperFlex cluster is going to be a member off.	
Local Administrator User Name	Local Administrative username on the Hyper-V Hosts	Administrator
Local Administrator Password	Password for the local administrative user on the Hyper-V hosts	Cisco 123

 If you have not changed the Administrator password for the Windows Hyper-V in the previous step, the default value is shown.

## Solution Configuration

The screenshot shows the 'UCS Manager Credentials' and 'Configuration' sections of the HX Data Platform Installer. The 'UCS Manager Credentials' section includes fields for 'UCS Manager Host Name' (ucs-hx-eap.ciscolab.dk), 'UCS Manager User Name' (admin), and 'Password' (masked). The 'Domain Information' section has a 'Domain Name' field (ciscolab.dk). The 'Hypervisor Credentials' section includes 'Local Administrator User Name' (Administrator) and 'Local Administrator Password' (masked). The 'Configuration' section on the right has a dashed box for configuration files with a 'Select a File' button. At the bottom are 'Back' and 'Continue' buttons.

The HX Data Platform Installer now connects to UCSM and lists the relevant servers for the HX Cluster. The HX Data Platform Installer will validate UCS Firmware, etc.

6. Validate the selected servers and click Continue.
7. Complete the network information as you have done in section Cisco UCS Manager Configuration using HX Data Platform Installer and make sure the data is the same. Click Continue.

The screenshot shows the 'VLAN Configuration' and 'Configuration' sections. The 'VLAN Configuration' section includes fields for 'VLAN for Hypervisor and HyperFlex management' (VLAN Name: hxc-hypervisors, VLAN ID: 2486), 'VLAN for HyperFlex storage traffic' (VLAN Name: hxc-storage-data, VLAN ID: 2597), 'VLAN for VM Live Migration' (VLAN Name: hxc-livemigrate, VLAN ID: 2586), and 'VLAN for VM Network' (VLAN Name: vmm-network, VLAN ID: 2599). It also includes a 'MAC Pool' section (MAC Pool Prefix: 02:25:00:51) and an 'IP Block' section (IP Block: 10.101.2.232-247, Subnet Mask: 255.255.255.0, Gateway: 10.101.2.1). The 'Configuration' section on the right includes 'Credentials' (UCS Manager Host Name: ucs-hx-eap.ciscolab.dk, UCS Manager User Name: admin, Domain Name: ciscolab.dk, Time Zone: CEST, Local Administrator User Name: Administrator), 'Server Selection' (listing four servers with their IP addresses), and 'UCSM Configuration' (listing various VLANs and their IDs, including hxc-hypervisors, hxc-storage-data, hxc-livemigrate, vmm-network, and hxc-external-mgmt). At the bottom are 'Back' and 'Continue' buttons.

8. Configure Hypervisor Settings. Input the values for the Hypervisor configuration as show below:



Field	Description	Example Value
Configure common Hypervisor Settings		
Subnet Mask	Subnet mask for the hypervisor hosts management network	255.255.255.0
Gateway	Default gateway for the hypervisor hosts management network	10.101.251.1
DNS Servers	Comma separated list for the DNS Servers in the AD that the hypervisor hosts are going to be member of.	10.99.2.200,10.99.2.201
Hypervisor Settings		
Static IP address	Management IP address for each host	10.101.251.41
Hostname	Hostname for each host	HX-Hypv-01



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.

### Configure common Hypervisor Settings

Subnet Mask

255.255.255.0

Gateway

10.101.251.1

DNS Server(s)

10.99.2.200,10.99.2.201

### Hypervisor Settings

☒ Make IP Addresses and Hostnames Sequential

#	Name	Serial	Static IP Address	Hostname
1	Server 1	WZP214807SY	10.101.251.41	HX-EAP-01
2	Server 2	WZP214807RI	10.101.251.42	HX-EAP-02
3	Server 3	WZP214807RE	10.101.251.43	HX-EAP-03
4	Server 4	WZP214807SC	10.101.251.44	HX-EAP-04

Primary DNS Suffix

CiscoLab.dk

Additional DNS Suffixes

### Configuration

#### Credentials

UCS Manager Host Name

ucs-hx-eap.ciscolab.dk

UCS Manager User Name

admin

Domain Name

Ciscolab.dk

Time Zone

CST

Local Administrator User Name

Administrator

#### Server Selection

Server 2

WZP214807RI / HXAF220C-M55X

Server 3

WZP214807RE / HXAF220C-M55X

Server 1

WZP214807SY / HXAF220C-M55X

Server 4

WZP214807SC / HXAF220C-M55X

#### UCSM Configuration

VLAN Name

hx-inband-mgmt

VLAN ID

2986

VLAN Name

hx-storage-data

VLAN ID

2987

VLAN Name

hx-ovmigrate

VLAN ID

2988

VLAN Name

vm-network

VLAN ID(s)

2989

MAC Pool Prefix

0E:25:90:51

IP Blocks

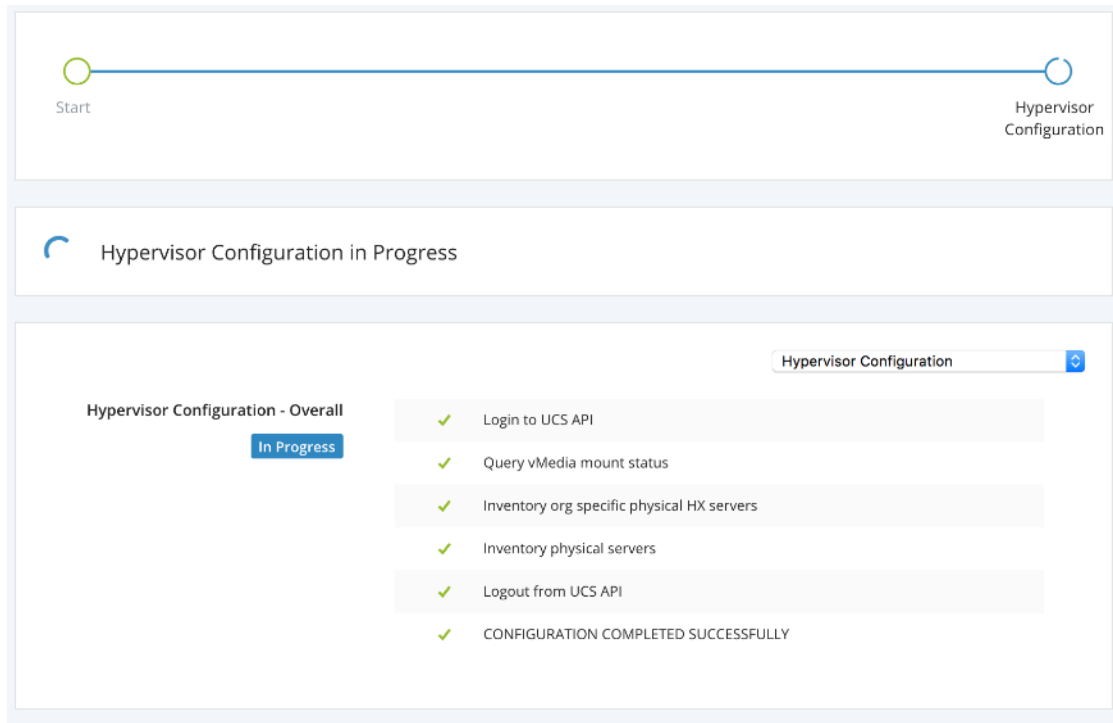
10.101.3.243-247

< Back

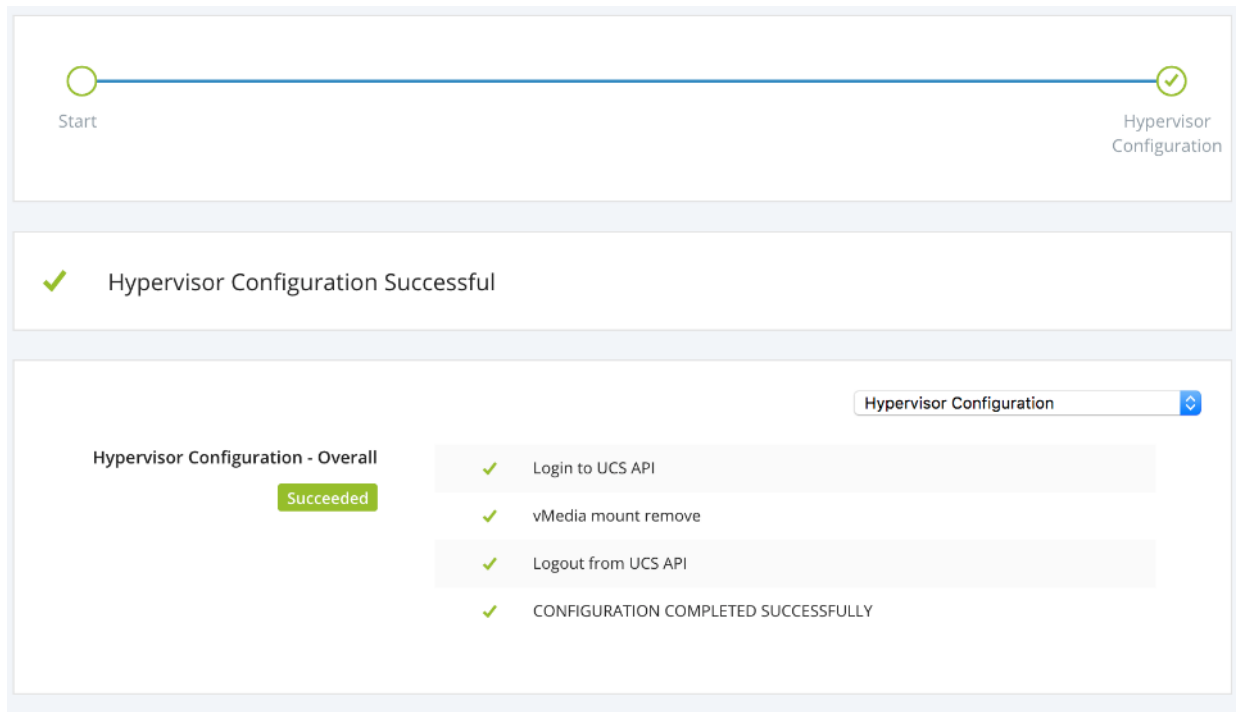
Start

9. Click Start to begin the Hypervisor Configuration.

The installation continues and configures the Hypervisor hosts.



Be aware that if the steps are completed as shown above, the Hypervisor configuration is not completed. The servers are working in the background until the installer reports an overall completion as shown below.



## Deploying HX Data Platform Installer and Cluster Configuration

1. Log into HX Data Platform Installer.

2. In the top right corner of the install, select Start Over, and confirm that you wish to start over. From the HX Data Platform Installer Workflow page, select I know what I'm doing, let me customize my workflow.
3. Check both Deploy HX Software and Create HX Cluster and Continue.
4. Click Confirm and Proceed at the Warning message.
5. Enter Domain Information, Hypervisor credentials, etc.

Field	Description	Example Value
Domain name	Active Directory domain that the cluster is going to a part of.	contoso.com
HX Service Account	The HX Service account that was created during the Pre-installation checklist	hxadmin
Password	Password for the HX Service account	
<b>Configure Constrained Delegation Now</b>		Checked
Use HX Service Account	Uses the HX Service Account for the Constrained delegation. The user must be a Domain admin.	Checked
<b>Advanced Attributes (optional)</b>		
Domain controller	FQDN for the Domain controller that you want to use specifically for the installation.	dc.contoso.com
Organization Unit	The OU that was created during the Pre-Installation can be filled out here, and then this OU will be home for the HX Nodes in the Active directory	OU=HyperFlex,DC=contoso,DC=com.
<b>Hypervisor Credentials</b>		
Local Administrator User Name	Local Administrative username on the Hyper-V Hosts	Administrator
Local Administrator Password	Password for the local administrative user on the Hyper-V hosts	Cisco123

6. Enter IP Addresses.
7. Click the Add server button for you to have the amount of server entries that you need for your cluster.
8. Enter the hostnames for the Hyper-Hosts and the Storage Controllers running on the Hyper-V hosts. These hostnames must be added to forward and reverse lookup prior to this step and remember that only Windows AD Integrated DNS is supported.



That the management VLAN uses DNS to resolve the addresses and the Data VLAN does not.

Table 11 Cluster Addresses

Field	Description	Example Value
<b>Management</b>		
Cluster Address	Hostname for the HX Connect UI	HX-EAP-01-MGMT
Subnet Mask	Subnet mask for the management VLAN	255.255.255.0
Gateway	Gateway address for the management VLAN	10.101.251.1
<b>Data</b>		
Cluster Address	IP address for the HX cluster on the data VLAN	10.101.252.50
Subnet Mask	Subnet mask for the HX data VLAN	255.255.255.0
Gateway	Gateway for the HX data VLAN	10.101.252.1

Table 12 Cisco HX Cluster

Field	Description	Example Value
<b>Cisco HX Cluster</b>		
Cluster Name (SMB Access Point	The cluster name will be used as the FQDN For the datastores.	HX-EAP-01
Replication Factor	How many copies of data you want.	3 (default)
Failover Cluster Name	This name is used for the windows failover cluster.	HX-EAP-CLU01
<b>Controller VM</b>		
Create Admin Password	The Installer will automatically change the default password on the controllers to a user defined password*	
Confirm Admin Password	Confirm the previous password.	
<b>System Services</b>		
DNS Servers	Comma separated lists of DNS Servers.	10.99.2.200,10.99.2.201
NTP Servers	The Controller VM's needs to be in sync with the Windows AD therefore you need to point to your AD domain controllers for time synchronization.	dc1.contoso.com,dc2.contoso.com
DNS Domain Name	The Domain name for the Active Directory.	contoso.com
Time Zone	The time zone in which you want the HX controllers to report.	
<b>Auto Support</b>		
Enable Connected Services	Auto Support will ship telemetry data of the HX cluster to Cisco Support.	Checked

Field	Description	Example Value
Send Service ticket to	Add your own email address or alias so you can received a copy of the ticket sent to Cisco.	<a href="mailto:con@contoso.com">con@contoso.com</a>
<b>Advance Network</b>		
Management VLAN Tag	VLAN used for the Management network, it must be the same as used earlier in the installation process for the management network.	Management VLAN tag
Data VLAN Tag	VLAN used for the data network, it must be the same as used earlier in the installation process for the data network.	Data VLAN tag
<b>Advance Configuration</b>		
Enable Jumbo Frames on Data Network	Ensures that we are running Jumbo frames on the links connected to the storage VM's.	Checked
Disk Partitions	If you want to clean up you environment during a re-installation check this box.	Unchecked
VDI	Optimize for VDI only deployment	Checked
<b>Hypervisor Settings</b>		
Primary DNS suffix	Already filled out by some of the earlier steps.	
Additional DNS Suffixes	If you need more suffixes appended on your Hyper-V hosts.	

9. Fill in the fields on the page according to the following information:



The password must be complex, have a minimum of 10 characters, and include at least 1 uppercase letter, 1 digit, 1 special character.

10. Click Continue.

Cisco HX Cluster

Cluster Name (SMB Access Point)

hx-eap-01

Replication Factor

3

Fallover Cluster Name ⓘ

HX-EAP-CLU01

Controller VM

Create Admin Password

\*\*\*\*\*

Confirm Admin Password

\*\*\*\*\*

System Services

DNS Server(s)

10.99.2.200,10.99.2.201

NTP Server(s)

CiscoLab.dk

DNS Domain Name

ciscoLab.dk

Time Zone

(UTC+01:00) Brussels, Copenhagen, Madrid, Paris ⓘ

Auto Support

Auto Support

☒ Enable Connected Services (Recommended) ⓘ

Send service ticket notifications to

lagranbo@cisco.com

Advanced Networking

Management VLAN Tag

2996

Management vSwitch

vswitch-hx-inband-mgmt

Data VLAN Tag

2997

Data vSwitch

vswitch-hx-storage-data

Advanced Configuration

Jumbo Frames

☒ Enable Jumbo Frames on Data Network ⓘ

Disk Partitions

☒ Clean up disk partitions

Virtual Desktop (VDI)

☐ Optimize for VDI only deployment

Configuration

Credentials

Domain Name

CiscoLab.dk

HX Service Account

hxadmin

Time Zone

Romance Standard Time

Local Administrator User Name

Administrator

IP Addresses

Cluster Name (SMB Access Point)

hx-eap-01

Management Cluster

HX-EAP-01-MGMT

Data Cluster

10.101.252.50

Management Subnet Mask

255.255.255.0

Data Subnet Mask

255.255.255.0

Management Gateway

10.101.251.1

Data Gateway

10.101.252.1

Server 0

Management Hypervisor

HX-EAP-1.CiscoLab.dk

Management Storage Controller

HX-EAP-1-CNTL.CiscoLab.dk

Data Hypervisor

10.101.252.41

Data Storage Controller

10.101.252.51

Server 1

Management Hypervisor

HX-EAP-2.CiscoLab.dk

Management Storage Controller

HX-EAP-2-CNTL.CiscoLab.dk

Data Hypervisor

10.101.252.42

Data Storage Controller

10.101.252.52

Server 2

Management Hypervisor

HX-EAP-3.CiscoLab.dk

Management Storage

HX-EAP-3-

< Back

Start



After clicking Start, the installation and configuration of HX will begin and can take up to 2 hours to complete.

Start

Deploy Validation

Deploy

Create Validation

Cluster Creation

Deploy in Progress

Deploy - Overall

In Progress

HX-EAP-1.Ciscolab.dk

In Progress

✓

HyperV Host Add

Getting Ready to Copy Required Powershell, Task and XML files to Windows Host  
Chmod Keyless Password File (later copy to VM)

HX-EAP-2.Ciscolab.dk

In Progress

✓

HyperV Host Add

Getting Ready to Copy Required Powershell, Task and XML files to Windows Host  
Chmod Keyless Password File (later copy to VM)

HX-EAP-3.Ciscolab.dk

In Progress

✓

HyperV Host Add

Getting Ready to Copy Required Powershell, Task and XML files to Windows Host  
Chmod Keyless Password File (later copy to VM)

HX-EAP-4.Ciscolab.dk

In Progress

✓

HyperV Host Add

Getting Ready to Copy Required Powershell, Task and XML files to Windows Host  
Chmod Keyless Password File (later copy to VM)

Config

Credent

Domain P

HX Servic

Time Zon

Local Adr

IP Address

Cluster N

Managen

Data Clus

Managen

Data Sub

Managen

Data Gat

Server 0

Managen

Managen

Controlle

Data Hyp

Data Stor

Server 1

Managen

Managen

Controlle

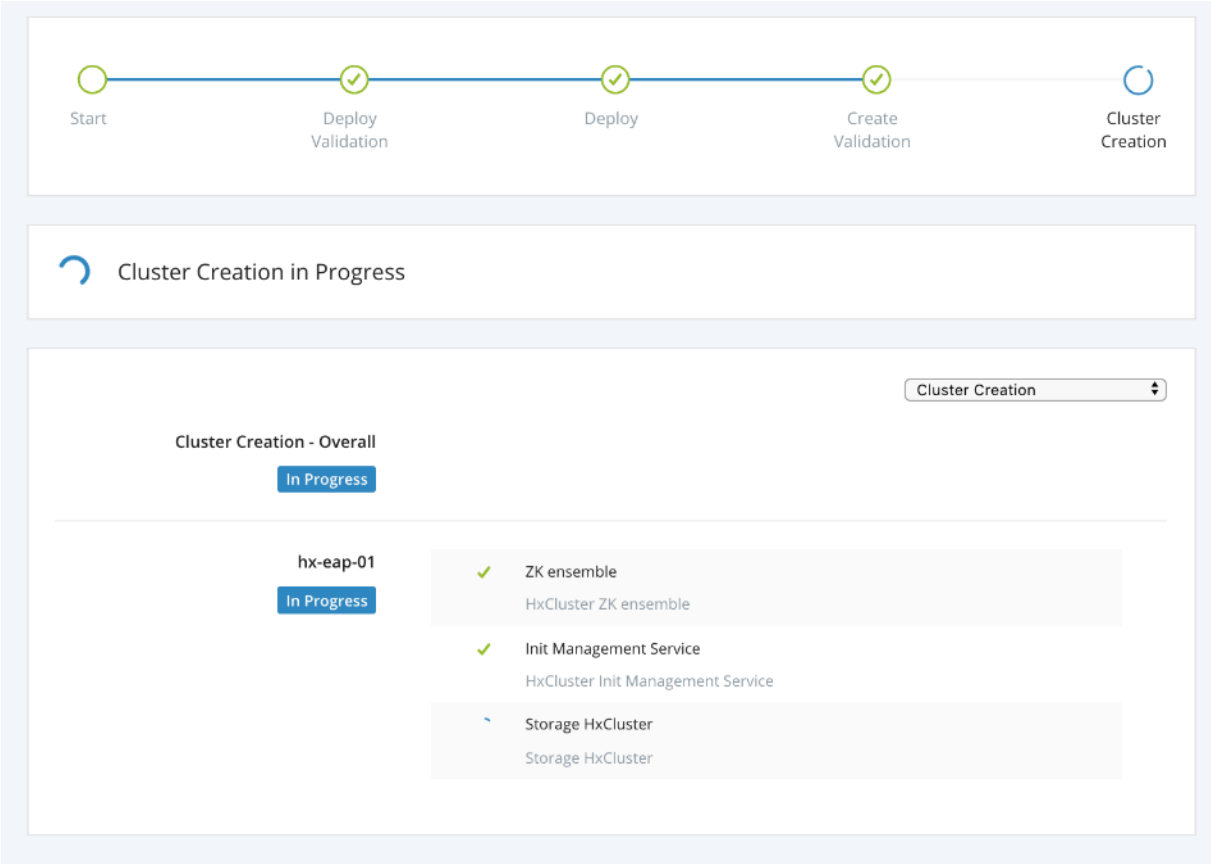
Data Hyp

Data Stor

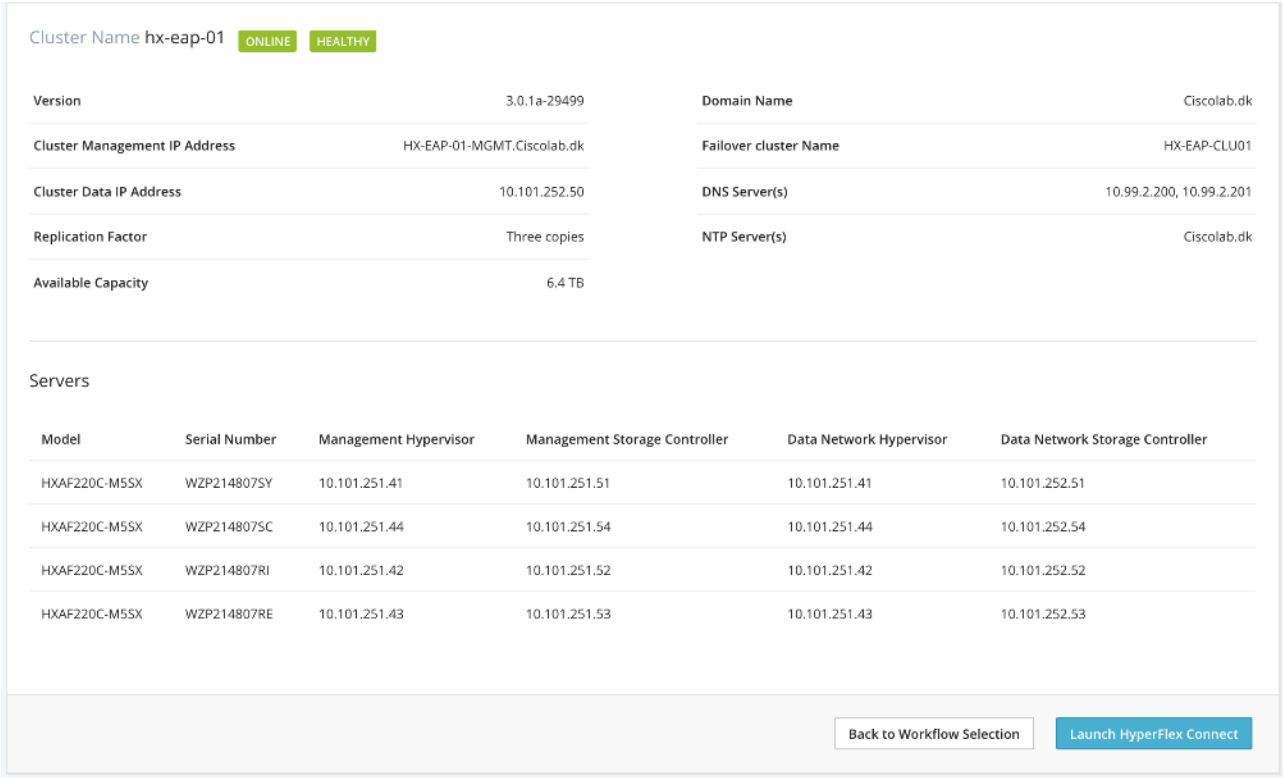
Server 2

Managen

Configuring the hosts:



The cluster is created and on completion you will see the following screen:

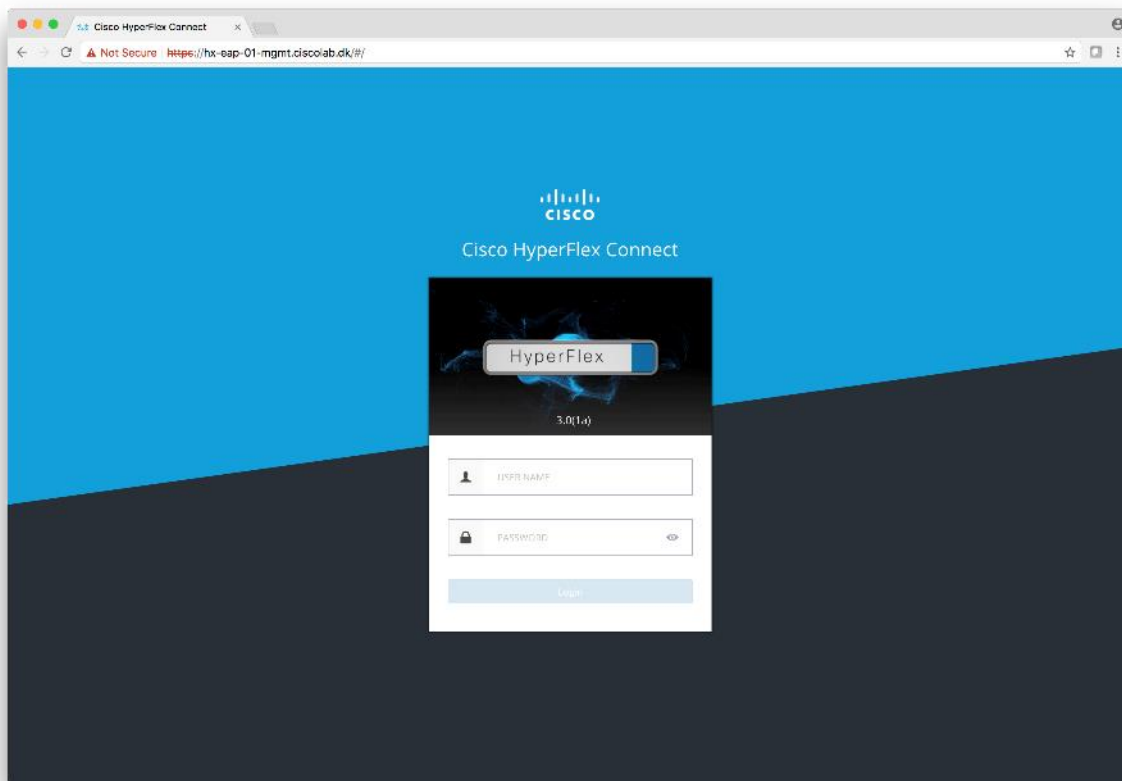




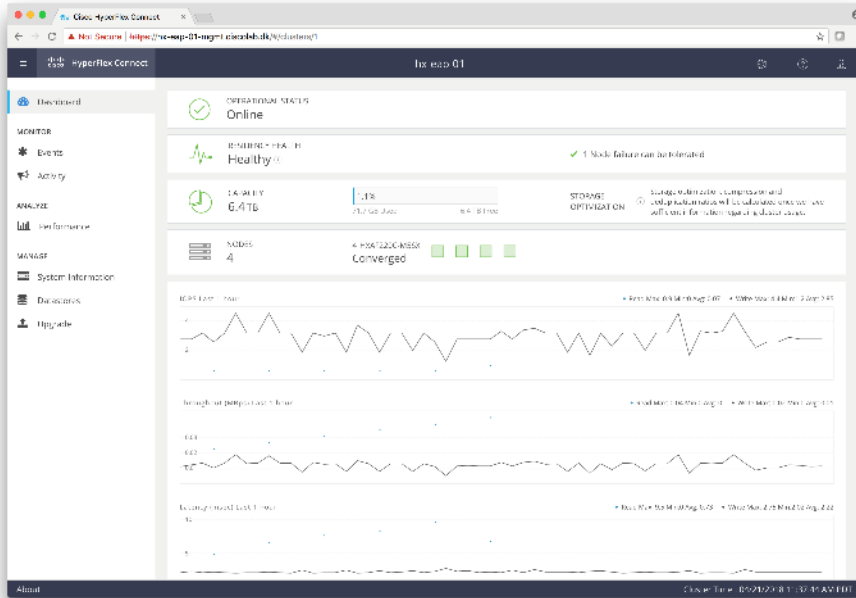
## Cluster Validation

To verify that the cluster is healthy, complete the following steps:

1. Log in to the HyperFlex Connect UI and verify that it reports cluster as healthy.
2. Click Launch HyperFlex Connect on the screen to access the HyperFlex Connect UI.



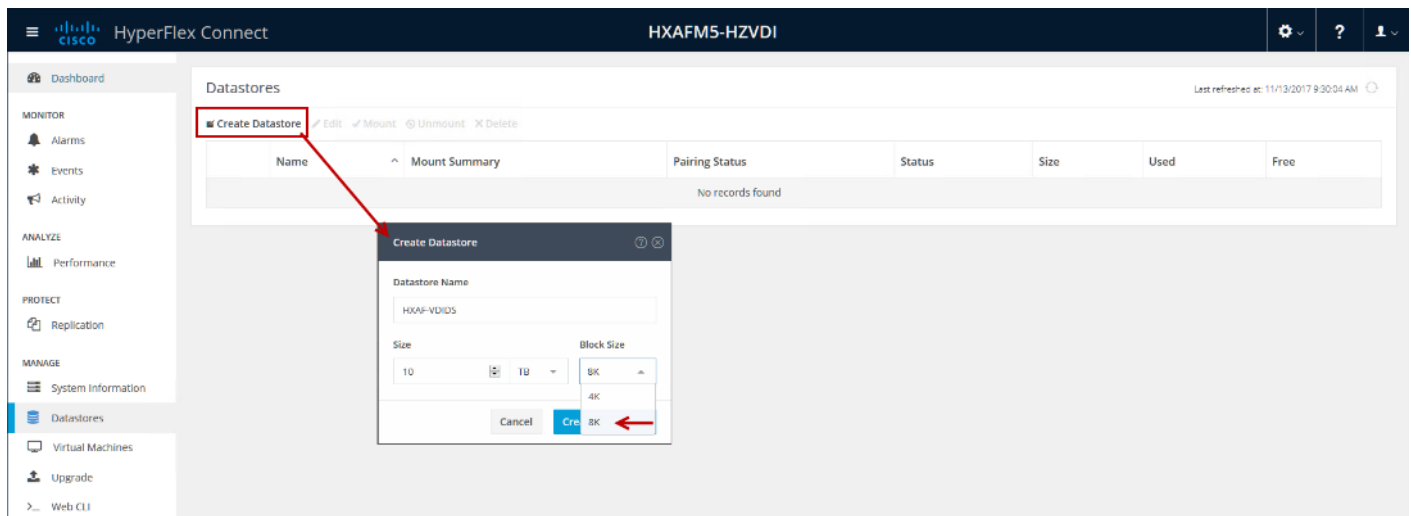
3. HyperFlex Connect UI displays. Log in with your newly created hxadmin user and password that you provided during the installation. If the cluster is healthy, it will appear as shown in the image below:



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



4. Use the HyperFlex connect UI to create a datastore. While using HyperFlex Connect UI to create a datastore there is an option to select the block size. 8K is the default.



## Post HyperFlex Cluster Installation for Hyper-V 2016

---

The HyperFlex installer configures all components of the environment from the 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 XenDesktop to use the environment. Below are the steps to prepare the environment for VDI use.

### Create Run As Account in VMM

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.

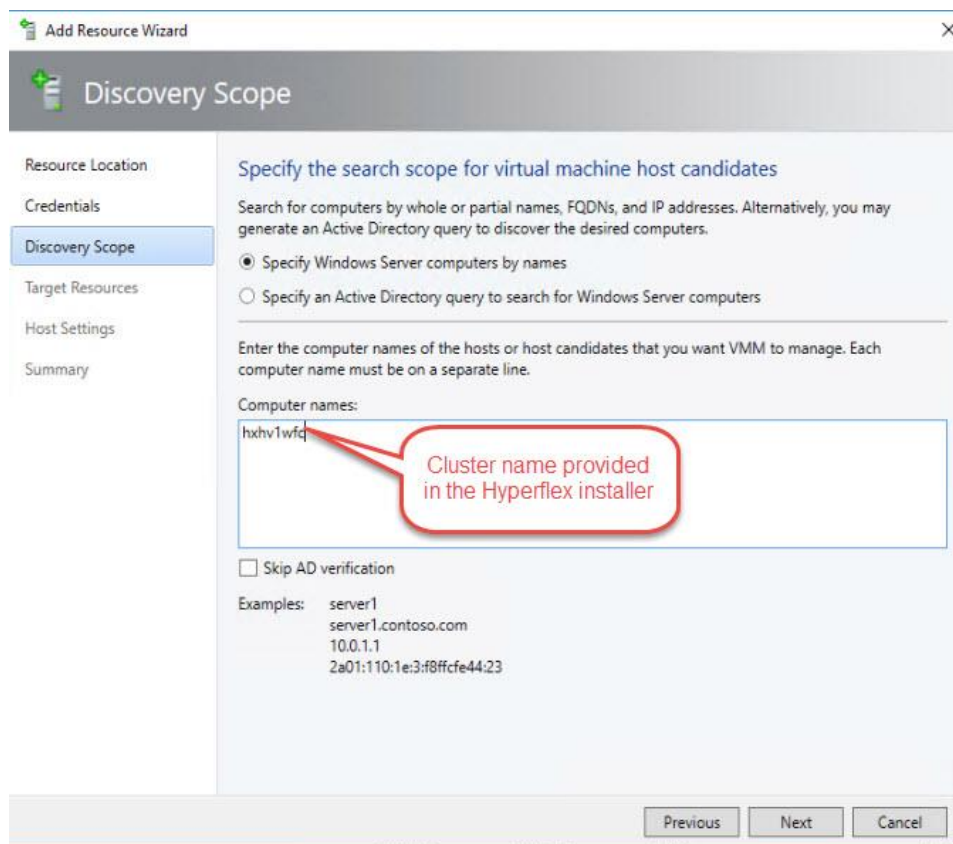
To create a Run As account, complete the following steps:

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.

### Servers

**When running the Hyperflex installer, we specified the cluster hostname of 'HXHV1WFC'. The installer created and configured the cluster and it needs to be added to the System Center VMM server to use with XenDesktop.**

- Add Cluster created by HX Installer to SCVMM



## Networking

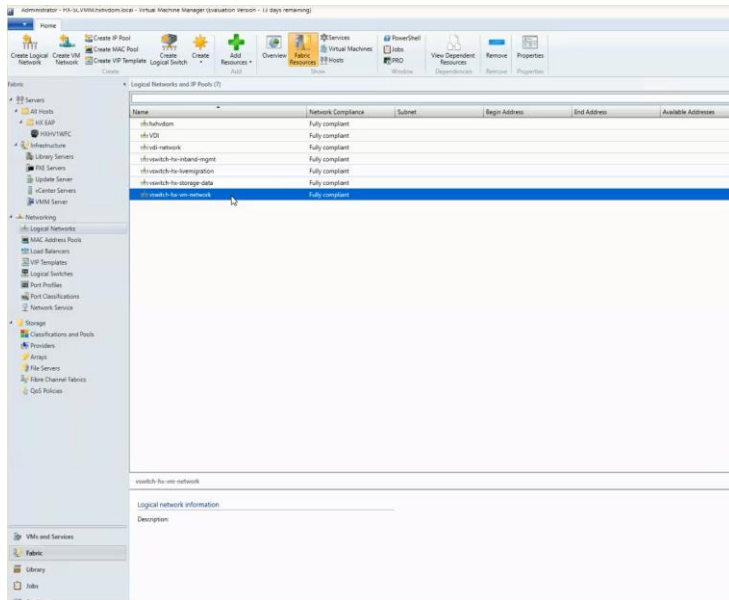
Network switches and interfaces are created by the HyperFlex installer. A network team will be created for the Management, Storage, VM Network and Live Migration networks specified during the installer.

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


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



In this example, it is 'vswitch-hx-vm-network'.



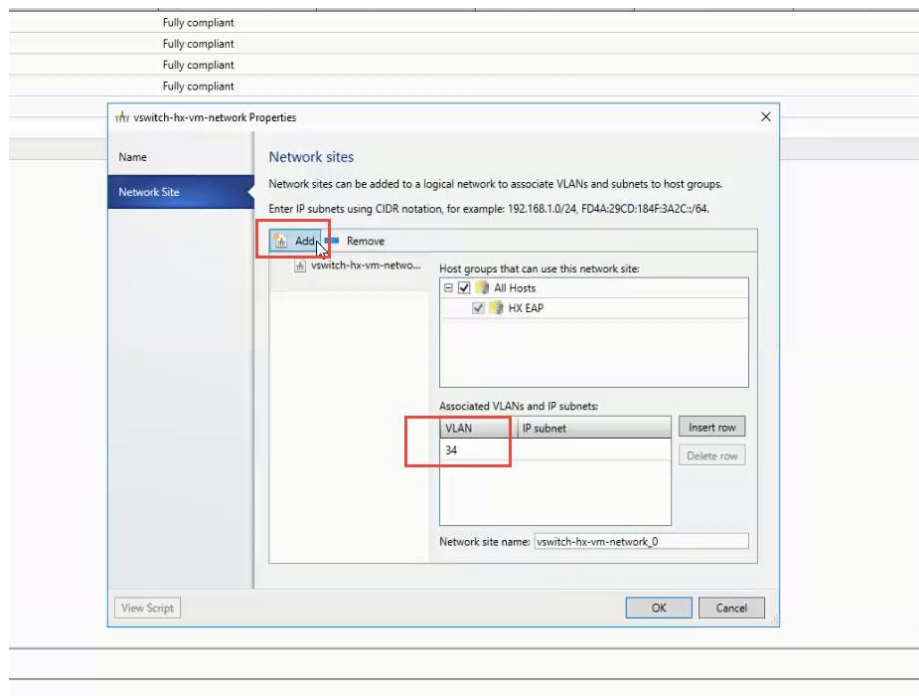
2. Right-click and select 'Properties' of the logical network.



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 VM network.



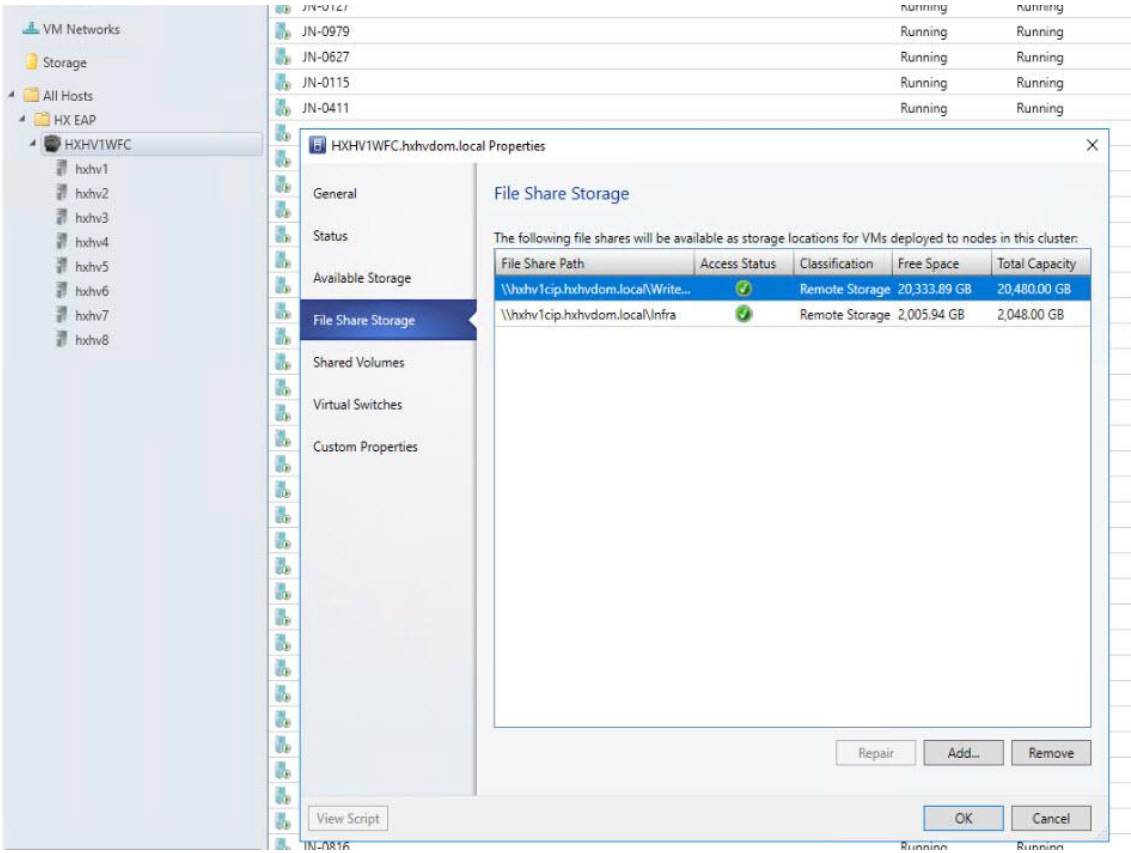
- When the network site is created, make sure each host in the cluster has the proper VLAN checked in its properties. This can be found under the properties of each host, under Hardware -> and scroll to the **'team-hx-vm-network'**

## Storage

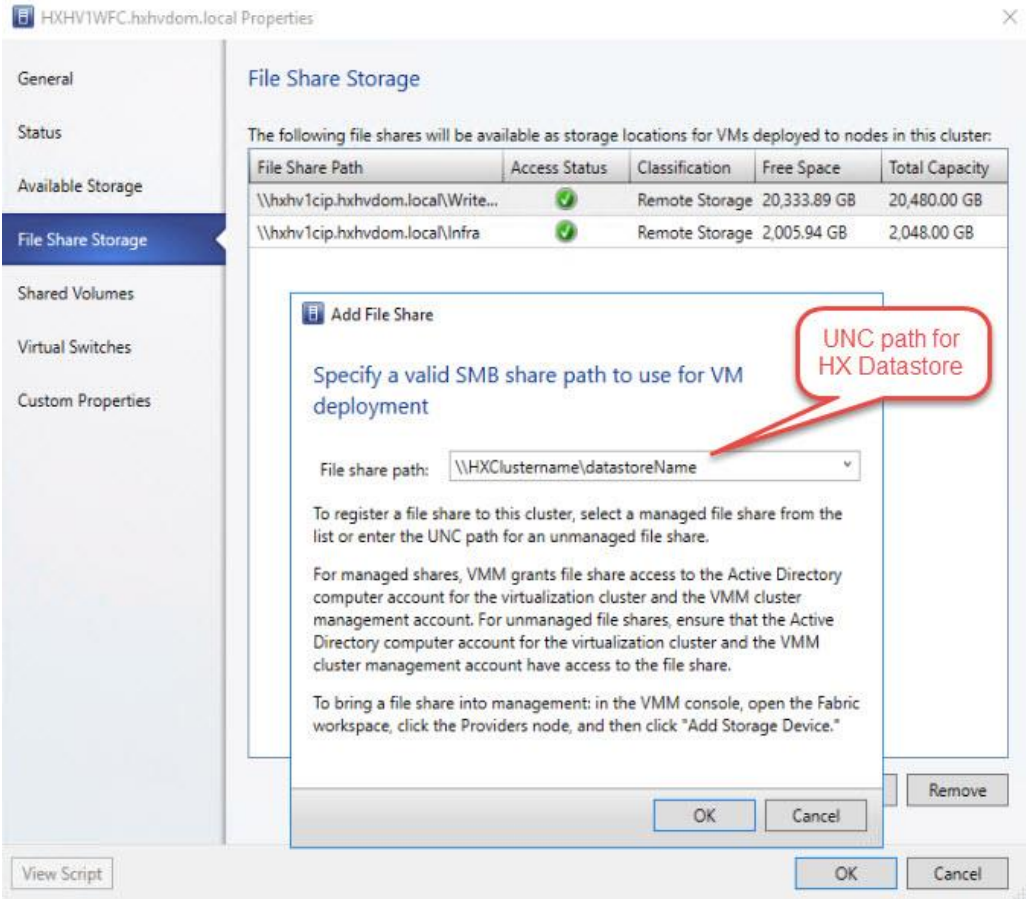
The data stores created in the HyperFlex Connect interface, creates an SMB share to use to place Virtual Machines. The naming convention is '\\hxClustername\DatastoreName'.

To add the HX Datastores to the HX cluster, complete the following steps:

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



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



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 13

Table 13 Test Infrastructure Virtual Machine Configuration

Configuration	Citrix XenDesktop 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 InBand-Mgmt	VMNIC InBand-Mgmt
Disk-1 (OS) size and location	40 GB Infra-DS volume	40 GB Infra-DS volume
Disk-2 size and location	–	
Configuration	Microsoft Active Directory DCs Virtual Machines	
Operating system	Microsoft Windows Server 2012 R2	
Virtual CPU amount	4	
Memory amount	4 GB	
Network	VMNIC InBand-Mgmt	
Disk size and location	40 GB	
Configuration	Microsoft SQL Server Virtual Machine	Citrix StoreFront Virtual Machines
Operating system	Microsoft Windows Server 2016 Microsoft SQL Server 2016	Microsoft Windows Server 2016
Virtual CPU amount	4	4
Memory amount	16 GB	8 GB

Network	VMNIC InBand-Mgmt	VMNIC InBand-Mgmt
Disk-1 (OS) size and location	40 GB Infra-DS volume	40 GB Infra-DS volume
Disk-2 size and location	200 GB Infra-DS volume SQL Logs	–
<b>Configuration</b>	<b>Citrix License Server Virtual Machines</b>	<b>NetScaler VPX Appliance Virtual Machine</b>
Operating system	Microsoft Windows Server 2016	NS11.1 52.13.nc
Virtual CPU amount	4	2
Memory amount	4 GB	2 GB
Network	VMNIC InBand-Mgmt	VMNIC InBand-Mgmt Infra-Mgmt
Disk size and location	40 GB	20 GB

## Prepare the Master Images

This section detail how to create the golden (or master) images for the environment. VMs 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 VMs 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 VMs were configured as listed in Table 14 :

Table 14 HVD and HSD Configurations

<b>Configuration</b>	<b>HVDI Virtual Machines</b>	<b>HSD (Not used in this study) Virtual Machines</b>
Operating system	Microsoft Windows 10 64-bit	Microsoft Windows Server 2016
Virtual CPU amount	2	6
Memory amount	2.0 GB (reserved)	24 GB (reserved)

Configuration	HVDI Virtual Machines	HSD (Not used in this study) Virtual Machines
Network	VMNIC vm-network	VMNIC vm-network
Citrix PVS vDisk size and location	24 GB (thick) VDI WriteCache Volume	100 GB (thick) Infra-DS volume
Citrix PVS write cache  Disk size	6 GB	24 GB
Additional software used for testing	Microsoft Office 2016  Login VSI 4.1.32 (Knowledge Worker Workload)	Microsoft Office 2016  Login VSI 4.1.32 (Knowledge Worker Workload)

## Install and Configure XenDesktop Delivery Controller, Citrix Licensing, and StoreFront


This section details the installation of the core components of the XenDesktop/XenApp 7.16 system. This CVD provides the process to install two XenDesktop Delivery Controllers to support hosted shared desktops (HSD), non-persistent virtual desktops (VDI), and persistent virtual desktops (VDI).

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

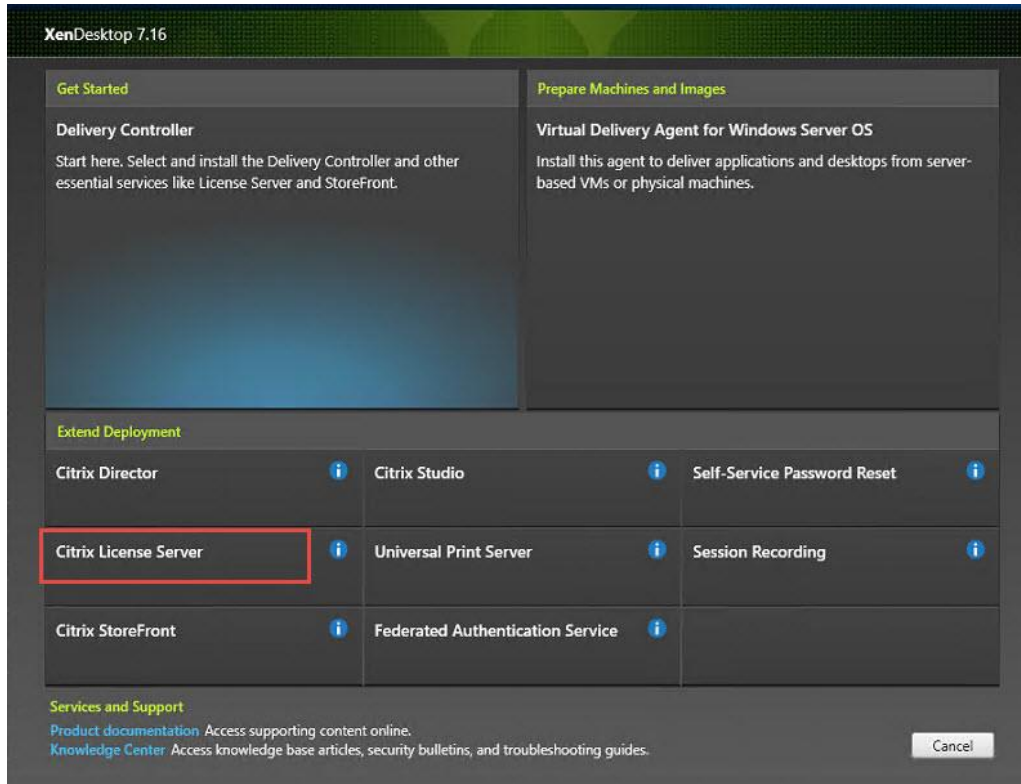
### Install Citrix License Server

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

1. To begin the installation, connect to the first Citrix License server and launch the installer from the Citrix XenDesktop 7.16 ISO.
2. Click Start.

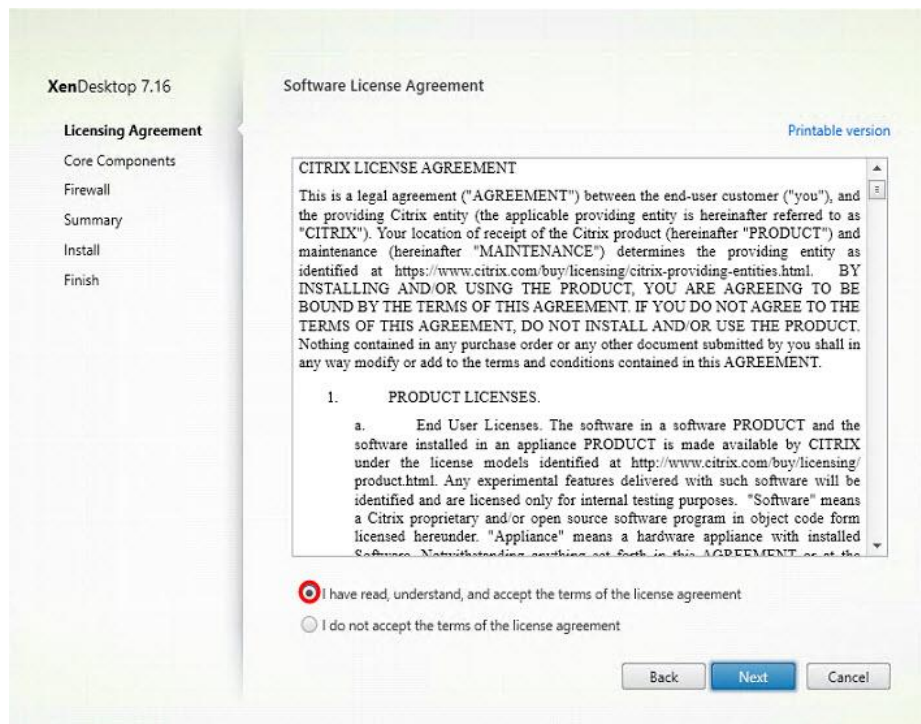


3. Click "Extend Deployment – Citrix License Server."

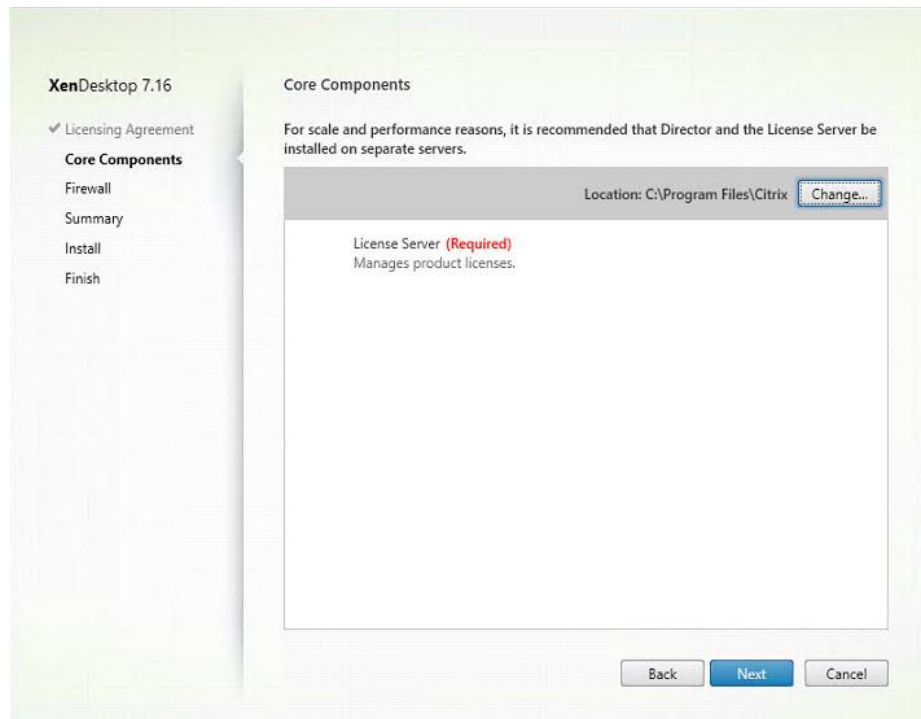


4. Read the Citrix License Agreement.

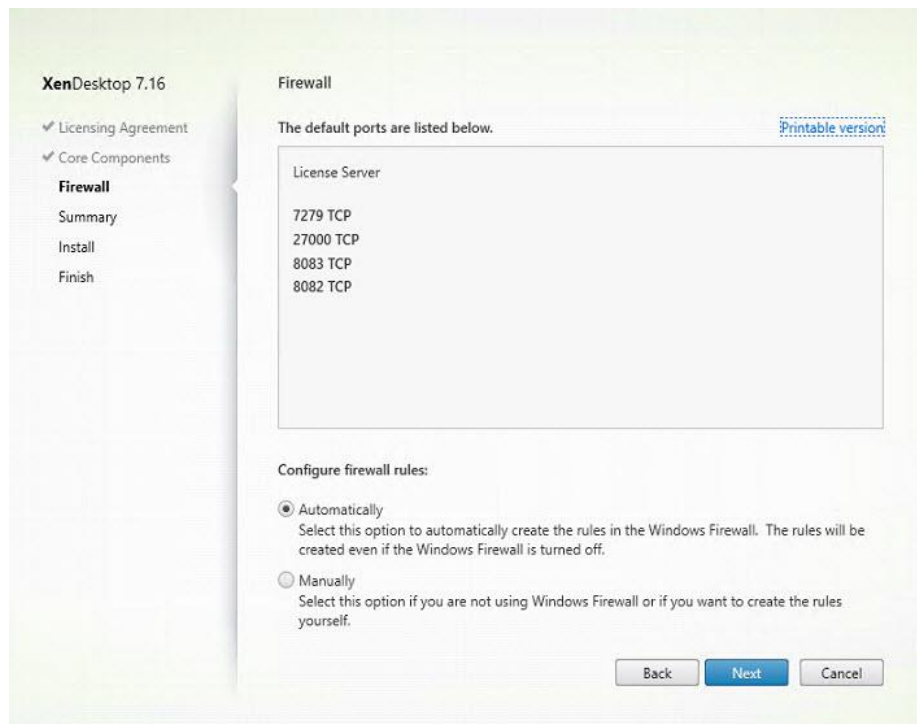
5. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.
6. Click Next.



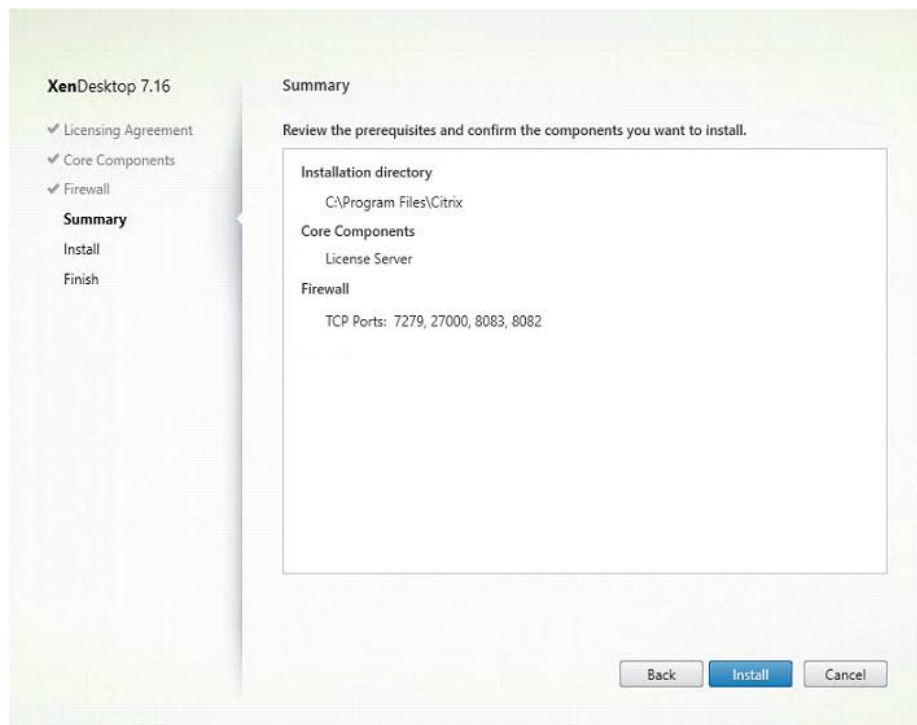
7. Click Next.



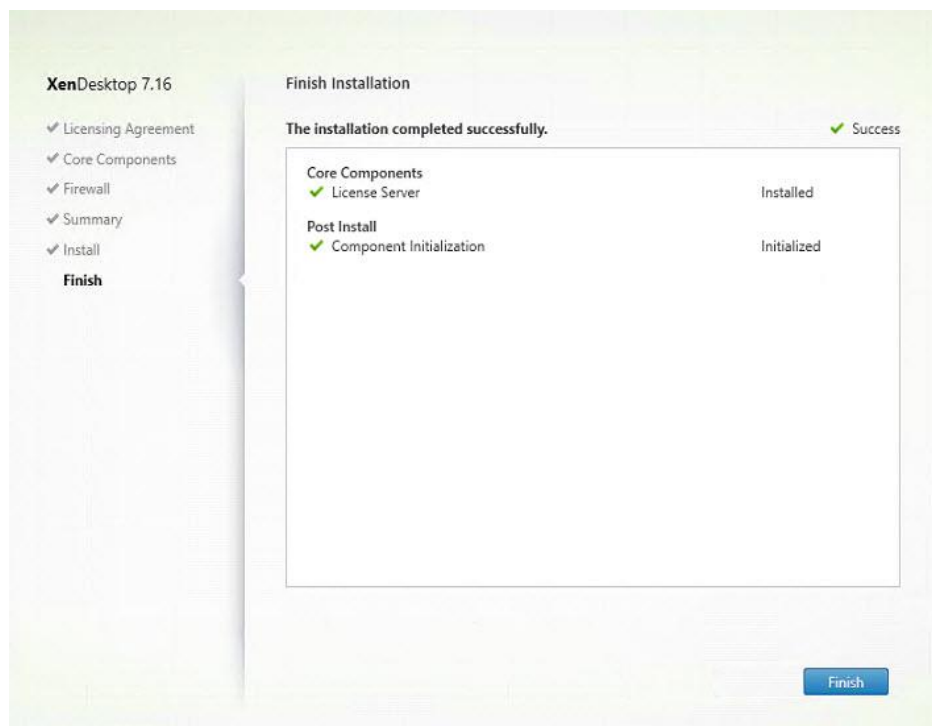
8. Select the default ports and automatically configured firewall rules.
9. Click Next



10. Click Install.



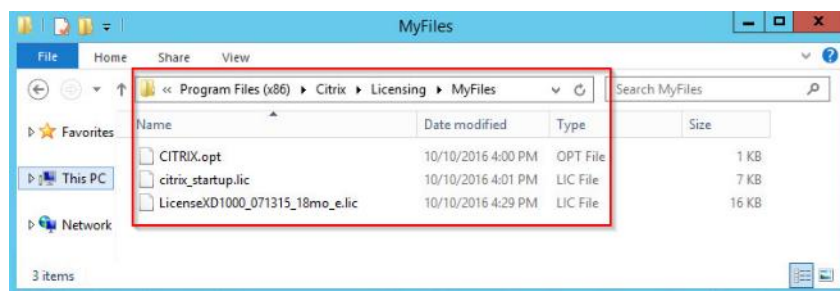
11. Click Finish to complete the installation.



## Install Citrix Licenses

To install the Citrix Licenses, complete the following steps:

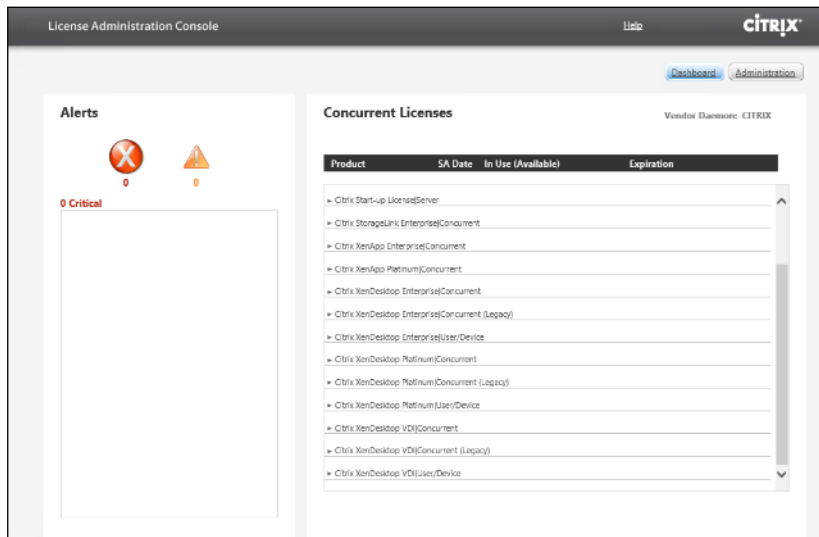
1. Copy the license files to the default location (C:\Program Files (x86)\Citrix\Licensing\ MyFiles) on the li-  
cense server.



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



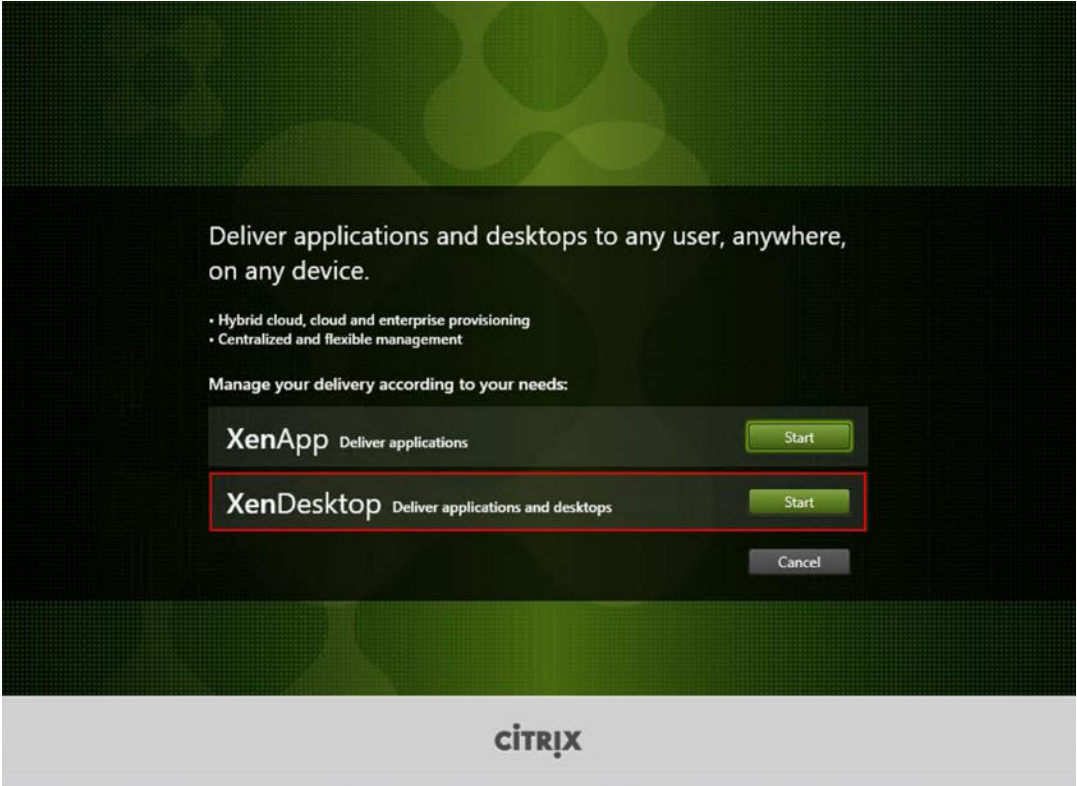
4. Confirm that the license files have been read and enabled correctly.



## Install the XenDesktop

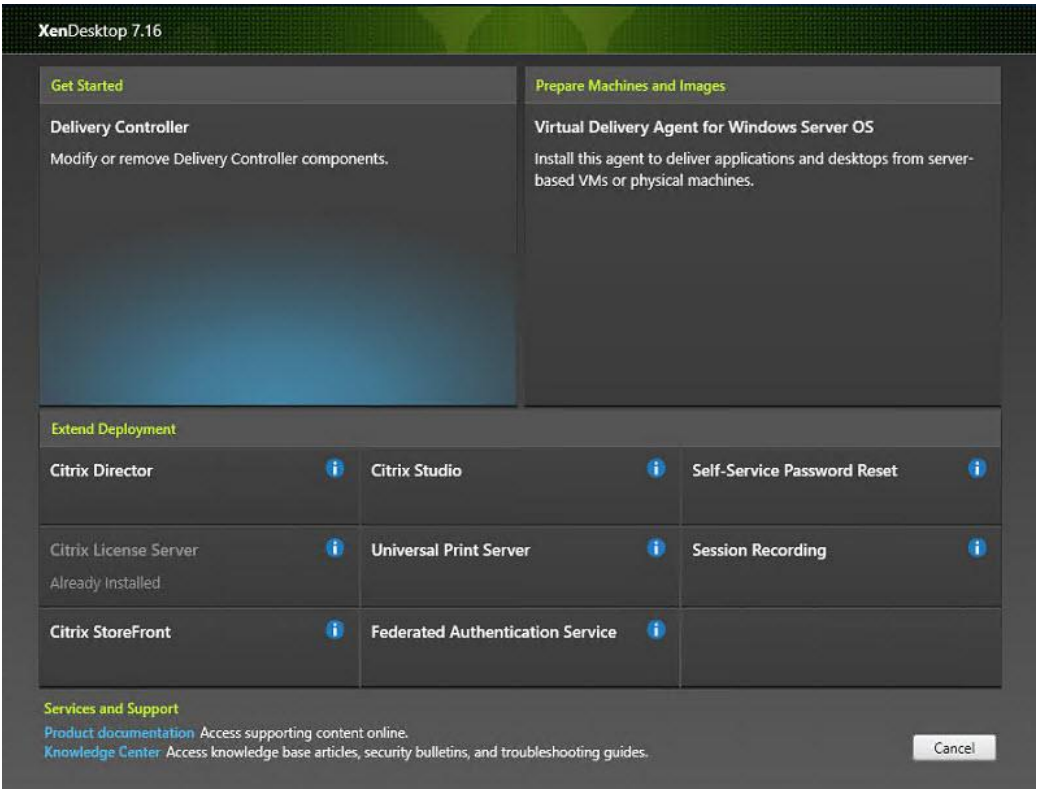
1. To begin the installation, connect to the first XenDesktop server and launch the installer from the Citrix XenDesktop 7.16 ISO.
2. Click Start.



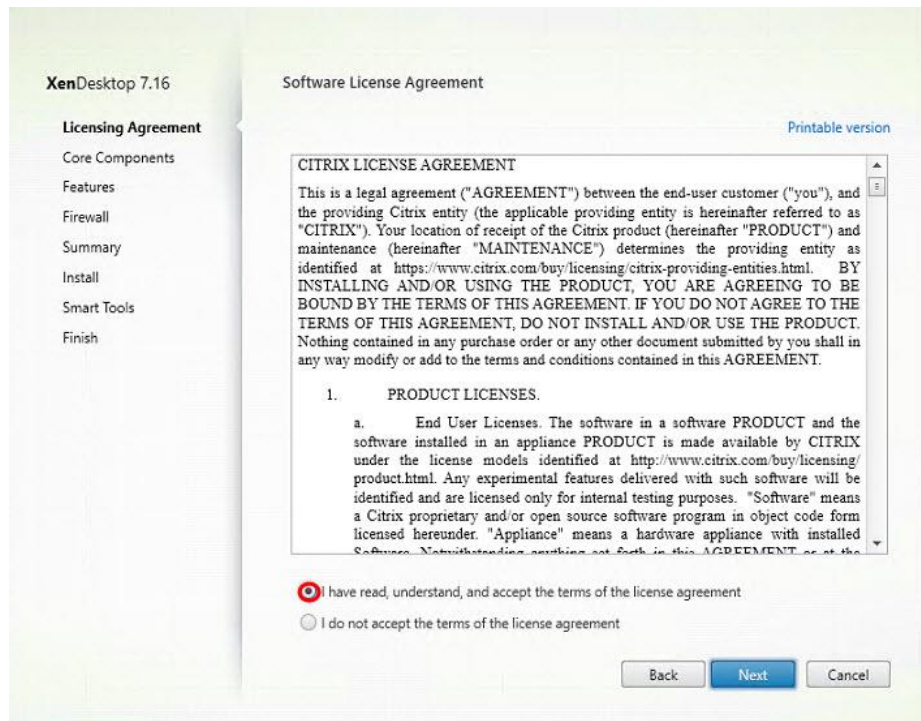


The installation wizard presents a menu with three subsections.

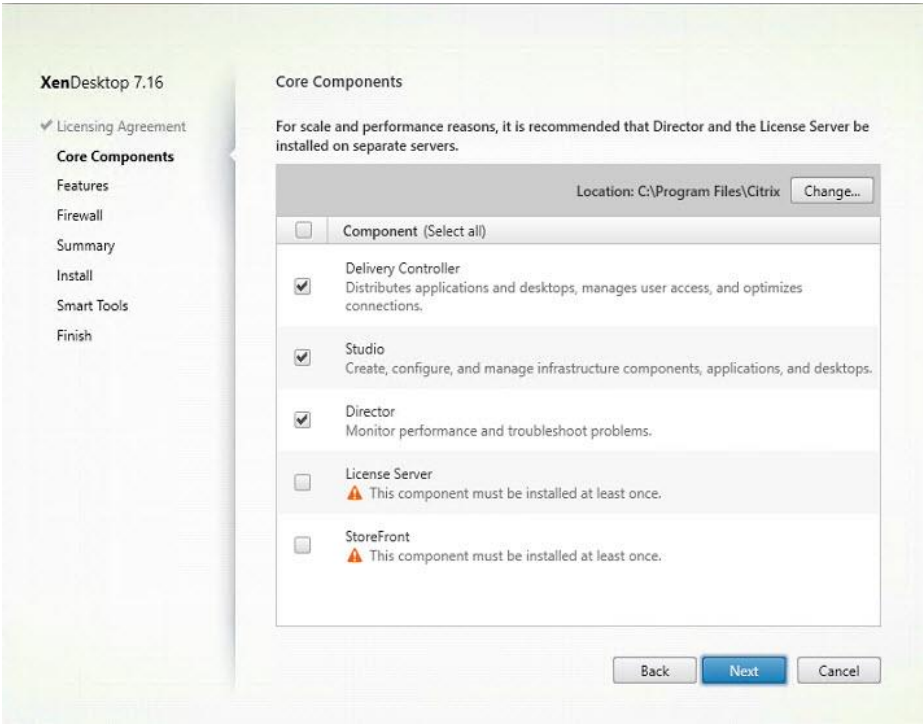
3. Click “Get Started - Delivery Controller.”



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

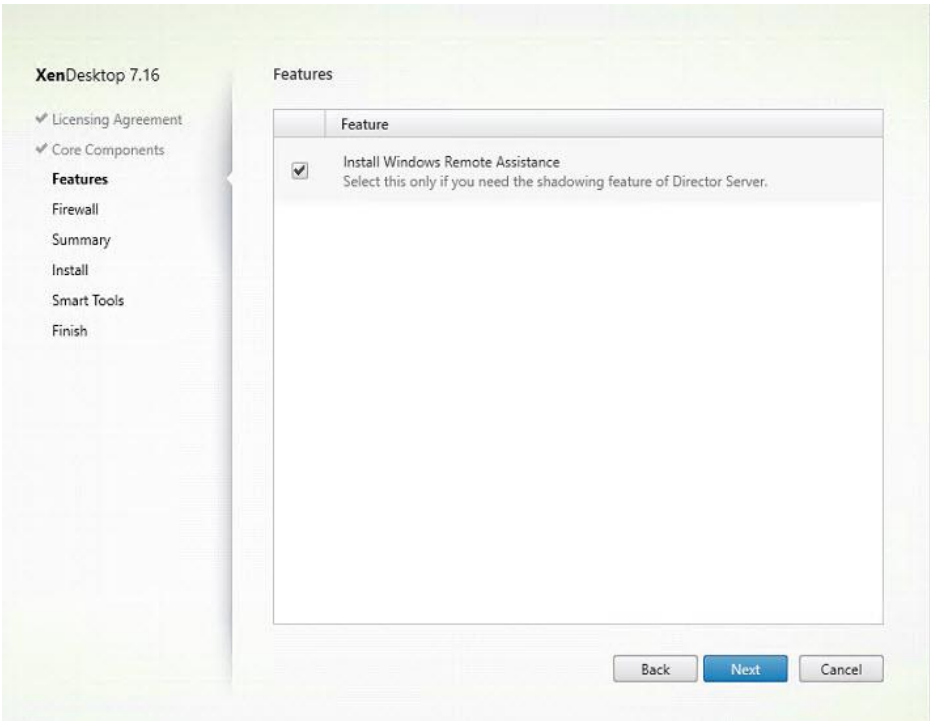


7. Select the components to be installed on the first Delivery Controller Server:
  - a. Delivery Controller
  - b. Studio
  - c. Director
8. Click Next.



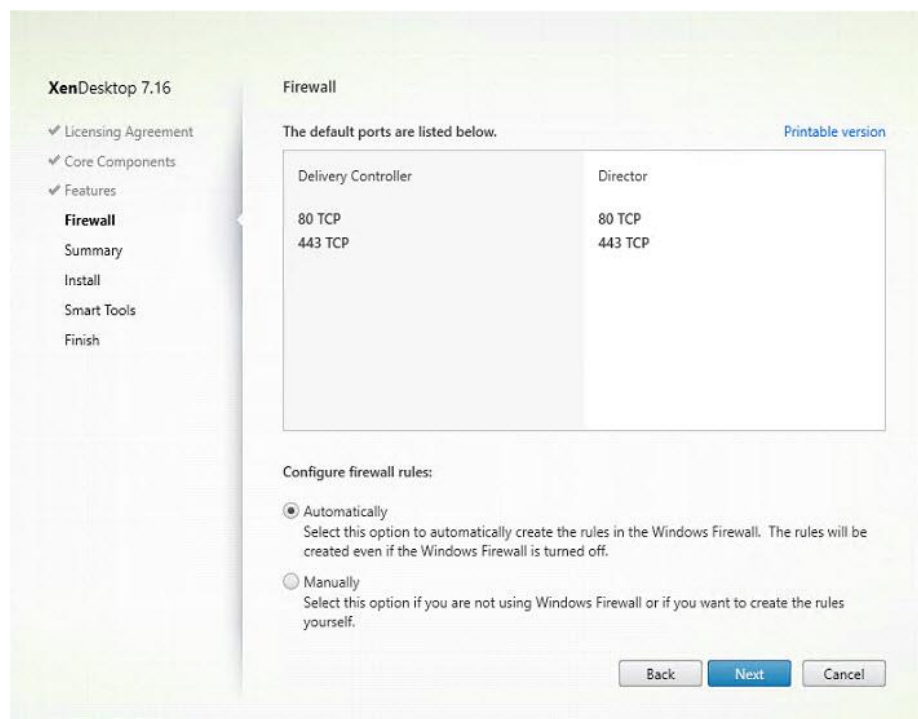
Dedicated StoreFront and License servers should be implemented for large scale deployments.

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

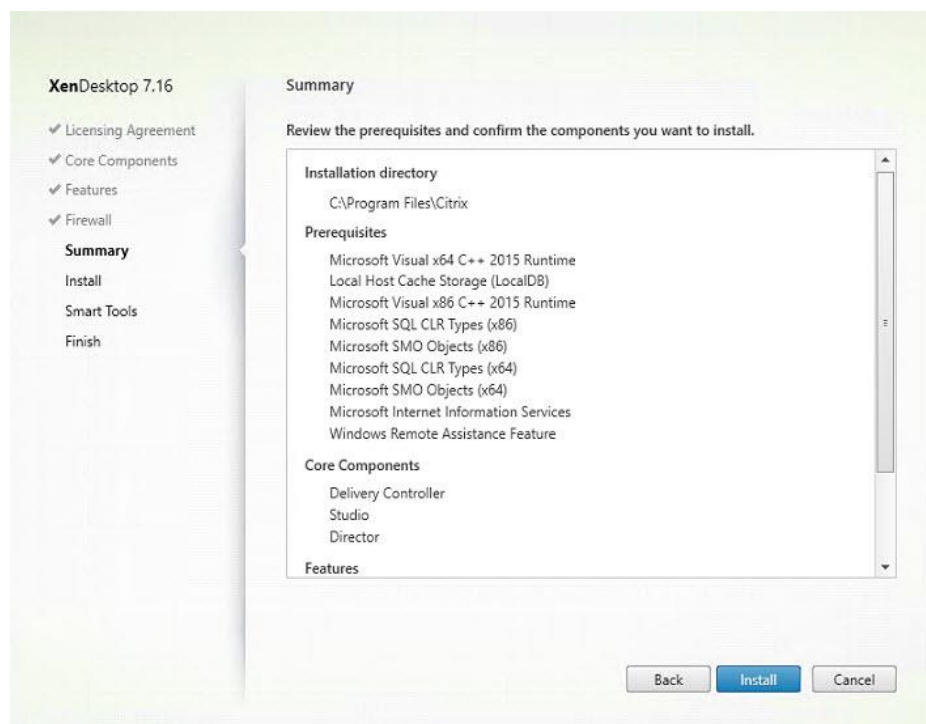


11. Select the default ports and automatically configured firewall rules.

12. Click Next.

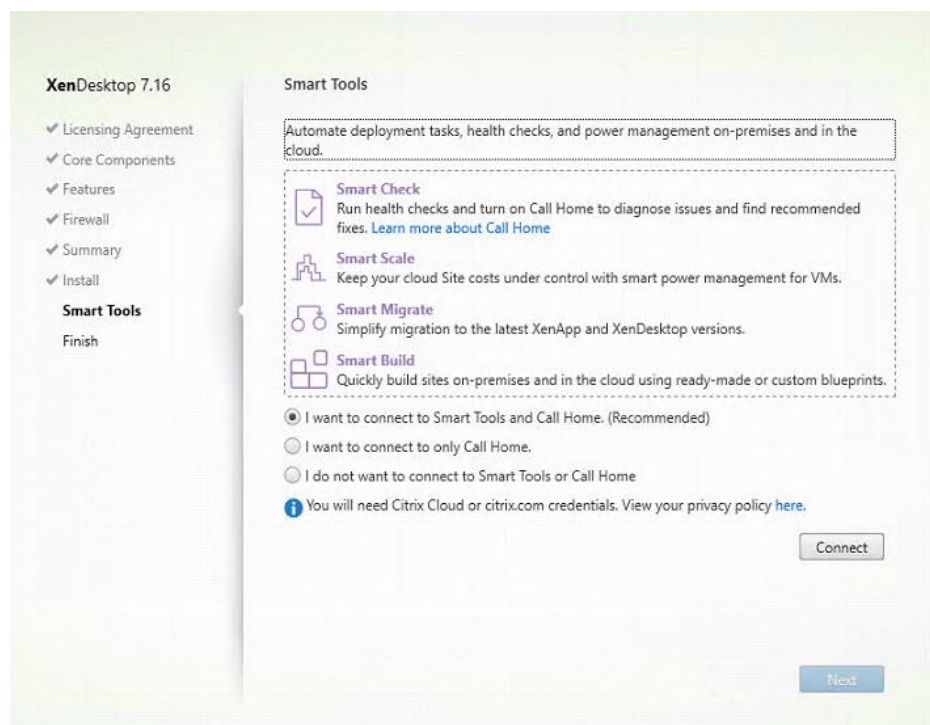


13. Click Install to begin the installation.



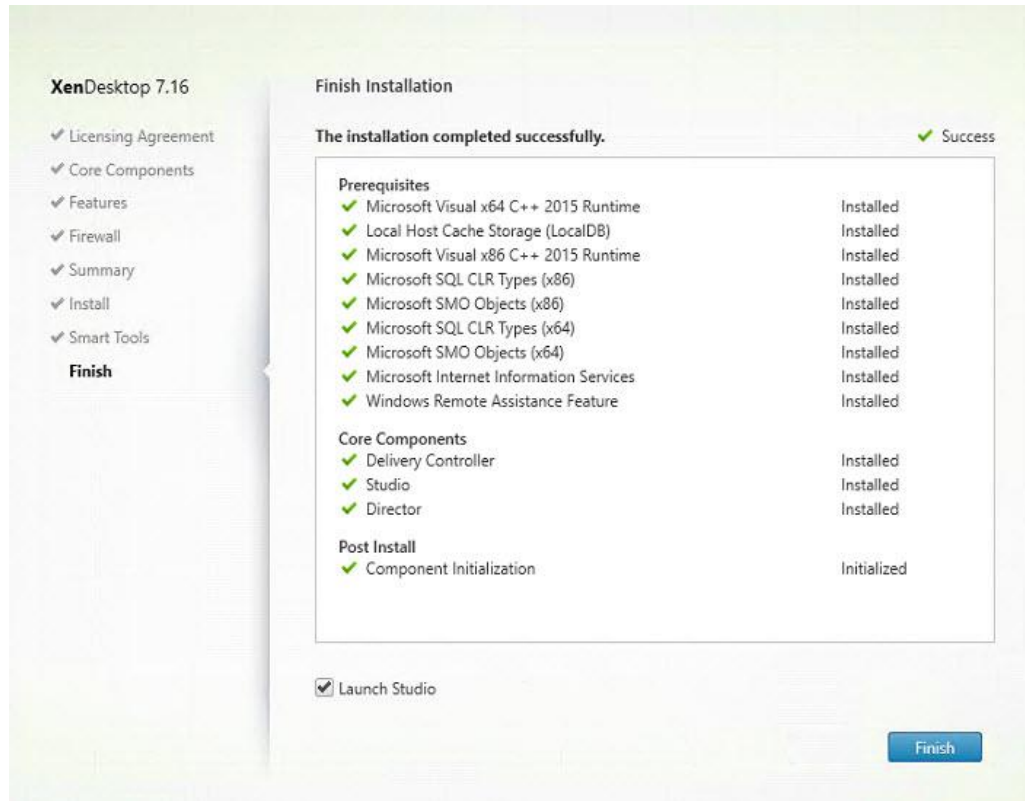
14. (Optional) Click the Call Home participation.

15. Click Next.



16. Click Finish to complete the installation.

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



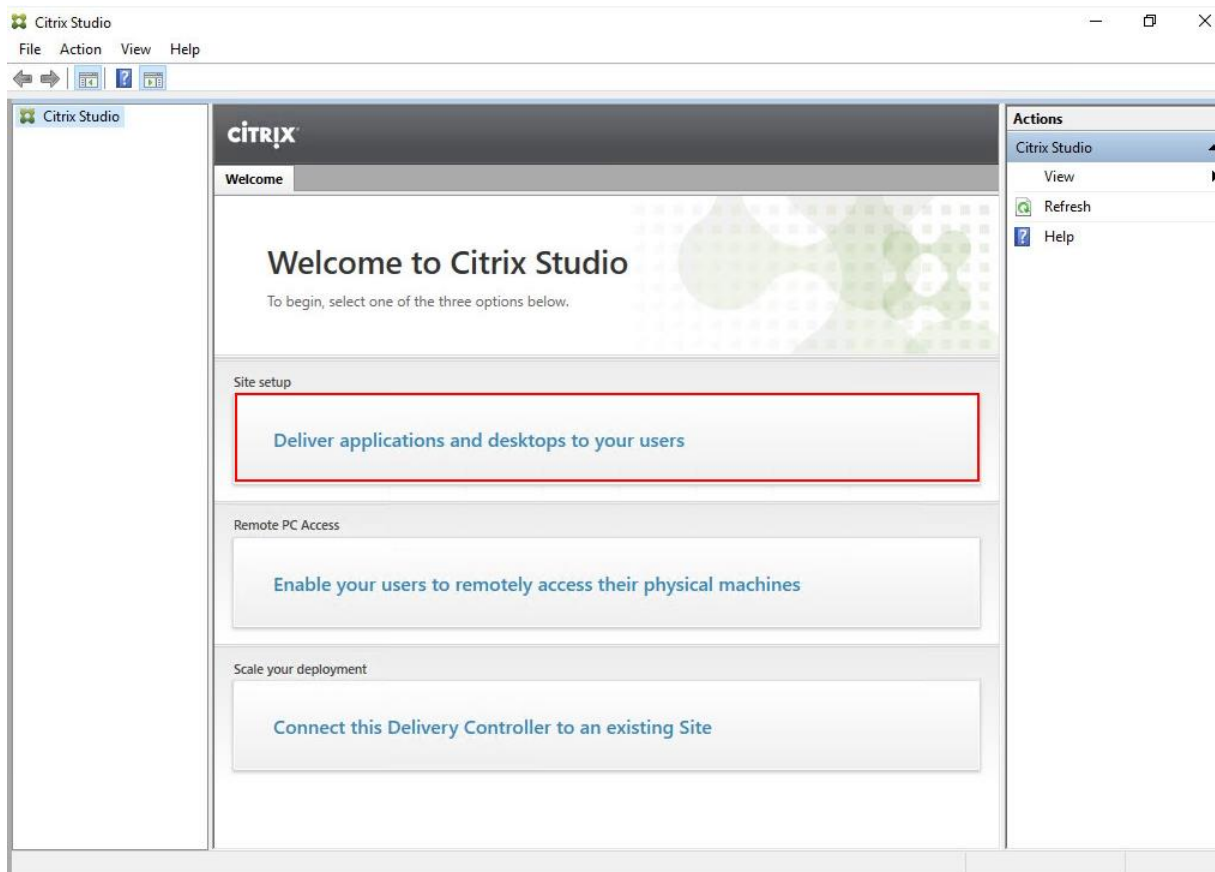
## Configure the XenDesktop Site

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 XenDesktop Delivery Controller installation, or if necessary, it can be launched manually. Citrix Studio is used to create a Site, which is the core XenDesktop 7.16 environment consisting of the Delivery Controller and the Database.

To configure XenDesktop, complete the following steps:

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



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



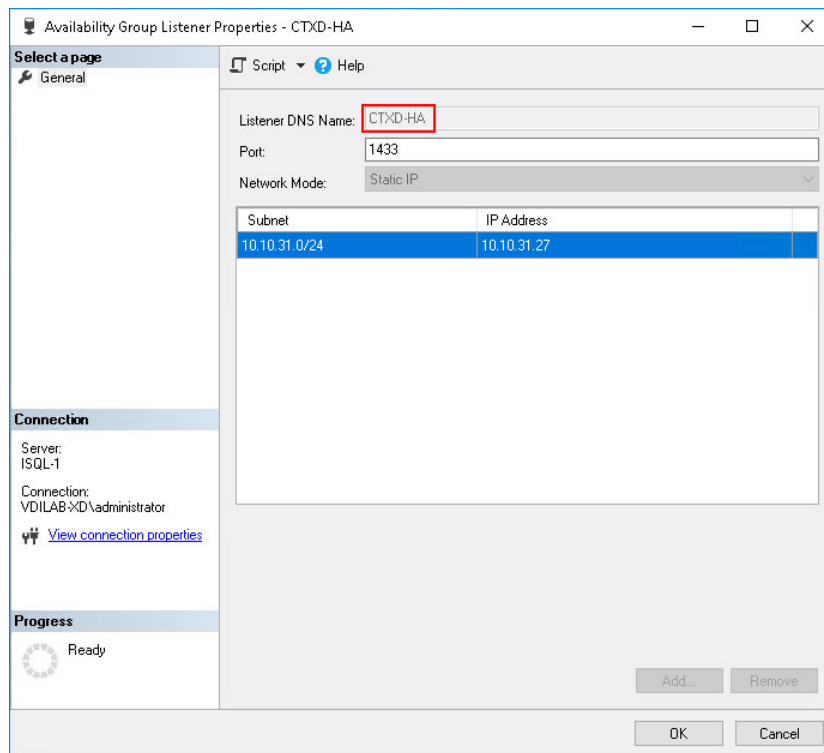
The screenshot shows the 'Site Setup' wizard in the 'Introduction' step. On the left, a 'Studio' sidebar lists navigation options: Introduction (selected), Databases, Licensing, Connection, Network, Additional Features, and Summary. The main content area is titled 'Introduction' and contains the following text: 'You have two options when creating a new Site. The simplest option is to automatically create a fully configured, production-ready Site. The second, more advanced option is to create an empty Site, which you must configure yourself.' Below this, a question asks 'What kind of Site do you want to create?' with two radio button options: 'A fully configured, production-ready Site (recommended for new users)' (which is selected) and 'An empty, unconfigured Site'. A 'Site name:' label is followed by a text input field containing 'CTXD'. At the bottom right, there are three buttons: 'Back', 'Next' (highlighted with a red box), and 'Cancel'.

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

The screenshot shows the 'Site Setup' wizard in the 'Databases' step. The 'Studio' sidebar on the left now has 'Databases' selected. The main content area is titled 'Databases' and includes the text: 'Databases store information about Site setup, configuration logging and monitoring. Choose how you want to set up the databases. [Learn more](#)'. There are two radio button options: 'Create and set up databases from Studio (You can provide details of existing empty databases)' (selected) and 'Generate scripts to manually set up databases on the database server'. Below this is a section titled 'Provide database details' containing a table with three columns: 'Data type', 'Database name', and 'Location (formats)'. The table has three rows: 'Site' with 'CTXD713\_site' and 'CTXD-HA'; 'Monitoring' with 'CTXD713\_monitoring' and 'CTXD-HA'; and 'Logging' with 'CTXD713\_logging' and 'CTXD-HA'. The 'CTXD-HA' entries in the 'Location' column are highlighted with a red box. Below the table, an information icon and text state: 'For an AlwaysOn Availability Group, specify the group's listener in the location.' At the bottom, there is a label 'Specify additional Delivery Controllers for this Site [Learn more](#)' and a 'Select...' button. At the bottom right, there are three buttons: 'Back', 'Next' (highlighted with a red box), and 'Cancel'.

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



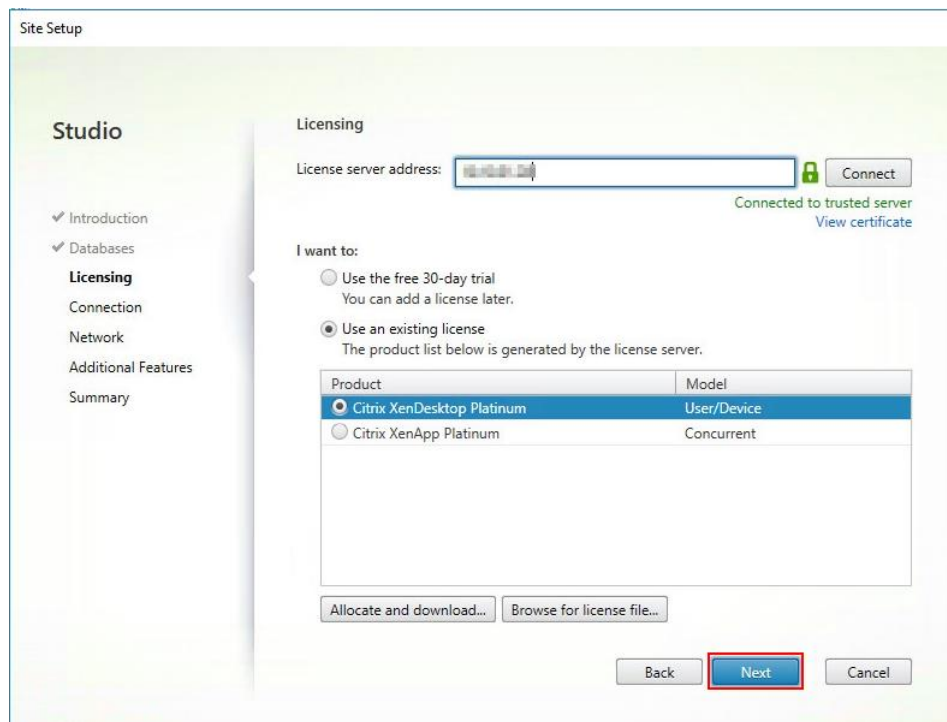


7. Provide the FQDN of the license server.
8. Click Connect to validate and retrieve any licenses from the server.



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 System Center Virtual Machine Manager.
12. Enter the FQDN of the SCVMM server (in Server\_FQDN/sdk format).
13. Enter the username (in domain\username format) for the Hyper-V account.
14. Provide the password for the Domain Admin account.
15. Provide a connection name.
16. Select the Other tools radio button.

The screenshot shows a 'Site Setup' window with a left-hand navigation pane and a main configuration area. The navigation pane includes links for Introduction, Databases, Licensing, Connection (which is highlighted), Storage Management, Storage Selection, Network, Additional Features, and Summary. The main area is titled 'Connection' and contains the following fields and options:

- Connection type:** A dropdown menu showing 'Microsoft® System Center Virtual Machine Mana...'.
- Connection address:** A text box containing '10.10.31.23'.
- User name:** A text box containing 'administrator'.
- Password:** A text box filled with dots.
- Connection name:** A text box containing 'SCVMM'.
- Create virtual machines using:** Two radio buttons. The first, 'Studio tools (Machine Creation Services)', is selected and has a sub-note: 'Select this option when using AppDisks, even if you are using Provisioning Services.' The second is 'Other tools'.

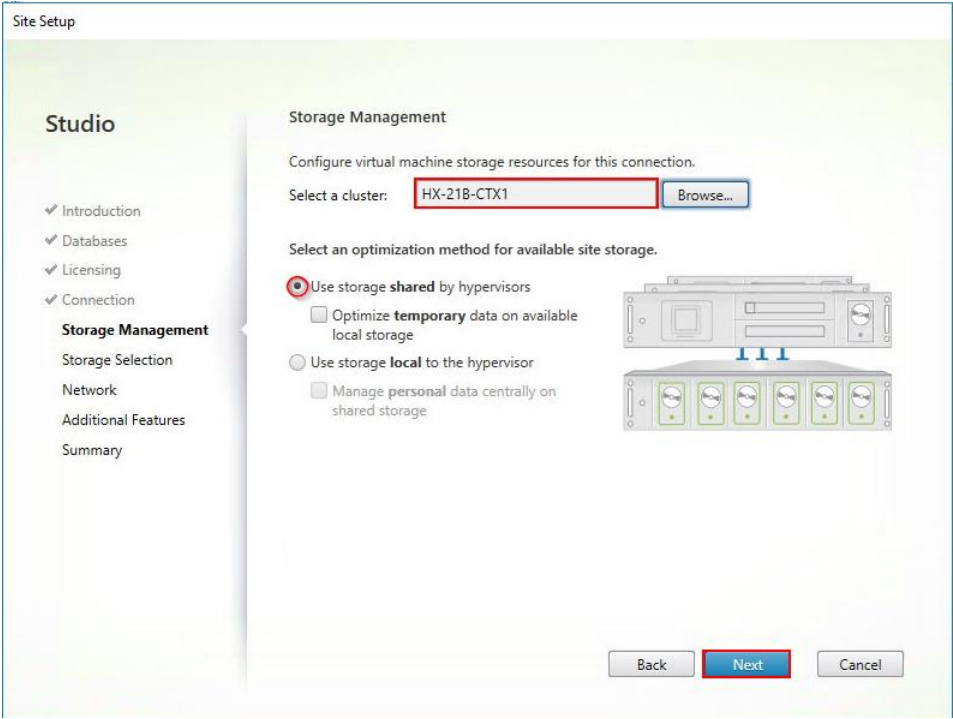
At the bottom right, there are three buttons: 'Back', 'Next' (which is highlighted with a mouse cursor), and 'Cancel'.

17. Click Next.

18. Select HyperFlex Cluster that will be used by this connection.

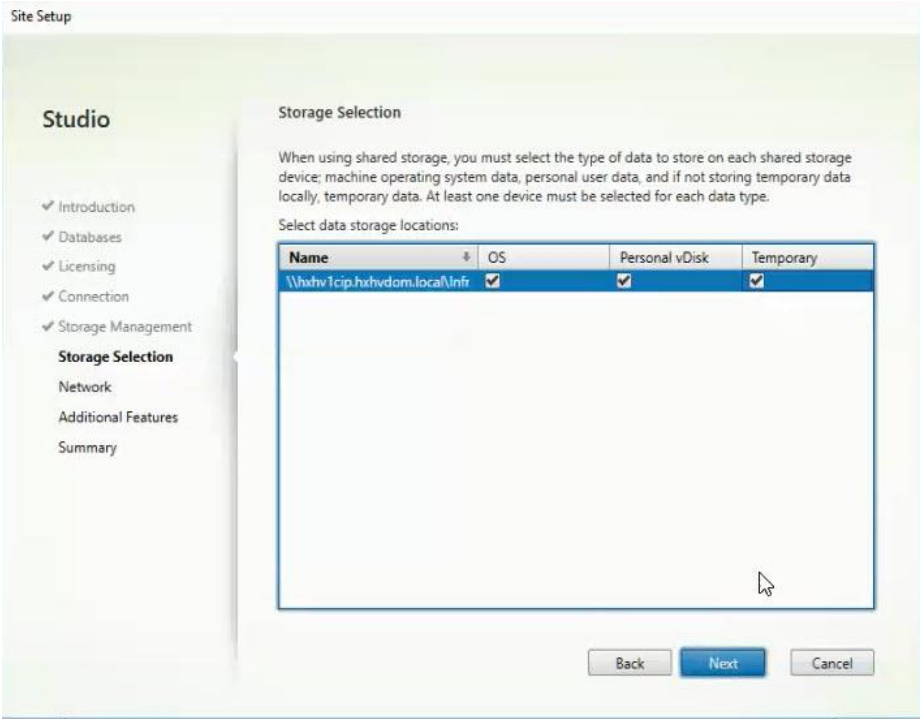
19. Check Studio Tools radio button required to support desktop provisioning task by this connection.

20. Click Next.



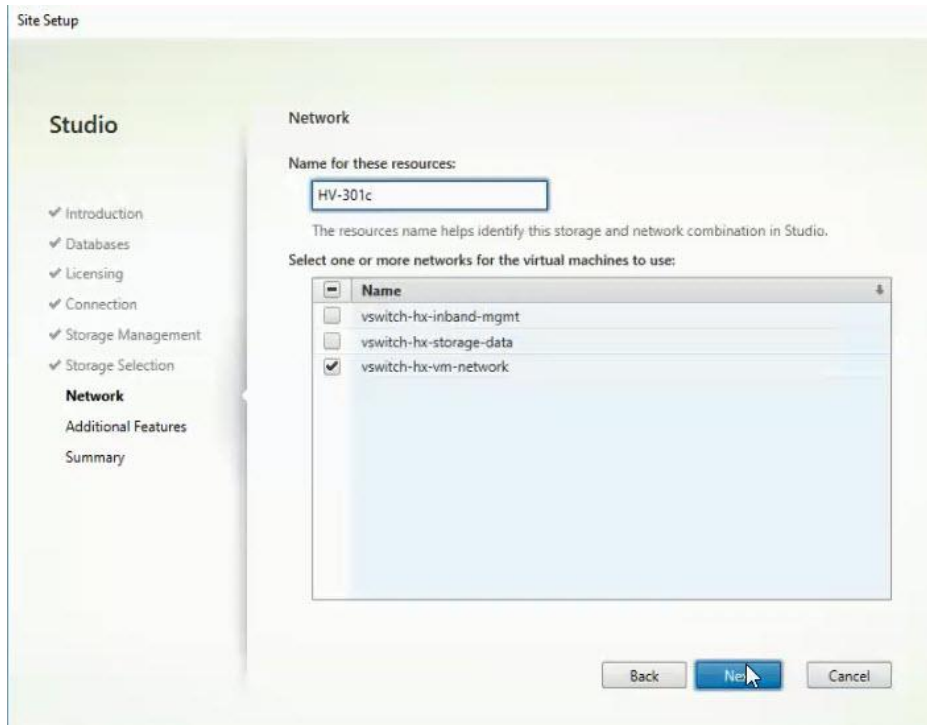
21. Make Storage selection to be used by this connection.

22. Click Next.



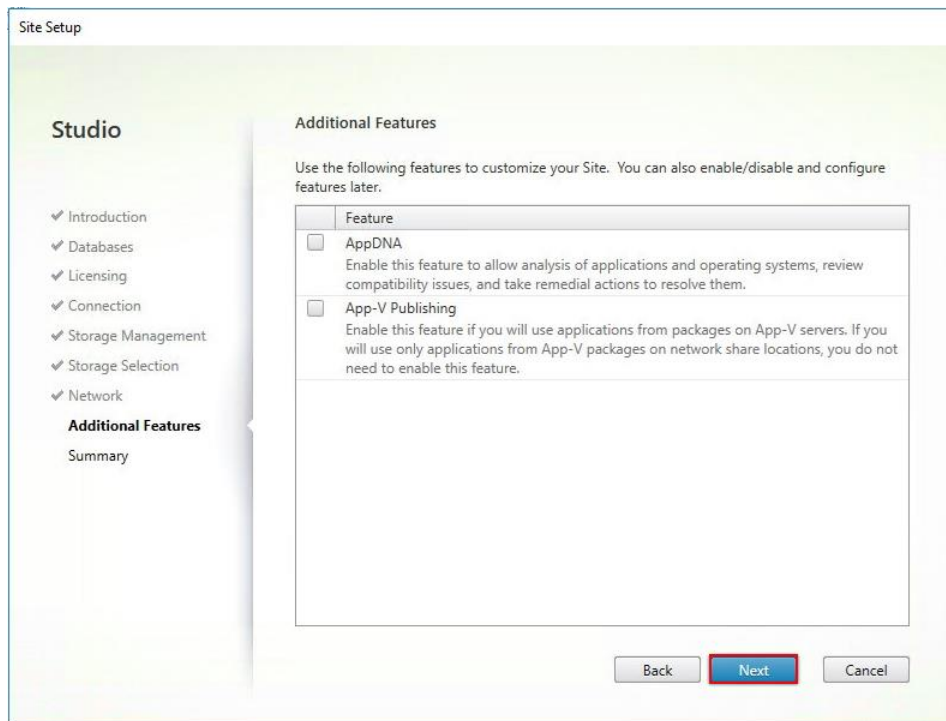
23. Make Network selection to be used by this connection.

24. Click Next.

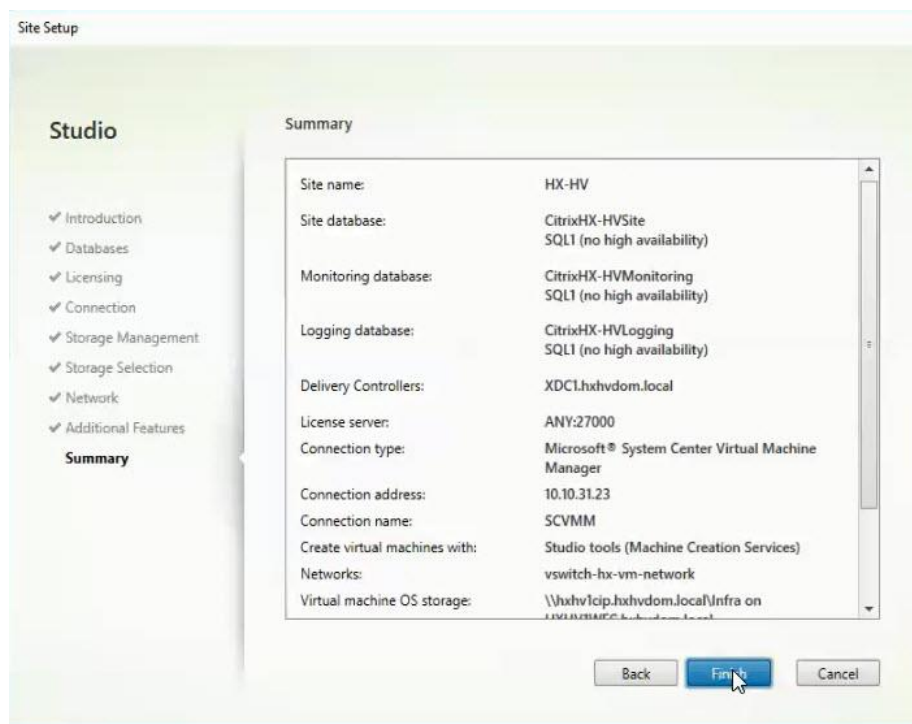


25. Select Additional features.

26. Click Next.




27. Review Site configuration Summary and click Finish.



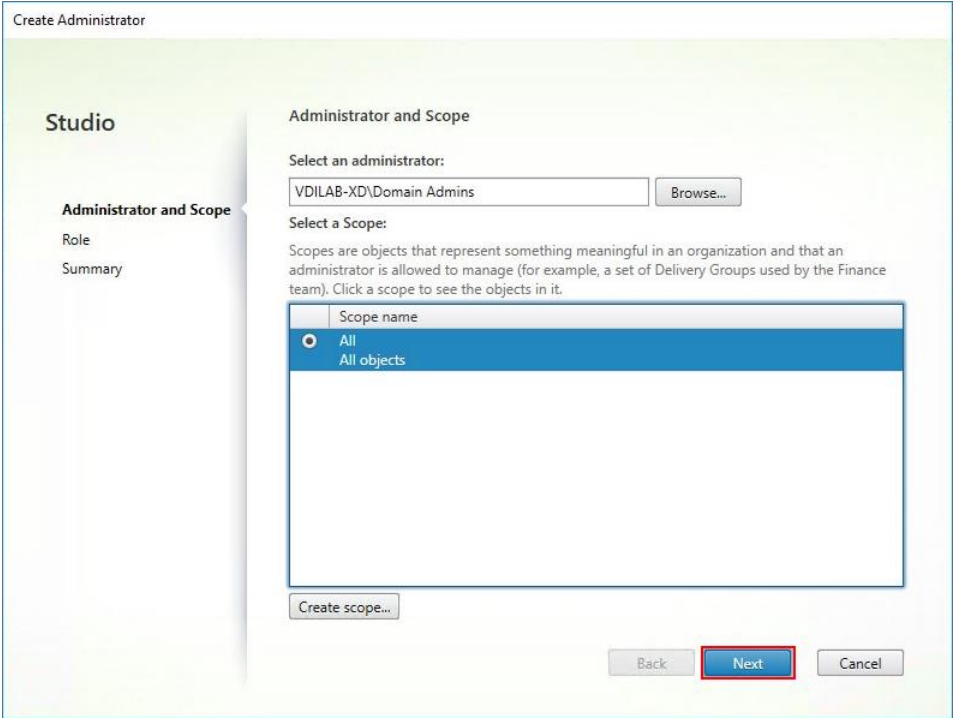
## Configure the XenDesktop Site Administrators

To configure the XenDesktop site administrators, complete the following steps:

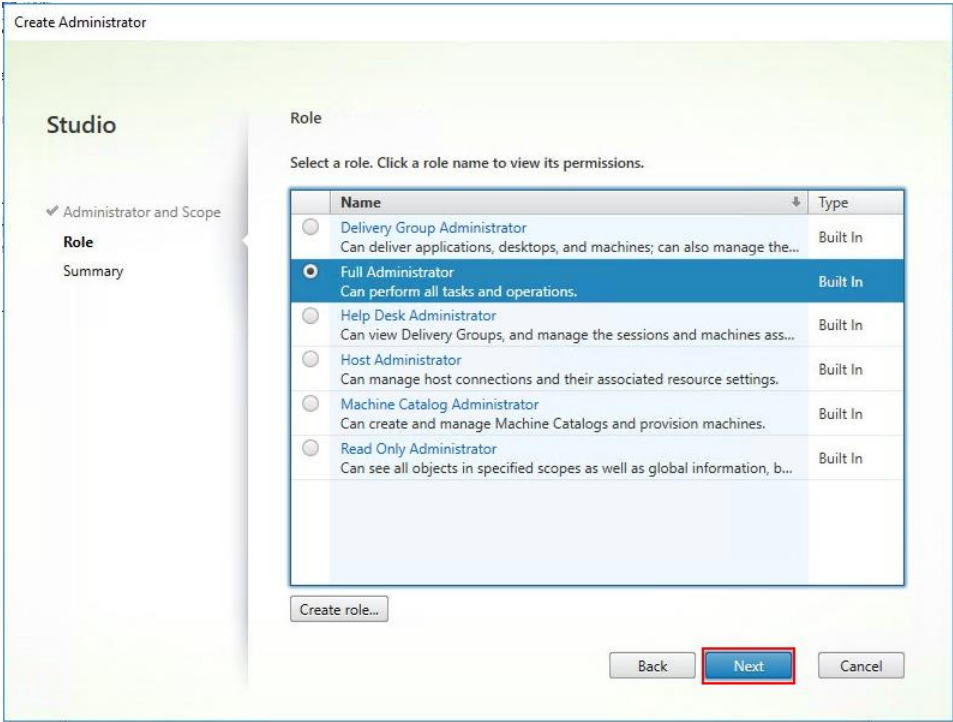
1. Connect to the XenDesktop server and open Citrix Studio Management console.
2. From the Configuration menu, right-click Administrator and select Create Administrator from the drop-down list.



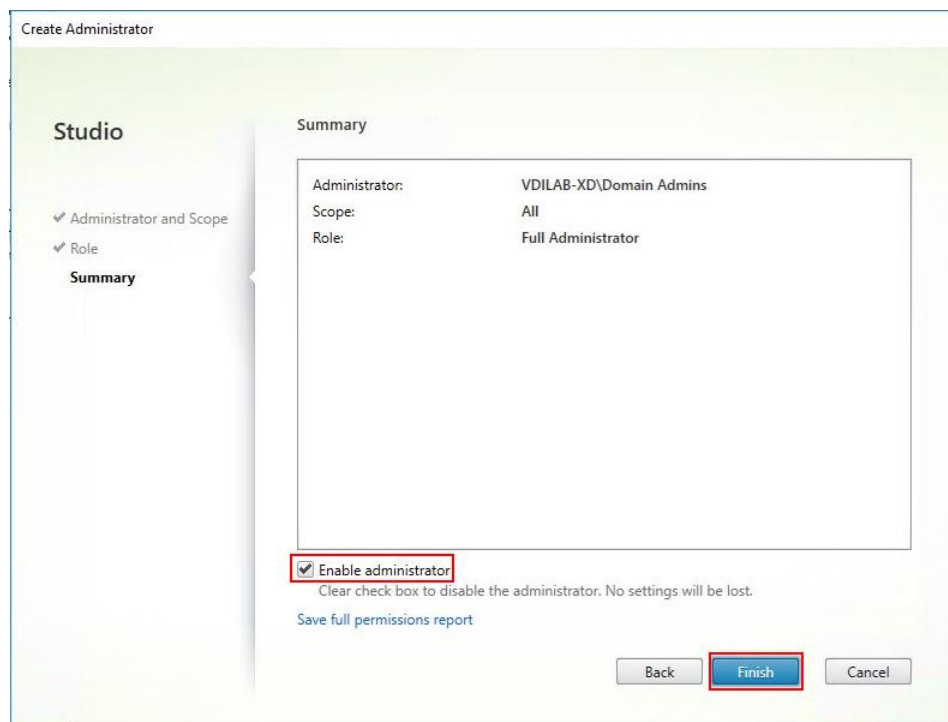
3. Select/Create appropriate scope and click Next.



4. Choose an appropriate Role.



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



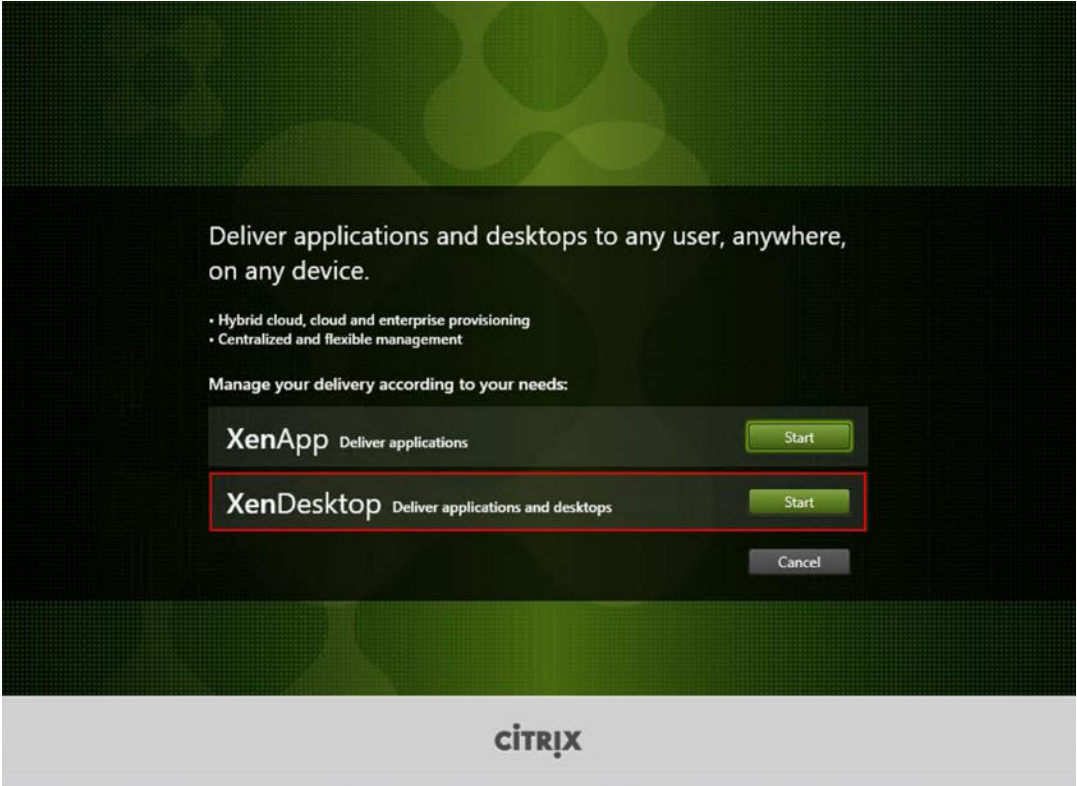
## Configure Additional XenDesktop Controller

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.

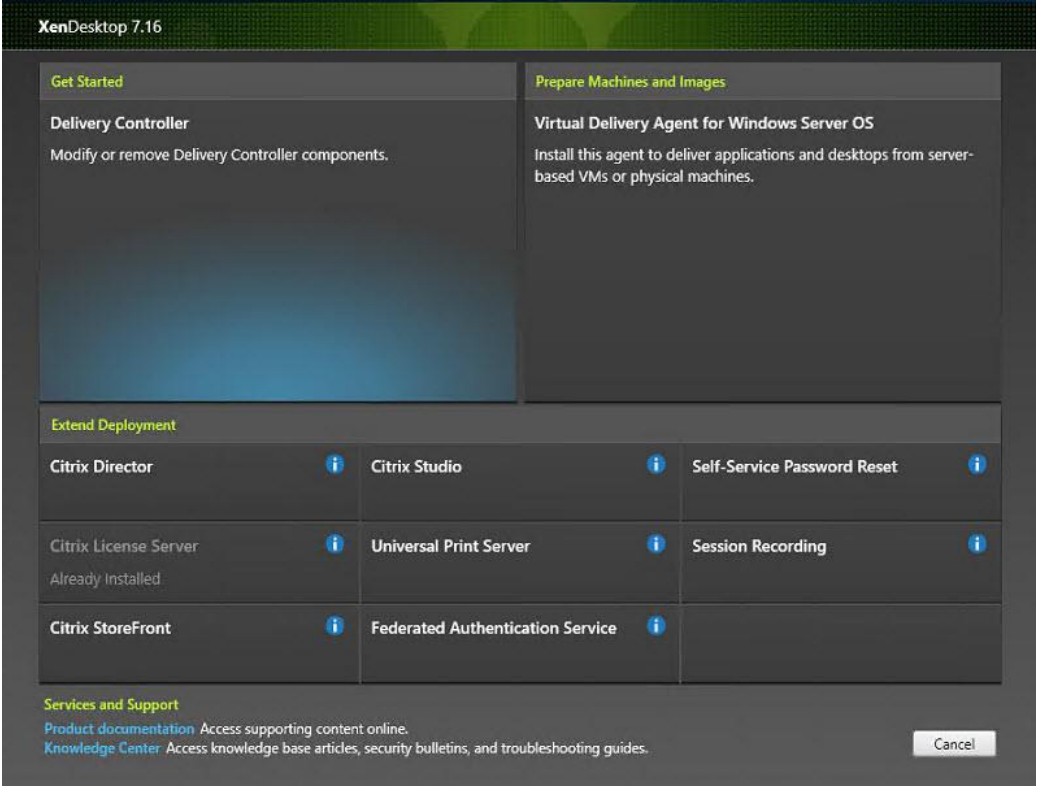
To configure additional XenDesktop controllers, complete the following steps:

1. To begin the installation of the second Delivery Controller, connect to the second XenDesktop server and launch the installer from the Citrix XenDesktop 7.16 ISO.
2. Click Start.

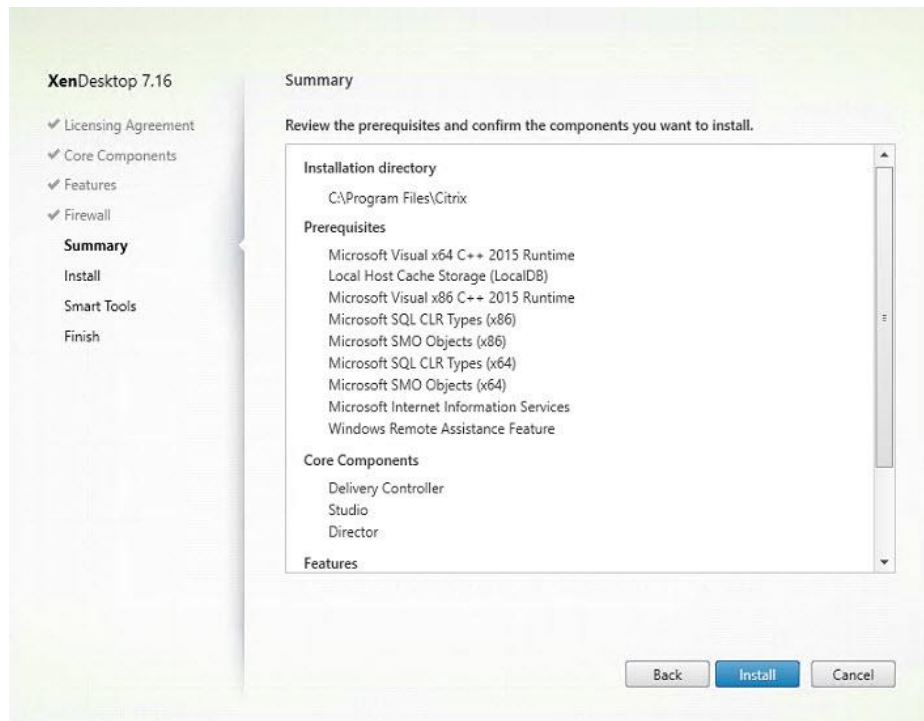




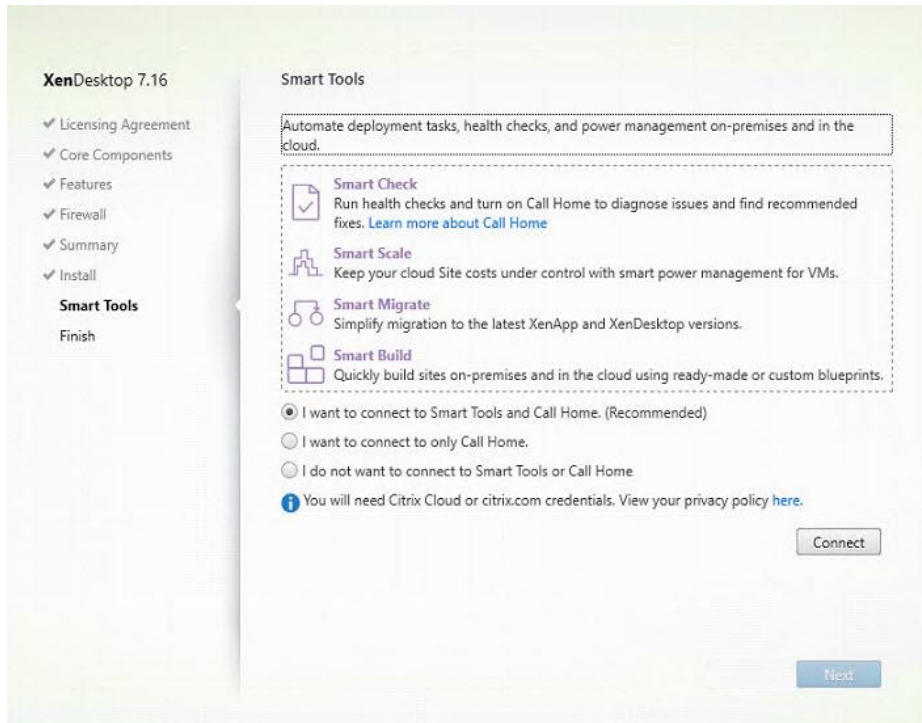
3. Click Delivery Controller.



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

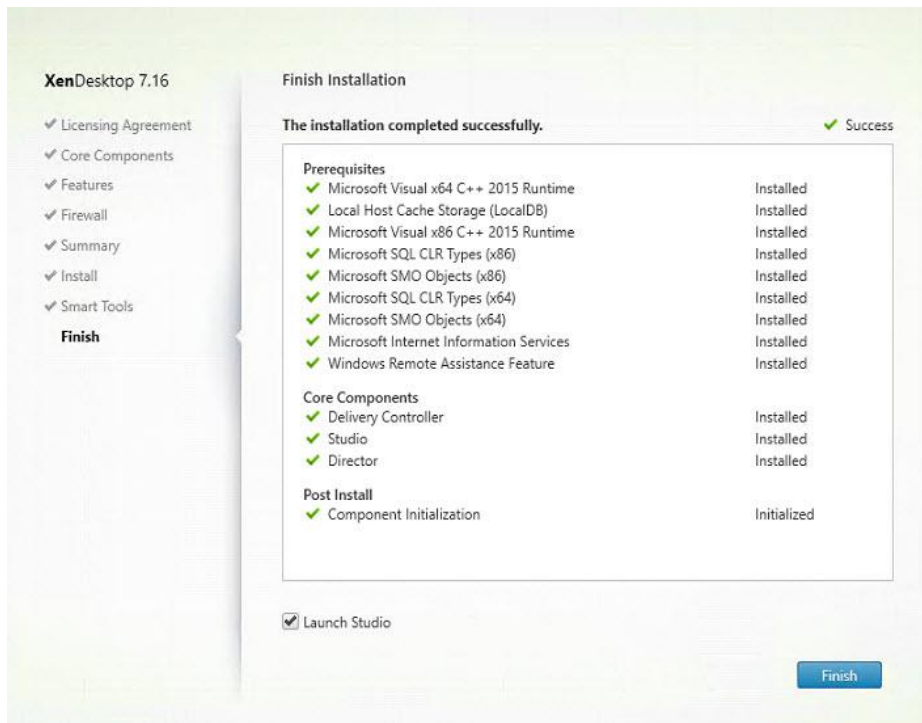


7. (Optional) Click the “I want to participate in Call Home.”
8. Click Next.



9. Verify the components installed successfully.

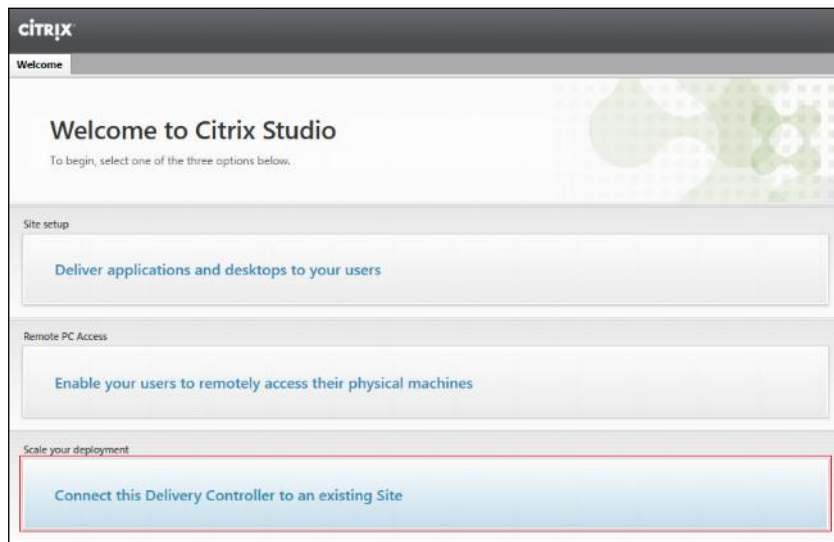
10. Click Finish.



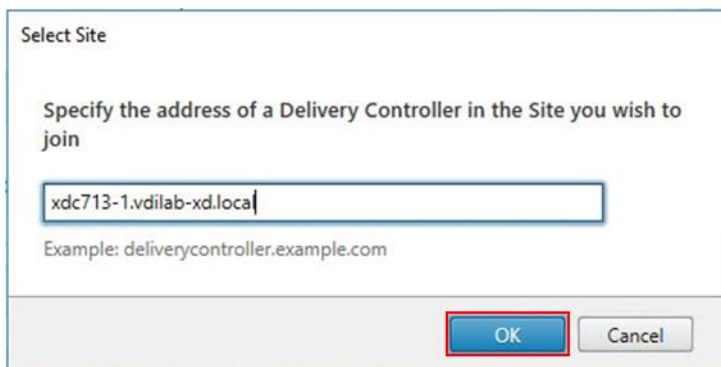
## Add the Second Delivery Controller to the XenDesktop Site

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

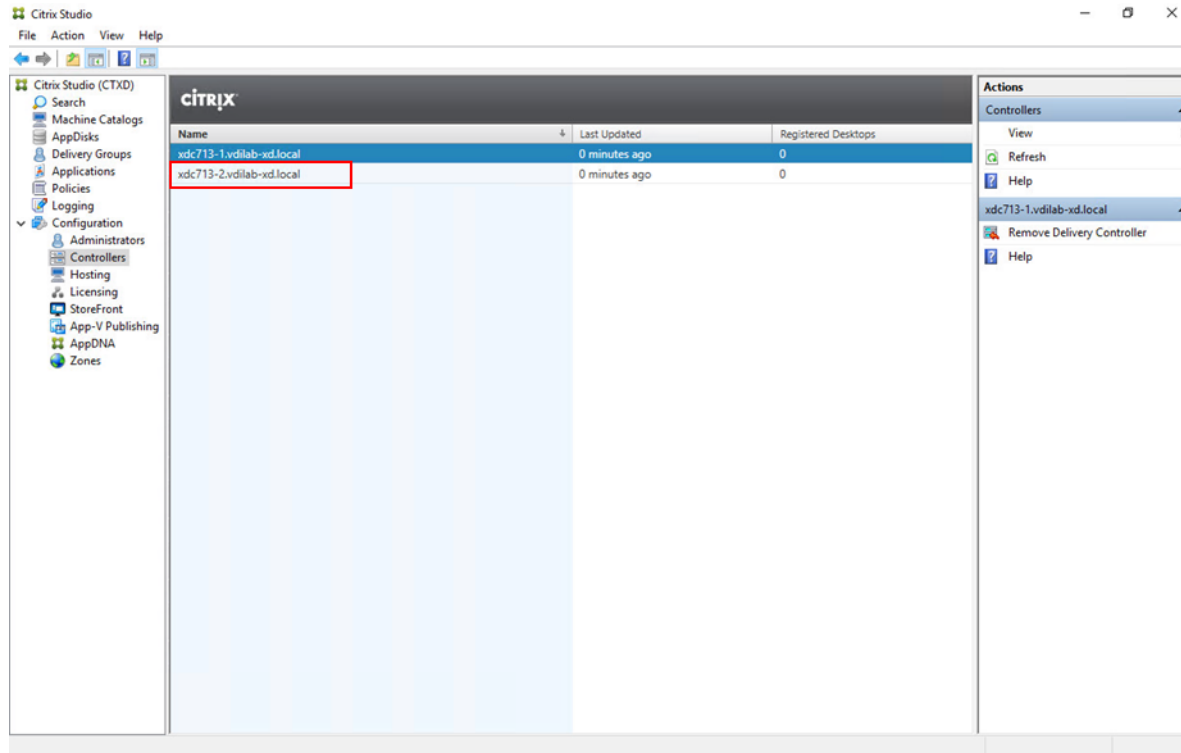
1. In Desktop Studio click the “Connect this Delivery Controller to an existing Site” button.



2. Enter the FQDN of the first delivery controller.
3. Click OK.



4. Click Yes to allow the database to be updated with this controller's information automatically.
5. When complete, test the site configuration and verify the Delivery Controller has been added to the list of Controllers.



## Install and Configure StoreFront

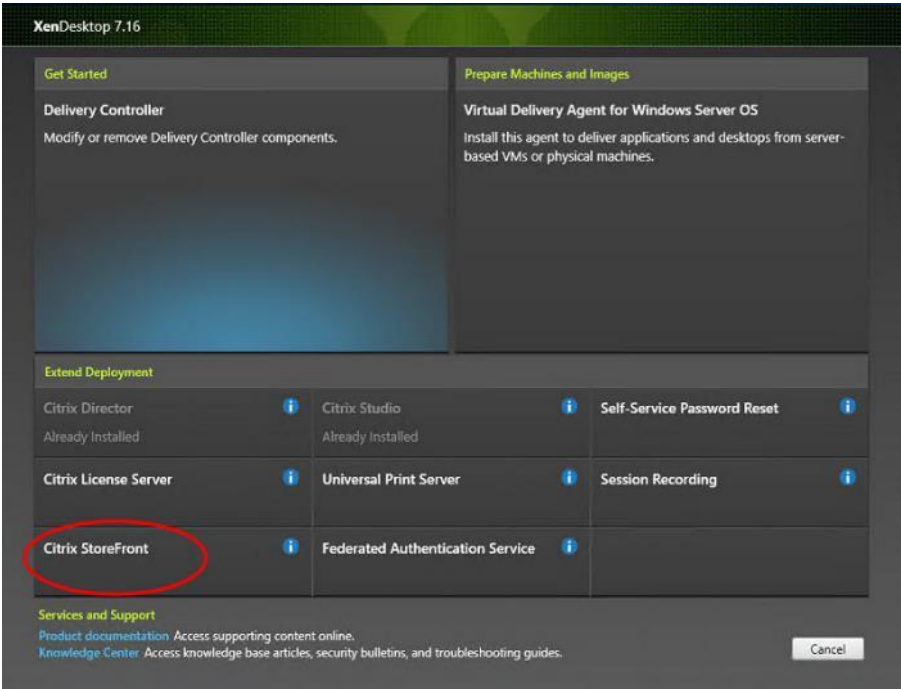
Citrix StoreFront stores aggregate desktops and applications from XenDesktop sites, making resources readily available to users. In this CVD, we created two StoreFront servers on dedicated virtual machines.

To install and configure StoreFront, complete the following steps:

1. To begin the installation of the StoreFront, connect to the first StoreFront server and launch the installer from the Citrix XenDesktop 7.16 ISO.
2. Click Start.

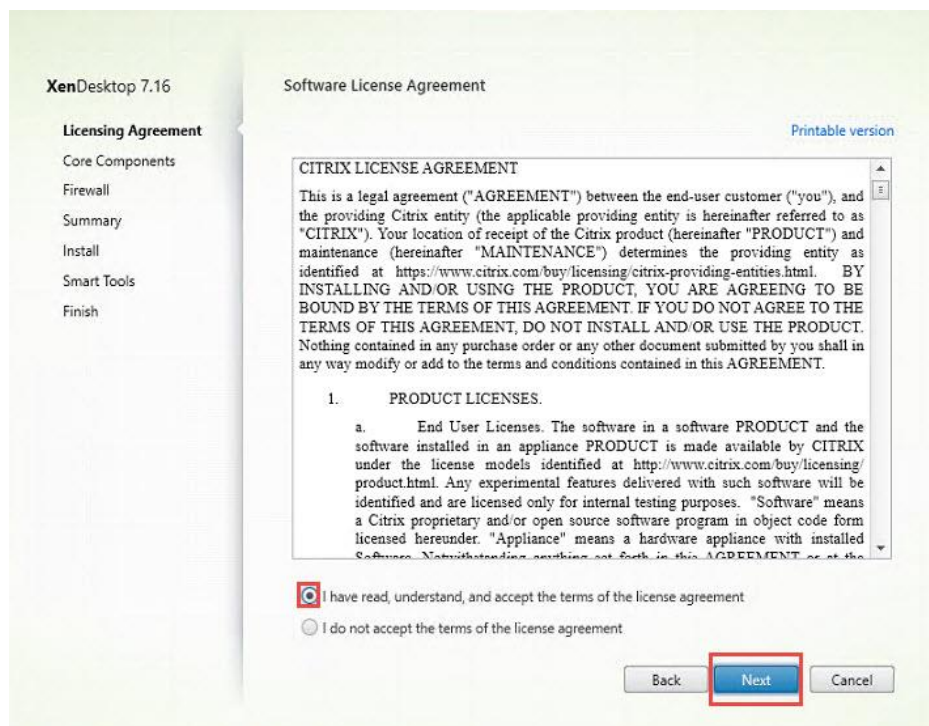


3. Click Extend Deployment Citrix StoreFront.

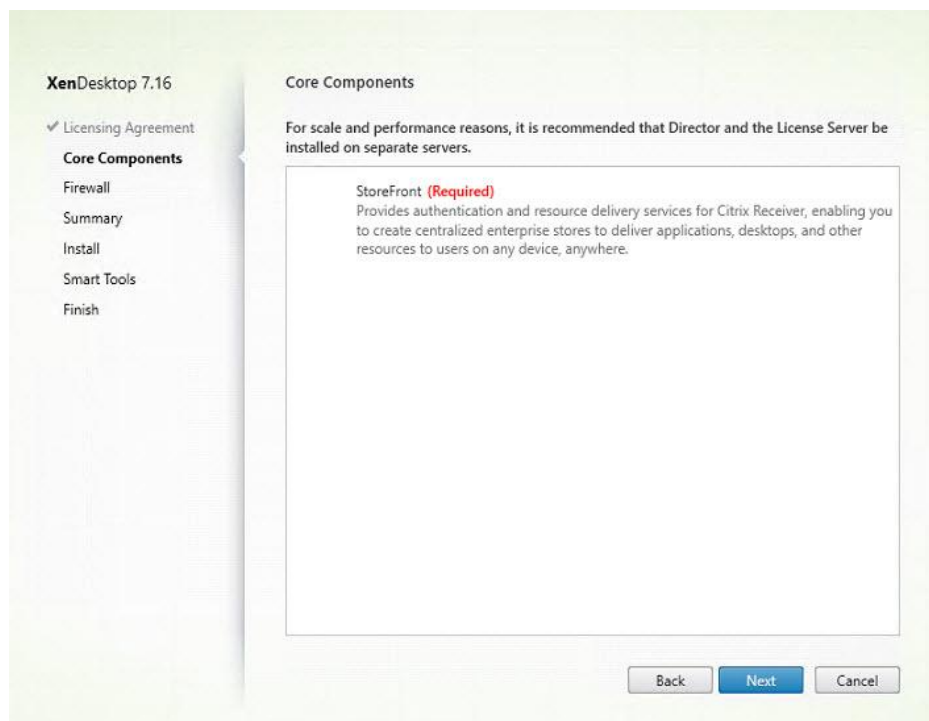


4. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.
5. Click Next.



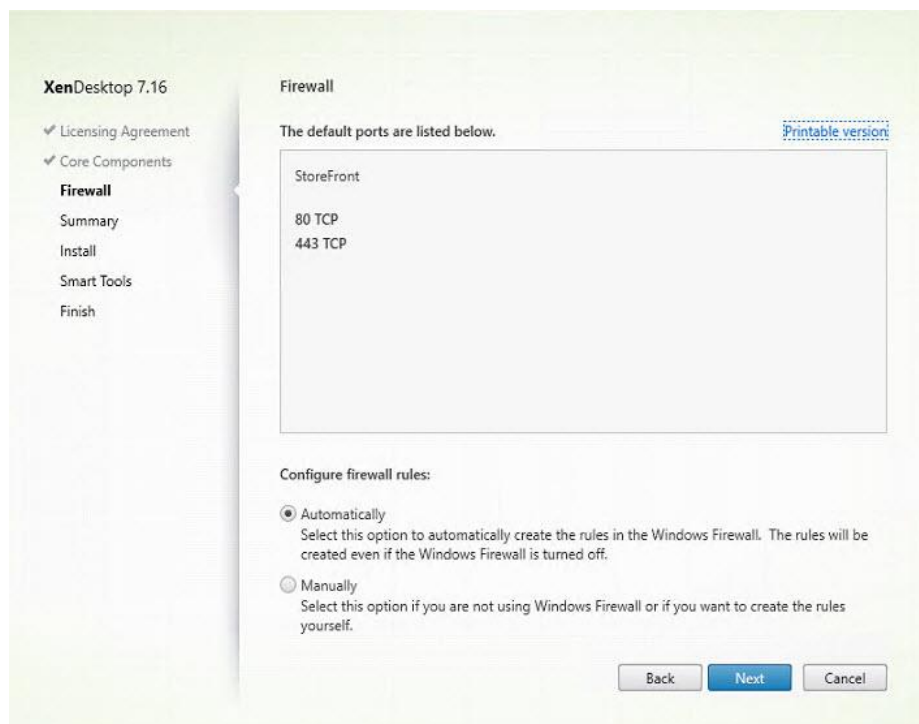


6. Click Next.

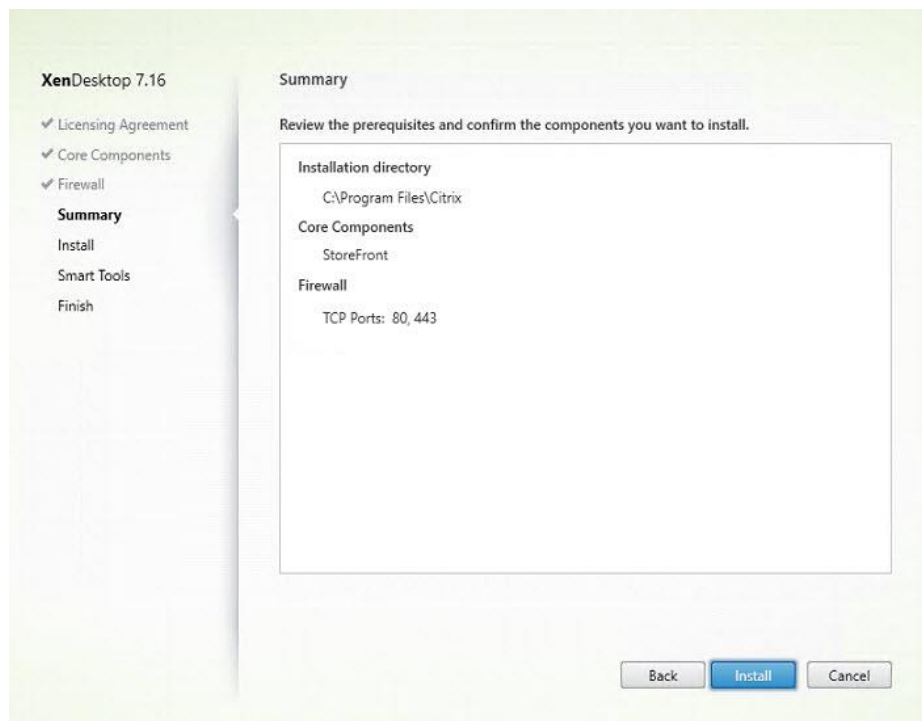


7. Select the default ports and automatically configured firewall rules.

8. Click Next.



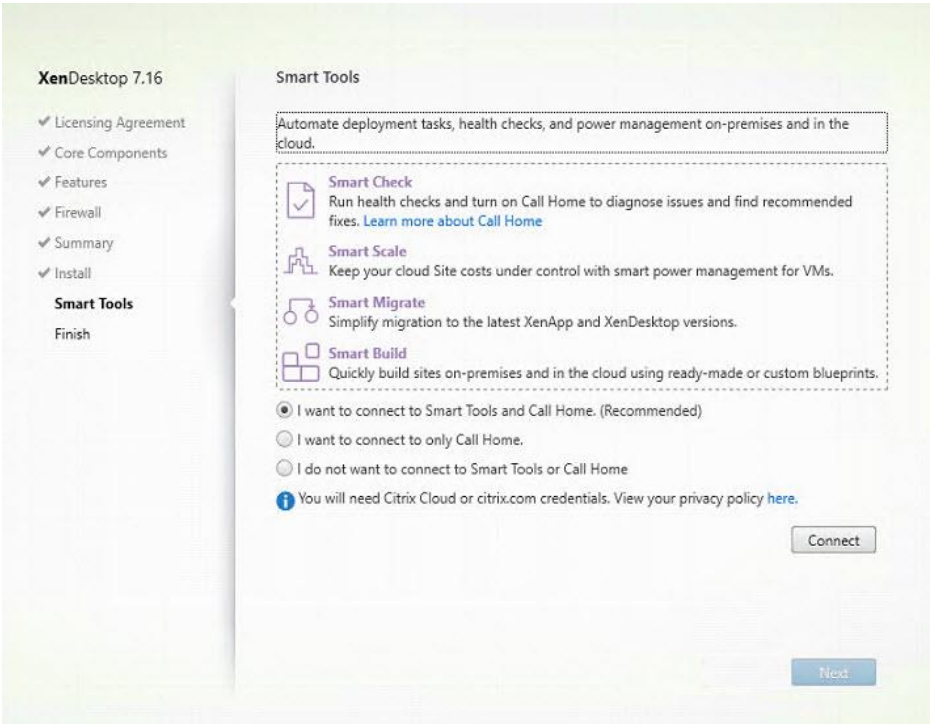
9. Click Install.



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

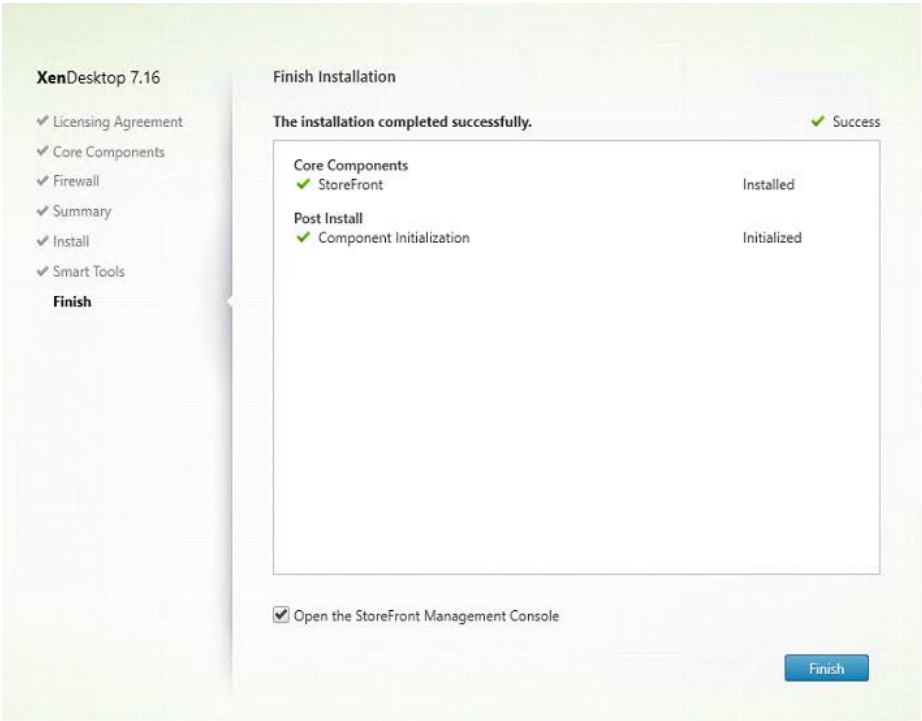
11. Click Next.



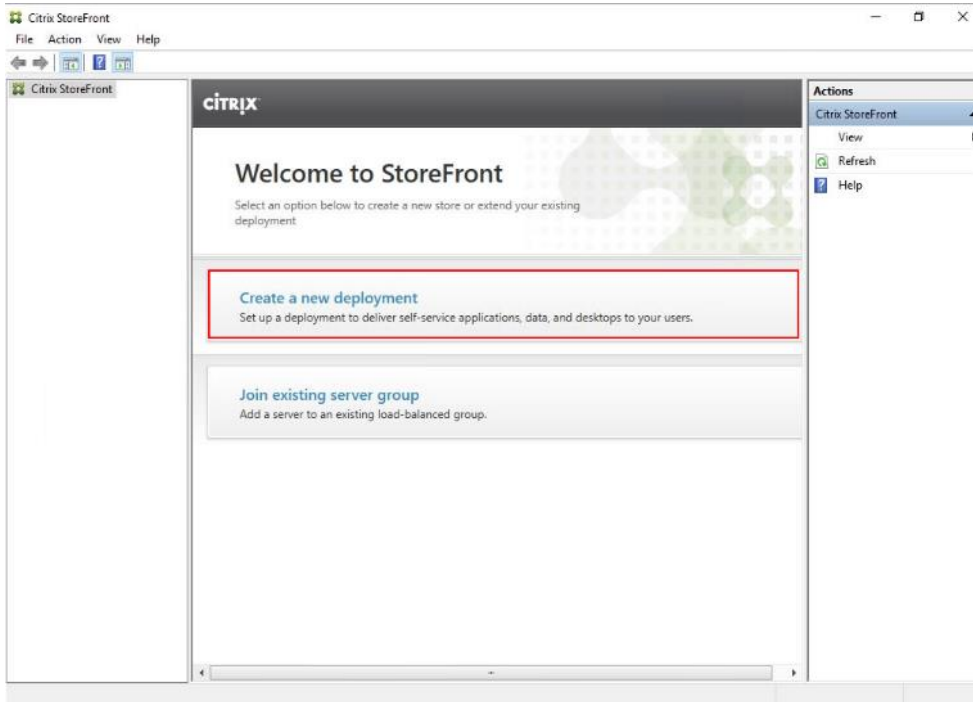


12. Check “Open the StoreFront Management Console.”

13. Click Finish.



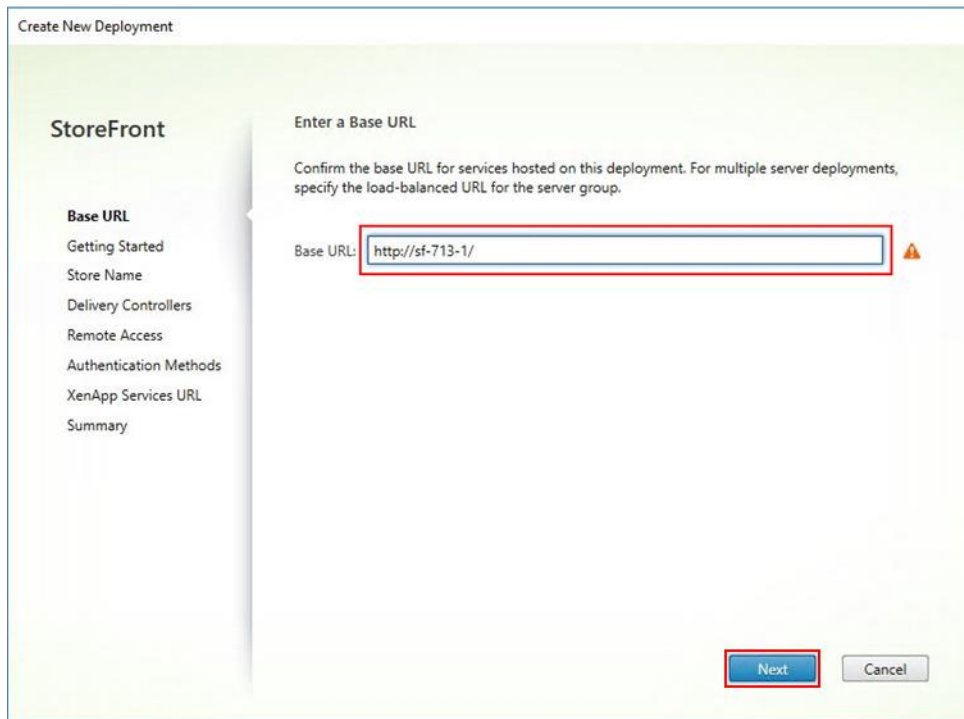
14. Click Create a new deployment.



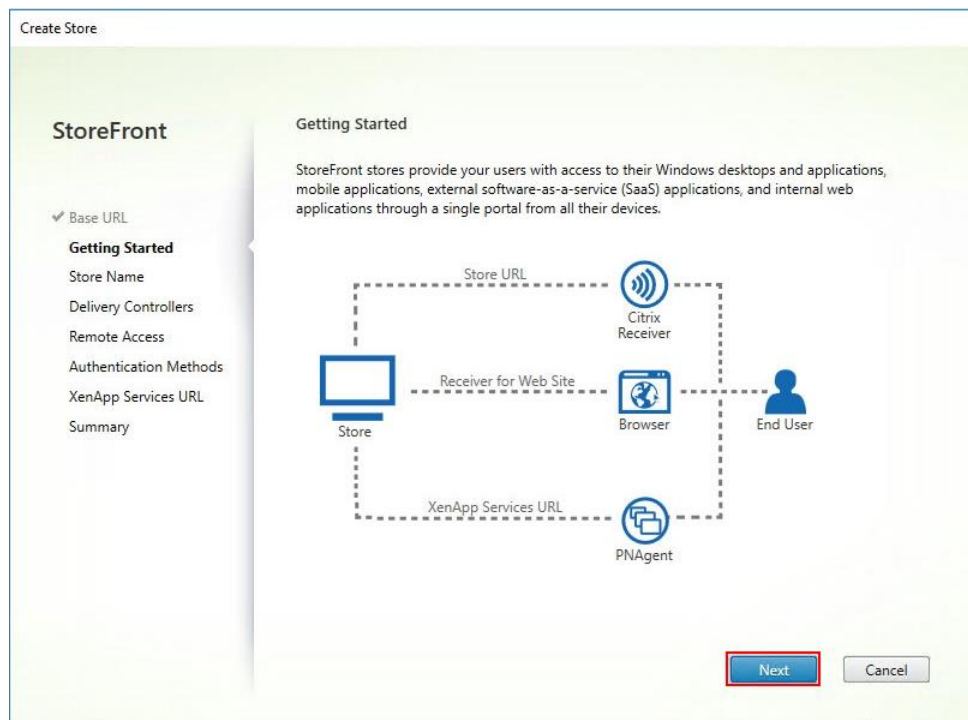
15. Specify the URL of the StoreFront server and click Next.



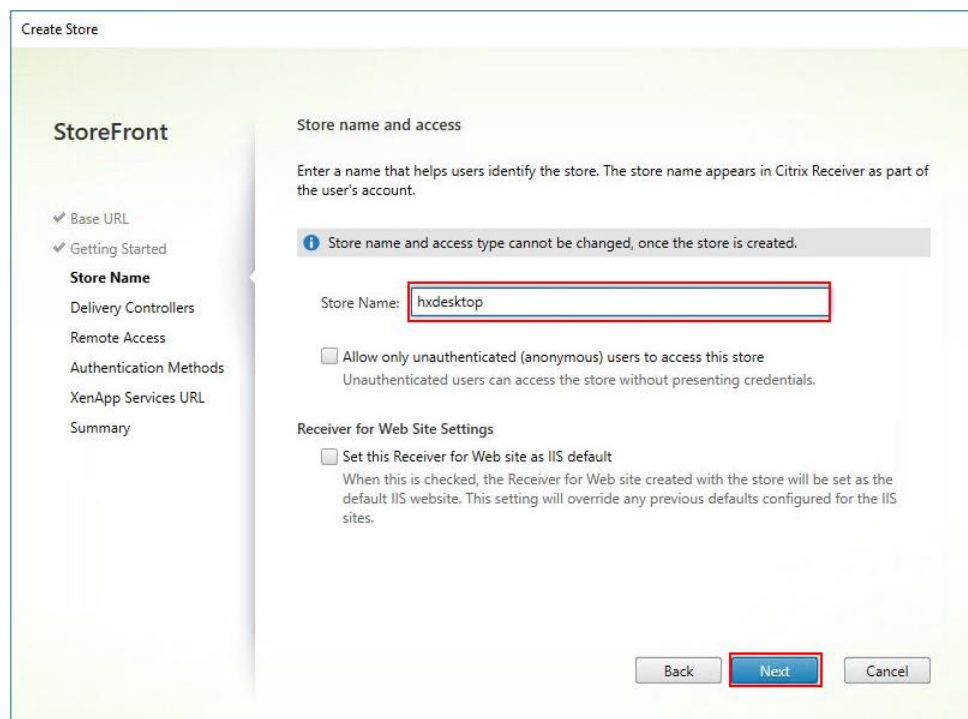
For a multiple server deployment use the load balancing environment in the Base URL box.



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.

The screenshot shows the 'Create Store' wizard in the Citrix StoreFront console. The left sidebar shows the navigation menu with 'Delivery Controllers' selected. The main area is titled 'Delivery Controllers' and contains instructions: 'Specify the XenDesktop delivery controllers, XenApp servers and XenMobile App Controller instances for this store. Citrix recommends grouping delivery controllers based on deployments (sites/farms).' Below this is a table with three columns: 'Name', 'Type', and 'Servers'. The table contains one entry: 'HX' of type 'XenDesktop' with servers 'XD-713-1.vdilab-x...'. Below the table are 'Add...', 'Edit...', and 'Remove' buttons. At the bottom right are 'Back', 'Next' (highlighted with a red box), and 'Cancel' buttons.

Name	Type	Servers
HX	XenDesktop	XD-713-1.vdilab-x...

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.

The screenshot shows the 'Create Store' wizard in the Citrix StoreFront console, now at the 'Remote Access' step. The left sidebar shows 'Remote Access' selected. The main area is titled 'Remote Access' and contains instructions: 'Enabling remote access will allow users outside the firewall to access resources securely. You need to add a NetScaler Gateway once remote access is enabled.' Below this is a checkbox 'Enable Remote Access' which is checked. Underneath is a section 'Select the permitted level of access to internal resources' with two radio button options: 'Allow users to access only resources delivered through StoreFront (No VPN tunnel)' (selected) and 'Allow users to access all resources on the internal network (Full VPN tunnel)'. Below these is a text box for 'NetScaler Gateway appliances' with an 'Add...' button. At the bottom is a 'Default appliance:' dropdown menu. At the bottom right are 'Back', 'Next' (highlighted with a red box), and 'Cancel' buttons.

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:

**Create Store**

**StoreFront**

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

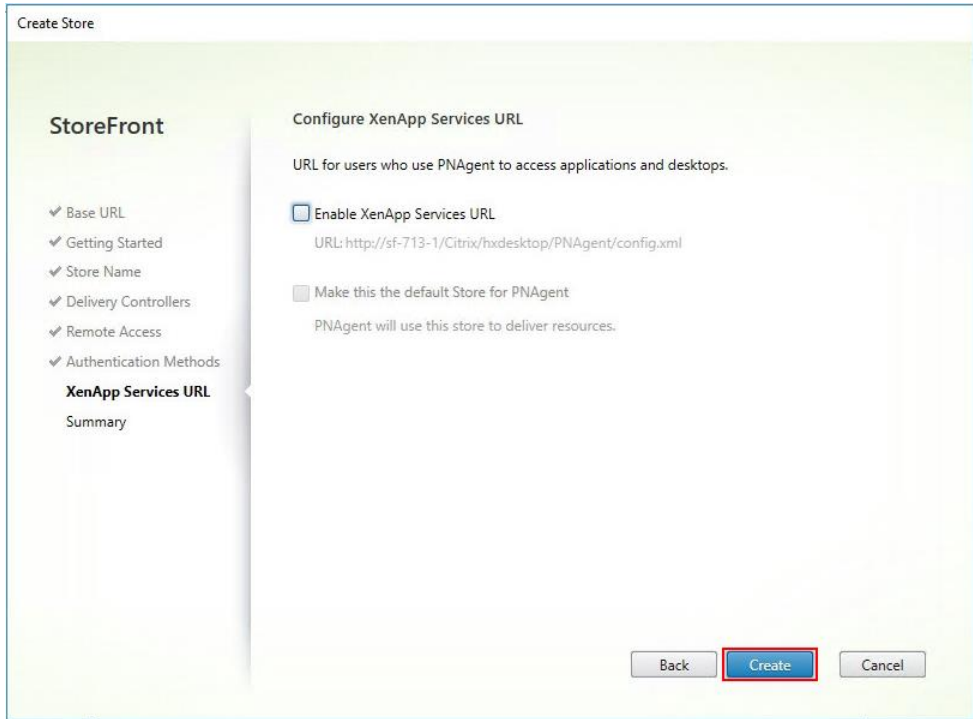
**Configure Authentication Methods**

Select the methods which users will use to authenticate and access resources.

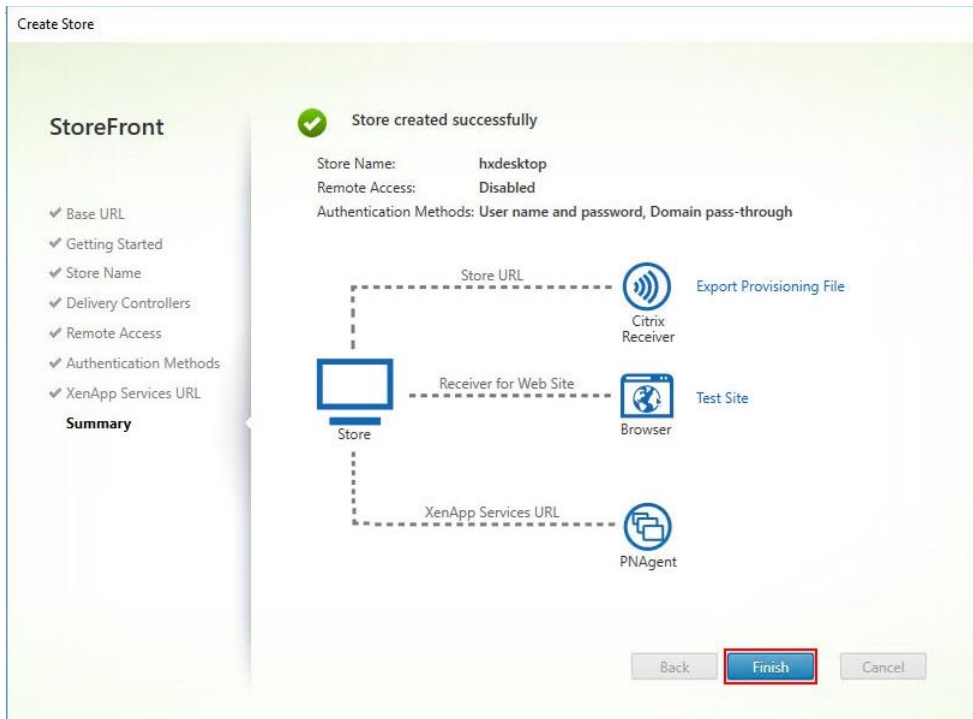
Method
<input checked="" type="checkbox"/> User name and password
<input type="checkbox"/> SAML Authentication
<input checked="" type="checkbox"/> Domain pass-through <small>Can be enabled / disabled separately on Receiver for Web sites</small>
<input type="checkbox"/> Smart card <small>Can be enabled / disabled separately on Receiver for Web sites</small>
<input type="checkbox"/> HTTP Basic
<input type="checkbox"/> Pass-through from NetScaler Gateway

Back Next Cancel

21. Username and password: Users enter their credentials and are authenticated when they access their stores.
22. Domain pass-through: Users authenticate to their domain-joined Windows computers and their credentials are used to log them on automatically when they access their stores.
23. Configure the XenApp Service URL for users who use PNAgent to access the applications and desktops and click Create.



24. After creating the store click Finish.

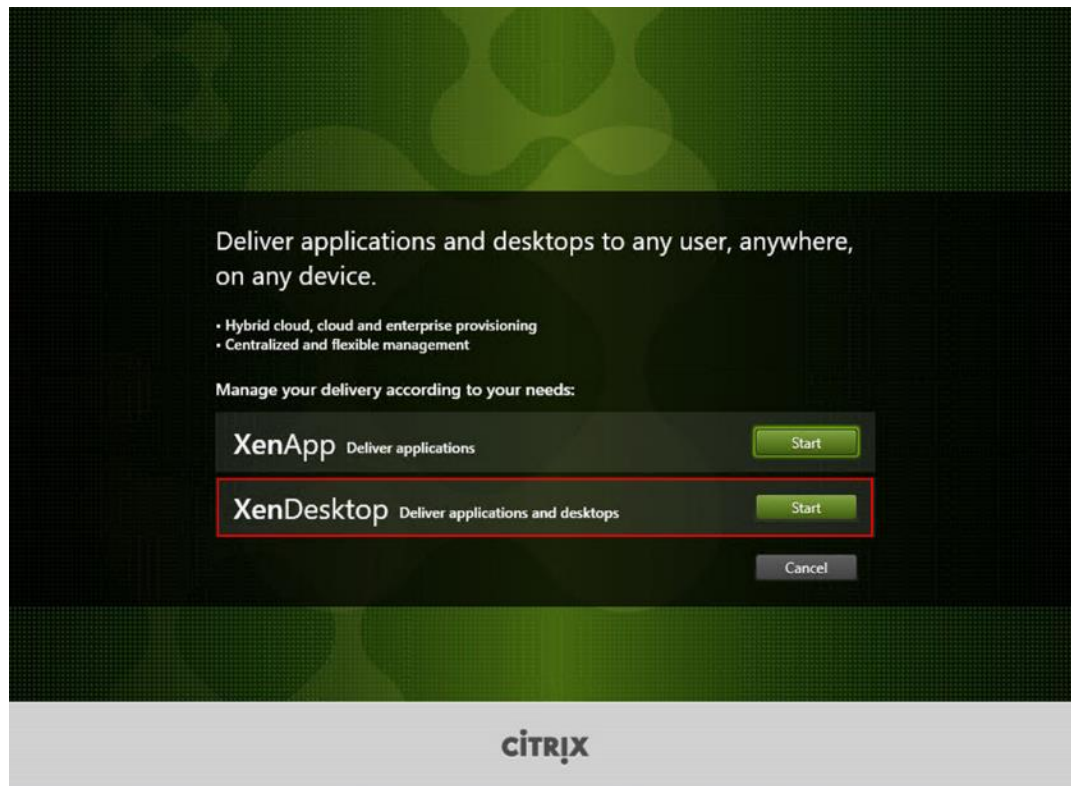


## Additional StoreFront Configuration

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

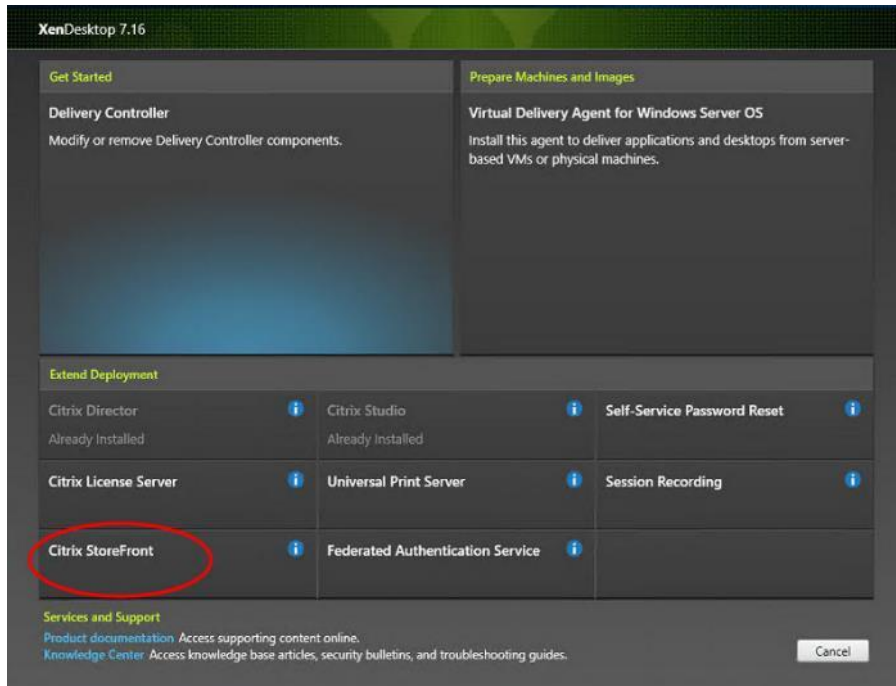
To configure additional StoreFront server, complete the following steps:

1. To begin the installation of the second StoreFront, connect to the second StoreFront server and launch the installer from the Citrix XenDesktop 7.16 ISO.
2. Click Start.

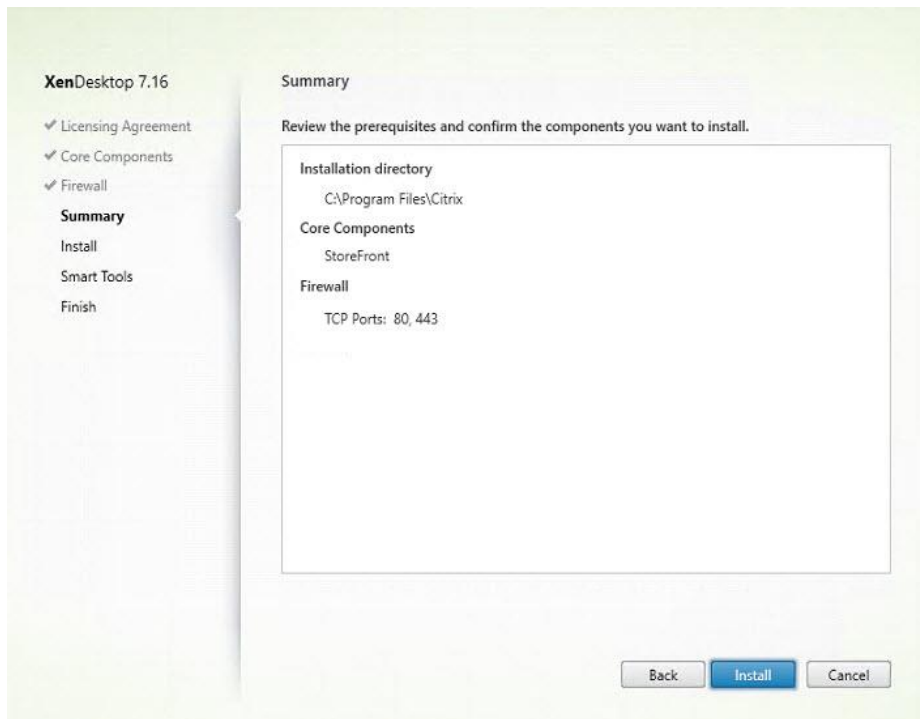


3. Click Extended Deployment Citrix StoreFront.





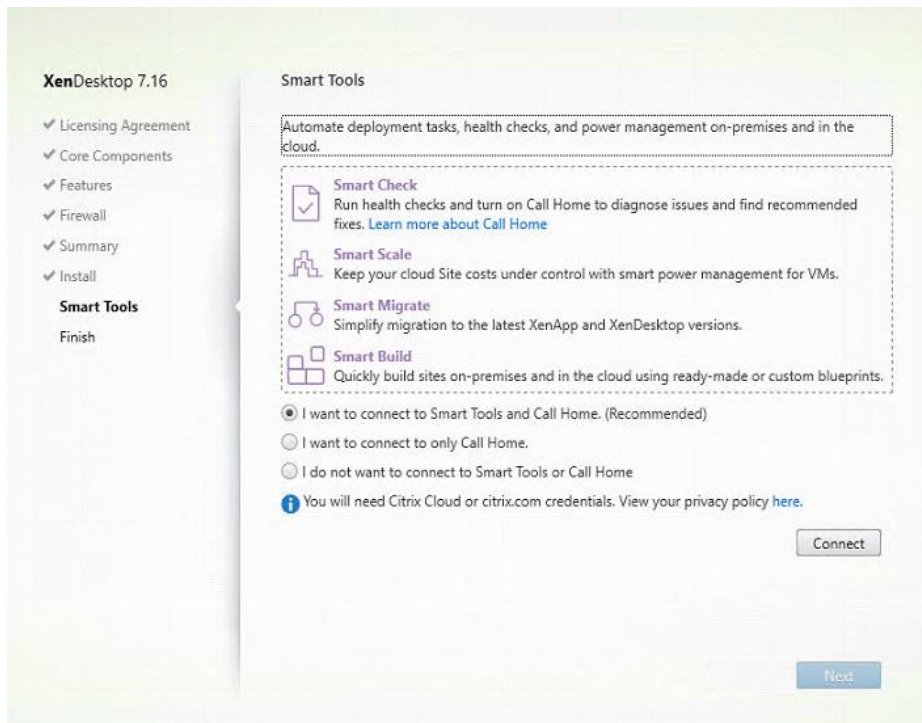
4. Repeat the same steps used to install the first StoreFront.
5. Review the Summary configuration.
6. Click Install.



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

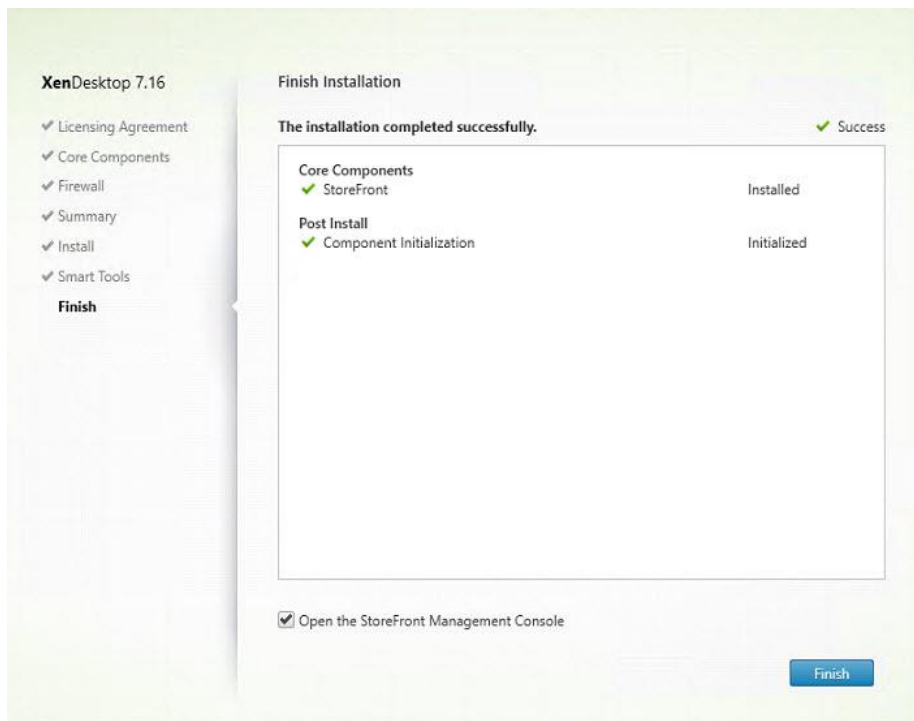


8. Click Next.



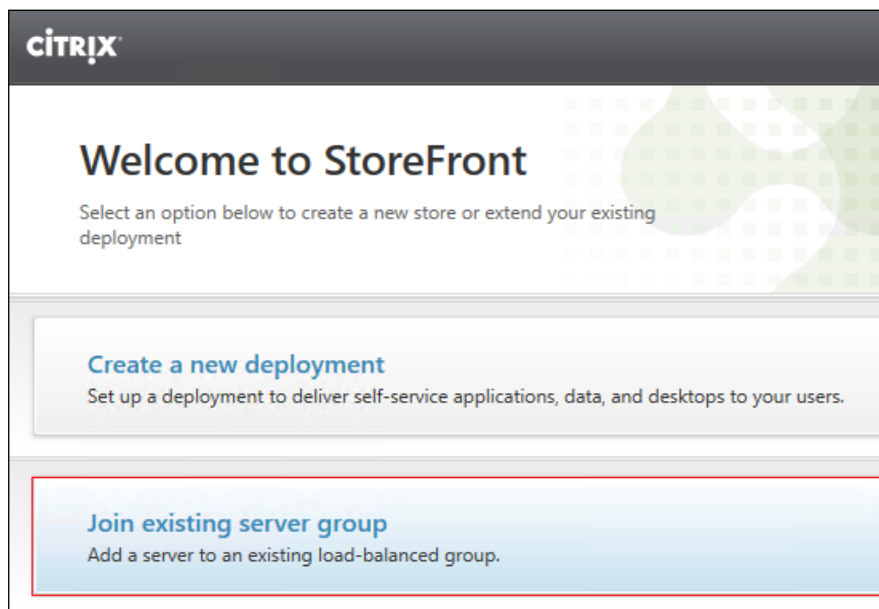
9. Check "Open the StoreFront Management Console."

10. Click Finish.

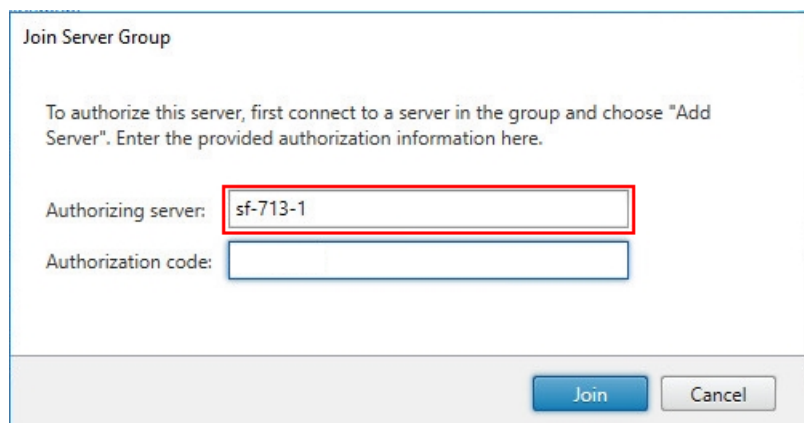


To configure the second StoreFront if used, complete the following steps:


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



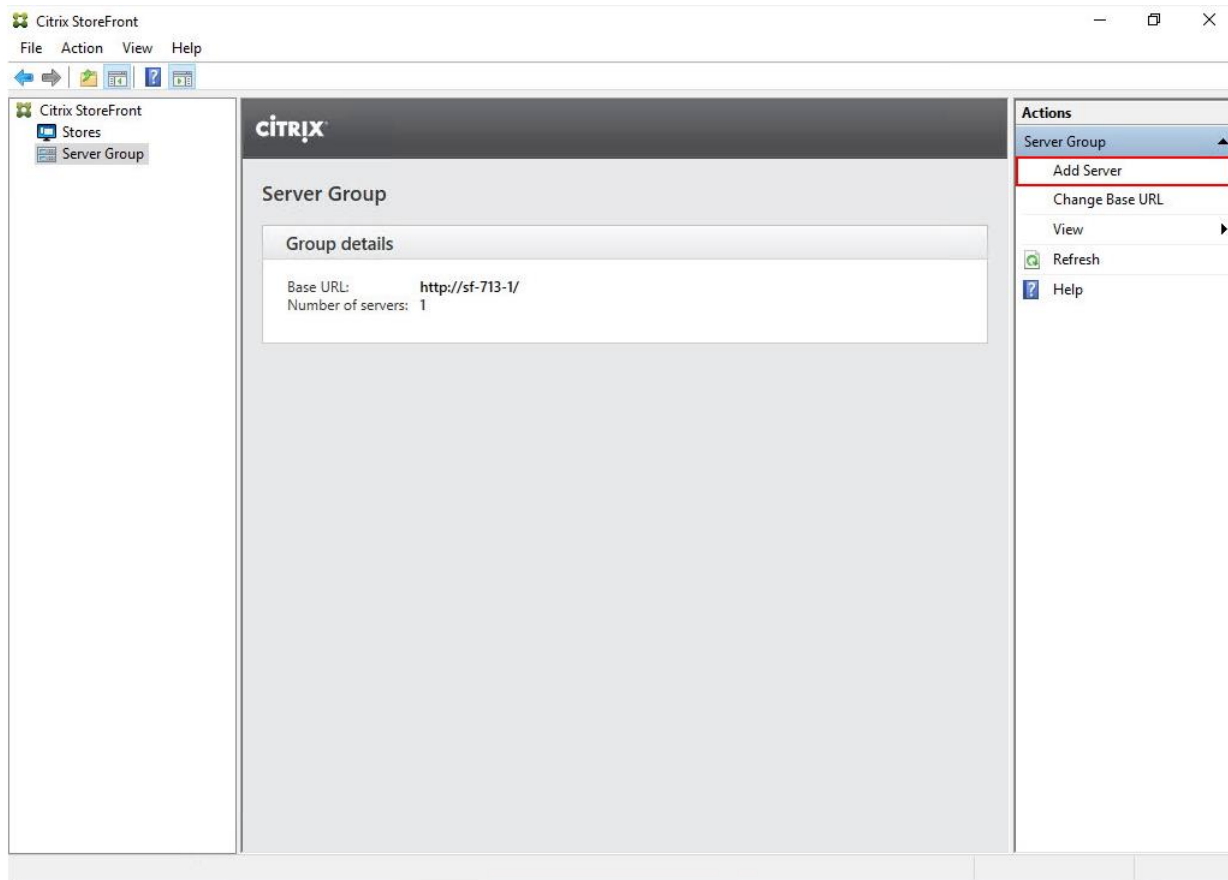
2. In the Join Server Group dialog, enter the name of the first Storefront server.




3. Before the additional StoreFront server can join the server group, you must connect to the first Storefront server, add the second server, and obtain the required authorization information.
4. Connect to the first StoreFront server.
5. Using the StoreFront menu on the left, you can scroll through the StoreFront management options.
6. Select Server Group from the menu.



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

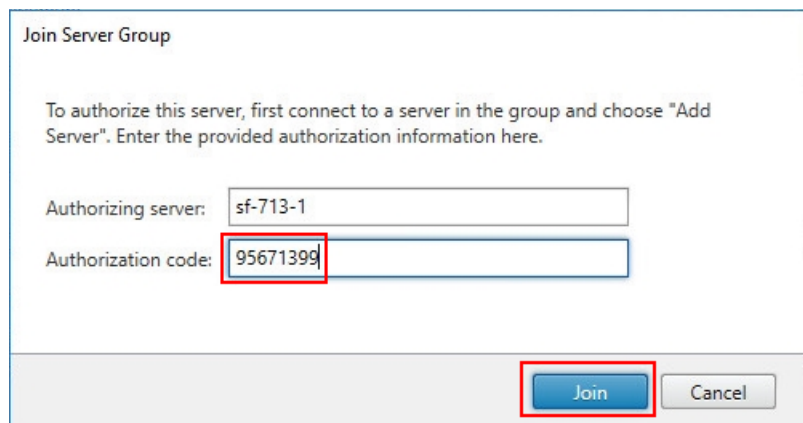


8. Copy the Authorization code from the Add Server dialog.



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

10. Click Join.

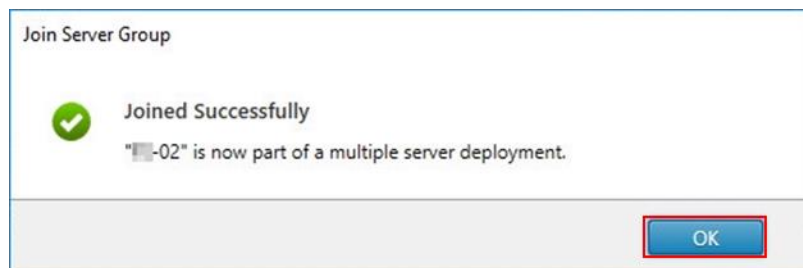


The 'Join Server Group' dialog box contains the following elements:

- Title: Join Server Group
- Instruction: To authorize this server, first connect to a server in the group and choose "Add Server". Enter the provided authorization information here.
- Field 1: Authorizing server: sf-713-1
- Field 2: Authorization code: 95671399 (highlighted with a red box)
- Buttons: Join (highlighted with a red box) and Cancel

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

12. Click OK.



The 'Join Server Group' dialog box displays the following success message:

- Icon: Green checkmark
- Text: Joined Successfully
- Text: "sf-02" is now part of a multiple server deployment.
- Button: OK (highlighted with a red box)

The second StoreFront is now in the Server Group.

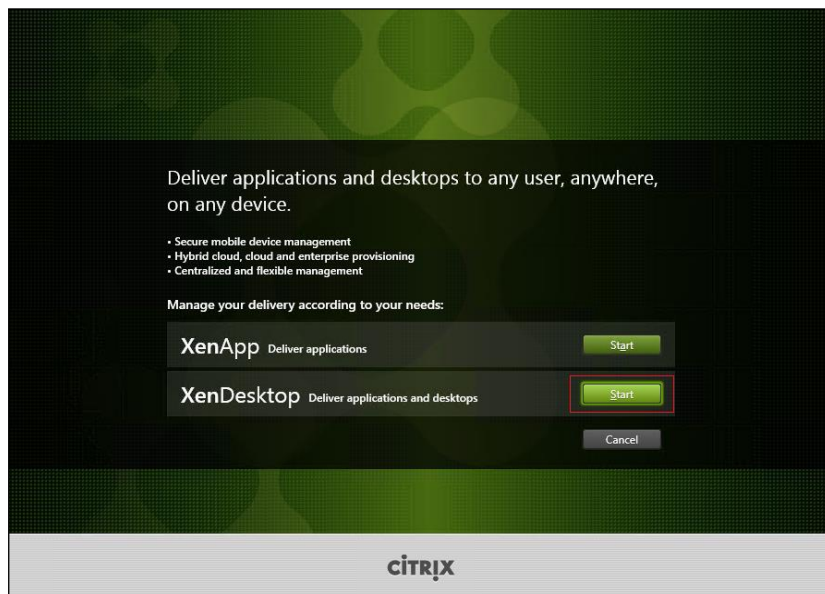
## Install XenDesktop Virtual Desktop Agents

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.

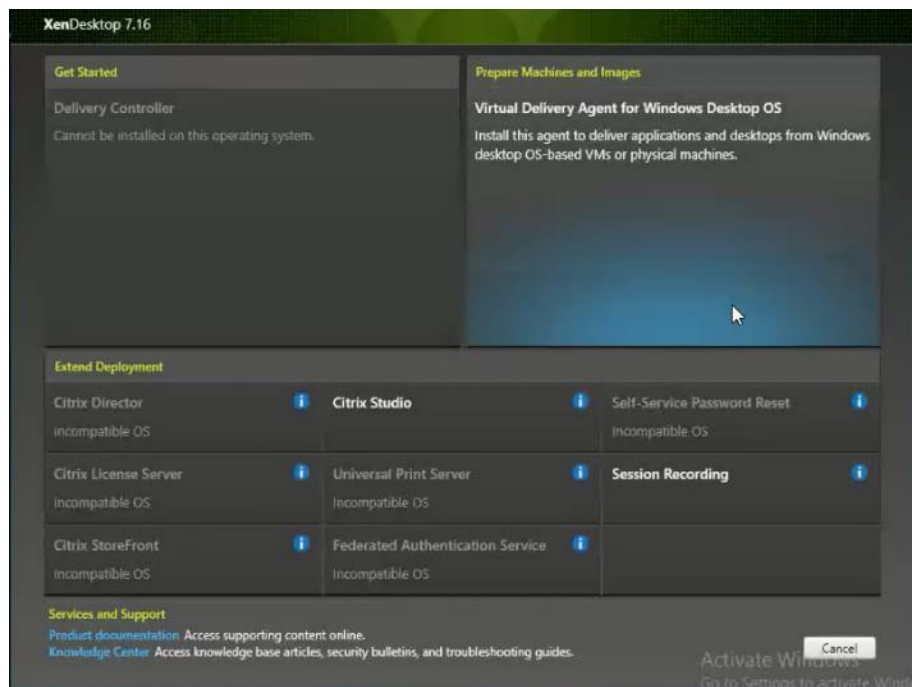
By default, when you install the Virtual Delivery Agent, Citrix User Profile Management is installed silently on master images. (Using profile management as a profile solution is optional but was used for this CVD, and is described in a later section.)

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

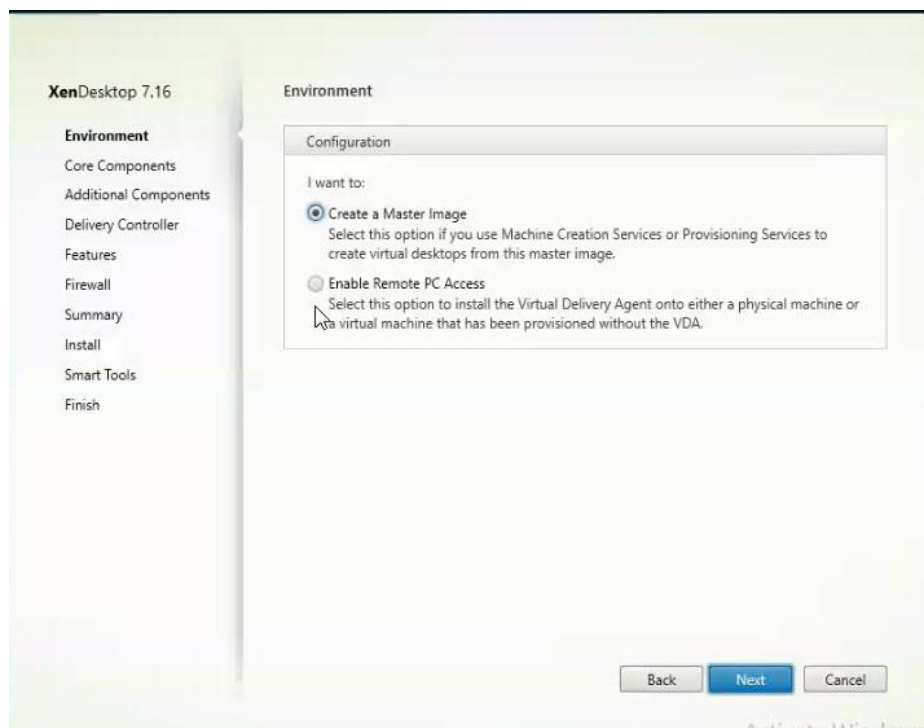
1. Launch the XenDesktop installer from the XenDesktop 7.16 ISO.
2. Click Start on the Welcome Screen.



- 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.

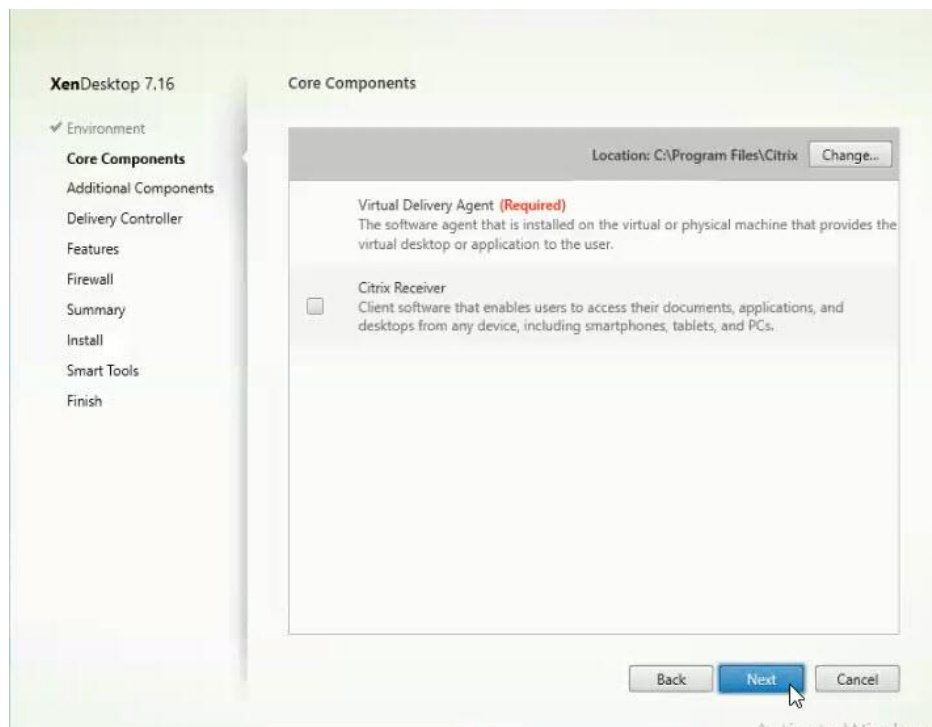


- Select "Create a Master Image."
- Click Next.

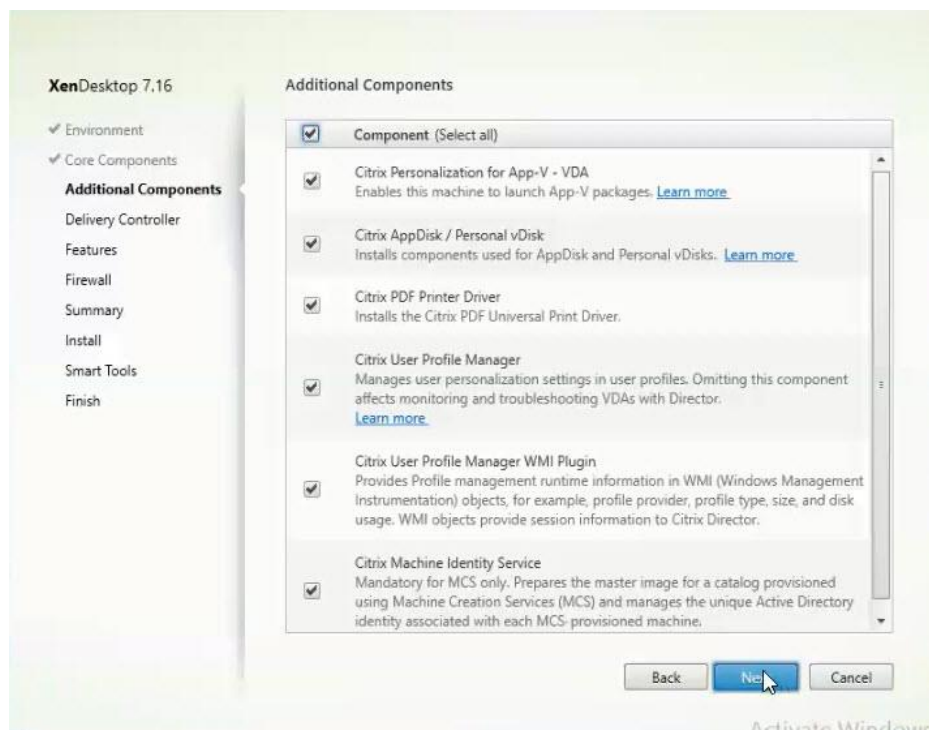


6. Optional: Select Citrix Receiver.

7. Click Next.

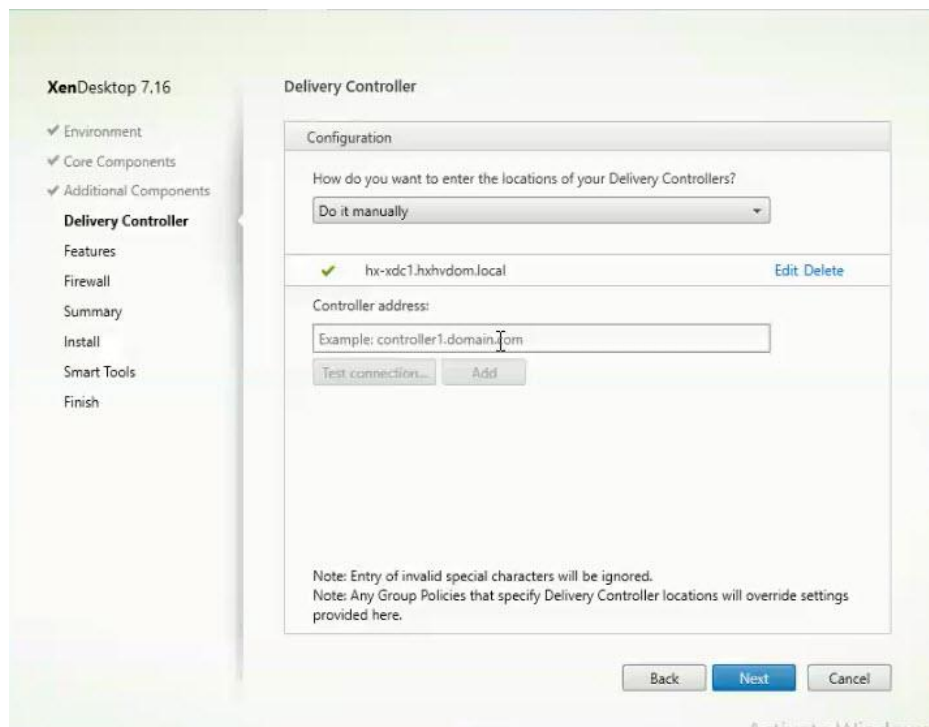


8. Click Next.



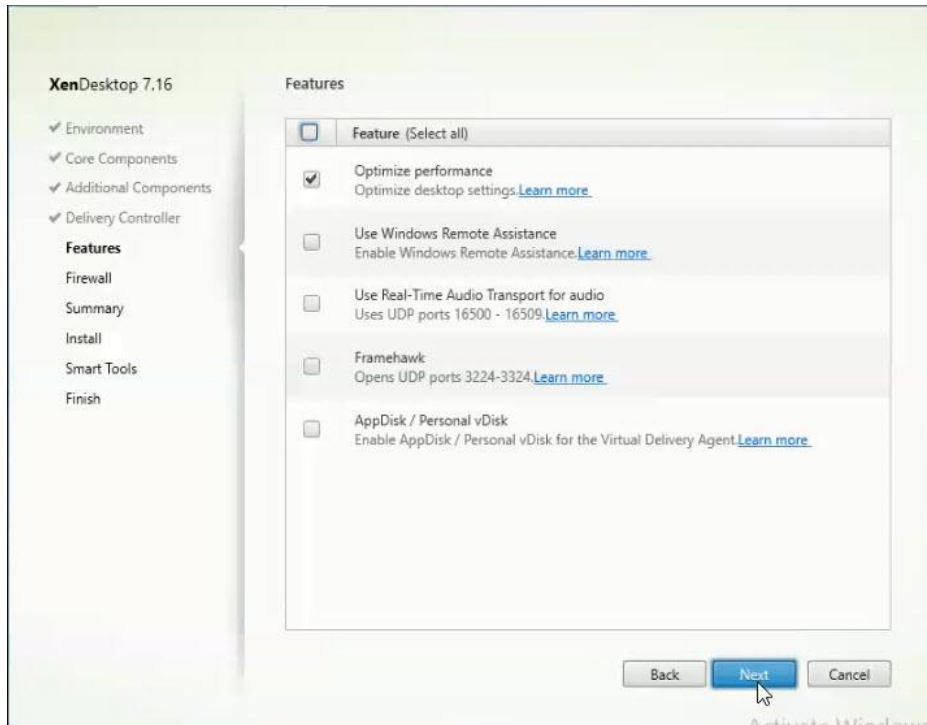
9. Select “Do it manually” and specify the FQDN of the Delivery Controllers.

10. Click Next.



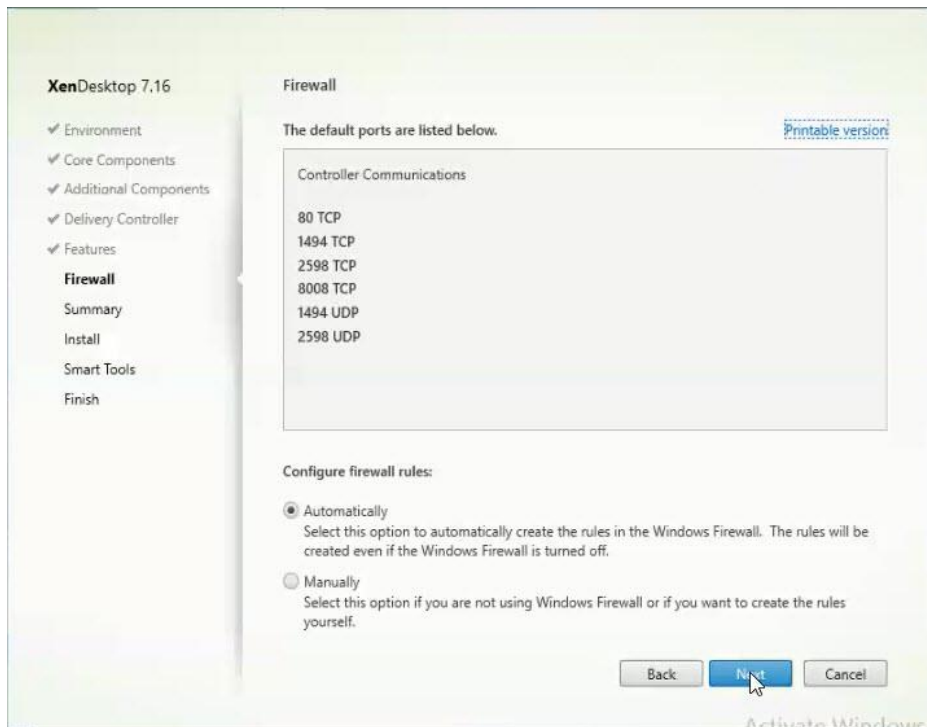
11. Accept the default features.

12. Click Next.



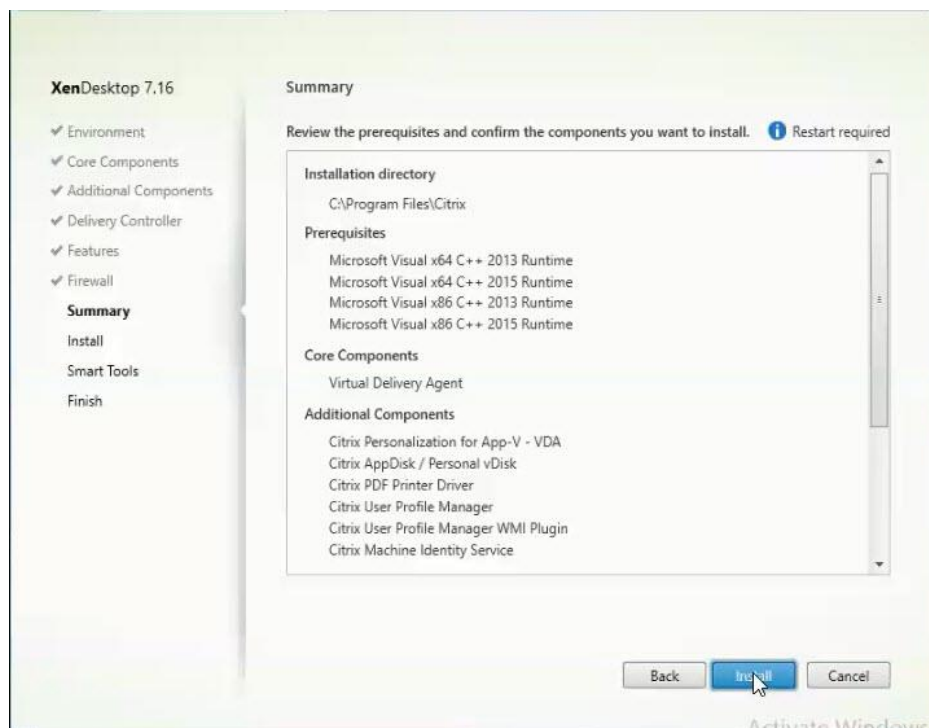
13. Allow the firewall rules to be configured automatically.

14. Click Next.

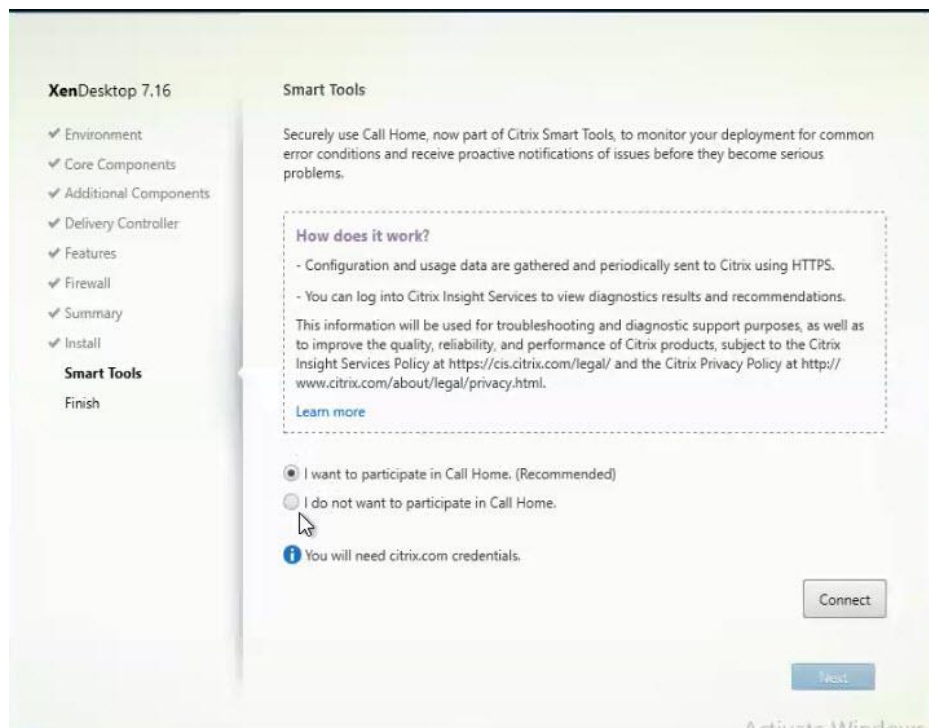


15. Verify the Summary and click Install.



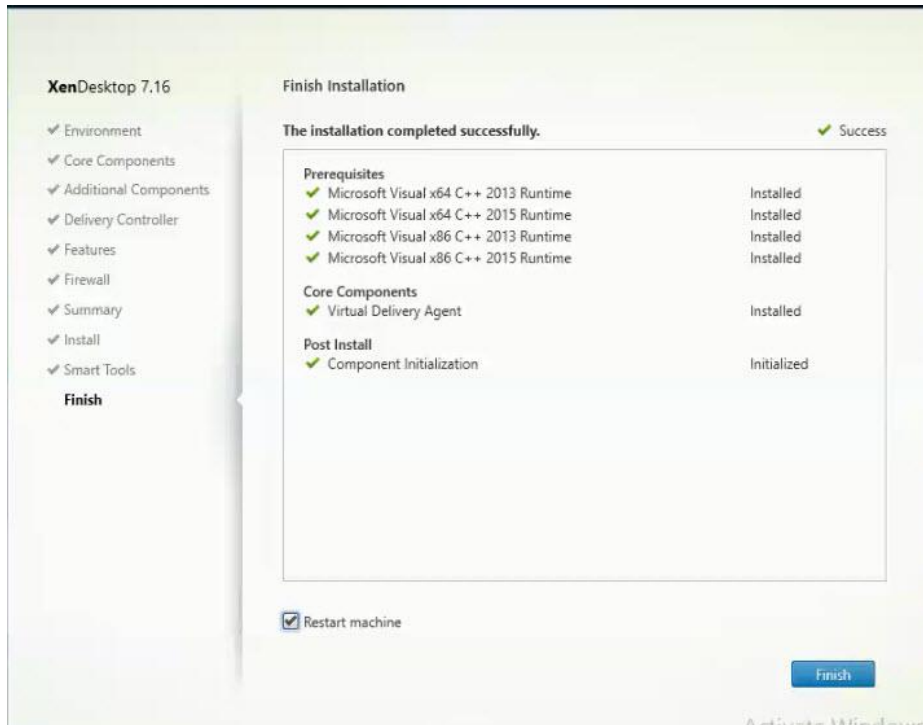


16. (Optional) Select Call Home participation.

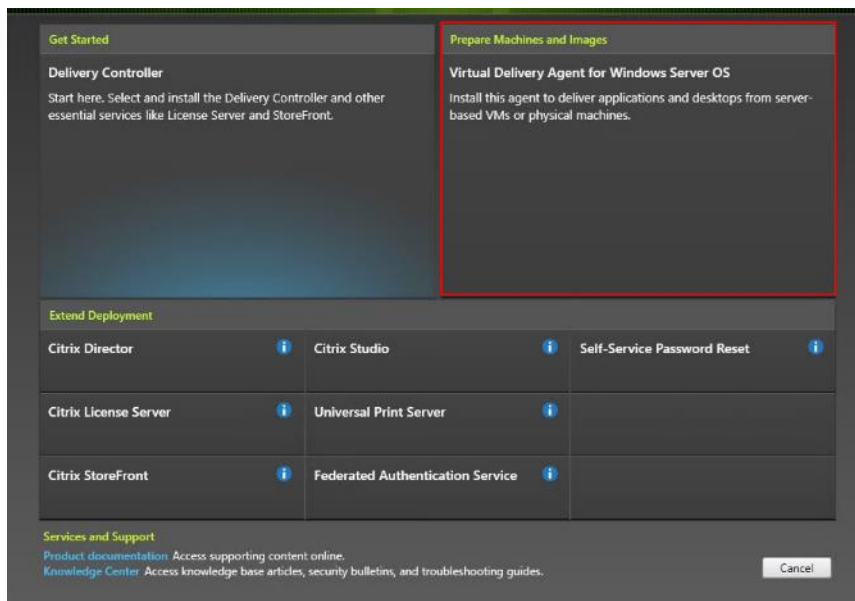


17. (Optional) check “Restart Machine.”

18. Click Finish.




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

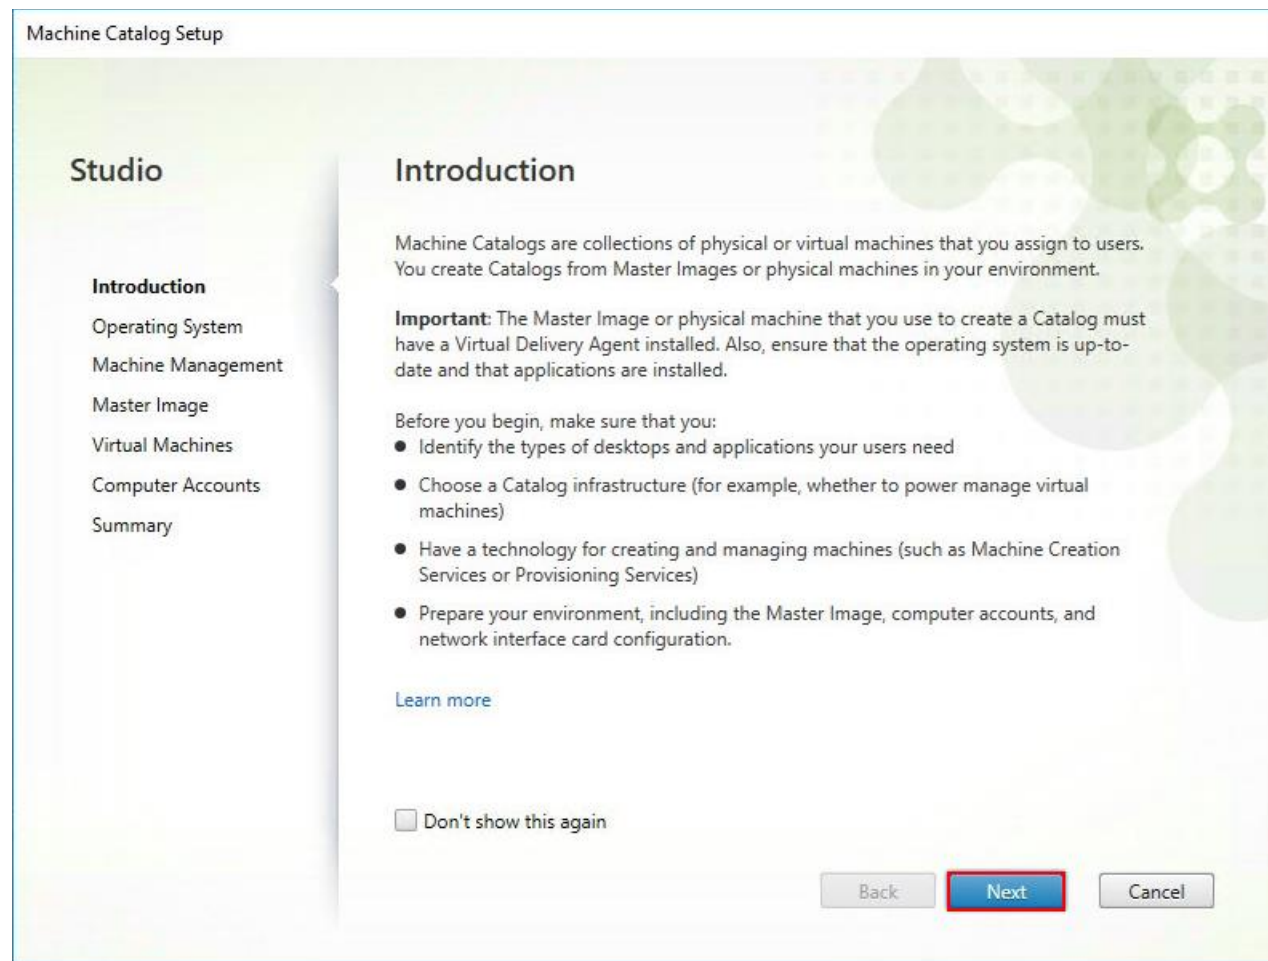


## Persistent Static Provisioned with MCS

1. Connect to a XenDesktop server and launch Citrix Studio.
2. Choose Create Machine Catalog from the drop-down list.

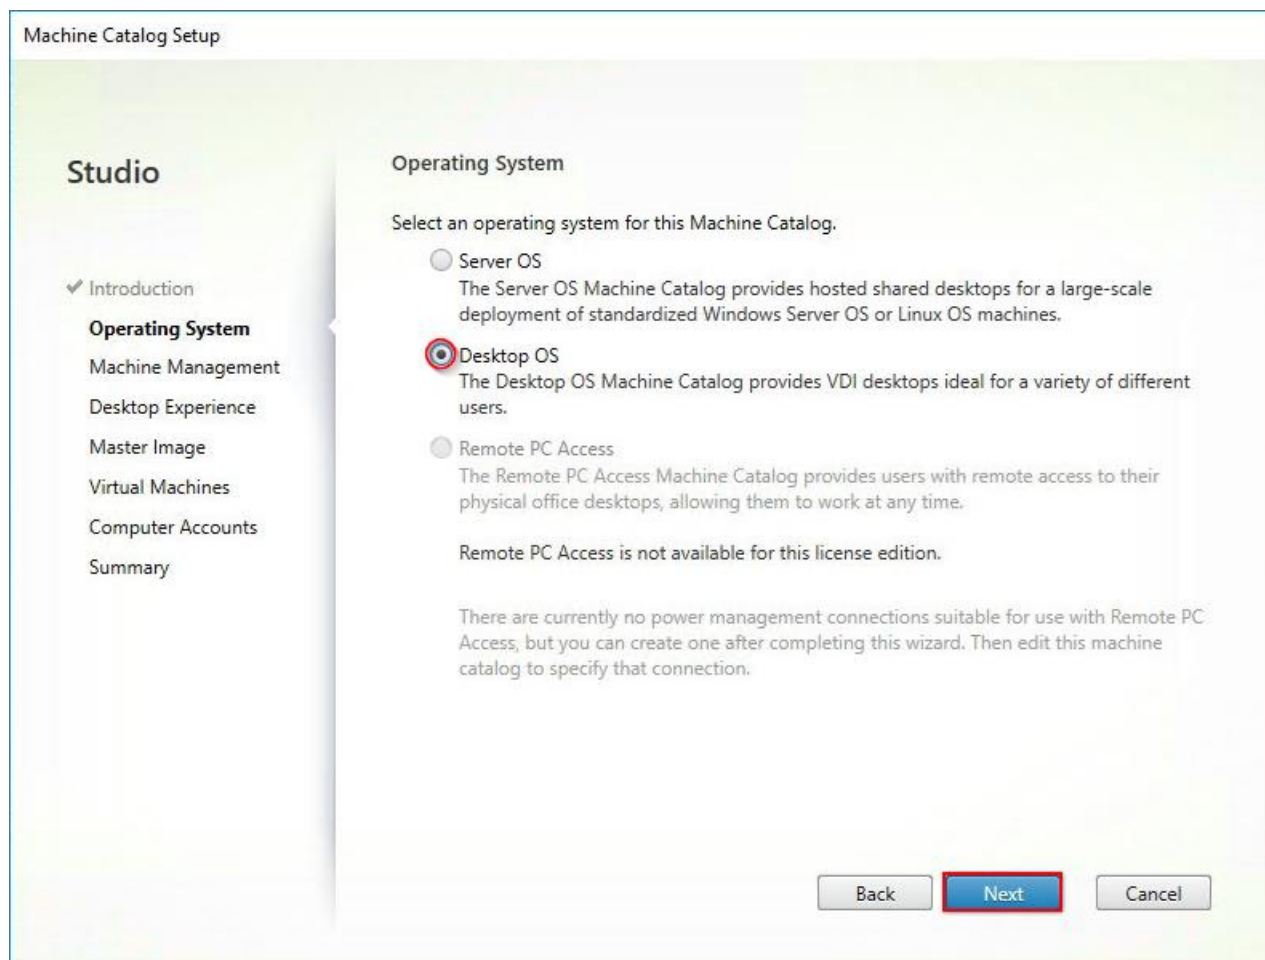


3. Click Next.



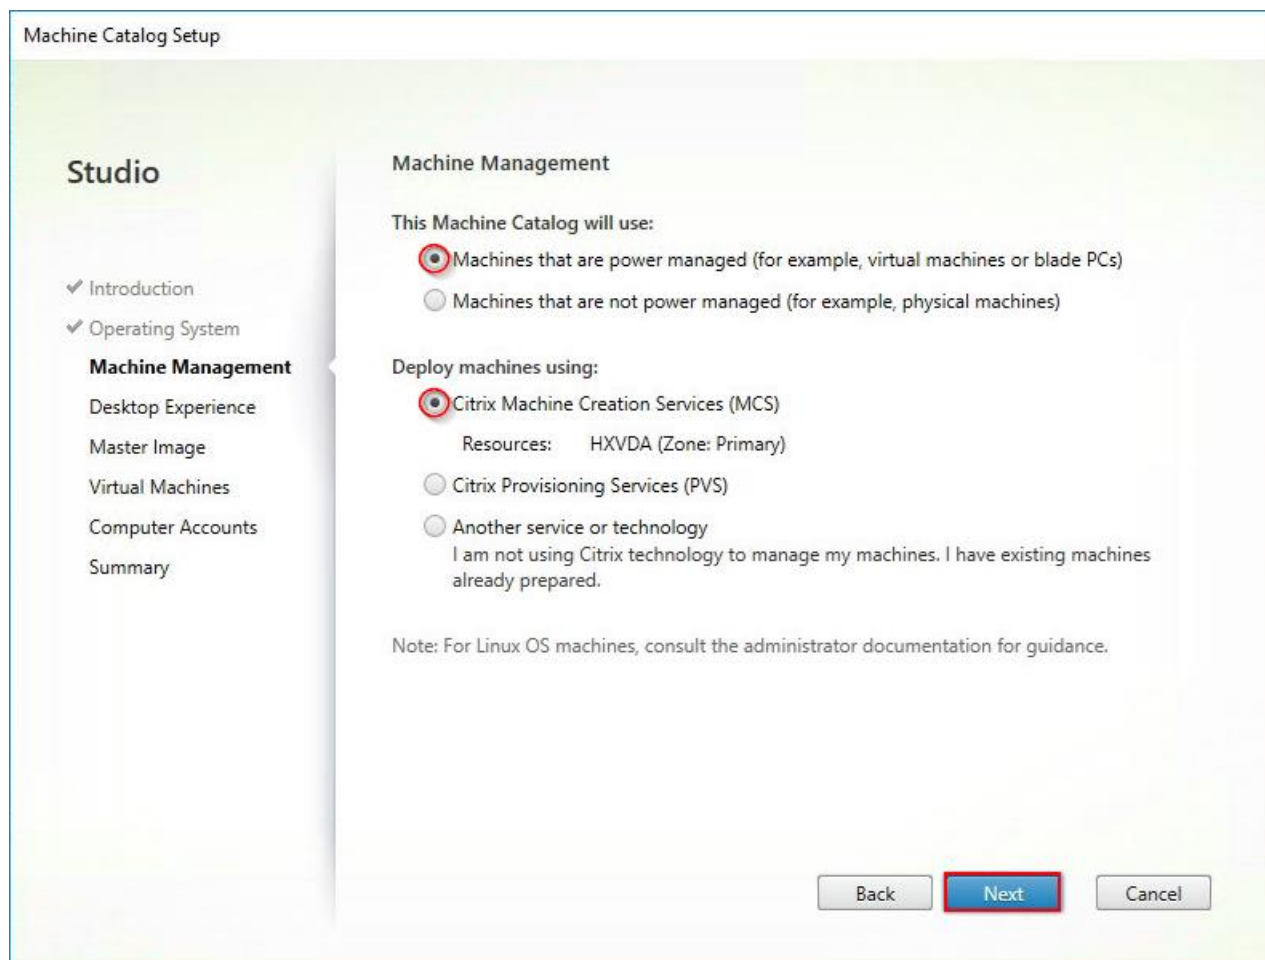
4. Select Desktop OS.

5. Click Next.



6. Select appropriate machine management.

7. Click Next.



The image shows the 'Machine Catalog Setup' wizard, specifically the 'Machine Management' step. On the left, a 'Studio' sidebar lists the steps: Introduction, Operating System, Machine Management (highlighted), Desktop Experience, Master Image, Virtual Machines, Computer Accounts, and Summary. The main area is titled 'Machine Management' and contains two sections. The first section, 'This Machine Catalog will use:', has two radio button options: 'Machines that are power managed (for example, virtual machines or blade PCs)' (which is selected and circled in red) and 'Machines that are not power managed (for example, physical machines)'. The second section, 'Deploy machines using:', has three radio button options: 'Citrix Machine Creation Services (MCS)' (selected and circled in red), 'Citrix Provisioning Services (PVS)', and 'Another service or technology'. Below the third option is a text box with the placeholder text 'I am not using Citrix technology to manage my machines. I have existing machines already prepared.' Below these sections is a note: 'Note: For Linux OS machines, consult the administrator documentation for guidance.' At the bottom right, there are three buttons: 'Back', 'Next' (highlighted with a red border), and 'Cancel'.

Machine Catalog Setup

**Studio**

- ✓ Introduction
- ✓ Operating System
- Machine Management**
- Desktop Experience
- Master Image
- Virtual Machines
- Computer Accounts
- Summary

**Machine Management**

This Machine Catalog will use:

- ☒ Machines that are power managed (for example, virtual machines or blade PCs)
- ☐ Machines that are not power managed (for example, physical machines)

Deploy machines using:

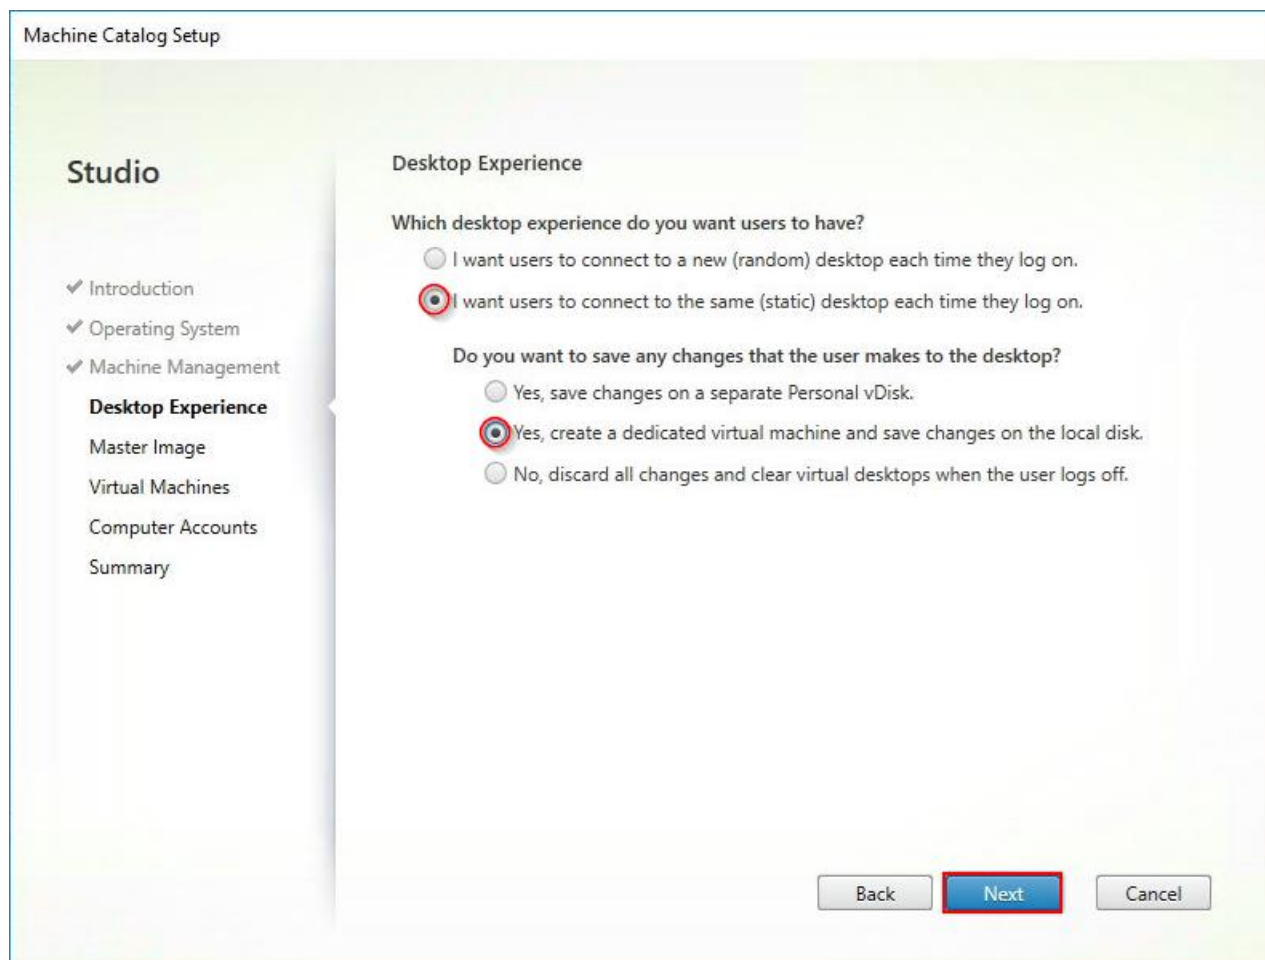
- ☒ Citrix Machine Creation Services (MCS)  
Resources: HXVDA (Zone: Primary)
- ☐ Citrix Provisioning Services (PVS)
- ☐ Another service or technology  
I am not using Citrix technology to manage my machines. I have existing machines already prepared.

Note: For Linux OS machines, consult the administrator documentation for guidance.

Back Next Cancel

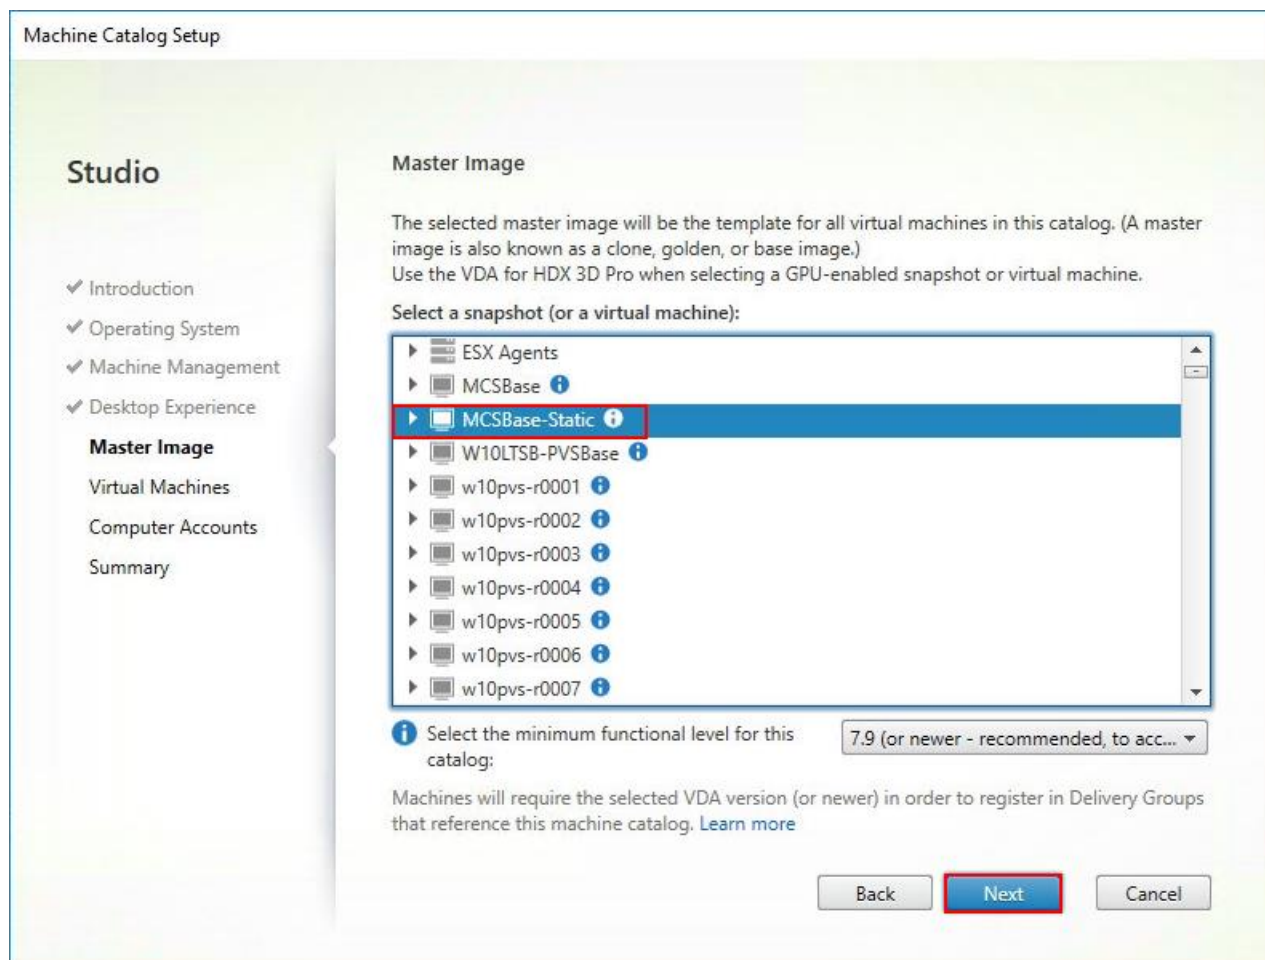
8. Select Static, Dedicated Virtual Machine for Desktop Experience.

9. Click Next.



10. Select a Virtual Machine to be used for Catalog Master image.

11. Click Next.



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

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

14. Select Full Copy for machine copy mode.

15. Click Next.



Machine Catalog Setup

## Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- Virtual Machines**
- Computer Accounts
- Summary

### Virtual Machines

How many virtual machines do you want to create?

450 - +

Configure your machines.

Total memory (MB) on each machine: 2048 - +

Select a virtual machine copy mode.

☐ Use fast clone for more efficient storage use and faster machine creation.

☒ Use full copy for better data recovery and migration support, with potentially reduced IOPS after the machines are created.

Back Next Cancel

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

17. Click Next.



Machine Catalog Setup

## Studio

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Virtual Machines
- Computer Accounts**
- Summary

### Active Directory Computer Accounts

Each machine in a Machine Catalog needs a corresponding Active Directory computer account.

Select an Active Directory account option:

☒ Create new Active Directory accounts

☐ Use existing Active Directory accounts

Active Directory location for computer accounts:

Domain:

- Computers
- Domain Controllers
- ForeignSecurityPrincipals
- InfraSrv
- LoginVSI
- Computers**

Selected location:

Account naming scheme:

w10mcs-s0123

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

**Machine Catalog Setup**

**Studio**

- ✓ Introduction
- ✓ Operating System
- ✓ Machine Management
- ✓ Desktop Experience
- ✓ Master Image
- ✓ Virtual Machines
- ✓ Computer Accounts
- Summary**

**Summary**

Machine type:	Desktop OS
Machine management:	Virtual
Provisioning method:	Machine creation services (MCS)
Desktop experience:	Users connect to the same desktop each time they log on Save changes on the local disk
Resources:	HXVDA
Master Image name:	MCSBase-Static A snapshot of the Master Image VM will be created
VDA version:	7.9 (or newer)
Number of VMs to create:	1000

Machine Catalog name:  
**Win10MCS-Static**

Machine Catalog description for administrators: (Optional)  
*Example: Windows 7 SP1 desktops for the London Sales office*

To complete the deployment, assign this Machine Catalog to a Delivery Group by selecting Delivery Groups and then Create or Edit a Delivery Group.

Back Finish Cancel

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

- Delivery Controller > Citrix Studio > Machine Catalogs
- Domain Controller > Active Directory Users and Computers

## Create Delivery Groups

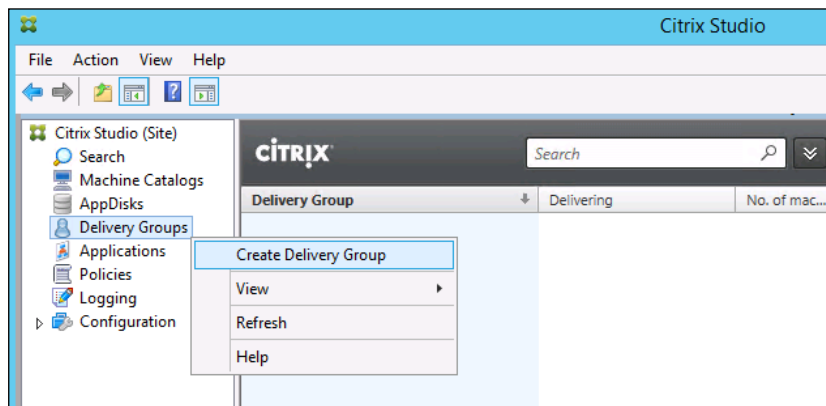
Delivery Groups are collections of machines that control access to desktops and applications. With Delivery Groups, you can specify which users and groups can access which desktops and applications.

To create delivery groups, complete the following steps:

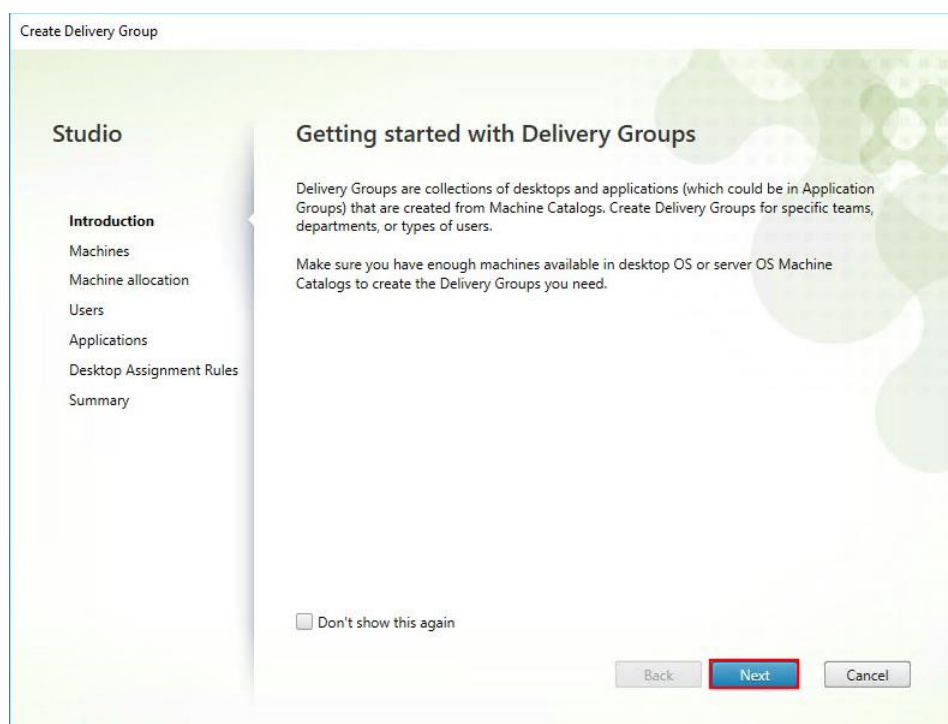


The instructions below outline the procedure to create a Delivery Group for VDI desktops. When you have completed these steps, repeat the procedure to a Delivery Group for HVD desktops.

1. Connect to a XenDesktop 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.

Create Delivery Group

**Studio**

- ✓ Introduction
- Machines**
- Machine allocation
- Users
- Applications
- Desktop Assignment Rules
- Summary

**Machines**

Select a Machine Catalog.

Catalog	Type	Machines
<input checked="" type="radio"/> Win10MCS-Random	VDI MCS Random	1000
<input type="radio"/> Win10MCS-Static	VDI MCS Static Local Disk	1000
<input type="radio"/> Win10PVS-Random	VDI PVS Random	1000
<input type="radio"/> Win2016-HSD	RDS PVS Random	72

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

Back Next Cancel

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

Create Delivery Group

**Studio**

- ✓ Introduction
- ✓ Machines
- Users**
- Applications
- Summary

**Users**

Specify who can use the applications and desktops in this Delivery Group. You can assign users and user groups who log on with valid credentials.

☒ Allow any authenticated users to use this Delivery Group.

☐ Restrict use of this Delivery Group to the following users:

Add users and groups

Add... Remove

☐ Sessions must launch in a user's home zone, if configured.

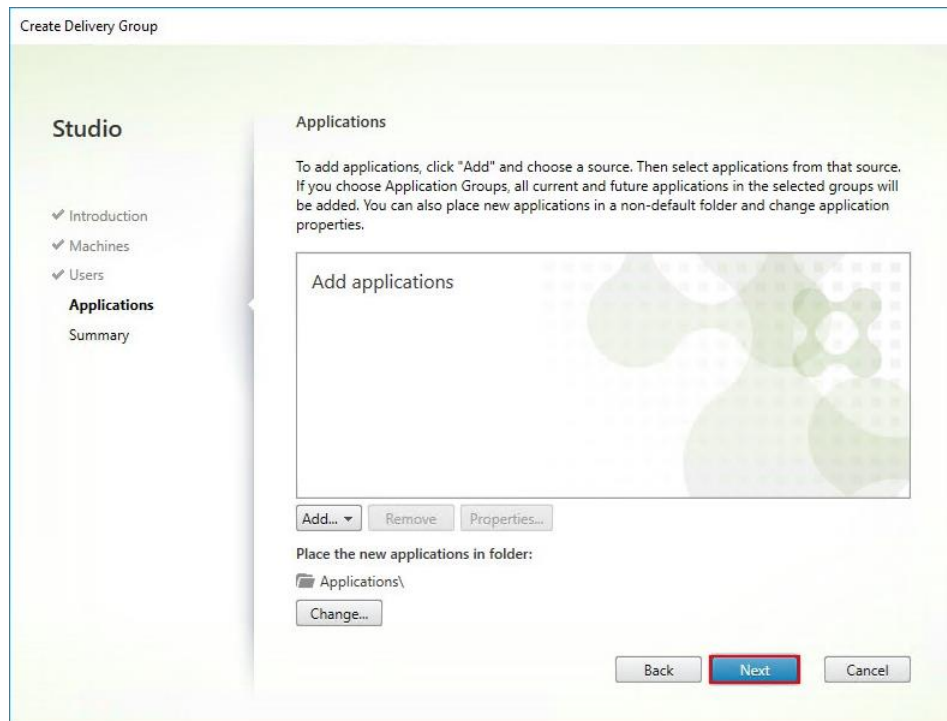
Back Next Cancel



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.

Create Delivery Group

**Studio**

- ✓ Introduction
- ✓ Machines
- ✓ Users
- ✓ Applications
- Summary**


**Summary**

Machine Catalog:	Win10MCS-Random
Machine type:	Desktop OS
Allocation type:	Random
Machines added:	VDILAB-XD\w10mcs-r0001 VDILAB-XD\w10mcs-r0002 VDILAB-XD\w10mcs-r0003 VDILAB-XD\w10mcs-r0004 VDILAB-XD\w10mcs-r0005 VDILAB-XD\w10mcs-r0006 VDILAB-XD\w10mcs-r0007 VDILAB-XD\w10mcs-r0008 VDILAB-XD\w10mcs-r0009 VDILAB-XD\w10mcs-r0010 VDILAB-XD\w10mcs-r0011 VDILAB-XD\w10mcs-r0012

Delivery Group name:

Delivery Group description, used as label in Receiver (optional):

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.”**



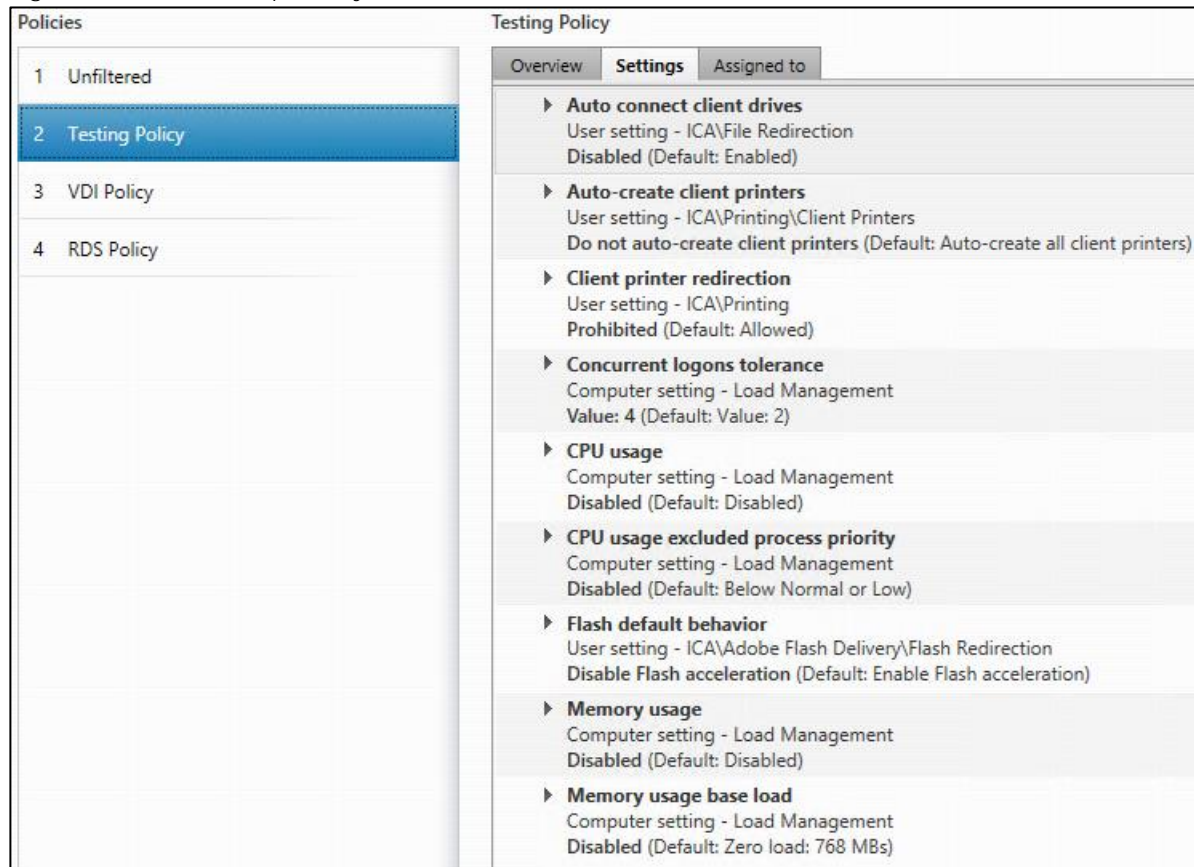
## Citrix XenDesktop Policies and Profile Management

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

## Configure Citrix XenDesktop Policies

Citrix XenDesktop 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). The screenshot below shows policies for Login VSI testing in this CVD.

Figure 41 XenDesktop Policy



## Configuring User Profile Management

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration, and provides users with fast logons and logoffs. A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Examples of settings that can be customized are:

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



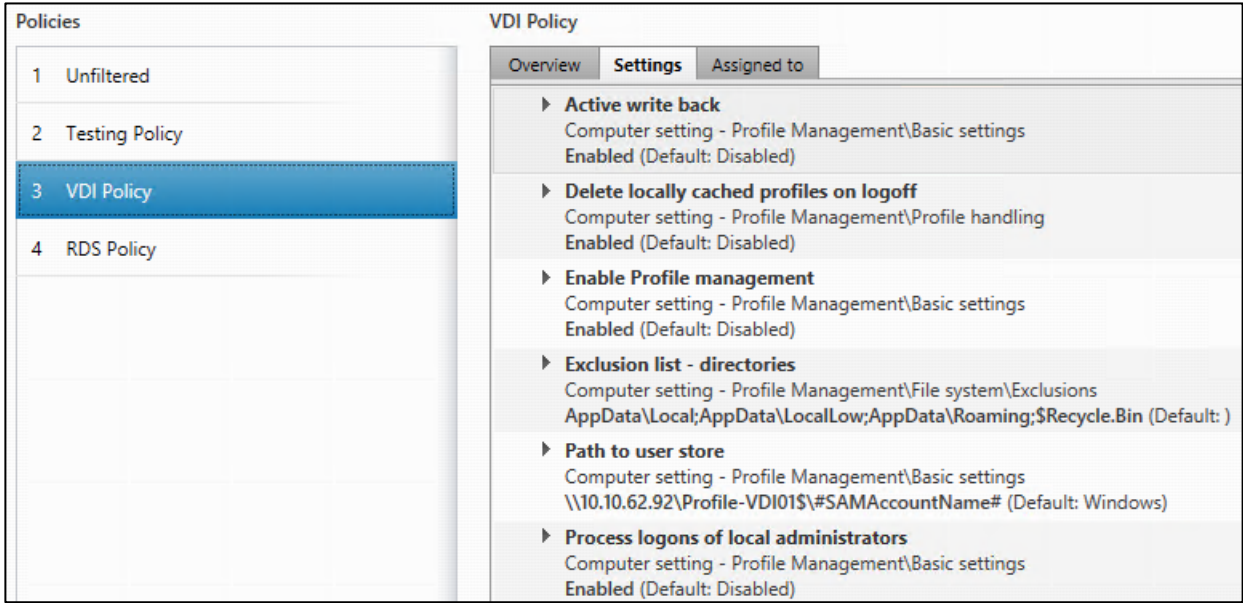
- Microsoft Outlook signature
- Printers

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

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature simplifies some of this decision-making for XenDesktop 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:

<http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-11.html>

Figure 42 VDI User Profile Manager Policy

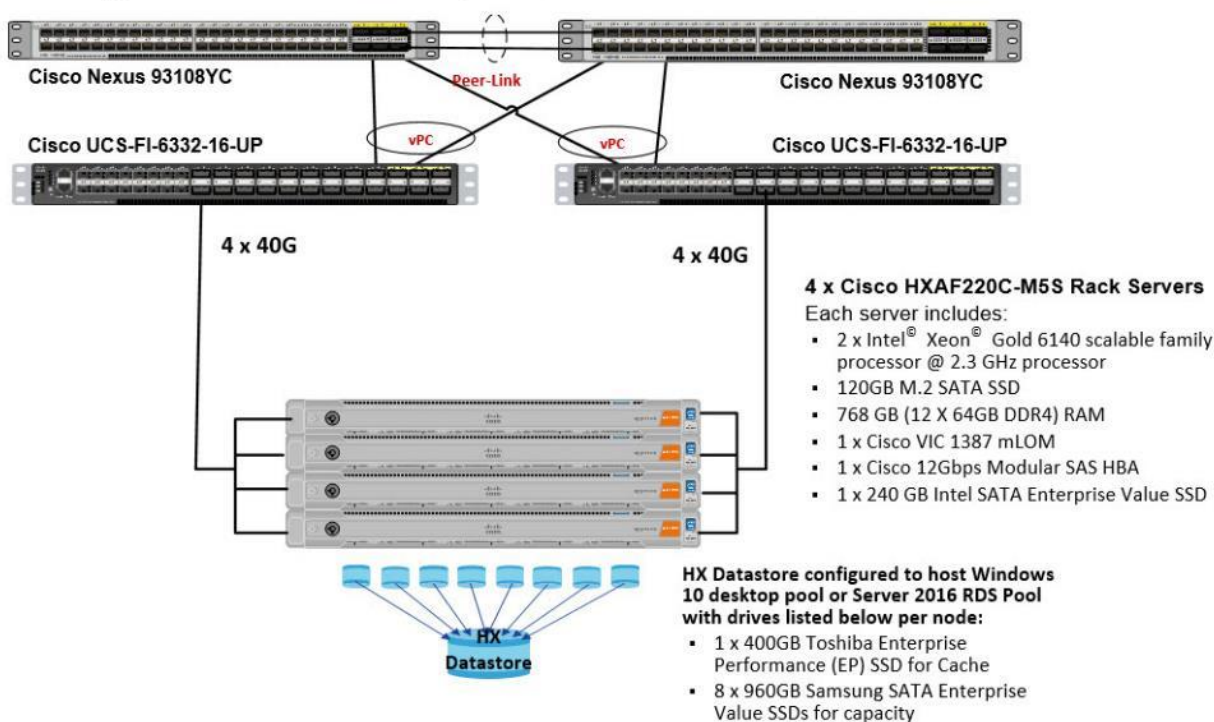




## Test Setup and Configurations

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.

**Cisco HyperFlex and Citrix XenDesktop 7.17, Reference Architecture**



Hardware Components:

- 2 x Cisco UCS 6332-16UP Fabric Interconnects
- 2 x Cisco Nexus 93108YCPX Access Switches
- 4 x Cisco UCS HXAF220c-M5SX Rack Servers (2 Intel Xeon Gold 6140 scalable family processor at 2.3 GHz, with 768 GB of memory per server [32 GB x 24 DIMMs at 2666 MHz])
- Cisco VIC 1387 mLOM
- 12G modular SAS HBA Controller
- 240GB M.2 SATA SSD drive (Boot and HyperFlex Data Platform controller VM)
- 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 3.2(3d)
- Cisco HyperFlex Data Platform 3.0.1a
- Microsoft Hyper-V 2016
- Citrix XenDesktop 7.16
- Citrix User Profile Management
- Citrix NetScaler VPX NS11.1 52.13.nc
- Microsoft SQL Server 2016
- Microsoft Windows 10
- Microsoft Windows 2016
- Microsoft Office 2016
- Login VSI 4.1.25.6

## Testing 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>.

## Testing 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 XenDesktop 7.16 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. Complete the following steps:

1. Time 0:00:00 Start esxtop Logging on the following systems:
  - Infrastructure and VDI Host Blades used in test run
  - All Infrastructure VMs used in test run (AD, SQL, View Connection brokers, image mgmt., etc.)
2. Time 0:00:10 Start Storage Partner Performance Logging on Storage System.
3. Time 0:05: Boot RDS Machines using Citrix XenDesktop 7.16 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 XenDesktop 7.16 Administrator Console dashboard. Typically a 20-30 minute rest period for Windows 10 desktops and 10 minutes for RDS VMs 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 VMs utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).
7. Time 2:23 Single Server or Scale target number of desktop VMs 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 XenDesktop 7.16 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 XenDesktop 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 XenDesktop 7.16 Hosted Shared Desktop with Citrix XenDesktop 7.16 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 amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSImax. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time it will be clear the response times escalate at saturation point.

**This VSImax is the “Virtual Session Index (VSI)”.** With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

## Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts **on every target system, and are initiated at logon within the simulated user’s desktop session context.**

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent 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 43 Sample of a VSI Max Response Time Graph, Representing a Normal Test





Figure 44 Sample of a VSI Test Response Time Graph with a Clear Performance Issue



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

The response times are very different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. This response time of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

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

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

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

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

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

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

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

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

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

VSImax v4.1.x is reached when the VSImax + 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 an individual user better performance than the 10 core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

However, the server with the slower 10 core CPU will easily have a larger capacity than the faster dual core system. This is indicated by 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 450 desktops and measure the time it takes for the 450<sup>th</sup> virtual machine to register as available in the XenDesktop Administrator console.

The Cisco HyperFlex HXAF220c-M5SX based All-Flash cluster running Data Platform version 3.0.1a software can accomplish this task in 5 minutes.

## Recommended Maximum Workload and Configuration Guidelines

### Four Node Cisco HXAF220c-M5S Rack Server and HyperFlex All-Flash Cluster

For Citrix XenApp 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 insure that end-user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.

---

 Memory should never be oversubscribed for Desktop Virtualization workloads.

---

---

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

---

Test Phase	Description
Boot	Start all RDS and/or VDI virtual machines at the same time.
Login	The Login VSI phase of test is where sessions are launched and start executing the workload over a 48 minutes duration.

Test Phase	Description
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 HXAF2240c-M5SX with Intel Xeon Gold 6140 scalable family processors and 768GB of RAM for Windows Server 2016 Hosted Sessions is 600 sessions with Office 2016 virtual desktops respectively.


#### 450 Windows 10 Citrix MCS Persistent Testing on Four Node Cisco HyperFlex Cluster

Floating assigned automated Linked-Clone desktop pool with 450 Windows 10 VMs hosting 450 User Sessions on four HXAF220c-M5SX HyperFlex cluster.

Test result highlights include:

- 0.676 second baseline response time
- 0.839 second average response time with 450 desktops running
- Average CPU utilization of 45 percent during steady state
- Average of 320 GB of RAM used out of 768 GB available
- 1000Mbps peak network utilization per host.
- Average Read Latency 0.7ms/Max Read Latency 1.4ms
- Average Write Latency 1.9ms/Max Write Latency 4.4ms
- 3500 peak I/O operations per second (IOPS) per cluster at steady state
- 80MBps peak throughput per cluster at steady state

Figure 45 Login VSI Analyzer Chart for 450 Windows 10 Citrix MCS Persistent Virtual Desktops



450-02RC1

Successfully completed Login VSI test with **450 knowledgeworker** sessions. VSImax (system saturation) was not reached. All Login VSI users completed the test.



PASS

LoginVSI Results

VSI Baseline: 676ms  
VSI Average: 832ms  
VSI Threshold: 1676ms  
Stuck Sessions: 0

Test result review

450 sessions were configured to be launched in 2880 seconds.

In total 0 sessions failed during the test:

- 0 sessions was/were not successfully launched
- 0 launched sessions failed to become active
- 450 sessions were active during the test
- 0 sessions got stuck during the test (before VSI max threshold)

With 450 sessions the maximum capacity VSI max (v4.1) knowledgeworker was not reached with a Login VSI baseline performance score of 676

Login VSI index average score is 693 lower than threshold. It might be possible to launch more sessions in this configuration.

Baseline performance of 676 is: **Very good**

VSI Baseline	Performance
0-799	Very Good
800-1299	Good
1200-1999	Reasonable
2000-9999	Bad

Figure 46 Three Consecutive Login VSI Analyzer Chart for 450 Windows 10 Citrix MCS Persistent Virtual Desktops




Figure 47 Sample 4x Hyper-V Hosts CPU Core Utilization Running 450 Windows 10 Citrix MCS Persistent Virtual Desktops


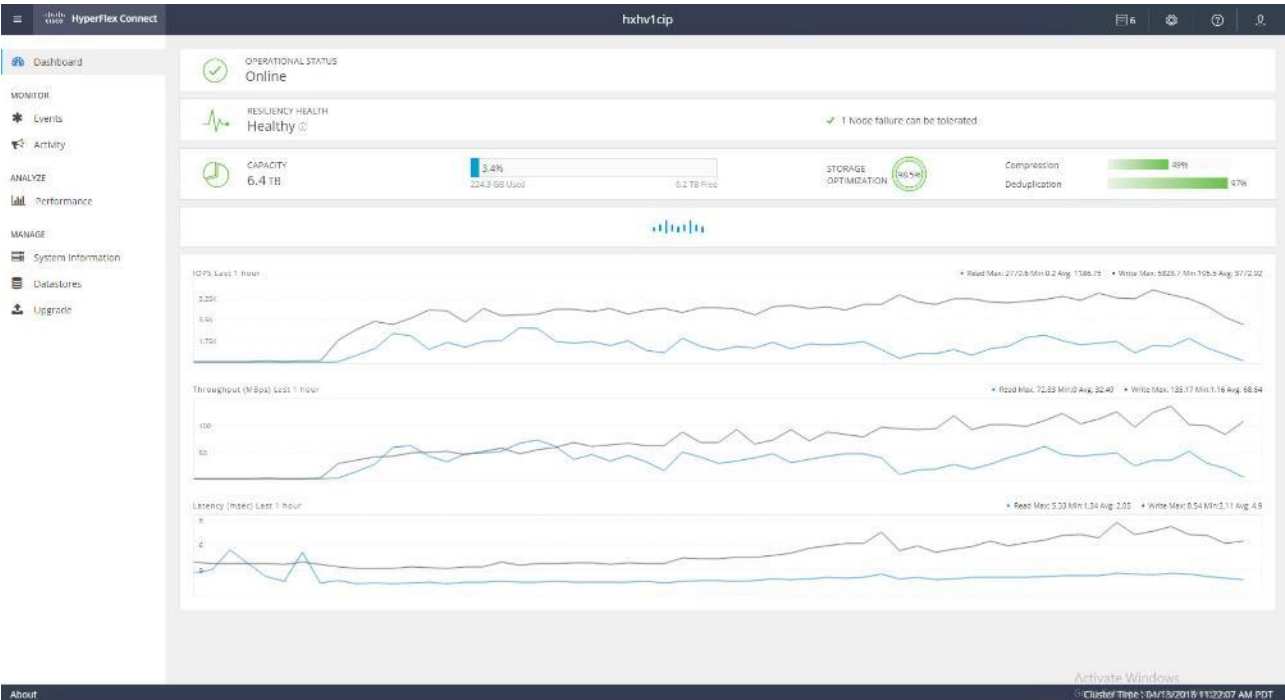


Figure 48 HyperFlex Cluster UI Performance Chart for Knowledge Worker Workload Running 450 User Test on Citrix MCS Persistent Windows 10



## 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 XenDesktop 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 XenDesktop.

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 Authors

---

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

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

## Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge their contribution and expertise that resulted in developing this document:

- Mike Brennan, Product Manager, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

## Appendix A – Cisco Nexus 93108YC Switch Configuration

---

### Switch A Configuration

!Command: show running-config

```
version 7.0(3)I2(2d)

switchname XXXXXXXXXXXX

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 XXXXXXXXXXXX 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$Bo0lrVnESUVxLcbRHg86j1 role network-admin
ip domain-lookup
no service unsupported-transceiver
class-map type qos match-all class-fcoe
policy-map type qos jumbo
  class class-default
    set qos-group 0
copp profile strict
snmp-server user admin network-admin auth md5 0x71d6a9cf1ea007cd3166e91a6f3807e5
priv 0x71d6a9cf1ea007cd3166e91a6f3807e5 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.10.50.2
ntp peer 10.10.50.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8

vlan 1,50-54
```

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

interface Vlan50

no shutdown

```
ip address 10.10.50.2/24
```

```
hsrp version 2
```

```
hsrp 50
```

```
preempt
```

```
priority 110
```

```
ip 10.10.50.1
```

```
ip dhcp relay address 10.10.51.21
```

```
ip dhcp relay address 10.10.51.22
```

```
interface Vlan51
```

```
no shutdown
```

```
ip address 10.10.51.2/24
```

```
hsrp version 2
```

```
hsrp 51
```

```
preempt
```

```
priority 110
```

```
ip 10.10.51.1
```

```
interface Vlan52
```

```
no shutdown
```

```
ip address 10.10.52.2/24
```

```
hsrp version 2
```

```
hsrp 52
```

```
preempt
```

```
priority 110
```

```
ip 10.10.52.1
```

```
interface Vlan53
```

```
no shutdown
```

```
ip address 10.10.53.2/24
```

```
hsrp version 2
```

```
hsrp 53
```

```
preempt
```

```
priority 110
```

```
ip 10.10.53.1
```

```
interface Vlan54
```

```
no shutdown
```

```
ip address 10.54.0.2/20
```

```
hsrp version 2
```

```
hsrp 54
```

```
preempt
```

```
priority 110
```

```
ip 10.54.0.1
```

```
ip dhcp relay address 10.10.51.21
```

```
ip dhcp relay address 10.10.51.22
```

```
interface port-channel10
```

```
description vPC-PeerLink
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
spanning-tree port type network
```

```
service-policy type qos input jumbo
```

```
vpc peer-link
```

```
interface port-channel11
  description FI-Uplink-K22
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 11
```

```
interface port-channel12
  description FI-Uplink-K22
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  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 1,50-54
  channel-group 10 mode active
```

```
interface Ethernet1/2
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  channel-group 10 mode active
```

```
interface Ethernet1/3
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  channel-group 10 mode active
```

```
interface Ethernet1/4
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  channel-group 10 mode active
```

```
interface Ethernet1/5
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  mtu 9216
  channel-group 11 mode active
```

```
interface Ethernet1/6
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  mtu 9216
  channel-group 11 mode active
```

```
interface Ethernet1/7
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  mtu 9216
  channel-group 12 mode active
```

```
interface Ethernet1/8
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
mtu 9216
```

```
channel-group 12 mode active
```

```
interface Ethernet1/9
```

```
interface Ethernet1/10
```

```
interface Ethernet1/11
```

```
interface Ethernet1/12
```

```
interface Ethernet1/13
```

```
interface Ethernet1/14
```

```
interface Ethernet1/15
```

```
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 1,50-54
```

```
spanning-tree port type edge trunk
```

```
interface Ethernet1/26
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
spanning-tree port type edge trunk
```

```
interface Ethernet1/27
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
spanning-tree port type edge trunk
```

```
interface Ethernet1/28
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
spanning-tree port type edge trunk
```



```
interface Ethernet1/29
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  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 XXXXXXXXXXXX
```

```
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 XXXXXXXXXXXX id 1
```

```
limit-resource vlan minimum 16 maximum 4094
limit-resource vrf minimum 2 maximum 4096
limit-resource port-channel minimum 0 maximum 511
limit-resource u4route-mem minimum 248 maximum 248
limit-resource u6route-mem minimum 96 maximum 96
limit-resource m4route-mem minimum 58 maximum 58
limit-resource m6route-mem minimum 8 maximum 8
```

```
feature telnet
cfs eth distribute
```

```
feature interface-vlan
```

```
feature hsrp
```

```
feature lacp
```

```
feature dhcp
```

```
feature vpc
```

```
feature lldp
```

```
clock protocol ntp vdc 1
```

```
no password strength-check
```

```
username admin password 5 $1$jEwHqUvM$gpOec2hramkyX09KD3/Dn. role network-admin
```

```
ip domain-lookup
```

```
no service unsupported-transceiver
```

```
class-map type qos match-all class-fcoe
```

```
policy-map type qos jumbo
```

```
class class-default
```

```
set qos-group 0
```

```
copp profile strict
```

```
snmp-server user admin network-admin auth md5 0x9046c100ce1f4ecdd74ef2f92c4e83f9
```

```
priv 0x9046c100ce1f4ecdd74ef2f92c4e83f9 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp peer 10.10.50.2
ntp server 10.10.50.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8

vlan 1,50-54
vlan 50
    name InBand-Mgmt-C1
vlan 51
    name Infra-Mgmt-C1
vlan 52
    name StorageIP-C1
vlan 53
    name 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 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 Vlan50
  no shutdown
  ip address 10.10.50.3/24
  hsrp version 2
  hsrp 50
    preempt
    priority 110
    ip 10.10.50.1
  ip dhcp relay address 10.10.51.21
  ip dhcp relay address 10.10.51.22
```

```
interface Vlan51
  no shutdown
  ip address 10.10.51.3/24
  hsrp version 2
  hsrp 51
```

```
preempt
priority 110
ip 10.10.51.1
```

```
interface Vlan52
no shutdown
ip address 10.10.52.3/24
hsrp version 2
hsrp 52
preempt
priority 110
ip 10.10.52.1
```

```
interface Vlan53
no shutdown
ip address 10.10.53.3/24
hsrp version 2
hsrp 53
preempt
priority 110
ip 10.10.53.1
```

```
interface Vlan54
no shutdown
ip address 10.54.0.3/20
hsrp version 2
hsrp 54
preempt
priority 110
```

```
ip 10.54.0.1
```

```
ip dhcp relay address 10.10.51.21
```

```
ip dhcp relay address 10.10.51.22
```

```
interface port-channel10
```

```
description vPC-PeerLink
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
spanning-tree port type network
```

```
service-policy type qos input jumbo
```

```
vpc peer-link
```

```
interface port-channel11
```

```
description FI-Uplink-K22
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
spanning-tree port type edge trunk
```

```
mtu 9216
```

```
service-policy type qos input jumbo
```

```
vpc 11
```

```
interface port-channel12
```

```
description FI-Uplink-K22
```

```
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-54
```

```
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 1,50-54

channel-group 10 mode active

interface Ethernet1/2

switchport mode trunk

switchport trunk allowed vlan 1,50-54

channel-group 10 mode active

interface Ethernet1/3

switchport mode trunk

switchport trunk allowed vlan 1,50-54

channel-group 10 mode active

interface Ethernet1/4

switchport mode trunk

switchport trunk allowed vlan 1,50-54

channel-group 10 mode active

interface Ethernet1/5

switchport mode trunk

switchport trunk allowed vlan 1,50-54

mtu 9216

channel-group 11 mode active

interface Ethernet1/6

```
switchport mode trunk
switchport trunk allowed vlan 1,50-54
mtu 9216
channel-group 11 mode active
```

```
interface Ethernet1/7
switchport mode trunk
switchport trunk allowed vlan 1,50-54
mtu 9216
channel-group 12 mode active
```

```
interface Ethernet1/8
switchport mode trunk
switchport trunk allowed vlan 1,50-54
mtu 9216
channel-group 12 mode active
```

```
interface Ethernet1/9
```

```
interface Ethernet1/10
```

```
interface Ethernet1/11
```

```
interface Ethernet1/12
```

```
interface Ethernet1/13
```

```
interface Ethernet1/14
```

```
interface Ethernet1/15
```

```
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 1,50-54
```

```
    spanning-tree port type edge trunk
```

```
interface Ethernet1/26
```

```
    switchport mode trunk
```

```
    switchport trunk allowed vlan 1,50-54
```

```
    spanning-tree port type edge trunk
```

```
interface Ethernet1/27
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/28
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54

  spanning-tree port type edge trunk
```

```
interface Ethernet1/29
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/30
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/31
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
```

```
interface Ethernet1/32
  switchport mode trunk
```

switchport trunk allowed vlan 1,50-54

spanning-tree port type edge trunk

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44

interface Ethernet1/45

```
interface Ethernet1/46
```

```
interface Ethernet1/47
```

```
interface Ethernet1/48
```

```
    switchport access vlan 50
```

```
interface Ethernet1/49
```

```
interface Ethernet1/50
```

```
interface Ethernet1/51
```

```
interface Ethernet1/52
```

```
interface Ethernet1/53
```

```
interface Ethernet1/54
```

```
interface mgmt0
```

```
    vrf member management
```

```
    ip address 10.29.132.20/24
```

```
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
```