

# Cisco All-Flash HyperFlex 3.5 Hyperconverged System with up to 4400 VMware Horizon 7 Users

Design and Deployment of Cisco HyperFlex for Virtual  
Desktop Infrastructure with VMware Horizon 7.6

Published: December 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, visit:

<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 .....	7
Solution Overview .....	8
Introduction.....	8
Audience .....	8
Purpose of this Document.....	8
What's New in this Release? .....	8
Enhancements for Version 3.5 .....	8
Documentation Roadmap .....	10
All-Flash Versus Hybrid .....	12
Cisco HyperFlex Compute-Only Nodes .....	13
Cisco HyperFlex Data Platform Software.....	14
Cisco Desktop Virtualization Solutions: Data Center .....	15
The Evolving Workplace .....	15
Cisco Desktop Virtualization Focus.....	16
Use Cases .....	17
Physical Topology.....	19
Fabric Interconnects.....	20
HX-Series and C-Series Rack-Mount Servers .....	21
Cisco UCS B-Series Blade Servers .....	21
Logical Network Design.....	21
Logical Availability Zones.....	23
Configuration Guidelines.....	25
Solution Design .....	26
Cisco Unified Computing System .....	26
Cisco Unified Computing System Components .....	26
Enhancements for Version 3.5.1a .....	28
Cisco UCS Fabric Interconnect.....	30
Cisco HyperFlex HX-Series Nodes .....	30
Cisco VIC 1387 MLOM Interface Card.....	33
Cisco HyperFlex Compute Nodes .....	33
Cisco UCS B200-M5 Blade Server .....	33
Cisco VIC1340 Converged Network Adapter .....	34
Cisco UCS 5108 Blade Chassis.....	34
Features and Benefits .....	35
Cisco UCS 2304XP Fabric Extender.....	35
Cisco UCS C220-M5 Rack Server .....	36

Cisco HyperFlex HX Data Platform Administration Plug-in .....	36
Cisco HyperFlex Connect HTML5 Management Web Page .....	37
Replication Factor .....	39
Data Distribution.....	39
Data Operations.....	40
Data Optimization.....	41
Data Deduplication.....	41
Inline Compression.....	42
Log-Structured Distributed Objects.....	42
Encryption.....	43
Data Services .....	43
Thin Provisioning.....	43
Snapshots .....	43
Fast, Space-Efficient Clones .....	44
Data Replication and Availability.....	44
Data Rebalancing.....	44
Online Upgrades.....	45
Cisco Nexus 93180 Switches.....	45
VMware vSphere 6.5.....	46
VMware vCenter Server.....	46
VMware ESXi 6.5 Hypervisor .....	47
VMware Horizon .....	47
Advantages of Using VMware Horizon .....	47
What are VMware RDS Hosted Sessions?.....	51
Farms, RDS Hosts, Desktop and Application Pools .....	52
Architecture and Design of VMware Horizon on Cisco Unified Computing System and Cisco HyperFlex System Design	
Fundamentals .....	53
Understanding Applications and Data.....	54
Project Planning and Solution Sizing Sample Questions .....	55
Desktop Virtualization Design Fundamentals.....	56
VMware Horizon Design Fundamentals .....	56
Horizon VDI Pool and RDSH Servers Pool.....	56
Designing a VMware Horizon Environment for Various Workload Types .....	58
Deployment Hardware and Software .....	61
Products Deployed .....	61
Hardware Deployed .....	62
Software Deployed.....	63
Logical Architecture .....	63



VLANs .....	64
Jumbo Frames .....	64
VMware Clusters .....	65
ESXi Host Design .....	65
Solution Configuration .....	70
Cisco UCS Compute Platform .....	70
Physical Infrastructure .....	70
Cisco Unified Computing System Configuration .....	75
Deploy and Configure HyperFlex Data Platform .....	76
Prerequisites .....	76
Deploy Cisco HyperFlex Data Platform Installer VM .....	80
Cisco HyperFlex Cluster Configuration .....	85
Cisco HyperFlex Cluster Expansion .....	102
Building the Virtual Machines and Environment for Workload Testing .....	128
Horizon 7 Infrastructure Components Installation .....	128
Install VMware Horizon Composer Server .....	128
Install Horizon Connection/Replica Servers .....	134
Create a Microsoft Management Console Certificate Request .....	138
Configure the Horizon 7 Environment .....	138
Configure Event Database .....	138
Configure Horizon 7 Licenses .....	139
Configure vCenter .....	140
Configure Instant Clone Domain Admins .....	144
Horizon Persona Manager Installation .....	144
Master Image Creation for Tested Horizon Deployment Types .....	147
Prepare Microsoft Windows 10 and Server 2016 R2 with Microsoft Office 2016 .....	148
Optimization of Base Windows 10 or Server 2016 Guest OS .....	148
Virtual Desktop Agent Software Installation for Horizon .....	149
Install Additional Software .....	154
Create a Native Snapshot for Automated Desktop Pool Creation .....	154
Create Customization Specification for Virtual Desktops .....	155
RDSH Farm Creation .....	160
Create the Horizon 7 RDS Published Desktop Pool .....	164
VMware Horizon Linked-Clone Windows 10 Desktop Pool Creation .....	168
VMware Horizon Instant-Clone Windows 10 Desktop Pool Creation .....	173
VMware Horizon Persistent Windows 10 Desktop Pool Creation .....	178
Test Setup and Configurations .....	185
Testing Methodology and Success Criteria .....	187

Testing Procedure .....	187
Pre-Test Setup for Testing .....	187
Test Run Protocol.....	187
Success Criteria.....	188
Test Results .....	194
Boot Storms .....	194
Recommended Maximum Workload and Configuration Guidelines .....	195
Sixteen Node Cisco HXAF220c-M5S Rack Server, Eight Node Cisco UCS C220 M5 and Eight Node Cisco UCS B200 M5 HyperFlex Cluster.....	195
4400 User Full-Scale Testing on Thirty Two-Node Cisco HyperFlex Cluster .....	195
Summary .....	200
About the Authors.....	201
Acknowledgements .....	201
Appendix A – Cisco Nexus 93180 Switch Configuration .....	202
Switch A Configuration .....	202
Switch B Configuration .....	211
Appendix B – Cisco HyperFlex HXAF220c-M5, Cisco UCS C220 M5 and Cisco UCS B200 M5 32-Node Hyperflex Horizon 7 Cluster Deployed 4400 Scale Test: In-Flight Performance Metrics.....	221

## Executive Summary

---

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

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

This document provides an architectural reference and design guide for up to a 4400 user mixed workload on a 32-node (16 Cisco HyperFlex HXAF220C-M5SX servers plus 8 Cisco UCS C220 M5 Rack Servers and 8 B200 M5 blade servers) Cisco HyperFlex system. We provide deployment guidance and performance data for VMware Horizon 7.6 virtual desktop sessions running Windows Server 2016 RDSH server-based Remote Desktop Server sessions and Microsoft Windows 10 with Office 2016, highlighting provisioning via Instant-Clone, Linked-Clone and Persistent virtual desktops on vSphere 6.5.

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

The solution payload is 100 percent virtualized on Cisco HyperFlex HXAF220c-M5SX hyperconverged nodes, Cisco UCS C220 M5 rack servers and Cisco UCS B200 M5 blade servers booting via on-board Flex-Flash controller SD cards running VMware vSphere 6.5 U2 hypervisor. The virtual desktops are configured with VMware Horizon 7.6 which incorporates both traditional persistent and non-persistent virtual Windows 7/8/10 desktops, hosted applications and remote desktop service (RDS) server 2008 R2, server 2012 R2 or server 2016 based desktops. The solution provides unparalleled scale and management simplicity. VMware Horizon RDSH server based desktop sessions (1550,) instant-clone floating assignment Windows 10 desktops (950), Composer-based linked-clone floating assignment Windows 10 desktops (950) and full clone desktops (950) are provisioned on a thirty-two node Cisco HyperFlex cluster. Where applicable, this document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

The solution boots 4400 RDSH virtual server and virtual desktops machines in 15 minutes or less, insuring that users will not experience delays in accessing their virtual workspace on HyperFlex.

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

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

# Solution Overview

---

## Introduction

A current industry trend in data center design is towards small, granularly expandable hyperconverged infrastructures. By using virtualization along with pre-validated IT platforms, customers of all sizes have embarked on the journey to “just in time capacity” using this new technology. The Cisco Hyper Converged Solution can be quickly deployed, thereby increasing agility and reducing costs. Cisco HyperFlex uses best of breed storage, server and network components to serve as the foundation for desktop virtualization workloads, enabling efficient architectural designs that can be quickly and confidently deployed and scaled out.

## Audience

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

## Purpose of this Document

This document provides a systematic design, configuration, and implementation guide for the Cisco Validated Design for a Cisco HyperFlex All-Flash system running four different VMware Horizon 7 workloads with Cisco UCS 6300 series Fabric Interconnects and Cisco Nexus 93000 series networking switches.

## What’s New in this Release?

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

- Validation of Cisco Nexus 93000 YC series with Cisco HyperFlex Support
- Support for the Cisco UCS 4.0(1b) release and Cisco HyperFlex v 3.5(1a)
- VMware vSphere 6.5 U2 Hypervisor
- VMware Horizon 7.6 Remote Desktop Sever Hosted Sessions
- VMware Horizon 7.6 Horizon Instant Clones, Linked Clones and Persistent Desktops

## Enhancements for Version 3.5

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

- Native Disaster Recovery Enhancements— A simple to use Planned Migration workflow for Disaster Recovery, VM migration, and resuming replication. In addition, native support for Replication, Planned Migration and Disaster Recovery in Stretched Cluster deployments. For more information, see

[Cisco HyperFlex Systems Administration Guide, Release 3.5](#)

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

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

## Documentation Roadmap

For the comprehensive documentation suite, refer to the [Cisco HyperFlex Systems Documentation Roadmap](#).




---

### A login is required for the Documentation Roadmap.

---

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

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

The data center market segment is shifting toward heavily virtualized private, hybrid and public cloud computing models running on industry-standard systems. These environments require uniform design points that can be repeated for ease 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 Horizon Microsoft Windows 10 virtual desktops and Horizon RDSH 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®), VMware Horizon and VMware vSphere software in a single package. The design is efficient such that the networking, computing, and storage components occupy 18-rack units footprint in an industry standard 42U rack. Port density on the Cisco Nexus switches and Cisco UCS Fabric Interconnects enables the networking components to accommodate multiple HyperFlex clusters in a single Cisco UCS domain.

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

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

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

- More power, same size. Cisco HX-series nodes, dual 18-core 2.3 GHz Intel Xeon (Gold 6140) Scalable Family processors with 768GB of 2666Mhz memory with VMware Horizon 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 4400 user mixed hosted virtual desktop and hosted shared desktop environment booted and registered with the Horizon 7 in under 15 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.
- Stress-tested to the limits during simulated login storms. All 4400 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 4400 user system is 26 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects. Incremental seat Cisco HyperFlex clusters can be added one at a time to a total of 64 nodes.
- 100 percent virtualized This CVD presents a validated design that is 100 percent virtualized on VMware ESXi 6.5. All of the virtual desktops, user data, profiles, and supporting infrastructure components, including Active Directory, SQL Servers, VMware Horizon components, Horizon VDI desktops and RDSH 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 data center management: Cisco maintains industry leadership with the new Cisco UCS Manager 4.0(1b) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director insure that

customer environments are consistent locally, across Cisco UCS Domains and across the globe. Cisco UCS software suite offers increasingly simplified operational and deployment management, and it continues to widen the span of control for customer organizations' subject matter experts in compute, storage and network.

- Cisco 40G Fabric: Our 40G unified fabric story gets additional validation on 6300 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- Cisco HyperFlex Connect (HX Connect): An all-new HTML 5 based Web UI Introduced with HyperFlex v2.5 or later is available for use as the primary management tool for Cisco HyperFlex. Through this centralized point of control for the cluster, administrators can create volumes, monitor the data platform health, and manage resource use. Administrators can also use this data to predict when the cluster will need to be scaled.
- Cisco HyperFlex storage performance: Cisco HyperFlex provides industry-leading hyper converged storage performance that efficiently handles the most demanding I/O bursts (for example, login storms), high write throughput at low latency, delivers simple and flexible business continuity and helps reduce storage cost per desktop.
- Cisco HyperFlex agility: Cisco HyperFlex System enables users to seamlessly add, upgrade or remove storage from the infrastructure to meet the needs of the virtual desktops.
- Cisco HyperFlex vCenter integration: Cisco HyperFlex plugin for VMware vSphere provides easy-button automation for key storage tasks such as storage provisioning and storage resize, cluster health status and performance monitoring directly from the vCenter web client in a single pane of glass. Experienced vCenter administrators have a near zero learning curve when HyperFlex is introduced into the environment.
- VMware Horizon 7 advantage: VMware Horizon 7 follows a new unified product architecture that supports both hosted-shared desktops and applications (RDS) and complete virtual desktops (VDI). This new Horizon release simplifies tasks associated with large-scale VDI management. This modular solution supports seamless delivery of Windows apps and desktops as the number of user increase. In addition, PCoIP and Blast extreme enhancements help to optimize performance and improve the user experience across a variety of endpoint device types, from workstations to mobile devices including laptops, tablets, and smartphones.
- Optimized for performance and scale. For hosted shared desktop sessions, the best performance was achieved when the number of vCPUs assigned to the Horizon 7 RDSH virtual machines did not exceed the number of hyper-threaded (logical) cores available on the server. In other words, maximum performance is obtained when not overcommitting the CPU resources for the virtual machines running virtualized RDS systems.
- Provisioning desktop machines made easy: VMware Horizon 7 provisions hosted virtual desktops as well as hosted shared desktop virtual machines for this solution using a single method for both, the "Automated floating assignment desktop pool." "Dedicated user assigned desktop pool" for persistent desktops was provisioned in the same Horizon 7 administrative console. Horizon 7 introduces a new provisioning technique for non-persistent virtual desktops called "Instant-clone." The new method greatly reduces the amount of life-cycle spend and the maintenance windows for the guest OS.

## All-Flash Versus Hybrid

The initial HyperFlex product release featured hybrid converged nodes, which use a combination of solid-state disks (SSDs) for the short-term storage caching layer, and hard disk drives (HDDs) for the long-term storage

capacity layer. The hybrid HyperFlex system is an excellent choice for entry-level or midrange storage solutions, and hybrid solutions have been successfully deployed in many non-performance sensitive virtual environments. Meanwhile, there is significant growth in deployment of highly performance sensitive and mission critical applications. The primary challenge to the hybrid HyperFlex system from these highly performance sensitive applications, is their increased sensitivity to high storage latency. Due to the characteristics of the spinning hard disks, it is unavoidable that their higher latency becomes the bottleneck in the hybrid system. Ideally, if all of the storage operations were to occur in the caching SSD layer, the hybrid system's performance will be excellent. However, in several scenarios, the amount of data being written and read exceeds the caching layer capacity, placing larger loads on the HDD capacity layer, and the subsequent increases in latency will naturally result in reduced performance.

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

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

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

## Cisco HyperFlex Compute-Only Nodes

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

- Cisco UCS B200 M3 Blade Server

- Cisco UCS B200 M4 Blade Server
- Cisco UCS B200 M5 Blade Server
- Cisco UCS B260 M4 Blade Server
- Cisco UCS B420 M4 Blade Server
- Cisco UCS B460 M4 Blade Server
- Cisco UCS B480 M5 Blade Server
- Cisco UCS C220 M3 Rack-Mount Servers
- Cisco UCS C220 M4 Rack-Mount Servers
- Cisco UCS C220 M5 Rack-Mount Servers
- Cisco UCS C240 M3 Rack-Mount Servers
- Cisco UCS C240 M4 Rack-Mount Servers
- Cisco UCS C240 M5 Rack-Mount Servers
- Cisco UCS C460 M4 Rack-Mount Servers
- Cisco UCS C480 M5 Rack-Mount Servers

## Cisco HyperFlex Data Platform Software

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

- Data protection creates multiple copies of the data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- Stretched clusters allow nodes to be evenly split between two physical locations, keeping a duplicate copy of all data in both locations, thereby providing protection in case of an entire site failure.
- Logical availability zones provide multiple logical grouping of nodes and distributes the data across these groups in such a way that no single group has more than one copy of the data. This enables enhanced protection from node failures, allowing for more nodes to fail while the overall cluster remains online.
- Deduplication is always on, helping reduce storage requirements in virtualization clusters in which multiple operating system instances in guest virtual machines result in large amounts of replicated data.
- Compression further reduces storage requirements, reducing costs, and the log-structured file system is designed to store variable-sized blocks, reducing internal fragmentation.
- Replication copies virtual machine level snapshots from one Cisco HyperFlex cluster to another, to facilitate recovery from a cluster or site failure, via a failover to the secondary site of all VMs.

- Encryption stores all data on the caching and capacity disks in an encrypted format, to prevent accidental data loss or data theft. Key management can be done using local Cisco UCS Manager managed keys, or third-party Key Management Systems (KMS) via the Key Management Interoperability Protocol (KMIP).
- Thin provisioning allows large volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a “pay as you grow” proposition.
- Fast, space-efficient clones rapidly duplicate virtual storage volumes so that virtual machines can be cloned simply through metadata operations, with actual data copied only for write operations.
- Snapshots help facilitate backup and remote-replication operations, which are needed in enterprises that require always-on data availability.

## Cisco Desktop Virtualization Solutions: Data Center

### The Evolving Workplace

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

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

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

**Figure 1 Cisco Data Center Partner Collaboration**



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

## Cisco Desktop Virtualization Focus

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

### Simplified

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

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

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

### Secure

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

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

### Scalable

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

Cisco HyperFlex servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability to up to 3.0



terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco HyperFlex helps maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs based on VMware Horizon, Cisco HyperFlex solutions have demonstrated scalability and performance, with up to 4400 hosted virtual desktops and hosted shared desktops up and running in ~10 minutes.

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

### Savings and Success

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

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

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

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

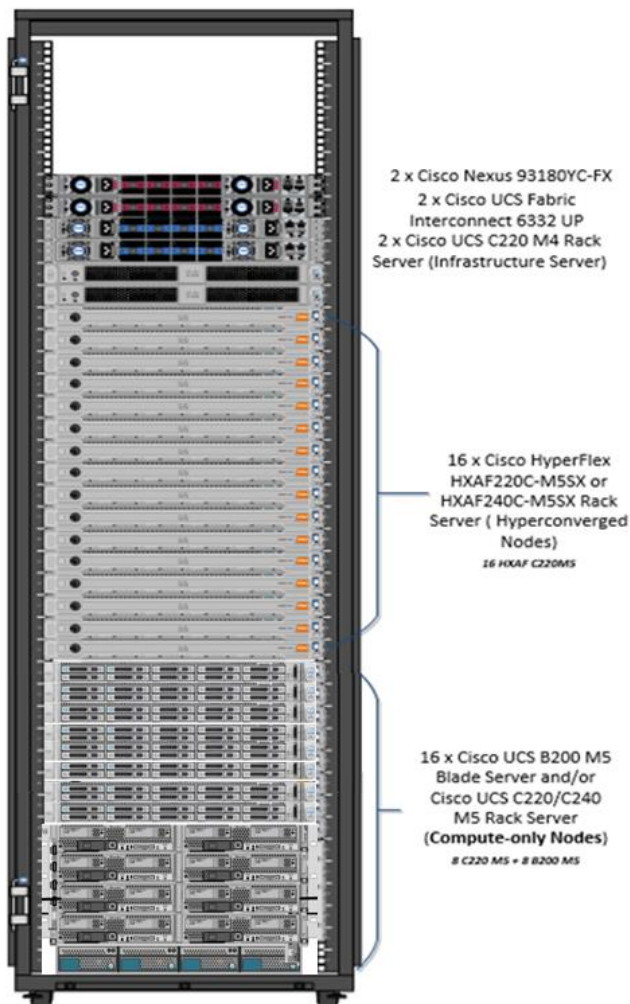
### Use Cases

- Healthcare: Mobility between desktops and terminals, compliance, and cost
- Federal government: Teleworking initiatives, business continuance, continuity of operations (COOP), and training centers
- Financial: Retail banks reducing IT costs, insurance agents, compliance, and privacy
- Education: K-12 student access, higher education, and remote learning

- State and local governments: IT and service consolidation across agencies and interagency security
- Retail: Branch-office IT cost reduction and remote vendors
- Manufacturing: Task and knowledge workers and offshore contractors
- Microsoft Windows 10 migration
- Graphic intense applications
- Security and compliance initiatives
- Opening of remote and branch offices or offshore facilities
- Mergers and acquisitions

Figure 2 shows the VMware Horizon 7 on vSphere 6.5 built on Cisco UCS components and the network connections. The reference architecture reinforces the "wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Figure 2 VMware Horizon 7 on vSphere 6.5 Built on Cisco Unified Computing System



## Physical Topology

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



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

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

Figure 3 Cisco HyperFlex Standard Topology

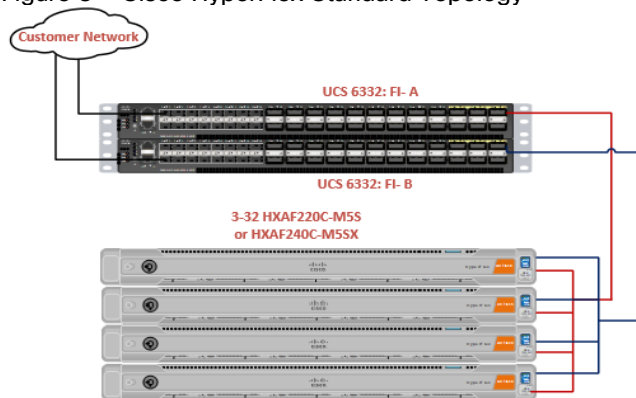
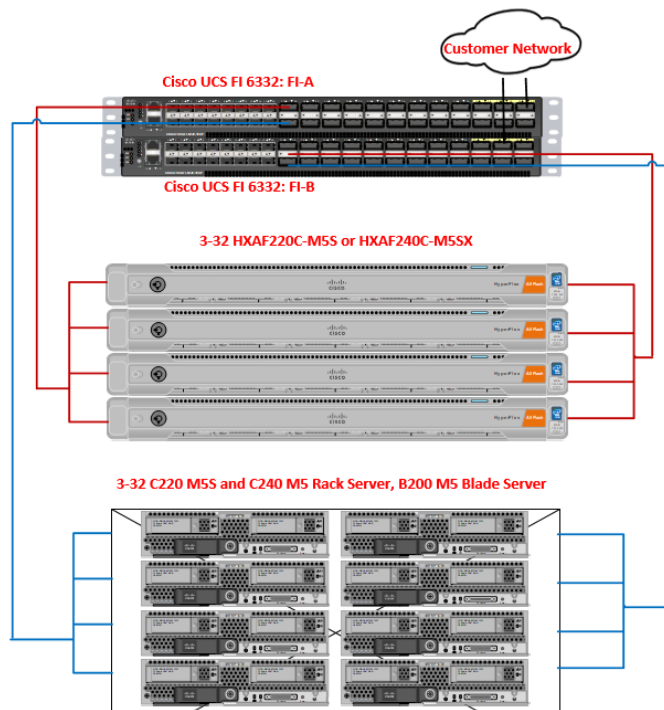


Figure 4 Cisco HyperFlex Hyperconverged Plus Compute Only Node Topology



## Fabric Interconnects

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

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

- **Mgmt:** A 10/100/1000 Mbps port for managing the Fabric Interconnect and the Cisco UCS domain 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 and C-Series Rack-Mount Servers

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



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


**Figure 5 HX-Series and C-Series Server Connectivity**



## Cisco UCS B-Series Blade Servers

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

**Figure 6 Cisco UCS 5108 Chassis Connectivity**



## Logical Network Design


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

- Management Zone: This zone comprises the connections needed to manage the physical hardware, the hypervisor hosts, and the storage platform controller virtual machines (SCVM). These interfaces and IP addresses need to be available to all staff who will administer the HX system, throughout the LAN/WAN. This zone must provide access to Domain Name System (DNS) and Network Time Protocol (NTP) services and allow Secure Shell (SSH) communication. In this zone are multiple physical and virtual components:
  - Fabric Interconnect management ports.
  - Cisco UCS external management interfaces used by the servers and blades, which answer through the FI management ports.
  - ESXi host management interfaces.
  - Storage Controller 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, ESXi 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 vmkernel interface used for storage traffic for each ESXi host in the HX cluster.
  - Storage Controller VM storage interfaces.
  - A roaming HX cluster storage interface.
- VMotion Zone: This zone comprises the connections used by the ESXi hosts to enable vMotion of the guest 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



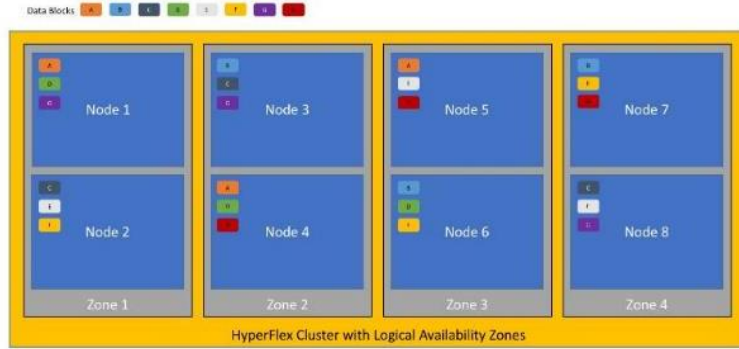
## Logical Availability Zones

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

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

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

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



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

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

The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-FX switches
- Two Cisco UCS 6332 fabric interconnects
- Sixteen Cisco HXAF C220M5 rack servers running HyperFlex data platform version 3.5.1a
- Eight Cisco UCS C220 M5 rack server running HyperFlex data platform version 3.5.1a as compute-only nodes.
- Eight Cisco UCS B200 M5 blade server running HyperFlex data platform version 3.5.1a as compute-only nodes.

For desktop virtualization, the deployment includes VMware Horizon 7 running on VMware vSphere 6.5. The design is intended to provide a large-scale building block for both RDSH and persistent/non-persistent desktops with following density per thirty-node configuration:

- 1550 Horizon 7 RDSH server desktop sessions
- 950 Horizon 7 Windows 10 non-persistent instant clone virtual desktops
- 950 Horizon 7 Windows 10 non-persistent composer clone virtual desktops
- 950 Horizon 7 Windows 10 persistent full clone virtual desktops



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

---

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

## Configuration Guidelines

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

## Solution Design

---

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

### Cisco Unified Computing System

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

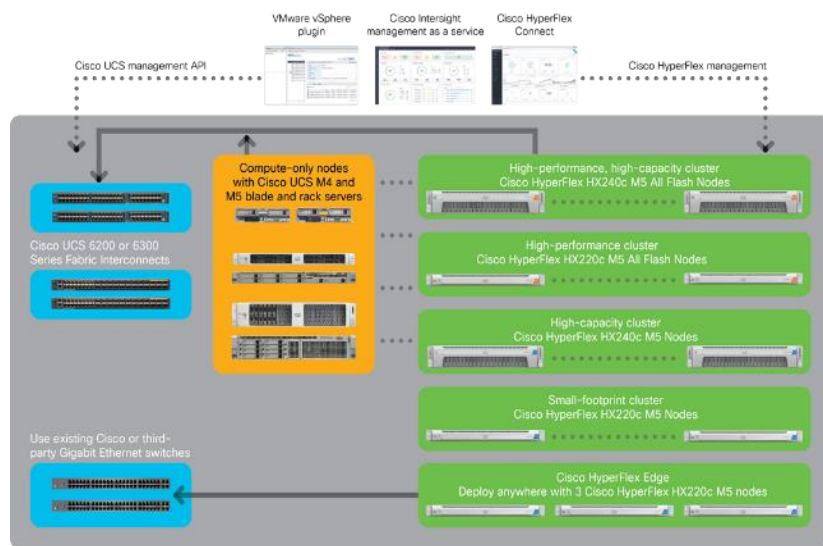
Cisco UCS is a next-generation data center platform that unites computing, networking, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system 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. Storage resources are managed from the familiar VMware vCenter web client, extending the capability of vCenter administrators.
- **Management:** Cisco UCS uniquely integrates all system components, enabling the entire solution to be managed as a single entity by Cisco UCS Manager. The manager has an intuitive GUI, a CLI, and a robust API for managing all system configuration processes and operations.

**Figure 9 Cisco HyperFlex Family Overview**  
Choice of management point for hardware and software



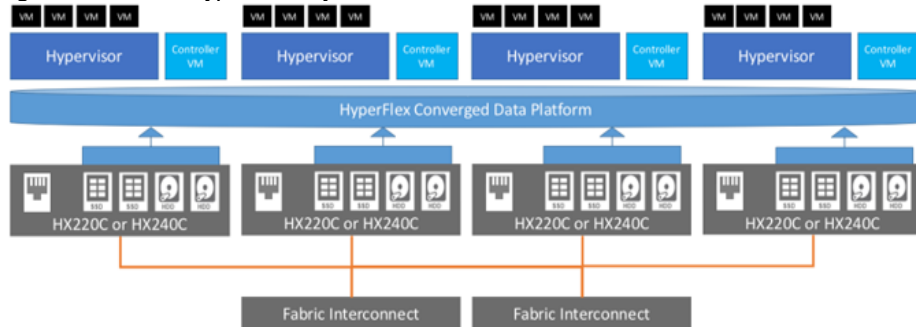
Cisco UCS and Cisco HyperFlex are designed to deliver:

- 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.

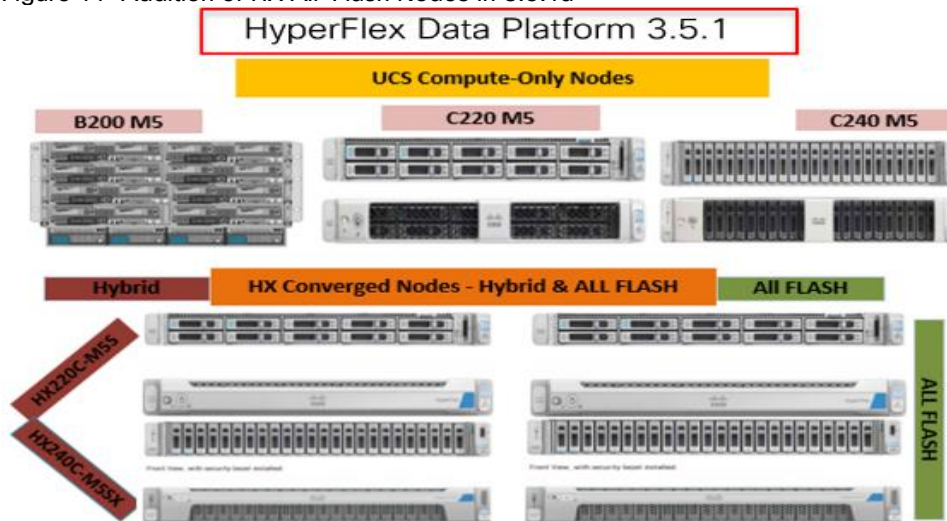
**Figure 10 Cisco HyperFlex System Overview**



## Enhancements for Version 3.5.1a

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

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



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



- REST APIs - [Cisco HyperFlex Systems REST API Getting Started Guide](#) on Cisco DevNet
- New All-Flash and Hybrid HX M5 server models are added to the Cisco HyperFlex product family
- Cisco Smart Licensing—Support for Cisco Smart Software Manager satellite. Please refer to the [https://www.cisco.com/c/en/us/td/docs/hyperconverged\\_systems/HyperFlex\\_HX\\_DataPlatformSoftware/Installation\\_VMWare\\_ESXi/3\\_5/b\\_HyperFlexSystems\\_Installation\\_Guide\\_for\\_VMware\\_ESXi\\_3\\_5.html](https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/Installation_VMWare_ESXi/3_5/b_HyperFlexSystems_Installation_Guide_for_VMware_ESXi_3_5.html) for more details.
- [M5 Servers](#)
- Key release highlights:
  - Same software feature set as HX 3.5.1a
  - Support for M5 servers in HyperFlex.
  - Enablement for Cisco HX240c M5 and HXAF240c M5 servers:
    - Dual CPU—Intel Xeon processor scalable family
    - Up to 3TB DRAM—Recommended minimum of 256 GB DRAM
    - M.2 Drive—For ESX Boot and for Storage Controller VM
    - Up to 2 GPUs—M10, P40, AMD 7150 x 2
    - Dedicated rear slots for caching
  - Enablement for Cisco HX220c M5 and HXAF220c M5 servers:
    - Dual CPU (Except Edge)—Intel Xeon processor scalable family
    - Up to 3TB DRAM—Recommended minimum of 256 GB DRAM
    - 8 x Data Drives (SATA/SAS)
    - M.2 Drive—For ESX Boot and for Storage Controller VM
- M4/M5 support in the same cluster.
  - A mixed cluster is defined by having both M4 and M5 HX converged nodes within the same storage cluster.
  - HyperFlex Edge does not support mixed clusters.
  - SED SKUs do not support mixed clusters.
- Peripherals
  - Option for 6-8 drives in HX220C-M5S and HXAF220C-M5S nodes.
  - Up to two GPUs for HX240C-M5SX and HXAF240C-M5SX nodes

## Cisco UCS Fabric Interconnect

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

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

For networking, the Cisco UCS 6300 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 40 Gigabit Ethernet on all ports, 2.56-terabit (Tb) switching capacity, and 320 Gbps of bandwidth per chassis, independent of packet size and enabled services. The product series supports Cisco low-latency, lossless, 40 Gigabit Ethernet unified network fabric capabilities, increasing the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnects support multiple traffic classes over a lossless Ethernet fabric, from the blade server through the interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

**Figure 12 Cisco UCS 6332 Series Fabric Interconnect**

Front View



Rear View



**Figure 13 Cisco UCS 6332-16UP Fabric Interconnect**

Front View



Rear View



## 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 being replicated across at least two of these nodes, and a third node is required for continuous operation in the event of a single-node failure. Each node that has disk storage is equipped with at least one high-performance SSD drive

for data caching and rapid acknowledgment of write requests. Each node is also equipped with the platform's physical capacity of either spinning disks or enterprise-value SSDs for maximum data capacity.

### Cisco UCS HXAF220c-M5S Rack Server

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

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

**Table 1 HXAF220c-M5SX Server Options**

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

Figure 14 Cisco UCS HXAF220c-M5SX Rack Server Front View





Figure 15 Cisco UCS HXAF220c-M5SX Rack Server Rear View



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

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

## Cisco VIC 1387 MLOM Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1387 is a dual-port Enhanced Small Form-Factor Pluggable (QSFP+) 40-Gbps Ethernet and Fibre Channel over Ethernet (FCoE) in a modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS HX-Series Rack Servers (Figure 5). The mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases. The card enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile.

**Figure 16 Cisco VIC 1387 mLOM Card**



**Table 2 Supported Physical Connectivity**

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

## Cisco HyperFlex Compute Nodes

### Cisco UCS B200-M5 Blade Server

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

Figure 17 Cisco UCS B200 M5 Blade Server



### Cisco VIC1340 Converged Network Adapter


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

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

Figure 18 Cisco UCS VIC 1340



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



### Cisco UCS 5108 Blade Chassis

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

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

Four hot-swappable power supplies are accessible from the front of the chassis, and single-phase 2500 W AC, 2500 W -48 VDC, and 2500 W 200 - 380 VDC power supplies and chassis are available. These power supplies



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

**Figure 20 Cisco UCS 5108 Blade Chassis Front and Rear Views**



## Features and Benefits

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

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

## Cisco UCS 2304XP Fabric Extender

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

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

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

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

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

**Figure 21 Cisco UCS 2304XP Fabric Extender**



## Cisco UCS C220-M5 Rack Server

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

**Figure 22 Cisco UCS C220 M5 Rack Server**



## Cisco HyperFlex HX Data Platform Administration Plug-in

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

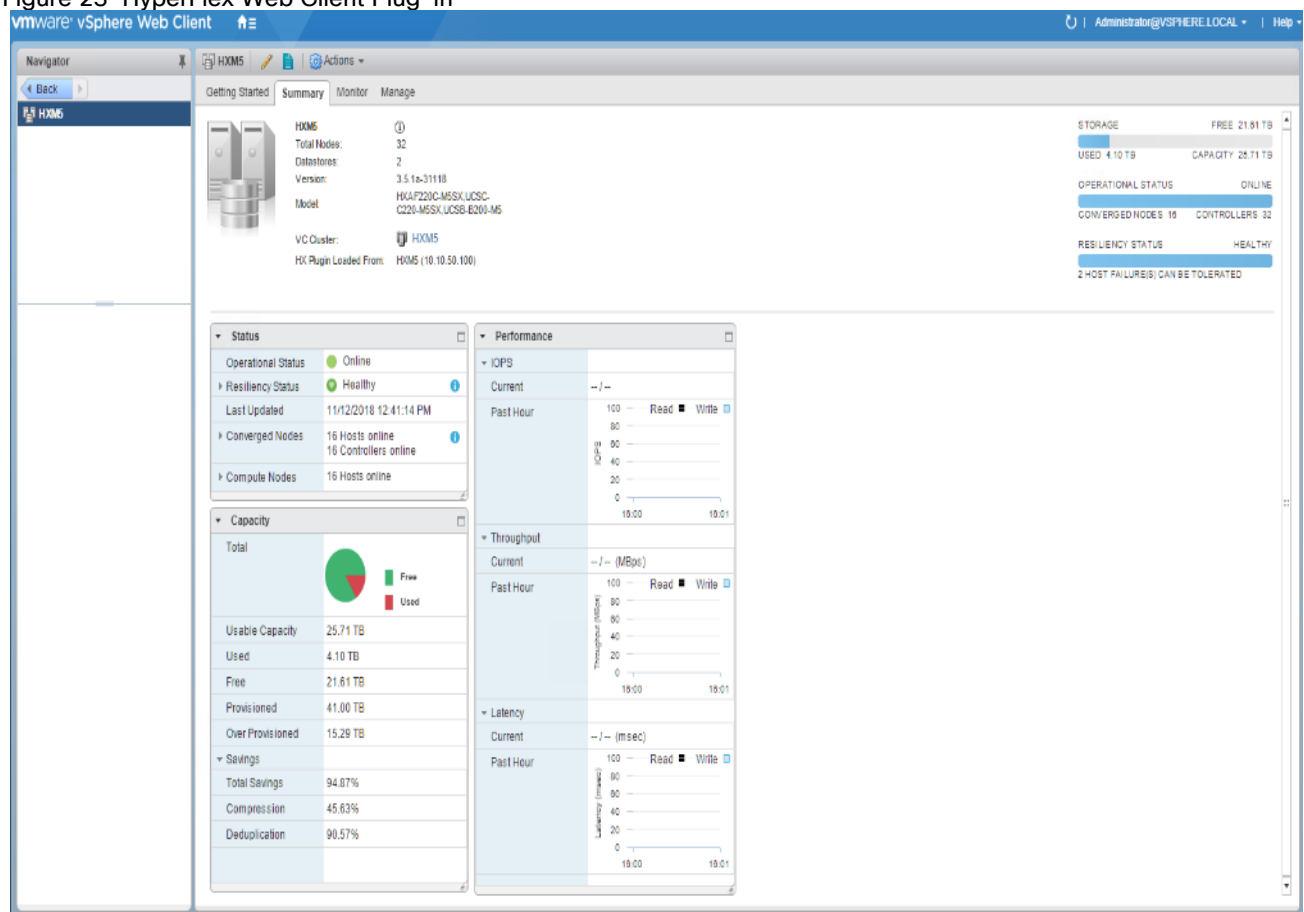
- Replication replicates data across the cluster so that data availability is not affected if single or multiple components fail (depending on the replication factor configured).
- Deduplication is always on, helping reduce storage requirements in virtualization clusters in which multiple operating system instances in client virtual machines result in large amounts of replicated data.
- Compression further reduces storage requirements, reducing costs, and the log-structured file system is designed to store variable-sized blocks, reducing internal fragmentation.



- Thin provisioning allows large volumes to be created without requiring storage to support them until the need arises, simplifying data volume growth and making storage a “pay as you grow” proposition.
- Fast, space-efficient clones rapidly replicate storage volumes so that virtual machines can be replicated simply through metadata operations, with actual data copied only for write operations.
- Snapshots help facilitate backup and remote-replication operations: needed in enterprises that require always-on data availability.

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

**Figure 23 HyperFlex Web Client Plug-in**



## Cisco HyperFlex Connect HTML5 Management Web Page

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

For the Tech Preview Web UI, connect to HX controller cluster IP: [http://hx\\_controller\\_cluster\\_ip/ui](http://hx_controller_cluster_ip/ui)

Figure 24 HyperFlex Web GUI Preview

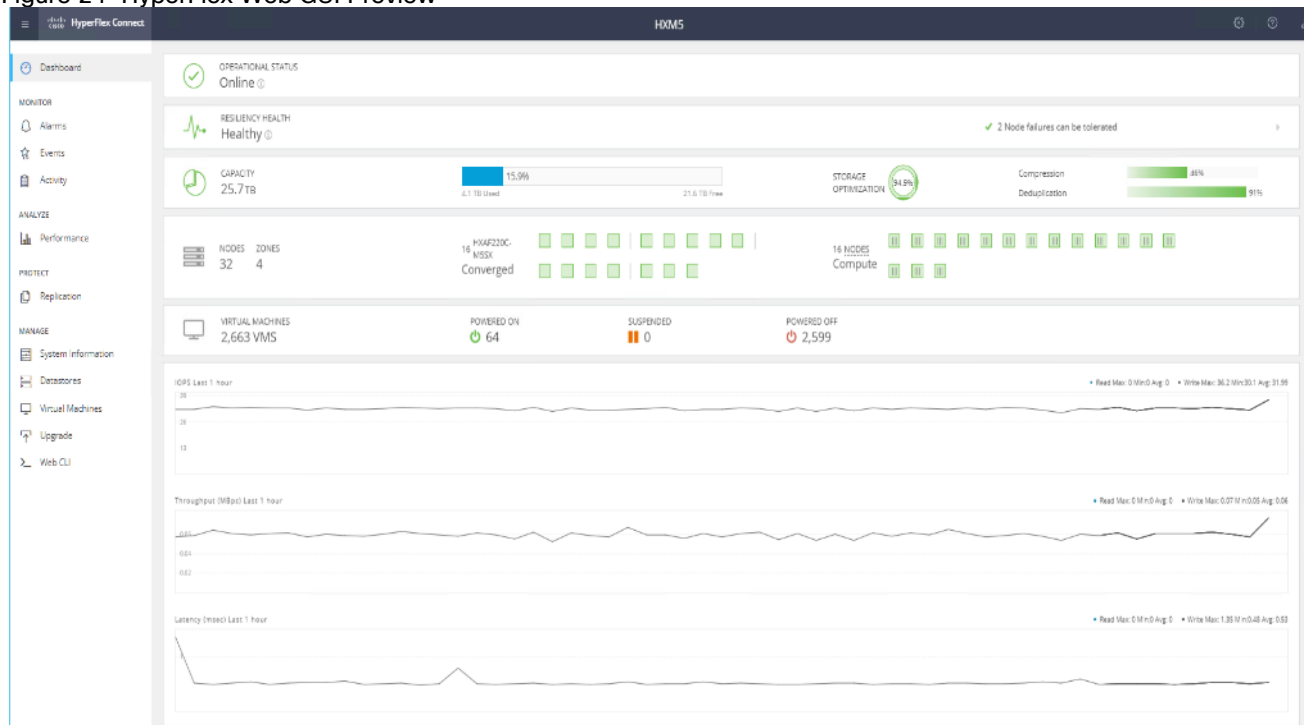
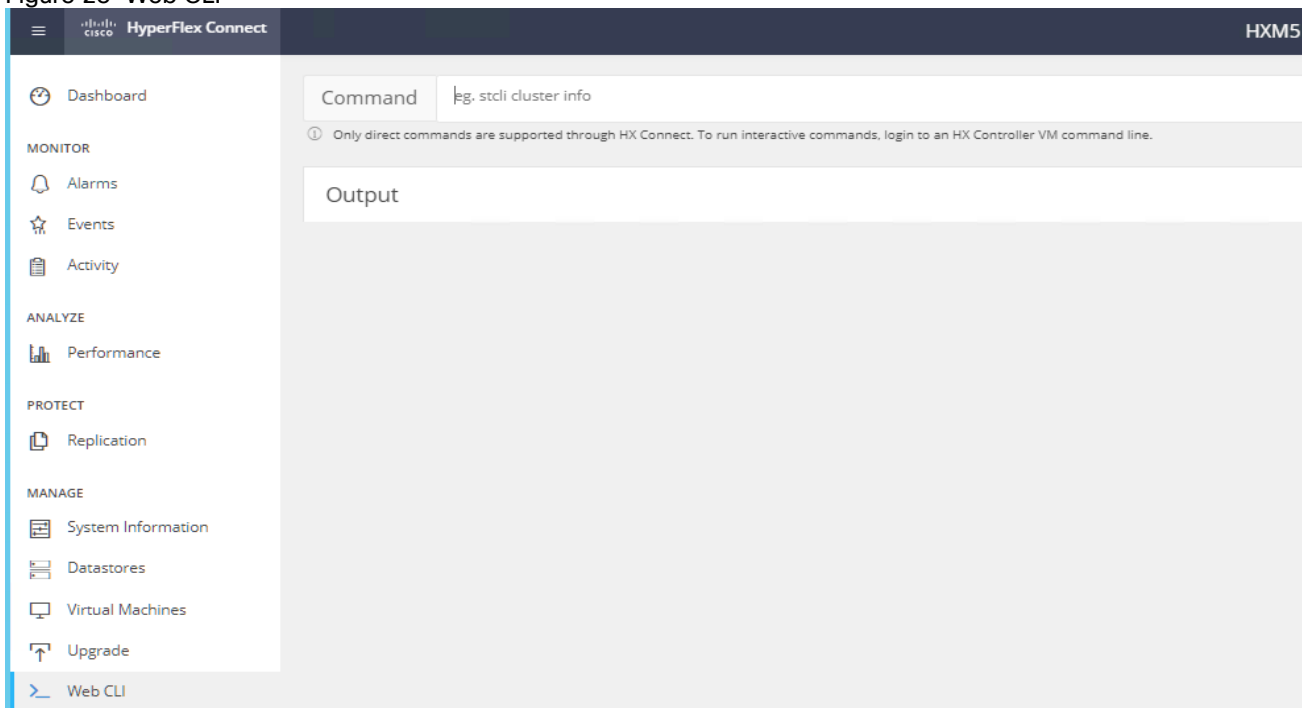


Figure 25 Web CLI



### Cisco HyperFlex HX Data Platform Controller

A Cisco HyperFlex HX Data Platform controller resides on each node and implements the distributed file system. The controller runs in user space within a virtual machine and intercepts and handles all I/O from guest virtual

machines. The platform controller VM uses the VMDirectPath I/O feature to provide PCI pass-through control of the physical server's SAS disk controller. This method gives the controller VM full control of the physical disk resources, utilizing the SSD drives as a read/write caching layer, and the HDDs as a capacity layer for distributed storage. The controller integrates the data platform into VMware software through the use of two preinstalled VMware ESXi vSphere Installation Bundles (VIBs):

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

## Replication Factor

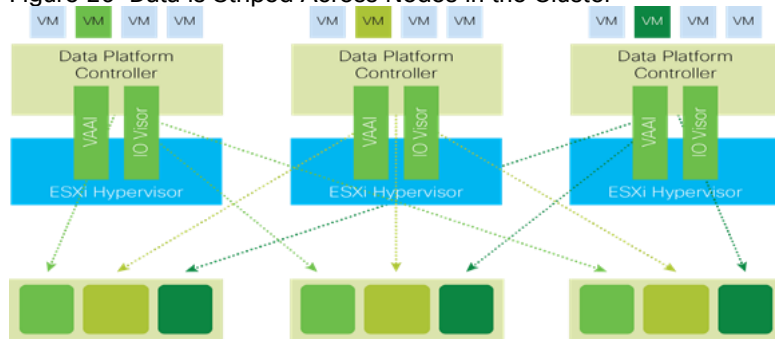
The policy for the number of duplicate copies of each storage block is chosen during cluster setup and is referred to as the replication factor (RF).

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

## Data Distribution

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

Figure 26 Data is Striped Across Nodes in the Cluster



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

## Data Operations

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

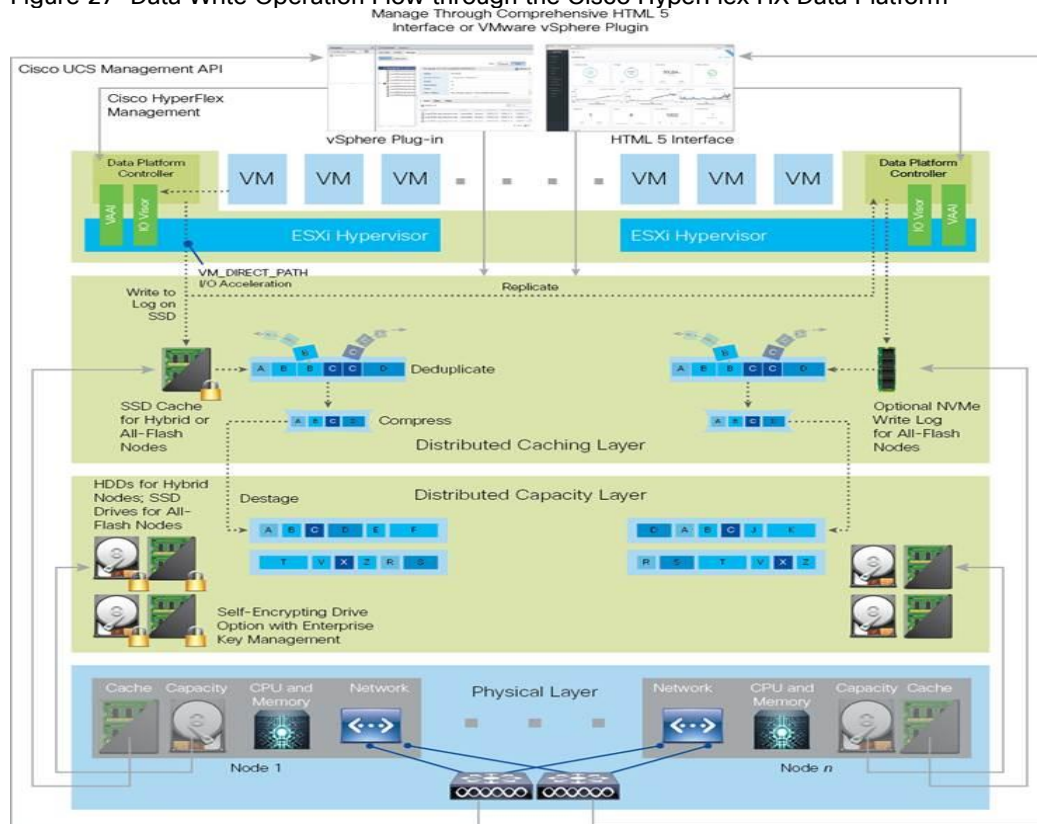
In the all-flash-memory configuration, the data platform uses a caching layer in SSDs to accelerate write responses, and it implements the capacity layer in SSDs. Read requests are fulfilled directly from data obtained from the SSDs in the capacity layer. A dedicated read cache is not required to accelerate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

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

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

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

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



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

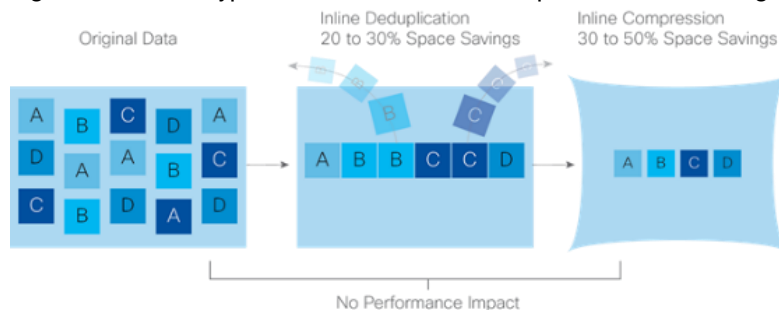
## Data Optimization

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

## Data Deduplication

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

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



## Inline Compression

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

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


## Log-Structured Distributed Objects

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

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

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

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



## Encryption

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

## Data Services

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

## Thin Provisioning

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

## Snapshots

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

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

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

Full featured backup applications, such as [Veeam Backup and Replication v9.5](#), have the ability to limit the amount of throughput the backup application can consume which can protect latency sensitive applications during the production hours. With the release of v9.5 update 2, Veeam is the first partner to [integrate HX native snapshots](#) into the product. HX Native snapshots do not suffer the performance penalty of delta-disk snapshots, and do not require heavy disk IO impacting consolidation during snapshot deletion.

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

## Fast, Space-Efficient Clones

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

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

## Data Replication and Availability

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

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

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

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

## Data Rebalancing

A distributed file system requires a robust data rebalancing capability. In the Cisco HyperFlex HX Data Platform, no overhead is associated with metadata access, and rebalancing is extremely efficient. Rebalancing is a non-disruptive online process that occurs in both the caching and persistent layers, and data is moved at a fine level of specificity to improve the use of storage capacity. The platform automatically rebalances existing data when nodes



and drives are added or removed or when they fail. When a new node is added to the cluster, its capacity and performance is made available to new and existing data. The rebalancing engine distributes existing data to the new node and helps ensure that all nodes in the cluster are used uniformly from capacity and performance perspectives. If a node fails or is removed from the cluster, the rebalancing engine rebuilds and distributes copies of the data from the failed or removed node to available nodes in the clusters.

## Online Upgrades

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

## Cisco Nexus 93180 Switches

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

### Specifications at-a-Glance

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

### Architectural Flexibility

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

### Feature Rich

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

### Highly Available and Efficient Design

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

### Simplified Operations

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

### Top-notch Security

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

### Real-time Visibility and Telemetry

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

### Investment Protection

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

### Resources

- [Cisco Nexus 9300-EX and 9300-FX Platform Leaf](#)
- [Switches for Cisco Application Centric Infrastructure Data Sheet](#)

Figure 30 Cisco Nexus 93180YC-FX Switch



## VMware vSphere 6.5

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

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

### VMware vCenter Server

- Migration Tool
- Improved appliance management

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

## VMware ESXi 6.5 Hypervisor

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

## VMware Horizon

VMware Horizon desktop virtualization solutions built on a unified architecture so they are simple to manage and flexible enough to meet the needs of all your organization's users. You use the same architecture and management tools to manage public, private, and hybrid cloud deployments as you do for on premises deployments

- VMware Horizon Virtual machines and RDSH known as server-based hosted sessions: These are applications hosted from Microsoft Windows servers to any type of device, including Windows PCs, Macs, smartphones, and tablets. Some VMware 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.
- VMware Horizon RDSH session users 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.

## Advantages of Using VMware Horizon

VMware Horizon 7 version 7.6 provides the following new features and enhancements:

- Instant Clones
  - A new type of desktop virtual machines that can be provisioned significantly faster than the traditional View Composer linked clones.
  - A fully functional desktop can be provisioned in two seconds or less.

- Recreating a desktop pool with a new OS image can be accomplished in a fraction of the time it takes a View Composer desktop pool because the parent image can be prepared well ahead of the scheduled time of pool recreation.
- Clones are automatically rebalanced across available datastores.
- View storage accelerator is automatically enabled.
- You can use NVIDIA GRID vGPUs with instant-clone desktop pools. Configuring PCoIP as the display protocol with NVIDIA GRID vGPU is a technical preview feature.
- You can select multiple vLAN networks to create a larger instant-clone desktop pool. Only the static port group is supported.
- You can use the internal VM debug mode to troubleshoot internal virtual machines in an instant-clone desktop pool or in an instant-clone farm.
- Administrators can perform a restart or reset of the virtual desktops managed by the vCenter Server.
- You can perform maintenance on instant-clone virtual machines by putting the ESXi hosts into maintenance mode. Use vSphere Web Client to put the ESXi host into maintenance mode. The ESXi host maintenance operation automatically deletes the parent virtual machines from that ESXi host.
- VMware Blast Extreme
  - VMware Blast Extreme is now fully supported on the Horizon platform.
  - Connections to physical machines that have no monitors attached are supported with NVIDIA graphics cards. This is a technical preview feature for Horizon 7 version 7.1.
  - The Blast Secure Gateway includes Blast Extreme Adaptive Transport (BEAT) networking, which dynamically adjusts to network conditions such as varying speeds and packet loss.
  - Administrators can select the VMware Blast display protocol as the default or available protocol for pools, farms, and entitlements.
  - End users can select the VMware Blast display protocol when connecting to remote desktops and applications.
  - VMware Blast Extreme features include:
    - TCP and UDP transport support
    - H.264 support for the best performance across more devices
    - Reduced device power consumption for longer battery life
    - NVIDIA GRID acceleration for more graphical workloads per server, better performance, and a superior remote user experience
- True SSO
  - For VMware Identity Manager Integration, True SSO streamlines the end-to-end login experience. After users log in to VMware Identity Manager using a smart card or an RSA SecurID or RADIUS token, users are not required to also enter Active Directory credentials in order to use a remote desktop or application.
  - Uses a short-lived Horizon virtual certificate to enable a password-free Windows login.

- Supports using either a native Horizon Client or HTML Access.
- System health status for True SSO appears in the Horizon Administrator dashboard.
- Can be used in a single domain, in a single forest with multiple domains, and in a multiple-forest, multiple-domain setup.
- Smart Policies
  - Control of the clipboard cut-and-paste, client drive redirection, USB redirection, and virtual printing desktop features through defined policies.
  - PCoIP session control through PCoIP profiles.
  - Conditional policies based on user location, desktop tagging, pool name, and Horizon Client registry values.
- Configure the Clipboard Memory Size for VMware Blast and PCoIP Sessions

Horizon administrators can configure the server clipboard memory size by setting GPOs for VMware Blast and PCoIP sessions. Horizon Client 4.1 users on Windows, Linux, and Mac OS X systems can configure the client clipboard memory size. The effective memory size is the lesser of the server and client clipboard memory size values.

- VMware Blast Network Recovery Enhancements

Network recovery is now supported for VMware Blast sessions initiated from iOS, Android, Mac OS X, Linux, and Chrome OS clients. Previously, network recovery was supported only for Windows client sessions. If you lose your network connection unexpectedly during a VMware Blast session, Horizon Client attempts to reconnect to the network and you can continue to use your remote desktop or application. The network recovery feature also supports IP roaming, which means you can resume your VMware Blast session after switching to a WiFi network.

- Configure Horizon Administrator to not remember the login name

Horizon administrators can configure not to display the Remember user name check box and therefore not remember the administrator's login name.

- Allow Mac OS X Users to Save Credentials

Horizon administrators can configure Connection Server to allow Horizon Client Mac OS X systems to remember a user's user name, password, and domain information. If users choose to have their credentials saved, the credentials are added to the login fields in Horizon Client on subsequent connections.

- Windows 10

- Windows 10 is supported as a desktop guest operating system
- Horizon Client runs on Windows 10
- Smart card is supported on Windows 10
- The Horizon User Profile Migration tool migrates Windows 7, 8/8.1, Server 2008 R2, or Server 2012 R2 user profiles to Windows 10 user profiles.

- RDS Desktops and Hosted Apps

- View Composer. View Composer and linked clones provide automated and efficient management of RDS server farms.
- Graphics Support. Existing 3D vDGA and GRID vGPU graphics solutions on VDI desktops have been extended to RDS hosts, enabling graphics-intensive applications to run on RDS desktops and Hosted Apps.
- Enhanced Load Balancing. A new capability provides load balancing of server farm applications based on memory and CPU resources.
- One-Way AD Trusts  
One-way AD trust domains are now supported. This feature enables environments with limited trust relationships between domains without requiring Horizon Connection Server to be in an external domain.
- Cloud Pod Architecture (CPA) Enhancements
  - Hosted App Support. Support for application remoting allows applications to be launched using global entitlements across a pod federation.
  - HTML Access (Blast) Support. Users can use HTML Access to connect to remote desktops and applications in a Cloud Pod Architecture deployment.
- Access Point Integration
  - Access Point is a hardened Linux-based virtual appliance that protects virtual desktop and application resources to allow secure remote access from the Internet. Access Point provides a new authenticating DMZ gateway to Horizon Connection Server. Smart card support on Access Point is available as a Tech Preview. Security server will continue to be available as an alternative configuration. For more information, see [Deploying and Configuring Access Point](#).
- FIPS
  - Install-time FIPS mode allows customers with high security requirements to deploy Horizon 6.
- Graphics Enhancements
  - AMD vDGA enables vDGA pass-through graphics for AMD graphics hardware.
  - 4K resolution monitors (3840x2160) are supported.
- Horizon Administrator Enhancements
  - Horizon Administrator shows additional licensing information, including license key, named user and concurrent connection user count.
  - Pool creation is streamlined by letting Horizon administrators clone existing pools.
- Horizon 7 for Linux Desktop Enhancements
  - Support for managed virtual machines
  - Support for smart card redirection with SSO
  - Support for Horizon Client for iOS
  - Support for SLES 12 SP1

- Support for H.264 encoder software
- True SSO support on the RHEL/CentOS 7 desktops  
True single sign-on (SSO) is now supported on RHEL 7 and CentOS desktops.
- Additional supported platforms  
Support the Ubuntu 18.04, RHEL/CentOS 6.10, and RHEL/CentOS 7.5 platforms have been added.
- Instant-clone support for RHEL 7.1 and later versions  
You can now create instant-clone floating desktop pool on systems with RHEL 7.1 or later installed.
- Additional Features
  - Support for IPv6 with VMware Blast Extreme on security servers.
  - Horizon Administrator security protection layer. See VMware Knowledge Base (KB) article 2144303 for more information.
  - Protection against inadvertent pool deletion.
  - RDS per-device licensing improvements.
  - Support for Intel vDGA.
  - Support for AMD Multiuser GPU Using vDGA.
  - More resilient upgrades.
  - Display scaling for Windows Horizon Clients.
  - DPI scaling is supported if it is set at the system level and the scaling level is greater than 100.
- vSphere Support
  - vSphere 6.5 U2 is supported.
- Virtual Desktops
  - vMotion support for vGPU-enabled virtual desktops

## What are VMware RDS Hosted Sessions?

The following describes the VMware RDS Hosted Sessions:

- An RDS host is a server computer that hosts applications and desktop sessions for remote access. An RDS host can be a virtual machine or a physical server.
- An RDS host has the Microsoft Remote Desktop Services role, the Microsoft Remote Desktop Session Host service, and Horizon Agent installed. Remote Desktop Services was previously known as Terminal Services. The Remote Desktop Session Host service allows a server to host applications and remote desktop sessions. With Horizon Agent installed on an RDS host, users can connect to applications and desktop sessions by using the display protocol PCoIP or Blast Extreme. Both protocols provide an optimized user experience for the delivery of remote content, including images, audio and video.
- The performance of an RDS host depends on many factors. For information on how to tune the performance of different versions of Windows Server, see <http://msdn.microsoft.com/library/windows/hardware/gg463392.aspx>.

- Horizon 7 supports at most one desktop session and one application session per user on an RDS host.
- When users submit print jobs concurrently from RDS desktops or applications that are hosted on the same RDS host, the ThinPrint server on the RDS host processes the print requests serially rather than in parallel. This can cause a delay for some users. Note that the print server does not wait for a print job to complete before processing the next one. Print jobs that are sent to different printers will print in parallel.
- If a user launches an application and also an RDS desktop, and both are hosted on the same RDS host, they share the same user profile. If the user launches an application from the desktop, conflicts may result if both applications try to access or modify the same parts of the user profile, and one of the applications may fail to run properly.
- The process of setting up applications or RDS desktops for remote access involves the following tasks:
  - Installing Applications
    - If you plan to create application pools, you must install the applications on the RDS hosts. If you want Horizon 7 to automatically display the list of installed applications, you must install the applications so that they are available to all users from the Start menu. You can install an application at any time before you create the application pool. If you plan to manually specify an application, you can install the application at any time, either before or after creating an application pool.
- Important
  - When you install an application, you must install it on all the RDS hosts in a farm and in the same location on each RDS host. If you do not, a health warning will appear on the View Administrator dashboard. In such a situation, if you create an application pool, users might encounter an error when they try to run the application.
  - When you create an application pool, Horizon 7 automatically displays the applications that are available to all users rather than individual users from the Start menu on all of the RDS hosts in a farm. You can choose any applications from that list. In addition, you can manually specify an application that is not available to all users from the Start menu. There is no limit on the number of applications that you can install on an RDS host.

## Farms, RDS Hosts, Desktop and Application Pools

With VMware Horizon, you can create desktop and application pools to give users remote access to virtual machine-based desktops, session-based desktops, physical computers, and applications. Horizon takes advantage of Microsoft Remote Desktop Services (RDS) and VMware PC-over-IP (PCoIP) technologies to provide high-quality remote access to users.

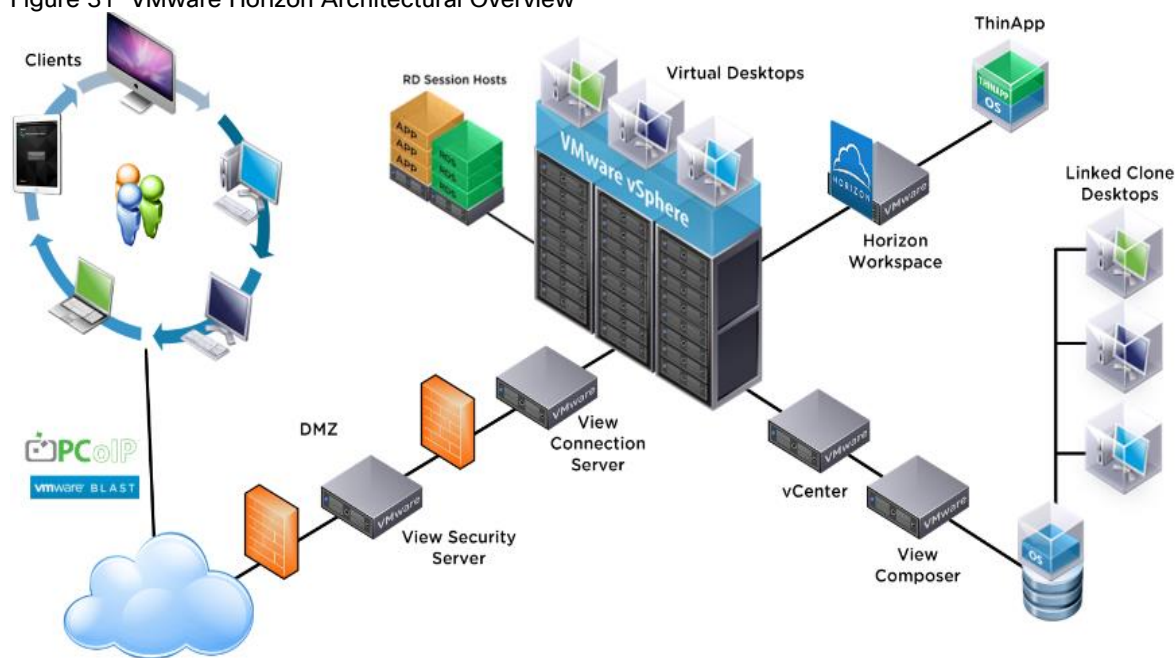
- RDS Hosts
  - RDS hosts are server computers that have Windows Remote Desktop Services and View Agent installed. These servers host applications and desktop sessions that users can access remotely. To use RDS desktop pools or applications, your end users must have access to Horizon Client 3.0 or later software.
- Desktop Pools
  - There are three types of desktop pools: automated, manual, and RDS. Automated desktop pools use a vCenter Server virtual machine template or snapshot to create a pool of identical virtual machines. Manual desktop pools are a collection of existing vCenter Server virtual machines, physical computers, or third-party virtual machines. In automated or manual pools, each machine is available for one user to



access remotely at a time. RDS desktop pools are not a collection of machines, but instead, provide users with desktop sessions on RDS hosts. Multiple users can have desktop sessions on an RDS host simultaneously.

- Application Pools
  - Application pools let you deliver applications to many users. The applications in application pools run on a farm of RDS hosts.
- Farms
  - Farms are collections of RDS hosts and facilitate the management of those hosts. Farms can have a variable number of RDS hosts and provide a common set of applications or RDS desktops to users. When you create an RDS desktop pool or an application pool, you must specify a farm. The RDS hosts in the farm provide desktop and application sessions to users.

Figure 31 VMware Horizon Architectural Overview



## Architecture and Design of VMware Horizon on Cisco Unified Computing System and Cisco HyperFlex System 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 sample user classifications are provided:

- Knowledge Workers today do not just work in their offices all day – they attend meetings, visit branch offices, work from home, and even coffee shops. These anywhere workers expect access to all of their same applications and data wherever they are.
- External Contractors are increasingly part of your everyday business. They need access to certain portions of your applications and data, yet administrators still have little control over the devices they use and the

locations they work from. Consequently, IT is stuck making trade-offs on the cost of providing these workers a device vs. the security risk of allowing them access from their own devices.

- Task Workers perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements from their PCs. However, since these workers are interacting with your customers, partners, and employees, they have access to your most critical data.
- Mobile Workers need access to their virtual desktop from everywhere, regardless of their ability to connect to a network. In addition, these workers expect the ability to personalize their PCs, by installing their own applications and storing their own data, such as photos and music, on these devices.
- Shared Workstation users are often found in state-of-the-art university and business computer labs, conference rooms or training centers. Shared workstation environments have the constant requirement to re-provision desktops with the latest operating systems and applications as the needs of the organization change, tops the list.

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

- Traditional PC: A traditional PC is what typically 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 or 2016, 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 VMware RDSH 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 synchronized with the data center when the device is connected to the network.

For the purposes of the validation represented in this document both Horizon Virtual Desktops and Remote Desktop sever Hosted 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 VMware RDSH for Remote Desktop Server Hosted Sessions used?
- If RDSH is used what is the desktop OS planned? Server 2008, Server 2012 or Server 2016?
- How many RDSH sessions will be deployed in the pilot? In production?
- 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 (such as 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?

## Desktop Virtualization Design Fundamentals

An ever growing and diverse base of user devices, complexity in management of traditional desktops, security, and even Bring Your Own Device (BYOD) to work programs are prime reasons for moving to a virtual desktop solution.

## VMware Horizon Design Fundamentals

VMware Horizon 7 integrates Remote Desktop Server Hosted sessions users and VDI desktop virtualization technologies into a unified architecture that enables a scalable, simple, efficient, mixed users 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. VMware Horizon delivers a native touch-optimized experience via PCoIP or Blast Extreme high-definition performance, even over mobile networks.

## Horizon VDI Pool and RDSH Servers Pool

Collections of identical Virtual Machines (VMs) or physical computers are managed as a single entity called a Desktop Pool. In this CVD, VM provisioning relies on VMware View Composer aligning with VMware Horizon View Connection Server and vCenter Server components. Machines in these Pools are configured to run either a Windows Server 2016 OS (for RDSH hosted shared sessions) or a Windows 10 Desktop OS (for linked clone, instant clone and persistent VDI desktops).



Server OS and Desktop OS Machines were configured in this CVD to support RDSH hosted shared desktops and a variety of VDI hosted virtual desktops.

Figure 32 VMware Horizon Design Overview

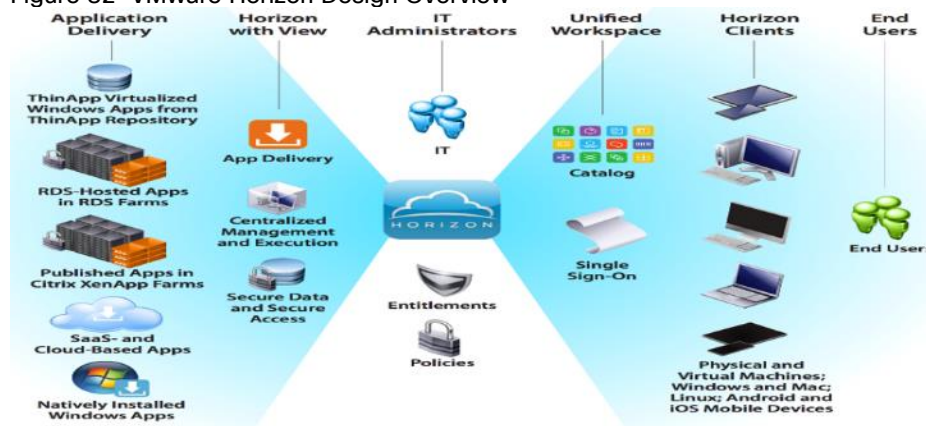
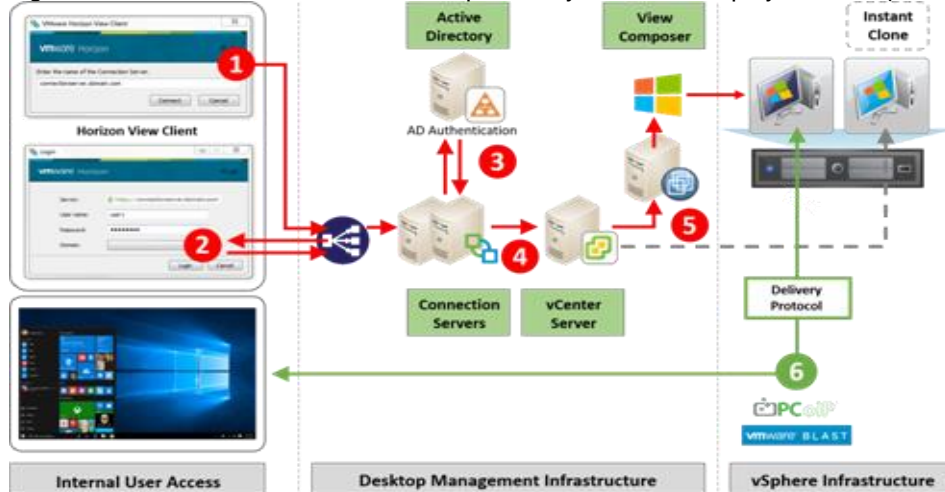


Figure 33 Horizon VDI and RDSH Desktop Delivery Based on Display Protocol (PCoIP/Blast/RDP)

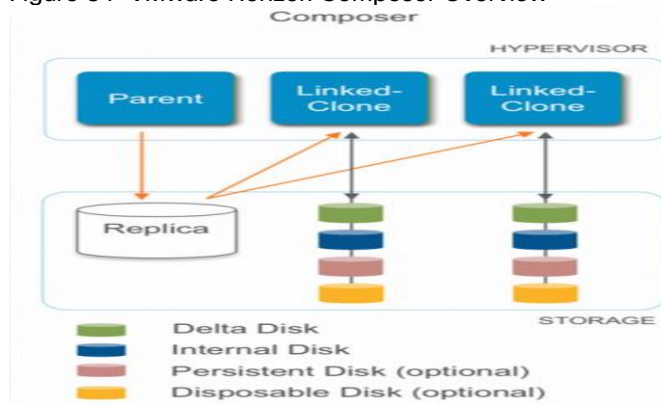


### VMware Horizon Composer

VMware Horizon Composer is a feature in Horizon that gives administrators the ability to manage virtual machine pools or the desktop pools that share a common [virtual disk](#). An administrator can update the [master image](#), then all desktops using [linked clones](#) of that master image can also be patched. Updating the master image will patch the cloned desktops of the users without touching their applications, data or settings.

The VMware View Composer pooled desktops solution’s infrastructure is based on software-streaming technology. After creating and configuring the Master Image for a virtual desktop pool, a snapshot is taken of the OS and applications that is accessible to host(s).

Figure 34 VMware Horizon Composer Overview



### VMware View Storage Accelerator

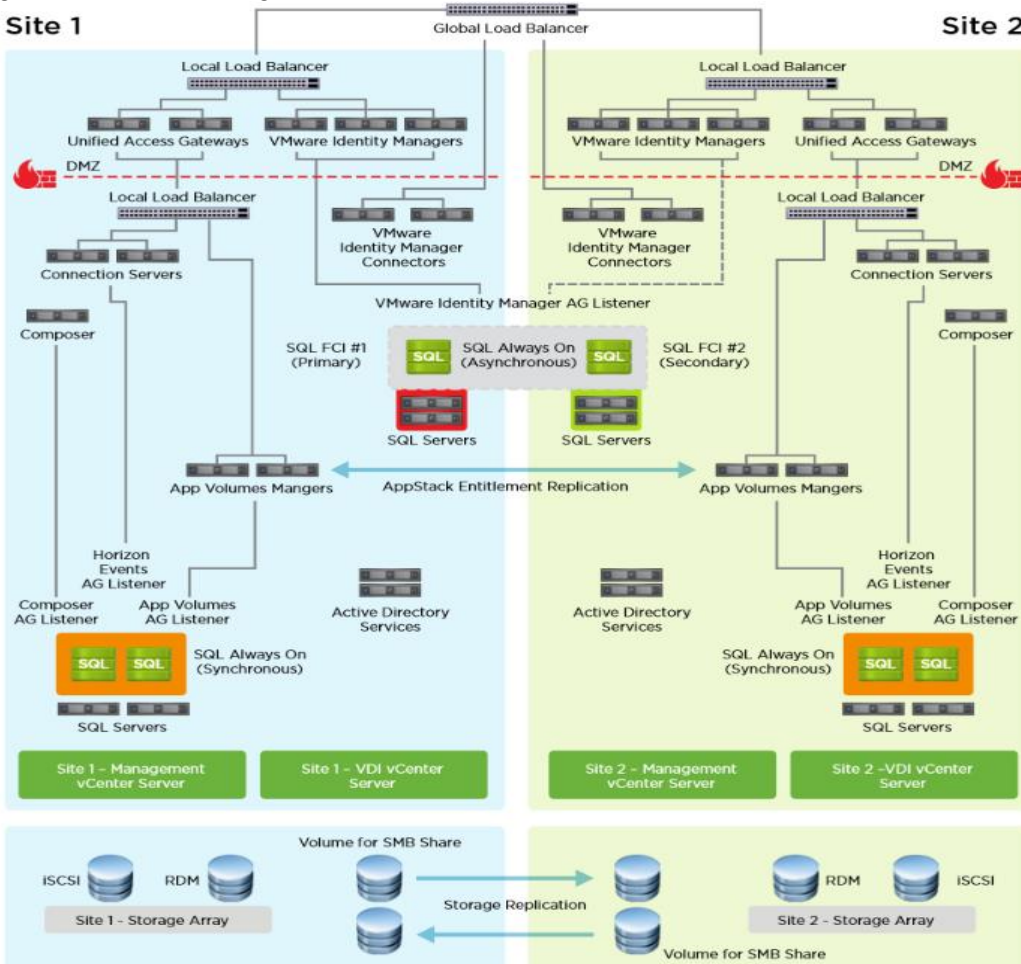
VMware View Storage Accelerator is an in-memory host caching capability that uses the content-based read cache (CBRC) feature in ESXi hosts. CBRC provides a per-host RAM-based solution for View desktops, which greatly reduces the number of read I/O requests that are issued to the storage layer. It also addresses boot storms—when multiple virtual desktops are booted at the same time—which can cause a large number of read operations. CBRC is beneficial when administrators or users load applications or data frequently. Note that CBRC was used in all tests that were performed on the solution described here: Horizon running pooled linked-clone desktops hosted on Cisco HyperFlex system.

### Multiple Site Configuration

If you have multiple regional sites, you can use any of the Load Balances Tools to direct the user connections to the most appropriate site to deliver the desktops and application to users.

Figure 35 illustrating sites, shows a site created in two data centers. Having two sites globally, rather than just one, minimizes the amount of unnecessary WAN traffic. Two Cisco blade servers host the required infrastructure services (Domain Controllers, DNS, DHCP, Profile, SQL, VMware Horizon View Connection Servers, View Composer server and web servers).

Figure 35 Multisite Configuration Overview



Based on the requirement and no of data centers or remote location, we can choose any of the available Load balancing software or tools accelerates the application performance, load balances servers, increases security, and optimizes the user experience.

---

 Multi-Site configuration is shown as the example.

---

### Designing a VMware Horizon Environment for Various Workload Types

With VMware Horizon 7, 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><b>You want:</b> 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><b>Your users:</b> Perform well-defined tasks and do not require personalization or off-line access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations.</p> <p><b>Application types:</b> Any application.  </p>
<p>Desktop OS machines</p>	<p><b>You want:</b> 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><b>Your users:</b> Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications.</p> <p><b>Application types:</b> 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><b>You want:</b> 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><b>Your users:</b> 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><b>Host:</b> The same as Desktop OS machines.</p> <p><b>Application types:</b> Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.</p>



For the Cisco Validated Design described in this document, individual configuration of Remote Desktop Server Hosted sessions (RDSH) using RDS-based Server OS machines and Hosted Virtual Desktops (HVDs) using Desktop OS machines via Instant-clone, Linked- clone automated pool and Full clone persistent desktops were configured and tested. The following sections discuss design decisions relative to the VMware Horizon deployment, including this CVD test environment.

# Deployment Hardware and Software

---

## Products Deployed

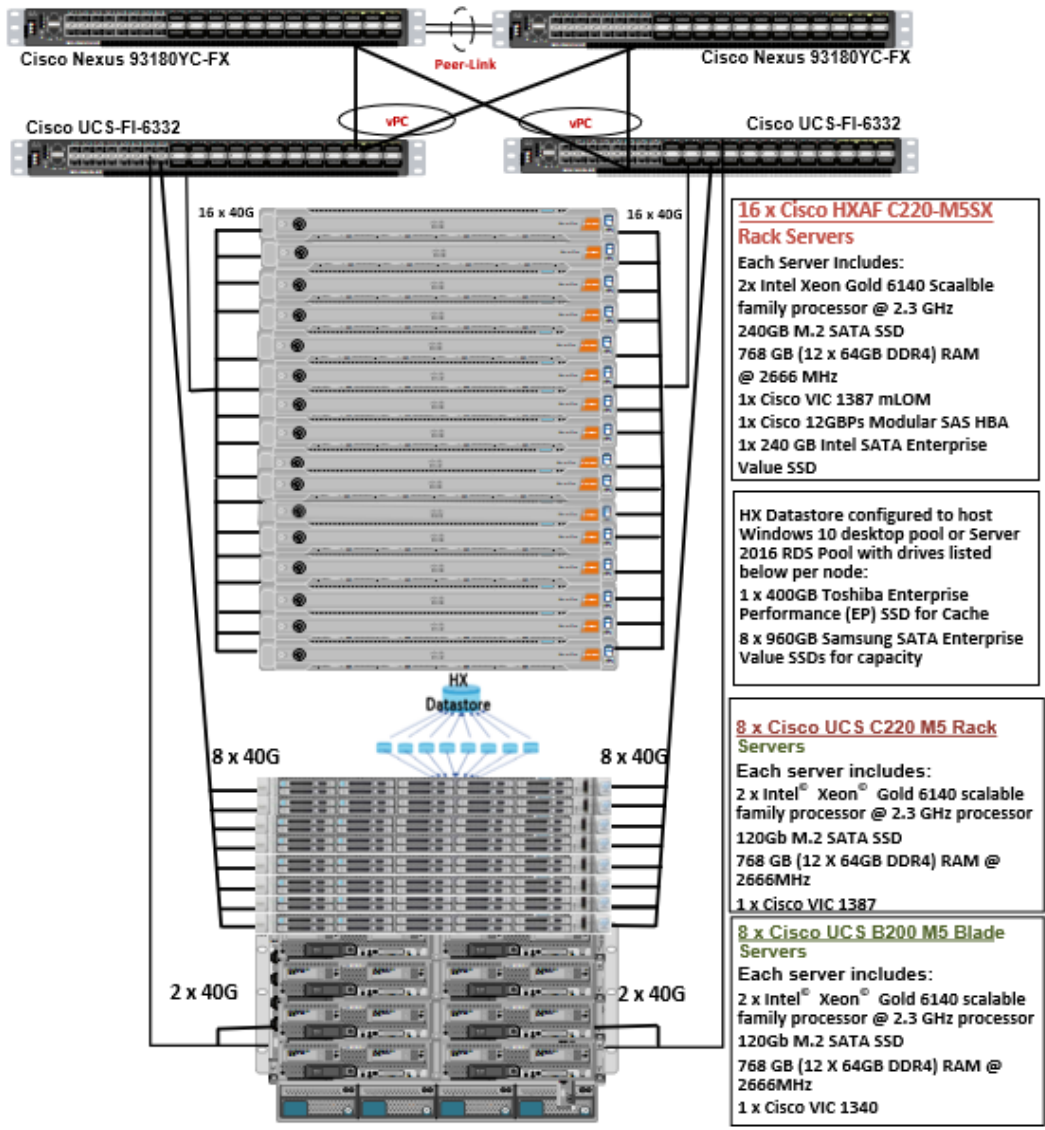
The architecture deployed is modular. While each customer's environment might vary in its exact configuration, the reference architecture contained in this document once built, can easily be scaled as requirements and demands change. This includes scaling both up (adding additional resources within existing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

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

This validated design document details the deployment of the multiple configurations extending to 4400 users for Horizon virtual desktop or Horizon RDSH published desktop workload respectively featuring the following software:

- VMware Horizon 7 Shared Remote Desktop Server Hosted (RDSH) sessions on Cisco HyperFlex
- VMware Horizon 7 Non-Persistent and persistent Virtual Desktops (VDI) 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
- VMware vSphere ESXi 6.5.2 (Update 2) Hypervisor
- Windows Server 2016 for RDSH Servers & Windows 10 64-bit Operating Systems for VDI virtual machines
- Microsoft SQL Server 2016
- Cisco HyperFlex data platform v3.5(1a)
- VMware Horizon 7 Connection Server and Replica Servers for redundancy and support up to 4400 seat scale
- VMware Horizon 7 Composer Server

Figure 36 Detailed Reference Architecture with Physical Hardware Cabling Configured to Enable the Solution  
**Cisco HyperFlex and VMware Horizon 7, Full Scale Single UCS Domain Reference Architecture**



### Hardware Deployed

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

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

- Eight Cisco UCS C220 M5 Rack servers with Intel Xeon Gold 6140 scalable family 2.3-GHz 18-core processors, 768GB RAM 2666-MHz and VIC1387 mLOM cards running Cisco HyperFlex data platform v3.5(1a) for the virtual desktop workloads with N+1 server fault tolerance
- Eight Cisco UCS B200 M5 blade servers with Intel Xeon Gold 6140 scalable family 2.3-GHz 18-core processors, 768GB RAM 2666-MHz and VIC1340 mLOM cards running Cisco HyperFlex data platform v3.5(1a) for the virtual desktop workloads with N+1 server fault tolerance.

## Software Deployed

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

**Table 3 Software and Firmware Versions**

Vendor	Product	Version
Cisco	UCS Component Firmware	4.0(1b) bundle release
Cisco	UCS Manager	4.0(1b) bundle release
Cisco	UCS HXAF220c-M5S rack server	4.0(1b) bundle release
Cisco	VIC 1387	4.2(2b)
Cisco	UCS B200 M5 blade server	4.0(1b) bundle release
Cisco	VIC 1340	4.2(2d)
Cisco	HyperFlex Data Platform	3.5.1a
Cisco	Cisco NENIC	2.1.2.71
Cisco	Cisco fNIC	1.6.0.37
Network	Cisco Nexus 9000 NX-OS	7.0(3)I7(2)
VMware	Horizon Connection Server	7.6.0-9823717
VMware	Horizon Composer Server	7.6.0-9491669
VMware	Horizon Agent	7.6.0-9539447
VMware	Horizon Client	4.9.0-9539668
VMware	vCenter Server Appliance	6.5.0-5973321
VMware	vSphere ESXi 6.5 Update 2	6.5.2.U2-8935087

## Logical Architecture

The logical architecture of this solution has designed to support up to 4400 RDSH hosted shared server desktop users and Hosted Virtual Microsoft Windows 10 Desktops within a sixteen node Cisco UCS HXAF220c-M5S, eight Cisco UCS C220-M5S and eight Cisco UCS B200 M5 HyperFlex cluster, which provides physical redundancy for each workload type.

Figure 37 Logical Architecture Design




Table 3 lists the software revisions for this solution.



This document is intended to allow you to fully configure your environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 4 through Table 8 lists the information you need to configure your environment.

## VLANs

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

**Table 4 Table 2 VLANs Configured in this Study**

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



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

## Jumbo Frames

All HyperFlex storage traffic traversing the hx-storage-data VLAN and subnet is configured to use jumbo frames, or to be precise all communication is configured to send IP packets with a Maximum Transmission Unit (MTU) size of 9000 bytes. Using a larger MTU value means that each IP packet sent carries a larger payload, therefore transmitting more data per packet, and consequently sending and receiving data faster. This requirement also

means that the Cisco UCS uplinks must be configured to pass jumbo frames. Failure to configure the Cisco UCS uplink switches to allow jumbo frames can lead to service interruptions during some failure scenarios, particularly when cable or port failures would cause storage traffic to traverse the northbound Cisco UCS uplink switches.

## VMware Clusters

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

- Infrastructure Cluster: Infrastructure VMs (vCenter, Active Directory, DNS, DHCP, SQL Server, VMware Horizon Connection Server, VMware Horizon Replica Servers, VMware Horizon Composer Server and HyperFlex Data Platform Installer, etc.).
- HyperFlex Cluster: VMware Horizon RDSH VMs (Windows Server 2016) or Persistent/Non-Persistent VDI VM Pools (Windows 10 64-bit).

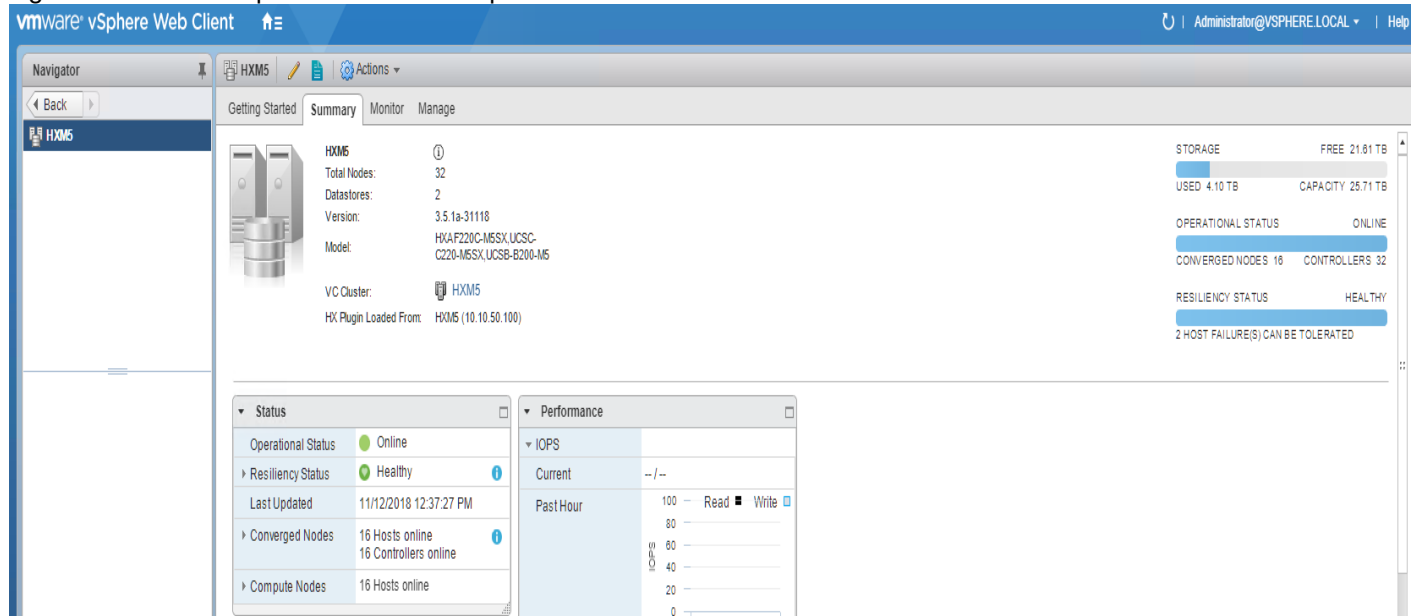


**HyperFlex release v3.0 or later supports 64 nodes in a single VMware cluster with 32 HXAF series HXAF220 or HXAF240 and 32 compute-only node. For more details, refer to:**

[https://www.cisco.com/c/en/us/td/docs/hyperconverged\\_systems/HyperFlex\\_HX\\_DataPlatformSoftware/Cisco\\_HXDataPlatform\\_RN\\_3\\_0.html](https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/Cisco_HXDataPlatform_RN_3_0.html).

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

**Figure 38 VMware vSphere Clusters on vSphere Web GUI**



## ESXi Host Design

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

### Virtual Networking Design

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

- vswitch-hx-inband-mgmt: This is the default vSwitch0 which is renamed by the ESXi kickstart file as part of the automated installation. The default vmkernel port, vmk0, is configured in the standard Management Network port group. The switch has two uplinks, active on fabric A and standby on fabric B, without jumbo frames. A second port group is created for the Storage Platform Controller VMs to connect to with their individual management interfaces. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere
- vswitch-hx-storage-data: This vSwitch is created as part of the automated installation. A vmkernel port, vmk1, is configured in the Storage Hypervisor Data Network port group, which is the interface used for connectivity to the HX Datastores 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 ESXi/vSphere
- vswitch-hx-vm-network: This vSwitch is created as part of the automated installation. The switch has two uplinks, active on both fabrics A and B, and without jumbo frames. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere
- vmotion: This vSwitch is created as part of the automated installation. The switch has two uplinks, active on fabric A and standby on fabric B, with jumbo frames required. The VLAN is not a Native VLAN as assigned to the vNIC template, and therefore assigned in ESXi/vSphere

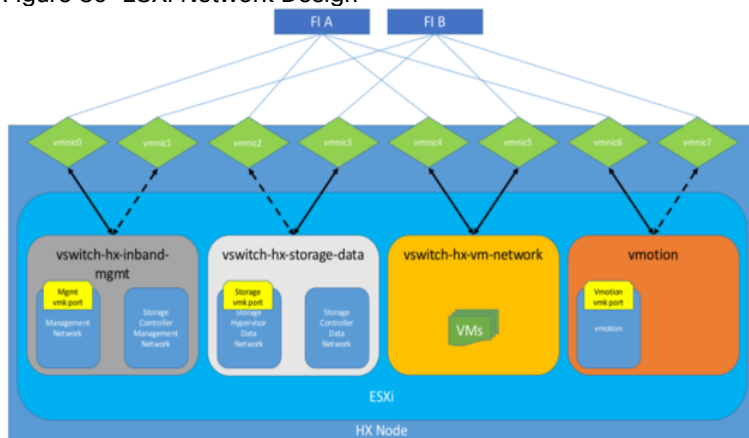
The following table and figures help give more details into the ESXi virtual networking design as built by the HyperFlex installer:

**Table 5 Table ESXi Host Virtual Switch Configuration**

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



Figure 39 ESXi Network Design



### VMDirectPath I/O Pass-through

VMDirectPath I/O allows a guest VM to directly access PCI and PCIe devices in an ESXi host as though they were physical devices belonging to the VM itself, also referred to as PCI pass-through. With the appropriate driver for the hardware device, the guest VM sends all I/O requests directly to the physical device, bypassing the hypervisor. In the Cisco HyperFlex system, the Storage Platform Controller VMs use this feature to gain full control of the Cisco 12Gbps SAS HBA cards in the Cisco HX-series rack-mount servers. This gives the controller VMs direct hardware level access to the physical disks installed in the servers, which they consume to construct the Cisco HX Distributed Filesystem. Only the disks connected directly to the Cisco SAS HBA or to a SAS extender, in turn connected to the SAS HBA are controlled by the controller VMs. Other disks, connected to different controllers, such as the SD cards, remain under the control of the ESXi hypervisor. The configuration of the VMDirectPath I/O feature is done by the Cisco HyperFlex installer and requires no manual steps.

### Storage Platform Controller Virtual Machines

A key component of the Cisco HyperFlex system is the Storage Platform Controller Virtual Machine running on each of the nodes in the HyperFlex cluster. The controller 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 vSphere ESXi agent, which is similar in concept to that of a Linux or Windows service. ESXi agents are tied to a specific host, they start and stop along with the ESXi hypervisor, and the system is not considered to be online and ready until both the hypervisor and the agents have started. Each ESXi hypervisor host has a single ESXi agent deployed, which is the controller VM for that node, and it cannot be moved or migrated to another host. The collective ESXi agents are managed via an ESXi agency in the vSphere 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 ESXi agents to the agency, therefore the ESXi hypervisors nor vCenter 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 vCenter server or appliance managing the vSphere cluster. The plugin communicates directly with the controller VMs to display the information requested, or make the configuration changes directed, all while operating within the same web-based interface of the vSphere Web Client. The deployment of the controller VMs, agents, agency, and vCenter plugin are all done by the Cisco HyperFlex installer and requires no manual steps.

### 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 an ESXi 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 VMDirectPath I/O. The configuration details of the models are as follows:

---

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

---

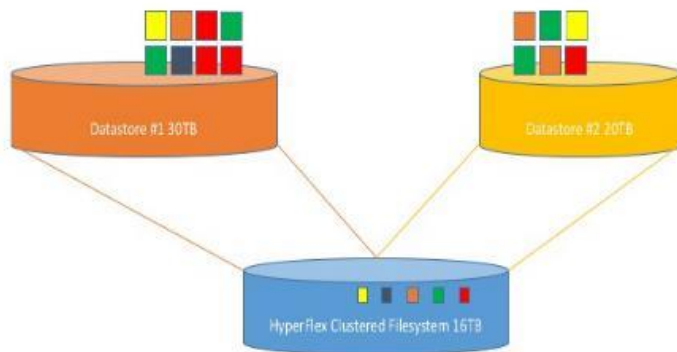
### Cisco HyperFlex Datastores

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

### Cisco HyperFlex Datastores

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

**Figure 40 Datastore Example**



### CPU Resource Reservations

Since the storage controller 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 ESXi hypervisor host are being heavily consumed by the guest VMs. Table 6 details the CPU resource reservation of the storage controller VMs:

**Table 6 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 ESXi hypervisor host are being heavily consumed by the guest VMs.

Table 7 lists the memory resource reservation of the storage controller VMs.

**Table 7 Controller VM Memory Reservations**

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

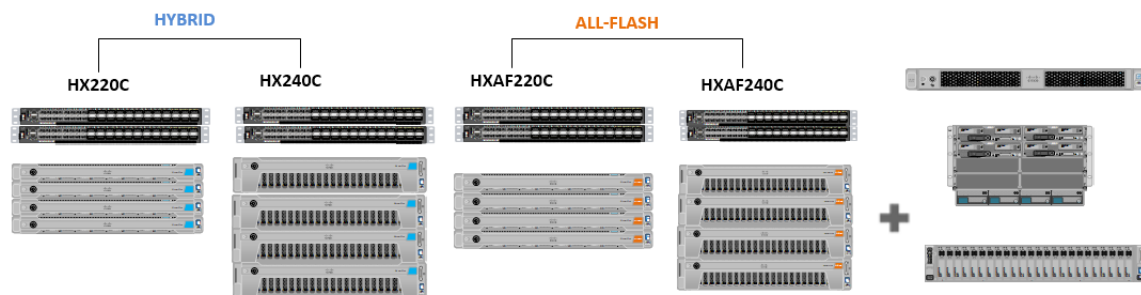


**The Cisco UCS compute-only Nodes have a lightweight storage controller VM; it is configured with only 1 vCPU and 512 MB of memory reservation.**

## Solution Configuration

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

**Figure 41 Configuration Topology for Scalable VMware Horizon 7 Workload with HyperFlex**



## Cisco UCS Compute Platform

The following subsections detail the physical connectivity configuration of the VMware Horizon 7 environment.

### Physical Infrastructure

#### Solution Cabling

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

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

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



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

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

**Table 8 Cisco Nexus 93180 A-Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 93180 A	Eth1/1	10GbE	Cisco Nexus 93180 B	Eth1/1
	Eth1/2	10GbE	Cisco Nexus 93180 B	Eth1/2
	Eth1/3	10GbE	Cisco Nexus 93180 B	Eth1/3

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/4	10GbE	Cisco Nexus 93180 B	Eth1/4
	Eth1/11	10GbE	Launcher-FI-A	Eth/29
	Eth1/12	10GbE	Launcher-FI-A	Eth/30
	Eth1/13	10GbE	Launcher-FI-B	Eth/29
	Eth1/14	10GbE	Launcher-FI-B	Eth/30
	MGMT0	GbE	GbE management switch	Any
	Eth1/15	10GbE	Infra-host-01	Port01
	Eth1/16	10GbE	Infra-host-02	Port01
	Eth1/49	40GbE	Cisco UCS fabric interconnect A	Eth1/35
	Eth1/50	40GbE	Cisco UCS fabric interconnect B	Eth1/35



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

**Table 9 Cisco Nexus 93180 B -Cabling Information**

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

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/49	40GbE	Cisco UCS fabric interconnect B	Eth1/36
	Eth1/50	40GbE	Cisco UCS fabric interconnect B	Eth1/36

**Table 10 Cisco UCS Fabric Interconnect A Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric interconnect A	Eth1/1	40GbE	Server 1	Port01
	Eth1/2	40GbE	Server 2	Port01
	Eth1/3	40GbE	Server 3	Port01
	Eth1/4	40GbE	Server 4	Port01
	Eth1/5	40GbE	Server 5	Port01
	Eth1/6	40GbE	Server 6	Port01
	Eth1/7	40GbE	Server 7	Port01
	Eth1/8	40GbE	Server 8	Port01
	Eth1/9	40GbE	Server 9	Port01
	Eth1/10	40GbE	Server 10	Port01
	Eth1/11	40GbE	Server 11	Port01
	Eth1/12	40GbE	Server 12	Port01
	Eth1/13	40GbE	Server 13	Port01
	Eth1/14	40GbE	Server 14	Port01
	Eth1/15	40GbE	Server 15	Port01
	Eth1/16	40GbE	Server 16	Port01
	Eth1/17	40GbE	Server 17	Port01
	Eth1/18	40GbE	Server 18	Port01
	Eth1/19	40GbE	Server 19	Port01

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/20	40GbE	Server 20	Port01
	Eth1/21	40GbE	Server 21	Port01
	Eth1/22	40GbE	Server 22	Port01
	Eth1/23	40GbE	Server 23	Port01
	Eth1/24	40GbE	Server 24	Port01
	Eth1/25	40GbE	5108 IOM Chassis-A	IOM-A/1/1
	Eth1/26	40GbE	5108 IOM Chassis-A	IOM-A/1/2
	Eth1/27	40GbE	Cisco Nexus 93180 A	Eth1/49
	Eth1/28	40GbE	Cisco Nexus 93180 B	Eth1/49
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2


**Table 11 Cisco UCS Fabric Interconnect B Cabling Information**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric interconnect B	Eth1/1	40GbE	Server 1	Port02
	Eth1/2	40GbE	Server 2	Port02
	Eth1/3	40GbE	Server 3	Port02
	Eth1/4	40GbE	Server 4	Port02
	Eth1/5	40GbE	Server 5	Port02
	Eth1/6	40GbE	Server 6	Port02
	Eth1/7	40GbE	Server 7	Port02
	Eth1/8	40GbE	Server 8	Port02
	Eth1/9	40GbE	Server 9	Port02



Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/10	40GbE	Server 10	Port02
	Eth1/11	40GbE	Server 11	Port02
	Eth1/12	40GbE	Server 12	Port02
	Eth1/13	40GbE	Server 13	Port02
	Eth1/14	40GbE	Server 14	Port02
	Eth1/15	40GbE	Server 15	Port02
	Eth1/16	40GbE	Server 16	Port02
	Eth1/17	40GbE	Server 17	Port02
	Eth1/18	40GbE	Server 18	Port02
	Eth1/19	40GbE	Server 19	Port02
	Eth1/20	40GbE	Server 20	Port02
	Eth1/21	40GbE	Server 21	Port02
	Eth1/22	40GbE	Server 22	Port02
	Eth1/23	40GbE	Server 23	Port02
	Eth1/24	40GbE	Server 24	Port02
	Eth1/25	40GbE	5108 IOM Chassis-B	IOM-B/1/1
	Eth1/26	40GbE	5108 IOM Chassis-B	IOM-B/1/2
	Eth1/35	40GbE	Cisco Nexus 93180 A	Eth1/50
	Eth1/36	40GbE	Cisco Nexus 93180 B	Eth1/50
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

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



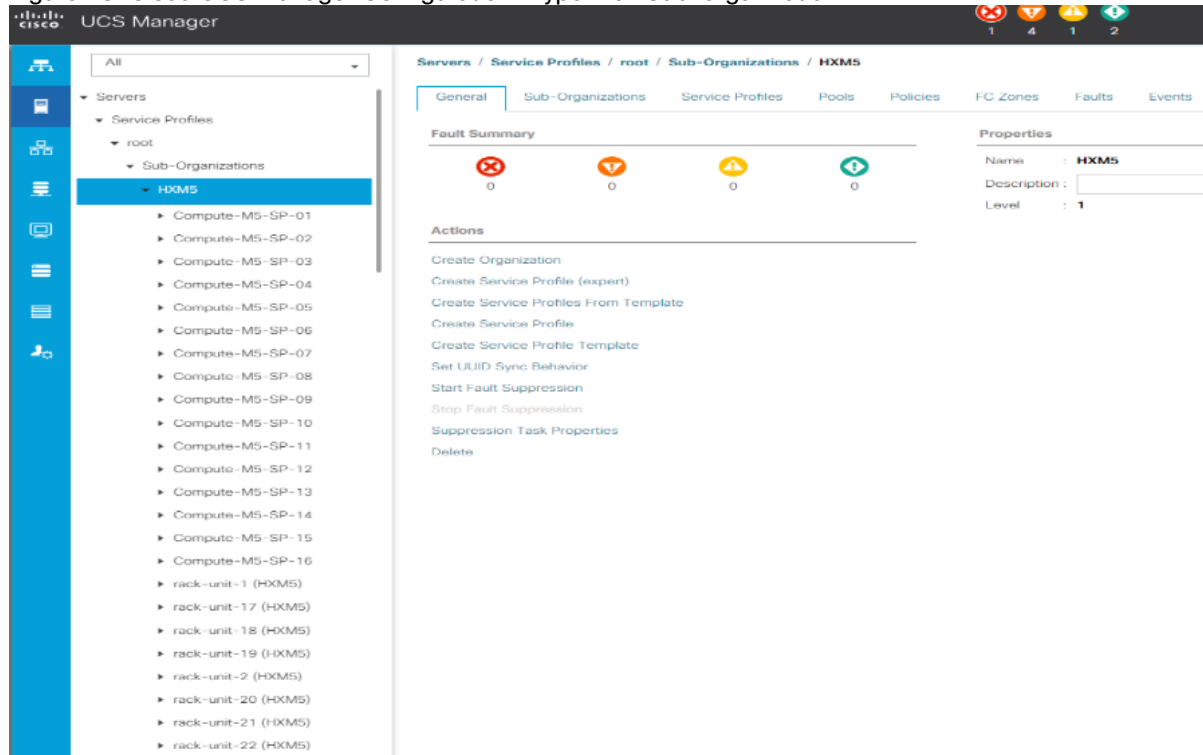
## Cisco Unified Computing System Configuration

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

For more information about racking, power, and installation of the chassis is described in the install guide (see [www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html](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 43 Cisco UCS Manager Configuration: HyperFlex Sub-organization

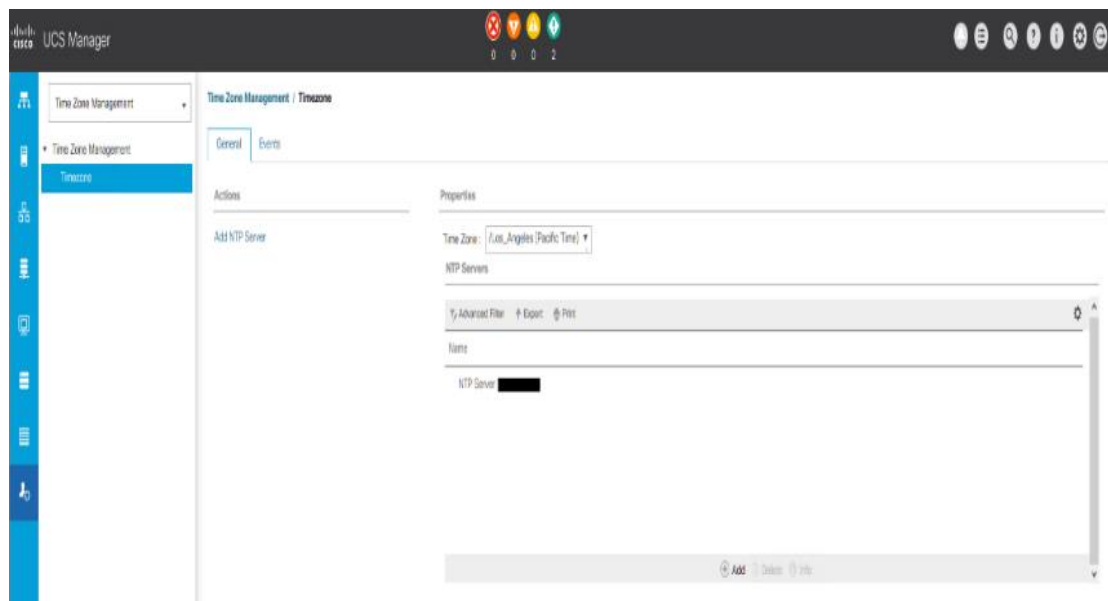


# Deploy and Configure HyperFlex Data Platform

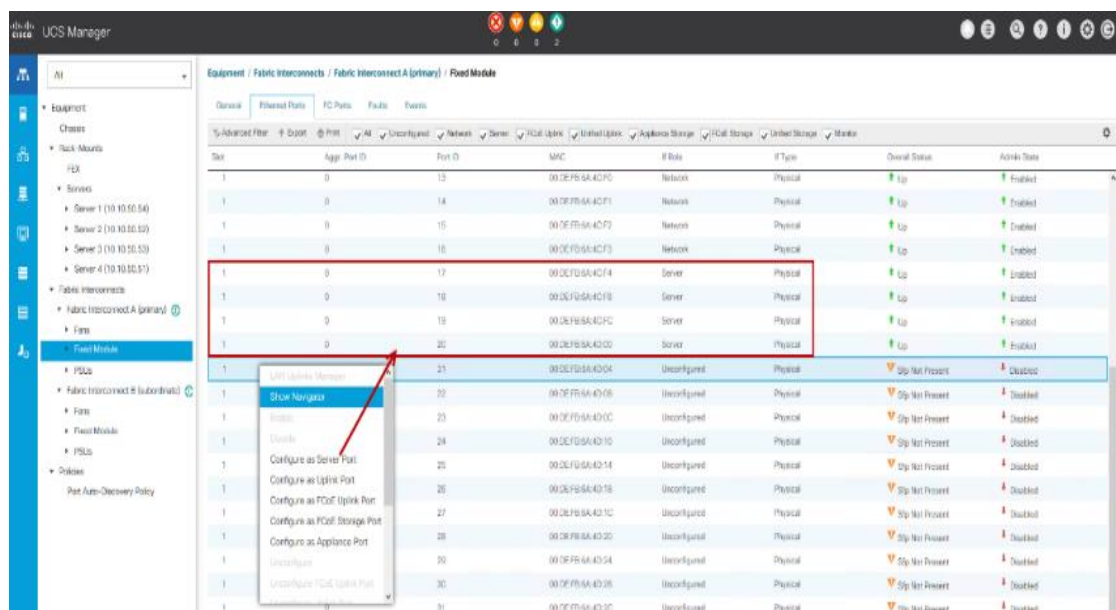
## Prerequisites

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

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

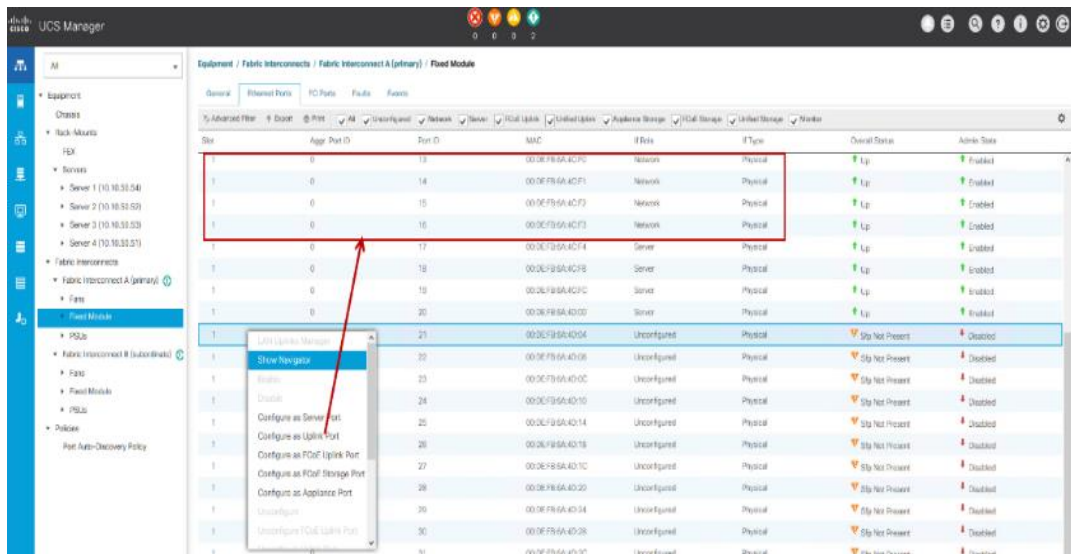


2. Configure Server Ports: Under the Equipment tab, Select Fabric A, select port to be configured as server port to manager HyperFlex rack server through Cisco UCS Manager.

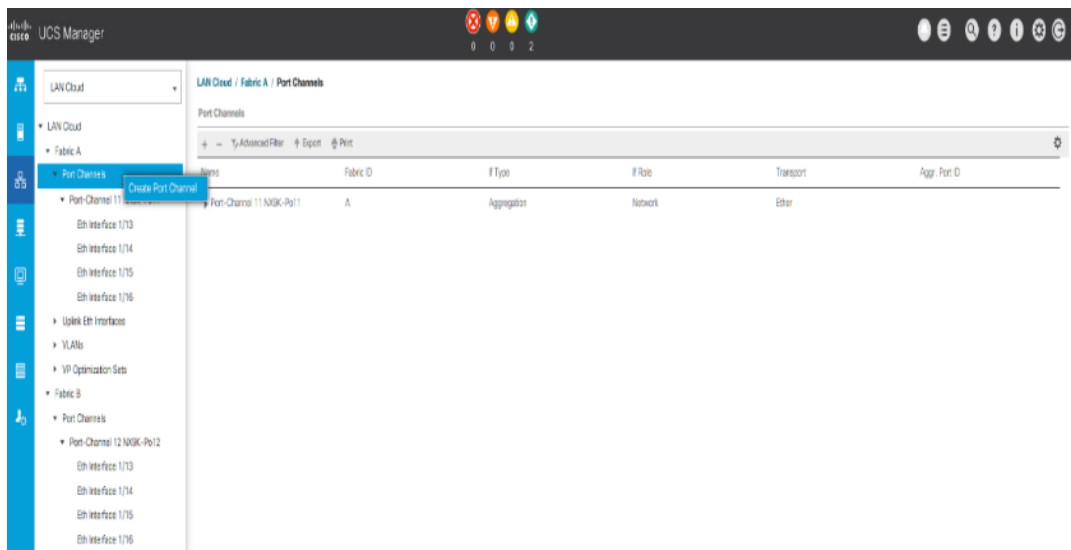


3. Repeat this step to configure server port on Fabric B.

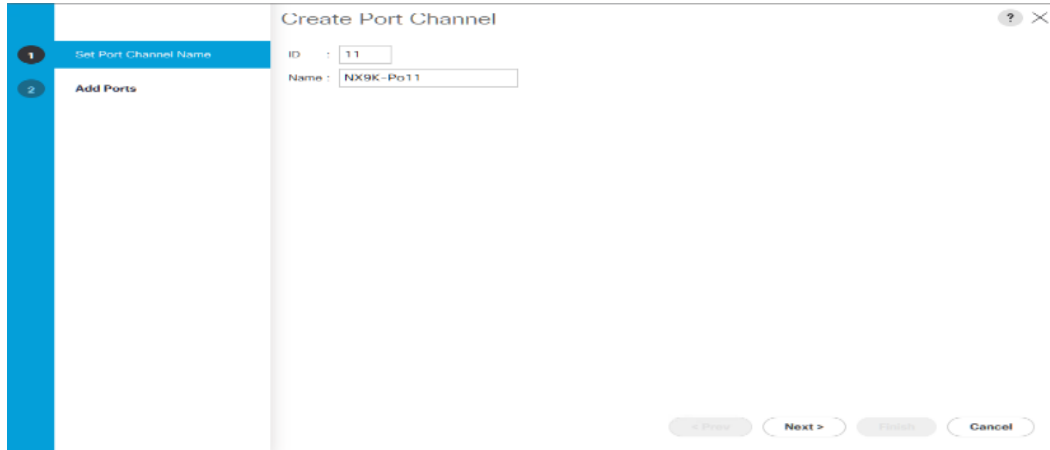
- Configure Uplink Ports: On Fabric A, Select port to be configured as uplink port for network connectivity to north bound switch.



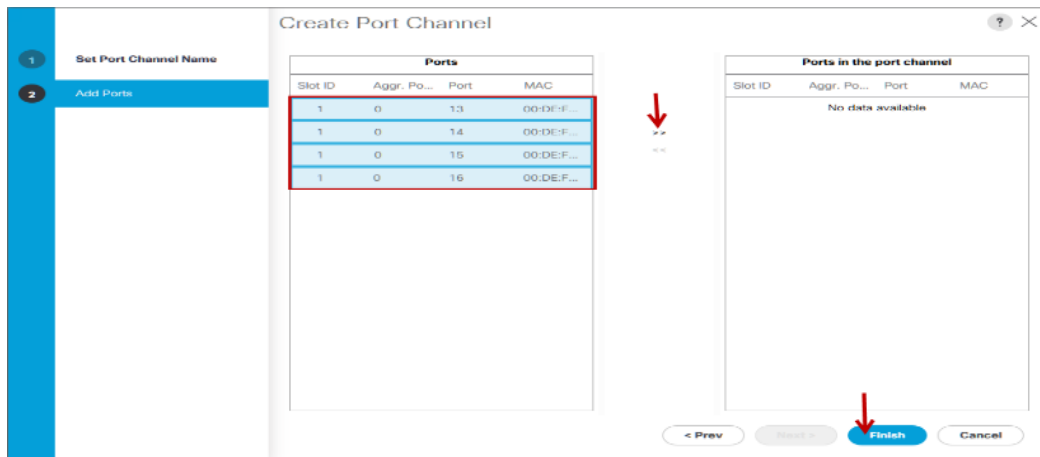
- Repeat this same on Fabric B.
- Create Port Channels: Under LAN tab, select expand LAN > LAN cloud > Fabric A. Right-click Port Channel.
- Select Create port-channel to connect with upstream switch as per Cisco UCS best practice. For our reference architecture, we connected a pair of Nexus 93180 switches.



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

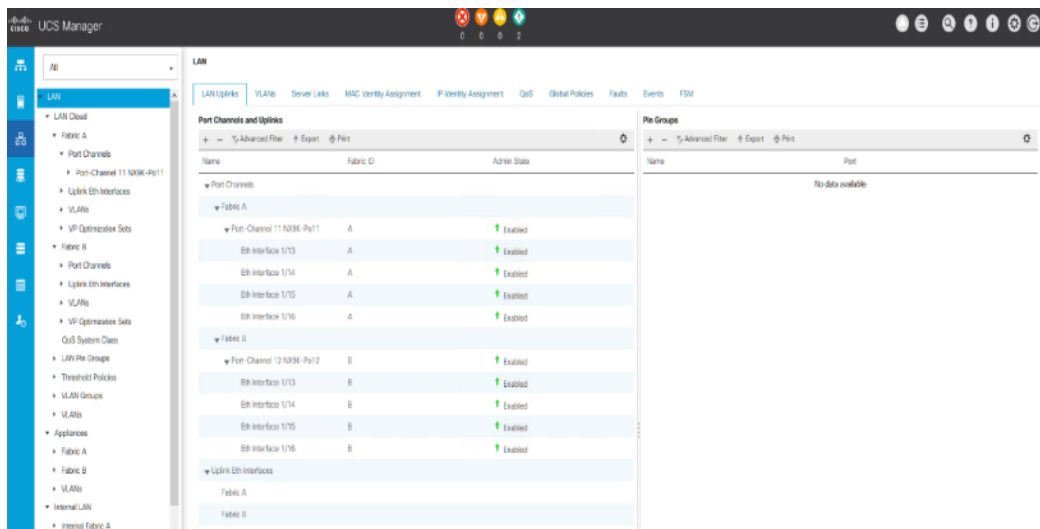


9. Select uplink ports to add as part of the port-channel.

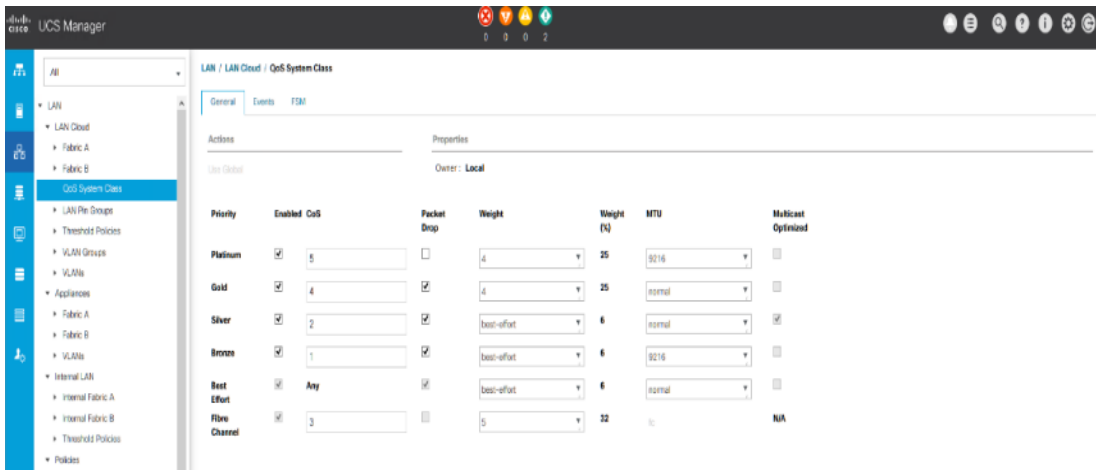


10. Click Finish.

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

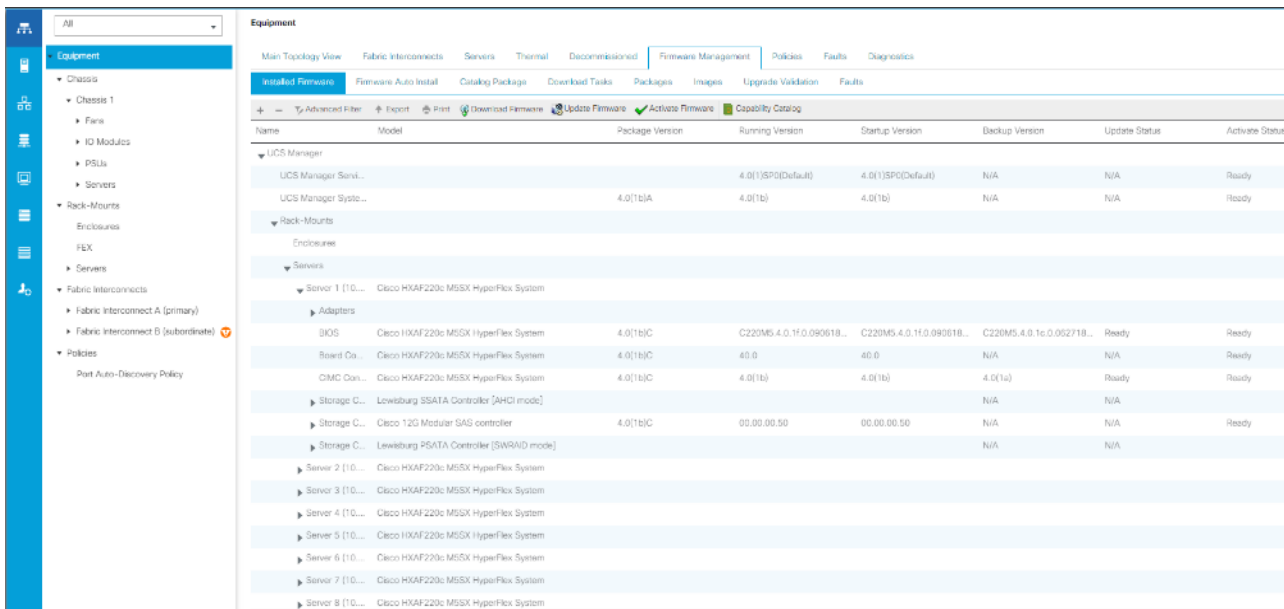


12. Configure QoS System Classes: From the LAN tab, under the Lan Cloud node, select QoS system class and configure the Platinum through Bronze system classes as shown in the following figure.
  - a. Set MTU to 9216 for Platinum (Storage data) and Bronze (vMotion)
  - b. Uncheck Enable Packet drop on the Platinum class
  - c. Set Weight for Platinum and Gold priority class to 4 and everything else as best-effort.
  - d. Enable multicast for silver class.



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

13. Verify the UCS Manager Software Version. In the Equipment tab, select Firmware Management > Installed Firmware.
14. Check and verify both Fabric Interconnects and Cisco USC Manager are configure with Cisco UCS Manager v4.0.1a.





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

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

## Deploy Cisco HyperFlex Data Platform Installer VM

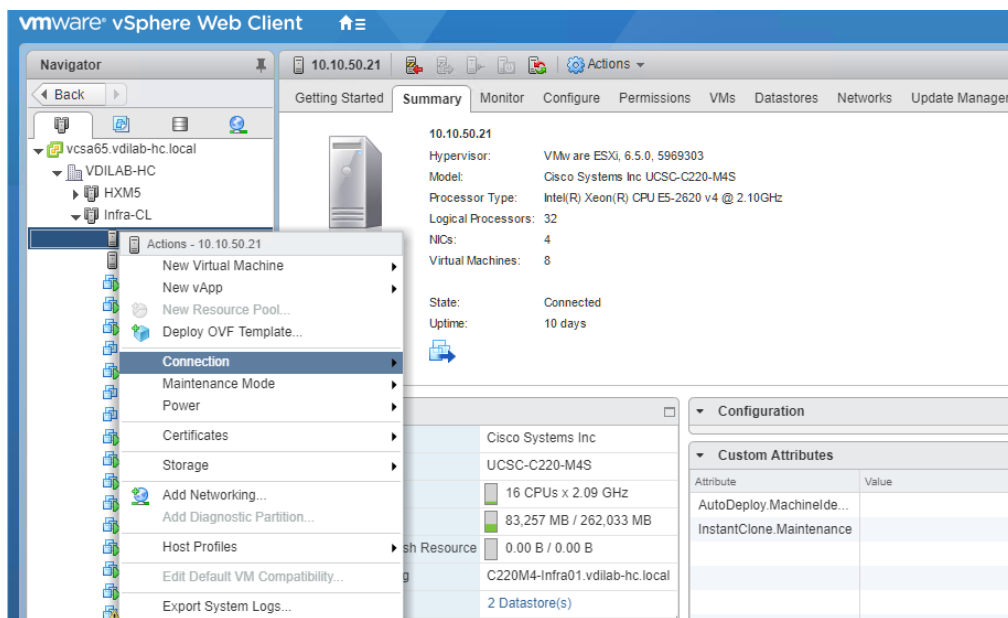
Download the latest installer OVA from Cisco.com. Software Download Link:

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

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

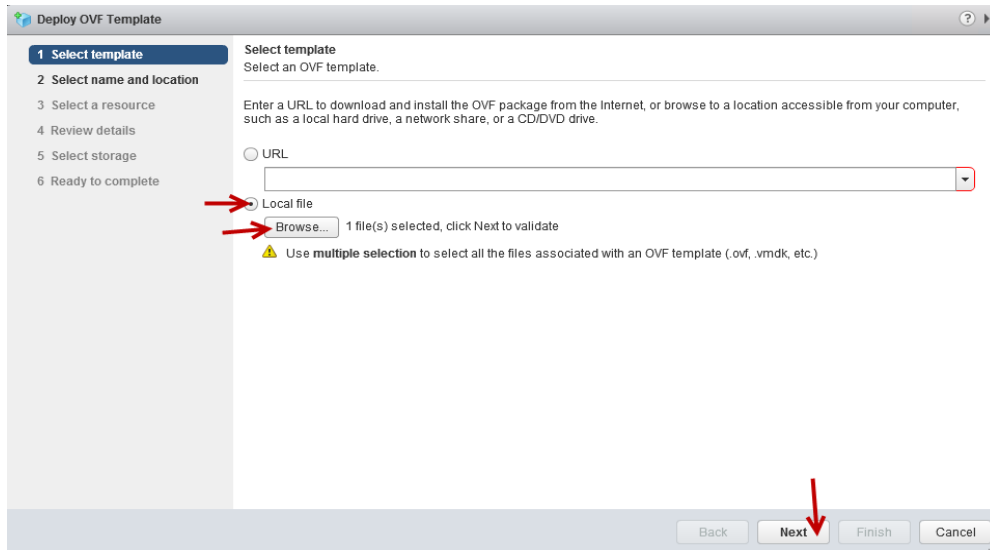
To deploy the OVA from the web client, complete the following steps:

- Log into vCenter web client via login to web browser with vCenter management IP address: <https://<FQDN>:9443/vcenter-client>
- Select ESXi host under hosts and cluster when HyperFlex data platform installer VM to deploy.
- Right-click ESXi host, select Deploy OVF Template.

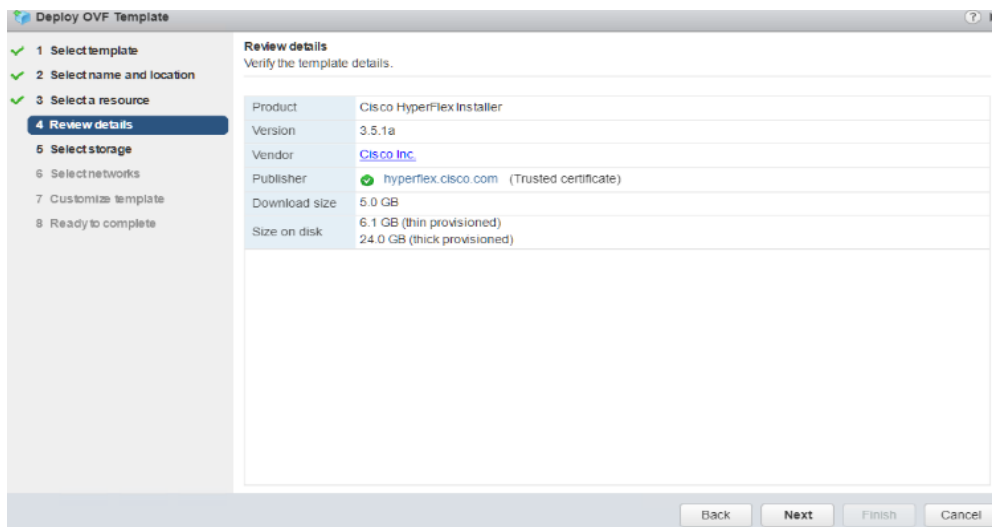
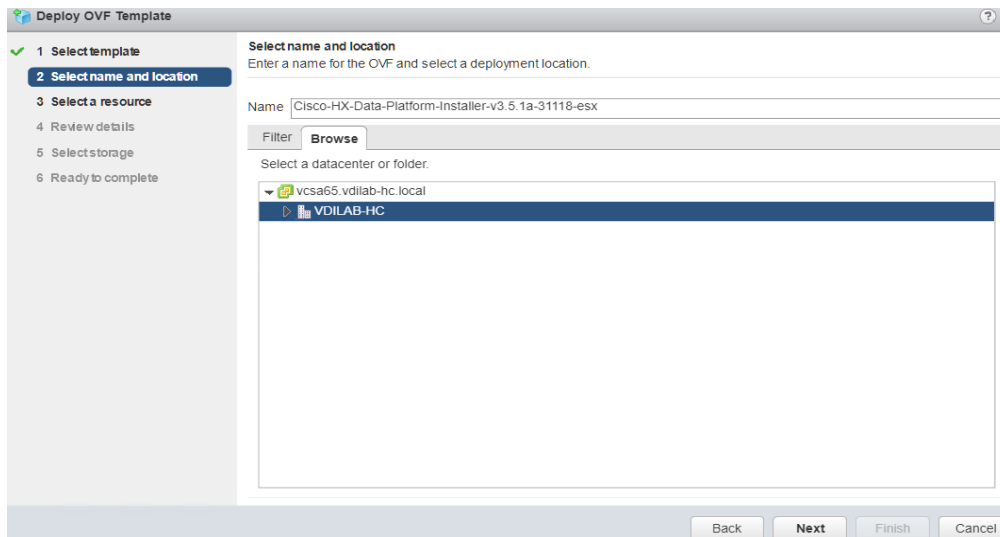


- Follow the deployment steps to configure HyperFlex data-platform installer VM deployment.
- Select OVA file to deploy, click Next.

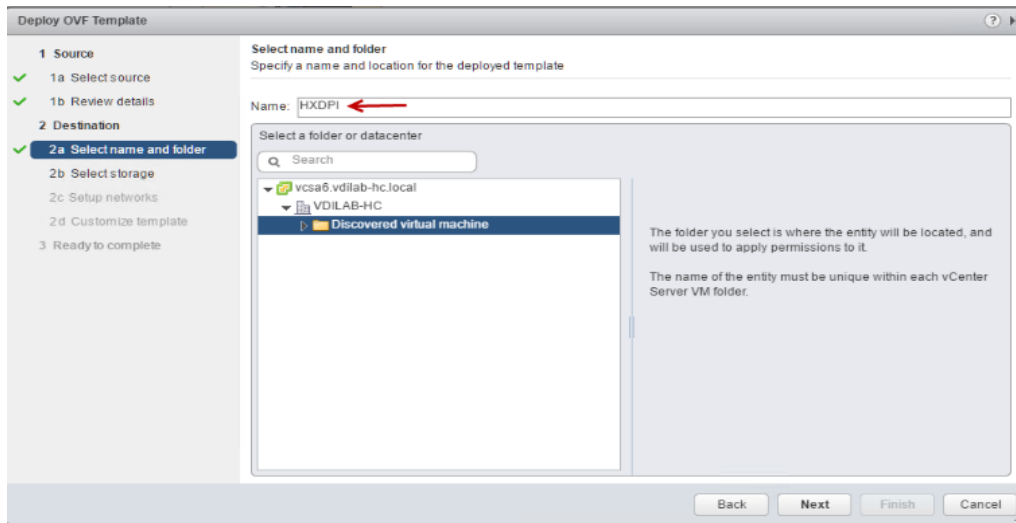




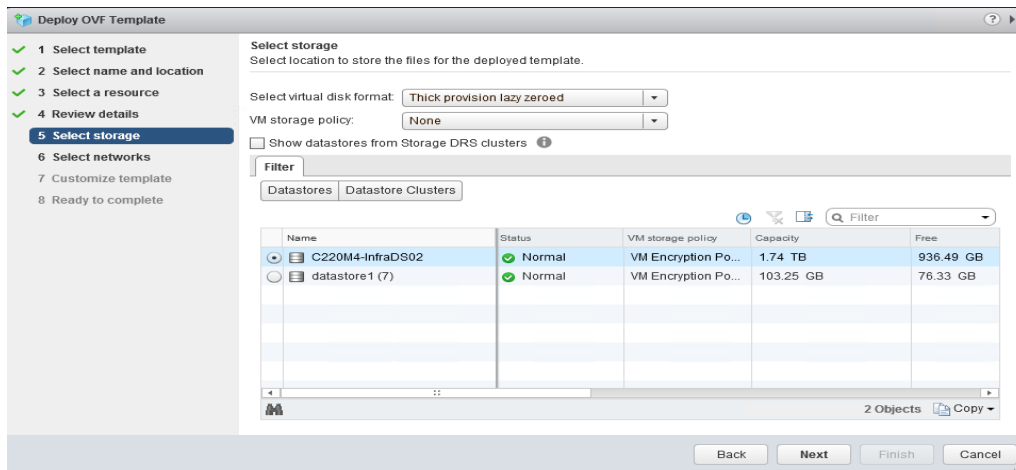
6. Review and verify the details for OVF template to deploy, click Next.



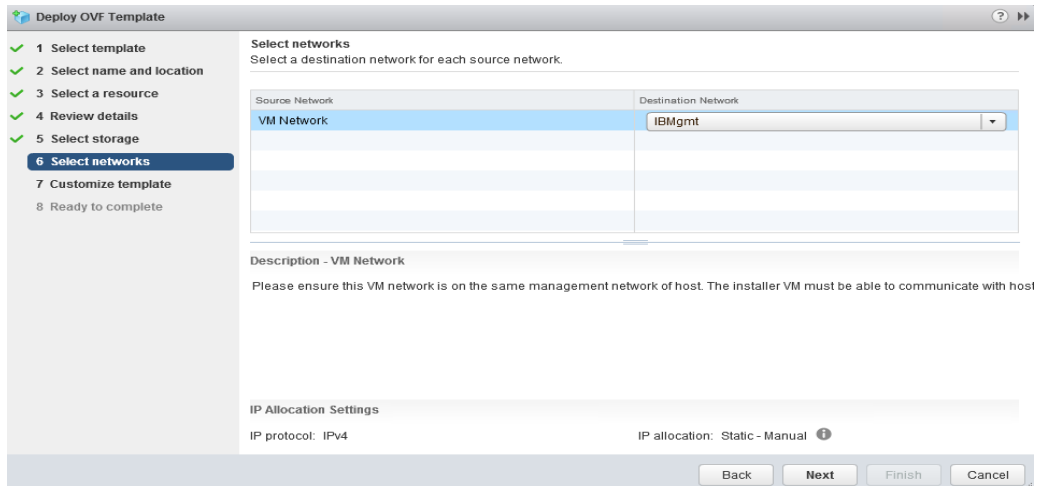
7. Enter name for OVF to template deploy, select datacenter and folder location. Click Next.



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



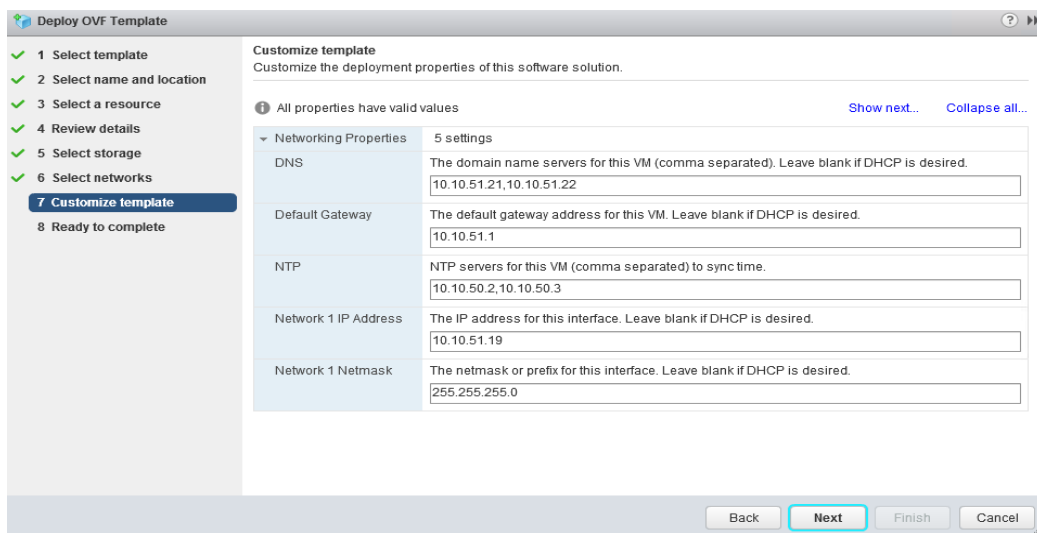
9. Select Network adapter destination port-group.



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



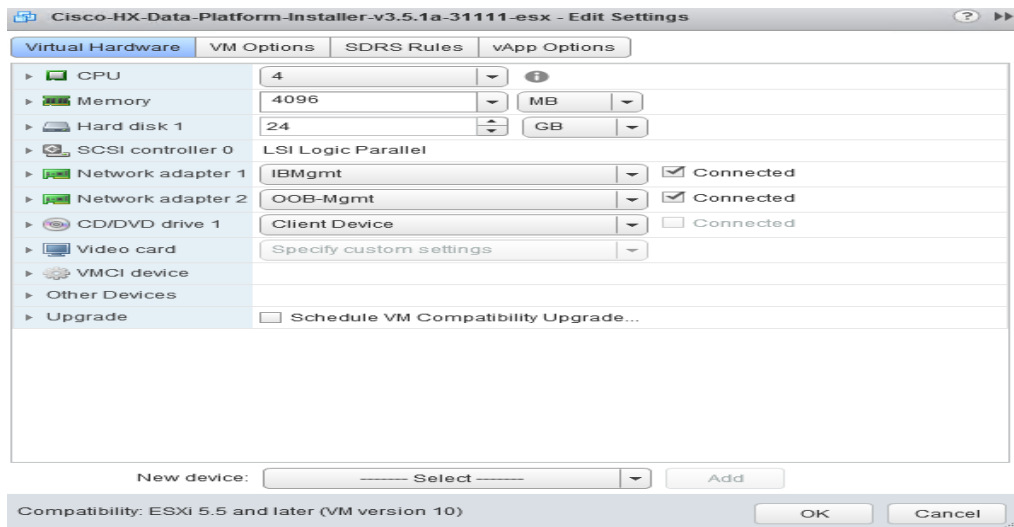
Provide a single DNS server only. Inputting multiple DNS servers will cause queries to fail. You must connect to vCenter to deploy the OVA file and provide the IP address properties. Deploying directly from an ESXi host will not allow you to set these values correctly.



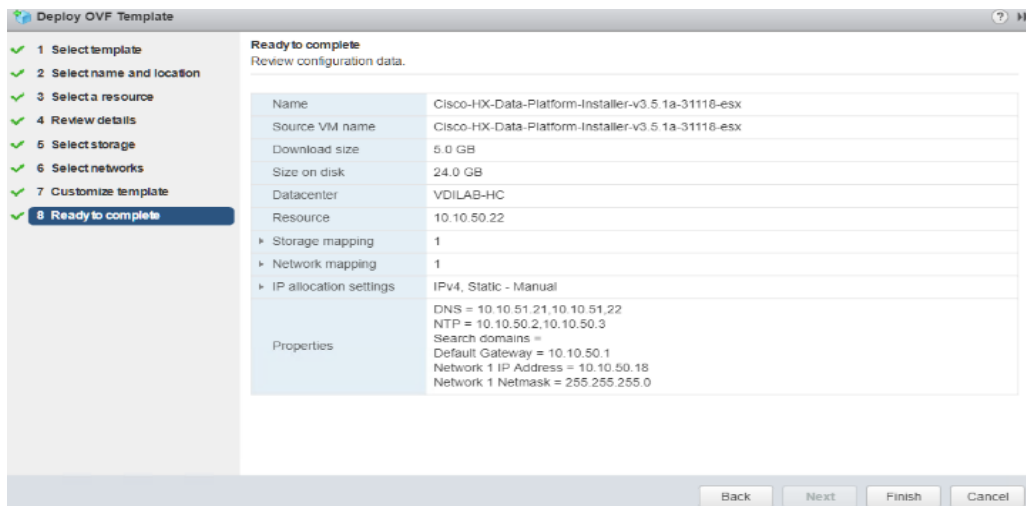
If you have internal firewall rules between these networks, please contact TAC for assistance.



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



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



The default credentials for the HyperFlex installer VM are: user name: root password: Cisco123

### Verify or Set DNS Resolution

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

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

```
root@Cisco-HX-Data-Platform-Installer:~# more /run/resolvconf/resolv.conf
# Dynamic resolv.conf(5) file for glibc resolver(3) generated by resolvconf(8)
```

```
# DO NOT EDIT THIS FILE BY HAND -- YOUR CHANGES WILL BE OVERWRITTEN
nameserver 10.10.51.21
nameserver 10.10.51.22
search vdilab-hc.local
```

## Cisco HyperFlex Cluster Configuration

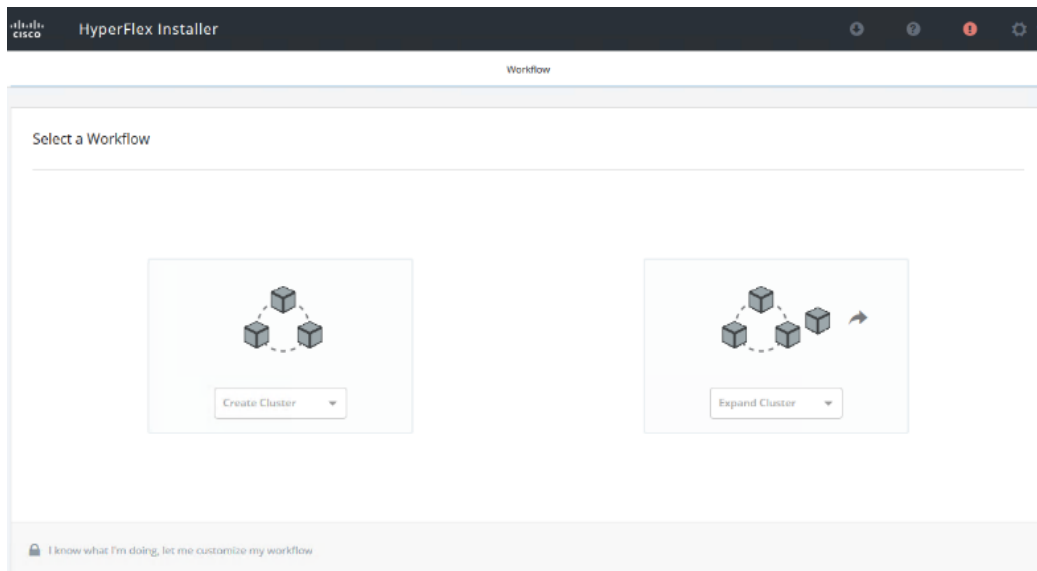
To configuring the Cisco HyperFlex Cluster, complete the following steps:

1. Login to HX Installer VM through a web browser: [http://<Installer\\_VM\\_IP\\_Address>](http://<Installer_VM_IP_Address>).

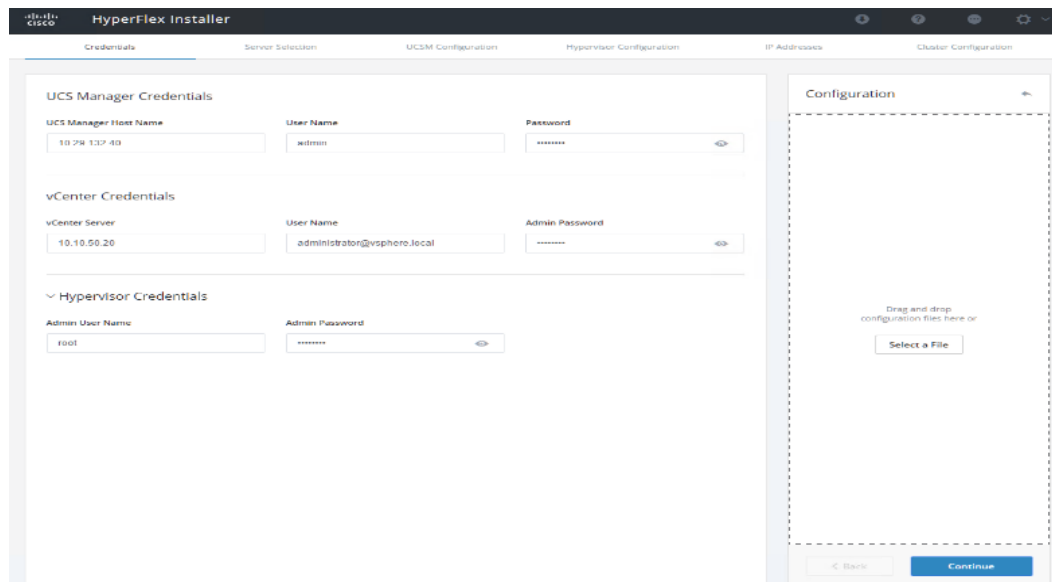


## Create a HyperFlex Cluster

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



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



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

The screenshot displays the HyperFlex Installer interface during the 'Server Selection' step. The main window shows a list of 24 servers, all currently 'unassociated'. The table includes columns for Server Name, Status, Model, Serial, Assoc State, and Actions. A 'Configure Server Ports' button is visible at the top right of the table area. To the right, a 'Configuration' panel is open, showing fields for 'Credentials' (UCS Manager Host Name, UCS Manager User Name, vCenter Server, User Name, Admin User name) and a 'Server Selection' list. The 'Server Selection' list shows a selection of servers with their respective IDs and models. At the bottom right of the configuration panel, there are 'Back' and 'Continue' buttons.



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


### Configure Server Ports (Optional)

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

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

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



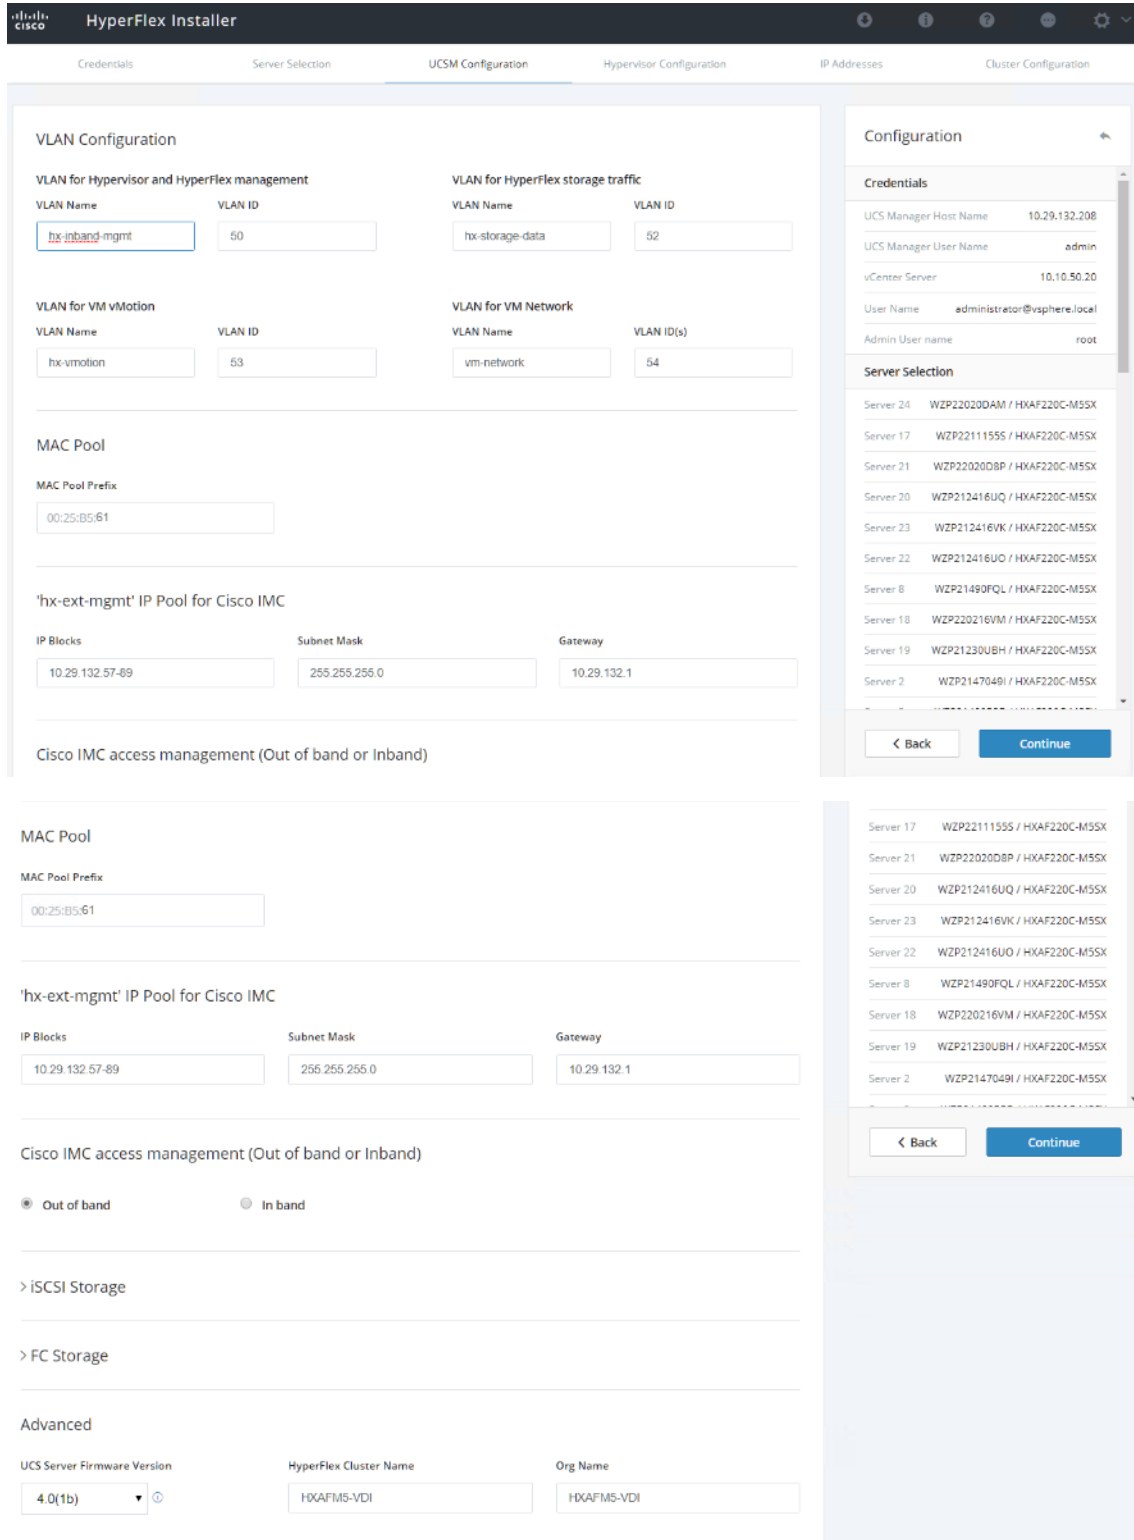


4. Enter the Details for the Cisco UCS Manager Configuration:
  - a. Enter VLAN ID for hx-inband-mgmt, hx-storage-data, hx-vmotion, vm-network.
  - b. MAC Pool Prefix: The prefix to use for each HX MAC address pool. Please select a prefix that does not conflict with any other MAC address pool across all Cisco UCS domains.
  - c. The blocks in the MAC address pool will have the following format:
    - $\{\text{prefix}\}:\{\text{fabric\_id}\}\{\text{vnic\_id}\}:\{\text{service\_profile\_id}\}$
    - The first three bytes should always be "00:25:B5".



**The first three bytes should always be "00:25:B5."**

5. Enter range of IP address to create a block of IP addresses for external management and access to CIMC/KVM.
6. Cisco UCS firmware version is set to 4.0 (1b) which is the required Cisco UCS Manager release for HyperFlex v3.5(1a) installation.
7. Enter HyperFlex cluster name.
8. Enter Org name to be created in Cisco UCS Manager.
9. Click Continue.



### Configure Hypervisor Settings

To configure the Hypervisor settings, complete the following steps:

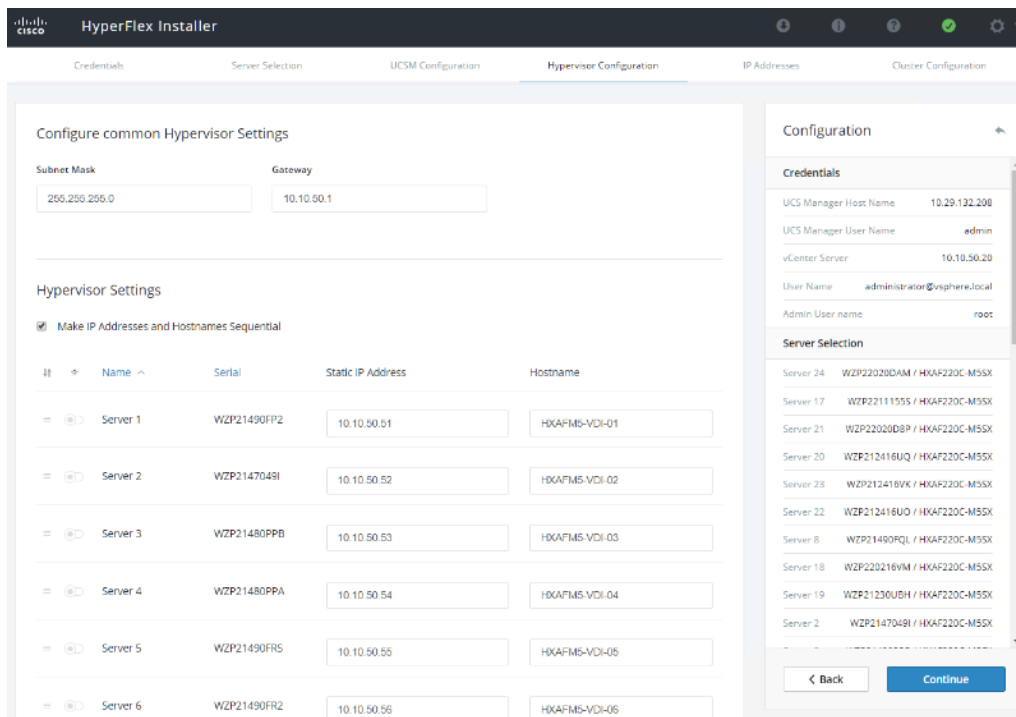
1. In the Configure common Hypervisor Settings section, enter:

- Subnet Mask
- Gateway
- DNS server(s)

2. In the Hypervisor Settings section:

- Select check box Make IP Address and Hostnames Sequential if they are following in sequence.
- Provide the starting IP Address.
- Provide the starting Host Name or enter Static IP address and Host Names manually for each node

3. Click Continue.



<input type="checkbox"/>	Server 6	WZP21490FR2	10.10.50.56	HXAFM5-VDI-06
<input type="checkbox"/>	Server 7	WZP21490FS4	10.10.50.57	HXAFM5-VDI-07
<input type="checkbox"/>	Server 8	WZP21490FQL	10.10.50.58	HXAFM5-VDI-08
<input type="checkbox"/>	Server 17	WZP2211155S	10.10.50.59	HXAFM5-VDI-09
<input type="checkbox"/>	Server 18	WZP220216VM	10.10.50.60	HXAFM5-VDI-10
<input type="checkbox"/>	Server 19	WZP21230UBH	10.10.50.61	HXAFM5-VDI-11
<input type="checkbox"/>	Server 20	WZP212416UQ	10.10.50.62	HXAFM5-VDI-12
<input type="checkbox"/>	Server 21	WZP22020D8P	10.10.50.63	HXAFM5-VDI-13
<input type="checkbox"/>	Server 22	WZP212416UO	10.10.50.64	HXAFM5-VDI-14
<input type="checkbox"/>	Server 23	WZP212416VK	10.10.50.65	HXAFM5-VDI-15
<input type="checkbox"/>	Server 24	WZP22020DAM	10.10.50.66	HXAFM5-VDI-16

Hypervisor Credentials

### IP Addresses

To add the IP addresses, complete the following steps:

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

Three additional columns appear on this page:

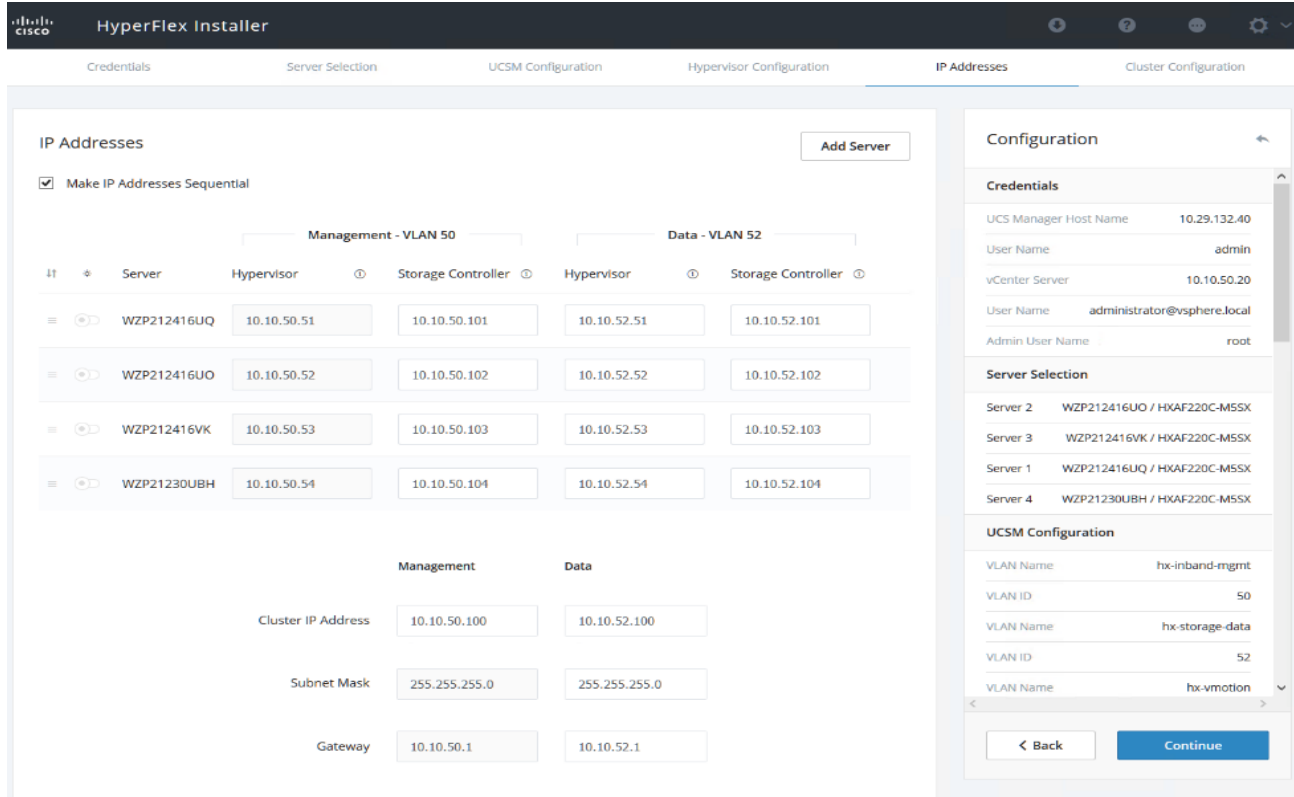
- Storage Controller/Management
- Hypervisor/Data
- Storage Controller/Data



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

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

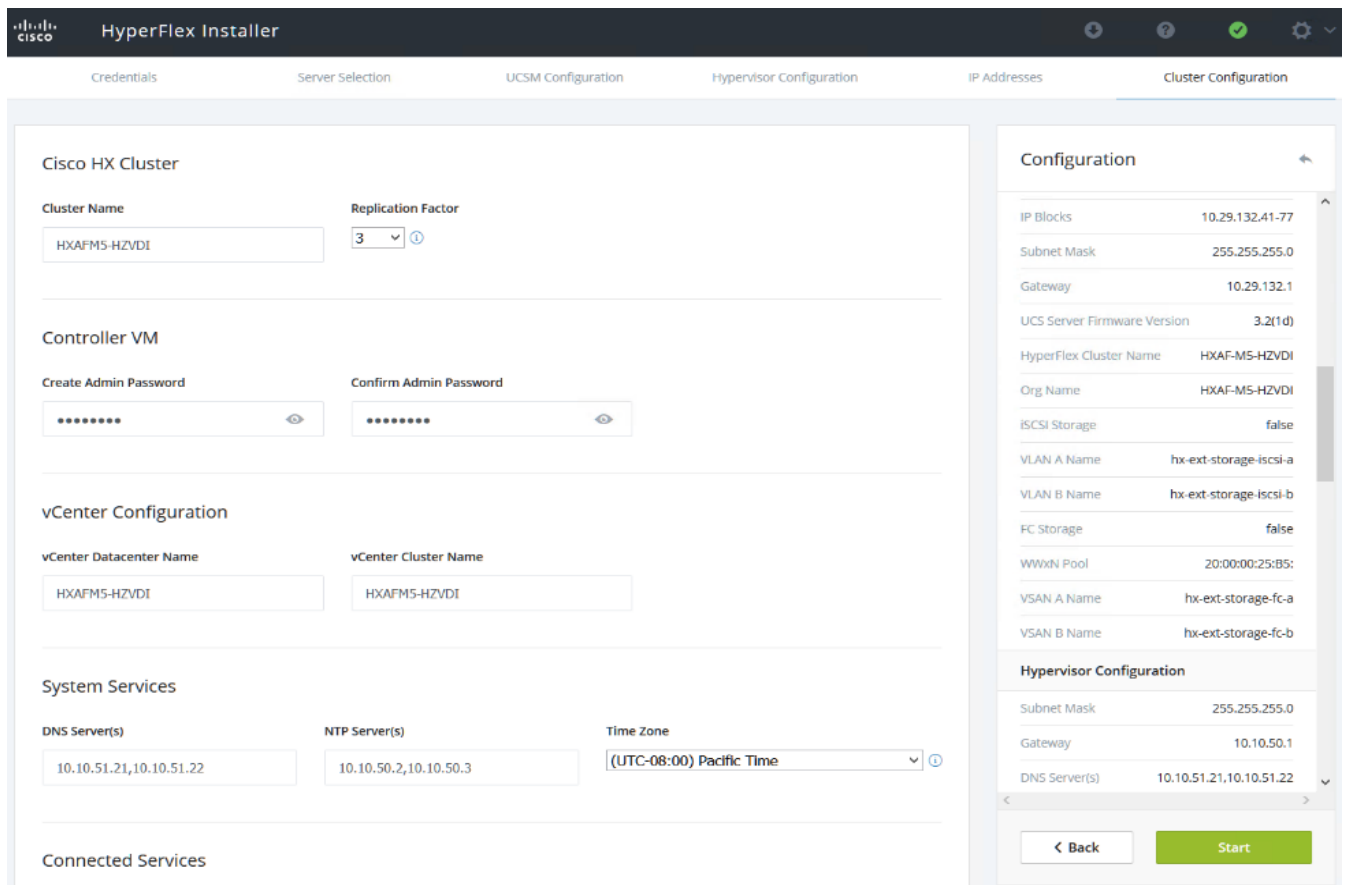
- Storage Controller/Data
2. Enter subnet and gateway details for the Management and Data subnets configured.
  3. Click Continue to proceed.



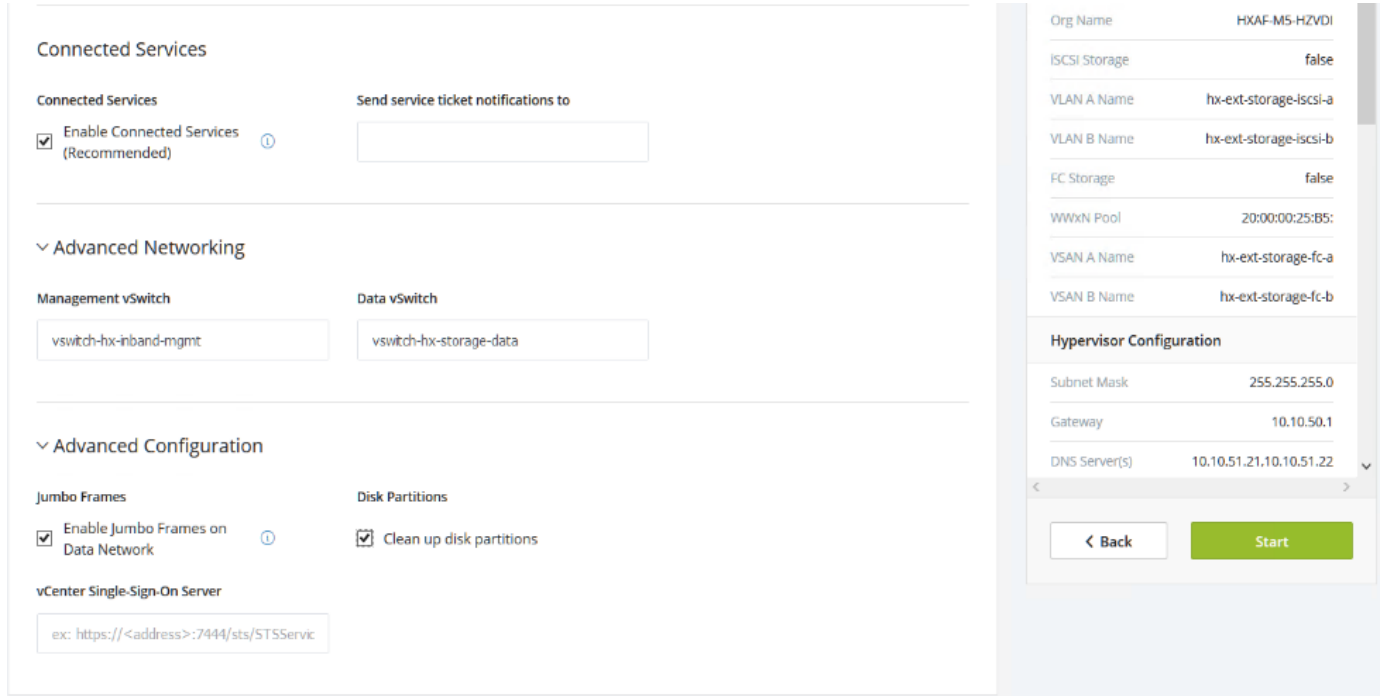
4. On the Cluster Configuration page, enter the following:

- Cluster Name
- Cluster management IP address
- Cluster data IP Address
- Set Replication Factor: 2 or 3
- Controller VM password
- vCenter configuration
  - vCenter Datacenter name
  - vCenter Cluster name
- System Services
  - DNS Server(s)
  - NTP Server(s)
  - Time Zone

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



- vCenter Single-Sign-On server



5. The configuration details can be exported to a JSON file by clicking the down arrow icon in the top right corner of the Web browser page as shown in the screenshot below.
6. Configuration details can be reviewed on Configuration page on right side section. Verify entered details for IP address entered in Credentials page, server selection for cluster deployment and creation workflow, Cisco UCS Manager configuration, Hypervisor Configuration, IP addresses.
7. Click Start after verifying details.

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

**Progress**

Start | **Config Installer** | Validations | UCSM Configuration | Hypervisor Configuration | Deploy Validation | Deploy | Create Validation | Cluster Creation

**Validations in Progress**

**Validations - Overall**  
In Progress

**UCSM Validation**

- ✓ Login to UCS API
- ✓ Inventorying physical servers
- ⌚ Validating the setup/environment

**Configuration**

**Credentials**

UCS Manager Host Name	10.29.132.208
UCS Manager User Name	admin
vCenter Server	10.10.50.20
User Name	administrator@vsphere.local
Admin User name	root

**Server Selection**

Server 24	WZP22020DAM / HXAF220C-M55X
Server 17	WZP22111555 / HXAF220C-M55X
Server 21	WZP22020D8P / HXAF220C-M55X
Server 20	WZP212416UQ / HXAF220C-M55X
Server 23	WZP212416VK / HXAF220C-M55X
Server 22	WZP212416UO / HXAF220C-M55X
Server 8	WZP21490FQL / HXAF220C-M55X
Server 18	WZP220216VM / HXAF220C-M55X
Server 19	WZP21230UBH / HXAF220C-M55X
Server 2	WZP2147049I / HXAF220C-M55X
Server 3	WZP21480PPB / HXAF220C-M55X
Server 1	WZP21490FP2 / HXAF220C-M55X
Server 6	WZP21490FR2 / HXAF220C-M55X



The screenshot displays the HyperFlex installer workflow. The progress bar shows the current step is 'Validations', which has a warning icon. Below the progress bar, a section titled 'Warnings found during Validations' contains a warning icon and two buttons: 'Retry Validations' and 'Skip Validations'. The main area shows a list of validation checks under 'Validations - Overall', with a 'Warning' label. All checks are marked with a green checkmark. At the bottom, a 'UCSM Validation' section shows a warning for 'QoS' with the message: 'QoS system class parameter(s) will be changed, which may require 6300 series Fabric Interconnect to reboot (both in cluster)'. On the right, a 'Configuration' panel lists various settings such as IP Blocks, Subnet Mask, Gateway, UCS Server Firmware Version, HyperFlex Cluster Name, Org Name, iSCSI Storage, VLAN A Name, VLAN B Name, FC Storage, WWxN Pool, VSAN A Name, and VSAN B Name. It also includes 'Hypervisor Configuration' for two servers, listing Subnet Mask, Gateway, DNS Server(s), Static IP Address, and Hostname. An 'Edit Configuration' button is located at the bottom of the configuration panel.



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

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

The screenshot displays the HyperFlex Installer interface. At the top, the title bar reads "HyperFlex Installer" with the Cisco logo on the left and navigation icons on the right. Below the title bar, a "Progress" section shows a sequence of steps: Start, Validations, UCSM Configuration, Hypervisor Configuration, Deploy Validation, Deploy, Create Validation, and Cluster Creation. The "UCSM Configuration" step is currently active and highlighted in blue.

Below the progress bar, a section titled "UCSM Configuration in Progress" contains a detailed list of tasks under the heading "UCSM Configuration - Overall". A blue "In Progress" button is visible. The tasks are as follows:

- ✓ Login to UCS API
- ✓ Inventory physical servers
- ✓ Validate UCS firmware version
- ✓ Setting flags for firmware validation
- ✓ Get inventory of firmware bundles
- ✓ Download firmware bundle
- ✓ Configure UCS Fabric Interconnect
- ✓ Configure FI Server Ports
- ✓ Configure QoS classes
- ✓ Configure org for the hx cluster
- ✓ Configure VLANs
- ✓ Configure Host Firmware policy
- ✓ Configure MAC address pools
- ✓ Configure QoS policies
- ✓ Configure Network Control policies
- Configure HyperFlex cluster
- Configure Adapter policies

On the right side of the interface, a "Configuration" panel is visible, containing several sections:

- Credentials:**
  - UCS Manager Host Name: 10.29.132.40
  - User Name: admin
  - vCenter Server: 10.10.50.20
  - User Name: administrator@vsphere.local
  - Admin User Name: root
- Server Selection:**
  - Server 2: WZP212416UO / HXAF220C-M5SX
  - Server 3: WZP212416VK / HXAF220C-M5SX
  - Server 1: WZP21230UBH / HXAF220C-M5SX
  - Server 4: WZP212416UQ / HXAF220C-M5SX
- UCSM Configuration:**
  - VLAN Name: hx-inband-mgmt, VLAN ID: 50
  - VLAN Name: hx-storage-data, VLAN ID: 52
  - VLAN Name: hx-vmotion, VLAN ID: 53
  - VLAN Name: vm-network, VLAN ID(s): 54
  - MAC Pool Prefix: 00:25:B5:23
  - IP Blocks: 10.29.132.41-77
  - Subnet Mask: 255.255.255.0
  - Gateway: 10.29.132.1
  - UCS Server Firmware Version: 3.2(2b)
  - HyperFlex Cluster Name: HXAF-M5-HZVDI

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

The screenshot shows the HyperFlex Installer interface. At the top, the title is "HyperFlex Installer" with the Cisco logo. Below the title bar, a "Progress" section shows a sequence of steps: Start, Validations, UCSM Configuration, Hypervisor Configuration (current), Deploy Validation, Deploy, Create Validation, and Cluster Creation. The "Hypervisor Configuration" step is highlighted with a blue circle and a progress indicator.

Below the progress bar, a section titled "Hypervisor Configuration in Progress" contains a sub-section "Hypervisor Configuration - Overall" with a status of "In Progress". A dropdown menu is set to "Hypervisor Configuration". The tasks listed are:

- ✓ Login to UCS API
- Configure static ip on the specified esxi servers
- Create threads to configure static ip on the esxi servers

On the right side, the "Configuration" panel is visible, containing two sections:

**Credentials**

UCS Manager Host Name	10.29.132.40
User Name	admin
vCenter Server	10.10.50.20
User Name	administrator@vsphere.local
Admin User Name	root

**Server Selection**

Server 16	FCH1938V085 / HXAF220C-M45
Server 10	FCH2033V1AD / HXAF220C-M45
Server 11	FCH1937V2JU / HXAF220C-M45
Server 14	FCH1842V1JG / HXAF220C-M45
Server 8	FCH1937V2TV / HXAF220C-M45
Server 9	FCH1937V2JV / HXAF220C-M45
Server 12	FCH2033V0LR / HXAF220C-M45
Server 15	FCH1937V2JT / HXAF220C-M45
Server 2	FCH2033V1E9 / HXAF220C-M45
Server 3	FCH2033V0HF / HXAF220C-M45

10. After a successful Hypervisor configuration, deploy validation task is performed which checks for required component and accessibility prior Deploy task is performed on Storage Controller VM.

**HyperFlex Installer**

Progress

Start | Validations | UCSM Configuration | Hypervisor Configuration | **Deploy Validation** | Deploy | Create Validation | Cluster Creation

Deploy Validation in Progress

Deploy Validation - Overall

10.10.50.60 **Succeeded**

- ✓ ESXi Management IP resolvability check
- ✓ ESXi Data IP resolvability check
- ✓ Controller Management IP resolvability check
- ✓ Controller Data IP resolvability check
- ✓ ESXi reachability check
- ✓ ESXi credential check
- ✓ Check for datastore inputs
- ✓ ESXi-Version
- ✓ Storage-HBA
- ✓ Storage-HBA-Count
- ✓ CPU-Threads
- ✓ HV-Support
- ✓ HyperThreading
- ✓ BootDisk-Adapter
- ✓ BootDisk-Size

**Configuration**

**Credentials**

UCS Manager Host Name	10.29.132.40
User Name	admin
vCenter Server	10.10.50.20
User Name	administrator@vSphere.local
Admin User Name	root

**Server Selection**

Server 16	FCH1938V085 / HXAF220C-M45
Server 10	FCH2033V1AD / HXAF220C-M45
Server 11	FCH1937V2JU / HXAF220C-M45
Server 14	FCH1842V1JG / HXAF220C-M45
Server 8	FCH1937V2TV / HXAF220C-M45
Server 9	FCH1937V2JV / HXAF220C-M45
Server 12	FCH2033V0LR / HXAF220C-M45
Server 15	FCH1937V2JT / HXAF220C-M45
Server 2	FCH2033V1E9 / HXAF220C-M45
Server 3	FCH2033V0HF / HXAF220C-M45
Server 13	FCH1937V2TS / HXAF220C-M45
Server 1	FCH2033V0BW / HXAF220C-M45
Server 6	FCH2031V054 / HXAF220C-M45
Server 7	FCH2033V0H8 / HXAF220C-M45
Server 4	FCH1936V0GE / HXAF220C-M45
Server 5	FCH2033V18F / HXAF220C-M45

**UCSM Configuration**

VLAN Name	hx-inband-mgmt
-----------	----------------

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

**HyperFlex Installer**

Progress

Start | Config Installer | Validations | UCSM Configuration | Hypervisor Configuration | Deploy Validation | **Deploy** | Create Validation | Cluster Creation

Deploy in Progress

Deploy - Overall Deploy

**10.10.50.51** In Progress

- ✓ Initializing Configuration
- ✓ Configuring CIMC server
- ⌚ Preparing ESXi Host for Installation  
Basic ESX Configuration.

**10.10.50.52** In Progress

- ✓ Initializing Configuration
- ✓ Configuring CIMC server
- ⌚ Preparing ESXi Host for Installation  
Get ESXi Version

**10.10.50.53** In Progress

- ✓ Initializing Configuration
- ✓ Configuring CIMC server
- ⌚ Preparing ESXi Host for Installation

**Configuration**

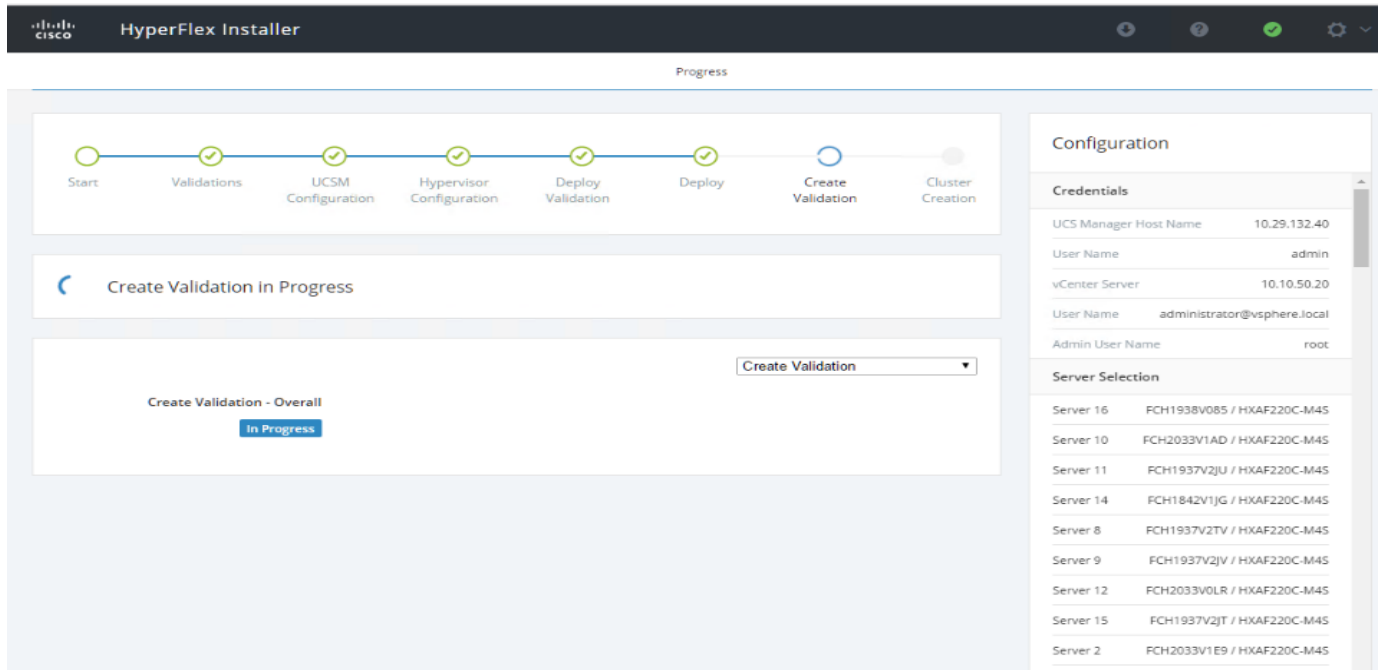
**Credentials**

UCS Manager Host Name	10.29.132.208
UCS Manager User Name	admin
vCenter Server	10.10.50.20
User Name	administrator@vsphere.local
Admin User name	root

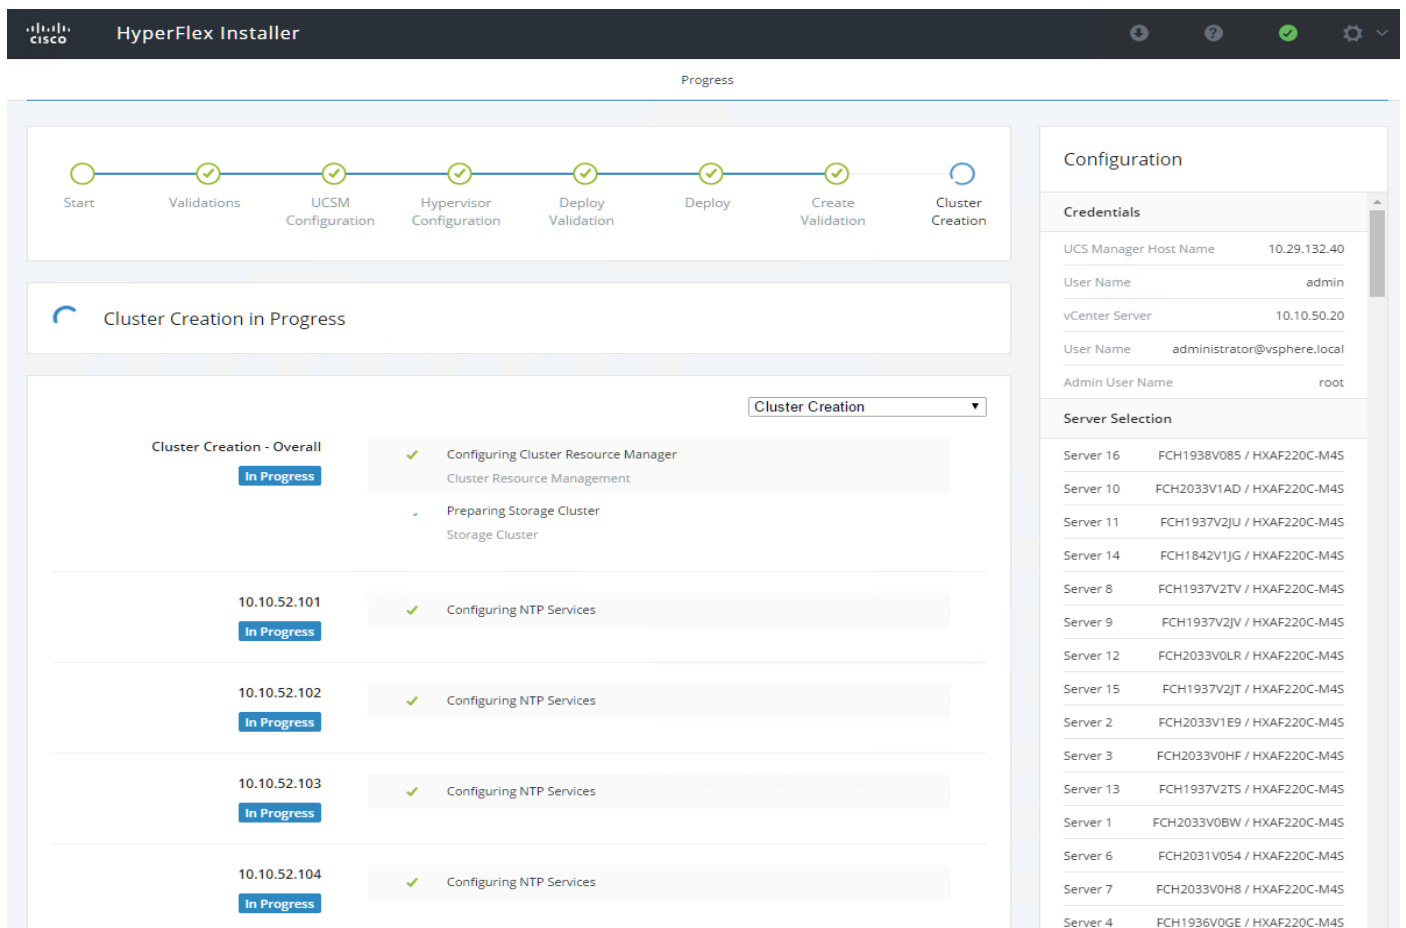
**Server Selection**

Server 24	WZP22020DAM / HXAF220C-M55X
Server 17	WZP22111555 / HXAF220C-M55X
Server 21	WZP22020D8P / HXAF220C-M55X
Server 20	WZP212416UQ / HXAF220C-M55X
Server 23	WZP212416VK / HXAF220C-M55X
Server 22	WZP212416UO / HXAF220C-M55X
Server 8	WZP21490FQL / HXAF220C-M55X
Server 18	WZP220216VM / HXAF220C-M55X
Server 19	WZP21230UBH / HXAF220C-M55X
Server 2	WZP2147049I / HXAF220C-M55X
Server 3	WZP21480PPB / HXAF220C-M55X
Server 1	WZP21490FP2 / HXAF220C-M55X
Server 6	WZP21490FR2 / HXAF220C-M55X
Server 7	WZP21490FS4 / HXAF220C-M55X
Server 4	WZP21480PPA / HXAF220C-M55X

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



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



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

The screenshot shows the Cisco HyperFlex Installer GUI in the Summary tab. The cluster name is HXMS, and its status is ONLINE and HEALTHY. The configuration parameters are as follows:

Version	3.5.1a-31118	vCenter Server	10.10.50.20
Cluster Management IP Address	10.10.50.100	vCenter Datacenter Name	VDILAB-HC
Cluster Data IP Address	10.10.52.100	vCenter Cluster Name	HXMS
Replication Factor	3	DNS Server(s)	10.10.51.21
Available Capacity	25.7 TB	NTP Server(s)	10.10.50.3, 10.10.50.2

Below the configuration parameters is a table of servers:

Model	Serial Number	Management Hypervisor	Management Storage Controller	Data Network Hypervisor	Data Network Storage Controller
HXAF220C-M5SX	WZP21490FP2	10.10.50.51	10.10.50.101	10.10.52.51	10.10.52.101
HXAF220C-M5SX	WZP2147049I	10.10.50.52	10.10.50.102	10.10.52.52	10.10.52.102
HXAF220C-M5SX	WZP21480PPB	10.10.50.53	10.10.50.103	10.10.52.53	10.10.52.103
HXAF220C-M5SX	WZP21480PPA	10.10.50.54	10.10.50.104	10.10.52.54	10.10.52.104
HXAF220C-M5SX	WZP21490FR5	10.10.50.55	10.10.50.105	10.10.52.55	10.10.52.105
HXAF220C-M5SX	WZP21490FR2	10.10.50.56	10.10.50.106	10.10.52.56	10.10.52.106
HXAF220C-M5SX	WZP21490FS4	10.10.50.57	10.10.50.107	10.10.52.57	10.10.52.107
HXAF220C-M5SX	WZP21490FQL	10.10.50.58	10.10.50.108	10.10.52.58	10.10.52.108
HXAF220C-M5SX	WZP22111555	10.10.50.59	10.10.50.109	10.10.52.59	10.10.52.109
HXAF220C-M5SX	WZP220216VM	10.10.50.60	10.10.50.110	10.10.52.60	10.10.52.110
HXAF220C-M5SX	WZP21230UBH	10.10.50.61	10.10.50.111	10.10.52.61	10.10.52.111
HXAF220C-M5SX	WZP212416UQ	10.10.50.62	10.10.50.112	10.10.52.62	10.10.52.112
HXAF220C-M5SX	WZP22020D8P	10.10.50.63	10.10.50.113	10.10.52.63	10.10.52.113
HXAF220C-M5SX	WZP212416UO	10.10.50.64	10.10.50.114	10.10.52.64	10.10.52.114
HXAF220C-M5SX	WZP212416VK	10.10.50.65	10.10.50.115	10.10.52.65	10.10.52.115
HXAF220C-M5SX	WZP22020DAM	10.10.50.66	10.10.50.116	10.10.52.66	10.10.52.116

## Cisco HyperFlex Cluster Expansion



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

### Prerequisites

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

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

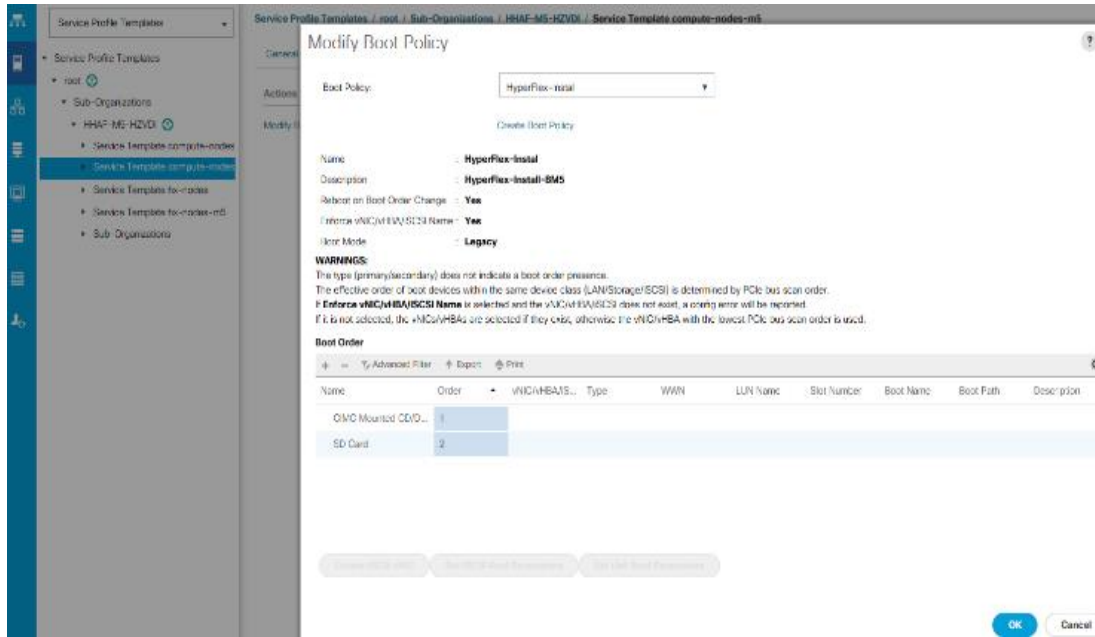
1. Login to Cisco UCS Manger.
2. Under “hx-cluster” sub-organization:
  - a. In the existing vMedia policy “HyperFlex” add vMedia mount details to boot ESXi image from data platform installer VM.
  - b. For Hostname/IP Address – Add IP address of data-platform installer VM which can also communicate with Cisco UCS Manager.

3. Change the existing service profile template to accommodate the new changes; install ESXi via vMedia policy.
4. In the existing service profile template “compute-nodes” select vMedia Policy tab.
5. Click Modify vMedia Policy.
6. From the drop-down list of vMedia Policy, select HyperFlex.

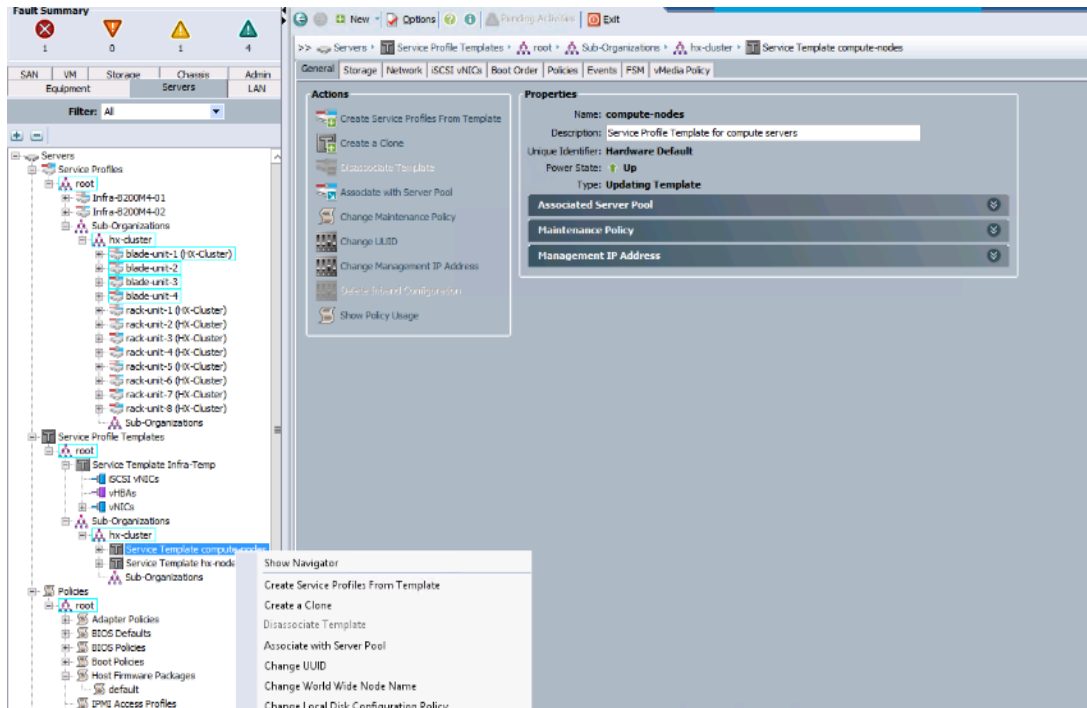
7. In the existing service profile template “compute-nodes” click Boot Order tab.
8. Click Modify Boot Policy.
9. From the drop-down list of Boot Policies, select HyperFlexInstall.




10. Save changes.



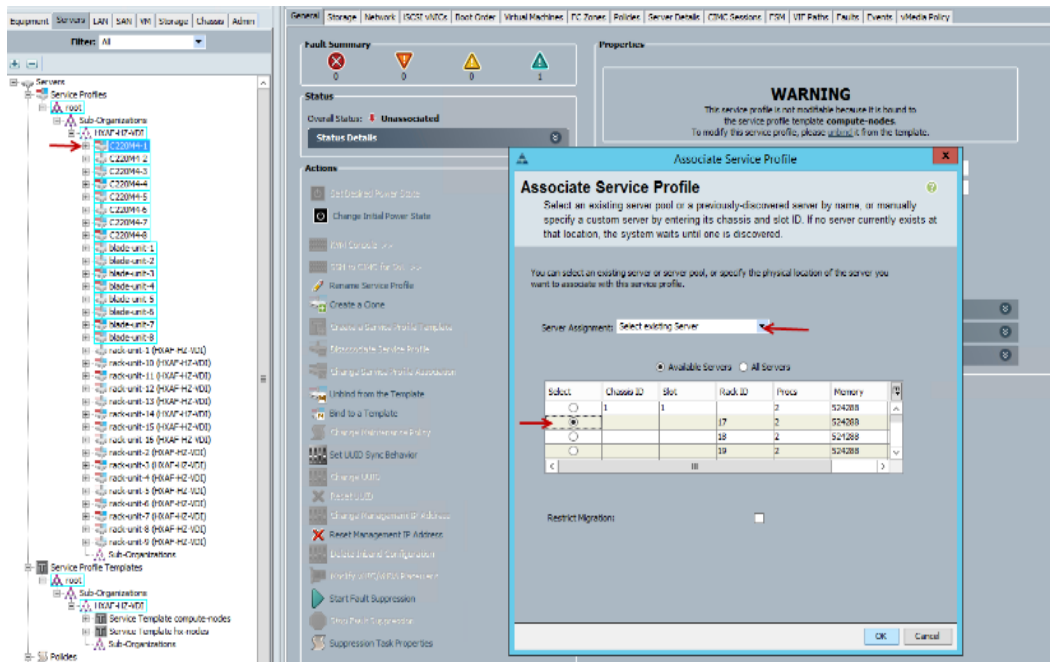
11. Create the service profile from the “compute-nodes” updating service profile template located in the HyperFlex cluster sub organization.



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



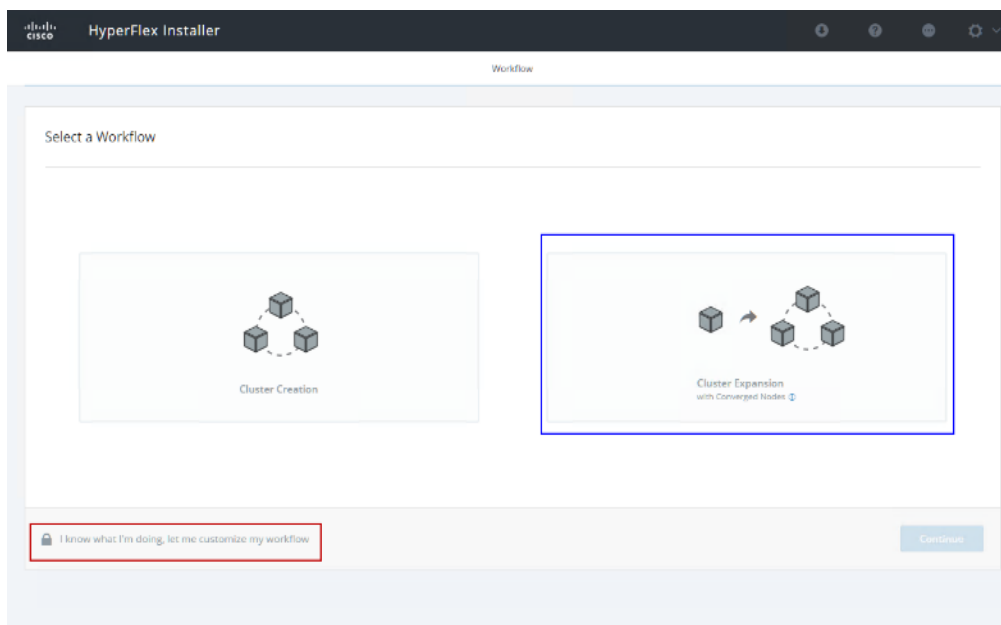
13. Click OK.



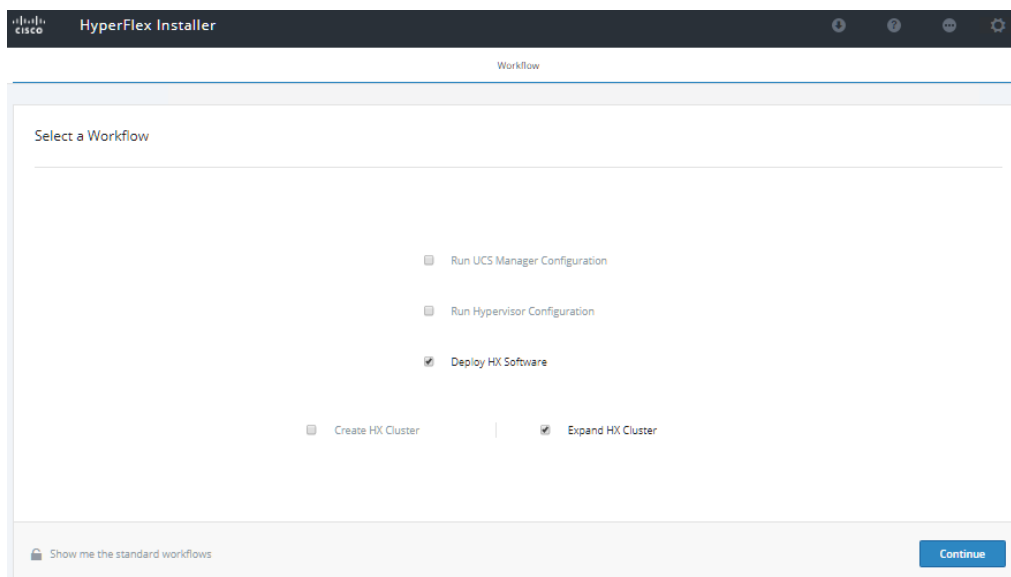
14. After the of ESXi install, assign the VLAN tag on the ESXi host; the static IP address configuration is located in the Configure Management Network section.



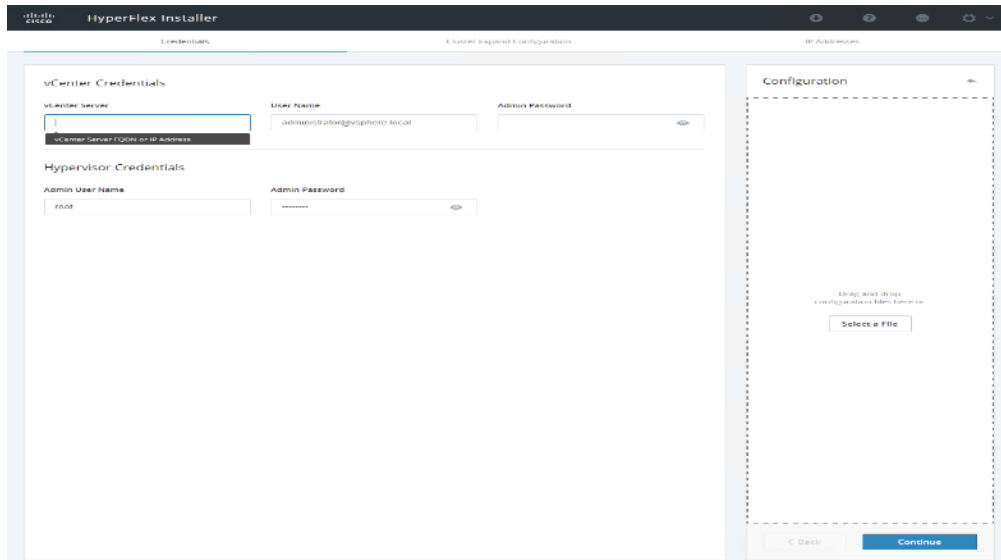
15. Log into the HyperFlex data platform installer WebUI. Click "I know what I'm doing, let me customize my workflow".



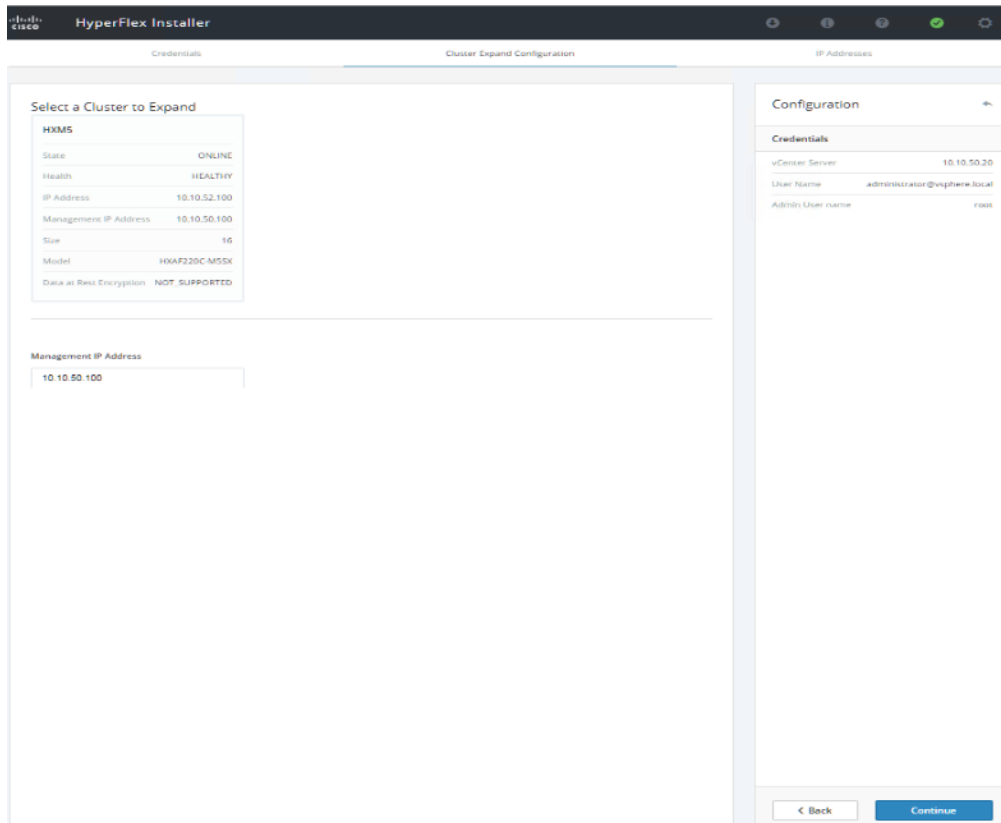
16. Select Deploy HX Software, Expand HX Cluster. Click Continue.



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



18. Select Cluster to expand, click Continue.



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

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

### IP Addresses

Add Compute Server   Add Converged Server

Make IP Addresses Sequential

Management - VLAN

Data - VLAN  
(FQDN or IP Address)

It	Name	Hypervisor	Storage Controller	Hypervisor	Storage Controller
=	Server 1 compute	10.10.50.75		10.10.52.75	X
=	Server 2 compute	10.10.50.76		10.10.52.76	X
=	Server 3 compute	10.10.50.77		10.10.52.77	X
=	Server 4 compute	10.10.50.78		10.10.52.78	X
=	Server 5 compute	10.10.50.79		10.10.52.79	X
=	Server 6 compute	10.10.50.80		10.10.52.80	X
=	Server 7 compute	10.10.50.81		10.10.52.81	X
=	Server 8 compute	10.10.50.82		10.10.52.82	X
=	Server 9 compute	10.10.50.83		10.10.52.83	X
=	Server 10 compute	10.10.50.84		10.10.52.84	X
=	Server 11 compute	10.10.50.85		10.10.52.85	X
=	Server 12 compute	10.10.50.86		10.10.52.86	X
=	Server 13 compute	10.10.50.87		10.10.52.87	X
=	Server 14 compute	10.10.50.88		10.10.52.88	X
=	Server 15 compute	10.10.50.89		10.10.52.89	X
=	Server 16 compute	10.10.50.90		10.10.52.90	X

Controller VM Password

### Configuration

#### Credentials

vCenter Server 10.10.50.20

User Name administrator@vsphere.local

Admin User name root

#### Cluster Expand Configuration

Management Cluster 10.10.50.100

← Back
Start

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

**Progress**

Start    Config Installer    **Deploy Validation**    Deploy    Expansion Validation    Cluster Expansion

Deploy Validation in Progress

Deploy Validation - Overall  
In Progress

**Configuration**

**Credentials**

vCenter Server	10.10.50.20
User Name	administrator@vsphere.local
Admin User name	root

**Cluster Expand Configuration**

Management Cluster	10.10.50.100
--------------------	--------------

**IP Addresses**

Server 1	
Management Hypervisor	10.10.50.75
Data Hypervisor	10.10.52.75
Server 2	
Management Hypervisor	10.10.50.76
Data Hypervisor	10.10.52.76
Server 3	
Management Hypervisor	10.10.50.85
Data Hypervisor	10.10.52.85
Server 4	
Management Hypervisor	10.10.50.86
Data Hypervisor	10.10.52.86
Server 5	
Management Hypervisor	10.10.50.87
Data Hypervisor	10.10.52.87
Server 6	
Management Hypervisor	10.10.50.88
Data Hypervisor	10.10.52.88
Server 7	
Management Hypervisor	10.10.50.89
Data Hypervisor	10.10.52.89

21. Performs deployment of HyperFlex controller VM create and deployment task.

**Progress**

Start    **Config Installer**    **Deploy Validation**    **Deploy**    Expansion Validation    Cluster Expansion

**Deploy in Progress**

**Deploy - Overall**    Deploy

**In Progress**

**10.10.50.75**    **In Progress**

- ✓ Initializing Configuration
- ✓ Configuring CIMC server
- ✓ Preparing ESXi Host for Installation
- ⚙ Configuring Hypervisor
  - Deploying Storage Controller VM

**10.10.50.76**    **In Progress**

- ✓ Initializing Configuration
- ✓ Configuring CIMC server
- ✓ Preparing ESXi Host for Installation
- ⚙ Configuring Hypervisor
  - Deploying Storage Controller VM

**10.10.50.77**    **In Progress**

- ✓ Initializing Configuration
- ✓ Configuring CIMC server
- ✓ Preparing ESXi Host for Installation
- ⚙ Configuring Hypervisor
  - Deploying Storage Controller VM

**Configuration**

Data Hypervisor 10.10.52.75

**Server 2**

Management Hypervisor 10.10.50.76

Data Hypervisor 10.10.52.76

**Server 3**

Management Hypervisor 10.10.50.85

Data Hypervisor 10.10.52.85

**Server 4**

Management Hypervisor 10.10.50.86

Data Hypervisor 10.10.52.86

**Server 5**

Management Hypervisor 10.10.50.87

Data Hypervisor 10.10.52.87

**Server 6**

Management Hypervisor 10.10.50.88

Data Hypervisor 10.10.52.88

**Server 7**

Management Hypervisor 10.10.50.89

Data Hypervisor 10.10.52.89

**Server 8**

Management Hypervisor 10.10.50.90

Data Hypervisor 10.10.52.90

**Server 9**

Management Hypervisor 10.10.50.77

Data Hypervisor 10.10.52.77

**Server 10**

Management Hypervisor 10.10.50.78

22. Performs expansion validation.

**HyperFlex Installer** Progress

Expansion Validation in Progress

Expansion Validation - Overall In Progress

**Configuration**

**Credentials**

vCenter Server	10.10.50.20
User Name	administrator@vsphere.local
Admin User Name	root

**Cluster Expand Configuration**

Management Cluster	10.10.50.100
--------------------	--------------

**IP Addresses**

Server 0	
Management Hypervisor	10.10.50.67
Data Hypervisor	10.10.52.27
Server 1	
Management Hypervisor	10.10.50.68
Data Hypervisor	10.10.52.28
Server 2	
Management Hypervisor	10.10.50.77

**HyperFlex Installer** Progress

Cluster Expansion in Progress

Cluster Expansion - Overall In Progress

VM Name	Role	Status
stCtIVM-FCH19377PVA	StNode Controller VM	In Progress
stCtIVM-FCH19057QG2	StNode Controller VM	In Progress
stCtIVM-FCH2033V1AN	StNode Controller VM	In Progress
stCtIVM-FCH2033V0EX	StNode Controller VM	In Progress
stCtIVM-FCH2033V0FP	StNode Controller VM	In Progress
stCtIVM-FCH2033V18H	StNode Controller VM	In Progress

**Configuration**

**Credentials**

vCenter Server	10.10.50.20
User Name	administrator@vsphere.local
Admin User Name	root

**Cluster Expand Configuration**

Management Cluster	10.10.50.100
--------------------	--------------

**IP Addresses**

Server 0	
Management Hypervisor	10.10.50.67
Data Hypervisor	10.10.52.27
Server 1	
Management Hypervisor	10.10.50.68
Data Hypervisor	10.10.52.28
Server 2	
Management Hypervisor	10.10.50.77
Data Hypervisor	10.10.52.37
Server 3	
Management Hypervisor	10.10.50.78
Data Hypervisor	10.10.52.38
Server 4	
Management Hypervisor	10.10.50.79
Data Hypervisor	10.10.52.39
Server 5	
Management Hypervisor	10.10.50.80
Data Hypervisor	10.10.52.40
Server 6	



23. Summary of Expansion cluster workflow performed.

The screenshot displays the Cisco HyperFlex Installer Summary page. At the top, the progress bar shows five steps: Start, Deploy Validation, Deploy, Expansion Validation, and Cluster Expansion, all marked with green checkmarks. Below the progress bar, a green checkmark and the text "Cluster Expansion Successful" are displayed, along with a "View Summary" button. The main content area is titled "Cluster Expansion - Overall" and shows a "Succeeded" status. A dropdown menu is set to "Cluster Expansion". Below this, four controller VMs are listed, each with a "Succeeded" status and a list of components that were successfully installed:

Controller VM Name	Component	Status
stCtlVM-FCH19377PVA	StNode Controller VM	✓ Succeeded
	Mount	✓ Succeeded
stCtlVM-FCH19057QG2	StNode Controller VM	✓ Succeeded
	Mount	✓ Succeeded
stCtlVM-FCH2033V1AN	StNode Controller VM	✓ Succeeded
	Mount	✓ Succeeded
stCtlVM-FCH2033V0EX	StNode Controller VM	✓ Succeeded
	Mount	✓ Succeeded

On the right side of the interface, the "Configuration" panel is visible, showing the following details:

- Credentials:** vCenter Server (10.10.50.20), User Name (administrator@vsphere.local), Admin User Name (root).
- Cluster Expand Configuration:** Management Cluster (10.10.50.100).
- IP Addresses:**
  - Server 0:** Management Hypervisor (10.10.50.67), Data Hypervisor (10.10.52.27).
  - Server 1:** Management Hypervisor (10.10.50.68), Data Hypervisor (10.10.52.28).
  - Server 2:** Management Hypervisor (10.10.50.77), Data Hypervisor (10.10.52.37).
  - Server 3:** Management Hypervisor (10.10.50.78), Data Hypervisor (10.10.52.38).
  - Server 4:** Management Hypervisor (10.10.50.79), Data Hypervisor (10.10.52.39).
  - Server 5:** Management Hypervisor (10.10.50.80).



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

24. Click Launch vSphere Web Client.

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



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

Figure 44 Cisco UCS Node Naming Convention

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem	EVC Mode	HA Protection
stCIVM-FCH2033V1AD	Powered On	Normal	2.42 GB	2.42 GB	804 MHz	49,335 MB		Protected
stCIVM-FCH1938V085	Powered On	Normal	2.42 GB	2.42 GB	830 MHz	49,335 MB		Protected
stCIVM-FCH1936V0GE	Powered On	Normal	2.42 GB	2.42 GB	804 MHz	49,335 MB		Protected
stCIVM-FCH1937V2JV	Powered On	Normal	2.42 GB	2.42 GB	1,037 MHz	49,335 MB		Protected
stCIVM-FCH2033V18F	Powered On	Normal	2.42 GB	2.42 GB	1,037 MHz	49,335 MB		Protected
stCIVM-FCH1937V2TS	Powered On	Normal	2.42 GB	2.42 GB	804 MHz	49,335 MB		Protected
stCIVM-FCH2033V0BW	Powered On	Normal	2.42 GB	2.42 GB	1,841 MHz	49,336 MB		Protected
stCIVM-FCH1937V2TV	Powered On	Normal	2.42 GB	2.42 GB	933 MHz	49,335 MB		Protected
stCIVM-FCH1937V2JU	Powered On	Normal	2.42 GB	2.42 GB	856 MHz	49,336 MB		Protected
stCIVM-FCH2033V1E9	Powered On	Normal	2.42 GB	2.42 GB	1,193 MHz	49,335 MB		Protected
stCIVM-FCH1937V2JT	Powered On	Normal	2.42 GB	2.42 GB	1,167 MHz	49,335 MB		Protected
stCIVM-FCH2033V0LR	Powered On	Normal	2.42 GB	2.42 GB	856 MHz	49,335 MB		Protected
stCIVM-FCH1842V1JG	Powered On	Normal	2.42 GB	2.42 GB	985 MHz	49,335 MB		Protected
stCIVM-FCH2033V0H8	Powered On	Normal	2.42 GB	2.42 GB	1,167 MHz	49,335 MB		Protected
stCIVM-FCH2033V0HF	Powered On	Normal	2.42 GB	2.42 GB	856 MHz	49,335 MB		Protected
stCIVM-FCH2031V054	Powered On	Normal	2.42 GB	2.42 GB	907 MHz	49,335 MB		Protected
stCIVM-FCH1838ZPL	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	467 MB		Protected
stCIVM-FLM1942AFAR	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	458 MB		Protected
stCIVM-FLM19379CJF	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	450 MB		Protected
stCIVM-FLM2033LLTP	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	467 MB		Protected
stCIVM-FLM2033LLYE	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	441 MB		Protected
stCIVM-FCH19377PVA	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	467 MB		Protected
stCIVM-FCH2033V0LF	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	445 MB		Protected
stCIVM-FCH2033V18H	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	445 MB		Protected
stCIVM-FCH2033V0EX	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	443 MB		Protected
stCIVM-FLM19379C8B	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	418 MB		Protected
stCIVM-FCH2033V0M0	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	446 MB		Protected
stCIVM-FCH2033V1AN	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	447 MB		Protected
stCIVM-FCH2033V0M2	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	444 MB		Protected
stCIVM-FCH2109V2WG	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	446 MB		Protected
stCIVM-FCH2033V0FP	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	445 MB		Protected
stCIVM-FCH19057QG2	Powered On	Normal	2.66 GB	2.66 GB	0 MHz	451 MB		Protected

Run Cluster Post Installation Script

After a successful installation of HyperFlex cluster, run the post\_install script by logging into the Data Platform Installer VM via SSH, using the credentials configured earlier.

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

```

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

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

Disable SSH warning? (y/n) y

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

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

Add VM network VLANs? (y/n) n

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

Send test email? (y/n) n

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

```

1. To run the script, first use your tool of choice to make a secure connection to the Cisco HyperFlex Data Platform installer using its IP address and port 22.
2. Authenticate with the credentials provided earlier. (user name: root with password Cisco 123 if you did not change the defaults.)

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

```

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

vCenter URL: 10.10.50.20

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

vCenter Password:

Found datacenter VDILAB-HX

Found cluster HX-VDI-CL

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

Disable SSH warning? (y/n) y

Add vmotion interfaces? (y/n) y

Netmask for vMotion: 255.255.255.0

VLAN ID: (0-4096) 53

vMotion IP for 10.10.50.27: 10.10.53.27

Adding vmotion to 10.10.50.27

Adding vmkernel to 10.10.50.27

vMotion IP for 10.10.50.28: 10.10.53.28

Adding vmotion to 10.10.50.28

Adding vmkernel to 10.10.50.28

vMotion IP for 10.10.50.29: 10.10.53.29

Adding vmotion to 10.10.50.29

Adding vmkernel to 10.10.50.29

vMotion IP for 10.10.50.30: 10.10.53.30

Adding vmotion to 10.10.50.30

Adding vmkernel to 10.10.50.30

vMotion IP for 10.10.50.31: 10.10.53.31

```

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

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

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

## Network Summary:

Host: 10.10.50.27

vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance\_srcid

vmnic0 - 1 - K22-HXVDI-A - active

vmnic1 - 1 - K22-HXVDI-B - standby

Portgroup Name - VLAN

Storage Controller Management Network - 50

Management Network - 50

vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance\_srcid

vmnic4 - 1 - K22-HXVDI-A - active

vmnic5 - 1 - K22-HXVDI-B - active

Portgroup Name - VLAN

vm-network-54 - 54

vswitch: vmotion - mtu: 9000 - policy: loadbalance\_srcid

vmnic6 - 1 - K22-HXVDI-A - active

vmnic7 - 1 - K22-HXVDI-B - standby

Portgroup Name - VLAN

vmotion - 53

vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance\_srcid

vmnic2 - 1 - K22-HXVDI-A - standby

vmnic3 - 1 - K22-HXVDI-B - active

Portgroup Name - VLAN

Storage Controller Data Network - 52

Storage Hypervisor Data Network - 52

Host: 10.10.50.28

vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance\_srcid

vmnic0 - 1 - K22-HXVDI-A - active

vmnic1 - 1 - K22-HXVDI-B - standby

Portgroup Name - VLAN

Storage Controller Management Network - 50

Management Network - 50

vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance\_srcid



```

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

```

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

```

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

```

```
vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
  vmnic6 - 1 - K22-HXVDI-A - active
  vmnic7 - 1 - K22-HXVDI-B - standby
  Portgroup Name - VLAN
  vmotion - 53

vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
  vmnic2 - 1 - K22-HXVDI-A - standby
  vmnic3 - 1 - K22-HXVDI-B - active
  Portgroup Name - VLAN
  Storage Controller Data Network - 52
  Storage Hypervisor Data Network - 52

Host: 10.10.50.33

vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
  vmnic0 - 1 - K22-HXVDI-A - active
  vmnic1 - 1 - K22-HXVDI-B - standby
  Portgroup Name - VLAN
  Storage Controller Management Network - 50
  Management Network - 50

vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
  vmnic4 - 1 - K22-HXVDI-A - active
  vmnic5 - 1 - K22-HXVDI-B - active
  Portgroup Name - VLAN
  vm-network-54 - 54

vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
  vmnic6 - 1 - K22-HXVDI-A - active
  vmnic7 - 1 - K22-HXVDI-B - standby
  Portgroup Name - VLAN
  vmotion - 53

vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
  vmnic2 - 1 - K22-HXVDI-A - standby
  vmnic3 - 1 - K22-HXVDI-B - active
  Portgroup Name - VLAN
```

Storage Controller Data Network - 52

Storage Hypervisor Data Network - 52

Host: 10.10.50.34

vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance\_srcid

vmnic0 - 1 - K22-HXVDI-A - active

vmnic1 - 1 - K22-HXVDI-B - standby

Portgroup Name - VLAN

Storage Controller Management Network - 50

Management Network - 50

vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance\_srcid

vmnic4 - 1 - K22-HXVDI-A - active

vmnic5 - 1 - K22-HXVDI-B - active

Portgroup Name - VLAN

vm-network-54 - 54

vswitch: vmotion - mtu: 9000 - policy: loadbalance\_srcid

vmnic6 - 1 - K22-HXVDI-A - active

vmnic7 - 1 - K22-HXVDI-B - standby

Portgroup Name - VLAN

vmotion - 53

vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance\_srcid

vmnic2 - 1 - K22-HXVDI-A - standby

vmnic3 - 1 - K22-HXVDI-B - active

Portgroup Name - VLAN

Storage Controller Data Network - 52

Storage Hypervisor Data Network - 52

Host: 10.10.50.27

No errors found

Host: 10.10.50.28

No errors found

Host: 10.10.50.29

No errors found

Host: 10.10.50.30

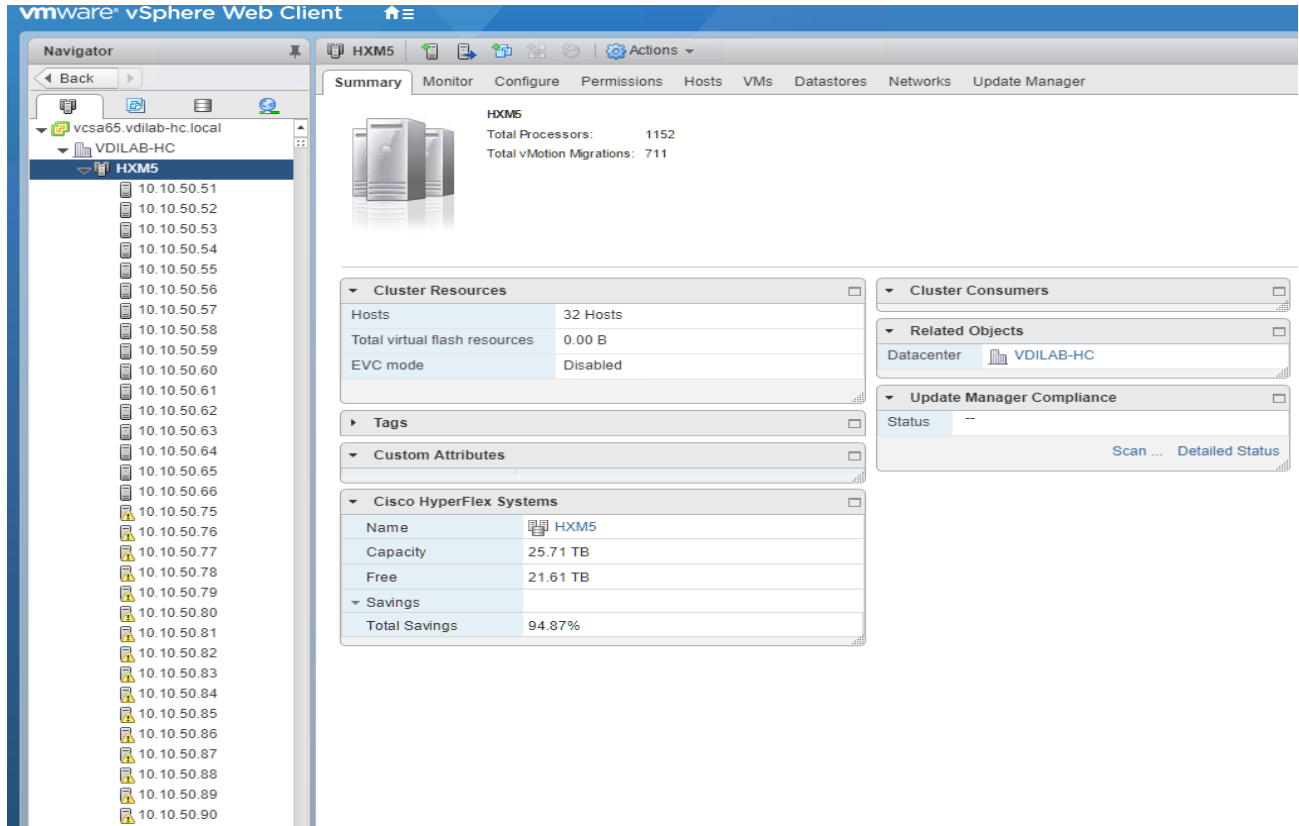
```
    No errors found
Host: 10.10.50.31
    No errors found
Host: 10.10.50.32
    No errors found
Host: 10.10.50.33
    No errors found

Host: 10.10.50.34
    No errors found
Controller VM Clocks:
    stCtlVM-FCH1937V2JV - 2018-10-22 05:32:09
    stCtlVM-FCH1937V2TV - 2018-10-22 05:32:25
    stCtlVM-FCH1842V1JG - 2018-10-22 05:32:41
    stCtlVM-FCH1936V0GE - 2018-10-22 05:32:57
    stCtlVM-FCH1937V2JT - 2018-10-22 05:33:14
    stCtlVM-FCH1938V085 - 2018-10-22 05:33:30
    stCtlVM-FCH1937V2TS - 2018-10-22 05:33:46
    stCtlVM-FCH1937V2JU - 2018-10-22 05:34:02

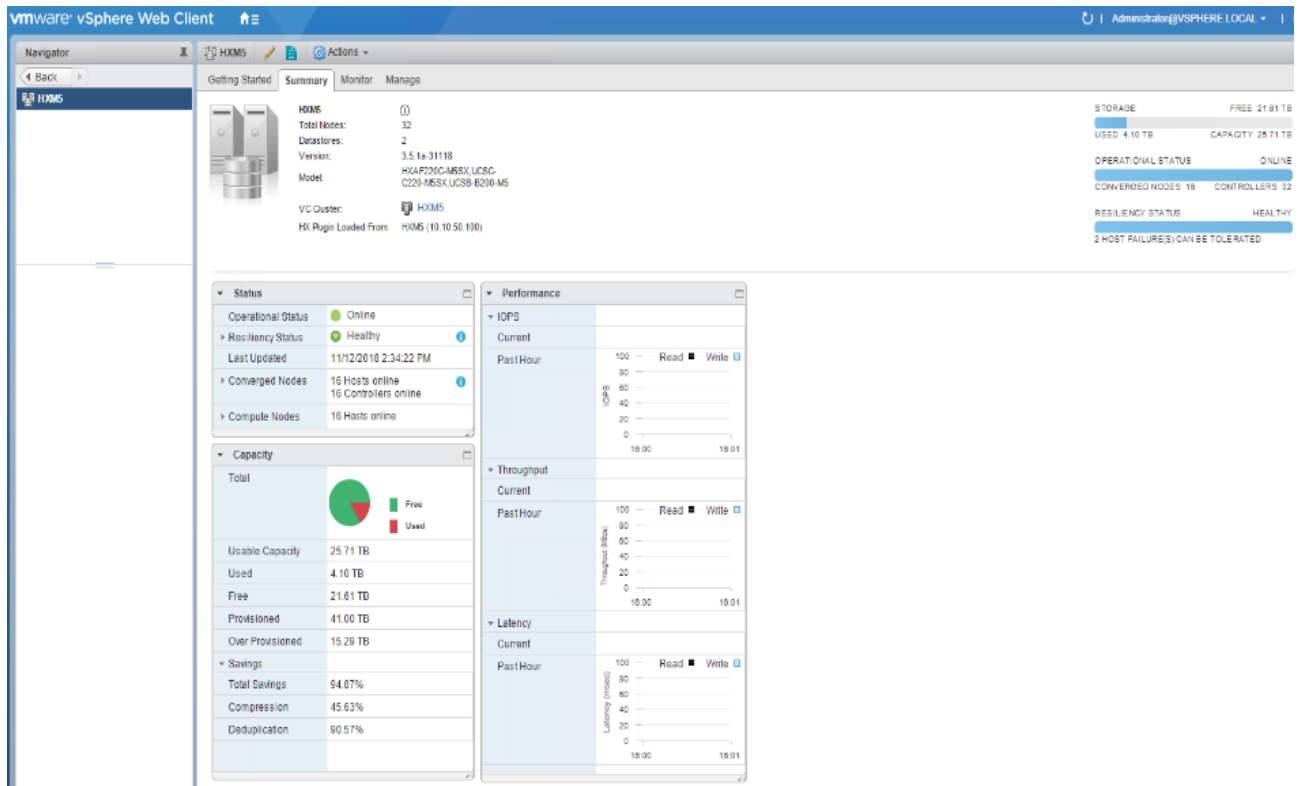
Cluster:
    Version - 3.5.1a-31118
    Model - HXAF220C-M5S
    Health - HEALTHY
    Access Policy - LENIENT
    ASUP enabled - False
    SMTP Server -

root@Cisco-HX-Data-Platform-Installer:~#
```

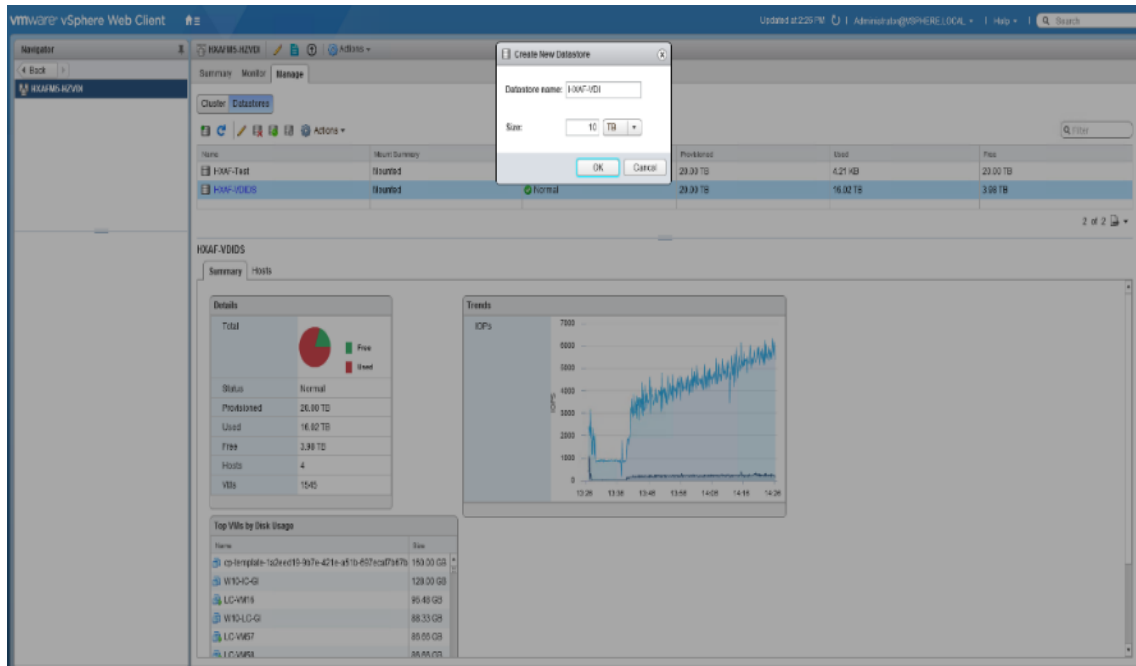
8. Login to vSphere WebClient to create additional shared datastore.
9. Go to the Summary tab on the cluster created via the HyperFlex cluster creation workflow.
10. On Cisco HyperFlex Systems click the cluster name.




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

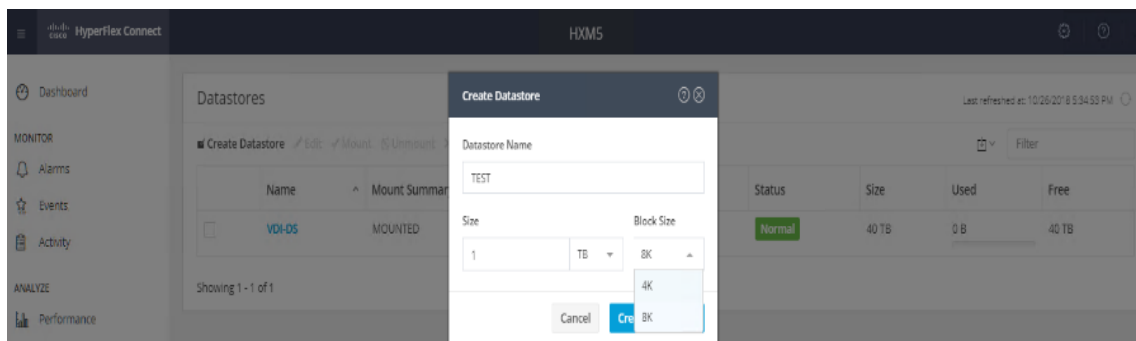


11. Click Manage, select Datasets. Click the Add dataset icon, select the dataset name and size to provision.



 You have created a 40TB dataset for the Horizon pooled, persistent/non-persistent, and RDSH server desktop performance test.

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





## Building the Virtual Machines and Environment for Workload Testing

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

### Horizon 7 Infrastructure Components Installation

The prerequisites for installing the view connection server, replica server(s) and composer server is to have Windows 2008, 2012 2012 R2 or 2016 virtual machines ready.



**In this study, we used Windows Server 2016 virtual machines for all Horizon infrastructure servers.**

Download the VMware Horizon 7 installation package from this link:

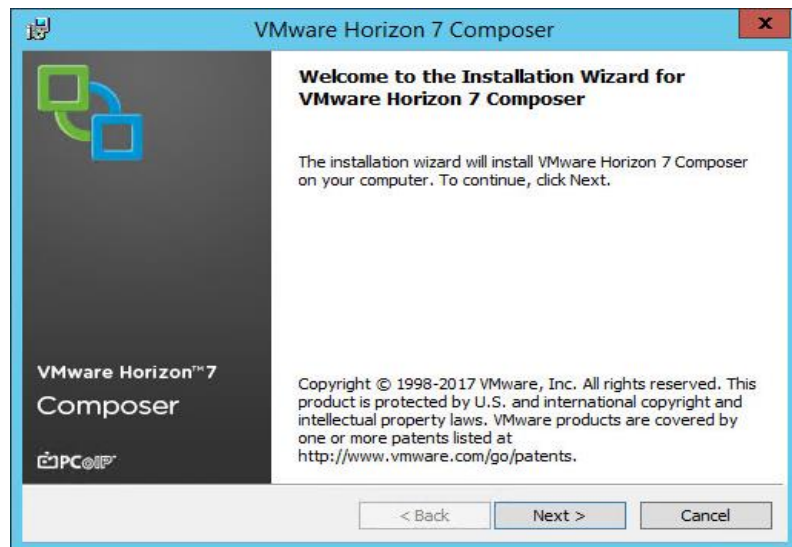
[https://my.vmware.com/web/vmware/info/slug/desktop\\_end\\_user\\_computing/vmware\\_horizon/7\\_6](https://my.vmware.com/web/vmware/info/slug/desktop_end_user_computing/vmware_horizon/7_6)

This following section provides a detailed, systematic installation process for Horizon 7 v7.6.0.

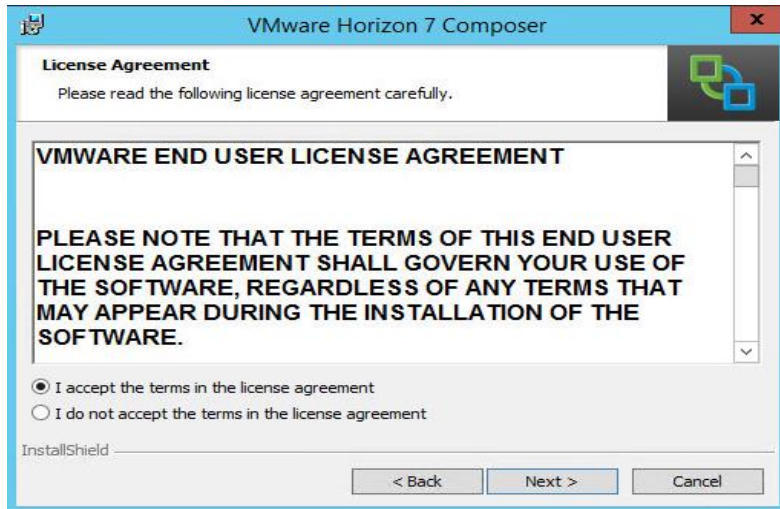
### Install VMware Horizon Composer Server

To install the VMware Horizon Composer Server, complete the following steps:


1. Open installer for Horizon composer. VMware-viewcomposer-7.6.0-9491669.exe
2. Click Next to continue.



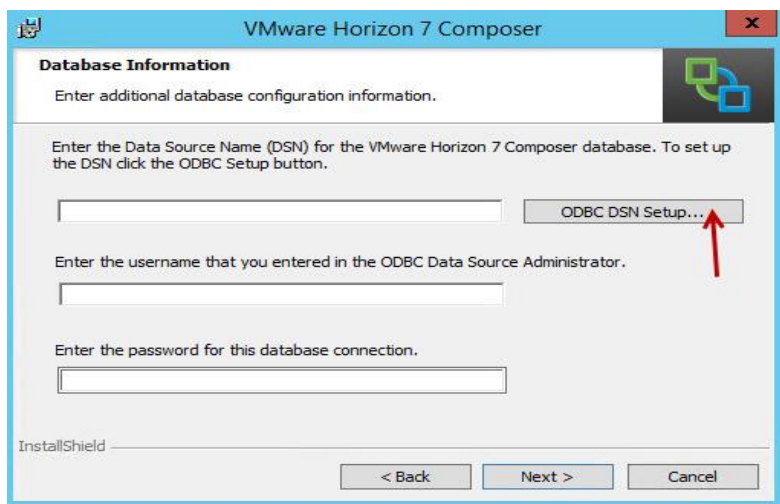
3. Accept the EULA. Click Next.



4. Click Next to accept the default installation folder.

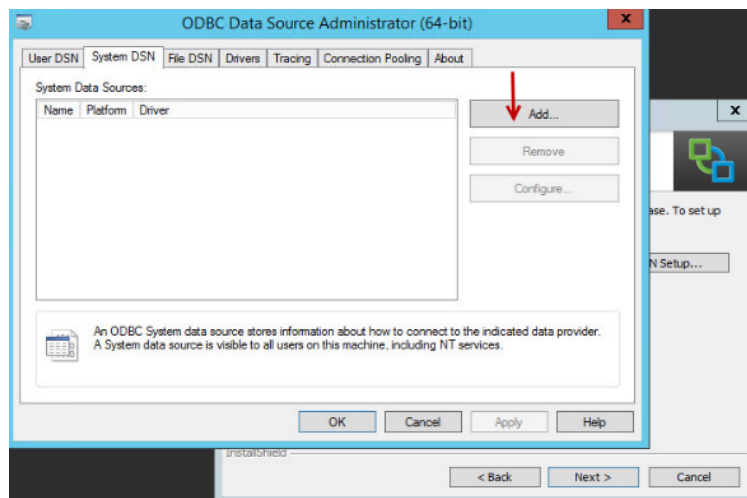


5. Enter the database information. The ODBC database can be configured during the installation by clicking ODBC DSN Setup.

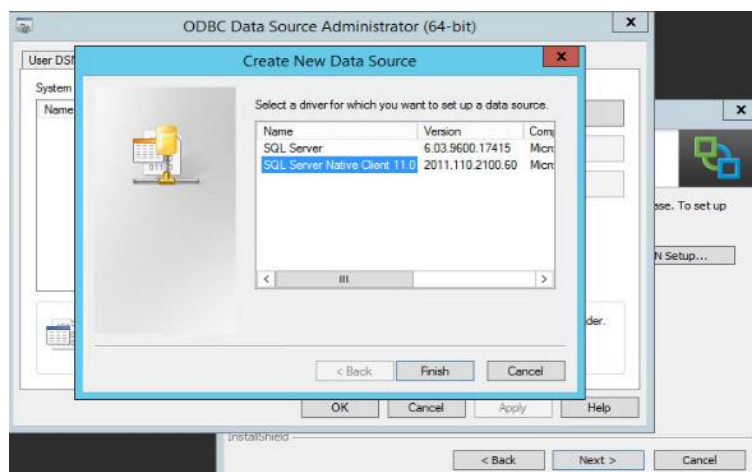


### Configure the ODBC Source Name

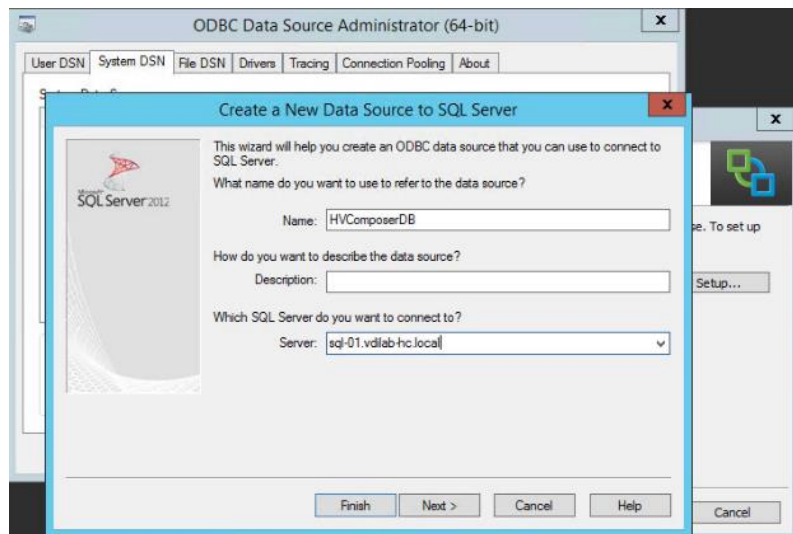
1. Open 64bit ODBC, select System DSN tab and click Add.



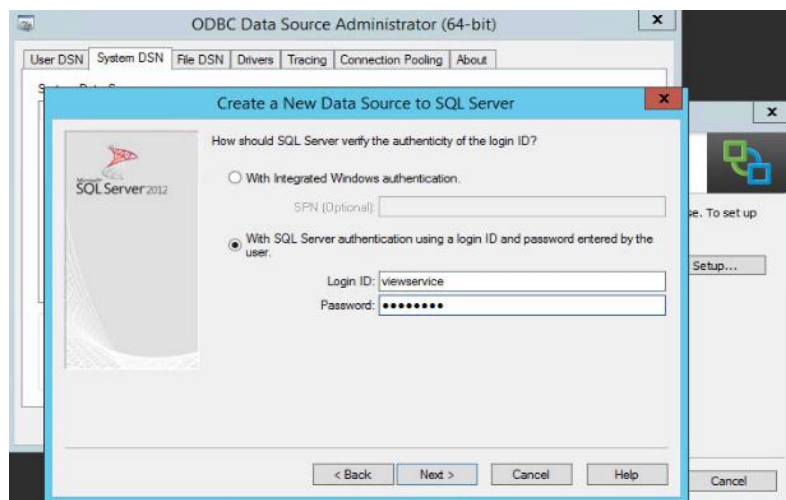
2. Create a new Data source and select SQL server native client. You will use an existing instance of the Microsoft SQL server 2016 for the current deployment. Click Finish.



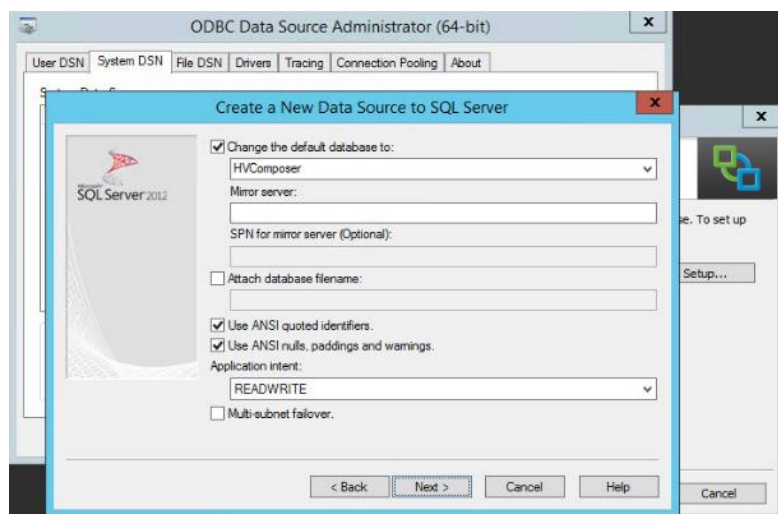
3. Create a name for data source, select SQL server, click Next.



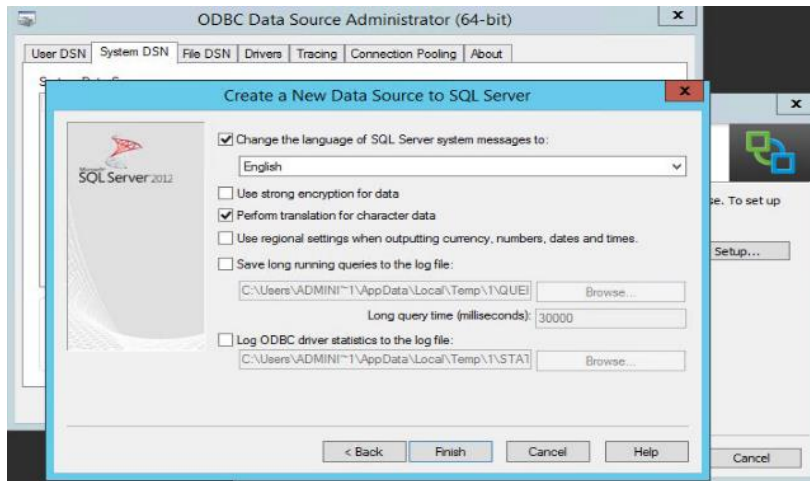
4. Enter the login credentials for the SQL server authentication or use Windows Authentication. Click Next.



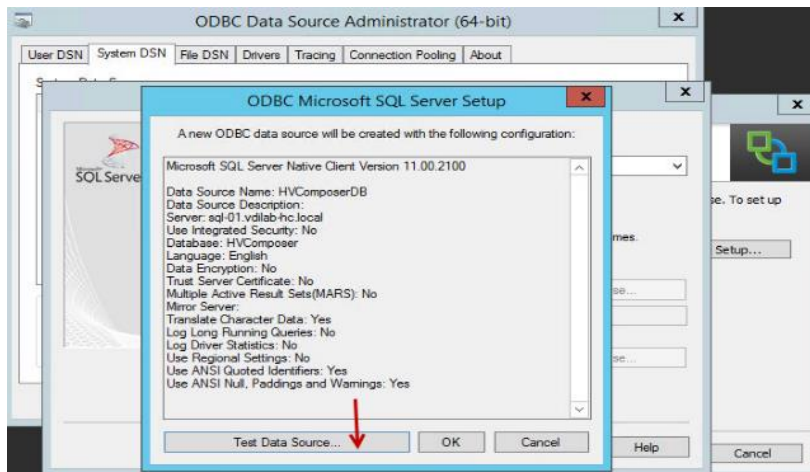
5. Select Default Database, click Next.



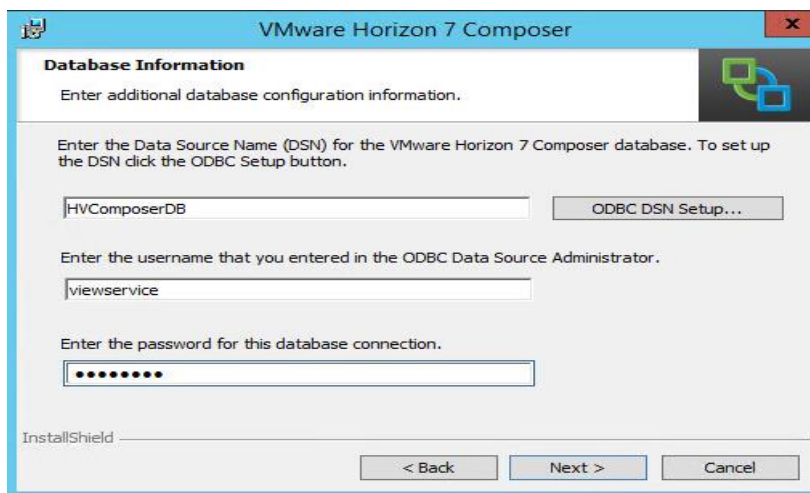
6. Check the box to select language for SQL server system messages. Click Finish.



7. Click Test datastore to verify connectivity between SQL server and newly create Data source.

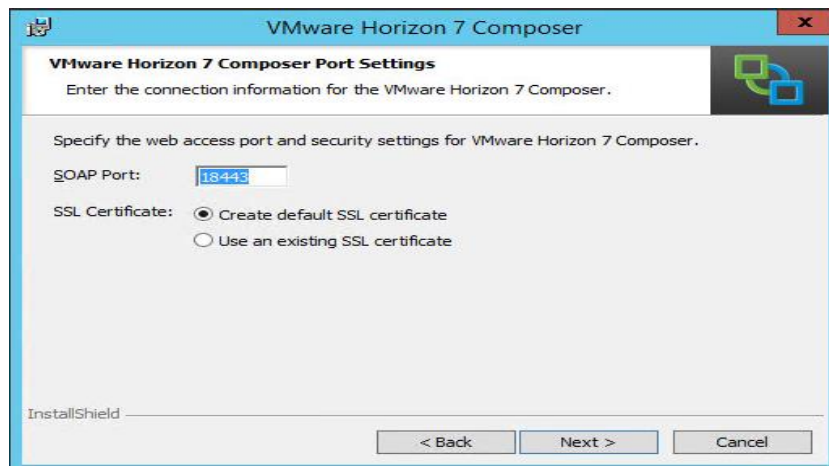


8. Since this a new instance of Composer server installation, a new SSL certificate will be created. In case of update or existing composer server installation either create new SSL certificate or use existing certificate.

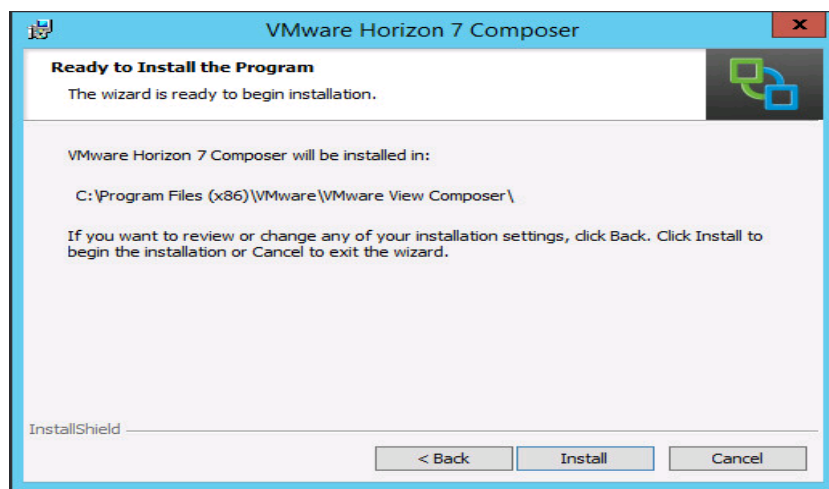


9. Leave default port configuration for SOAP port.


10. Click Next.



11. Click Install.



12. Click Finish.





## Install Horizon Connection/Replica Servers

To install the Horizon Connection/Replica Servers, complete the following steps:


1. Open view connection server installation, VMware-viewconnectionserver-x86\_64-7.6.0-9823717.exe.
2. Click Next.




3. Accept the EULA, click Next.



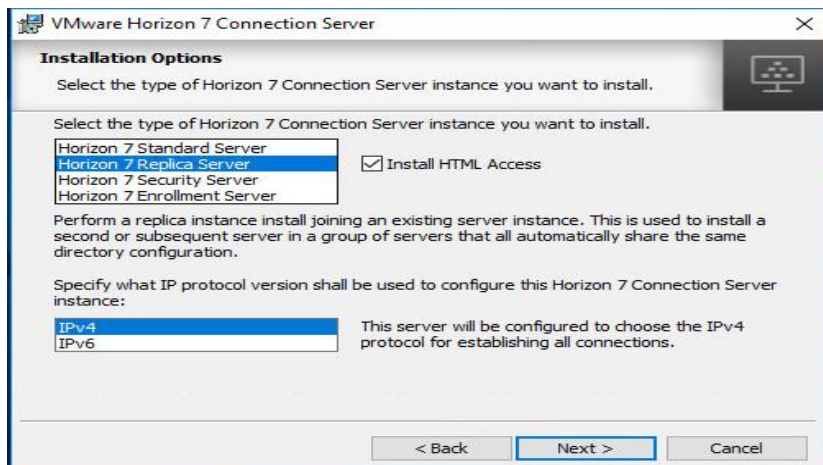
4. Leave default destination folder, click Next.



5. Select type of instance intended to install.
6. Select Standard Server instance for primary connection server installation.

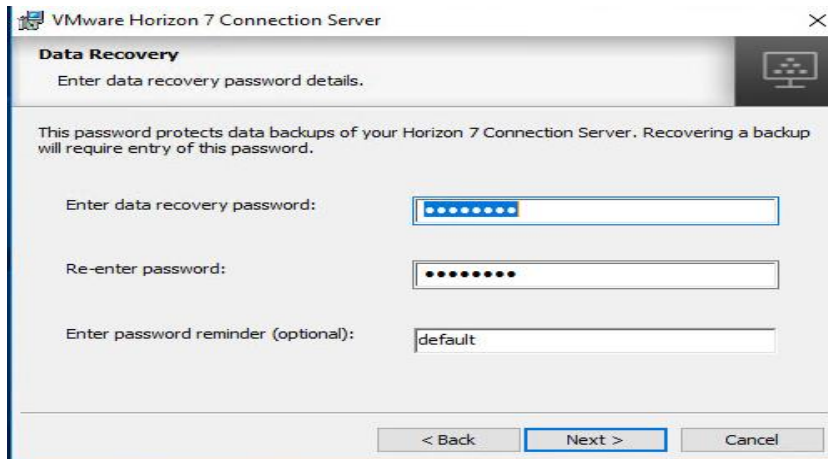


7. Select Replica server instance for fault tolerant connection server configuration after completion of Standard Server instance installation.

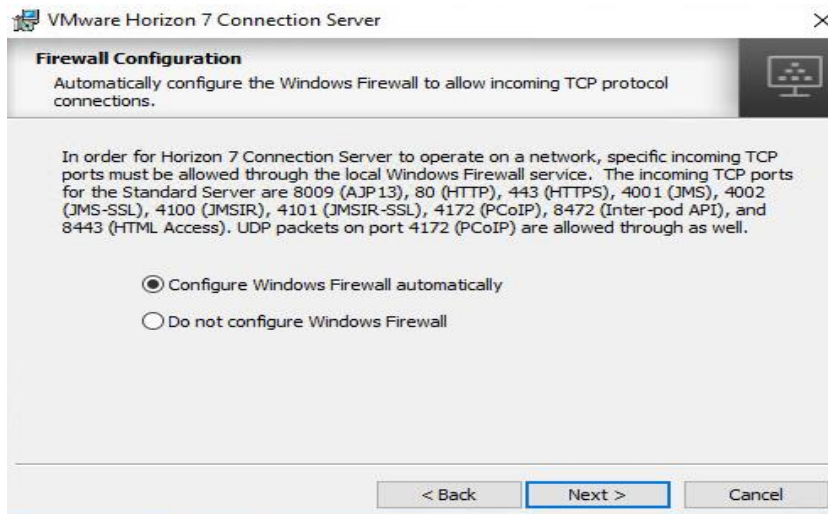




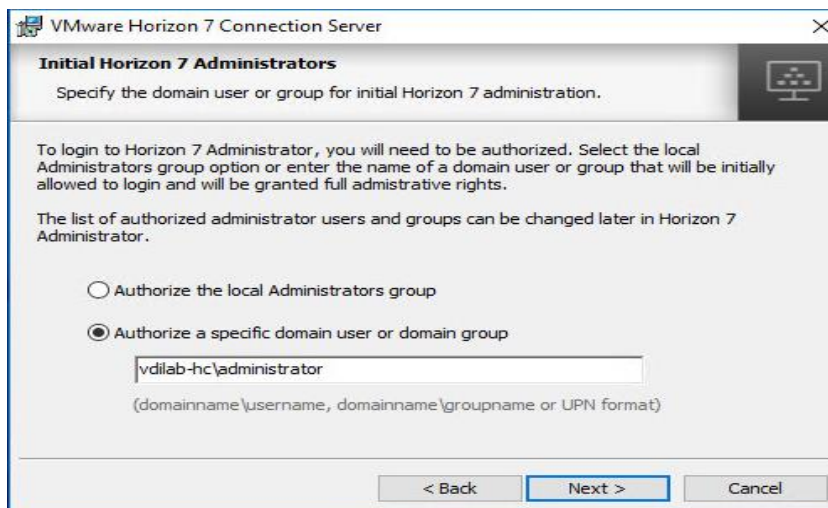
8. Enter the Data Recovery Password.




9. Click Next.



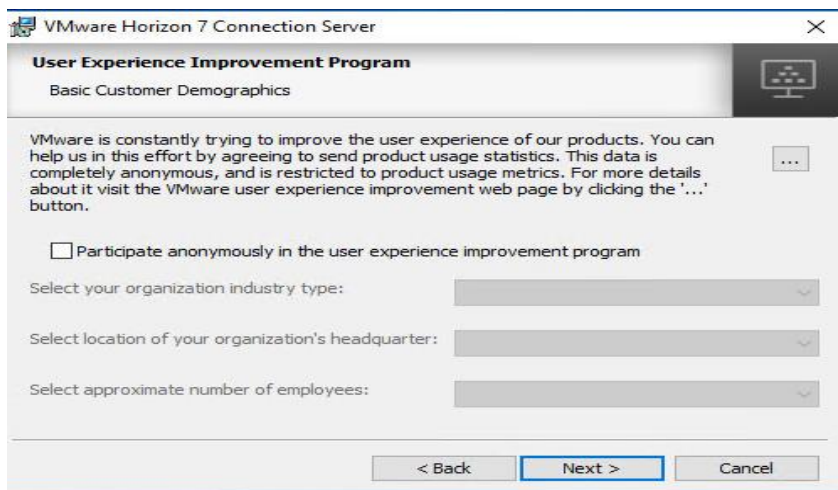
10. Select authorized users and group, click Next.




11. Enter domain credentials for previously specified domain user/group.



12. Opt-in or Opt-out of User Experience Improvement Program. Click Next.



13. Click Install.



14. Click Finish.

## Create a Microsoft Management Console Certificate Request

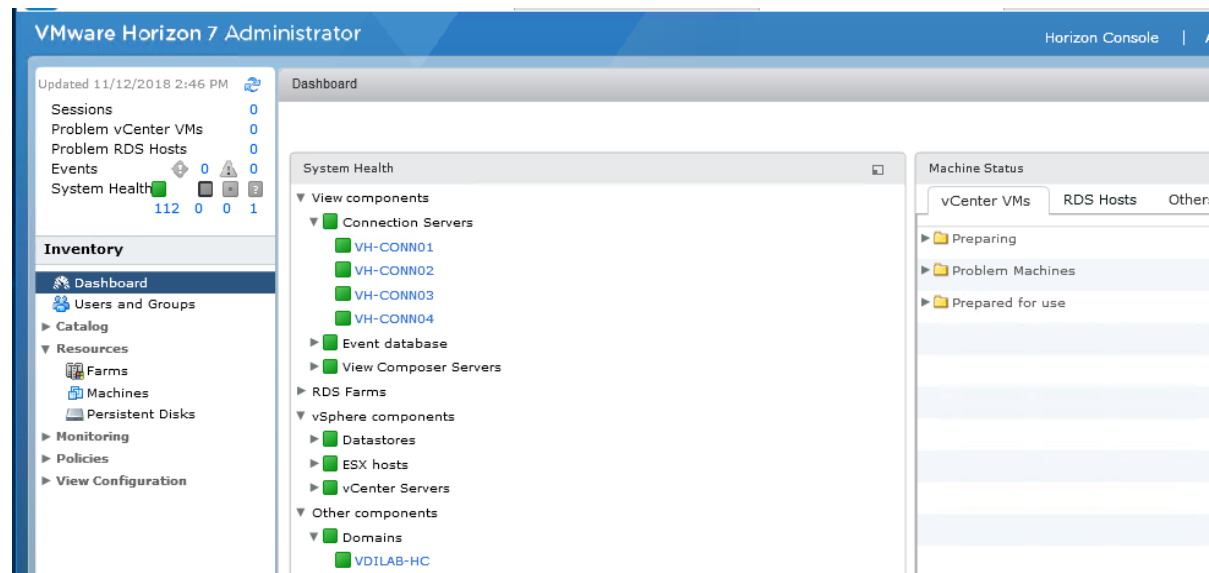
To generate a Horizon View SSL certificate request, use the Microsoft Management Console (MMC) Certificates snap-in:

[https://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2068666](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2068666)

## Configure the Horizon 7 Environment

To configure the Horizon 7 environment, complete the following steps:

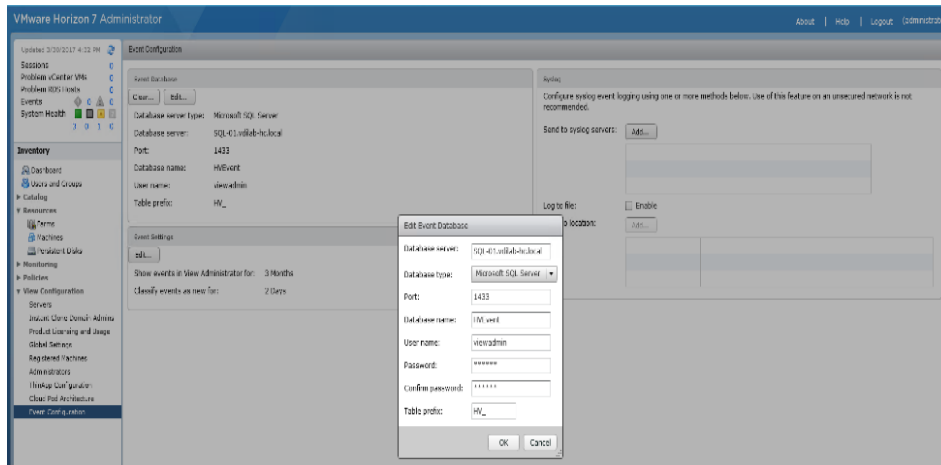
1. Open WebUI, Login to [https://<Horizon\\_Connection\\_server\\_Management\\_IP\\_Address>/admin](https://<Horizon_Connection_server_Management_IP_Address>/admin).



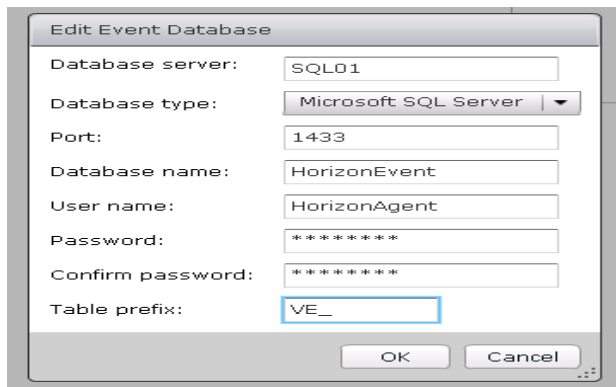
## Configure Event Database

To configure the Event Database, complete the following steps:

1. Configure the Event Database by adding Database Server, Database name, login credentials and prefix for the table from the Horizon 7 Administrator, View Configuration, Event Configuration node of the Inventory pane.
2. Click Edit in the action pane.



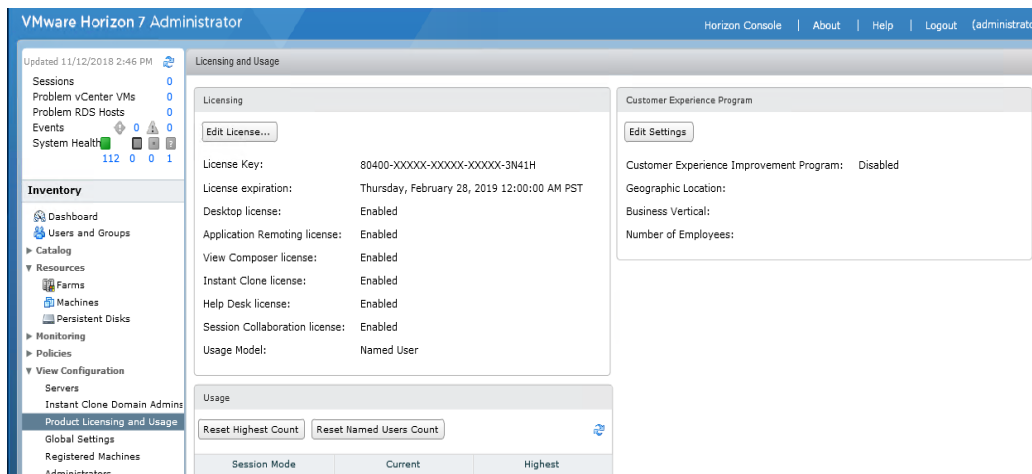
The details are shown below:



## Configure Horizon 7 Licenses

To configure the Horizon 7 licenses, complete the following steps:

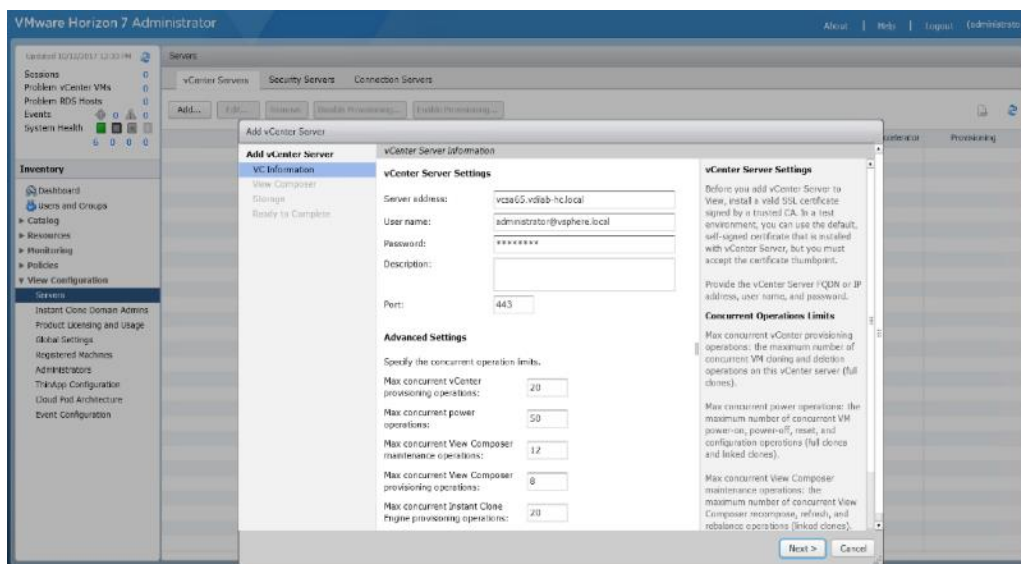
1. Click View Configuration.
2. Select Product Licensing and Usage.
3. Click Edit License in the action pane.
4. Add the License Serial Number.
5. Click OK.



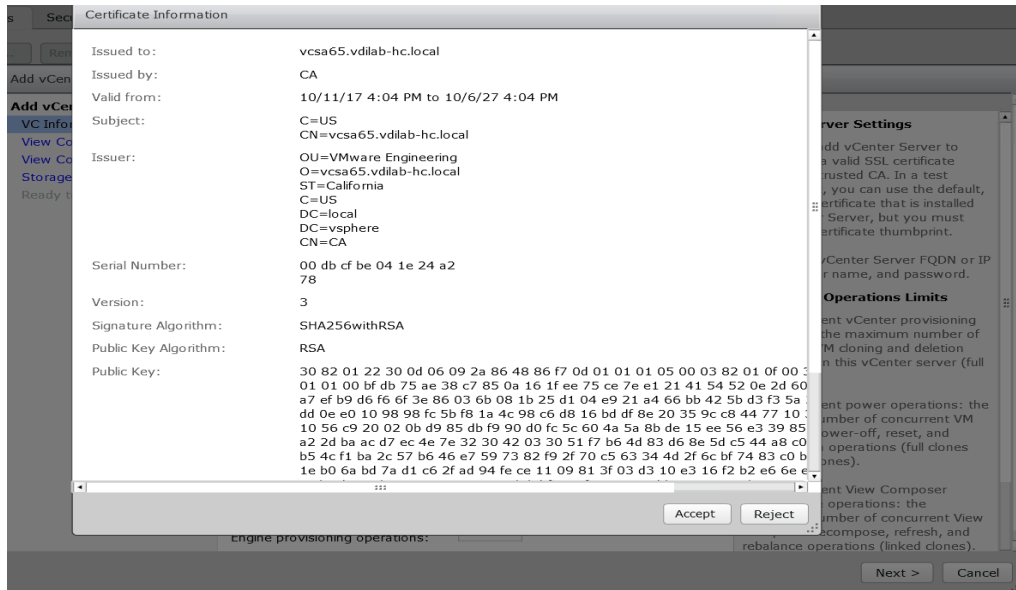
## Configure vCenter

To configure the vCenter, complete the following steps:

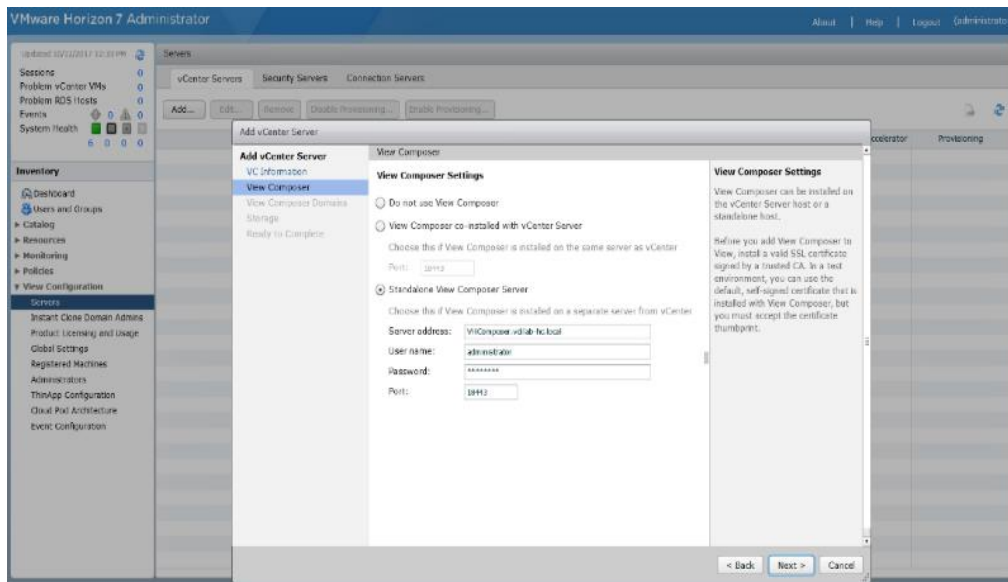
1. In View Configuration, Select Servers. Click Add vCenter Server tab.
2. Enter vCenter Server IP Address or FQDN, login credentials.
3. Advanced Settings options can be modified to change existing operations limit. Keep the advanced settings options as default.



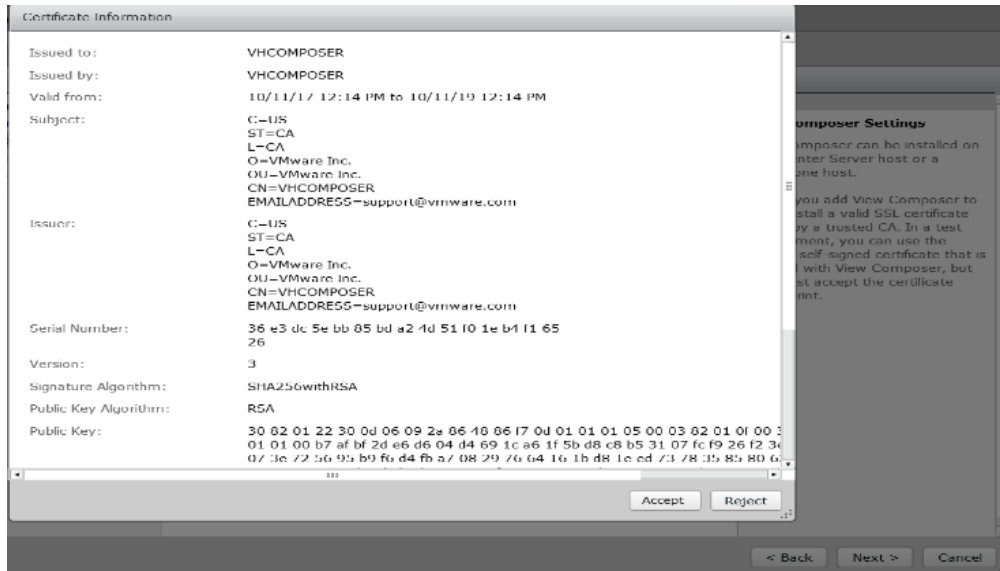
4. Click View certificate. Accept the certificate.



5. Add View composer settings, View composer server FQDN or Management IP address, login credentials.
6. Click Next.

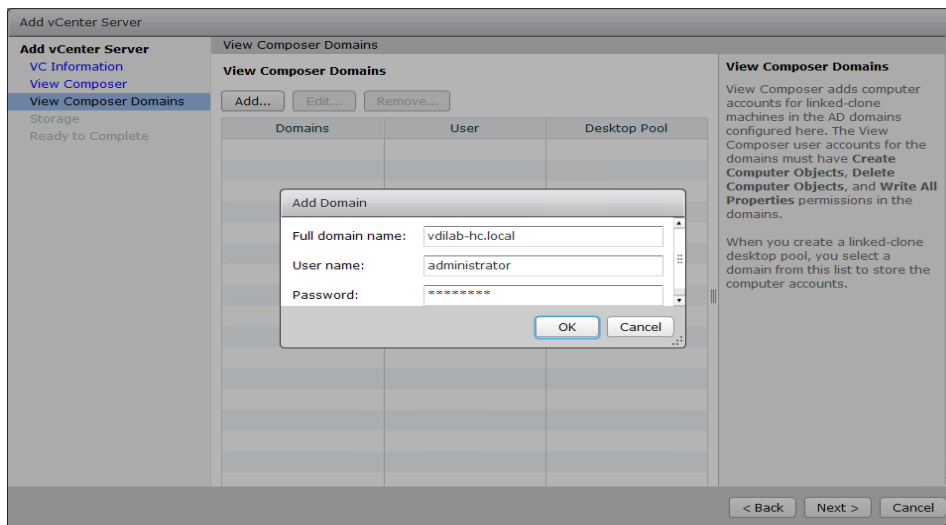


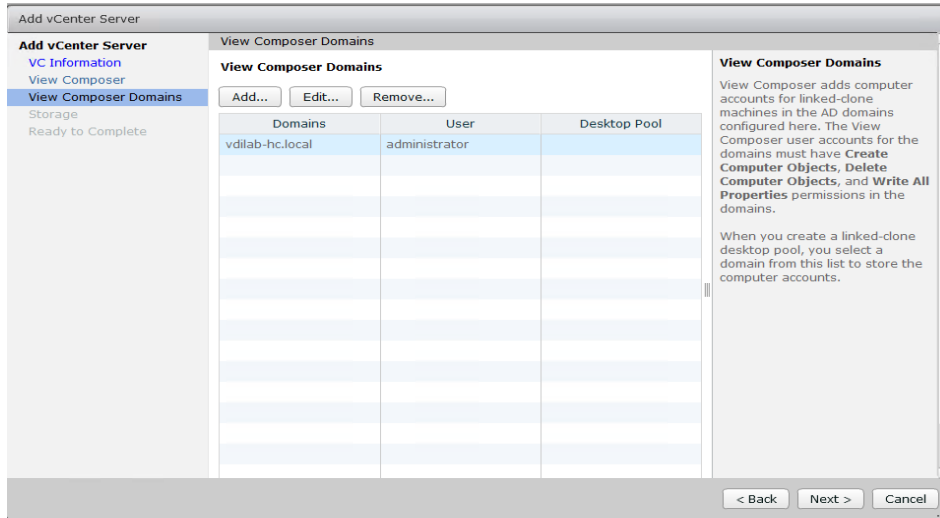
7. View and accept the certificate.



8. Click Add a new domain or Edit the existing domain.

9. Click Next.

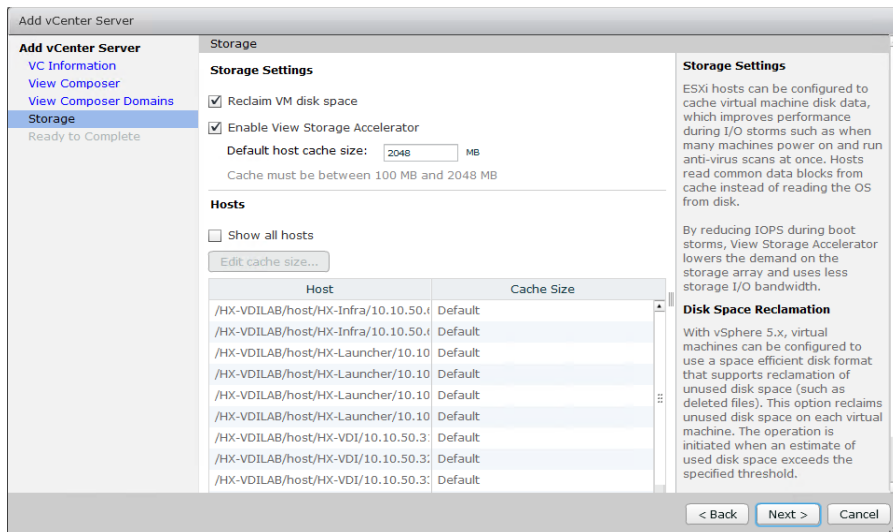




10. In Storage settings, select Reclaim VM disk space and View Storage Accelerator.

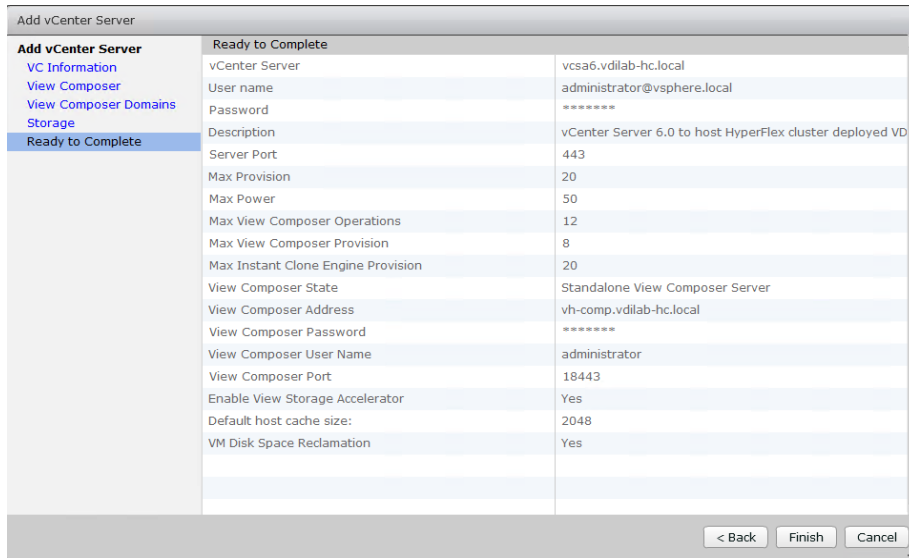
11. Configure default host cache size between 100MB and 2048MB. We configured the maximum, which is 2048MB.

12. Click Next.



13. Review Add vCenter Server settings and click Finish.

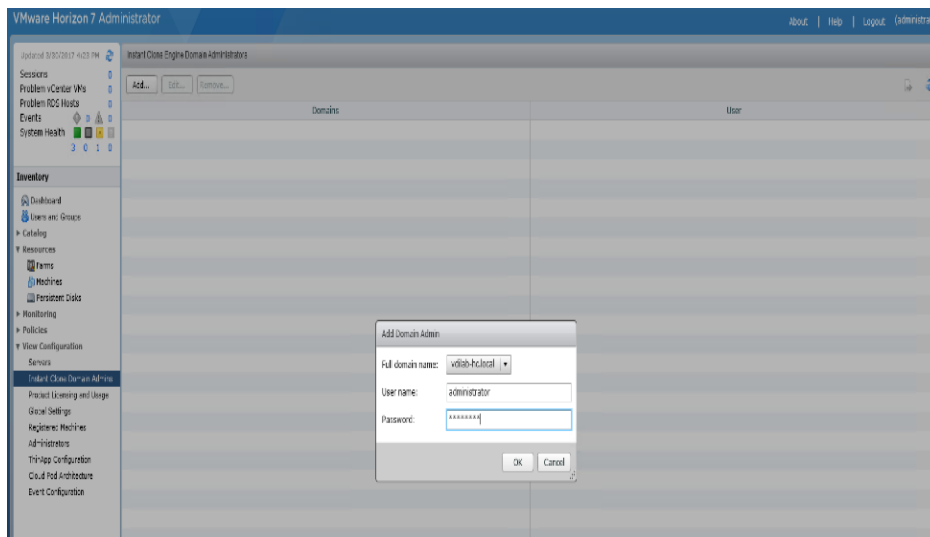




### Configure Instant Clone Domain Admins

To configure the instant clone domain admins, complete the following steps:

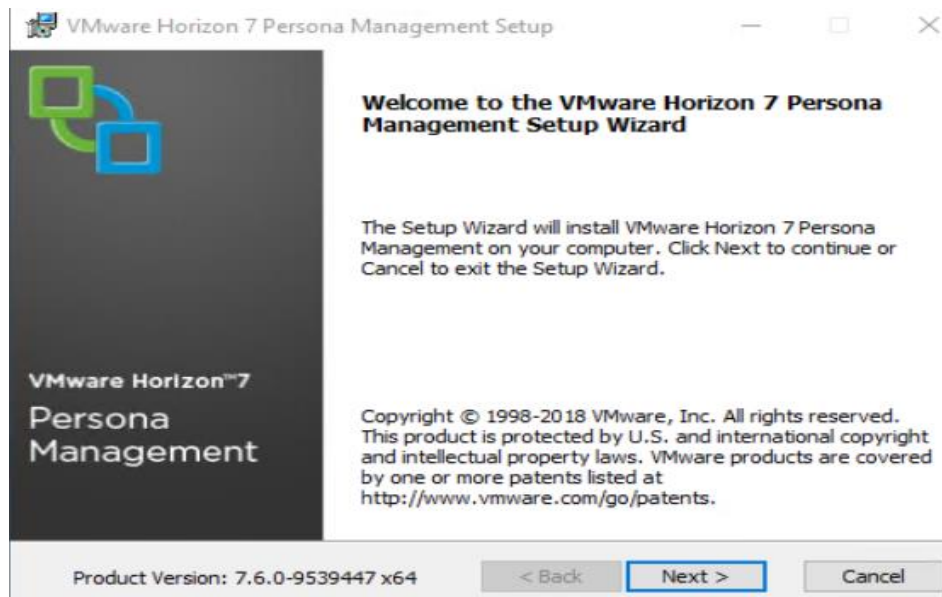
1. Under View Configuration, Click on Instant Clone Domain Admins.
2. Click Add button. Enter credentials for domain user/group.



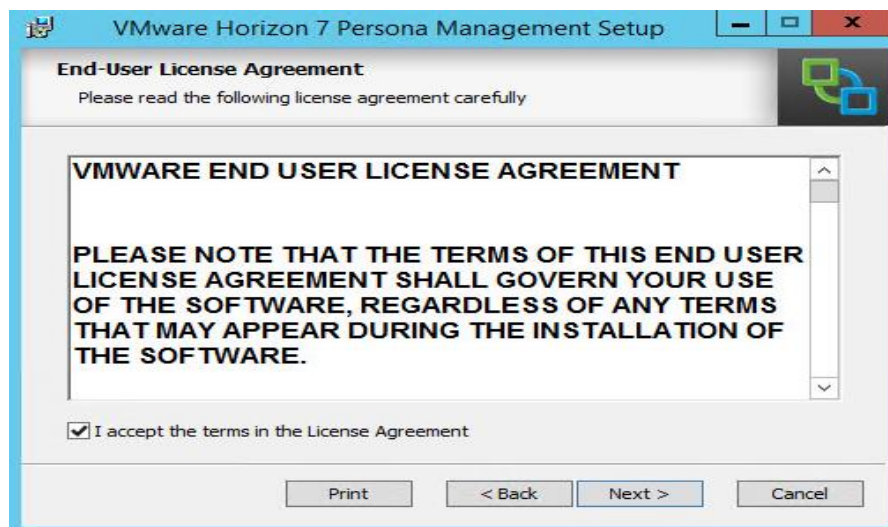
### Horizon Persona Manager Installation

To install Horizon Persona Manager, complete the following steps:


1. Open Horizon Persona Manager Installer, VMware-personamangement-x86\_64-7.6.0-9539447.exe.
2. Click Next.



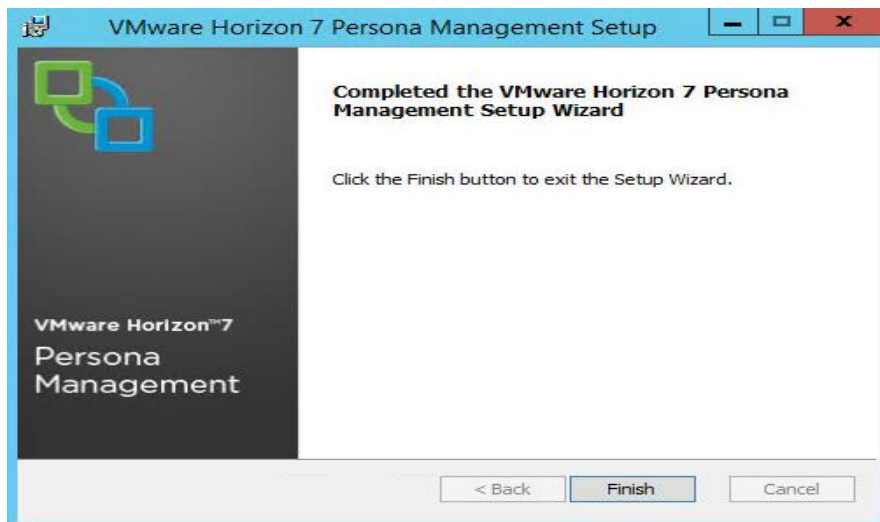
3. Accept the EULA and click Next.



4. Click Install.



5. Click Finish.



## Master Image Creation for Tested Horizon Deployment Types

To create the Master Image for the tested Horizon deployment types, complete the following steps:

1. Select an ESXi host in an existing infrastructure cluster and create the virtual machines to use as Golden Images with Windows 10 and Office 2016 for Instant Clone, Linked-Clone and Full Clone desktops.



**We used a 64-bit version of Microsoft Operating System and Office for our testing.**



**A fourth master image has been created using Microsoft Windows Server 2016 for RDSH (Remote Desktop Server Sessions) session host virtual machines.**

For the Master Image virtual machines, the following parameters were used (Table 12 ).

**Table 12 Golden Image Virtual Machine Parameters**

Attribute	Instant/Linkedclone	Persistent/Full Clone	RDSH server
Desktop operating system	Microsoft Windows 10 Enterprise (64-bit)	Microsoft Windows 10 Enterprise (64-bit)	Microsoft Windows Server 2016 standard (64-bit)
Hardware	VMware Virtual Hardware Version 13	VMware Virtual Hardware Version 13	VMware Virtual Hardware Version 13
vCPU	2	2	8
Memory	4096 MB	4096 MB*	24576MB
Memory reserved	4096 MB	4096 MB*	24576MB
Video RAM	35 MB	35 MB	4MB
3D graphics	Off	Off	Off
NIC	1	1	1
Virtual network adapter 1	VMXNet3 adapter	VMXNet3 adapter	VMXNet3 adapter
Virtual SCSI controller 0	Paravirtual	Paravirtual	Paravirtual
Virtual disk: VMDK 1	32 GB	100 GB	40 GB
Virtual disk: VMDK 2 (non-persistent disk for Linked-Clones)	6 GB	-	-
Virtual floppy drive 1	Removed	Removed	Removed
Virtual CD/DVD drive 1	-	-	-

Applications	Login VSI 4.1.32 application installation Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Microsoft Internet Explorer Microsoft Office 2016	Login VSI 4.1.32 application installation Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Microsoft Internet Explorer Microsoft Office 2016	Login VSI 4.1.32 application installation Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Microsoft Internet Explorer Microsoft Office 2016
VMware tools	Release 10.2.1.8267844	Release 10.2.1.8267844	Release 10.2.1.8267844
VMware View Agent	Release 7.6.0-9539447	Release 7.6.0-9539447	Release 7.6.0-9539447
Attribute	Linked-Clone/Instant-clone	Persistent/Full Clone	RDSH server

\* For Persistent Desktops, we configured 4GB of RAM as amount of memory allocated is sufficient to run LoginVSI Knowledge Worker workload. HyperFlex nodes and compute-only node were configured with 768GB of total memory for this performance study.

## Prepare Microsoft Windows 10 and Server 2016 R2 with Microsoft Office 2016

Prepare your master image for one or more of the following use cases:

- VMware Horizon 7 RDSH Virtual Machines
- VMware Horizon 7 Instant Clones
- VMware Horizon 7 Linked Clones
- VMware Horizon 7 Full clones

Include Microsoft Office 2016 and other applications used by all pool users in your organization into your master image.

Apply required Microsoft updates and patches to your master images.

For this study, we added Login VSI target software to enable the use the Login VSI Knowledge Worker workload to benchmark end user experience for each use case.

## Optimization of Base Windows 10 or Server 2016 Guest OS

Click the links below for information about how to optimize windows 10 for VDI deployment:

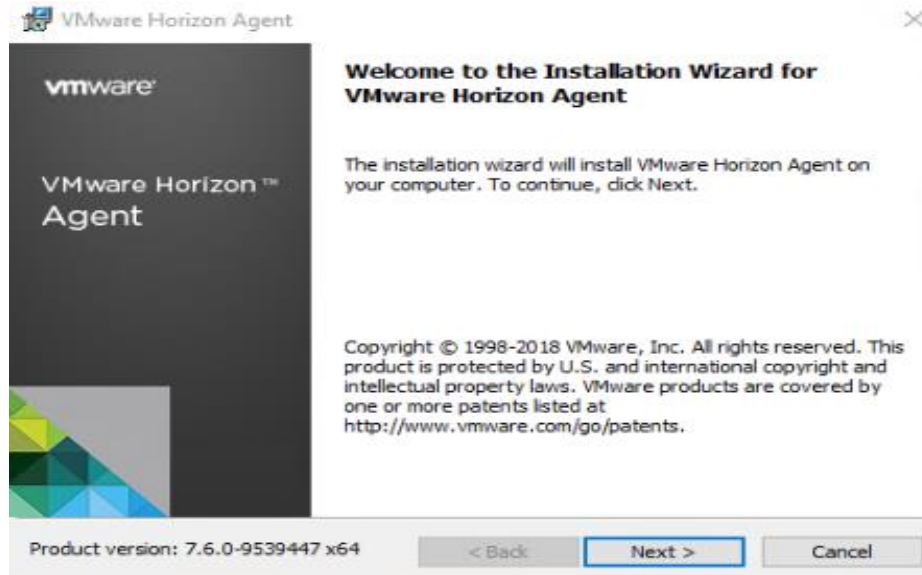
<http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/vmware-view-optimizationguidewindows7-en-white-paper.pdf>

VMware Optimization tool for HVD or HSD deployment: <https://labs.vmware.com/flings/vmware-os-optimization-tool>

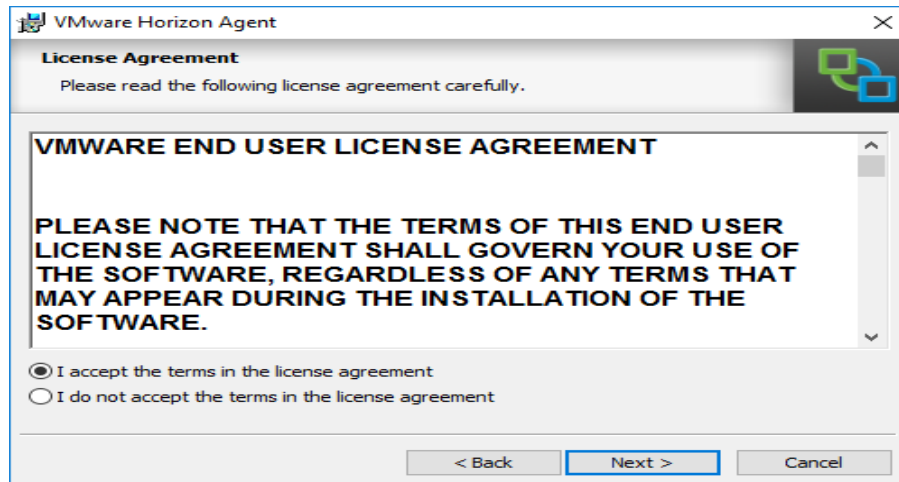
## Virtual Desktop Agent Software Installation for Horizon

To install the Virtual Desktop Agent software for Horizon, complete the following steps:

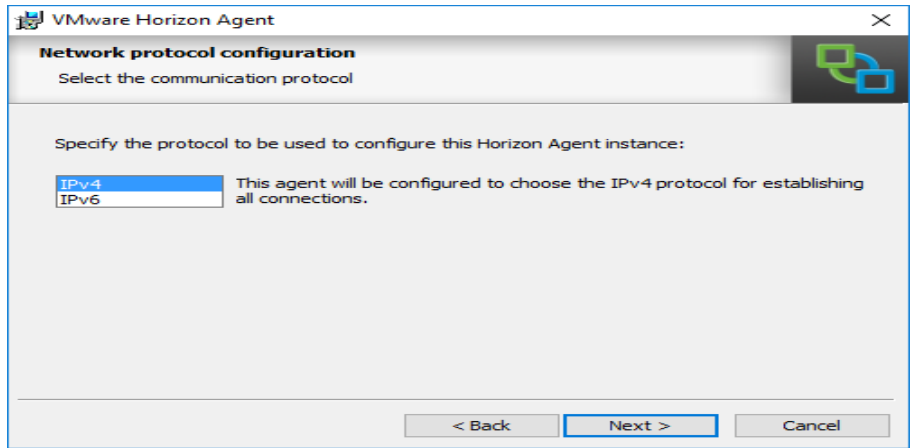
1. For each master image created, open the Horizon View Agent Installer, VMware-viewagent-7.6.0-9539447. Click Next to install.



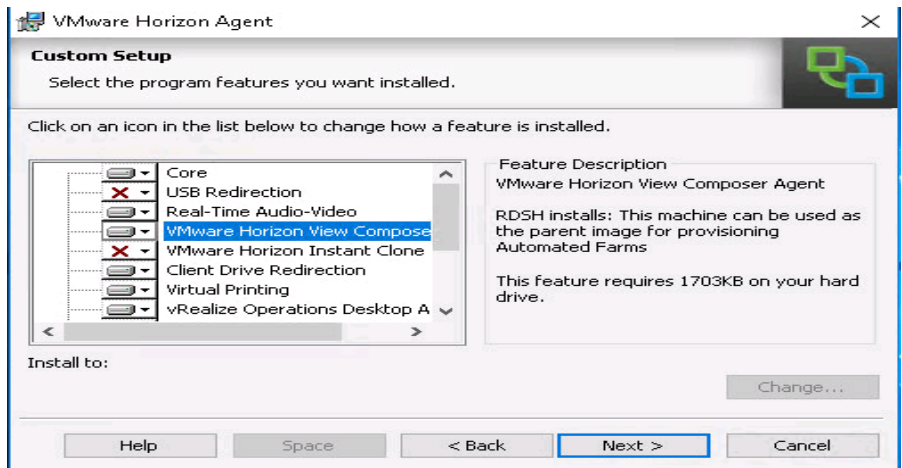
2. Review and accept the EULA Agreement. Click Next.



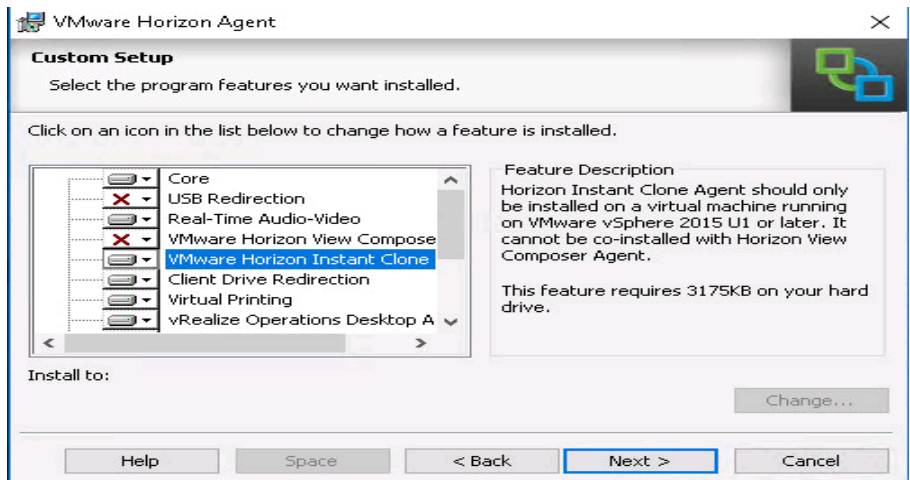
3. Select Network protocol configuration, click Next.



4. Based on the Desktop pool you want to create, select either View Composer Agent or Instant Clone Agent installation. Do not install both features on the same master image.
5. Enable installation of the VMware Horizon View Composer Agent for linked-clone VDI virtual machines.



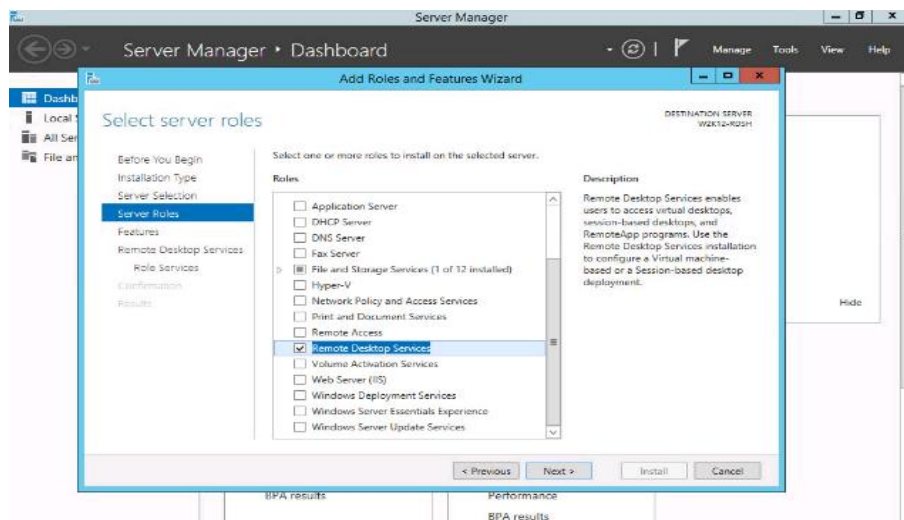
6. Disable the Horizon View Composer Agent and enable the Horizon Instant Clone Agent for Instant Clone floating assigned desktop pool creation.



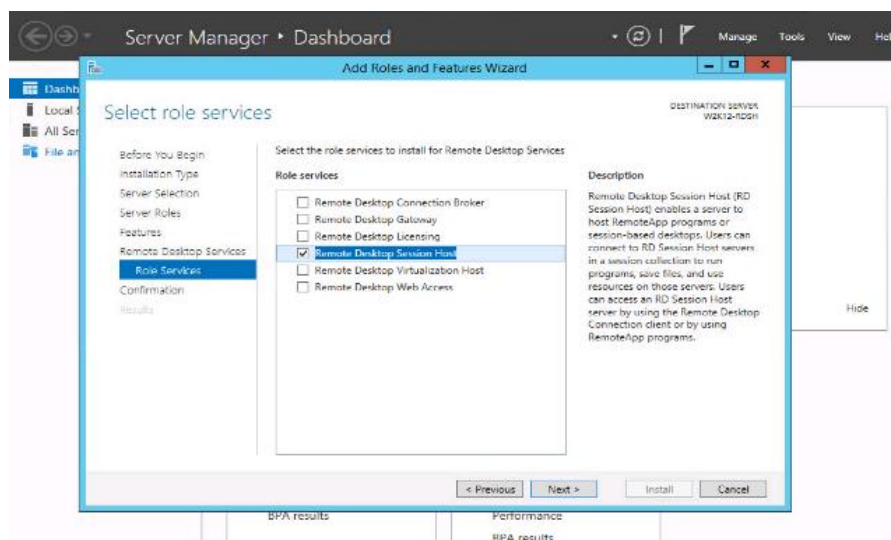


Prior to installing the Horizon View Agent on a Microsoft Server 2016 virtual machine, you must add the Remote Desktop Services role and the Remote Desktop Session Host role service.

- To add Remote Desktop Services role on Windows Server OS from the Server Manager, use the Add Roles and Features wizard:

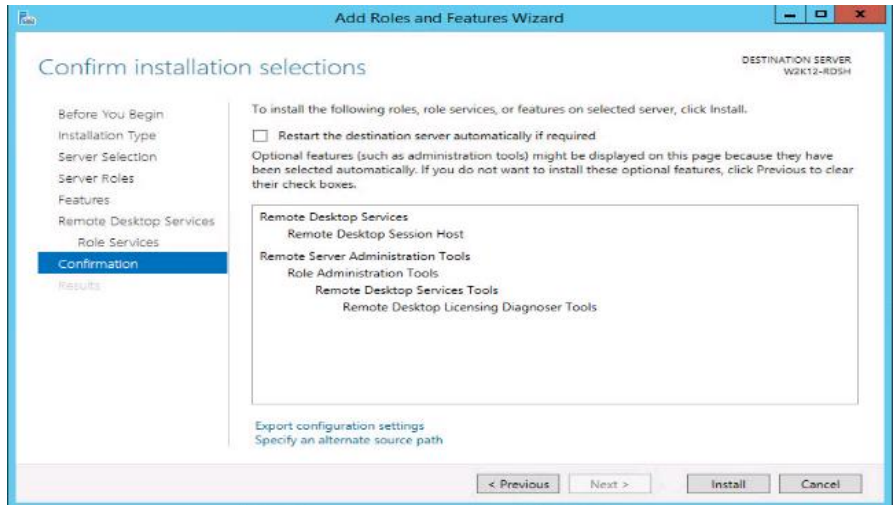


- Add Remote Desktop Session Host services.

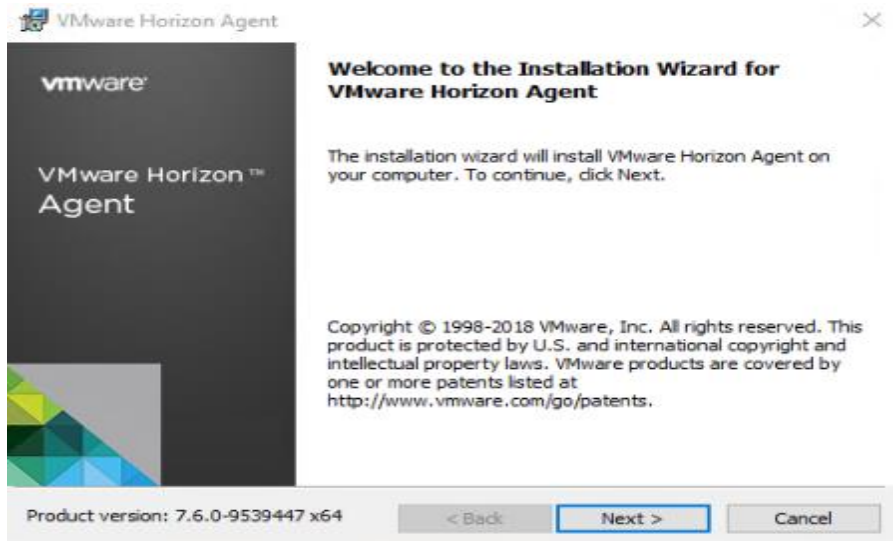


- Click Install.






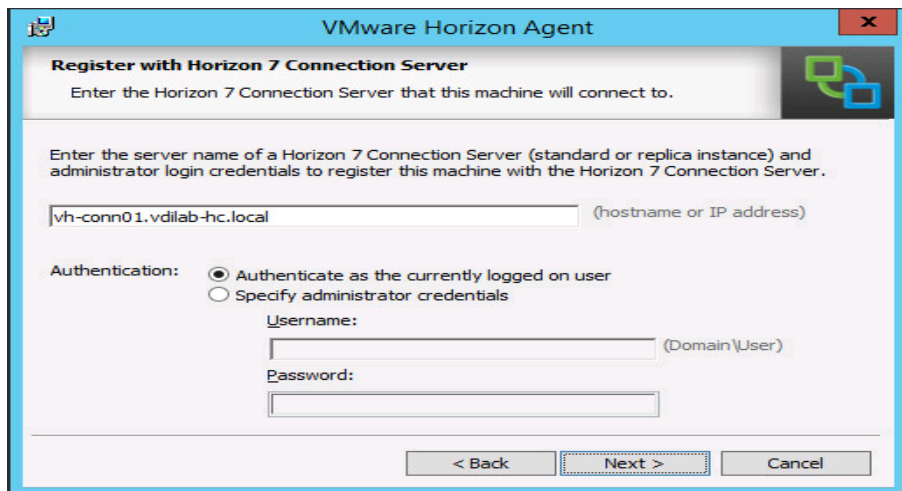
VMware Horizon Agent Installation:



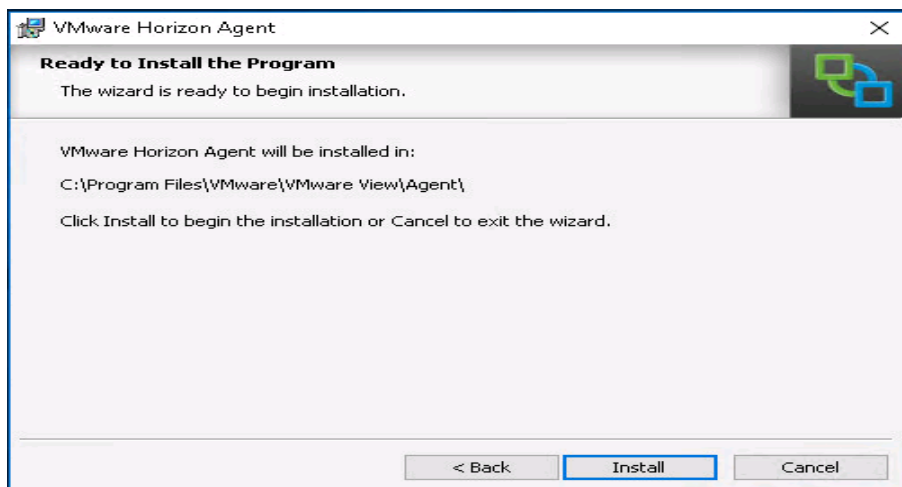
10. View Agent is will report as Install in "Desktop Mode" if Remote Desktop Services not installed.



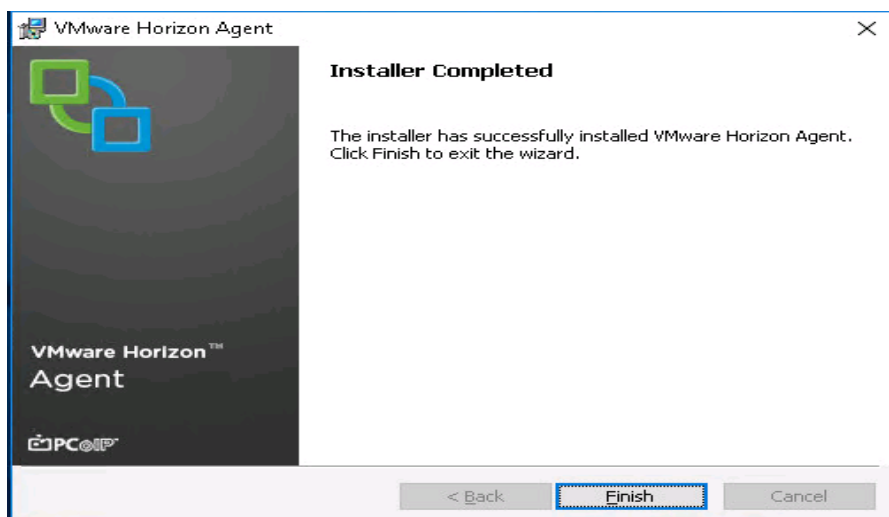
11. Add FQDN or IP address for Connection Server Instance to register the RDSH server.



12. Click Install.



13. Click Finish and restart the VM.



## Install Additional Software

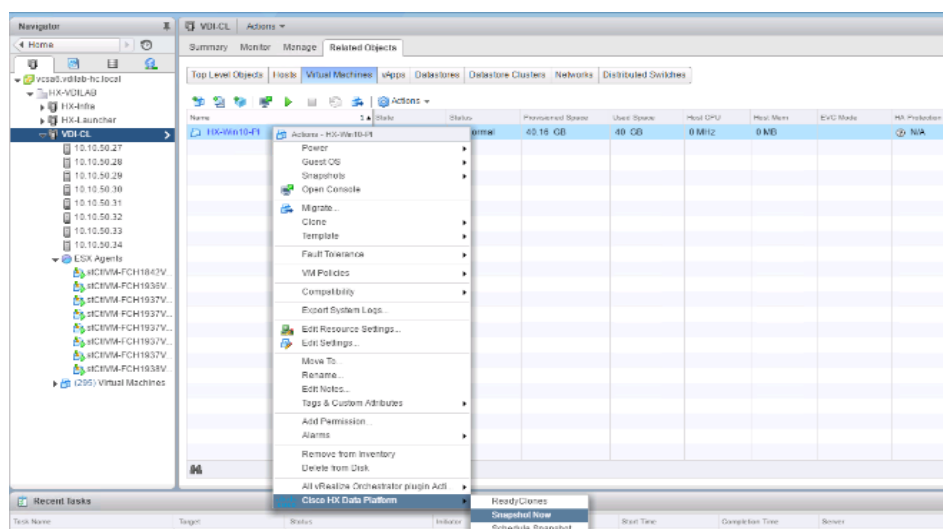
To install additional software required for your base windows image, complete the following steps:

1. For testing, we installed Microsoft Office 2016 64-bit version.
2. Log into the VSI Target software package to facilitate workload testing.
3. Install service packs and hot fixes required for the additional software components that are being added.
4. Reboot or shut down the VM as required.

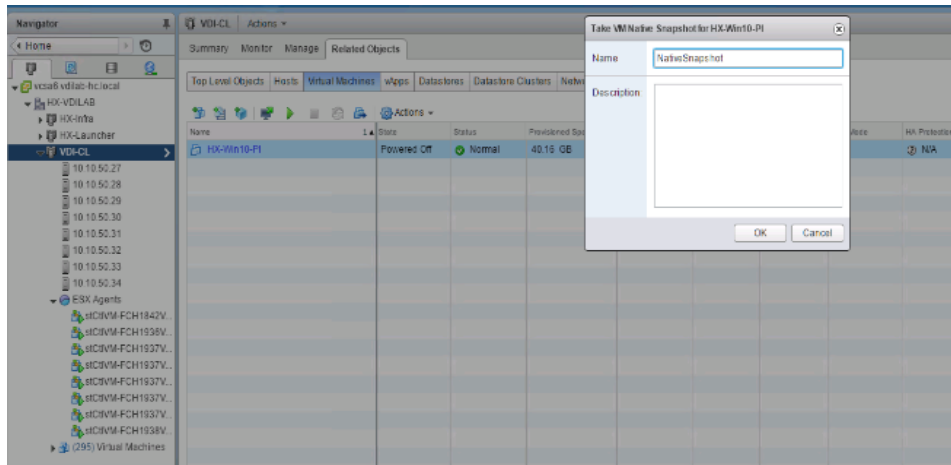
## Create a Native Snapshot for Automated Desktop Pool Creation

To create a native snapshot for the automated desktop pool, complete the following steps:

1. Log into vCenter WebUI.
2. Select the master image for the automated desktop pool creation.
3. Right-click, select Cisco HX Data Platform > SnapshotNow.



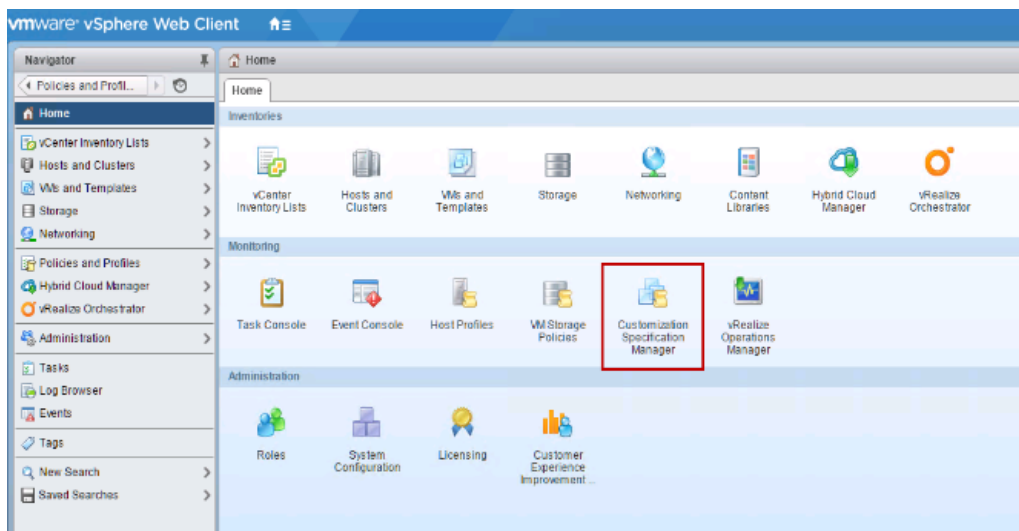
4. Enter a name for the HX native snapshot.



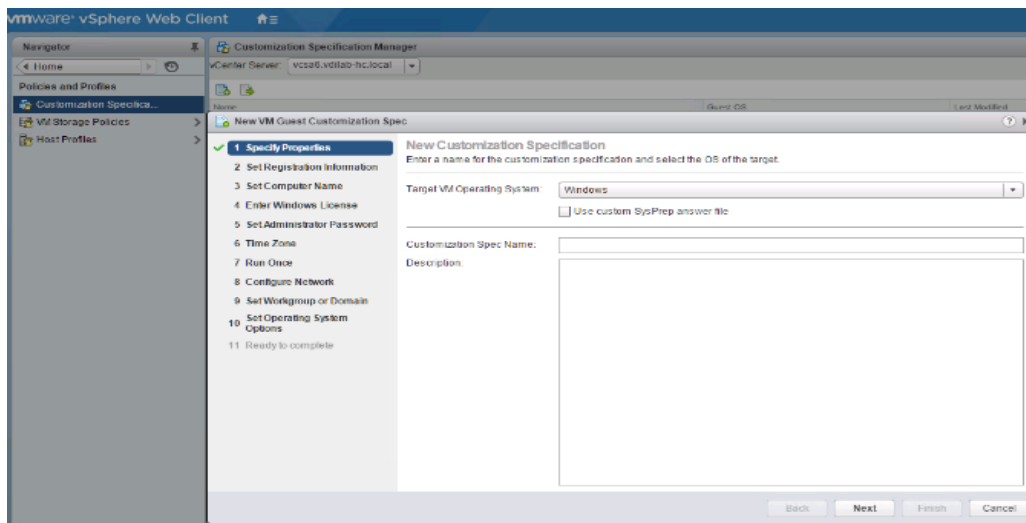
## Create Customization Specification for Virtual Desktops

To create Customization Specification for virtual desktops, complete the following steps:

1. On vCenter WebUI, select Customization Specification Manager.



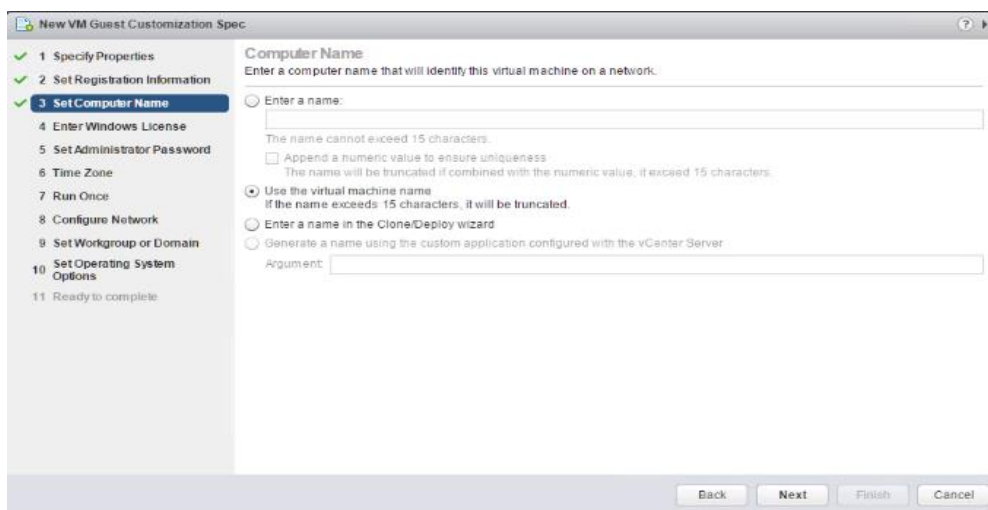
2. Select VM Operating System as Windows for Windows based guest OS optimization. Enter a name.



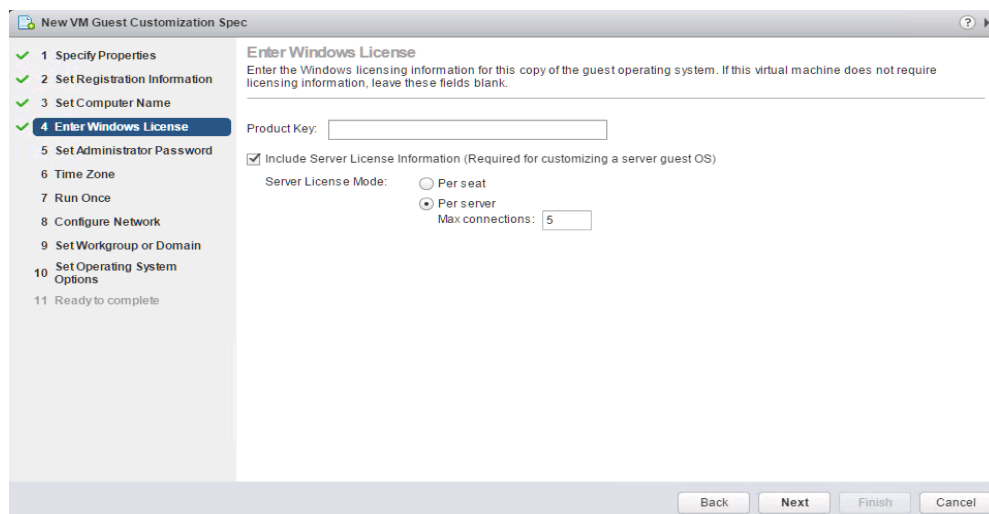
3. Provide name and organization details.



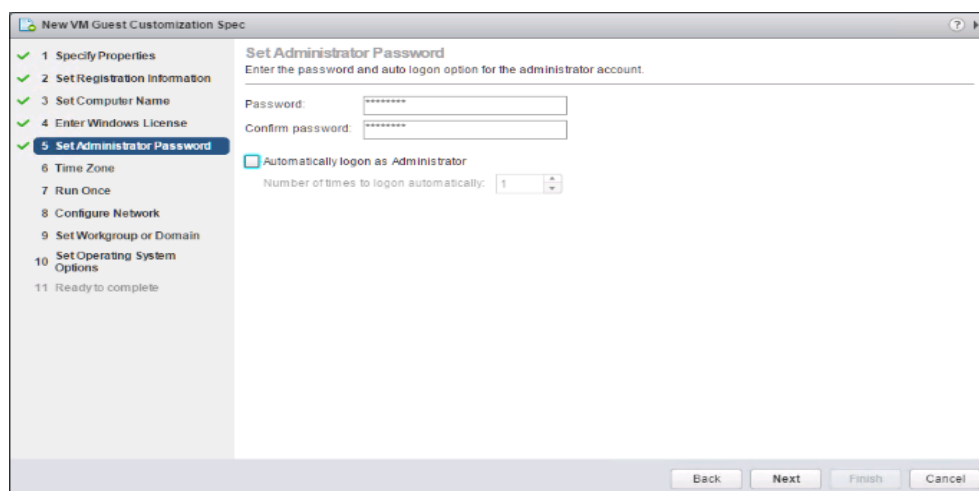
4. Provide a computer name. For this solution, we selected the radio button for Use the virtual machine name.




5. Provide the product License key if required.



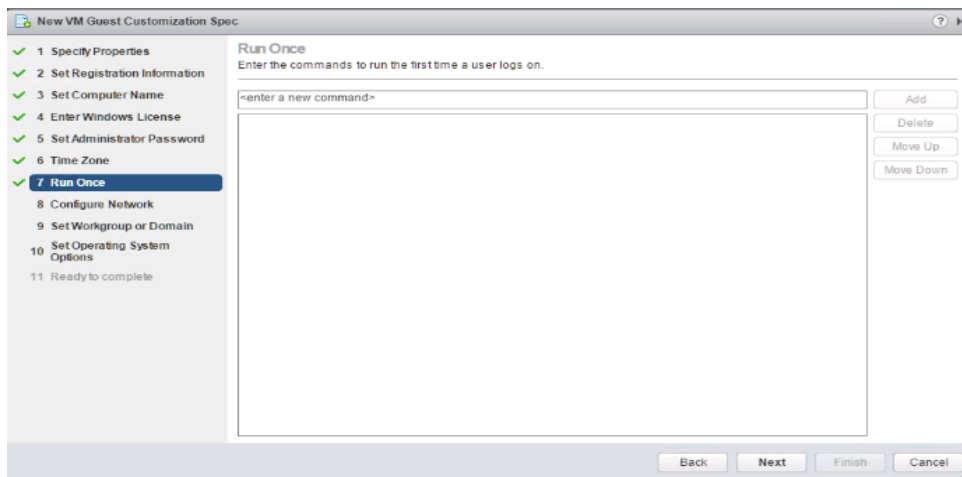
6. Provide Password credentials.



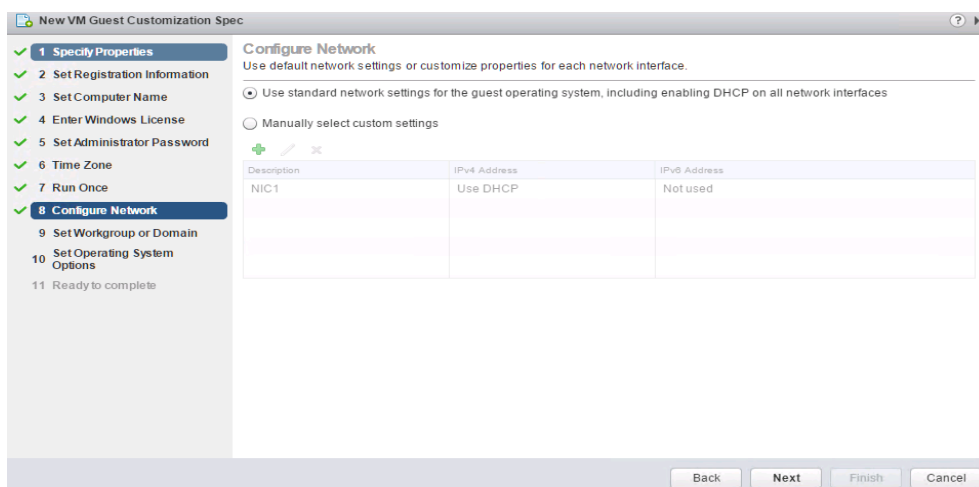
7. Select the Timezone.



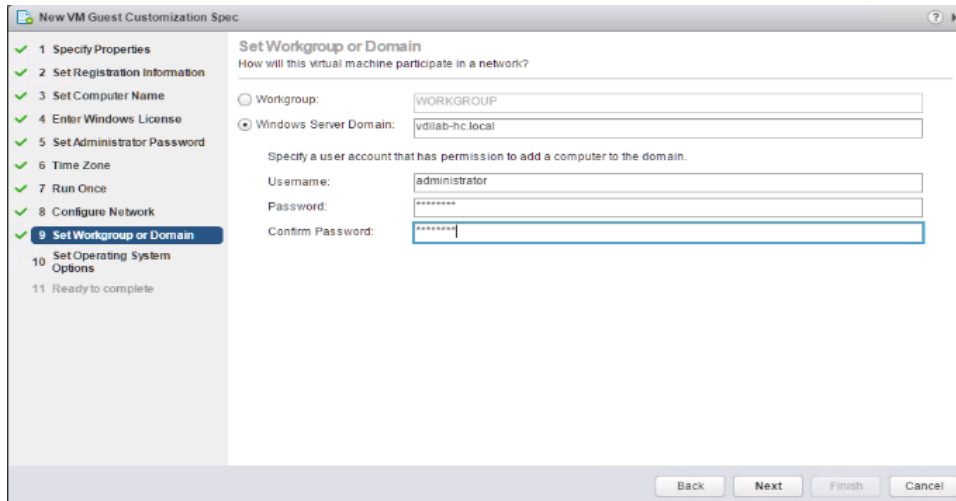
8. Add the commands to run when the first-time user logs in, if there are any.



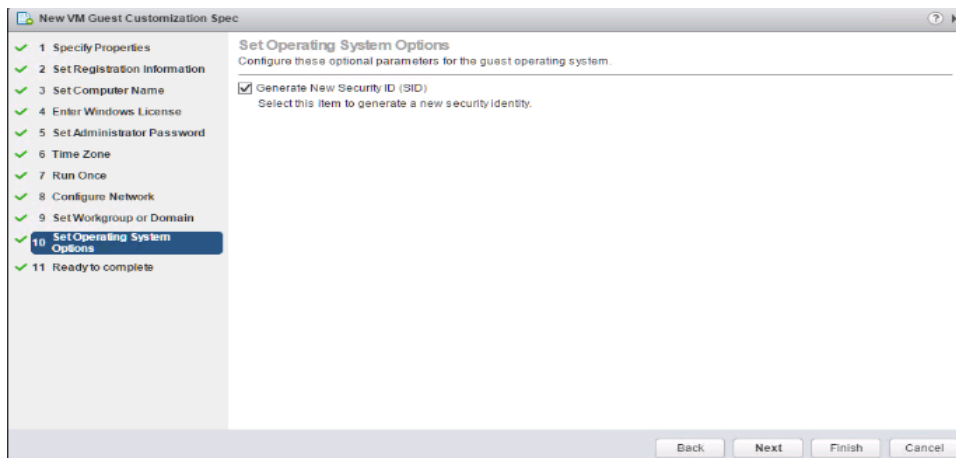
9. Provide the network information whether to use the DHCP server to assign IP address, or manual configuration.



10. Provide the domain name and user credentials.

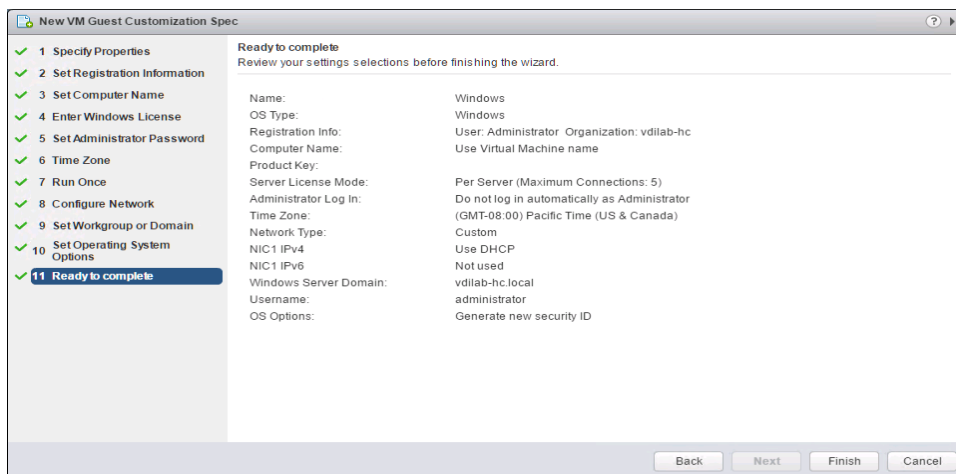


11. Select the checkbox Generate New Security ID (SID).



12. Review and click Next to complete creating the Customization Specs.

13. Click Finish.

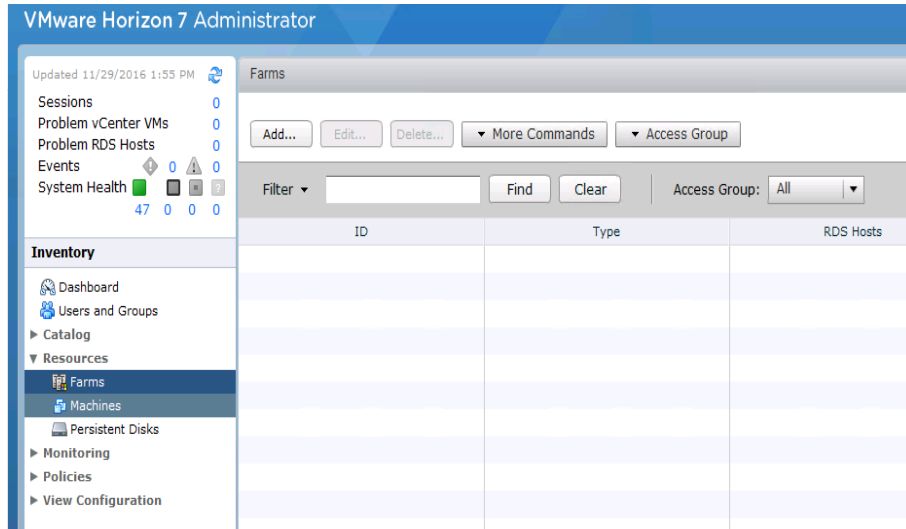




## RDSH Farm Creation

Before you can create an RDSH desktop pool, you must first create a RDSH Farm. To create a RDSH Farm, complete the following steps:

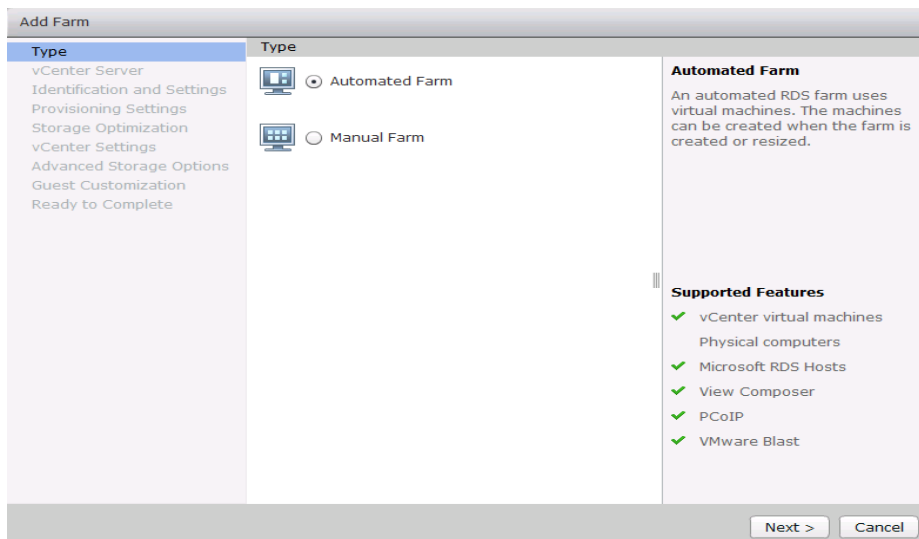
1. In the VMware Horizon Administration console, select Farms under the Resource node of the Inventory pane.
2. Click Add in the action pane to create a new RDSH Farm.



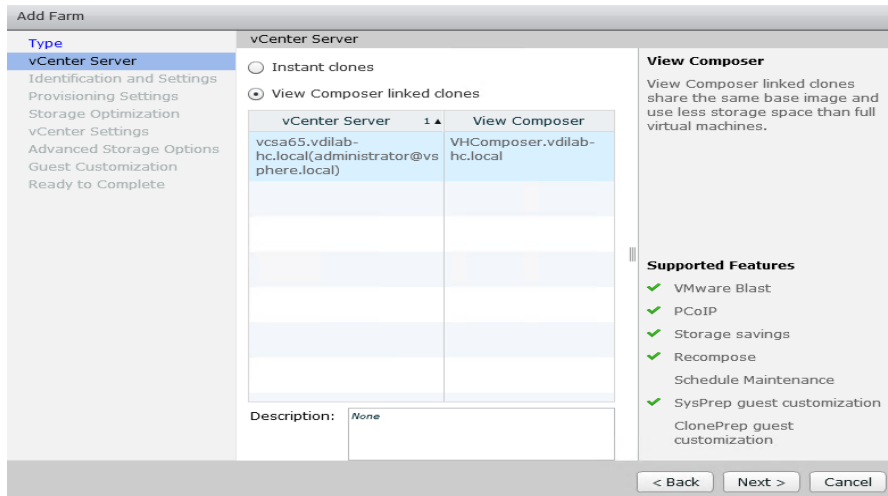
3. Select either to create an Automated or Manual Farm. In this solution, we selected Automated Farm.




**A Manual Farm requires a manual registration of each RDSH server to Horizon Connection or Replica Server instance.**

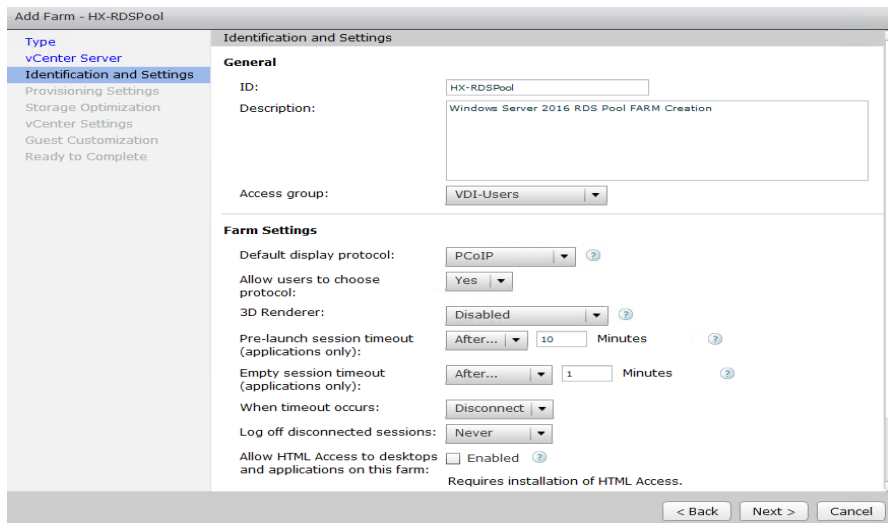


4. Select the vCenter Server and Horizon Composer server that you will use to deploy the Horizon RDSH Farm.
5. Click Next.



 You can choose to create either Instant clones or View Composer linked clones for the RDSH server FARM server VMs. Both have benefits and limitations, but detailing these differences are beyond the scope of this CVD. Please refer to your VMware documentation for more information.

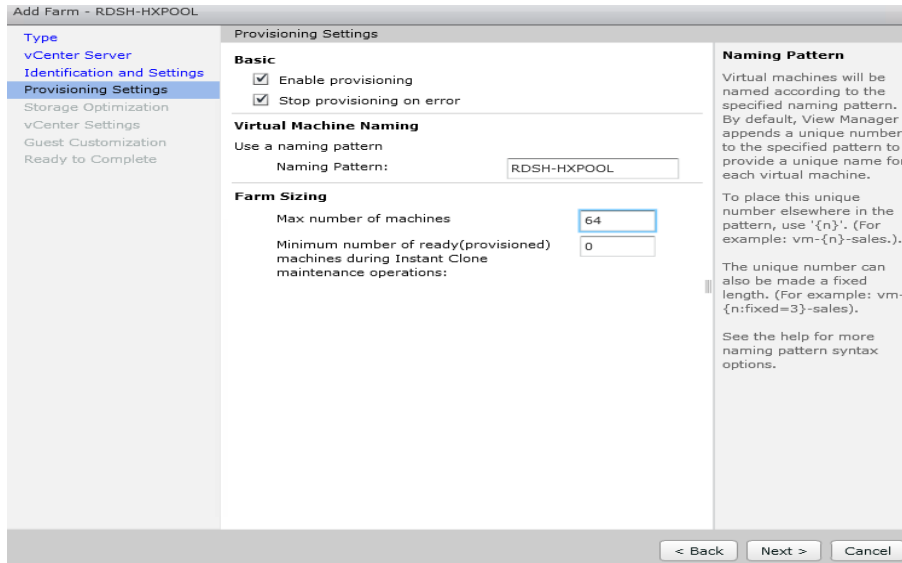
6. Enter the RDSH Farm ID, Access group, Default Display Protocol (Blast/PCoIP/RDP).
7. Select if users are allowed to change the default display protocol, Session timeout, Logoff Disconnected users, and select the checkbox to Enable HTML access.
8. Click Next.



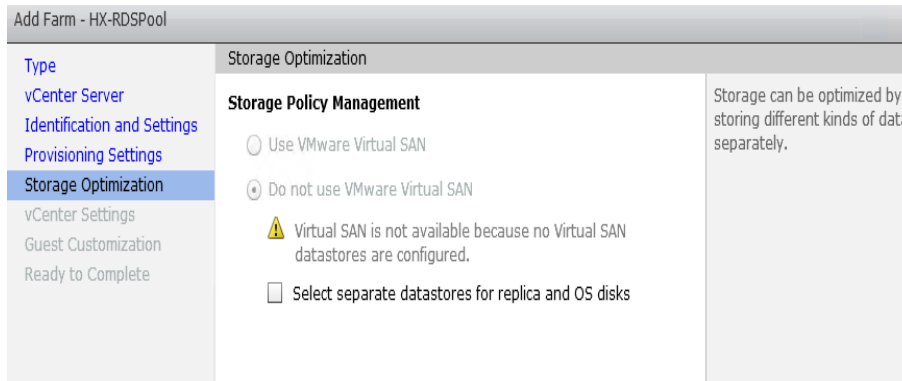
9. Select the provisioning settings, naming convention for RDSH server VM to deploy, and the number of VMs to deploy.

 In this study, we deployed 1032 RDSH virtual machines across our 16 node HyperFlex Cluster.

10. Click Next.

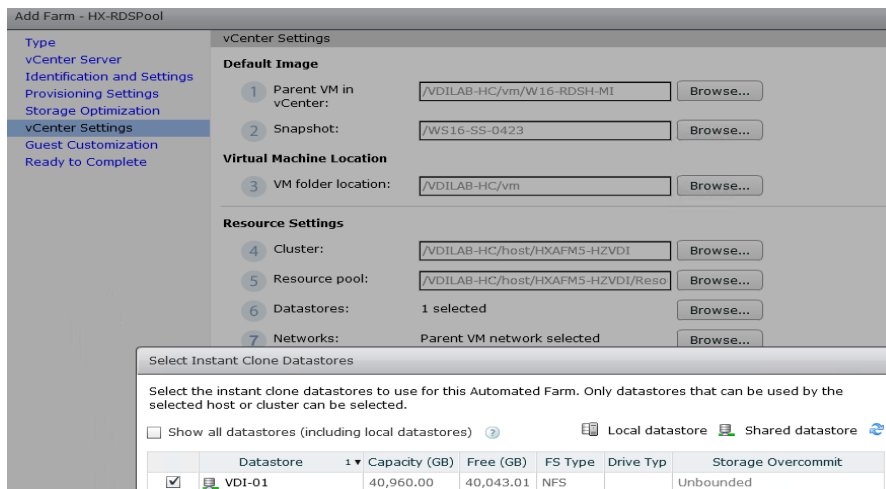


11. Click Next.



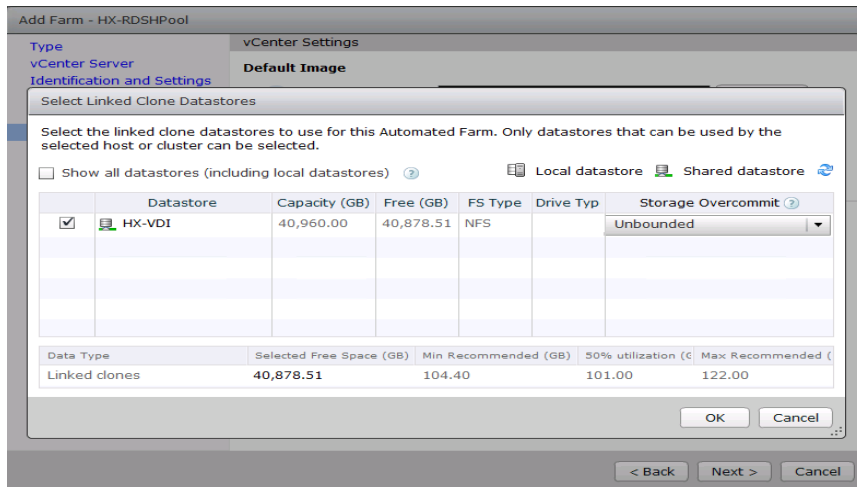
12. Select vCenter settings, for example; Master Image, snapshot, folder, Host or Cluster, resource pool, storage selection.

13. Click Next.



14. For Step 6 Datastores: Browse and choose Unbounded for the Storage Overcommit field.

15. Click OK.



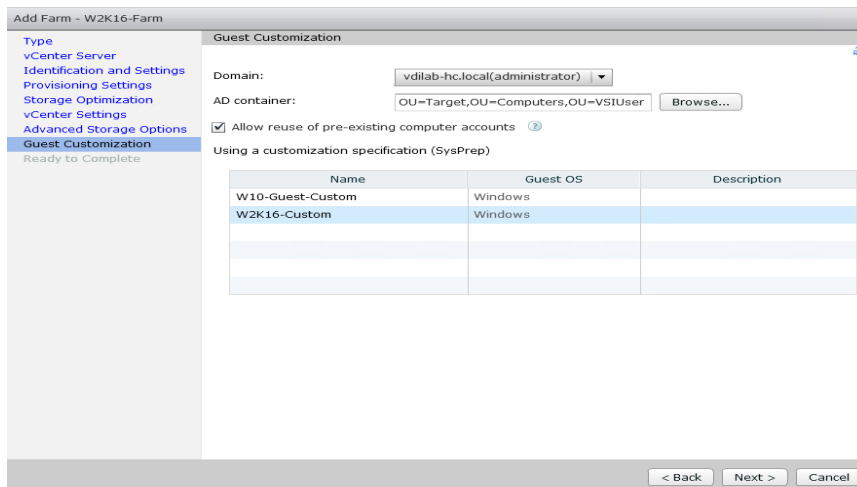
16. Click Next.

17. Select the Active Directory Domain, the Active Directory OU into which the RDSH machines will be provisioned, and the Sysprep file created as part of the customization specific configuration performed earlier.



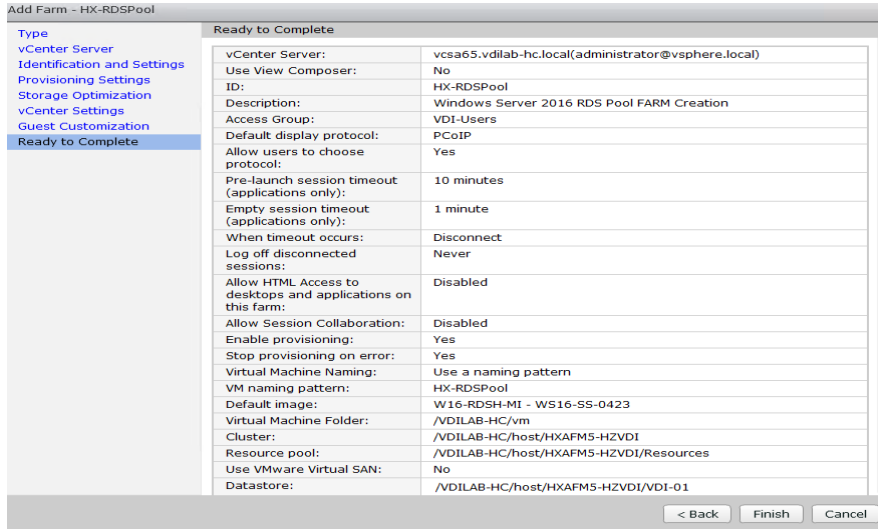
If you choose the instant clone pool for the RDSH FARM creation, you may not see the Sys prep guest customization step shown in the screenshot below.

18. Click Next.

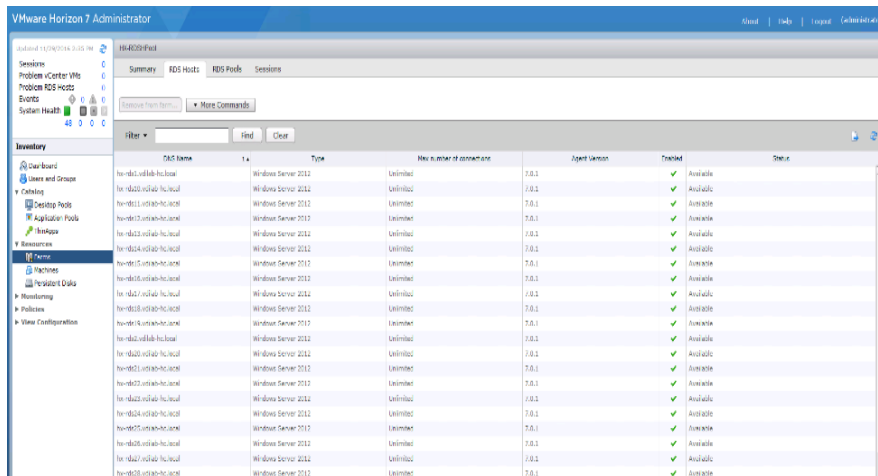
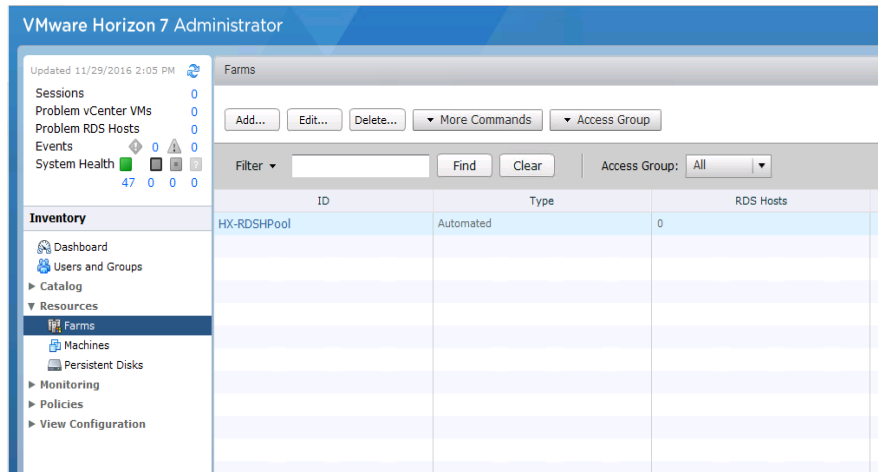


19. Review the pool creation information.

20. Click Finish.



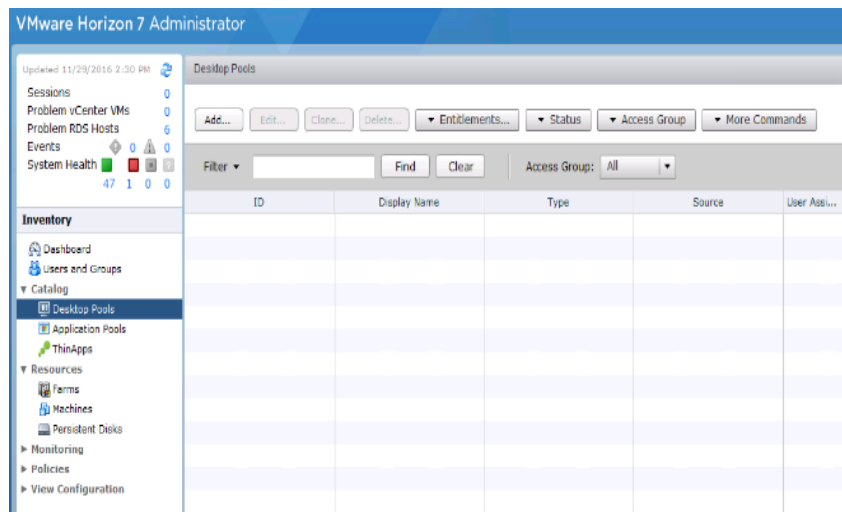
The VMware Horizon Administration console displays the status of the provisioning task and pool settings:



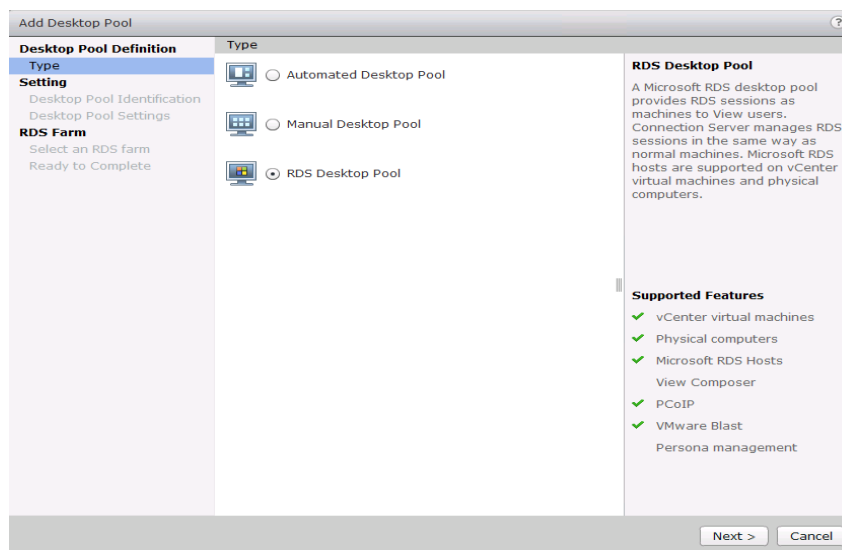
### Create the Horizon 7 RDS Published Desktop Pool

To create the Horizon 7 RDS Published Desktop Pool, complete the following steps:

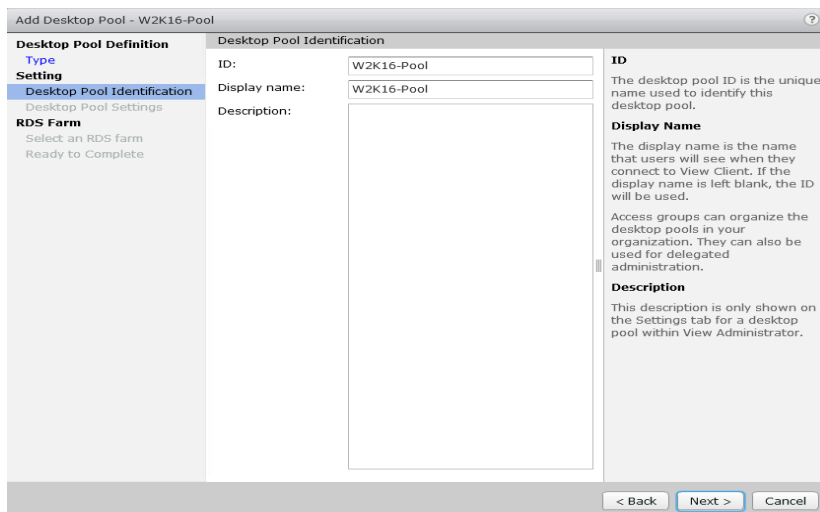
1. In the Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.
2. Click Add in the action pane.



3. Select RDS Desktop pool.
4. Click Next.

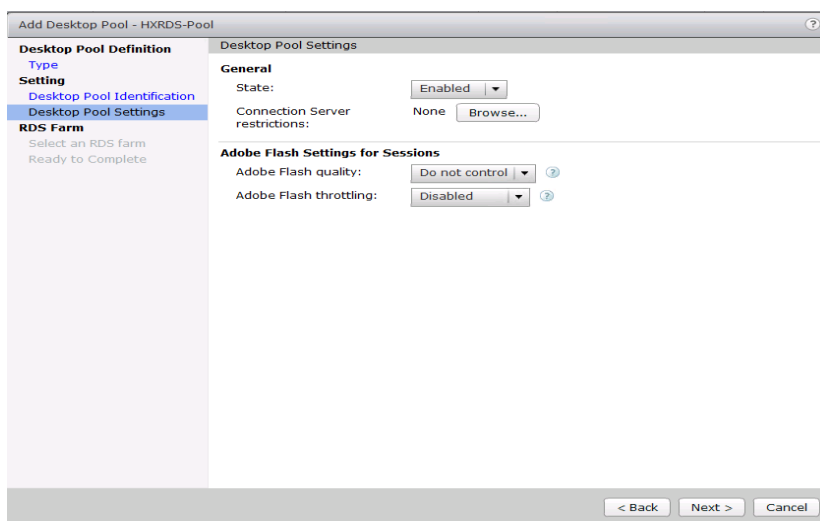


5. Enter Pool ID and Display name.
6. Click Next.



7. Accept the default settings on Desktop Pool Settings page.

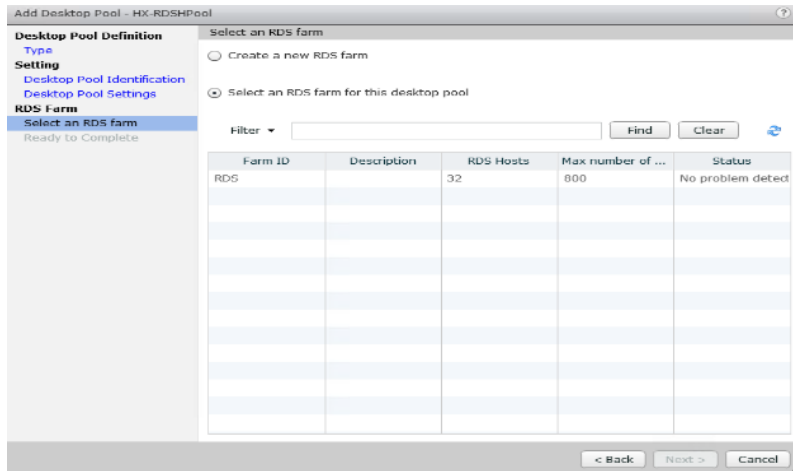
8. Click Next.



9. Click the “Select an RDS farm for this desktop pool” radio button.

10. Click the farm created in the previous section.

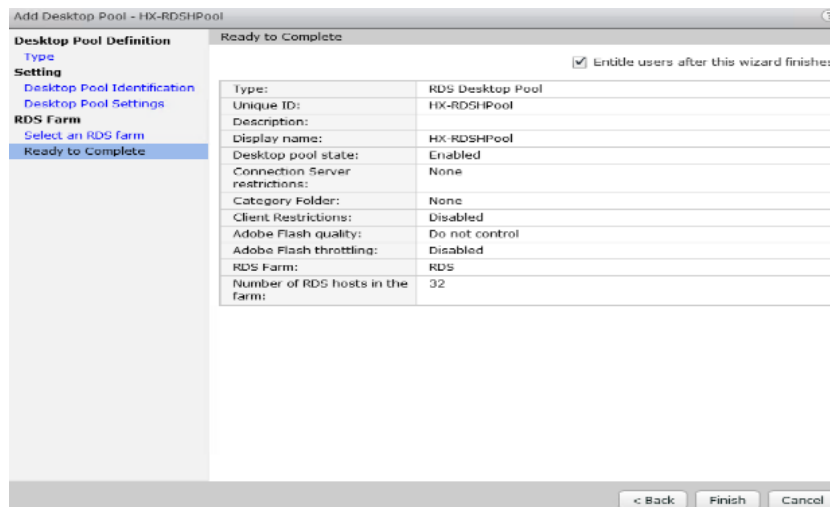
11. Click Next.



12. Review the pool settings.

13. Select the checkbox “Entitle users after this wizard finishes” to authorize users for the newly create RDSH desktop pool.

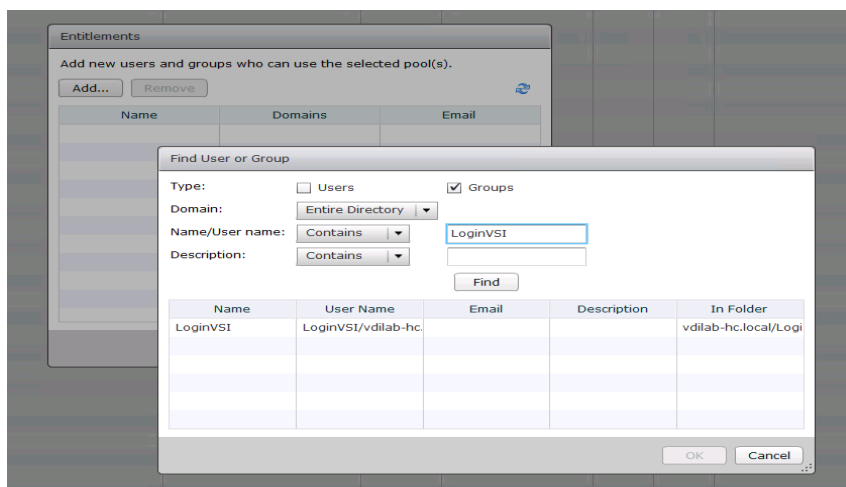
14. Click Finish.



15. Select the Users or Groups checkbox, use the search tools to locate the user or group to be authorized, highlight the user or group in the results box.

16. Click OK.



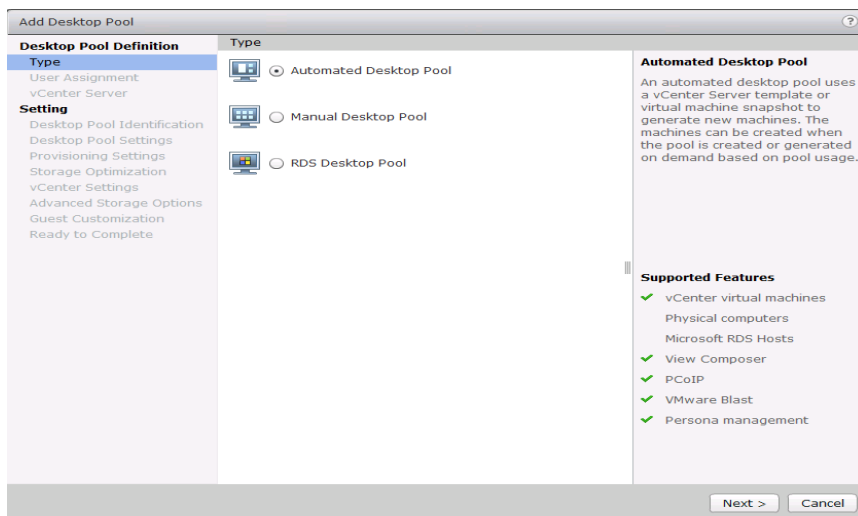


17. You now have a functional RDSH Farm and Desktop Pool with users identified who are authorized to utilize Horizon RDSH sessions.

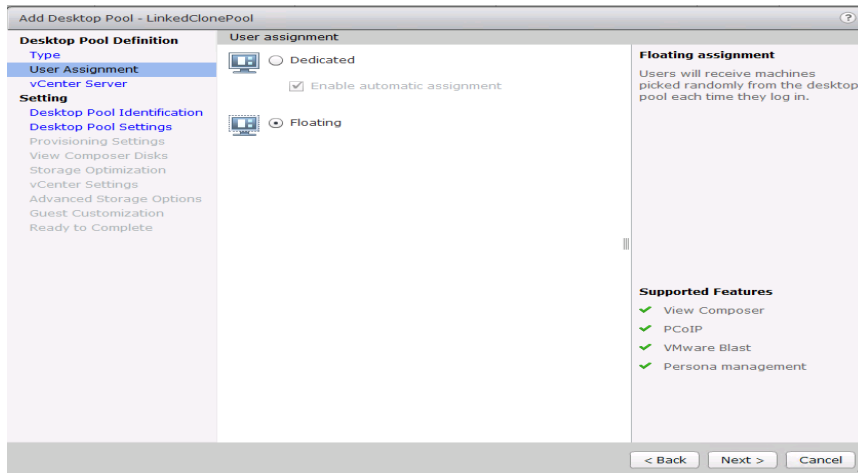
## VMware Horizon Linked-Clone Windows 10 Desktop Pool Creation

To create a VMware Horizon linked-clone Windows 10 Desktop Pool, complete the following steps:

1. In Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.
2. Click Add in the action pane.
3. Select assignment type for pool.
4. Click Next.

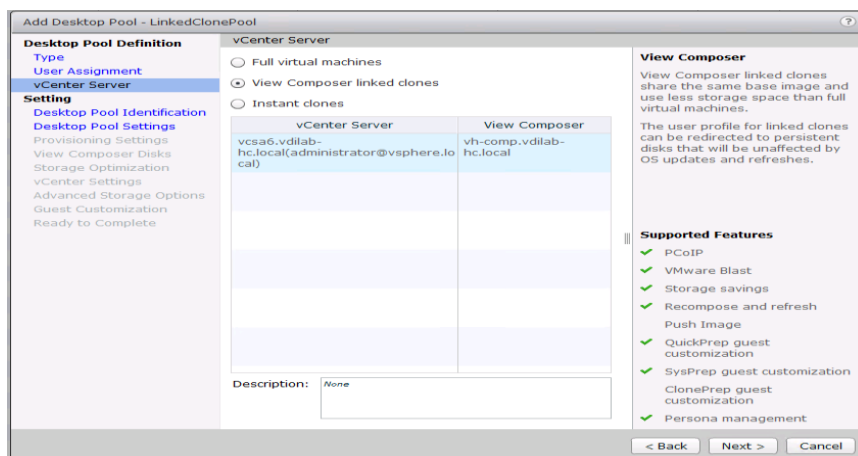


5. Select Floating or Dedicated user assignment.
6. Click Next.



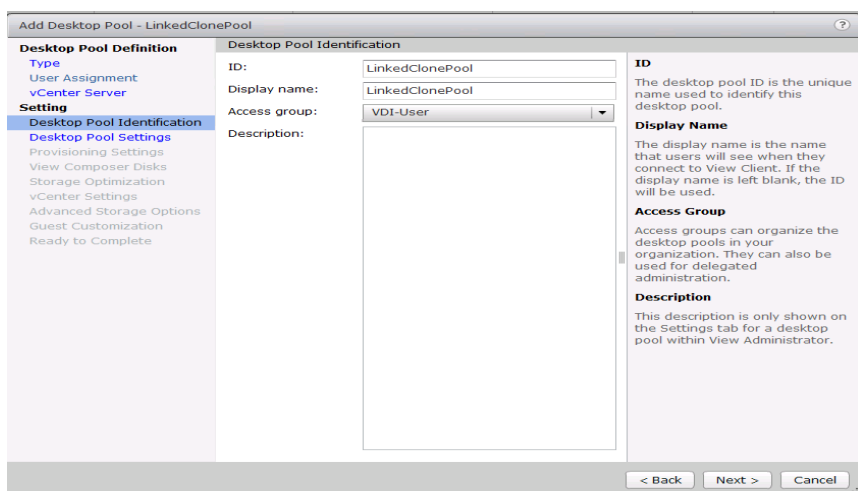
7. Select View Composer Linked Clones, highlight your vCenter and View Composer virtual machine.

8. Click Next.



9. Enter pool identification details.

10. Click Next.

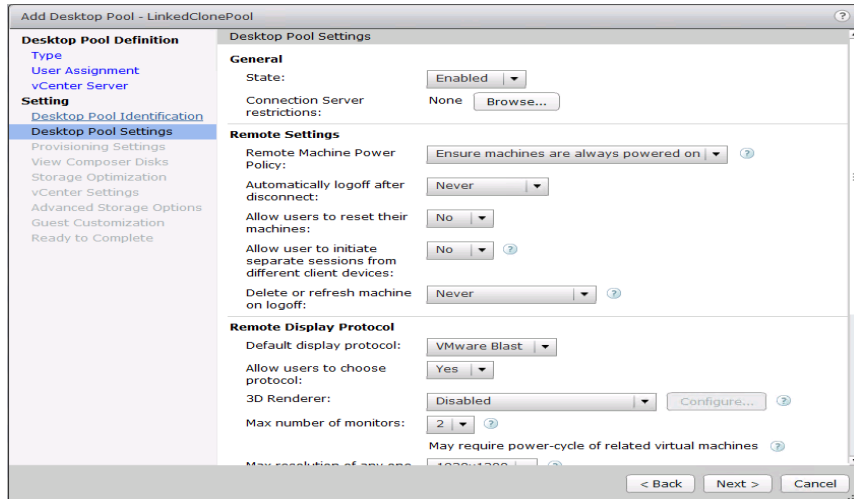


11. Select Desktop Pool settings.



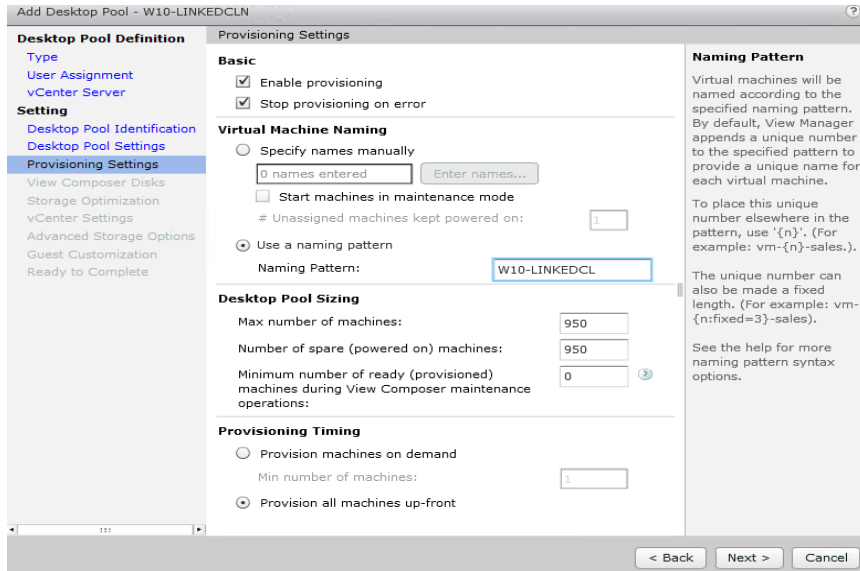
Be sure to scroll down in this dialogue to configure all options.

12. Click Next.



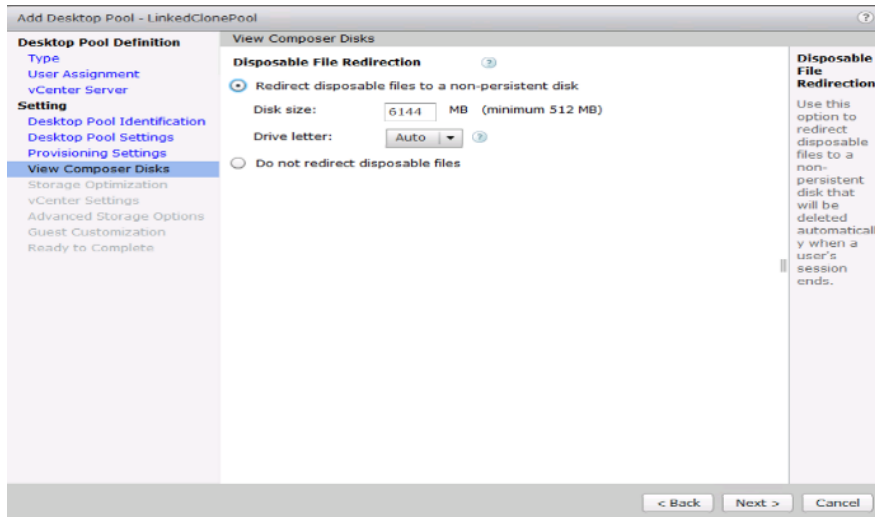
13. Select Provisioning Settings.

14. Click Next.

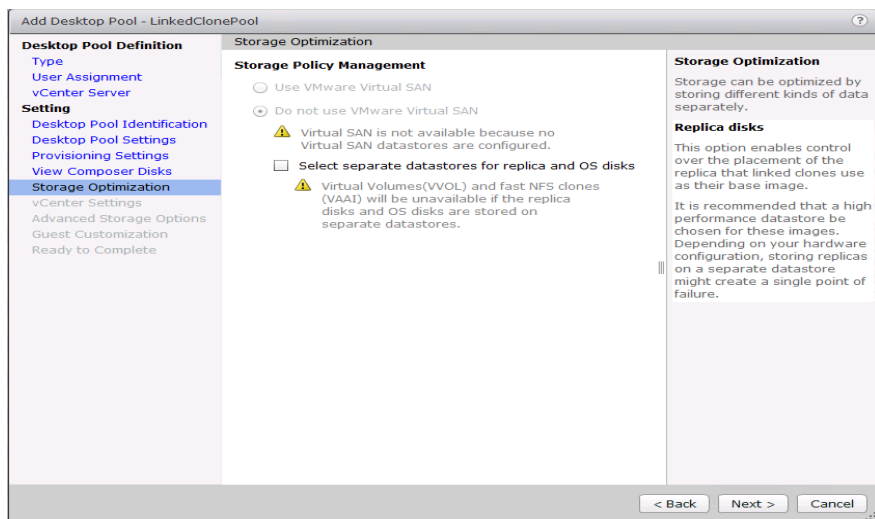


15. Select View Composer disk configuration.

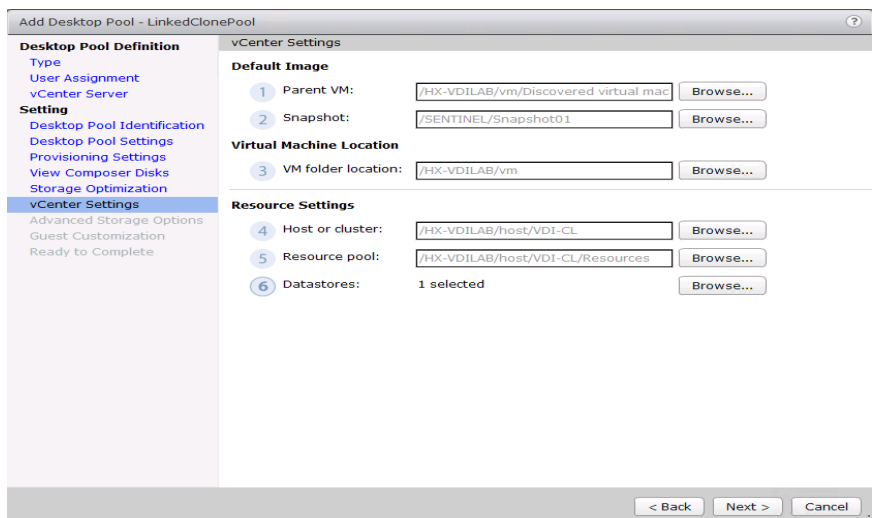
16. Click Next.



17. Click Next.




18. Select each of the six required vCenter Settings by using the Browse button next to each field.



19. For Datastore selection, select the correct datastore and set the Storage Overcommit as “Unbounded.”

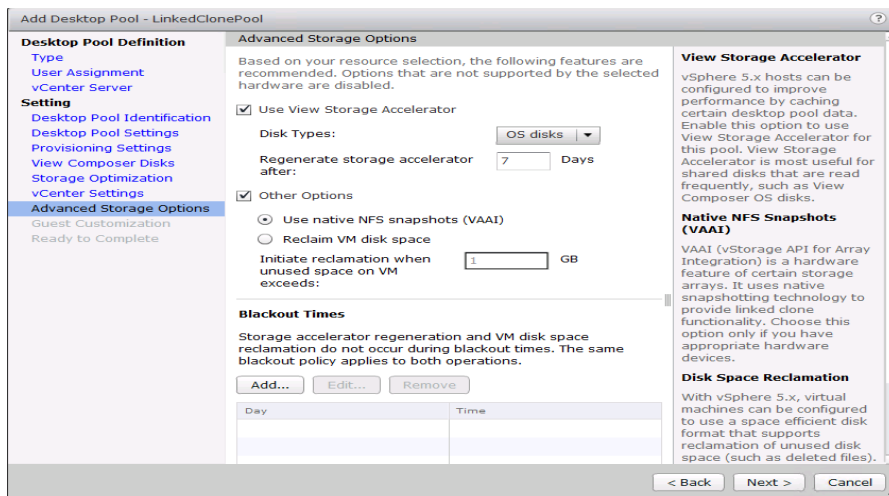
20. Click OK.

21. Click Next.



22. Set the Advanced Storage Options using the settings in the following screenshot.

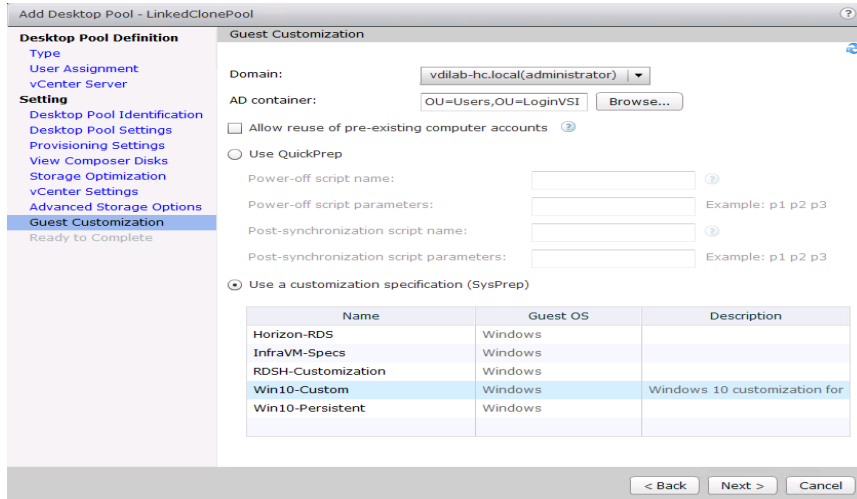
23. Click Next.



24. Select Guest optimization settings.

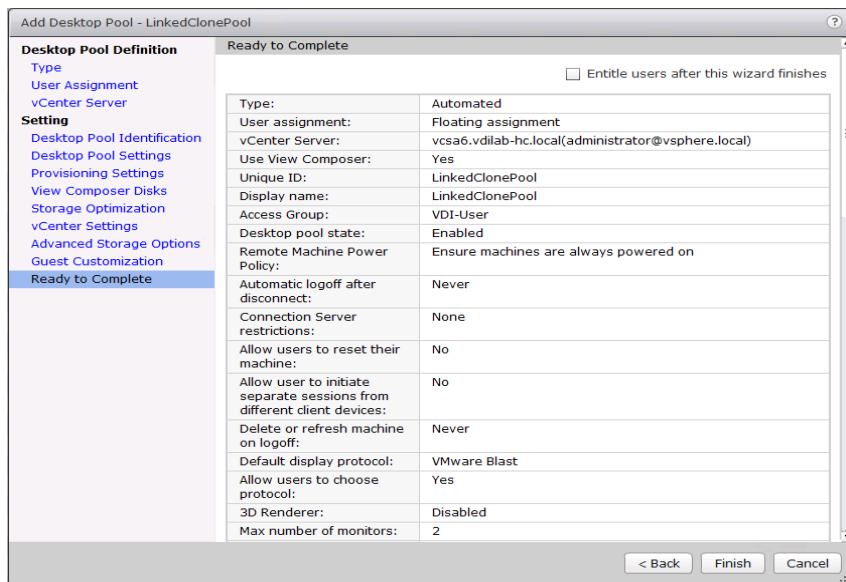
25. Select the Active Directory domain, browse to the Active Directory Container where the virtual machines will be provisioned and then choose either the QuickPrep or Sysprep option you would like to use. Highlight the Customization Spec previously prepared.

26. Click Next.



27. Select the checkbox “Entitle users after pool creation wizard completion” if you would like to authorize users as part of this process. Follow instructions provided in the Create Horizon 7 RDS Desktop Pool to authorize users for the Linked Clone Pool.

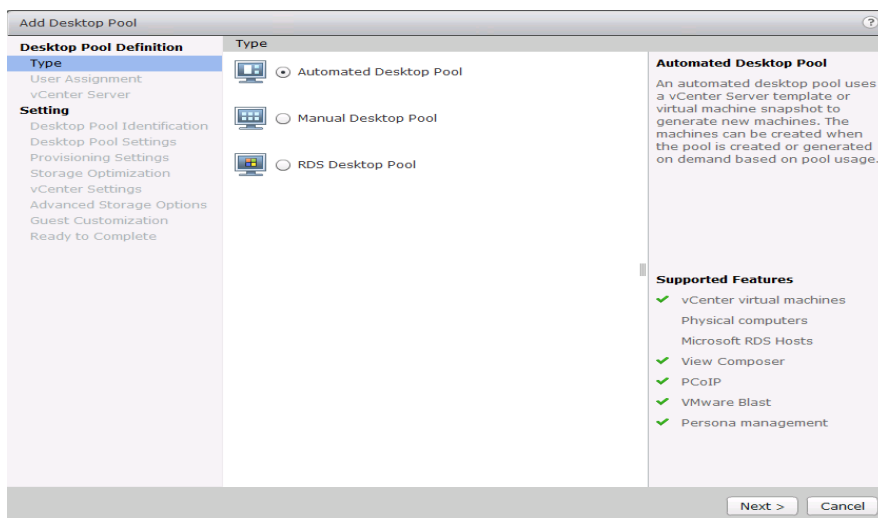
28. Click Finish to complete the Linked Clone Pool creation process.



## VMware Horizon Instant-Clone Windows 10 Desktop Pool Creation

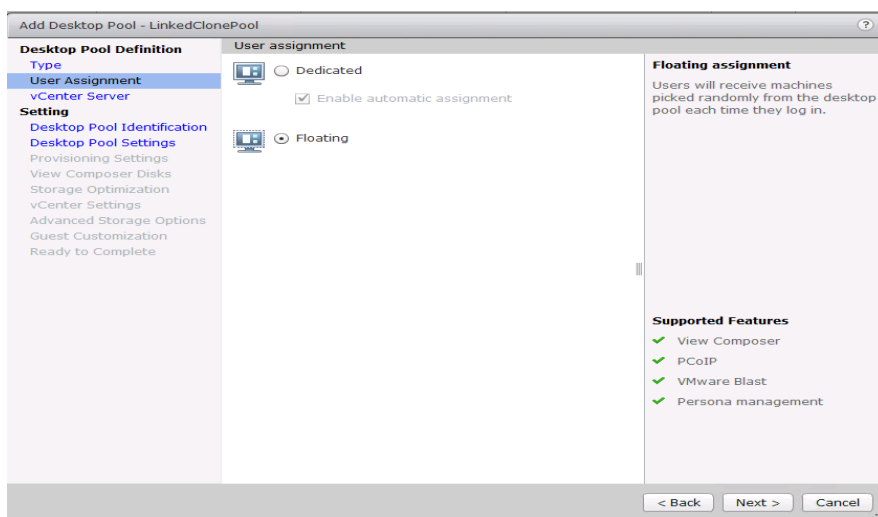
To create the VMware Horizon Instant-Clone Windows 10 Desktop Pool, complete the following steps:

1. In Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.
2. Click Add in the action pane.
3. Select Automated assignment type for pool.
4. Click Next.

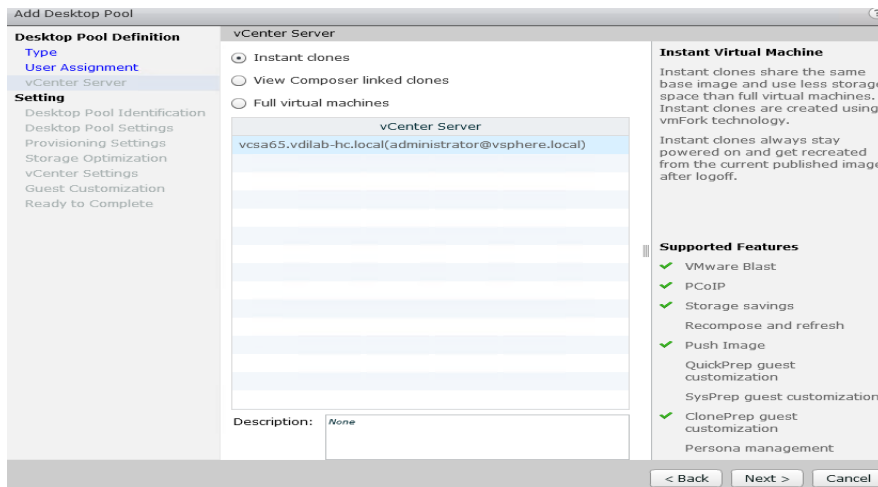


5. Select Floating or Dedicated user assignment.

6. Click Next.

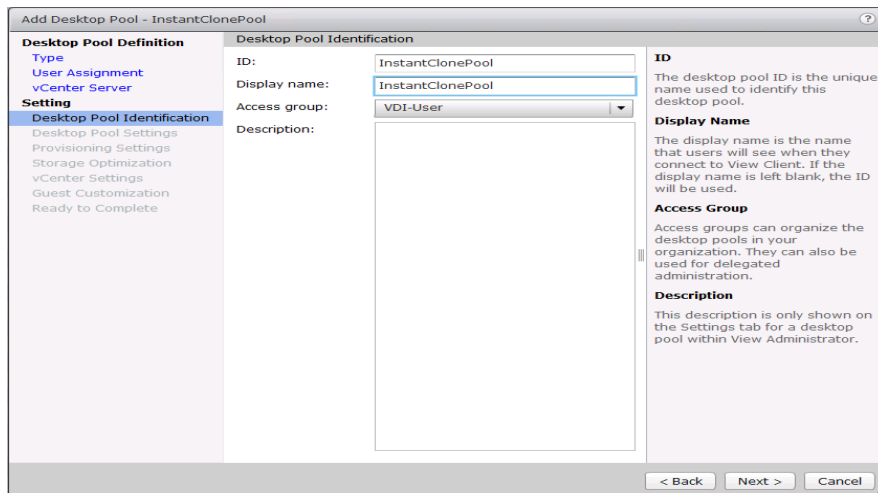


7. Select Instant Clones, highlight your vCenter server, then click Next.



8. Enter pool identification details.

9. Click Next.



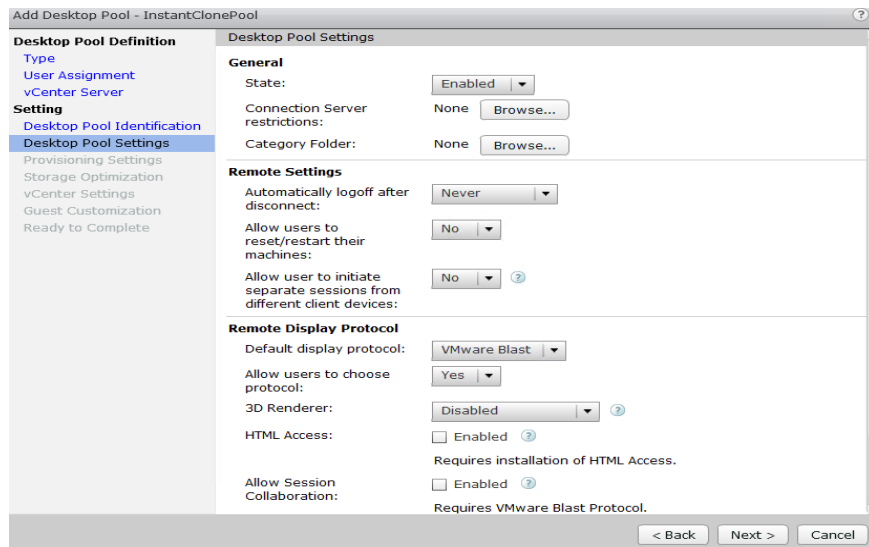
10. Select Desktop Pool settings.



Be sure to scroll down to choose the Acrobat Flash settings.

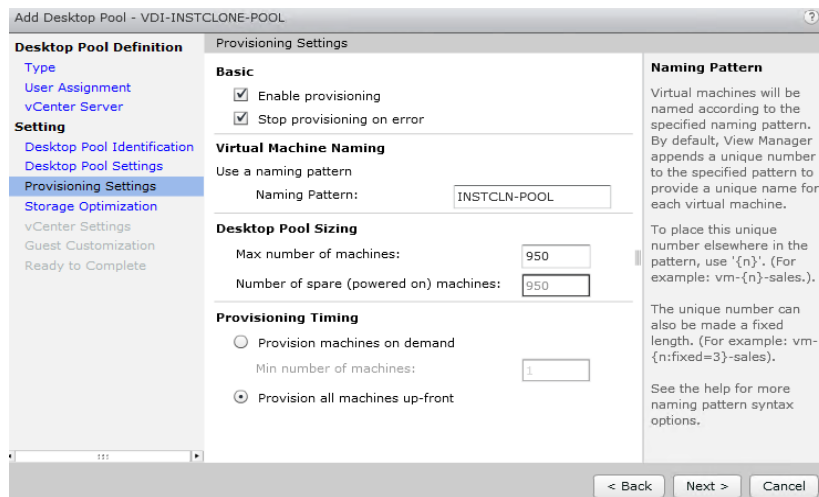
11. Click Next.



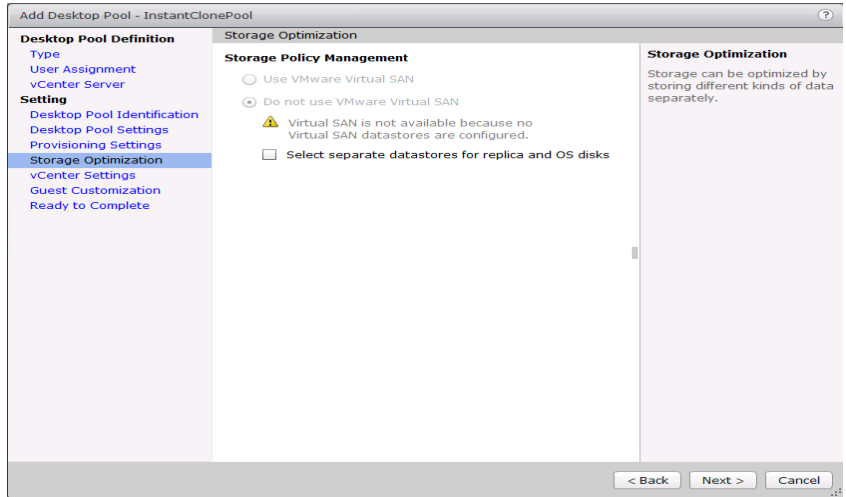


12. Select provisioning settings.

13. Click Next.

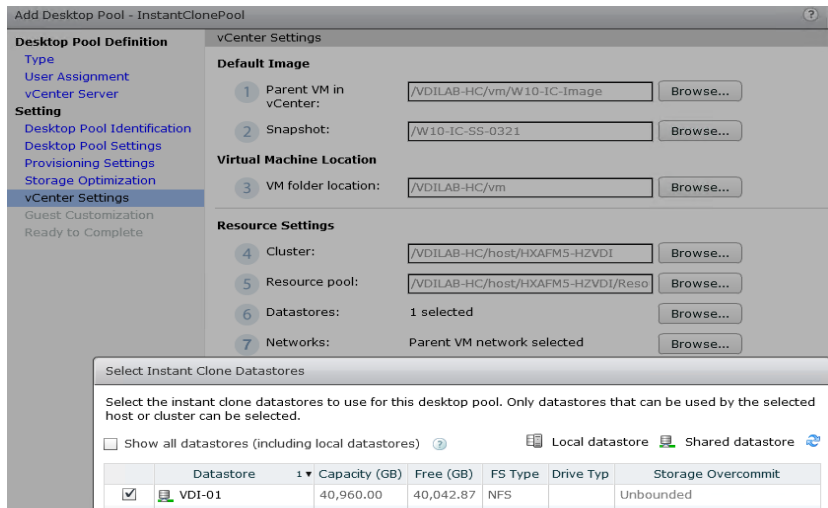


14. Click Next.



15. Select the vCenter Settings and browse for each of the six required inputs.

16. Click Next.



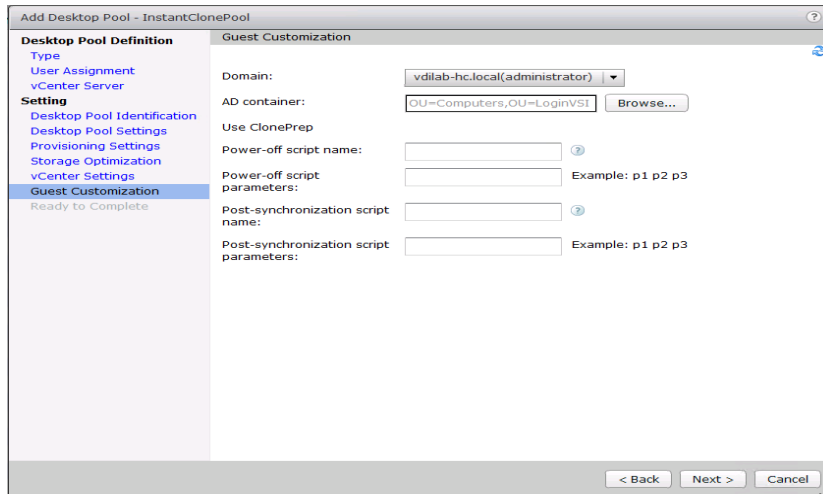
17. For Datastore selection, select the datastore with the storage overcommit as "Unbounded."

18. Click OK.

19. Select Guest Customization.

20. Browse to your Active Directory Domain and to the AD container into which you want your Instant Clone machines provisioned.

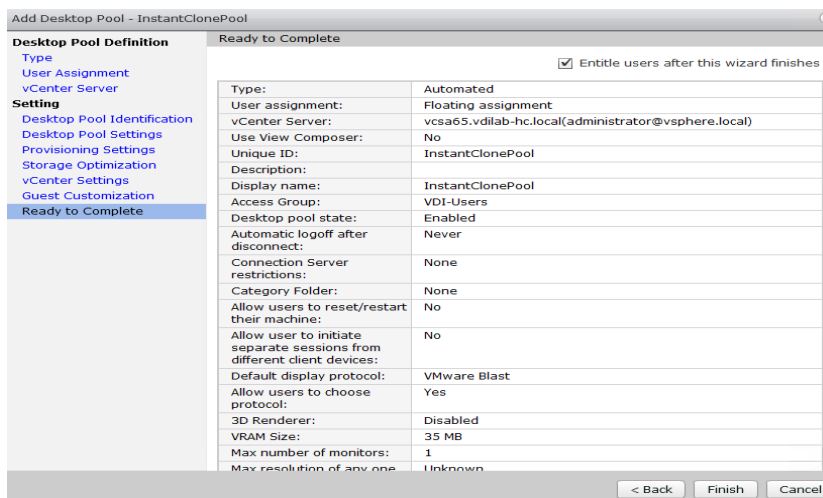
21. Click Next.



22. Review the summary of the pool configuration.

23. Select the checkbox “Entitle users after pool creation wizard completion” to authorize users or groups for the new pool.

24. Click Finish.

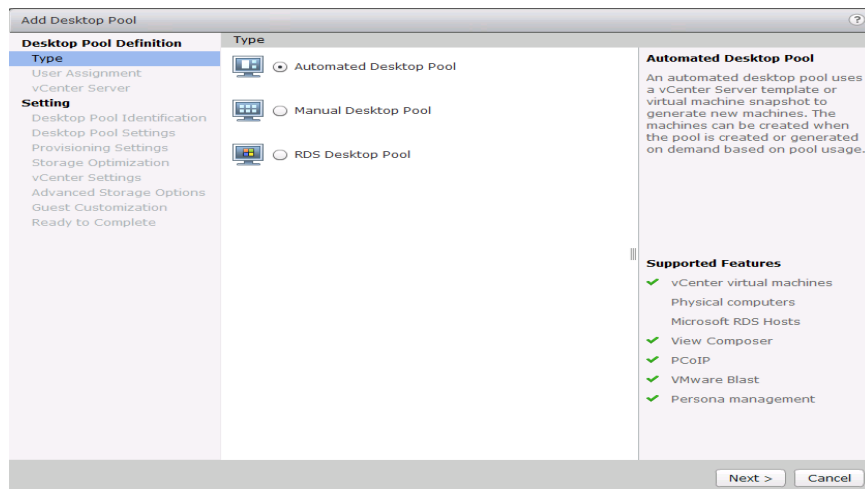


25. Follow the instructions provided in the Create Horizon 7 RDS Desktop Pool to authorize users for the Linked Clone Pool.

## VMware Horizon Persistent Windows 10 Desktop Pool Creation

To create the VMware Horizon Persistent Windows 10 Desktop Pool, complete the following steps:

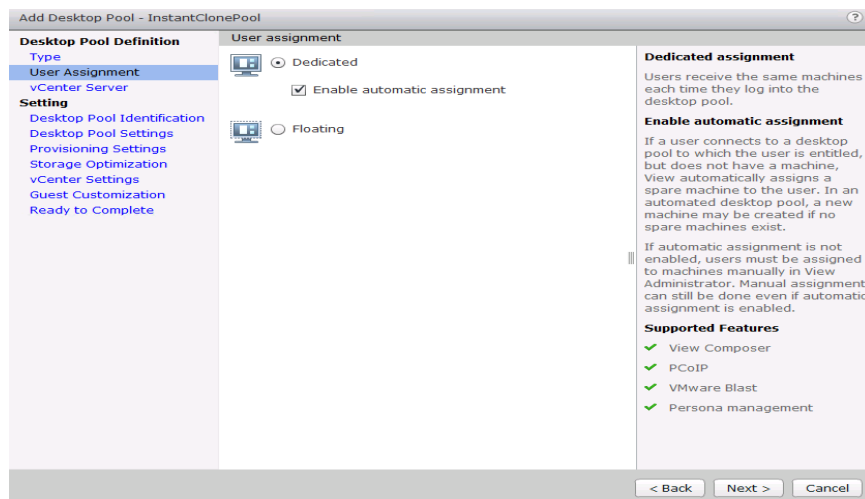
1. In Horizon Administrator console, select Desktop Pools in the Catalog node of the Inventory pane.
2. Click Add in the action pane.
3. Select assignment type for pool.
4. Click Next.



5. Select the Dedicated radio button.

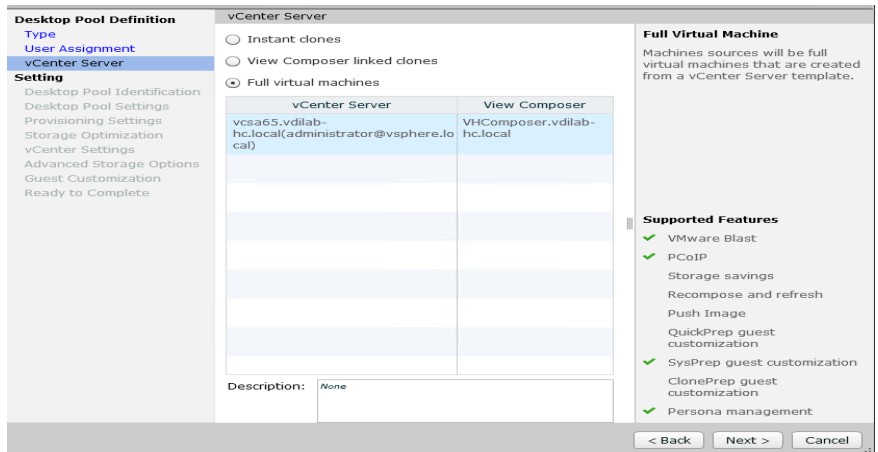
6. Select the Enable automatic assignment checkbox, if desired.

7. Click Next.



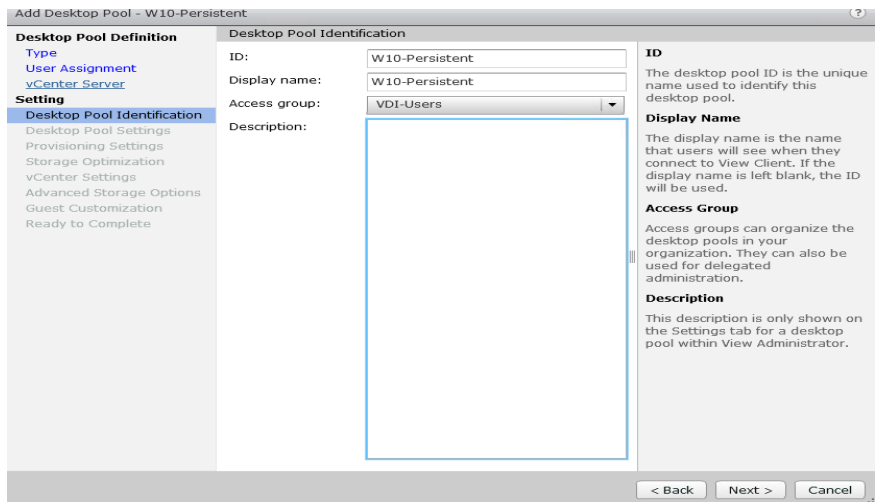
8. Select the Full Virtual Machines radio button and highlight your vCenter and Composer.

9. Click Next.



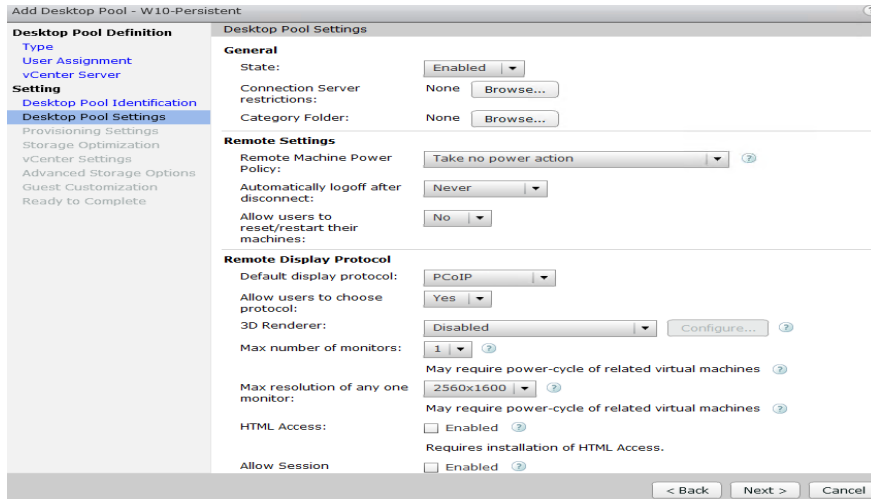
10. Enter the pool identification details.

11. Click Next.



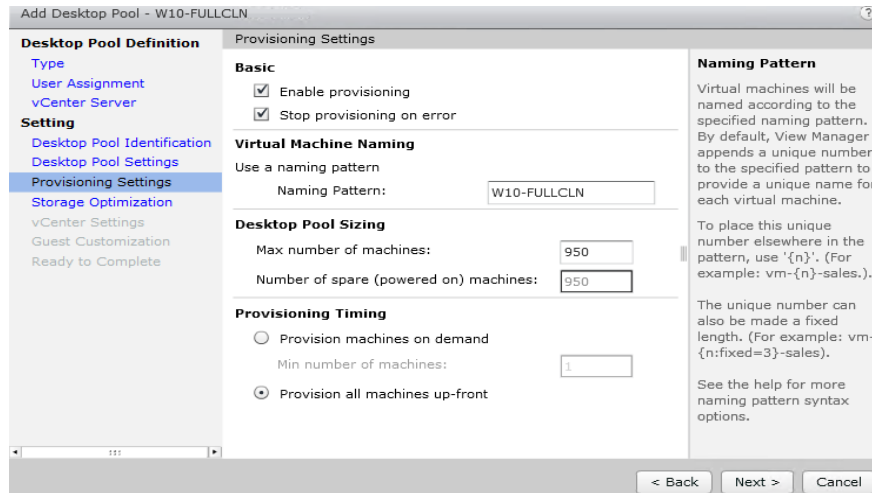
12. Select Desktop Pool settings.

13. Click Next.

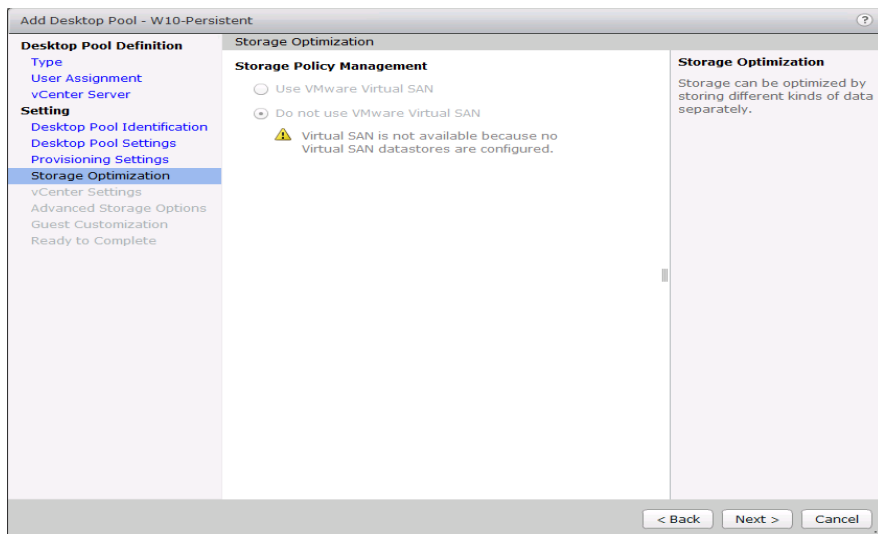


14. Select the provisioning settings to meet your requirements.

15. Click Next.

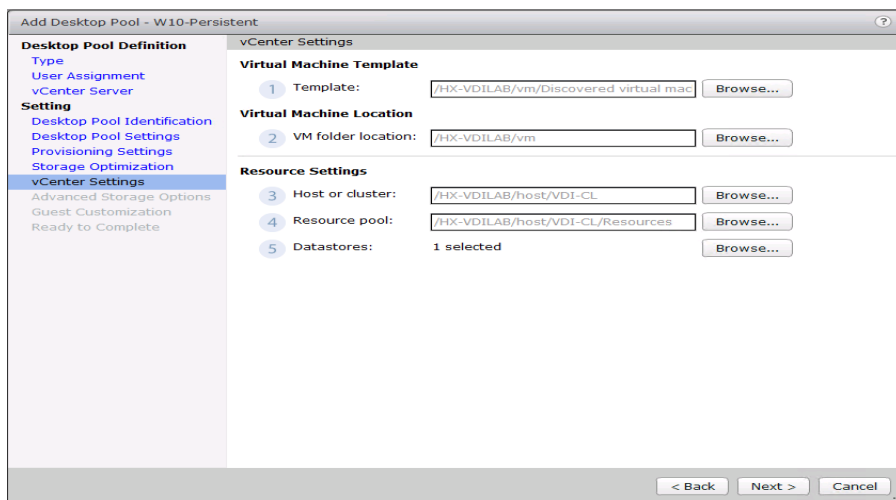


16. Click Next.




17. Select each of the five vCenter Settings.

18. Click Next.



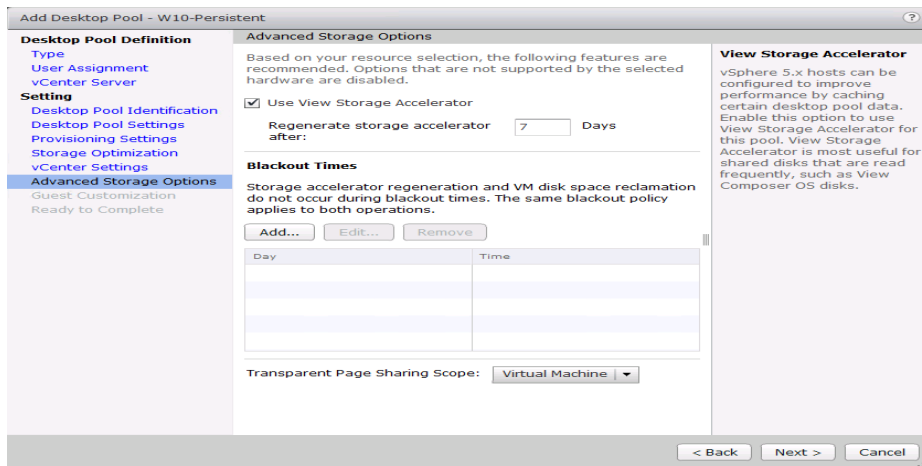
19. For Datastore selection, select the datastore with storage overcommit as "Unbounded."

20. Click OK.



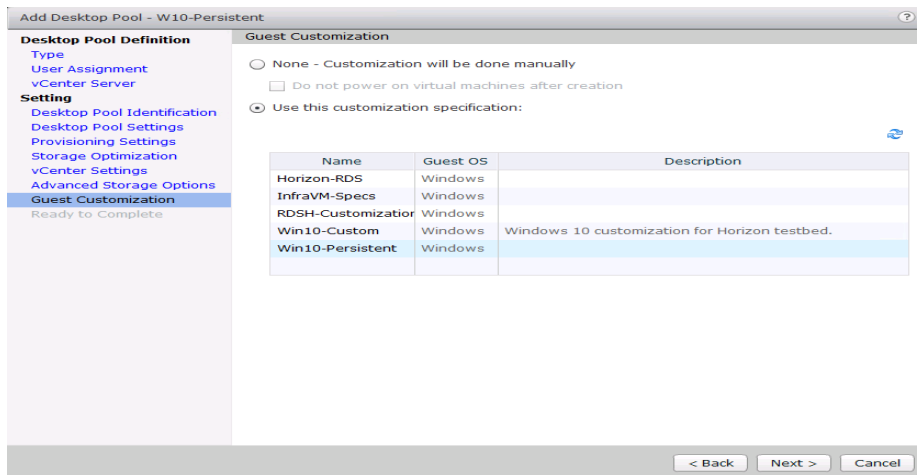
21. Select Advance Storage Options and enable the View Storage Accelerator.

22. Click Next.



23. Select Guest optimization settings.

24. Click Next.

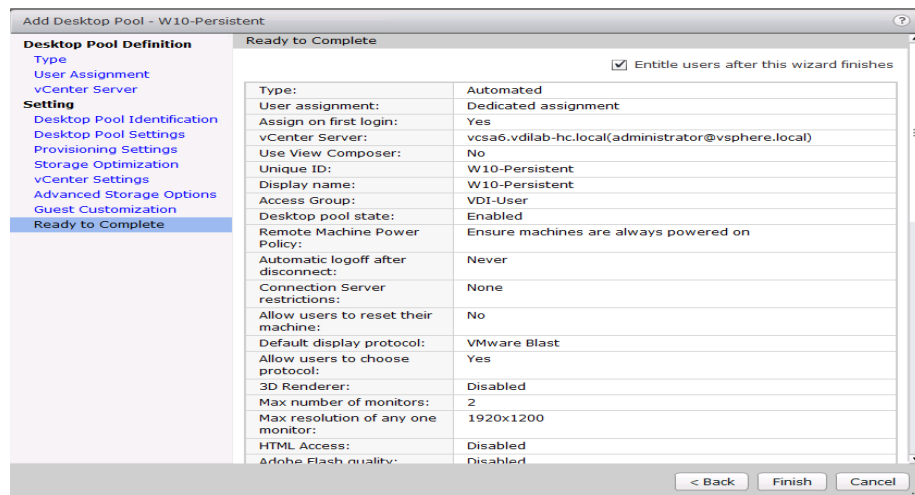




25. Review the summary of the pool you are creating.

26. Select the checkbox “Entitle users after pool creation wizard completion” to authorize users for the pool.

27. Click Finish.

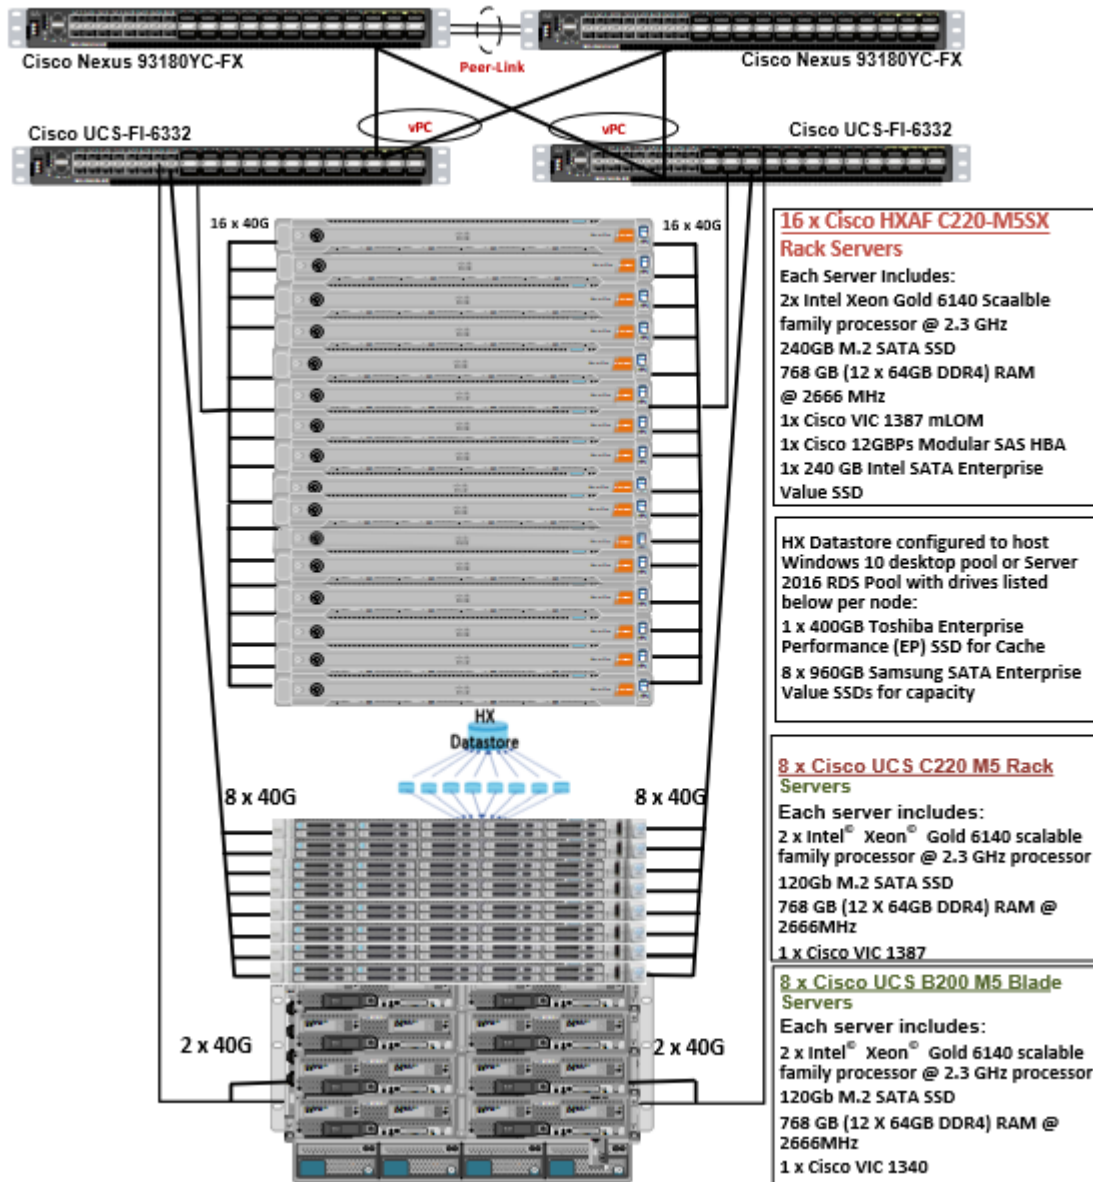


28. Follow the instructions provided in the Create Horizon 7 RDS Desktop Pool to authorize users for the Linked Clone Pool.

## Test Setup and Configurations

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

### Cisco HyperFlex and VMware Horizon 7, Full Scale Single UCS Domain Reference Architecture



Hardware Components:

- 2 x Cisco UCS 6332 UP Fabric Interconnects
- 2 x Cisco Nexus 93180YC-FX Access Switches

- 16 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)
- 8 x Cisco UCS C220 M5 Rack Servers (2 Intel Xeon processor 6140 CPUs at 2.3 GHz, with 768 GB of memory per server [32 GB x 24 DIMMs at 2666 MHz]).
- Cisco VIC 1387 mLOM
- 12G modular SAS HBA Controller
- 240GB M.2 SATA SSD drive (Boot and HyperFlex Data Platform controller VM)
- 8 x Cisco UCS B200 M5 Blade Servers (2 Intel Xeon processor 6140 CPUs at 2.3 GHz, with 768 GB of memory per server [32 GB x 24 DIMMs at 2666 MHz]).
- Cisco VIC 1340 mLOM
- 2 x 64GB SD card

### Software components:

- Cisco UCS firmware 4.0(1b)
- Cisco HyperFlex data platform 3.5.1a
- VMware vSphere 6.5
- VMware Horizon 7 Hosted Virtual Desktops and Hosted Shared Desktops
- VMware Horizon View Composer Server
- v-File Server for User Profiles
- Microsoft SQL Server 2016
- Microsoft Windows 10
- Microsoft Windows 2016
- Microsoft Office 2016
- Login VSI 4.1.32.1

## Testing Methodology and Success Criteria

All validation testing has been conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the RDSH Servers Session under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles have been conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

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

## Testing Procedure

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

### Pre-Test Setup for Testing

All virtual machines and RDSH Servers have been shut down utilizing the VMware Horizon 7 Administrator Console.

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

### Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, we require all sessions started, whether 60 single server users or 4400 full-scale test users to become active within two minutes after the last session is launched.

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

1. Time 0:00:00 Start PerfMon Logging on the following systems:
  - Infrastructure and VDI Host Blades used in test run
  - All Infrastructure 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 VMware Horizon 7 Administrator Console.
4. Time 0:06 First machines boot.
5. Time 0:35 Single Server or Scale target number of RDS Servers registered on Horizon.




---

**No more than 60 Minutes of rest time is allowed after the last desktop is registered and available on VMware Horizon 7 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 4.1.32.1 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 VMware Horizon 7 Administrator Dashboard must show that all desktops have been returned to the registered/available state as evidence of this condition being met.
12. Time 2:57 All logging terminated; Test complete.
13. Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shutdown all Windows 7 machines.
14. Time 3:30 Reboot all hypervisors.
15. Time 3:45 Ready for new test sequence.

## Success Criteria

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

The VMware Horizon Connection Server Dashboard will be monitored throughout the steady state to make sure of the following:

- All running sessions report In Use throughout the steady state
- No sessions move to unregistered, unavailable or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and the Login VSI Agent must have shut down. Cisco’s tolerance for Stuck Sessions is 0.5 percent (half of one percent.) If the Stuck Session count exceeds that value, we identify it as a test failure condition.

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

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing.

The purpose of this testing is to provide the data needed to validate VMware Horizon 7 Hosted Shared Desktop with VMware Horizon 7 Composer provisioning using Microsoft Windows Server 2016 sessions on Cisco UCS HXAF220C-M5SX.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here, and do not represent the full characterization of VMware products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish system performance and linear scalability.

### VSImax 4.1.x Description

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

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for SBC or VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well known desktop applications like Microsoft Office, Internet Explorer and Adobe PDF reader. By gradually increasing the numbers of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

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

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

### Server-Side Response Time Measurements

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

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

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solutions, for a benchmark like Login VSI.

### Calculating VSImax v4.1.x

The simulated desktop workload is scripted in a 48-minute loop when a simulated Login VSI user is logged on, performing generic Office worker activities. After the loop is finished it will restart automatically. Within each loop the response times of sixteen specific operations are measured in a regular interval: sixteen times in within each loop. The response times of these five operations are used to determine VSImax.

The five operations from which the response times are measured are:

- Notepad File Open (NFO)

Loading and initiating VSINotepad.exe and opening the openfile dialog. This operation is handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Notepad Start Load (NSLD)

Loading and initiating VSINotepad.exe and opening a file. This operation is also handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-user's point of view.

- Zip High Compression (ZHC)

This action copy's a random file and compresses it (with 7zip) with high compression enabled. The compression will very briefly spike CPU and disk IO.

- Zip Low Compression (ZLC)

This action copy's a random file and compresses it (with 7zip) with low compression enabled. The compression will very briefly disk IO and creates some load on the CPU.

- CPU

Calculates a large array of random data and spikes the CPU for a short period of time.

These measured operations within Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations become consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. If such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

Figure 45 Sample of a VSI Max Response Time Graph, Representing a Normal Test





Figure 46 Sample of a VSI Test Response Time Graph with a Clear Performance Issue



When the test is finished, VSI<sub>max</sub> can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSI<sub>max</sub> 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 VSI<sub>max</sub> models, this weighting much better represent system performance. All actions have very similar weight in the VSI<sub>max</sub> total. The following weighting of the response times are applied.

The following actions are part of the VSI<sub>max</sub> 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  $5 + 24 (=40 \text{ percent of } 60) = 31$  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. VSI<sub>max</sub> indicates what the total user capacity is for the system. These two are not automatically connected and related:

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

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

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

# Test Results

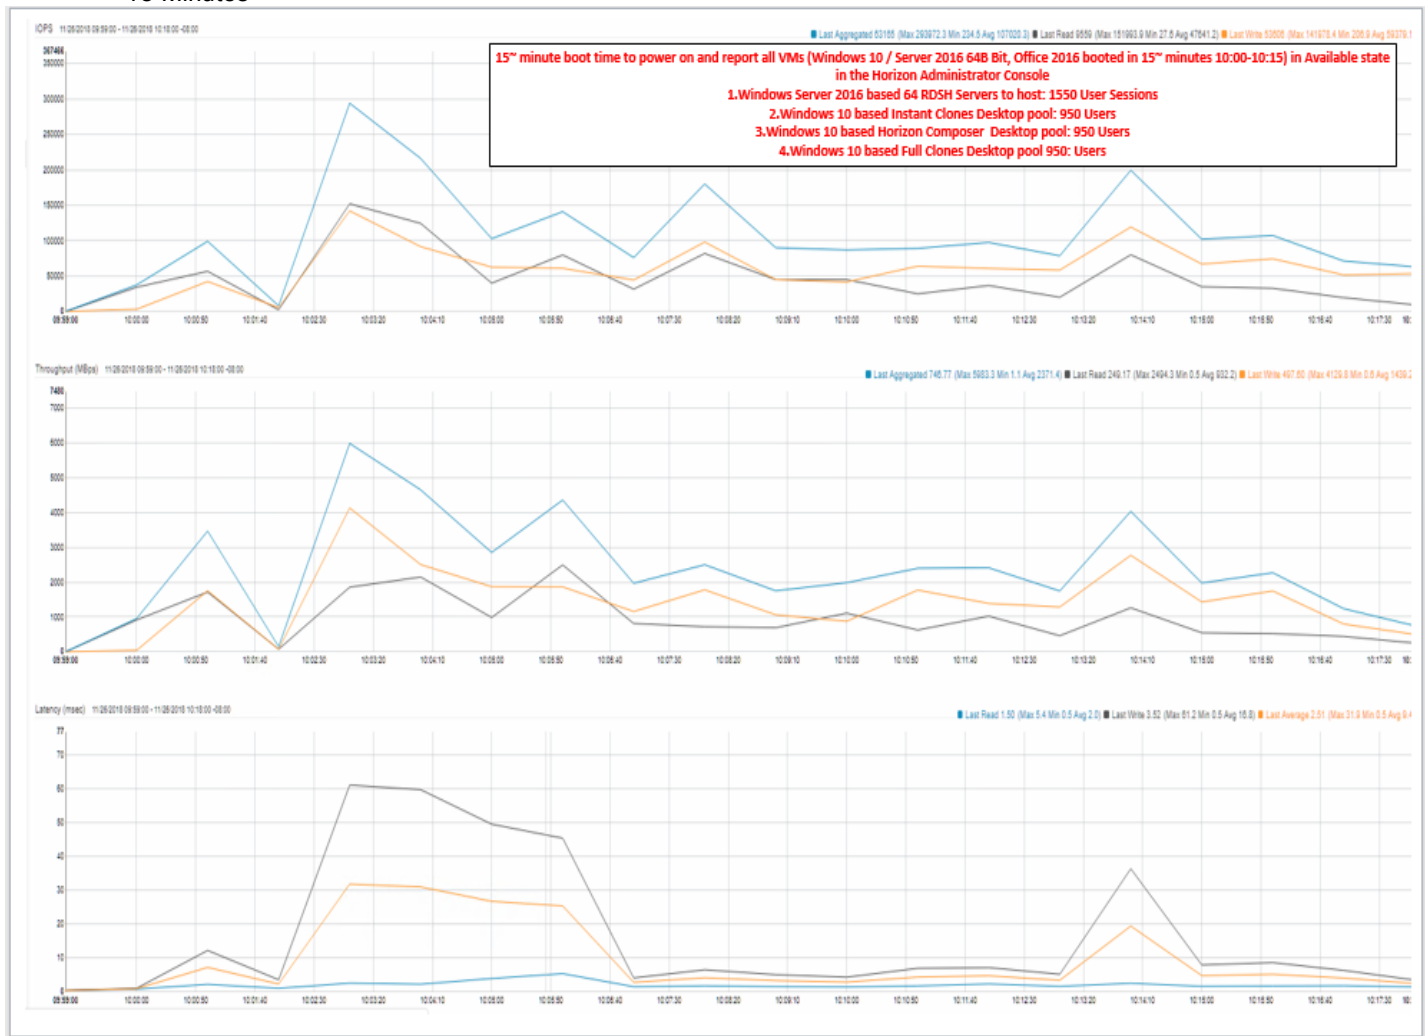
## Boot Storms

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 4400 desktops and measure the time it takes for the 4400<sup>th</sup> virtual machine to register as available in the Horizon Administrator console.

The Cisco HyperFlex HXAF220cM5SX, Cisco UCS C220M5 and B200 M5 cluster running Data Platform version 3.5(1a) software can accomplish this task in 15 minutes as shown in the following charts:

**Figure 47 4400 Horizon Server 2016 RDSH Sessions with Office 2016 Virtual Desktops 10 Instant-Clone, Linked-Clone, Persistent(Full Clone) Windows 10 Office 2016 Virtual Desktops Boot and Register as Available in Less Than 15 Minutes**



## Recommended Maximum Workload and Configuration Guidelines

### Sixteen Node Cisco HXAF220c-M5S Rack Server, Eight Node Cisco UCS C220 M5 and Eight Node Cisco UCS B200 M5 HyperFlex Cluster

For VMware Horizon 7 RDS Hosted Shared Desktop and Hosted Virtual Desktop use case, the recommended maximum workload was determined based on both Login VSI Knowledge Worker workload end user experience measures and HXAF220c-M5S and Cisco UCS C220 M5 and UCS B200 M5 server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the blade can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to insure that end-user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95 percent.



**Memory should never be oversubscribed for Desktop Virtualization workloads.**



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

Test Phase	Description
Boot	Start all RDS and/or VDI virtual machines at the same time.
Login	The Login VSI phase of test is where sessions are launched and start executing the workload over a 48 minutes duration.
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files.
Logoff	Sessions finish executing the Login VSI workload and logoff.



**The recommended maximum workload for a Cisco HyperFlex cluster configured on Cisco HXAF220c-M5S, Cisco US C220 M5 and Cisco UCS B200 M5 nodes with 6140 Gold processors and 768GB of RAM for Windows Server 2016 Hosted Sessions and persistent/non-persistent Hosted Virtual Desktop users is 4400 sessions with Office 2016 virtual desktops respectively.**

### 4400 User Full-Scale Testing on Thirty Two-Node Cisco HyperFlex Cluster

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

The full-scale testing with 4400 users comprised of 1550 RDS Hosted Server Sessions, 950 VDI Non-Persistent Instant clone, 950 VDI non-persistent Linked clone and 950 VDI persistent full clone virtual machines VMs.

Test result highlights include:

- 0.634 second baseline response time
- 1.091 second average response time with 4400 desktop sessions running
- Average CPU utilization of 85 percent during steady state
- Average of 356 GB of RAM used out of 768 GB available
- 1800 Mbps peak network utilization per host.
- Average Read Latency 0.15ms/Max Read Latency 0.79ms
- Average Write Latency 2.8ms/Max Write Latency 5.0ms
- 50000 peak I/O operations per second (IOPS) per cluster at steady state
- 1000Mbps peak throughput per cluster at steady state
- 89 percent Deduplication savings
- 44 percent Compression savings
- Total of 94 percent storage space savings

Figure 48 LoginVSI Analyzer Chart for 4400 Users Test




Figure 49 LoginVSI Analyzer Chart for Four Consecutive Test Running 4400 Knowledge Workload on 32 Node HyperFlex Cluster

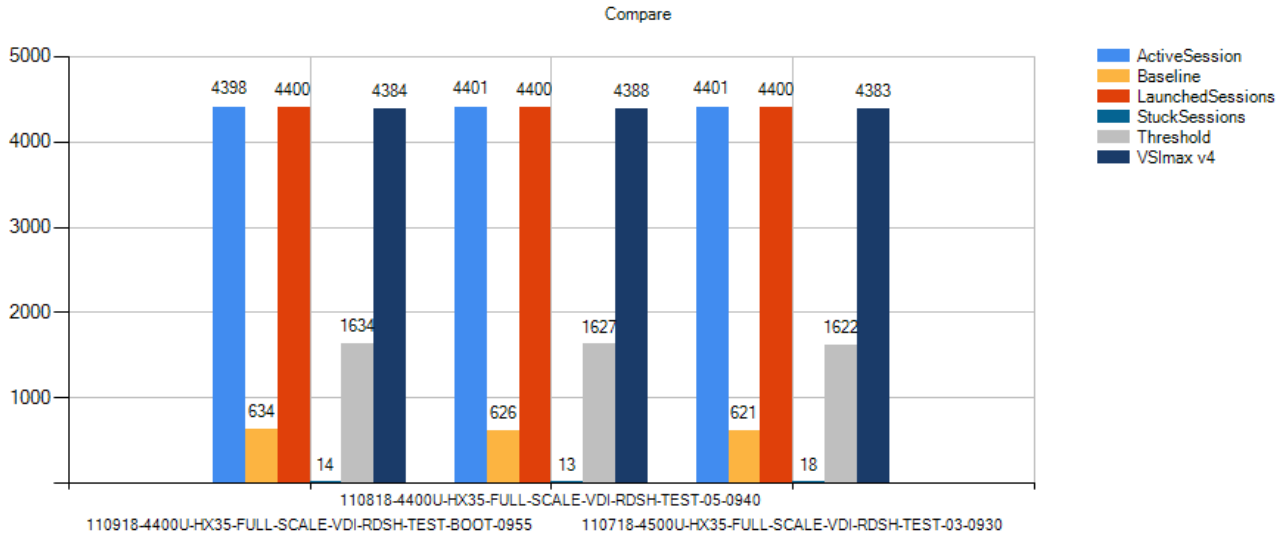


Figure 50 Sample ESXi host CPU Core Utilization Running 4400 User Test on 32 Nodes




Figure 51 Sample ESXi Host Memory Usage in Mbytes running 4400 User Test on 32 Node




Figure 52 Sample ESXi Host Network Adapter (VMNICs) Mbits Received/ Transmitted Per Sec Running 4400 User Test on 32 Node




Figure 53 HyperFlex Cluster WebUI Performance Chart for Knowledge Worker Workload Running 4400 User Test on 32 Node


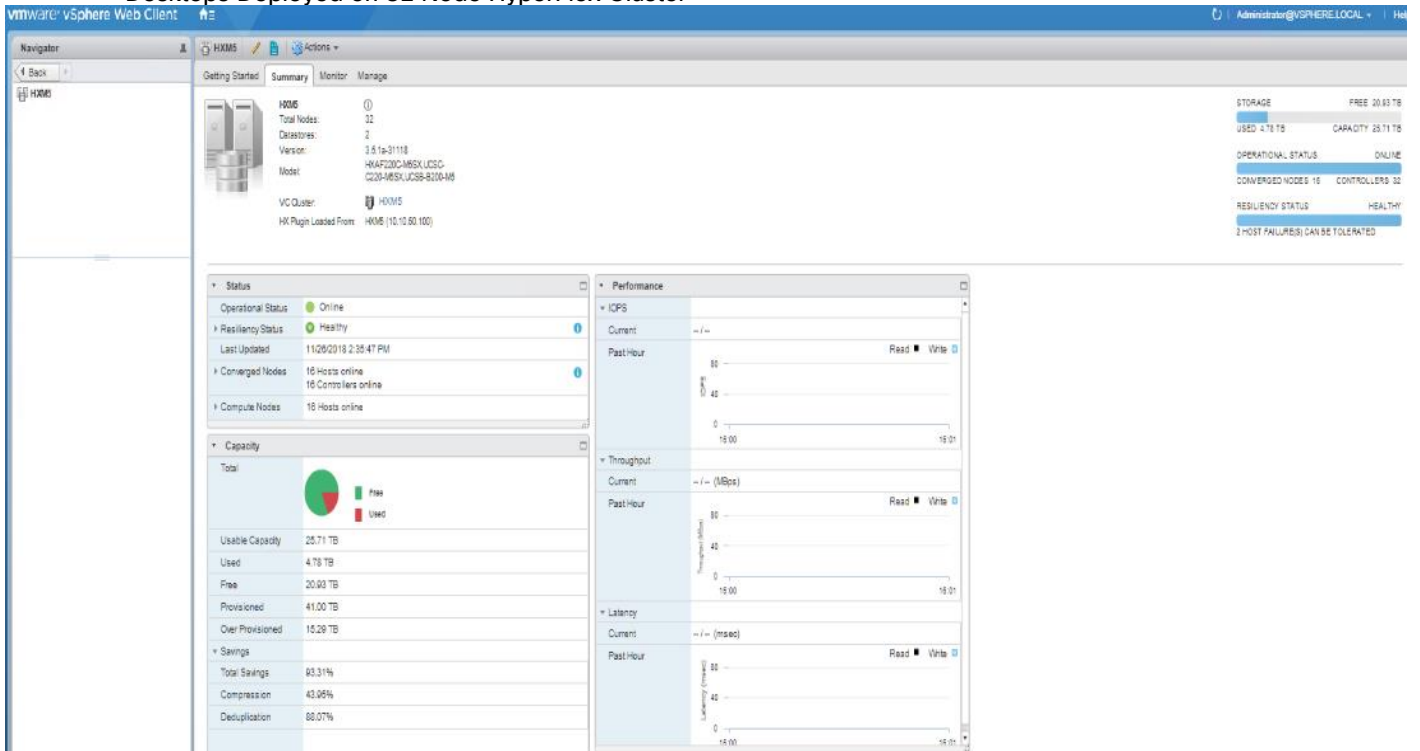


Figure 54 vCenter WebUI Reporting HyperFlex Cluster De-duplication and Compression Savings for 4400 User Sessions Supported on Windows Server 2016 Based Hosted Shared Sessions and Windows 10 Based Hosted Virtual Desktops Deployed on 32 Node HyperFlex Cluster





## 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 hyperconverged system.

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

Delivering responsive, resilient, high performance VMware Horizon 7 provisioned Microsoft Windows 10 Virtual Machines and Microsoft Windows Server 2016 for hosted Apps or desktops has many advantages for desktop virtualization administrators.

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

The solution 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 or up to six NVIDIA P4 cards. The Cisco UCS B200 M5 server supports up to two NVIDIA P6 cards for high density, high performance graphics workload support. See the [Cisco Graphics White Paper](#) for our fifth generation servers with NVIDIA GPUs and software for details on how to integrate this capability with VMware Horizon.

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 hyperconverged platform.

## About the Authors

---

Ramesh Guduru

Technical Marketing Engineer, Virtual Client Computing and Graphics Solutions, Cisco Systems, Inc.

Ramesh is a subject matter expert on Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, VMware vSphere and VMware Horizon end user computing. Ramesh is a member of the Cisco's Computing Systems Business Unit's solution team.

## Acknowledgements

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

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

## Appendix A - Cisco Nexus 93180 Switch Configuration

---

### Switch A Configuration

```
!Command: show running-config.
!Time: Fri Nov 13 17:17:40 2018
version 7.0(3)I7(2)
switchname 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$Bo0IrVnESUVxLcbRHg86j1 role network-admin
ip domain-lookup
no service unsupported-transceiver
class-map type qos match-all class-fcoe
policy-map type qos jumbo
class class-default
set qos-group 0
copp profile strict
snmp-server user admin network-admin auth md5 0x71d6a9cf1ea007cd3166e91a6f3807e5
priv 0x71d6a9cf1ea007cd3166e91a6f3807e5 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp server 10.10.50.2
ntp peer 10.10.50.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8
vlan 1,50-56
vlan 50
name InBand-Mgmt-C1
vlan 51
name Infra-Mgmt-C1
vlan 52
name StorageIP-C1
vlan 53
name vMotion-C1
vlan 54
name VM-Data-C1
vlan 55
```

```
name Launcher-C1
service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
ip route 0.0.0.0/0 10.29.132.1
vpc domain 50
role priority 1000
peer-keepalive destination 10.29.132.5 source 10.29.132.4
interface Vlan1
no shutdown
ip address 10.29.132.2/24
interface Vlan50
no shutdown
ip address 10.10.50.2/24
hsrp version 2
hsrp 50
preempt
priority 110
ip 10.10.50.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface Vlan51
no shutdown
ip address 10.10.51.2/24
hsrp version 2
hsrp 51
preempt
priority 110
ip 10.10.51.1
interface Vlan52
```

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

```
interface port-channel11
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 11
interface port-channel12
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 12
interface port-channel13
description FI-Uplink-K13
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 13
interface port-channel14
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 14
```

```
interface port-channel49
description FI-Uplink-K23
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
vpc 49

interface port-channel50
description FI-Uplink-K23
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
vpc 50

interface Ethernet1/1
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/2
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/3
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
interface Ethernet1/4
switchport mode trunk
switchport trunk allowed vlan 1,50-55
channel-group 10 mode active
```



```
interface Ethernet1/5
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 11 mode active
interface Ethernet1/6
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 11 mode active
interface Ethernet1/7
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 12 mode active
interface Ethernet1/8
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 12 mode active
interface Ethernet1/9
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 13 mode active
interface Ethernet1/10
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 13 mode active
interface Ethernet1/11
interface Ethernet1/12
```

```
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
description HX-Infra01
switchport mode trunk
switchport trunk allowed vlan 1,50-55
interface Ethernet1/16
description HX-Infra02
switchport mode trunk
switchport trunk allowed vlan 1,50-55
interface Ethernet1/17
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
switchport mode trunk
switchport trunk allowed vlan 1,50-54
spanning-tree port type edge trunk
interface Ethernet1/28
switchport mode trunk
switchport trunk allowed vlan 1,50-54
spanning-tree port type edge trunk
interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/31
interface Ethernet1/32
```

```
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
interface Ethernet1/47
interface Ethernet1/48
interface Ethernet1/49
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 49 mode active
interface Ethernet1/50
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 50 mode active
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
vrf member management
```

```

ip address 10.29.132.4/24
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin

```

## Switch B Configuration

```

!Command: show running-config
!Time: Fri Nov 13 17:18:36 2018
version 7.0(3)I7(2)
switchname 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 0x9046c100celf4ecdd74ef2f92c4e83f9
priv 0x9046c100celf4ecdd74ef2f92c4e83f9 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
ntp peer 10.10.50.2
ntp server 10.10.50.3
ntp server 171.68.38.66 use-vrf management
ntp logging
ntp master 8
vlan 1,50-54
vlan 50
name InBand-Mgmt-C1
vlan 51
name Infra-Mgmt-C1
vlan 52
name StorageIP-C1
vlan 53
```

```
name vMotion-C1
vlan 54
name VM-Data-C1
service dhcp
ip dhcp relay
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
ip route 0.0.0.0/0 10.29.132.1
vpc domain 50
role priority 2000
peer-keepalive destination 10.29.132.4 source 10.29.132.5
interface Vlan1
no shutdown
ip address 10.29.132.3/24
interface Vlan50
no shutdown
ip address 10.10.50.3/24
hsrp version 2
hsrp 50
preempt
priority 110
ip 10.10.50.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface Vlan51
no shutdown
ip address 10.10.51.3/24
hsrp version 2
hsrp 51
preempt
priority 110
```

```
ip 10.10.51.1
interface Vlan52
no shutdown
ip address 10.10.52.3/24
hsrp version 2
hsrp 52
preempt
priority 110
ip 10.10.52.1
interface Vlan53
no shutdown
ip address 10.10.53.3/24
hsrp version 2
hsrp 53
preempt
priority 110
ip 10.10.53.1
interface Vlan54
no shutdown
ip address 10.54.0.3/20
hsrp version 2
hsrp 54
preempt
priority 110
ip 10.54.0.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface Vlan55
no shutdown
ip address 10.10.55.3/24
hsrp version 2
hsrp 55
```

```
preempt
priority 110
ip 10.55.0.1
ip dhcp relay address 10.10.51.21
ip dhcp relay address 10.10.51.22
interface port-channel10
description vPC-PeerLink
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type network
service-policy type qos input jumbo
vpc peer-link
interface port-channel11
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 11
interface port-channel12
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
mtu 9216
service-policy type qos input jumbo
vpc 12
interface port-channel13
description FI-Uplink-K22
switchport mode trunk
switchport trunk allowed vlan 1,50-55
```



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

```
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 11 mode active
interface Ethernet1/7
switchport mode trunk
switchport trunk allowed vlan 1,50-55
mtu 9216
channel-group 12 mode active
interface Ethernet1/8
switchport mode trunk
switchport trunk allowed vlan 1,50-54
mtu 9216
channel-group 12 mode active
interface Ethernet1/9
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
interface Ethernet1/14
interface Ethernet1/15
description HX-Infra01
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/16
description HX-Infra02
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/17
interface Ethernet1/18
```

```
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
interface Ethernet1/24
interface Ethernet1/25
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/26
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/27
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/28
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/29
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/30
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/31
switchport mode trunk
```

```
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/32
switchport mode trunk
switchport trunk allowed vlan 1,50-55
spanning-tree port type edge trunk
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
interface Ethernet1/47
interface Ethernet1/48
switchport access vlan 50
interface Ethernet1/49
interface Ethernet1/50
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
interface Ethernet1/54
interface mgmt0
vrf member management
ip address 10.29.132.5/24
```


```
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I7.2.bin
no system default switchport shutdown
```

## Appendix B – Cisco HyperFlex HXAF220c-M5, Cisco UCS C220 M5 and Cisco UCS B200 M5 32-Node Hyperflex Horizon 7 Cluster Deployed 4400 Scale Test: In-Flight Performance Metrics

The following charts delineate performance parameters for the 32-node cluster during a Login VSI 4.1.32.1 Knowledge Worker workload test on 4400 Horizon 7 deployed user benchmark test.

The performance charts indicates that the HyperFlex All-Flash nodes and compute-only nodes in Hybrid configuration running Data Platform version v3.5.(1a) were operating consistently from node to node and well within normal operating parameters for hardware in this class. The data also supports the even distribution of the workload across all 32 servers.

**Figure 55 HXAF-VDI01: Memory Usage in Mbytes**



**Figure 56 HXAF-VDI01: Host CPU Core Utilization**




Figure 57 HXAF-VDI01: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 58 HXAF-VDI02: Memory Usage in Mbytes

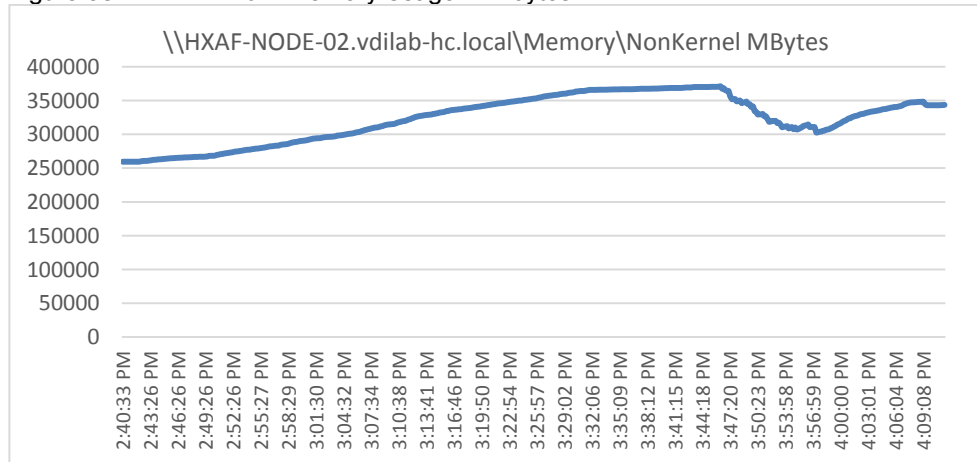


Figure 59 HXAF-VDI02: Host CPU Core Utilization




Figure 60 HXAF-VDI02: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 61 HXAF-VDI03: Memory Usage in Mbytes




Figure 62 HXAF-VDI03: Host CPU Core Utilization






Figure 63 HXAF-VDI03: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 64 HXAF-VDI04: Memory Usage in Mbytes




Figure 65 HXAF-VDI04: Host CPU Core Utilization




Figure 66 HXAF-VDI04: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 67 HXAF-VDI05: Memory Usage in Mbytes




Figure 68 HXAF-VDI05: Host CPU Core Utilization




Figure 69 HXAF-VDI05: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 70 HXAF-VDI06: Memory Usage in Mbytes




Figure 71 HXAF-VDI06: Host CPU Core Utilization




Figure 72 HXAF-VDI06: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 73 HXAF-VDI07: Memory Usage in Mbytes




Figure 74 HXAF-VDI07: Host CPU Core Utilization




Figure 75 HXAF-VDI07: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 76 HXAF-VDI08: Memory Usage in Mbytes




Figure 77 HXAF-VDI08: Host CPU Core Utilization




Figure 78 HXAF-VDI08: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 79 HXAF-VDI09: Memory Usage in Mbytes




Figure 80 HXAF-VDI09: Host CPU Core Utilization




Figure 81 HXAF-VDI09: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 82 HXAF-VDI10: Memory Usage in Mbytes




Figure 83 HXAF-VDI10: Host CPU Core Utilization






Figure 84 HXAF-VDI10: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 85 HXAF-VDI11: Memory Usage in Mbytes




Figure 86 HXAF-VDI11: Host CPU Core Utilization




Figure 87 HXAF-VDI11: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 88 HXAF-VDI12: Memory Usage in Mbytes




Figure 89 HXAF-VDI12: Host CPU Core Utilization




Figure 90 HXAF-VDI12: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 91 HXAF-VDI13: Memory Usage in Mbytes




Figure 92 HXAF-VDI13: Host CPU Core Utilization




Figure 93 HXAF-VDI13: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 94 HXAF-VDI14: Memory Usage in Mbytes




Figure 95 HXAF-VDI14: Host CPU Core Utilization




Figure 96 HXAF-VDI14: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 97 HXAF-VDI15: Memory Usage in Mbytes




Figure 98 HXAF-VDI15: Host CPU Core Utilization




Figure 99 HXAF-VDI15: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 100 HXAF-VDI16: Memory Usage in Mbytes




Figure 101 HXAF-VDI16: Host CPU Core Utilization




Figure 102 HXAF-VDI16: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 103 Cisco UCS C220 M5-VDI09: Memory Usage in Mbytes






Figure 104 Cisco UCS C220 M5-VDI09: Host CPU Core Utilization




Figure 105 Cisco UCS C220 M5-VDI09: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 106 Cisco UCS C220 M5-VDI10: Memory Usage in Mbytes




Figure 107 Cisco UCS C220 M5-VDI10: Host CPU Core Utilization




Figure 108 Cisco UCS C220 M5-VDI10: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 109 Cisco UCS C220 M5-VDI11: Memory Usage in Mbytes




Figure 110 Cisco UCS C220 M5-VDI11: Host CPU Core Utilization




Figure 111 Cisco UCS C220 M5-VDI11: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 112 Cisco UCS C220 M5-VDI12: Memory Usage in Mbytes




Figure 113 Cisco UCS C220 M5-VDI12: Host CPU Core Utilization




Figure 114 Cisco UCS C220 M5-VDI12: Network Adapter (vMNICs) Mbits Received /Transmitted Per Sec




Figure 115 Cisco UCS C220 M5-VDI13: Memory Usage in Mbytes




Figure 116 Cisco UCS C220 M5-VDI13: Host CPU Core Utilization




Figure 117 Cisco UCS C220 M5-VDI13: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 118 Cisco UCS C220 M5-VDI14: Memory Usage in Mbytes




Figure 119 Cisco UCS C220 M5-VDI14: Host CPU Core Utilization




Figure 120 Cisco UCS C220 M5-VDI14: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 121 Cisco UCS C220 M5-VDI15: Memory Usage in Mbytes




Figure 122 Cisco UCS C220 M5-VDI15: Host CPU Core Utilization




Figure 123 Cisco UCS C220 M5-VDI15: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 124 Cisco UCS C220 M5-VDI16: Memory Usage in Mbytes






Figure 125 Cisco UCS C220 M5-VDI16: Host CPU Core Utilization




Figure 126 Cisco UCS C220 M5-VDI16: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 127 Cisco UCS B200 M5-VDI01: Memory Usage in Mbytes




Figure 128 Cisco UCS B200 M5-VDI01: Host CPU Core Utilization




Figure 129 Cisco UCS B200 M5-VDI01: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 130 Cisco UCS B200 M5-VDI02: Memory Usage in Mbytes




Figure 131 Cisco UCS B200 M5-VDI02: Host CPU Core Utilization




Figure 132 Cisco UCS B200 M5-VDI02: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 133 Cisco UCS B200 M5-VDI03: Memory Usage in Mbytes




Figure 134 Cisco UCS B200 M5-VDI03: Host CPU Core Utilization




Figure 135 Cisco UCS B200M5-VDI03: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 136 Cisco UCS B200 M5-VDI04: Memory Usage in Mbytes




Figure 137 Cisco UCS B200 M5-VDI04: Host CPU Core Utilization




Figure 138 Cisco UCS B200 M5-VDI04: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 139 Cisco UCS B200 M5-VDI05: Memory Usage in Mbytes




Figure 140 Cisco UCS B200 M5-VDI05: Host CPU Core Utilization




Figure 141 Cisco UCS B200 M5-VDI05: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 142 Cisco UCS B200 M5-VDI06: Memory Usage in Mbytes






Figure 143 Cisco UCS B200 M5-VDI06: Host CPU Core Utilization




Figure 144 Cisco UCS B200 M5-VDI06: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 145 Cisco UCS B200 M5-VDI07: Memory Usage in Mbytes




Figure 146 Cisco UCS B200 M5-VDI07: Host CPU Core Utilization




Figure 147 Cisco UCS B200 M5-VDI07: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec




Figure 148 Cisco UCS B200 M5-VDI08: Memory Usage in Mbytes




Figure 149 Cisco UCS B200 M5-VDI08: Host CPU Core Utilization




Figure 150 Cisco UCS B200 M5-VDI08: Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec

