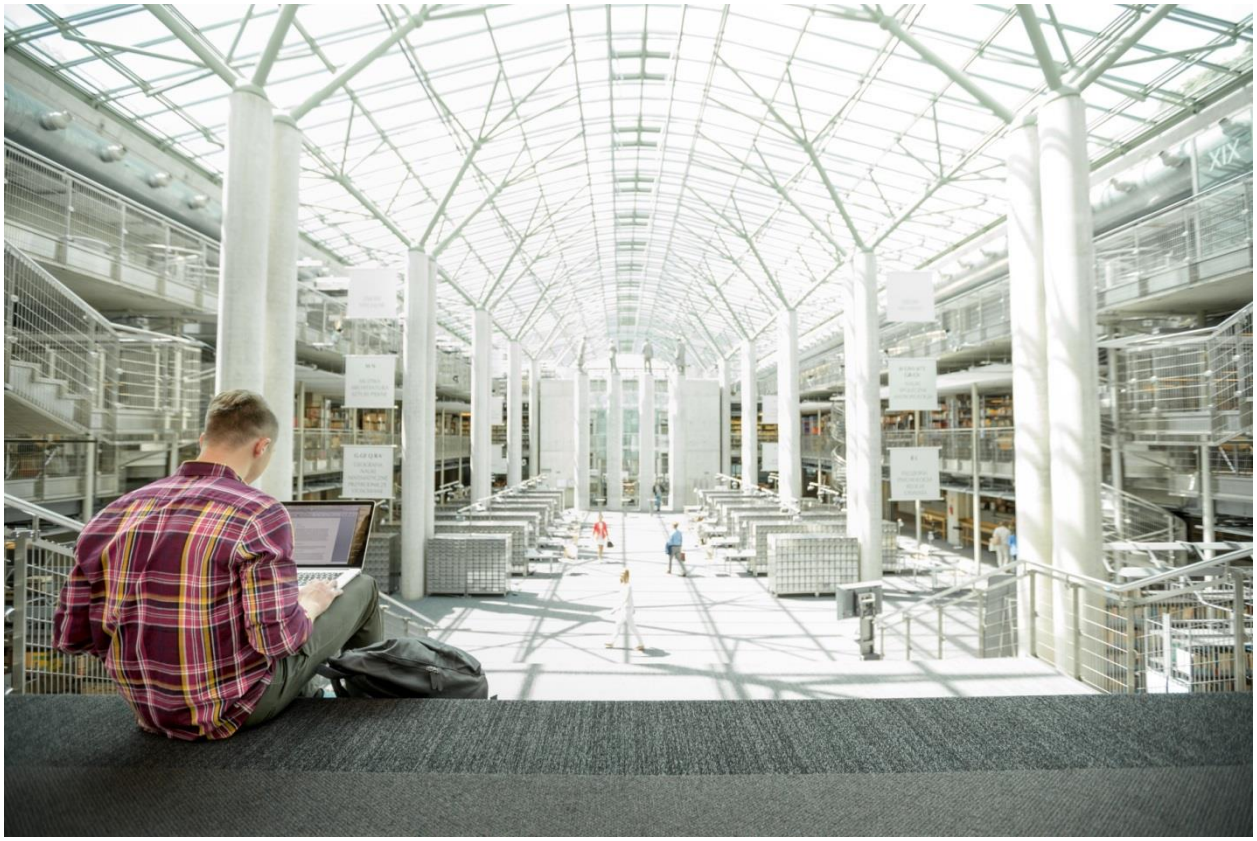

Cisco HyperFlex M5 All-Flash Hyperconverged System with up to 600 VMware Horizon 7 Users

Design and Deployment of Cisco HyperFlex for Virtual Desktop Infrastructure with VMware Horizon 7

Last Updated: January 8, 2018

CISCO VALIDATED DESIGN

About Cisco Validated Designs

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

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

Table of Contents

Executive Summary

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

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

This document provides an architectural reference and design guide for up to a 450 user mixed workload on a 4-node (4 Cisco HyperFlex HXAF220C-M5SX server) Cisco HyperFlex system. We provide deployment guidance and performance data for VMware Horizon 7.3 virtual desktops running Microsoft Windows 10 with Office 2016 Linked-Clone, Instant-Clone and Persistent virtual desktops as well as Windows Server 2016 RDS server-based sessions on VMware vSphere 6.5. The solution is a pre-integrated, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches and Cisco HyperFlex Data Platform software version 2.6.1a.

The solution payload is 100 percent virtualized on Cisco HyperFlex HXAF220C-M5SX hyperconverged nodes booting via on-board M.2 SATA SSD drive running VMware vSphere 6.5 U1 hypervisor and the Cisco HyperFlex Data Platform storage controller VM. The virtual desktops are configured with VMware Horizon 7.3.1, which incorporates both traditional persistent and non-persistent virtual Windows 7/8/10 desktops, hosted applications and remote desktop service (RDS) Microsoft Server 2008 R2, Server 2012 R2 or Server 2016 based desktops. The solution provides unparalleled scale and management simplicity. VMware Horizon Instant-Clone or Linked-Clone floating assignment Windows 10 desktops (450,) or full clone desktops (450) or RDSH server based desktops (600) can be provisioned on a four node Cisco HyperFlex cluster. Where applicable, this document provides best practice recommendations and sizing guidelines for customer deployment of this solution.

The solution boots 450 virtual desktops or 36 RDSH virtual server machines in 5 minutes or less, making sure that users will not experience delays in accessing their virtual workspace on HyperFlex.

Our past Cisco Validated Design studies with HyperFlex show linear scalability out to the cluster size limits of 16 HyperFlex hyperconverged nodes plus 16 Cisco UCS B200 M5, UCS C220 M5, or UCS C240 M5 compute only nodes. You can expect that our new HyperFlex all flash system running HX Data Platform 2.6 on Cisco HXAF220 M5 or Cisco HXAF240 M5 nodes will scale up to 4800 knowledge worker users per cluster with N+1 server fault tolerance.

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

The solution provides outstanding virtual desktop end user experience as measured by the Login VSI 4.1.25 Knowledge Worker workload running in benchmark mode. Index average end-user response times for all tested delivery methods is under 1 second, representing the best performance in the industry.

Solution Overview

Introduction

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

Audience

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

Purpose of this Document

This document provides a step-by-step design, configuration, and implementation guide for the Cisco Validated Design for a Cisco HyperFlex All-Flash system running four different VMware Horizon 7 workloads with Cisco UCS 6300 series Fabric Interconnects and Cisco Nexus 9300 series switches.

What's New?

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

- Validation of Cisco Nexus 9000 with Cisco HyperFlex Support for the Cisco UCS 3.2(2) release and Cisco HyperFlex Data Platform v 2.6.1a.
- VMware vSphere 6.5 U1 Hypervisor
- VMware Horizon 7.3.1 Instant Clones, Linked Clones, Persistent Desktops and RDSH shared server sessions

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 respective Cisco UCS Manager release version of Hardware and Software Interoperability for Cisco HyperFlex HX-Series.

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

These factors have led to the need predesigned computing, networking and storage building blocks optimized to lower the initial design cost, simply management, and enable 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 VMware Horizon Microsoft Windows 10 virtual desktops and VMware 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 an 8-rack unit footprint in an industry standard 42U rack. Port density on the Cisco Nexus switches and Cisco UCS Fabric Interconnects enables the networking components to accommodate multiple HyperFlex clusters in a single Cisco UCS domain.

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

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

The combination of technologies from Cisco Systems, Inc. and 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 450 user mixed hosted virtual desktop and 600 user hosted shared desktop environment booted and registered with the Horizon 7 in under 5 minutes, providing our customers with an extremely fast, reliable cold-start desktop virtualization system.
- Stress-tested to the limits during simulated login storms. All 450 users logged in and started running workloads up to steady state in 48-minutes without overwhelming the processors, exhausting memory or exhausting the storage subsystems, providing customers with a desktop virtualization system that can easily handle the most demanding login and startup storms.
- Ultra-condensed computing for the datacenter. The rack space required to support the initial 450 user system is 8 rack units, including Cisco Nexus Switching and Cisco Fabric interconnects.

Incremental VMware Horizon users can be added to the Cisco HyperFlex cluster up to the cluster scale limits, currently 16 hyper converged and 16 compute only nodes, by adding one or more nodes.

- 100 percent virtualized: This CVD presents a validated design that is 100 percent virtualized on 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 datacenter management: Cisco maintains industry leadership with the new Cisco UCS Manager 3.2(2) software that simplifies scaling, guarantees consistency, and eases maintenance. Cisco's ongoing development efforts with Cisco UCS Manager, Cisco UCS Central, and Cisco UCS Director insure that customer environments are consistent locally, across Cisco UCS Domains and across the globe. Cisco UCS software suite offers increasingly simplified operational and deployment management, and it continues to widen the span of control for customer organizations' subject matter experts in compute, storage and network.
- Cisco 40G Fabric: Our 40G unified fabric story gets additional validation on 6300 Series Fabric Interconnects as Cisco runs more challenging workload testing, while maintaining unsurpassed user response times.
- Cisco HyperFlex Connect (HX Connect): An all-new HTML 5 based Web UI Introduced with HyperFlex v2.5 is available for use as the primary management tool for Cisco HyperFlex. Through this centralized point of control for the cluster, administrators can create volumes, monitor the data platform health, and manage resource use. Administrators can also use this data to predict when the cluster will need to be scaled.
- Cisco HyperFlex storage performance: Cisco HyperFlex provides industry-leading hyper converged storage performance that efficiently handles the most demanding I/O bursts (for example, login storms), high write throughput at low latency, delivers simple and flexible business continuity and helps reduce storage cost per desktop.
- Cisco HyperFlex agility: Cisco HyperFlex System enables users to seamlessly add, upgrade or remove storage from the infrastructure to meet the needs of the virtual desktops.
- 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.

Cisco Desktop Virtualization Solutions: Data Center

The Evolving Workplace

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

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

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

Figure 1 Cisco Data Center Partner Collaboration

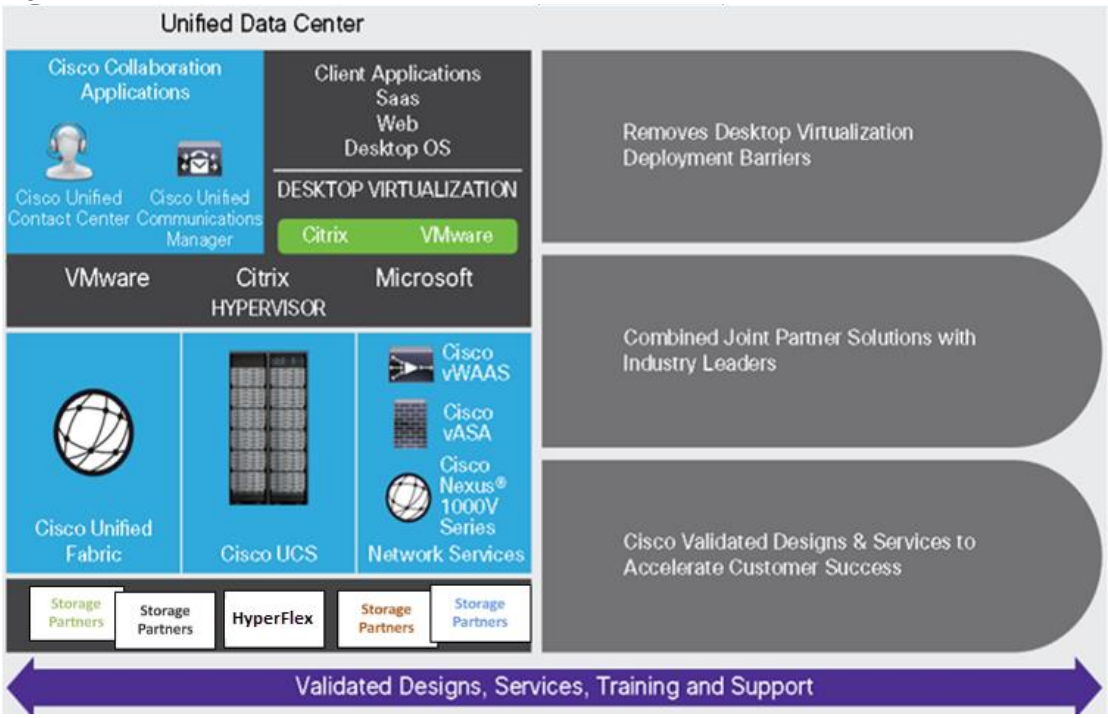

Some of the key drivers for desktop virtualization are increased data security, the ability to expand and contract capacity and reduced TCO through increased control and reduced management costs.

Cisco Desktop Virtualization Focus

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

Simplified

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

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

Simplification also leads to more successful desktop virtualization implementation. Cisco and its technology partners like 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 1.5 terabyte (TB) of memory with 2- and 4-socket servers). Using unified fabric technology as a building block, Cisco UCS server aggregate bandwidth can scale to up to 80 Gbps per server, and the northbound Cisco UCS fabric interconnect can output 2 terabits per second (Tbps) at line rate, helping prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency unified fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, Cisco HyperFlex helps maintain data availability and optimal performance during boot and login storms as part of the Cisco Desktop Virtualization Solutions. Recent Cisco Validated Designs based on VMware Horizon, Cisco HyperFlex solutions have demonstrated scalability and performance, with up to 450 hosted virtual desktops and hosted shared desktops up and running in 5 minutes.

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

Savings and Success

The simplified, secure, scalable Cisco data center infrastructure for desktop virtualization solutions saves time and money compared to alternative approaches. Cisco UCS enables faster payback and ongoing

savings (better ROI and lower TCO) and provides the industry's greatest virtual desktop density per server, reducing both capital expenditures (CapEx) and operating expenses (OpEx). The Cisco UCS architecture and Cisco Unified Fabric also enables much lower network infrastructure costs, with fewer cables per server and fewer ports required. In addition, storage tiering and deduplication technologies decrease storage costs, reducing desktop storage needs by up to 50 percent.

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

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

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

Use Cases

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

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

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

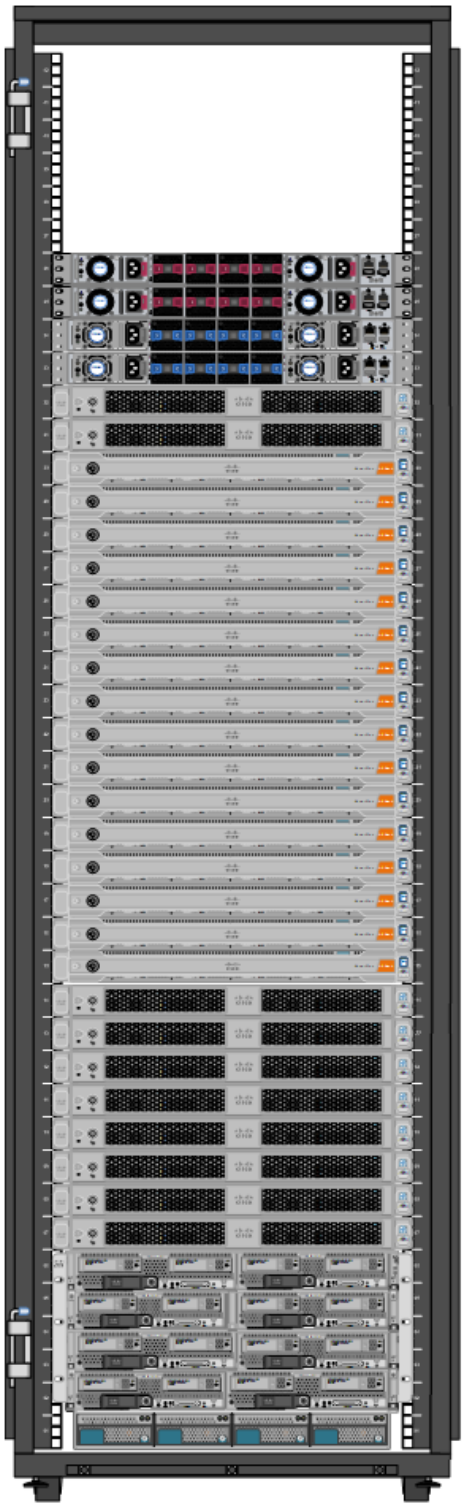

Physical Topology

The Cisco HyperFlex system is composed of a pair of Cisco UCS 6200/6300 series Fabric Interconnects, along with up to 16 HXAF-Series rack mount servers per cluster. In addition, up to 16 compute only servers can be added per cluster. Adding Cisco UCS 5108 Blade chassis allows use of Cisco UCS B200-M5 blade servers for additional compute resources in a hybrid cluster design. Cisco UCS C240 and C220 servers can also be used for additional compute resources. Up to 8 separate HX clusters can be installed under a single pair of Fabric Interconnects. The Fabric Interconnects both connect to every HX-Series rack mount server, and both connect to every Cisco UCS 5108 blade chassis. Upstream network connections, also referred to as “northbound” network connections are made from the Fabric Interconnects to the customer datacenter network at the time of installation.

For this study, we uplinked the Cisco 6332-16UP Fabric Interconnects to Cisco Nexus 9372PX switches.

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

Figure 3 Cisco HyperFlex Standard Topology

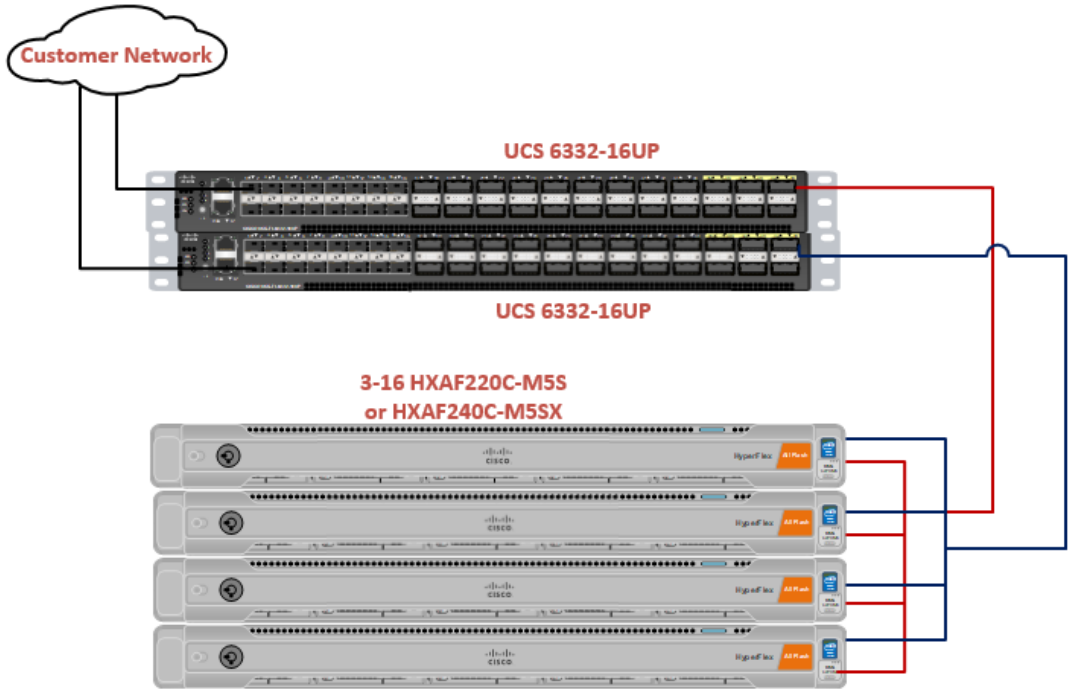

Figure 4 Cisco HyperFlex Hyperconverged plus Compute Only Node Topology

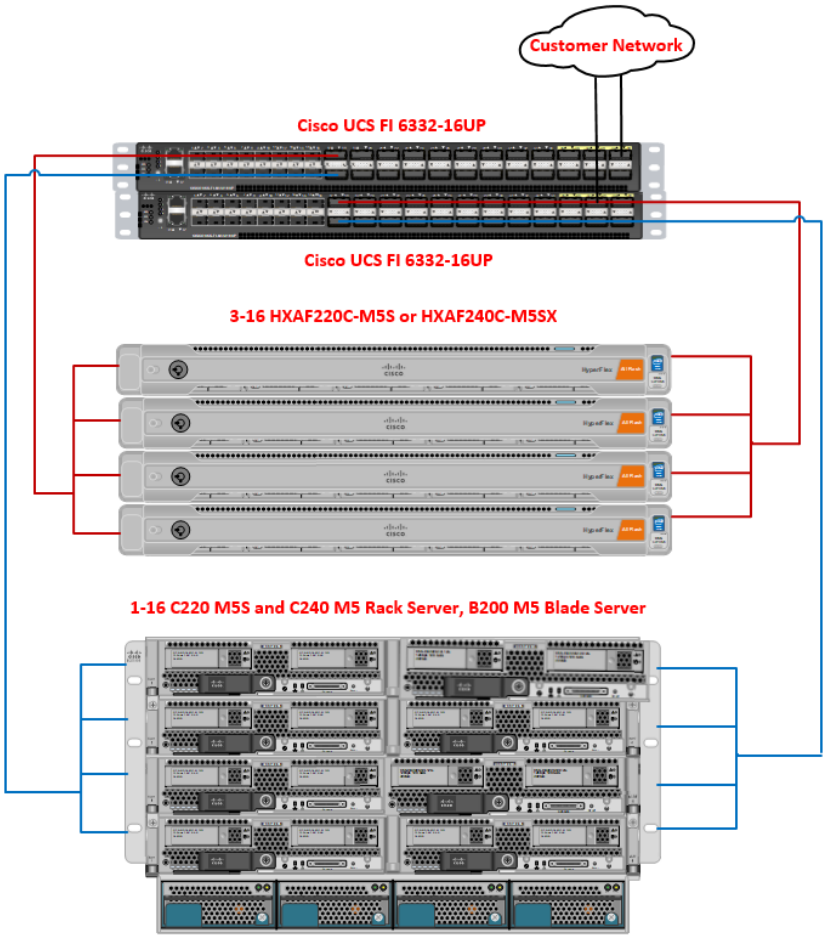

Fabric Interconnects

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

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

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

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

HX-Series Rack Mount Servers

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

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

Figure 5 HX-Series Server Connectivity

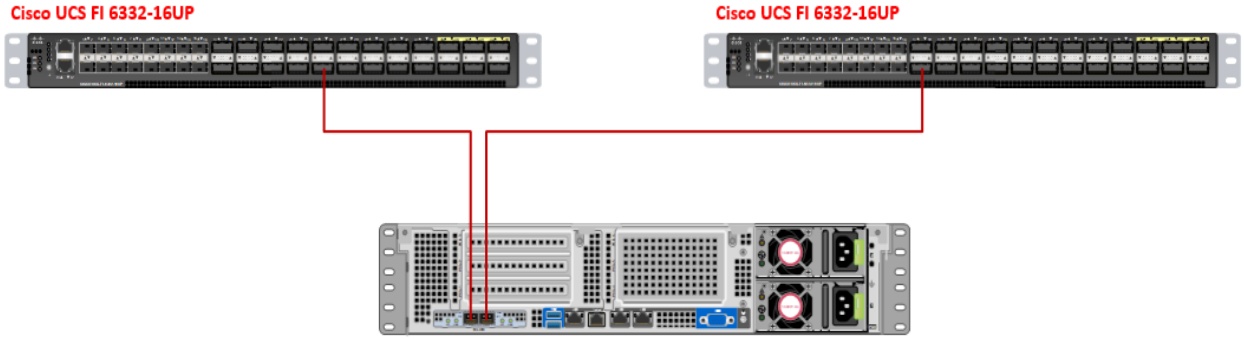

Cisco UCS B-Series Blade Servers

Hybrid HyperFlex clusters also incorporate 1-8 Cisco UCS B200 M5 blade servers for additional compute capacity. Like all other Cisco UCS B-series blade servers, the Cisco UCS B200 M5 must be installed within a Cisco UCS 5108 blade chassis. The blade chassis comes populated with 1-4 power supplies, and 8 modular cooling fans. In the rear of the chassis are two bays for installation of Cisco Fabric Extenders. The Fabric Extenders (also commonly called IO Modules, or IOMs) connect the chassis to the Fabric Interconnects. Internally, the Fabric Extenders connect to the Cisco VIC 1340 card installed in each blade server across the chassis backplane. The standard connection practice is to connect 1-4 10 GbE or 2 x 40 (native) GbE links from the left side IOM, or IOM 1, to FI A, and to connect the same number of 10 GbE links from the right side

IOM, or IOM 2, to FI B (Figure 6). All other cabling configurations are invalid, and can lead to errors, discovery failures, and loss of redundant connectivity.

Figure 6 Cisco UCS 5108 Chassis Connectivity

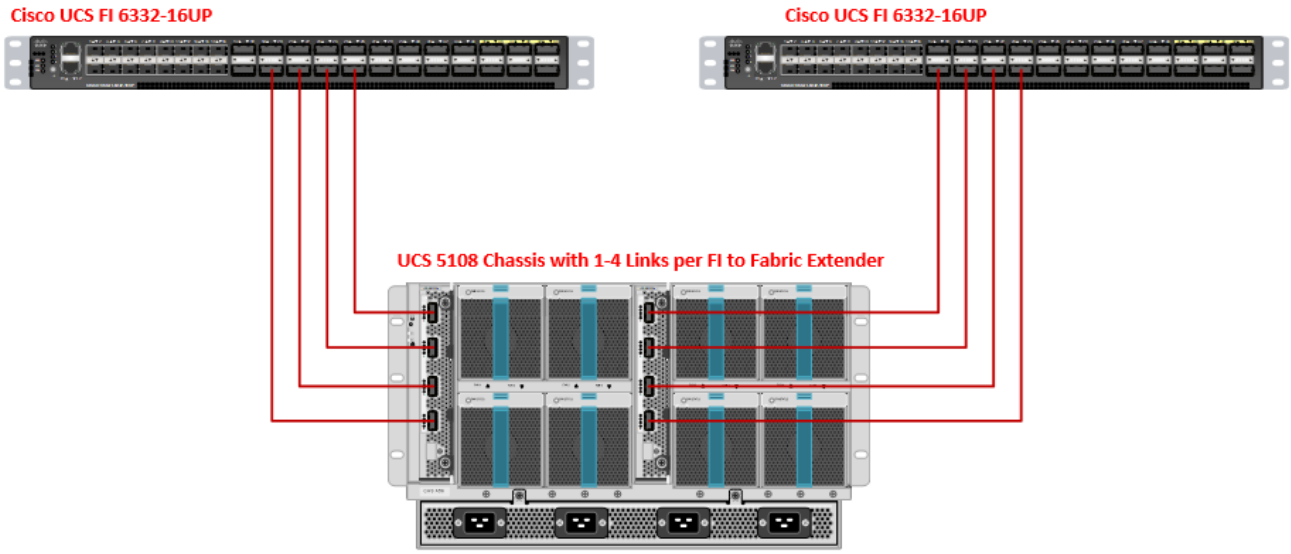

Logical Network Design

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

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

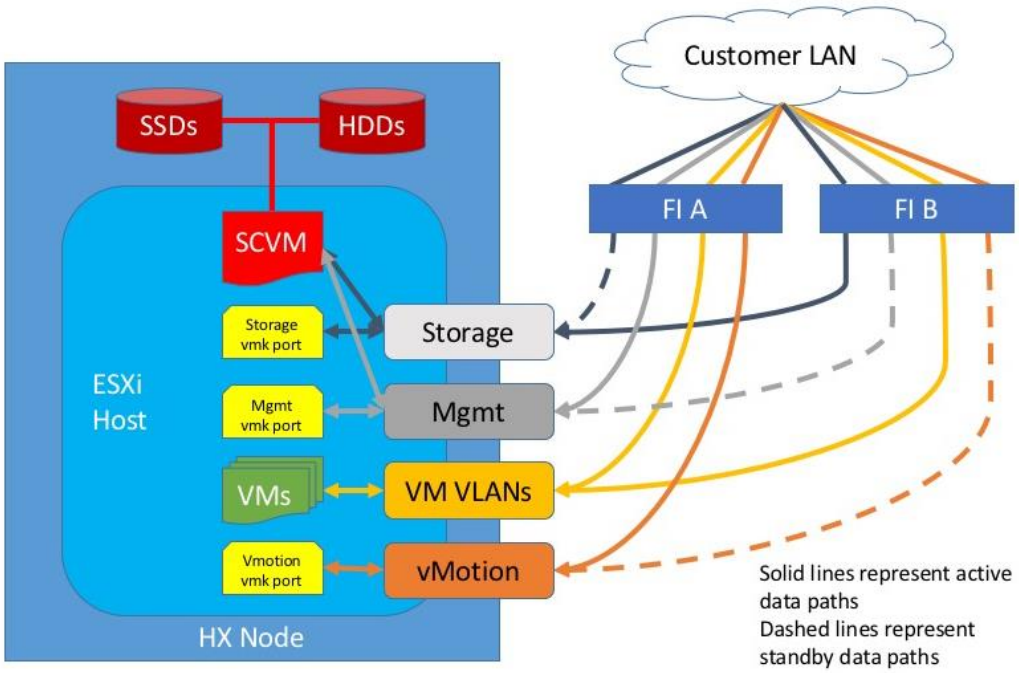

The reference hardware configuration includes:

- Two Cisco Nexus 9372PX switches

- Two Cisco UCS 6332-16UP fabric interconnects
- Four Cisco HX-series Rack server running HyperFlex data platform version 2.6.1a .

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 Four node configuration:

- 600 Horizon 7 RDSH server desktop sessions
- 450 Horizon 7 Windows 10 non-persistent virtual desktops
- 450 Horizon 7 Windows 10 persistent virtual desktops

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

Data provided here will allow customers to run 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 02. These procedures covers everything from physical cabling to network, compute and storage device configurations.

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a Cisco Validated Design for various type of Virtual Desktop workloads on Cisco HyperFlex. Configuration guidelines are provided that refer to which redundant component is being configured with each step. For example, Cisco Nexus A or Cisco Nexus B identifies a member in the pair of Cisco Nexus switches that are configured. Cisco UCS 6248UP 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 8 Cisco HyperFlex Family Overview

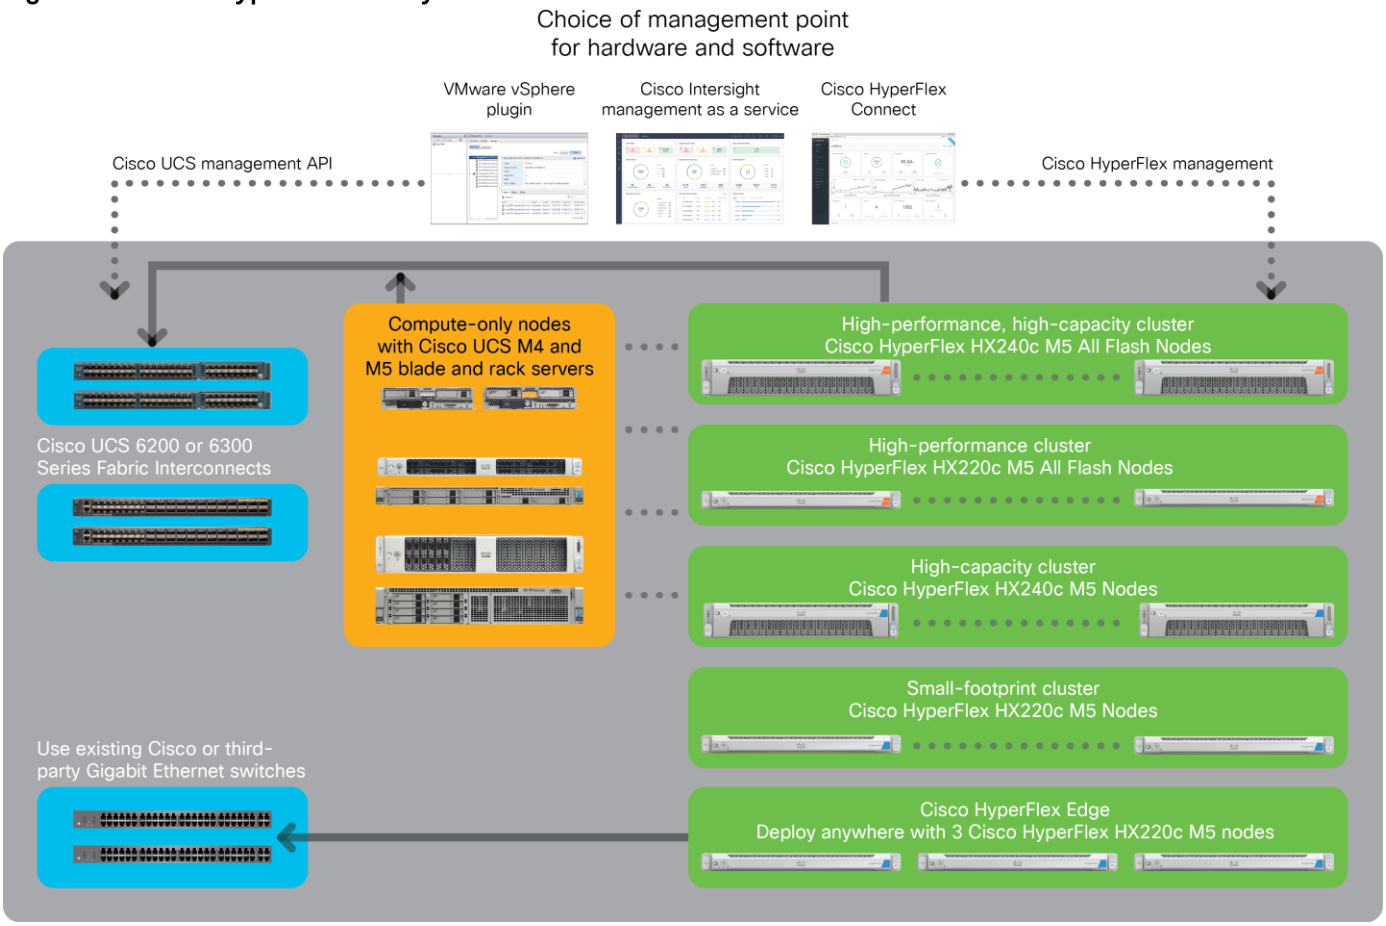

Cisco UCS and Cisco HyperFlex are designed to deliver:

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

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

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

Figure 9 Cisco HyperFlex System Overview

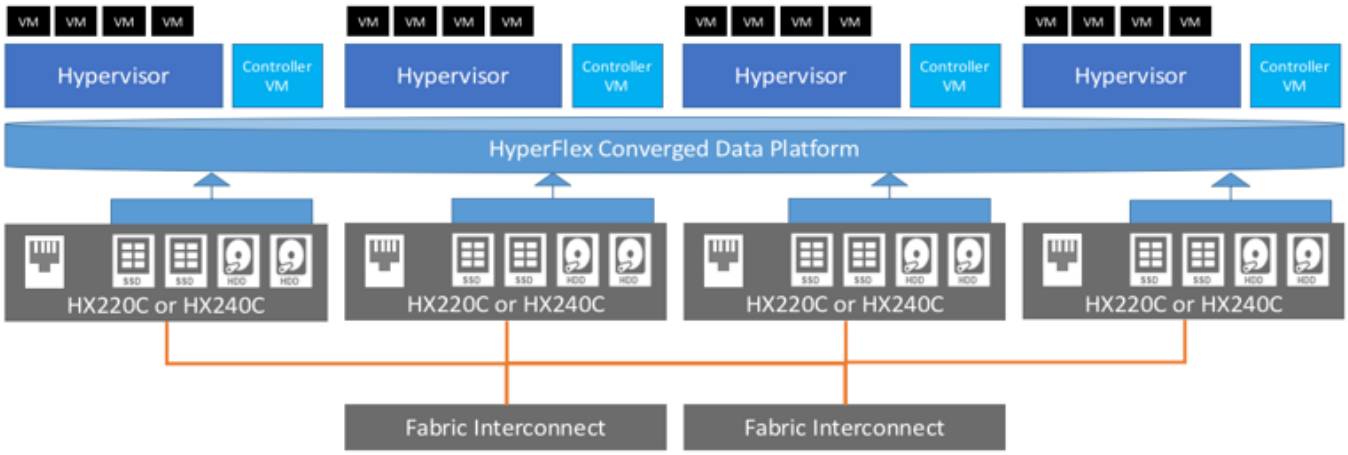

Enhancements for Version 2.6.1

The Cisco HyperFlex system has several new capabilities and enhancements in version 2.6.1:

- New All-Flash and Hybrid HX M5 server models are added to the Cisco HyperFlex product family
- Cisco HyperFlex now support the latest generation of Cisco UCS software, Cisco UCS Manager 3.2(2b) and beyond. For new All-Flash deployments on M5 servers, verify that Cisco UCS Manager 3.2(2b) or later is installed.
- Cisco Smart Licensing–Support for Cisco Smart Software Manager satellite. Please refer to the Cisco HyperFlex Getting Started Guide, Release 2.6, for more details.
- M5 Servers
- Key release highlights:
  - Same software feature set as HX 2.5.
  - 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 10 Cisco UCS 6332 Fabric Interconnect

Figure 11 Cisco UCS 6332-16UP Fabric Interconnect

Cisco HyperFlex HX-Series Nodes

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

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

equipped with the platform's physical capacity of either spinning disks or enterprise-value SSDs for maximum data capacity.

Cisco UCS HXAF220c-M5S Rack Server

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

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

Figure 12 Cisco UCS HXAF220c-M5SX Rack Server Front View

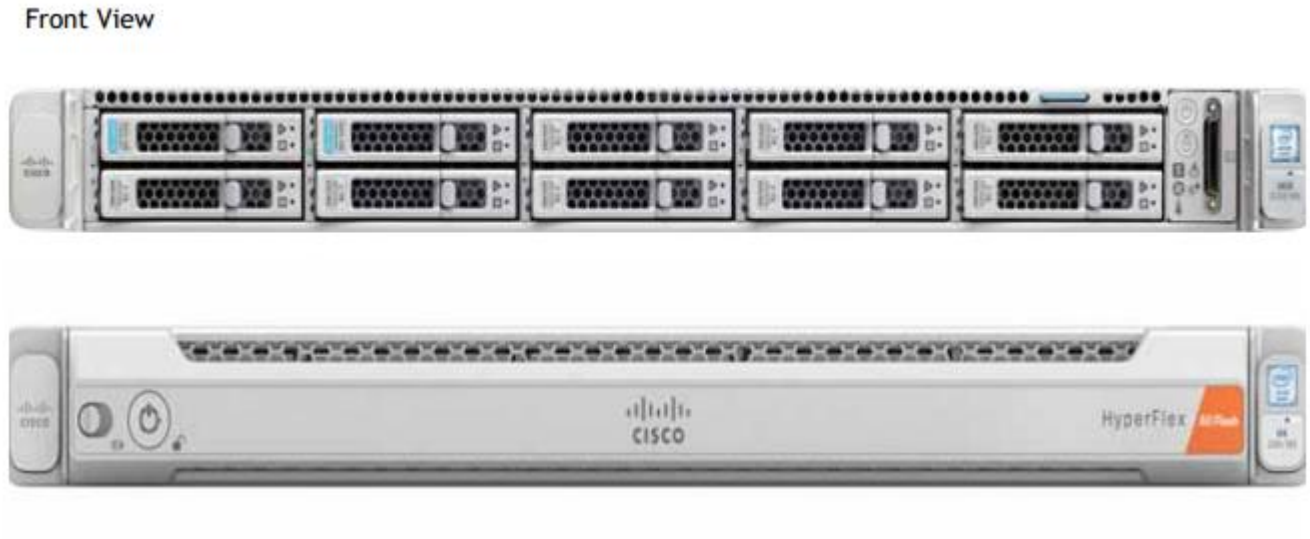

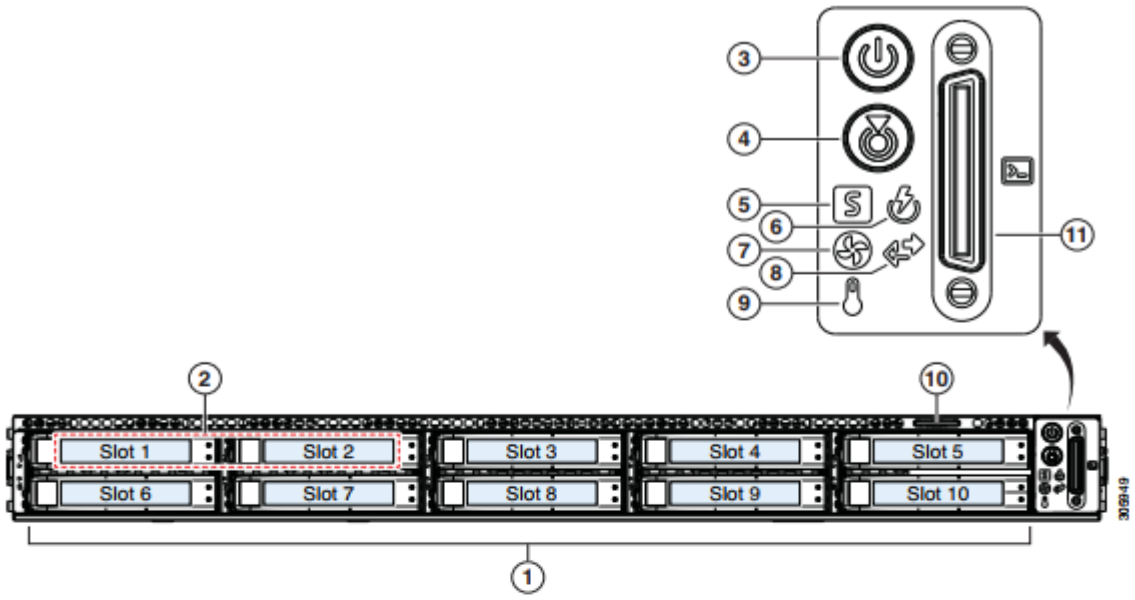

1	Drive Slots Slot 01 (For System/Log drive) • 1 x SATA SSD Slot 02 (For Cache drive) • 1 x NVMe SSD **OR** • 1 x SAS SSD **OR** • 1 x SED SAS SSD Slot 03 through 10 (For Capacity drives) • Upto 8 x SATA SSD **OR** • Upto 8 x SED SATA SSD **OR** • upto 8 x SED SAS SSD	7	Fan status LED
2	N/A	8	Network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

Figure 13 Cisco UCS HXAF220c-M5SX Rack Server Rear View

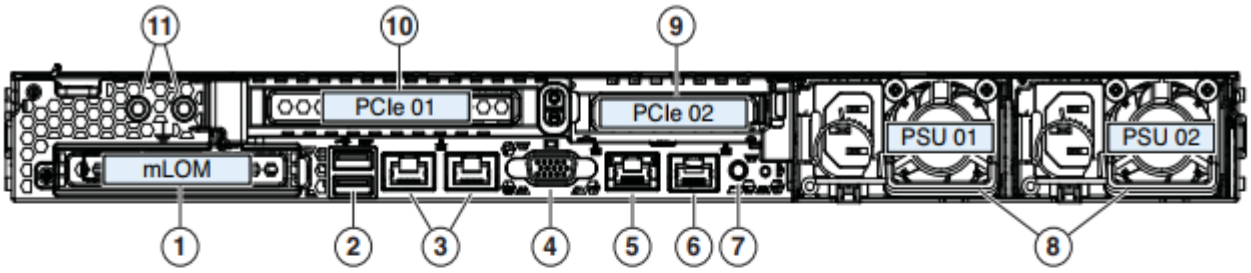

1	Modular LAN-on-motherboard (mLOM) card bay (x16)	7	Rear unit identification button/LED
2	USB 3.0 ports (two)	8	Power supplies (two, redundant as 1+1)
3	Dual 1/10-Gb Ethernet ports (LAN1 and LAN2). LAN1 is left connector and LAN2 is right connector	9	PCIe riser 2 (slot 2) (half-height, x16);
4	VGA video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

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

The Cisco UCS HXAF220c-M5S provides:

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

- Enterprise-class pass-through RAID controller
- Easily add, change, and remove Cisco FlexStorage modules

Cisco HyperFlex HXAF240c-M5SX Nodes

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

Figure 14 HXAF240c-M4SX Node

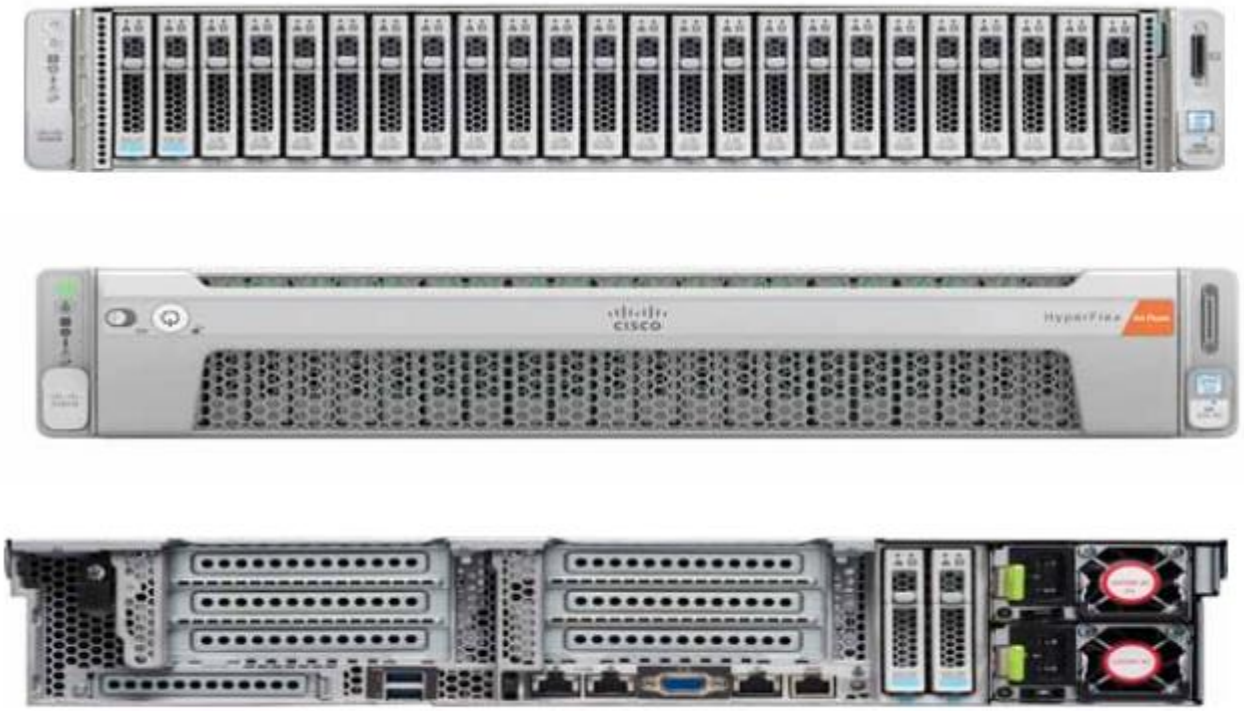

Cisco HyperFlex HX220c-M4S Hybrid Node

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

Figure 15 HX220c-M4S Node

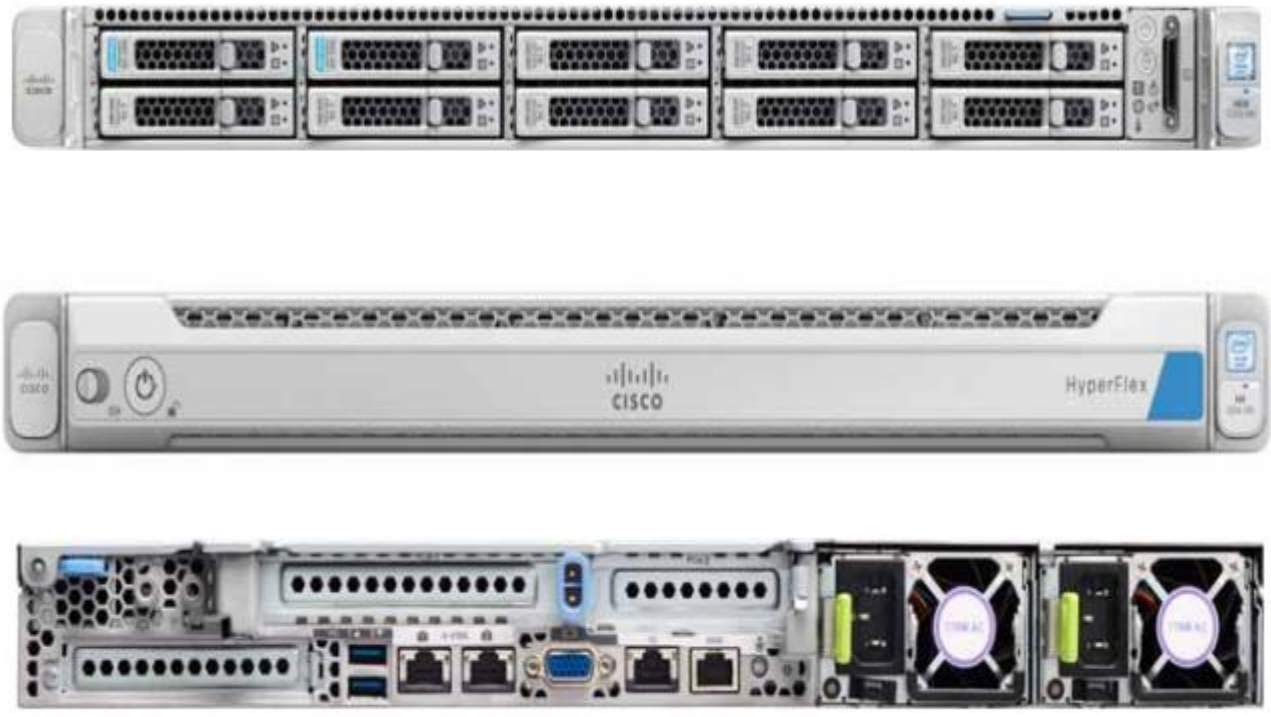

Cisco HyperFlex HX240c-M4SX Hybrid Node

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

Figure 16 HX240c-M5SX Node

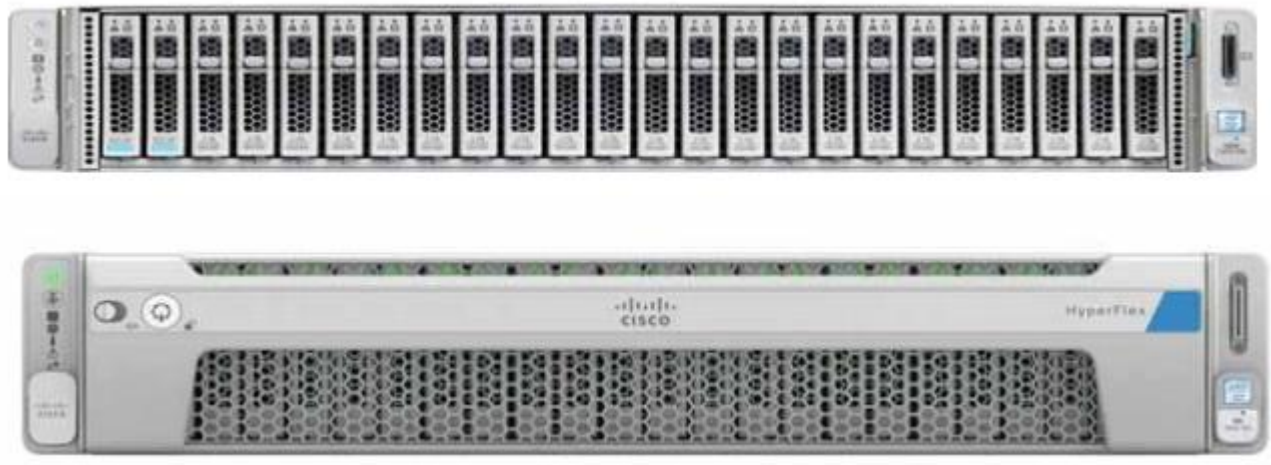

Cisco VIC 1387 MLOM Interface Card

The Cisco UCS Virtual Interface Card (VIC) 1387 is a dual-port Enhanced Small Form-Factor Pluggable (QSFP+) 40-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)) in a modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS HX-Series Rack Servers (Error! Reference source not found.). 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 17 Cisco VIC 1387 mLOM Card

Cisco HyperFlex Compute Nodes

Cisco UCS B200-M5 Blade

For workloads that require additional computing and memory resources, but not additional storage capacity, a compute-intensive hybrid cluster configuration is allowed. This configuration requires a minimum of three

(up to sixteen) HyperFlex converged nodes with one to sixteen Cisco UCS B200-M5 Blade Servers for additional computing capacity. The HX-series Nodes are configured as described previously, and the Cisco UCS B200-M5 servers are equipped with boot drives. Use of the Cisco UCS B200-M5 compute nodes also requires the Cisco UCS 5108 blade server chassis, and a pair of Cisco UCS 2300/2200 series Fabric Extenders. For detailed information, see the Cisco UCS B200 M5 Blade Server Spec Sheet.

Figure 18 Cisco UCS B200 M5 Server

Cisco VIC1340 Converged Network Adapter

The Cisco UCS Virtual Interface Card (VIC) 1340 (Figure 19) 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 19 Cisco UCS VIC 1340

Figure 20 illustrates the Cisco UCS VIC 1340 Virtual Interface Cards Deployed in the Cisco UCS B-Series B200 M4 Blade Servers.

Figure 20 Cisco UCS VIC 1340 Deployed in the Cisco UCS B200 M4

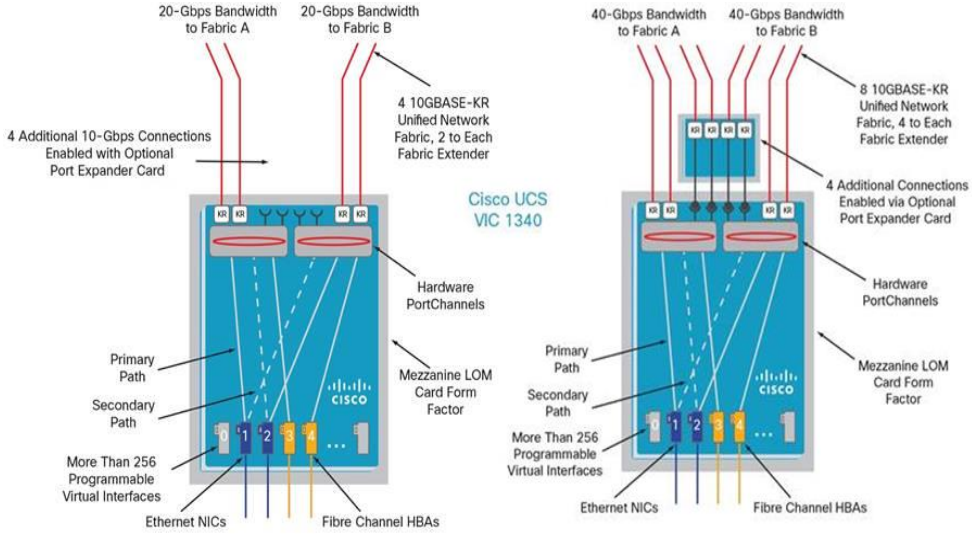

Cisco UCS 5108 Blade Chassis

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

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

Figure 21 Cisco UCS 5108 Blade Chassis Front and Rear Views

Cisco UCS 2304XP Fabric Extender

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

second-generation Cisco UCS 2200 Series Fabric Extenders and is backward compatible with the shipping Cisco UCS 5108 Blade Server Chassis.

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

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

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

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

Figure 22 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 scalable family processor product family, next-generation DDR4 memory, and 12-Gbps SAS throughput, delivering significant performance and efficiency gains. Cisco UCS C220 M5 Rack Server can be used to build a compute-intensive hybrid HX cluster, for an environment where the workloads require additional computing and memory resources but not additional storage capacity, along with the HX-series converged nodes. This configuration contains a minimum of three (up to sixteen) HX-series converged nodes with one to sixteen Cisco UCS C220-M4 Rack Servers for additional computing capacity.

Figure 23 Cisco UCS C220 M5 Rack Server

Cisco UCS C240-M5 Rack Server

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

Figure 24 Cisco UCS C240 M5 Rack Server

Cisco HyperFlex Converged Data Platform Software

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

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

Cisco HyperFlex HX Data Platform Administration Plug-in

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 that prefer a light weight 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 25 HyperFlex Web Client Plug-in

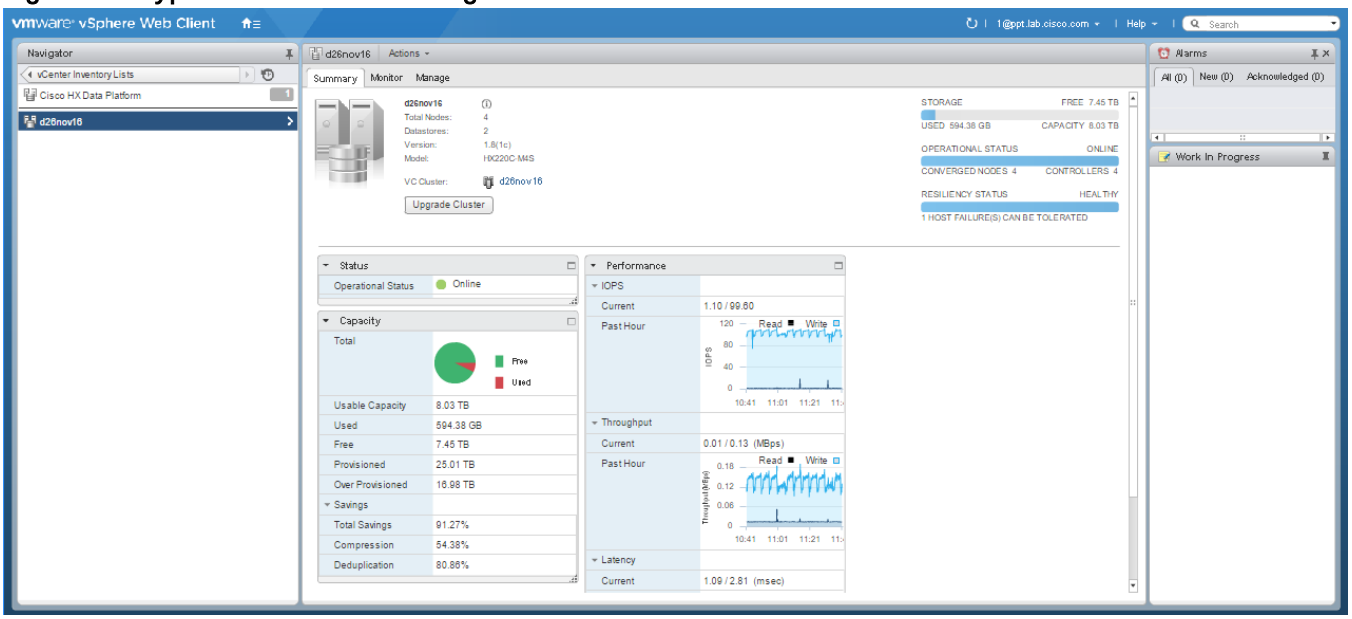

Cisco HyperFlex Connect HTML5 Management Web Page

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>.

Figure 26 HyperFlex Connect GUI

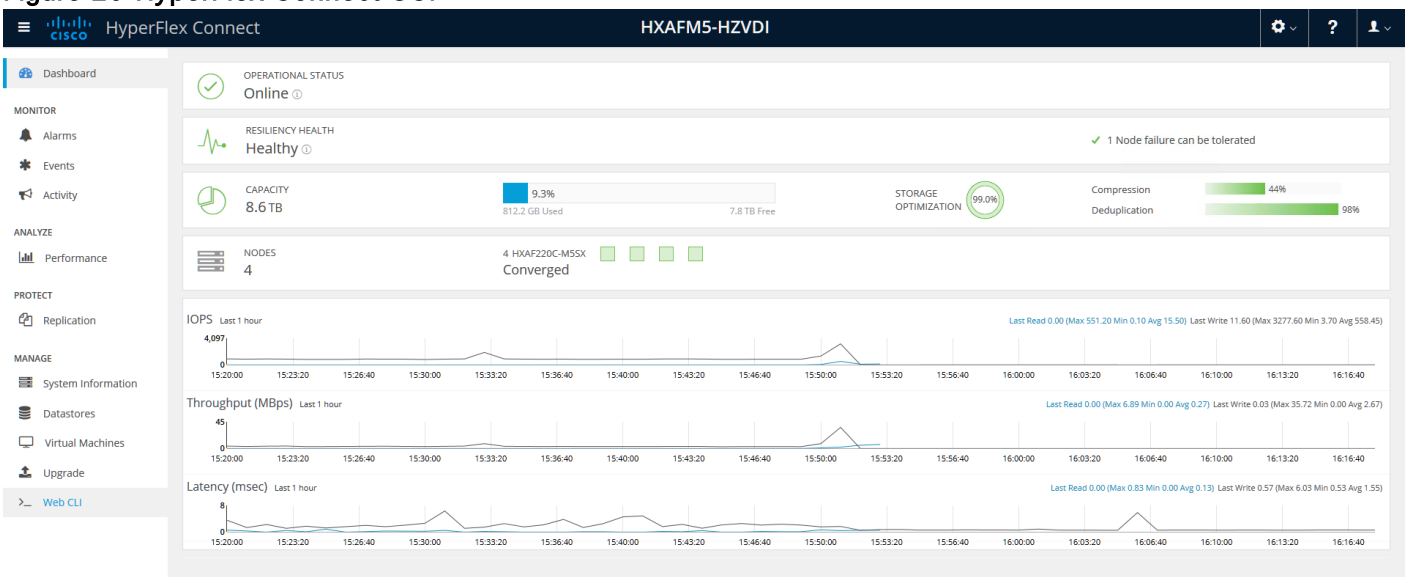

Cisco HyperFlex HX Data Platform Controller

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 27). 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 27 Data is Striped Across Nodes in the Cluster

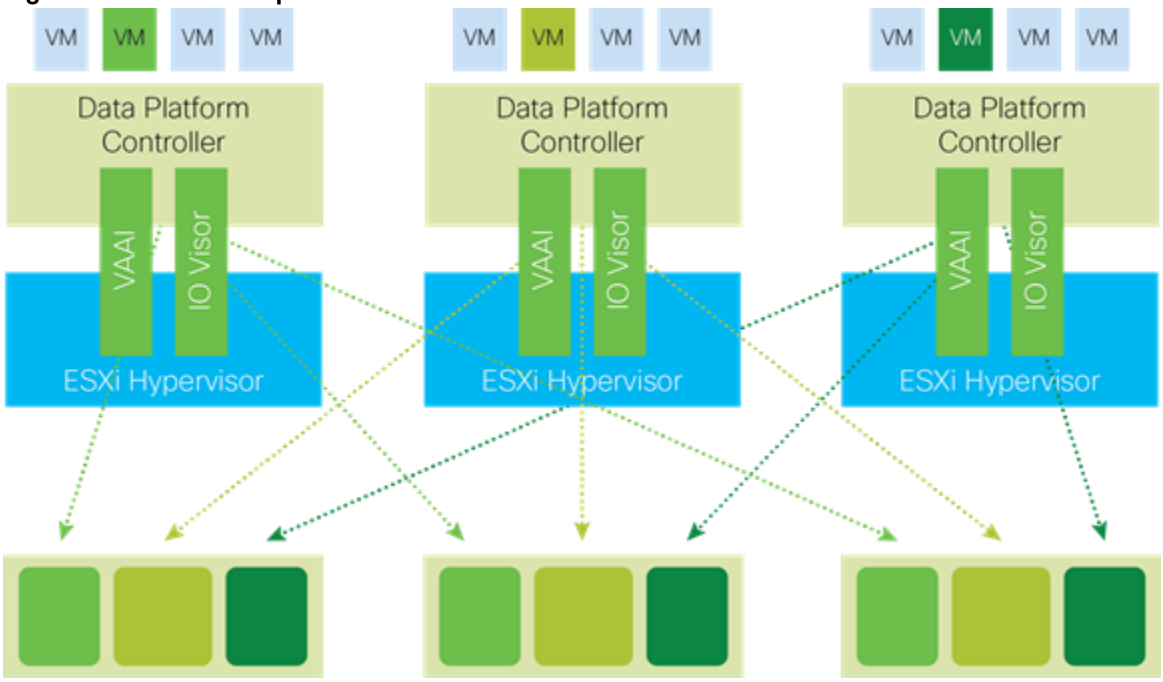

When moving a virtual machine to a new location using tools such as VMware Dynamic Resource Scheduling (DRS), the Cisco HyperFlex HX Data Platform does not require data to be moved. This approach significantly reduces the impact and cost of moving virtual machines among systems.

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 28).

Figure 28 Data Write Operation Flow Through the Cisco HyperFlex HX Data Platform

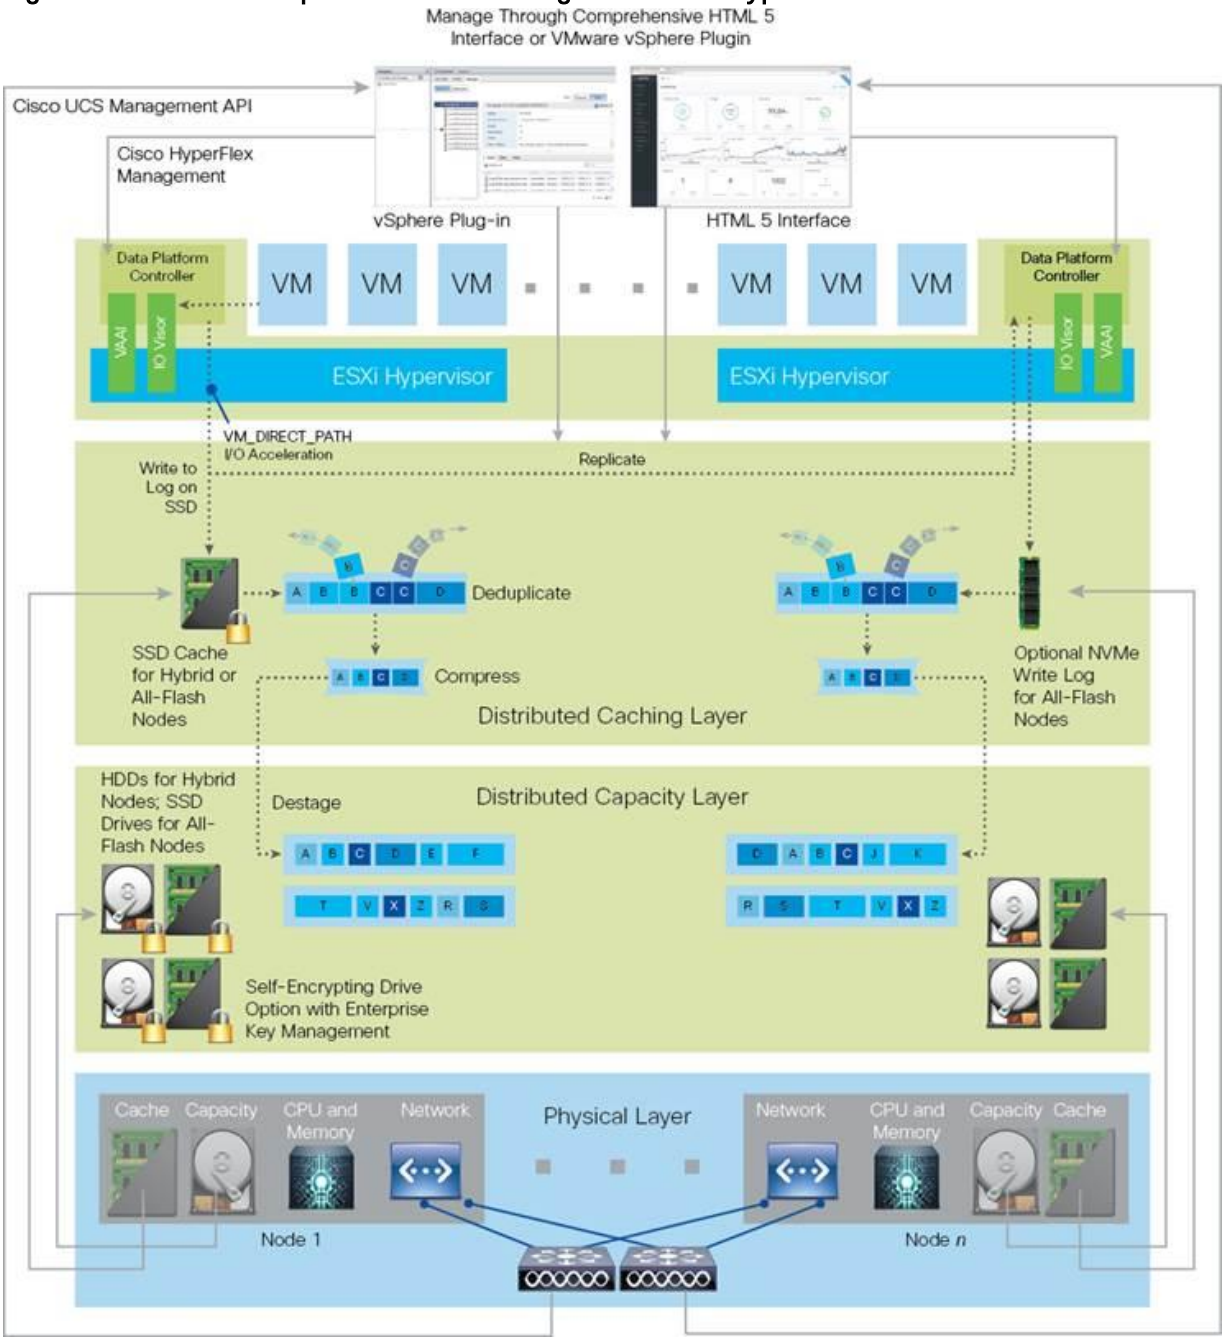

Hot data sets, data that are frequently or recently read from the capacity tier, are cached in memory. All-Flash configurations, however, does not use an SSD read cache since there is no performance benefit of such a cache; the persistent data copy already resides on high-performance SSDs. In these configurations,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

Data Optimization

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

Data Deduplication

Data deduplication is used on all storage in the cluster, including memory and SSD drives. Based on a patent-pending Top-K Majority algorithm, the platform uses conclusions from empirical research that show that most data, when sliced into small data blocks, has significant deduplication potential based on a minority of the data blocks. By fingerprinting and indexing just these frequently used blocks, high rates of deduplication can be achieved with only a small amount of memory, which is a high-value resource in cluster nodes (Figure 29).

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

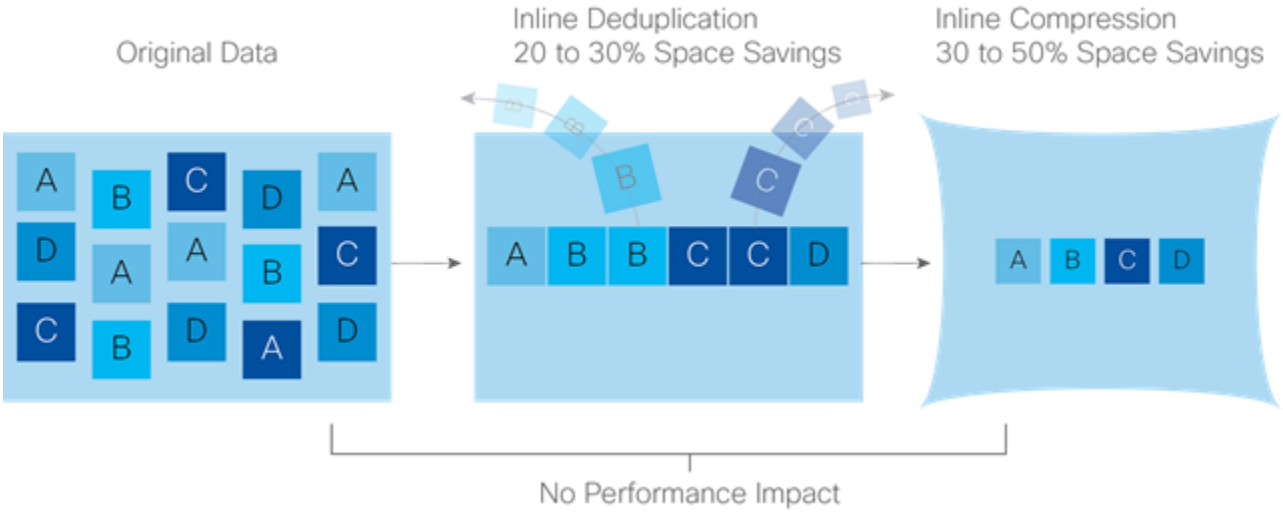

Inline Compression

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 30). 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 30 Cisco HyperFlex HX Data Platform Optimizes Data Storage with No Performance Impact

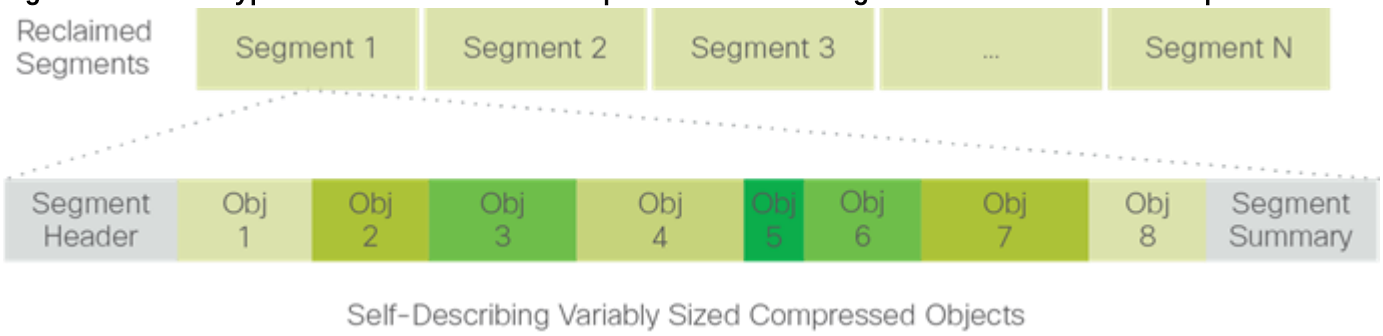

Encryption

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

Data Services

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

Thin Provisioning

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

storage-intensive nodes as your business requirements dictate, eliminating the need to purchase large amounts of storage before you need it.

Snapshots

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

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

Many basic backup applications, read the entire dataset, or the changed blocks since the last backup at a rate that is usually as fast as the storage, or the operating system can handle. This can cause performance implications since HyperFlex is built on Cisco UCS with 10GbE which could result in multiple gigabytes per second of backup throughput. These basic backup applications, such as Windows Server Backup, should be scheduled during off-peak hours, particularly the initial backup if the application lacks some form of change block tracking.

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

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

Fast, Space-Efficient Clones

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

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

Data Replication and Availability

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

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

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

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

Data Rebalancing

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

Online Upgrades

Cisco HyperFlex HX-Series systems and the HX Data Platform support online upgrades so that you can expand and update your environment without business disruption. You can easily expand your physical

resources; add processing capacity; and download and install BIOS, driver, hypervisor, firmware, and Cisco UCS Manager updates, enhancements, and bug fixes.

Cisco Nexus 9372PX Switches

The Cisco Nexus 9372PX/9372PX-E Switches has 48 1/10-Gbps Small Form Pluggable Plus (SFP+) ports and 6 Quad SFP+ (QSFP+) uplink ports. All the ports are line rate, delivering 1.44 Tbps of throughput in a 1-rack-unit (1RU) form factor. Cisco Nexus 9372PX benefits are listed below.

Architectural Flexibility

- Includes top-of-rack or middle-of-row fiber-based server access connectivity for traditional and leaf-spine architectures
- Leaf node support for Cisco ACI architecture is provided in the roadmap
- Increase scale and simplify management through Cisco Nexus 2000 Fabric Extender support

Feature Rich

- Enhanced Cisco NX-OS Software is designed for performance, resiliency, scalability, manageability, and programmability
- ACI-ready infrastructure helps users take advantage of automated policy-based systems management
- Virtual Extensible LAN (VXLAN) routing provides network services
- Cisco Nexus 9372PX-E supports IP-based endpoint group (EPG) classification in ACI mode

Highly Available and Efficient Design

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

Simplified Operations

- Power-On Auto Provisioning (POAP) support allows for simplified software upgrades and configuration file installation
- An intelligent API offers switch management through remote procedure calls (RPCs, JSON, or XML) over a HTTP/HTTPS infrastructure
- Python Scripting for programmatic access to the switch command-line interface (CLI)
- Hot and cold patching, and online diagnostics

Investment Protection

A Cisco 40 Gb bidirectional transceiver allows reuse of an existing 10 Gigabit Ethernet multimode cabling plant for 40 Gigabit Ethernet Support for 1 Gb and 10 Gb access connectivity for data centers migrating access switching infrastructure to faster speed. The following is supported:

- 1.44 Tbps of bandwidth in a 1 RU form factor
- 48 fixed 1/10-Gbps SFP+ ports
- 6 fixed 40-Gbps QSFP+ for uplink connectivity that can be turned into 10 Gb ports through a QSFP to SFP or SFP+ Adapter (QSA)
- Latency of 1 to 2 microseconds
- Front-to-back or back-to-front airflow configurations
- 1+1 redundant hot-swappable 80 Plus Platinum-certified power supplies
- Hot swappable 2+1 redundant fan tray

Figure 31 Cisco Nexus 9372PX Switch

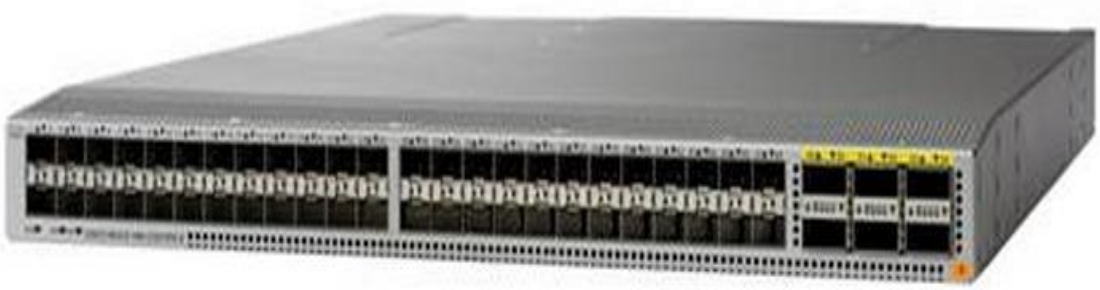

VMware vSphere 6.5

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 which 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 7.3.1 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 7.3.1.
  - 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 checkbox 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, Server 2012 R2, or Server 2016 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
  - Several new features are supported on Horizon 6 for Linux desktops, including NVIDIA GRID vGPU, vSGA, RHEL 7.1 and Ubuntu 14.04 guest operating systems, and View Agent installation of JRE 8 with no user steps required.
  - 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
- Additional Features
  - Support for lPv6 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.

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 32 VMware Horizon Architectural Overview

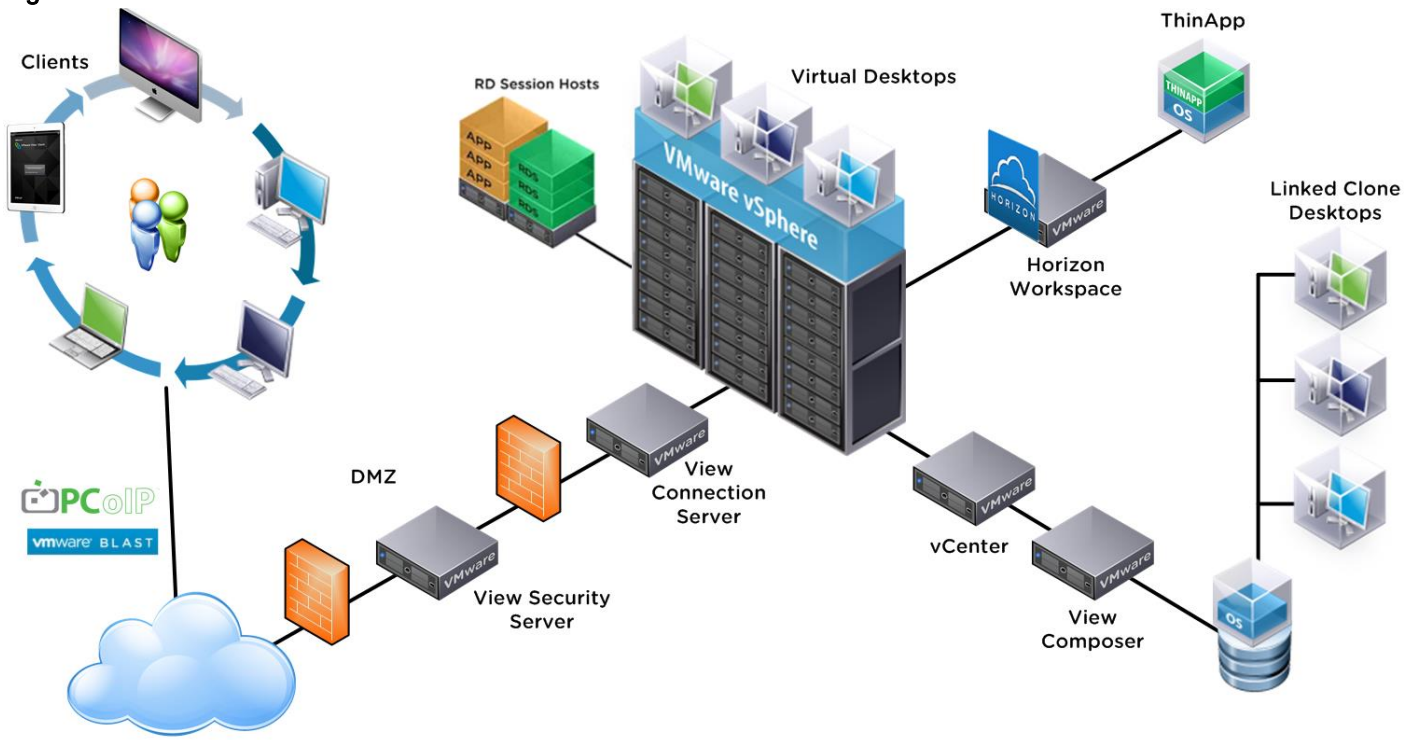

Architecture and Design of VMware Horizon on Cisco Unified Computing System and Cisco HyperFlex System Design Fundamentals

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 synced 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?
- 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 (e.g., 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 33 VMware Horizon Design Overview

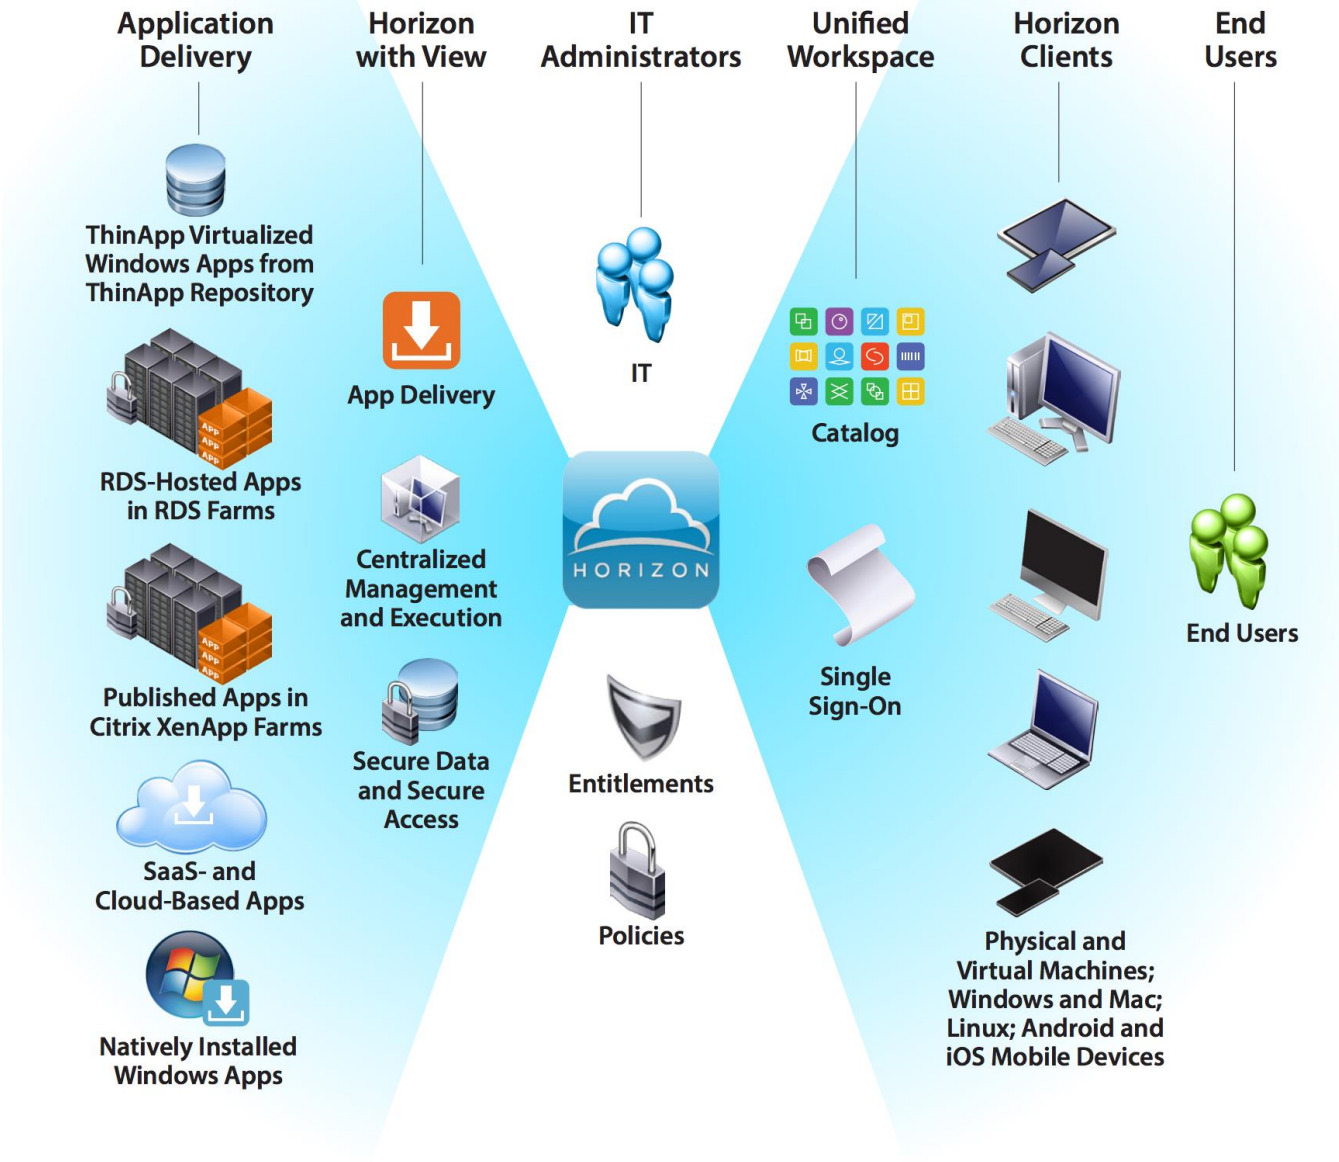

Figure 34 Horizon VDI and RDSH Desktop Delivery Based on Display Protocol (PCoIP/Blast/RDP)

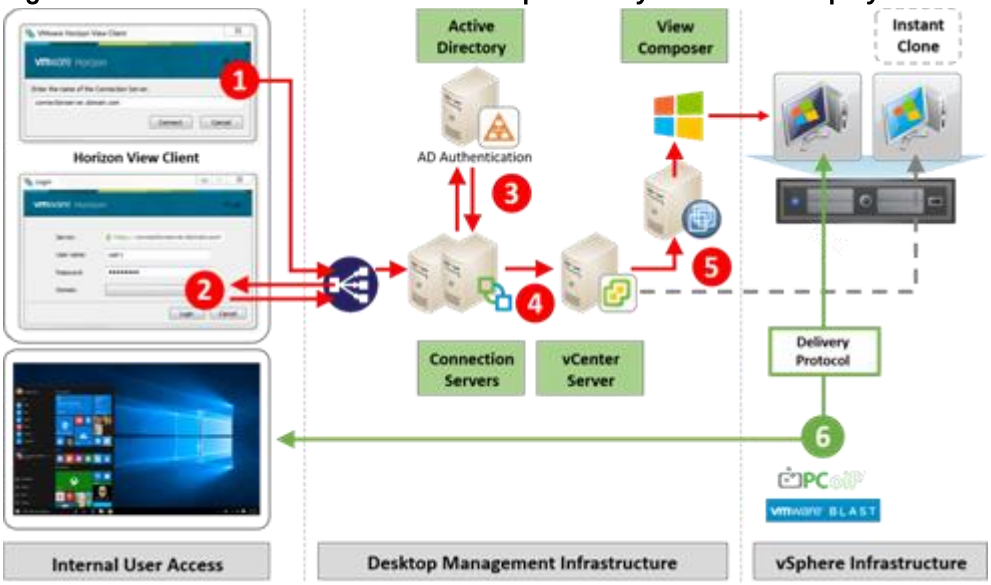

VMware Horizon Composer

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 35 VMware Horizon Composer Overview

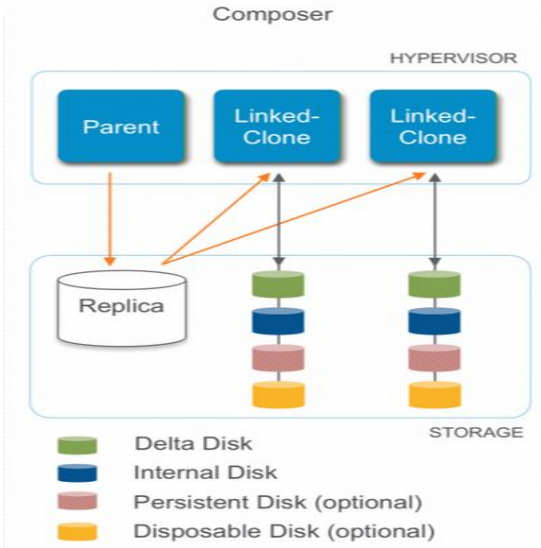

VMware View Storage Accelerator

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 (Ex:- Big-IP Global Traffic Manager) to direct the user connections to the most appropriate site to deliver the desktops and application to users.

In Figure 36, The image depicting sites, a site was 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 36 Multisite Configuration Overview

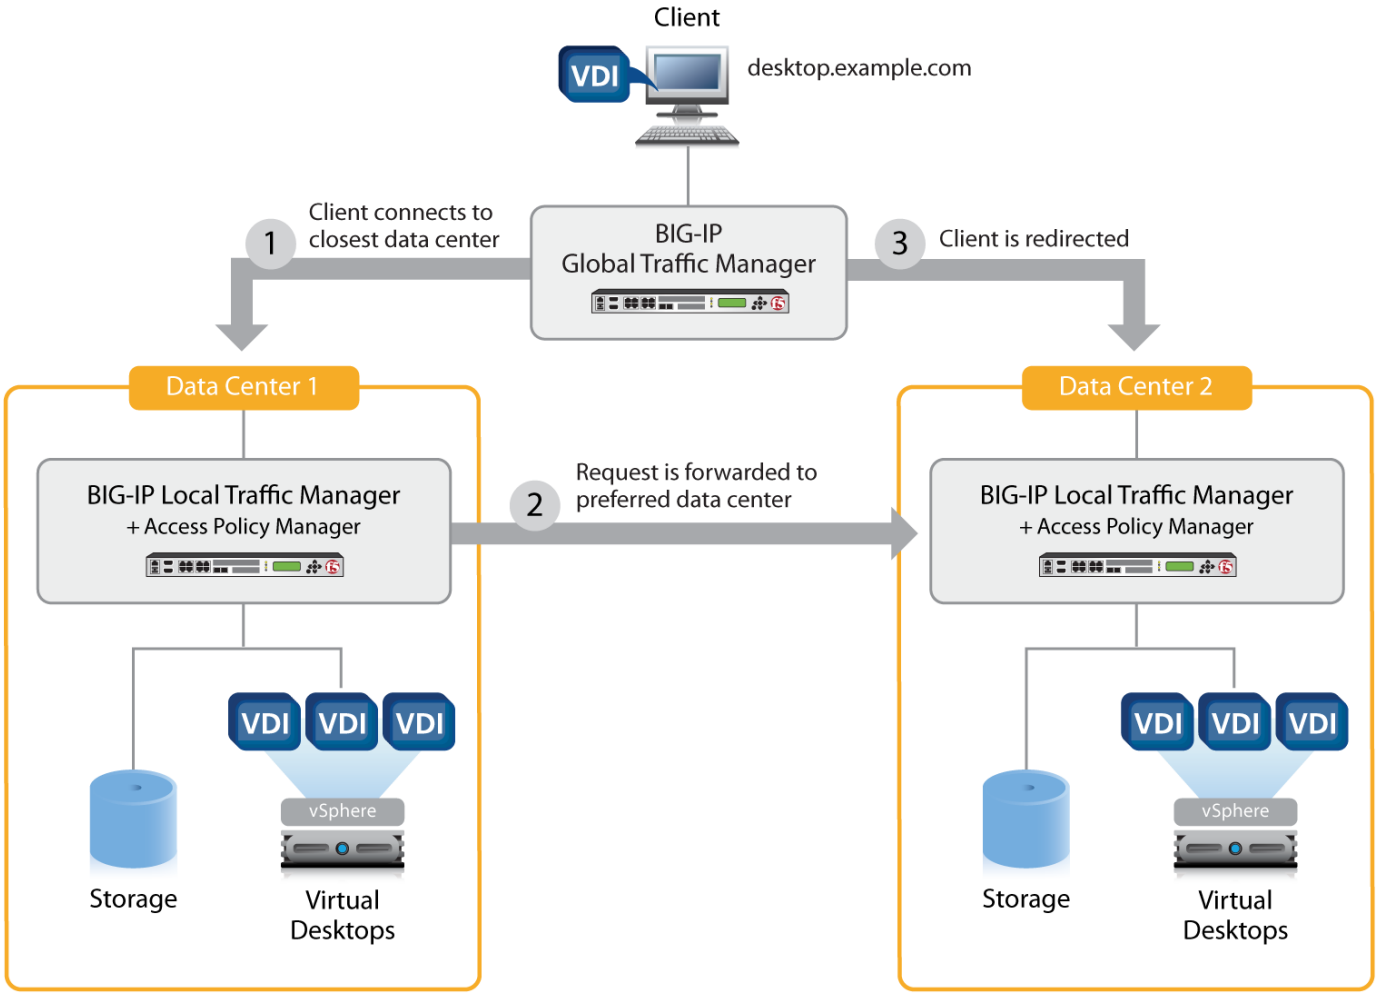

Based on the requirement and no of data centers or remote location, we can chose any of the available Load balancing software or tools accelerates the application performance, load balances servers, increases security, and optimizes the user experience. In this example, two Big-IP Local Traffic Manager are used to provide a high availability configuration.

 BIG-IP Local Traffic Manager has been shown as example for presentation purpose.

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.

Server OS machines	**You want**: Inexpensive server-based delivery to minimize the cost of delivering applications to a large number of users, while providing a secure, high-definition user experience. **Your users**: Perform well-defined tasks and do not require personalization or offline access to applications. Users may include task workers such as call center operators and retail workers, or users that share workstations. **Application types**: Any application.

Desktop OS machines	You want: A client-based application delivery solution that is secure, provides centralized management, and supports a large number of users per host server (or hypervisor), while providing users with applications that display seamlessly in high-definition. Your users: Are internal, external contractors, third-party collaborators, and other provisional team members. Users do not require off-line access to hosted applications. Application types: Applications that might not work well with other applications or might interact with the operating system, such as .NET framework. These types of applications are ideal for hosting on virtual machines. Applications running on older operating systems such as Windows XP or Windows Vista, and older architectures, such as 32-bit or 16-bit. By isolating each application on its own virtual machine, if one machine fails, it does not impact other users.

Remote PC Access	You want: Employees with secure remote access to a physical computer without using a VPN. For example, the user may be accessing their physical desktop PC from home or through a public WIFI hotspot. Depending upon the location, you may want to restrict the ability to print or copy and paste outside of the desktop. This method enables BYO device support without migrating desktop images into the datacenter. Your users: Employees or contractors that have the option to work from home, but need access to specific software or data on their corporate desktops to perform their jobs remotely. Host: The same as Desktop OS machines. Application types: Applications that are delivered from an office computer and display seamlessly in high definition on the remote user's device.

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 Linked-clone and Instant-clone automated pool were configured and tested. The following sections discuss design decisions relative to the VMware Horizon deployment, including the CVD test environment.

Deployment Hardware and Software

Products Deployed

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

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

This validated design document details the deployment of the multiple configurations extending to 450 or 600 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 Virtual Desktops (VDI) on Cisco HyperFlex
- Microsoft Windows Server 2016 for User Profile Manager
- Microsoft Windows 2016Server for Login VSI Management and data servers to simulate real world VDI workload
- VMware vSphere ESXi 6.5 Update 1 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 v2.6.1a
- VMware Horizon 7 Connection Server and Replica Servers for redundancy and support up to 600 seat scale
- VMware Horizon 7 View Composer Server

Figure 37 Detailed Reference Architecture with Physical Hardware Cabling Configured to Enable the Solution

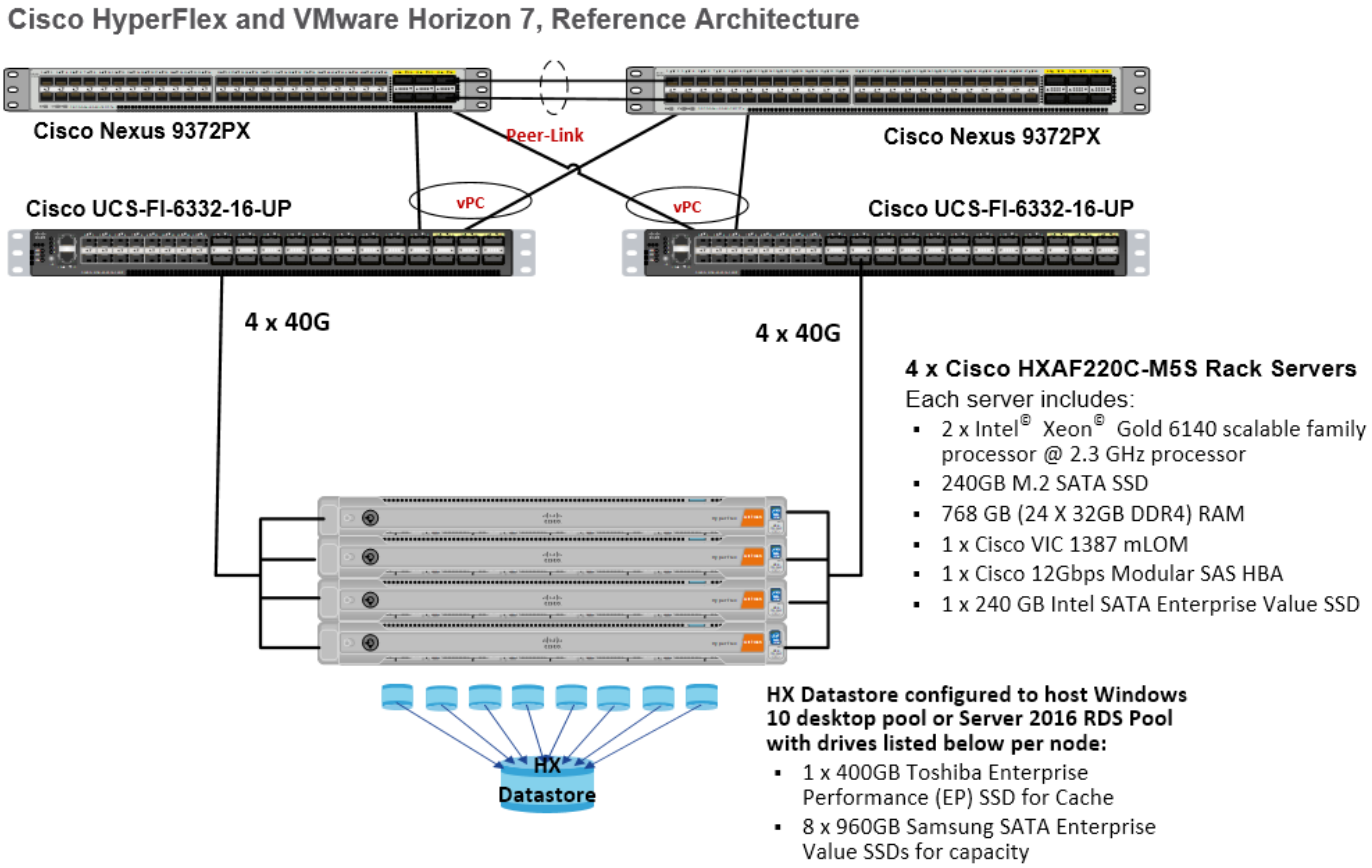

Figure 38 Detailed Reference Architecture with Physical Hardware Cabling Configured to Enable the Scale-Out Solution as Per the Current Cluster Limit

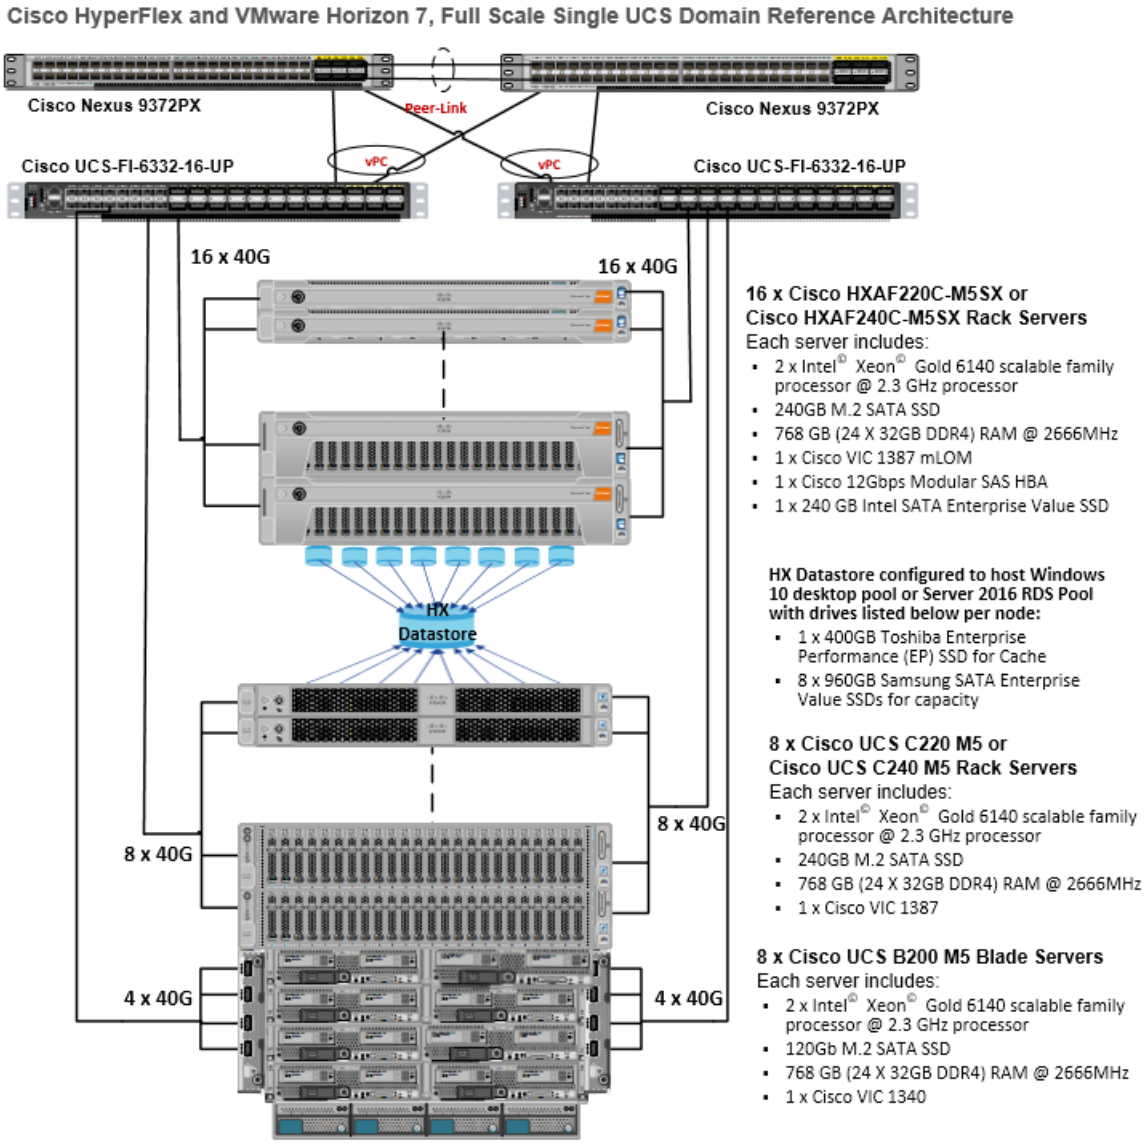

Hardware Deployed

The solution contains the following hardware as shown in Figure 37:

- Two Cisco Nexus 9372PX Layer 2 Access Switches
- Two Cisco UCS C220 M4 Rack servers with dual socket Intel Xeon E5-2620v4 2.1-GHz 8-core processors, 128GB RAM 2133-MHz and VIC1227 mLOM card for the hosted infrastructure with N+1 server fault tolerance. (Not show in the diagram).
- Four Cisco UCS HXAF220c-M5S Rack servers with Intel Xeon Gold 6140 scalable family 2.3-GHz 18-core processors, 768GB RAM 2666-MHz and VIC1387 mLOM cards running Cisco HyperFlex data platform v2.6.1a for the virtual desktop workloads with N+1 server fault tolerance

Software Deployed

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

Table 1 Software and Firmware Versions

Vendor	Product	Version
Cisco	UCS Component Firmware	3.2(2b) bundle release
Cisco	UCS Manager	3.2(2b) bundle release
Cisco	UCS HXAF220c-M5S rack server	3.2(2b) bundle release
Cisco	VIC 1387	4.2(2b)
Cisco	HyperFlex Data Platform	2.6.1a-26588
Cisco	Cisco NENIC	1.0.2.02
Cisco	Cisco fNIC	1.6.0.34
Network	Cisco Nexus 9000 NX-OS	7.0(3)I2(2d)
VMware	Horizon Connection Server	7.3.1-6760913
VMware	Horizon Composer Server	7.3.1-6744335
VMware	Horizon Agent	7.3.1-6761322
VMware	Horizon Client	4.4.0-5171611
VMware	vCenter Server Appliance	6.5.0-5973321
VMware	vSphere ESXi 6.5 Update 1	6.5 U1-5969303

Logical Architecture

The logical architecture of this solution is designed to support up to 4000 Hosted Virtual Microsoft Windows 10 Desktops and RDSH hosted shared server desktop users within a sixteen node Cisco UCS HXAF220c-M4S, eight Cisco UCS C220 M4 and eight Cisco UCS B200 M4 HyperFlex cluster, which provides physical redundancy for each workload type.

Figure 39 Logical Architecture Design

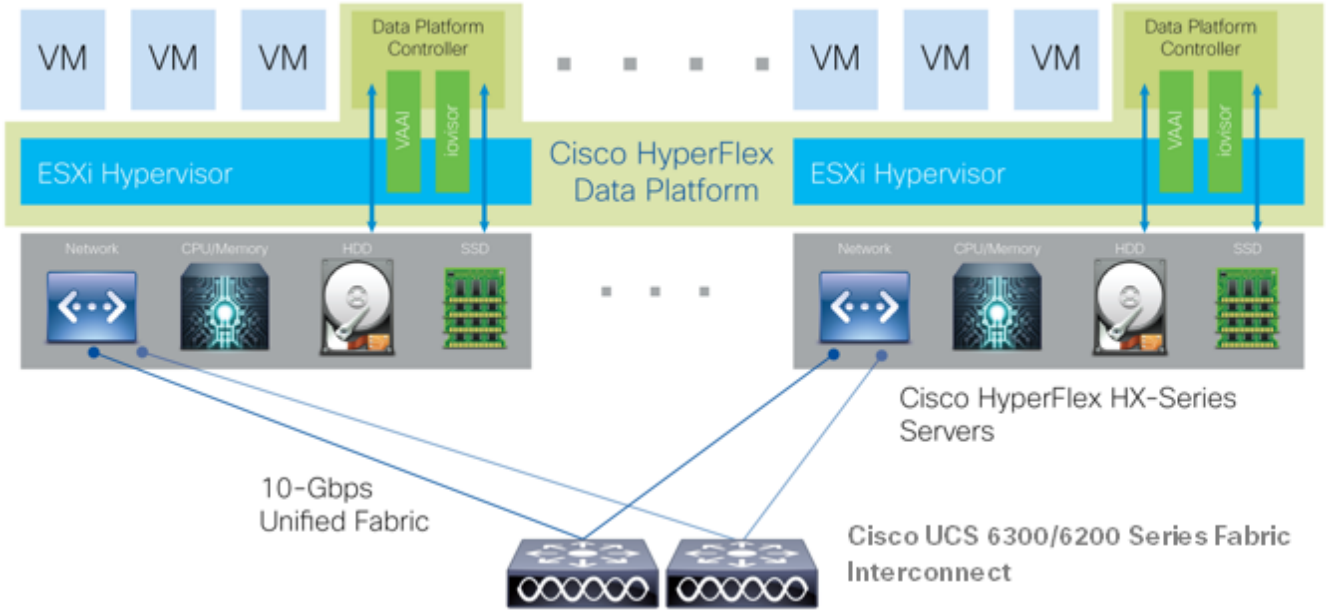

Table 1 lists the software revisions for this solution.

This document is intended to allow you to fully configure your environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 2 through Table 6 lists the information you need to configure your environment.

VLANs

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

Table 2 Table 2 VLANs Configured in this Study

VLAN Name	VLAN ID	VLAN Purpose
Default	1	Native VLAN
Hx-in-Band-Mgmt	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

VLAN Name	VLAN ID	VLAN Purpose
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 Connection Servers, VMware Replica Servers, View Composer Server, Cisco Nexus 1000v Virtual Supervisor Module, and VSMs, etc.)
- HyperFlex Cluster: VMware Horizon RDSH VMs (Windows Server 2016) or Persistent/Non-Persistent VDI VM Pools (Windows 10 64-bit)
- 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 40 VMware vSphere Clusters on vSphere Web GUI

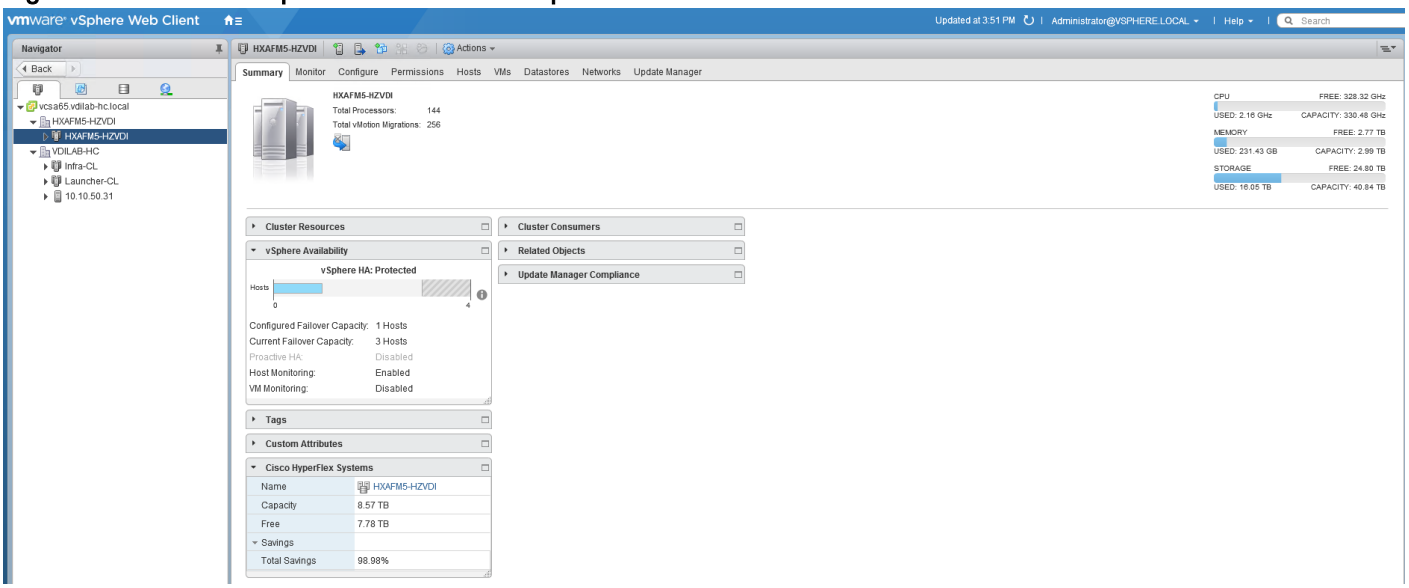

ESXi Host Design

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 3 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	vmnic1	hx-inband-mgmt	no
vswitch-hx-storage-data	Storage Controller Data Network Storage Hypervisor Data Network	vmnic3	vmnic2	hx-storage-data	yes
vswitch-hx-vm-network	none	vmnic4,vmnic5	none	vm-network	no
vmotion	none	vmnic6	vmnic7	hx-vmotion	yes

Figure 41 ESXi Network Design

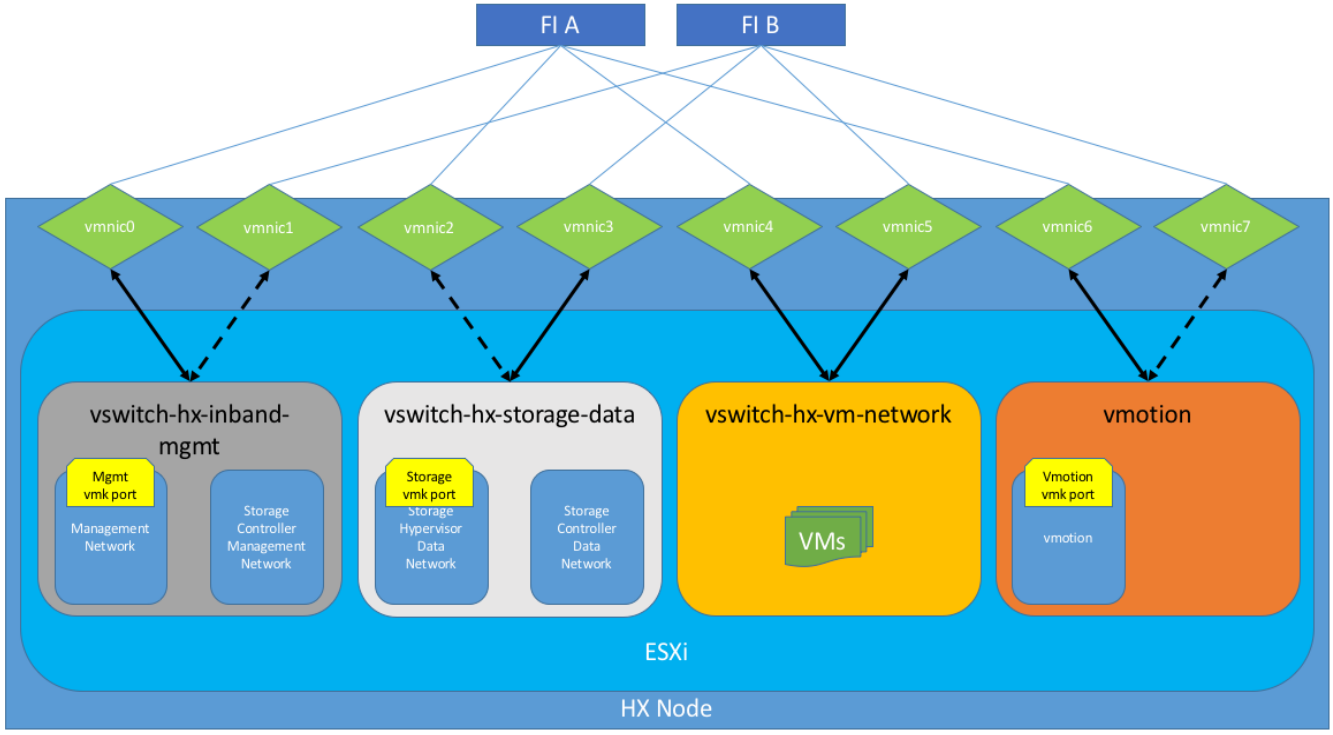

VMDirectPath I/O Pass-through

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.

Figure 42 Datastore Example

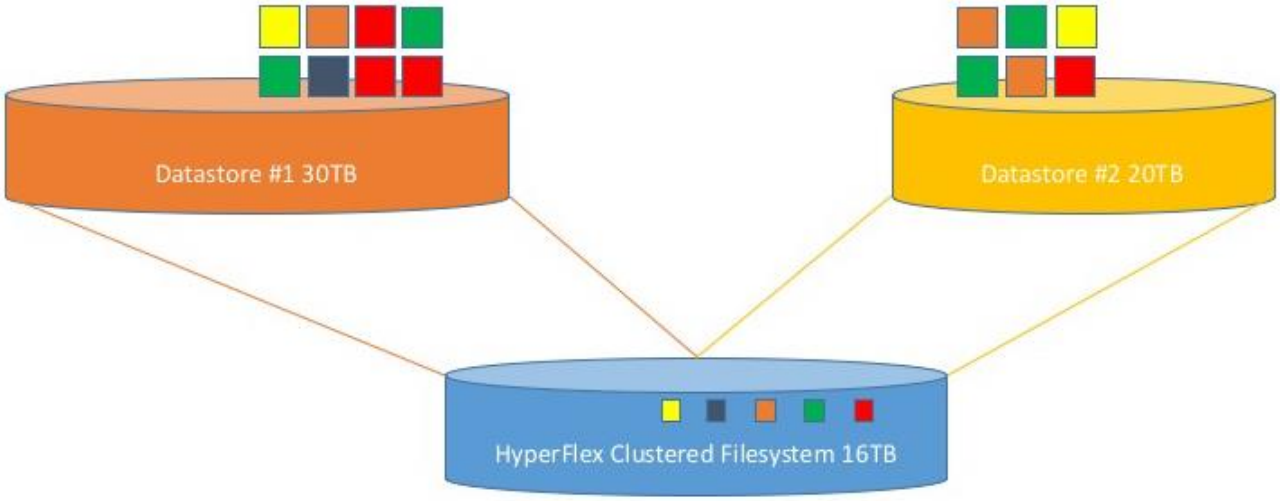

CPU Resource Reservations

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 4 details the CPU resource reservation of the storage controller VMs:

Table 4 Controller VM CPU Reservations

Number of vCPU	Shares	Reservation	Limit
8	Low	10800 MHz	unlimited

Memory Resource Reservations

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

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

Table 5 Controller VM Memory Reservations

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HX220c-M4S HXAF220c-M4S	48 GB	Yes
HX240c-M4SX	72 GB	Yes

Server Model	Amount of Guest Memory	Reserve All Guest Memory
HXAF240c-m4SX		

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 43 illustrates the configuration topology for this solution.

Figure 43 Configuration Topology for Scalable VMware Horizon 7 Workload with HyperFlex

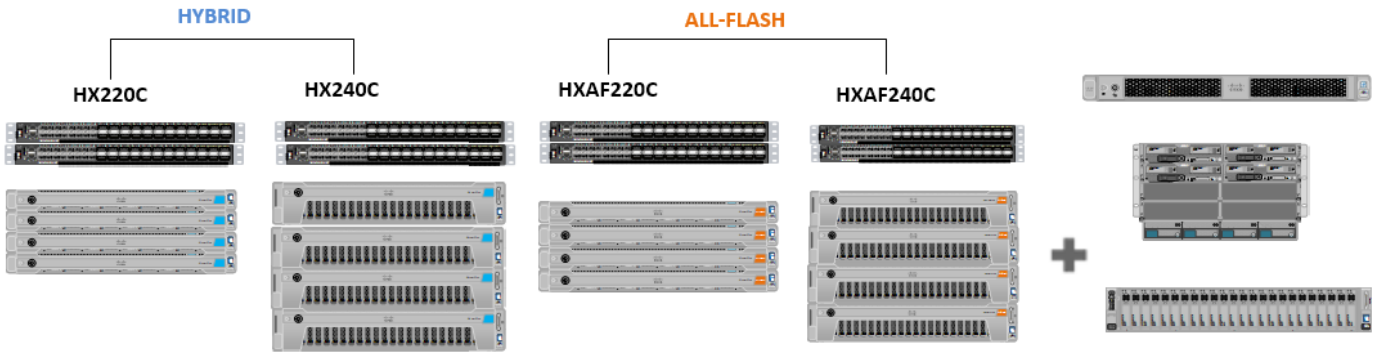

Cisco UCS Compute Platform

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

Table 6 Cisco Nexus 9372-Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 A	Eth1/1	10GbE	Cisco Nexus 9372 B	Eth1/1
	Eth1/2	10GbE	Cisco Nexus 9372 B	Eth1/2
	Eth1/3	10GbE	Cisco UCS fabric interconnect A	Eth1/13
	Eth1/4	10GbE	Cisco UCS fabric interconnect A	Eth1/14
	Eth1/5	10GbE	Cisco UCS fabric interconnect B	Eth1/13
	Eth1/6	10GbE	Cisco UCS fabric interconnect B	Eth1/14
	Eth1/25	10GbE	Infra-host-01	Port01
	Eth1/26	10GbE	Infra-host-02	Port01
	Eth1/27	10GbE	Launcher-host-01	Port01
	Eth1/28	10GbE	Launcher-host-02	Port01
	Eth1/29	10GbE	Launcher-host-03	Port01
	Eth1/30	10GbE	Launcher-host-04	Port01
	MGMT0	GbE	GbE management switch	Any

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

Table 7 Cisco Nexus 9372-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9372 B	Eth1/1	10GbE	Cisco Nexus 9372 A	Eth1/1
	Eth1/2	10GbE	Cisco Nexus 9372 A	Eth1/2
	Eth1/3	10GbE	Cisco UCS fabric interconnect A	Eth1/15
	Eth1/4	10GbE	Cisco UCS fabric interconnect A	Eth1/16
	Eth1/5	10GbE	Cisco UCS fabric interconnect B	Eth1/15
	Eth1/6	40GbE	Cisco UCS fabric interconnect B	Eth1/16
	Eth1/25	10GbE	Infra-host-01	Port02

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/26	10GbE	Infra-host-02	Port02
	Eth1/27	10GbE	Launcher-host-01	Port02
	Eth1/28	10GbE	Launcher-host-02	Port02
	Eth1/29	10GbE	Launcher-host-03	Port02
	Eth1/30	10GbE	Launcher-host-04	Port02
	MGMT0	GbE	GbE management switch	Any

Table 8 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/13	10GbE	Cisco Nexus 9372 A	Eth1/3
	Eth1/14	10GbE	Cisco Nexus 9372 A	Eth1/4
	Eth1/15	10GbE	Cisco Nexus 9372 B	Eth1/5
	Eth1/16	10 GbE	Cisco Nexus 9372 B	Eth 1/6
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 9 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect B	Eth1/13	10GbE	Cisco Nexus 9372 B	Eth1/3
	Eth1/14	10GbE	Cisco Nexus 9372 B	Eth1/4
	Eth1/15	10GbE	Cisco Nexus 9372 A	Eth1/5
	Eth1/16	10GbE	Cisco Nexus 9372 A	Eth 1/6
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect A	L1

Local Device	Local Port	Connection	Remote Device	Remote Port
	L2	GbE	Cisco UCS fabric interconnect A	L2

Figure 44 Cable Connectivity Between Cisco Nexus 9372 A and B to Cisco UCS 6248 Fabric A and B

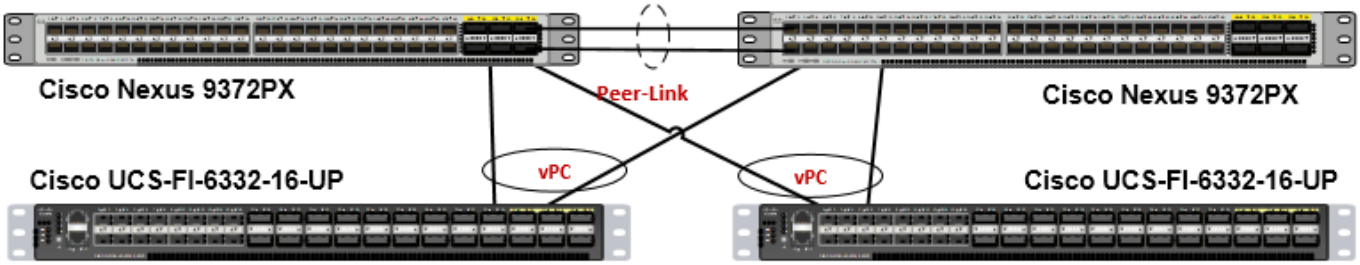

Cisco Unified Computing System Configuration

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

For complete detail on racking, power, and installation of the chassis is described in the install guide (see www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-guides-list.html) 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 45 Cisco UCS Manager Configuration: HyperFlex Sub-organization

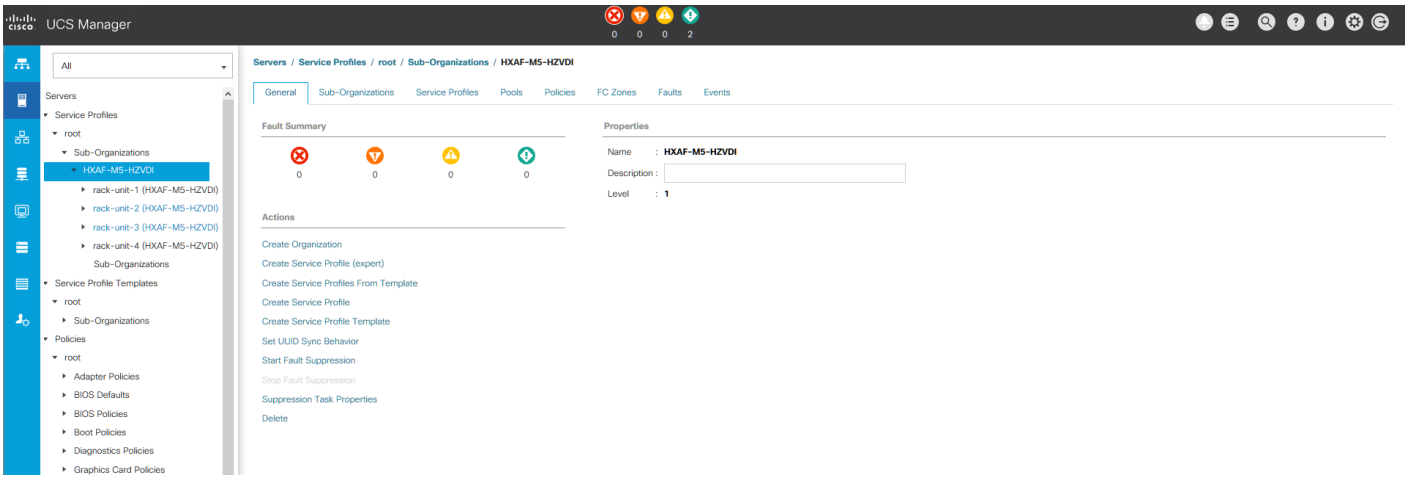

Deploy and Configure HyperFlex Data Platform

Prerequisites

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

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

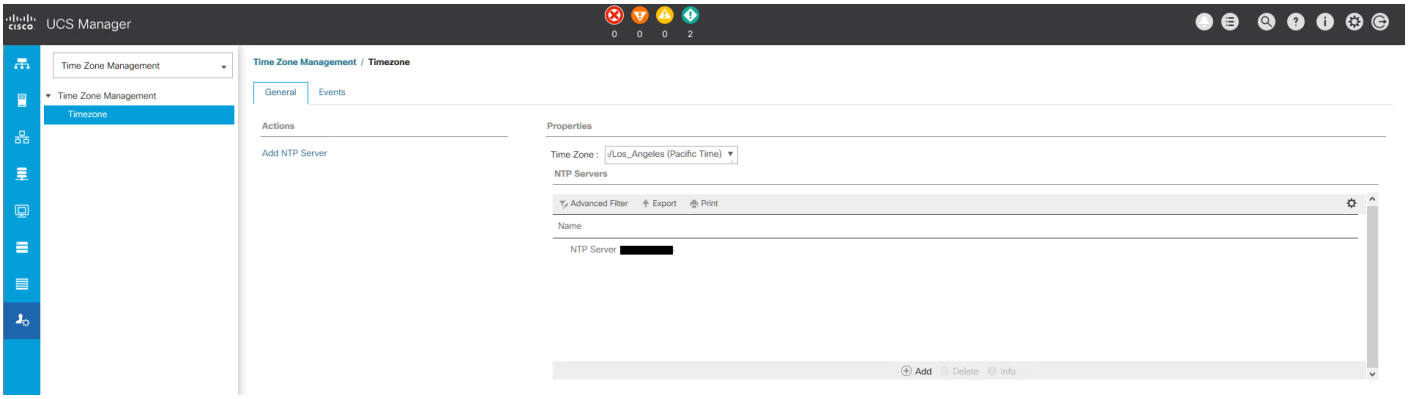

2. Configure Server Ports: Under the Equipment tab, Select Fabric A, select port to be configured as server port to manager HyperFlex rack server through Cisco UCS Manager.

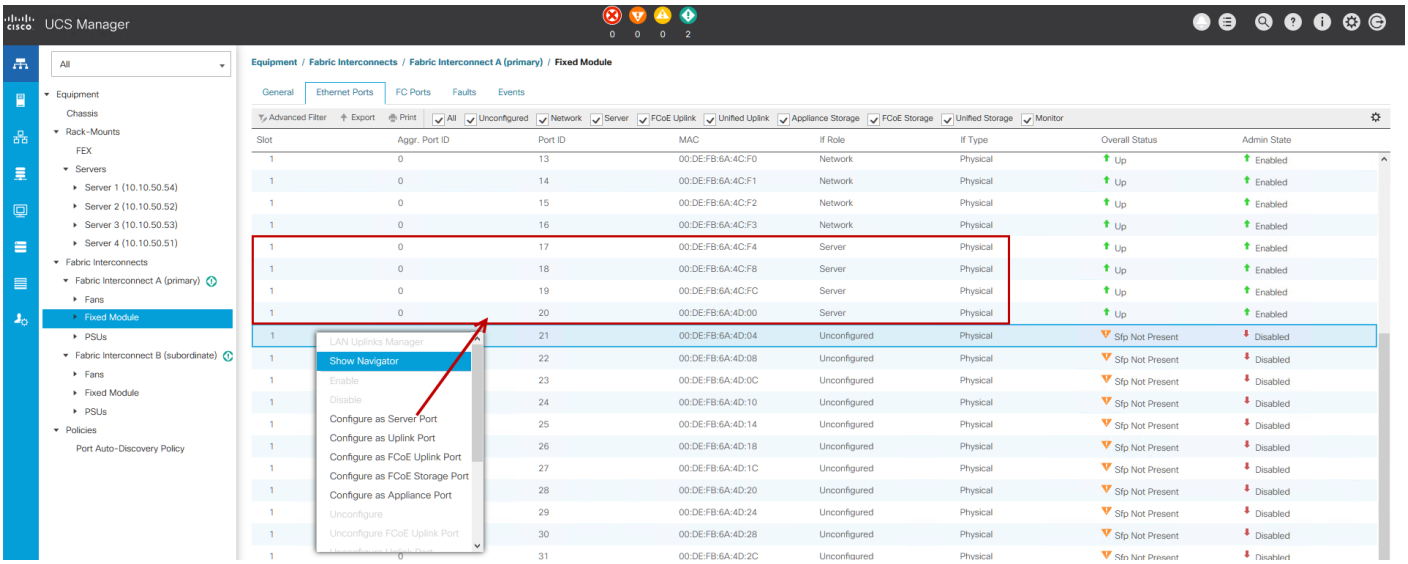

3. Repeat this step to configure server port on Fabric B.

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

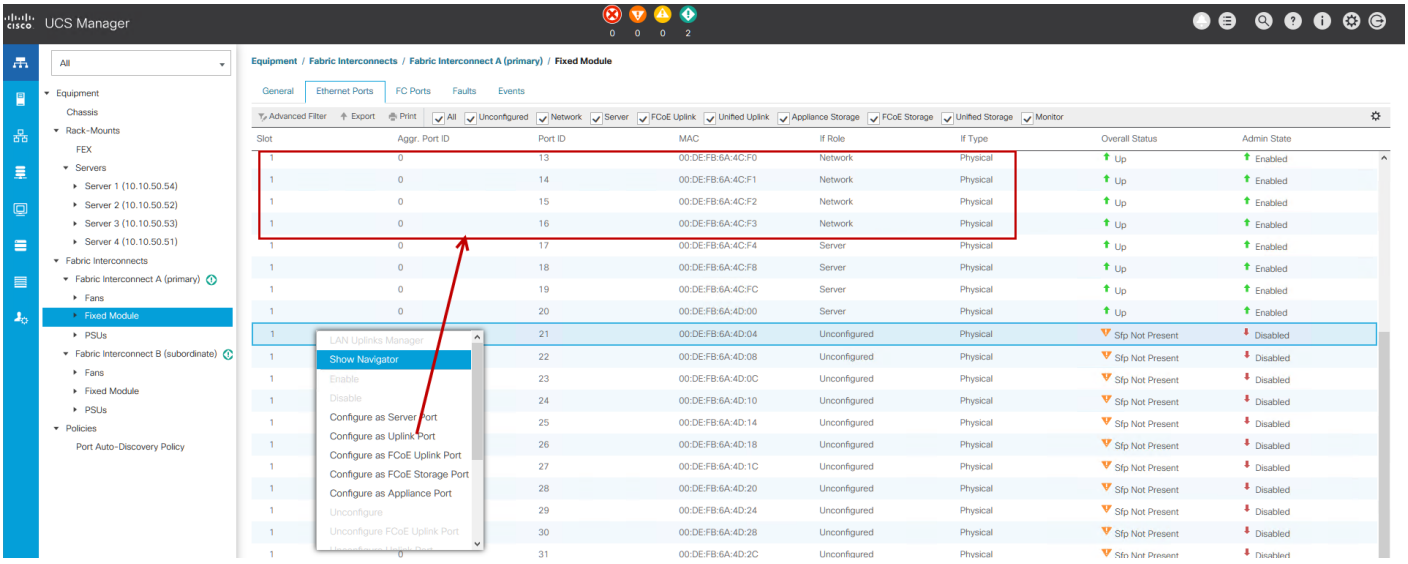

5. Repeat this same on Fabric B.

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

7. Select Create port-channel to connect with upstream switch as per Cisco UCS best practice. For our reference architecture, we connected a pair of Nexus 9372PX switches.

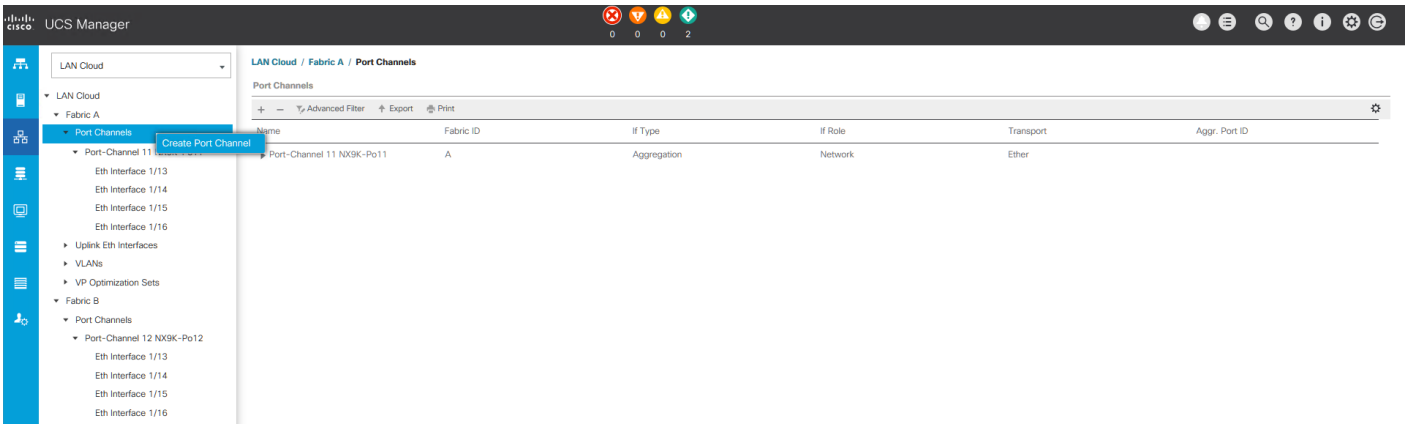

8. Enter port-channel ID number and name to be created, click Next.

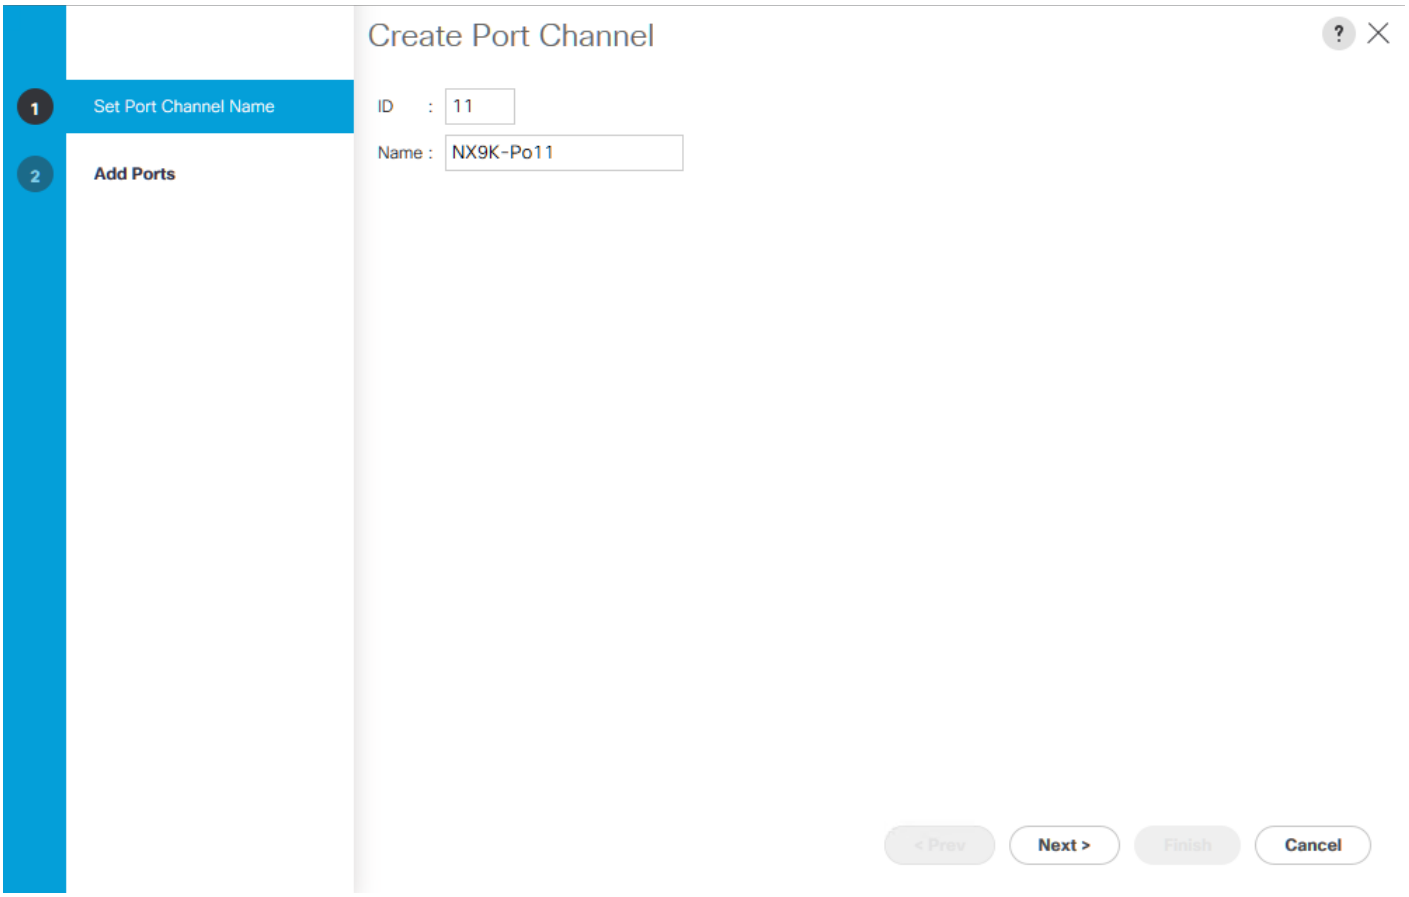

9. Select uplink ports to add as part of the port-channel.

10. Click Finish.

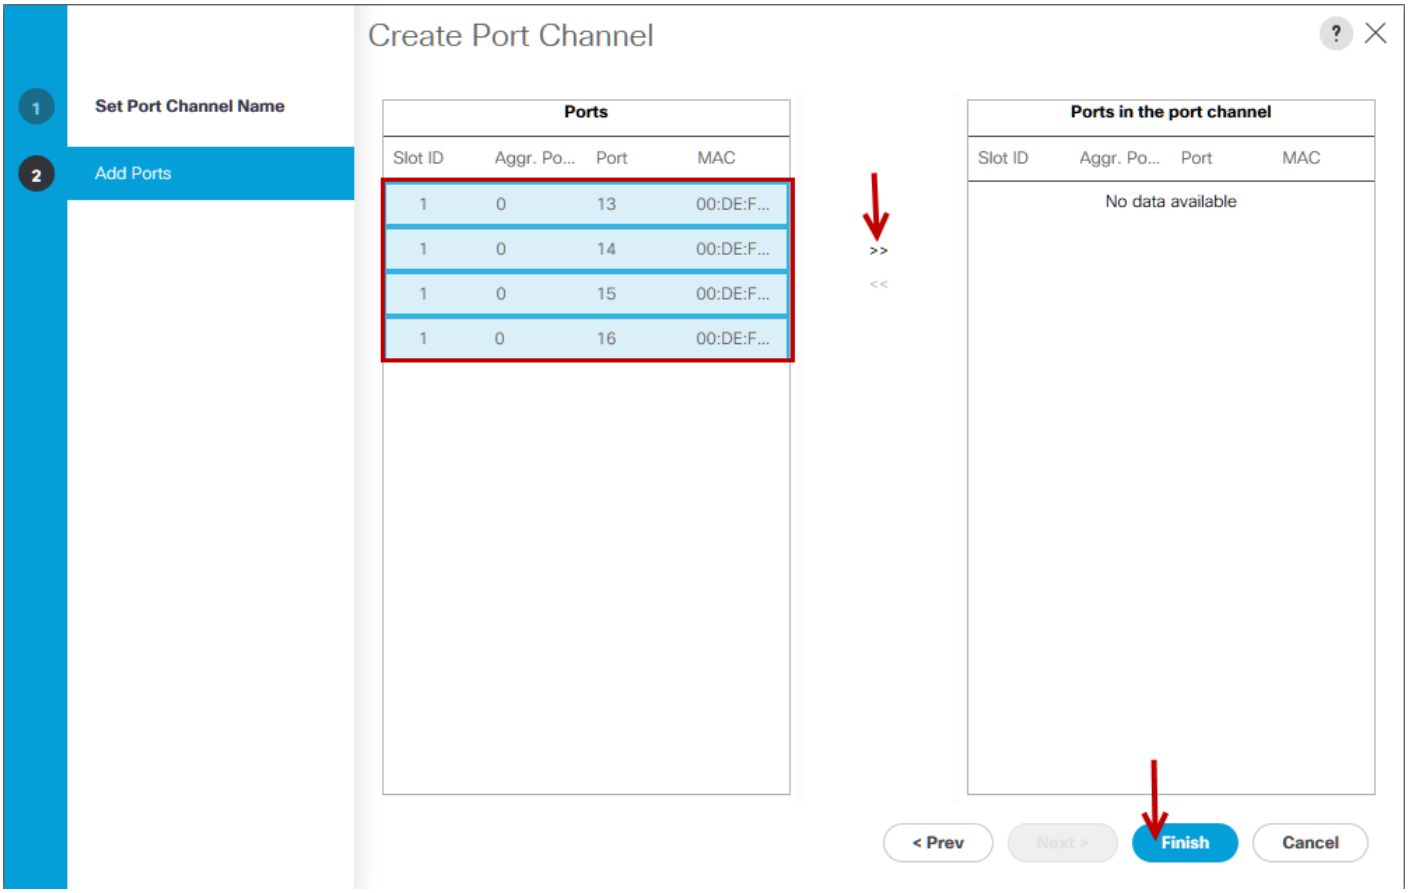

11. Follow the previous steps to create the port-channel on Fabric B, using a different port-channel ID.

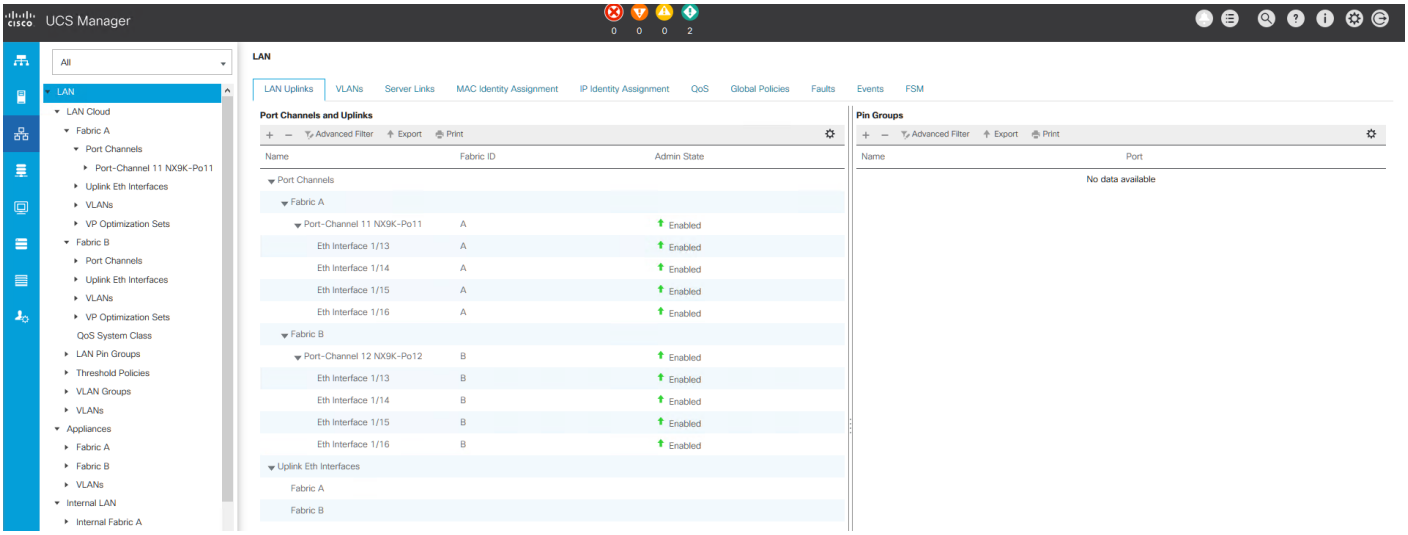

12. Configure QoS System Classes: From the LAN tab, below the Lan Cloud node, select QoS system class and configure the Platinum through Bronze system classes as shown in the following figure.

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

   – Uncheck Enable Packet drop on the Platinum class

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

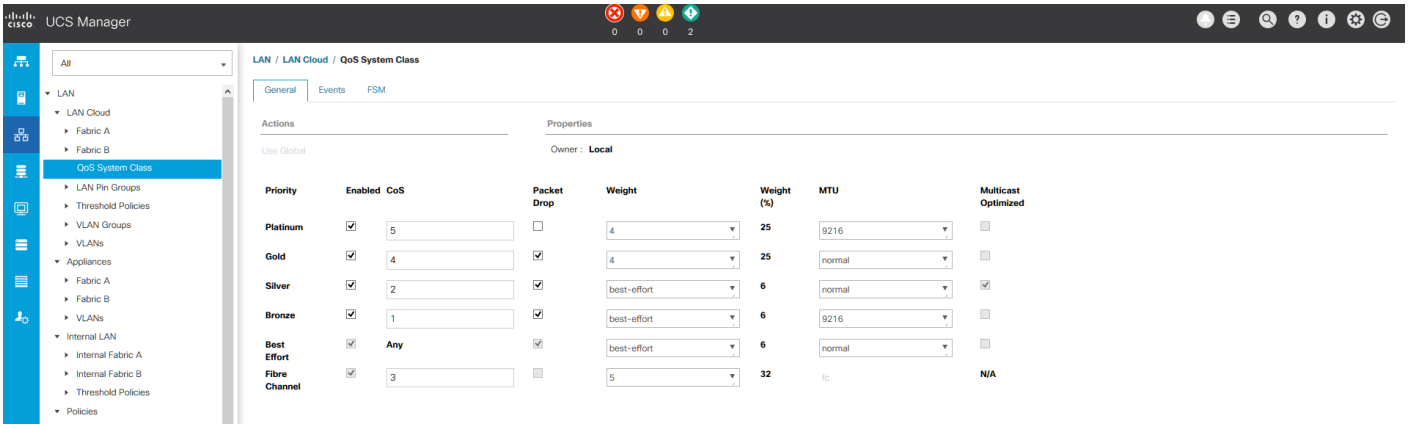

Changing QoS system class configuration on 6300 series Fabric Interconnect requires reboot of FIs.

13. Verify 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 v3.1.2g.

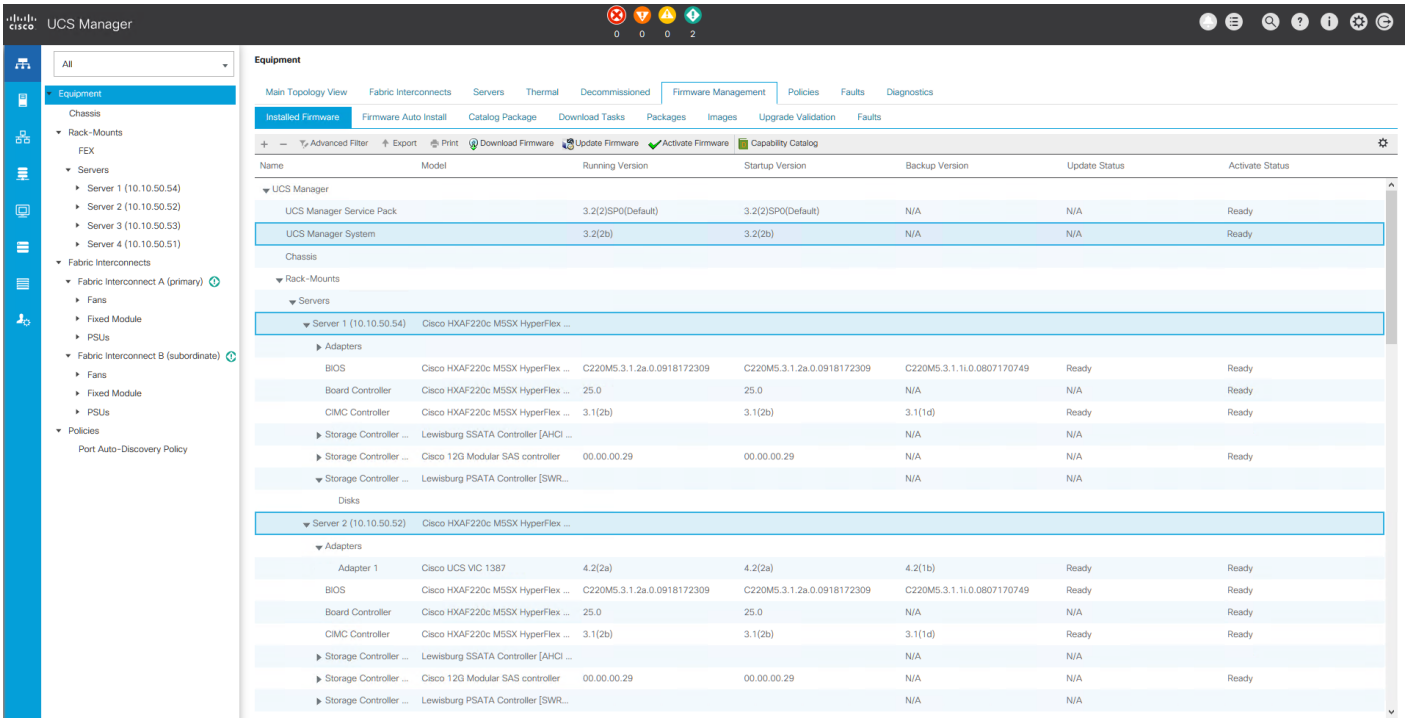

It is recommended to let the HX Installer handle upgrading the server firmware automatically as designed. This will occur once the service profiles are applied to the HX nodes during the automated deployment process.

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

Deploy Cisco HyperFlex Data Platform Installer VM

Download the latest installer OVA from Cisco.com:

https://software.cisco.com/download/release.html?mdfid=286305544&flowid=79522&softwareid=286305994&release=2.1(1c)&relind=AVAILABLE&rellifecycle=&reltype=latest

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

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

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

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

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

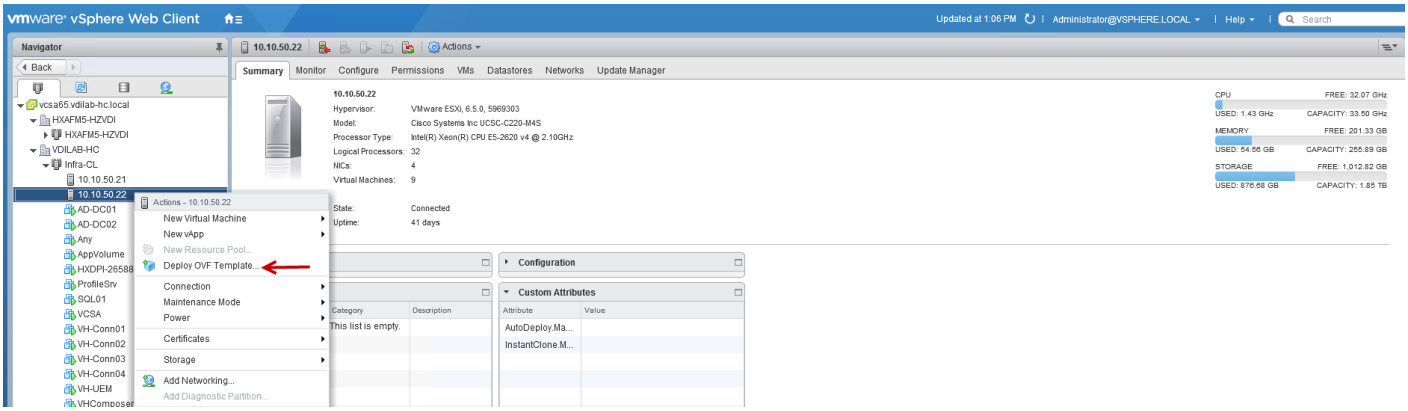

4. Follow the deployment steps to configure HyperFlex data-platform installer VM deployment.

5. Select OVA file to deploy, click Next.

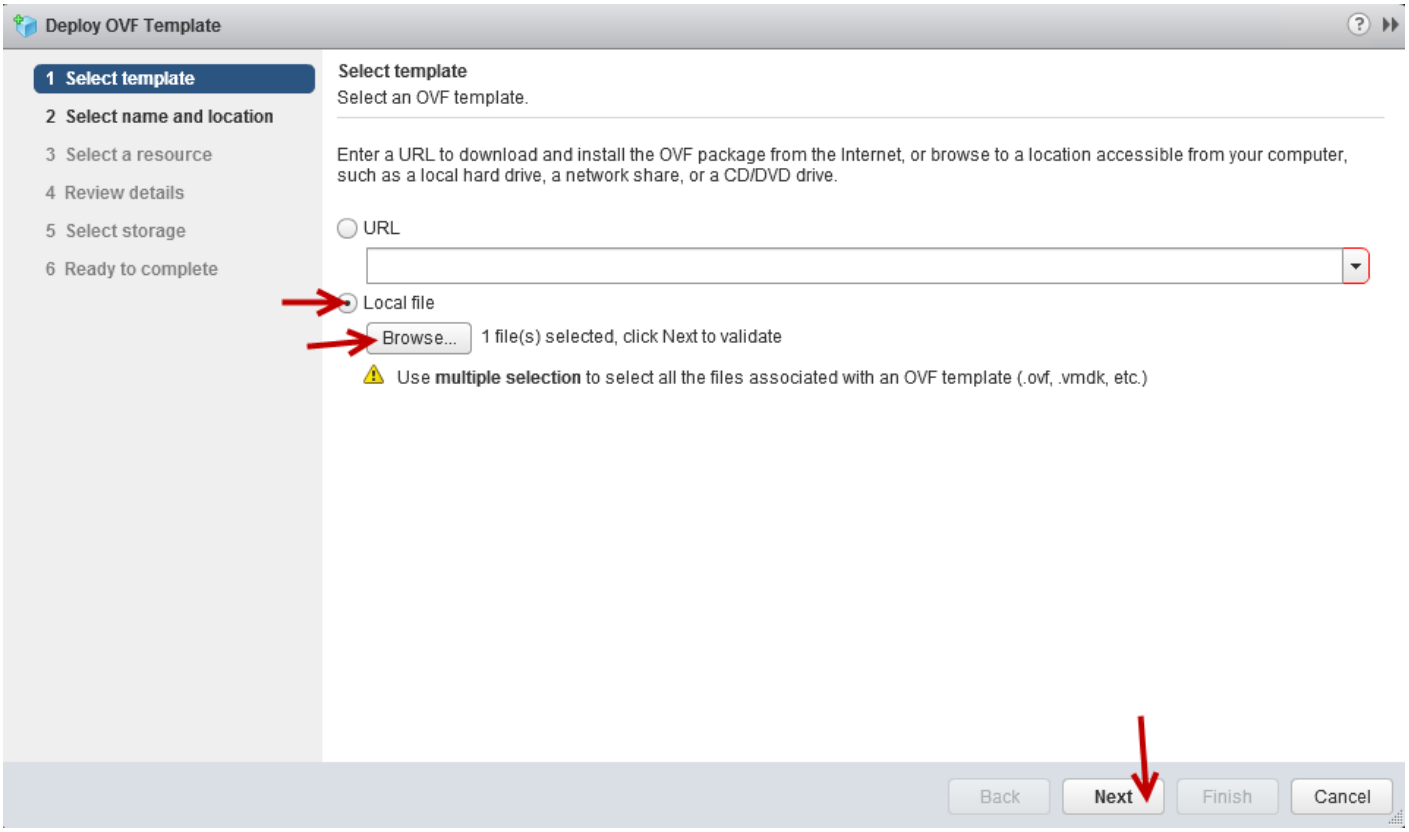

6. Enter name for OVF to template deploy, select datacenter and folder location. Click Next.

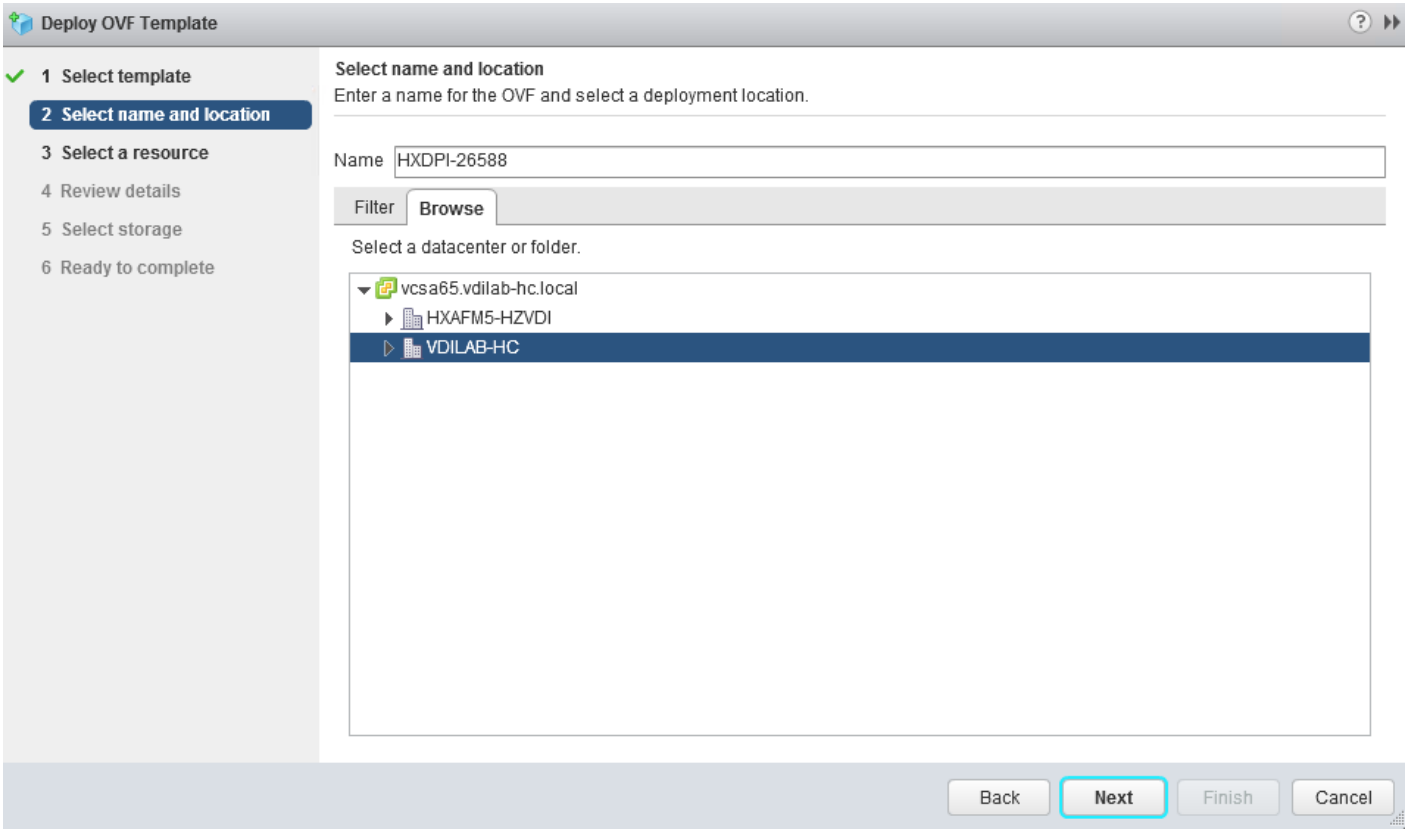

7. Review and verify the details for OVF template to deploy, click Next.

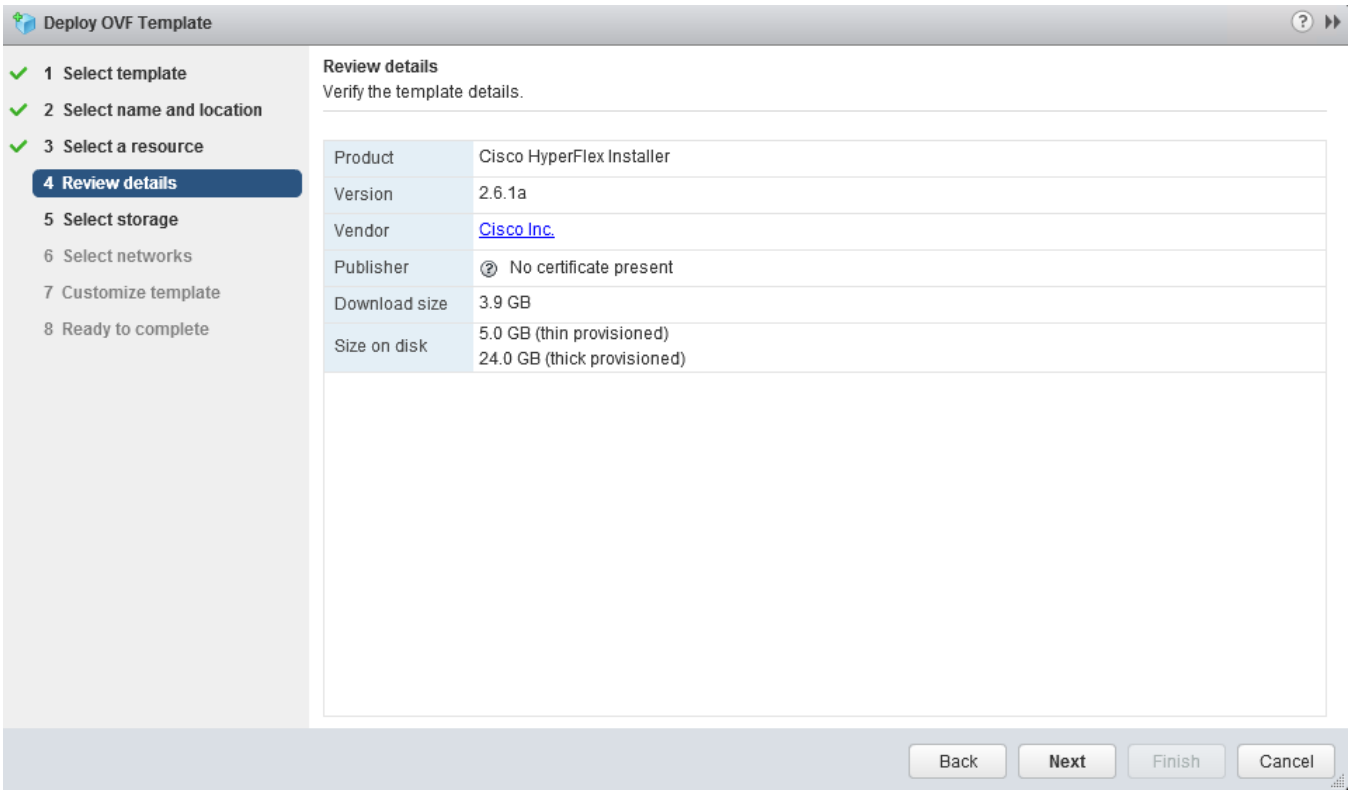

8. Select virtual disk format, VM storage policy set to datastore default, select datastore for OVF deployment. Click Next.

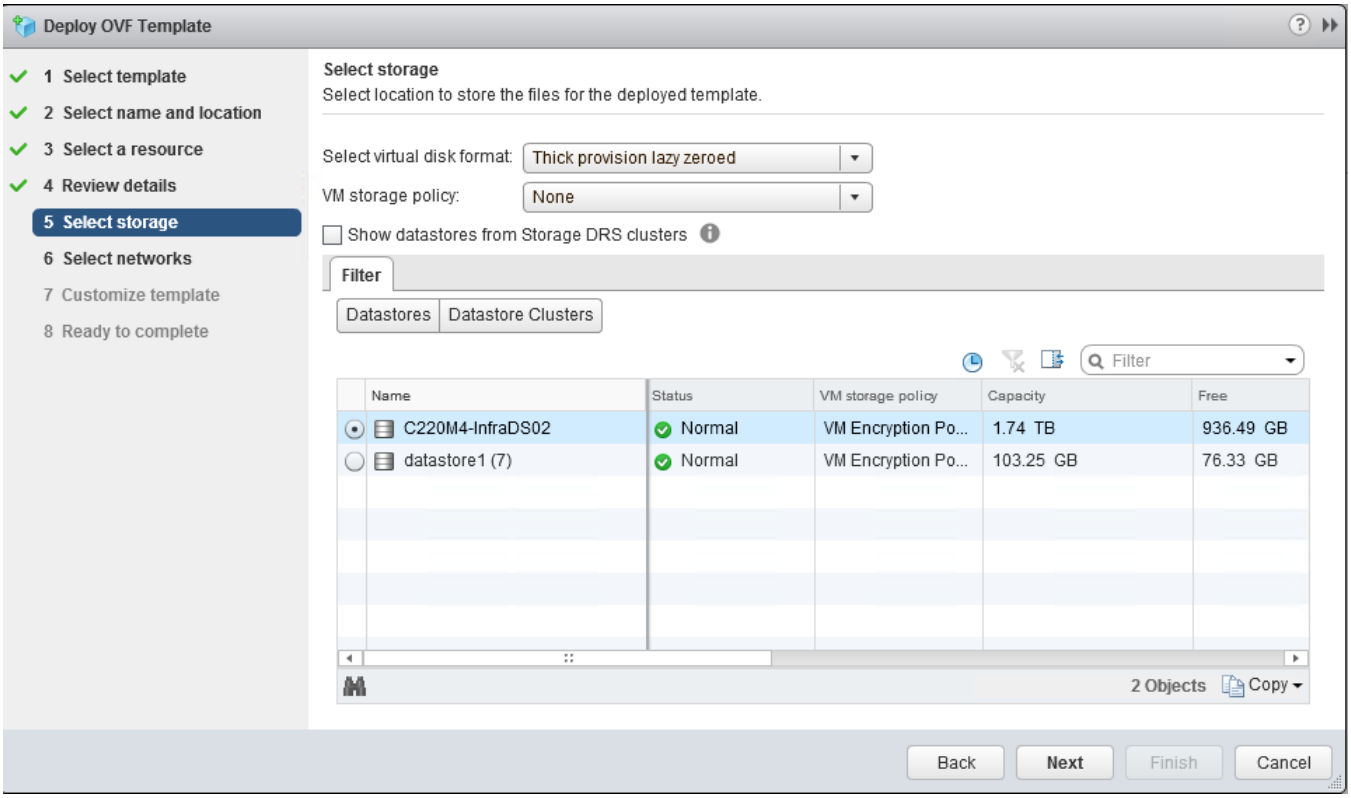

9. Select Network adapter destination port-group.

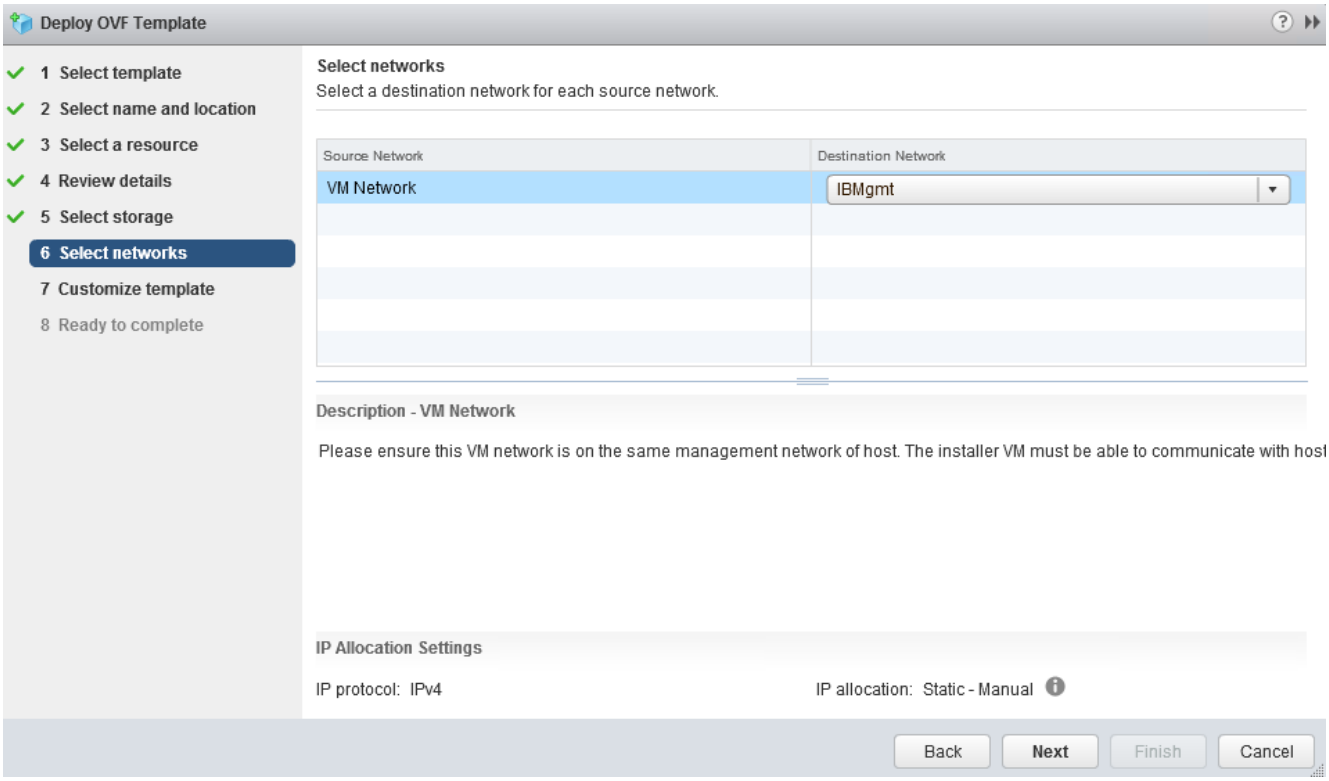

10. Fill out the parameters requested for hostname, gateway, DNS, IP address, and netmask. Alternatively, leave all blank for a DHCP assigned address.

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.

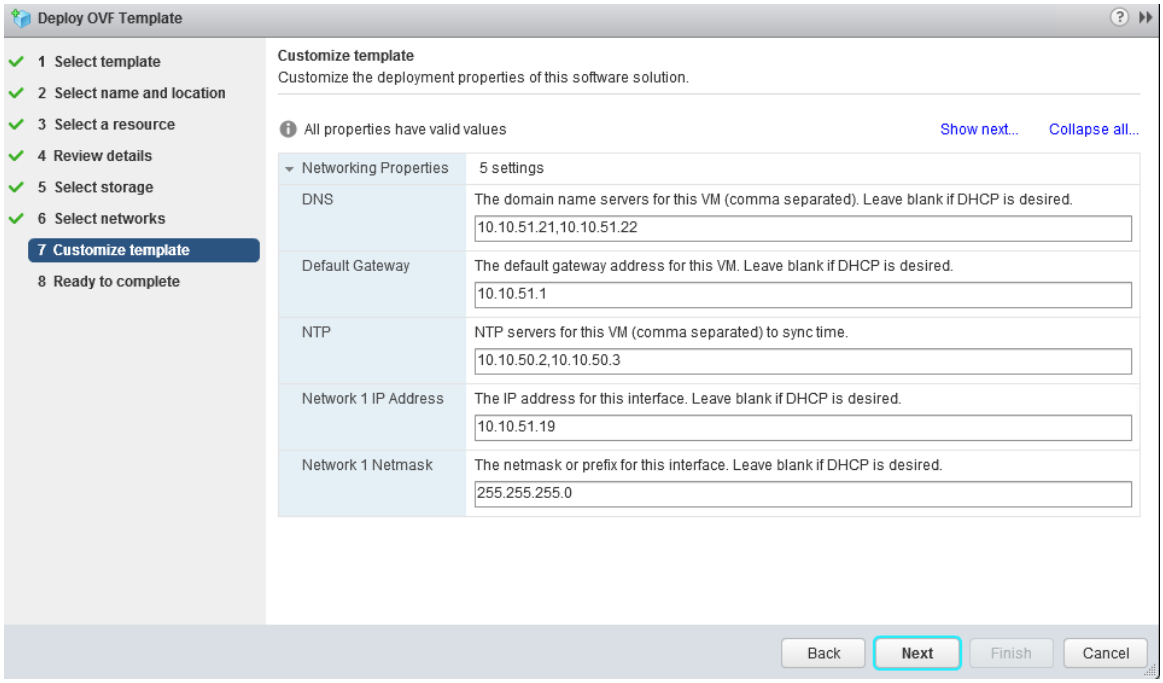

If you have internal firewall rules between these networks, please contact TAC for assistance.

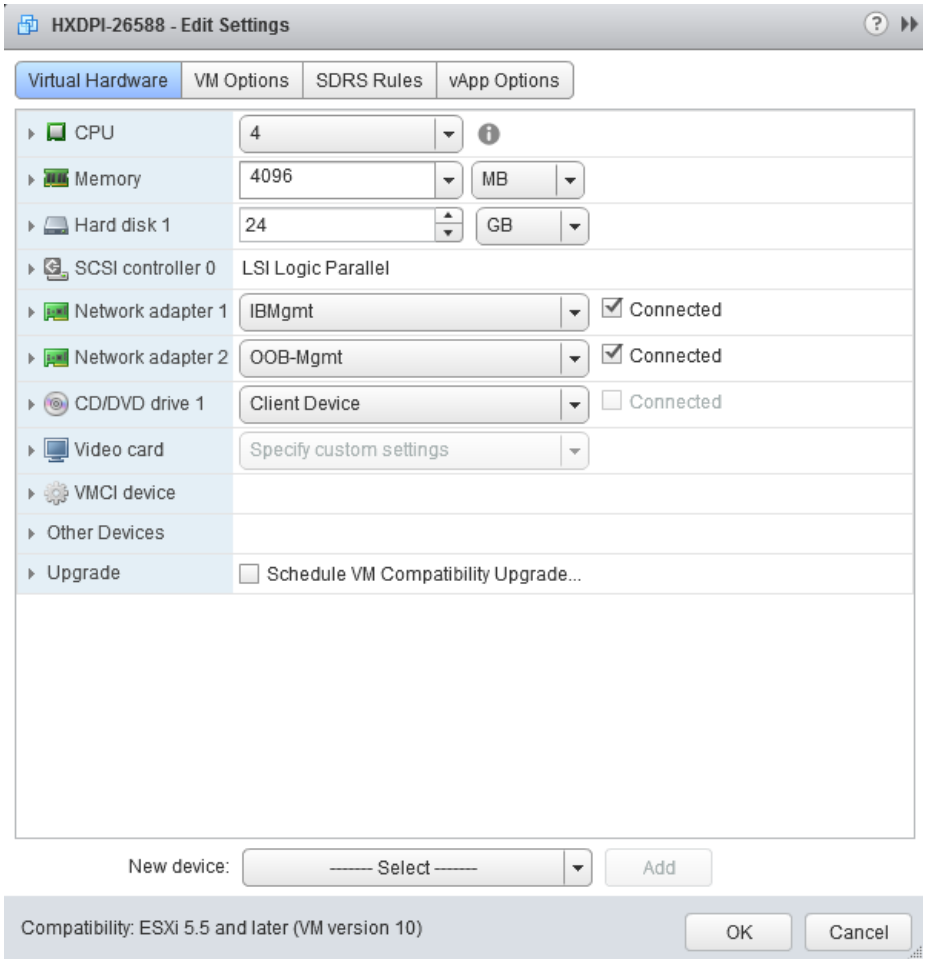

If required, an additional network adapter can be added to the HyperFlex Platform Installer VM after OVF deployment is completed successfully. For example, in case of a separate Inband and Out-Of-Mgmt network, see the screenshot below:

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 10.10.50.19
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 configure the Cisco HyperFlex Cluster, complete the following steps:

1. Login to HX Installer VM through a web browser: http://<Installer_VM_IP_Address>

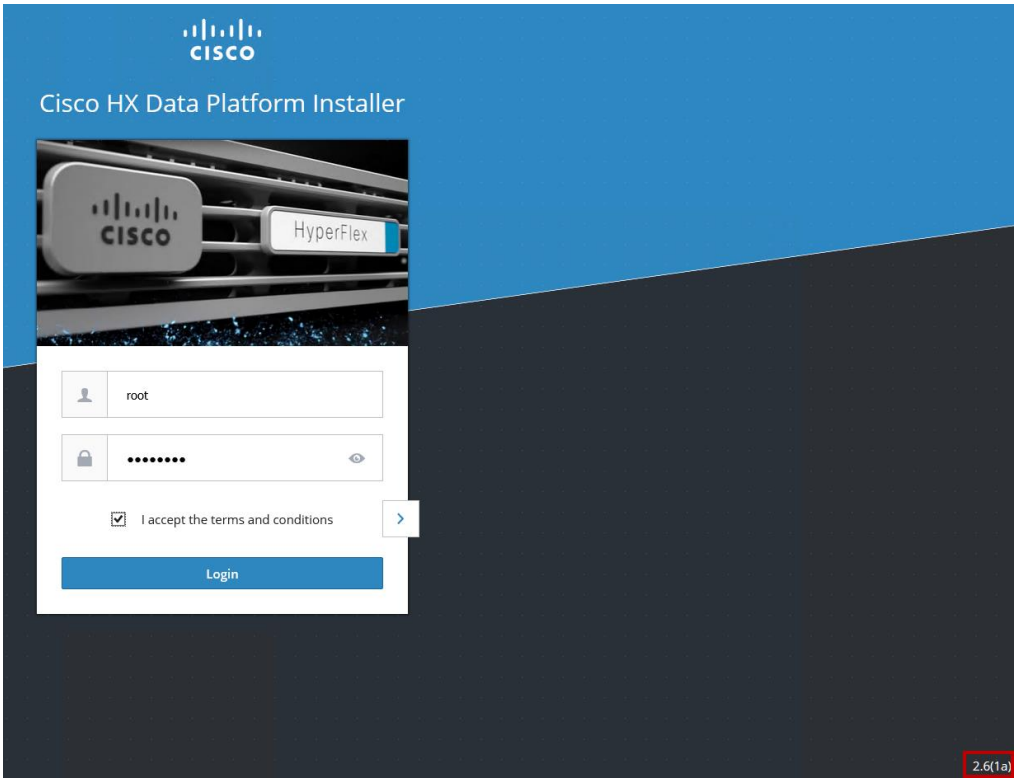

Create a HyperFlex Cluster

To create a HyperFlex Cluster, complete the following steps:

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

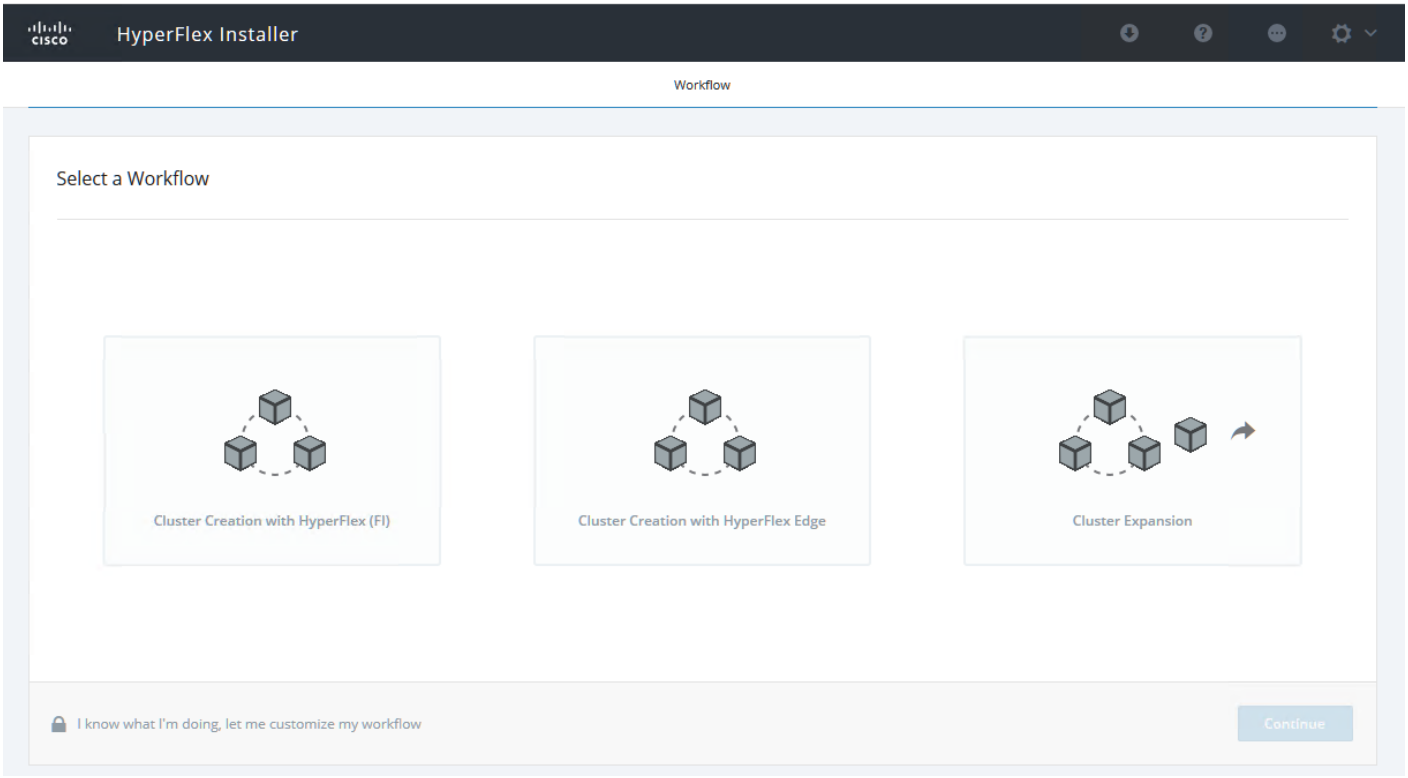

2. On the credentials page, enter the access details for Cisco UCS Manager, vCenter server, and Hypervisor. Click Continue.

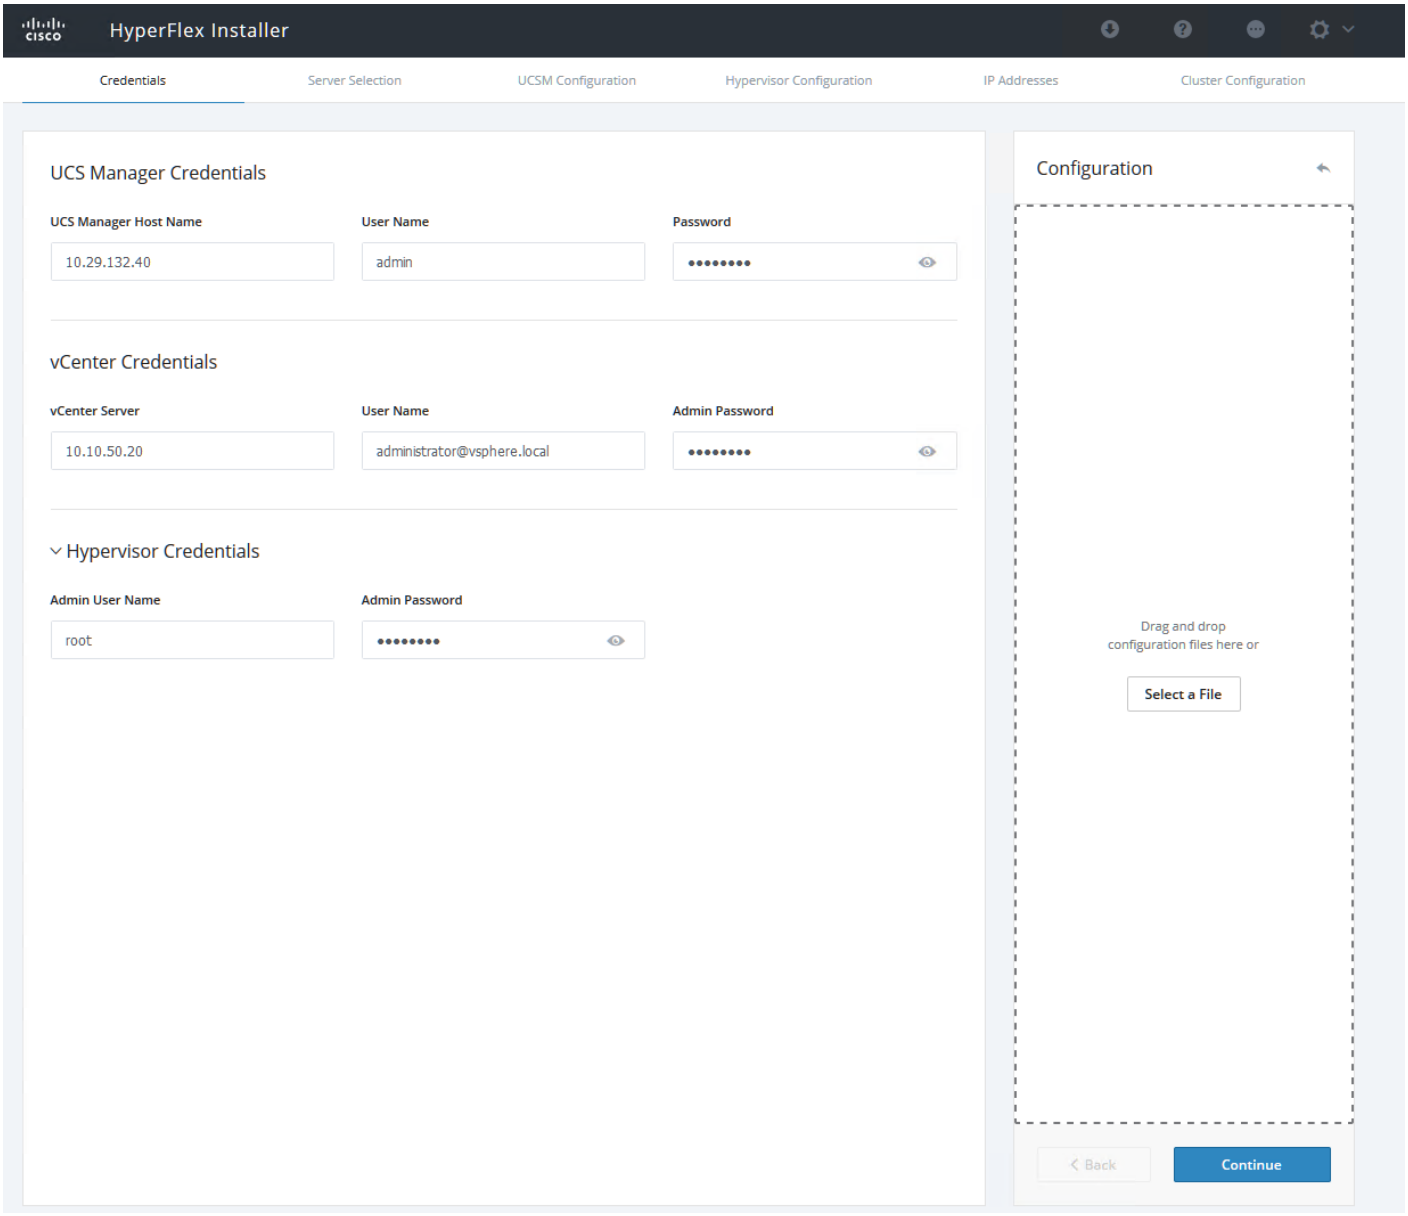

3. Select the top-most check box at the top right corner of the HyperFlex installer to select all unassociated servers. (To configure a subset of available of the HyperFlex servers, manually click the checkbox for individual servers.)

4. Click Continue after completing server selection.

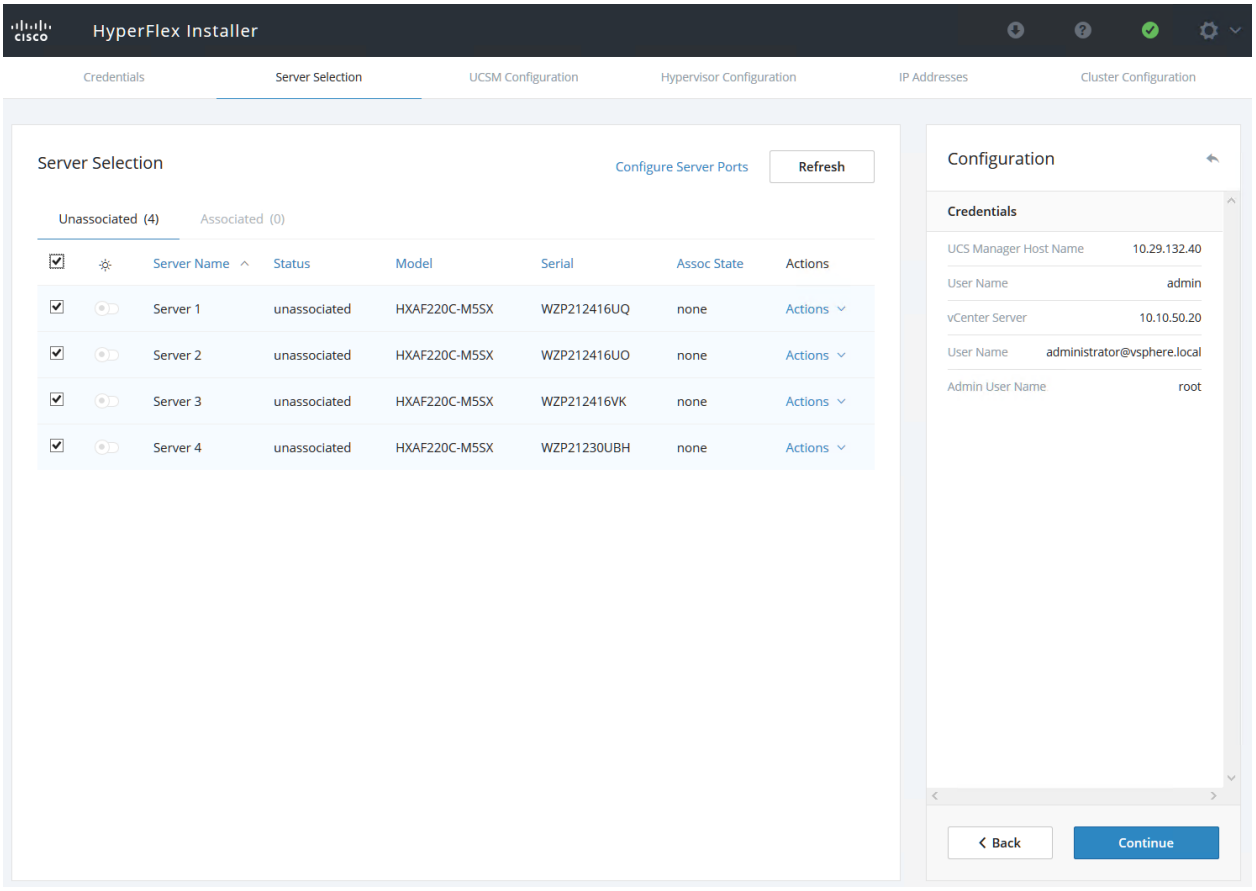

The required server ports can be configured from Installer workflow but it will extend the time to complete server discovery. Therefore, we recommend configuring the server ports and complete HX node discovery in Cisco UCS Manager as described in the Pre-requisites section above prior starting workflow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

Configure Server Ports (Optional)

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

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

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

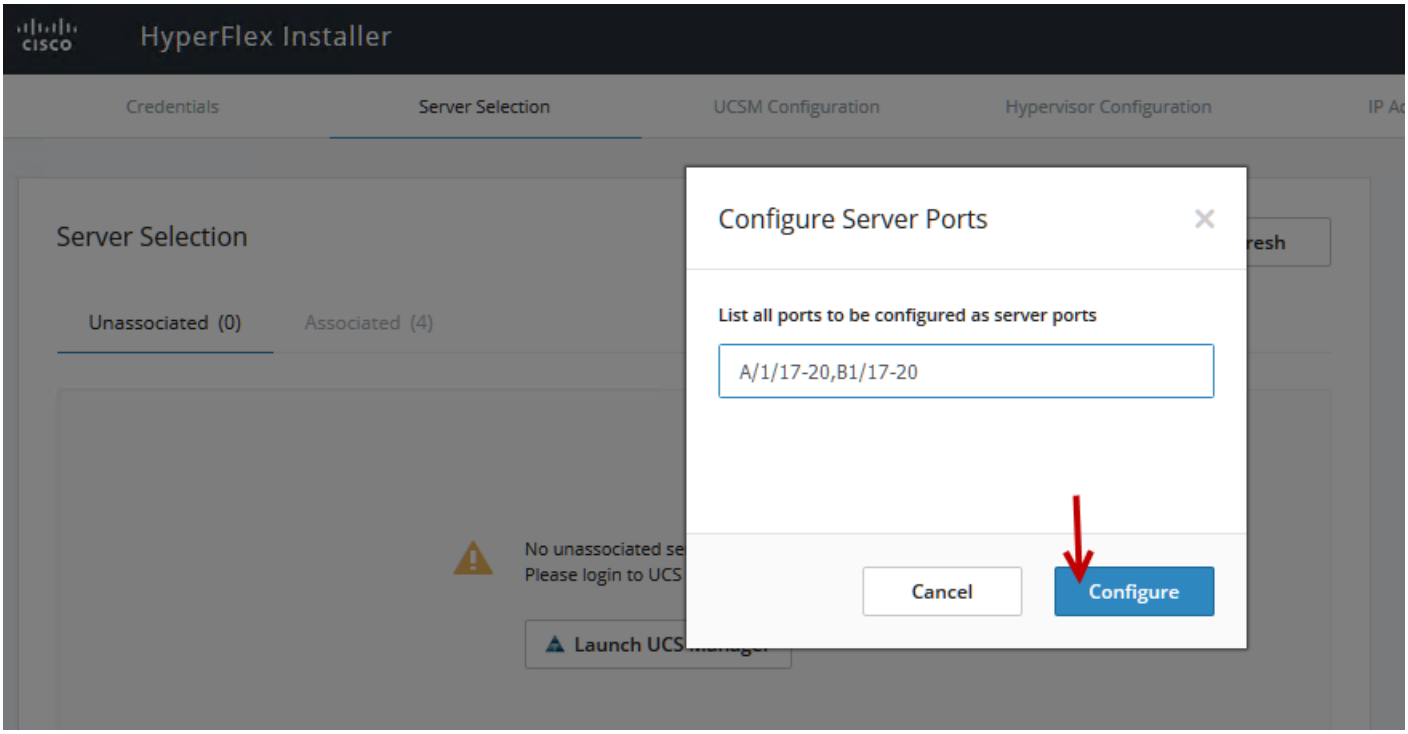

4. Enter the Details for the Cisco UCS Manager Configuration;

   a. 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:

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

      – 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 3.1(2g) which is the required Cisco UCS Manager release for HyperFlex v2.6.1a installation.

7. Enter HyperFlex cluster name.

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

9. Click Continue.

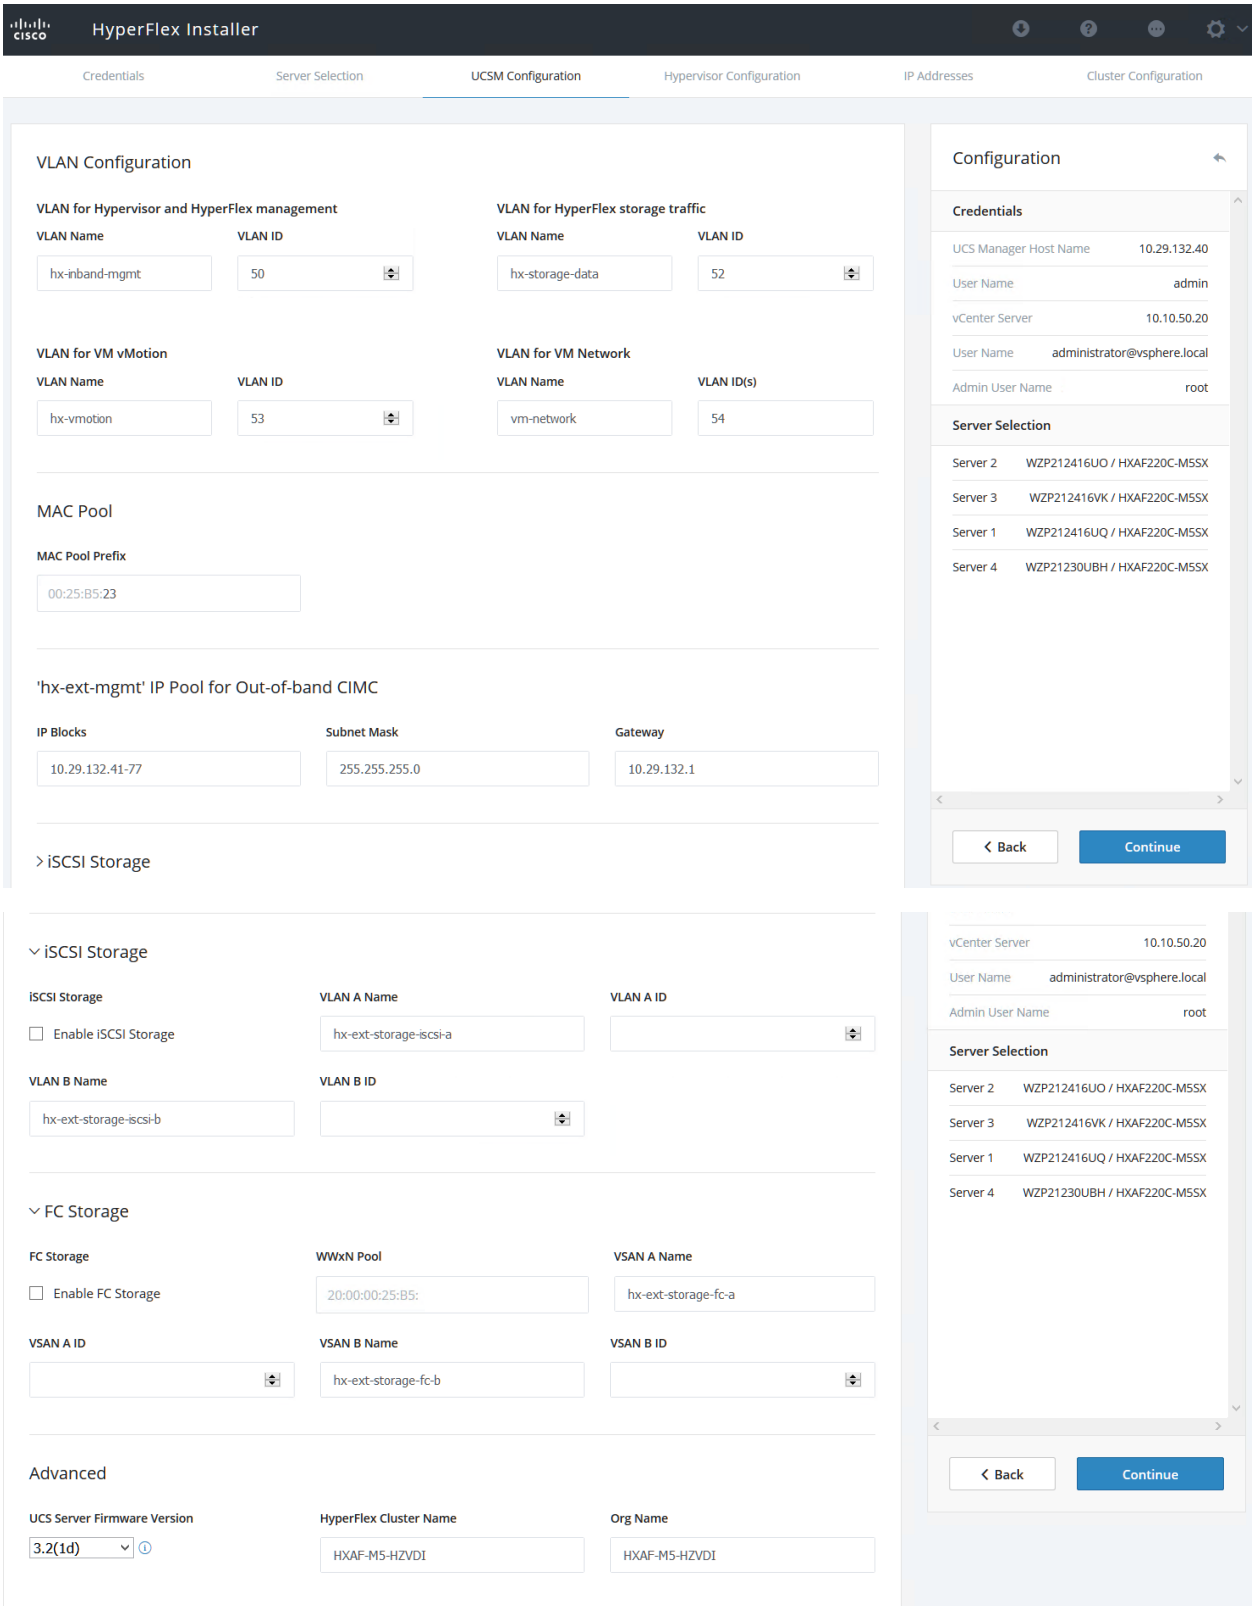

Configure Hypervisor Settings

To configure the Hypervisor settings, complete the following steps:

1. In the Configure common Hypervisor Settings section, enter:
   - Subnet Mask
   - Gateway
   - DNS server(s)
2. In the Hypervisor Settings section:
   - Select check box Make IP Address and Hostnames Sequential if they are following in sequence.
   - Provide the starting IP Address.
   - Provide the starting Host Name or enter Static IP address and Host Names manually for each node
3. Click Continue.

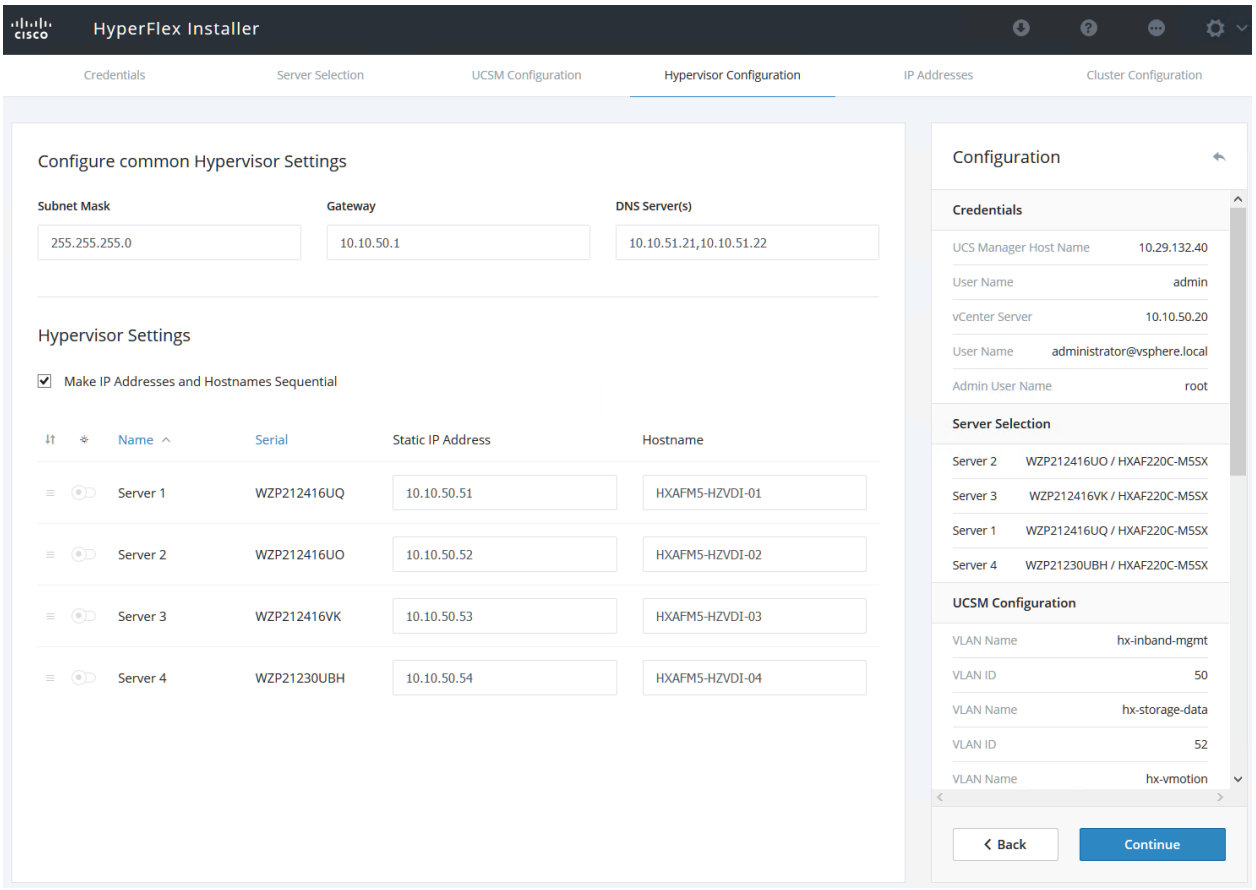

IP Addresses

To add the IP addresses, complete the following steps:

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

Three additional columns appear on this page:

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

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

---

1. On the IP Addresses page, check the box Make IP Addresses Sequential or enter the IP address manually for each node for the following requested values:
   - Storage Controller/Management
   - Hypervisor/Data
   - Storage Controller/Data
2. Enter subnet and gateway details for the Management and Data subnets configured.
3. Click Continue to proceed.

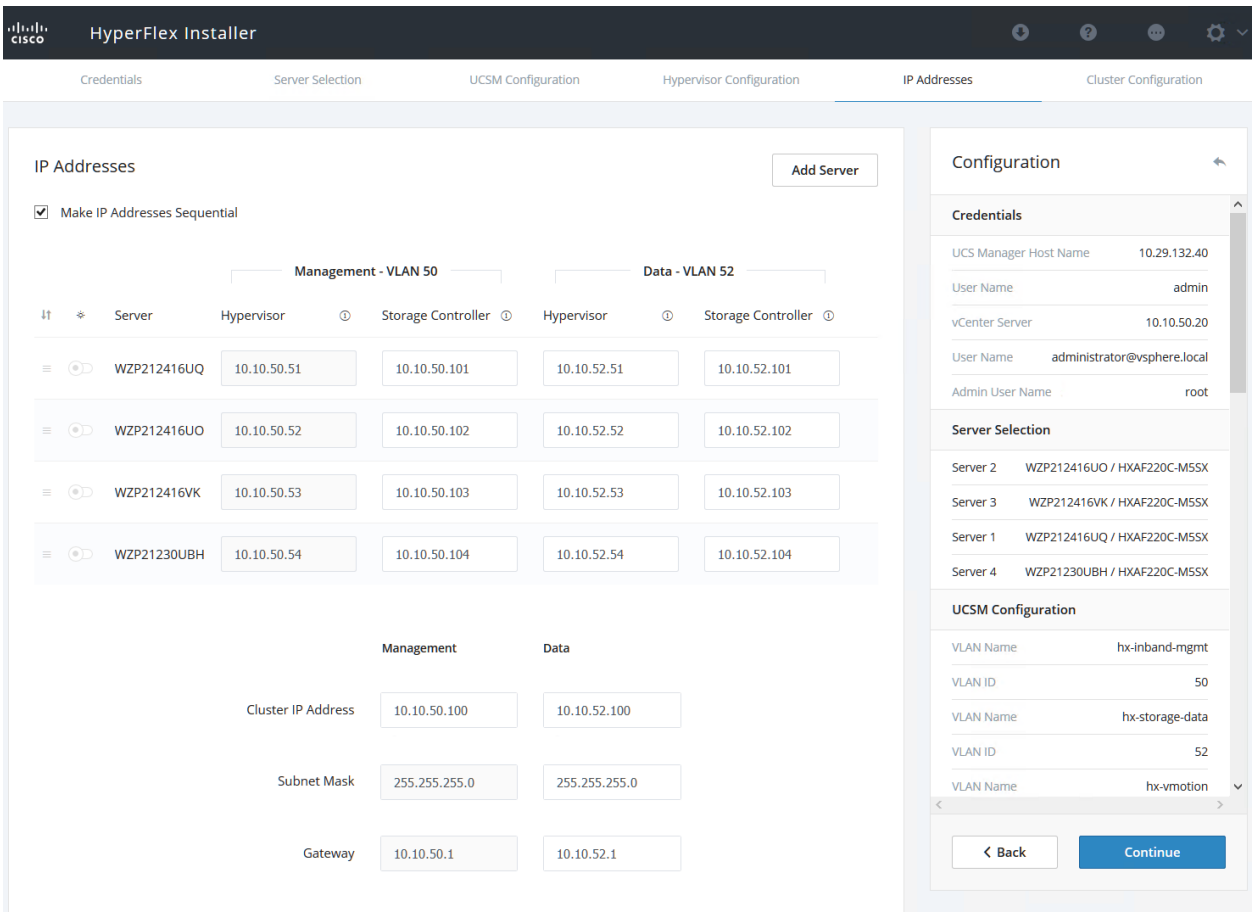

4. On the Cluster Configuration page, enter the following:
   - Cluster Name
   - Cluster management IP address
   - Cluster data IP Address
   - Set Replication Factor: 2 or 3
   - Controller VM password
   - vCenter configuration
     - vCenter Datacenter name
     - vCenter Cluster name
   - System Services
     - DNS Server(s)
     - NTP Server(s)
     - Time Zone
   - Auto Support
     - Click on check box for Enable Auto Support
     - Mail Server
     - Mail Sender
     - ASUP Recipient(s)
   - Advanced Networking
     - Management vSwitch
     - Data vSwitch
   - Advanced Configuration
     - Click on check box to Optimize for VDI only deployment
     - Enable jumbo Frames on Data Network
     - Clean up disk partitions (optional)
   - vCenter Single-Sign-On server

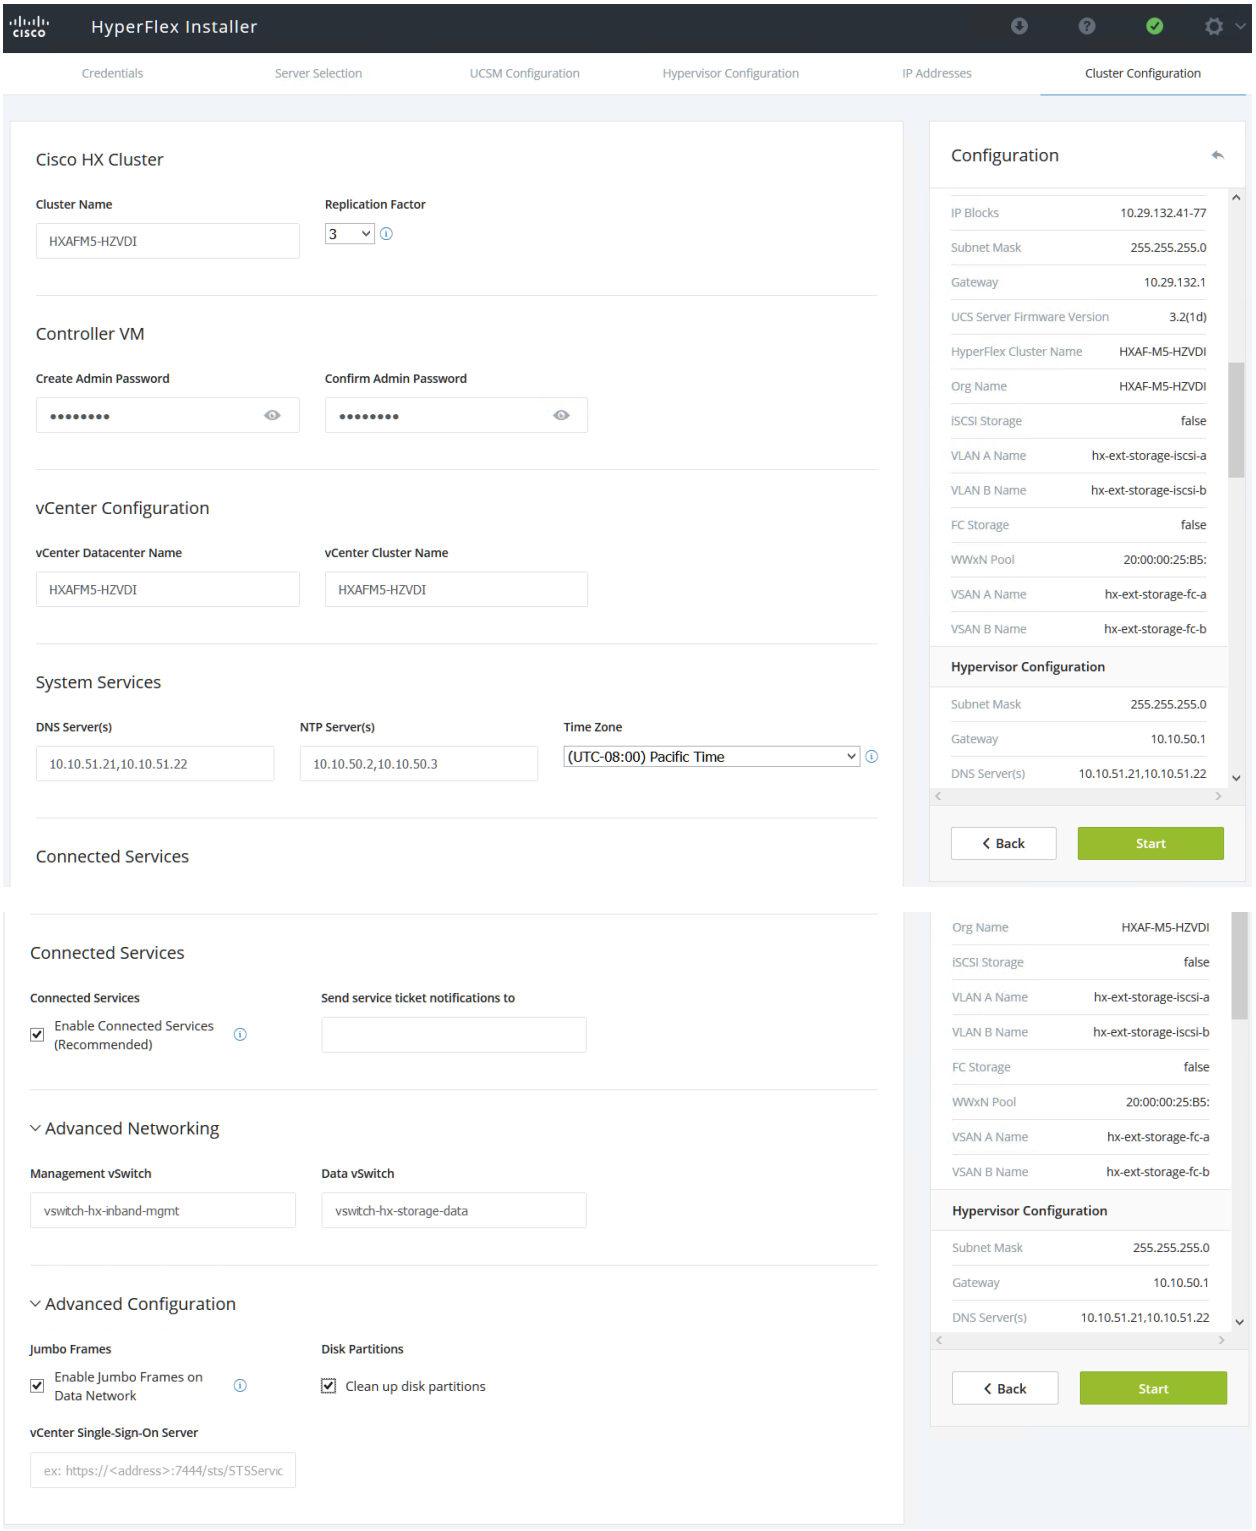

5. The configuration details can be exported to a JSON file by clicking the down arrow icon in the top right corner of the Web browser page as shown in the screenshot below.

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.

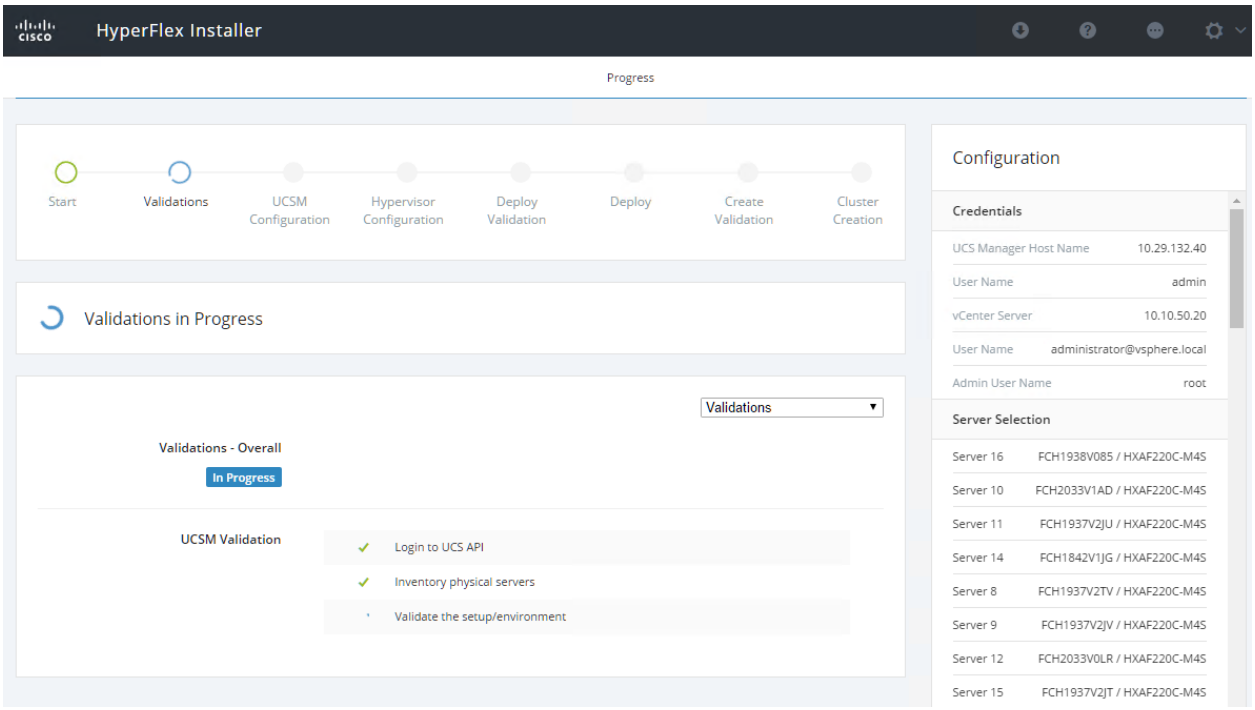

If QoS system class is not defined as per the requirement HyperFlex installer will go ahead and make required changes. There will be a warning generated accordingly in HyperFlex Installer workflow. For 6300 series Fabric Interconnect change in QoS system class requires reboot of FIs.

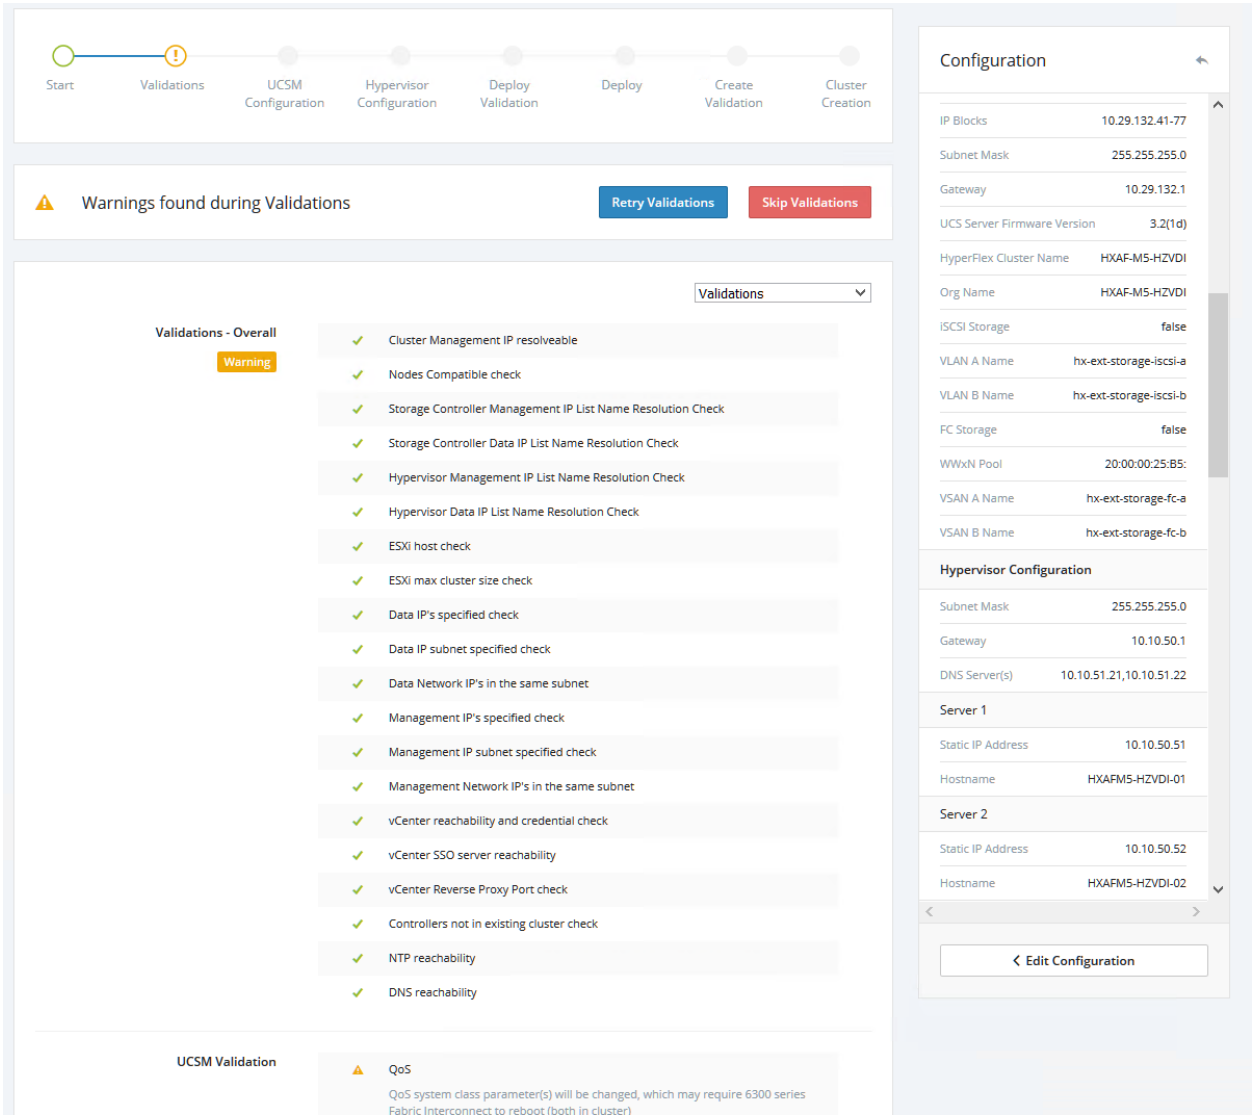

8. After a successful validation, the workflow continues with the Cisco UCS Manager configuration.

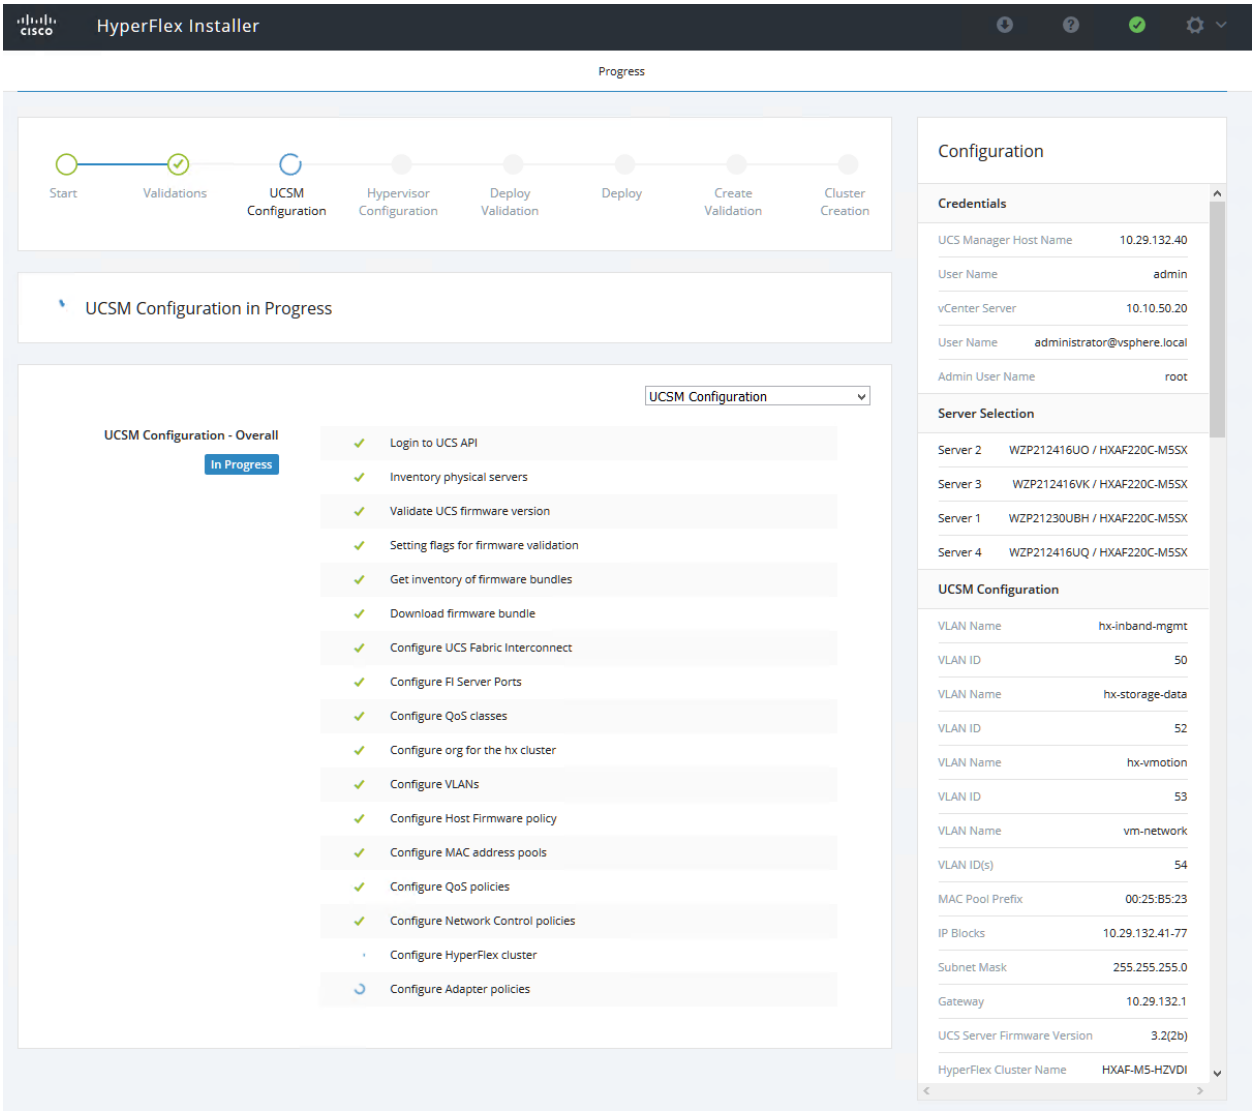

9. After a successful Cisco UCS Manager configuration, the installer proceeds with the Hypervisor configuration.

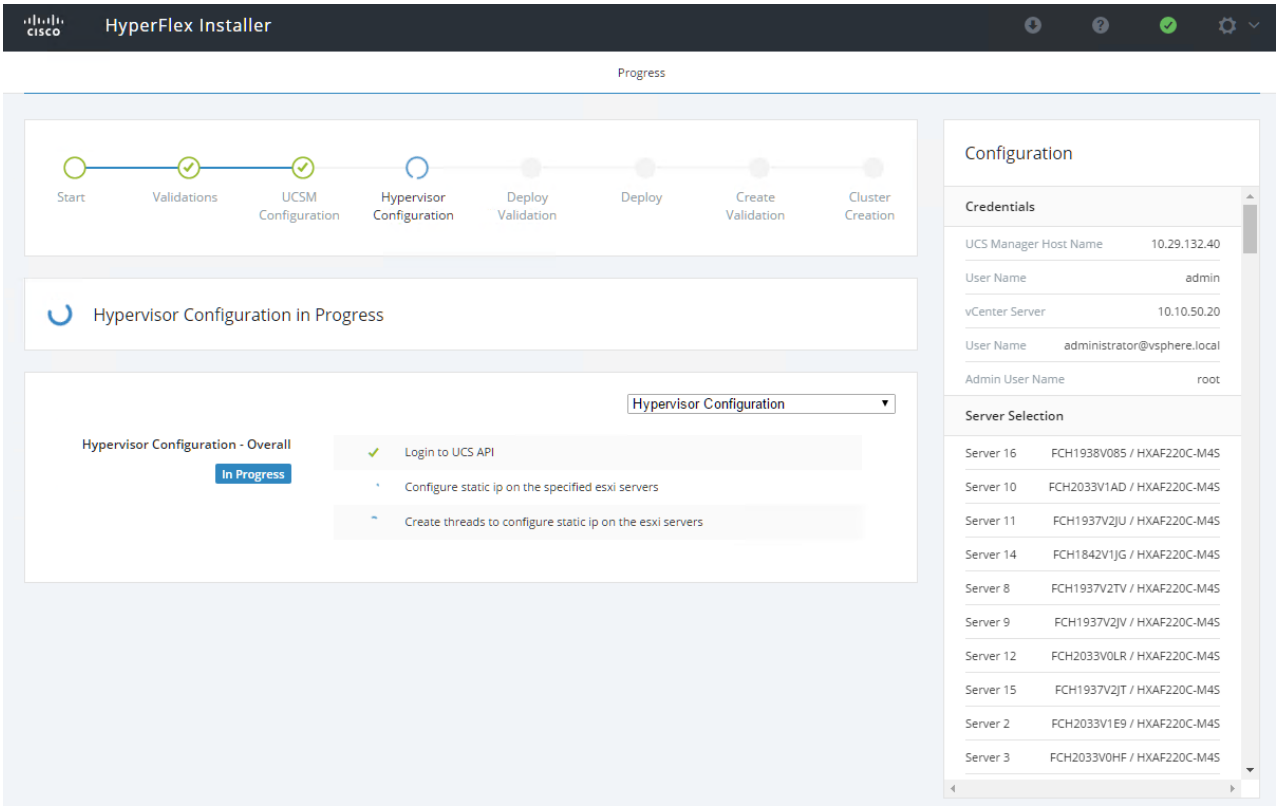

10. After a successful Hypervisor configuration, deploy validation task is performed which checks for required compoenent and accessibilty prior Deploy task is performed on Storage Controller VM.

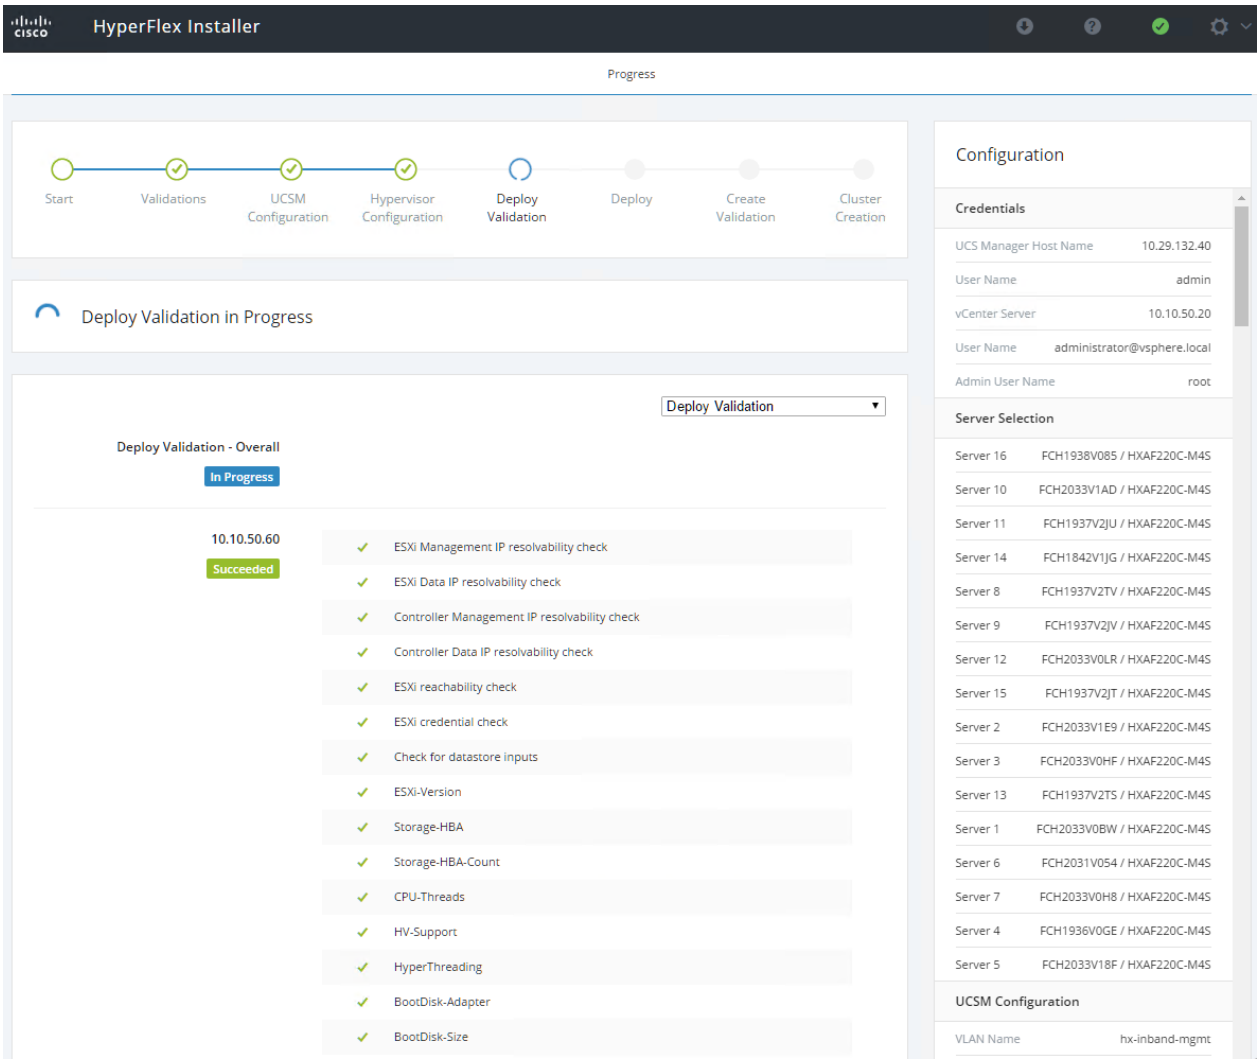

11. Installer performs deployment task after successfully validating Hypervisor configuration.

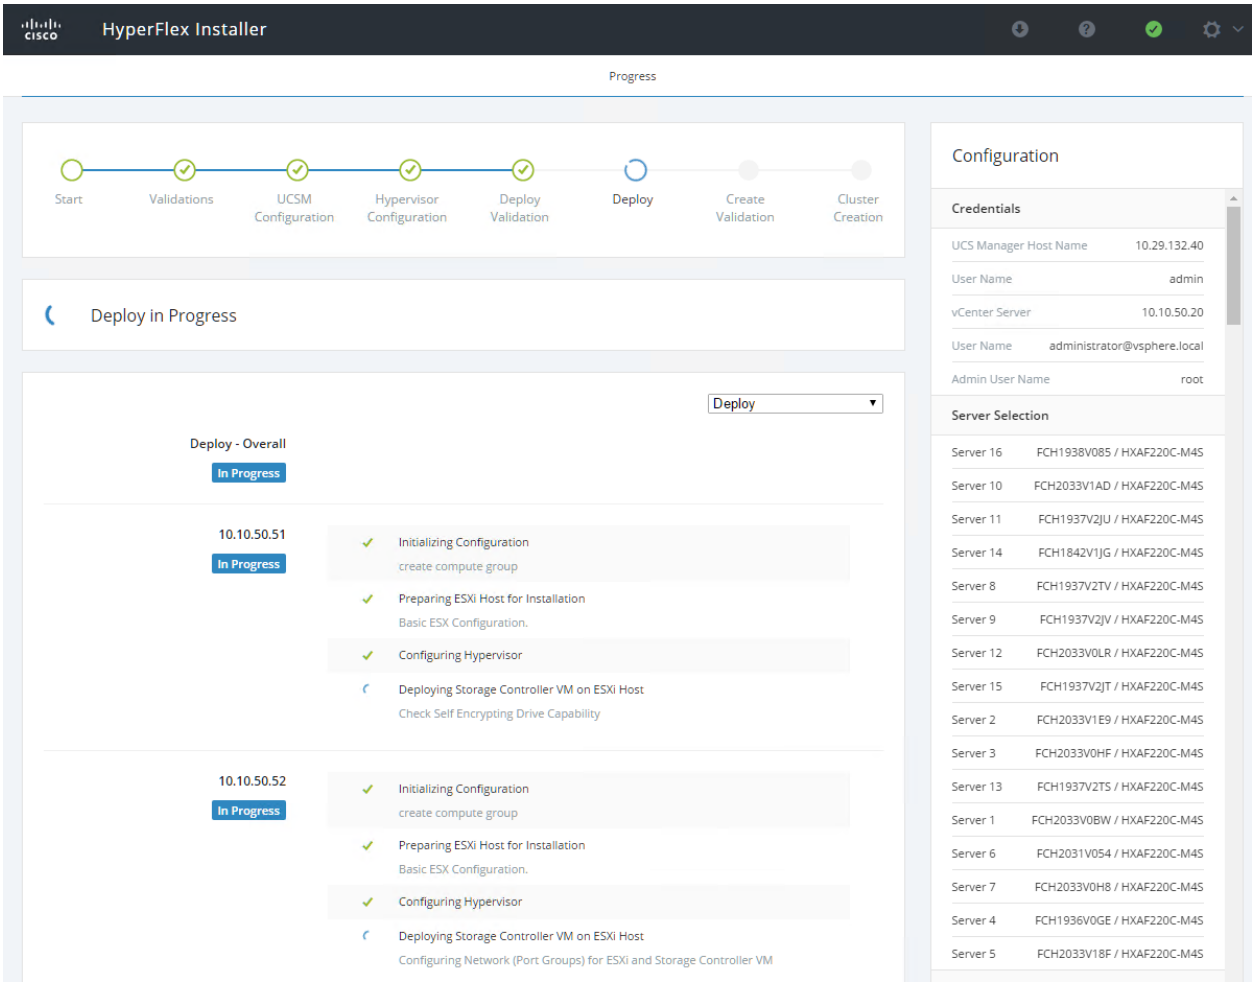

12. After a successful deployment of the ESXi hosts configuration, the Controller VM software components for HyperFlex installer checks for validation prior to creating the cluster.

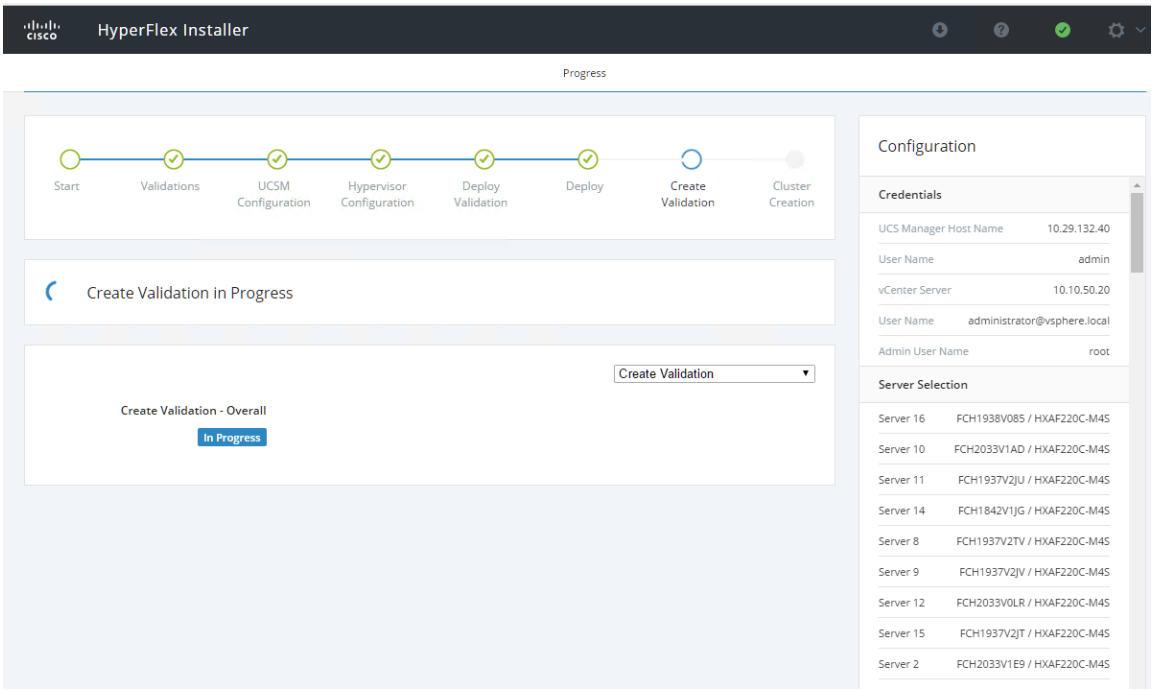

13. After a successful validation, the installer creates and starts the HyperFlex cluster service.

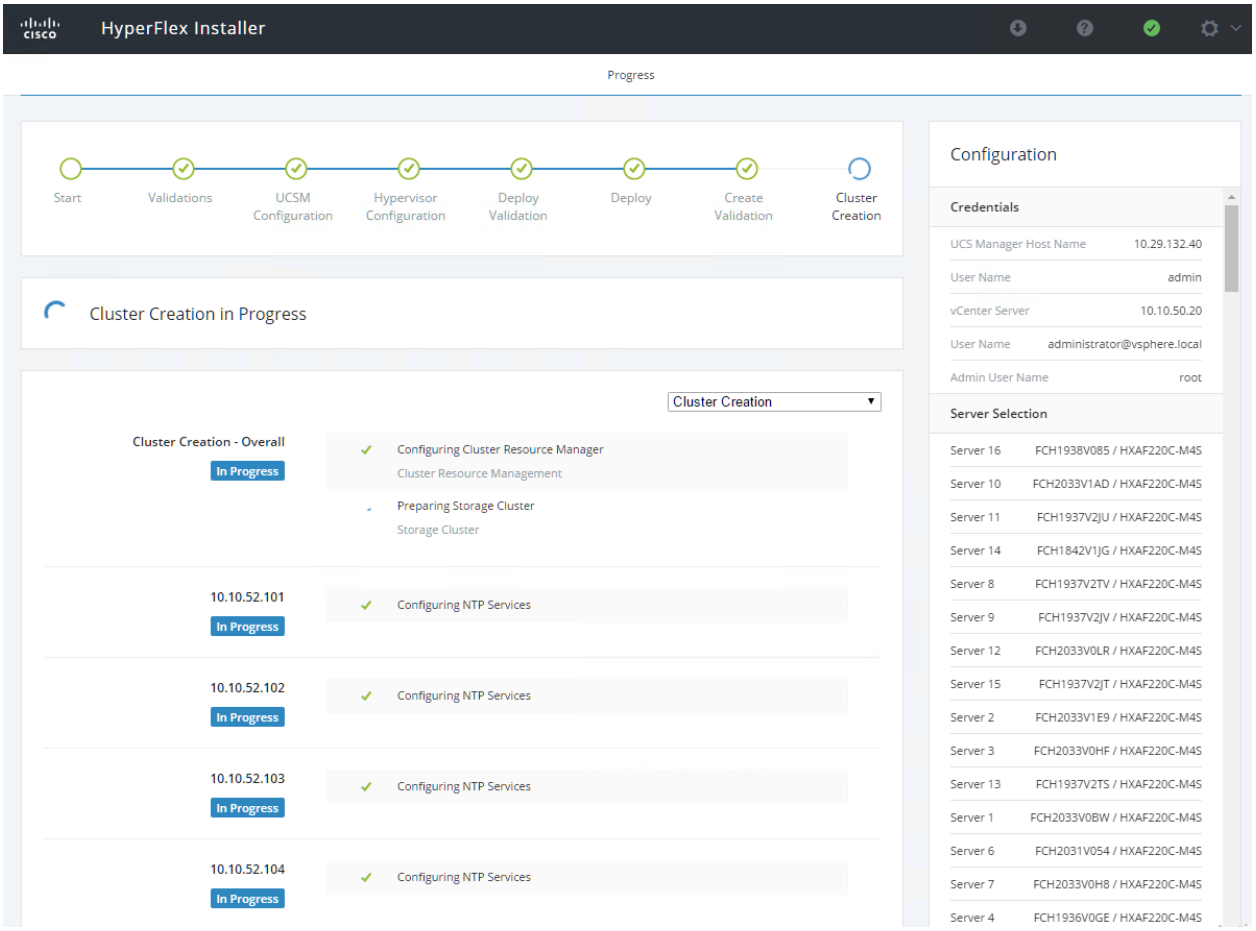

14. After a successful HyperFlex Installer VM workflow completion, the installer GUI provides a summary of the cluster that has been created.

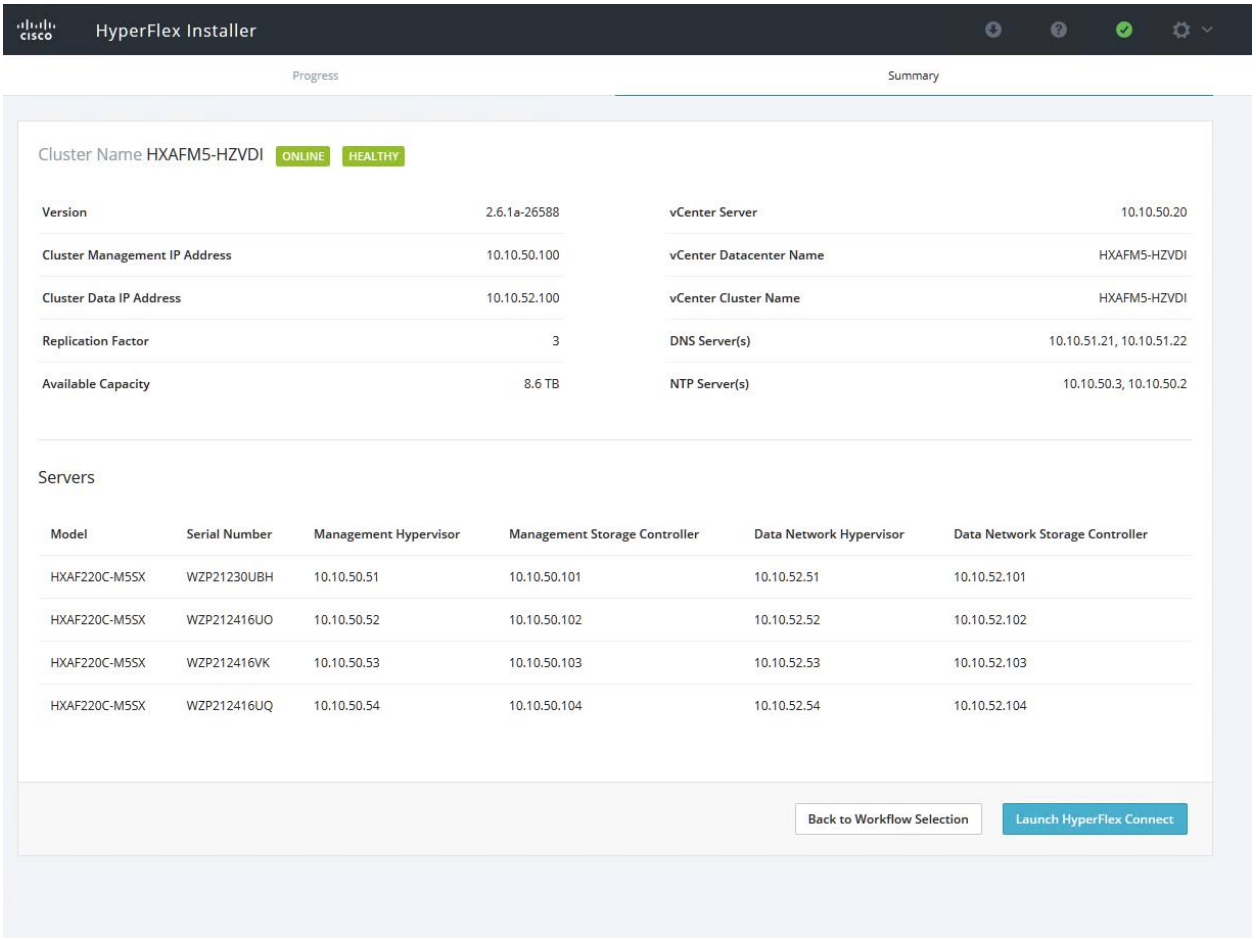

15. Click Launch vSphere Web Client.

Cisco HyperFlex installer creates and configured 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 46.

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

Figure 46 Cisco UCS Node Naming Convention

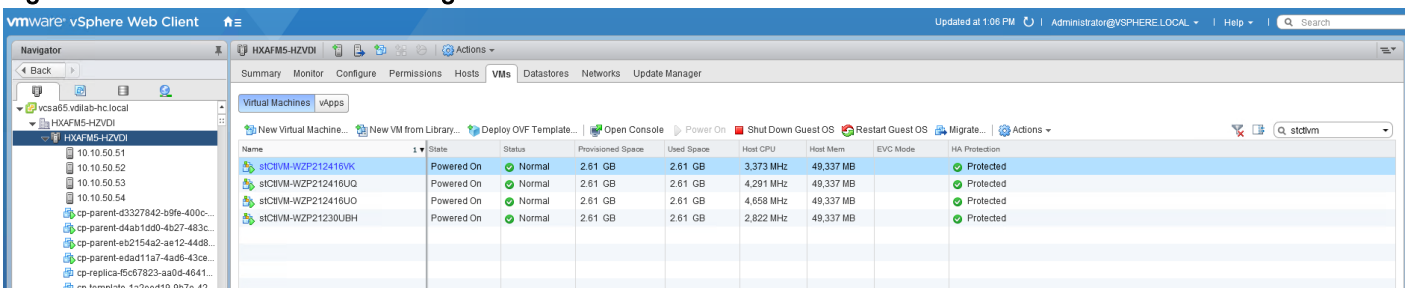

Run Cluster Post Installation Script

After a successful installation of HyperFlex cluster, run the post_install script by loging 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, use your tool of choice to make a secure connection to the Cisco HyperFlex Data Platform installer using it's IP address and port 22.

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

3. When authenticated, enter post_install at the command prompt, then press Enter.

4. Provide a valid vCenter administrator user name and password and the vCenter url IP address.

5. Type y for yes to each of the promts that follow except Add VM network VLANs? (y/n) where you can choose whether or not to send health status data 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 10.10.50.51: 10.10.53.51
 Adding vmotion-53 to 10.10.50.51
 Adding vmkernel to 10.10.50.51
 vMotion IP for 10.10.50.52: 10.10.53.52
 Adding vmotion-53 to 10.10.50.52
 Adding vmkernel to 10.10.50.52
 vMotion IP for 10.10.50.53: 10.10.53.53
 Adding vmotion-53 to 10.10.50.53
 Adding vmkernel to 10.10.50.53
 vMotion IP for 10.10.50.54: 10.10.53.54
 Adding vmotion-53 to 10.10.50.54
 Adding vmkernel to 10.10.50.54
Add VM network VLANs? (y/n) n
Send test email? (y/n) n
Validating cluster health and configuration...
 Found UCSM 10.29.132.40, logging with username admin.  Org is HXAF-M5-HZVDI
 UCSM Password:
Could not connect to UCSM at 10.29.132.40 - coercing to Unicode: need string
or buffer, NoneType found.  Skipping UCSM check
Checking MTU settings
 Pinging 169.254.254.2 from vmk1
 Pinging 10.10.50.52 from vmk0
 Pinging 10.10.50.51 from vmk0
 Pinging 10.10.50.53 from vmk0
 Pinging 10.10.50.54 from vmk0
 Setting vmnic1 to active and vmnic0 to standby
 Pinging 10.10.50.52 from vmk0
 Pinging 10.10.50.51 from vmk0
 Pinging 10.10.50.53 from vmk0
 Pinging 10.10.50.54 from vmk0
```

```
 Setting vmnic0 to active and vmnic1 to standby


Network Summary:
 Host: 10.10.50.51
    vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
       vmnic0 - 1 - AAK23-VDIHZ-A - active
       vmnic1 - 1 - AAK23-VDIHZ-B - standby
         Portgroup Name  -  VLAN
         VM Network - 0
         Storage Controller Management Network - 50
         Storage Controller Replication Network - 0
         Management Network - 50
    vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
       vmnic4 - 1 - AAK23-VDIHZ-A - active
       vmnic5 - 1 - AAK23-VDIHZ-B - active
         Portgroup Name  -  VLAN
         vm-network-54 - 54
    vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
       vmnic6 - 1 - AAK23-VDIHZ-A - active
       vmnic7 - 1 - AAK23-VDIHZ-B - standby
         Portgroup Name  -  VLAN
         vmotion-53 - 53
    vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
       vmnic2 - 1 - AAK23-VDIHZ-A - standby
       vmnic3 - 1 - AAK23-VDIHZ-B - active
         Portgroup Name  -  VLAN
         Storage Controller Data Network - 52
         Storage Hypervisor Data Network - 52


 Host: 10.10.50.52
    vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
       vmnic0 - 1 - AAK23-VDIHZ-A - active
```

```
      vmnic1 - 1 - AAK23-VDIHZ-B - standby
        Portgroup Name  -  VLAN
        VM Network - 0
        Storage Controller Management Network - 50
        Storage Controller Replication Network - 0
        Management Network - 50
   vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
      vmnic4 - 1 - AAK23-VDIHZ-A - active
      vmnic5 - 1 - AAK23-VDIHZ-B - active
        Portgroup Name  -  VLAN
        vm-network-54 - 54
   vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
      vmnic6 - 1 - AAK23-VDIHZ-A - active
      vmnic7 - 1 - AAK23-VDIHZ-B - standby
        Portgroup Name  -  VLAN
        vmotion-53 - 53
   vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
      vmnic2 - 1 - AAK23-VDIHZ-A - standby
      vmnic3 - 1 - AAK23-VDIHZ-B - active
        Portgroup Name  -  VLAN
        Storage Controller Data Network - 52
        Storage Hypervisor Data Network - 52

Host: 10.10.50.53
   vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
      vmnic0 - 1 - AAK23-VDIHZ-A - active
      vmnic1 - 1 - AAK23-VDIHZ-B - standby
        Portgroup Name  -  VLAN
        VM Network - 0
        Storage Controller Management Network - 50
        Storage Controller Replication Network - 0
        Management Network - 50
```

```
    vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
       vmnic4 - 1 - AAK23-VDIHZ-A - active
       vmnic5 - 1 - AAK23-VDIHZ-B - active
         Portgroup Name  -  VLAN
         vm-network-54 - 54
    vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
       vmnic6 - 1 - AAK23-VDIHZ-A - active
       vmnic7 - 1 - AAK23-VDIHZ-B - standby
         Portgroup Name  -  VLAN
         vmotion-53 - 53
    vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
       vmnic2 - 1 - AAK23-VDIHZ-A - standby
       vmnic3 - 1 - AAK23-VDIHZ-B - active
         Portgroup Name  -  VLAN
         Storage Controller Data Network - 52
         Storage Hypervisor Data Network - 52

Host: 10.10.50.54
    vswitch: vswitch-hx-inband-mgmt - mtu: 1500 - policy: loadbalance_srcid
       vmnic0 - 1 - AAK23-VDIHZ-A - active
       vmnic1 - 1 - AAK23-VDIHZ-B - standby
         Portgroup Name  -  VLAN
         VM Network - 0
         Storage Controller Management Network - 50
         Storage Controller Replication Network - 0
         Management Network - 50
    vswitch: vswitch-hx-vm-network - mtu: 1500 - policy: loadbalance_srcid
       vmnic4 - 1 - AAK23-VDIHZ-A - active
       vmnic5 - 1 - AAK23-VDIHZ-B - active
         Portgroup Name  -  VLAN
         vm-network-54 - 54
    vswitch: vmotion - mtu: 9000 - policy: loadbalance_srcid
```

```
        vmnic6 - 1 - AAK23-VDIHZ-A - active
        vmnic7 - 1 - AAK23-VDIHZ-B - standby
          Portgroup Name  -  VLAN
          vmotion-53 - 53
     vswitch: vswitch-hx-storage-data - mtu: 9000 - policy: loadbalance_srcid
        vmnic2 - 1 - AAK23-VDIHZ-A - standby
        vmnic3 - 1 - AAK23-VDIHZ-B - active
          Portgroup Name  -  VLAN
          Storage Controller Data Network - 52
          Storage Hypervisor Data Network - 52


 Host: 10.10.50.51
     Could not ping IP 169.254.254.2 from vmk1, verify network connectivity


 Host: 10.10.50.52
 Host: 10.10.50.53
 Host: 10.10.50.54
 Controller VM Clocks:
     stCtlVM-WZP212416UO  -  2017-11-13 17:57:21 - Have not recently synced
with NTP server
     stCtlVM-WZP21230UBH  -  2017-11-13 17:57:22 - Have not recently synced
with NTP server
     stCtlVM-WZP212416VK  -  2017-11-13 17:57:24 - Have not recently synced
with NTP server
     stCtlVM-WZP212416UQ  -  2017-11-13 17:57:25 - Have not recently synced
with NTP server
 Cluster:
     Version - 2.6.1a-26588
     Model - HXAF220C-M5SX
     Health - HEALTHY
     ASUP enabled - False
     SMTP Server -
root@Cisco-HX-Installer-Appliance: ~root@Cisco-HX-Installer-Appliance:~#
```

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.

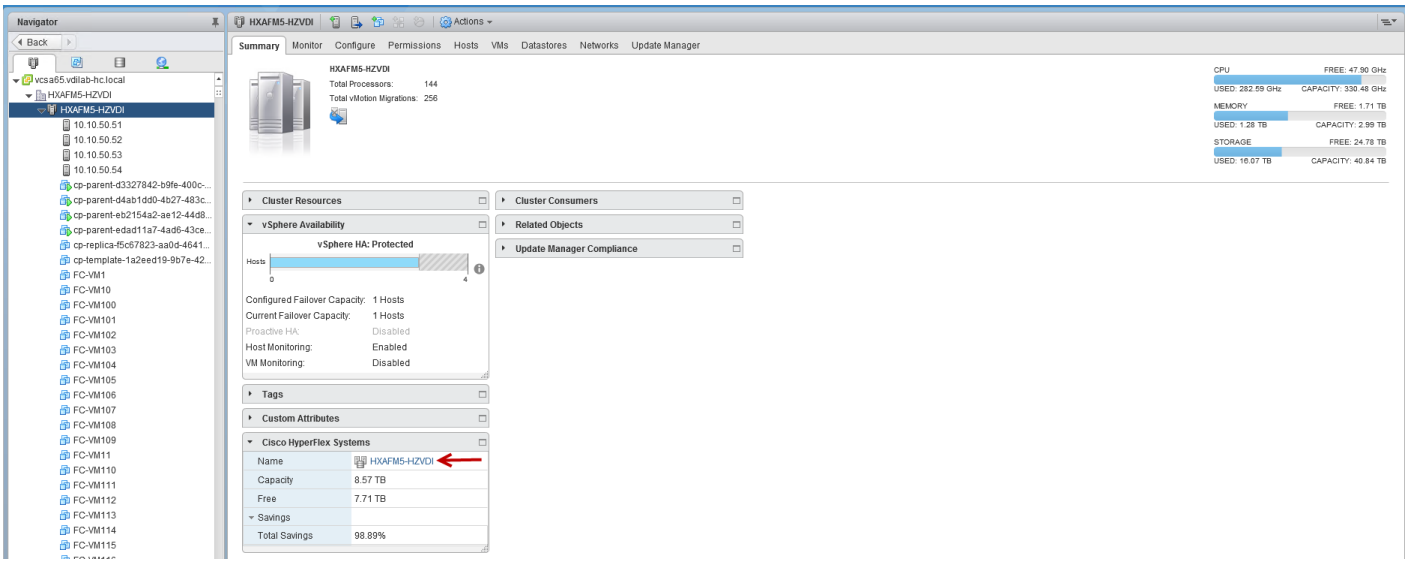

The Summary tab shows the details about the cluster status, capacity, and performance.

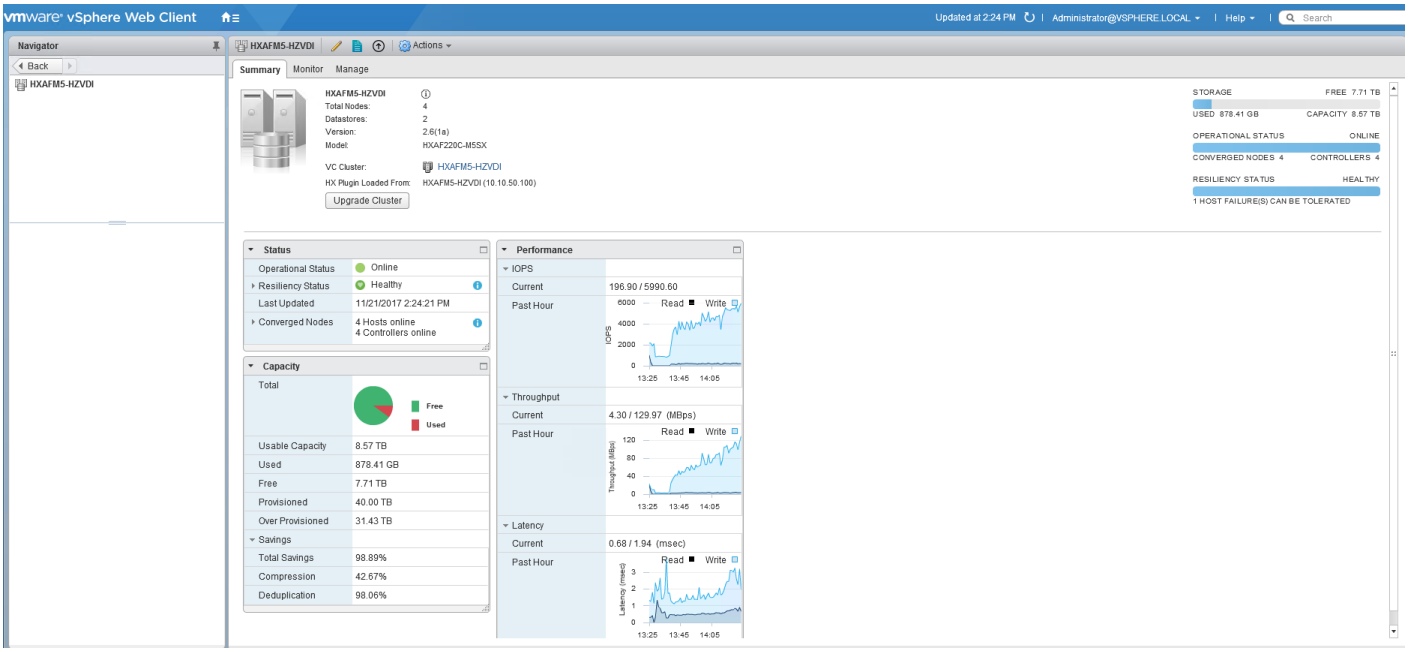

11. Click Manage, select Datastores. Click the Add datastore icon, select the datastore name and size to provision.

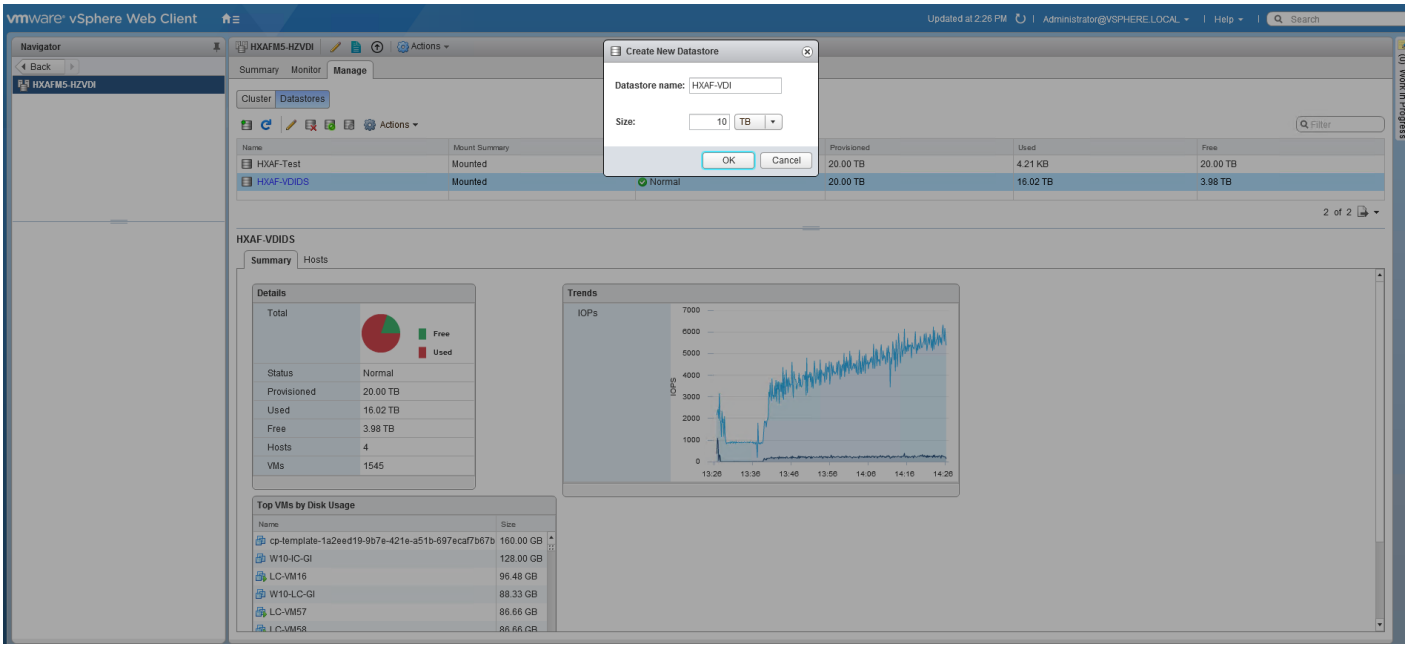

You have now created a 20TB datastore for the Horizon pooled, persistent/non-persistent, and RDSH server desktop performance test.

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

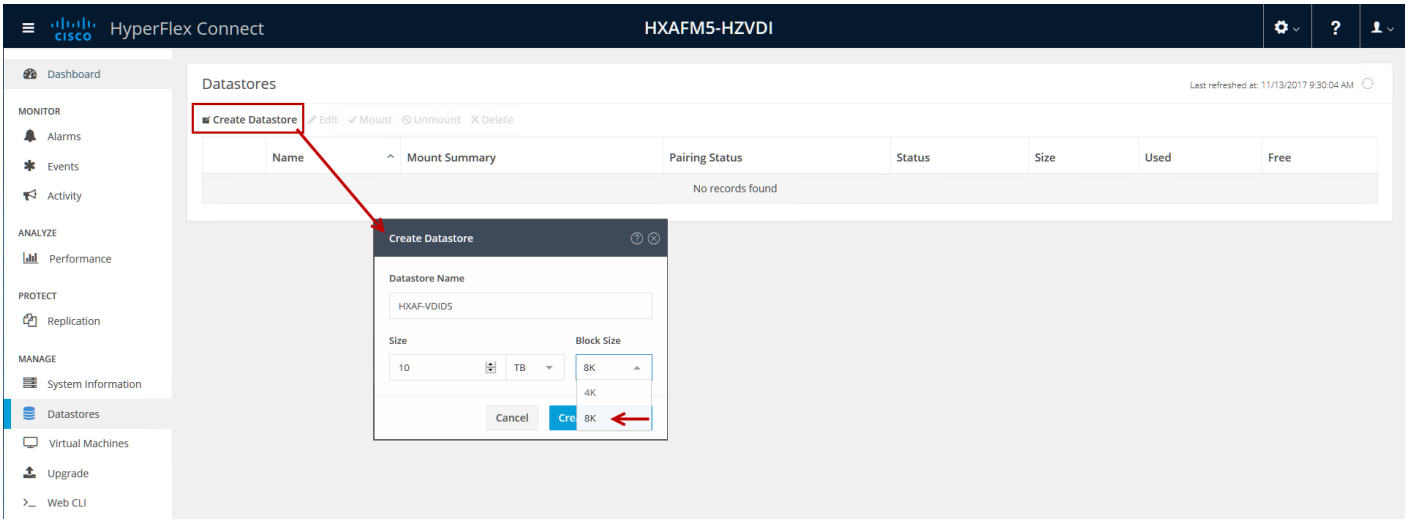

Building the Virtual Machines and Environment for Workload Testing

This section 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_3

Install VMware Horizon Composer Server

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

1. Open installer for Horizon composer. VMware-viewcomposer-7.3.1-6744335.exe
2. Click Next to continue.

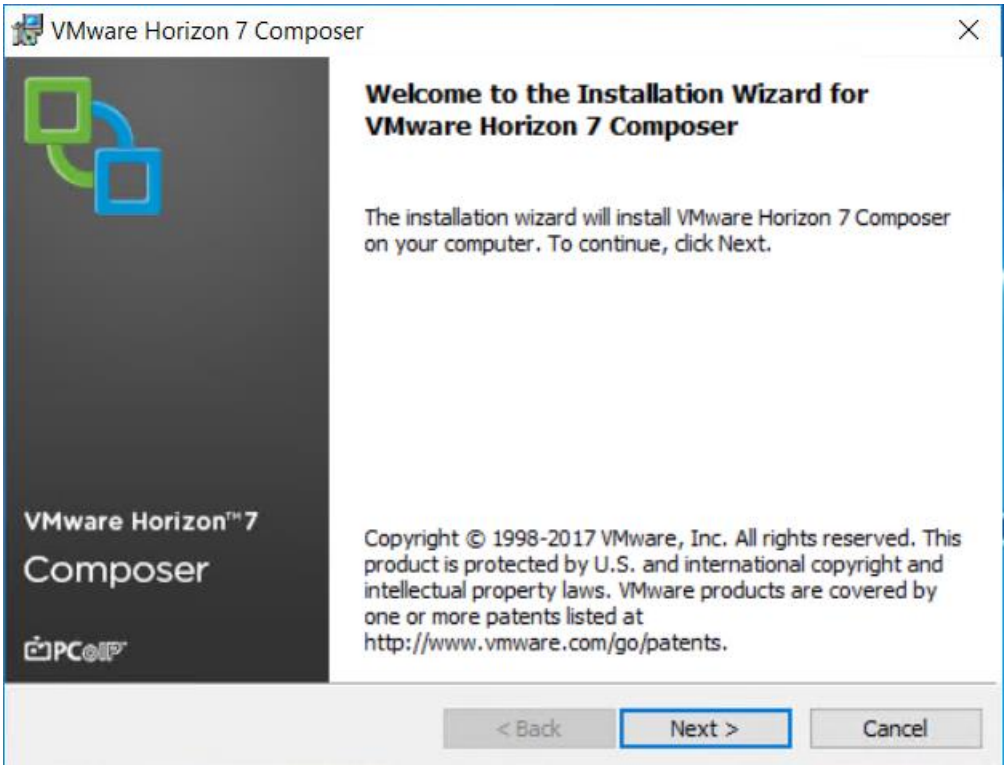

3. Accept the EULA. Click Next.

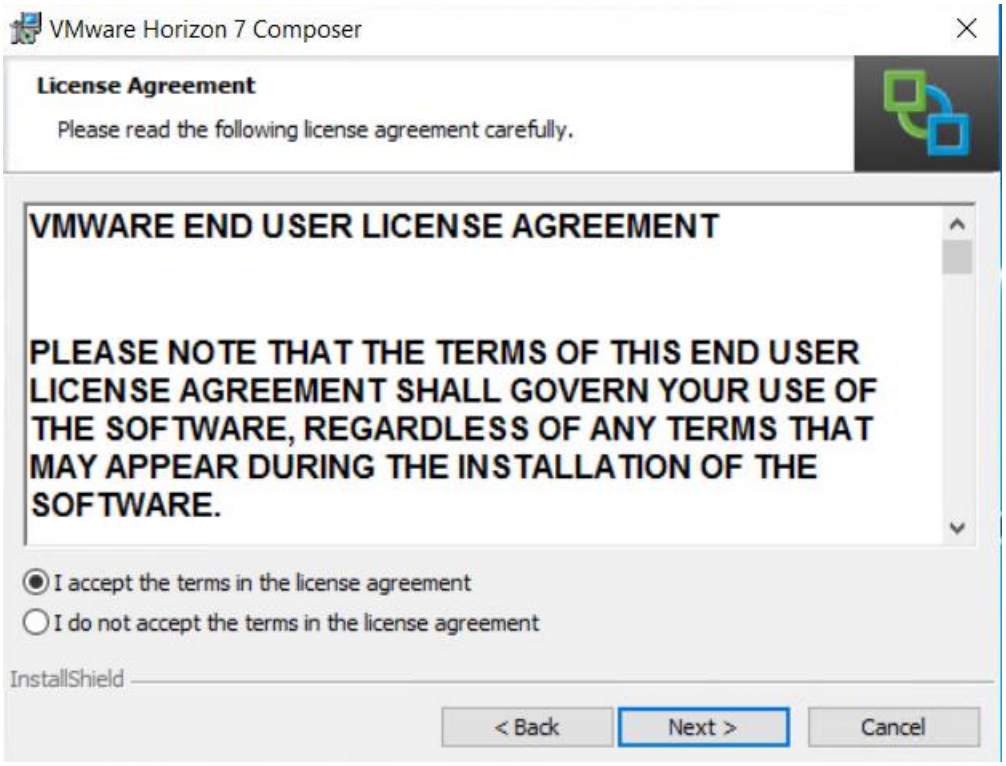

4. Click Next to accept the default installation folder.

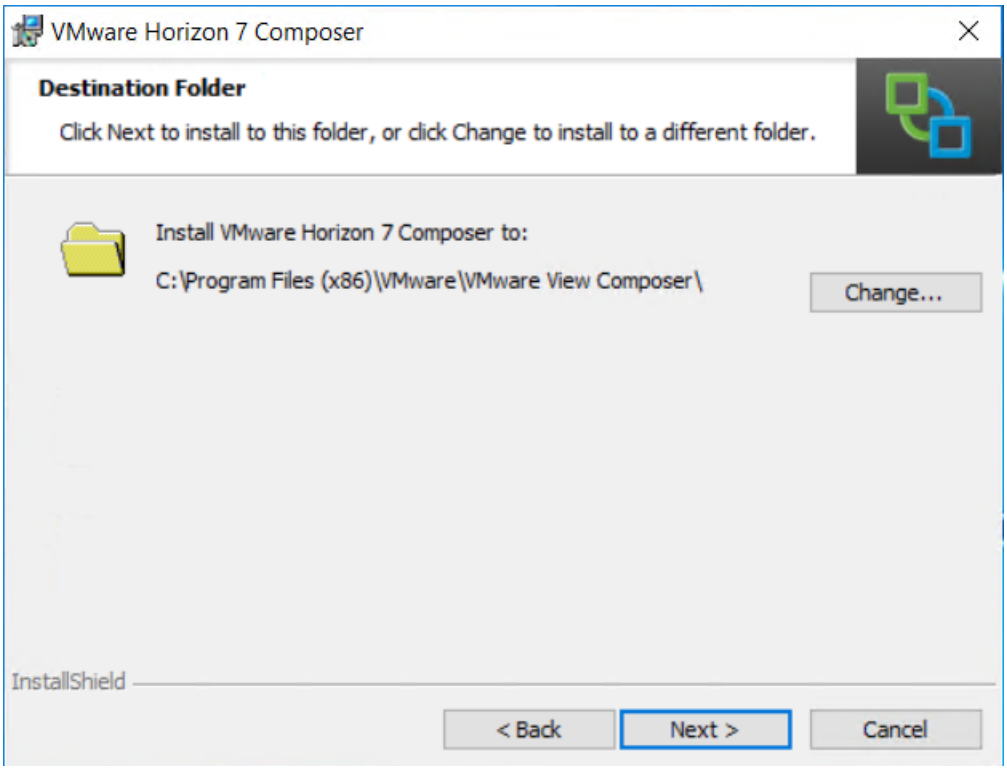

5. Enter the database information. The ODBC database can be configured during the installation by clicking ODBC DSN Setup.

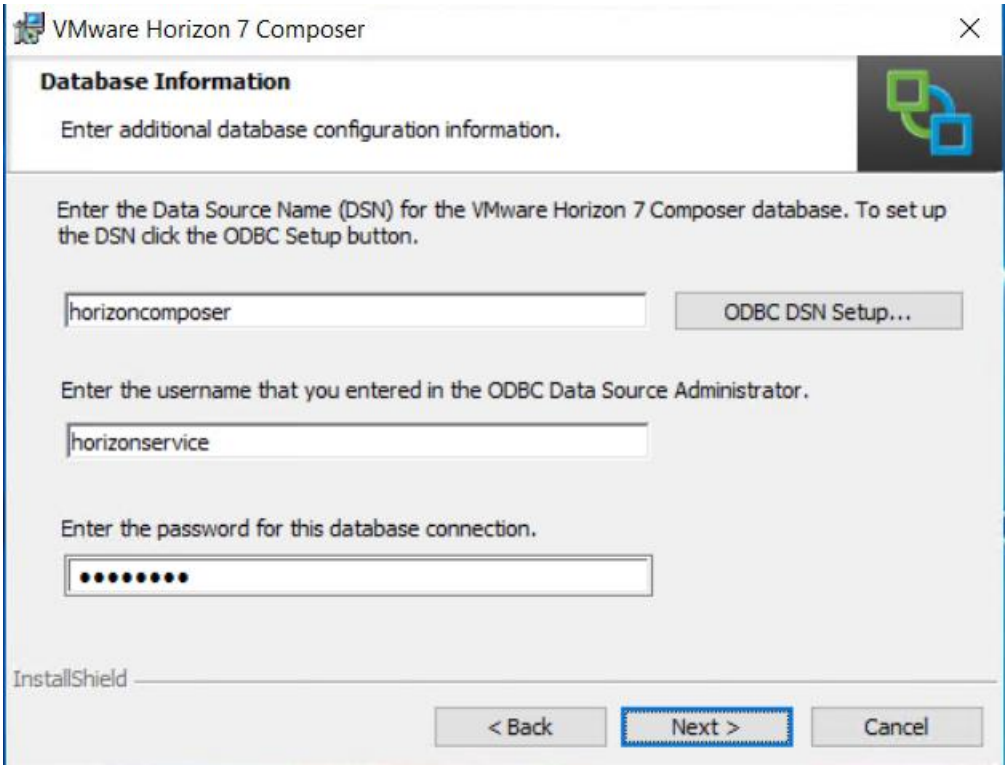

Configure the ODBC Source Name

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

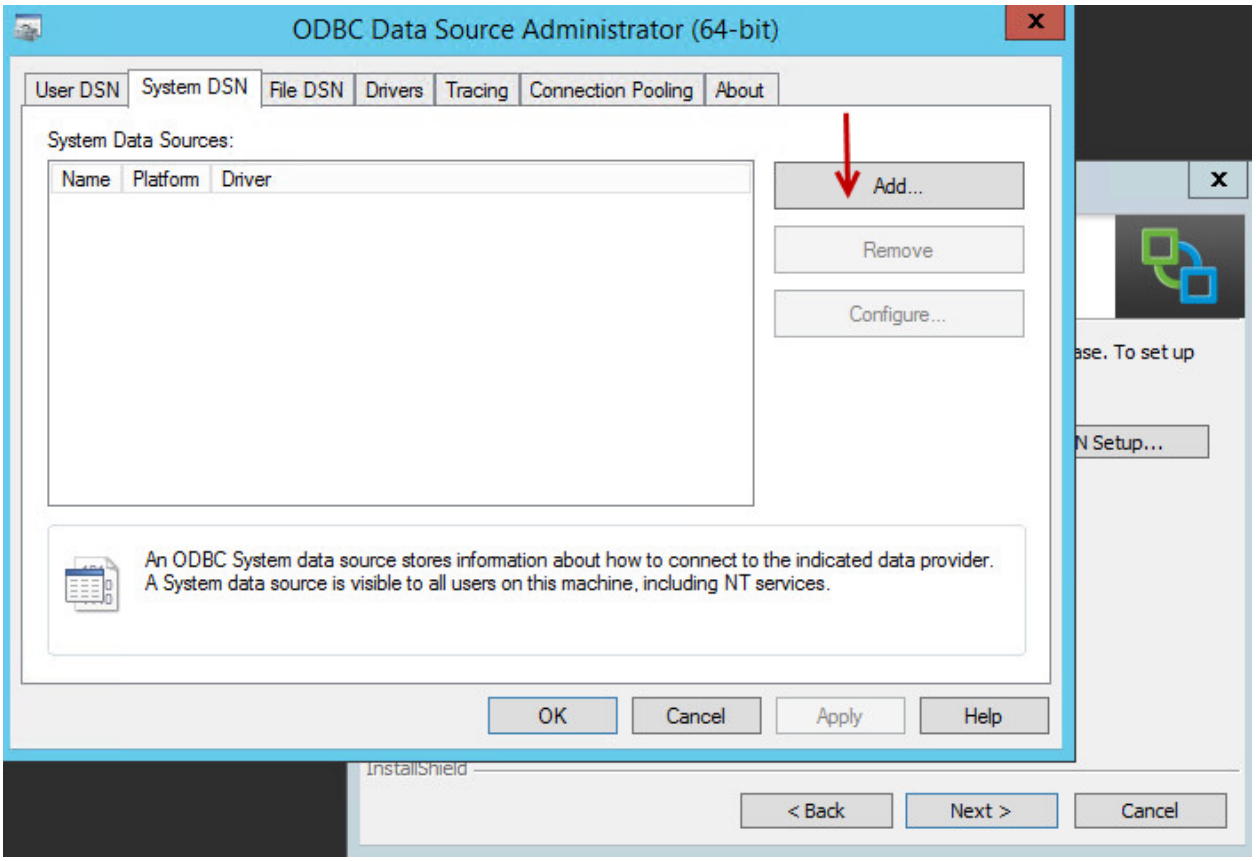

2. Create a new Data source and select SQL server native client. You will use an existing instanace of the Mircosoft SQL server 2016 for the current deployment. Click Finish.

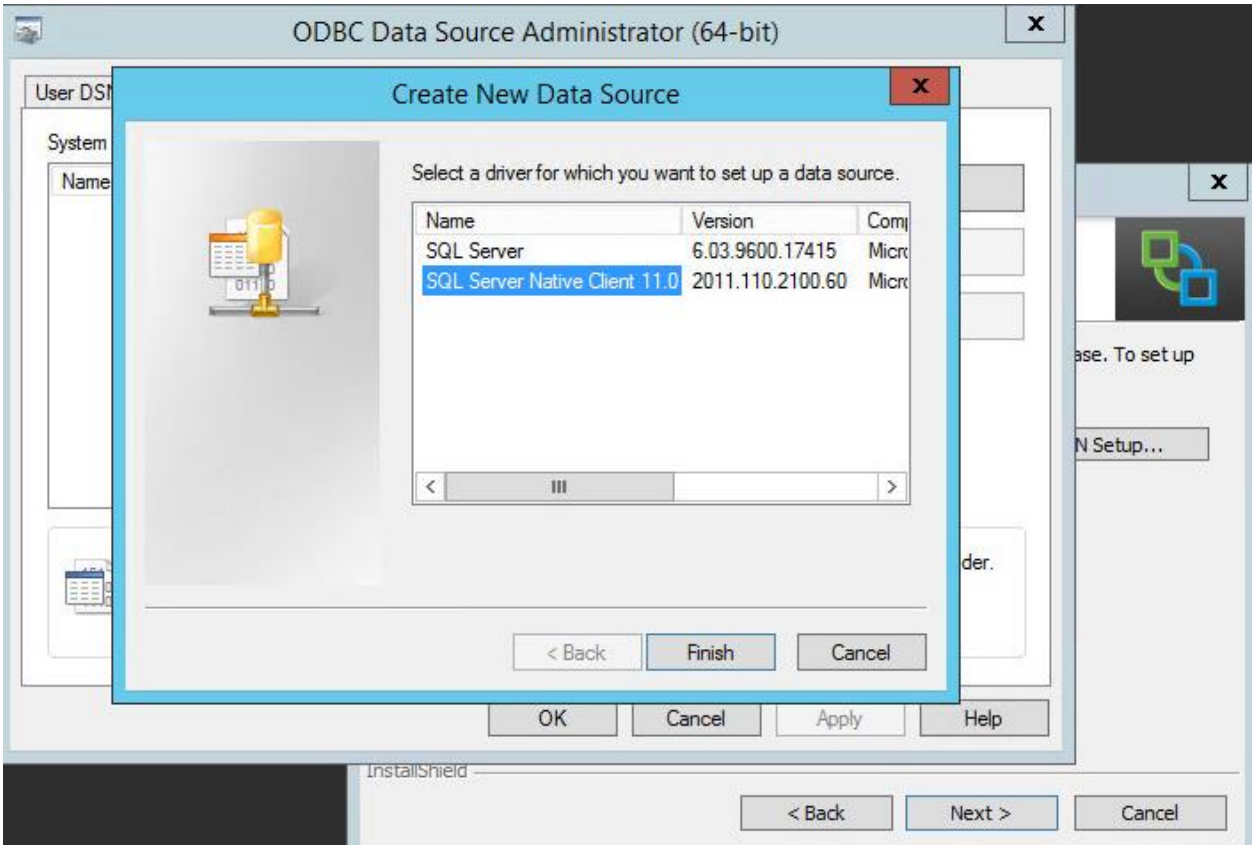

3. Create a name for data source, select SQL server, click Next.

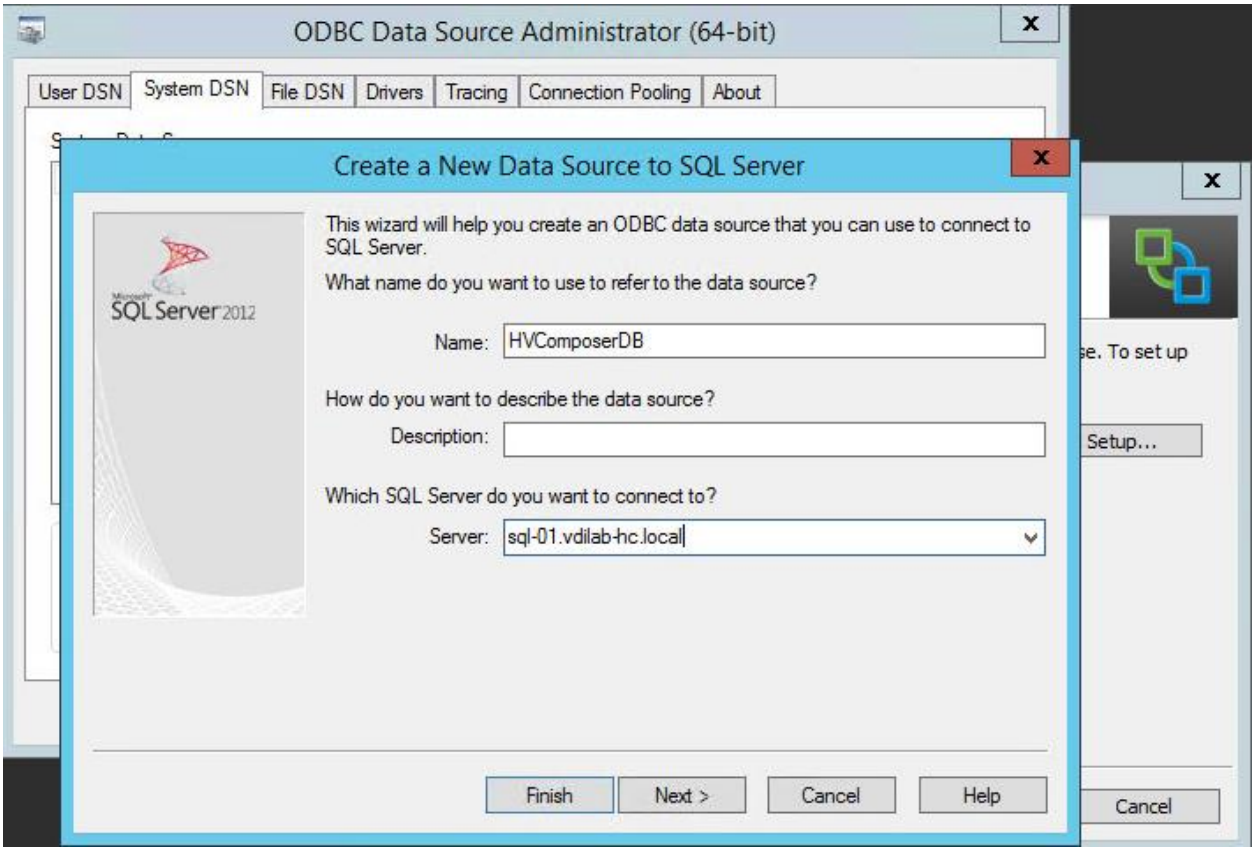

4. Enter the login credentails for the SQL server authentication or use Windows Authentication. Click Next.

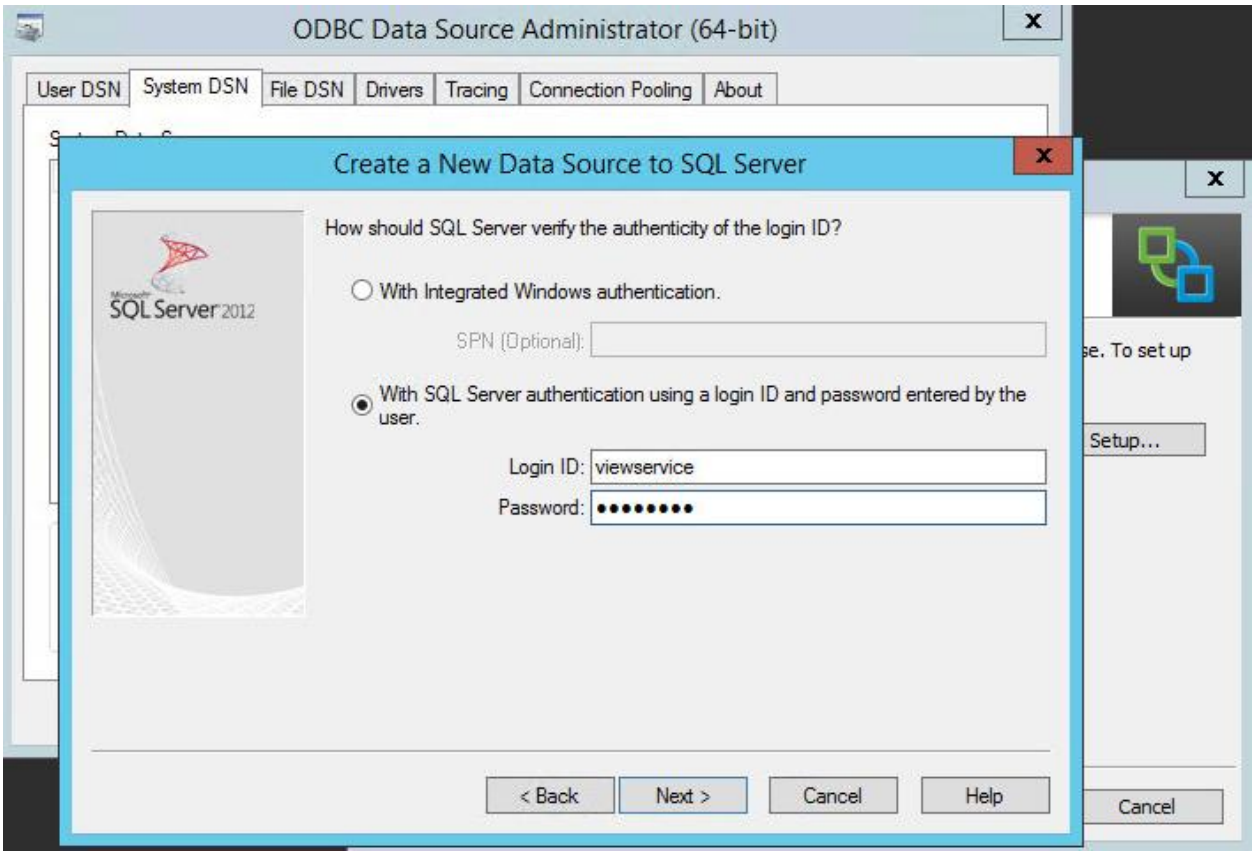

5. Select Default Database, click Next.

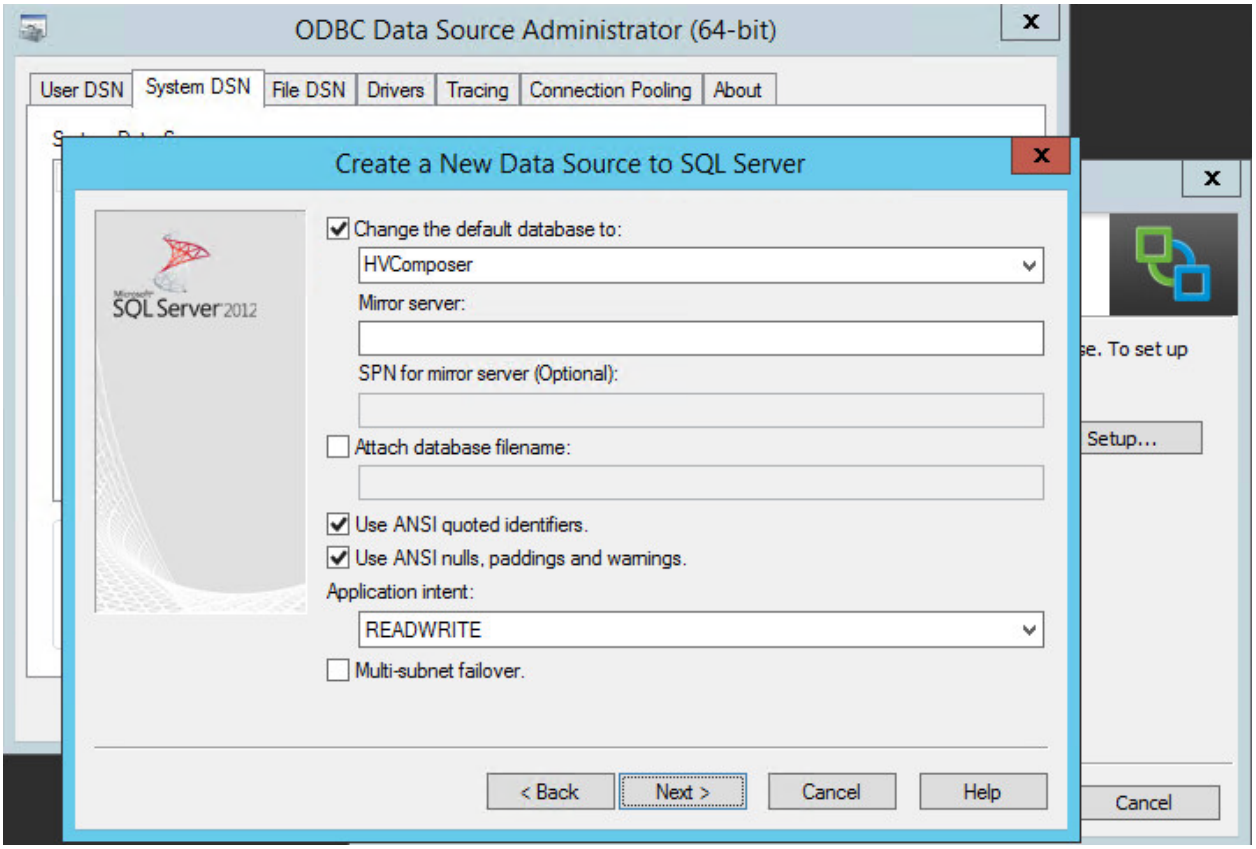

6. Check the box to select language for SQL server system messages. Click Finish.

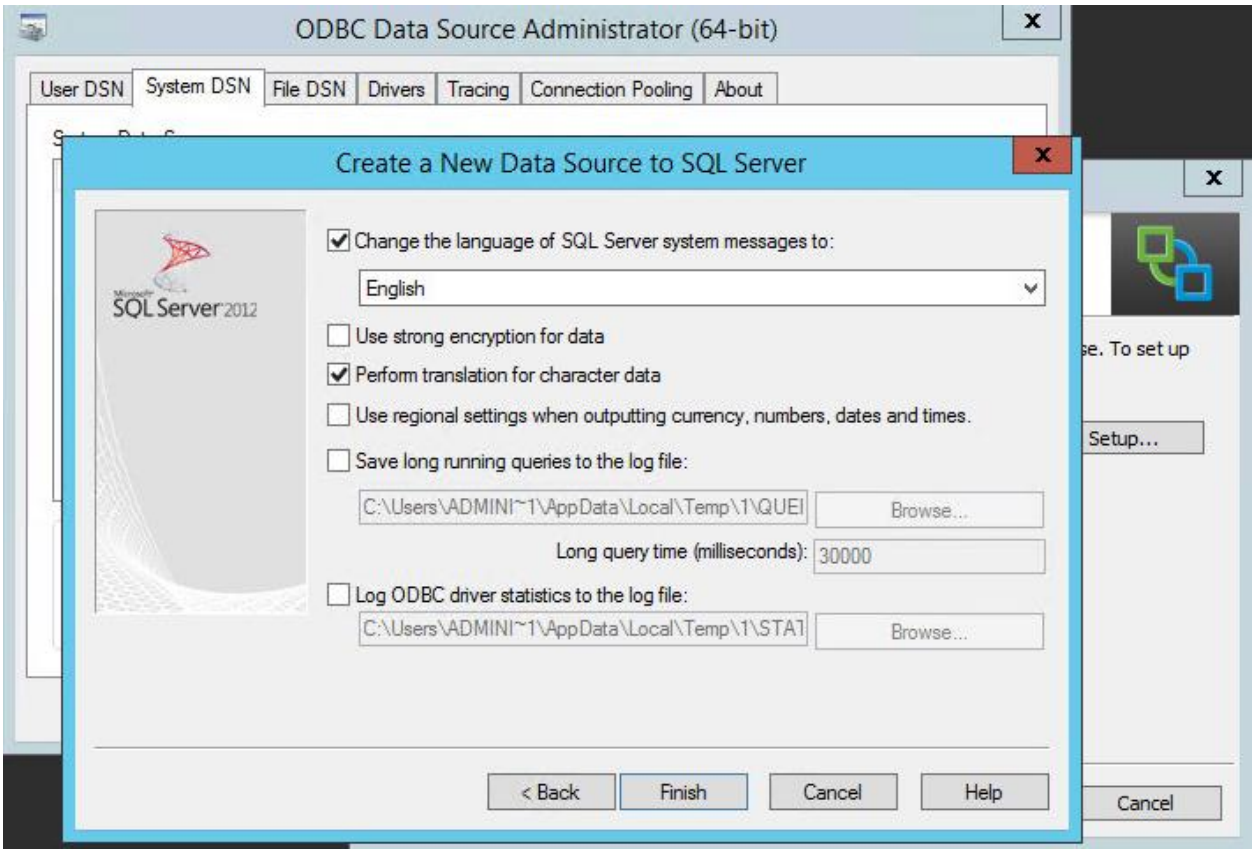

7. Click Test datastore to verify connectivity between SQL server and newly create Data source.

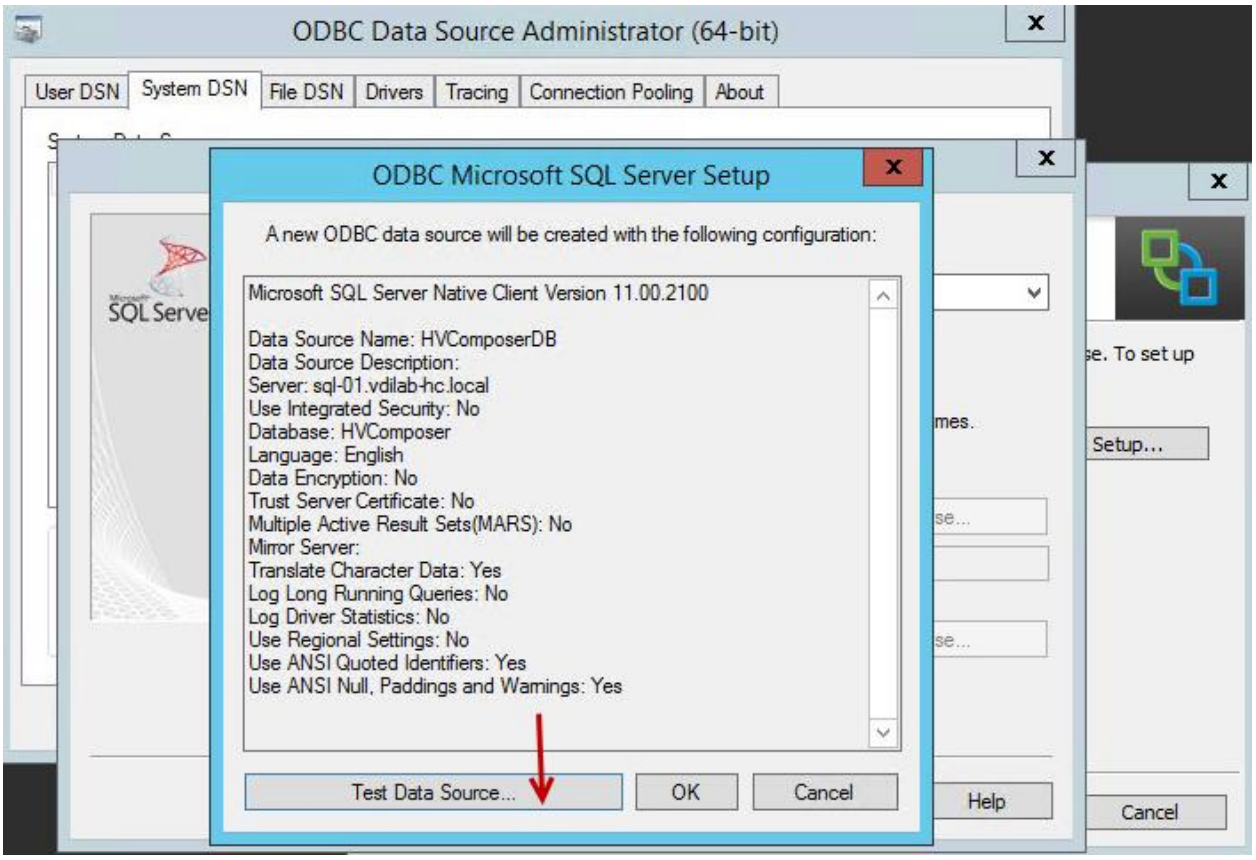
ODBC Data Source Administrator (64-bit)
User DSN
System DSN
File DSN
Drivers
Tracing
Connection Pooling
About
ODBC Microsoft SQL Server Setup
A new ODBC data source will be created with the following configuration:
Microsoft SQL Server Native Client Version 11.00.2100
Data Source Name: HVComposerDB
Data Source Description:
Server: sql-01.vdilab-hc.local
Use Integrated Security: No
Database: HVComposer
Language: English
Data Encryption: No
Trust Server Certificate: No
Multiple Active Result Sets(MARS): No
Mirror Server:
Translate Character Data: Yes
Log Long Running Queries: No
Log Driver Statistics: No
Use Regional Settings: No
Use ANSI Quoted Identifiers: Yes
Use ANSI Null, Paddings and Warnings: Yes
Test Data Source...
OK
Cancel
Help
Cancel
Setup...

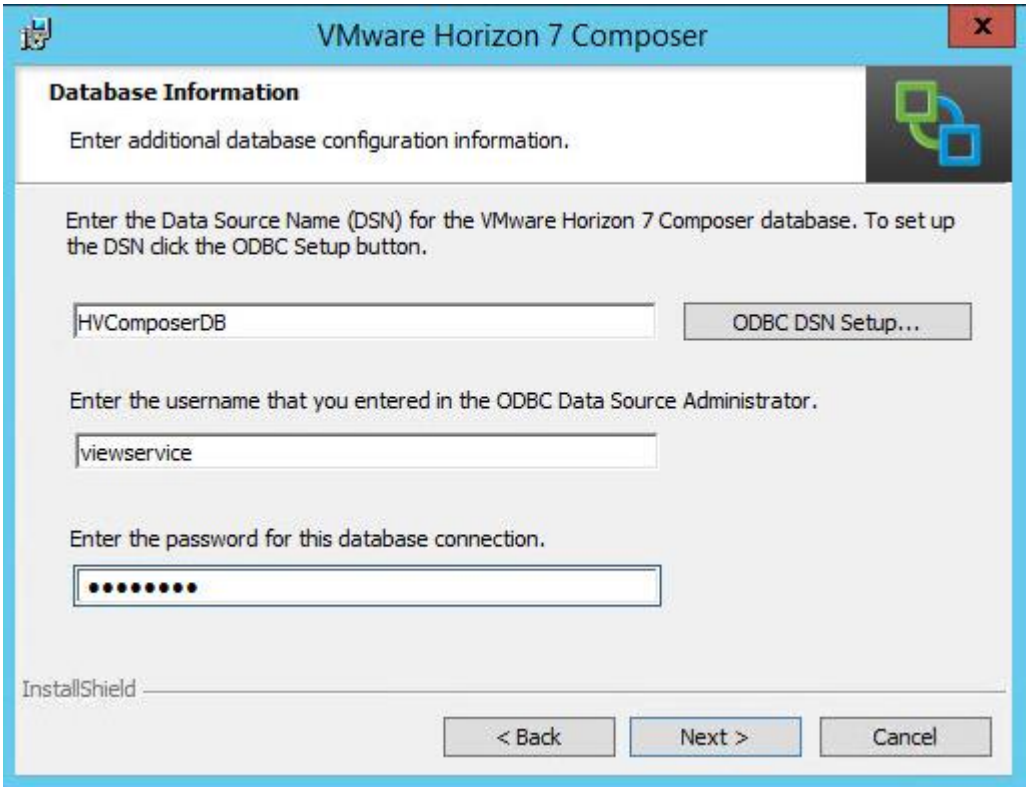
VMware Horizon 7 Composer
Database Information
Enter additional database configuration information.
Enter the Data Source Name (DSN) for the VMware Horizon 7 Composer database. To set up the DSN click the ODBC Setup button.
HVComposerDB
ODBC DSN Setup...
Enter the username that you entered in the ODBC Data Source Administrator.
viewservice
Enter the password for this database connection.
InstallShield
< Back
Next >
Cancel

8. Since this a new instanace 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.

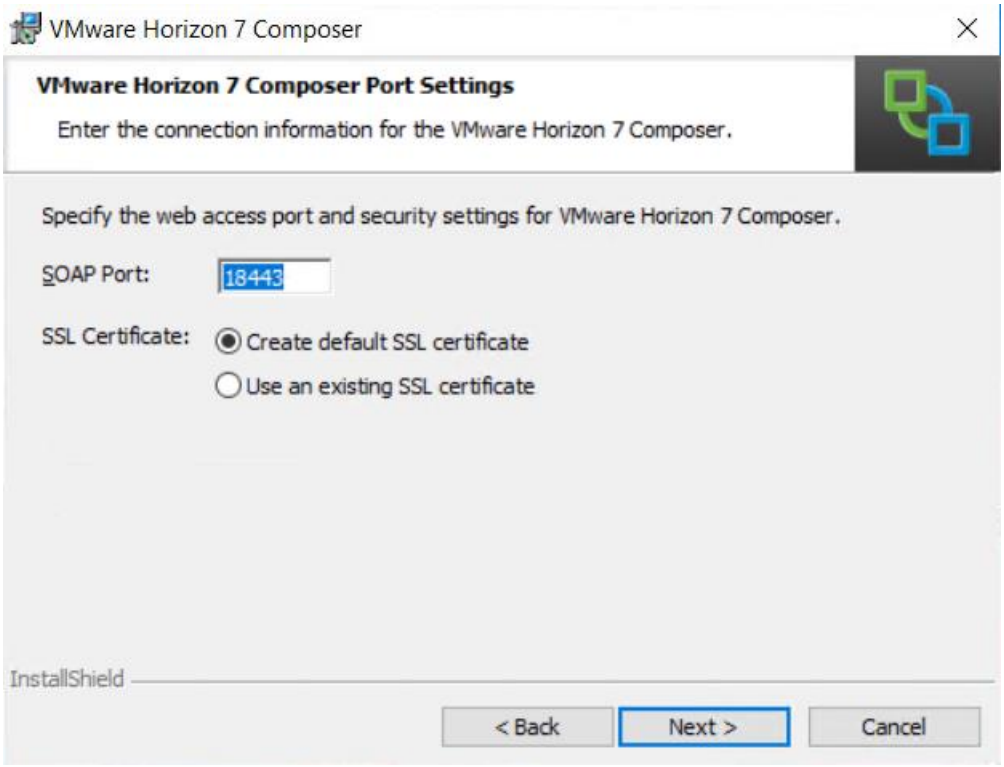

11. Click Install.

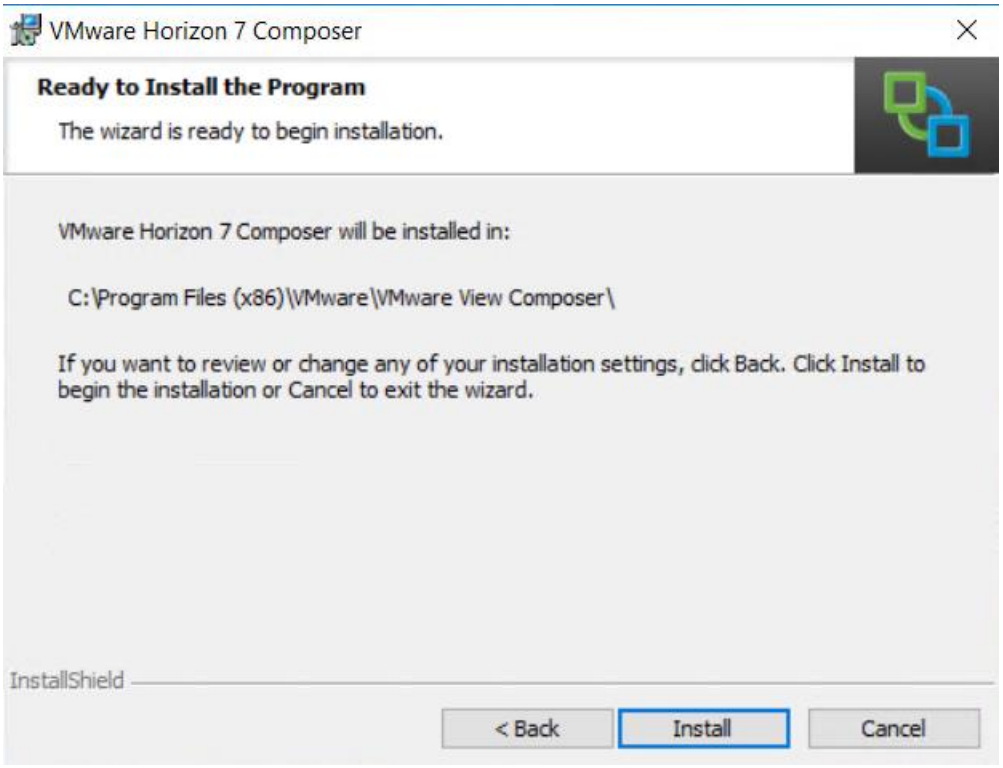

12. Click Finish.

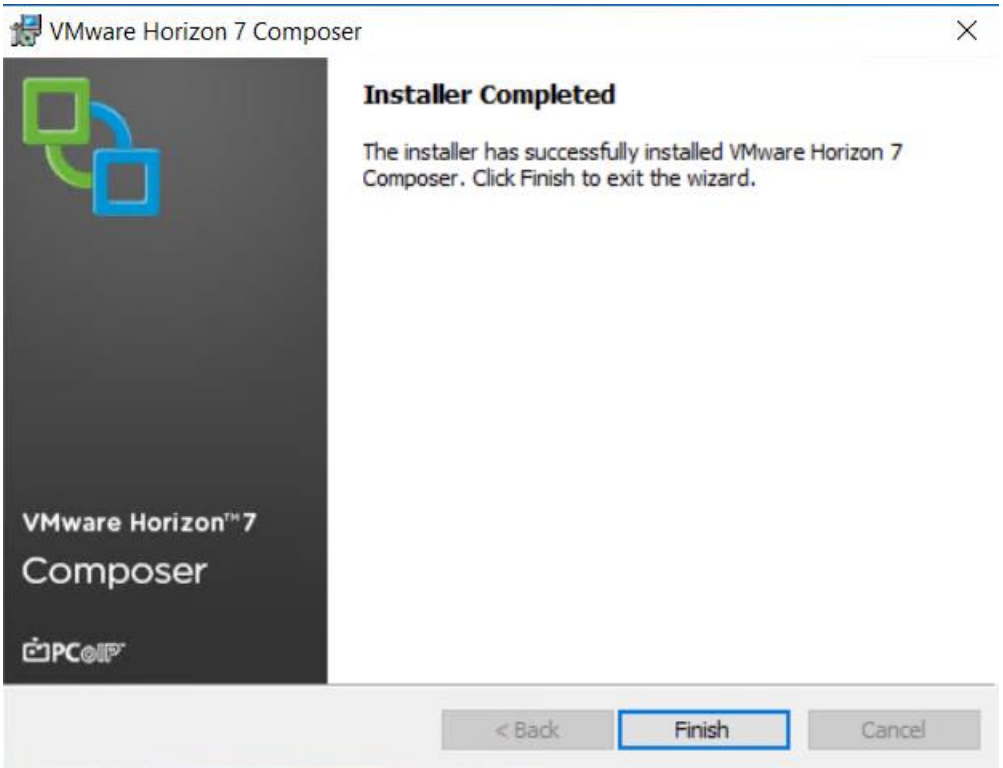

Install Horizon Connection/Replica Servers

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

1. Open view connection server installation,VMware-viewconnectionserver-x86_64-7.3.1-6760913.exe.
2. Click Next.

3. Accpet the EULA, click Next.

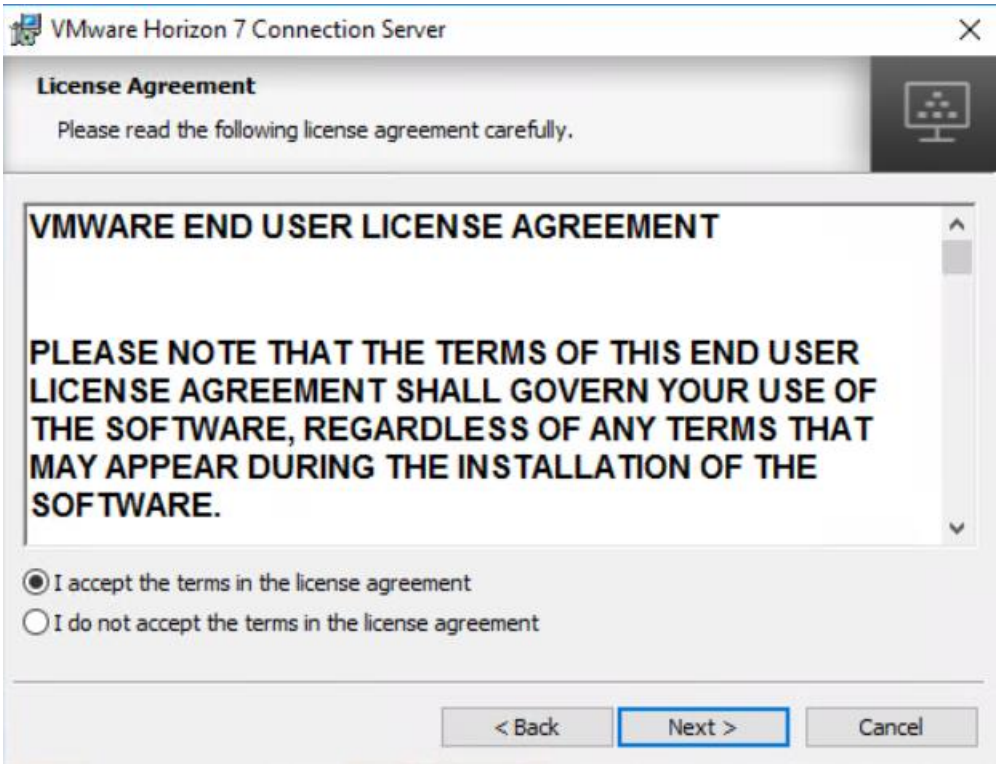

4. Leave default destination folder, click Next.

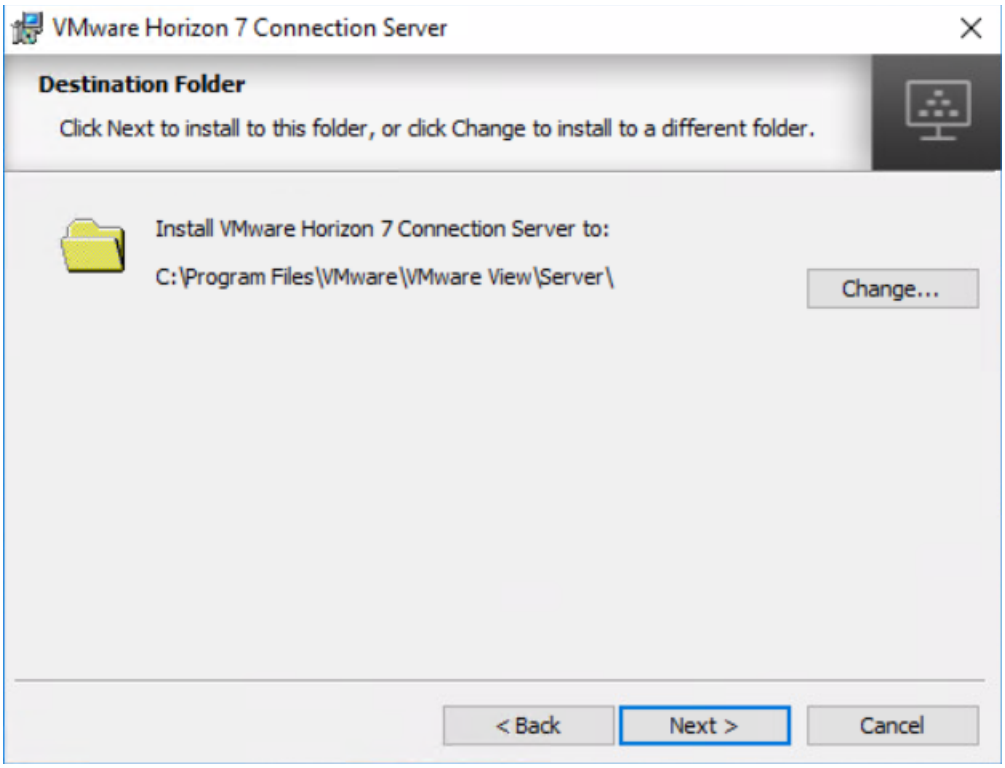

5. Select type of instance intended to install.

6. Select Standard Server instance for primary connection server installation.

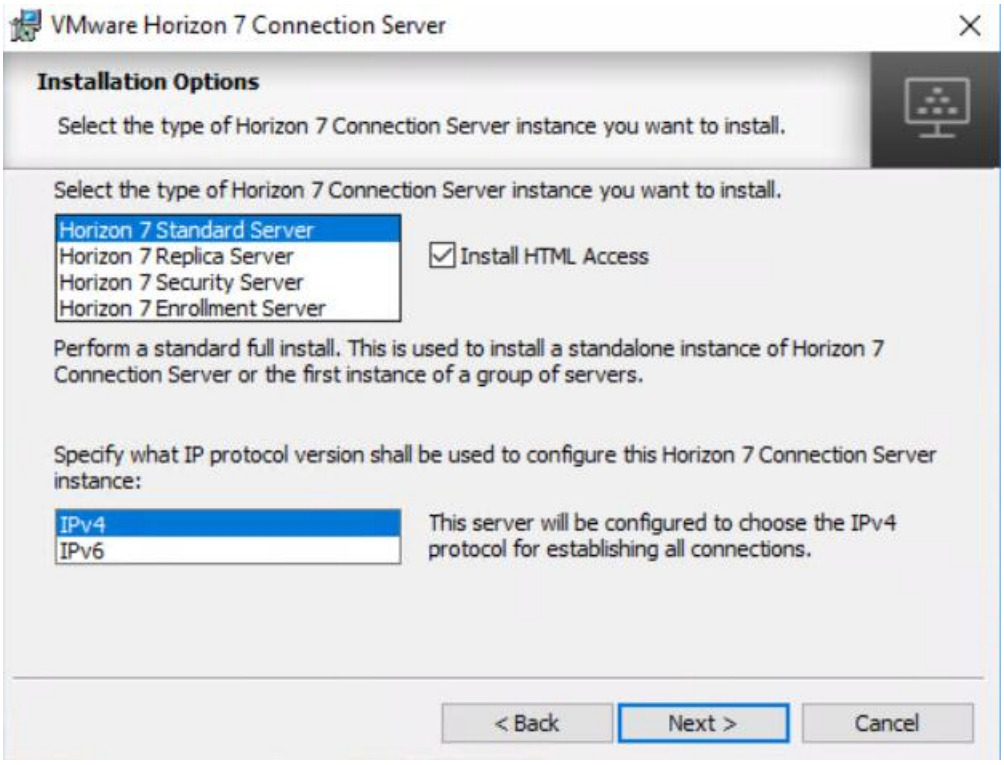

7. Select Replica server instance for fault tolerant connection server configration after completion of Standard Server instance installation.

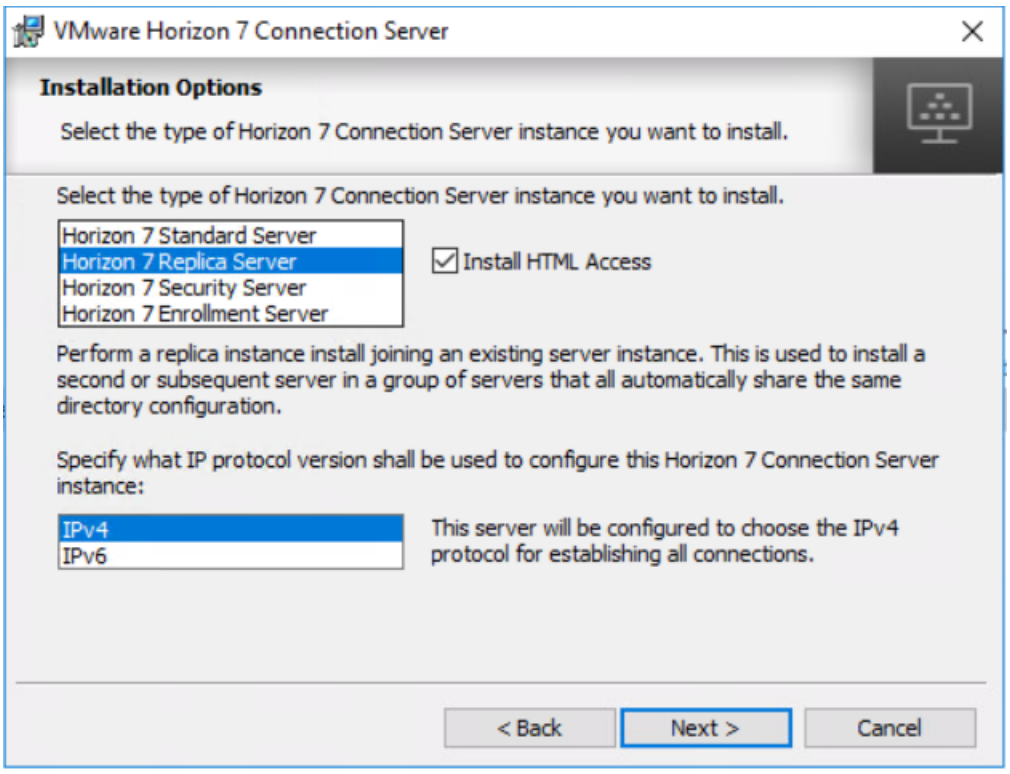

8. Enter the Data Recovery Password.

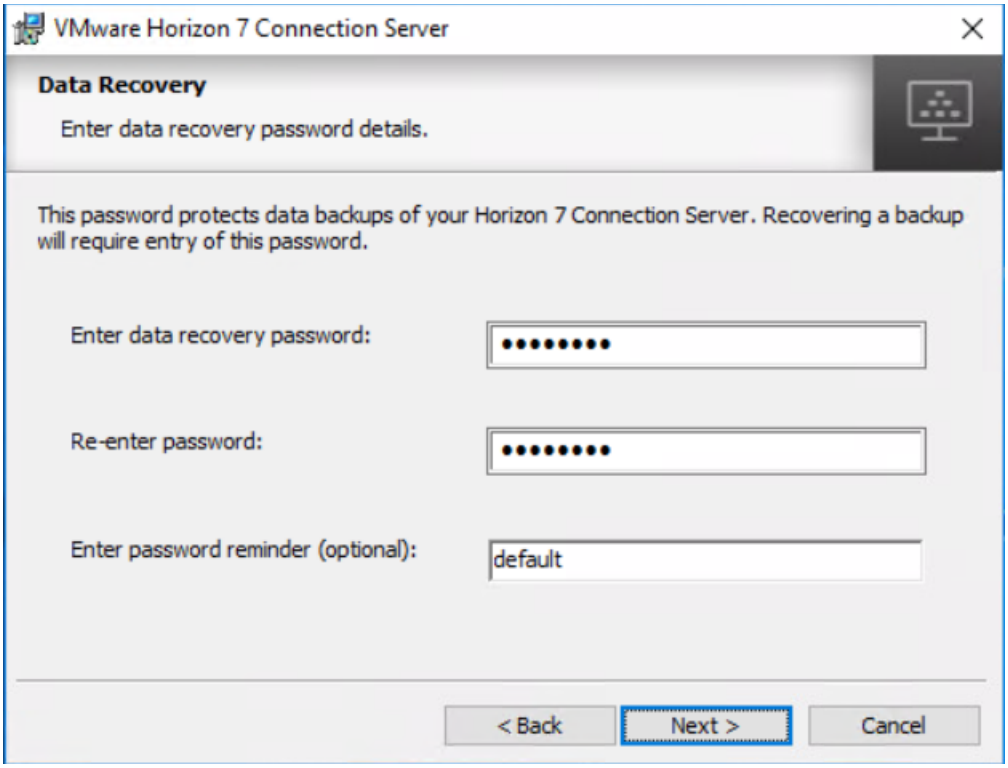

9. Click Next.

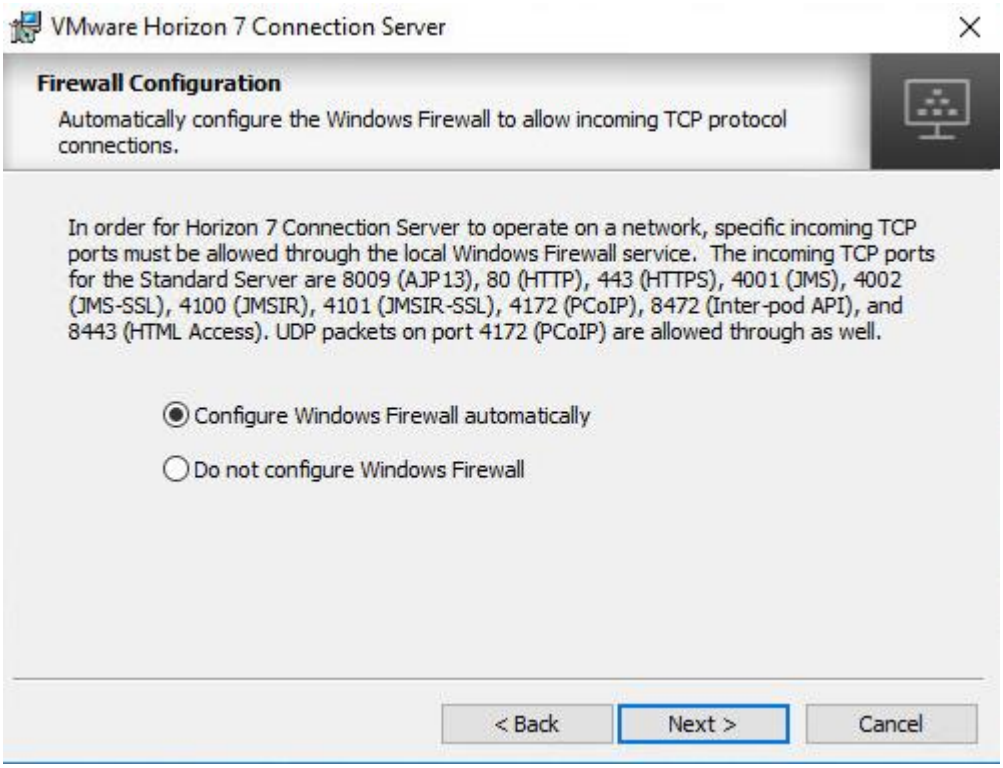

10. Select authorized users and group, click Next.

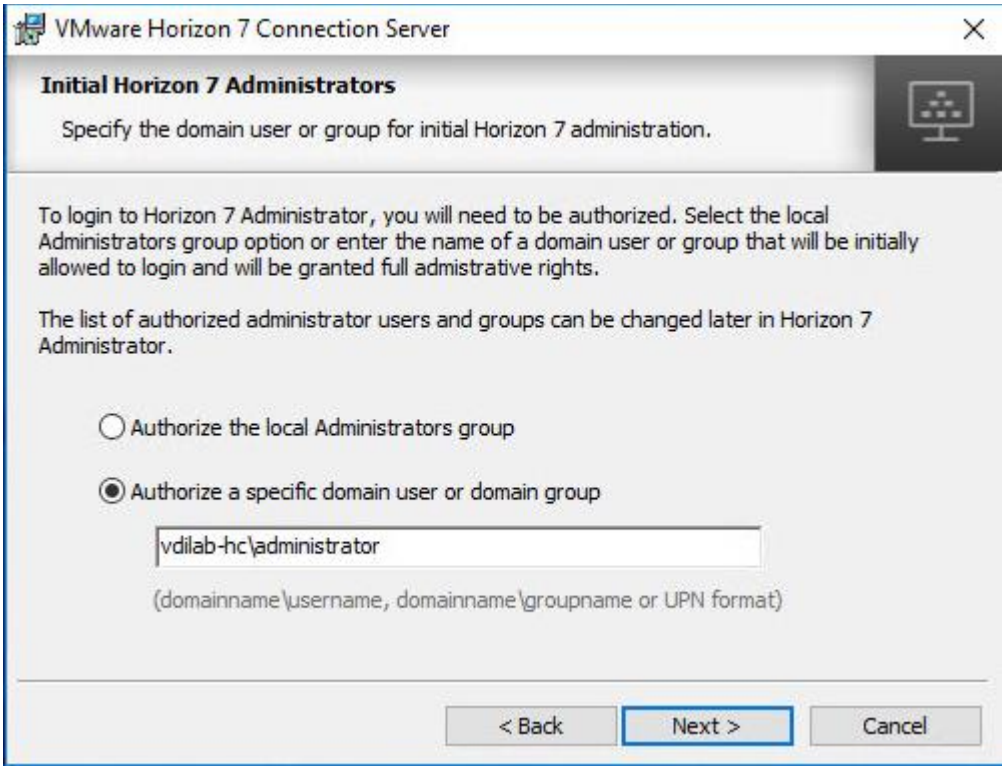

11. Enter domain credentials for previously specified domain user/group.

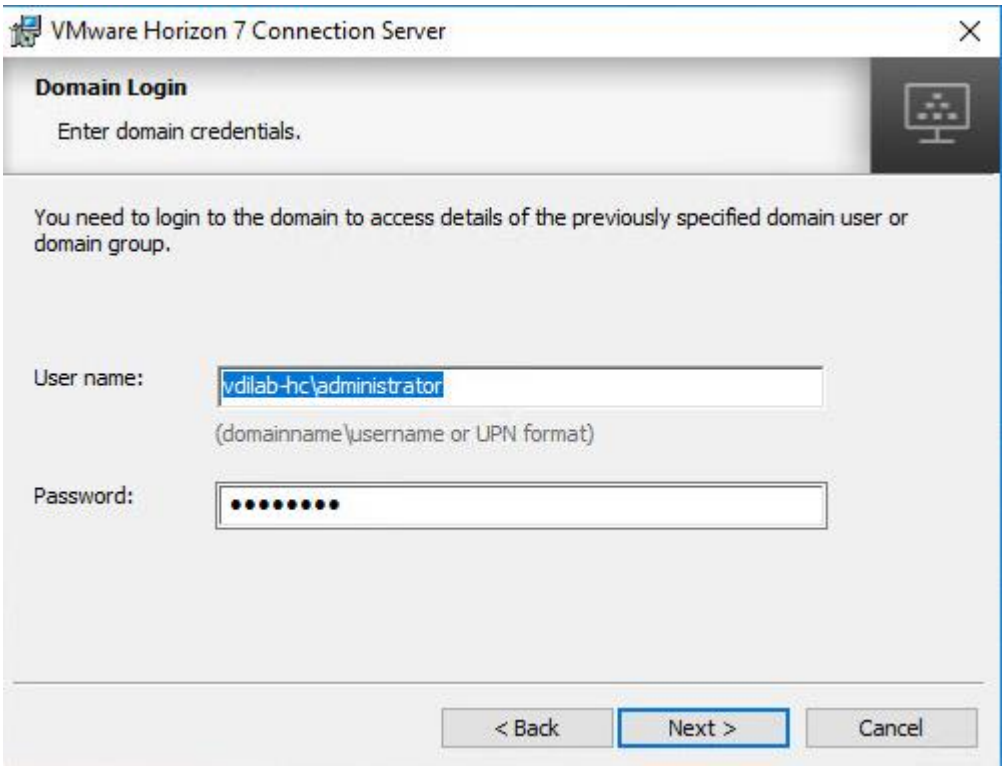

12. Opt-in or Opt-out of User Experience Improvement Program. Click Next.

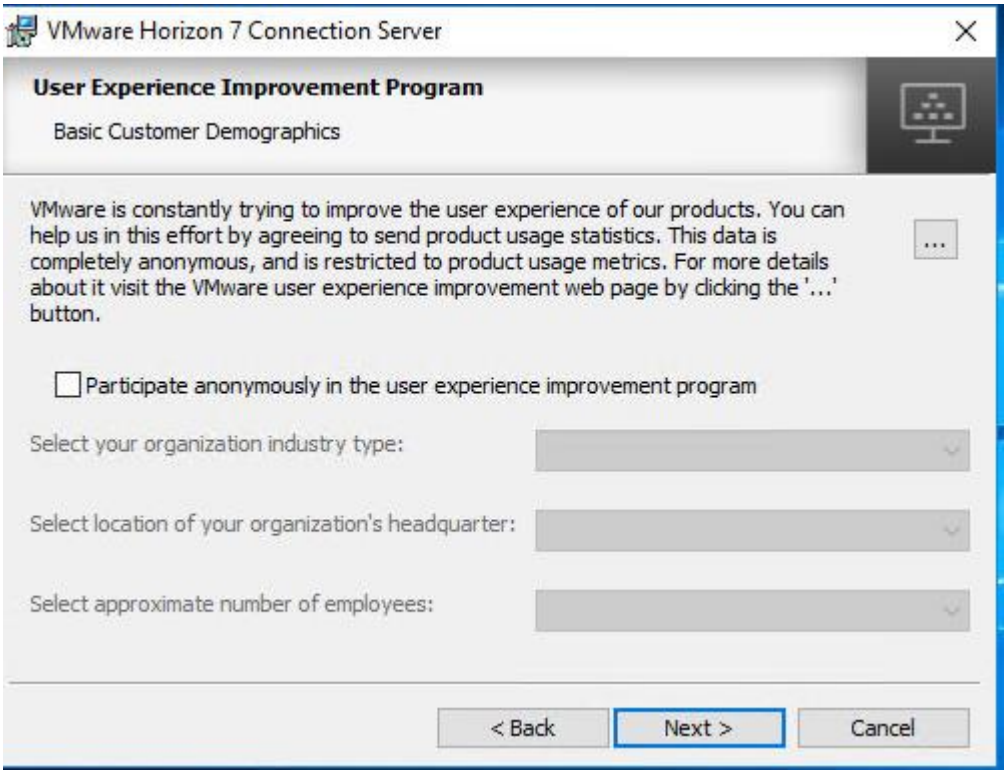

13. Click Install.

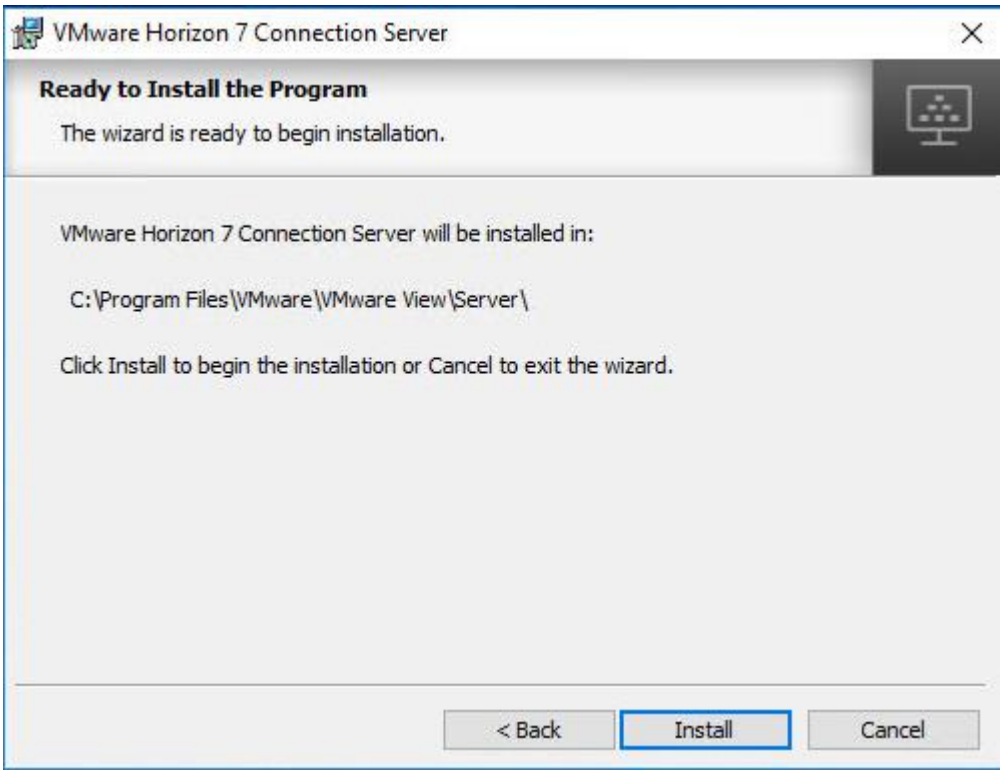

14. Click Finish.

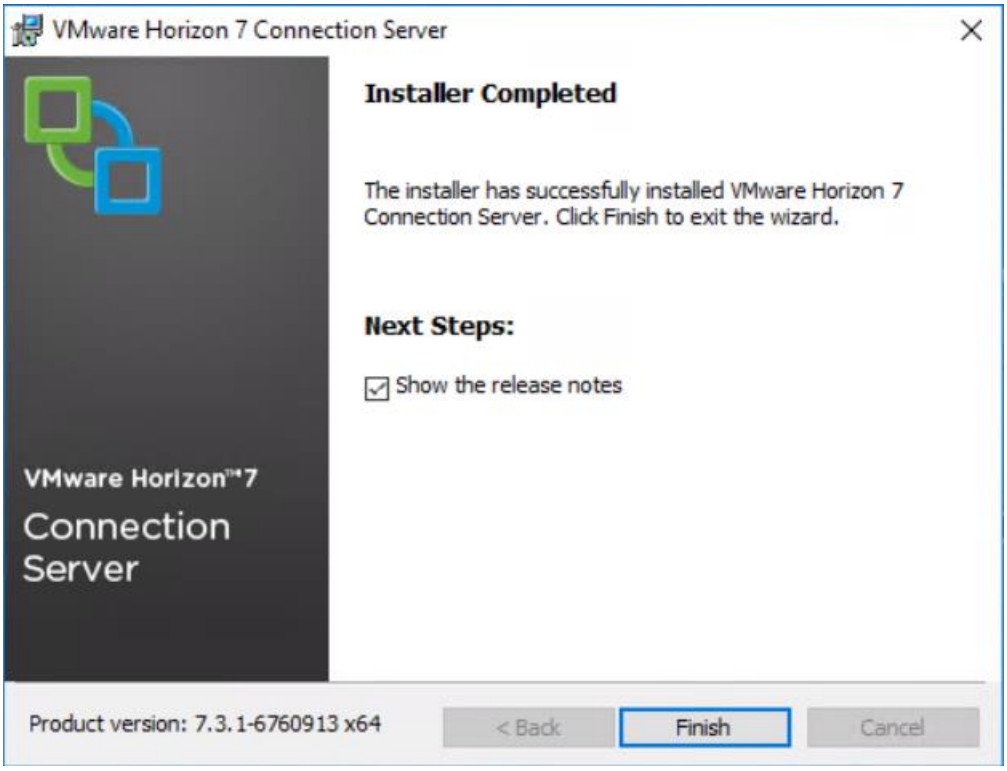

Create a Microsoft Management Console Certificate Request

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

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.

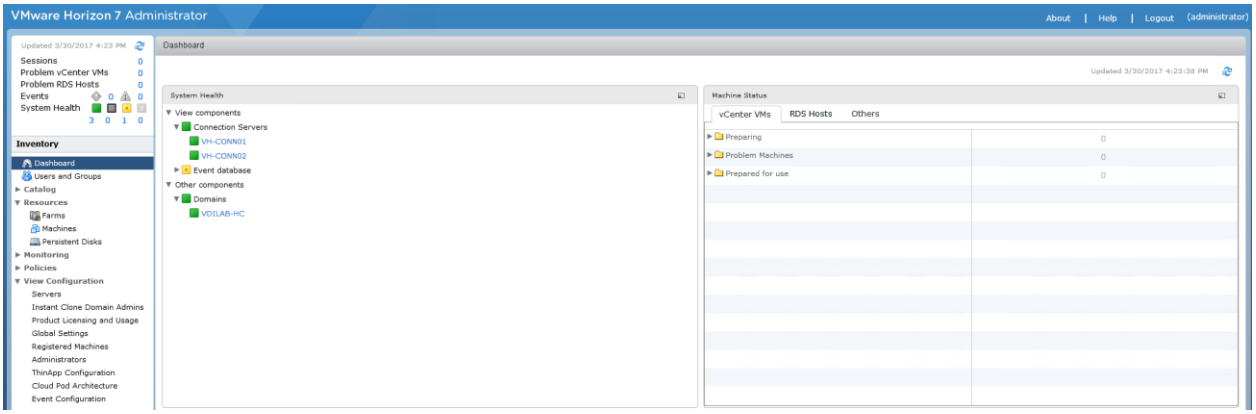

Configure Event Database

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.

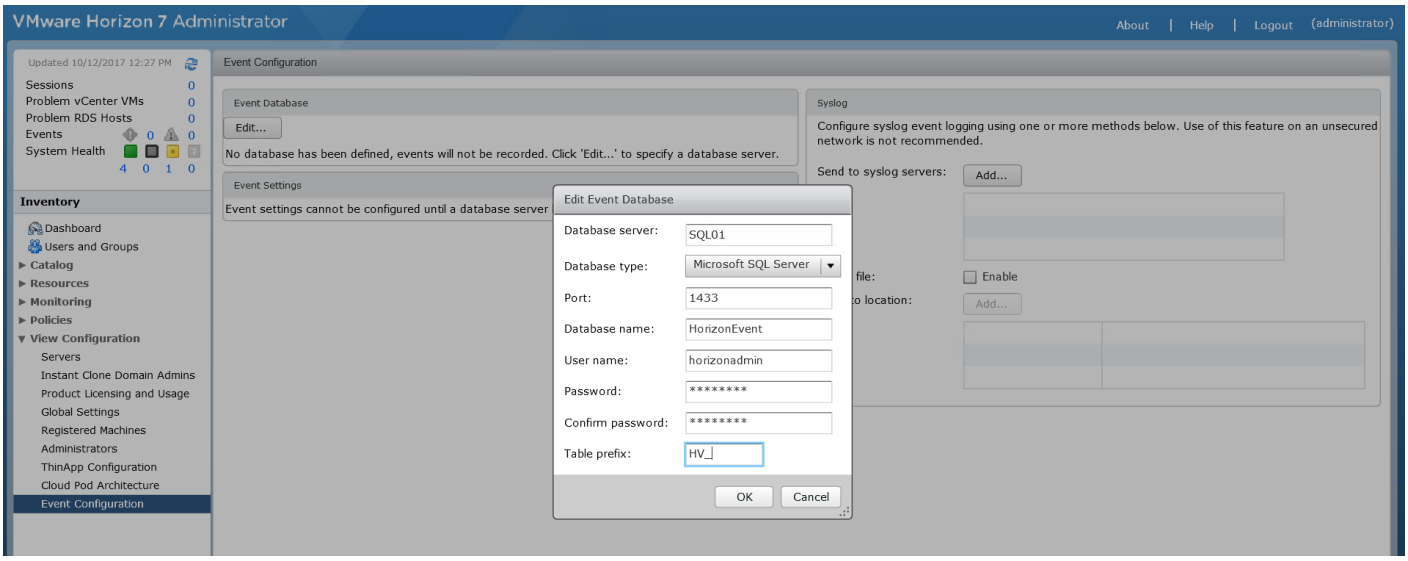

The details are shown below:

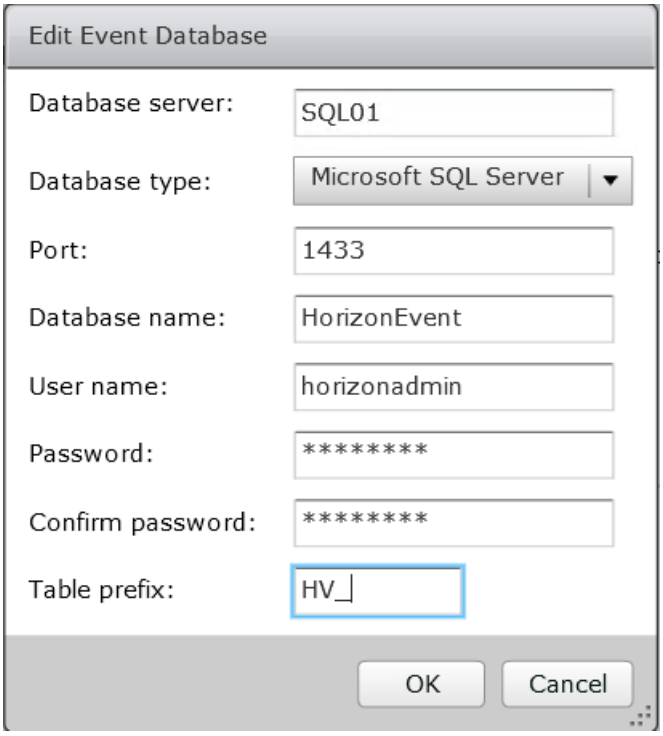

Configure Horizon 7 Licenses

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.

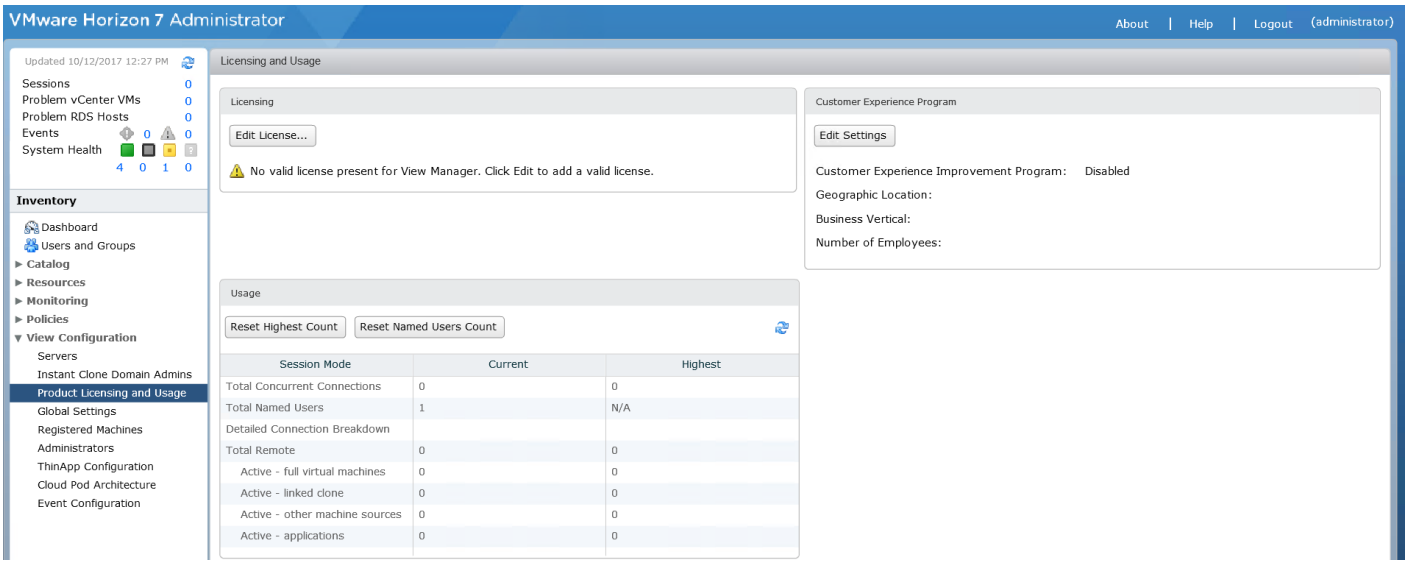

Configure vCenter

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.

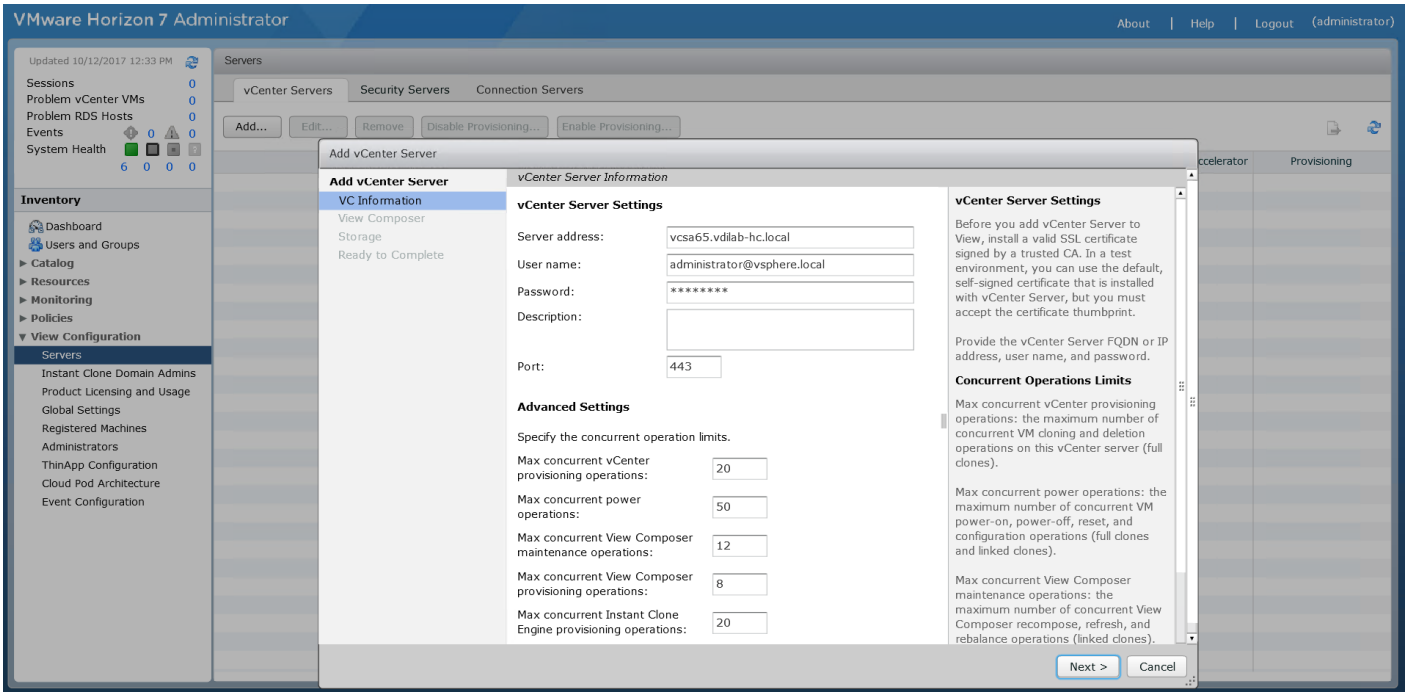

4. Click View certificate. Accept the certificate.

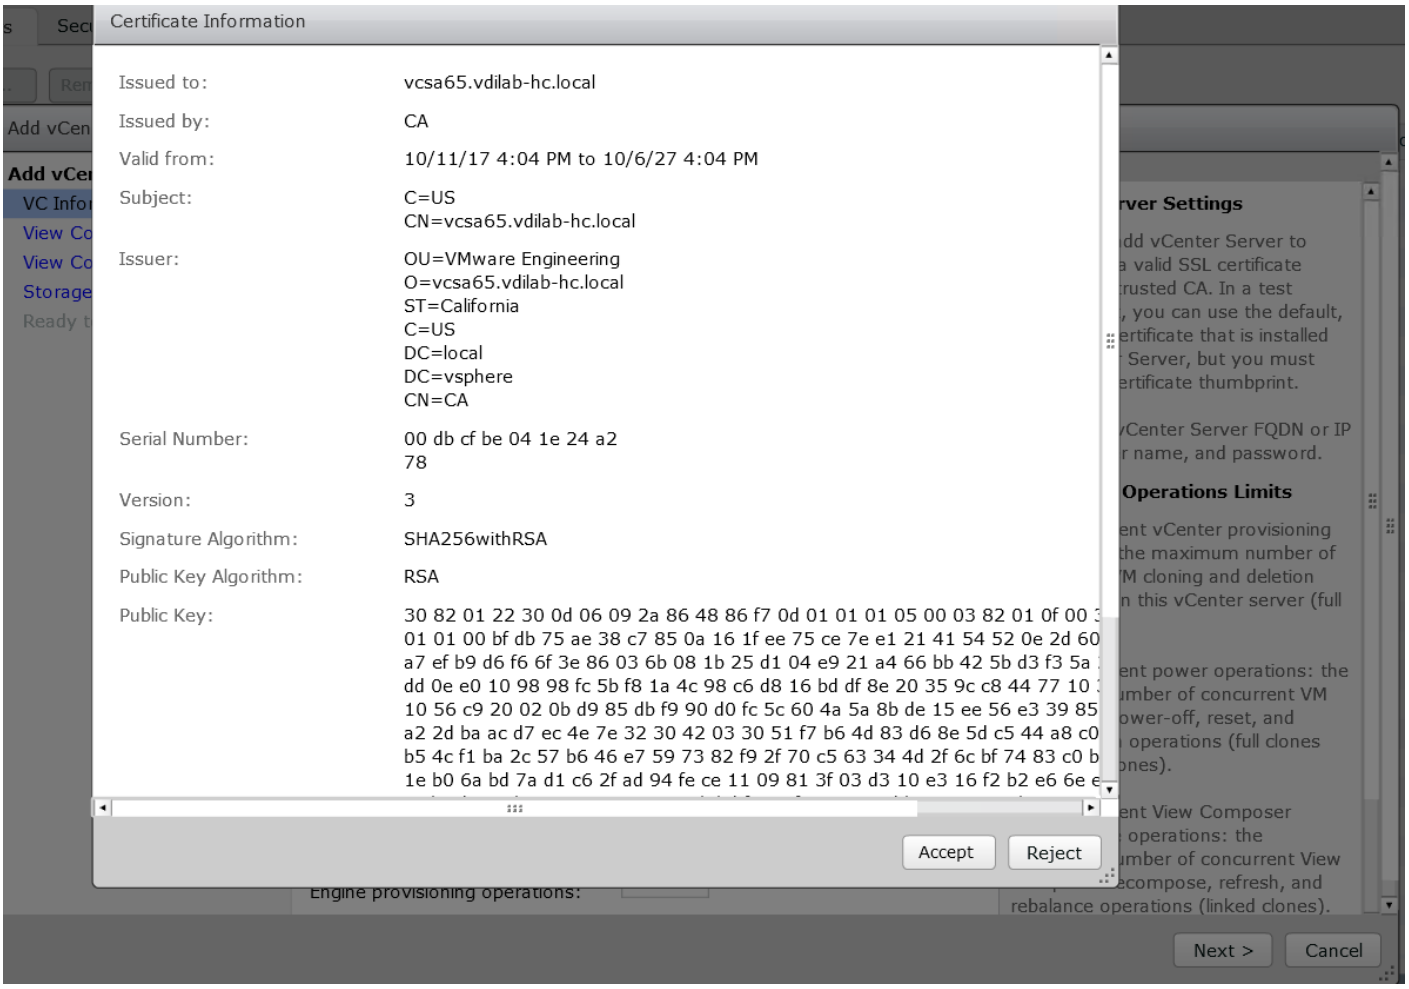

5. Add View composer settings, View composer server FQDN or Management IP address, login credentials. Click Next.

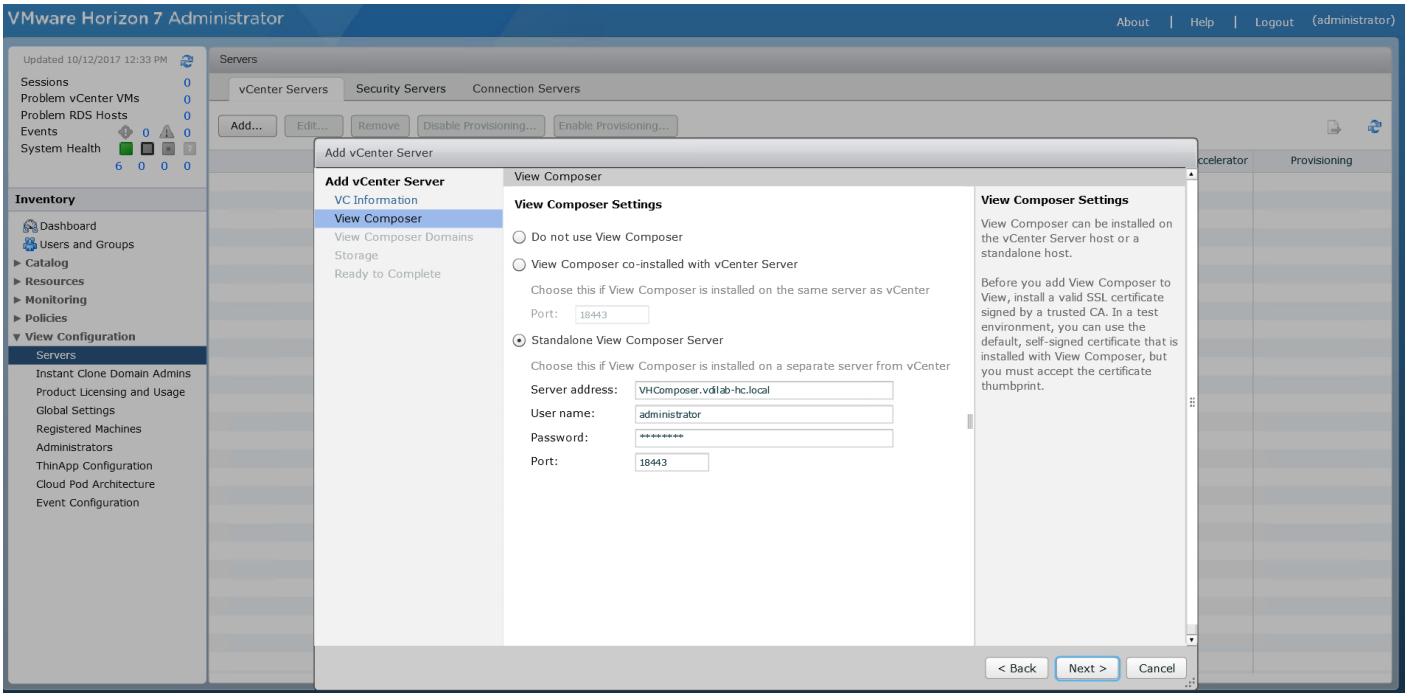

6. View and accept the certificate.

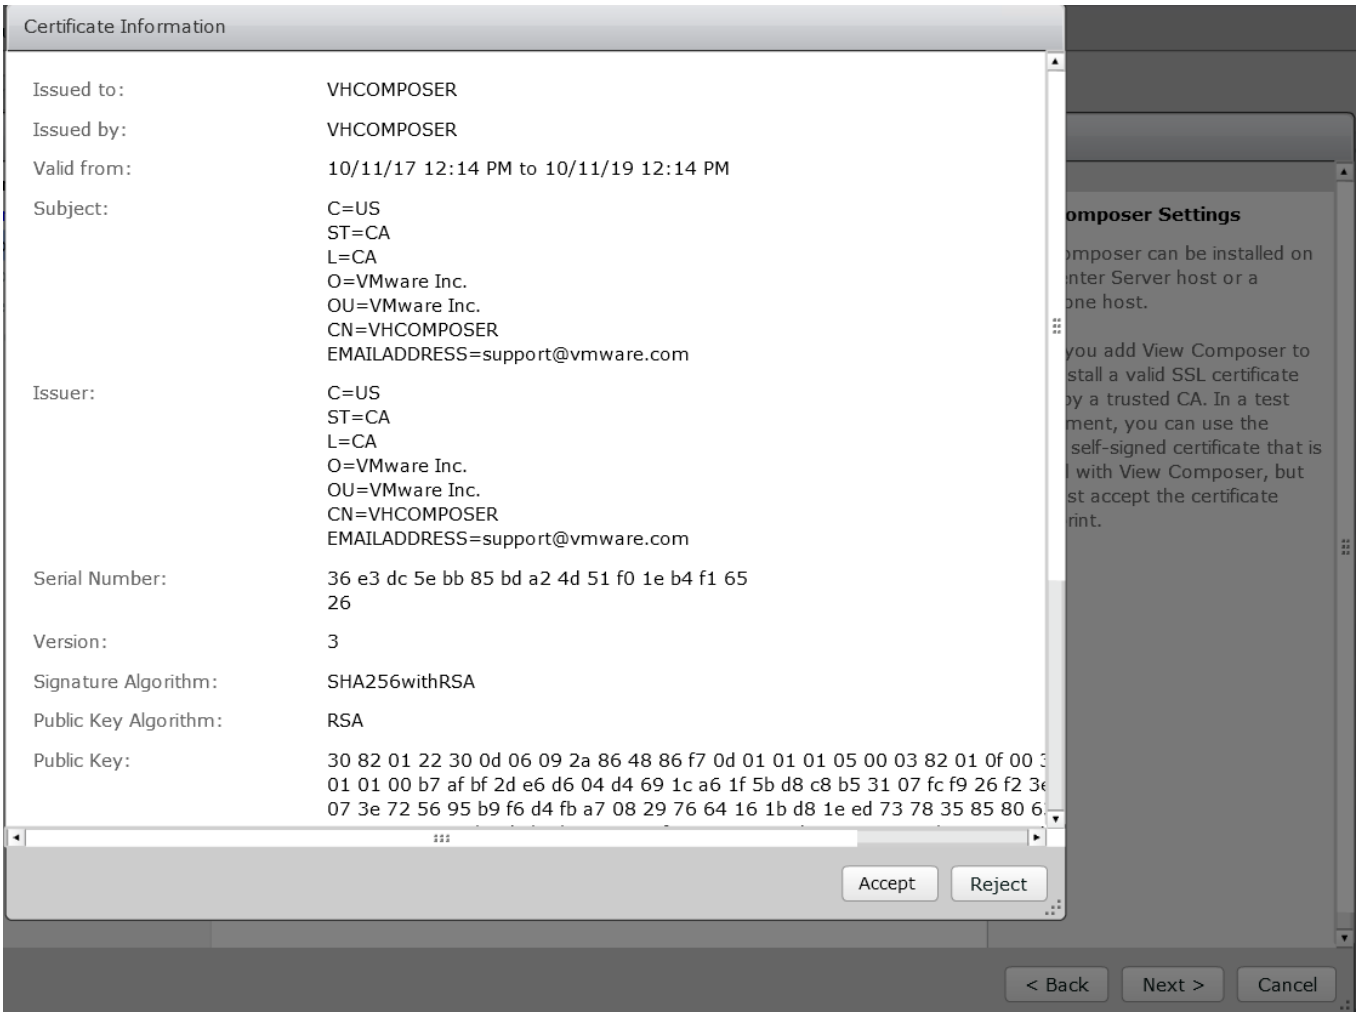

7. Click Add a new domain or Edit the existing domain.

8. Click Next.

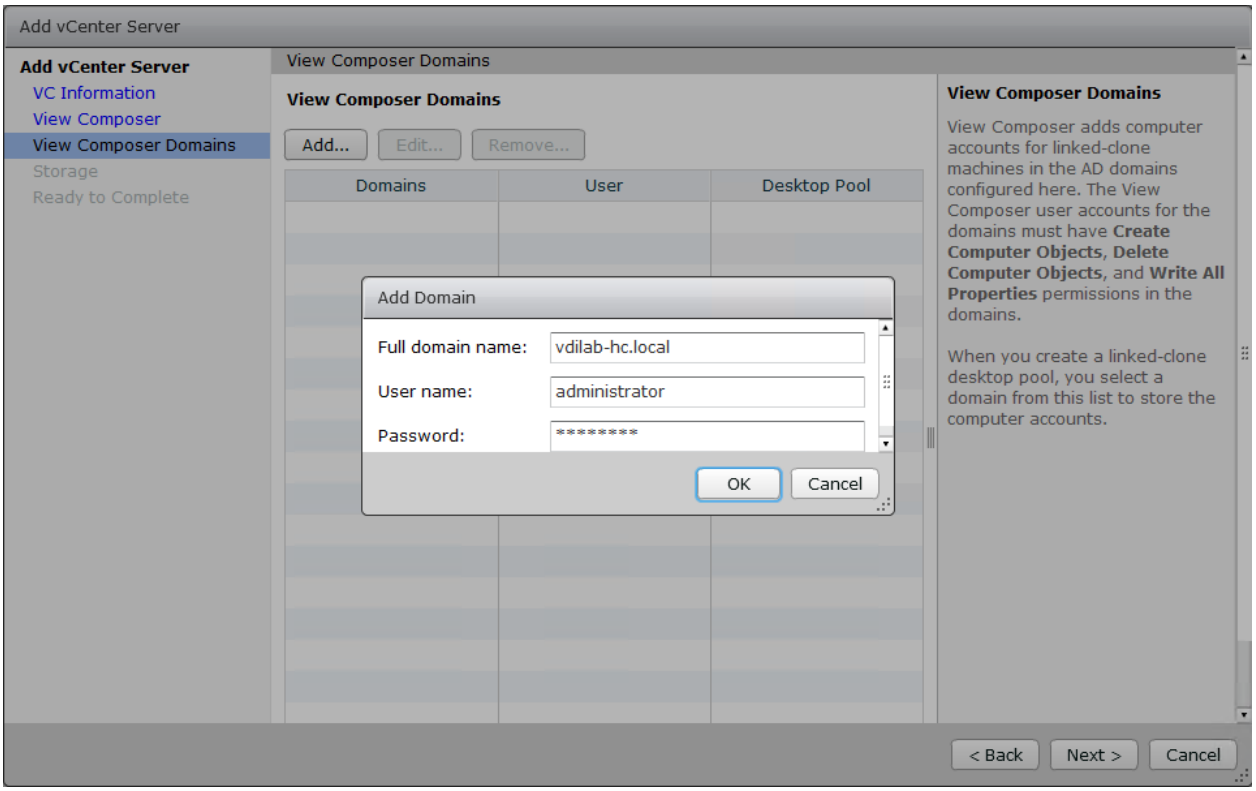

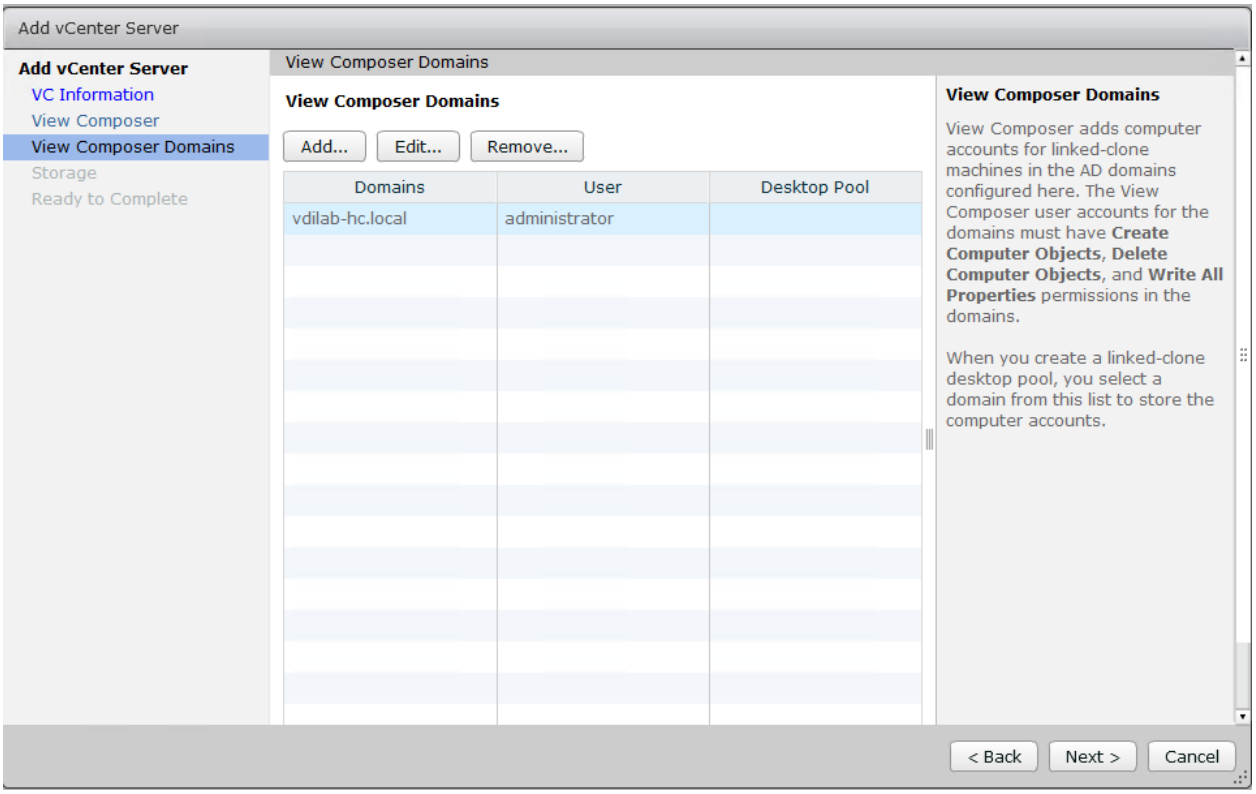

9. In Storage settings, select Reclaim VM disk space and View Storage Accelarator.

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

11. Click Next.

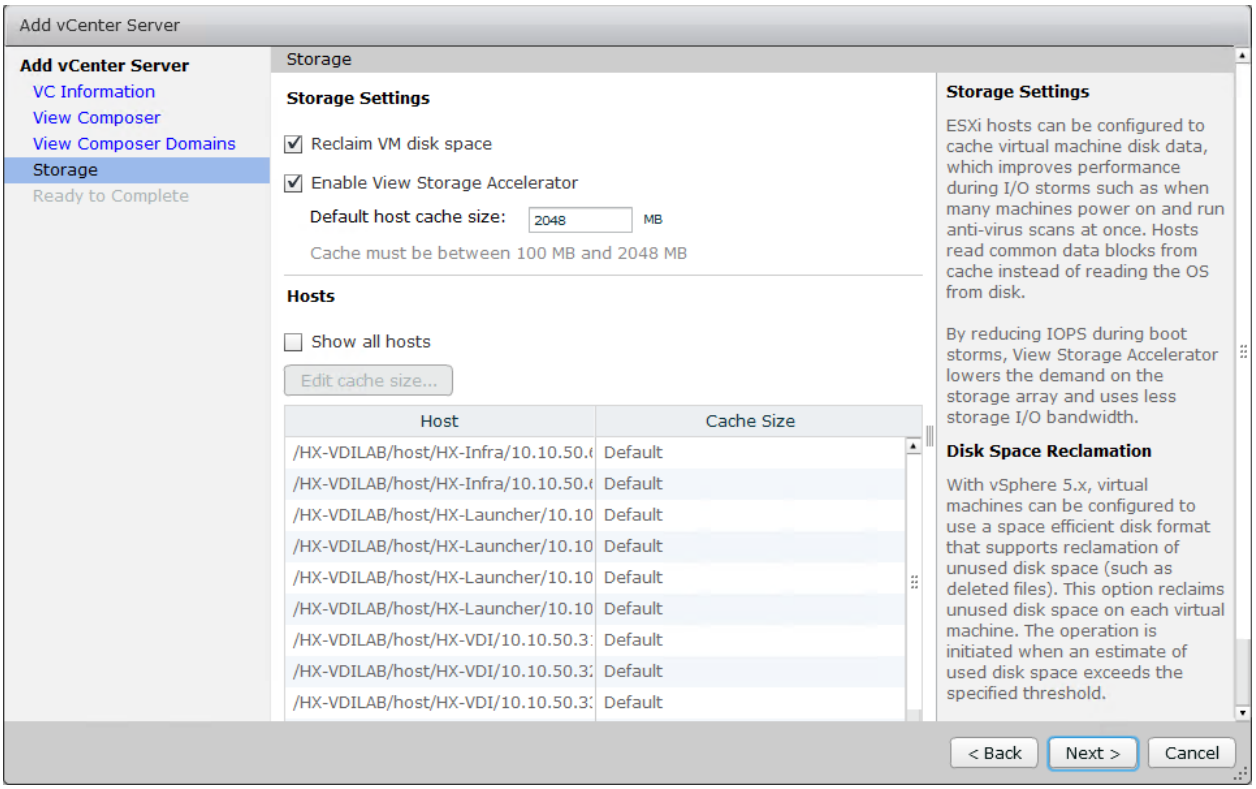

12. Review Add vCenter Server settings and click Finish.

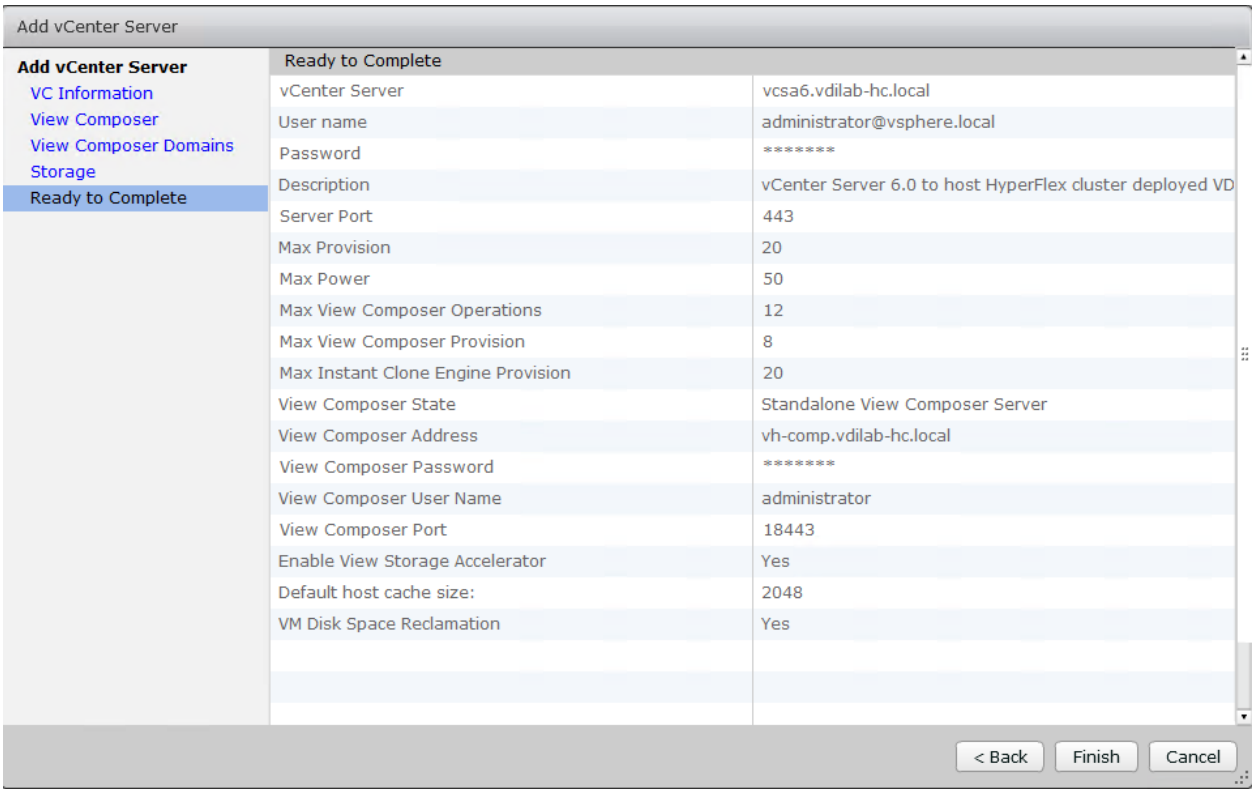

Configure Instant Clone Domain Admins

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

1. Under View Configuration, click Instant Clone Domain Admins.

2. Click Add. Enter credentials for domain user/group.

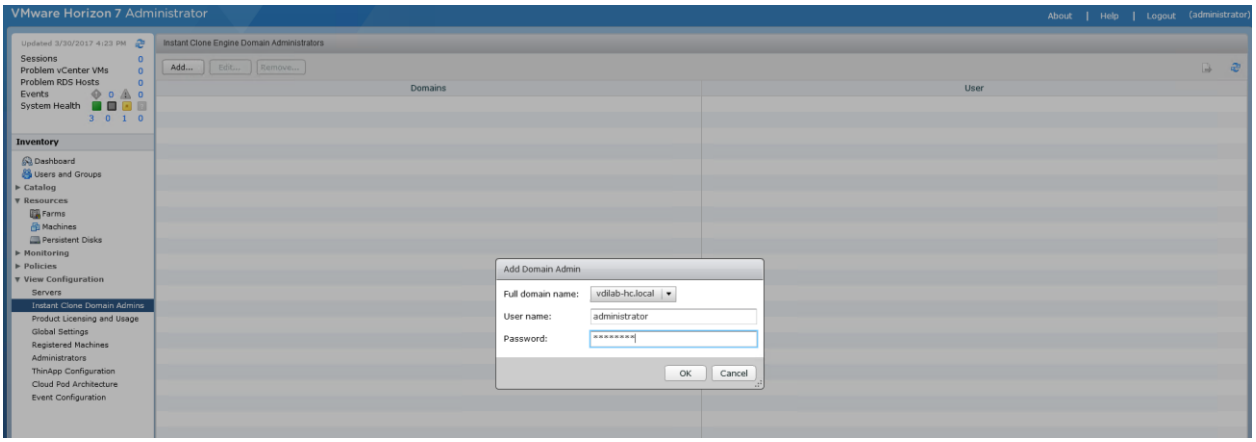

Master Image Creation for Tested Horizon Deployment Types

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 Linked-Clone, Instant Clone and Full Clone desktops.

We used a 64 bit version of OS and Office for our testing.

A fourth Golden Image was created using Microsoft Windows Server 2016 for RDSH session host virtual machines.

For the Golden Image virtual machines, the following parameters were used (Table 10).

Table 10 Golden Image Virtual Machine Parameters

Attribute	Linked-Clone/Instant-clone	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-	6 GB	-	-

Attribute	Linked-Clone/Instant-clone	Persistent/Full Clone	RDSH server
persistent disk for Linked-Clones)			
Virtual floppy drive 1	Removed	Removed	Removed
Virtual CD/DVD drive 1	–	–	-
Applications	• Login VSI 4.1.25 application installation • Adobe Acrobat 11 • Adobe Flash Player 16 • Doro PDF 1.82 • FreeMind • Microsoft Internet Explorer • Microsoft Office 2016	• Login VSI 4.1.25 application installation • Adobe Acrobat 11 • Adobe Flash Player 16 • Doro PDF 1.82 • FreeMind • Microsoft Internet Explorer • Microsoft Office 2016	• Login VSI 4.1.25 application installation • Adobe Acrobat 11 • Adobe Flash Player 16 • Doro PDF 1.82 • FreeMind • Microsoft Internet Explorer • Microsoft Office 2016
VMware tools	Release 10.1.7-5541682	Release 10.1.7-5541682	Release 10.1.7-5541682
VMware View Agent	Release 7.3.1-6761322	Release 7.3.1-6761322	Release 7.3.1-6761322

* For Persistent Desktops, we configured 4GB of RAM as amount of memory allocated is sufficient to run LoginVSI Knowledge Worker workload. HyperFlex nodes were configured with 768GB of total memory for this performance study. By adding memory to each HyperFlex node, for example twenty-four 64GB DIMMs per node, we could allocate up to 10GB of RAM per VM at the same user density.

Prepare Microsoft Windows 10 and Server 2016 with Microsoft Office 2016

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

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

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

Apply Microsoft updates 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 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-x86_64-7.3.1-6761322.exe. Click Next to install.

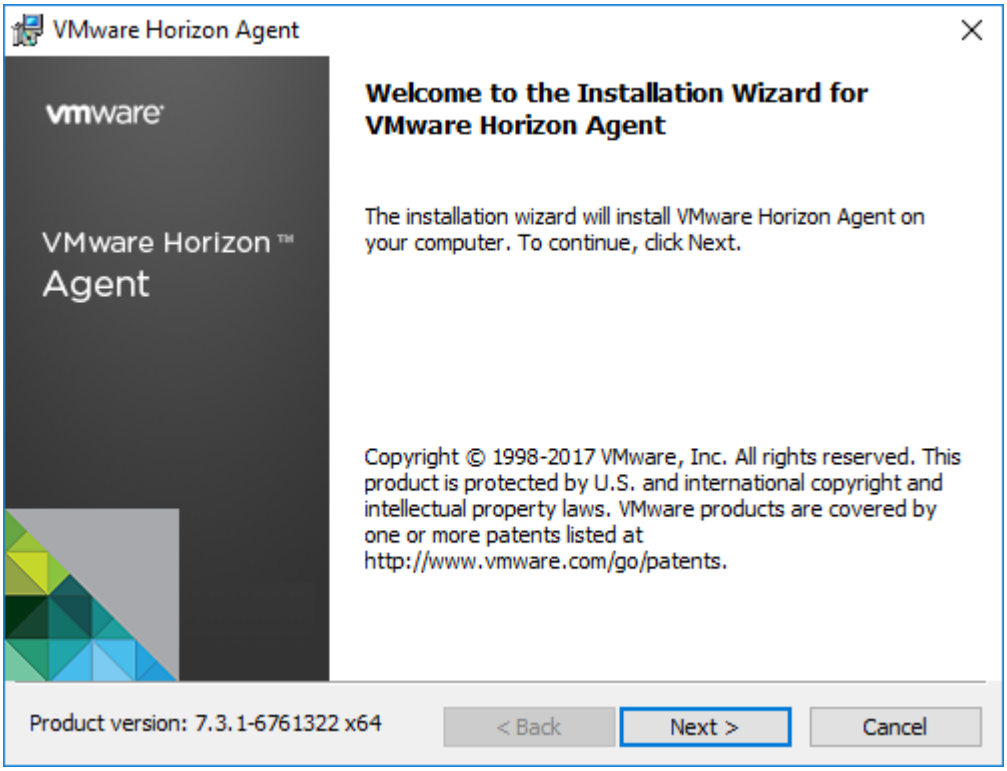

2. Review and accept the EULA Agreement. Click Next.

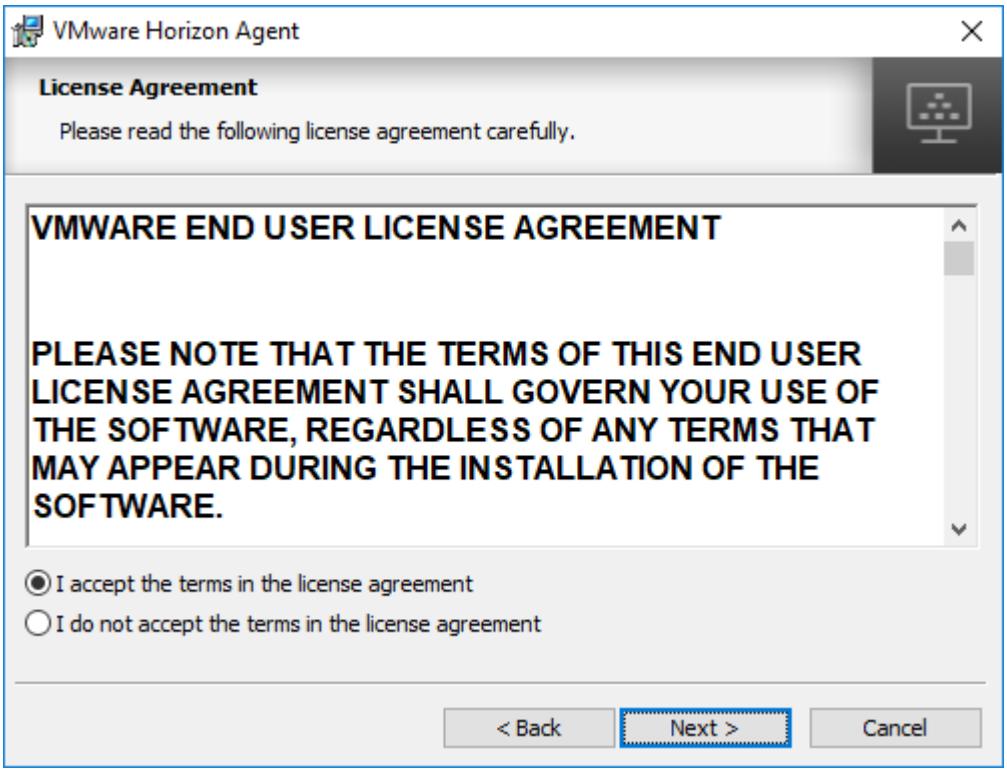

3. Select Network protocol configuration, click Next.

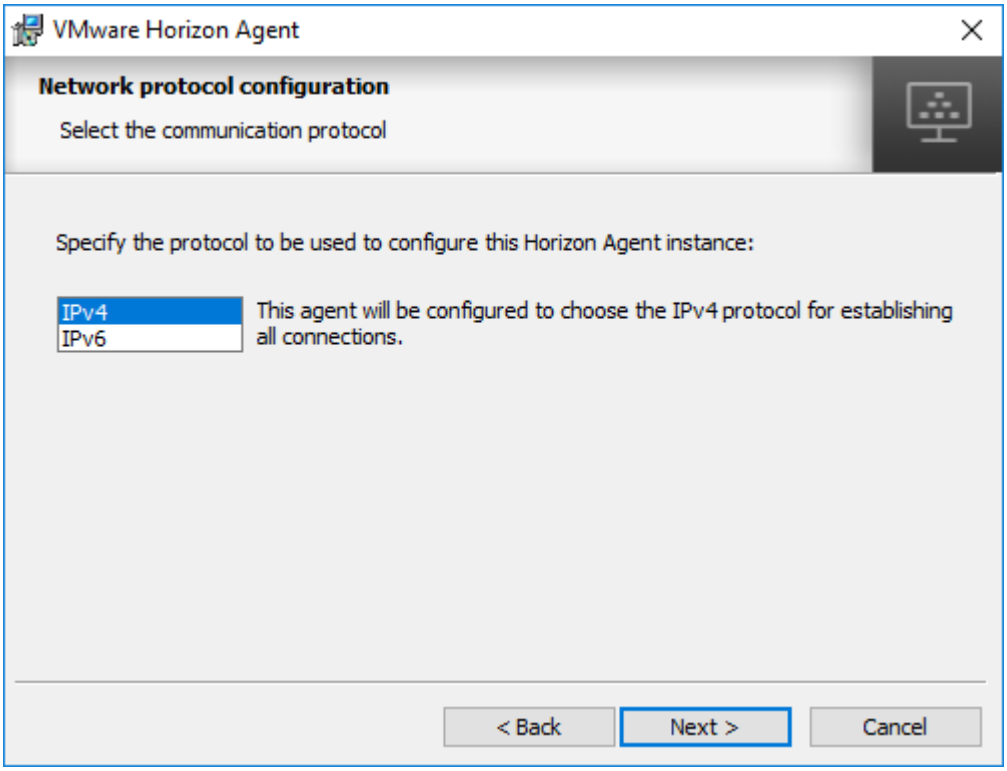

4. Based on the Desktop pool you want to create, select either View Composer Agent or Instant Clone Agent installation. Do not install both features on the same master image.

5. Enable installation of the VMware Horizon View Composer Agent for linked-clone VDI virtual machines.

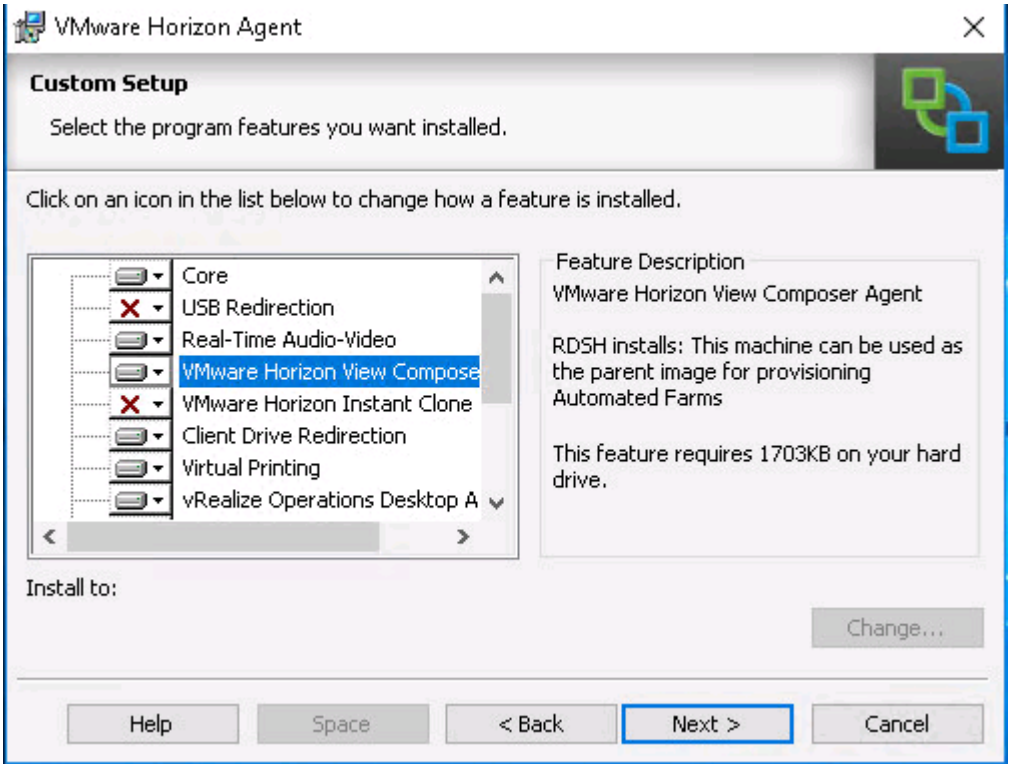

6. Disable the Horizon View Composer Agent and enable the Horizon Instant Clone Agent for Instant Clone floating assigned desktop pool creation.

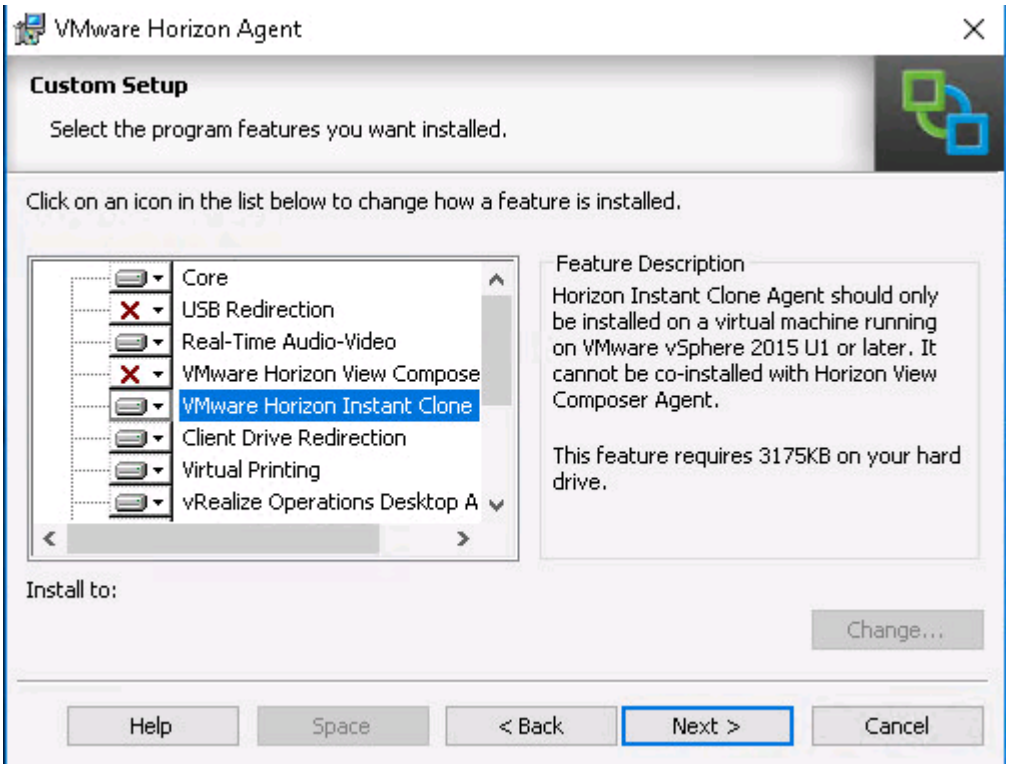

Prior to installing the Horizon View Agent on a Microsoft Server 2016 virtual machine, you must add the Remote Desktop Services role and the Remote Desktop Session Host role service.

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

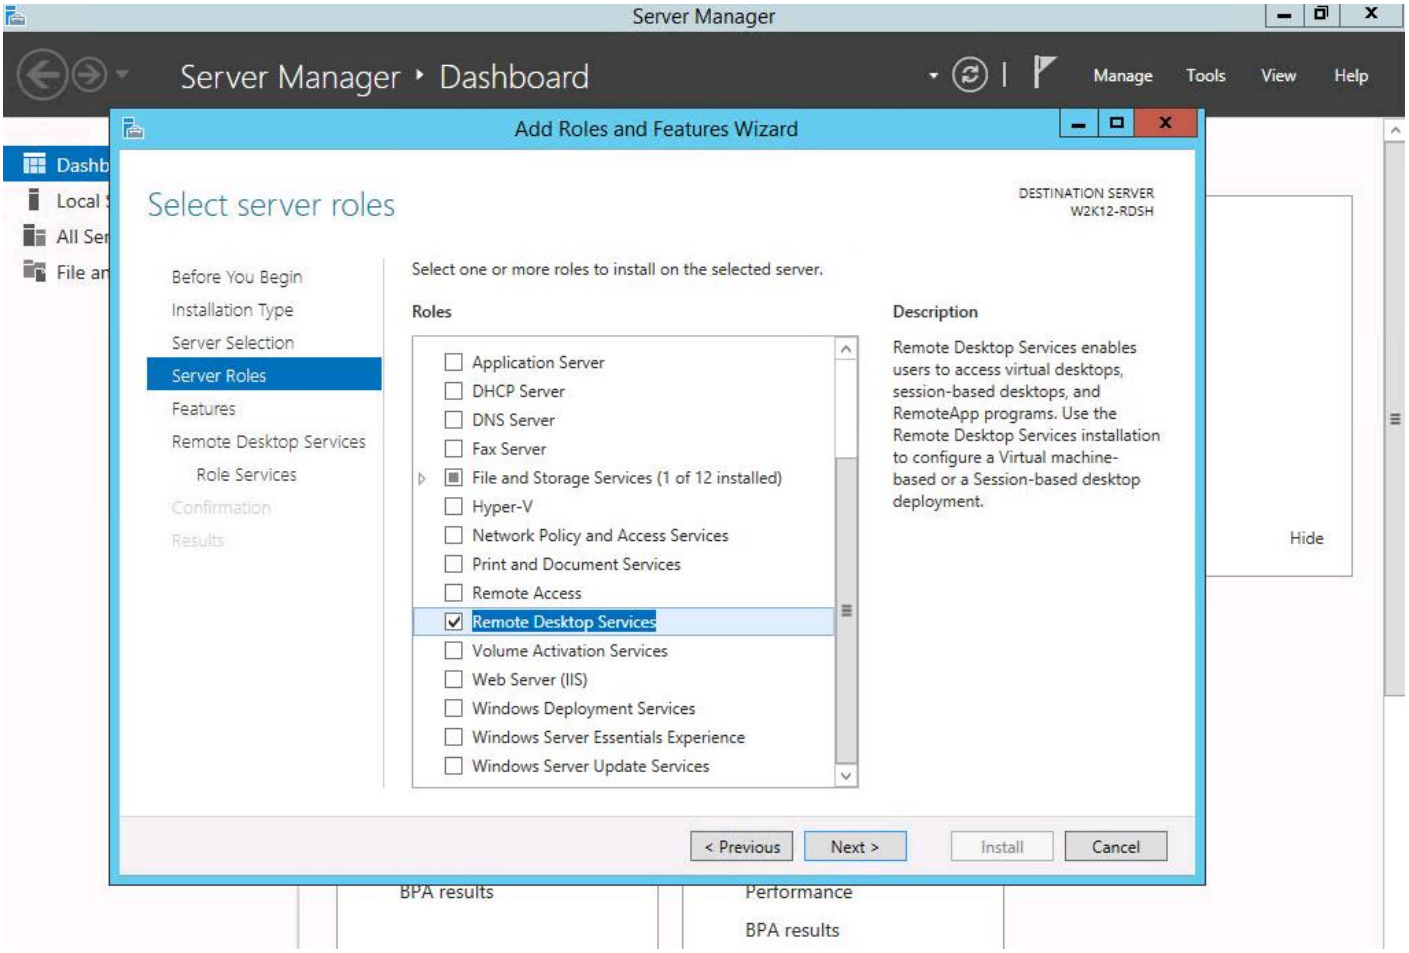

8. Add Remote Desktop Session Host services.

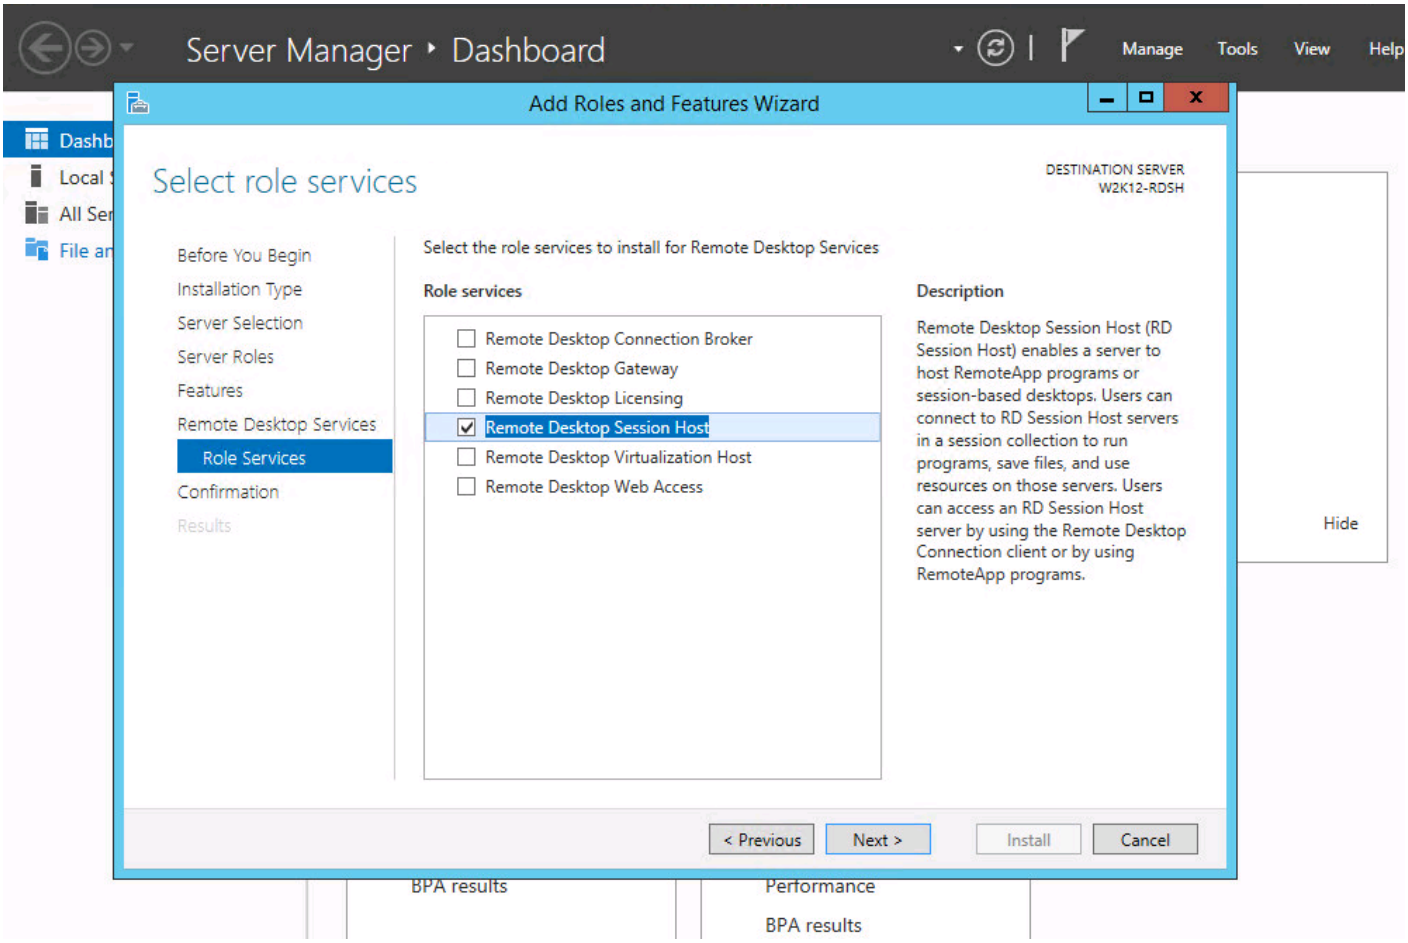

9. Click Install.

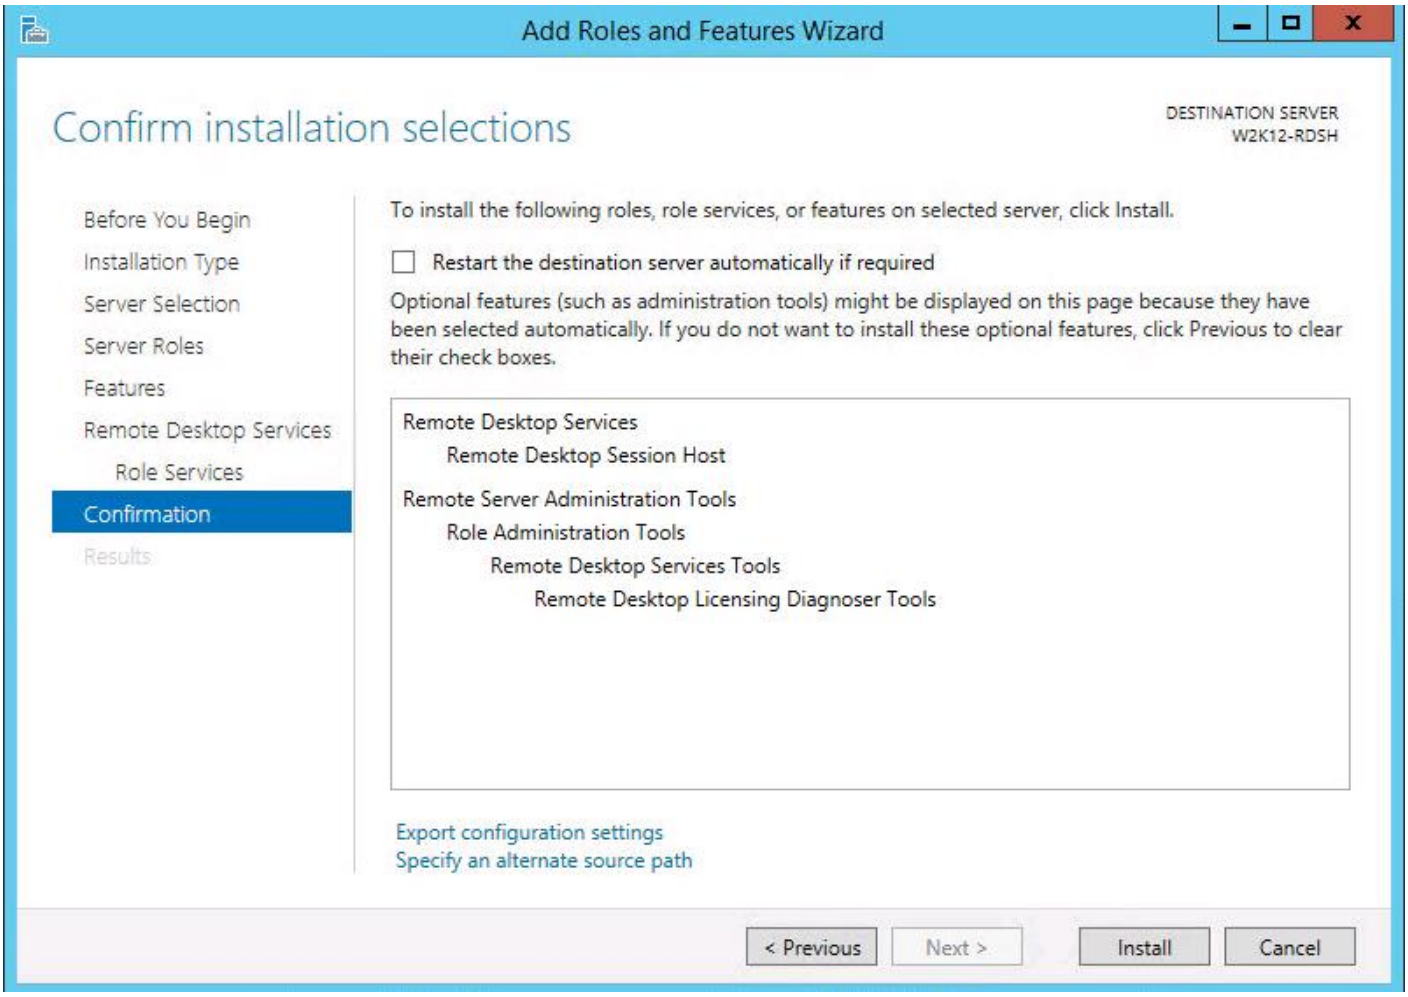
Add Roles and Features Wizard
Confirm installation selections
DESTINATION SERVER
W2K12-RDSH
Before You Begin
Installation Type
Server Selection
Server Roles
Features
Remote Desktop Services
Role Services
Confirmation
Results
To install the following roles, role services, or features on selected server, click Install.
Restart the destination server automatically if required
Optional features (such as administration tools) might be displayed on this page because they have been selected automatically. If you do not want to install these optional features, click Previous to clear their check boxes.
Remote Desktop Services
Remote Desktop Session Host
Remote Server Administration Tools
Role Administration Tools
Remote Desktop Services Tools
Remote Desktop Licensing Diagnoser Tools
Export configuration settings
Specify an alternate source path
< Previous
Next >
Install
Cancel

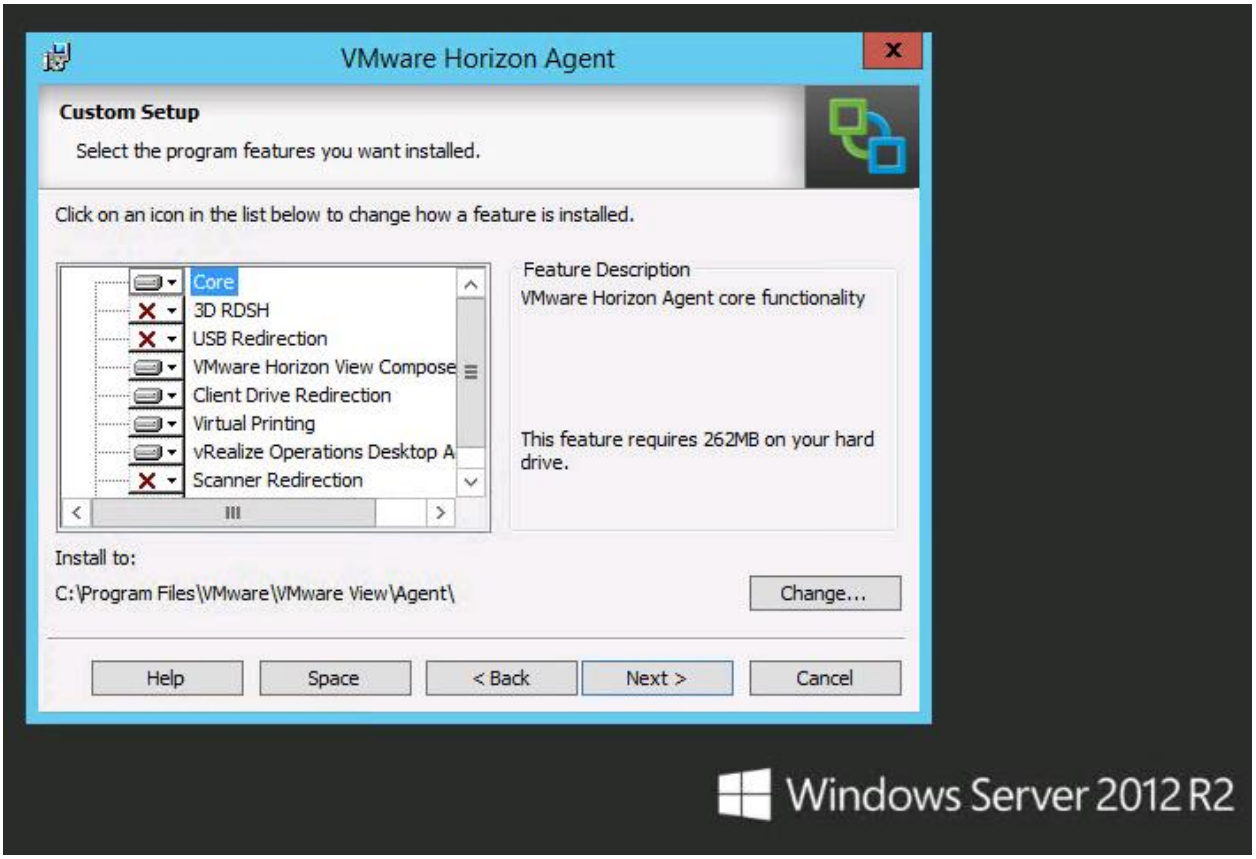

10. View Agent is will report as Install in “Desktop Mode” if Remote Desktop Services not installed.

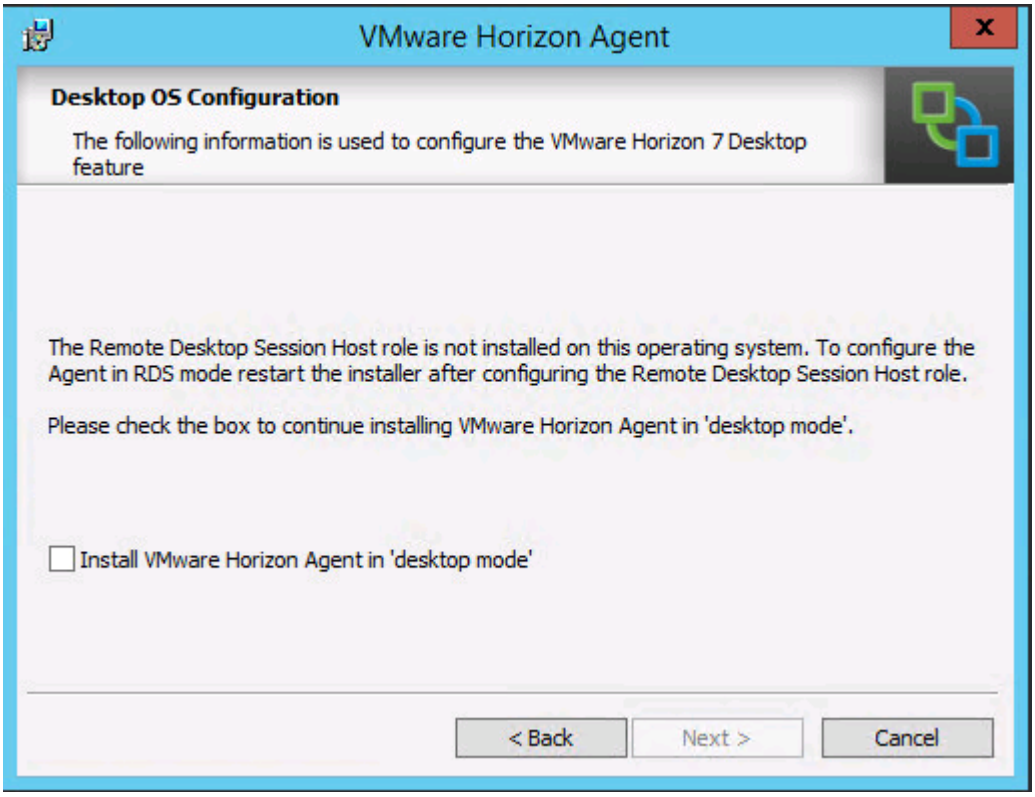

11. Add FQDN or IP address for Connection Server Instance to register the RDSH server.

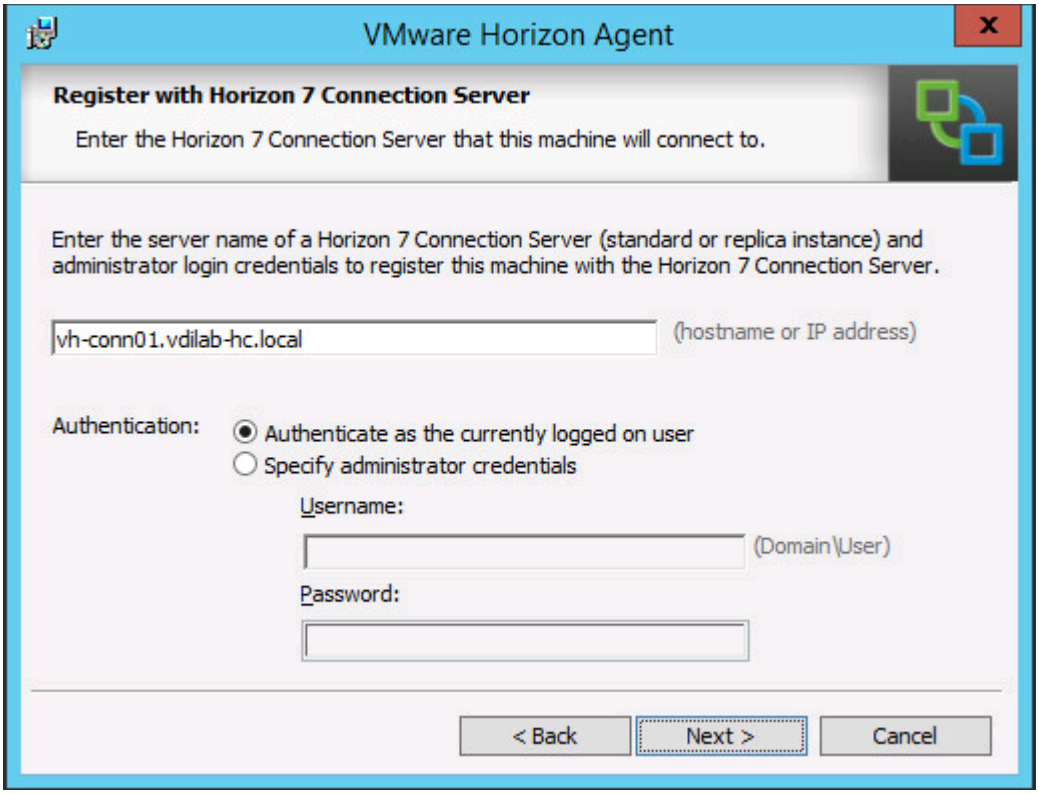

12. Click Install.

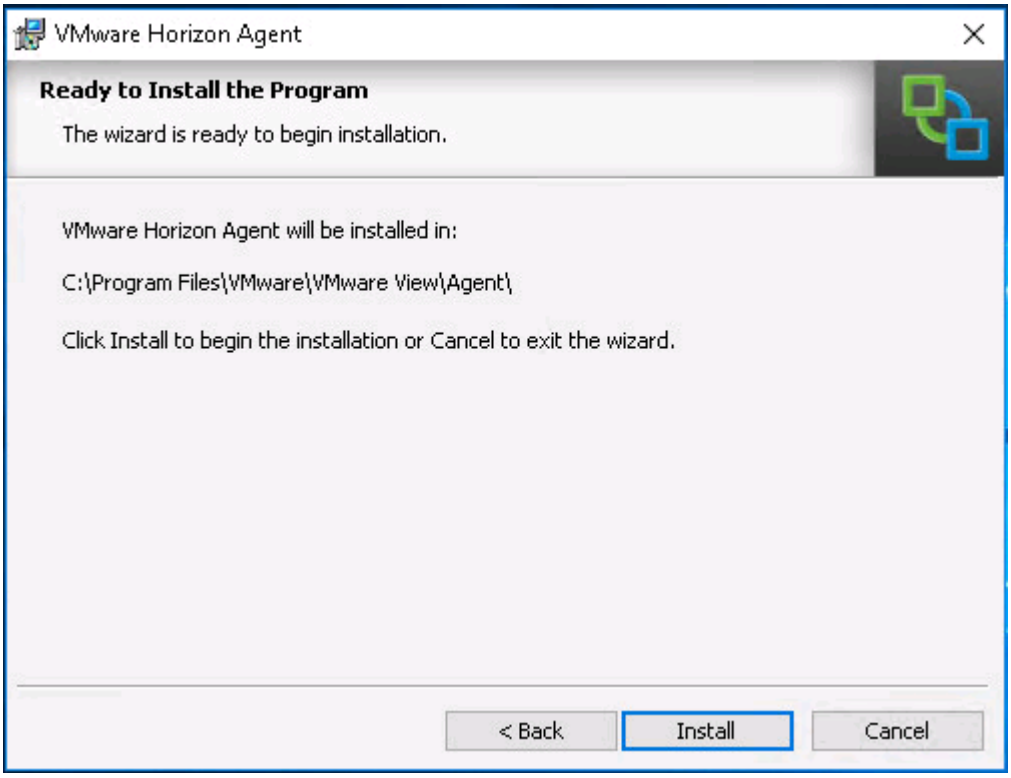

13. Click Finish and restart the VM.

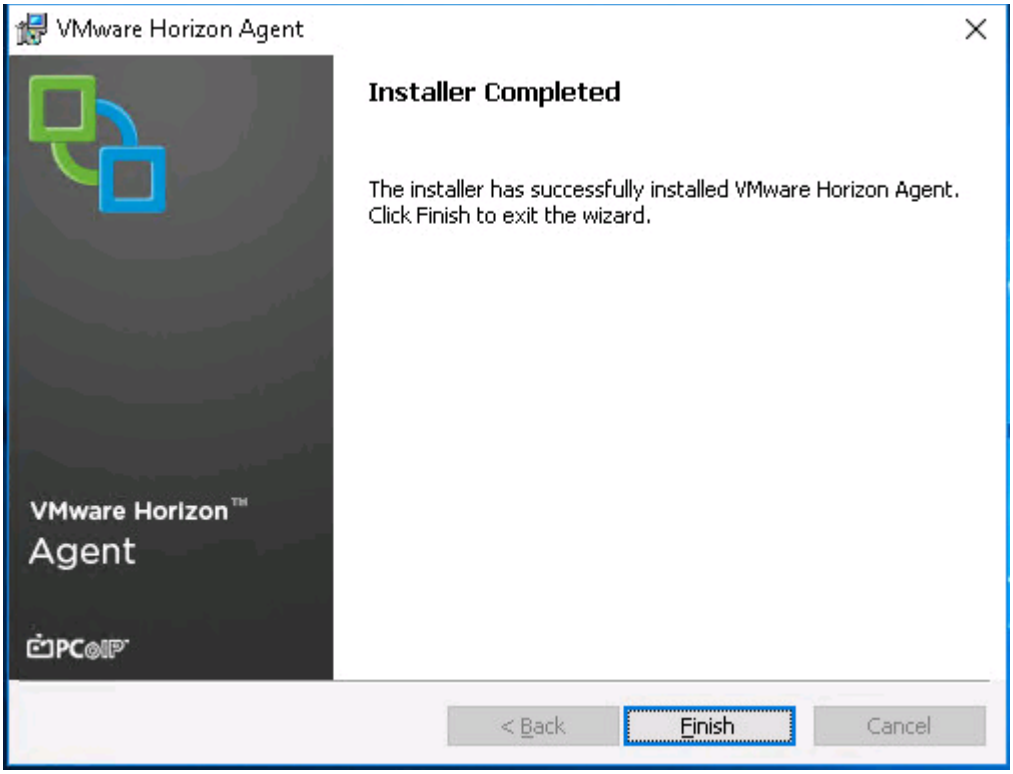

Install Additional Software

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

1. For testing, we installed Office 2016 64bit 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 were 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.

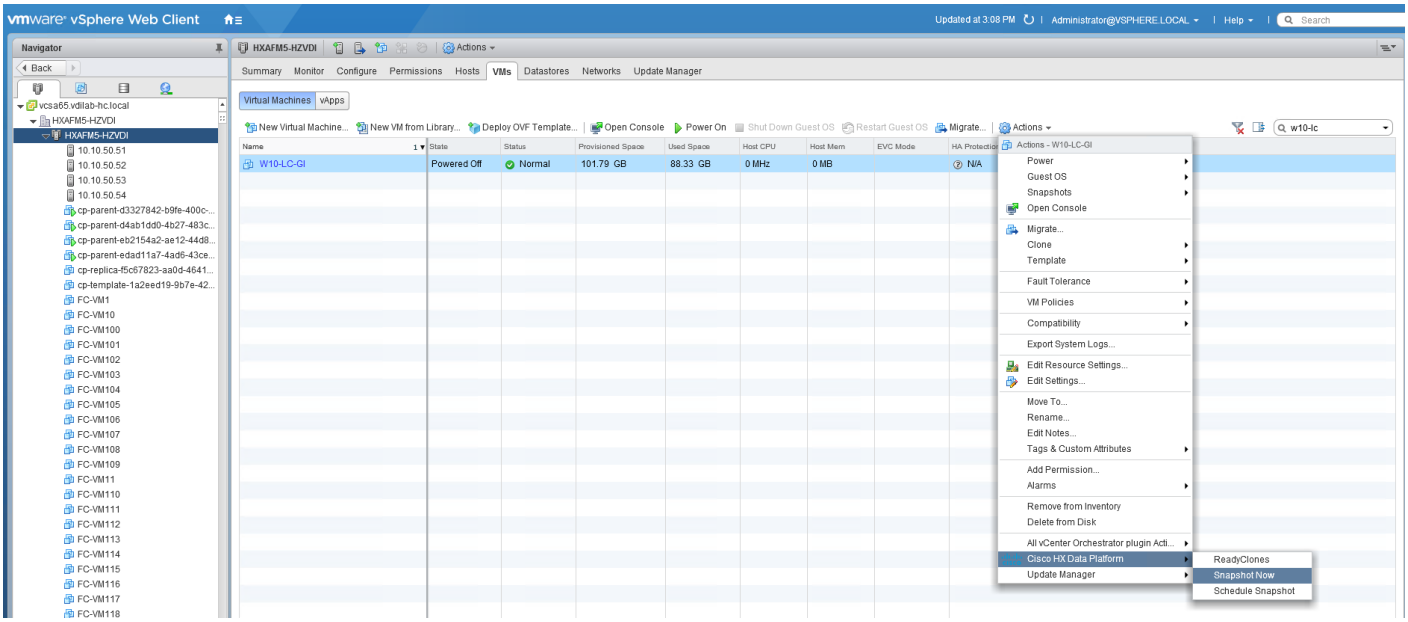

4. Enter a name for the HX native snapshot.

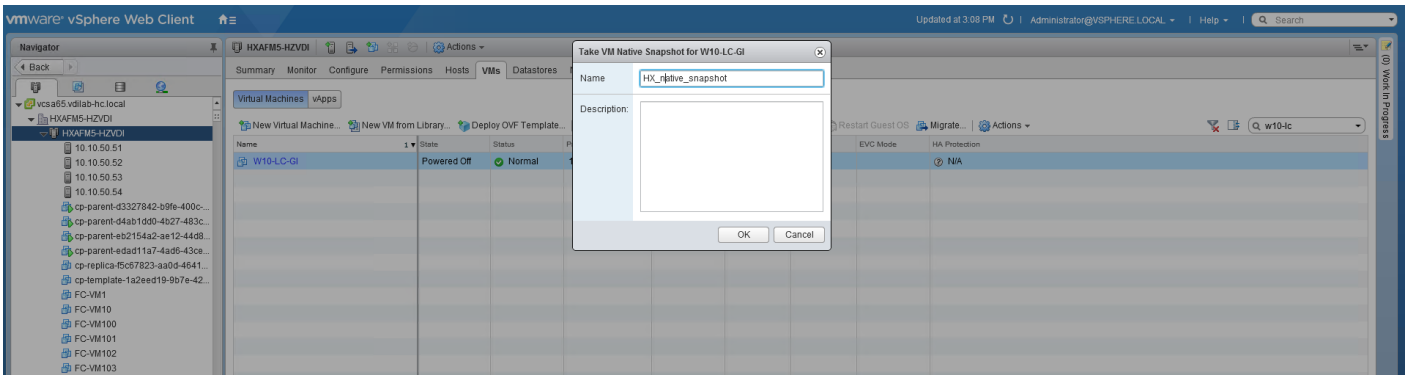

Create Customization Specification for Virtual Desktops

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

1. On vCenter WebUI, select Customization Specification Manager.

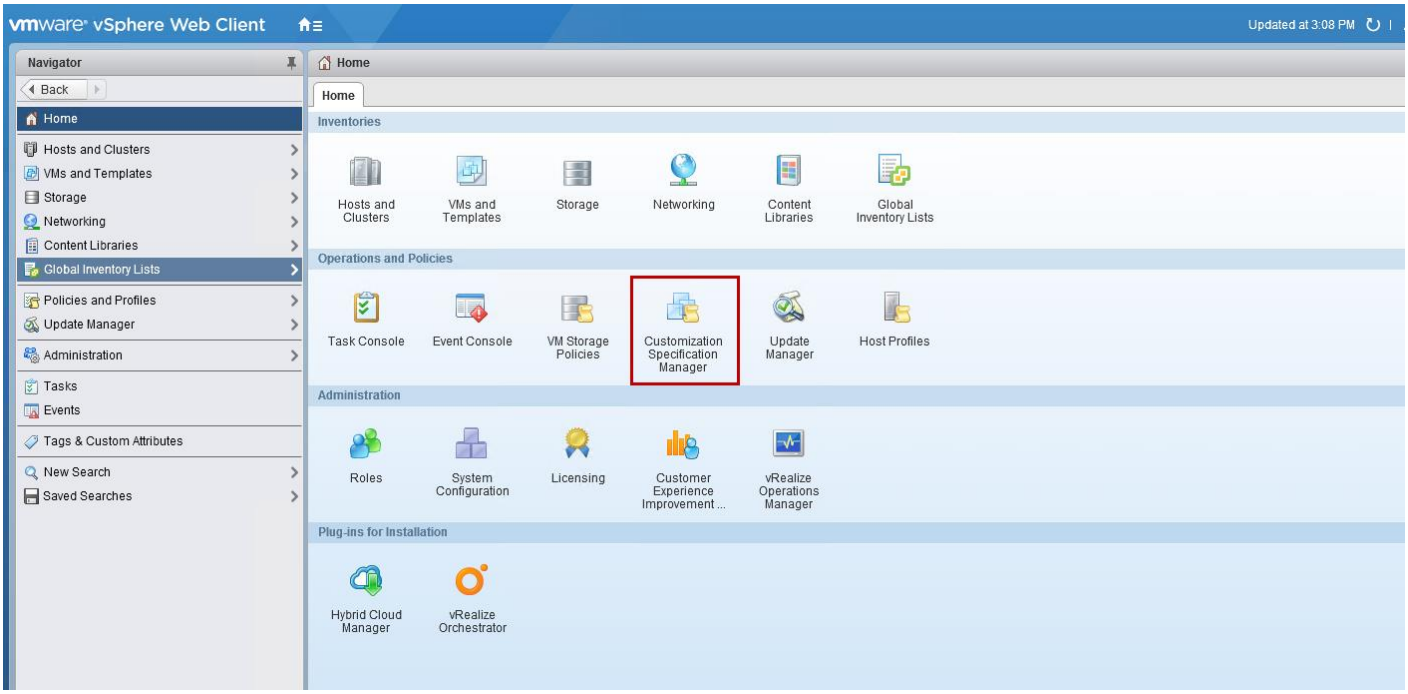

2. Select VM Operating System as Windows for Windows based guest OS optimization. Enter a name.

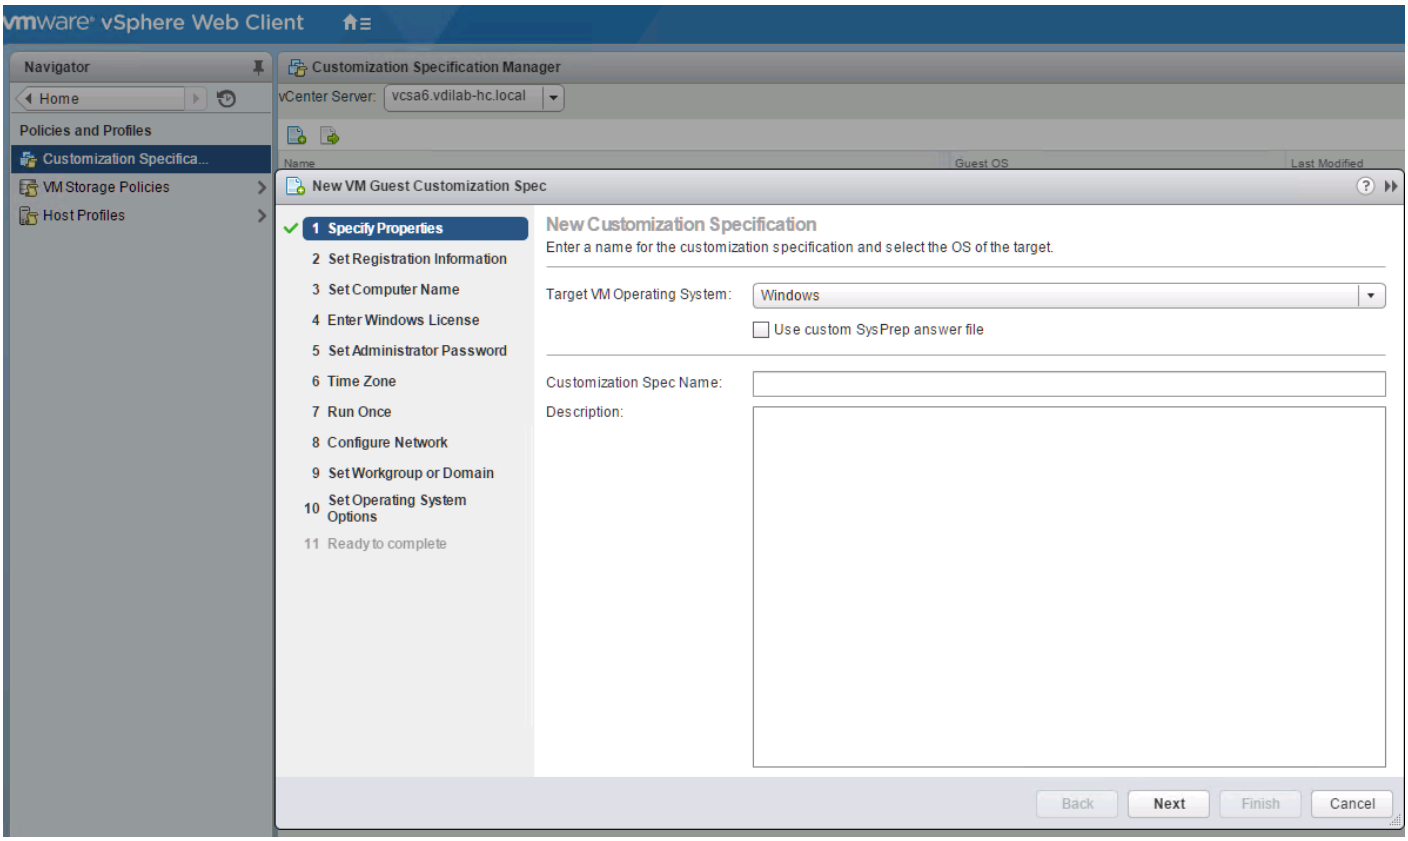

3. Provide name and organization details.

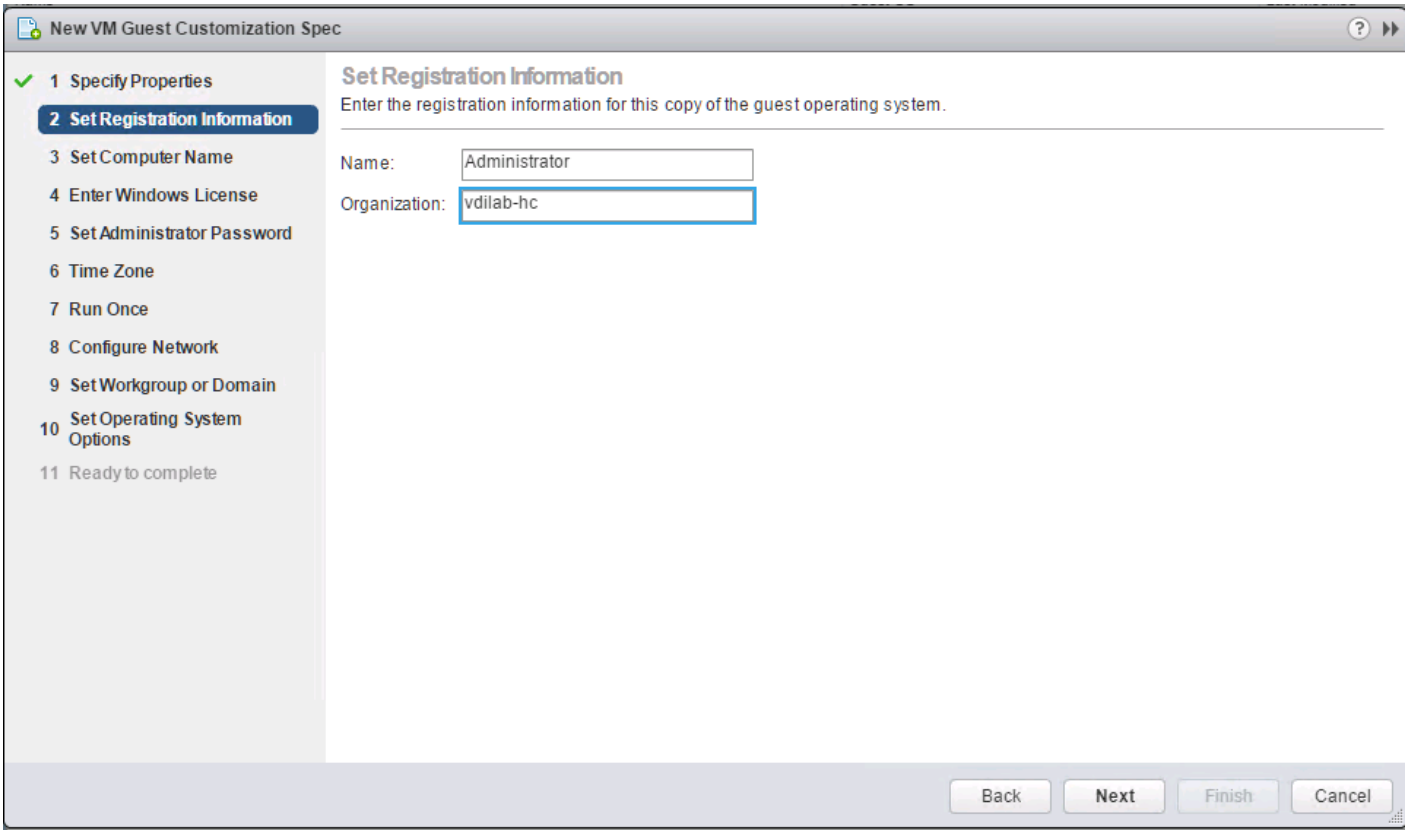

4. Provide a computer name. For this solution, we selected the radio button for Use the virtual machine name.

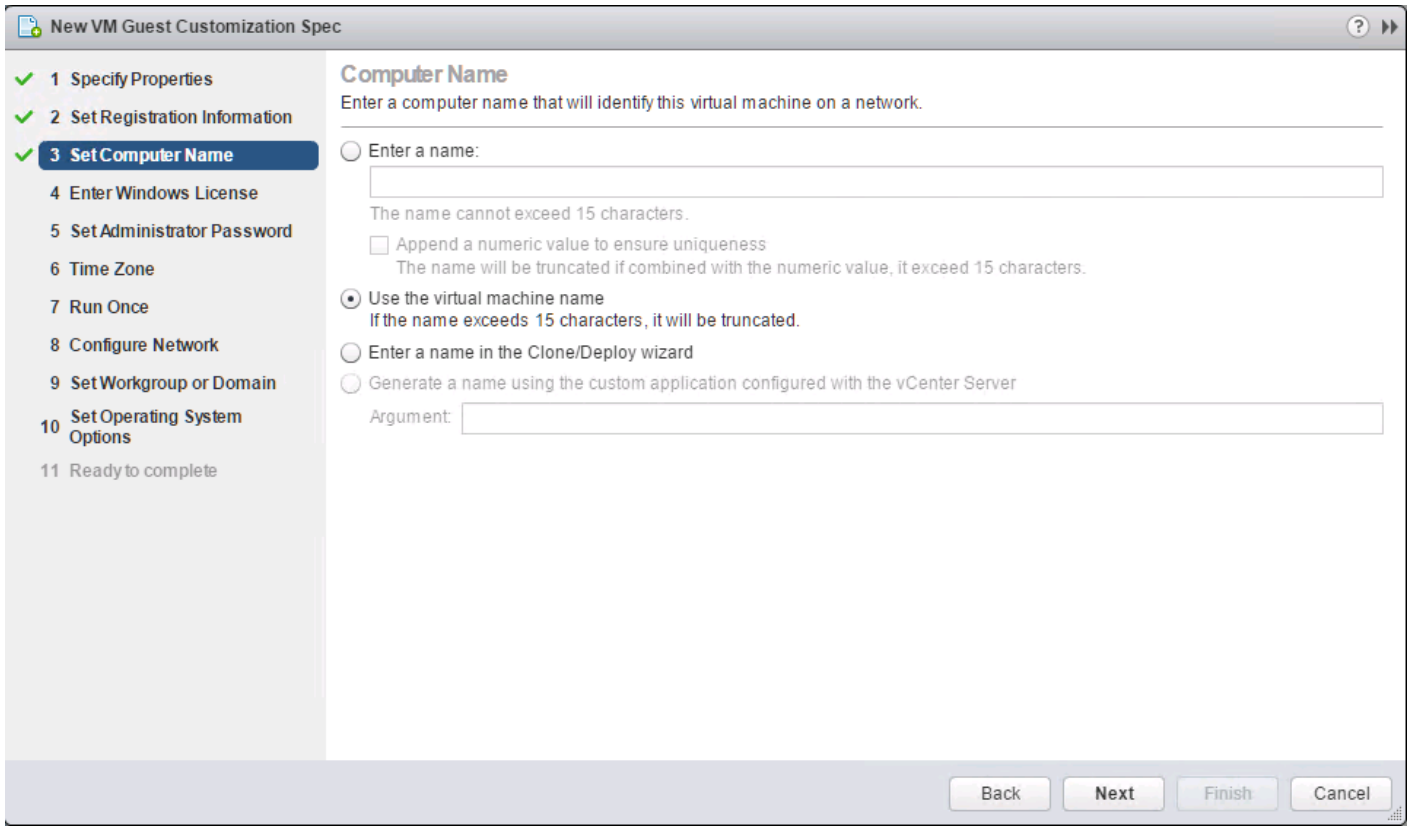

5. Provide the product License key if required.

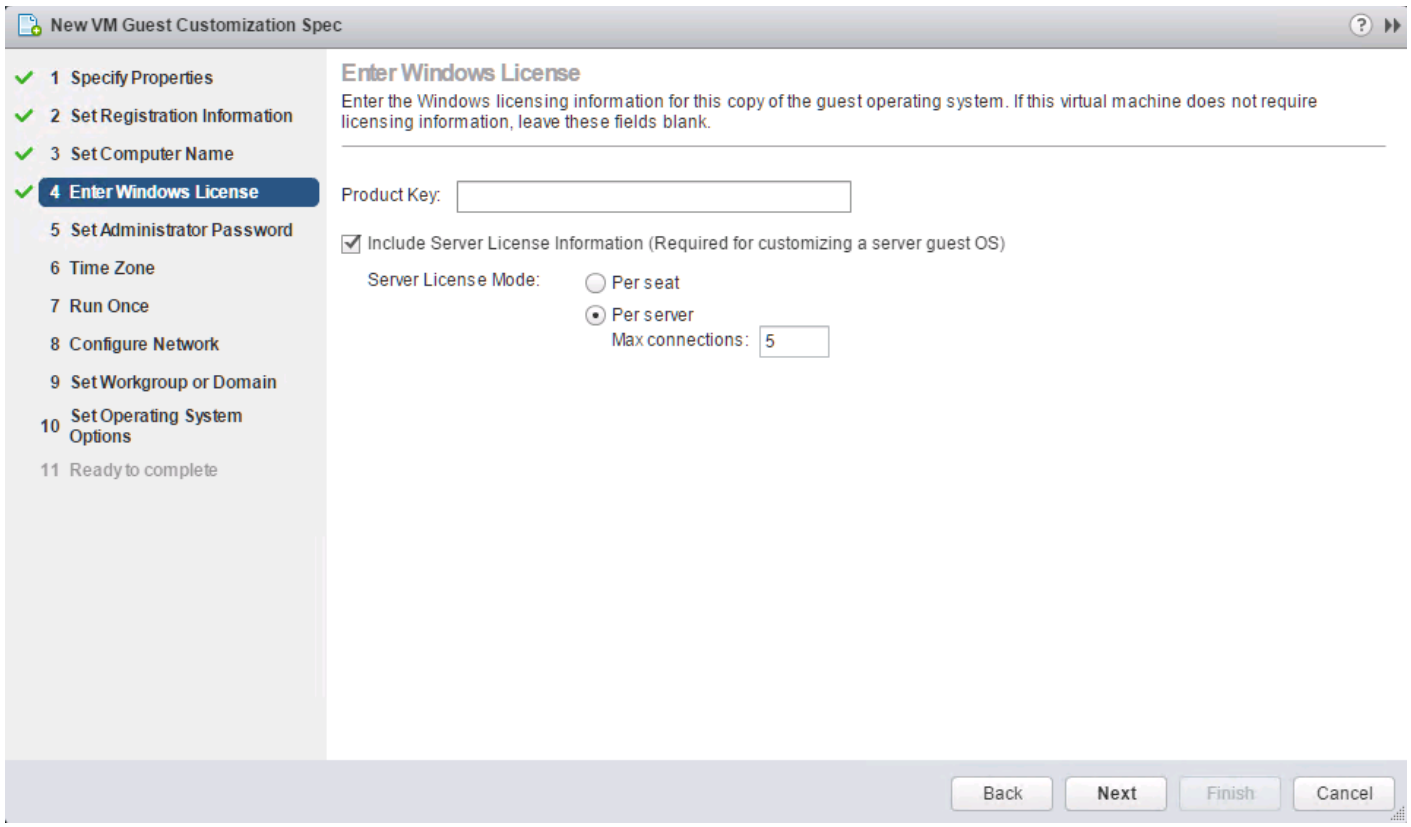

6. Provide Password credentials.

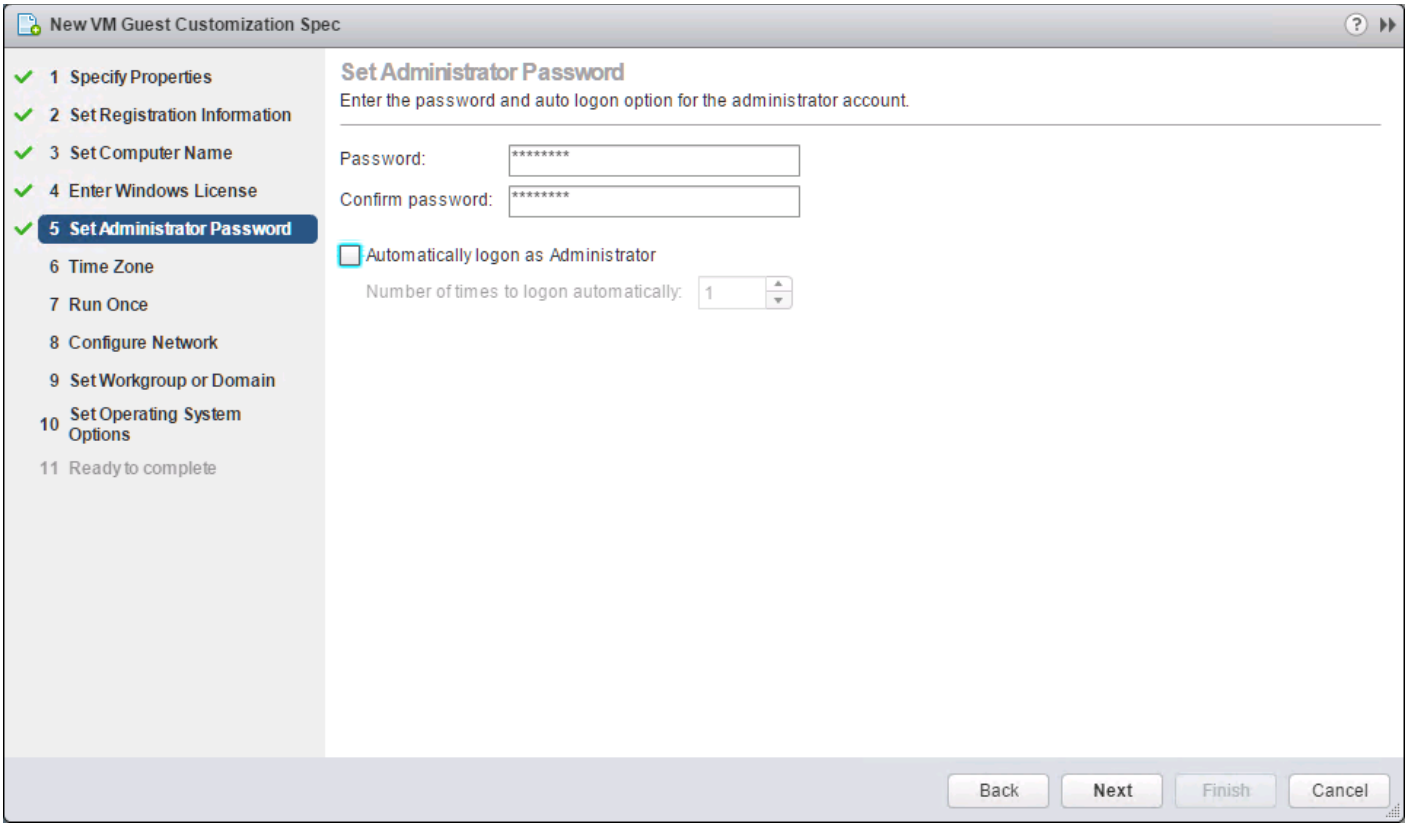

7. Select the Timezone.

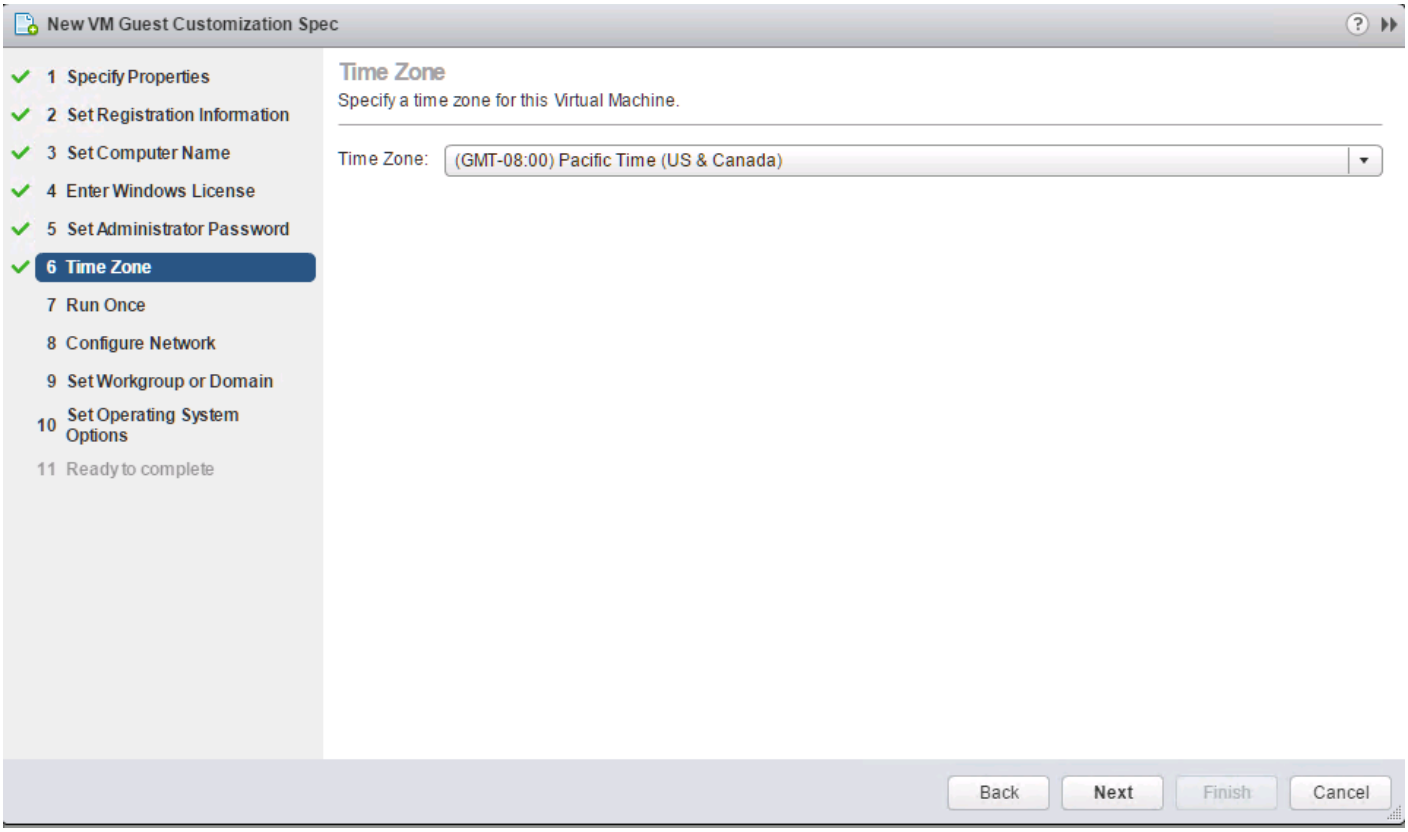

8. Add the commands to run when the first-time user logs in, if there are any.

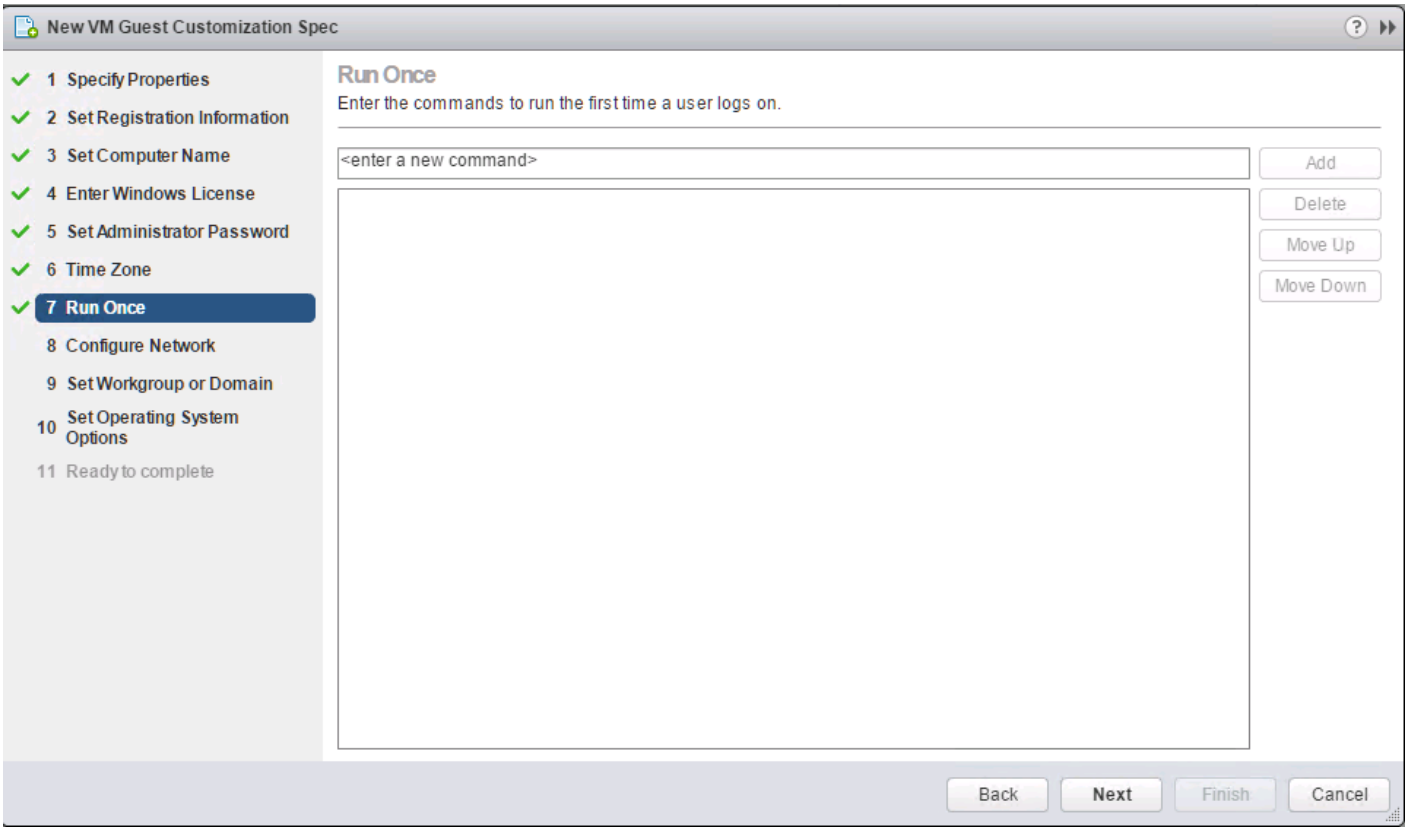

9. Provide the network information whether to use the DHCP server to assign IP address, or manual configuration.

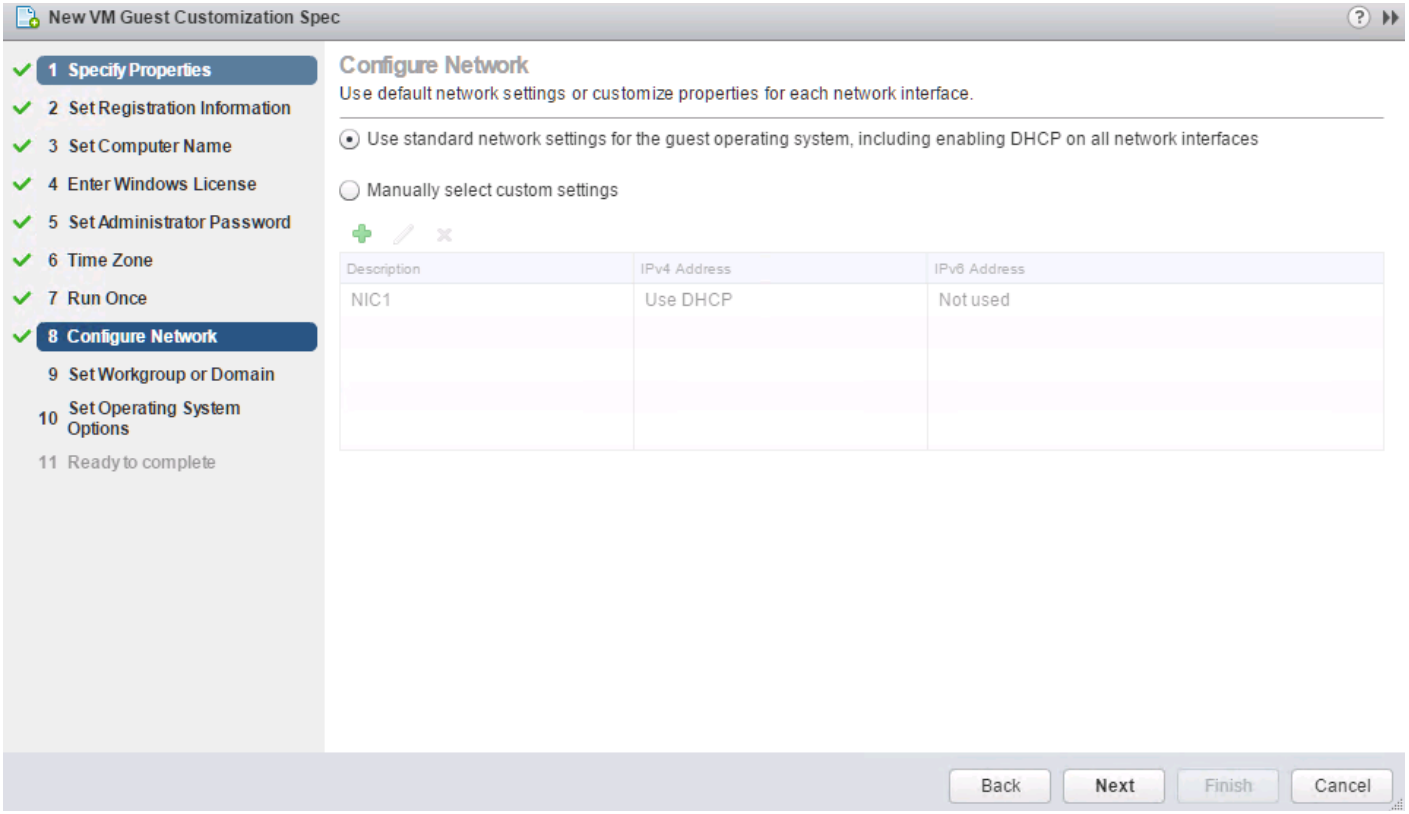

10. Provide the domain name and user credentials.

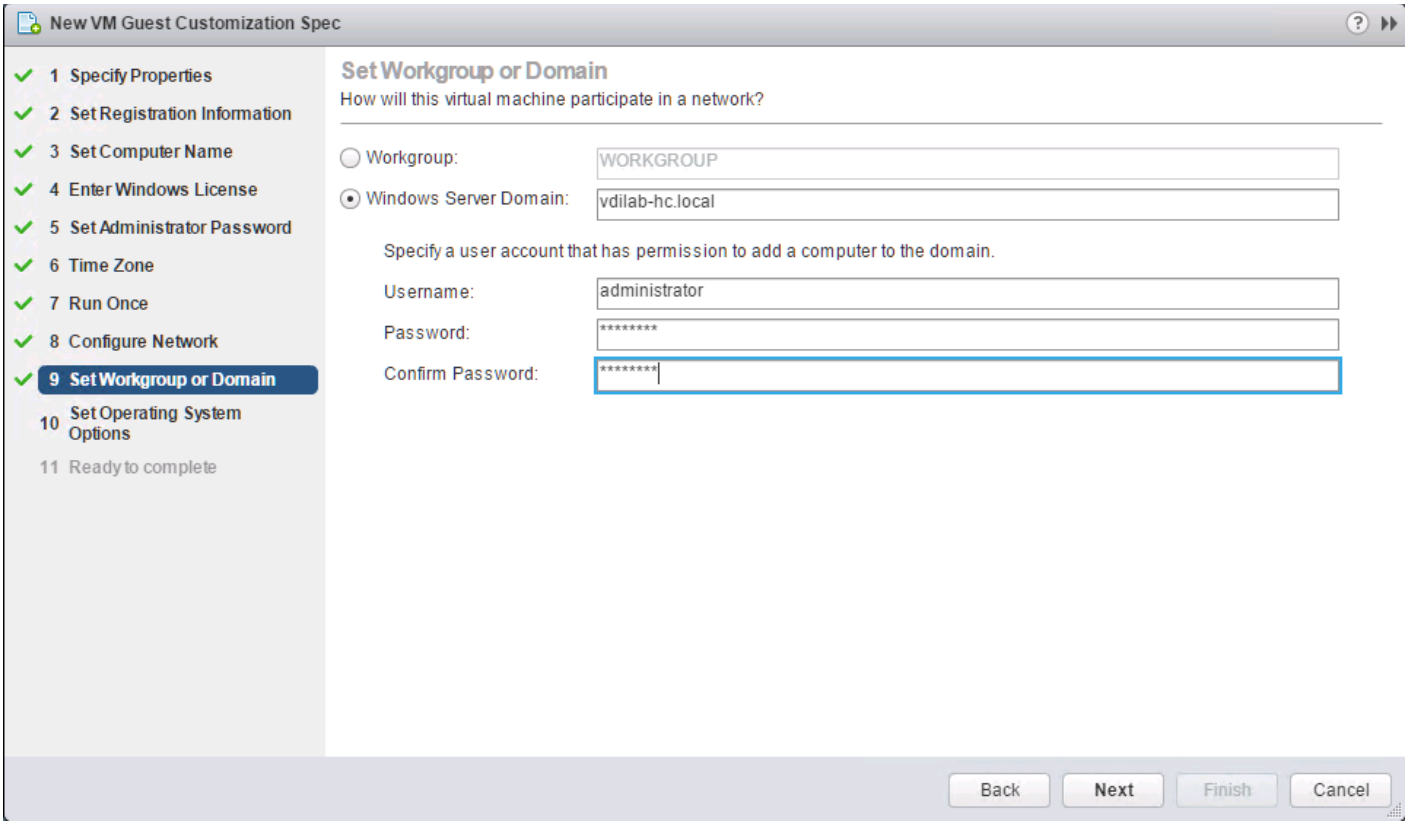

11. Select the checkbox Generate New Security ID (SID).

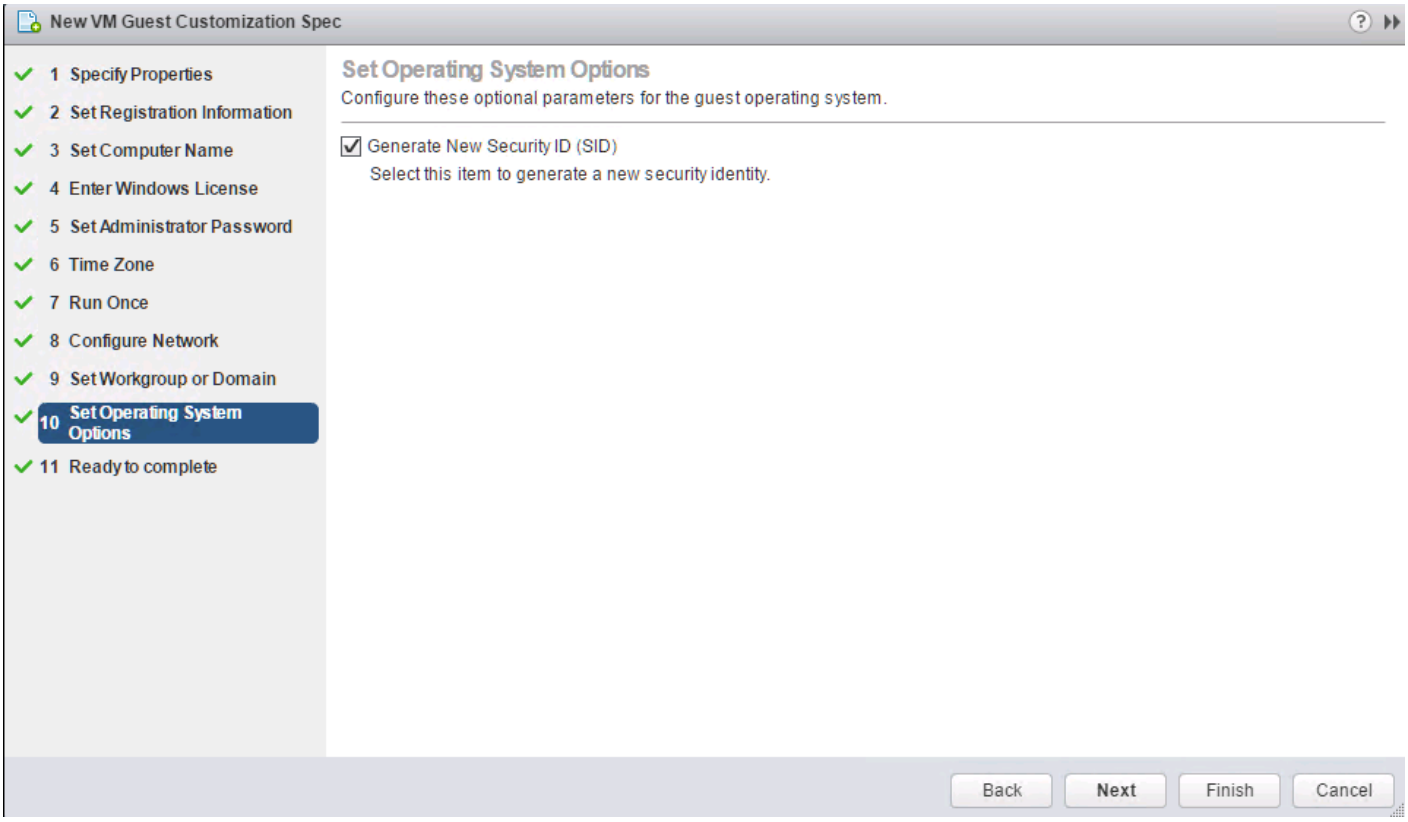

12. Review and click Next to complete creating the Customization Specs.

13. Click Finish.

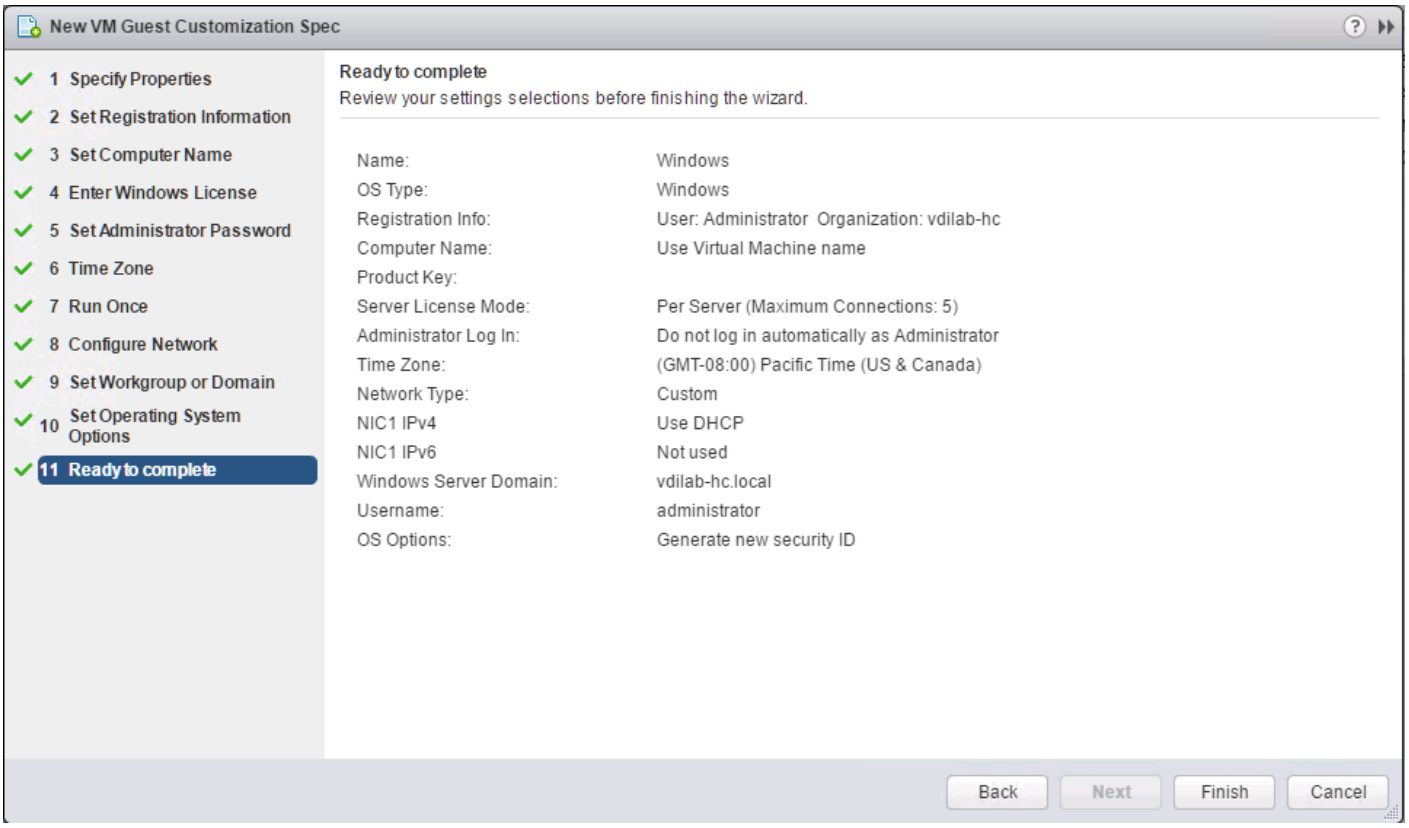
New VM Guest Customization Spec
1 Specify Properties
2 Set Registration Information
3 Set Computer Name
4 Enter Windows License
5 Set Administrator Password
6 Time Zone
7 Run Once
8 Configure Network
9 Set Workgroup or Domain
10 Set Operating System Options
11 Ready to complete
Ready to complete
Review your settings selections before finishing the wizard.
Name: Windows
OS Type: Windows
Registration Info: User: Administrator Organization: vdilab-hc
Computer Name: Use Virtual Machine name
Product Key:
Server License Mode: Per Server (Maximum Connections: 5)
Administrator Log In: Do not log in automatically as Administrator
Time Zone: (GMT-08:00) Pacific Time (US & Canada)
Network Type: Custom
NIC1 IPv4 Use DHCP
NIC1 IPv6 Not used
Windows Server Domain: vdilab-hc.local
Username: administrator
OS Options: Generate new security ID
Back
Next
Finish
Cancel

VMware Horizon Farm and Pool Creation

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.

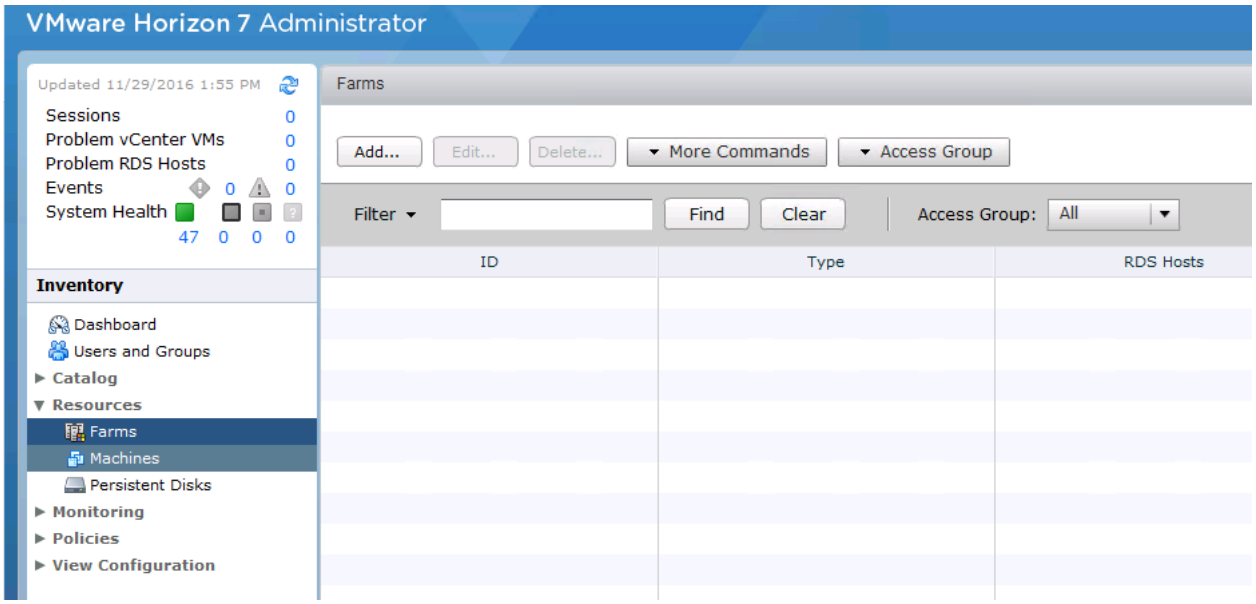

3. Select either to create an Automated or Manual Farm. In this solution, we selected Automated Farm.

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

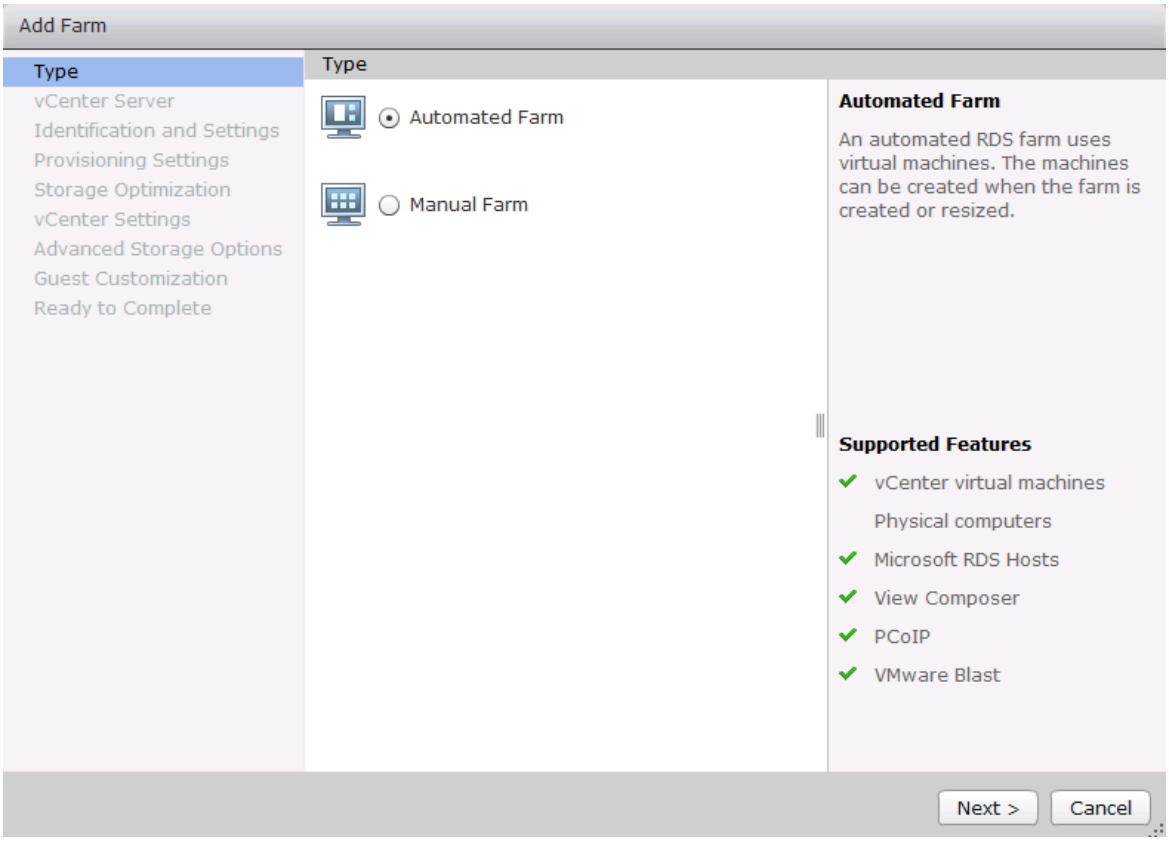

4. Select deployment type whether Instant-clone or Linked-Clone and vCenter server for RDSH farm.

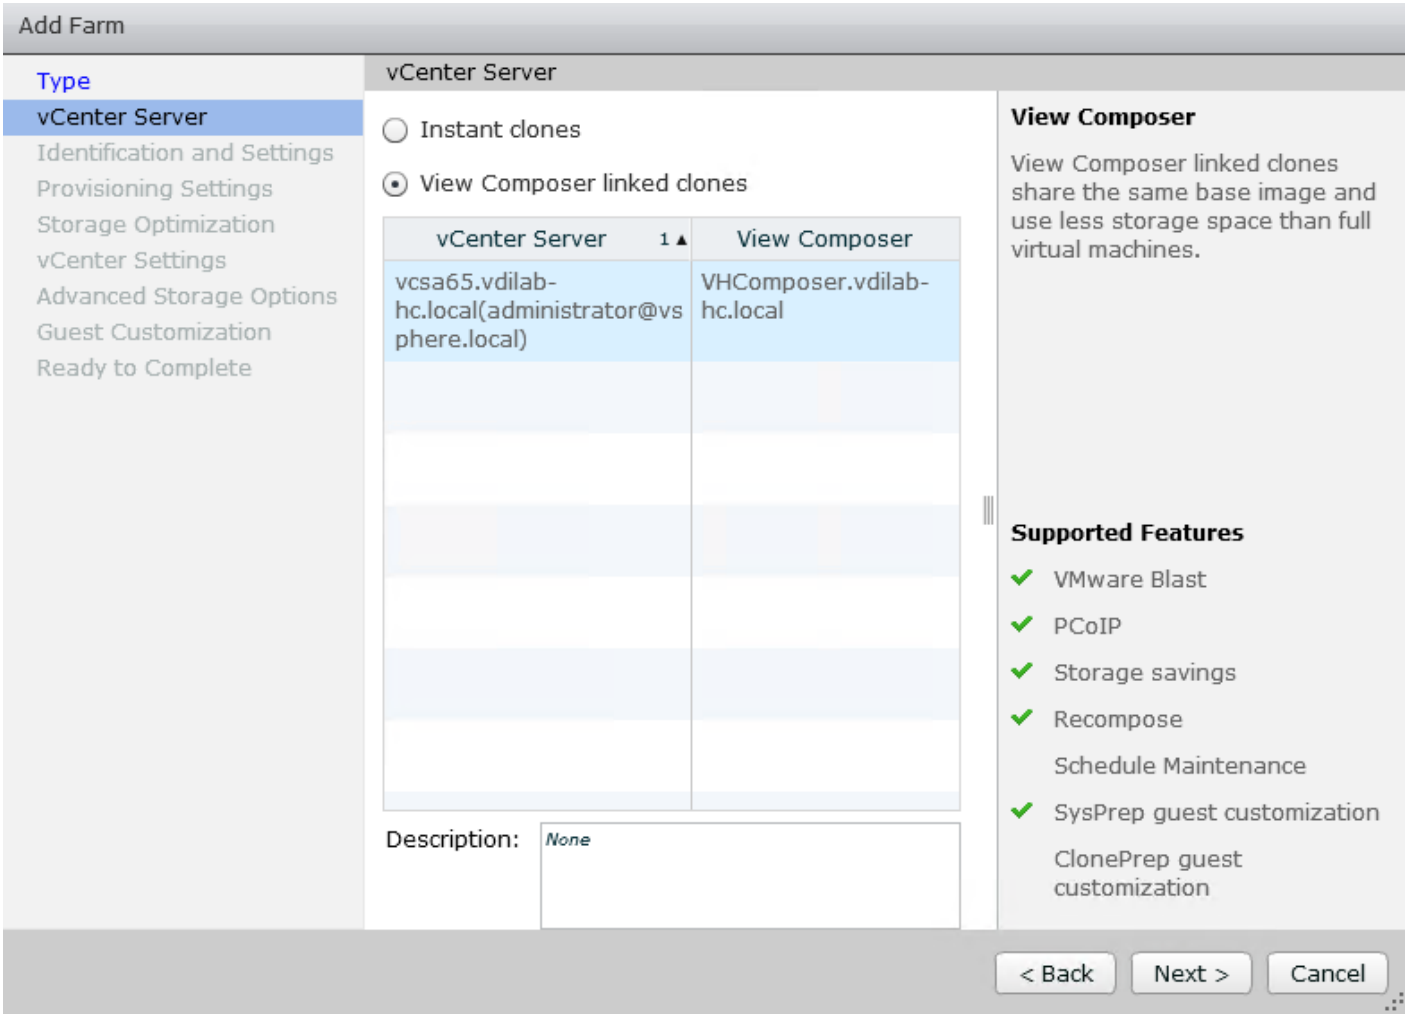

5. Enter the RDSH Farm ID, Access group, Default Display Protocol (Blast/PCoIP/RDP).

6. Select if users are allowed to change the default display protocol, Session timeout, Logoff Disconnected users, and select the checkbox to Enable HTML access.

7. Click Next.

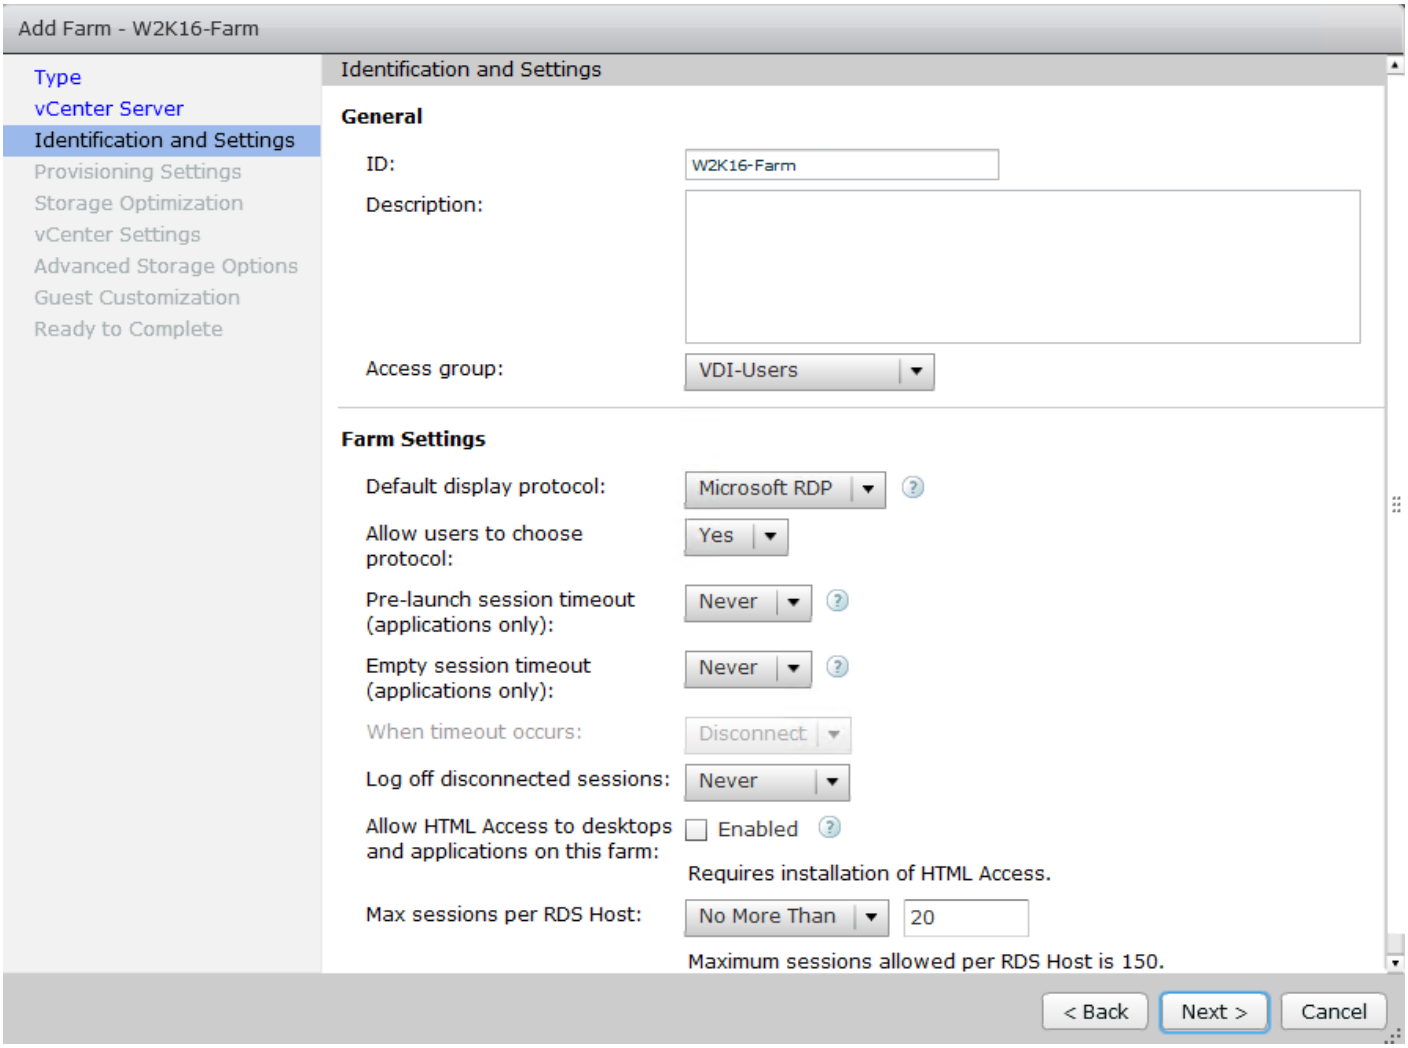

8. Select the provisioning settings, naming convention for RDSH server VM to deploy, and the number of VMs to deploy.

In this study, we deployed 32 RDSH virtual machines across our 4 node HyperFlex Cluster.

9. Click Next.

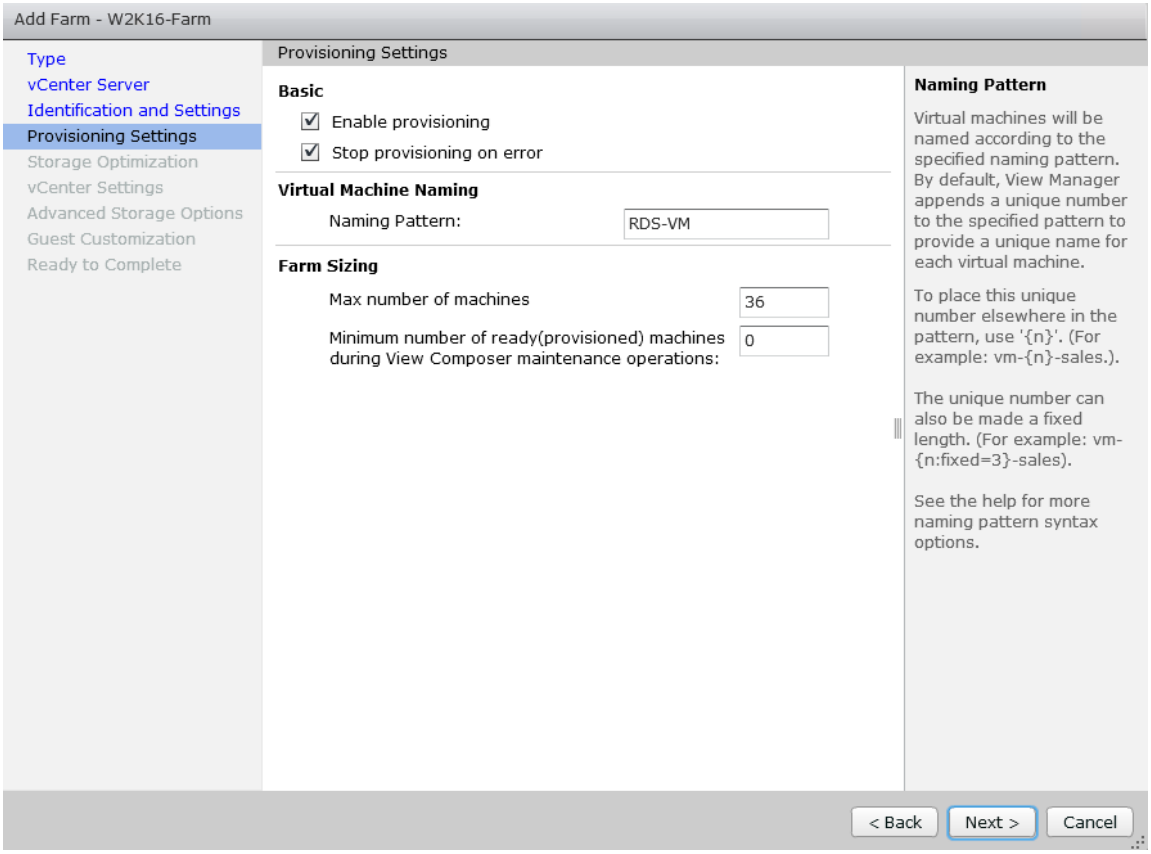

10. Click Next.

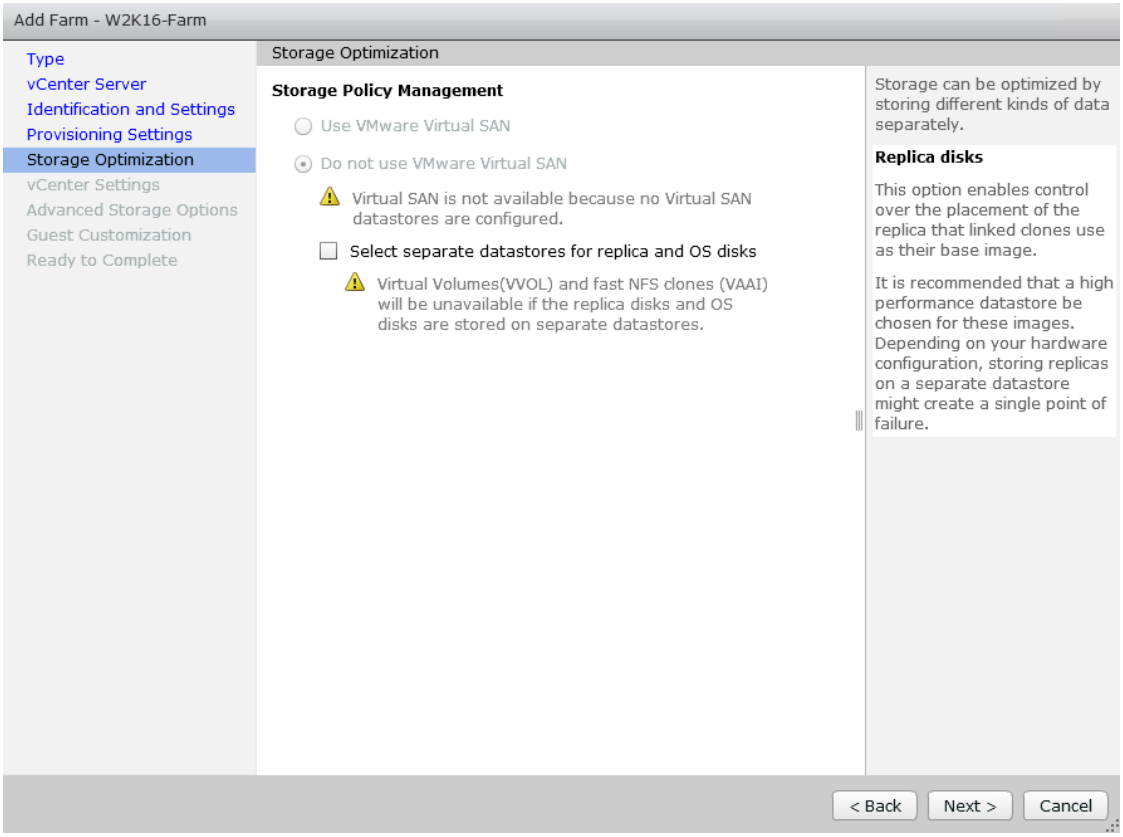

11. Select vCenter settings, for example; Master Image, snapshot, folder, Host or Cluster, resource pool, storage selection.

12. Click Next.

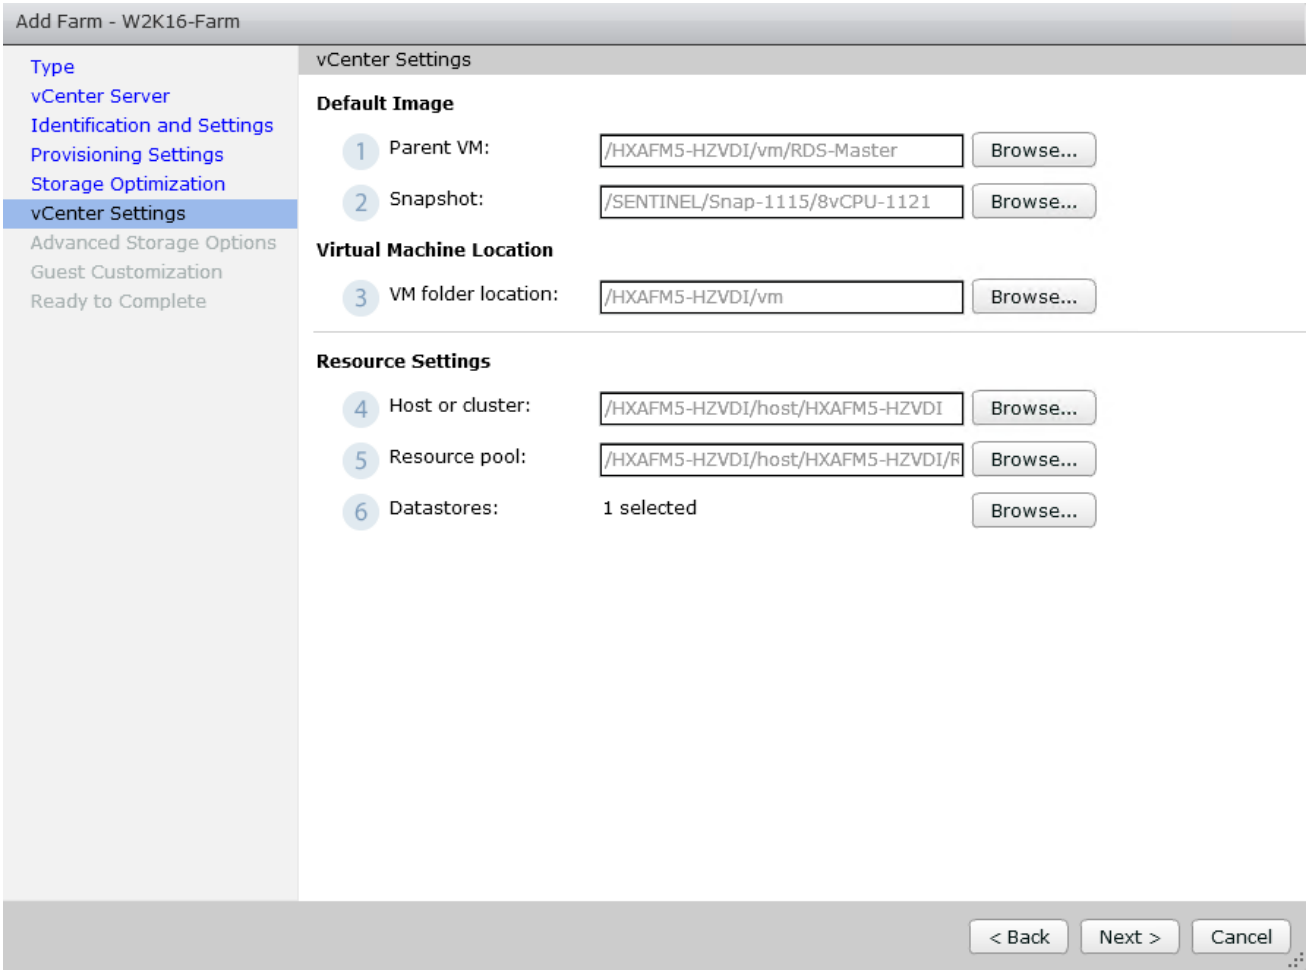

13. For Step 6 Datastores: Browse and choose Unbounded for the Storage Overcommit field.

14. Click OK.

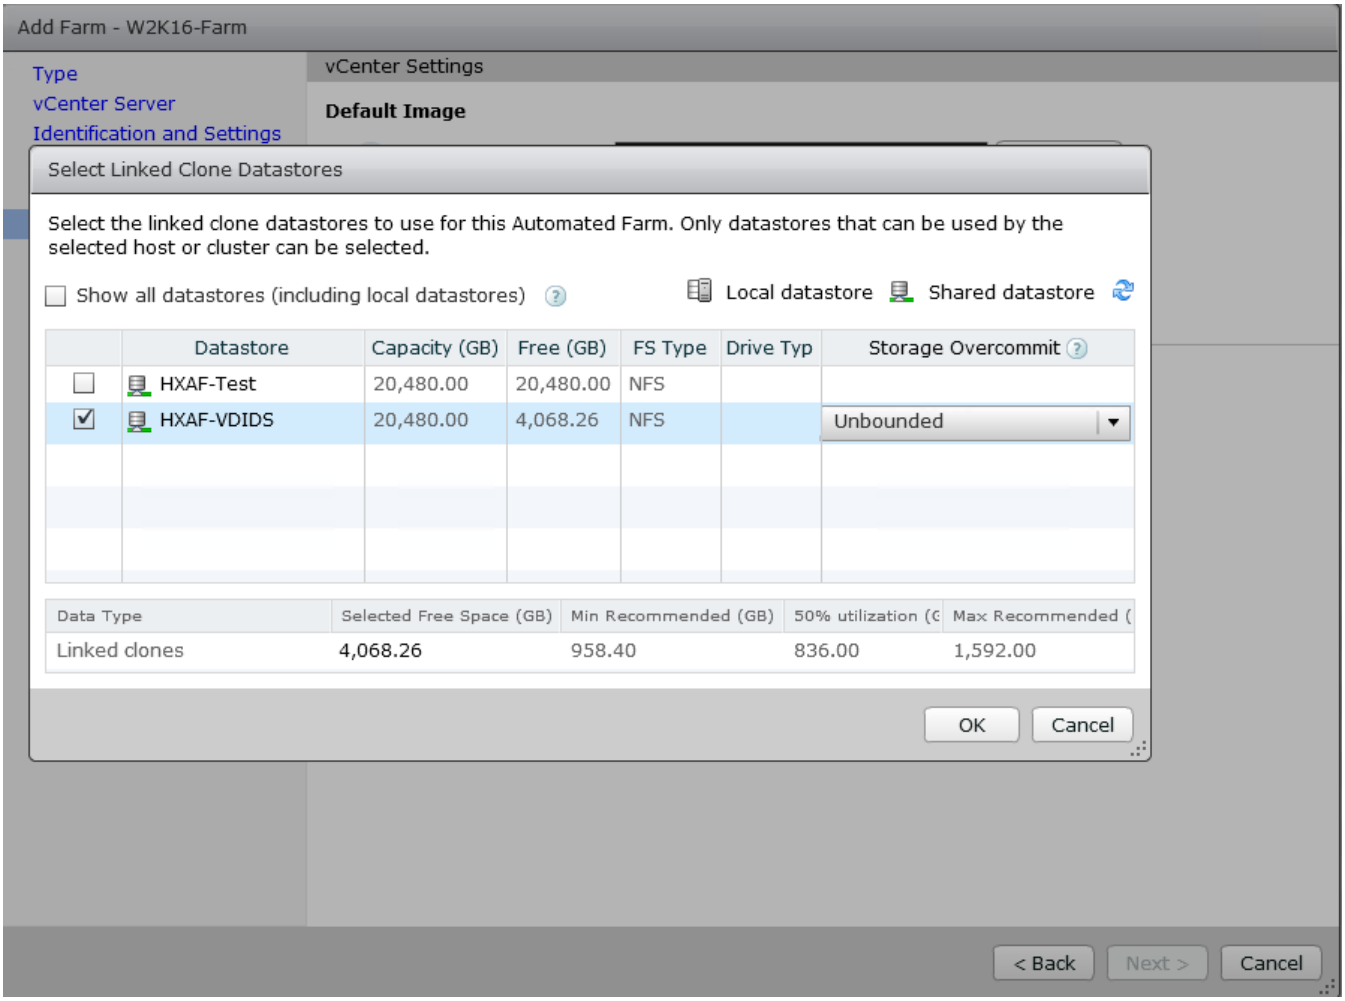

15. In the Advanced Storage Options, select Use native NFS snapshot (VAAI).

16. Click Next.

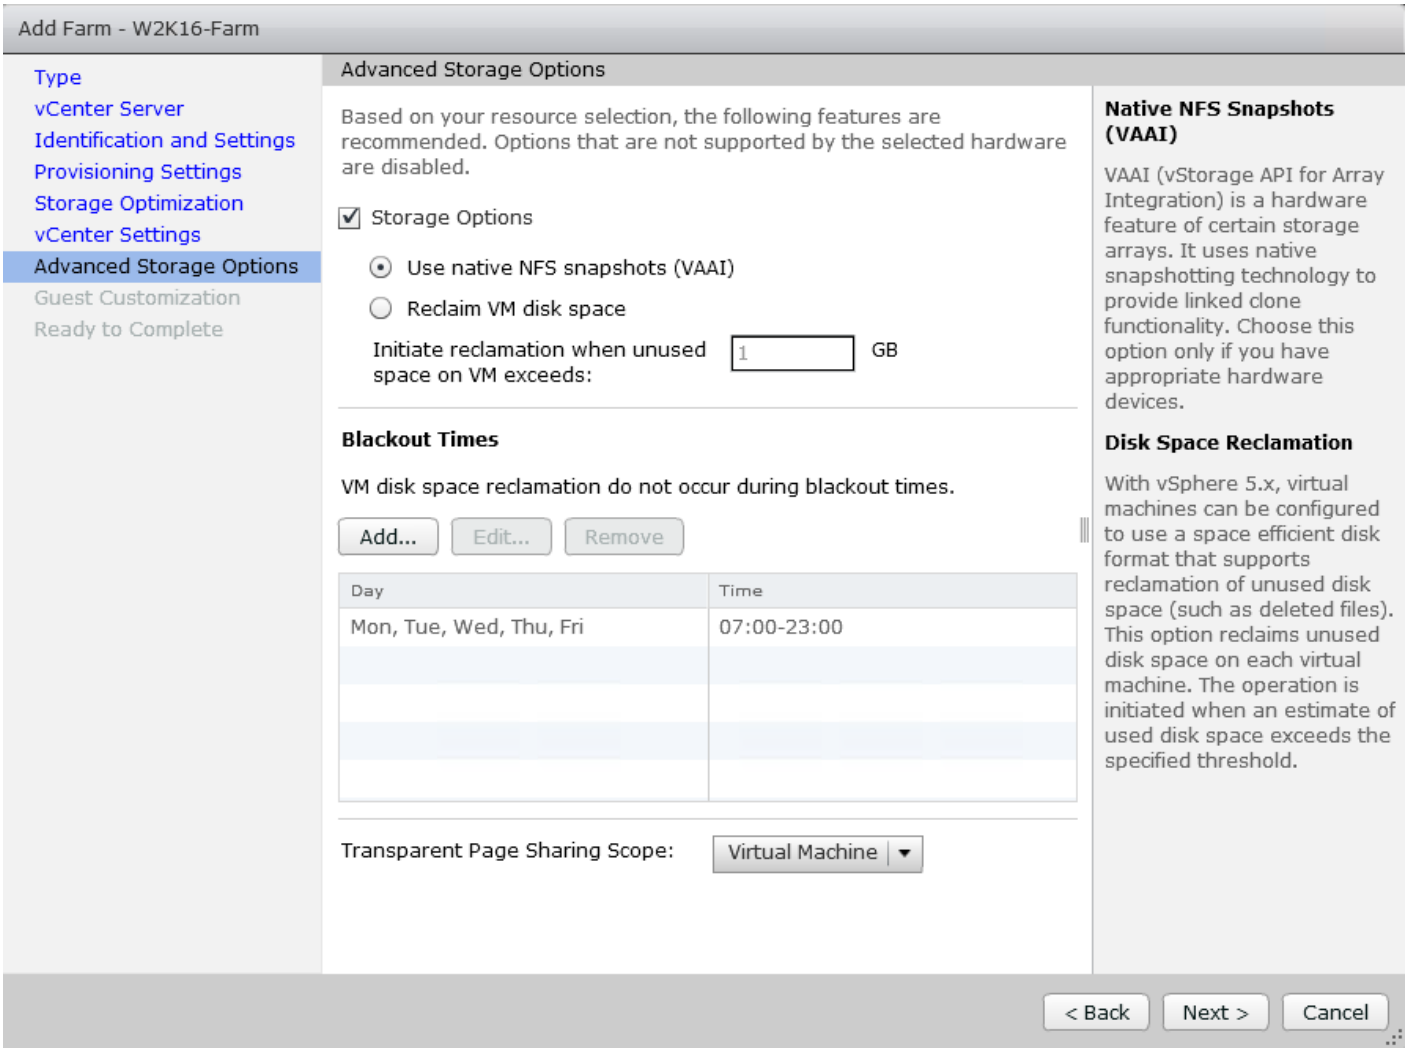

17. Select the Active Directory Domain, the Active Directory OU into which the RDSH machines will be provisioned, and the Sysprep file created as part of the customization specific configuration performed earlier.

18. Click Next.

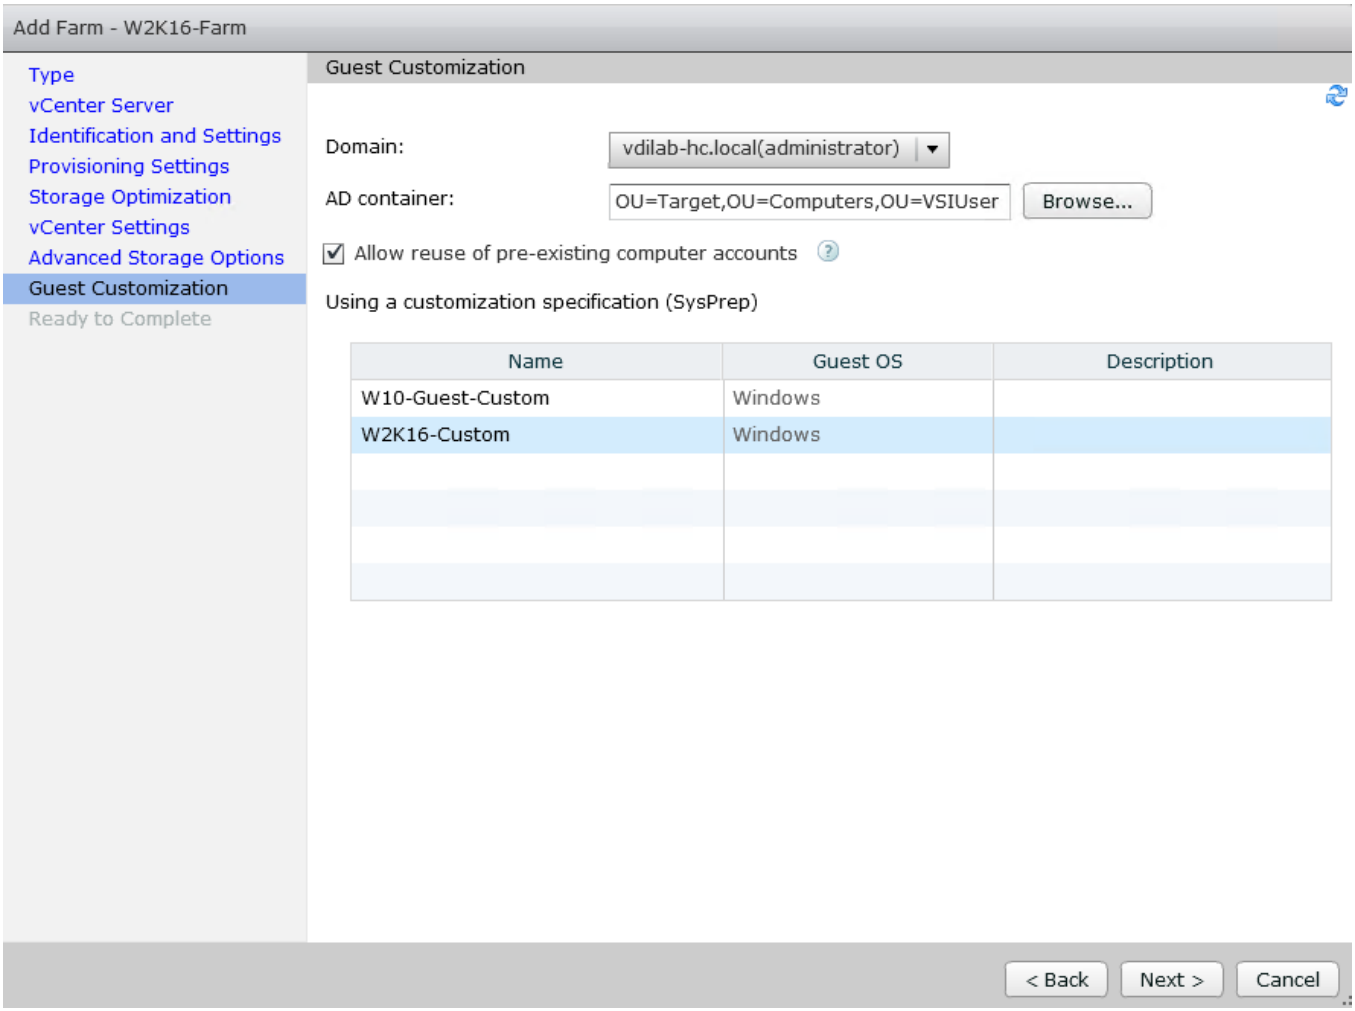

19. Review the pool creation information.

20. Click Finish.

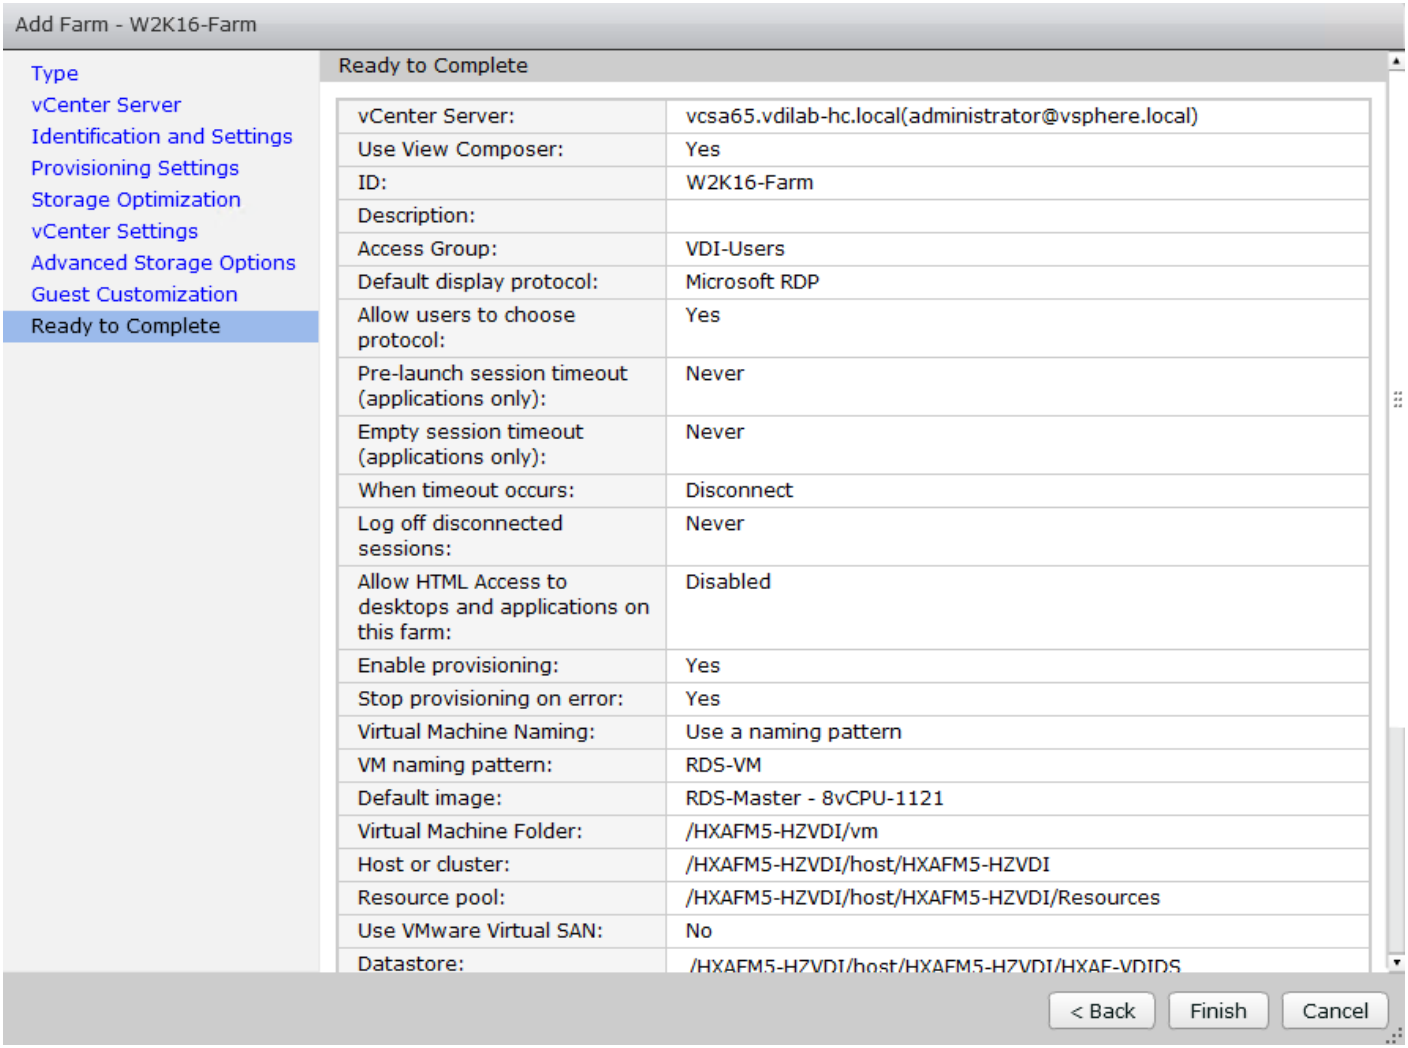

The VMware Horizon Administration console displays the status of the provisioning task and pool settings:

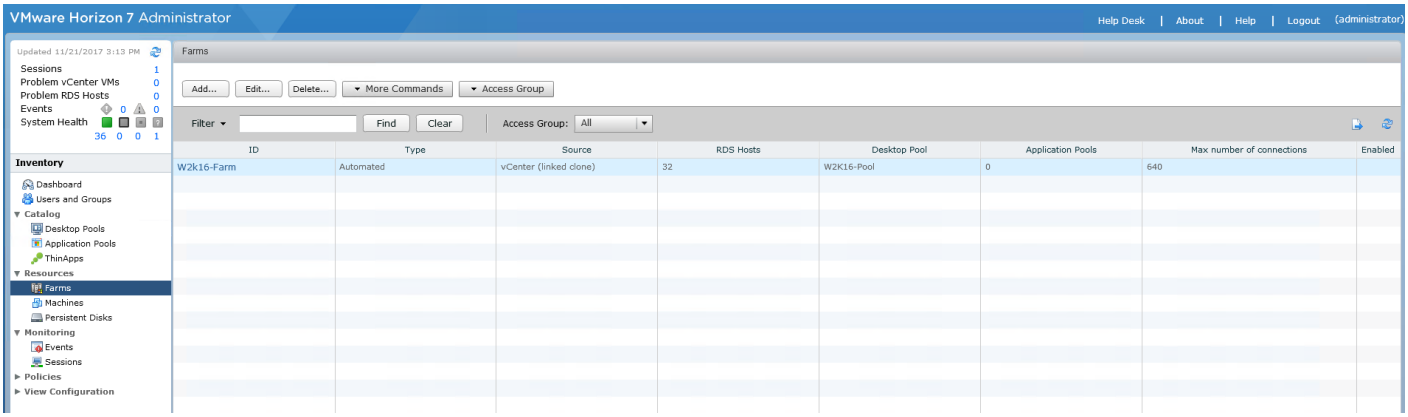

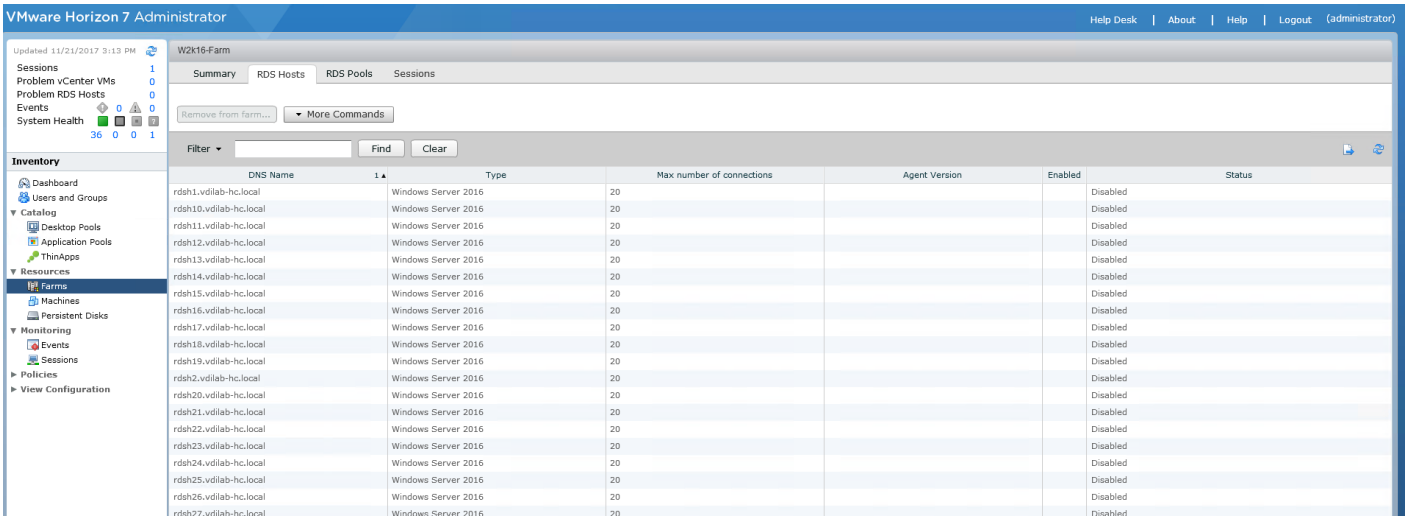

Create the Horizon 7 RDS Published Desktop Pool

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.

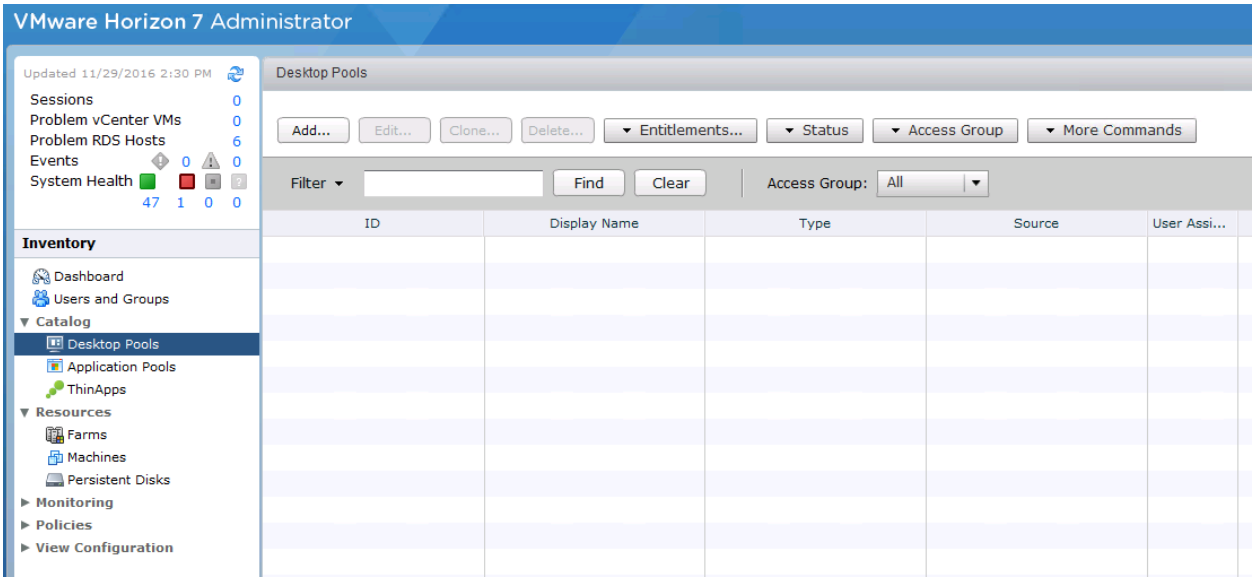

3. Select RDS Desktop pool.

4. Click Next.

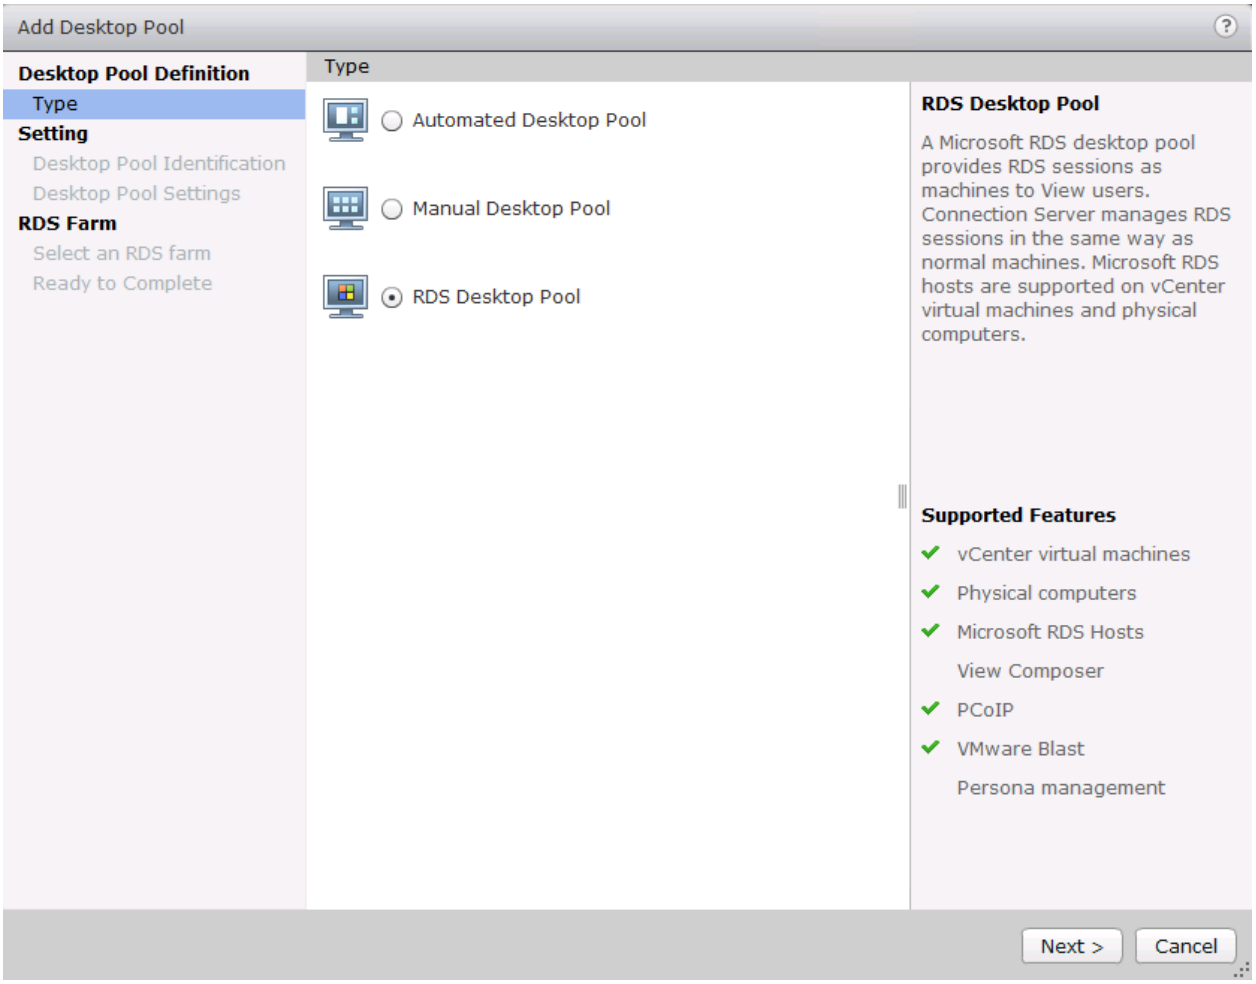

5. Enter Pool ID and Display name.

6. Click Next.

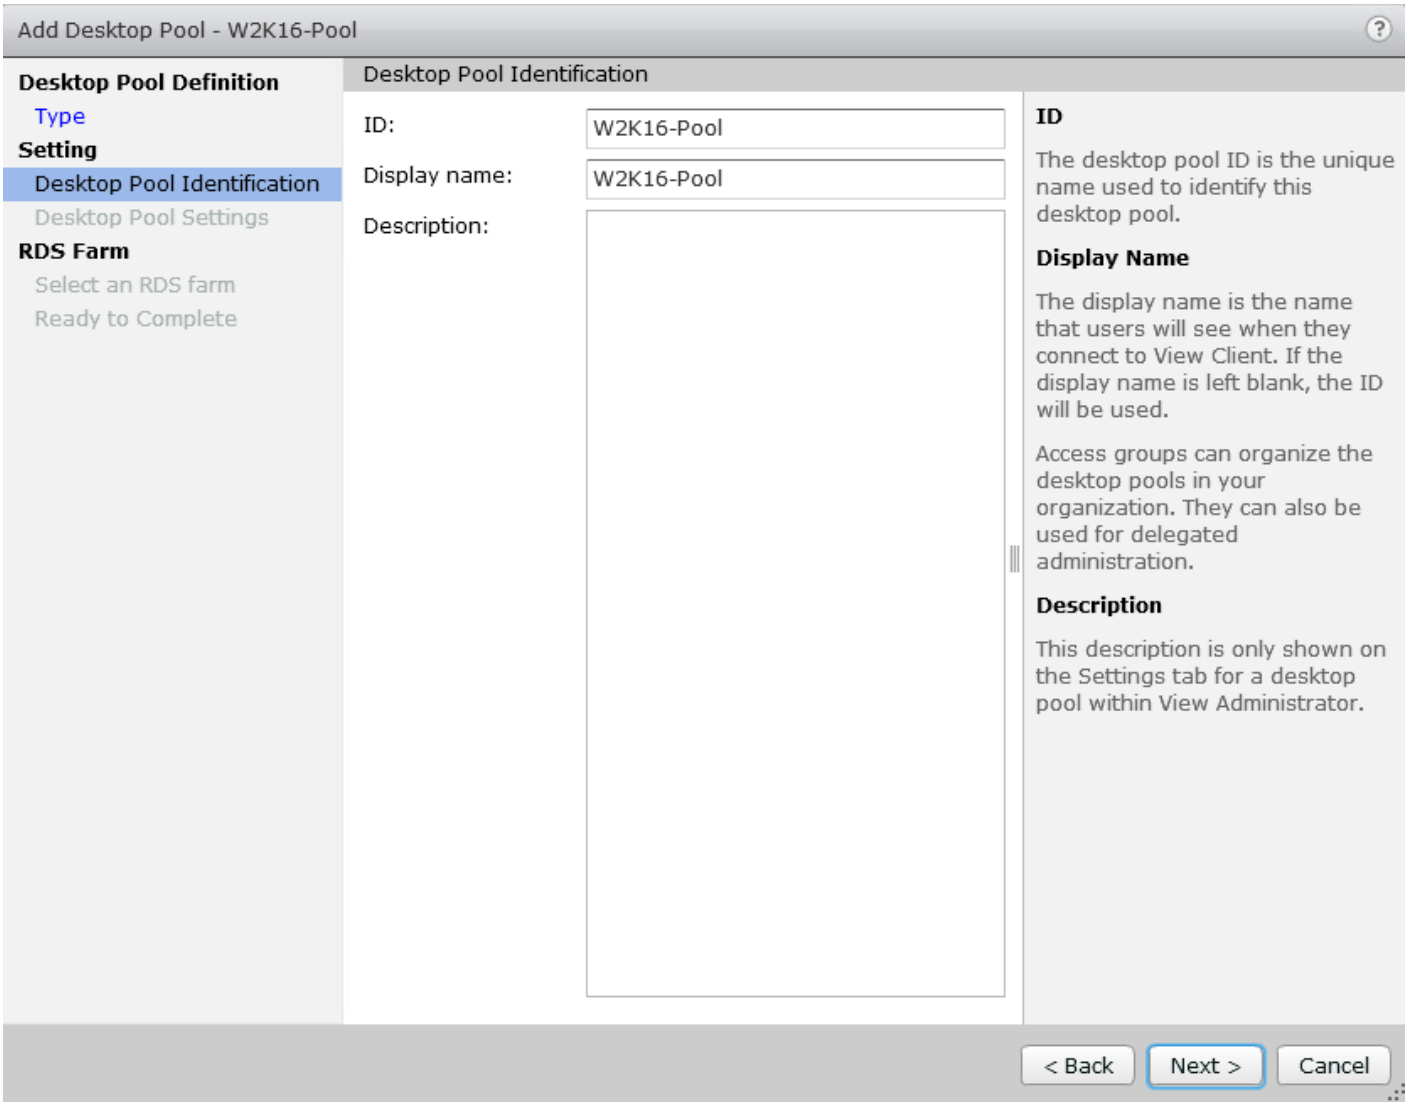

7. Accept the default settings on Desktop Pool Settings page.

8. Click Next.

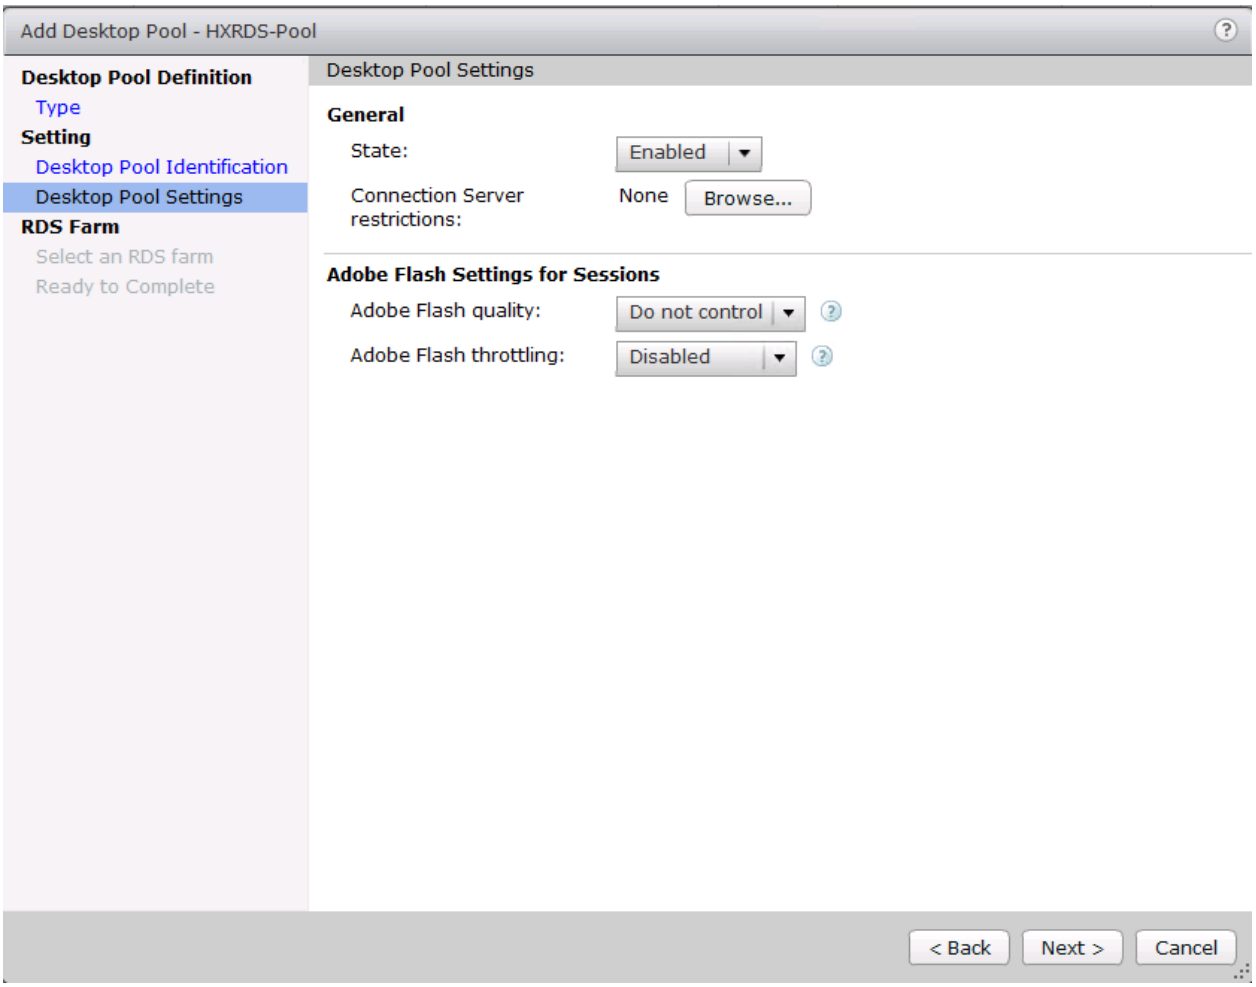

9. Click the “Select an RDS farm for this desktop pool” radio button.

10. Click the farm created in the previous section.

11. Click Next.

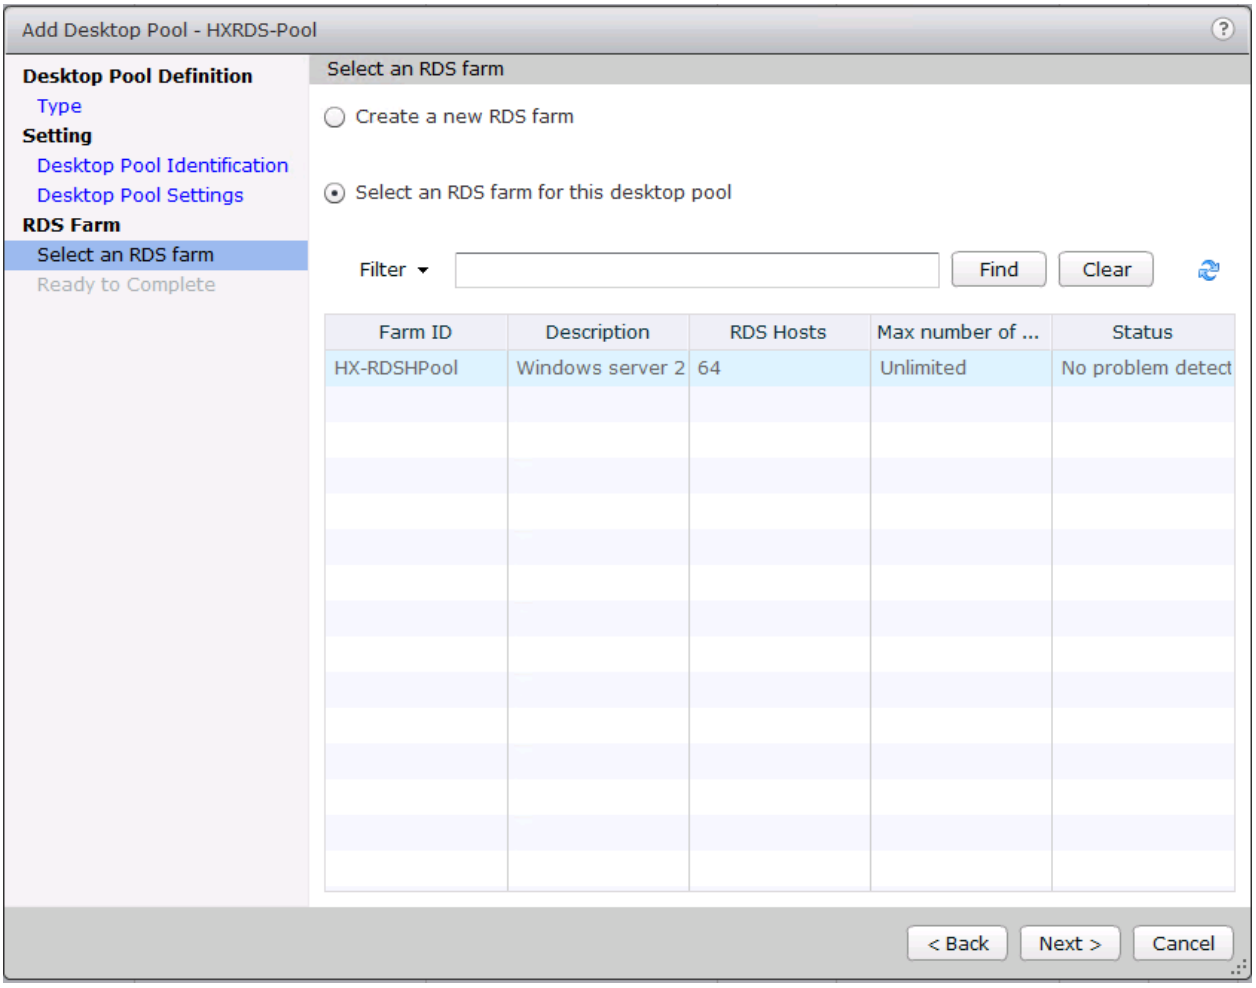

12. Review the pool settings.

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

14. Click Finish.

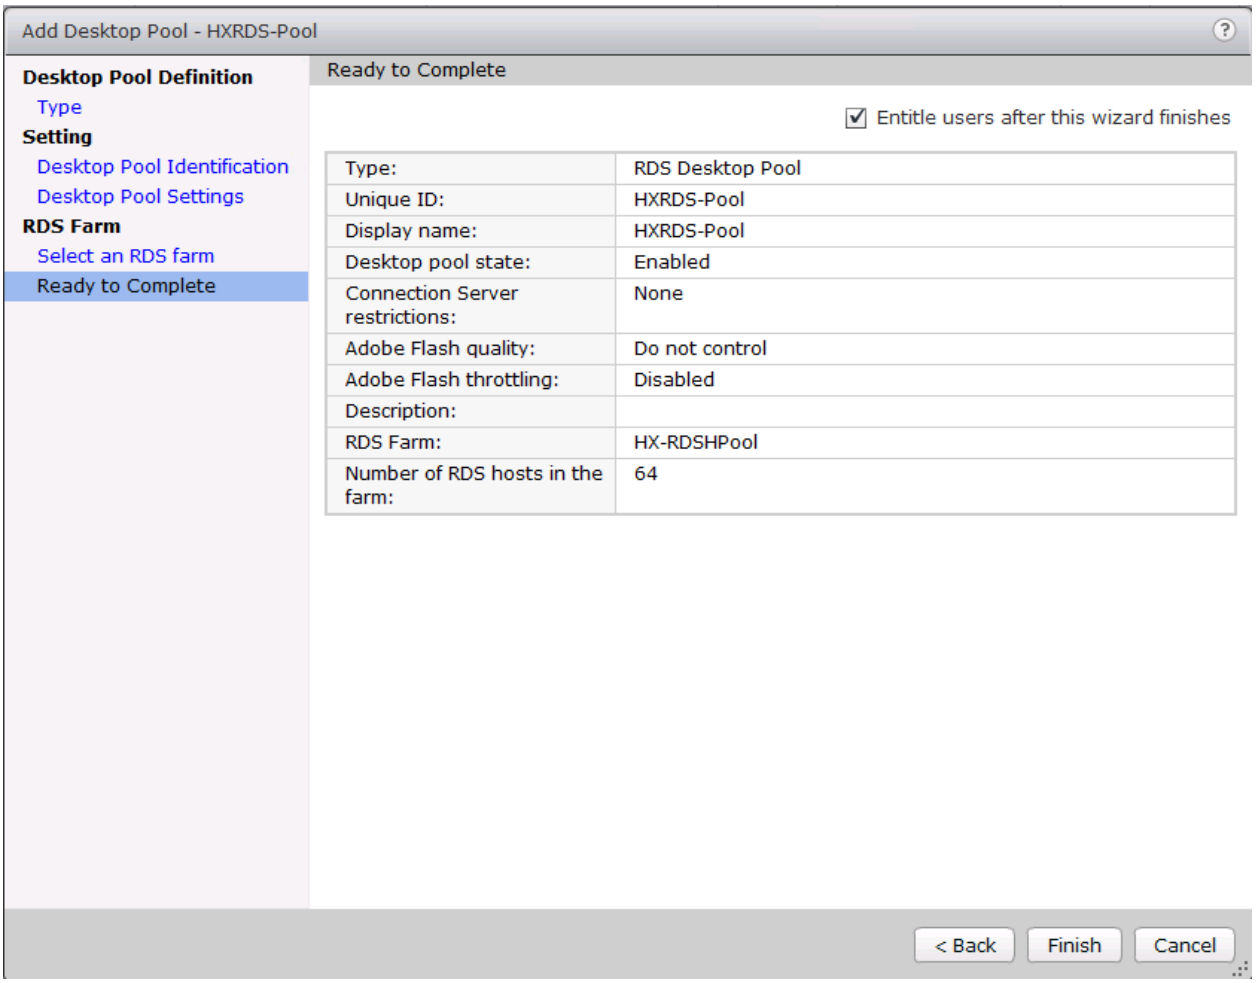

15. Select the Users or Groups checkbox, use the search tools to locate the user or group to be authorized, highlight the user or group in the results box.

16. Click OK.

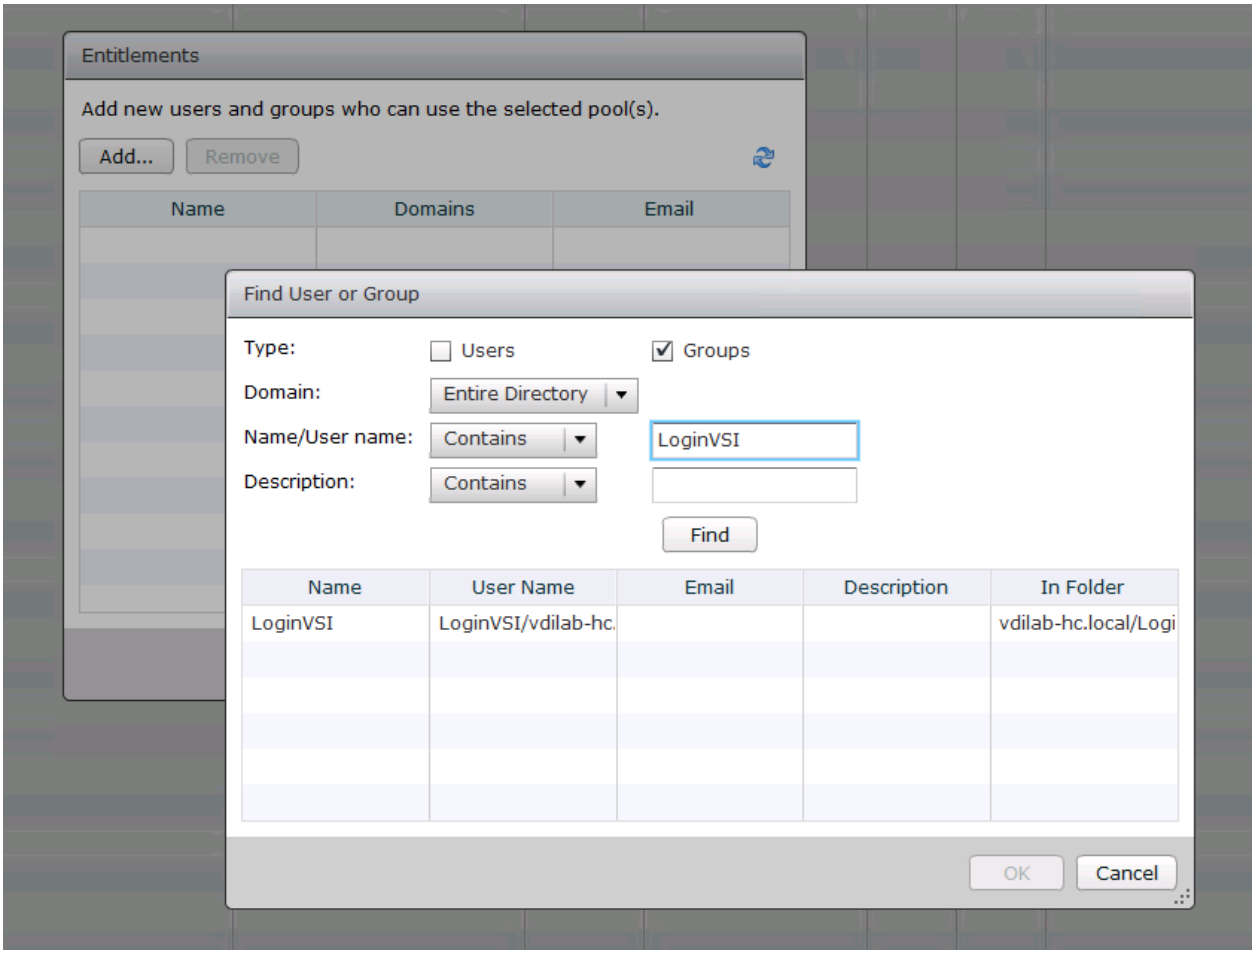

17. You now have a functional RDSH Farm and Desktop Pool with users identified who are authorized to utilize Horizon RDSH sessions.

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.

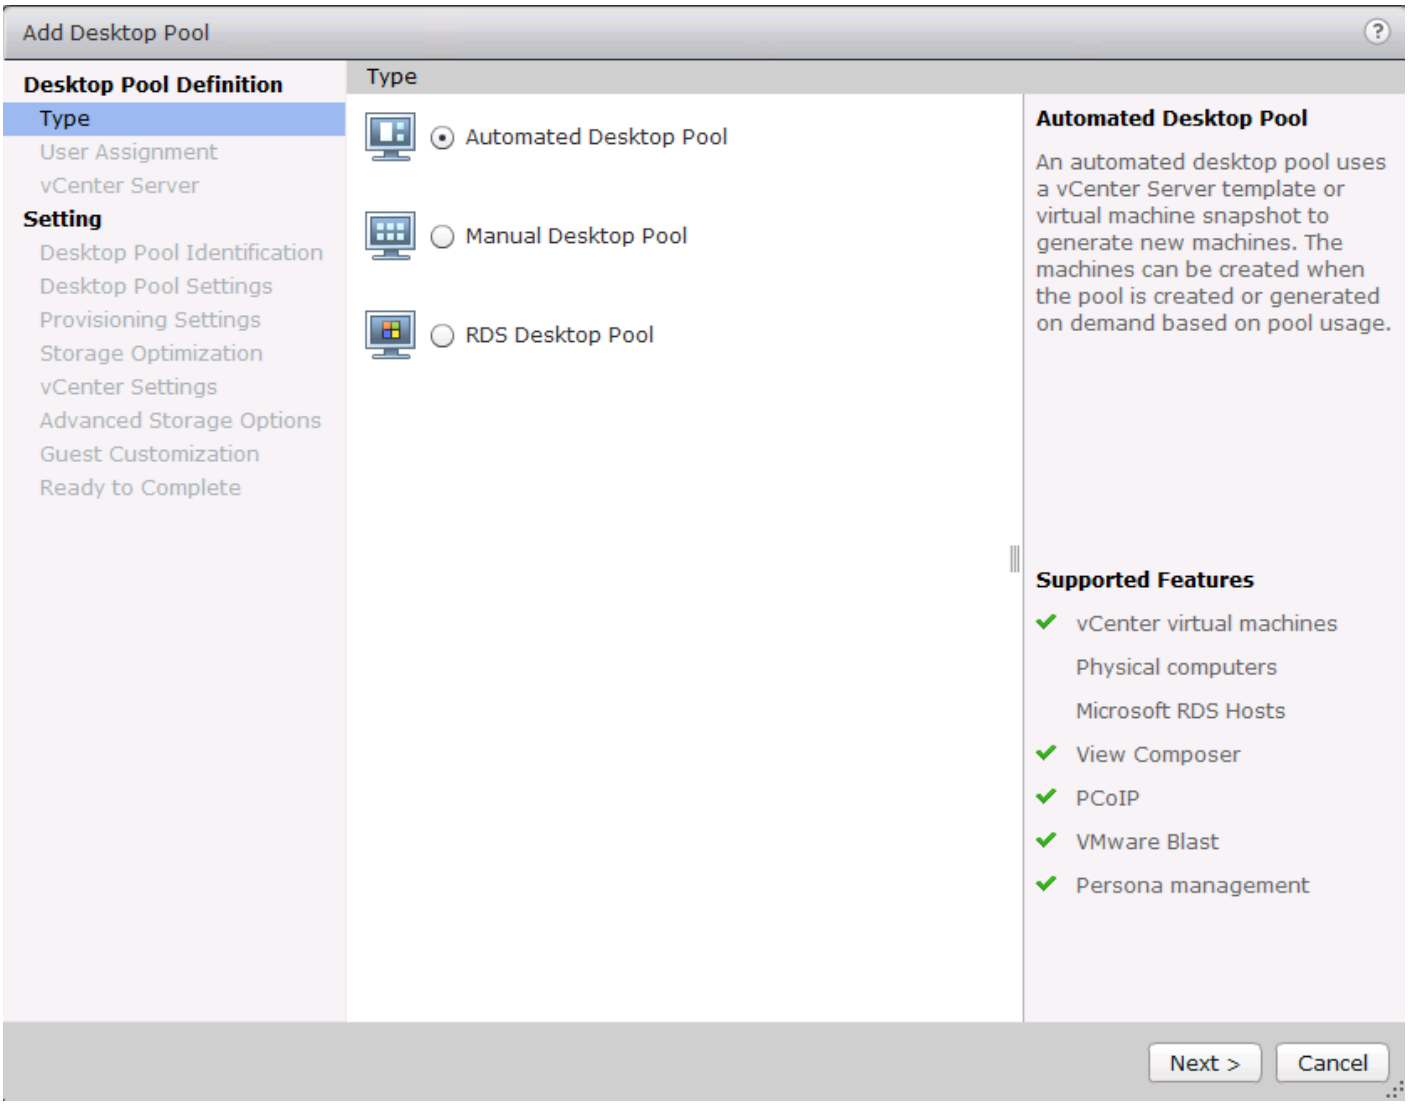

5. Select Floating or Dedicated user assignment.

6. Click Next.

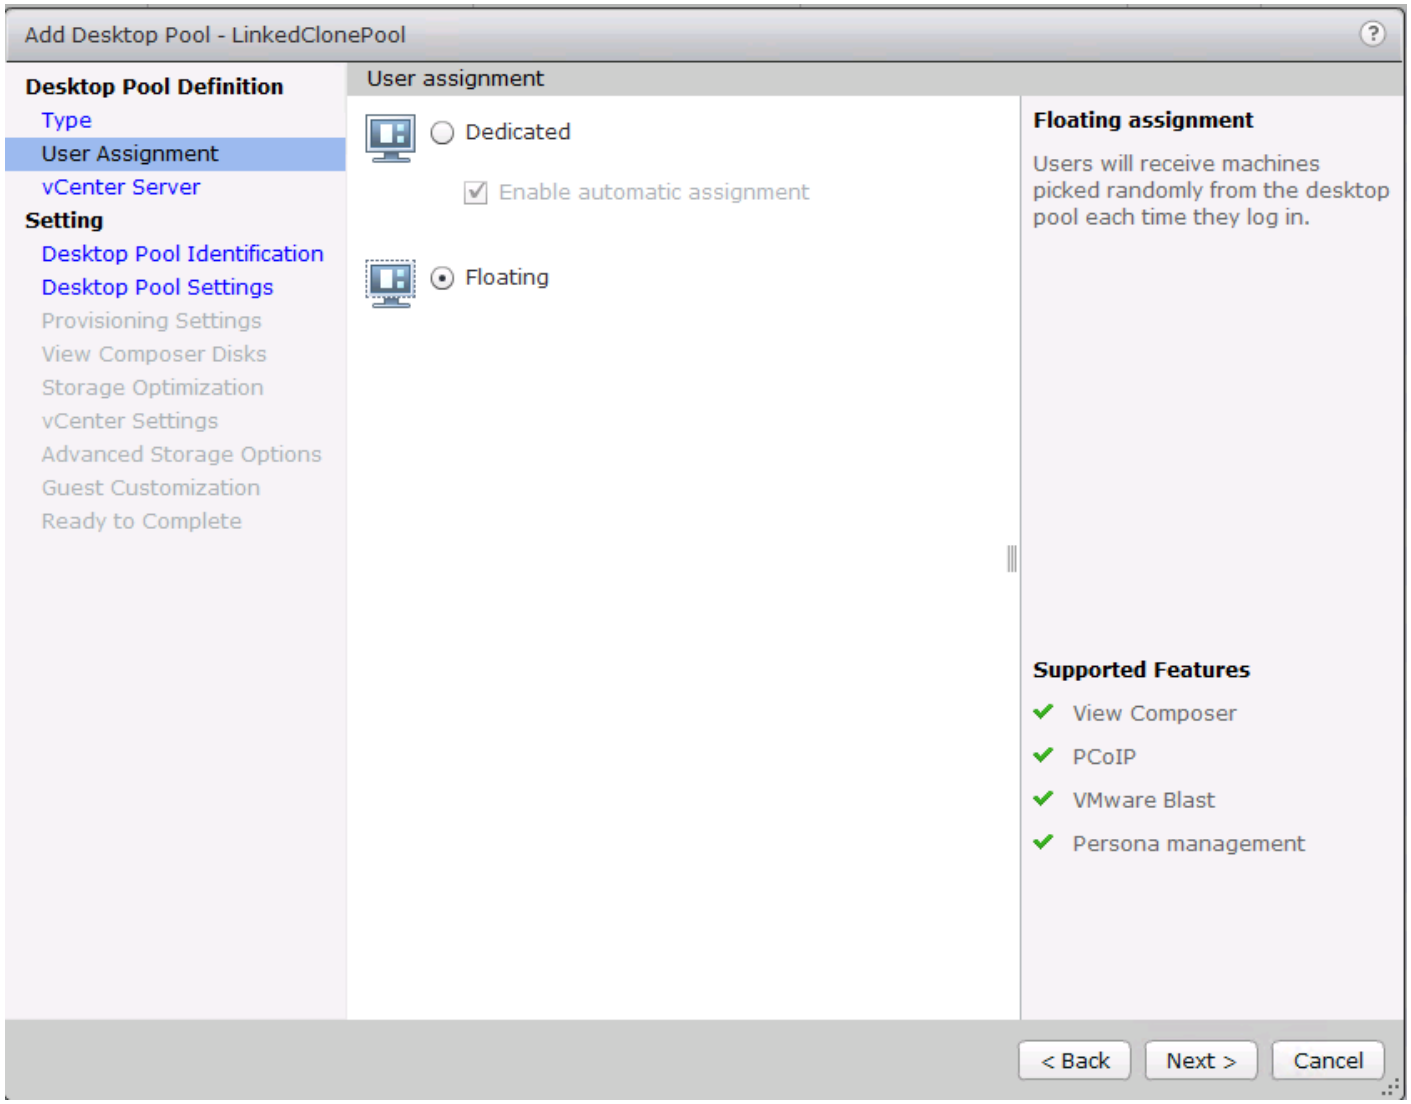

7. Select View Composer Linked Clones, highlight your vCenter and View Composer virtual machine.

8. Click Next.

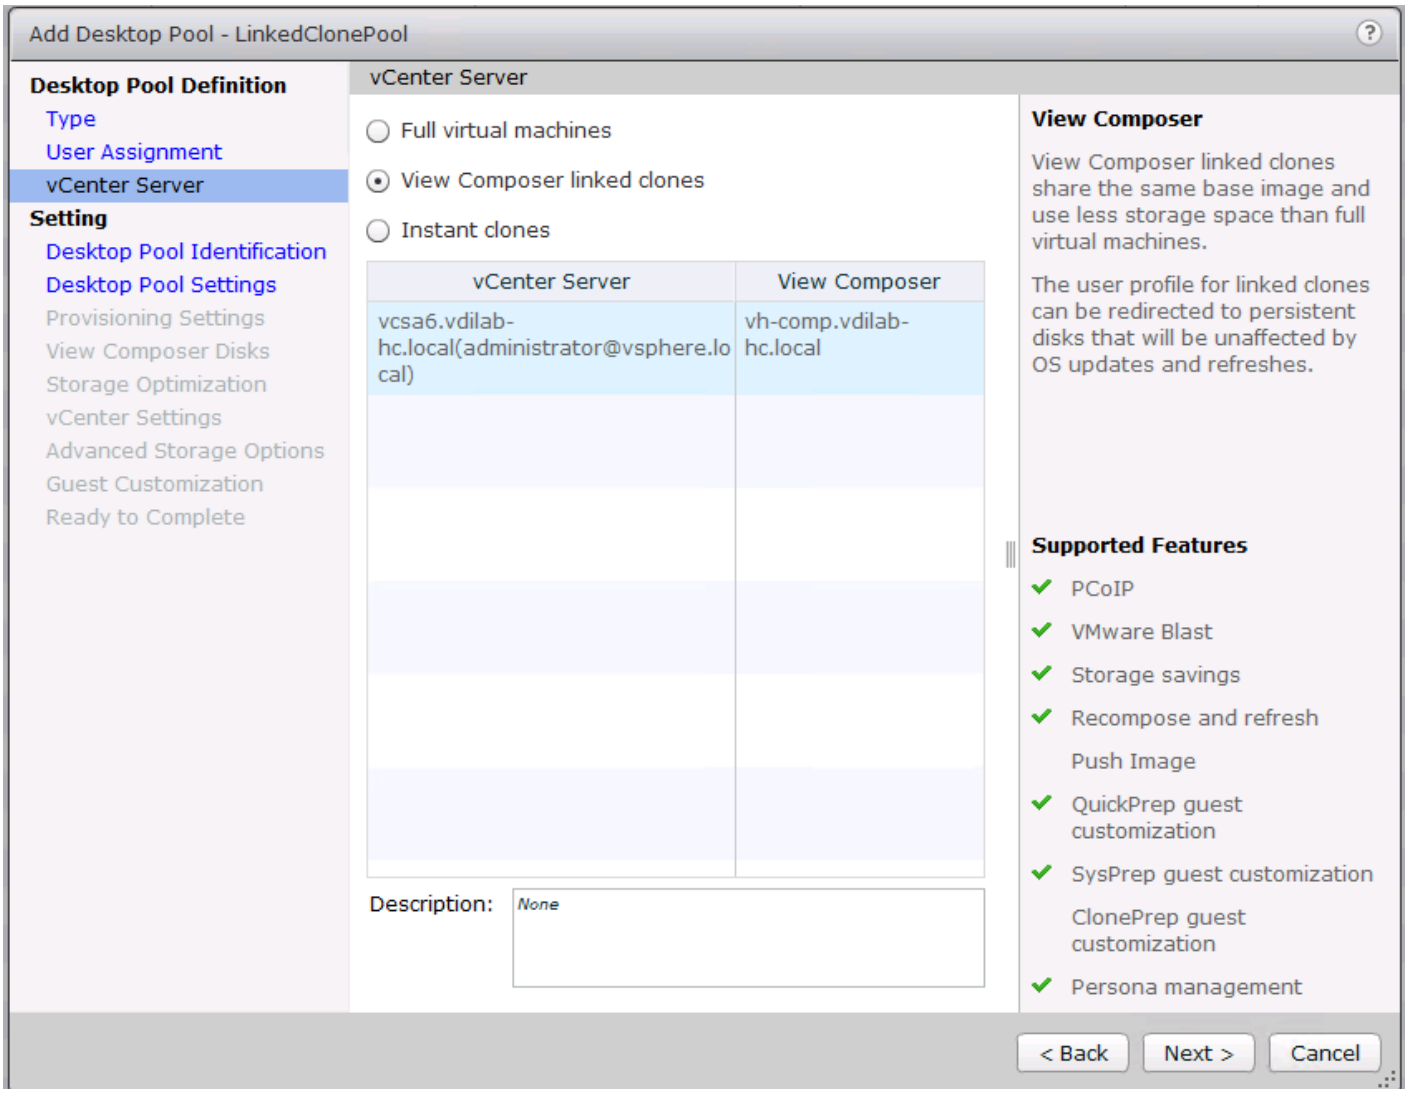

9. Enter pool identification details.

10. Click Next.

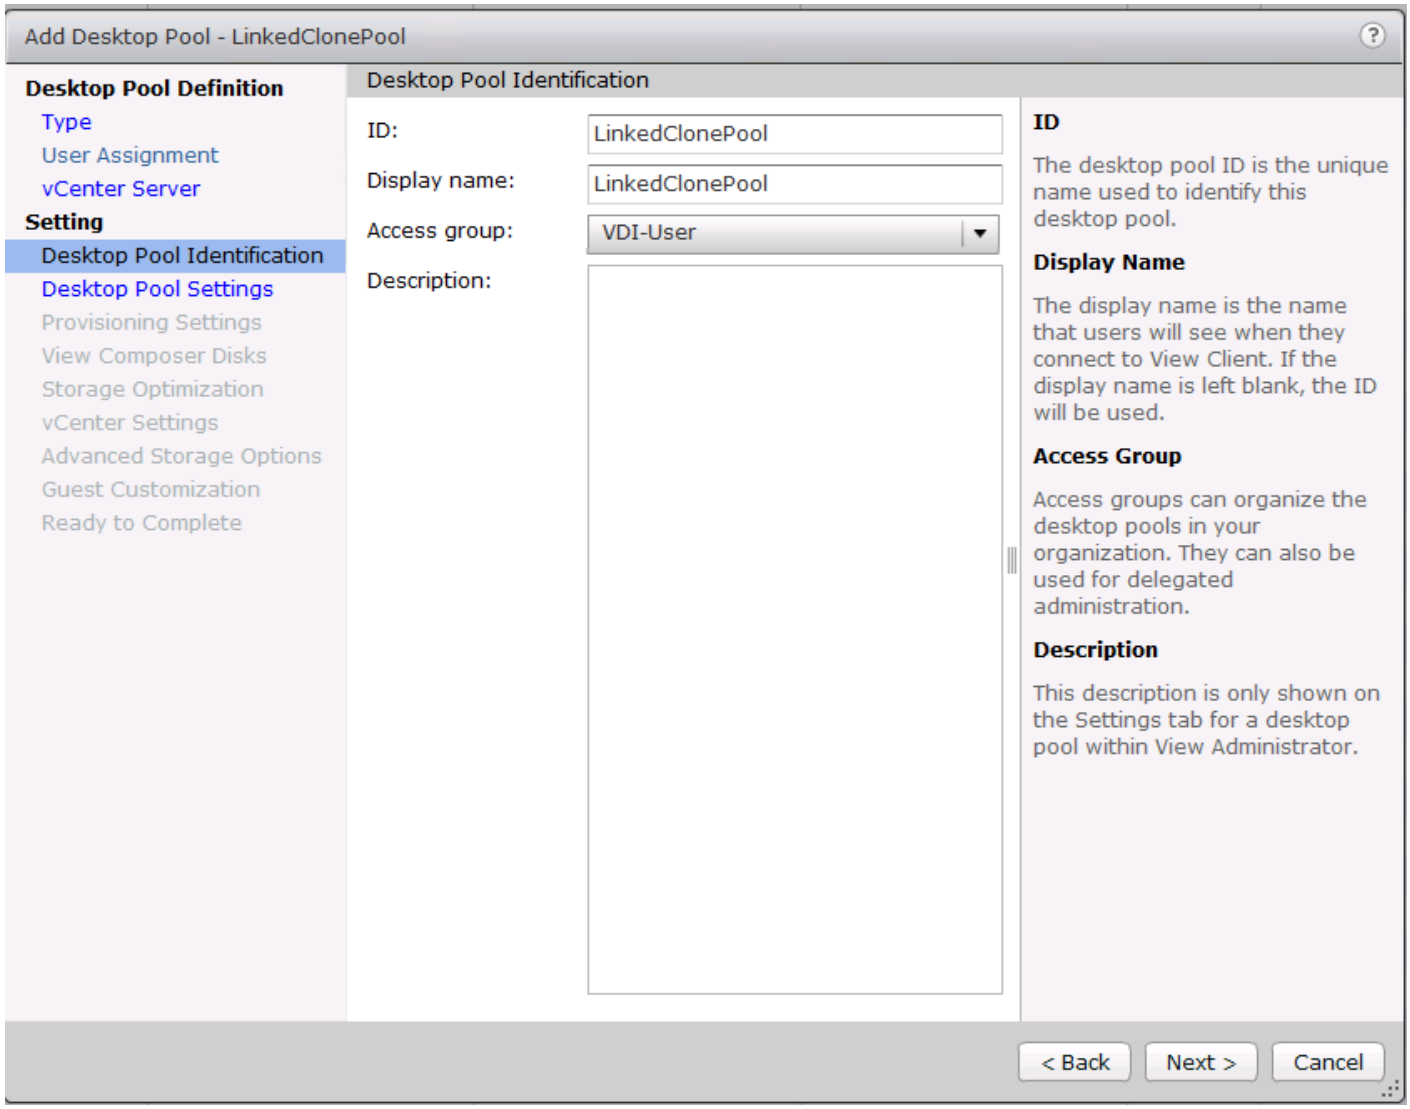

11. Select Desktop Pool settings.

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

12. Click Next.

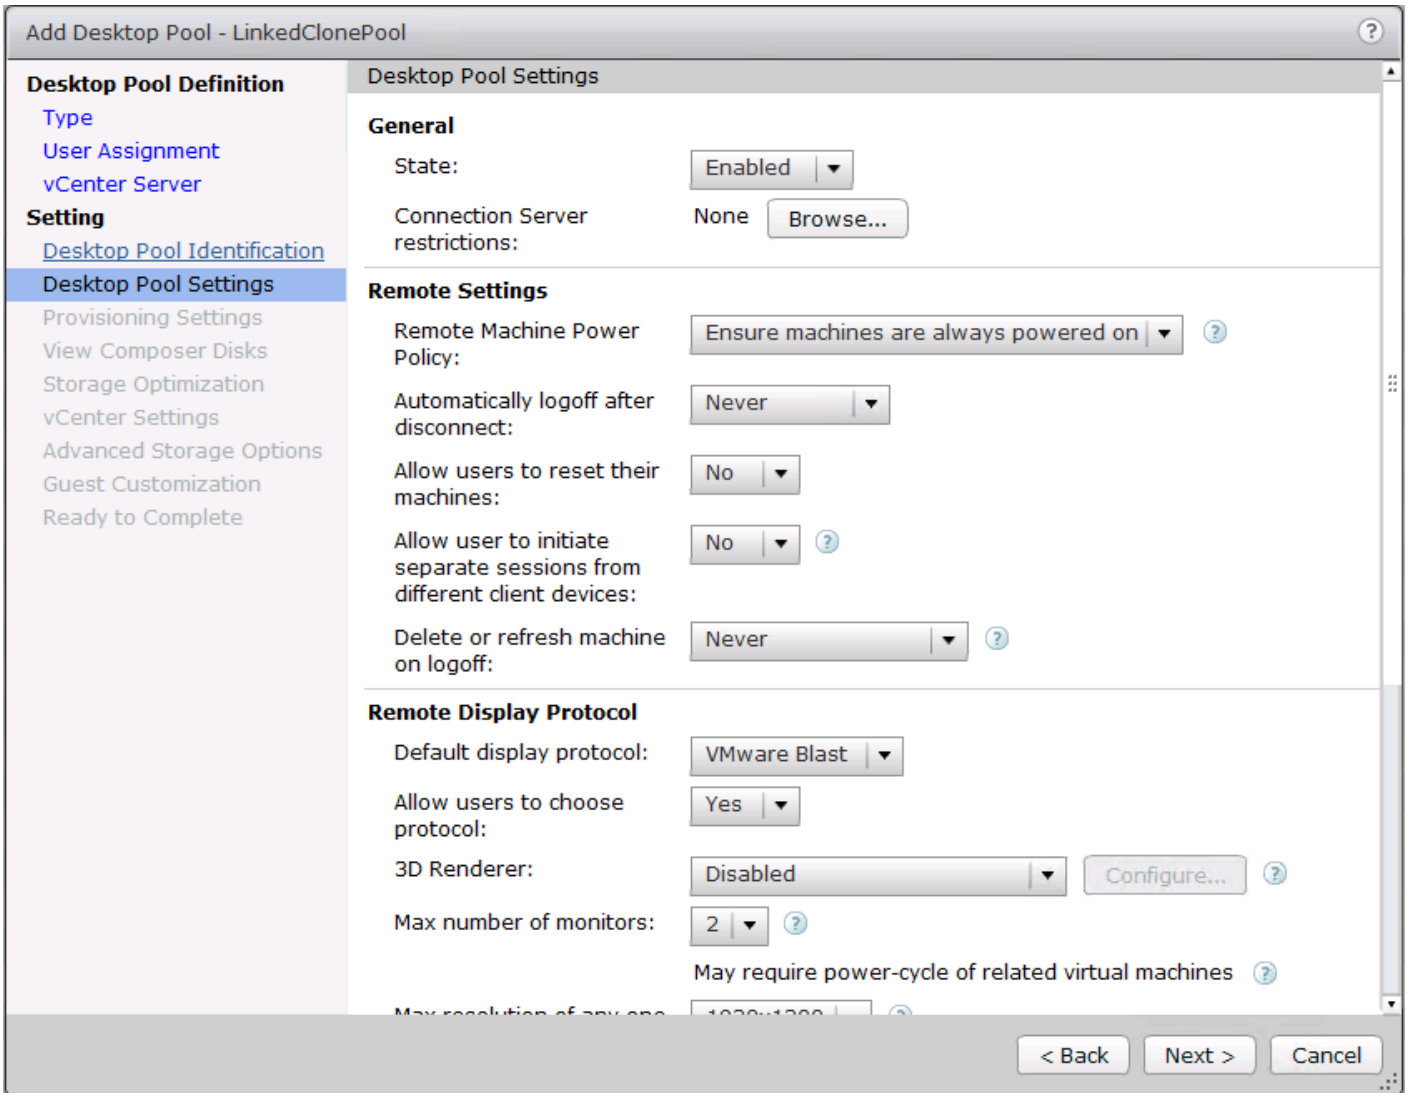

13. Select Provisioning Settings.

14. Click Next.

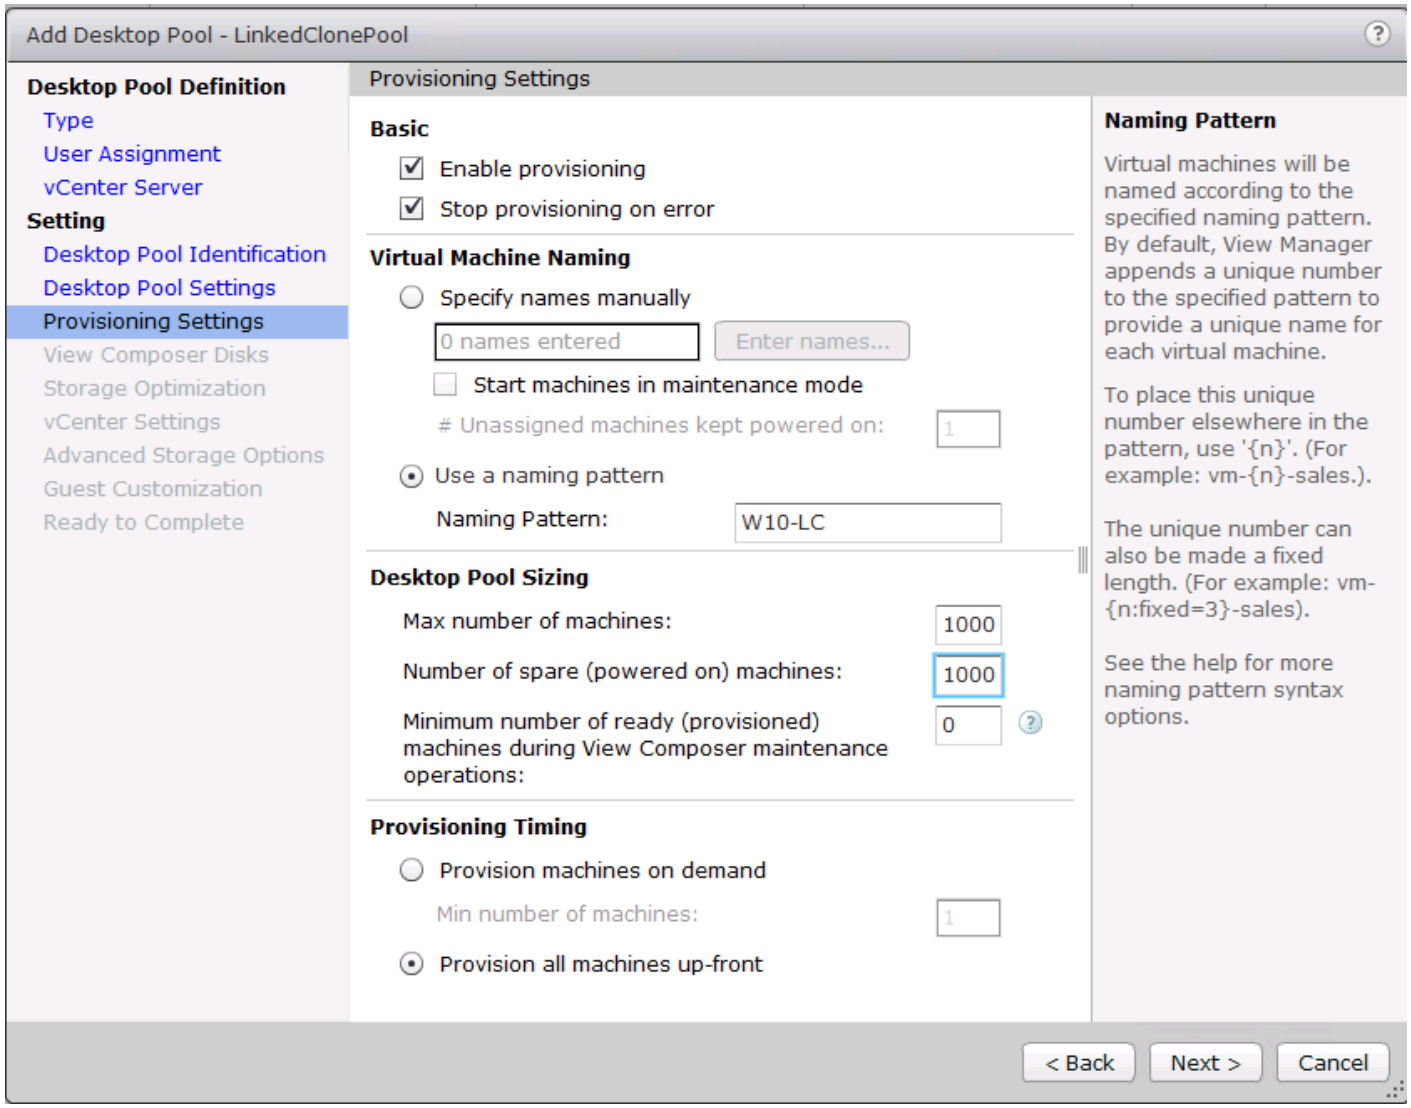

15. Select View Composer disk configuration.

16. Click Next.

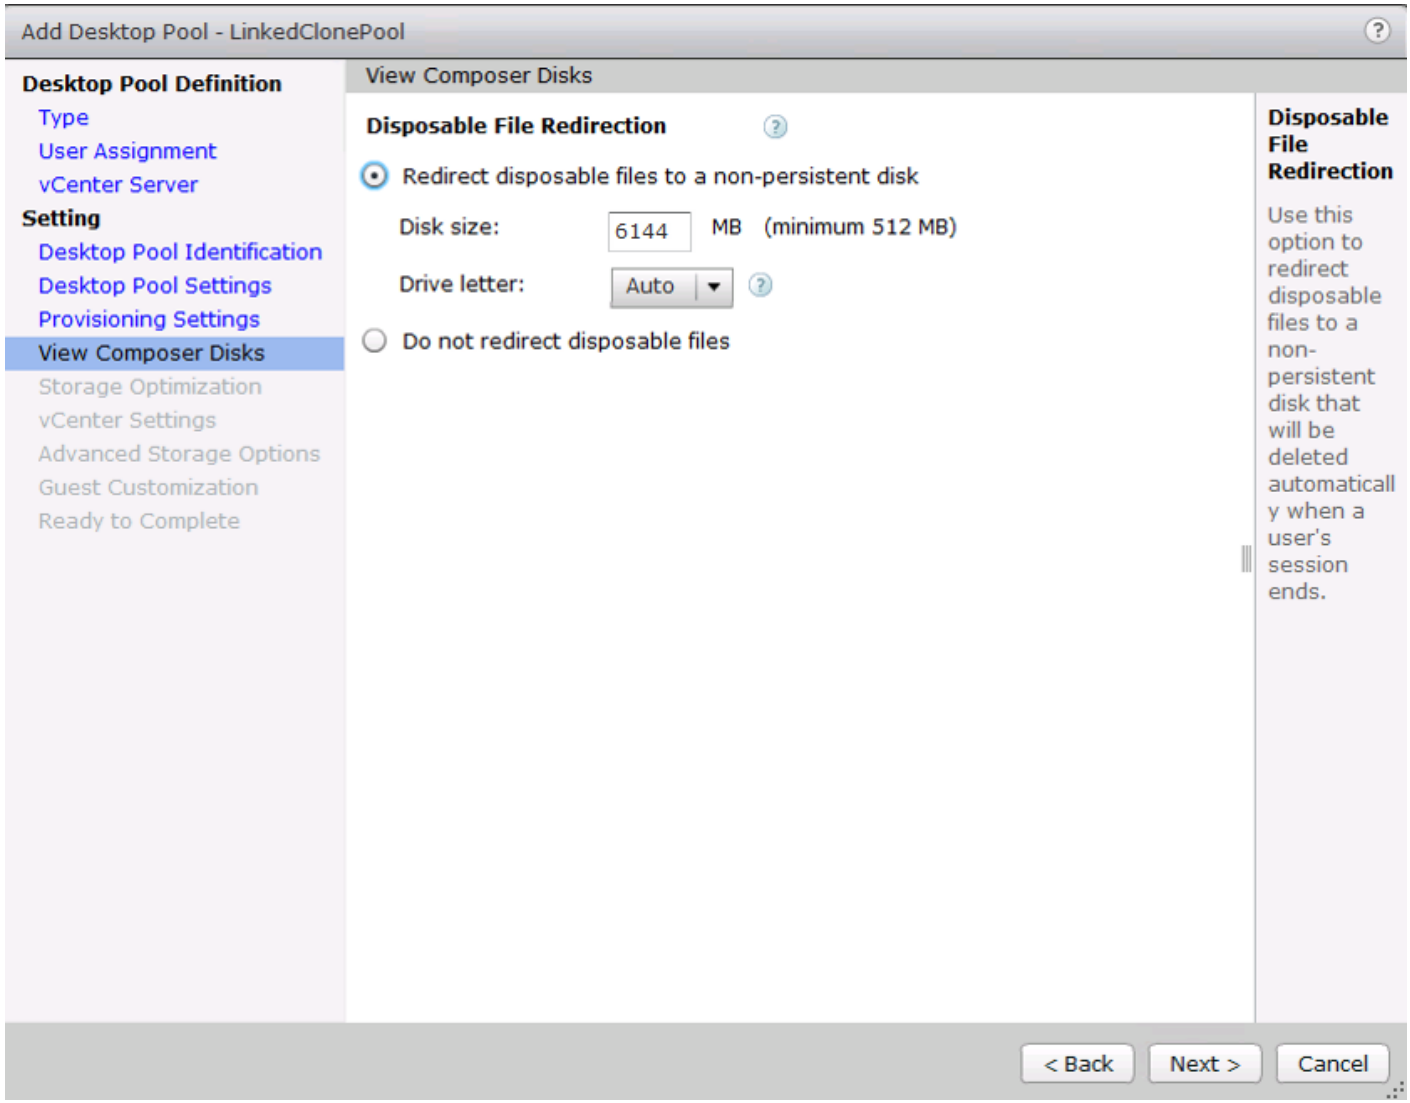

17. Click Next.

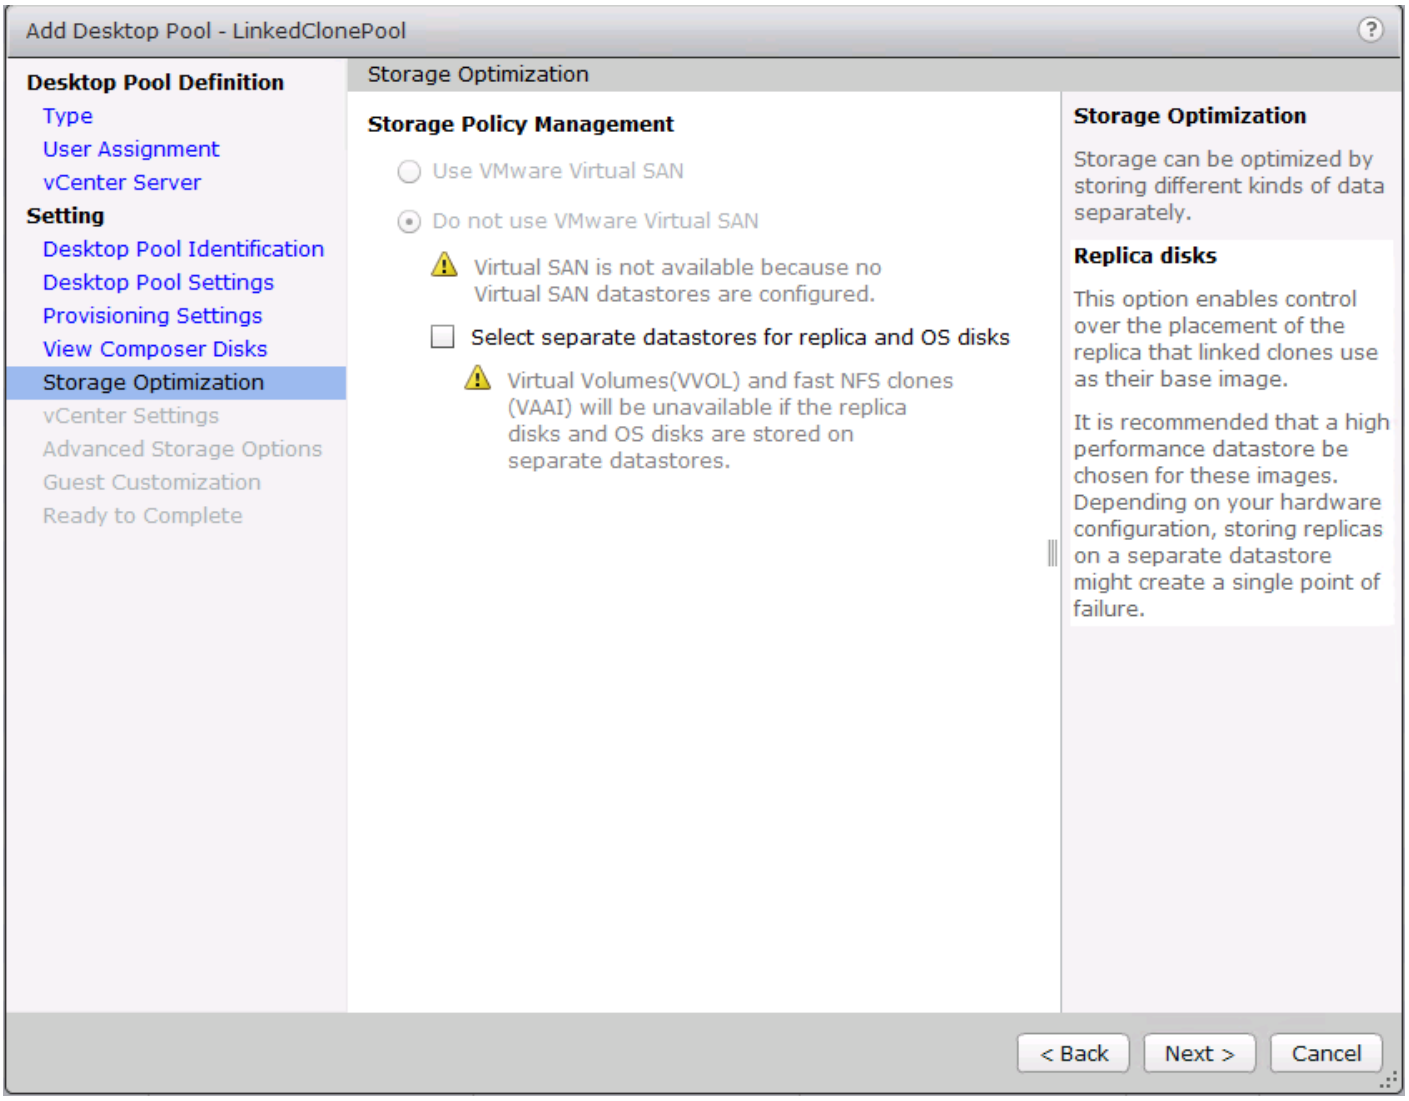

18. Select each of the six required vCenter Settings by using the Browse button next to each field.

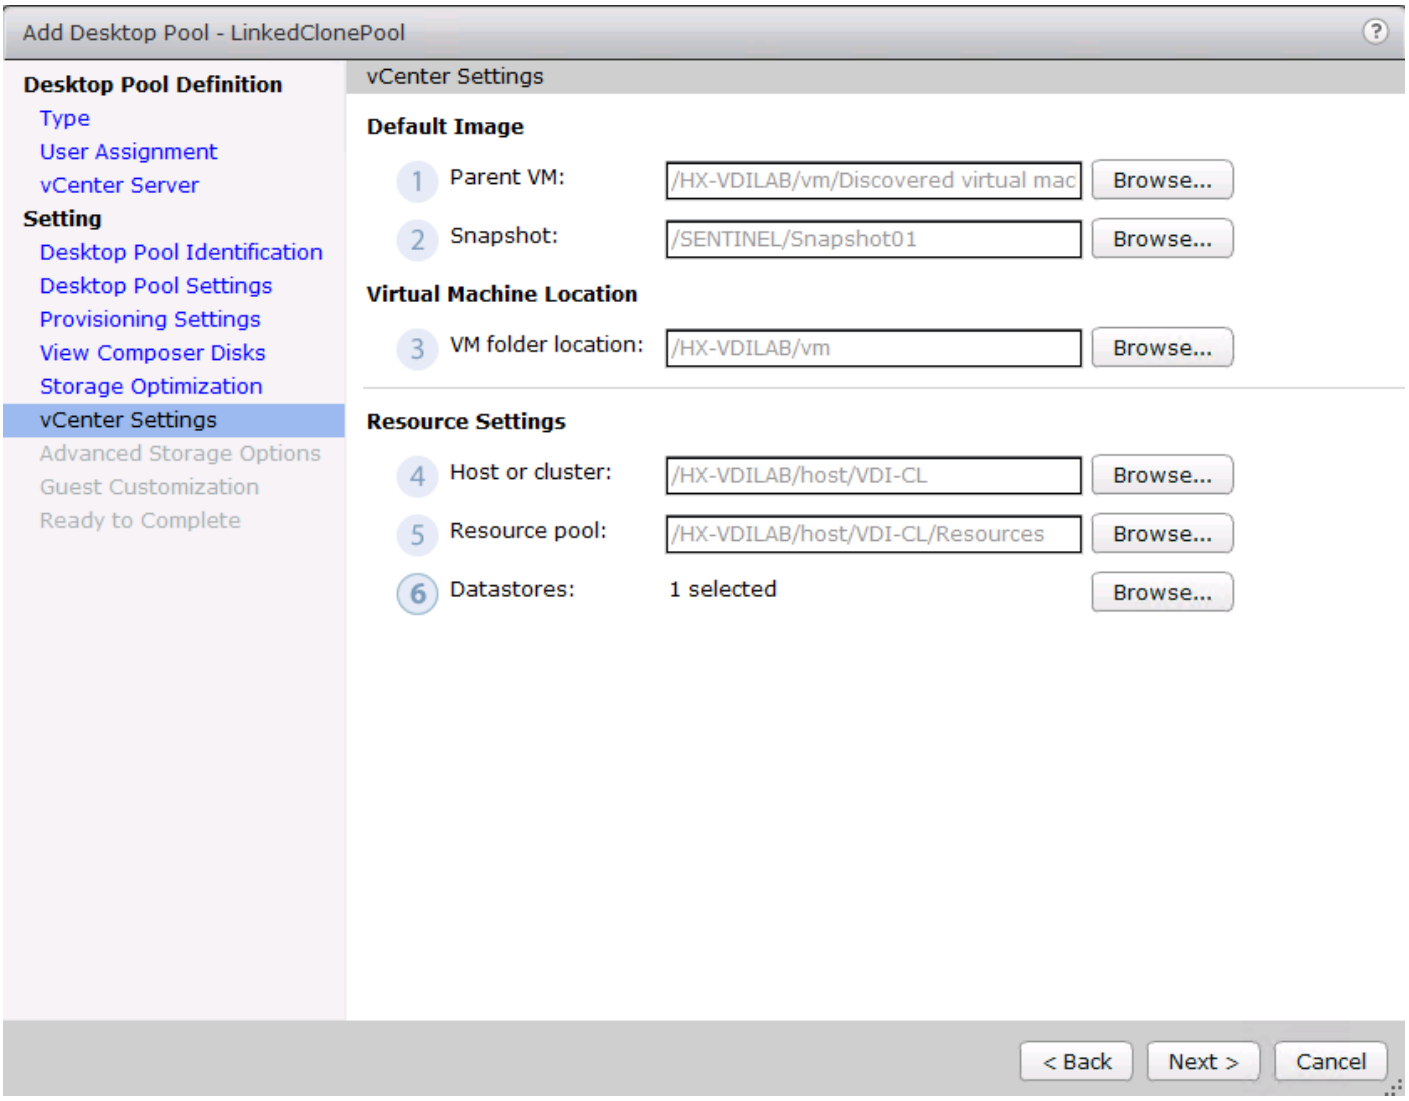

19. For Datastore selection, select the correct datastore and set the Storage Overcommit as “Unbounded.”

20. Click OK.

21. Click Next.

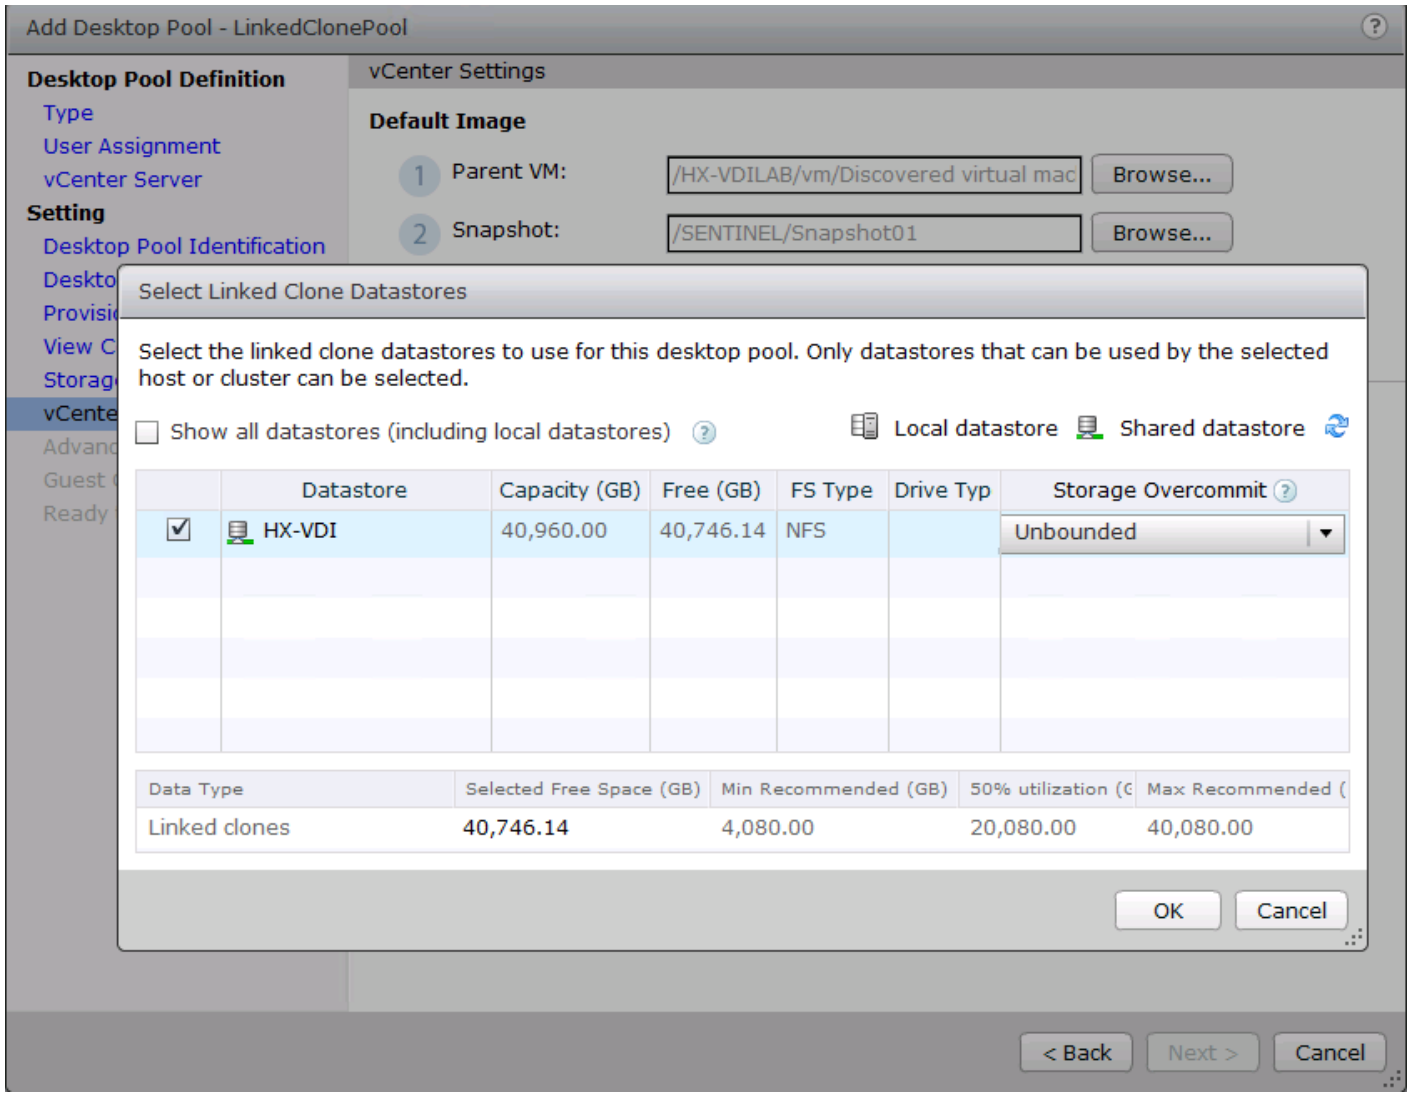

22. Set the Advanced Storage Options using the settings in the following screenshot.

23. Click Next.

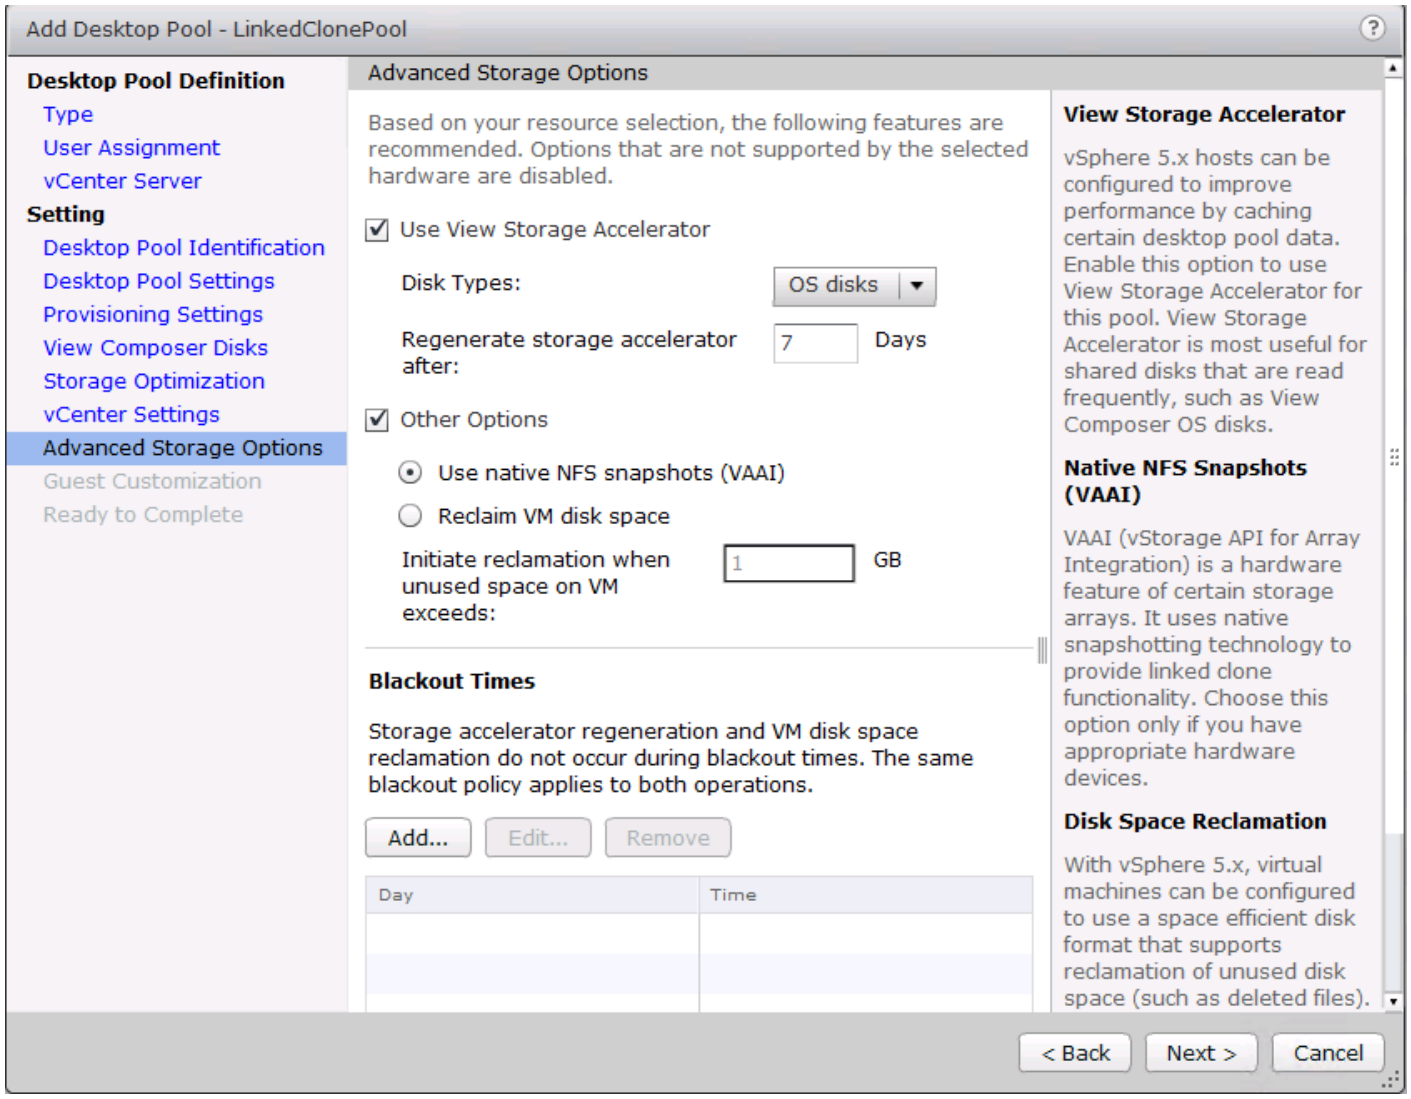

24. Select Guest optimization settings.

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.

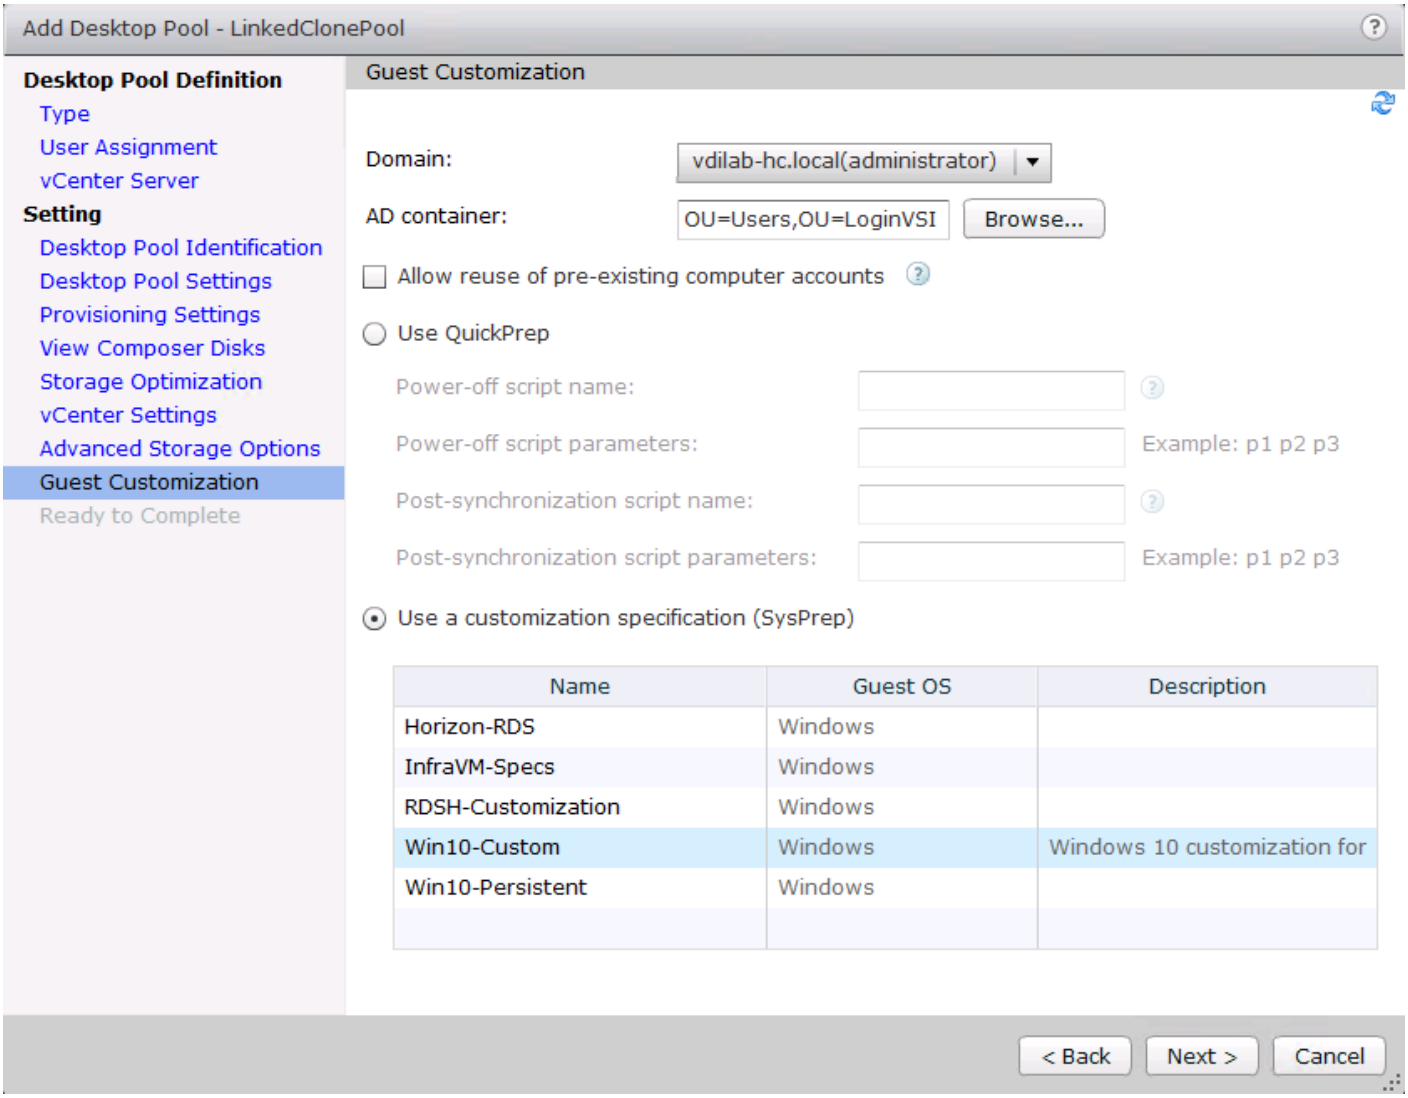

27. Select the checkbox “Entitle users after pool creation wizard completion” if you would like to authorize users as part of this process. Follow instructions provided in the Create Horizon 7 RDS Desktop Pool to authorize users for the Linked Clone Pool.

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

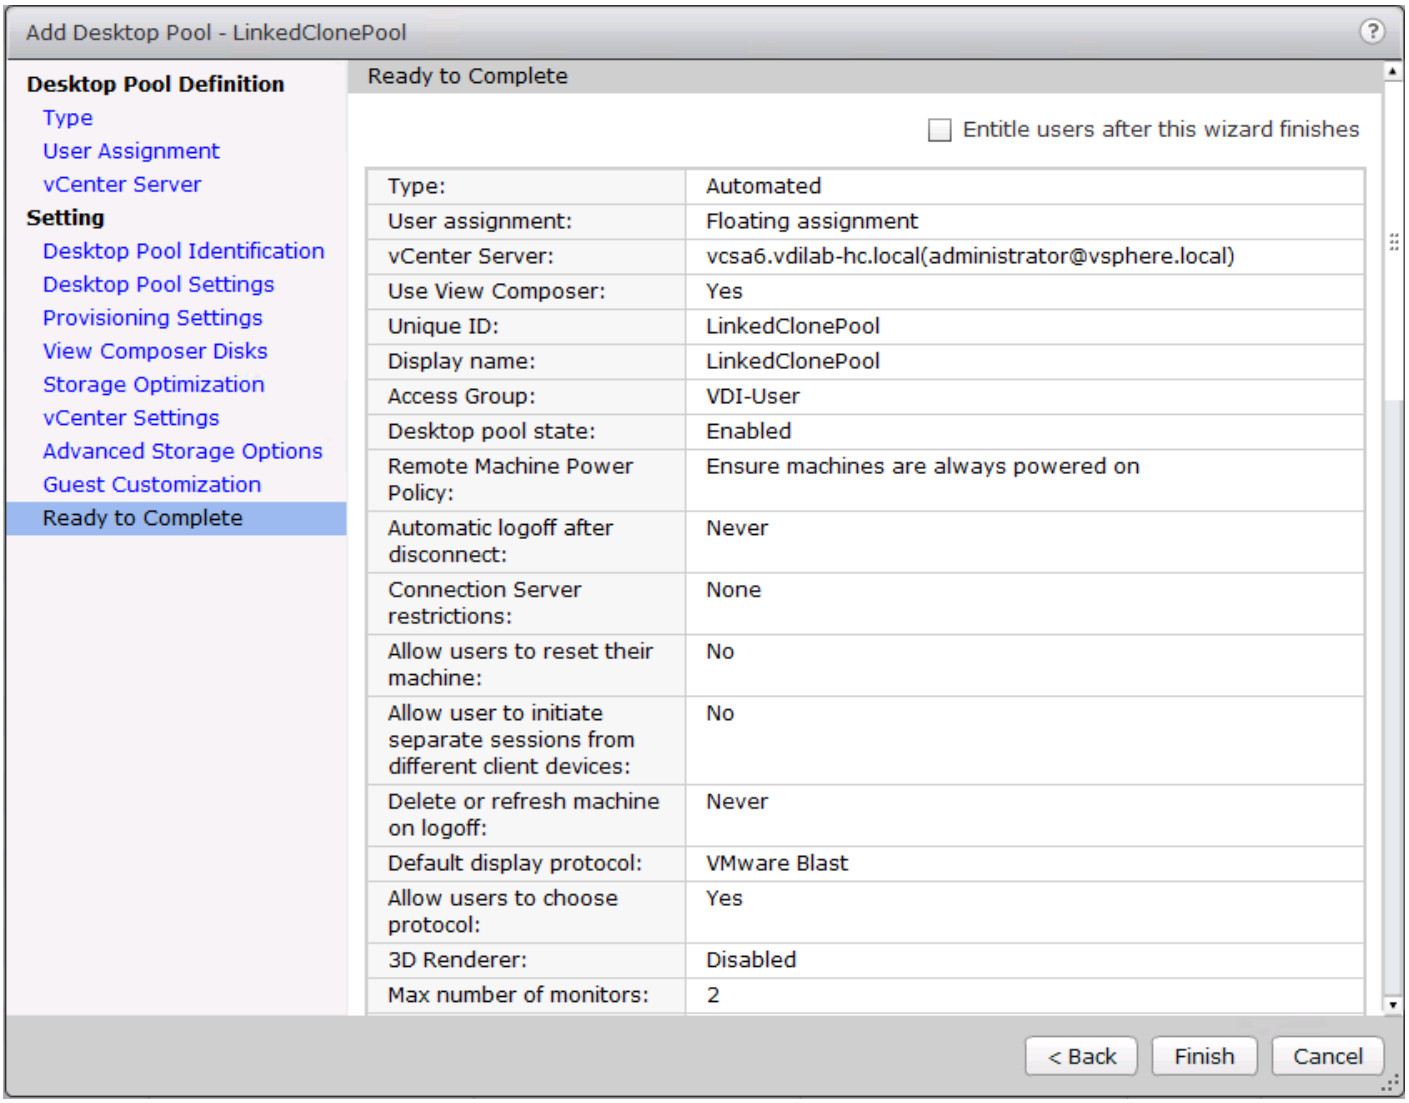

VMware Horizon Instant-Clone Windows 10 Desktop Pool Creation

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.

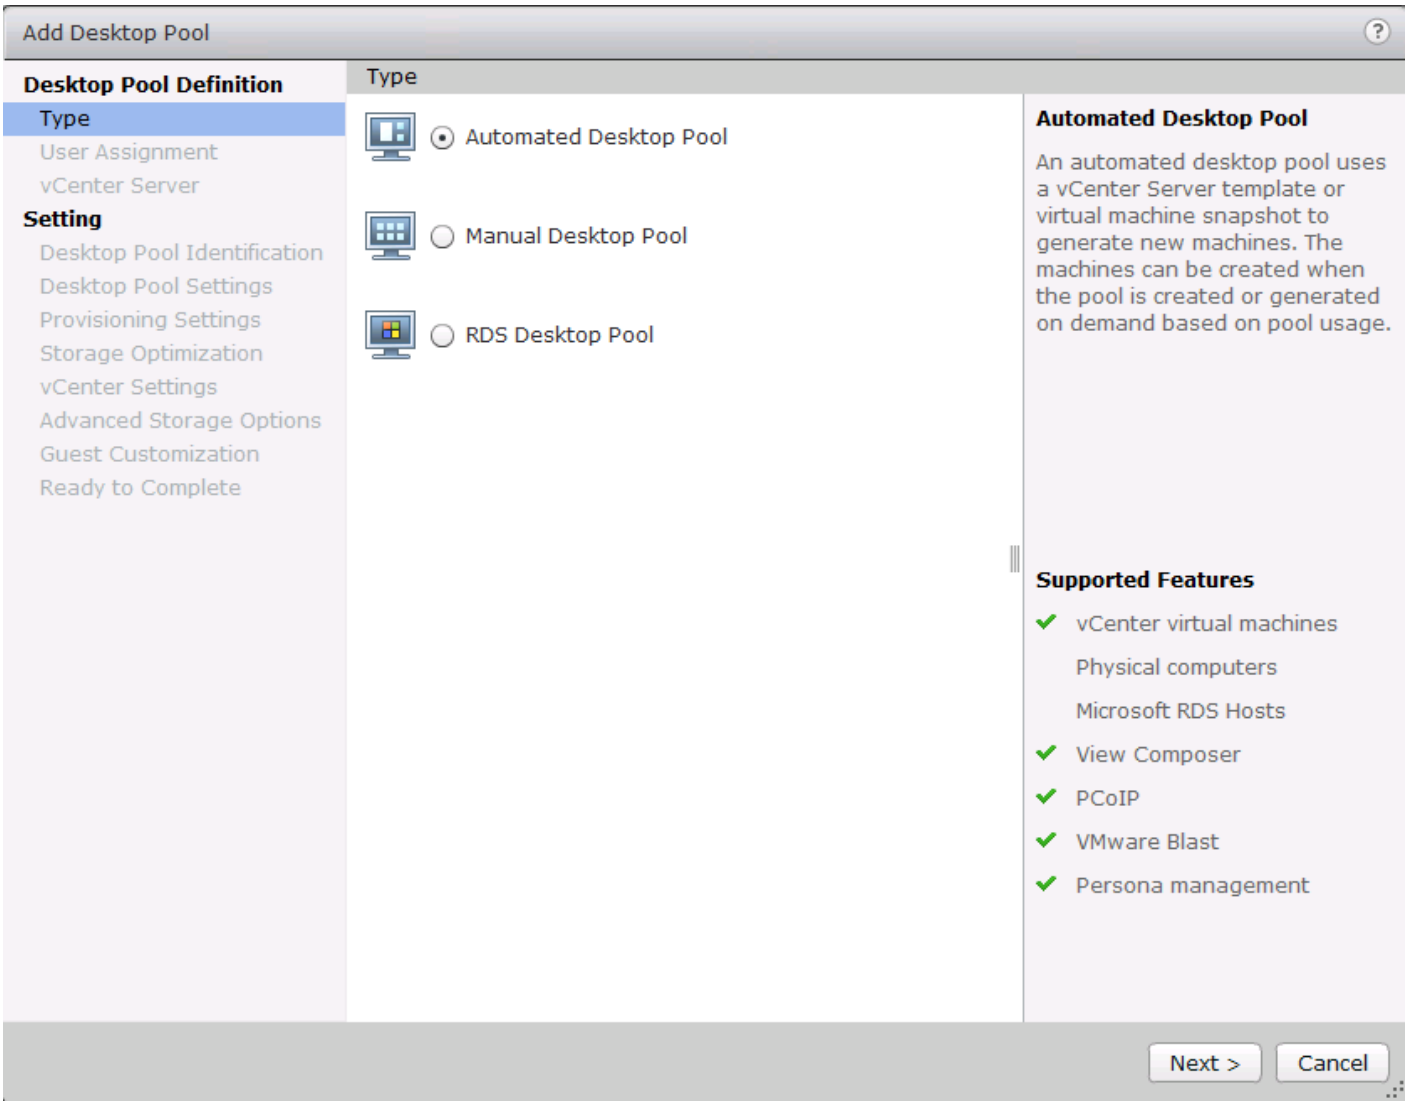

5. Select Floating or Dedicate user assignment.

6. Click Next.

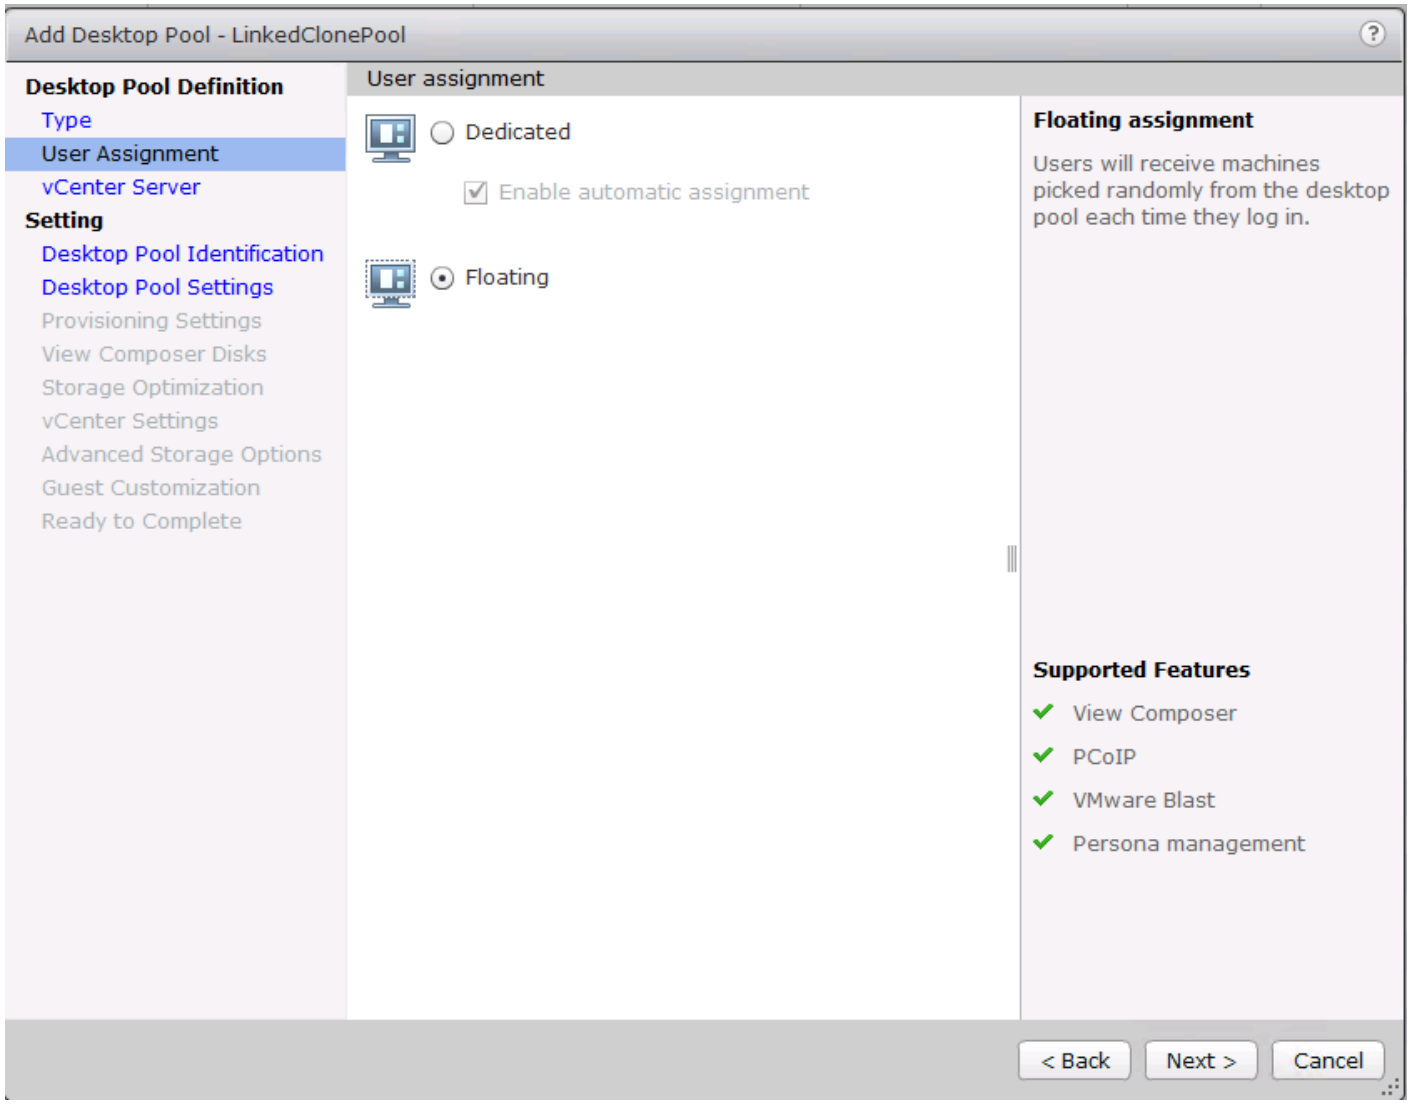

7. Select Instant Clones, highlight your vCenter server, then click Next.

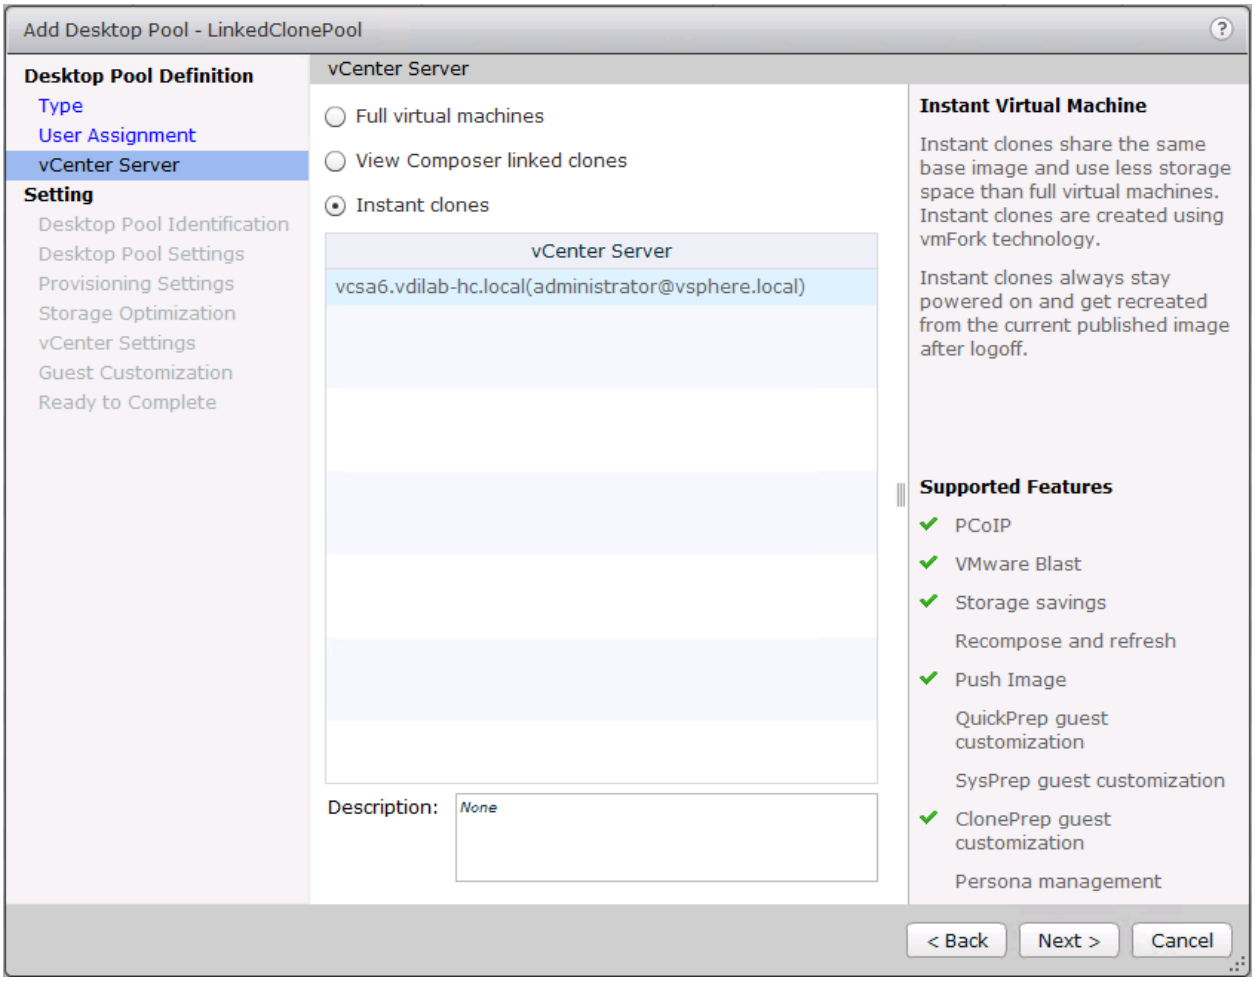

8. Enter pool identification details.

9. Click Next.

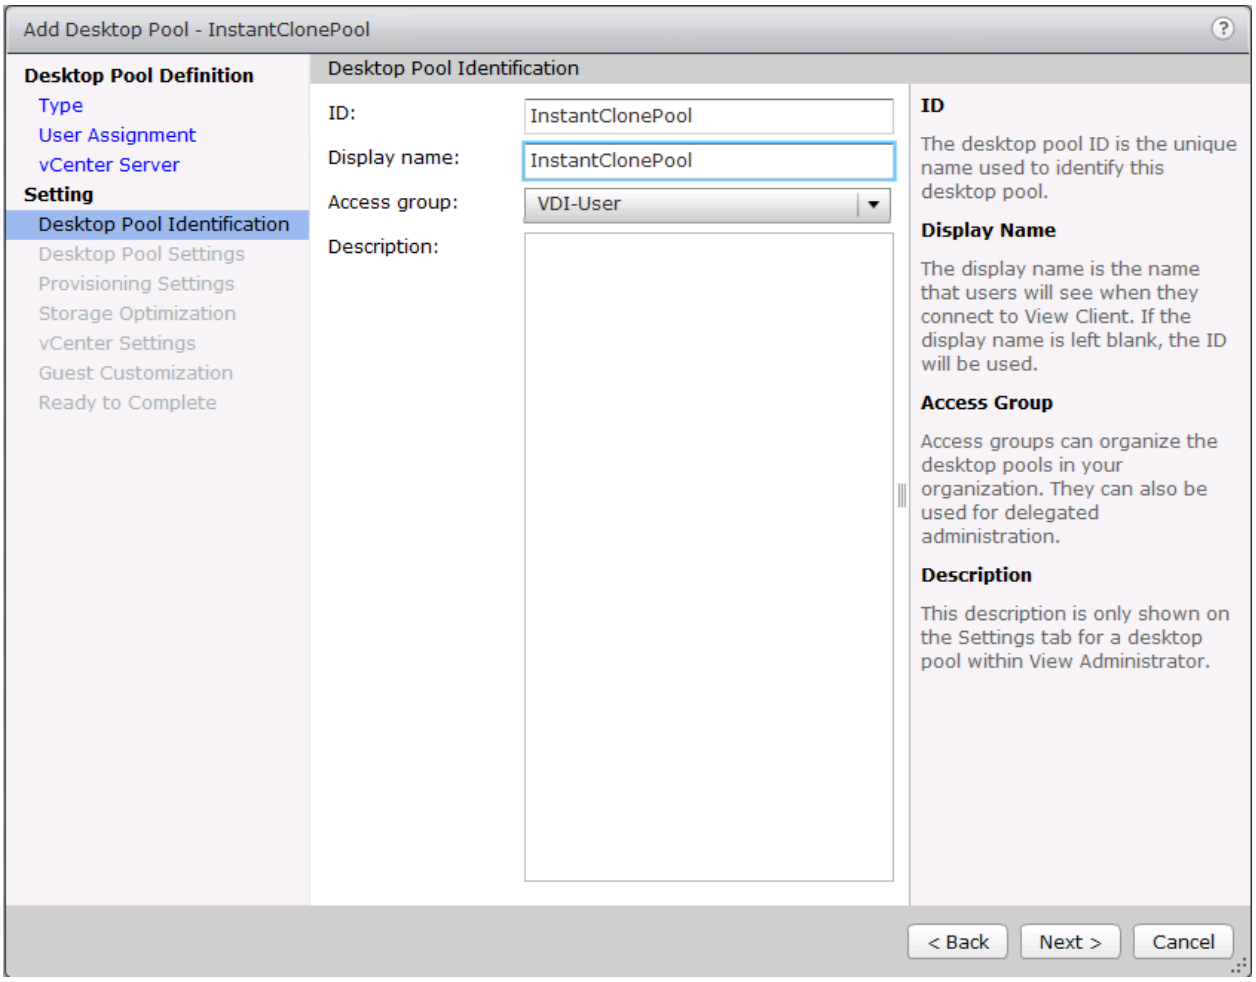

10. Select Desktop Pool settings.

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

11. Click Next.

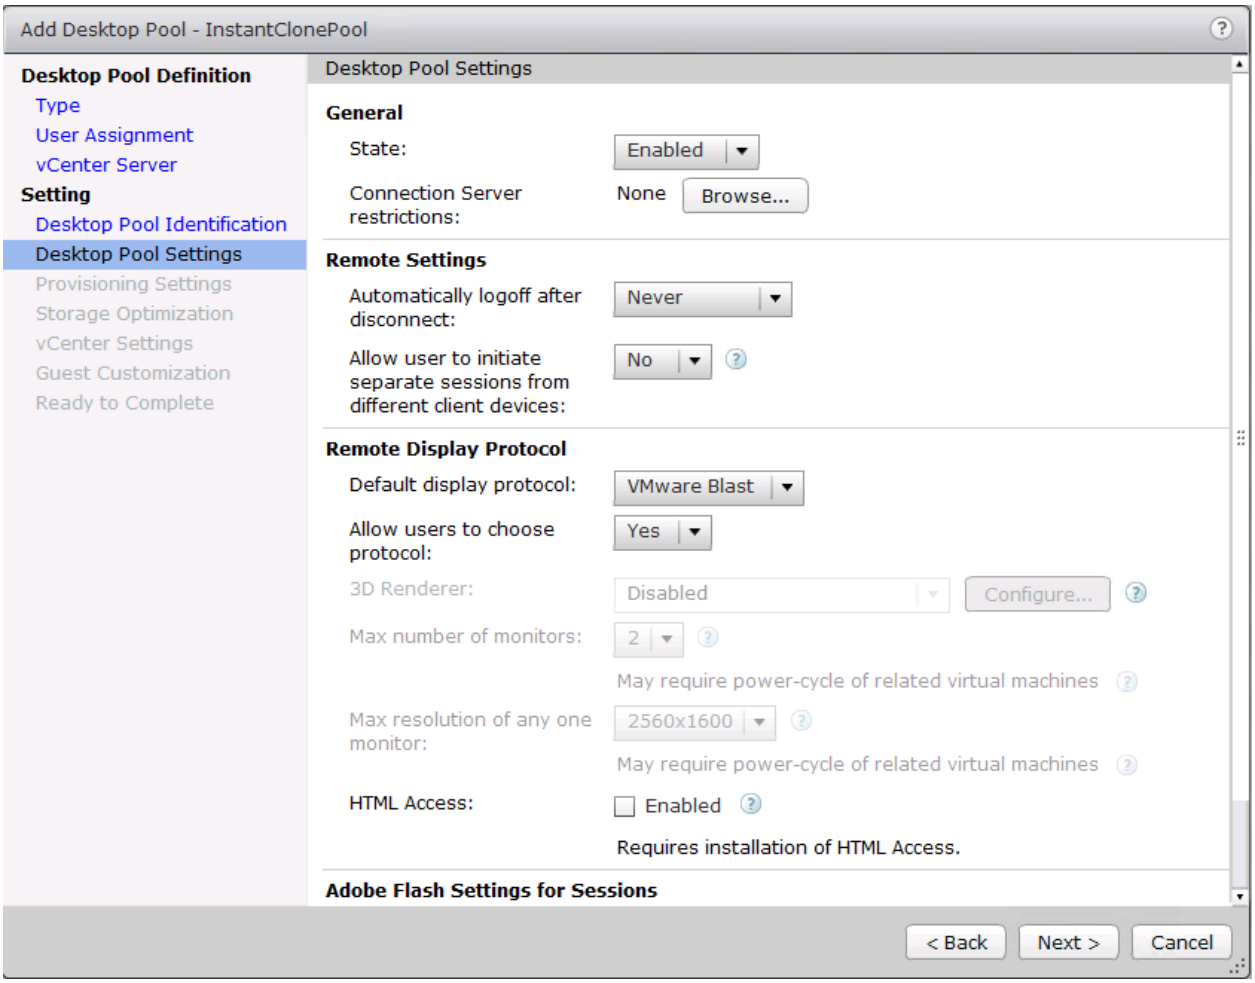

12. Select provisioning settings.

13. Click Next.

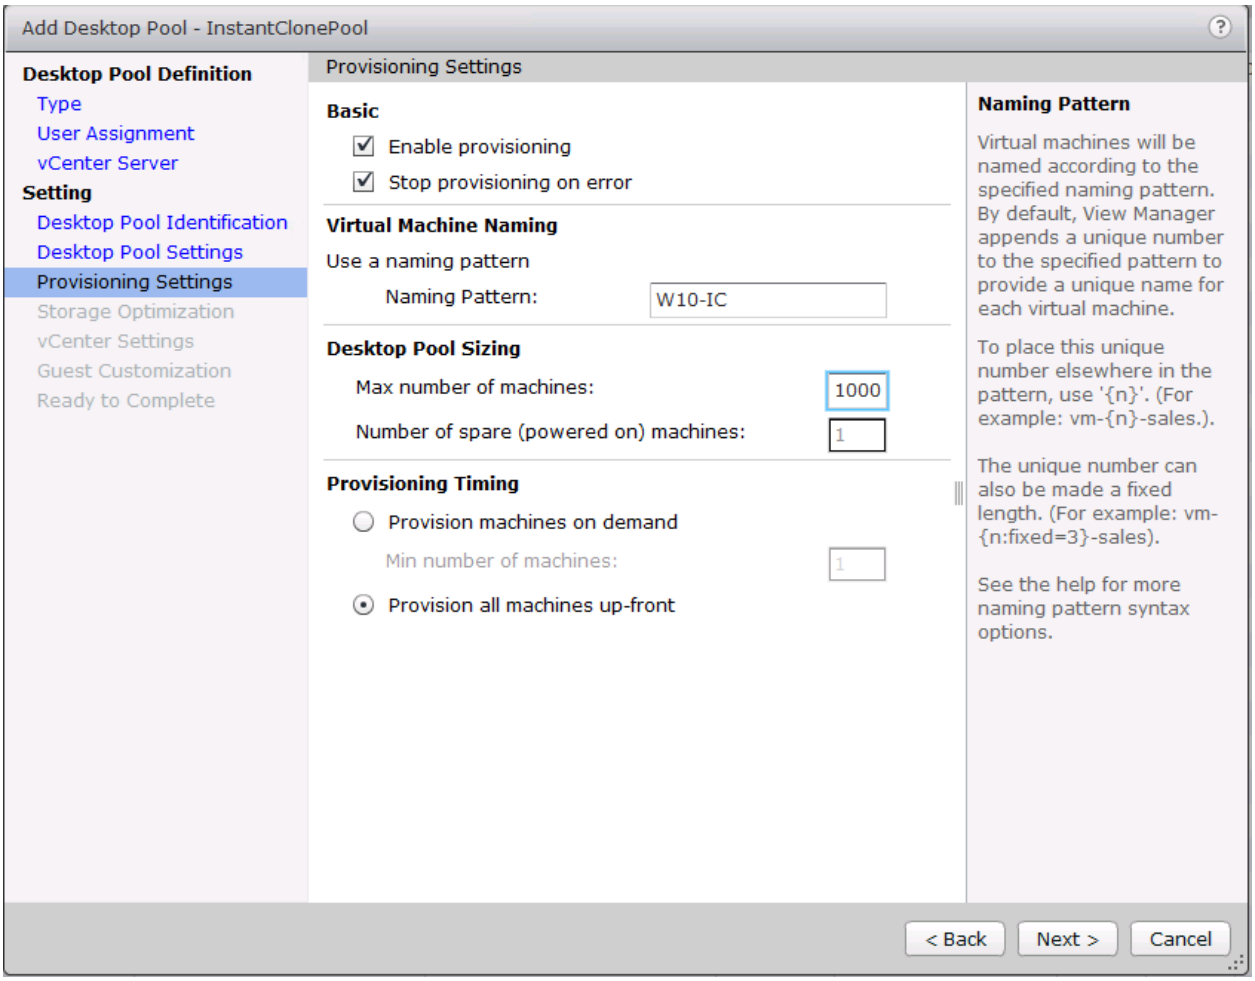

14. Click Next.

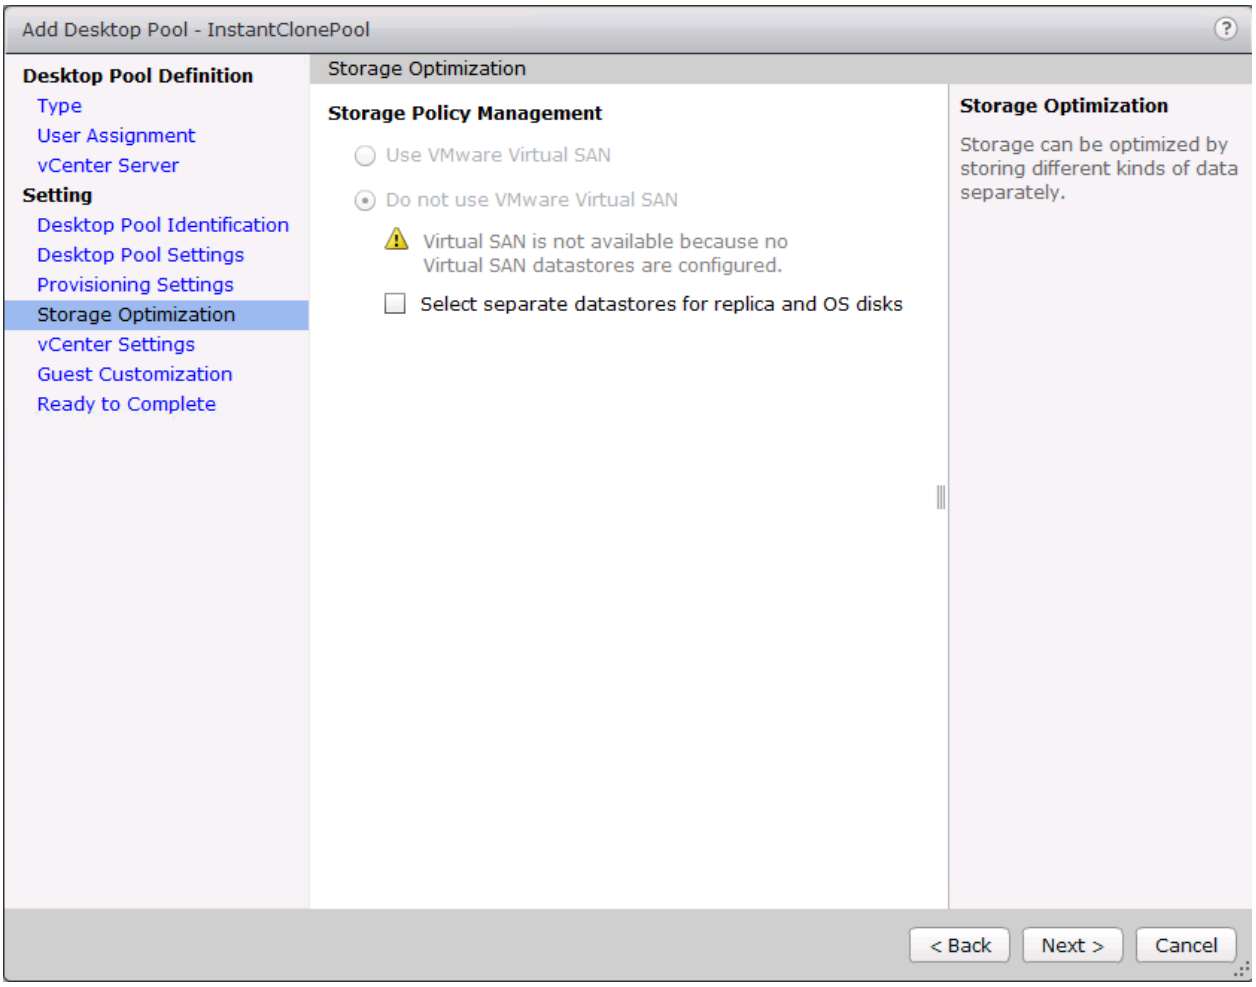

15. Select the vCenter Settings and browse for each of the six required inputs.

16. Click Next.

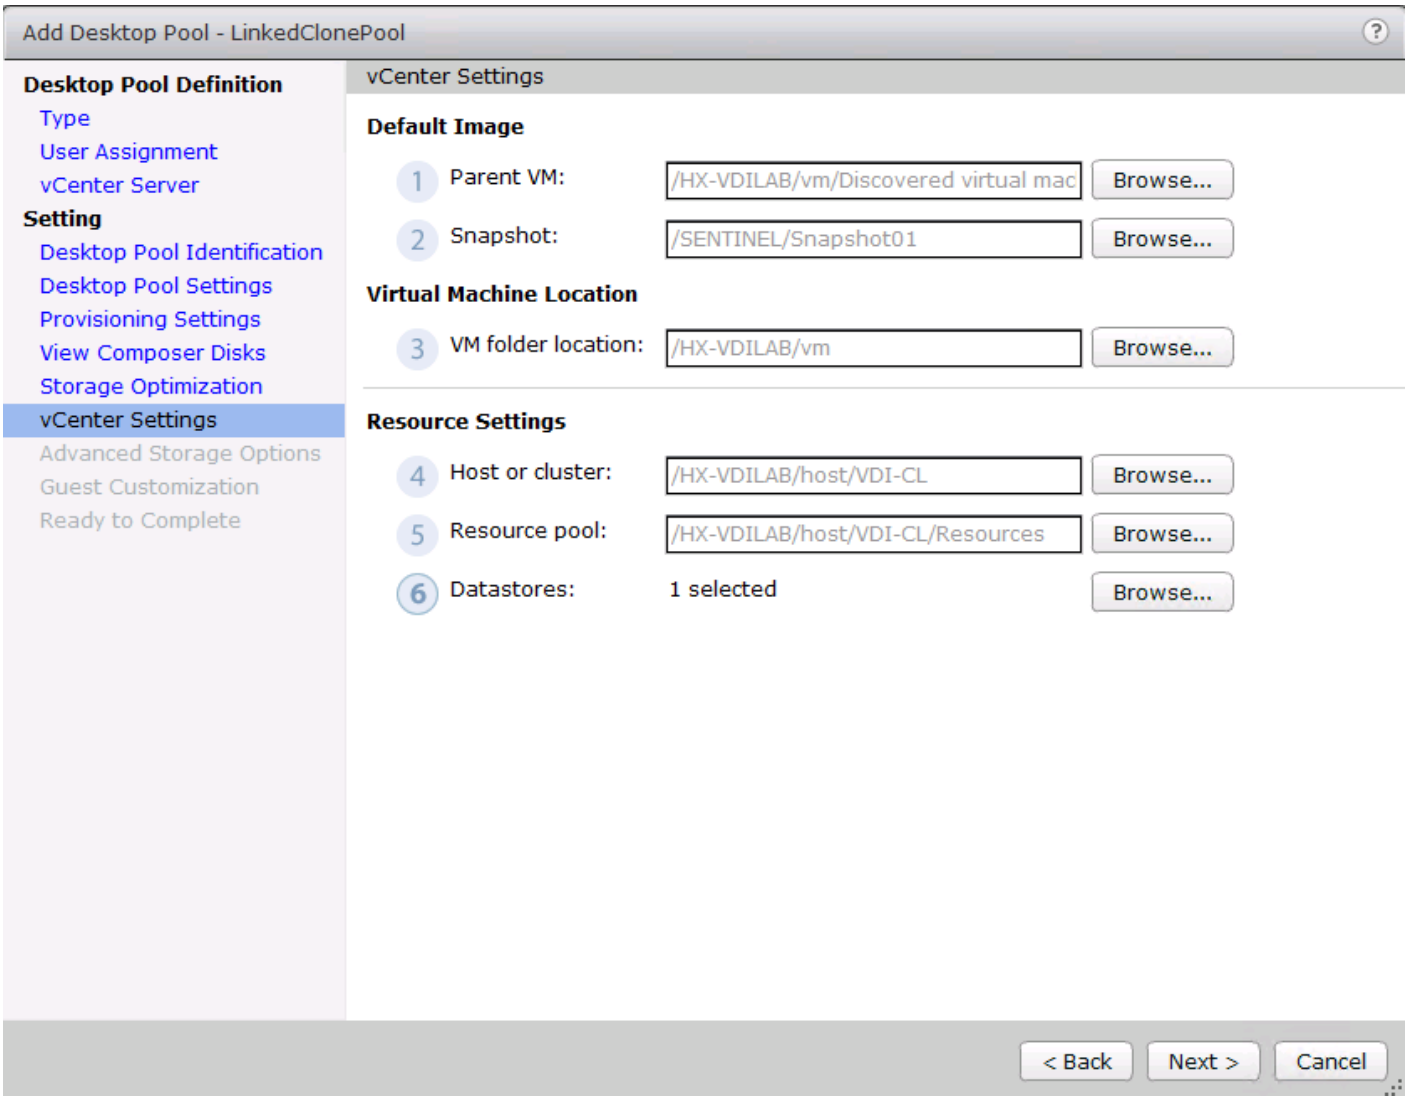

17. For Datastore, select the datastore with the storage overcommit as “Unbounded”.

18. Click OK.

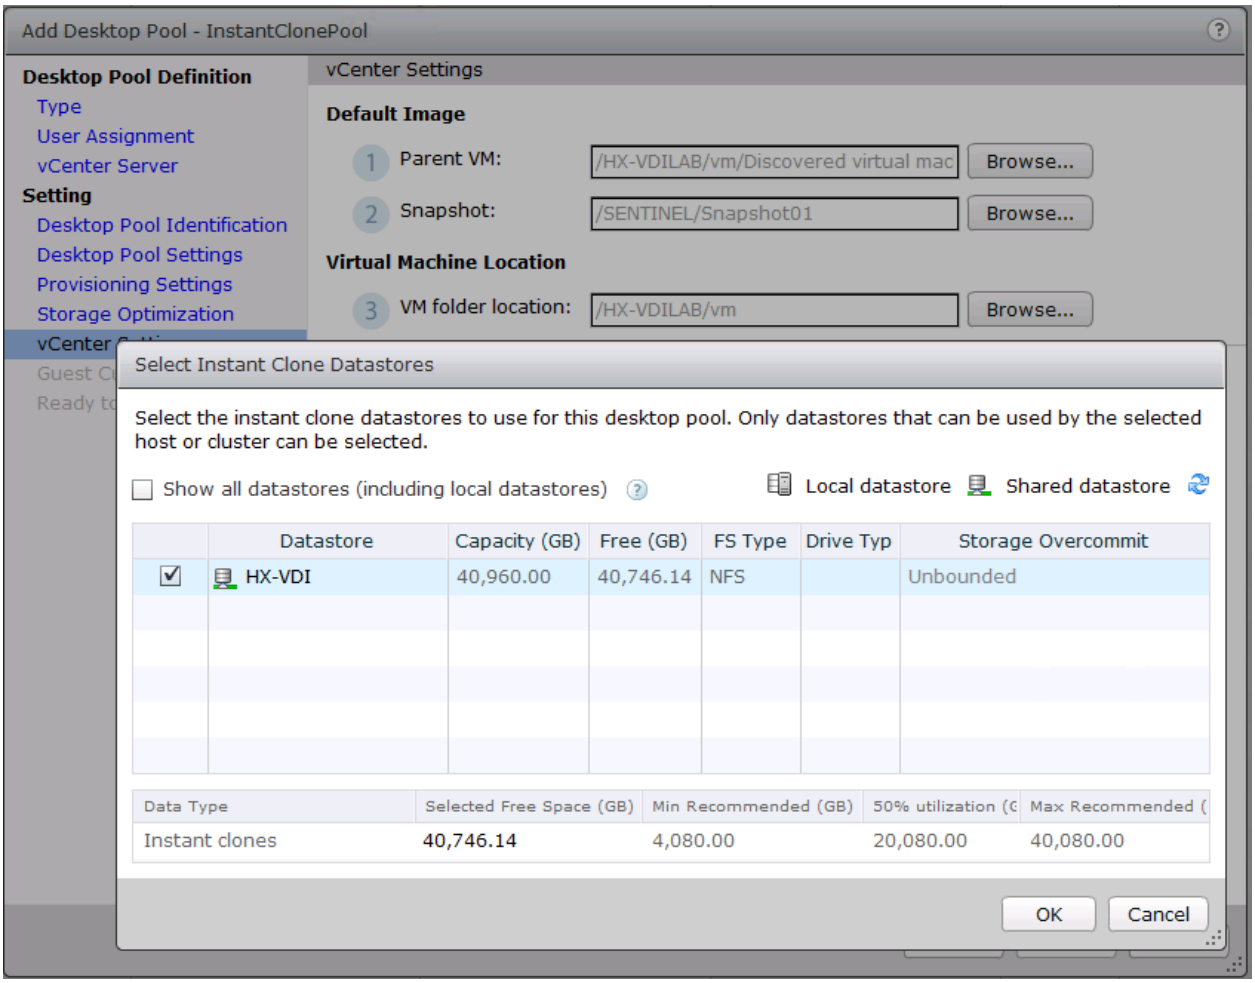

19. Select Guest Customization.

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

21. Click Next.

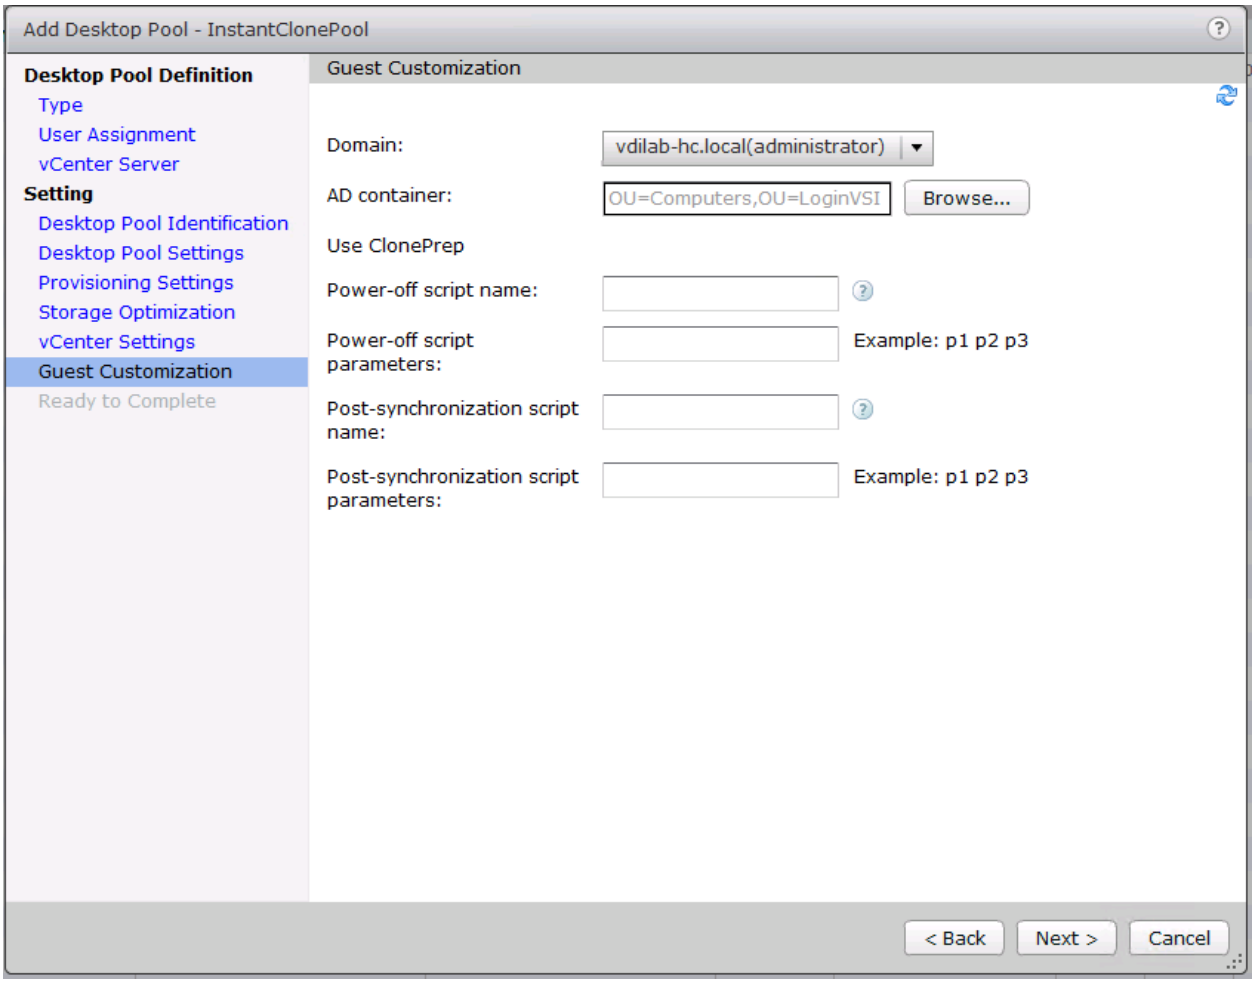

22. Review the summary of the pool configuration.

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

24. Click Finish.

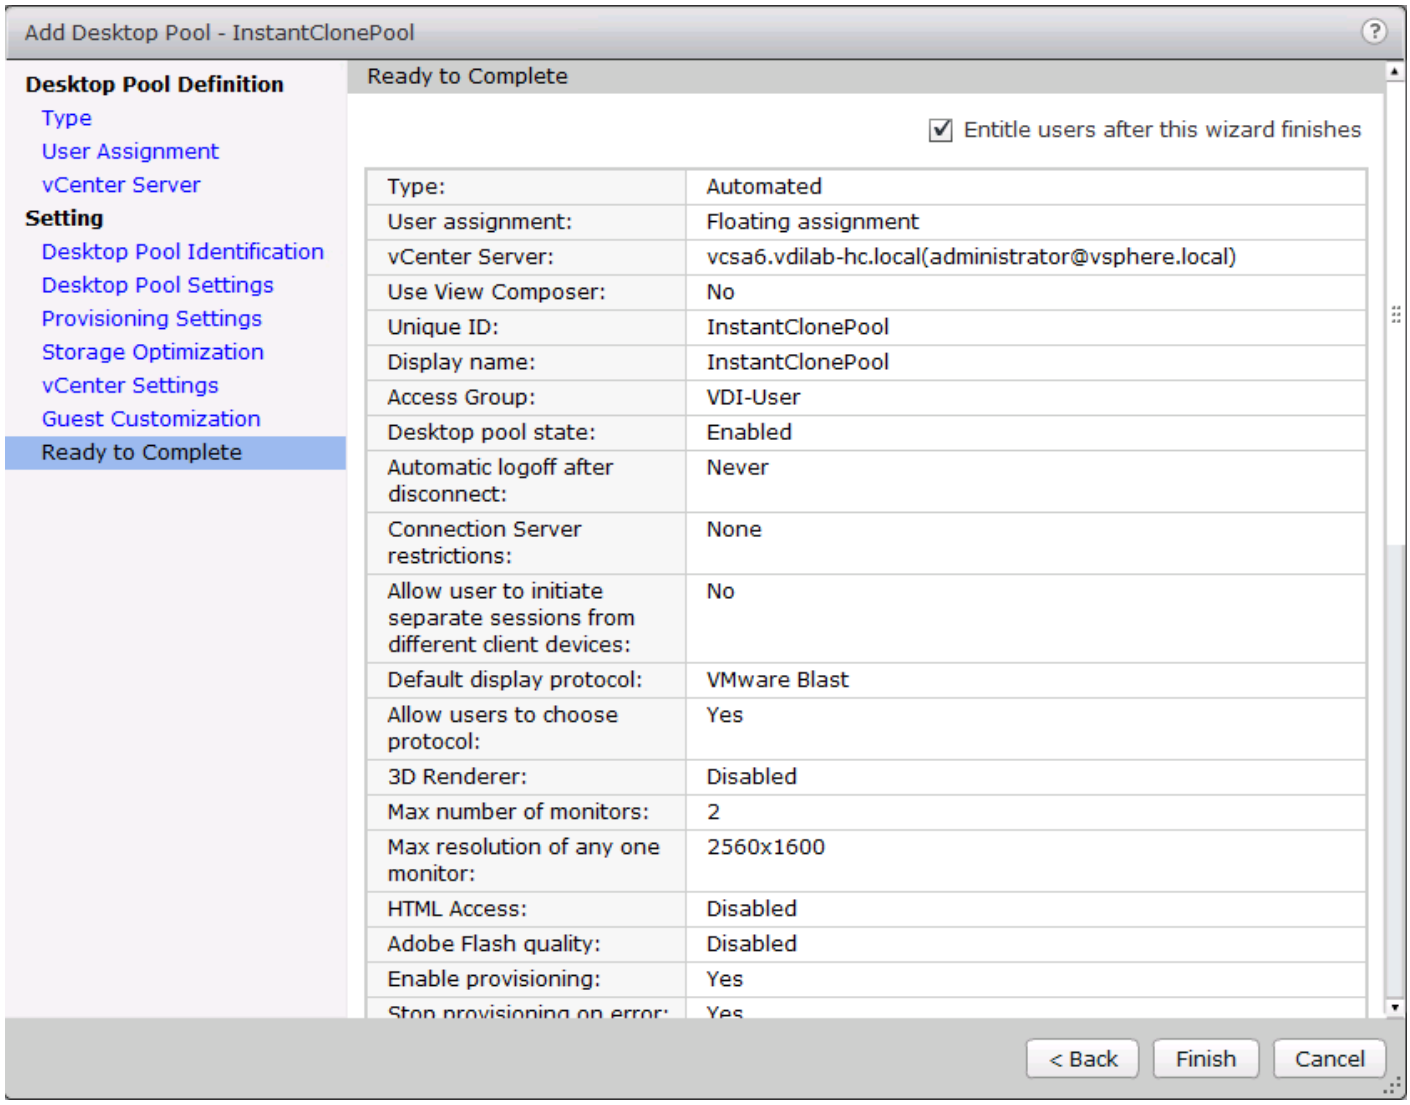

25. Follow the instructions provided in the Create Horizon 7 RDS Desktop Pool to authorize users for the Linked Clone Pool.

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.

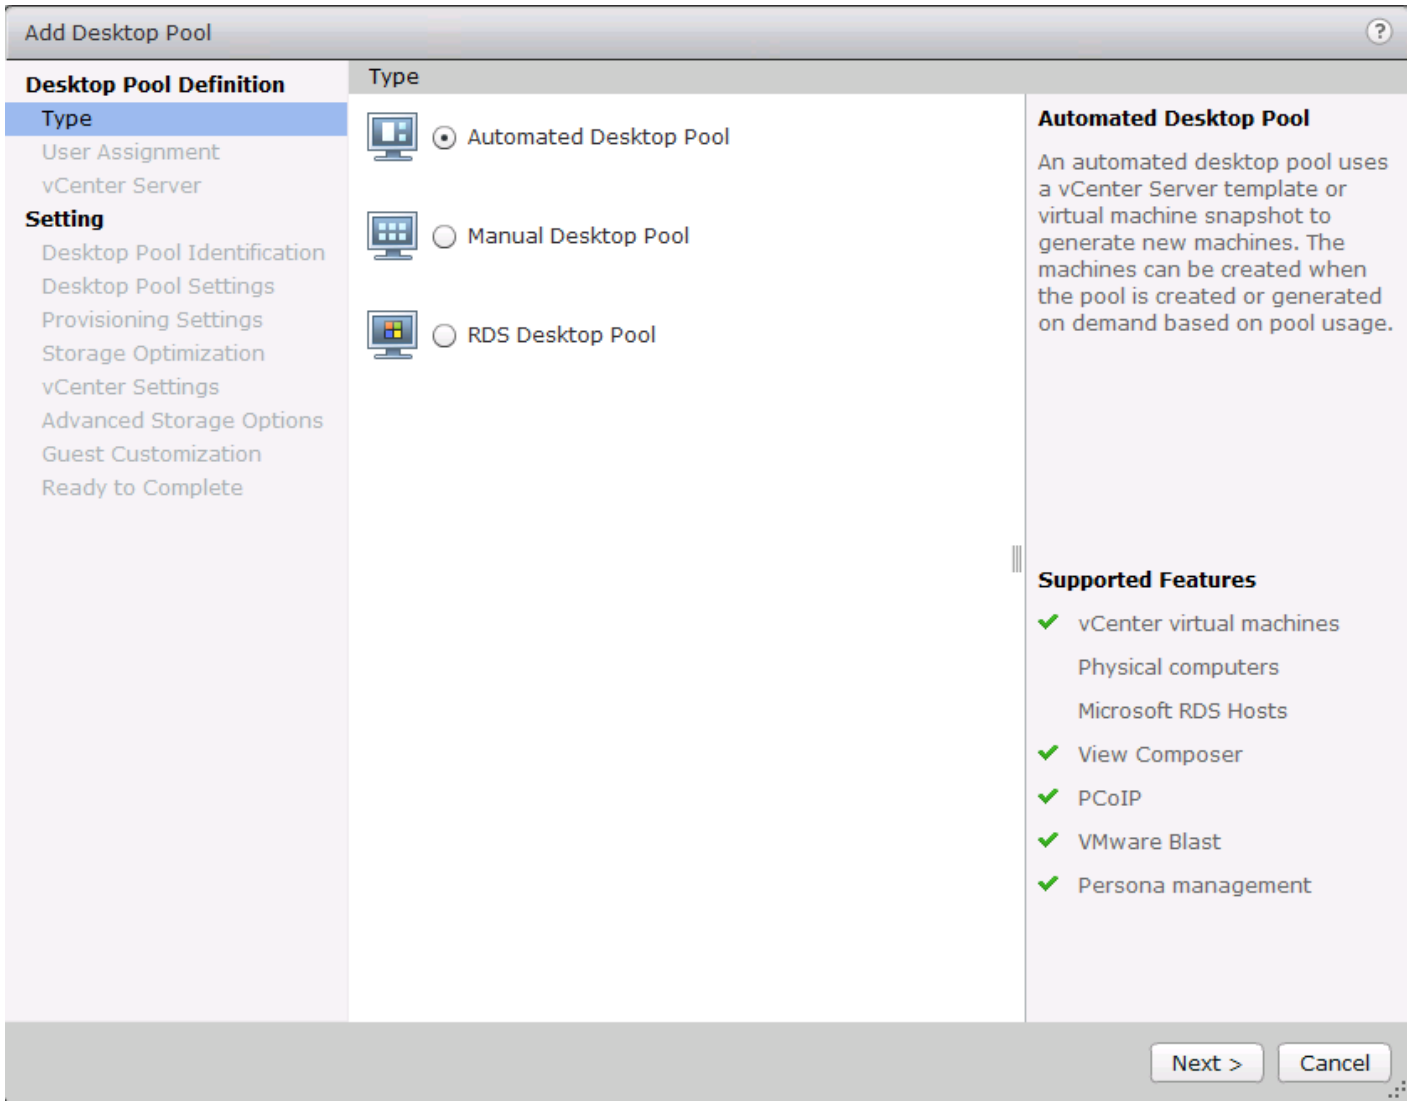

5. Select the Dedicated radio button.
6. Select the Enable automatic assignment checkbox if desired.
7. Click Next.

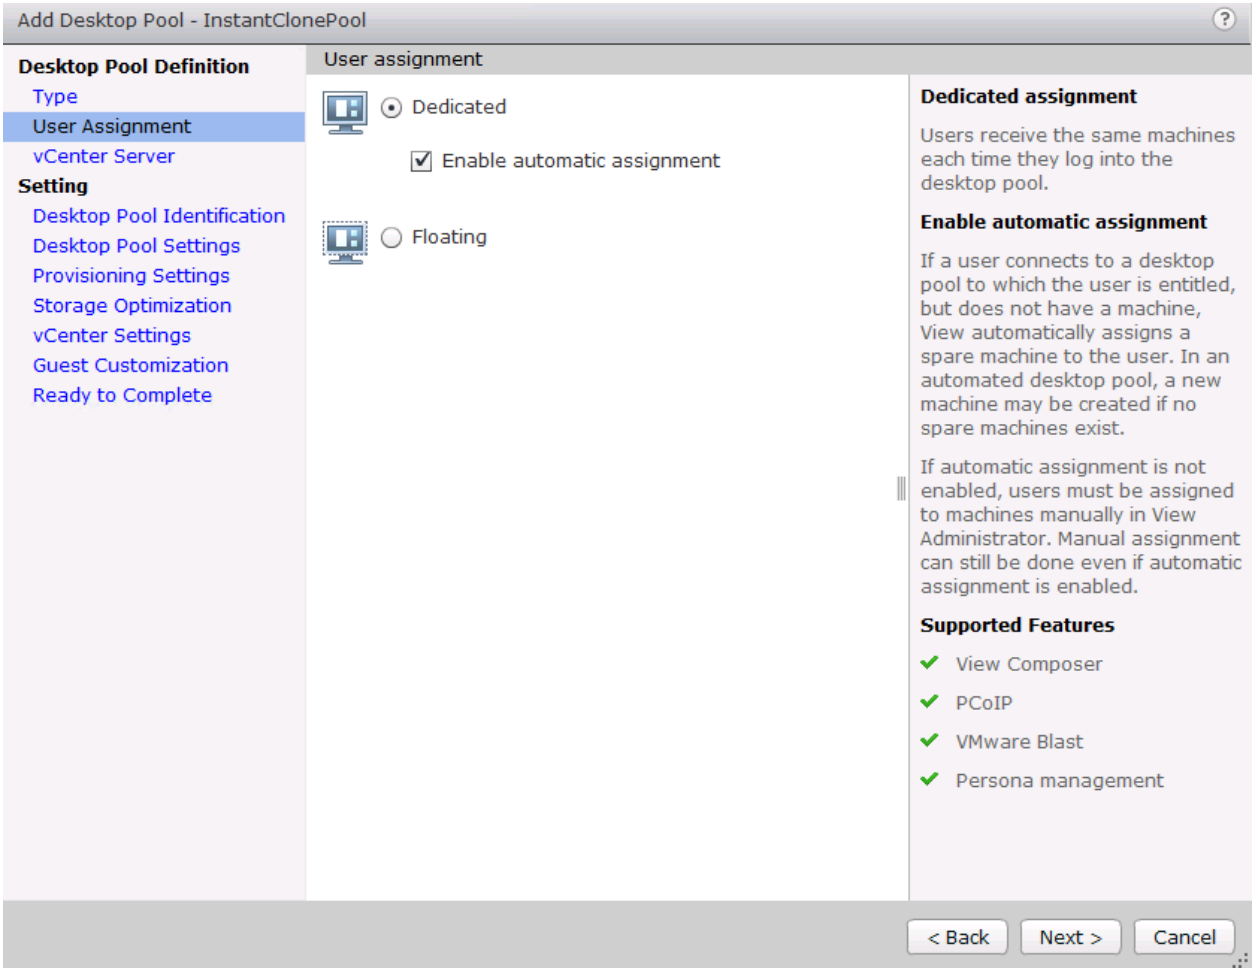

8. Select the Full Virtual Machines radio button and highlight your vCenter and Composer.

9. Click Next.

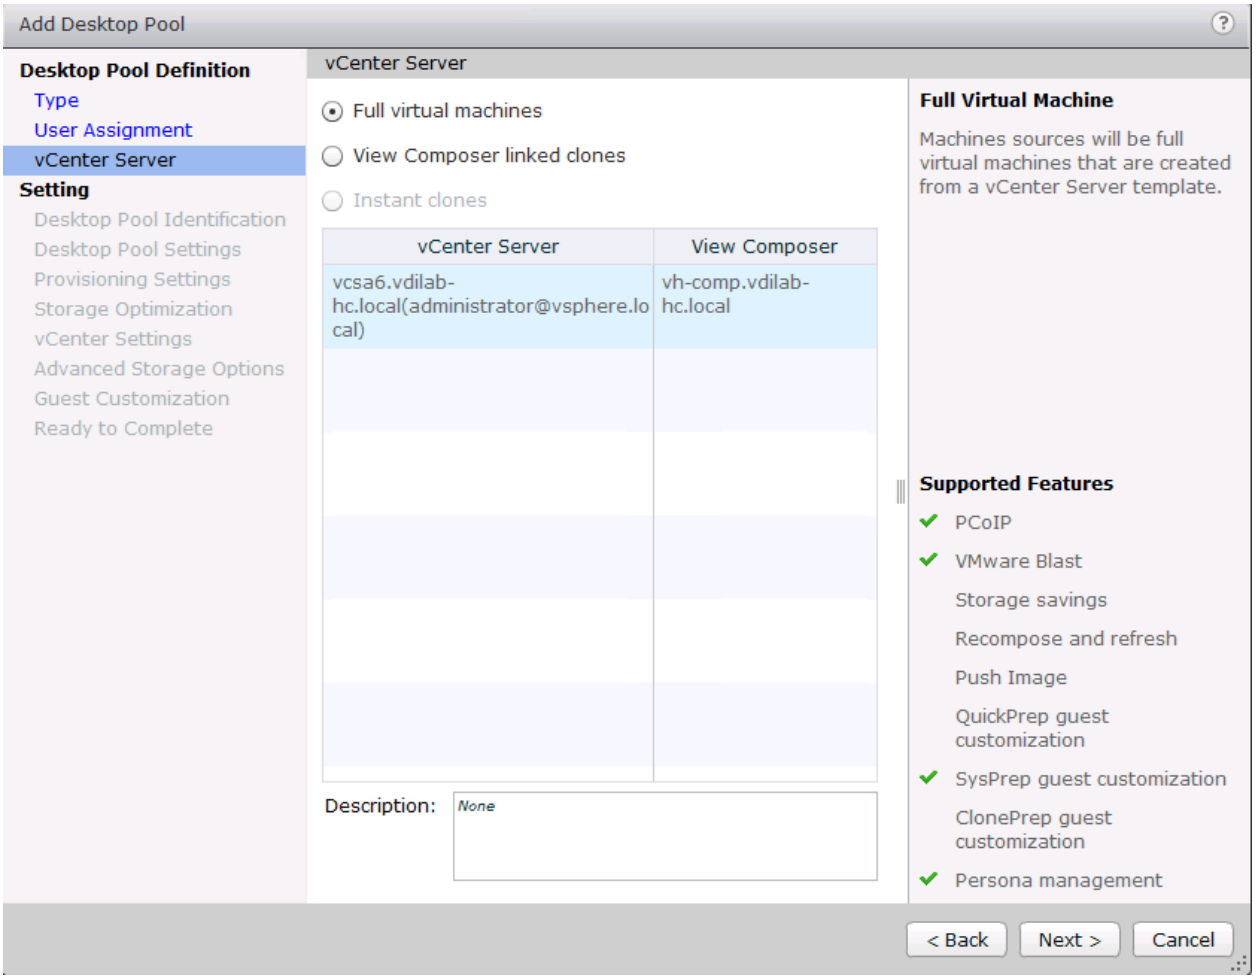

10. Enter the pool identification details.

11. Click Next.

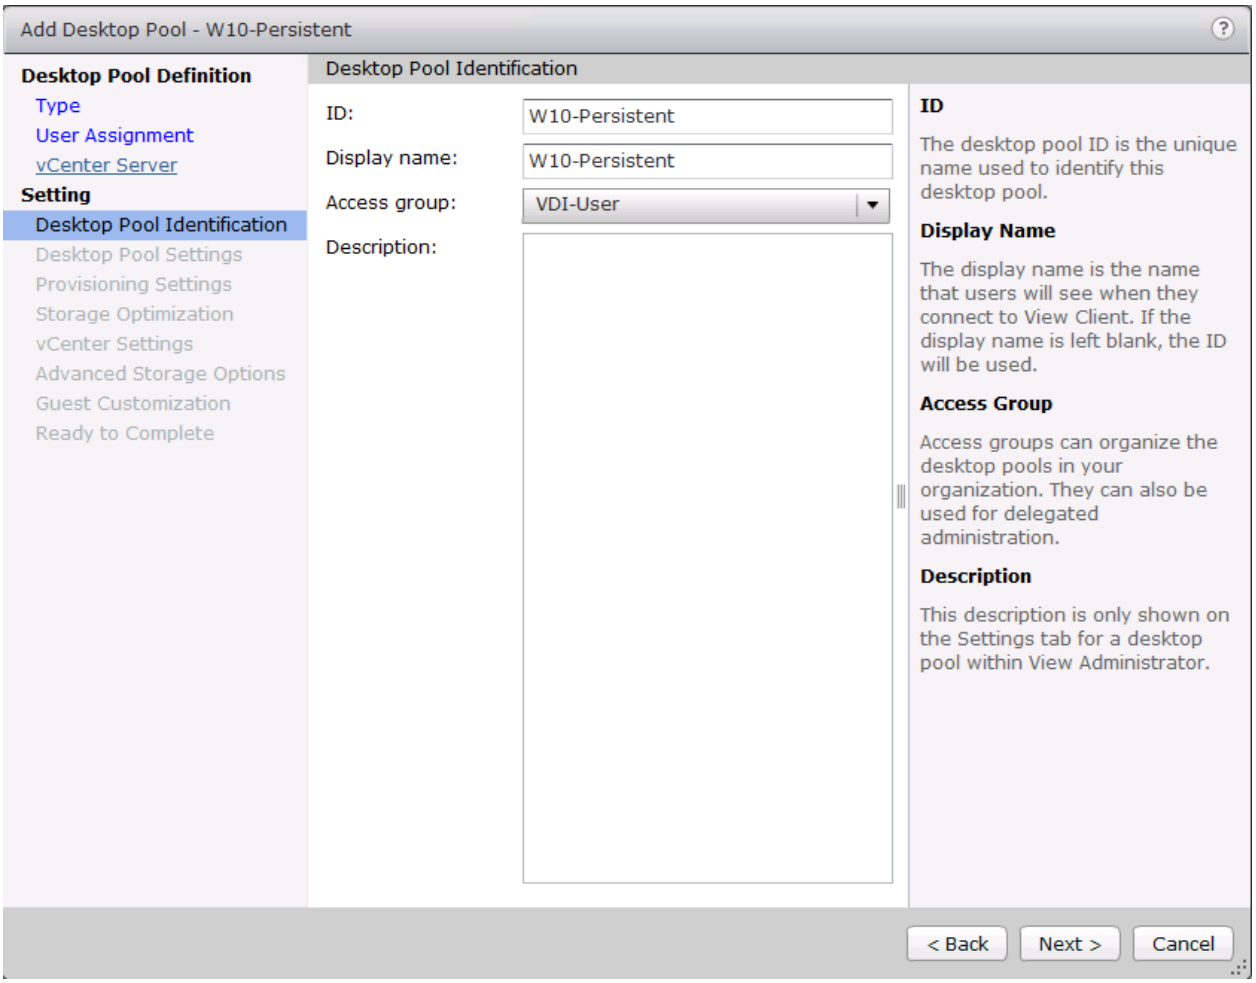

12. Select Desktop Pool settings.

13. Click Next.

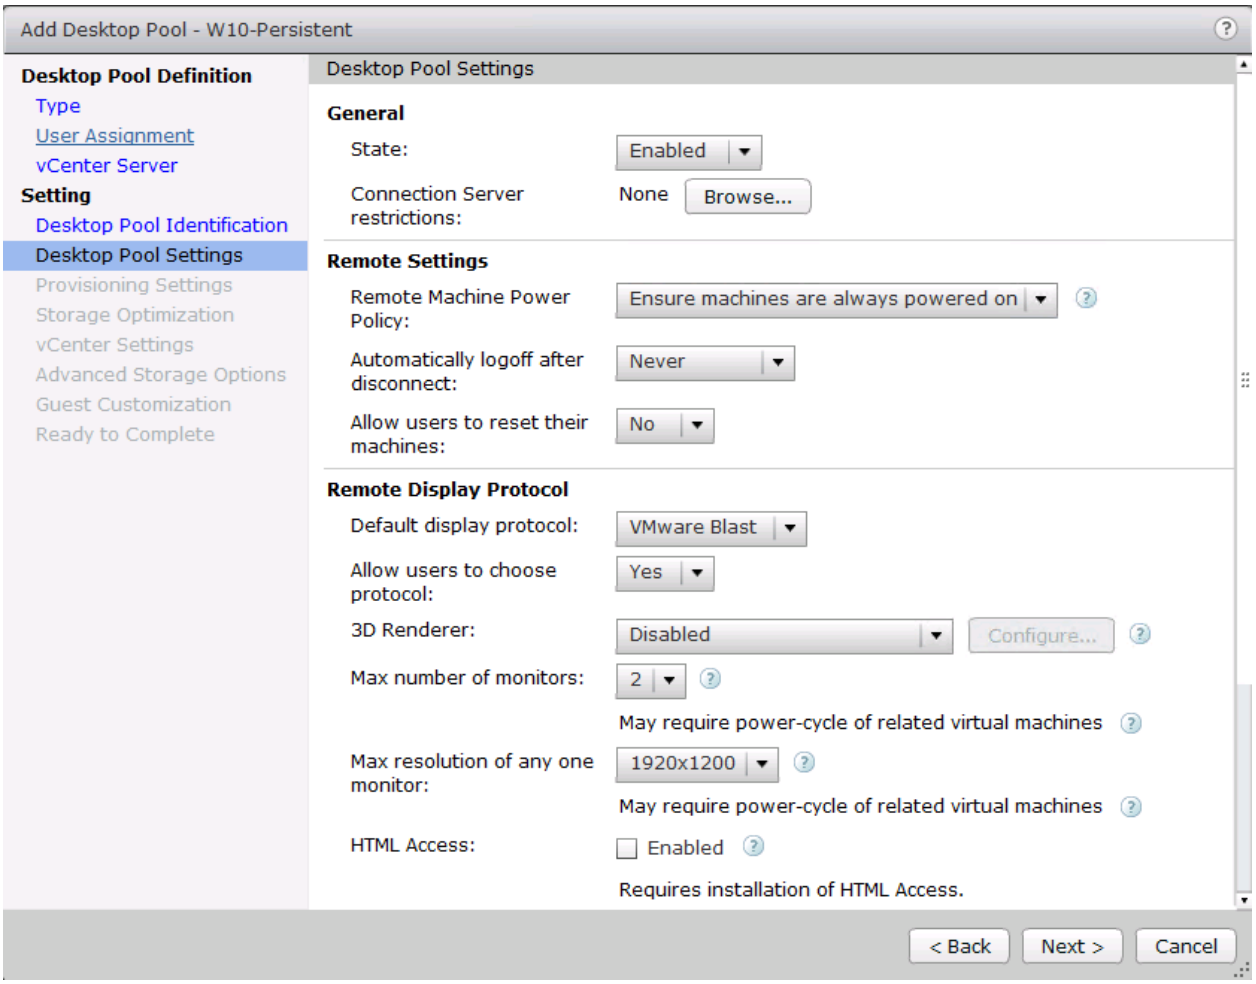

14. Select the provisioning settings to meet your requirements.

15. Click Next.

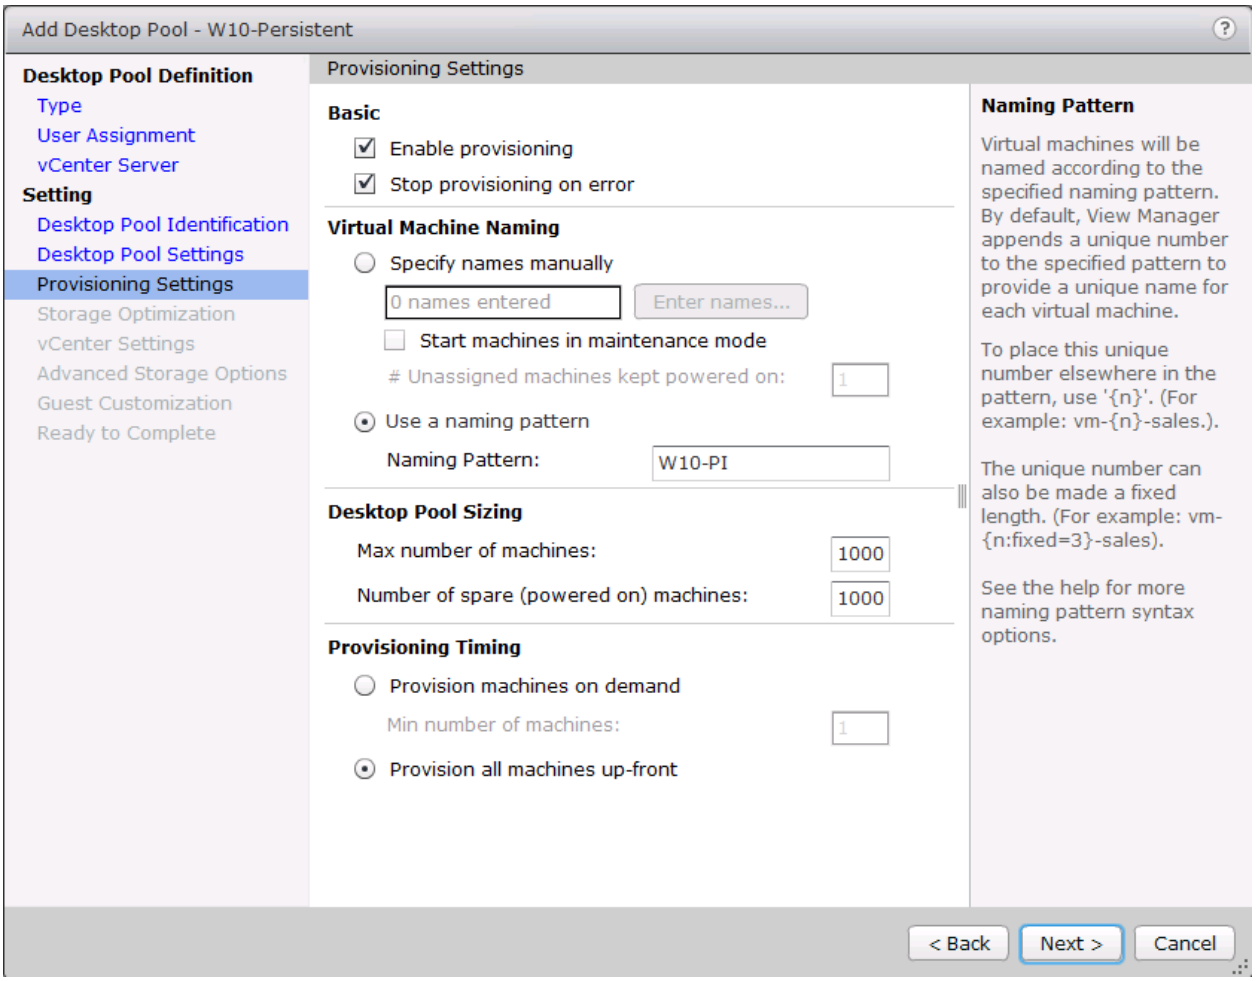

16. Click Next.

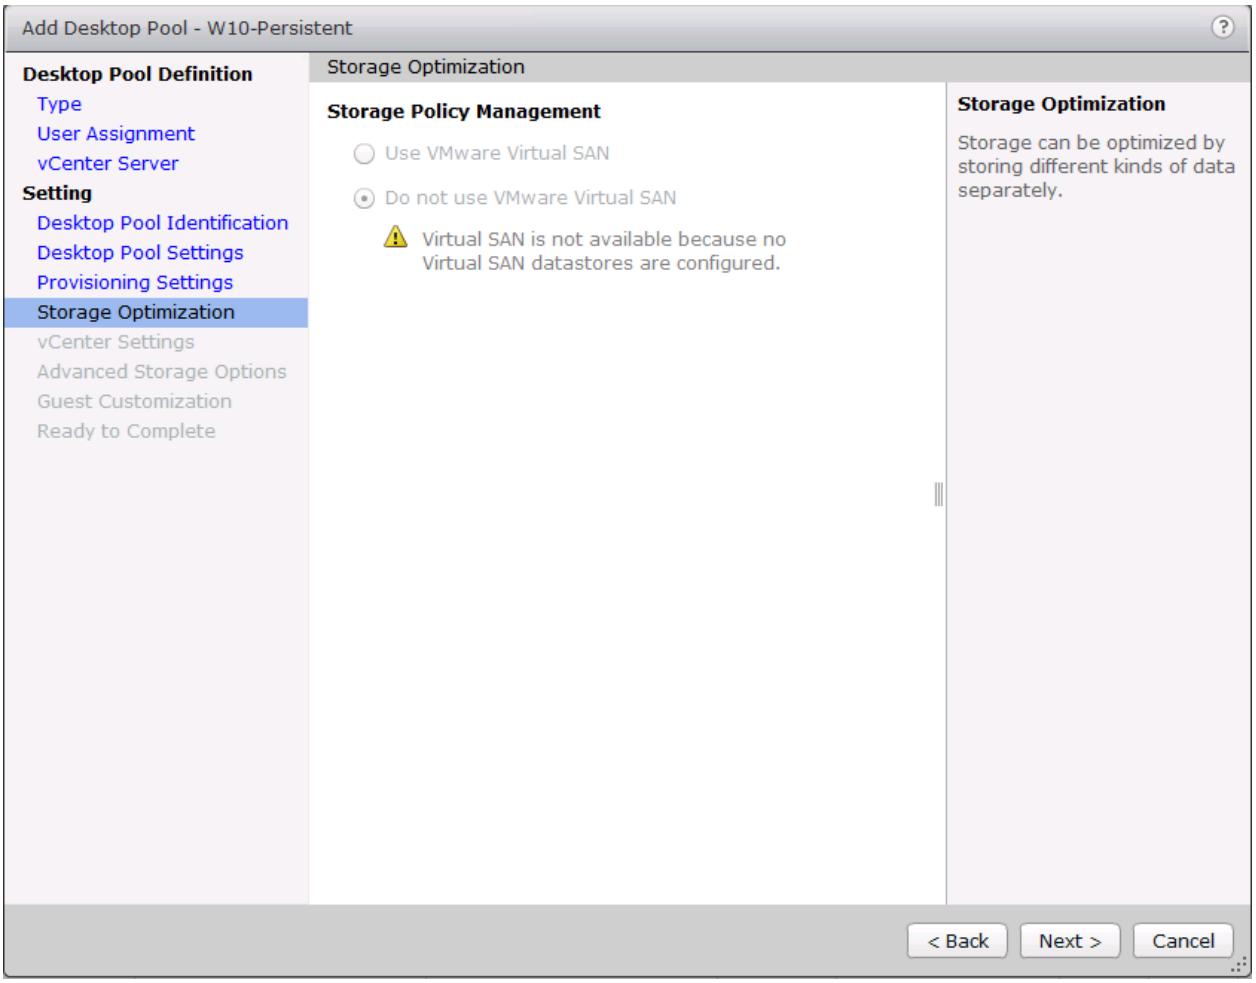

17. Select each of the five vCenter Settings.

18. Click Next.

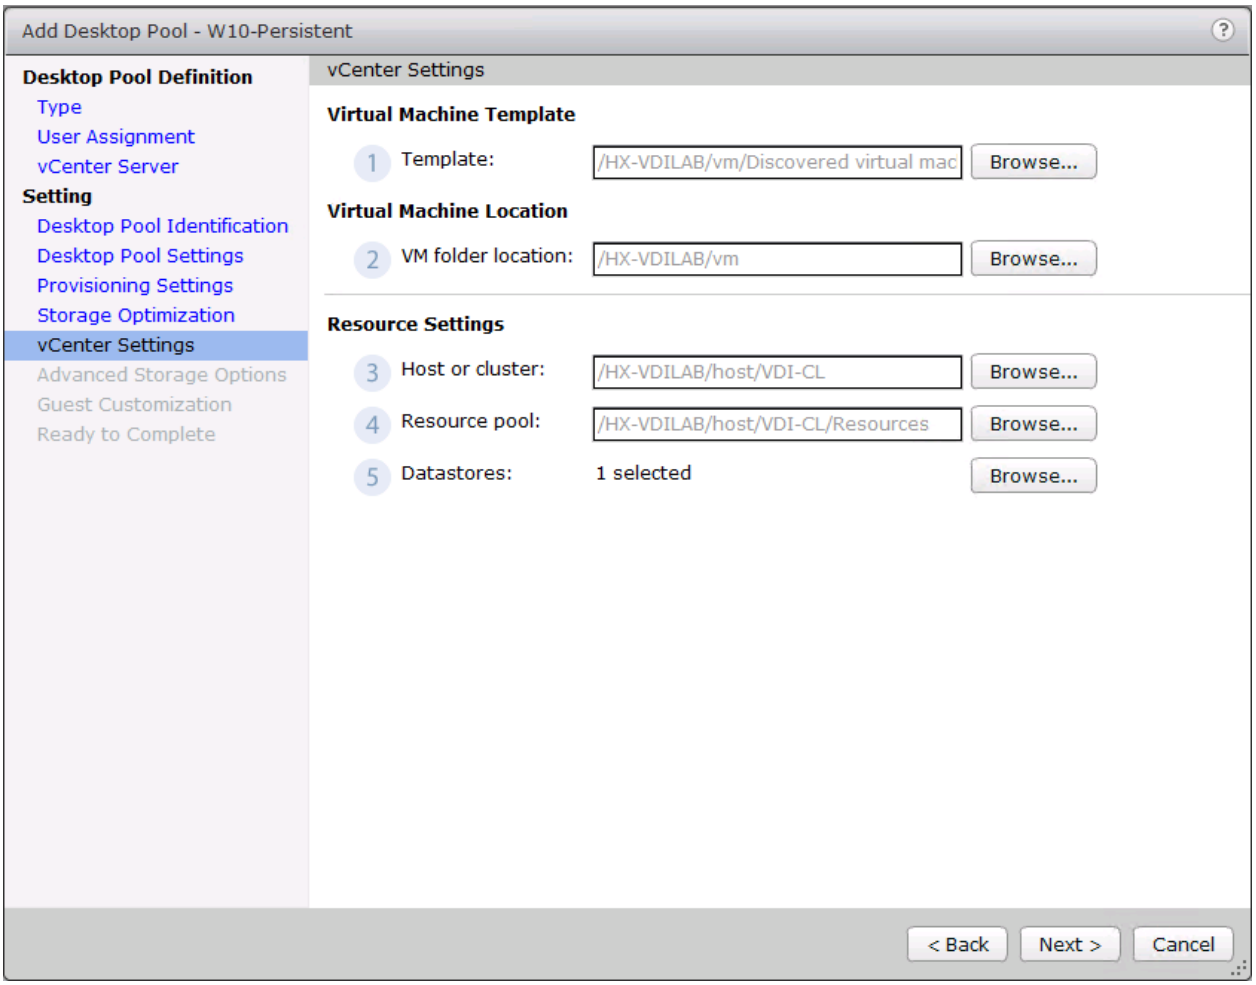

19. For Datastore selection, select the datastore with storage overcommit as “Unbounded.”

20. Click OK.

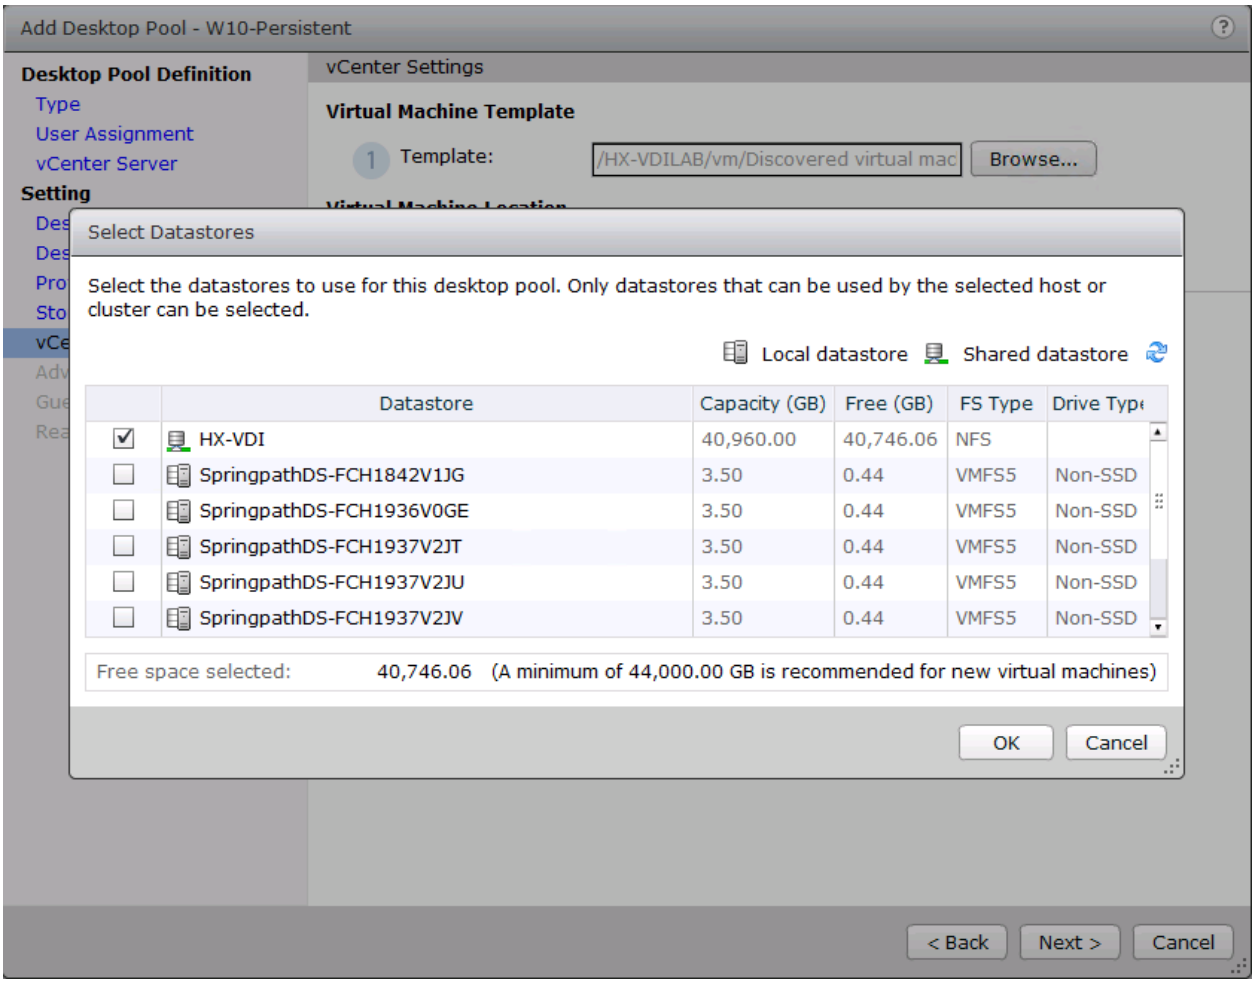

21. Select Advance Storage Options and enable the View Storage Accelerator.

22. Click Next.

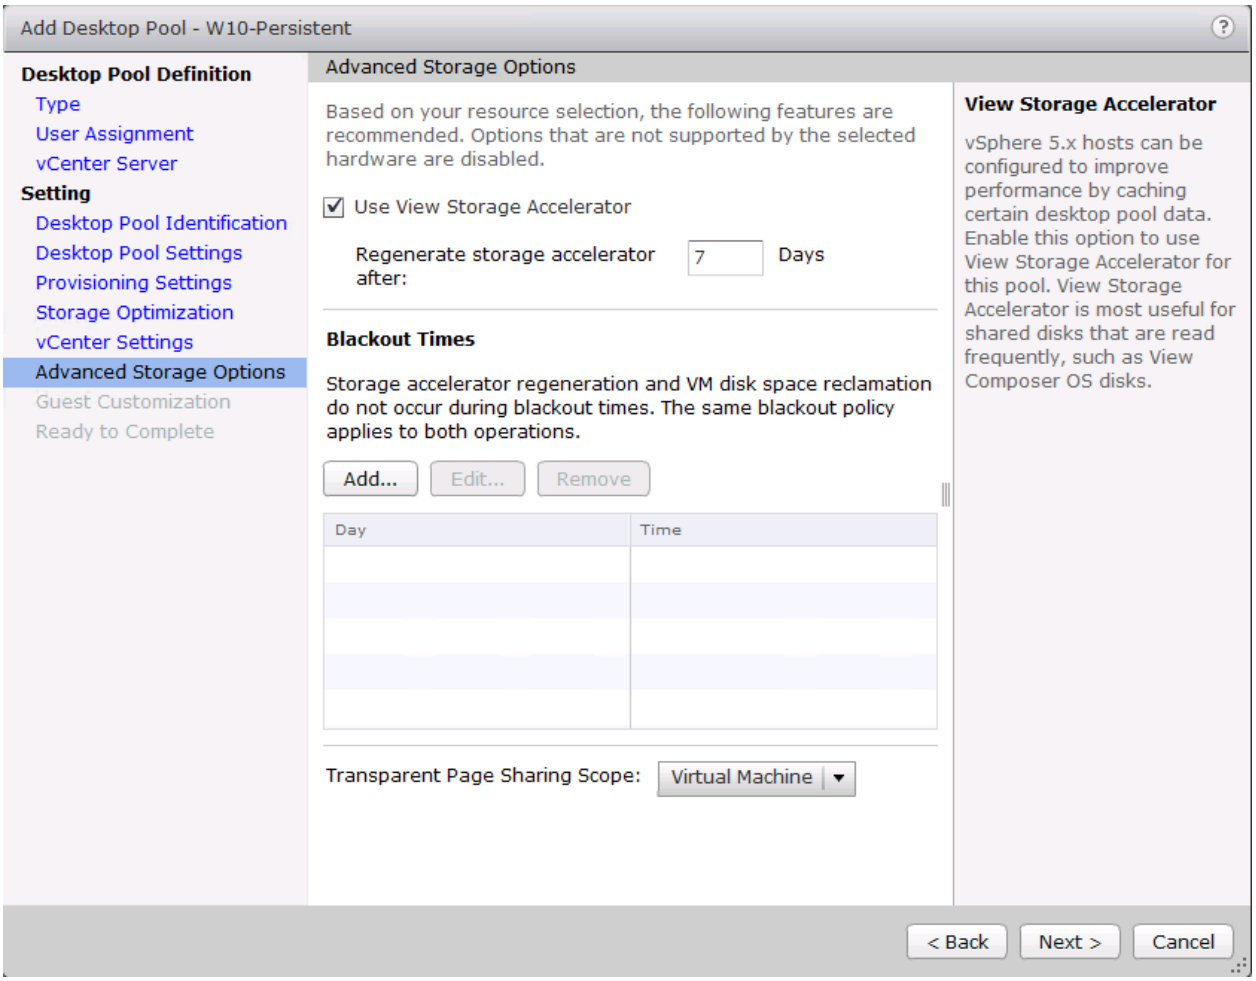

23. Select Guest optimization settings.

24. Click Next.

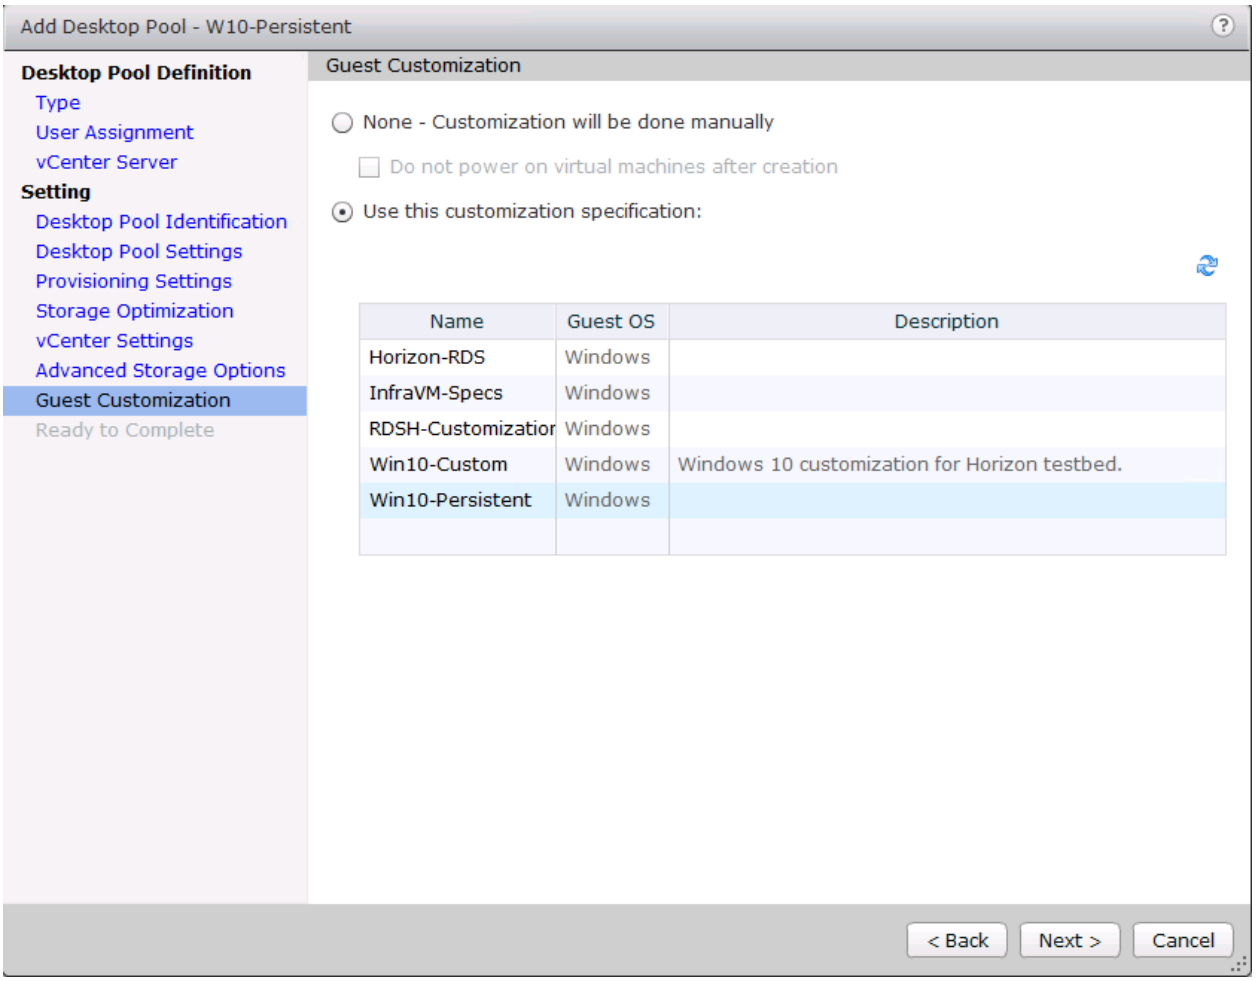

25. Review the summary of the pool you are creating.

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

27. Click Finish.

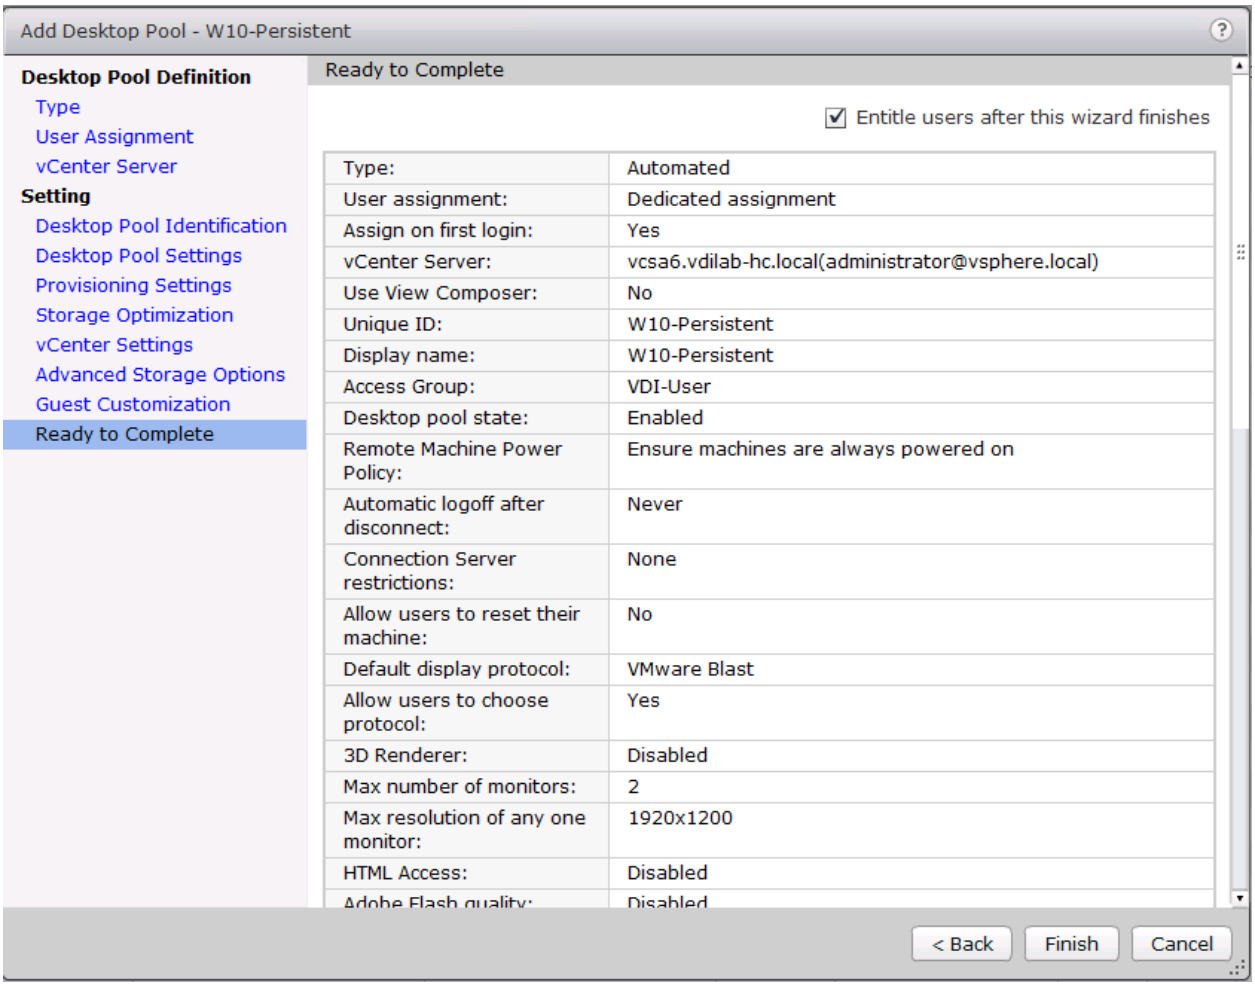

28. Follow the instructions provided in the Create Horizon 7 RDS Desktop Pool to authorize users for the Linked Clone Pool.

Test Setup and Configurations

In this project, we tested a single Cisco HyperFlex cluster running four Cisco UCS HXAF220C-M5SX Rack Servers in a single Cisco UCS domain. This solution is tested to illustrate linear scalability for each workload studied.

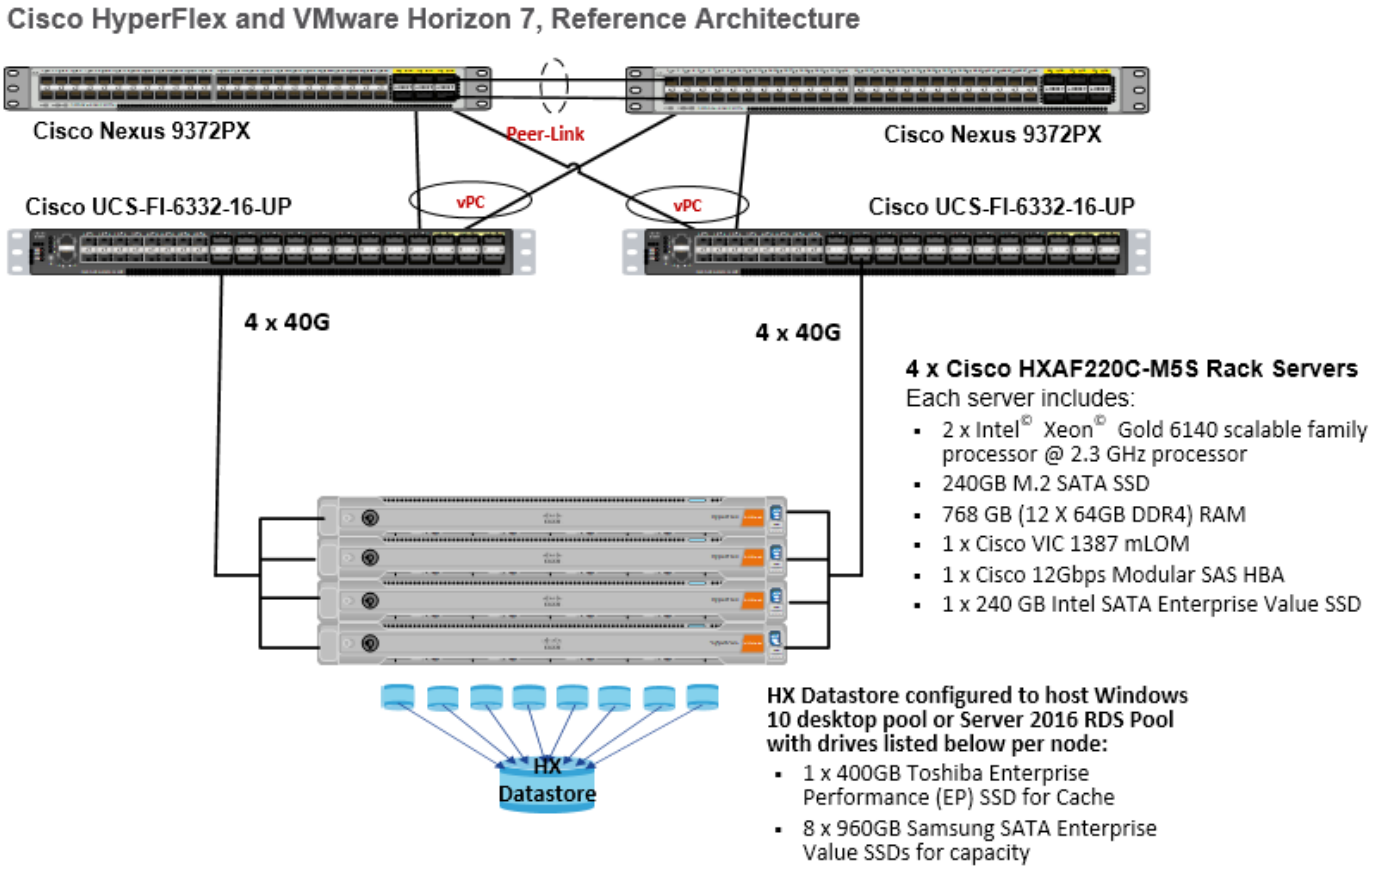

Hardware Components:

- 2 x Cisco UCS 6332-16UP Fabric Interconnects
- 2 x Cisco Nexus 9372PX Access Switches
- 4 x Cisco UCS HXAF220c-M5SX Rack Servers (2 Intel Xeon Gold 6140 scalable family processor at 2.3 GHz, with 768 GB of memory per server [32 GB x 24 DIMMs at 2666 MHz]).
- Cisco VIC 1387 mLOM
- 12G modular SAS HBA Controller
- 240GB M.2 SATA SSD drive (Boot and HyperFlex Data Platform controller VM)
- 240GB 2.5” 6G SATA SSD drive (Housekeeping)
- 400GB 2.5” 6G SAS SSD drive (Cache)

- 8 x 960GB 2.5” SATA SSD drive (Capacity)
- 1 x 32GB mSD card (Upgrades temporary cache)

Software Components:

- Cisco UCS firmware 3.2(2b)
- Cisco HyperFlex Data Platform 2.6.1a
- VMware vSphere 6.5 U1
- 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 4.1.25.6

Testing Methodology and Success Criteria

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

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the 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 were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

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

Testing Procedure

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

Pre-Test Setup for Testing

All machines were shut down utilizing the VMware Horizon 7 Administrator Console.

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

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, we require all sessions started, whether 60 single server users or 600 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 XD.

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 4.1.25 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 4.1.25 to launch Knowledge Worker workload sessions. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The 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% (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% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with 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 amount of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times show a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

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

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

Server-Side Response Time Measurements

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

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

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

Calculating VSImax v4.1.x

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

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

- Notepad File Open (NFO)

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

- Notepad Start Load (NSLD)

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

- Zip High Compression (ZHC)

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

- Zip Low Compression (ZLC)

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

- CPU

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

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

Figure 47 Sample of a VSI Max Response Time Graph, Representing a Normal Test

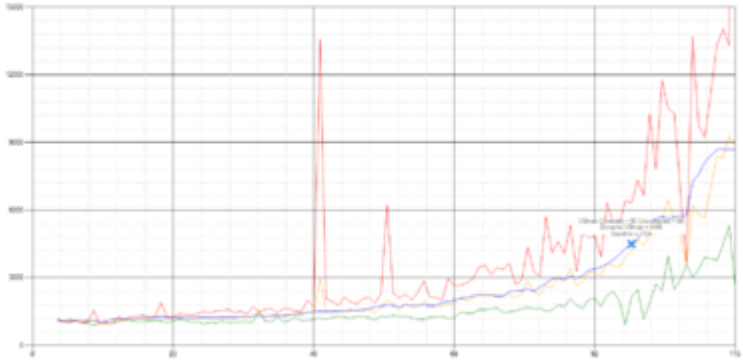

Figure 48 Sample of a VSI Test Response Time Graph with a Clear Performance Issue

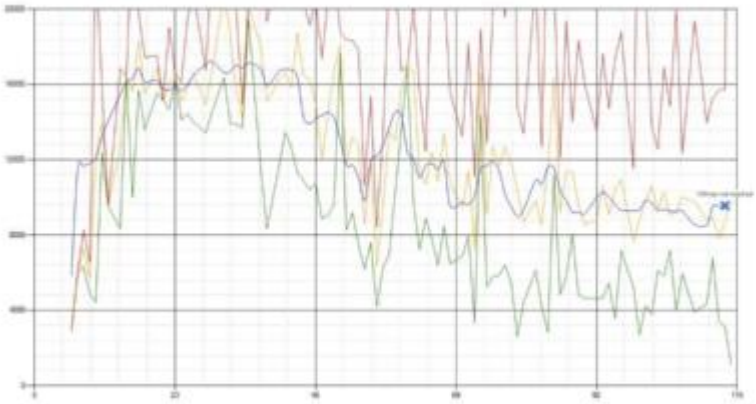

When the test is finished, VSImax can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSImax is not reached and the amount of sessions ran successfully.

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

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

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

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

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

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

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

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

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

To remove noise (accidental spikes) from the calculation, the top 5% and bottom 5% 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% of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

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

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

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

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

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

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

When a server with a very fast dual core CPU, running at 3.6 GHZ, is compared to a 10 core CPU, running at 2,26 GHZ, the dual core machine will give and individual user better performance than the 10 core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

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

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

Test Results

Boot Storms

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

As part of Cisco's virtual desktop test protocol, we shut down each virtual machine at the conclusion of a benchmark test. When we run a new test, we cold boot all 450 desktops and measure the time it takes for the 450th virtual machine to register as available in the Horizon Administrator console.

The Cisco HyperFlex HXAF220c-M5SX based All-Flash cluster running Data Platform version 2.6(1a) software can accomplish this task in 5 minutes as shown in the following charts:

Figure 49 450 Horizon 7 Linked-Clone Windows 10 Sessions with Office 2016 Virtual Desktops Boot and Register as Available in Less Than 5 Minutes

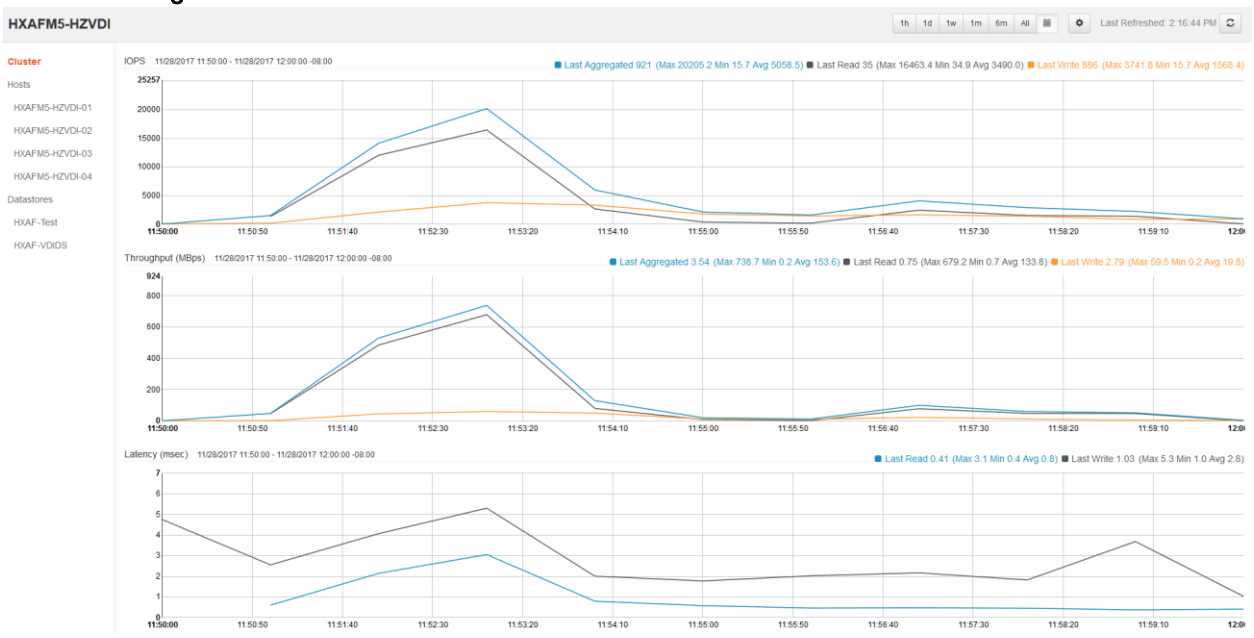

Figure 50 450 Horizon 7 Persistent(Full Clone) Windows 10 Sessions with Office 2016 Virtual Desktops Boot and Register as Available in Less Than 5 Minutes

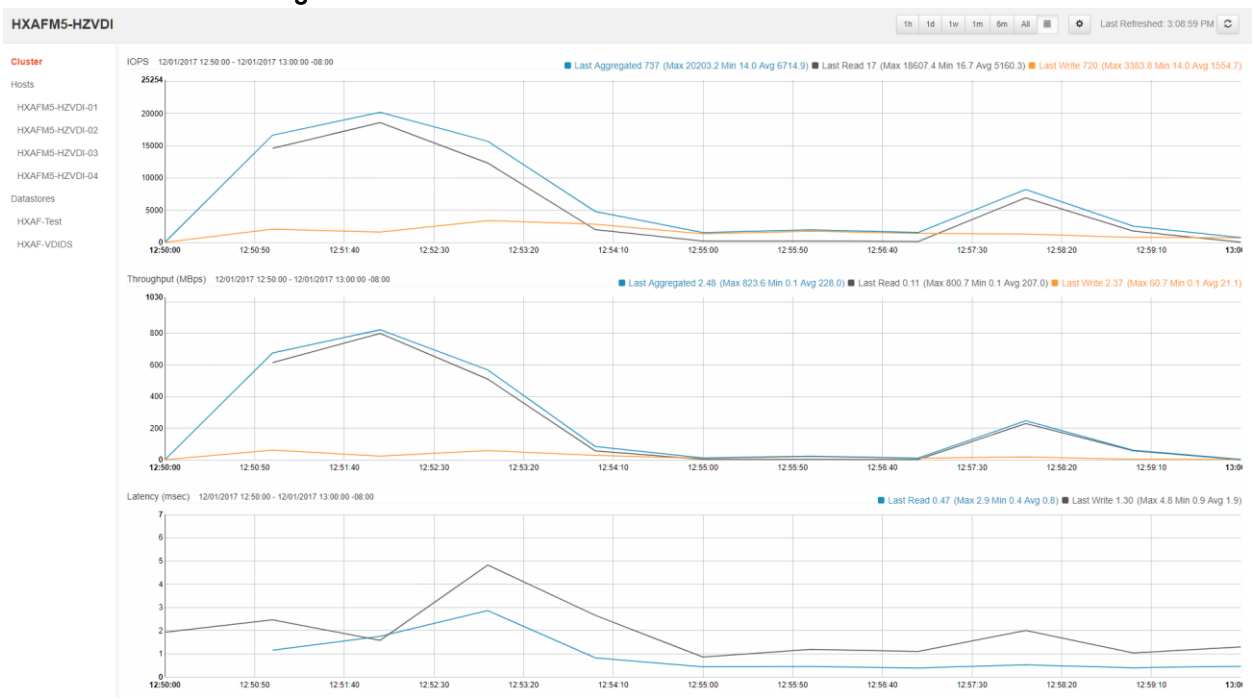

Recommended Maximum Workload and Configuration Guidelines

Four Node Cisco HXAF220c-M5S Rack Server, HyperFlex All-Flash 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-M5SX server operating parameters.

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

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

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

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

RDSH Server Pool Testing on Four Node Cisco HyperFlex Cluster

This section shows the key performance metrics that were captured on the Cisco UCS HyperFlex storage cluster configured with four HXAF220c-M5SX converged node running RDSH VMs. The full-scale testing with 600 user session on 36 Windows Server 2016 RDSH VMs on four HXAF220c-M5SX HyperFlex cluster.

Test result highlights include:

- 0.610 second baseline response time
- 0.832 second average response time with 4000 desktop sessions running
- Average CPU utilization of 70 percent during steady state
- Average of 250 GB of RAM used out of 768 GB available
- 3000Mbps peak network utilization per host.
- Average Read Latency 0.5ms/Max Read Latency 1.8ms
- Average Write Latency 4.5ms/Max Write Latency 8.7ms
- 2800 peak I/O operations per second (IOPS) per cluster at steady state
- 125MBps peak throughput per cluster at steady state

Figure 51 LoginVSI Analyzer Chart for 600 Users on RDSH Server Desktop Test

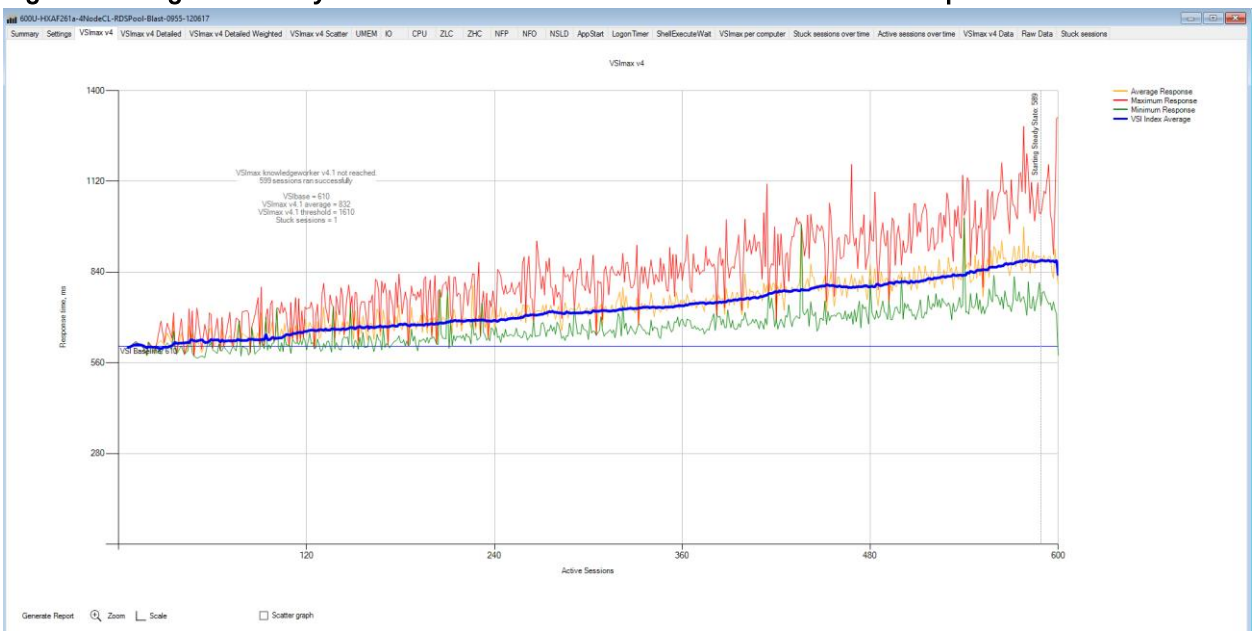

Figure 52 LoginVSI Analyzer Chart for Three Consecutive Test Running 600 Knowledge Worker Workload on Four Node HyperFlex Cluster

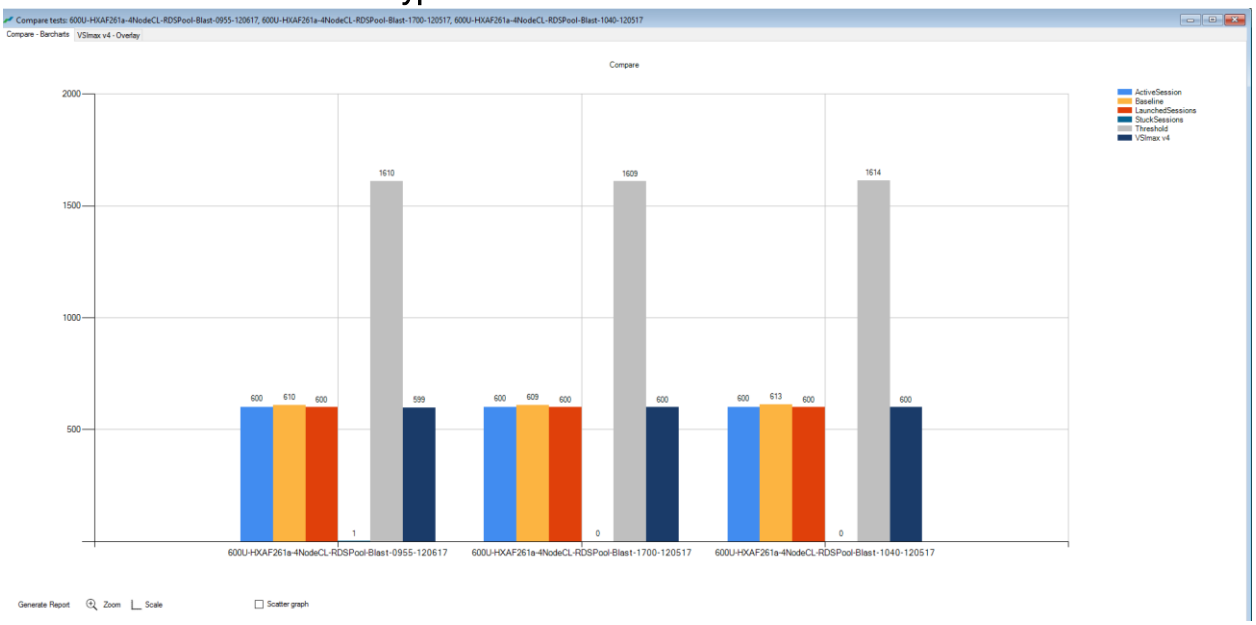

Figure 53 Sample ESXi Host CPU Core Utilization Running 600 User Test with 36 RDSH Server VMs on Four Nodes

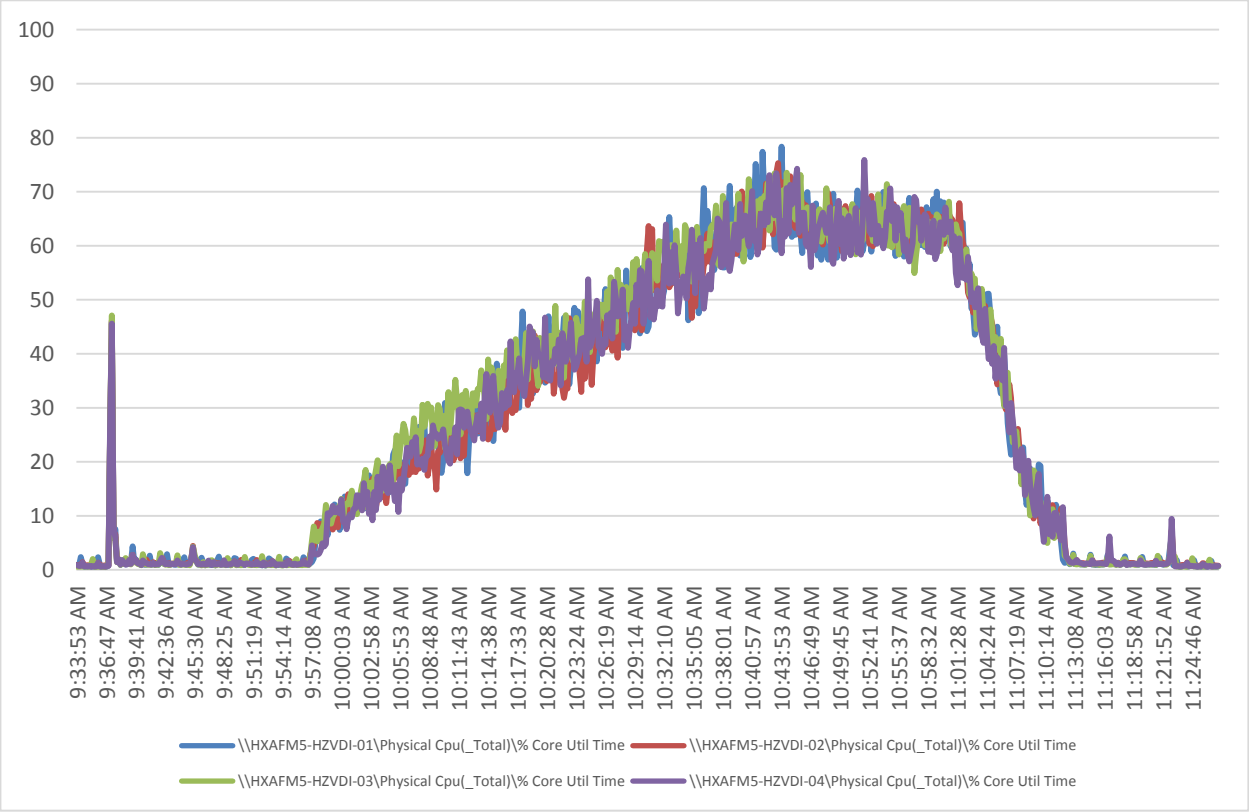

Figure 54 Sample ESXi Host Memory Usage in Mbytes Running 600 User Test with 36 RDSH Server VMs on Four Nodes

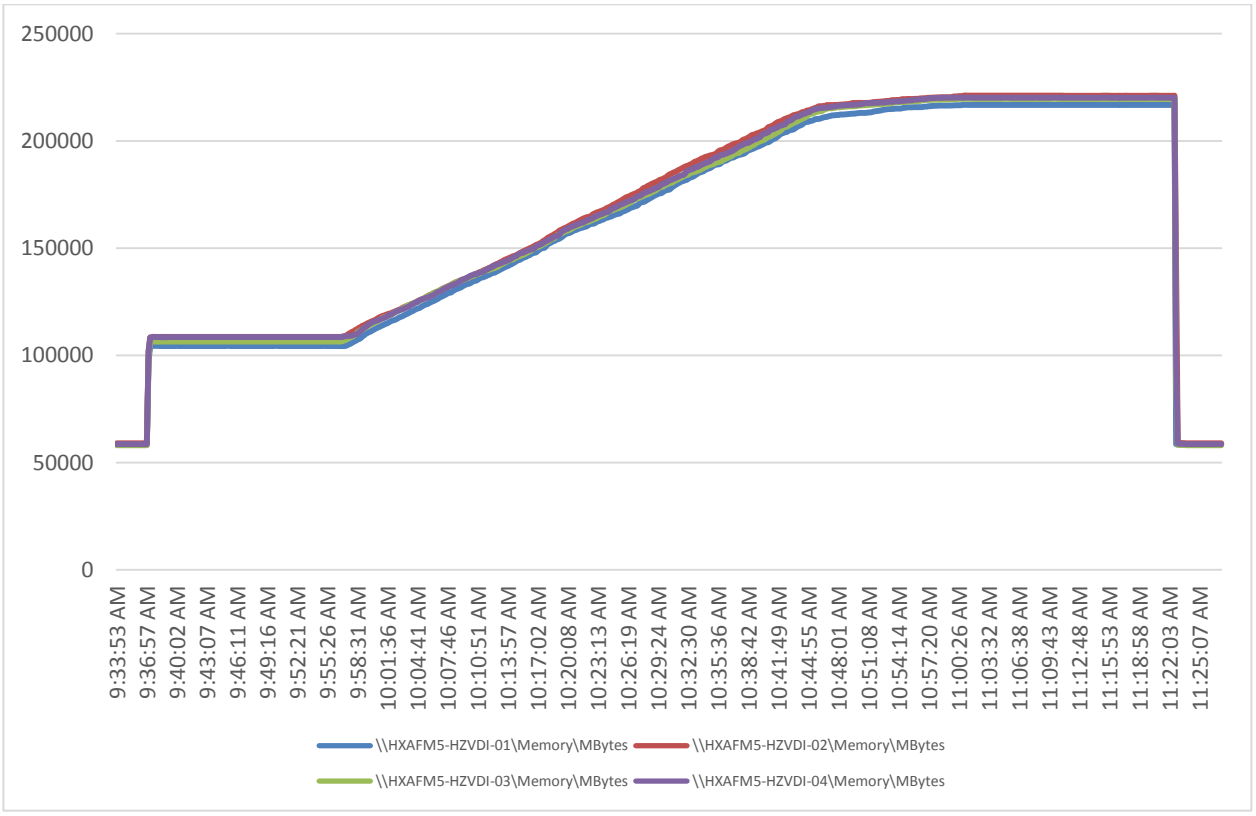

Figure 55 Sample ESXi Host Network Adapter (VMNICs) Mbits Received/ Transmitted Per Sec Running 600 User Test with 36 RDSH Server VMs on Four Nodes

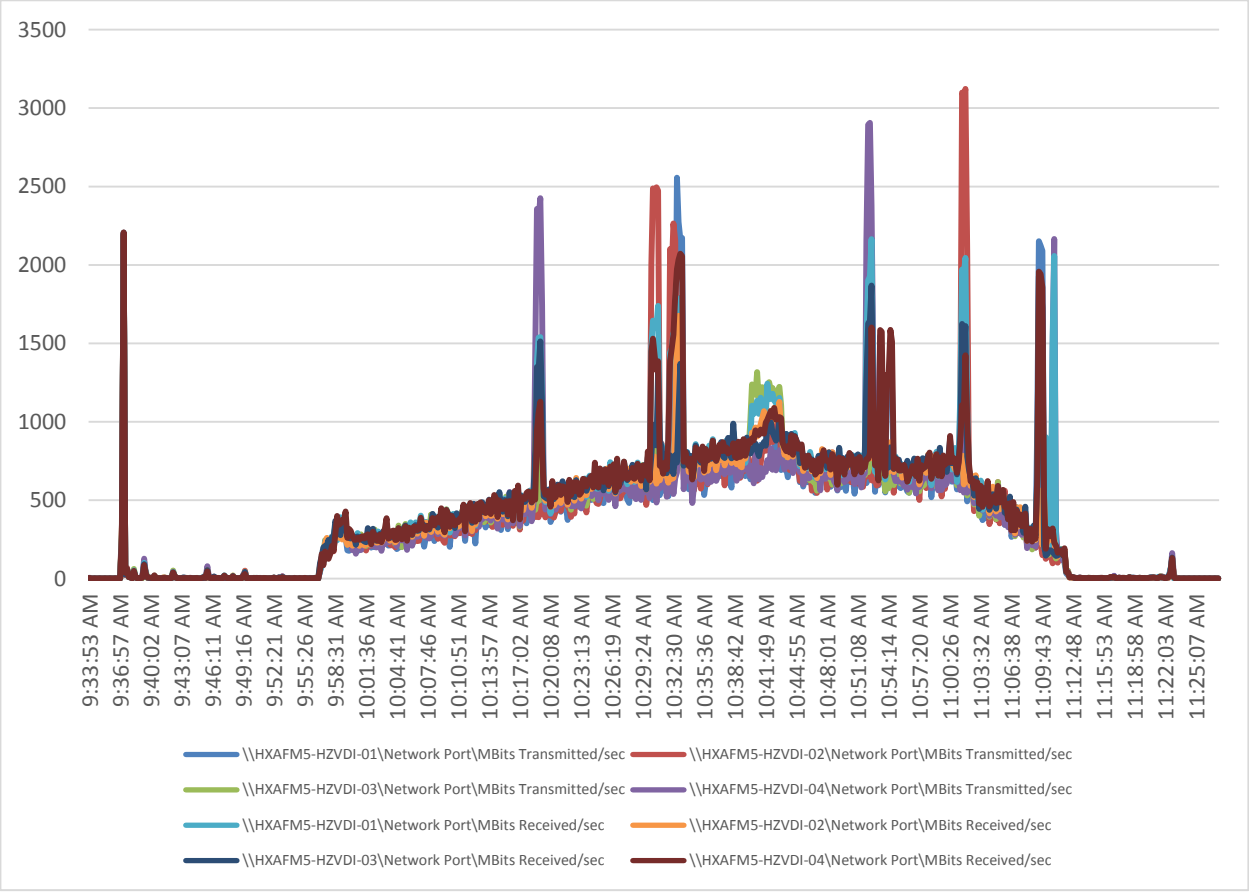

Figure 56 HyperFlex Cluster WebUI Performance Chart for Knowledge Worker Workload Running 600 User Test with 36 RDSH Server VMs on Four Nodes

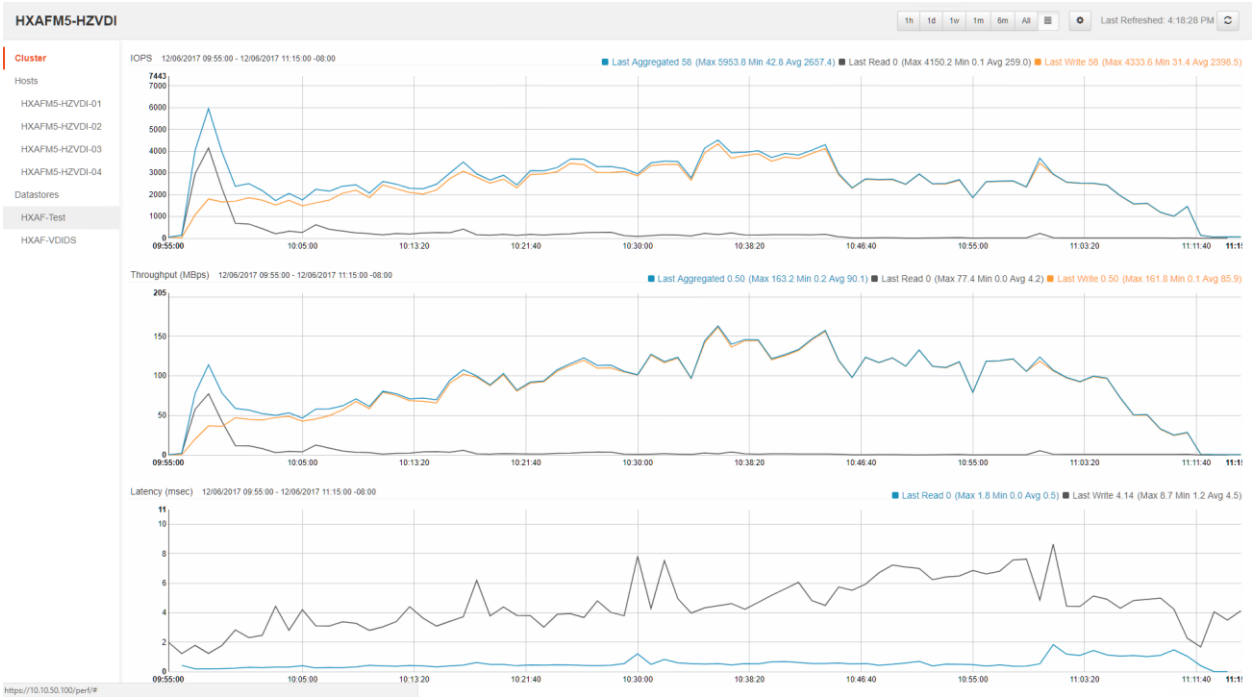

Figure 57 vCenter WebUI Reporting HyperFlex Cluster De-duplication and Compression Savings for 600 User Sessions Supported on Windows Server 2016 Based Hosted Shared Sessions Deployed on 32 Node HyperFlex Cluster

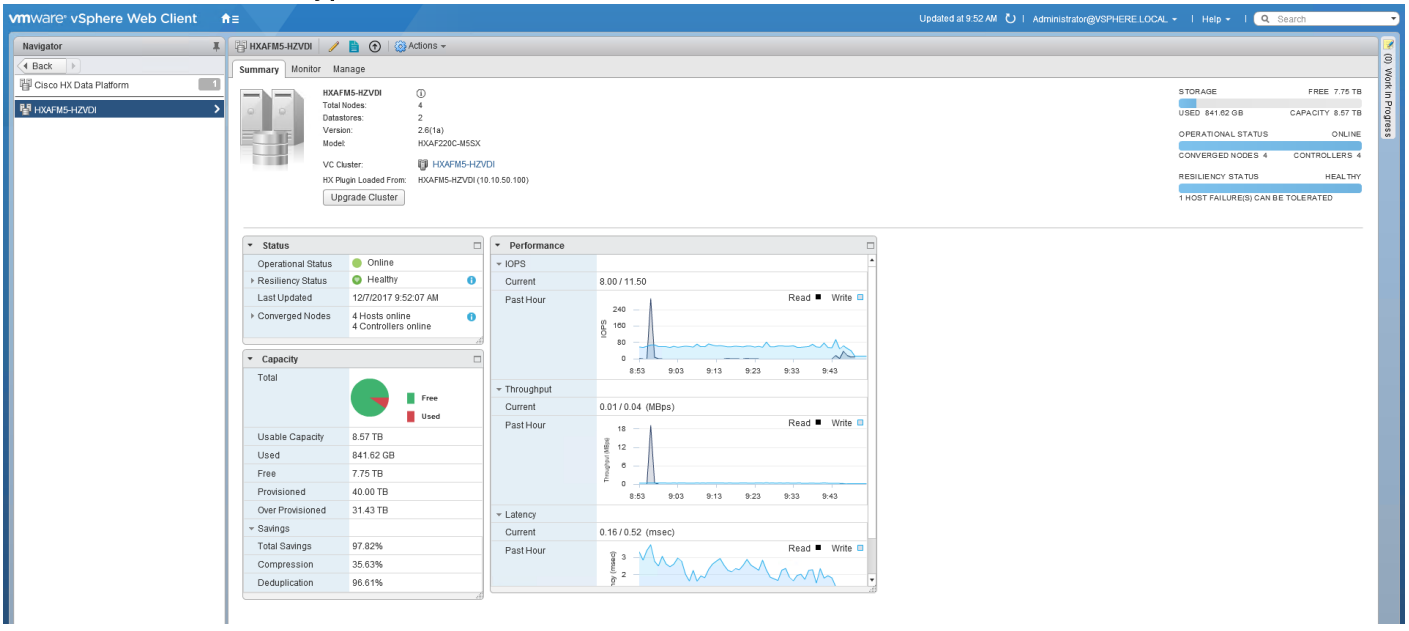

Figure 58 Horizon Administrator Console Reporting RDSH Pool Properties and 600 Active Sessions

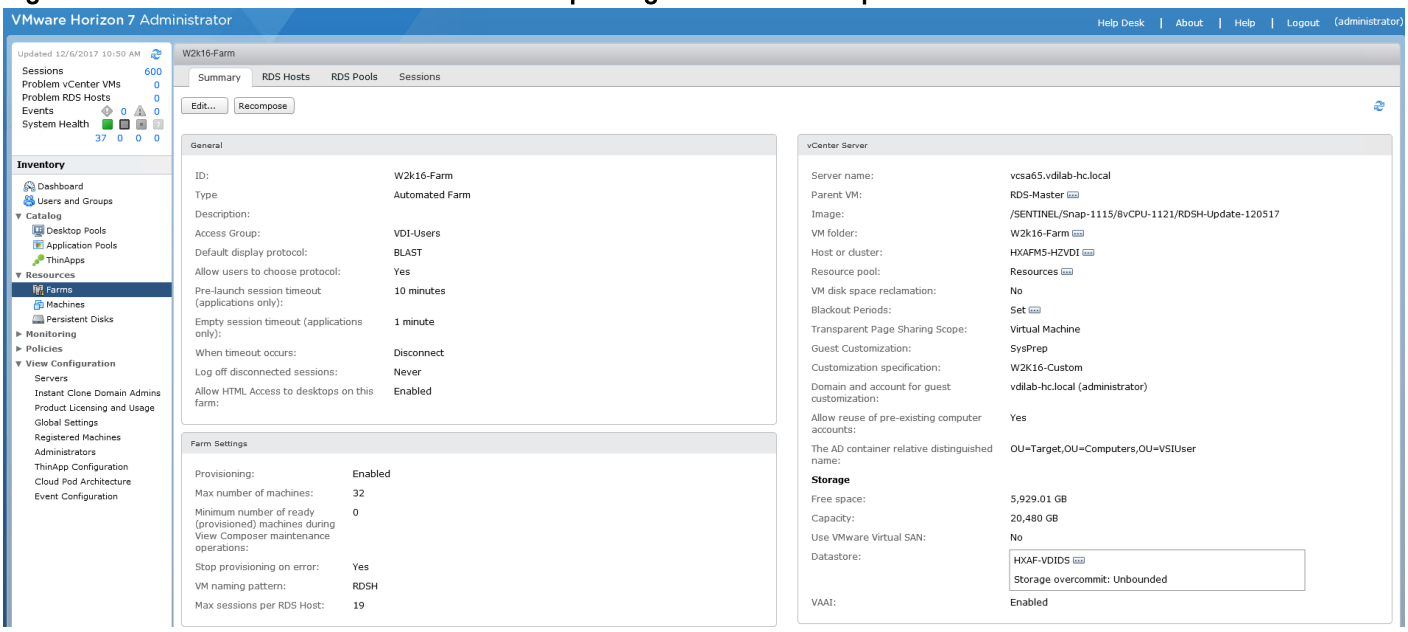

450 Windows 10 Instant-clone Desktop Pool Testing on Four Node Cisco HyperFlex Cluster

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

Test result highlights include:

- 0.688 second baseline response time
- 0.996 second average response time with 2000 desktops running
- Average CPU utilization of 80 percent during steady state
- Average of 342 GB of RAM used out of 768 GB available
- 1500Mbps peak network utilization per host.
- Average Read Latency 0.4ms/Max Read Latency 0.7ms
- Average Write Latency 2.0ms/Max Write Latency 5.0ms
- 6000 peak I/O operations per second (IOPS) per cluster at steady state
- 130MBps peak throughput per cluster at steady state

Figure 59 Login VSI Analyzer Chart for 450 Windows 10 Instant-Clone Virtual Desktops

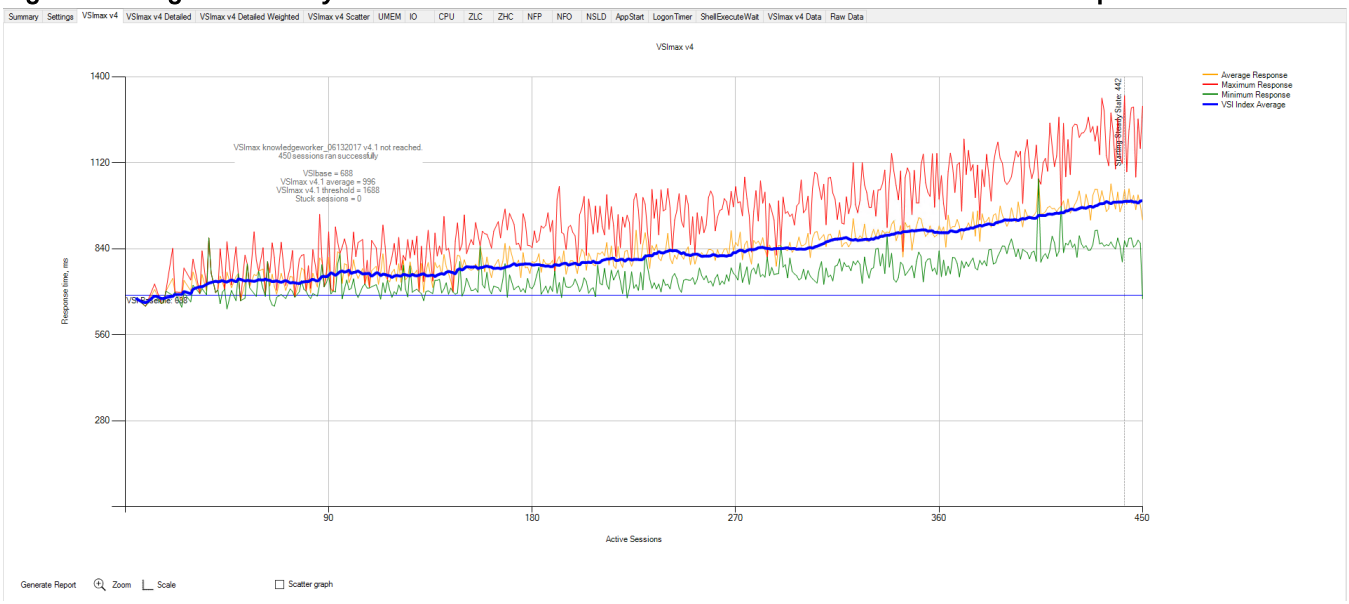

Figure 60 Three Consecutive Login VSI Analyzer Chart for 450 Windows 10 Instant-Clone Virtual Desktops

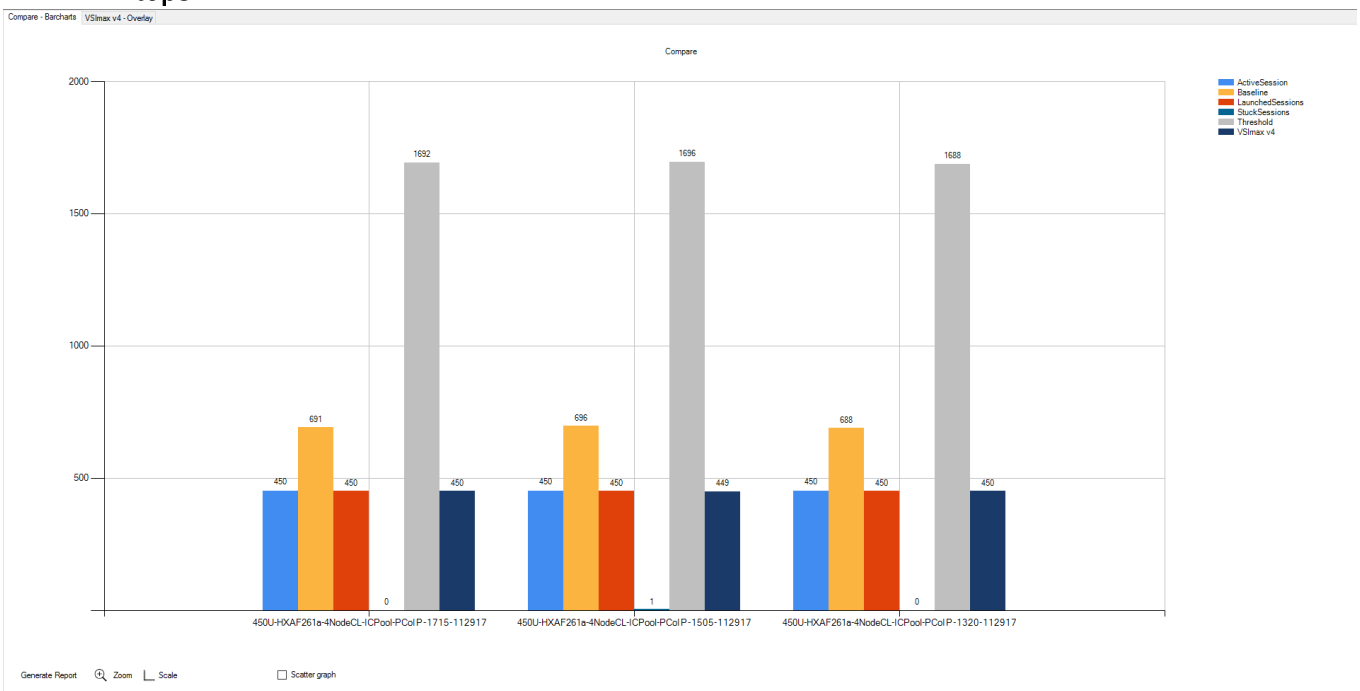

Figure 61 Sample ESXi Host CPU Core Utilization Running 450 Windows 10 Instant-Clone Virtual Desktops

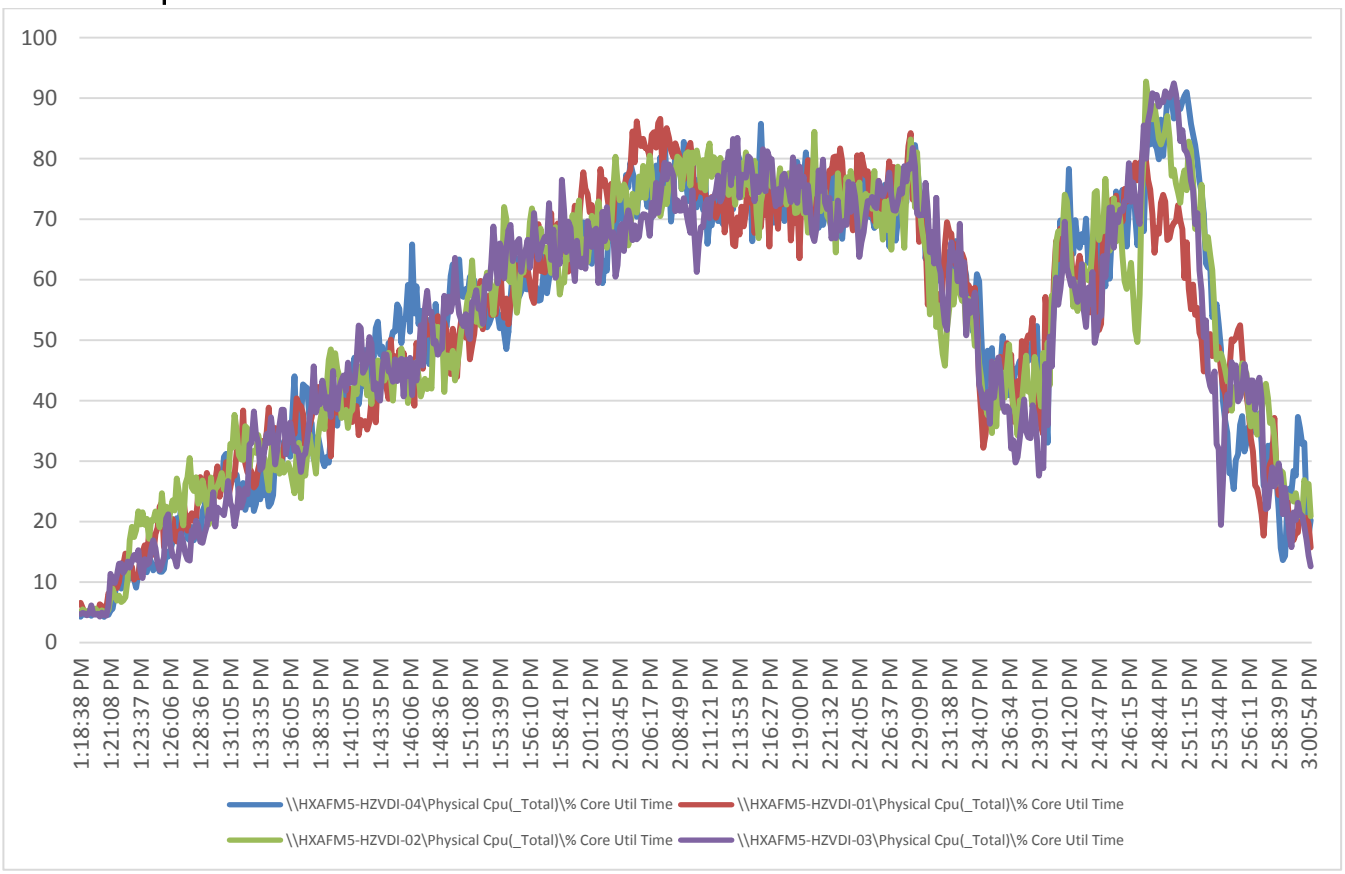

Figure 62 ESXi Host Memory Usage in Mbytes Running 450 Windows 10 Instant-Clone Virtual Desktops

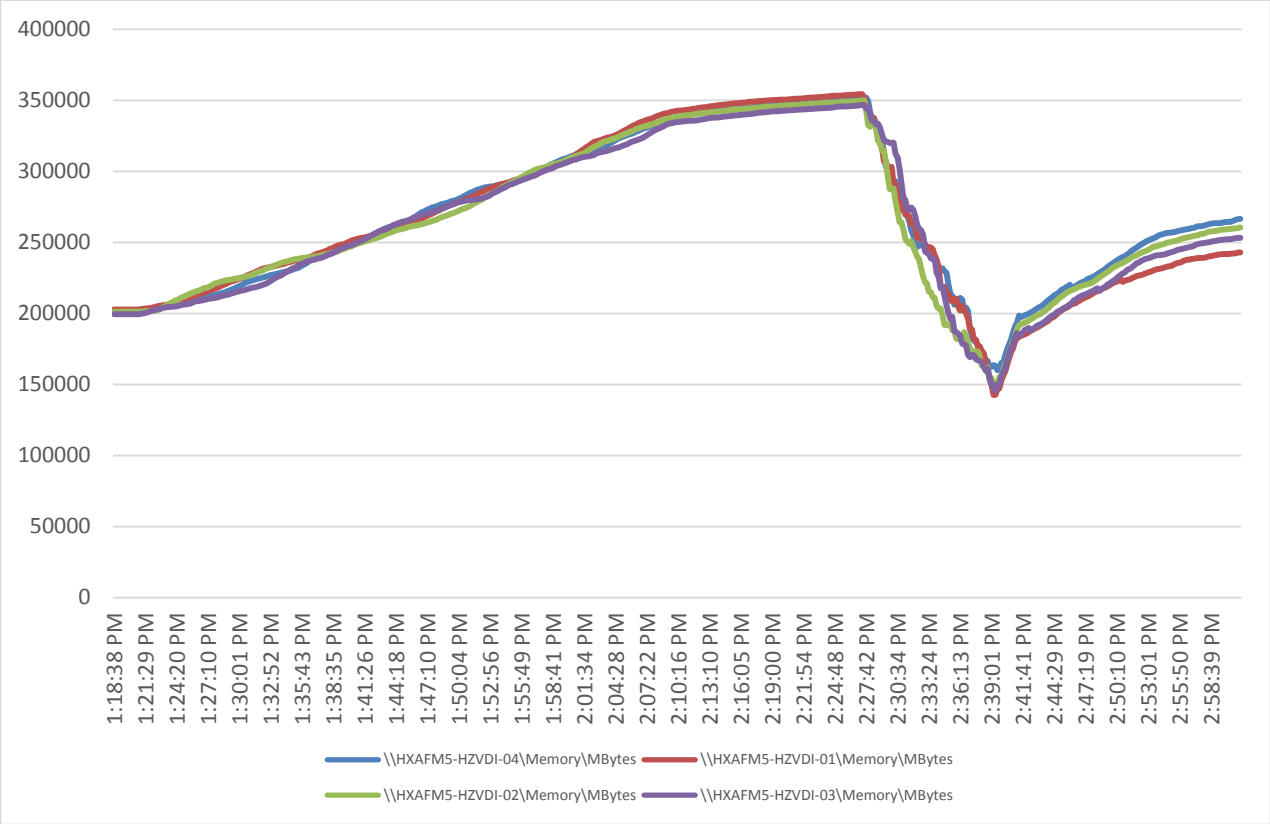

Figure 63 ESXi Host Network Adapter (VMNICs) Mbits Received/Transmitted Per Sec Running 450 Windows 10 Instant-Clone Virtual Desktops

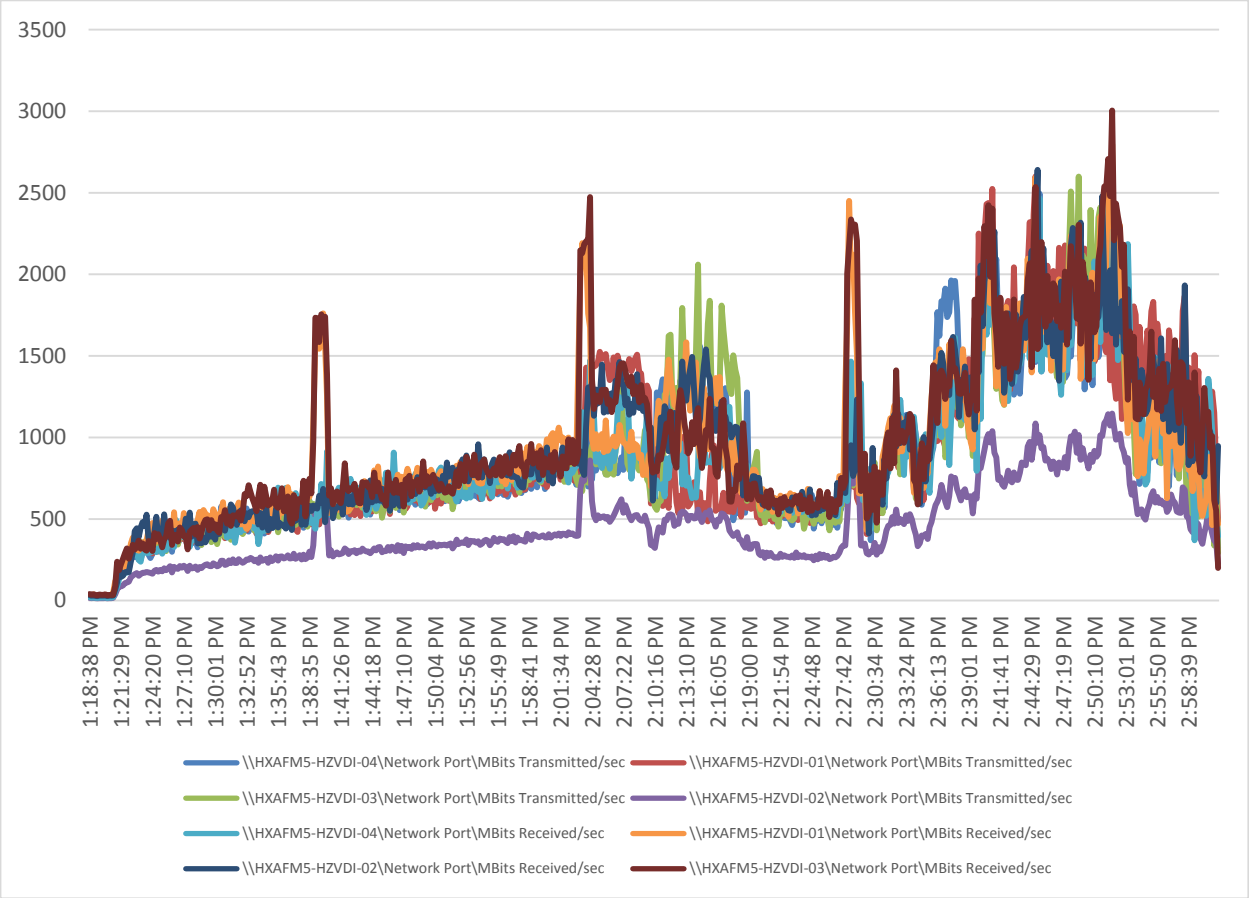

Figure 64 HyperFlex Cluster Performance Chart for Knowledge Worker Workload Running 450 User Test on Instant-Clone Windows 10

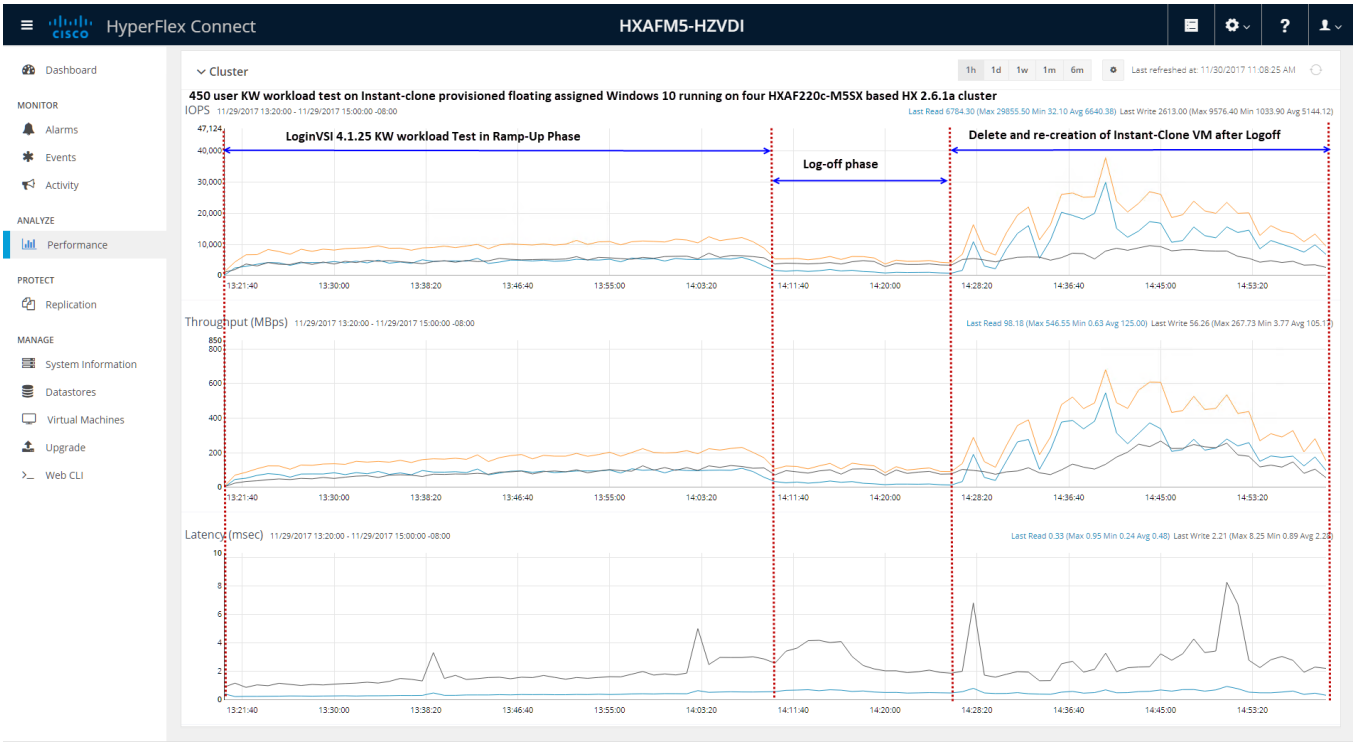

Figure 65 vCenter WebUI Reporting HyperFlex Cluster Deduplication and Compression Savings for 450 Instant-Clone VMs Running Windows 10/Office 2016 Supporting 450 Users

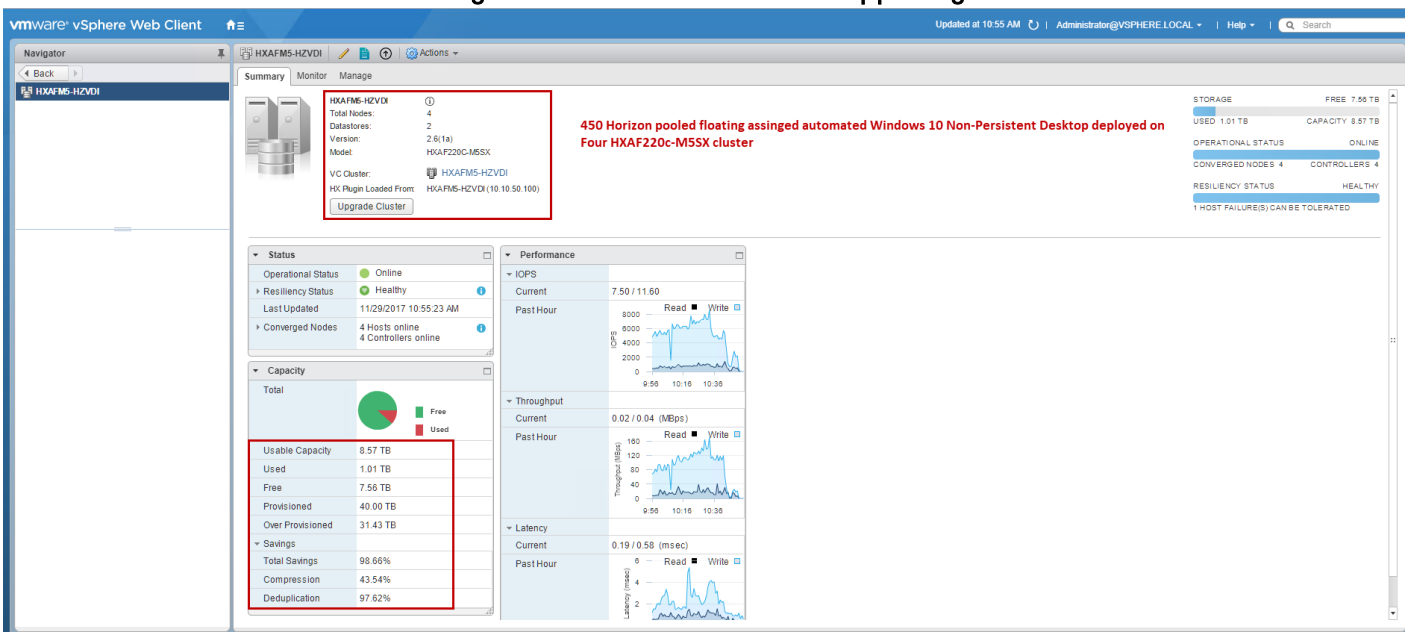

Figure 66 Horizon Administrator Console Reporting Instant Clone Desktop Pool Configuration and 450 Active Sessions

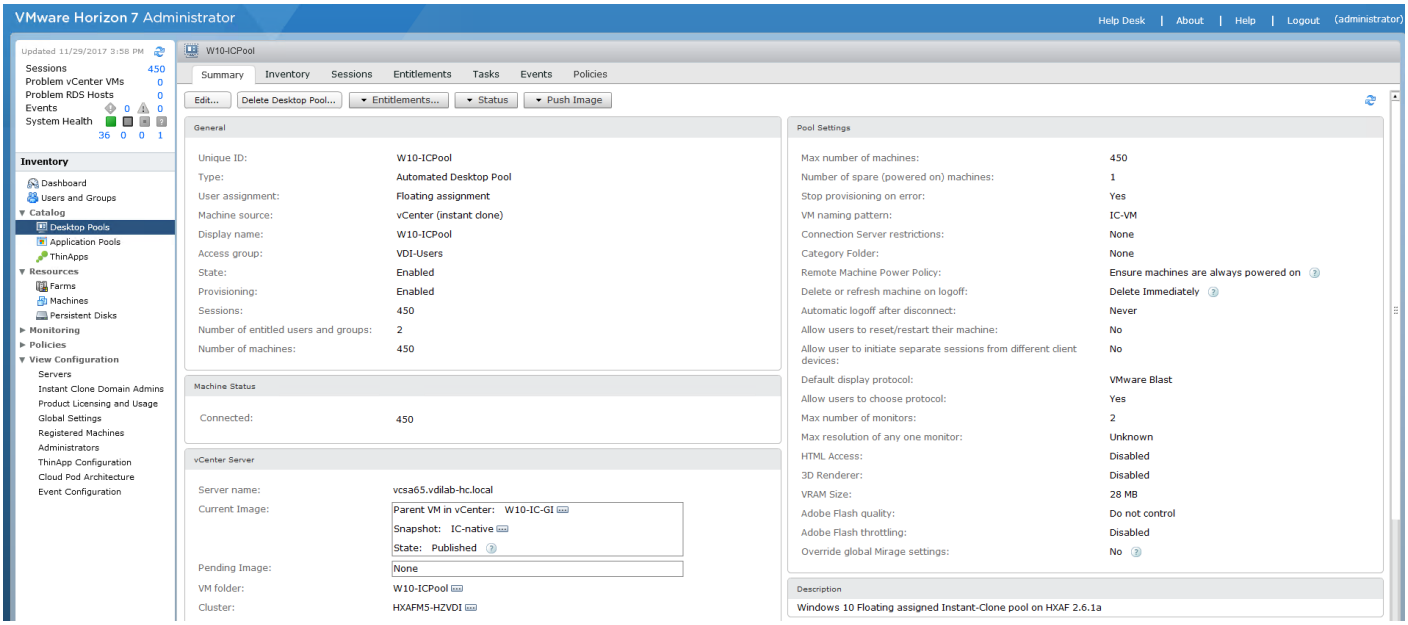

450 Windows 10 Linked-Clone Desktop Pool Testing on Four Node Cisco HyperFlex Cluster

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

Test result highlights include:

- 0.691 second baseline response time
- 0.941 second average response time with 2000 desktops running
- Average CPU utilization of 80 percent during steady state
- Average of 320 GB of RAM used out of 768 GB available
- 1000Mbps peak network utilization per host.
- Average Read Latency 0.7ms/Max Read Latency 1.4ms
- Average Write Latency 1.9ms/Max Write Latency 4.4ms
- 3500 peak I/O operations per second (IOPS) per cluster at steady state
- 80MBps peak throughput per cluster at steady state

Figure 67 Login VSI Analyzer Chart for 450 Windows 10 Linked-Clone Virtual Desktops

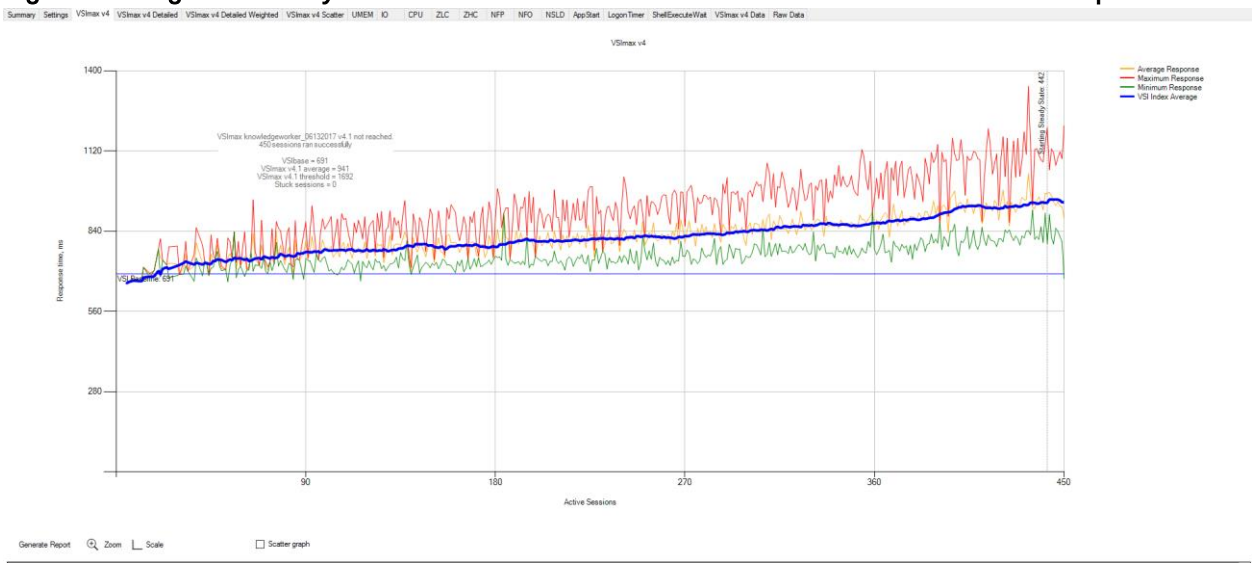

Figure 68 Three Consecutive Login VSI Analyzer Chart for 450 Windows 10 Linked-Clone Virtual Desktops

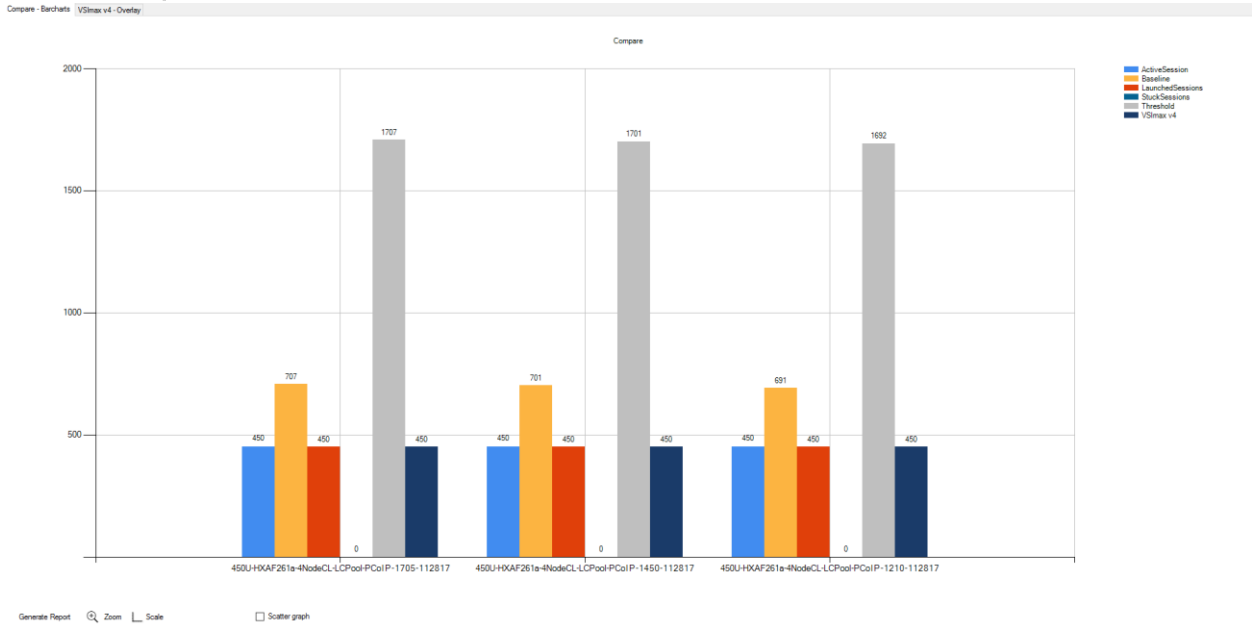

Figure 69 Sample ESXi host CPU Core Utilization Running 450 Windows 10 Linked-Clone Virtual Desktops

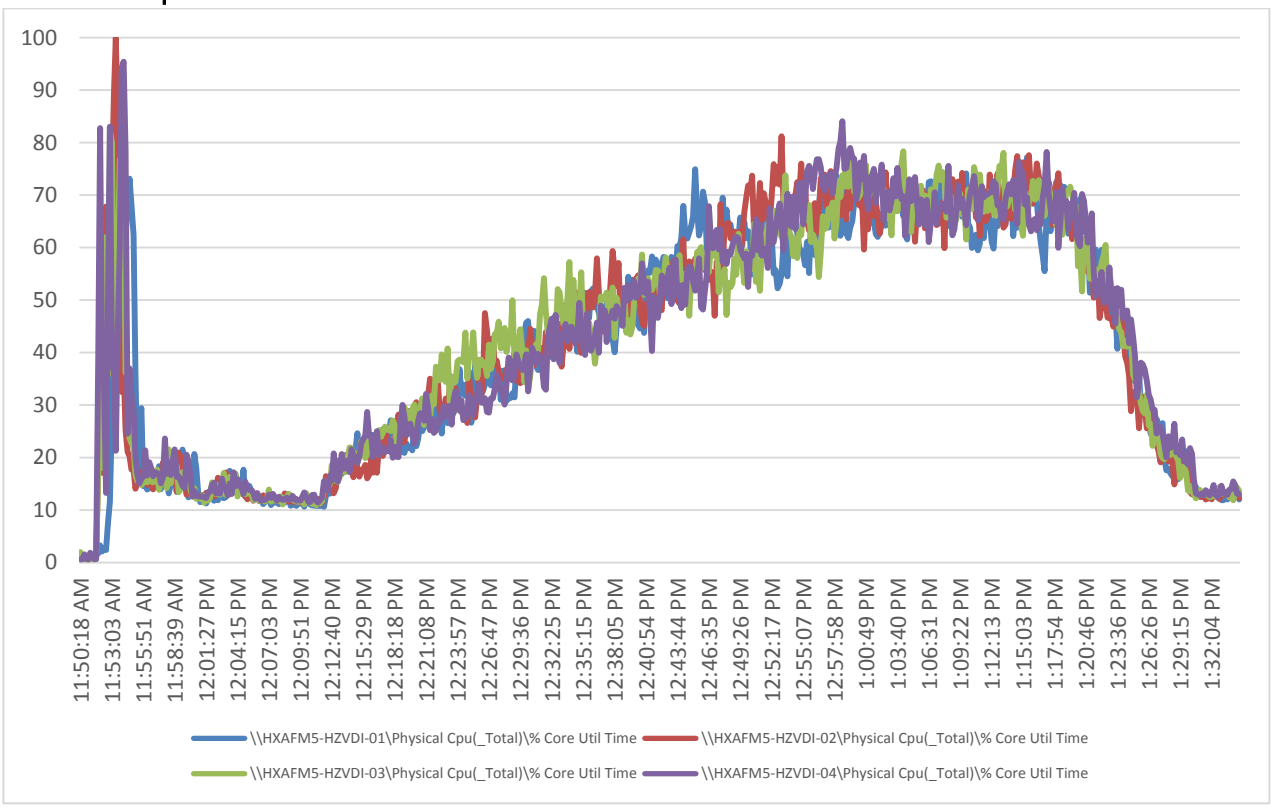

Figure 70 ESXi Host Memory Usage in Mbytes Running 450 Windows 10 Linked-Clone Virtual Desktops

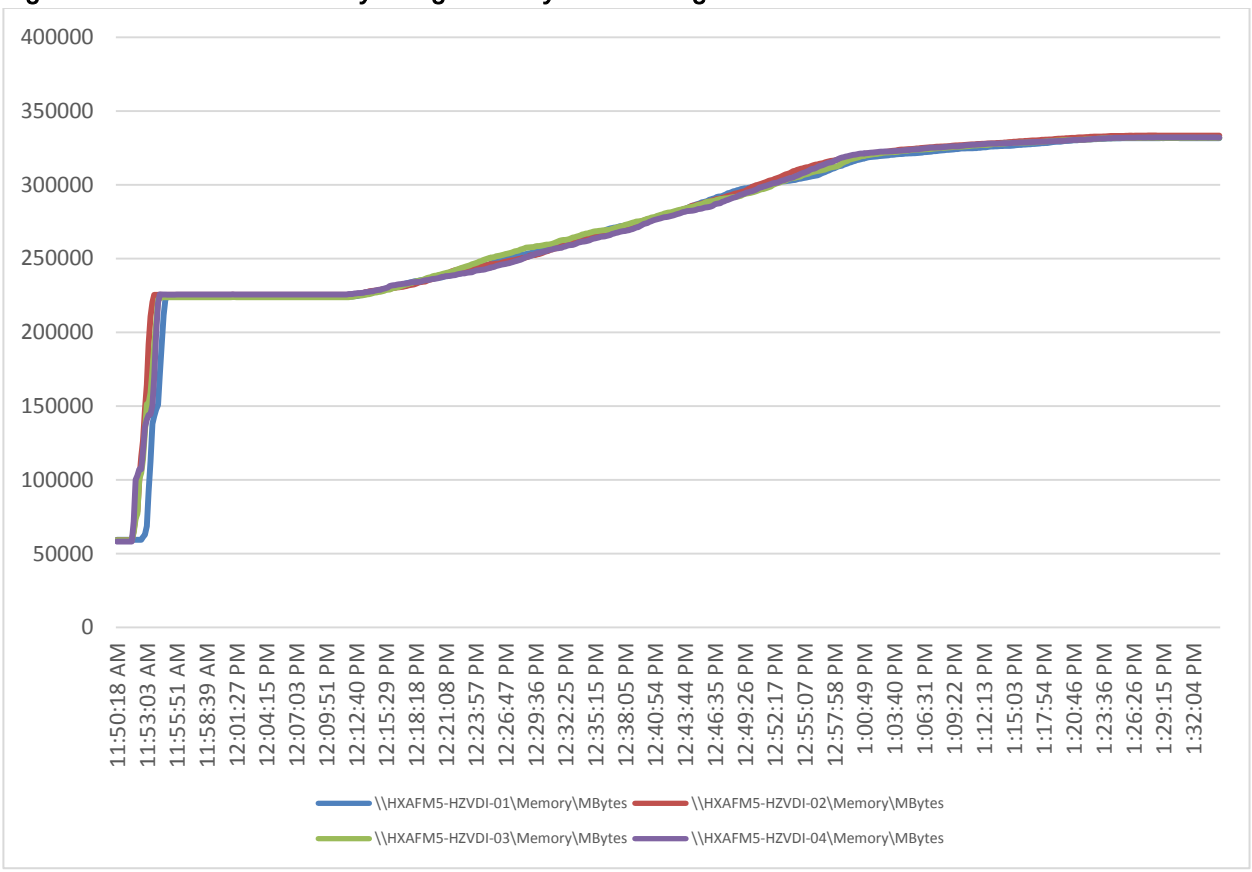

Figure 71 ESXi Host Network Adapter (VMNICs) Mbits Received/Transmitted Per Sec Running 450 Windows 10 Linked-Clone Virtual Desktops

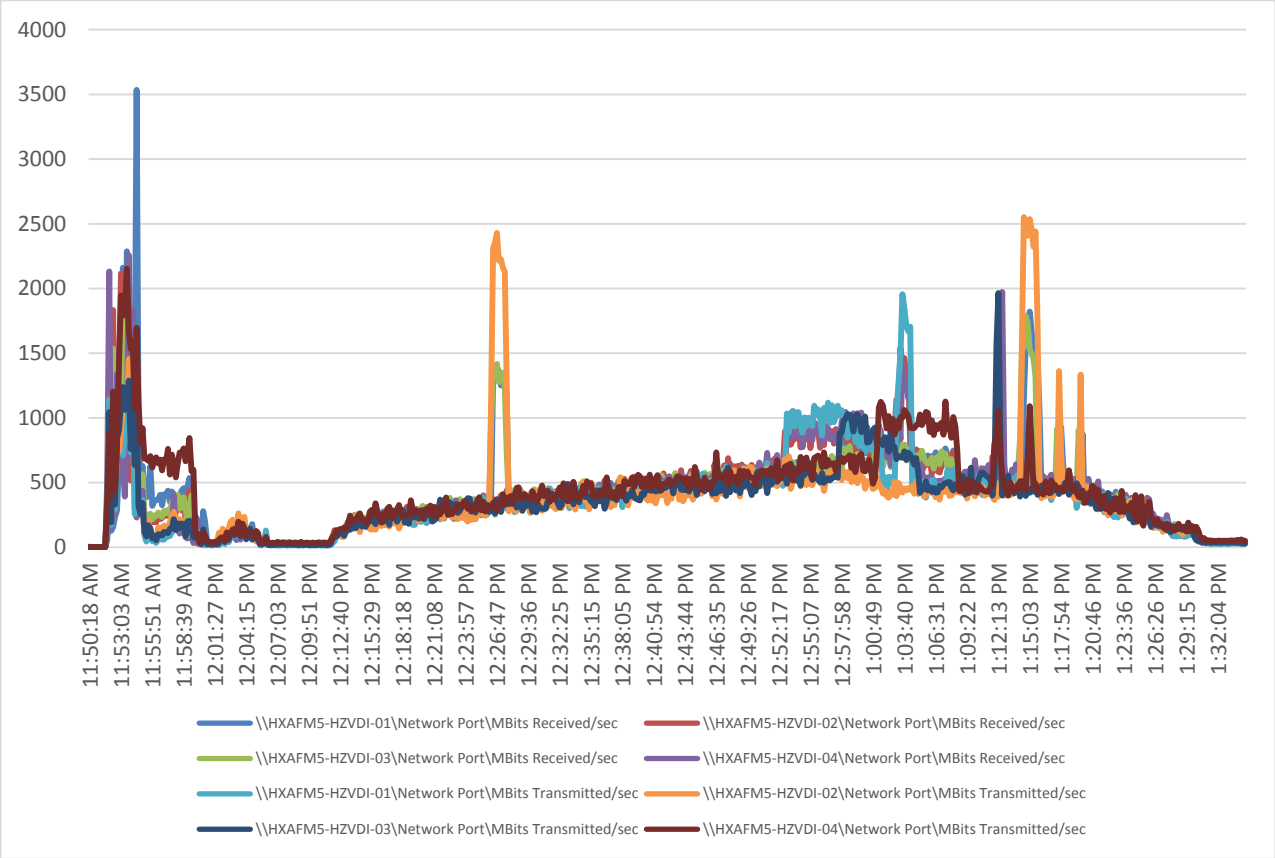

Figure 72 HyperFlex Cluster WebUI Performance Chart for Knowledge Worker Workload Running 450 User Test on Linked-Clone Windows 10

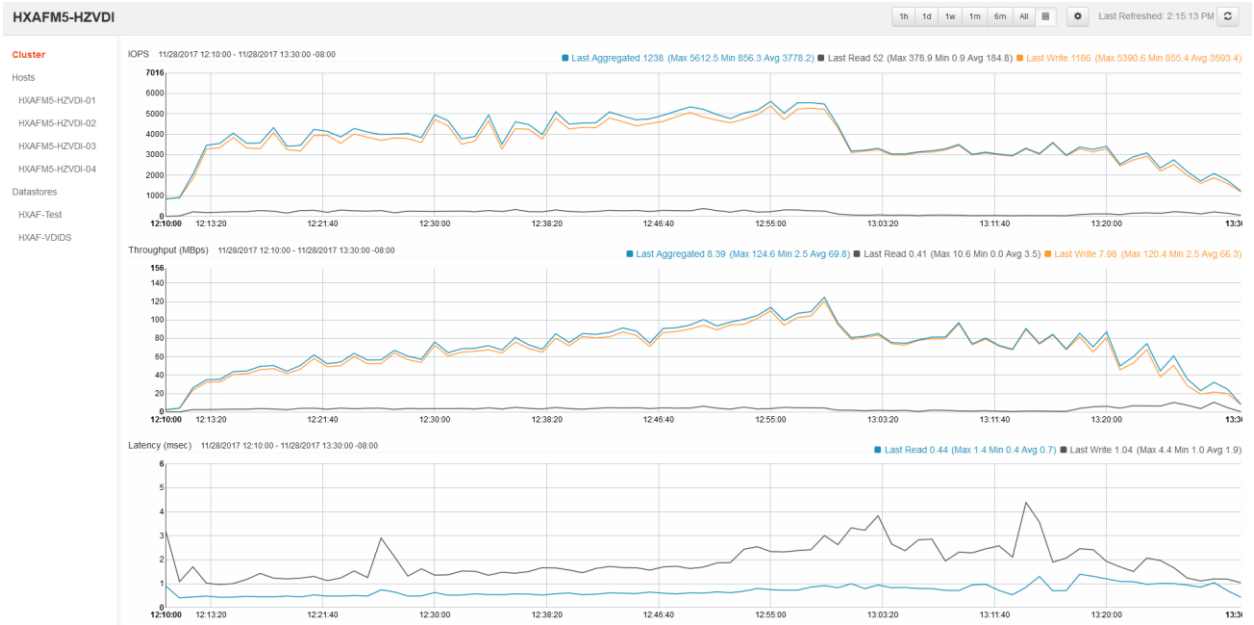

Figure 73 vCenter WebUI Reporting HyperFlex Cluster Deduplication and Compression Savings for 450 Linked-Clone VMs Running Windows 10/Office 2016 Supporting 450 Users

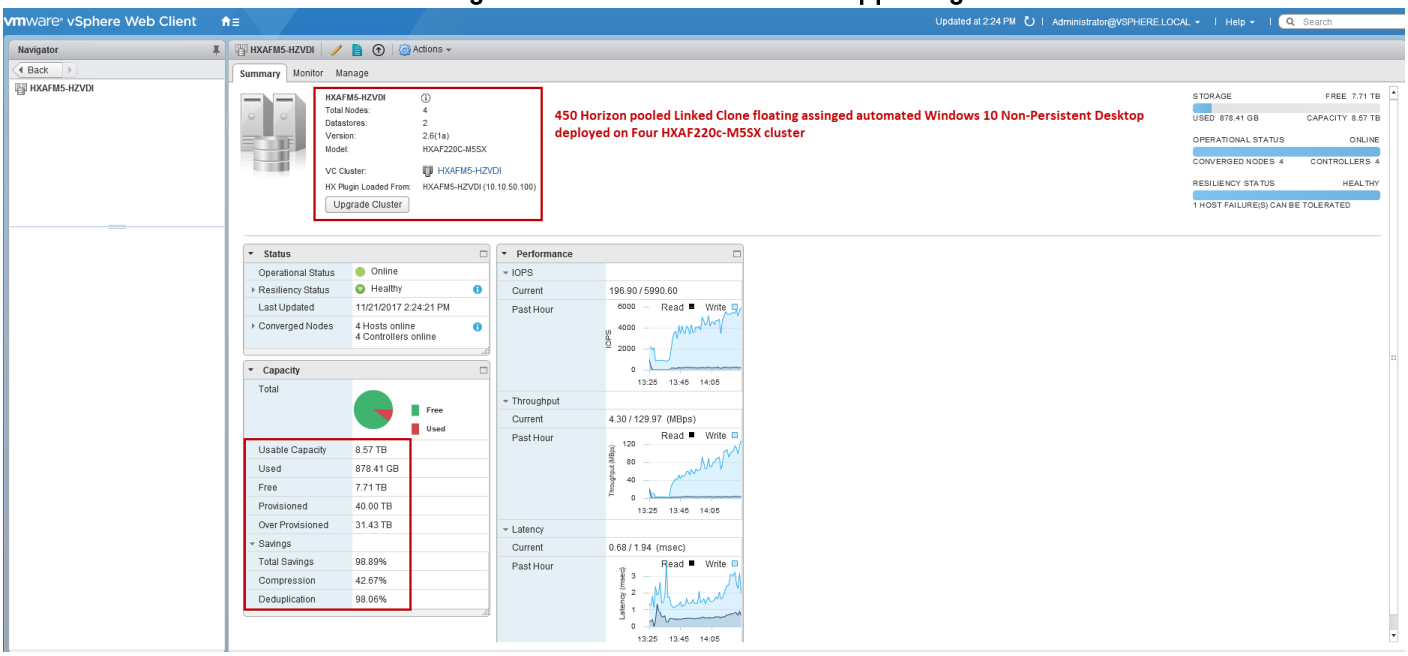

450 Windows 10 Full Clone Desktop Pool Testing on Four Node Cisco HyperFlex Cluster

450 user dedicated assignment automated pool, Windows 10 with Office 2016 full clone desktops on four HXAF220c-M5SX HyperFlex Cluster.

Test result highlights include:

- 0.676 second baseline response time
- 0.986 second average response time with 2000 desktops running
- Average CPU utilization of 80 percent during steady state
- Average of 340GB of RAM used out of 768 GB available per node
- 1000Mbps peak network utilization per host.
- Average Write Latency 1.8ms/Max Write Latency 4.7ms
- Average Read Latency 0.8ms/Max Read Latency 1.4ms
- 3000 peak I/O operations per second (IOPS) at steady state
- 117MBps peak throughput at steady state

Figure 74 Login VSI Analyzer Chart for 450 User Full-Clone Windows 10 Virtual Desktops

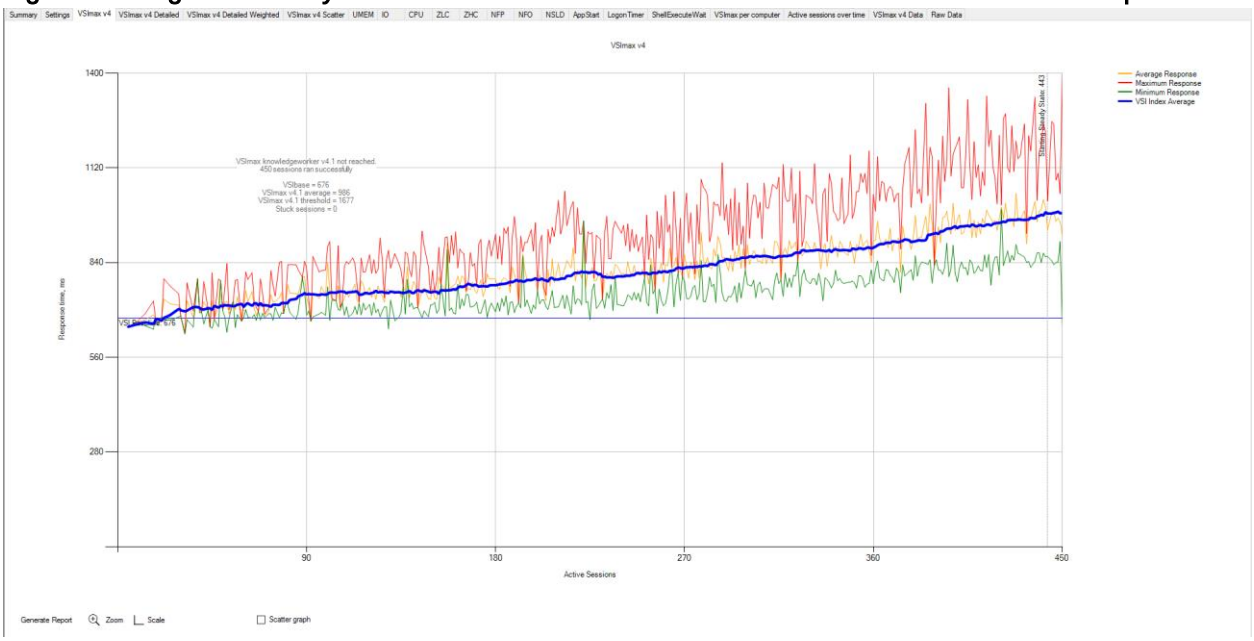

Figure 75 Three Consecutive Test Login VSI Analyzer Chart for 450 User Full-Clone Windows 10 Virtual Desktops

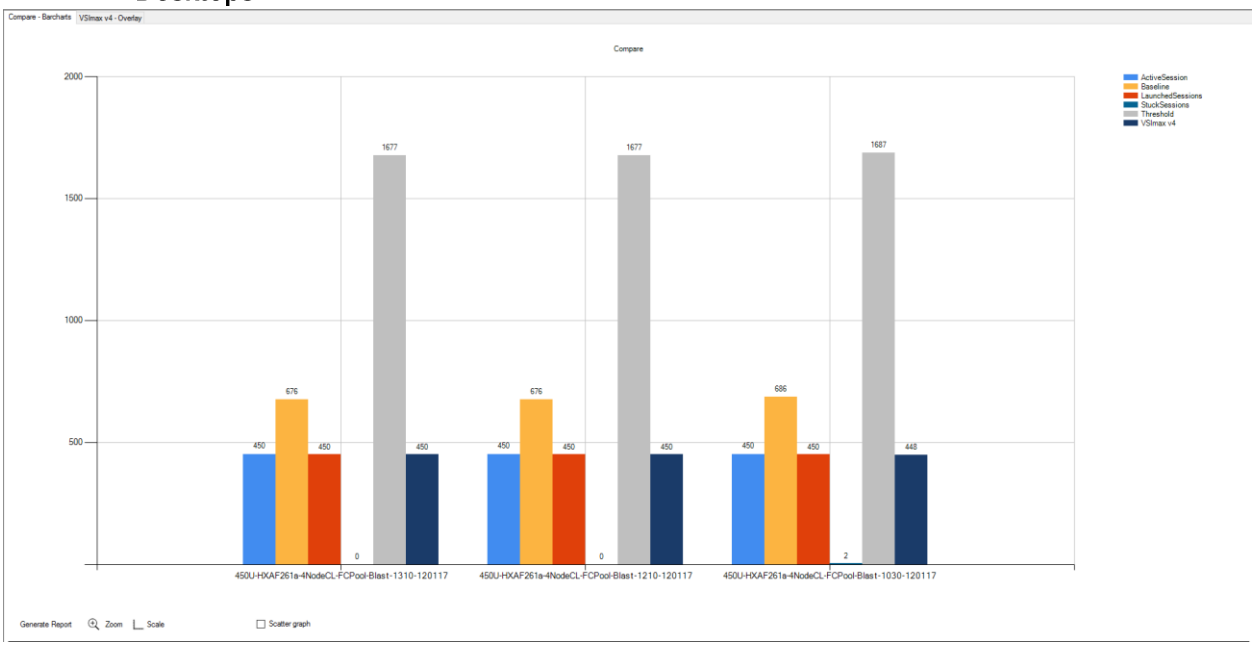

Figure 76 Sample ESXi Host CPU Core Utilization Running 450 User Full-Clone Windows 10 Virtual Desktops

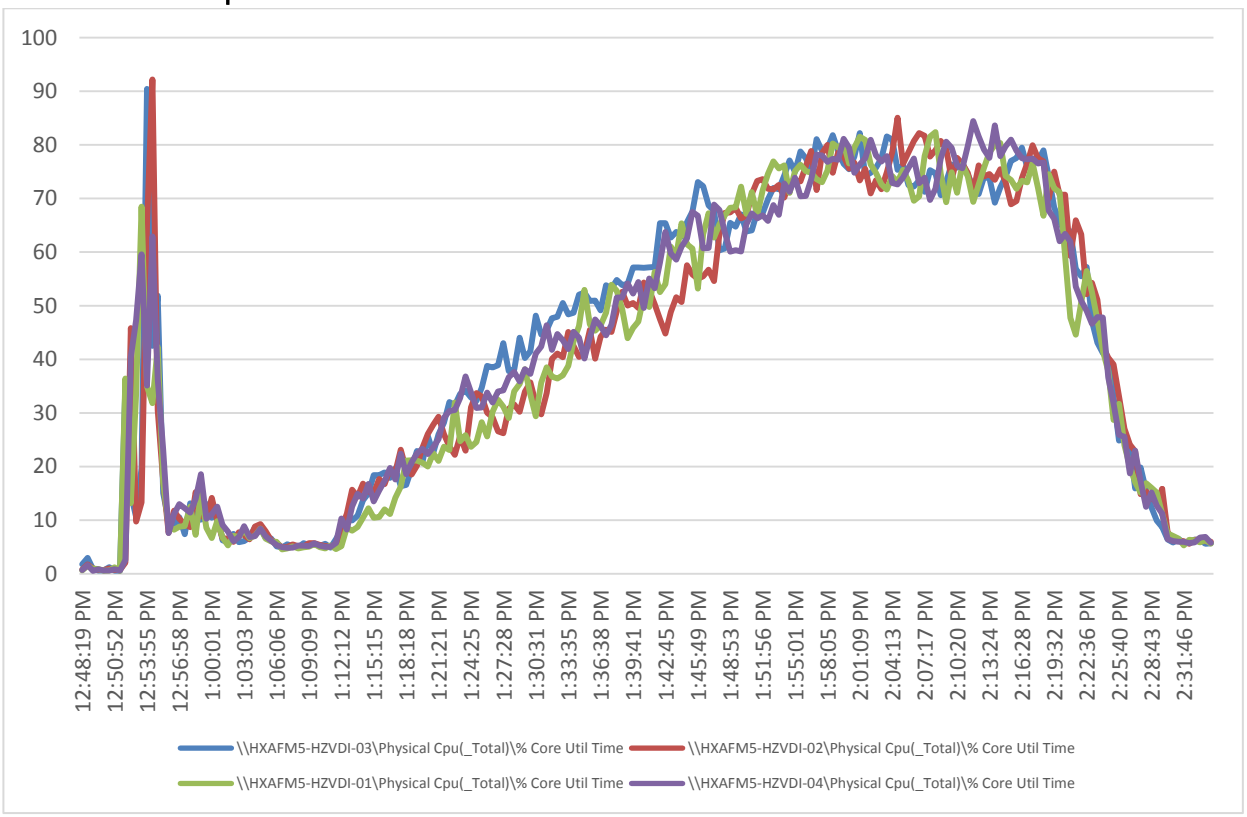

Figure 77 Sample ESXi Host Memory Usage in Mbytes Running 450 User Full-Clone Windows 10 Virtual Desktops

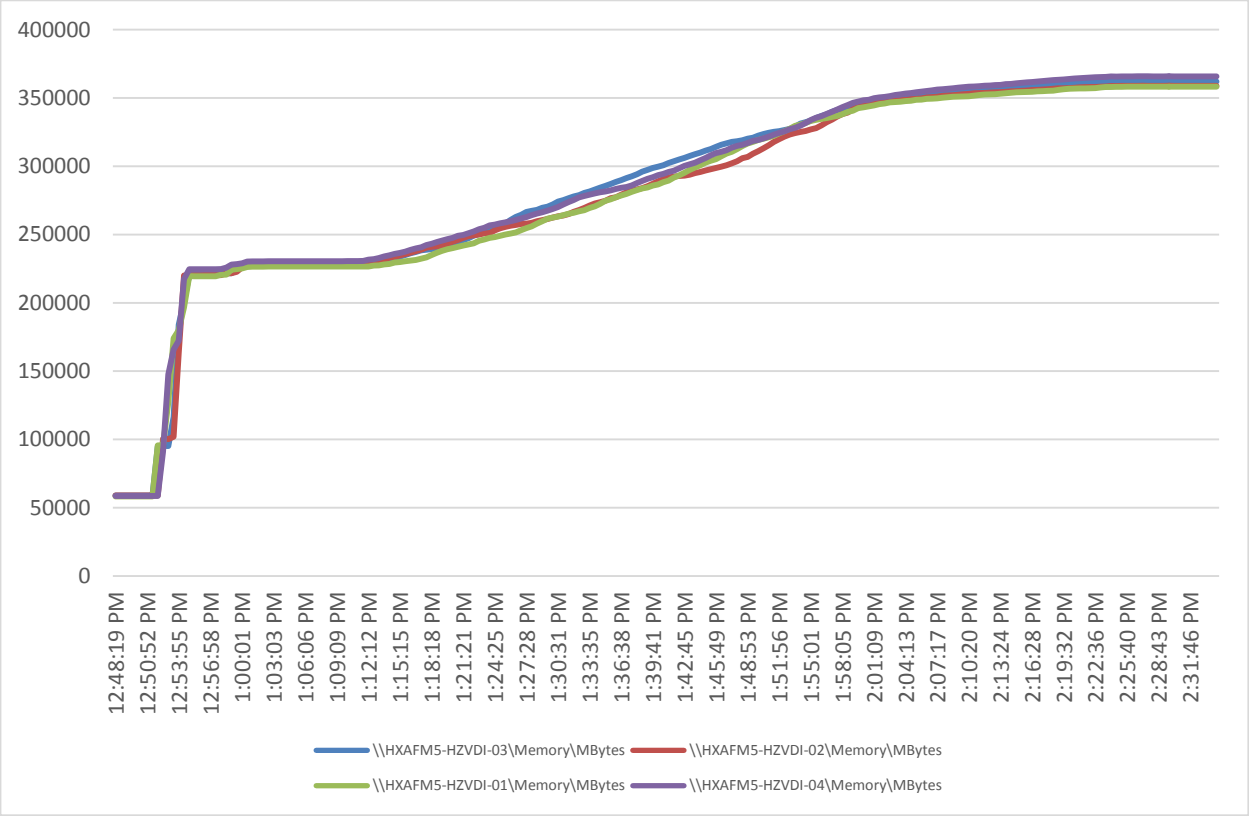

Figure 78 Sample ESXi Host Network Adapter (VMNICs) Mbits Received /Transmitted Per Sec Running 450 User Full-Clone Windows 10 Virtual Desktops

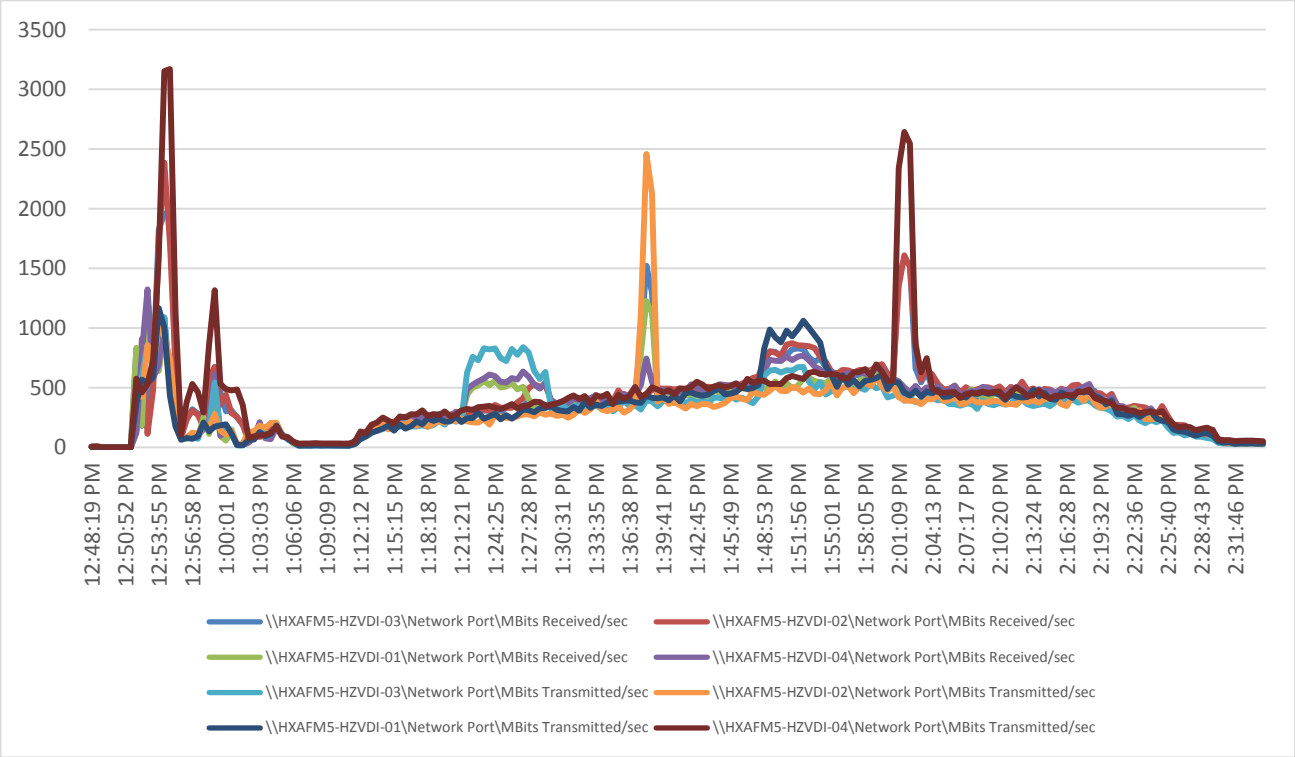

Figure 79 HyperFlex Cluster WebUI Performance Chart for Knowledge Worker Workload Running 450 User Test on Full Clone Windows 10 Virtual Desktops

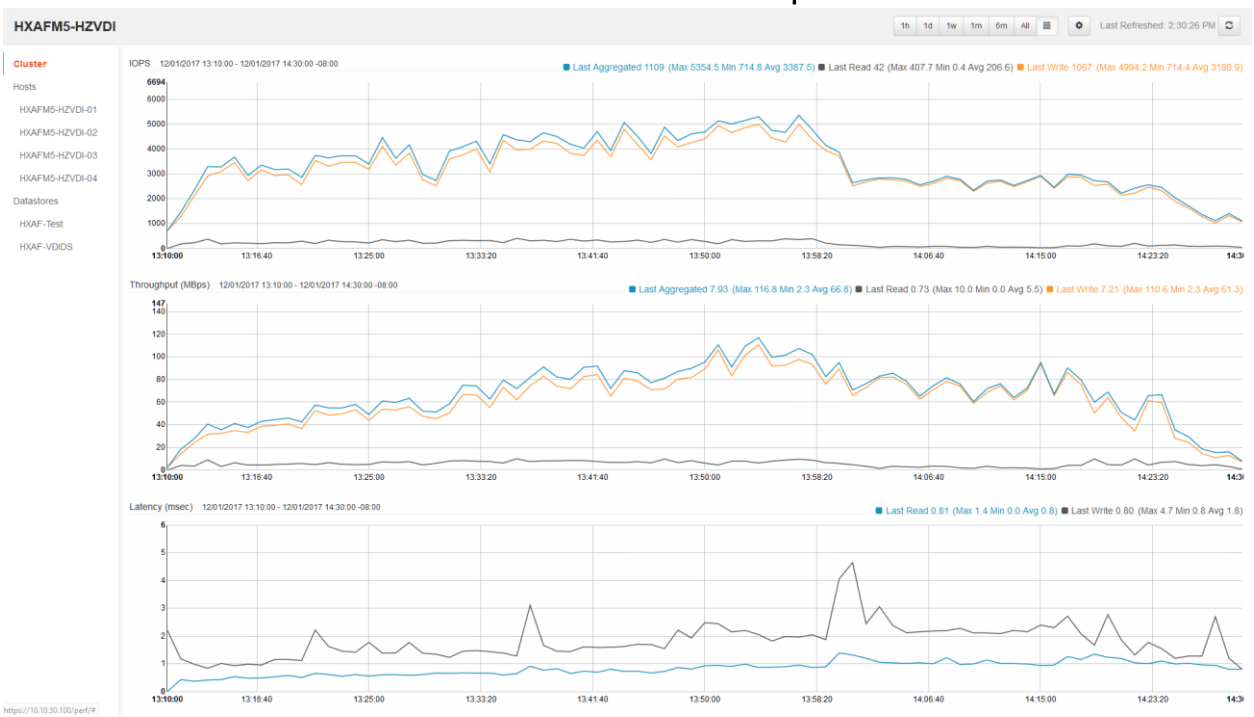

Figure 80 vCenter WebUI Reporting HyperFlex Cluster De-duplication and Compression Savings for 450 Full-Clone VMs running Windows 10/Office 2016 Supporting 450 Users.

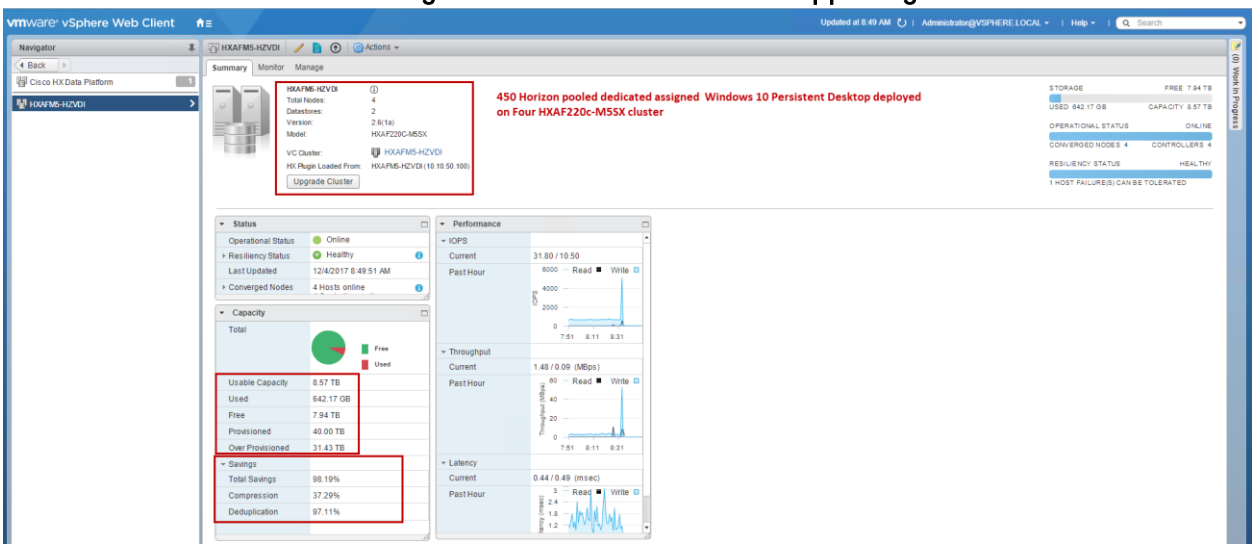

Figure 81 Horizon Administrator Console Reporting Full Clone Desktop Pool Property and 450 Active Sessions

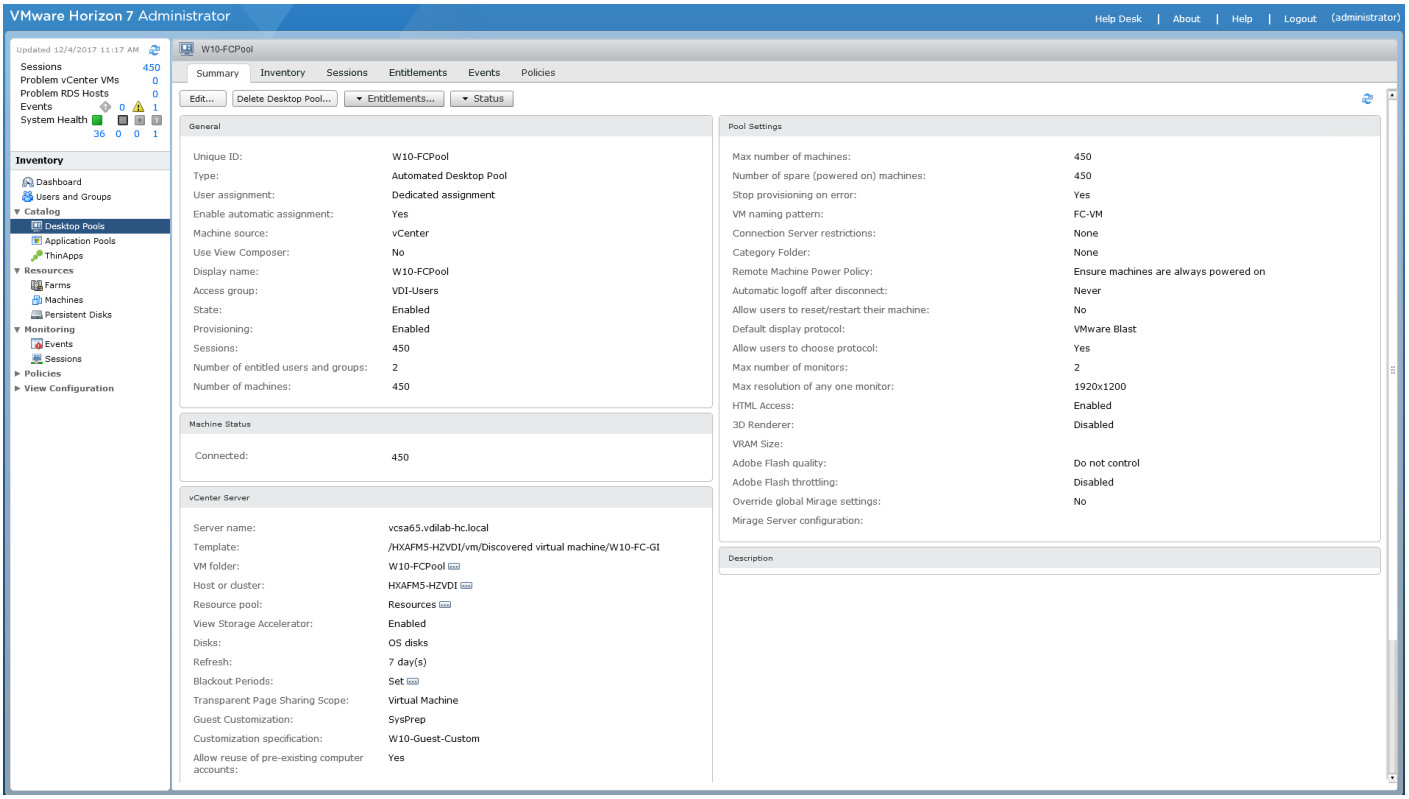

Summary

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 for hosted Apps or desktops has many advantages for desktop virtualization administrators.

The four node tested system can be expanded to 32 nodes (16 hyper converged plus 16 compute only nodes) for an expected user capacity of 4800 knowledge worker users.

The solution if fully capable of supporting graphics accelerated workloads. Each Cisco HyperFlex HXAF240c M5 node and each Cisco UCS C240 M5 server can support up to two NVIDIA M10 or P40 cards. 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.

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

About the Authors

Hardik Patel, Senior Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Hardik is a subject matter expert on Cisco HyperFlex, Cisco Unified Computing System, Cisco Nexus Switching, VMware vSphere and VMware Horizon end user computing. Hardik is a member of the Cisco's Computer Systems Product Group team.

Acknowledgements

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

- Mike Brennan, Product Manager, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.
- Shyam Palakodety, Technical Marketing Engineer, Springpath, Inc.
- Bhumik Patel, VMware Alliance Team, VMware, Inc.

Appendix A – Cisco Nexus 9372 Switch Configuration

Switch A Configuration

```
!Command: show running-config
!Time: Fri Dec 15 17:17:40 2017

version 7.0(3)I2(2d)
switchname XXXXXXXXXXX
class-map type network-qos class-fcoe
match qos-group 1
class-map type network-qos class-all-flood
match qos-group 2
class-map type network-qos class-ip-multicast
match qos-group 2
vdc XXXXXXXXXX id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs eth distribute

feature interface-vlan
feature hsrp
```

```
feature lacp
feature dhcp
feature vpc
feature lldp
clock protocol ntp vdc 1

no password strength-check
username admin password 5 $1$MSJwTJtn$Bo0lrVnESUVxLcbRHg86j1  role network-admin
ip domain-lookup
no service unsupported-transceiver
class-map type qos match-all class-fcoe
policy-map type qos jumbo
  class class-default
    set qos-group 0
copp profile strict
snmp-server user admin network-admin auth md5 0x71d6a9cf1ea007cd3166e91a6f3807e5
 priv 0x71d6a9cf1ea007cd3166e91a6f3807e5 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

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

```
vlan 1,50-54
vlan 50
  name InBand-Mgmt-C1
vlan 51
  name Infra-Mgmt-C1
vlan 52
  name StorageIP-C1
vlan 53
  name 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 1000
  peer-keepalive destination 10.29.132.20 source 10.29.132.19


interface Vlan1
  no shutdown
  ip address 10.29.132.2/24

interface Vlan50
```

```
  no shutdown
  ip address 10.10.50.2/24
  hsrp version 2
  hsrp 50
    preempt
    priority 110
    ip 10.10.50.1

  ip dhcp relay address 10.10.51.21
  ip dhcp relay address 10.10.51.22

interface Vlan51
  no shutdown
  ip address 10.10.51.2/24
  hsrp version 2
  hsrp 51
    preempt
    priority 110
    ip 10.10.51.1

interface Vlan52
  no shutdown
  ip address 10.10.52.2/24
  hsrp version 2
  hsrp 52
    preempt
    priority 110
    ip 10.10.52.1
```

```
interface Vlan53
  no shutdown

  ip address 10.10.53.2/24
  hsrp version 2
  hsrp 53
    preempt
    priority 110
    ip 10.10.53.1

interface Vlan54
  no shutdown
  ip address 10.54.0.2/20
  hsrp version 2
  hsrp 54
    preempt
    priority 110
    ip 10.54.0.1
  ip dhcp relay address 10.10.51.21
  ip dhcp relay address 10.10.51.22

interface port-channel10
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type network

  service-policy type qos input jumbo
  vpc peer-link
```

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

interface port-channel12
  description FI-Uplink-K22
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
  mtu 9216
  service-policy type qos input jumbo
  vpc 12

interface Ethernet1/1
  switchport mode trunk

  switchport trunk allowed vlan 1,50-54
  channel-group 10 mode active

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

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

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

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

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

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

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

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13

interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20
```

```
interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk

interface Ethernet1/26
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk

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

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

  spanning-tree port type edge trunk
```

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

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

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

interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36
```

```
interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44


interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49

interface Ethernet1/50
```

```
interface Ethernet1/51

interface Ethernet1/52

interface Ethernet1/53

interface Ethernet1/54

interface mgmt0
  vrf member management
  ip address 10.29.132.19/24

clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
line vty
boot nxos bootflash:/nxos.7.0.3.I2.2d.bin
```

Switch B Configuration

```
!Command: show running-config
!Time: Fri Dec 15 17:18:36 2017

version 7.0(3)I2(2d)
switchname XXXXXXXXXX
class-map type network-qos class-fcoe
match qos-group 1
class-map type network-qos class-all-flood
match qos-group 2
class-map type network-qos class-ip-multicast
match qos-group 2
```

```
vdc XXXXXXXXXX id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature telnet
cfs eth distribute

feature interface-vlan
feature hsrp
feature lacp
feature dhcp
feature vpc
feature lldp
clock protocol ntp vdc 1

no password strength-check
username admin password 5 $1$jEwHqUvM$gpOec2hramkyX09KD3/Dn.  role network-admin
ip domain-lookup
no service unsupported-transceiver
class-map type qos match-all class-fcoe
policy-map type qos jumbo
  class class-default
    set qos-group 0
copp profile strict
```

```
snmp-server user admin network-admin auth md5 0x9046c100ce1f4ecdd74ef2f92c4e83f9
 priv 0x9046c100ce1f4ecdd74ef2f92c4e83f9 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

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


vlan 1,50-54
vlan 50
  name InBand-Mgmt-C1
vlan 51
  name Infra-Mgmt-C1
vlan 52
  name StorageIP-C1
vlan 53
  name vMotion-C1
vlan 54
  name VM-Data-C1


service dhcp
ip dhcp relay
```

```
ip dhcp relay information option
ipv6 dhcp relay
vrf context management
  ip route 0.0.0.0/0 10.29.132.1
vpc domain 50
  role priority 2000
  peer-keepalive destination 10.29.132.19 source 10.29.132.20


interface Vlan1
  no shutdown
  ip address 10.29.132.3/24

interface Vlan50
  no shutdown
  ip address 10.10.50.3/24
  hsrp version 2
  hsrp 50
    preempt
    priority 110
    ip 10.10.50.1

  ip dhcp relay address 10.10.51.21
  ip dhcp relay address 10.10.51.22

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

```
  hsrp 51
    preempt
    priority 110
    ip 10.10.51.1

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

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

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

```
    priority 110
    ip 10.54.0.1
  ip dhcp relay address 10.10.51.21
  ip dhcp relay address 10.10.51.22

interface port-channel10
  description vPC-PeerLink
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type network

  service-policy type qos input jumbo
  vpc peer-link

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

interface port-channel12
  description FI-Uplink-K22
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk
  mtu 9216
```

```
  service-policy type qos input jumbo
  vpc 12

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

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

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

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

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

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

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

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

interface Ethernet1/9

interface Ethernet1/10

interface Ethernet1/11

interface Ethernet1/12

interface Ethernet1/13
```

```
interface Ethernet1/14

interface Ethernet1/15

interface Ethernet1/16

interface Ethernet1/17

interface Ethernet1/18

interface Ethernet1/19

interface Ethernet1/20

interface Ethernet1/21

interface Ethernet1/22

interface Ethernet1/23

interface Ethernet1/24

interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk

interface Ethernet1/26
  switchport mode trunk
```

```
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk

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

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

  spanning-tree port type edge trunk

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

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

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

```
interface Ethernet1/32
  switchport mode trunk
  switchport trunk allowed vlan 1,50-54
  spanning-tree port type edge trunk

interface Ethernet1/33

interface Ethernet1/34

interface Ethernet1/35

interface Ethernet1/36

interface Ethernet1/37

interface Ethernet1/38

interface Ethernet1/39

interface Ethernet1/40

interface Ethernet1/41

interface Ethernet1/42

interface Ethernet1/43

interface Ethernet1/44
```

```
interface Ethernet1/45

interface Ethernet1/46

interface Ethernet1/47

interface Ethernet1/48
  switchport access vlan 50

interface Ethernet1/49

interface Ethernet1/50

interface Ethernet1/51

interface Ethernet1/52

interface Ethernet1/53

interface Ethernet1/54

interface mgmt0
  vrf member management

  ip address 10.29.132.20/24
clock timezone PST -8 0
clock summer-time PDT 2 Sunday March 02:00 1 Sunday November 02:00 60
line console
```

```
line vty
boot nxos bootflash:/nxos.7.0.3.I2.2d.bin
```