



Cisco HyperFlex 4.0 for Virtual Server Infrastructure with Microsoft Hyper-V

Deployment Guide for Cisco HyperFlex 4.0(1b) for Virtual Server Infrastructure using Microsoft Hyper-V Hypervisor, Cisco UCS 6000 Fabric Interconnect, and Cisco HyperFlex Data Platform Software

Published: July 2020



About the Cisco Validated Design Program

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

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

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

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

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

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

© 2020 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	7
Solution Overview	8
Introduction.....	8
Audience	9
Purpose of this Document.....	9
Microsoft Hyper-V enhancements for HX Version 4.0.x	9
Documentation Roadmap	10
Solution Summary	10
Technology Overview	12
Cisco Unified Computing System	12
Cisco UCS Fabric Interconnect	13
Cisco UCS 6248UP Fabric Interconnect.....	13
Cisco UCS 6296UP Fabric Interconnect.....	14
Cisco UCS 6332 Fabric Interconnect	14
Cisco UCS 6332-16UP Fabric Interconnect	14
Cisco UCS 6454 Fabric Interconnect	14
Cisco HyperFlex HX-Series Nodes	15
Cisco HyperFlex HXAF220c-M5SX All-Flash Node.....	15
Cisco HyperFlex HXAF240c-M5SX All-Flash Node.....	15
Cisco HyperFlex HX220c-M5SX Hybrid Node.....	16
Cisco HyperFlex HX240c-M5SX Hybrid Node.....	16
Cisco HyperFlex HX240c-M5L Hybrid Node	16
Cisco VIC 1387 MLOM Interface Cards.....	17
Cisco VIC 1457 MLOM Interface Cards.....	17
All-Flash versus Hybrid.....	18
Cisco HyperFlex Compute-Only Nodes	19
Cisco HyperFlex Data Platform Software.....	19
Cisco HyperFlex Connect HTML5 Management Web Page	20
Cisco Intersight Cloud Based Management	21
Cisco HyperFlex HX Data Platform Controller	22
Data Operations and Distribution	23
Solution Design	27
Requirements	27
Physical Components.....	27
Software Components	31
Licensing	31
Considerations.....	32

Version Control	32
Microsoft Windows Active Directory.....	33
Scale	34
Capacity	35
Physical Topology.....	37
Topology Overview	37
Fabric Interconnects.....	39
HX-Series Rack-Mount Servers	40
Cisco UCS B-Series Blade Servers	42
Cisco UCS C-Series Rack-Mount Servers	42
Logical Topology	43
Logical Network Design	43
Design Elements.....	45
Network Design.....	45
Cisco UCS Design	47
Cisco UCS Organization.....	48
Cisco UCS LAN Policies.....	48
Cisco UCS Servers Policies	56
Cisco UCS Service Profile Templates	62
Microsoft Hyper-V Host Design.....	63
Virtual Networking Design	63
Discrete Device Assignment (I/O Passthrough)	66
Storage Platform Controller Virtual Machines	66
Installation	71
Prerequisites.....	71
IP Addressing	71
DHCP versus Static IP.....	74
Configure the Active Directory for Constrained Delegation	74
Prepopulate AD DNS with Records	76
NTP	78
VLANs	80
Network Uplinks	80
Usernames and Passwords	81
Physical Installation.....	82
Cabling	82
Cisco UCS Installation	86
Cisco UCS Fabric Interconnect A.....	86
Cisco UCS Fabric Interconnect B.....	87
Cisco UCS Manager.....	88

Cisco UCS Configuration	89
Cisco UCS Firmware	89
NTP	90
Uplink Ports	90
Uplink Port Channels	91
Server Ports	92
Server Discovery	94
Deploy HX Data Platform Installer on Hyper-V Infrastructure	94
Assign a Static IP Address to the HX Data Platform Installer Virtual Machine	99
HyperFlex Installer Web Page	100
HyperFlex Installation	102
Post Installation Tasks	119
Create Datastores	119
Constrained Delegation (Optional)	121
Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces	124
Rename the Cluster Network in Windows Failover Cluster - Optional	127
Configure the Windows Failover Cluster Network Roles	127
Configure the Windows Failover Cluster Network for Live Migration	128
Create Folders on the HX Datastore	129
Configure the Default Folder to Store Virtual Machine Files on Hyper-V	130
Validate the Windows Failover Cluster Configuration	130
Configure Quorum in Windows Server Failover Cluster	131
Initial Tasks and Testing	132
Ready Clones	132
Auto-Support and Notifications	134
Smart Licensing	136
HyperFlex Cluster Expansion	138
Expansion with Converged Nodes	138
Expansion with Compute-Only Nodes	149
Management	173
HyperFlex Connect	173
Dashboard	174
Monitor	174
Analyze	175
Manage	176
Configure Remote Management Stations for Managing VMs in HX Cluster	177
Microsoft Hyper-V Manager	179
Change the Default Location to Store the Virtual Machine Files using Hyper-V Manager	180
Create Virtual Machines using Hyper-V Manager	182

Windows Failover Cluster Manager	184
Microsoft System Center Virtual Machine Manager 2019	190
Create Run-As Account for Managing the Hyper-V Cluster	190
Manage Servers and Clusters	191
Networking	195
Storage	196
Create a Virtual Machine using SCVMM	198
Microsoft Windows Admin Center (WAC)	205
Connect to Managed Nodes and Clusters	205
Manage Servers with WAC	208
Manage a Failover Cluster with WAC	210
Appendix	213
A: Cluster Capacity Calculations	213
B: Install Microsoft Windows Server 2016/2019	213
Configure Cisco UCS Manager using HX Installer	213
Configure Cisco UCS vMedia and Boot Policies	214
Install Microsoft Windows Server 2019 OS	221
Undo vMedia and Boot Policy Changes	224
Undo vMedia and Boot Policy Changes (for Compute-only Nodes)	225
C: Live Migration of Virtual Machines Between a Standalone Hyper-V Host and HyperFlex Hyper-V Host	226
Prerequisites	226
Configure HX Storage Access to the Standalone Hyper-V Host	226
Configure the Source and Destination Computers for Live Migration	228
Configure Performance Options for Live Migration - Optional	230
Move a Running Virtual Machine	230
D: Delegate HX Service Account with Least Privileges for Administrative Tasks	233
E: Common Resiliency and High Availability Scenario for HX Hyper-V	237
Hyper-V Failover Clustering Resiliency	237
F: Antivirus Best practices for Hyper-V and Windows Server Failover Cluster	238
About the Authors	239
Acknowledgements	239



Executive Summary

With the proliferation of virtualized environments across most IT landscapes, other technology stacks which have traditionally not offered the same levels of simplicity, flexibility, and rapid deployment as virtualized compute platforms have come under increasing scrutiny. In particular, networking devices and storage systems have lacked the agility of hypervisors and virtual servers. With the introduction of Cisco HyperFlex, Cisco has brought the dramatic enhancements of hyperconvergence to the modern data center.

Cisco HyperFlex systems are based on the Cisco UCS platform, combining Cisco HX-Series x86 servers and integrated networking technologies through the Cisco UCS Fabric Interconnects, into a single management domain, along with industry leading virtualization hypervisor software from Microsoft, and next-generation software defined storage technology. The combination creates a complete virtualization platform, which provides the network connectivity for the guest virtual machine connections, and the distributed storage to house the virtual machines, spread across all of the Cisco UCS x86 servers, versus using specialized storage or networking components. The unique storage features of the HyperFlex log-based filesystem enable rapid cloning of virtual machines, snapshots without the traditional performance penalties, and inline data deduplication and compression. All configuration, deployment, management, and monitoring of the solution can be done with existing tools for Cisco UCS and Microsoft, such as Cisco UCS Manager and Microsoft Hyper-V Manager, PowerShell, SCVMM, and new integrated HTML based management tools, such as Cisco HyperFlex Connect and Cisco Intersight. This powerful linking of advanced technology stacks into a single, simple, rapidly deployed solution makes Cisco HyperFlex a true second generation hyperconverged platform. Customers can choose to deploy SSD-only All-Flash HyperFlex clusters for improved performance, increased density, and reduced latency, or use HyperFlex hybrid clusters which combine high-performance SSDs and low-cost, high-capacity HDDs to optimize the cost of storing data. Further enhancements include improvements and customization capabilities in the HyperFlex Connect management tool, larger scale 32-node clusters, large form-factor disks for larger storage capacity.

Solution Overview

Introduction

The Cisco HyperFlex System provides an all-purpose virtualized server platform, with hypervisor hosts, networking connectivity, and virtual server storage across a set of Cisco UCS HX-Series x86 rack-mount servers. Legacy datacenter deployments have relied on a disparate set of technologies, each performing a distinct and specialized function, such as network switches connecting endpoints and transferring Ethernet network traffic, and Fibre Channel (FC) storage arrays providing block based storage via a dedicated storage array network (SAN). Each of these systems had unique requirements for hardware, connectivity, management tools, operational knowledge, monitoring, and ongoing support. A legacy virtual server environment was often divided up into areas commonly referred to as silos, within which only a single technology operated, along with their correlated software tools and support staff. Silos could often be divided between the x86 computing hardware, the networking connectivity of those x86 servers, SAN connectivity and storage device presentation, the hypervisors and virtual platform management, and finally the guest virtual machine themselves along with their OS and applications. This model proves to be inflexible, difficult to navigate, and is susceptible to numerous operational inefficiencies.

A more modern datacenter model was developed called a converged infrastructure. Converged infrastructures attempt to collapse the traditional silos by combining these technologies into a more singular environment, which has been designed to operate together in pre-defined, tested, and validated designs. A key component of the converged infrastructure was the revolutionary combination of x86 rack and blade servers, along with converged Ethernet and Fibre Channel networking offered by the Cisco UCS platform. Converged infrastructures leverage Cisco UCS, plus new deployment tools, management software suites, automation processes, and orchestration tools to overcome the difficulties deploying traditional environments and do so in a much more rapid fashion. These new tools place the ongoing management and operation of the system into the hands of fewer staff, with more rapid deployment of workloads based on business needs, while still remaining at the forefront of flexibility to adapt to workload needs and offering the highest possible performance. Cisco has had incredible success in these areas with our various partners, developing leading solutions such as Cisco FlexPod, FlashStack, VersaStack, and VxBlock architectures. Despite these advances, because these converged infrastructures contained some legacy technology stacks, particularly in the storage subsystems, there often remained a division of responsibility amongst multiple teams of administrators. There is also a recognition that these converged infrastructures can still be a somewhat complex combination of components, where a simpler system would suffice to serve the workloads being requested.

Significant changes in the storage marketplace have given rise to the software defined storage (SDS) system. Legacy FC storage arrays often contained a specialized subset of hardware, such as Fibre Channel Arbitrated Loop (FC-AL) based controllers and disk shelves along with optimized Application Specific Integrated Circuits (ASIC), read/write data caching modules and cards, plus highly customized software to operate the arrays. With the rise of Serial Attached SCSI (SAS) bus technology and its inherent benefits, storage array vendors began to transition their internal hardware architectures to SAS, and with dramatic increases in processing power from recent x86 processor architectures, they also used fewer or no custom ASICs at all. As disk physical sizes shrank, x86 servers began to have the same density of storage per rack unit (RU) as the arrays themselves, and with the proliferation of NAND based flash memory solid state disks (SSD), they also now had access to input/output (IO) devices whose speed rivaled that of dedicated caching devices. If servers themselves now contained storage devices and technology to rival many dedicated arrays on the market, then the major differentiator between them was the software providing allocation, presentation, and management of the storage, plus the advanced features many vendors offered. This has led to the rise of software defined storage, where the x86 servers with the storage devices ran software to effectively turn one or more of them, working cooperatively, into a storage array much the same as the traditional arrays were. In a somewhat unexpected turn of events, some of the major storage array vendors themselves were pioneers in this field, recognizing the technological shifts in the market,

and attempting to profit from the software features they offered versus their specialized hardware, as had been done in the past.

Some early uses of SDS systems simply replaced the traditional storage array in the converged architectures as described earlier. That configuration still had a separate storage system from the virtual server hypervisor platform, and depending on the solution provider, still remained separate from the network devices. If the servers that hosted the virtual machines, and also provided the SDS environment were in fact the same model of server, could they simply do both things at once and collapse the two functions into one? This ultimate combination of resources becomes what the industry has given the moniker of a hyperconverged infrastructure. Hyperconverged infrastructures coalesce the computing, memory, hypervisor, and storage devices of servers into a single platform for virtual servers. There is no longer a separate storage system, as the servers running the hypervisors also provide the software defined storage resources to store the virtual servers, effectively storing the virtual machines on themselves. Now nearly all the silos are gone, and a hyperconverged infrastructure becomes something almost completely self-contained, simpler to use, faster to deploy, easier to consume, yet still flexible and with very high performance. Many hyperconverged systems still rely on standard networking components, such as on-board network cards in the x86 servers, and top-of-rack switches. The Cisco HyperFlex system combines the convergence of computing and networking provided by Cisco UCS, along with next-generation hyperconverged storage software, to uniquely provide the compute resources, network connectivity, storage, and hypervisor platform to run an entire virtual environment, all contained in a single uniform system.

Some key advantages of hyperconverged infrastructures are the simplification of deployment, day to day management operations, as well as increased agility, thereby reducing the amount operational costs. Since hyperconverged storage can be easily managed by an IT generalist, this can also reduce technical debt going forward that is often accrued by implementing complex systems that need dedicated management teams and skillsets.

Audience

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

Purpose of this Document

This document describes the steps required to deploy, configure, and manage a Cisco HyperFlex system using the Microsoft Hyper-V hypervisor. The document is based on all known best practices using the software, hardware and firmware revisions specified in the document. As such, recommendations and best practices can be amended with later versions. This document showcases the installation and configuration of Cisco HyperFlex with Hyper-V in a typical customer datacenter environment. While readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to the deployment of this solution are provided in this CVD.

Microsoft Hyper-V enhancements for HX Version 4.0.x

The Cisco HyperFlex system has several new capabilities and enhancements for Hyper-V in version 4.0.1x:

- Support for second generation Intel Xeon scalable processors
- Windows Server 2019 with Hyper-V–Support has been added in this release for the Windows Server 2019 operating system for Hyper-V based HyperFlex deployments.

- New Cache and increased scale for Hyper-V –NVMe & Optane SSDs are now supported as cache drives for Hyper-V deployments.
- Scale limits have been increased to 16+16 (Converged+Compute-only) for both SFF (AF and Hybrid) & LFF (Hybrid) clusters.
- Support for Cisco UCS VIC 1400 series and 6400 series Fabric Interconnect
- Support for Cisco UCS C240 M5 servers as Compute-only node.

Documentation Roadmap

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



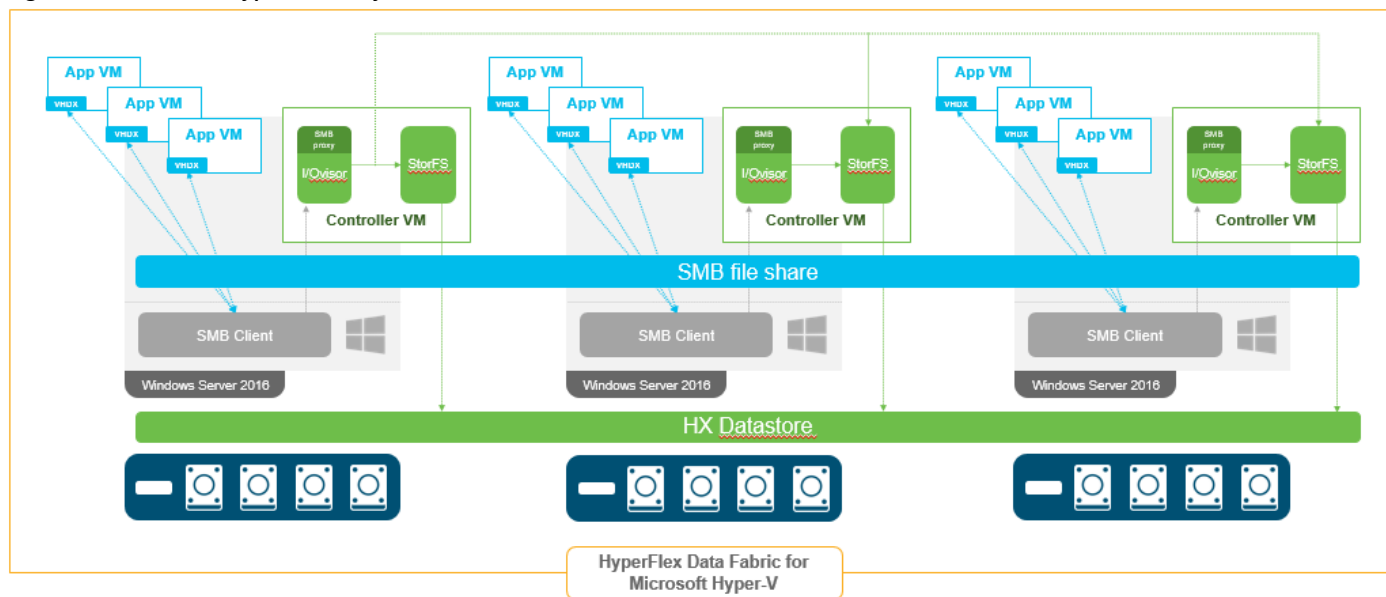
A login is required for the Documentation Roadmap.

This is the Hyperconverged Infrastructure web link: <http://hyperflex.io>

Solution Summary

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

Figure 1 Cisco HyperFlex System Overview



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

- One pair of Cisco UCS Fabric Interconnects, choose from the following models:
 - Cisco UCS 6248UP Fabric Interconnect
 - Cisco UCS 6296UP Fabric Interconnect

- Cisco UCS 6332 Fabric Interconnect
- Cisco UCS 6332-16UP Fabric Interconnect
- Cisco UCS 6454 Fabric Interconnect
- Three to Eight Cisco HyperFlex HX-Series Rack-Mount Servers, choose from the following models:
 - Cisco HyperFlex HX220c-M5SX Rack-Mount Servers
 - Cisco HyperFlex HX240c-M5SX Rack-Mount Servers
 - Cisco HyperFlex HXAF220c-M5SX All-Flash Rack-Mount Servers
 - Cisco HyperFlex HXAF240c-M5SX All-Flash Rack-Mount Servers
 - Cisco HyperFlex HX240c-M5L Rack-Mount Servers
- Cisco HyperFlex Data Platform Software
- Microsoft Windows Server 2016/2019 Hyper-V Hypervisor
- Microsoft Windows Active Directory and DNS services, RSAT tools (end-user supplied)
- SCVMM – optional (end-user supplied)

Optional components for additional compute-only resources are:

- Cisco UCS 5108 Chassis
- Cisco UCS 2204XP, 2208XP or 2304 model Fabric Extenders
- Cisco UCS B200-M4, or B200-M5 blade servers
- Cisco UCS C220-M5 Rack-Mount servers

Technology Overview

Cisco Unified Computing System

Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that unites compute, network, 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 10 Gigabit Ethernet or 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.

The main components of Cisco Unified Computing System are:

- **Computing:** The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers based on Intel Xeon Processors.
- **Network:** The system is integrated onto a low-latency, lossless, 10-Gbps, 25-Gbps or 40-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are often separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, 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 access:** The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying storage access, the Cisco Unified Computing System can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This provides customers with their choice of storage protocol and physical architecture, and enhanced investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.
- **Management:** The system uniquely integrates all system components which enable the entire solution to be managed as a single entity by the Cisco UCS Manager (UCSM). The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations.

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system which 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 Fabric Interconnect

The Cisco UCS Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. Depending on the model chosen, the Cisco UCS Fabric Interconnect offers line-rate, low-latency, lossless 10 Gigabit or 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel connectivity. Cisco UCS Fabric Interconnects provide the management and communication backbone for the Cisco UCS C-Series, S-Series, and HX-Series Rack-Mount Servers, Cisco UCS B-Series Blade Servers, and Cisco UCS 5100 Series Blade Server Chassis. All servers and chassis, and therefore all blades, attached to the Cisco UCS Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabrics, the Cisco UCS Fabric Interconnects provide both the LAN and SAN connectivity for all servers within its domain.

From a networking perspective, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low latency, line rate 10 Gigabit Ethernet on all ports, up to 1.92 Tbps switching capacity and 160 Gbps bandwidth per chassis, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over the Ethernet fabric from the servers to the uplinks. 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.

The Cisco UCS 6300 Series offers the same features while supporting even higher performance, low latency, lossless, line rate 40 Gigabit Ethernet, with up to 2.56 Tbps of switching capacity. Backward compatibility and scalability are assured with the ability to configure 40 Gbps quad SFP (QSFP) ports as breakout ports using 4x10GbE breakout cables. Existing Cisco UCS servers with 10GbE interfaces can be connected in this manner, although Cisco HyperFlex nodes must use a 40GbE VIC adapter in order to connect to a Cisco UCS 6300 Series Fabric Interconnect.

The Cisco UCS 6454 uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10/25/40/100 Gigabit Ethernet ports, a switching capacity of 3.82 Tbps, and 320 Gbps bandwidth between FI 6454 and IOM 2208 per 5108 blade chassis, independent of packet size and enabled services. The product family supports Cisco® low-latency, lossless 10/25/40/100 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over a lossless Ethernet fabric from the server through the Fabric Interconnect.

Cisco UCS 6248UP Fabric Interconnect

The Cisco UCS 6248UP Fabric Interconnect is a one-rack-unit (1RU) 10 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 960 Gbps throughput and up to 48 ports. The switch has 32 1/10-Gbps fixed Ethernet, FCoE, or 1/2/4/8 Gbps FC ports, plus one expansion slot.

Figure 2 Cisco UCS 6248UP Fabric Interconnect



Cisco UCS 6296UP Fabric Interconnect

The Cisco UCS 6296UP Fabric Interconnect is a two-rack-unit (2RU) 10 Gigabit Ethernet, FCoE, and native Fibre Channel switch offering up to 1920 Gbps of throughput and up to 96 ports. The switch has 48 1/10-Gbps fixed Ethernet, FCoE, or 1/2/4/8 Gbps FC ports, plus three expansion slots.

Figure 3 Cisco UCS 6296UP Fabric Interconnect



Cisco UCS 6332 Fabric Interconnect

The Cisco UCS 6332 Fabric Interconnect is a one-rack-unit (1RU) 40 Gigabit Ethernet and FCoE switch offering up to 2560 Gbps of throughput. The switch has 32 40-Gbps fixed Ethernet and FCoE ports. Up to 24 of the ports can be reconfigured as 4x10Gbps breakout ports, providing up to 96 10-Gbps ports.

Figure 4 Cisco UCS 6332 Fabric Interconnect



Cisco UCS 6332-16UP Fabric Interconnect

The Cisco UCS 6332-16UP Fabric Interconnect is a one-rack-unit (1RU) 10/40 Gigabit Ethernet, FCoE, and native Fibre Channel switch offering up to 2430 Gbps of throughput. The switch has 24 40-Gbps fixed Ethernet and FCoE ports, plus 16 1/10-Gbps fixed Ethernet, FCoE, or 4/8/16 Gbps FC ports. Up to 18 of the 40-Gbps ports can be reconfigured as 4x10Gbps breakout ports, providing up to 88 total 10-Gbps ports.

Figure 5 Cisco UCS 6332-16UP Fabric Interconnect



When used for a Cisco HyperFlex deployment, due to mandatory QoS settings in the configuration, the 6332 and 6332-16UP will be limited to a maximum of four 4x10Gbps breakout ports, which can be used for other non-HyperFlex servers.

Cisco UCS 6454 Fabric Interconnect

The Cisco UCS 6454 54-Port Fabric Interconnect is a One-Rack-Unit (1RU) 10/25/40/100 Gigabit Ethernet, FCoE and Fibre Channel switch offering up to 3.82 Tbps throughput and up to 54 ports. The switch has 36 10/25-Gbps Ethernet ports, 4 1/10/25-Gbps Ethernet ports, 6 40/100-Gbps Ethernet uplink ports and 8 unified

ports that can support 8 10/25-Gbps Ethernet ports or 8/16/32-Gbps Fibre Channel ports. All Ethernet ports are capable of supporting FCoE.

Figure 6 Cisco UCS 6454 Fabric Interconnect



Cisco HyperFlex HX-Series Nodes

A HyperFlex cluster requires a minimum of three HX-Series “converged” 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 also is equipped with additional disks, up to the platform’s physical limit, for long term storage and capacity.



In the following Cisco UCS HX server models, SED and NVMe cache drives are not supported on HyperFlex systems with Microsoft Hyper-V at the time of publishing this document.

Cisco HyperFlex HXAF220c-M5SX All-Flash Node

This small footprint Cisco HyperFlex all-flash model contains a 240 GB M.2 form factor solid-state disk (SSD) that acts as the boot drive, a 240 GB housekeeping SSD drive, either a single 375 GB Optane NVMe SSD, a 1.6 TB NVMe SSD or 400GB SAS SSD write-log drive, and six to eight 960 GB or 3.8 TB SATA SSD drives for storage capacity. For configurations requiring self-encrypting drives, the caching SSD is replaced with an 800 GB SAS SED SSD, and the capacity disks are also replaced with either 800 GB, 960 GB or 3.8 TB SED SSDs.

Figure 7 HXAF220c-M5SX All-Flash Node



Cisco HyperFlex HXAF240c-M5SX All-Flash Node

This capacity optimized Cisco HyperFlex all-flash model contains a 240 GB M.2 form factor solid-state disk (SSD) that acts as the boot drive, a 240 GB housekeeping SSD drive, either a single 375 GB Optane NVMe SSD, a 1.6 TB NVMe SSD or 400GB SAS SSD write-log drive installed in a rear hot swappable slot, and six to twenty-three 960 GB or 3.8 TB SATA SSD drives for storage capacity. For configurations requiring self-encrypting drives, the caching SSD is replaced with an 800 GB SAS SED SSD, and the capacity disks are also replaced with either 800 GB, 960 GB or 3.8 TB SED SSDs.

Figure 8 HXAF240c-M5SX Node



Cisco HyperFlex HX220c-M5SX Hybrid Node

This small footprint Cisco HyperFlex hybrid model contains a minimum of six, and up to eight 1.8 terabyte (TB) or 1.2 TB SAS hard disk drives (HDD) that contribute to cluster storage capacity, a 240 GB SSD housekeeping drive, a 480 GB or 800 GB SSD caching drive, and a 240 GB M.2 form factor SSD that acts as the boot drive. For configurations requiring self-encrypting drives, the caching SSD is replaced with an 800 GB SAS SED SSD, and the capacity disks are replaced with 1.2TB SAS SED HDDs.

Figure 9 HX220c-M5SX Node



Either a 480 GB or 800 GB caching SAS SSD may be chosen. This option is provided to allow flexibility in ordering based on product availability, pricing, and lead times. There is no performance, capacity, or scalability benefit in choosing the larger disk.

Cisco HyperFlex HX240c-M5SX Hybrid Node

This capacity optimized Cisco HyperFlex hybrid model contains a minimum of six and up to twenty-three 1.8 TB or 1.2 TB SAS small form factor (SFF) hard disk drives (HDD) that contribute to cluster storage, a 240 GB SSD housekeeping drive, a single 1.6 TB SSD caching drive installed in a rear hot swappable slot, and a 240 GB M.2 form factor SSD that acts as the boot drive. For configurations requiring self-encrypting drives, the caching SSD is replaced with a 1.6 TB SAS SED SSD, and the capacity disks are replaced with 1.2TB SAS SED HDDs.

Figure 10 HX240c-M5SX Node



Cisco HyperFlex HX240c-M5L Hybrid Node

This density optimized Cisco HyperFlex hybrid model contains a minimum of six and up to twelve 6 TB or 8 TB SAS large form factor (LFF) hard disk drives (HDD) that contribute to cluster storage, a 240 GB SSD housekeeping drive and a single 3.2 TB SSD caching drive, both installed in the rear hot swappable slots, and a

240 GB M.2 form factor SSD that acts as the boot drive. Large form factor nodes cannot be configured with self-encrypting disks and are limited to a maximum of eight nodes in a cluster in the initial release of HyperFlex 3.0.

Figure 11 HX240c-M5L Node



Cisco VIC 1387 MLOM Interface Cards

The Cisco UCS VIC 1387 Card is a dual-port Enhanced Quad Small Form-Factor Pluggable (QSFP+) 40-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter installed in the Cisco UCS HX-Series Rack Servers. The VIC 1387 is used in conjunction with the Cisco UCS 6332 or 6332-16UP model Fabric Interconnects.

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, each dynamically configured as either a network interface card (NICs) or host bus adapter (HBA). The personality of the interfaces is set programmatically using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and Worldwide Name [WWN]), failover policy, adapter settings, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all specified using the service profile.

Figure 12 Cisco VIC 1387 mLOM Card



Hardware revision V03 or later of the Cisco VIC 1387 card is required for the Cisco HyperFlex HX-series servers.


Cisco VIC 1457 MLOM Interface Cards

The Cisco UCS VIC 1400 platform extends the network fabric directly to both servers and virtual machines so that a single connectivity mechanism can be used to connect both physical and virtual servers with the same level of visibility and control. Cisco VICs provide complete programmability of the Cisco UCS I/O infrastructure, with the number and type of I/O interfaces configurable on demand with a zero-touch model. Cisco VICs support Cisco SingleConnect technology, which provides an easy, intelligent, and efficient way to connect and manage computing in your data center. Cisco SingleConnect unifies LAN, SAN, and systems management into one

simplified link for rack servers, blade servers, and virtual machines. This technology reduces the number of network adapters, cables, and switches needed and radically simplifies the network, reducing complexity. Cisco VICs can support 256 PCIe virtual devices, either virtual NICs (vNICs) or virtual HBAs (vHBAs), a high rate of I/O operations per second (IOPS), support for lossless Ethernet, and 10- and 40-Gbps connection to servers. The PCIe 3.0 x16 interface helps ensure optimal bandwidth to the host for network-intensive applications, with a redundant path to the fabric interconnect. Cisco VICs support NIC teaming, with fabric failover for increased reliability and availability. In addition, it provides a policy-based, stateless, agile server infrastructure for your data center. The VIC 1400 series is designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers, C-Series Rack Servers, and S-Series Storage Servers. The adapters are capable of supporting 10-, 25, and 40-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE). The VIC incorporates Cisco's next-generation converged network adapter (CNA) technology and offers a comprehensive feature set, providing investment protection for future feature software releases. In addition, the VIC supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology. This technology extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

Cisco UCS VIC 1457 is based on the most recent generation of the Cisco UCS VIC 1400 platform. The Cisco UCS VIC 1457 (Figure 13) is a quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M5 generation of Cisco UCS C-Series Rack Servers. The card supports 10/25-Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.

Figure 13 Cisco VIC 1457 mLOM Card



All-Flash versus Hybrid

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

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

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

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

Cisco HyperFlex Compute-Only Nodes

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

- Cisco UCS B200 M4 Blade Server
- Cisco UCS B200 M5 Blade Server
- Cisco UCS C220 M5 Rack-Mount Servers
- Cisco UCS C240 M5 Rack-Mount Servers

Cisco HyperFlex Data Platform Software

The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology, exceeding the boundaries of first-generation hyperconverged infrastructures. The data platform has all the features expected in an enterprise shared storage system, eliminating the need to configure and maintain

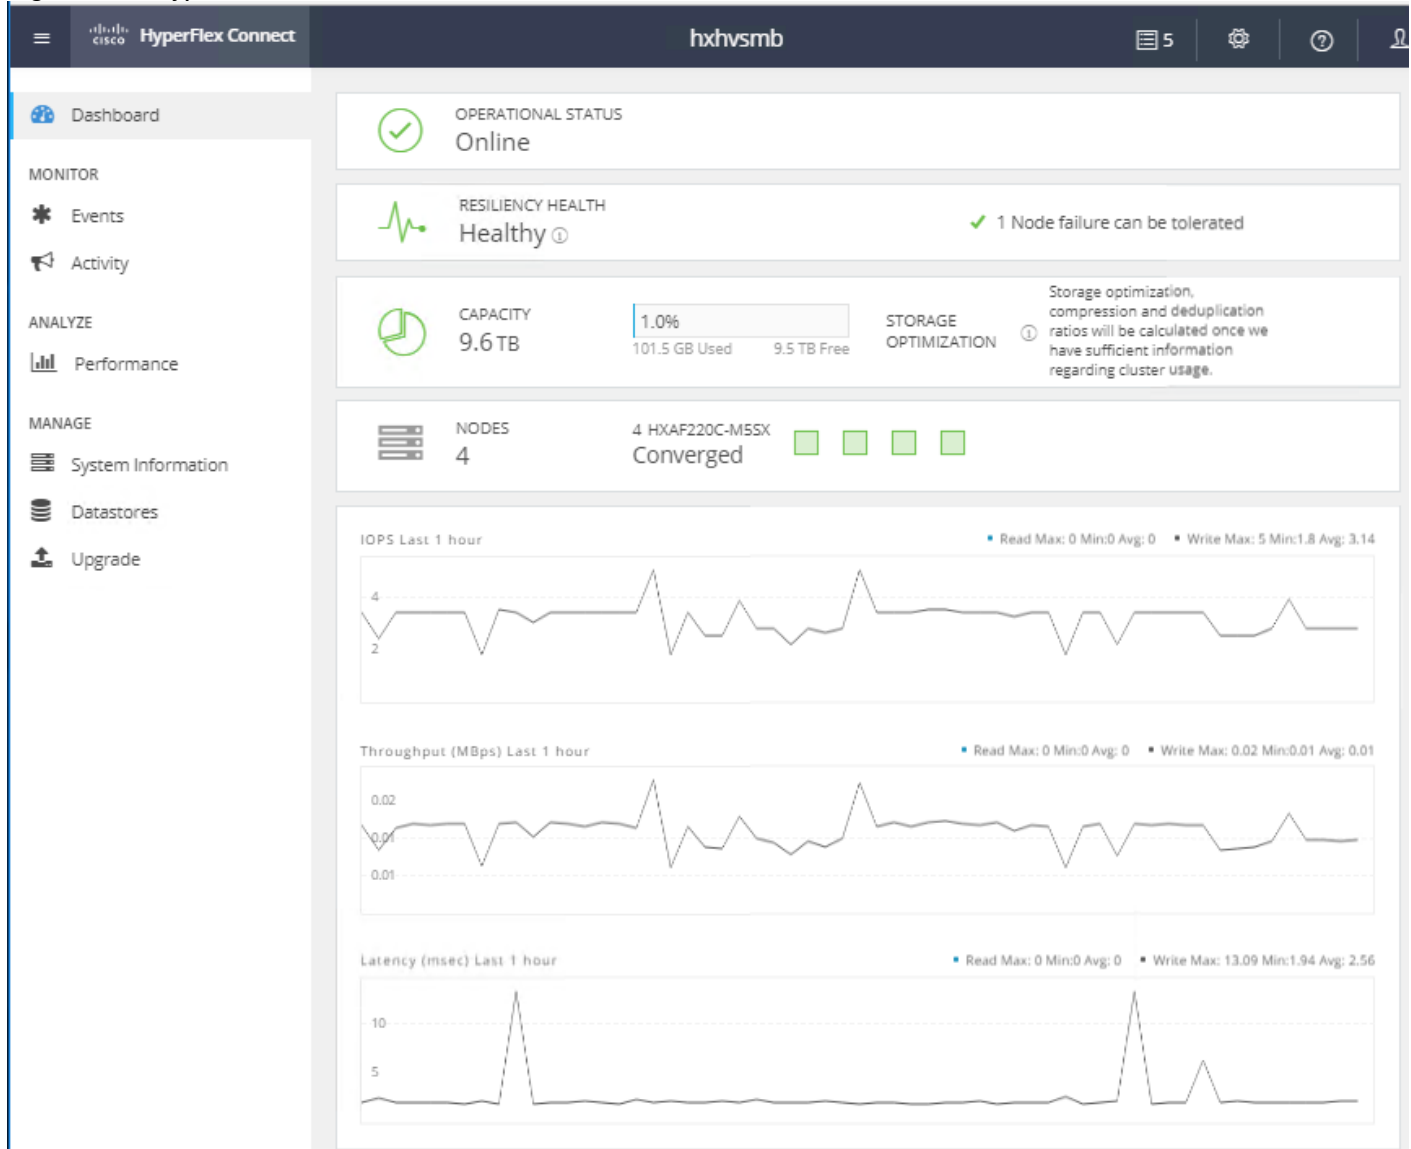
complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:

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

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>.

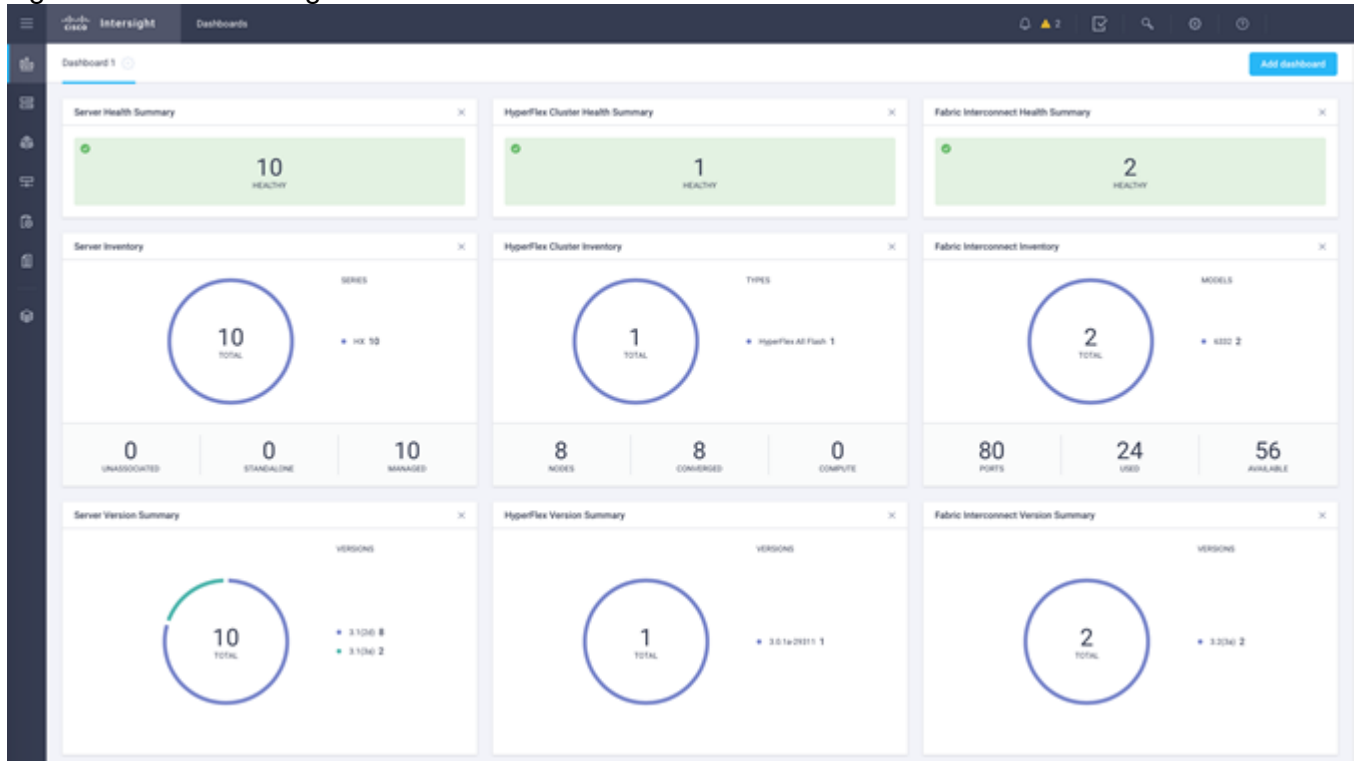
Figure 14 HyperFlex Connect GUI



Cisco Intersight Cloud Based Management

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

Figure 15 Cisco Intersight



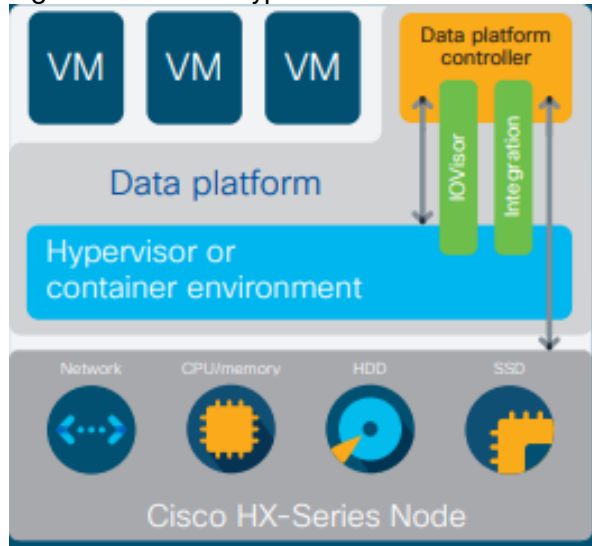
Cisco HyperFlex HX Data Platform Controller

A Cisco HyperFlex HX Data Platform controller resides on each node and implements the distributed file system. The controller runs as software in user space within a virtual machine, and intercepts and handles all I/O from the guest virtual machines. Dedicated CPU cores and memory allow the controller to deliver consistent performance without affecting performance of the other virtual machines in the cluster.

The data platform has modules to support the specific hypervisor or container platform in use. The controller accesses all of the node's disk storage through hypervisor bypass mechanisms (Discrete Device Assignment feature in Windows Server 2016 Hyper-V) for excellent performance. It uses the node's memory and dedicated SSD drives as part of a distributed caching layer, and it uses the node's HDDs, or SSD drives for distributed storage. The controller VM exposes the distributed storage as SMB file share to each Hyper-V node. The data platform controller interfaces with the hypervisor in two ways:

- IOvisor: The data platform controller intercepts all I/O requests and routes them to the nodes responsible for storing or retrieving the blocks. The IOvisor makes the existence of the hyperconvergence layer transparent to the hypervisor.
- Hypervisor agent: A module uses the hypervisor APIs to support advanced storage system operations such as snapshots and cloning. These are accessed through the hypervisor so that the hyperconvergence layer appears just as if it were enterprise shared storage. The controller accelerates operations by manipulating metadata rather than actual data copying, providing rapid response, and thus rapid deployment of new application environments.

Figure 16 Cisco HyperFlex Data Platform Controller Plugs into the Hypervisor in Each Node



Data Operations and Distribution

The Cisco HyperFlex HX Data Platform controllers handle all read and write operation requests from the guest virtual machines to their virtual disks (VHD/VHDX) stored in the distributed datastores in the cluster. The data platform distributes the data across multiple nodes of the cluster, and also across multiple capacity disks of each node, according to the replication level policy selected during the cluster setup. This method avoids storage hotspots on specific nodes, and on specific disks of the nodes, and thereby also avoids networking hotspots or congestion from accessing more data on some nodes versus others.

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 of 2 entire nodes in a cluster of 5 nodes or greater, without losing data and resorting to restore from backup or other recovery processes. RF3 is recommended for all production systems.
- 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 of 1 entire node without losing data and resorting to restore from backup or other recovery processes. RF2 is suitable for non-production systems, or environments where the extra data protection is not needed.

Data Write and Compression Operations

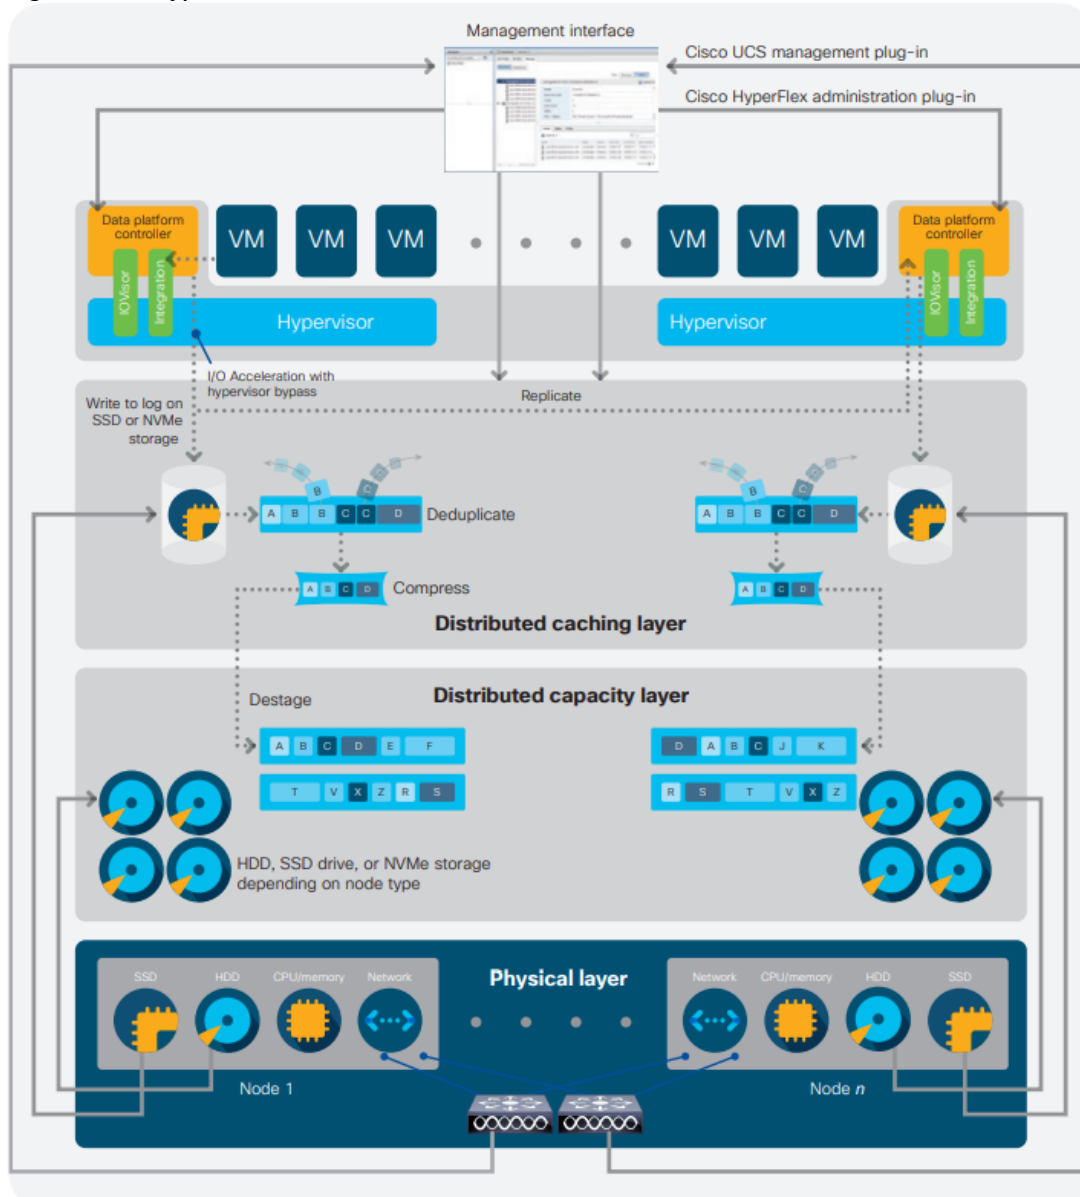
Internally, the contents of each virtual disk are subdivided and spread across multiple servers by the HXDP software. For each write operation, the data is intercepted by the IO Visor module on the node where the virtual machine is running, a primary node is determined for that particular operation via a hashing algorithm, and then sent to the primary node via the network. The primary node compresses the data in real time, writes the compressed data to its caching SSD, and replica copies of that compressed data are written to the caching SSD of the remote nodes in the cluster, according to the replication factor setting. For example, at RF=3 a write will be

written to the primary node for that virtual disk address, and two additional writes will be committed in parallel on two other nodes. Because the virtual disk contents have been divided and spread out via the hashing algorithm, this method results in all writes being spread across all nodes, avoiding the problems with data locality and “noisy” virtual machines consuming all the IO capacity of a single node. The write operation will not be acknowledged until all three copies are written to the caching layer SSDs. Written data is also cached in a write log area resident in memory in the controller virtual machine, along with the write log on the caching SSDs. This process speeds up read requests when reads are requested of data that has recently been written.

Data Destaging and Deduplication

The Cisco HyperFlex HX Data Platform constructs multiple write caching segments on the caching SSDs of each node in the distributed cluster. As write cache segments become full and based on policies accounting for I/O load and access patterns, those write cache segments are locked and new writes roll over to a new write cache segment. The data in the now locked cache segment is destaged to the HDD capacity layer of the nodes for the Hybrid system or to the SDD capacity layer of the nodes for the All-Flash system. During the destaging process, data is deduplicated before being written to the capacity storage layer, and the resulting data can now be written to the HDDs or SDDs of the server. On hybrid systems, the now deduplicated and compressed data is also written to the dedicated read cache area of the caching SSD, which speeds up read requests of data that has recently been written. When the data is destaged to an HDD, it is written in a single sequential operation, avoiding disk head seek thrashing on the spinning disks and accomplishing the task in the minimal amount of time. Since the data is already deduplicated and compressed before being written, the platform avoids additional I/O overhead often seen on competing systems, which must later do a read/dedupe/compress/write cycle. Deduplication, compression and destaging take place with no delays or I/O penalties to the guest virtual machines making requests to read or write data, which benefits both the HDD and SDD configurations.

Figure 17 HyperFlex HX Data Platform Data Movement



Data Read Operations

For data read operations, data may be read from multiple locations. For data that was very recently written, the data is likely to still exist in the write log of the local platform controller memory, or the write log of the local caching layer disk. If local write logs do not contain the data, the distributed filesystem metadata will be queried to see if the data is cached elsewhere, either in write logs of remote nodes, or in the dedicated read cache area of the local and remote caching SSDs of hybrid nodes. Finally, if the data has not been accessed in a significant amount of time, the filesystem will retrieve the requested data from the distributed capacity layer. As requests for reads are made to the distributed filesystem and the data is retrieved from the capacity layer, the caching SSDs of hybrid nodes populate their dedicated read cache area to speed up subsequent requests for the same data. This multi-tiered distributed system with several layers of caching techniques, ensures that data is served at the highest possible speed, leveraging the caching SSDs of the nodes fully and equally. All-flash configurations, however, do not employ a dedicated read cache because such caching does not provide any performance benefit; the persistent data copy already resides on high-performance SSDs.

In summary, the Cisco HyperFlex HX Data Platform implements a distributed, log-structured file system that performs data operations via two configurations:

- In a Hybrid configuration, the data platform provides a caching layer using SSDs to accelerate read requests and write responses, and it implements a storage capacity layer using HDDs.
- In an All-Flash configuration, the data platform provides a dedicated caching layer using high endurance SSDs to accelerate write responses, and it implements a storage capacity layer also using SSDs. Read requests are fulfilled directly from the capacity SSDs, as a dedicated read cache is not needed to accelerate read operations.

Solution Design

Requirements

The following sections detail the physical hardware, software revisions, and firmware versions required to install a single cluster of the Cisco HyperFlex system. A maximum of 8 converged nodes are supported on Cisco HyperFlex with Microsoft Hyper-V.

Physical Components

Table 1 HyperFlex System Components

Component	Hardware Required
Fabric Interconnects	Two Cisco UCS 6248UP Fabric Interconnects, or Two Cisco UCS 6296UP Fabric Interconnects, or Two Cisco UCS 6332 Fabric Interconnects, or Two Cisco UCS 6332-16UP Fabric Interconnects, or Two Cisco UCS 6454 Fabric Interconnects
Servers	Three to Sixteen Cisco HyperFlex HXAF220c-M5SX All-Flash rack servers, or Three to Sixteen Cisco HyperFlex HXAF240c-M5SX All-Flash rack servers, or Three to Sixteen Cisco HyperFlex HX220c-M5SX Hybrid rack servers, or Three to Sixteen Cisco HyperFlex HX240c-M5SX Hybrid rack servers, or Three to Sixteen Cisco HyperFlex HX240c-M5L Hybrid rack servers,

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

- [Compare Models](#)
- [HXAF220c-M5SX Spec Sheet](#)
- [HXAF240c-M5SX Spec Sheet](#)
- [HX220c-M5SX Spec Sheet](#)
- [HX240c-M5SX Spec Sheet](#)
- [HX240c-M5L Spec Sheet](#)

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

Table 2 HXAF220c-M5SX Server Options

HXAF220c-M5SX options	Hardware Required
-----------------------	-------------------

HXAF220c-M5SX options		Hardware Required
Processors		One or two Intel® Xeon® scalable family CPUs or one or two 2nd Generation Intel® Xeon® scalable family CPUs
Memory		192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4, 2666 MHz or 2933 MHz depending on CPU type, 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA
SSDs	Standard	<ul style="list-style-type: none"> • One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD • One 375GB 2.5-inch Intel Optane Drive, Extreme Perf & Endurance, or one 1.6TB 2.5 inch Enterprise performance 12G SAS SSD (3X endurance) Six to eight 3.8 TB 2.5 Inch Enterprise Value 6G SATA SSDs, or six to eight 960 GB 2.5 Inch Enterprise Value 6G SATA SSDs
Network		Cisco UCS VIC1387 VIC MLOM Cisco UCS VIC1457 VIC MLOM
Boot Device		One 240 GB M.2 form factor SATA SSD
microSD Card		One 32GB microSD card for local host utilities storage
Optional		Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

Table 3 lists the hardware component options for the HXAF240c-M5SX server model.

Table 3 HXAF240c-M5SX Server Options

HXAF240c-M5SX Options		Hardware Required
Processors		Chose a matching pair of Intel Xeon Processor Scalable Family CPUs
Memory		192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 MHz or 2933 MHz depending on CPU type, 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA
SSD	Standard	<ul style="list-style-type: none"> • One 240 GB 2.5 Inch Enterprise Value 6G SATA SSDOne 375GB 2.5-inch Intel Optane Drive, Extreme Perf & Endurance, or one 1.6TB 2.5 inch Enterprise performance 12G SAS SSD(3X endurance) • Six to twenty-three 3.8 TB 2.5 Inch Enterprise Value 6G SATA SSDs, or six to twenty-three 960 GB 2.5 Inch Enterprise Value 6G SATA SSDs
Network		Cisco UCS VIC1387 VIC MLOM Cisco UCS VIC1457 VIC MLOM

HXAF240c-M5SX Options	Hardware Required
Boot Device	One 240 GB M.2 form factor SATA SSD
microSD Card	One 32GB microSD card for local host utilities storage
Optional	Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

Table 4 lists the hardware component options for the HX220c-M5SX server model.

Table 4 HX220c-M5SX Server Options

HX220c-M5SX Options		Hardware Required
Processors		Chose a matching pair of Intel Xeon Processor Scalable Family CPUs
Memory		192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 MHz or 2933 MHz depending on CPU type, 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA
SSDs	Standard	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One 480 GB 2.5 Inch Enterprise Performance 6G SATA SSD, or one 800 GB 2.5 Inch Enterprise Performance 12G SAS SSD
HDDs	Standard	Six to eight 1.8 TB or 1.2 TB SAS 12Gbps 10K rpm SFF HDD
Network		Cisco UCS VIC1387 VIC MLOM Cisco UCS VIC1457 VIC MLOM
Boot Device		One 240 GB M.2 form factor SATA SSD
microSD Card		One 32GB microSD card for local host utilities storage
Optional		Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

Table 5 lists the hardware component options for the HX240c-M5SX server model.

Table 5 HX240c-M5SX Server Options

HX240c-M5SX Options		Hardware Required
Processors		Chose a matching pair of Intel Xeon Processor Scalable Family CPUs
Memory		192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 MHz or 2933 MHz depending on CPU type, 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA

HX240c-M5SX Options		Hardware Required
SSDs	Standard	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One 1.6 TB 2.5 Inch Enterprise Performance 12G SAS SSD
HDDs	Standard	Six to twenty-three 1.8 TB or 1.2 TB SAS 12Gbps 10K rpm SFF HDD
Network		Cisco UCS VIC1387 VIC MLOM Cisco UCS VIC1457 VIC MLOM
Boot Device		One 240 GB M.2 form factor SATA SSD
microSD Card		One 32GB microSD card for local host utilities storage
Optional		Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

Table 6 lists the hardware component options for the HX240c-M5L server model.

Table 6 HX240c-M5L Server Options

HX240c-M5L Options		Hardware Required
Processors		Chose a matching pair of Intel Xeon Processor Scalable Family CPUs
Memory		192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 MHz or 2933 MHz depending on CPU type, 1.2v modules
Disk Controller		Cisco 12Gbps Modular SAS HBA
SSDs	Standard	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One 3.2 TB 2.5 Inch Enterprise Performance 12G SAS SSD
HDDs	Standard	Six to twelve 6 TB or 8 TB or 12 TB SAS 7.2K rpm LFF HDD
Network		Cisco UCS VIC1387 VIC MLOM Cisco UCS VIC1457 VIC MLOM
Boot Device		One 240 GB M.2 form factor SATA SSD
microSD Card		One 32GB microSD card for local host utilities storage
Optional		Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

Software Components

Table 7 lists the software components and the versions required for the Cisco HyperFlex system for Microsoft Hyper-V.

Table 7 Software Components

Component	Software Required
Hypervisor	Hyper-V - Microsoft Windows Server 2016 Datacenter (Recommended UBR version above 1884), Or Hyper-V - Microsoft Windows Server 2019 Datacenter (Recommended UBR version above 107) Note: Microsoft Windows Server with Hyper-V will NOT be installed in Cisco Factory. Customers need to bring their own Windows Server ISO image that needs to be installed at deployment site
Active Directory	A Windows 2012 or later domain and forest functionality level with AD integrated DNS server.
Management Server	Windows 10 or Windows Server 2016 with PowerShell and RSAT tools installed. System Center VMM 2016 (optional) Windows Admin Center (Optional)
Cisco HyperFlex Data Platform	Cisco HyperFlex HX Data Platform Installer for Microsoft Hyper-V 4.0(1b) (Cisco-HX-Data-Platform-Installer-v4.0.1b-33133-hyperv.vhdx.zip)
HX Driver Image for System Preparation Script	File (name - 'latest.img') available inside the HXDP installer virtual machine
Ready Clone PowerShell Script	Cisco HyperFlex Data Platform Hyper-V ReadyClone PowerShell Script (HxClone-HyperV-v4.0.1b-33133.ps1)
Cisco UCS Firmware	Recommended Cisco UCS Infrastructure software, Cisco UCS B-Series and C-Series bundles, revisions 4.0(4d) for HXDP 4.0(1b) with 1 st and 2 nd Gen Intel Xeon Processors

Licensing

Cisco HyperFlex systems must be properly licensed using Cisco Smart Licensing, which is a cloud-based software licensing management solution used to automate many manual, time consuming and error prone licensing tasks. Cisco HyperFlex 2.5 and later communicate with the Cisco Smart Software Manager (CSSM) online service via a Cisco Smart Account, to check out or assign available licenses from the account to the Cisco HyperFlex cluster resources. Communications can be direct via the internet, they can be configured to communicate via a proxy server, or they can communicate with an internal Cisco Smart Software Manager satellite server, which caches and periodically synchronizes licensing data. In a small number of highly secure environments, systems can be provisioned with a Permanent License Reservation (PLR) which does not need to communicate with CSSM. Contact your Cisco sales representative or partner to discuss if your security requirements will necessitate use of these permanent licenses. New HyperFlex cluster installations will operate for

90 days without licensing as an evaluation period, thereafter the system will generate alarms and operate in a non-compliant mode. Systems without compliant licensing will not be entitled to technical support.

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



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

Table 8 lists the licensing editions and the features available with each type of license.

Table 8 HyperFlex System License Editions

HyperFlex Licensing Edition	Standard
Features Available	32 Converged Nodes standard cluster with Fabric Interconnects (16 converged + 16 compute-only) All Cisco UCS M5 server models Replication Factor 3 Compute-only nodes 10 GbE, 25GbE or 40 GbE Ethernet

Considerations

Version Control

The software revisions listed in Table 7 are the only valid and supported configuration at the time of the publishing of this validated design. Special care must be taken not to alter the revision of the hypervisor, Active Directory server, Cisco HX platform software, or the Cisco UCS firmware without first consulting the appropriate release notes and compatibility matrixes to ensure that the system is not being modified into an unsupported configuration.



The UBR # should be greater than 1884 for Windows Server 2016 and 107 for Windows Server 2019. If not, upgrade the Hyper-V servers to the latest version.



The remote management server should also have a version UBR # greater than 1884 and 107 for Windows Server 2019. You must upgrade the Hyper-V management server if the version is 1884 or lower.

To find out the Windows server version on the Hyper-V host for HyperFlex, run one of the following commands:

1. Type "winver" in the run command.



2. Open PowerShell and run the following command:

```
Get-ItemProperty 'HKLM:\SOFTWARE\Microsoft\Windows NT\CurrentVersion' | select
  ProductName,ReleaseID,CurrentBuild,CurrentBuildNumber,UBR
```

```
PS C:\Users\administrator.HXHVDOM2> Get-ItemProperty 'HKLM:\SOFTWARE\Microsoft\Windows NT\CurrentVersion'
| select ProductName,ReleaseID,CurrentBuild,CurrentBuildNumber,UBR

ProductName       : Windows Server 2019 Datacenter
ReleaseId        : 1809
CurrentBuild      : 17763
CurrentBuildNumber : 17763
UBR               : 107
```

Microsoft Windows Active Directory

The Microsoft Windows Active Directory 2012 or later is required due to the requirement of Cisco HyperFlex for Microsoft Hyper-V. The Active Directory with integrated DNS server must be installed and operational prior to the installation of the Cisco HyperFlex HX Data Platform software. The following best practice guidance applies to installations of HyperFlex 4.0:

- Do not modify the default TCP port settings. Using non-standard ports can lead to failures during the installation.
- It is recommended to build the Microsoft Active Directory Domain Controller and DNS servers on a physical server or in a virtual environment outside of the HyperFlex cluster. Building the Microsoft Active Directory Domain Controller and DNS servers as a virtual machine inside the HyperFlex cluster environment is highly discouraged as it creates cyclic dependency.
- Avoid creating single point of failure when you plan your virtual domain controller deployment
 - Run at least two virtualized domain controllers per domain on different virtualization hosts, which reduces the risk of losing all domain controllers if a single virtualization host fails.

- Maintain physical domain controllers in each of your domains. This mitigates the risk of a virtualization platform malfunction that affects all host systems that use that platform.

For best practices and guidance about virtualizing domain controllers on Hyper-V, see the following Microsoft articles:

<https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/get-started/virtual-dc/virtualized-domain-controllers-hyper-v>

<https://support.microsoft.com/en-in/help/888794/things-to-consider-when-you-host-active-directory-domain-controllers-i>

For best practices and guidance about virtualizing domain controllers on VMware vSphere environment, see the following VMware article:

https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/Virtualizing_Windows_Active_Directory.pdf



This document does not explain the installation and configuration of Microsoft Windows Active Directory and DNS server and assumes it to be available at the customer site up and running.

Scale

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



At the time of the publication of this document, Cisco HyperFlex for Microsoft Hyper-V does not support Cisco UCS M4 server models and SED drives.

Table 9 lists the Cisco HX Data platform storage cluster specifications for Microsoft Hyper-V.

Table 9 Cisco HX Data Platform Storage Cluster Specifications

HyperFlex with Hyper-V		
HX Servers	HX220c M5 HX240c M5 HX220c AF M5 HX240c AF M5	HX240c-M5L
Cisco UCS B-Series/C-Series Rack Servers	B200 M4, B200 M5	B200 M4, B200 M5
Supported Nodes	Converged and Compute-only nodes	Converged and Compute-only nodes
HXDP-S Licensed Node Limits	Converged nodes: 3-16	Converged nodes: 3-16

HyperFlex with Hyper-V		
1:1 ratio of HXDP-S to Compute only nodes (Min–Max)	Compute-only nodes: 0-16	Compute-only nodes: 0-16
HXDP-P Licensed Node Limits 1:2 ratio of HXDP-P to Compute only nodes (Min–Max)	Converged nodes: 3-16 Compute-only nodes: 0-16	Converged nodes: 3-16 Compute-only nodes: 0-16
Max Cluster Size	32	32
Max Compute to Converged ratio	1:01	1:01
Expansion	✓	✓



HyperFlex on Hyper-V supports a maximum of 32 nodes (16 Converged nodes + 16 Compute-Only nodes) in a cluster, it is recommended to create smaller clusters for the following reasons; creates smaller failure domains, a single big cluster is always a challenge to update/upgrade, allows room for expansion in the future, and easier to manage and operate.

Capacity

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

Disk drive manufacturers have adopted a size reporting methodology using calculation by powers of 10, also known as decimal prefix. As an example, a 120 GB disk is listed with a minimum of 120 x 10⁹ bytes of usable addressable capacity, or 120 billion bytes. However, many operating systems and filesystems report their space based on standard computer binary exponentiation, or calculation by powers of 2, also called binary prefix. In this example, 2¹⁰ or 1024 bytes make up a kilobyte, 2¹⁰ kilobytes make up a megabyte, 2¹⁰ megabytes make up a gigabyte, and 2¹⁰ gigabytes make up a terabyte. As the values increase, the disparity between the two systems of measurement and notation get worse, at the terabyte level, the deviation between a decimal prefix value and a binary prefix value is nearly 10 percent.

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

Table 10 SI Unit Values (Decimal Prefix)

Value	Symbol	Name
1000 bytes	kB	Kilobyte

Value	Symbol	Name
1000 kB	MB	Megabyte
1000 MB	GB	Gigabyte
1000 GB	TB	Terabyte

The [International Organization for Standardization](#) (ISO) and the [International Electrotechnical Commission](#) (IEC) defines values and binary prefix by powers of 2 in ISO/IEC 80000-13:2008 Clause 4 as follows:

Table 11 IEC Unit Values (binary prefix)

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

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

[Table 12](#) lists a set of HyperFlex HX Data Platform cluster usable capacity values, using binary prefix, for an array of cluster configurations. These values provide an example of the capacity calculations, for determining the appropriate size of HX cluster to initially purchase, and how much capacity can be gained by adding capacity disks. The calculations for these values are listed in [Appendix A: Cluster Capacity Calculations](#).

Table 12 Cluster Usable Capacities

HX-Series Server Model	Node Quantity	Capacity Disk Size (each)	Capacity Disk Quantity (per node)	Cluster Usable Capacity at RF=2	Cluster Usable Capacity at RF=3
HXAF220c-M5SX	8	3.8 TB	8	102.8 TiB	68.6 TiB
		960 GB	8	25.7 TiB	17.1 TiB
HXAF240c-M5SX	8	3.8 TB	6	77.1 TiB	51.4 TiB
			15	192.8 TiB	128.5 TiB
			23	295.7 TiB	197.1 TiB

HX-Series Server Model	Node Quantity	Capacity Disk Size (each)	Capacity Disk Quantity (per node)	Cluster Usable Capacity at RF=2	Cluster Usable Capacity at RF=3
		960 GB	6	19.3 TiB	12.9 TiB
			15	48.2 TiB	32.1 TiB
			23	73.9 TiB	49.3 TiB



Calculations will be based upon the number of nodes, the number of capacity disks per node, and the size of the capacity disks. The above table is not a comprehensive list of all capacities and models available.

Physical Topology

Topology Overview

The Cisco HyperFlex for Microsoft Hyper-V system is composed of a pair of Cisco UCS Fabric Interconnects along with up to sixteen HX-Series rack-mount servers as converged nodes per cluster. Up to sixteen compute-only servers can also be added per HyperFlex cluster. Adding Cisco UCS rack-mount servers and/or Cisco UCS 5108 Blade chassis, which house Cisco UCS blade servers, allows for additional compute resources in an extended cluster design. Up to eight separate HX clusters can be installed under a single pair of Fabric Interconnects. The two Fabric Interconnects both connect to every HX-Series rack-mount server, and both connect to every Cisco UCS 5108 blade chassis, and Cisco UCS rack-mount server. Upstream network connections, also referred to as “northbound” network connections are made from the Fabric Interconnects to the customer data center network at the time of installation.

Figure 18 HyperFlex Standard Cluster Topology

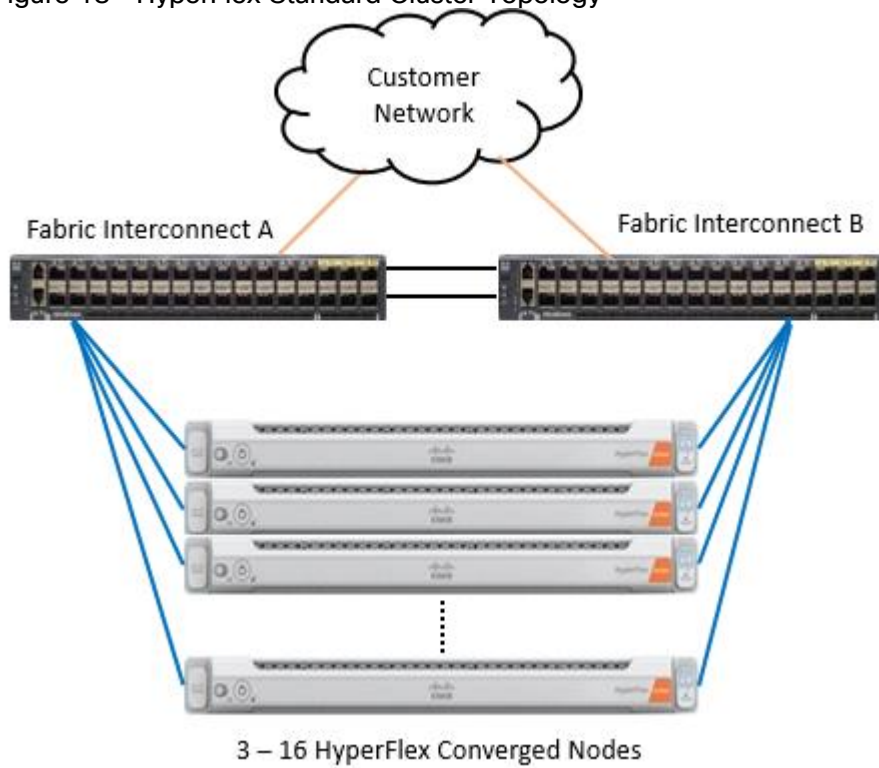
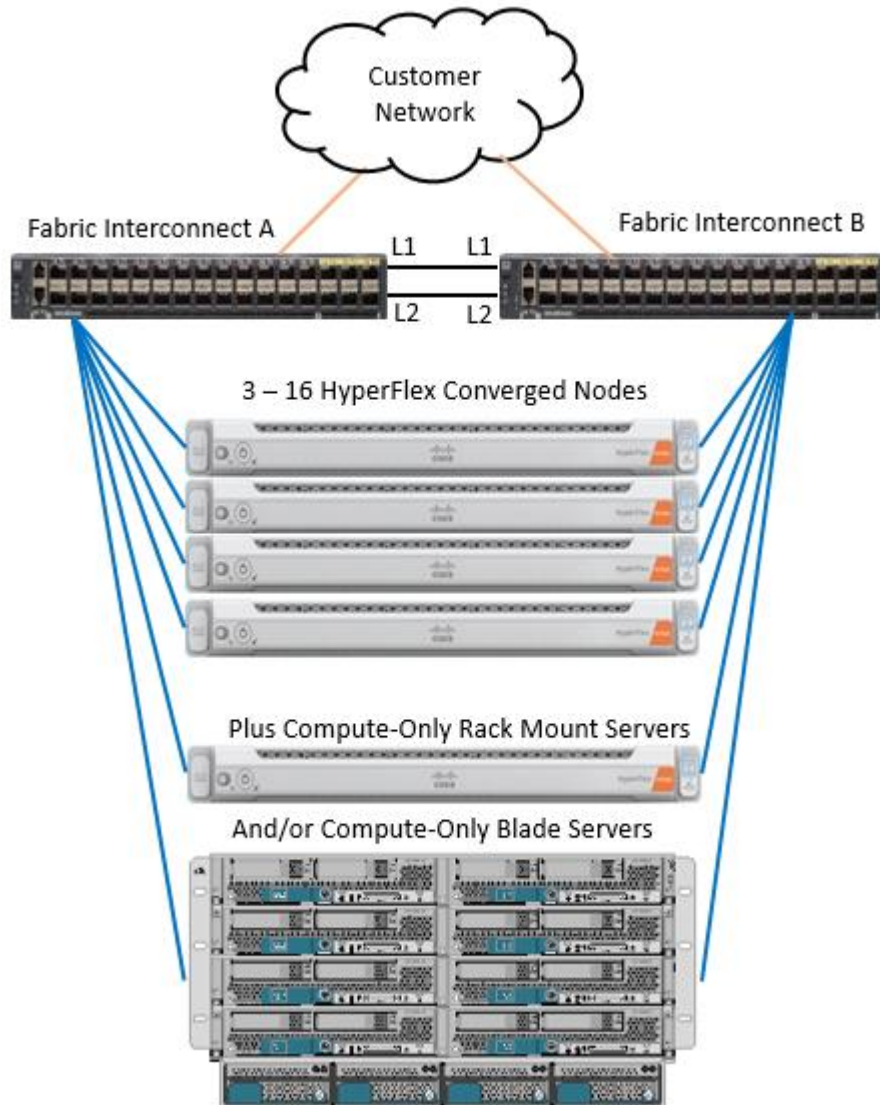


Figure 19 HyperFlex Extended Cluster Topology



Fabric Interconnects

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

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

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

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

HX-Series Rack-Mount Servers

The HX-Series converged servers are connected directly to the Cisco UCS Fabric Interconnects in Direct Connect mode. This option enables Cisco UCS Manager to manage the HX-Series Rack-Mount Servers using a single cable for both management traffic and data traffic. Cisco HyperFlex M5 generation servers can be configured only with the Cisco VIC 1387 card. The standard and redundant connection practice is to connect port 1 of the VIC card (the right-hand port) to a port on FI A, and port 2 of the VIC card (the left-hand port) to a port on FI B (Table 13). The HyperFlex installer checks for this configuration, and that all servers’ cabling matches. Failure to follow this cabling practice can lead to errors, discovery failures, and loss of redundant connectivity.

Various combinations of physical connectivity between the Cisco HX-series servers and the Fabric Interconnects are possible, but only specific combinations are supported. For example, use of the Cisco QSA module to convert a 40 GbE QSFP+ port into a 10 GbE SFP+ port is allowed for M5 generation servers in order to configure M5 generation servers along with model 6248 or 6296 Fabric Interconnects. Table 13 lists the possible connections, and which of these methods is supported.

Table 13 Supported Physical Connectivity

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

Figure 20 HX-Series Server Connectivity




Table 14 describes the topologies supported with Cisco UCS VIC 1400 Series and UCS Fabric Interconnects 6200, 6300 and 6400 Series:

Table 14 Cisco VIC 1400 Series Support for Cisco UCS Fabric Interconnects

M5 Cisco UCS VIC 1400 Series Connectivity	Cisco UCS VIC 1400 Series Adapters for both B-Series and C-Series		
	6400 Series	6300 Series	6200 Series
1 x 10G	Supported starting Release 3.5(2a)	Not Supported	Supported starting Release 3.5(1a)
2 x 10G	Supported starting Release 3.5(2a)	Not Supported	Supported starting Release 3.5(1a)
1 x 25G	Supported starting Release 3.5(2a)	Not Supported	Not Supported
2 x 25G	Supported starting Release 3.5(2a)	Not Supported	Not Supported

Table 15 Cisco UCS Fabric Interconnects Matrix


MLOM VIC	Interfaces	HX-FI-6454 or UCS-FI-6454
HX-MLOM-25Q-04 (VIC 1457)	2 or 4 ports, 10-Gbps Ethernet	10Gbps support, starting with Release 3.5(2a)
HX-MLOM-25Q-04 (VIC 1457)	2 or 4 ports, 25-Gbps Ethernet	25Gbps support, starting with Release 3.5(2a)
HX-MLOM-C40Q-03 (VIC 1387)	2 ports, 10-Gbps Ethernet (with QSA Adapter)	10Gbps support, starting with Release 3.5(2a)
HX-MLOM-C40Q-03 (VIC 1387)	2 ports, 40-Gbps Ethernet	Not Supported
HX-MLOM-CSC-02 (VIC 1227)	2 ports, 10-Gbps Ethernet	10Gbps support, starting with Release 3.5(2a)

For more information, refer to [Cisco HyperFlex Systems–Networking Topologies](#).

Cisco UCS B-Series Blade Servers

HyperFlex extended clusters also incorporate 1–16 Cisco UCS blade servers for additional compute capacity. 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 card installed in each blade server across the chassis backplane. The standard connection practice is to connect 1–8 10 GbE links, or 1–4 40 GbE links (depending on the IOMs and FIs purchased) from the left-side IOM, or IOM 1, to FI A, and to connect the same number of 10 GbE or 40 GbE links from the right-side IOM, or IOM 2, to FI B (Figure 21). All other cabling configurations are invalid, and can lead to errors, discovery failures, and loss of redundant connectivity.


Figure 21 Cisco UCS 5108 Chassis Connectivity



Cisco UCS C-Series Rack-Mount Servers

HyperFlex extended clusters also incorporate 1–32 Cisco UCS Rack-Mount Servers for additional compute capacity. The Cisco UCS C-Series Rack-Mount Servers are connected directly to the Cisco UCS Fabric Interconnects in Direct Connect mode. Internally the Cisco UCS C-Series servers are configured with the Cisco VIC 1227 or Cisco VIC 1387 network interface card (NIC) installed in a modular LAN on motherboard (MLOM) slot, which has dual 10 Gigabit Ethernet (GbE) ports or 40 Gigabit Ethernet (GbE) ports. The standard and redundant connection practice is to connect port 1 of the VIC card to a port on FI A, and port 2 of the VIC card to a port on FI B. Failure to follow this cabling practice can lead to errors, discovery failures, and loss of redundant connectivity.

Figure 22 Cisco UCS C-Series Server Connectivity



Logical Topology

Logical Network Design

The Cisco HyperFlex system has communication pathways that fall into four defined zones (Figure 23):

- 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, which answer via the FI management ports.
 - Hyper-V host management interfaces.
 - Storage Controller virtual machine management interfaces.
 - A roaming HX cluster management interface.
 - Storage Controller virtual machine Management interfaces.
- VM Zone: This zone comprises the connections needed to service network IO to the guest virtual machines that will run inside the HyperFlex hyperconverged system. This zone typically contains multiple VLANs that are trunked to the Cisco UCS Fabric Interconnects via the network uplinks and tagged with 802.1Q VLAN IDs. These interfaces and IP addresses need to be available to all staff and other computer endpoints which need to communicate with the guest virtual machines in the HX system, throughout the LAN/WAN.
- Storage Zone: This zone comprises the connections used by the Cisco HX Data Platform software, Hyper-V hosts, and the storage controller virtual machines to service the HX Distributed Data Filesystem. These interfaces and IP addresses need to be able to communicate with each other at all times for proper operation. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX storage traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B, and vice-versa. This zone is primarily jumbo

frame traffic therefore jumbo frames must be enabled on the Cisco UCS uplinks. In this zone are multiple components:

- A teamed interface is used for storage traffic on each Hyper-V host in the HX cluster.
- Storage Controller virtual machine storage interfaces.
- A roaming HX cluster storage interface.
- Live Migration Zone: This zone comprises the connections used by the Hyper-V hosts to enable live migration of the guest virtual machines from host to host. During normal operation, this traffic all occurs within the Cisco UCS domain, however there are hardware failure scenarios where this traffic would need to traverse the network northbound of the Cisco UCS domain. For that reason, the VLAN used for HX live migration traffic must be able to traverse the network uplinks from the Cisco UCS domain, reaching FI A from FI B, and vice-versa.

Figure 23 and Figure 24 illustrates the logical network design at the HX node and cluster level.

Figure 23 Logical Network Design - HX Node

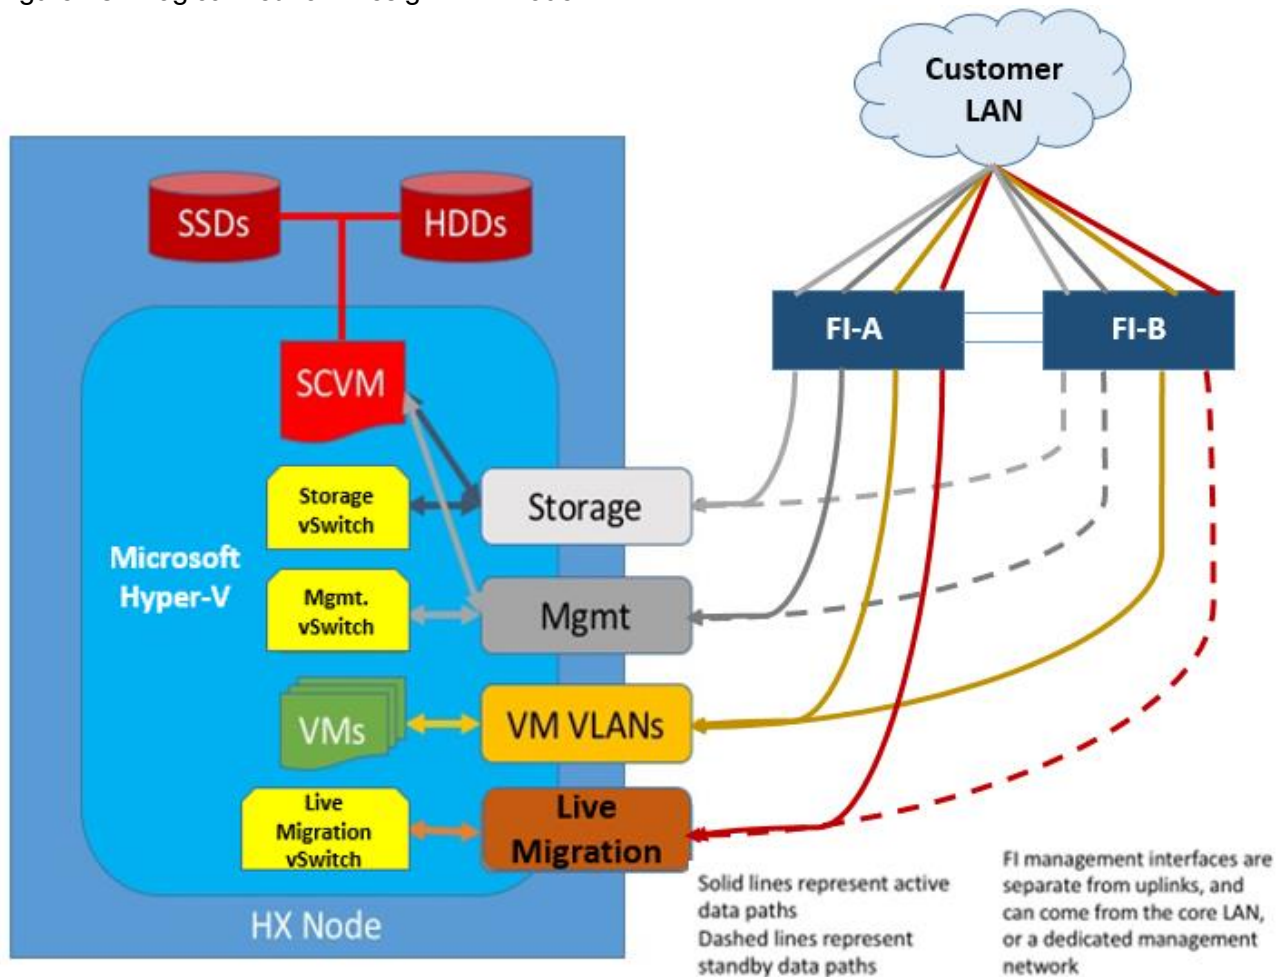
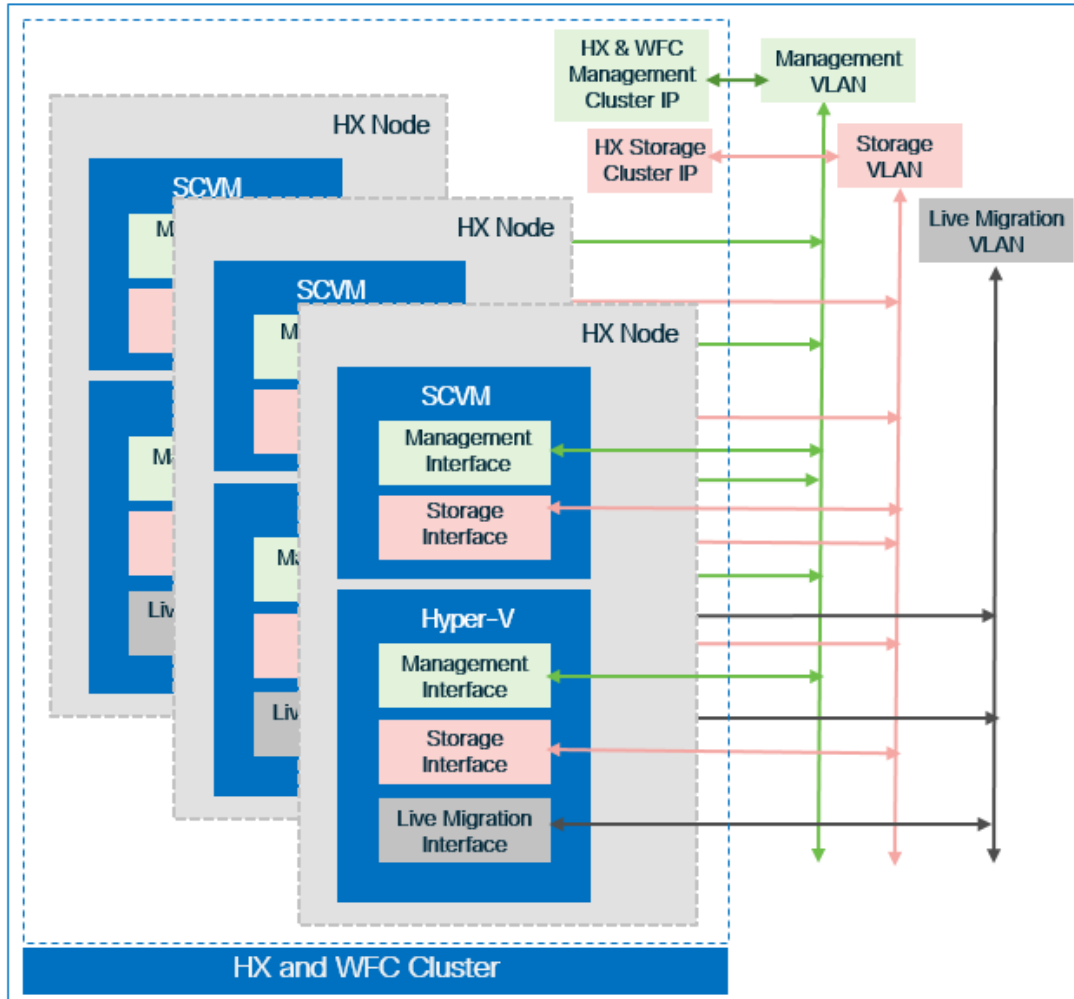


Figure 24 Logical Network Design - HX Cluster Node



Design Elements

Installing the HyperFlex system is primarily done through a deployable HyperFlex installer virtual machine, available for download at cisco.com as an OVA file. The installer virtual machine performs most of the Cisco UCS configuration work, it can be leveraged to simplify the installation of Windows Server 2016 on the HyperFlex hosts, and also performs significant portions of the configuration. Finally, the installer virtual machine is used to install the HyperFlex HX Data Platform software and create the HyperFlex cluster. Because this simplified installation method has been developed by Cisco, this CVD will not give detailed manual steps for the configuration of all the elements that are handled by the installer. Instead, the elements configured will be described and documented in this section, and the subsequent sections will guide you through the manual steps needed for installation, and how to utilize the HyperFlex Installer for the remaining configuration steps.

Network Design

Cisco UCS Uplink Connectivity

Cisco UCS network uplinks connect “northbound” from the pair of Cisco UCS Fabric Interconnects to the LAN in the customer datacenter. All Cisco UCS uplinks operate as trunks, carrying multiple 802.1Q VLAN IDs across the

uplinks. The default Cisco UCS behavior is to assume that all VLAN IDs defined in the Cisco UCS configuration are eligible to be trunked across all available uplinks.

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


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

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

vPC to Multiple Switches

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

Figure 25 Connectivity with vPC



VLANs and Subnets

For the base HyperFlex system configuration, multiple VLANs need to be carried to the Cisco UCS domain from the upstream LAN, and these VLANs are also defined in the Cisco UCS configuration. The hx-storage-data VLAN must be a separate VLAN ID from the remaining VLANs. Table 16 lists the VLANs created by the HyperFlex installer in Cisco UCS, and their functions:

Table 16 VLANs

VLAN Name	VLAN ID	Purpose
hx-inband-mgmt	Customer supplied	Hyper-V host management interfaces HX Storage Controller virtual machine management interfaces HX Storage Cluster roaming management interface
hx-storage-data	Customer supplied	Hyper-V host storage interfaces HX Storage Controller storage network interfaces HX Storage Cluster roaming storage interface
vm-network	Customer supplied	Guest virtual machine network interfaces
hx-livemigrate	Customer supplied	Hyper-V host live migration interfaces



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

Jumbo Frames

All HyperFlex storage traffic traversing the hx-storage-data VLAN and subnet is configured by default 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. In addition, the default MTU for the hx-livemigrate VLAN is also set to use jumbo frames. 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 configuration 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, including Cisco UCS firmware upgrades, or when a cable or port failure would cause storage traffic to traverse the northbound Cisco UCS uplink switches.

HyperFlex clusters can be configured to use standard size frames of 1500 bytes, however Cisco recommends that this configuration only be used in environments where the Cisco UCS uplink switches are not capable of passing jumbo frames, and that jumbo frames be enabled in all other situations.

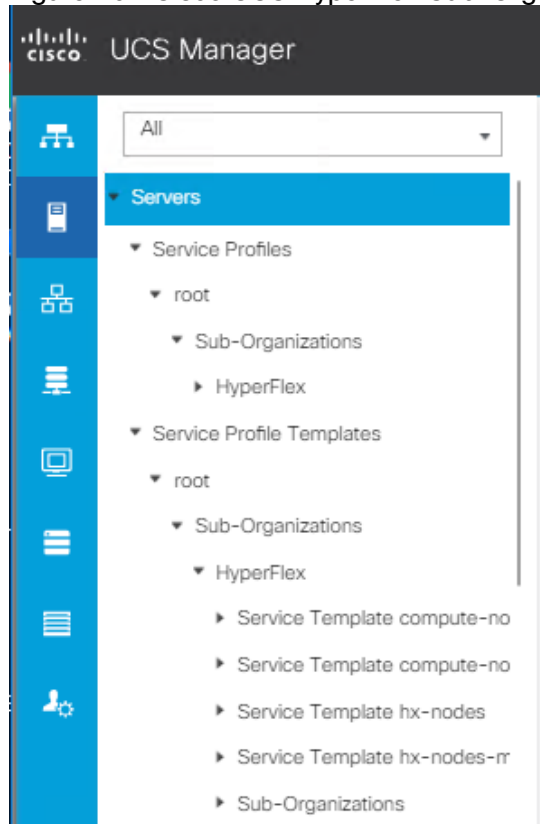
Cisco UCS Design

This section describes the elements within Cisco UCS Manager that are configured by the Cisco HyperFlex installer. Many of the configuration elements are fixed in nature, meanwhile the HyperFlex installer does allow for some items to be specified at the time of creation, for example VLAN names and IDs, external management IP pools and more. Where the elements can be manually set during the installation, those items will be noted in << >> brackets.

Cisco UCS Organization

During the HyperFlex installation a new Cisco UCS Sub-Organization is created. You must specify a unique Sub-Organization name for each cluster during the installation, for example “hx1hybrid”, or “hx2sed”. The sub-organization is created underneath the root level of the Cisco UCS hierarchy, and is used to contain all policies, pools, templates, and service profiles used by HyperFlex, which prevents problems from overlapping settings across policies and pools. This arrangement also allows for organizational control using Role-Based Access Control (RBAC) and administrative locales within Cisco UCS Manager 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 26 Cisco UCS HyperFlex Sub-Organization



Cisco UCS LAN Policies

QoS System Classes

Specific Cisco UCS Quality of Service (QoS) system classes are defined for a Cisco HyperFlex system. These classes define Class of Service (CoS) values that can be used by the uplink switches north of the Cisco UCS domain, plus which classes are active, along with whether packet drop is allowed, the relative weight of the different classes when there is contention, the maximum transmission unit (MTU) size, and if there is multicast optimization applied. QoS system classes are defined for the entire Cisco UCS domain, the classes that are enabled can later be used in QoS policies, which are then assigned to Cisco UCS vNICs. [Table 17](#) and [Figure 27](#) list the QoS System Class settings configured for HyperFlex:

Table 17 QoS System Classes

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

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

Figure 27 QoS System Classes

LAN / LAN Cloud / QoS System Class

General

Events

FSM

Actions

Use Global

Properties

Owner: **Local**

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



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

QoS Policies

In order to apply the settings defined in the Cisco UCS QoS System Classes, specific QoS Policies must be created, and then assigned to the vNICs, or vNIC templates used in Cisco UCS Service Profiles. [Table 18](#) details the QoS Policies configured for HyperFlex, and their default assignment to the vNIC templates created:

Table 18 HyperFlex QoS Policies

Policy	Priority	Burst	Rate	Host Control	Used by vNIC Template
Platinum	Platinum	10240	Line-rate	None	storage-data-a

Policy	Priority	Burst	Rate	Host Control	Used by vNIC Template
					storage-data-b
Gold	Gold	10240	Line-rate	None	vm-network-a vm-network-b
Silver	Silver	10240	Line-rate	None	hv-mgmt-a hv-mgmt-b
Bronze	Bronze	10240	Line-rate	None	hv-livemigrate-a hv- livemigrate -b
Best Effort	Best Effort	10240	Line-rate	None	N/A

Multicast Policy

A Cisco UCS Multicast Policy is configured by the HyperFlex installer, which is referenced by the VLANs that are created. The policy allows for future flexibility if a specific multicast policy needs to be created and applied to other VLANs that may be used by non-HyperFlex workloads in the Cisco UCS domain. [Table 19](#) and [Figure 28](#) details the Multicast Policy configured for HyperFlex:

Table 19 Multicast Policy

Name	IGMP Snooping State	IGMP Snooping Queries State
HyperFlex	Enabled	Disabled

Figure 28 Multicast Policy

[LAN](#) / [Policies](#) / [root](#) / [Multicast Policies](#) / [HyperFlex](#)

General

Events

Actions

[Delete](#)

[Use Global](#)

Properties

Name : **HyperFlex**

IGMP Snooping State : Enabled Disabled

IGMP Snooping Querier State : Enabled Disabled

Owner : **Local**

VLANs

VLANs are created by the HyperFlex installer to support a base HyperFlex system, with a VLAN for live migrate, and a single or multiple VLANs defined for guest virtual machine traffic. Names and IDs for the VLANs are defined in the Cisco UCS configuration page of the HyperFlex installer web interface. The VLANs listed in Cisco UCS must already be present on the upstream network, and the Cisco UCS FIs do not participate in VLAN Trunk Protocol (VTP). [Table 20](#) and [Figure 29](#) list the VLANs configured for HyperFlex.

Table 20 Cisco UCS VLANs

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

Figure 29 Cisco UCS VLANs

[LAN](#) / [LAN Cloud](#) / [VLANs](#)

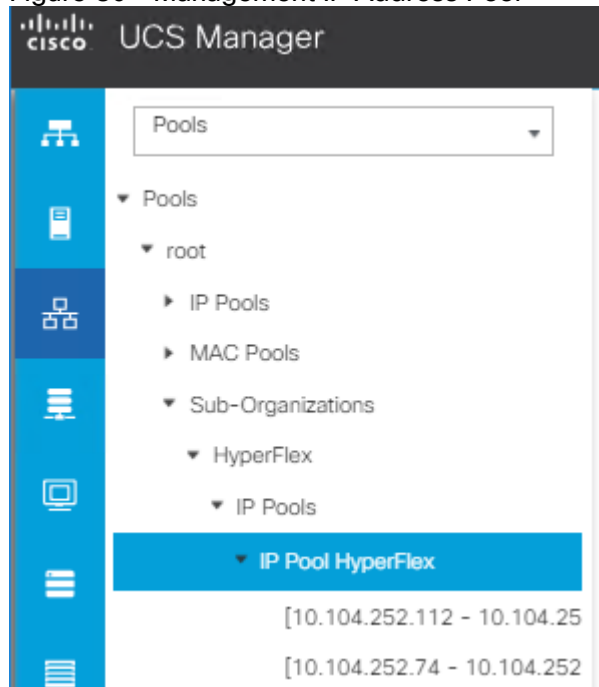
VLANs

Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN ...	Multicast Polic...
VLAN vm-network (3174)	3174	Lan	Ether	No	None		HyperFlex
VLAN hx-livemigrate (3173)	3173	Lan	Ether	No	None		HyperFlex
VLAN hx-storage-data (3172)	3172	Lan	Ether	No	None		HyperFlex
VLAN hx-inband-cimc (613)	613	Lan	Ether	No	None		HyperFlex
VLAN hx-inband-mgmt (613)	613	Lan	Ether	No	None		HyperFlex
VLAN default (1)	1	Lan	Ether	Yes	None		

Management IP Address Pool

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

Figure 30 Management IP Address Pool



MAC Address Pools

One of the core benefits of the Cisco UCS and Virtual Interface Card (VIC) technology is the assignment of the personality of the card via Cisco UCS Service Profiles. The number of virtual NIC (vNIC) interfaces, their VLAN associations, MAC addresses, QoS policies and more are all applied dynamically as part of the service profile association process. Media Access Control (MAC) addresses use 6 bytes of data as a unique address to identify the interface on the layer 2 network. All devices are assigned a unique MAC address, which is ultimately used for all data transmission and reception. The Cisco UCS and VIC technology picks a MAC address from a pool of addresses and assigns it to each vNIC defined in the service profile when that service profile is created.

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

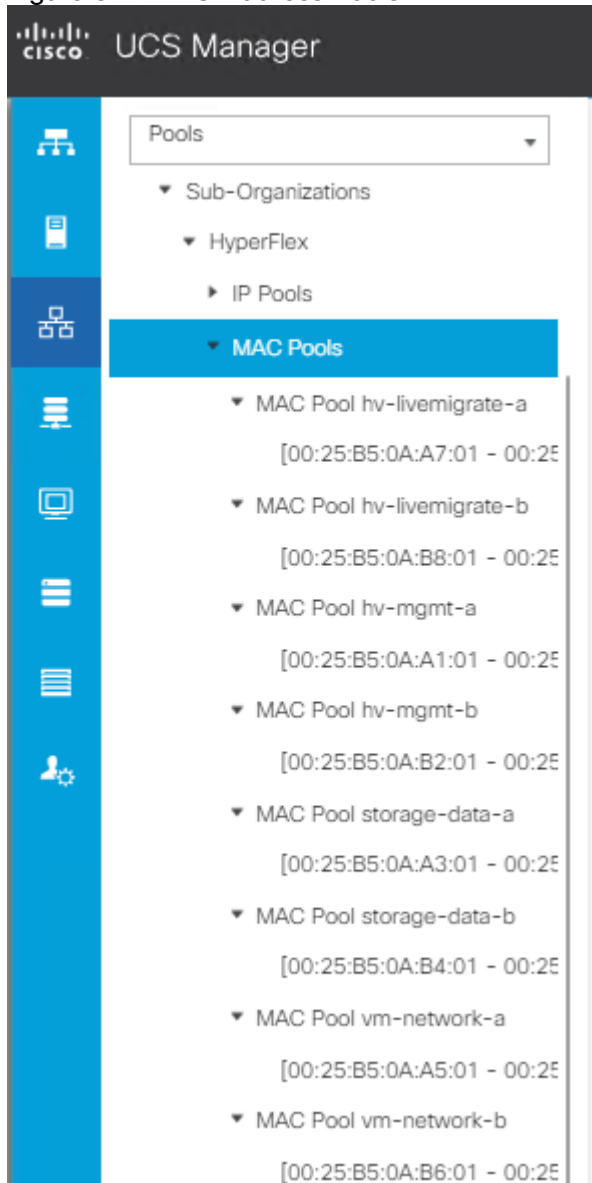
Table 21 lists the MAC Address Pools configured for HyperFlex and their default assignment to the vNIC templates created.

Table 21 MAC Address Pools

Name	Block Start	Size	Assignment Order	Used by vNIC Template
hv-mgmt-a	00:25:B5:<xx>:A1:01	100	Sequential	hv-mgmt-a
hv-mgmt-b	00:25:B5:<xx>:B2:01	100	Sequential	hv-mgmt-b

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

Figure 31 MAC Address Pools



Network Control Policies

Cisco UCS Network Control Policies control various aspects of the behavior of vNICs defined in the Cisco UCS Service Profiles. Settings controlled include enablement of Cisco Discovery Protocol (CDP), MAC address registration, MAC address forging, and the action taken on the vNIC status if the Cisco UCS network uplinks are failed. Two policies are configured by the HyperFlex Installer, HyperFlex-infra is applied to the “infrastructure” vNIC interfaces of the HyperFlex system, and HyperFlex-vm, which is only applied to the vNIC interfaces carrying guest virtual machine traffic. This allows for more flexibility, even though the policies are currently configured with the same settings. [Table 22](#) details the Network Control Policies configured for HyperFlex, and their default assignment to the vNIC templates created:

Table 22 Network Control Policy

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

Figure 32 Network Control Policy

Properties

Name : **HyperFlex-infra**

Description :

Owner : **Local**

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge: Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

vNIC Templates

Cisco UCS Manager has a feature to configure vNIC templates, which can be used to simplify and speed up configuration efforts. VNIC templates are referenced in service profiles and LAN connectivity policies, versus configuring the same vNICs individually in each service profile, or service profile template. VNIC templates contain

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

Table 23 vNIC Template hv-mgmt-a/b

vNIC Template Name:	hv-mgmt-a	hv-mgmt-b	storage-data-a	storage-data-b	hv-livemigrate-a	hv-livemigrate-b	vm-network-a	vm-network-b
Setting	Value	Value	Value	Value	Value	Value	Value	Value
Fabric ID	A	B	A	B	A	B	A	B
Fabric Failover	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled	Disabled
Target	Adapter	Adapter	Adapter	Adapter	Adapter	Adapter	Adapter	Adapter
Type	Updating Template	Updating Template	Updating Template	Updating Template	Updating Template	Updating Template	Updating Template	Updating Template
MTU	1500	1500	9000	9000	9000	9000	1500	1500
MAC Pool	hv-mgmt-a	hv-mgmt-b	storage-data-a	storage-data-b	hv-livemigrate-a	hv-livemigrate-b	vm-network-a	vm-network-b
QoS Policy	silver	silver	platinum	platinum	bronze	bronze	gold	gold
Network Control Policy	HyperFlex-infra	HyperFlex-infra	HyperFlex-infra	HyperFlex-infra	HyperFlex-infra	HyperFlex-infra	HyperFlex-vm	HyperFlex-vm
VLANs	<<hx-inband-mgmt>>	<<hx-inband-mgmt>>	<<hx-storage-data>>	<<hx-storage-data>>	<<hx-livemigrate>>	<<hx-livemigrate>>	<<vm-network>>	<<vm-network>>
Native VLAN	No	No	No	No	No	No	No	No
Connection Policies	Dynamic vNIC - Not Set	Dynamic vNIC - Not Set	VMQ - HyperFlex	VMQ - HyperFlex	VMQ - HyperFlex	VMQ - HyperFlex	VMQ - HyperFlex	VMQ - HyperFlex

LAN Connectivity Policies

Cisco UCS Manager has a feature for LAN Connectivity Policies, which aggregates all of the vNICs or vNIC templates desired for a service profile configuration into a single policy definition. This simplifies configuration efforts by defining a collection of vNICs or vNIC templates once and using that policy in the service profiles or service profile templates. The HyperFlex installer configures a LAN Connectivity Policy named HyperFlex, which contains all of the vNIC templates defined in the previous section, along with an Adapter Policy named HyperFlex, also configured by the HyperFlex installer. [Table 24](#) lists the LAN Connectivity Policy configured for HyperFlex:

Table 24 LAN Connectivity Policy

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

Cisco UCS Servers Policies

Adapter Policies

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

Figure 33 Cisco UCS Adapter Policy Resources

Resources	
Pooled	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Transmit Queues	: <input type="text" value="1"/> [1-1000]
Ring Size	: <input type="text" value="256"/> [64-4096]
Receive Queues	: <input type="text" value="4"/> [1-1000]
Ring Size	: <input type="text" value="512"/> [64-4096]
Completion Queues	: <input type="text" value="5"/> [1-2000]
Interrupts	: <input type="text" value="8"/> [1-1024]

Figure 34 Cisco UCS Adapter Policy Options

Options	
Transmit Checksum Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Receive Checksum Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
TCP Segmentation Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
TCP Large Receive Offload	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Receive Side Scaling (RSS)	: <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled
Accelerated Receive Flow Steering	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Network Virtualization using Generic Routing Encapsulation	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Virtual Extensible LAN	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Failback Timeout (Seconds)	: <input type="text" value="5"/> [0-600]
Interrupt Mode	: <input checked="" type="radio"/> MSI X <input type="radio"/> MSI <input type="radio"/> IN Tx
Interrupt Coalescing Type	: <input checked="" type="radio"/> Min <input type="radio"/> Idle
Interrupt Timer (us)	: <input type="text" value="125"/> [0-65535]
RoCE	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Advance Filter	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Interrupt Scaling	: <input checked="" type="radio"/> Disabled <input type="radio"/> Enabled

BIOS Policies

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

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

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

Boot Policies

Cisco UCS Boot Policies define the boot devices used by rack-mount servers, and the order that they are attempted to boot from. Cisco HX-Series M5 generation rack-mount servers have their Hyper-V hypervisors installed to an internal M.2 SSD boot drive, therefore they require a unique boot policy defining that the servers should boot from that location. The HyperFlex installer configures a boot policy named “HyperFlex-m5” specifying boot from the M.2 SSDs, referred to as “Embedded Disk” which is used by the HyperFlex M5 converged nodes, and should not be modified.

Figure 35 details the HyperFlex Boot Policies for Cisco HX-Series M5 generation rack-mount servers.

Figure 35 Cisco UCS M5 Boot Policy

Actions

[Delete](#)

[Show Policy Usage](#)

[Use Global](#)

Properties

Name : **HyperFlex-m5**

Description : Recommended boot policy for HyperFlex servers

Owner : **Local**

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Warning

The type (primary/secondary) does not indicate a boot order presence.
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If Enforce vNIC/vHBA/iSCSI Name is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
 If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

+ Local Devices

+ CIMC Mounted vMedia

+ vNICs

+ vHBAs

Boot Order

+ - ▼ Advanced Filter ↑ Export 🖨 Print

Name	Order ▲	vNIC/vH...	Type	LUN Name	WWN
CD/DVD	1				
Embedded Disk	2				

Host Firmware Packages

Cisco UCS Host Firmware Packages represent one of the most powerful features of the Cisco UCS platform; the ability to control the firmware revision of all the managed blades and rack-mount servers via a policy specified in the service profile. Host Firmware Packages are defined and referenced in the service profiles. Once a service profile is associated to a server, the firmware of all the components defined in the Host Firmware Package are automatically upgraded or downgraded to match the package. The HyperFlex installer creates a Host Firmware Packages named “HyperFlex-m5” which use the simple package definition method, applying firmware revisions to all components that matches a specific Cisco UCS firmware bundle, versus defining the firmware revisions part by part. Figure 36 shows the Host Firmware Package configured by the HyperFlex installer for Cisco HX-Series M5 generation rack-mount servers.

Figure 36 Cisco UCS M5 Host Firmware Package

Servers / Policies / root / Sub-Organizations / HXHV1 / Host Firmw... / HyperFlex-m5

General Events

Actions	Properties
Delete	Name : HyperFlex-m5
Show Policy Usage	Description : Recommended Host Firmware Packages for M5 Hyp
Use Global	Owner : Local
Modify Package Versions	Blade Package : 4.0(4d)B Blade Backup Package :
Modify Backup Package Versions	Rack Package : 4.0(4d)C Rack Backup Package :
	Service Pack :

Local Disk Configuration Policies

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

Figure 37 shows the Local Disk Configuration Policies configured by the HyperFlex installer.

Figure 37 Cisco UCS M5 Local Disk Configuration Policy

Properties

Name : **HyperFlex-m5**

Description : Recommended Local Disk policy for M5 HyperFlex s

Owner : **Local**

Mode : Any Configuration

Protect Configuration :

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State : Disable Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately. Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State: Disable Enable

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily. Please ensure SD cards are not in use before changing the FlexFlash Removable State.

FlexFlash Removable State : Yes No No Change

Maintenance Policies

Cisco UCS Maintenance Policies define the behavior of the attached blades and rack-mount servers when changes are made to the associated service profiles. The default Cisco UCS Maintenance Policy setting is “Immediate” meaning that any change to a service profile that requires a reboot of the physical server will result in an immediate reboot of that server. The Cisco best practice is to use a Maintenance Policy set to “user-ack”, which requires a secondary acknowledgement by a user with the appropriate rights within Cisco UCS Manager, before the server is rebooted to apply the changes. The HyperFlex installer creates a Maintenance Policy named “HyperFlex” with the setting changed to “user-ack”. In addition, the On Next Boot setting is enabled, which will automatically apply changes the next time the server is rebooted, without any secondary acknowledgement. [Figure 38](#) shows the Maintenance Policy configured by the HyperFlex installer:

Figure 38 Cisco UCS Maintenance Policy

Properties

Name : **HyperFlex**

Description : Recommended maintenance policy for HyperFlex se

Owner : **Local**

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy: Immediate User Ack

Reboot Policy : Immediate User Ack Timer Automatic

On Next Boot (Apply pending changes at next reboot.)

Power Control Policies

Cisco UCS Power Control Policies allow administrators to set priority values for power application to servers in environments where power supply may be limited, during times when the servers demand more power than is available. The HyperFlex installer creates a Power Control Policy named “HyperFlex” with all power capping

disabled, and fans allowed to run at full speed when necessary. [Figure 39](#) shows the Power Control Policy configured by the HyperFlex installer.

Figure 39 Cisco UCS Power Control Policy

Properties

Name : **HyperFlex**

Description : Recommended Power control policy for HyperFlex s

Owner : **Local**

Fan Speed Policy : Any

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its with 1 being the highest priority. If you choose **no-cap**, the server is exempt from

No Cap cap

Scrub Policies

Cisco UCS Scrub Policies are used to scrub, or erase data from local disks, BIOS settings and FlexFlash SD cards. If the policy settings are enabled, the information is wiped when the service profile using the policy is disassociated from the server. The HyperFlex installer creates a Scrub Policy named “HyperFlex” which has all settings disabled, therefore all data on local disks, SD cards and BIOS settings will be preserved if a service profile is disassociated. [Figure 40](#) shows the Scrub Policy configured by the HyperFlex installer.

Figure 40 Cisco UCS Scrub Policy

Properties

Name : **HyperFlex**

Description : Recommended Scrub policy for HyperFlex servers

Owner : **Local**

Disk Scrub : No Yes

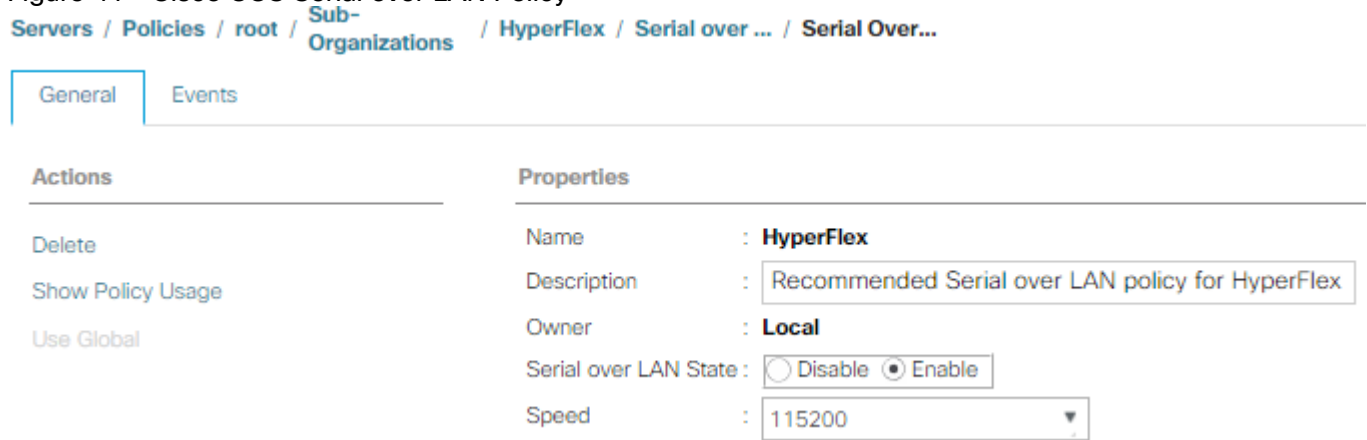
BIOS Settings Scrub : No Yes

FlexFlash Scrub : No Yes

Serial over LAN Policies

Cisco UCS Serial over LAN (SoL) Policies enable console output which is sent to the serial port of the server, to be accessible via the LAN. For many Linux based operating systems, such as VMware ESXi, the local serial port can be configured as a local console, where users can watch the system boot, and communicate with the system command prompt interactively. Since many blade servers do not have physical serial ports, and often administrators are working remotely, the ability to send and receive that traffic via the LAN is very helpful. Connections to a SoL session can be initiated from Cisco UCS Manager. The HyperFlex installer creates a SoL named “HyperFlex” to enable SoL sessions and uses this feature to configure the ESXi hosts’ management networking configuration. [Figure 41](#) shows the SoL Policy configured by the HyperFlex installer:

Figure 41 Cisco UCS Serial over LAN Policy



vMedia Policies

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

Cisco UCS Service Profile Templates

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

Table 25 Cisco UCS Service Profile Template Settings and Values

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

Service Profile Template Name:	hx-nodes-m5
Setting	Value
Firmware Policy	HyperFlex-m5
Power Control Policy	HyperFlex
Scrub Policy	HyperFlex
Serial over LAN Policy	HyperFlex
vMedia Policy	Not defined
Service Profile Template Name:	compute-nodes-m5
Setting	Value
UUID Pool	Hardware Default
Associated Server Pool	None
Maintenance Policy	HyperFlex
Management IP Address Policy	hx-ext-mgmt
Local Disk Configuration Policy	hx-compute-m5
LAN Connectivity Policy	HyperFlex
Boot Policy	hx-compute-m5
BIOS Policy	HyperFlex-m5
Firmware Policy	HyperFlex-m5
Power Control Policy	HyperFlex
Scrub Policy	HyperFlex
Serial over LAN Policy	HyperFlex
vMedia Policy	Not defined

Microsoft Hyper-V Host Design

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

Virtual Networking Design

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

The vSwitches created are:

- vswitch-hx-inband-mgmt: This vSwitch is created as part of the automated installation. As shown in the below figure, it is configured to use a teamed interface named “team-hx-inband-mgmt” with “hv-mgmt-a” and “hv-mgmt-b” as member adapters, without jumbo frames. The teaming and load balancing mode are configured for ‘Switch Independent’ and ‘Hyper-V Port’ respectively with “hv-mgmt-b” as Standby adapter.

The management interfaces of Hyper-V host and Storage Platform Controller virtual machines connect to this vSwitch. The VLANs are not Native VLANs as assigned to the vNIC templates, and therefore they are defined in Hyper-V virtual switch manager.

- vswitch-hx-storage-data: This vSwitch is created as part of the automated installation. As shown in the below figure, it is configured to use a teamed interface named “team-hx-storage-data” with “storage-data-a” and “storage-data-b” as member adapters, with jumbo frames highly recommended. The teaming and load balancing mode are configured for ‘Switch Independent’ and ‘Hyper-V Port’ respectively with “storage-data -b” adapter in Standby mode. The storage interfaces of Hyper-V host connected to this vSwitch is used for connecting to the HX Datastore via SMB. The VLANs are not Native VLANs as assigned to the vNIC templates, and therefore they are defined in Hyper-V virtual switch manager.
- vswitch-hx-vm-network: This vSwitch is created as part of the automated installation. As shown in the below figure, it is configured to use a teamed interface named “team-vm-network-data” with “vm-network-a” and “vm-network-b” as member adapters, without jumbo frames. The teaming and load balancing mode are configured for ‘Switch Independent’ and ‘Hyper-V Port’ respectively with both adapters in active mode. The VLANs are not Native VLANs as assigned to the vNIC templates, and therefore they are defined in Hyper-V virtual switch manager.
- vswitch-hx-livemigration: This vSwitch is created as part of the automated installation. As shown in the below figure, it is configured to use a teamed interface named “team-hx-livemigration” with “hv-livemigrate-a” and “hv-livemigrate-b” as member adapters, with jumbo frames highly recommended. The teaming and load balancing mode are configured for ‘Switch Independent’ and ‘Hyper-V Port’ respectively with “hv-livemigrate-b” adapter in Standby mode. The VLANs are not Native VLANs as assigned to the vNIC templates, and therefore they are defined in Hyper-V virtual switch manager. The IP addresses to this interface are assigned post installation.

Table 26 and Figure 42 provide more details into the Hyper-V virtual networking design as built by the HyperFlex installer by default for a converged node.

Table 26 Virtual Switches

Virtual Switch	Interfaces Connected	Active adapter	Passive adapter	VLAN IDs	Jumbo
vswitch-hx-inband-mgmt	Management Network Storage Controller Management Network	hv-mgmt-a	hv-mgmt-b	<<hx-inband-mgmt>>	no
vswitch-hx-storage-data	Storage Controller Data Network Storage Hypervisor Data Network	storage-data -a	storage-data -b	<<hx-storage-data>>	yes
vswitch-hx-vm-network	vm-network-<<VLAN ID>>	vm-network-a vm-network-b		<<vm-network>>	no
vswitch-hx-livemigration	livemigrate-<<VLAN ID>>	hv-livemigrate-a	hv-livemigrate-b	<<hx-livemigrate>>	yes

Figure 42 Hyper-V Network Design

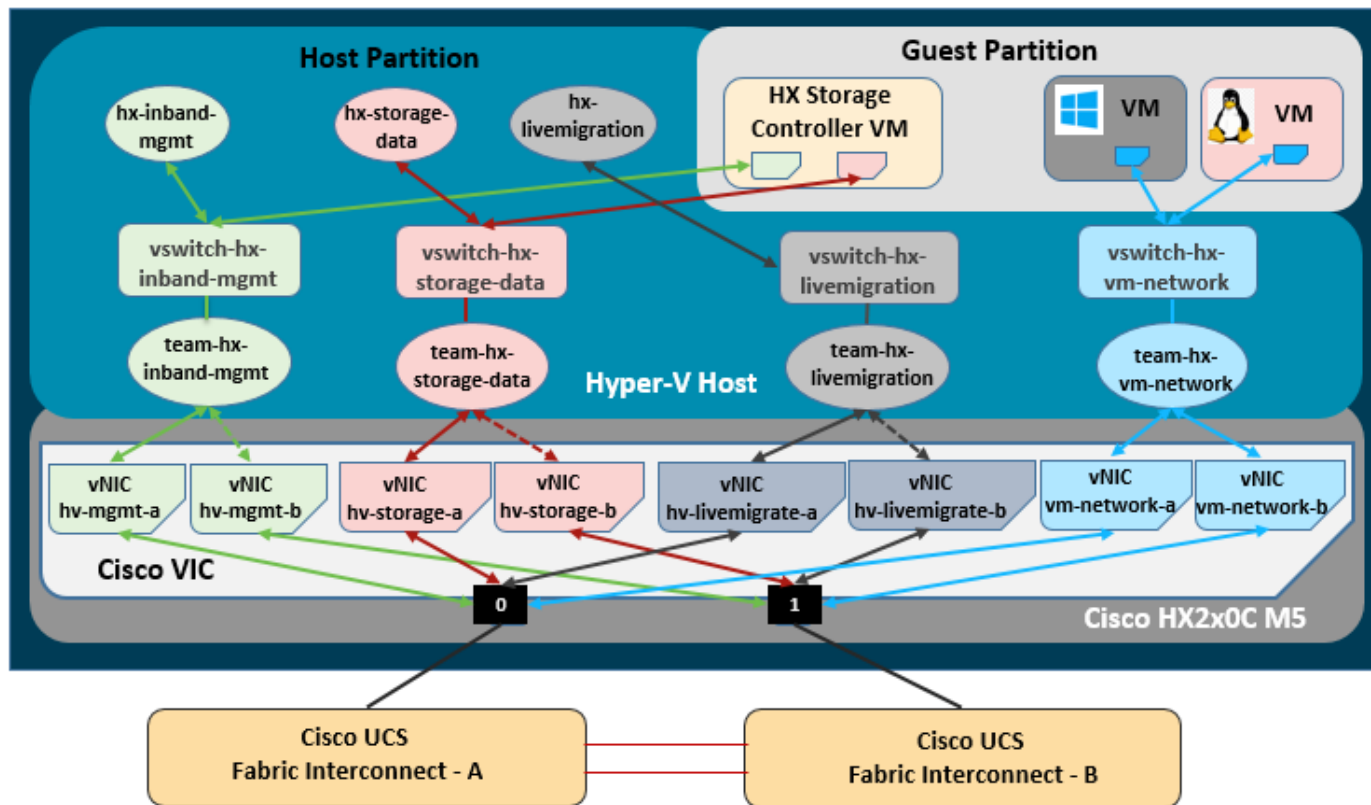
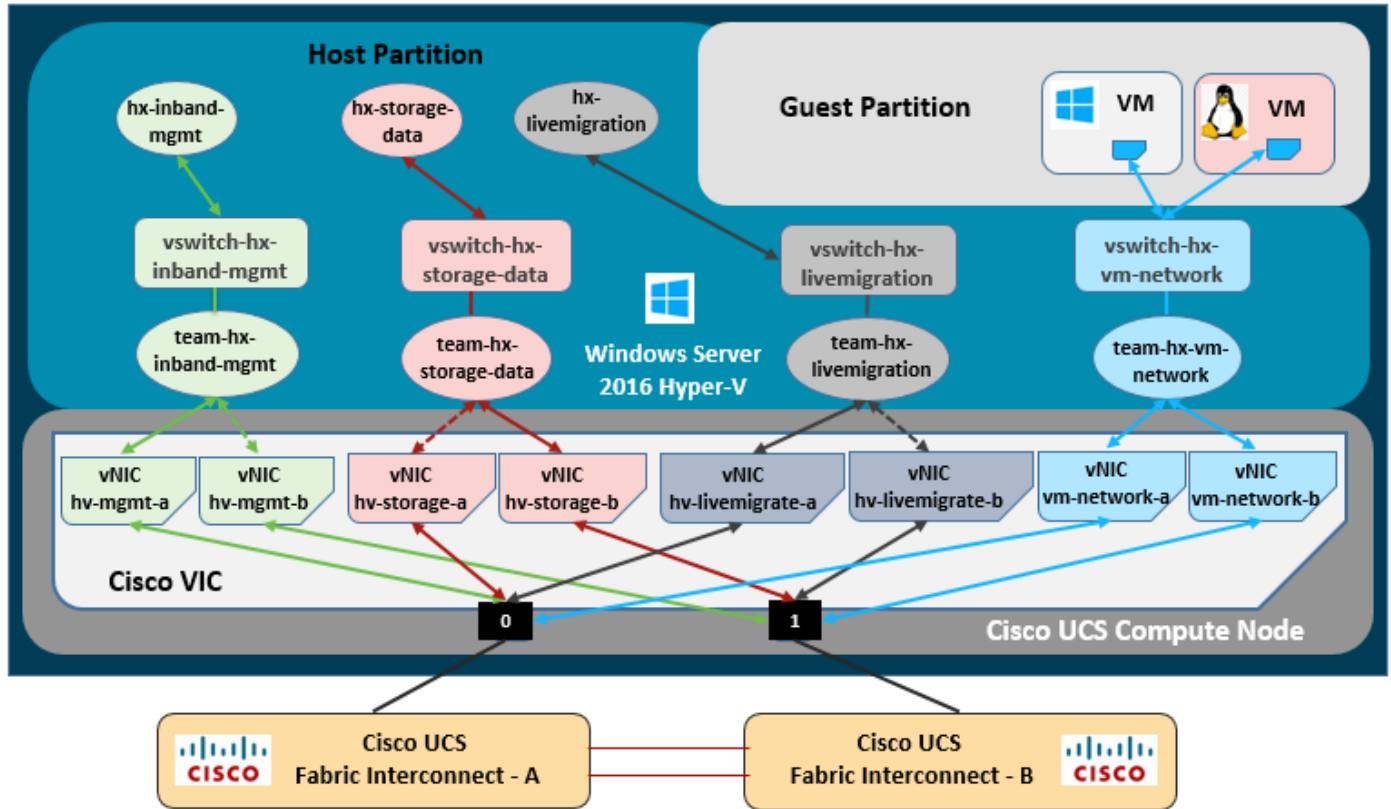


Figure 43 illustrates the Hyper-V virtual networking design as built by the HyperFlex installer by default for compute nodes. There is no storage controller virtual machine running on a compute node where the IOvisor is redirected and compute is assigned to a controller virtual machine running on a converged node.

Figure 43 Hyper-V Virtual Networking Design



Discrete Device Assignment (I/O Passthrough)

Discrete Device Assignment (DDA) is a feature introduced in Windows Server 2016 Hyper-V allowing access to an entire PCIe device into a virtual machine as though they were physical devices belonging to the virtual machine itself. With the appropriate driver for the hardware device, the guest virtual machine sends all I/O requests directly to the physical device, bypassing the hypervisor. In the Cisco HyperFlex system, the Storage Platform Controller virtual machines 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 virtual machines 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 are controlled by the controller virtual machines. Other disks, connected to different controllers, such as the M.2 drives, remain under the control of the Hyper-V hypervisor.



Configuring the DDA feature is done by the Cisco HyperFlex installer and requires no manual steps.

Storage Platform Controller Virtual Machines

A key component of the Cisco HyperFlex system is the Storage Platform Controller Virtual Machine running on each of the nodes in the HyperFlex cluster. The controller virtual machines cooperate to form and coordinate the Cisco HX Distributed Filesystem, and service all the guest virtual machine IO requests. The controller virtual machines deployed on the Hyper-V host are tied to a specific host, they start and stop along with the Hyper-V hypervisor, and the system is not considered to be online and ready until both the hypervisor and the Controller virtual machines have started.



Each Hyper-V host has a single Controller virtual machine deployed, and it cannot be moved or migrated to another host, nor should its settings be manually modified in any way.

The storage controller virtual machine runs custom software and services that manage and maintain the Cisco HX Distributed Filesystem. The services and processes that run within the controller virtual machines are not exposed, therefore the Hyper-V hosts have any direct knowledge of the storage services provided by the controller virtual machines. Management and visibility into the function of the controller virtual machines, and the Cisco HX Distributed Filesystem is done via the HyperFlex Connect HTML management webpage.



Deploying the controller virtual machines is done by the Cisco HyperFlex installer and requires no manual steps.

Controller Virtual Machine Locations

The physical storage location of the controller virtual machines across all the Cisco HX-Series M5 generation rack servers (HX220c M5, HXAF220c M5, HX240c M5 and HXAF240c M5) is same due to the physical disk location and connections on those server models. The storage controller virtual machine is operationally no different from any other typical virtual machines in a Hyper-V environment. The virtual machine must have a virtual disk with the bootable root filesystem available in a location separate from the SAS HBA that the virtual machine is controlling via DDA feature of Windows Server 2016/2019 Hyper-V.

- The server boots the Windows Server 2016/2019 Hyper-V from the internal M.2 form factor SSD.
- The Windows installer creates three partitions on the M.2 SSD - 500 MB for recovery partition, 90 GB for Windows OS boot partition and 30 GB partition where controller virtual machine's root filesystem is stored on a 2.5 GB virtual disk, /dev/sda. Both 90 GB and 30 GB partitions are formatted using NTFS file system. The controller virtual machine has full control of all the front and rear facing hot-swappable disks via DDA control of the SAS HBA. The controller virtual machine operating system sees the 240 GB SSD, also commonly called the "housekeeping" disk as /dev/sdb, and places HyperFlex binaries and logs on this disk. The remaining disks seen by the controller virtual machine OS are used by the HX Distributed filesystem for caching and capacity layers.

The following figures detail the Storage Platform Controller virtual machine placement on the Hyper-V hosts for Cisco HX M5 generation servers.

Figure 44 HX220c M5 Controller Virtual Machine Placement

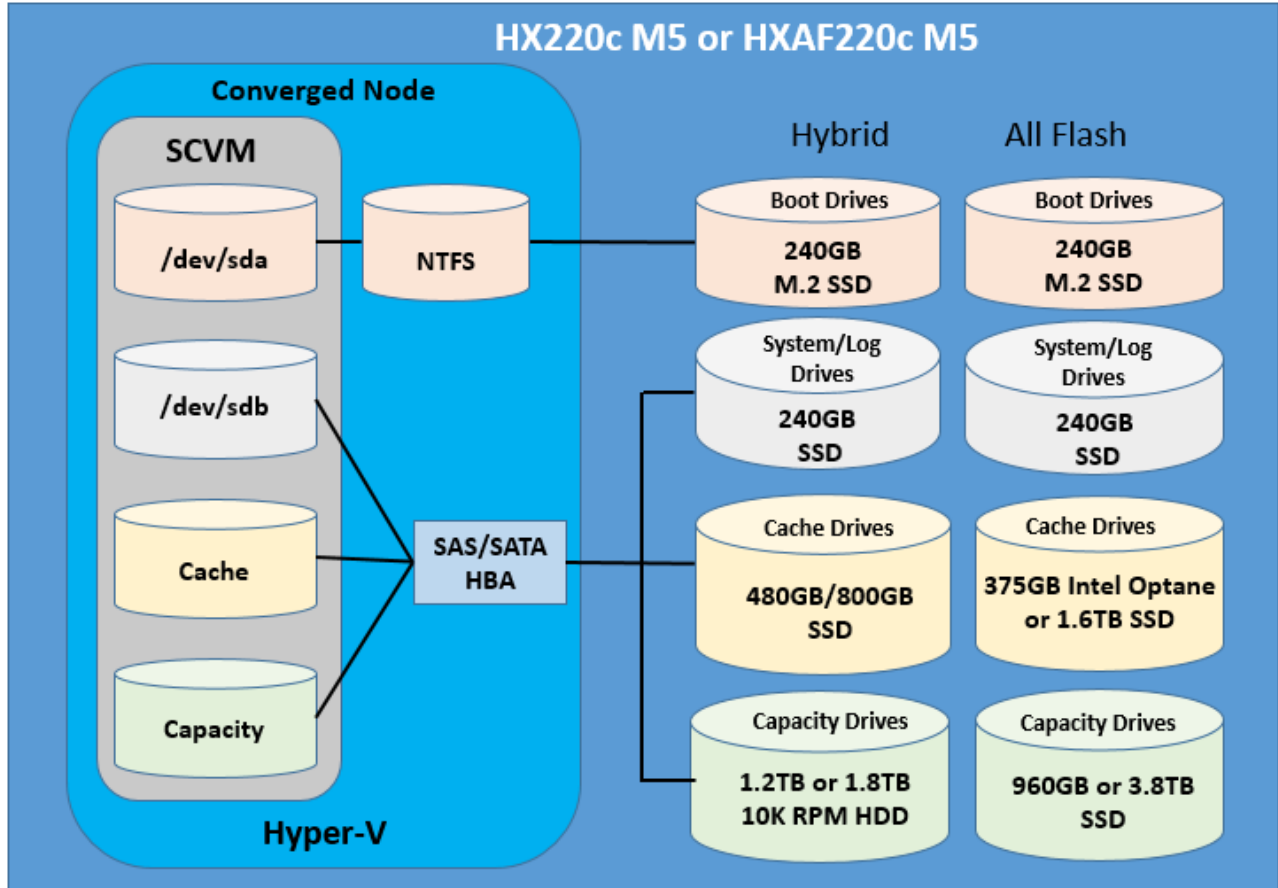
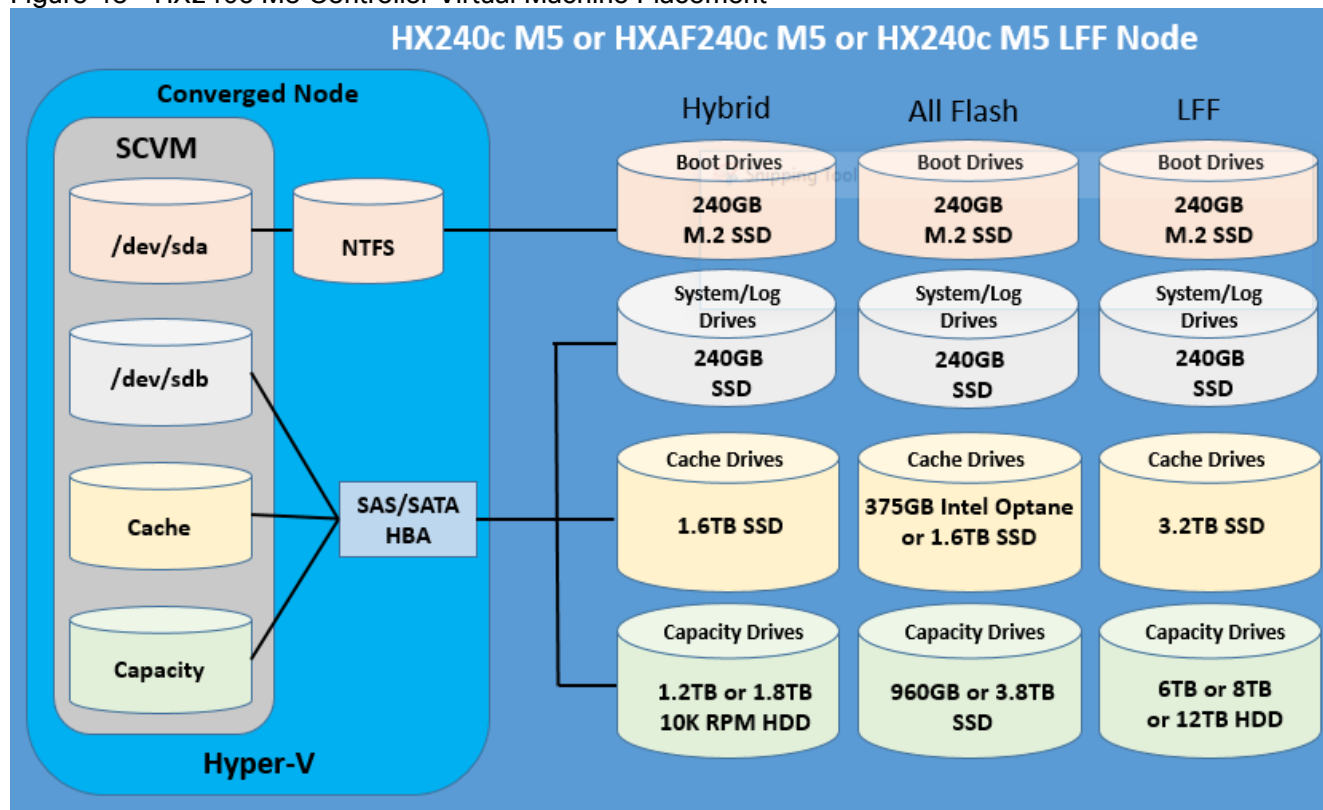



Figure 45 HX240c M5 Controller Virtual Machine Placement



HyperFlex Datastores

A 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 the HyperFlex Connect GUI. 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 HyperFlex Connect 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 46 Datastore Example



CPU Resource Control

Since the storage controller virtual machines provide critical functionality of the Cisco HX Distributed Data Platform, the HyperFlex installer will configure CPU resource control for the controller virtual machines. This resource control guarantees that the controller virtual machines will have CPU resources at a minimum level, in situations where the physical CPU resources of the Hyper-V hypervisor host are being heavily consumed by the guest virtual machines. [Table 27](#) details the CPU resource control settings of the storage controller virtual machines.

Table 27 Controller Virtual Machine CPU Reservations

Number of vCPU	Virtual Machine Reserve (percentage)	Virtual Machine Limit (percentage)	Relative Weight
10	100	100	100

Memory Resource Reservations

Since the storage controller virtual machines provide critical functionality of the Cisco HX Distributed Data Platform, the HyperFlex installer will configure by allocating static amount of memory for the controller virtual machines. This guarantees that the controller virtual machines will have access to all the memory resources allocated to it at power on. [Table 28](#) details the memory resource reservation of the storage controller virtual machines.

Table 28 Controller Virtual Machine Memory Reservations

Server Models	Amount of Static Memory for Virtual Machine
HX220c-M5SX HXAF220c-M5SX	48 GB
HX240c-M5SX HXAF240c-M5SX	72 GB
HX240c-M5L	78 GB

Installation

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

Prerequisites

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

IP Addressing

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

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

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

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

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



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

Table 29 HyperFlex Standard Cluster IP Addressing

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

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

Table 30 HyperFlex Extended Cluster IP Addressing


Address Group:	UCS Management	HyperFlex and Hyper-V Management	HyperFlex Storage	Live Migration
VLAN ID:				

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

Table 31 HyperFlex Standard Cluster Example IP Addressing for a 4-Node Cluster

Address Group:	UCS Management	HyperFlex and Hyper-V Management		HyperFlex Storage		Live Migration
VLAN ID:	121	613		3172		3173
Subnet:	10.65.121.0	10.104.252.0		192.168.11.0		192.168.73.0
Subnet Mask:	255.255.255.0	255.255.255.0		255.255.255.0		255.255.255.0
Gateway:	10.65.121.1	10.104.252.1				
Device	UCS Management Addresses	UCS Management Addresses	Hyper-V Management Interfaces	Hyper-V Hypervisor Storage Interfaces	Storage Controller Virtual Machine Storage Interfaces	Live Migration Interfaces
Fabric Interconnect A	10.65.121.241					

Address Group:	UCS Management	HyperFlex and Hyper-V Management	HyperFlex Storage	Live Migration		
Fabric Interconnect B	10.65.121.242					
UCS Manager	10.65.121.240					
HyperFlex Cluster		10.104.252.35	192.168.11.35			
Windows Failover Cluster		10.104.252.36				
HyperFlex Node #1	10.104.252.11	10.104.252.19	10.104.252.27	192.168.11.11	192.168.11.19	192.168.73.11
HyperFlex Node #2	10.104.252.12	10.104.252.20	10.104.252.28	192.168.11.12	192.168.11.20	192.168.73.12
HyperFlex Node #3	10.104.252.13	10.104.252.21	10.104.252.29	192.168.11.13	192.168.11.21	192.168.73.13
HyperFlex Node #4	10.104.252.14	10.104.252.22	10.104.252.30	192.168.11.14	192.168.11.22	192.168.73.14
HyperFlex Node #5	10.104.252.15	10.104.252.23	10.104.252.31	192.168.11.15	192.168.11.23	192.168.73.15
HyperFlex Node #6	10.104.252.16	10.104.252.24	10.104.252.32	192.168.11.16	192.168.11.24	192.168.73.16
HyperFlex Node #7	10.104.252.17	10.104.252.25	10.104.252.33	192.168.11.17	192.168.11.25	192.168.73.17
HyperFlex Node #8	10.104.252.18	10.104.252.26	10.104.252.34	192.168.11.18	192.168.11.26	192.168.73.18

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

DHCP versus Static IP

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

Configure the Active Directory for Constrained Delegation

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

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

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

- Hyper-V host
- HyperFlex Storage Controller virtual machine


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

To configure constrained delegation, follow these steps:

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

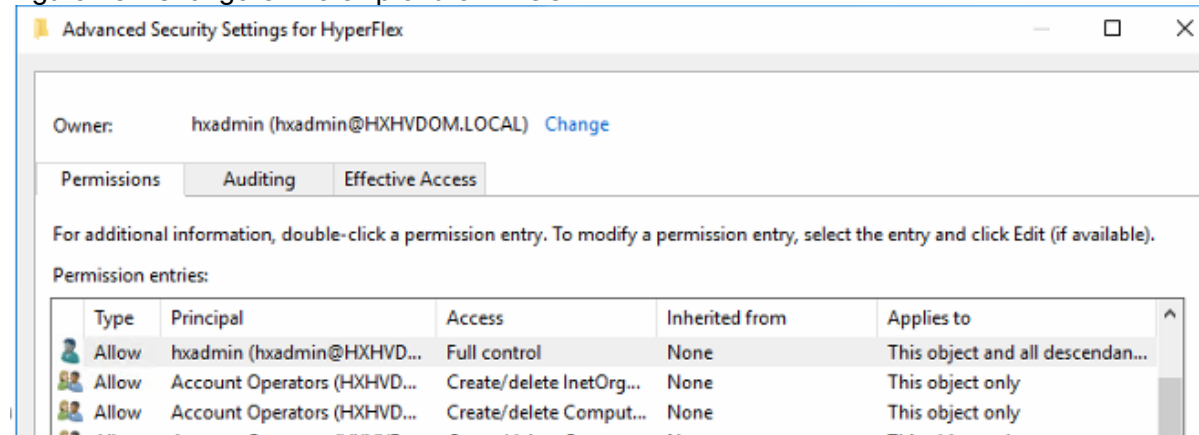
```
Get-ADOrganizationalUnit -Filter 'Name -like "Hyp*"' | Format-Table Name, DistinguishedName
```

Figure 47 PowerShell Get-ADOrganizationalUnit



5. Use Active Directory Users and Computers management tool to grant full permissions for the hxadmin user for the newly created OU. Make sure that Advanced features are enabled. If not, go back to Step 2.
 - a. Select the OU that you created. For example, HyperFlex > Properties > Security > Advance.
 - d. Click Change Owner and choose your hxadmin user.
 - e. Click Add in the Advanced view.
 - f. Select the principal and choose the hxadmin user. Choose Full Control and click OK.

Figure 48 Change Ownership of the AD OU





To configure the HX service account with the least privileges, refer Appendix section 'D' - Delegating HX service account with least privileges for administrative tasks. With this configuration, one-time domain admin credentials are required during the deployment of HyperFlex clusters

Prepopulate AD DNS with Records

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

To create DNS records, follow these steps:

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



Do not put in any DNS entry for SMB namespace. SMB namespace resolves differently on each host.



Standalone and non-Windows DNS servers are not supported.

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

Table 32 DNS Server Information

Item	Value
DNS Server #1	
AD DNS Domain	
UCS Domain Name	
HX Hyper-V Server #1	
HX Hyper-V Server #2	
HX Hyper-V Server #3	
HX Hyper-V Server #4	
HX Storage Controller VM #1	
HX Storage Controller VM #2	


Item	Value
HX Storage Controller VM #3	
HX Storage Controller VM #4	
Compute Node 1 Name	
Compute Node 2 Name	
HX Cluster CIP Name	
Windows Failover Cluster	

Table 33 DNS Server Example Information

Item	Value
DNS Server #1	Hxhv2dc.hxhvdome2.local
UCS Domain Name	HXHV-FI
AD DNS Domain	Hxhvdome2.local
HX Hyper-V Server #1	hxhv11. Hxhvdome2.local
HX Hyper-V Server #2	hxhv12. Hxhvdome2.local
HX Hyper-V Server #3	hxhv13. Hxhvdome2.local
HX Hyper-V Server #4	hxhv14. Hxhvdome2.local
HX Storage Controller VM #1	hxhv11scvm. Hxhvdome2.local
HX Storage Controller VM #2	hxhv12scvm. Hxhvdome2.local
HX Storage Controller VM #3	hxhv13scvm. Hxhvdome2.local
HX Storage Controller VM #4	hxhv14scvm. Hxhvdome2.local
Compute Node 1 Name	hxhv1co1. Hxhvdome2.local
HX Cluster CIP Name	hxhv1cip. Hxhvdome2.local
Windows Failover Cluster	hxhv1wfc. Hxhvdome2.local

Figure 49 shows pre-populated DNS with records for testing and validating of this HyperFlex document

Figure 49 DNS Manager with Pre-Populated DNS Records for HyperFlex



NTP

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

In an Active Directory domain, it is very important for all clocks to be within 5 minutes of each other (by default) due to the implementation of the Kerberos protocol for authentication. Also, Active Directory uses multi-master replication model between Domain Controllers.

Network Time Protocol (NTP) is the default time synchronization protocol used by the Windows Time Service (W32Time) in Windows servers and workstations. NTP uses UDP port 123 for all time synchronization communication and hence should be unblocked by the firewalls, in both directions.

In Active Directory deployment, the only computer configured with a time server explicitly should be computer holding the PDC Emulator FSMO role in the forest root domain. This is because the Forest root domain PDC emulator is the one and only one-time source for all the Domain Controllers, member servers and windows-based workstations for the entire forest.

All domain controllers in the forest root domain synchronize time with the PDC Emulator FSMO role-holder.

All Domain Controllers in child Domains synchronize time with any Domain Controller with Parent Domain or with PDC Emulator of its own Domain.

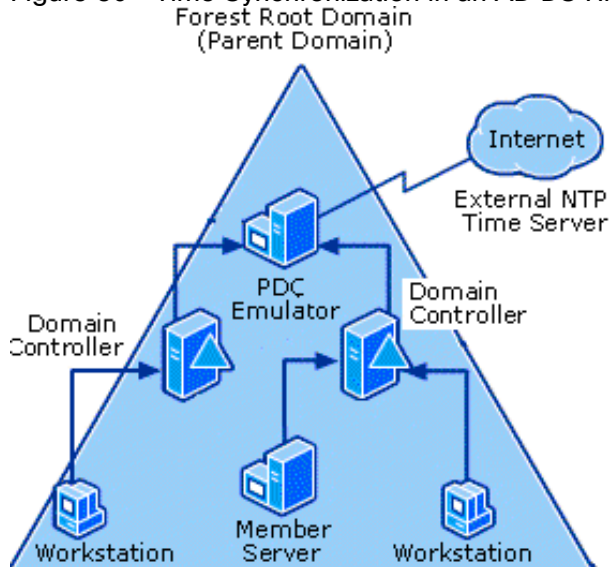
All PDC Emulator FSMO role-holders in child domains synchronize their time with domain controllers in their parent domain (including, potentially, the PDC Emulator FSMO role-holder in the forest root domain).

All domain member computers (Servers / Workstations/ any other devices) synchronize time with domain controller computers in their respective domains.

Additionally, virtual machines should not sync time with their host.

Figure 50 shows default time synchronization hierarchy in Active Directory Domain Services.

Figure 50 Time Synchronization in an AD DS Hierarchy



In a Windows domain, the time skew is set in Group Policy under Computer Configuration → Windows Settings → Account Policies → Kerberos Policy → Maximum tolerance for computer clock synchronization, and it is five minutes by default.

For more details, refer to the following Microsoft links:

Active Directory: Time Synchronization

<https://social.technet.microsoft.com/wiki/contents/articles/50924.active-directory-time-synchronization.aspx>

Windows Time Service Tools and Settings

<https://docs.microsoft.com/en-us/windows-server/networking/windows-time-service/Windows-Time-Service-Tools-and-Settings?redirectedfrom=MSDN>

VLANs

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

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

Table 34 VLAN Information

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

Table 35 VLAN Example Information

Name	ID
hx-inband-mgmt	3175
hx-storage-data	3172
vm-network	3174, 3175
hx-livemigrate	3173

Network Uplinks

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

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

Table 36 Network Uplink Configuration

Fabric Interconnect Port	Port Channel	Port Channel Type	Port Channel ID	Port Channel Name
--------------------------	--------------	-------------------	-----------------	-------------------

Fabric Interconnect Port		Port Channel	Port Channel Type	Port Channel ID	Port Channel Name
A		<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> LACP		
		<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> vPC		
		<input type="checkbox"/> Yes <input type="checkbox"/> No			
		<input type="checkbox"/> Yes <input type="checkbox"/> No			
B		<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> LACP		
		<input type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> vPC		
		<input type="checkbox"/> Yes <input type="checkbox"/> No			
		<input type="checkbox"/> Yes <input type="checkbox"/> No			

Table 37 Network Uplink Example Configuration

Fabric Interconnect Port		Port Channel	Port Channel Type	Port Channel ID	Port Channel Name
A	1/25	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> LACP	10	vpc-10
	1/26	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	<input checked="" type="checkbox"/> vPC		
		<input type="checkbox"/> Yes <input type="checkbox"/> No			
		<input type="checkbox"/> Yes <input type="checkbox"/> No			
B	1/25	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	<input type="checkbox"/> LACP	20	vpc-20
	1/26	<input checked="" type="checkbox"/> Yes <input type="checkbox"/> No	<input checked="" type="checkbox"/> vPC		
		<input type="checkbox"/> Yes <input type="checkbox"/> No			
		<input type="checkbox"/> Yes <input type="checkbox"/> No			

Username and Passwords

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

Table 38 Usernames and Passwords

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

Table 39 Example Usernames and Passwords

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

Physical Installation

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

- [Cisco UCS 6454 Series Fabric Interconnect Hardware Installation Guide](#)
- [Cisco UCS 6200 Series Fabric Interconnect Hardware Installation Guide](#)
- [Cisco UCS 6300 Series Fabric Interconnect Hardware Installation Guide](#)
- [Cisco HX220c M5 HyperFlex Node Installation Guide \(Hybrid and All-Flash Models\)](#)
- [Cisco HX240c M5 HyperFlex Node \(Hybrid and All-Flash Models\) Installation Guide](#)

Cabling

The physical layout of the HyperFlex system is described in the Physical Topology section. The Fabric Interconnects and HX-series rack-mount servers need to be cabled properly before beginning the installation activities. The information in this section is provided as a reference for cabling the physical equipment in this Cisco Validated Design environment.



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.

The following tables provide an example cabling map to install a Cisco HyperFlex system with four HyperFlex converged servers.

Table 40 Cisco Nexus 9396PX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9396PX-A	Eth1/15	10GbE	Cisco Nexus 9396PX-B	Eth1/15
	Eth1/16	10GbE	Cisco Nexus 9396PX-B	Eth1/16
	Eth2/5	40GbE	Cisco UCS fabric interconnect B	Eth1/50
	Eth2/6	40GbE	Cisco UCS fabric interconnect A	Eth1/49
	Eth1/31	10GbE	Infra-host-01	Port01
	MGMT0	GbE	GbE management switch	Any

Table 41 Cisco Nexus 9396PX-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9396PX-B	Eth1/15	10GbE	Cisco Nexus 9396PX-A	Eth1/15
	Eth1/16	10GbE	Cisco Nexus 9396PX-A	Eth1/16
	Eth2/5	40GbE	Cisco UCS fabric interconnect A	Eth1/50
	Eth2/6	40GbE	Cisco UCS fabric interconnect B	Eth1/49
	Eth1/31	10GbE	Infra-host-01	Port02
	MGMT0	GbE	GbE management switch	Any

Table 42 Cisco UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/49	40GbE	Cisco Nexus 9396PX-A	Eth2/6
	Eth1/50	40GbE	Cisco Nexus 9396PX-B	Eth2/5
	Eth1/1	25GbE	HX Server #1	mLOM Port 0
	Eth1/2	25GbE	HX Server #1	mLOM Port 1

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/3	25GbE	HX Server #2	mLOM Port 0
	Eth1/4	25GbE	HX Server #2	mLOM Port 1
	Eth1/5	25GbE	HX Server #3	mLOM Port 0
	Eth1/6	25GbE	HX Server #3	mLOM Port 1
	Eth1/7	25GbE	HX Server #4	mLOM Port 0
	Eth1/8	25GbE	HX Server #4	mLOM Port 1
	Eth1/9	25GbE	UCS C240 Server #1	mLOM Port 0
	Eth1/10	25GbE	UCS C240 Server #1	mLOM Port 1
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 43 Cisco UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/49	40GbE	Cisco Nexus 9396PX-B	Eth2/6
	Eth1/50	40GbE	Cisco Nexus 9396PX-A	Eth2/5
	Eth1/1	25GbE	HX Server #1	mLOM Port 0
	Eth1/2	25GbE	HX Server #1	mLOM Port 1
	Eth1/3	25GbE	HX Server #2	mLOM Port 0
	Eth1/4	25GbE	HX Server #2	mLOM Port 1
	Eth1/5	25GbE	HX Server #3	mLOM Port 0
	Eth1/6	25GbE	HX Server #3	mLOM Port 1
	Eth1/7	25GbE	HX Server #4	mLOM Port 0
	Eth1/8	25GbE	HX Server #4	mLOM Port 1
	Eth1/9	25GbE	UCS C240 Server #1	mLOM Port 0

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/10	25GbE	UCS C240 Server #1	mLOM Port 1
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

The following figures show a sample of direct connect mode physical connectivity for Cisco UCS C-Series Rack-Mount Server with Cisco UCS VIC 1457/1455.

Figure 51 Direct Connect Cabling Configuration with Cisco VIC 1400 Series (4-Port Linking)





Figure 52 Direct Connect Cabling Configuration with Cisco VIC 1400 Series (2-Port Linking)



The following restrictions apply: Ports 1 and 2 must connect to same Fabric Interconnect, for example Fabric-A. Ports 3 and 4 must connect to same Fabric Interconnect, for example Fabric-B. This is due to the internal port-channeling architecture inside the card. Ports 1 and 3 are used because the connections between ports 1 and 2 (also 3 and 4) form an internal port-channel



Do not connect port 1 to Fabric Interconnect A, and port 2 to Fabric Interconnect B. Use ports 1 and 3 only. Using ports 1 and 2 results in discovery and configuration failures.

Cisco UCS Installation

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

Cisco UCS Fabric Interconnect A

To configure Fabric Interconnect A, follow these steps:

1. Make sure the Fabric Interconnect cabling is properly connected, including the L1 and L2 cluster links, and power the Fabric Interconnects on by inserting the power cords.
2. Connect to the console port on the first Fabric Interconnect, which will be designated as the A fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.
3. Start your terminal emulator software.
4. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, and 1 stop bit.
5. Open the connection just created. You may have to press ENTER to see the first prompt.
6. Configure the first Fabric Interconnect, using the following example as a guideline:

```
---- Basic System Configuration Dialog ----
```

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

```
Type Ctrl-C at any time to abort configuration and reboot system.
```

```
To back track or make modifications to already entered values, complete input till
end of section and answer no when prompted to apply configuration.
```

```
Enter the configuration method. (console/gui) ? console
```

```
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
```

```
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
```

```
Enforce strong password? (y/n) [y]: y
```

```
Enter the password for "admin":
```

```
Confirm the password for "admin":
```

```
Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no)
[n]: yes
```

```
Enter the switch fabric (A/B) []: A
```

```
Enter the system name: HXHV-FI-A
```

```
Physical Switch Mgmt0 IP address : 10.29.149.203
```

Installation

```
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 10.29.149.1
Cluster IPv4 address : 10.29.149.205
Configure the DNS Server IP address? (yes/no) [n]: yes
  DNS IP address : 10.29.149.222
Configure the default domain name? (yes/no) [n]: yes
  Default domain name : hxxhvd.com.local
Join centralized management environment (UCS Central)? (yes/no) [n]: no
Following configurations will be applied:
  Switch Fabric=A
  System Name=HXHV-FI-A
  Enforced Strong Password=no
  Physical Switch Mgmt0 IP Address=10.29.149.203
  Physical Switch Mgmt0 IP Netmask=255.255.255.0
  Default Gateway=10.29.149.1
  Ipv6 value=0
  DNS Server=10.29.149.222
  Domain Name=hx.lab.cisco.com
  Cluster Enabled=yes
  Cluster IP Address=10.29.149.205
  NOTE: Cluster IP will be configured only after both Fabric Interconnects are ini-
tialized
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.
Configuration file - Ok
```

Cisco UCS Fabric Interconnect B

To configure Fabric Interconnect B, follow these steps:

1. Connect to the console port on the first Fabric Interconnect, which will be designated as the B fabric device. Use the supplied Cisco console cable (CAB-CONSOLE-RJ45=), and connect it to a built-in DB9 serial port, or use a USB to DB9 serial port adapter.
2. Start your terminal emulator software.
3. Create a connection to the COM port of the computer's DB9 port, or the USB to serial adapter. Set the terminal emulation to VT100, and the settings to 9600 baud, 8 data bits, no parity, and 1 stop bit.
4. Open the connection just created. You may have to press ENTER to see the first prompt.
5. Configure the second Fabric Interconnect, using the following example as a guideline:

```
---- Basic System Configuration Dialog ----
```

This setup utility will guide you through the basic configuration of the system. Only minimal configuration including IP connectivity to

Installation

the Fabric interconnect and its clustering mode is performed through these steps. Type Ctrl-C at any time to abort configuration and reboot system. To back track or make modifications to already entered values, complete input till end of section and answer no when prompted to apply configuration.

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

Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS Manager Cluster IP address, for exam-
ple <https://10.29.149.205>

Figure 53 Cisco UCS Manager



2. Click the “Launch UCS Manager” HTML link to open the Cisco UCS Manager web client.
3. At the login prompt, enter “admin” as the username, and enter the administrative password that was set during the initial console configuration.
4. Click No when prompted to enable Cisco Smart Call Home. This feature can be enabled at a later time.

Cisco UCS Configuration

Configure the following ports, settings, and policies in the Cisco UCS Manager interface prior to beginning the HyperFlex installation.

Cisco UCS Firmware

Your Cisco UCS firmware version should be correct as shipped from the factory, as documented in [Table 6](#) , that lists the hardware component options for the HX240c-M5L server model.

Table 44 HX240c-M5L Server Options

HX240c-M5L Options	Hardware Required
Processors	Chose a matching pair of Intel Xeon Processor Scalable Family CPUs
Memory	192 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 MHz 1.2v modules
Disk Controller	Cisco 12Gbps Modular SAS HBA
SSDs Standard	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD

HX240c-M5L Options	Hardware Required
HDDs Standard	One 3.2 TB 2.5 Inch Enterprise Performance 12G SAS SSD
Network	Cisco UCS VIC1457 VIC MLOM
Boot Device	One 240 GB M.2 form factor SATA SSD
microSD Card	One 32GB microSD card for local host utilities storage
Optional	Cisco QSA module to convert 40 GbE QSFP+ to 10 GbE SFP+

This document is based on Cisco UCS infrastructure, Cisco UCS B-series bundle, and Cisco UCS C-Series bundle software versions 4.0(4d). If the firmware version of the Fabric Interconnects is older than this version, the firmware must be upgraded to match the requirements prior to completing any further steps.

To upgrade the Cisco UCS Manager version, the Fabric Interconnect firmware, and the server bundles, refer to the [Cisco UCS Manager Firmware Management Guide, Release 4.0](#).

NTP

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

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

Uplink Ports

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

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

5. Select Fabric Interconnects > Fabric Interconnect B > Fixed Module or Expansion Module as appropriate > Ethernet Ports.
6. Select the ports that are to be uplink ports, right-click them, and click Configure as Uplink Port.
7. Click Yes to confirm the configuration and click OK.
8. Verify all the necessary ports are now configured as uplink ports, where their role is listed as "Network."

Figure 54 Uplinks Ports

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Stat...	Admin State
1	0	49	00:3A:9C:3...	Network	Physical	↑ Up	↑ Enabled
1	0	50	00:3A:9C:3...	Network	Physical	↑ Up	↑ Enabled

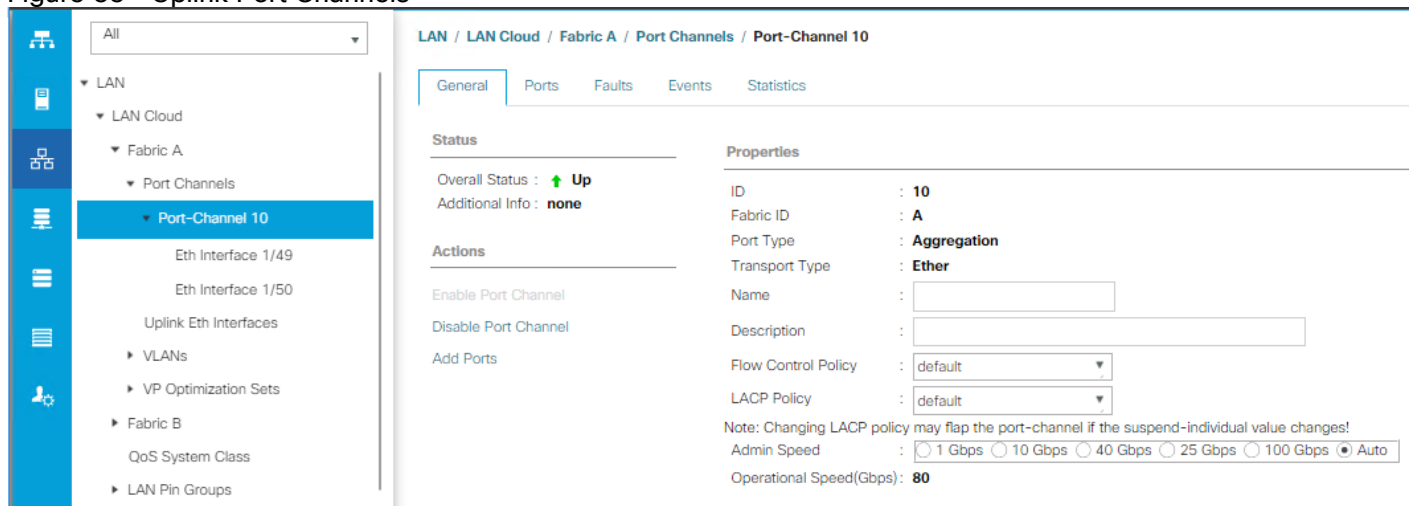
Uplink Port Channels

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

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

12. Enter the port channel ID number as the unique ID of the port channel (this does not have to match the port-channel ID on the upstream switch).
13. Enter the name of the port channel.
14. Click Next.
15. Click each port from Fabric Interconnect B that will participate in the port channel and click the >> button to add them to the port channel.
16. Click Finish.
17. Click OK.
18. Verify the necessary port channels have been created. It can take a few minutes for the newly formed port channels to converge and come online.

Figure 55 Uplink Port Channels



Server Ports

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

Auto Configuration

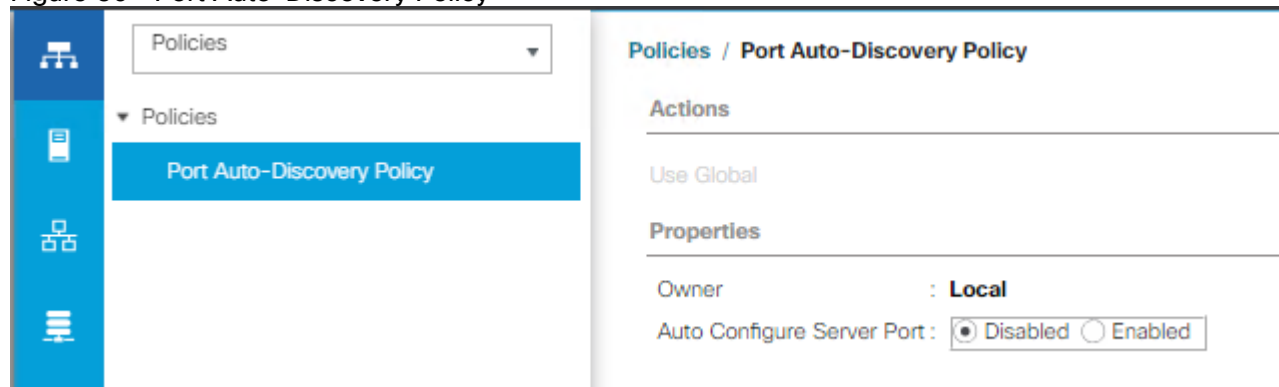
A new feature in Cisco UCS Manager 3.1(3a) and later is Server Port Auto-Discovery, which automates the configuration of ports on the Fabric Interconnects as server ports when a Cisco UCS rack-mount server is connected to them. The firmware on the rack-mount servers must already be at version 3.1(3a) or later in order for this feature to function properly. Enabling this policy eliminates the manual steps of configuring each server port, however it does configure the servers in a somewhat random order. For example, the rack-mount server at

the bottom of the stack, which you may refer to as server #1, and you may have plugged into port 1 of both Fabric Interconnects, could be discovered as server 2, or server 5, and so on. In order to have fine control of the rack-mount server or chassis numbering and order, the manual configuration steps listed in the next section must be followed.

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

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

Figure 56 Port Auto-Discovery Policy



Manual Configuration

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

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

7. Click Yes to confirm the configuration and click OK.
8. Wait for a brief period, until the rack-mount server appears in the Equipment tab underneath Equipment > Rack Mounts > Servers, or the chassis appears underneath Equipment > Chassis.
9. Repeat steps 1-8 for each pair of server ports, until all rack-mount servers and chassis appear in the order desired in the Equipment tab.

Server Discovery

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

Deploy HX Data Platform Installer on Hyper-V Infrastructure

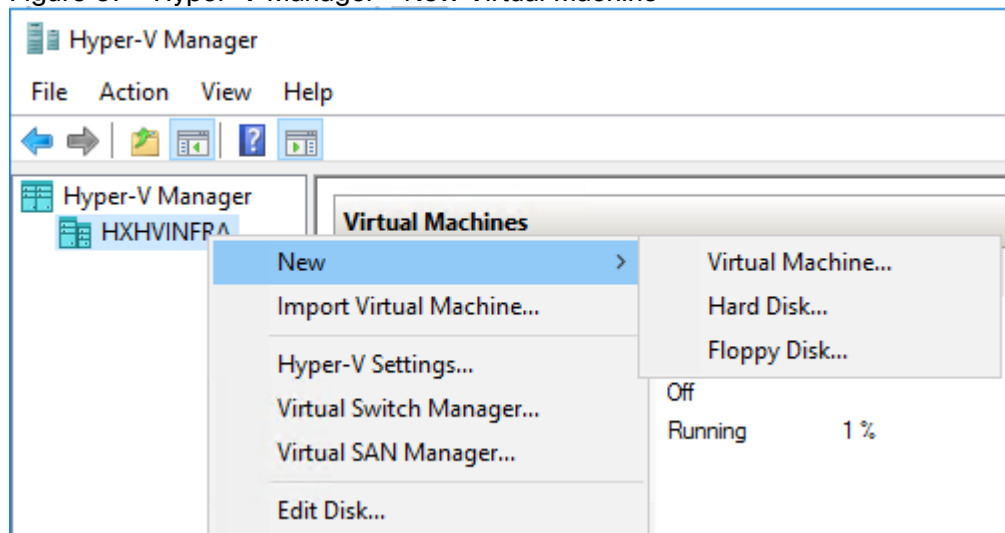
To deploy HX Data Platform Installer using Microsoft Hyper-V Manager to create a HX Data Platform Installer virtual machine, follow these steps:

1. Locate and download the HX Data Platform Installer.vhdx zipped file (for example, Cisco-HX-Data-Platform-Installer-v4.0.1b-33133-hyperv.vhdx.zip) from the Cisco Software Downloads site.
2. Extract the zipped folder to your local computer and copy the .vhdx file to the Hyper-V host where you want to host the HX Data Platform Installer. For example,

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

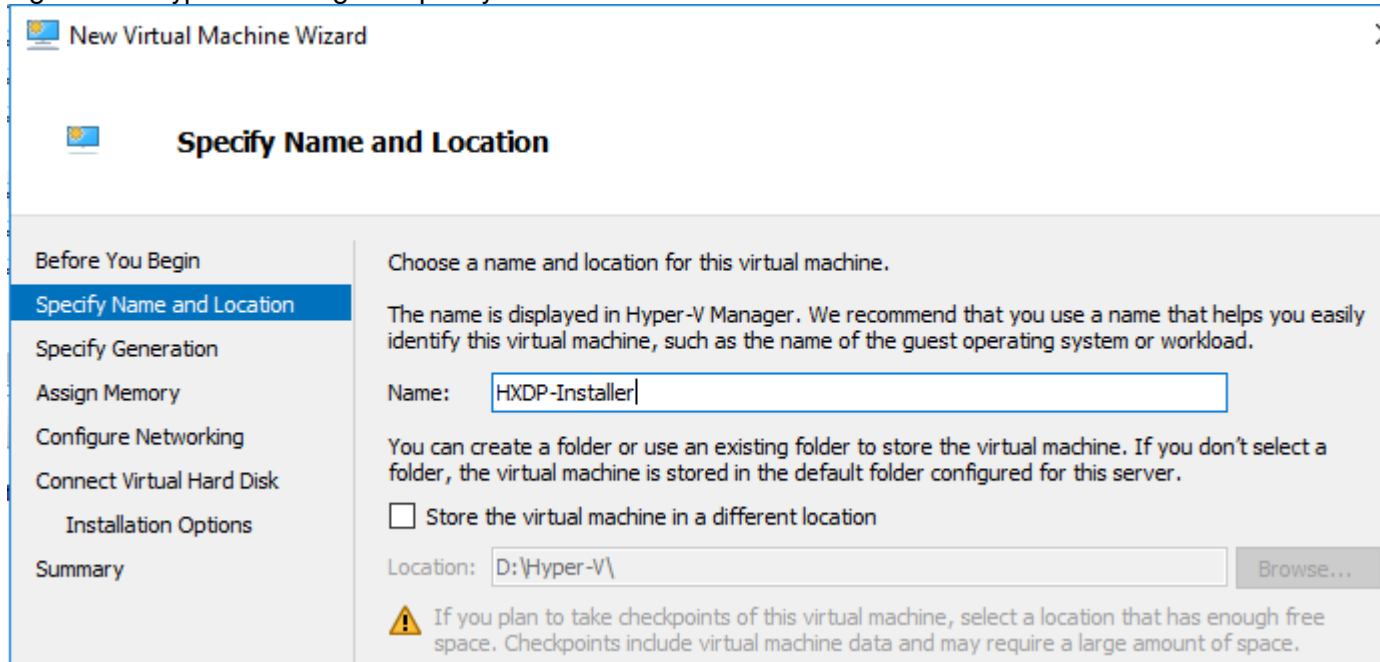
3. In Hyper-V Manager, navigate to one of the Hyper-V servers.
4. Select the Hyper-V server, and right-click and select New > Create a virtual machine. The Hyper-V Manager New Virtual Machine Wizard displays.

Figure 57 Hyper-V Manager - New Virtual Machine



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

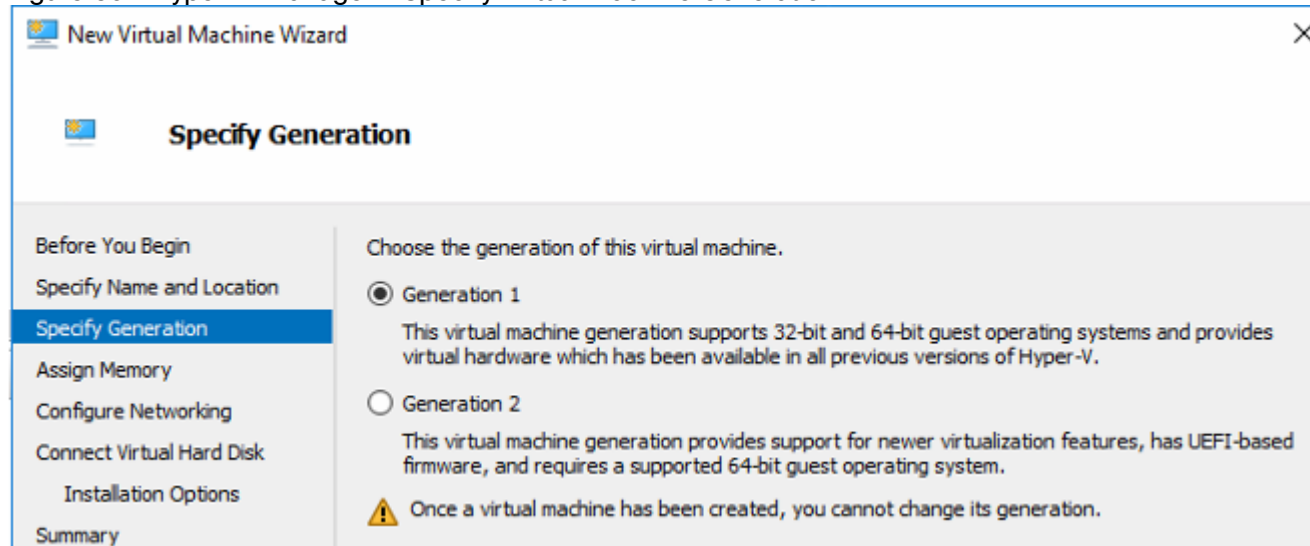
Figure 58 Hyper-V Manager – Specify Virtual Machine Name




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

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

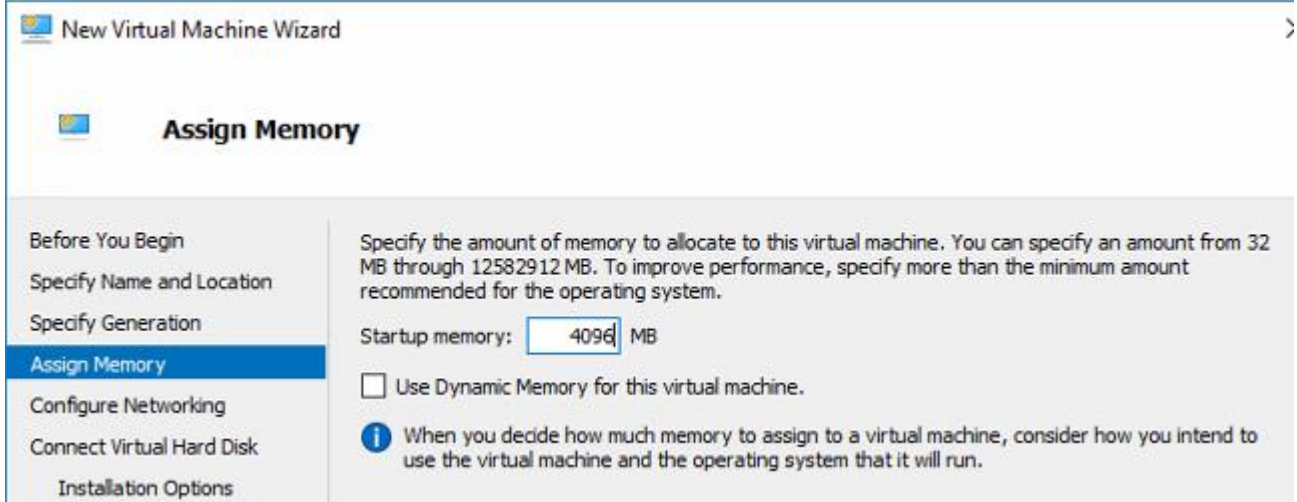
Figure 59 Hyper-V Manager – Specify Virtual Machine Generation



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


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

Figure 60 Hyper-V Manager – Assign Virtual Machine Memory



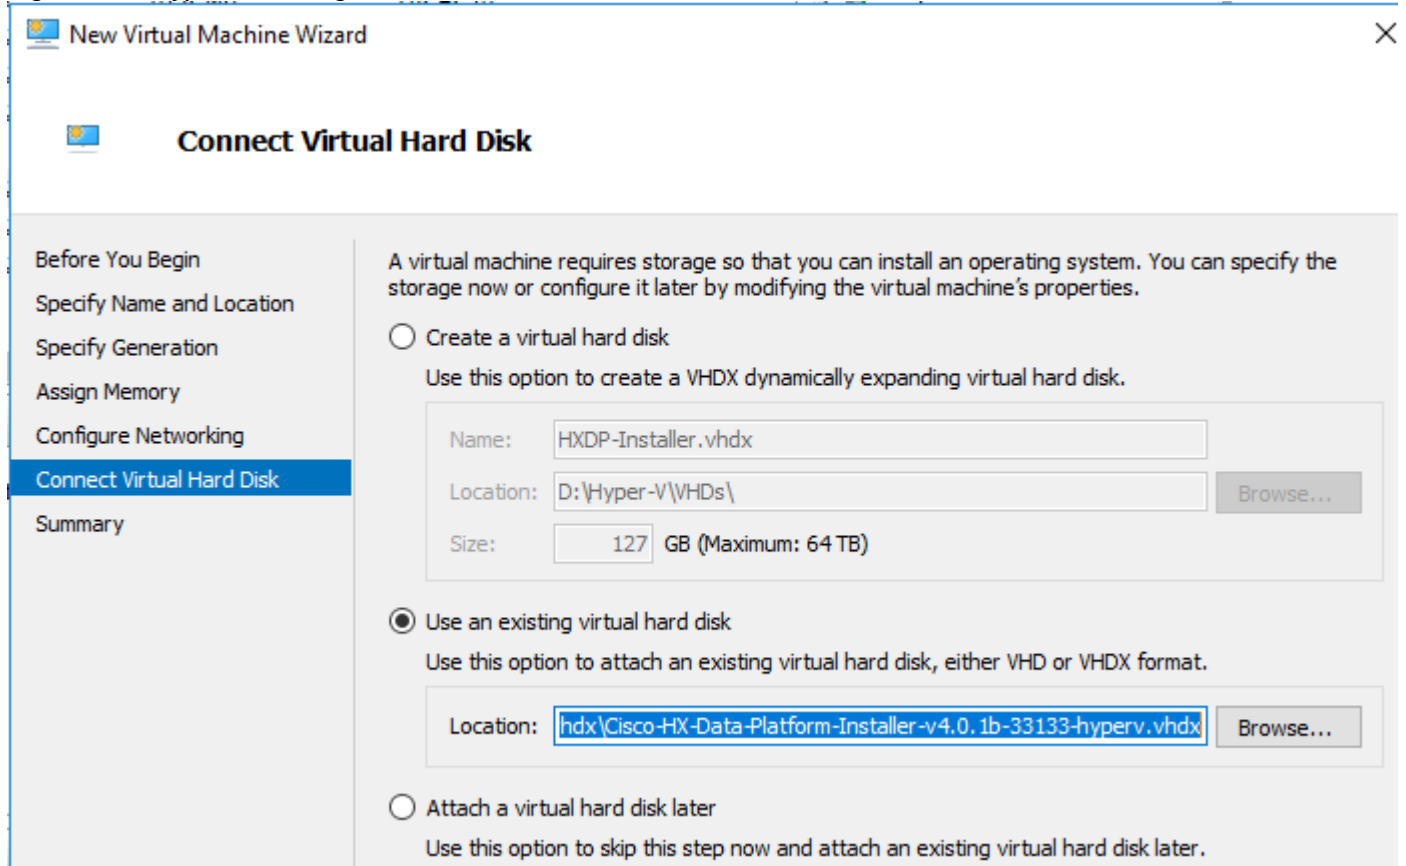
9. In the Configure Networking page, select a network connection for the virtual machine to use from a list of existing virtual switches. Click Next.

Figure 61 Hyper-V Manager – Configure Virtual Machine Networking



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

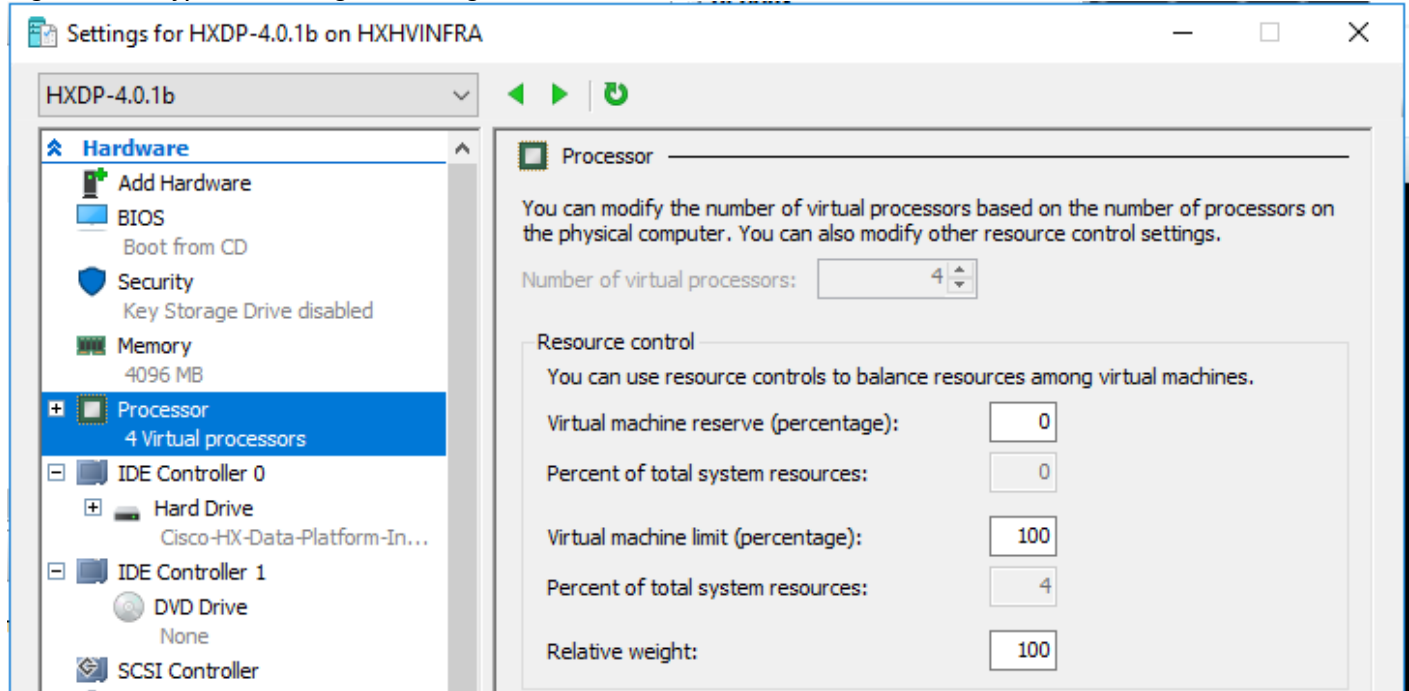
Figure 62 Hyper-V Manager - Connect Virtual Hard Disk



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


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

Figure 63 Hyper-V Manager - Assign Virtual CPUs



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

Figure 64 HXDP Installer



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

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

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

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

```
auto eth0
iface eth0 inet static
metric 100
address 10.104.252.48
netmask 255.255.255.0
gateway 10.104.252.1
dns-nameservers 10.104.252.48
dns-search hxhvd0m2.local
```

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

Figure 65 Show Interfaces

```

root@HyperFlex-Installer:~# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:15:5d:fc:33:1e
          inet addr:10.104.252.48  Bcast:10.104.252.255  Mask:255.255.255.0
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:59075803  errors:0  dropped:14  overruns:0  frame:0
          TX packets:5030442  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:1000
          RX bytes:23680472001 (23.6 GB)  TX bytes:34042337833 (34.0 GB)

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:3794673  errors:0  dropped:0  overruns:0  frame:0
          TX packets:3794673  errors:0  dropped:0  overruns:0  carrier:0
          collisions:0  txqueuelen:1
          RX bytes:19882985351 (19.8 GB)  TX bytes:19882985351 (19.8 GB)

```

Figure 66 Routing Table

```

root@HyperFlex-Installer:~# route -n
Kernel IP routing table
Destination      Gateway         Genmask        Flags Metric Ref    Use Iface
0.0.0.0          10.104.252.1  0.0.0.0        UG    0      0      0 eth0
10.104.252.0     0.0.0.0        255.255.255.0  U     0      0      0 eth0
239.255.255.253 0.0.0.0        255.255.255.255 UH    0      0      0 eth0
root@HyperFlex-Installer:~# _

```

Figure 67 /etc/resolv.conf file for Nameserver and Search Domain

```

root@HyperFlex-Installer:~# cat /etc/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 10.104.252.138
nameserver 10.104.252.44
nameserver 172.18.1.44
search HXHVDOM2.LOCAL hxhvd0m1.local
root@HyperFlex-Installer:~#


```

HyperFlex Installer Web Page

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

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

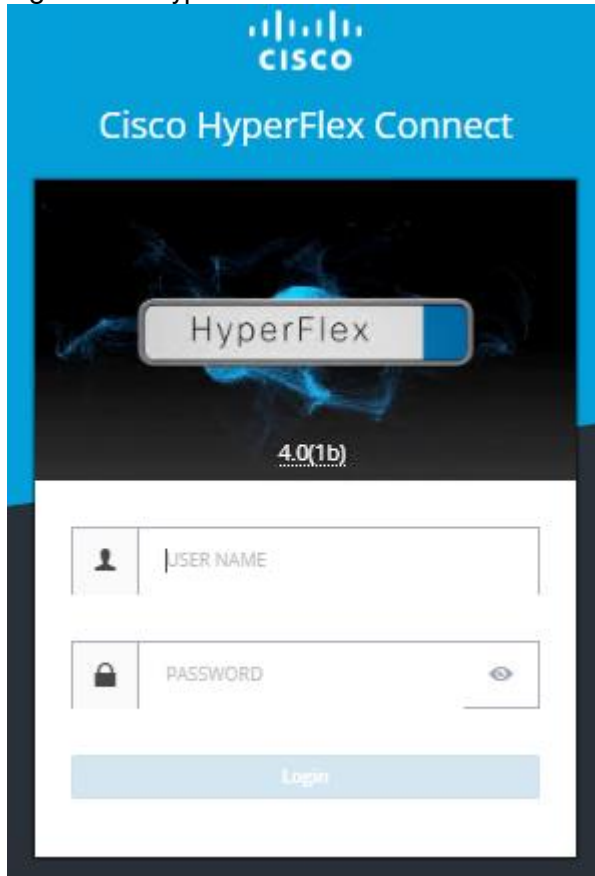
Figure 68 HyperFlex Installer Virtual Machine IP Address



To access the HyperFlex installer webpage, follow these steps:

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

Figure 69 HyperFlex Connect



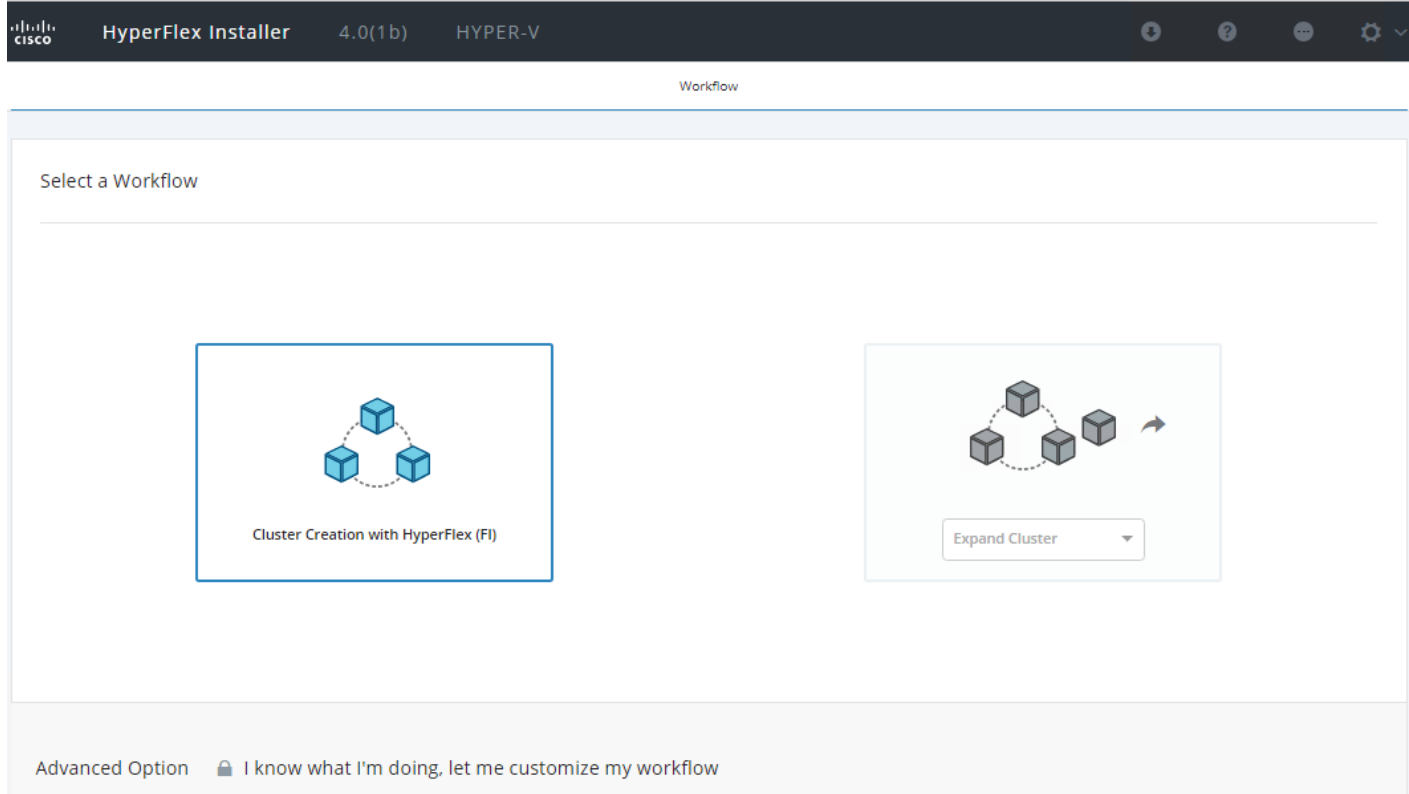
HyperFlex Installation

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

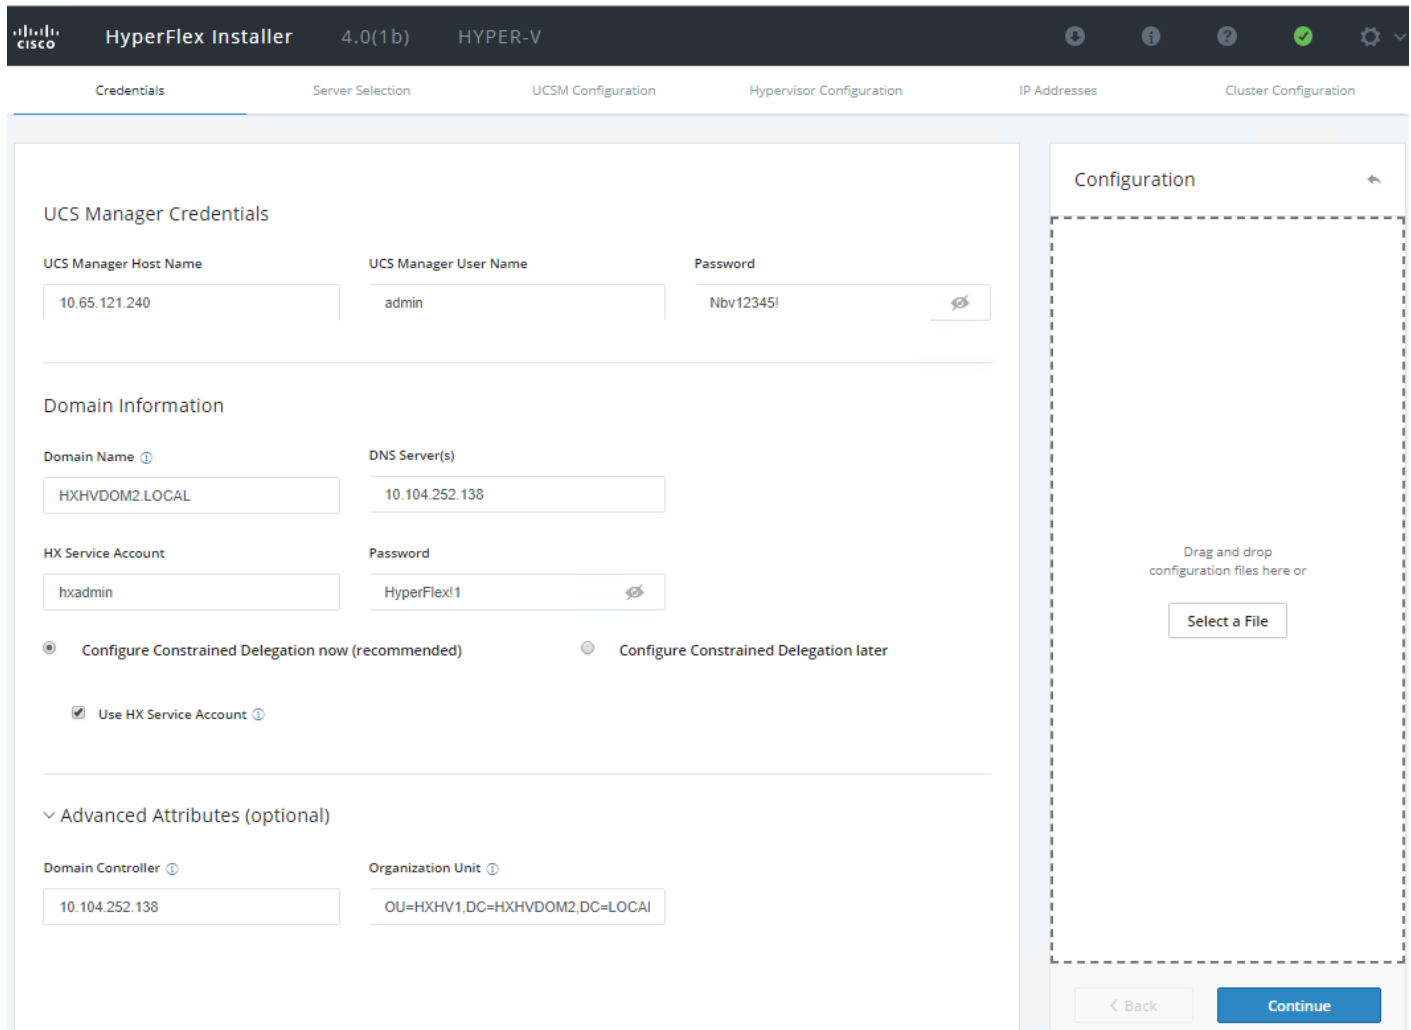
To install and configure a HyperFlex cluster, follow these steps:

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

Figure 70 HyperFlex Installer - Workflow



2. On the Credentials page:
 - a. Under the “UCS Manager Credentials”, enter the Cisco UCS Manager Host Name or IP Address, UCS Manager User Name and Password.
 - g. Under the “Domain Information”, enter the AD Domain Name, DNS Server IP Address, HX Service Account user name and password. It is recommended to select “Configure Constrained Delegation now” and select “Use HX Service Account” if HX service account is member of AD Domain Admin group, else provide Domain Admin credentials (which is a one-time requirement). To configure Constrained Delegation later, refer the appendix section of this document.
 - h. And, under “Advanced Attributes (optional)”, enter the Domain Controller IP address and the distinguished name of Organization Unit (OU). Providing distinguished name of the OU is required, if you want the Computer objects created to be placed under a specific OU instead of the default built-in “Computers” OU.
3. Optionally, you can import a JSON file that has the configuration information, except for the appropriate passwords.



4. Click Continue.

5. In the Server Selection page:

- a. Select the Unassociated HX server models that are to be used in the new HX cluster and click Continue. We have selected three nodes to demonstrate deployment of 3-node hx cluster.
- i. If the Fabric Interconnect server ports were not enabled in the earlier step, you have the option to enable them here to begin the discovery process by clicking the Configure Server Ports link.



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

The screenshot shows the HyperFlex Installer 4.0(1b) HYPERS-V interface. The 'Server Selection' tab is active, displaying a table of unassociated servers. The 'Configuration' panel on the right shows various settings like UCS Manager Host Name, User Name, Domain Name, etc.

Server Name	Status	Model	Serial	Assoc State	Actions
Server 1	unassociated	HXAF240C-M5SX	WZP22020L9E	none	none
Server 2	unassociated	HXAF240C-M5SX	WZP220216WY	none	none
Server 3	unassociated	HXAF240C-M5SX	WZP22020L96	none	none
Server 4	unassociated	HXAF240C-M5SX	WZP22020L9B	none	none

The Configuration panel shows the following settings:

- UCS Manager Host Name: 10.65.121.240
- UCS Manager User Name: admin
- Domain Name: HXHVDOM2.LOCAL
- HX Service Account: hxadmin
- Constrained Delegation: true
- Time Zone: India Standard Time
- DNS Server(s): 10.104.252.138
- Domain Controller: 10.104.252.138
- Organization Unit: OU=HXHV1,DC=HXHVDOM2,DC=LOCAL
- Local Administrator User Name: Administrator



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

6. On the UCSM Configuration page:
 - a. VLAN Configuration – HyperFlex needs to have at least 4 VLANs to function; each VLAN needs to be on different IP subnets and extended from the fabric interconnects to the connecting uplink switches, to make sure that traffic can flow from Primary Fabric Interconnect (Fabric A) to Subordinate Fabric Interconnect (Fabric B).
 - b. Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma separated VLAN IDs for different guest virtual machine networks are allowed here.



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



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



Renaming the 4 core networks is not supported.

- c. MAC Pool - Enter the MAC Pool prefix, only enter the 4th byte value, for example: 00:25:B5:0A.
 - d. 'hx' IP Pool for Cisco IMC - Enter the IP address range, subnet mask and gateway to be used by the CIMC interfaces of the servers in this HX cluster.
 - e. Cisco IMC access Management (Out of band or inband) - Select the recommended 'in band' option for faster installation of hypervisor OS on all the hx nodes.
 - f. The Out-Of-Band network needs to be on the same subnet as the Cisco UCS Manager. You can add multiple blocks of addresses as a comma separated line.
 - g. VLAN for Inband Cisco IMC Connectivity - Enter a VLAN name and ID.
 - h. Advanced - If multiple firmware packages exist on the Fabric Interconnect, choose the version to be installed on the servers that will comprise this cluster. Note that for HXDP 4.0(1b) release with M5 generation servers running on 2nd Generation Intel® Xeon® Scalable Processors, the supported and recommended version of UCS FI firmware's is 4.0(4d).
 - i. Enter a unique Org name for the HyperFlex Cluster.
-



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

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



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

The screenshot displays the HyperFlex Installer 4.0(1b) HYPER-V interface during the UCSM Configuration step. The main configuration area is divided into several sections:

- VLAN Configuration:**
 - VLAN for Hypervisor and HyperFlex management:** VLAN Name: hx-inband-mgmt, VLAN ID: 613
 - VLAN for HyperFlex storage traffic:** VLAN Name: hx-storage-data, VLAN ID: 3172
 - VLAN for VM Live Migration:** VLAN Name: hx-livemigrate, VLAN ID: 3173
 - VLAN for VM Network:** VLAN Name: vm-network, VLAN ID(s): 3174,3175
- MAC Pool:** MAC Pool Prefix: 00:25:B5:0A
- 'hx' IP Pool for Cisco IMC:** IP Blocks: 10.104.252.74-76, Subnet Mask: 255.255.255.0, Gateway: 10.104.252.1
- Cisco IMC access management (Out of band or Inband):** In band (recommended) is selected.
- VLAN for inband Cisco IMC connectivity:** VLAN Name: hx-inband-cimc, VLAN ID: 613
- Advanced:** UCS Server Firmware Version: 4.0(2d), HyperFlex Cluster Name: HXCLUS, Org Name: HXHV1

The right-hand Configuration sidebar shows the following details:

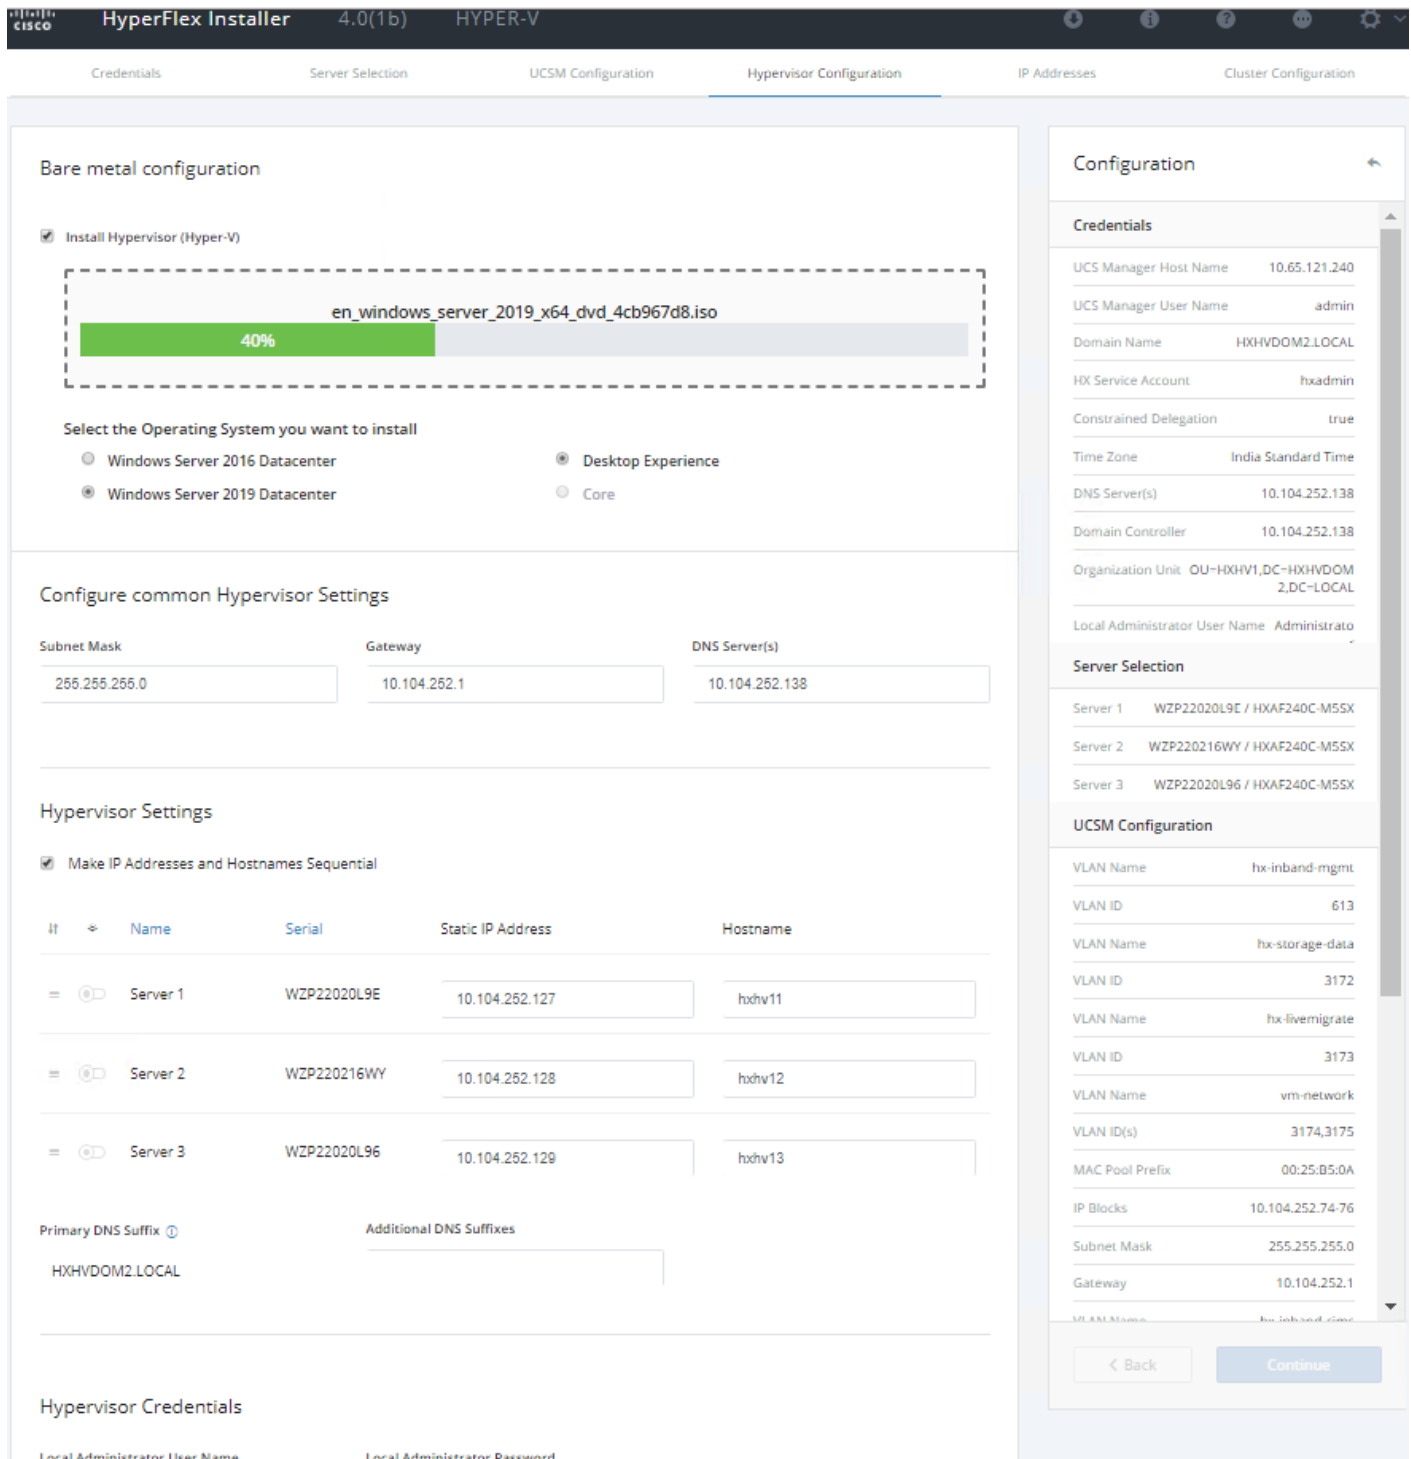
- Credentials:** UCS Manager Host Name: 10.65.121.240, UCS Manager User Name: admin, Domain Name: HXHVDOM2.LOCAL, HX Service Account: hxadmin, Constrained Delegation: true, Time Zone: India Standard Time, DNS Server(s): 10.104.252.138, Domain Controller: 10.104.252.138, Organization Unit: OU=HXHV1,DC=HXHVDOM2,DC=LOCAL, Local Administrator User Name: Administrator
- Server Selection:** Server 1: WZP22020L9E / HXAF240C-M55X, Server 2: WZP220216WY / HXAF240C-M55X, Server 3: WZP22020L96 / HXAF240C-M55X

Navigation buttons include '< Back' and 'Continue'.

7. Click Continue.

8. On the Hypervisor Configuration page:

- a. Bare Metal Configuration - If Windows Server 2016/2019 is not installed on the nodes, select "Install Hypervisor (Hyper-V)", drag or click browse to upload OS media file in the box and select the radio button to choose OS you want to install.
 - b. Configure Common Hypervisor Settings - Enter the subnet mask, gateway, DNS Server IP Address.
 - c. Hypervisor Settings - Enter IP addresses and hostnames for the Hypervisors that were created in the pre-installation section phase. The IP addresses will be assigned via Serial over Lan (SoL) through Cisco UCS Manager to the Hyper-V host systems as their management IP addresses.
 - d. Primary DNS Suffix - Add any additional DNS suffixes.
9. Click Continue.



10. On the IP Addresses page:

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



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

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



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

The screenshot shows the 'IP Addresses' configuration page in the HyperFlex Installer. The main area contains a table for server configurations and a configuration sidebar on the right.

		Management - VLAN 613 (HXHVDOM2.LOCAL)		Data - VLAN 3172 (Hostname or IP Address)	
It	Name	Hypervisor	Storage Controller	Hypervisor	Storage Controller
Server 1	hxhv11	hxhv11scvm	192.168.11.127	192.168.11.131	
Server 2	hxhv12	hxhv12scvm	192.168.11.128	192.168.11.132	
Server 4	hxhv14	hxhv14scvm	192.168.11.130	192.168.11.134	

	Management	Data
Cluster Address	hxhv1cip	192.168.11.135
Subnet Mask	255.255.255.0	255.255.255.0
Gateway	10.104.252.1	

Configuration Sidebar:

- Credentials:**
 - UCS Manager Host Name: 10.65.121.240
 - UCS Manager User Name: admin
 - Domain Name: HXHVDOM2.LOCAL
 - HX Service Account: hxadmin
 - Constrained Delegation: true
 - Time Zone: India Standard Time
 - DNS Server(s): 10.104.252.138
 - Domain Controller: 10.104.252.138
 - Organization Unit: OU=HXHV1,DC=HXHVDOM2,DC=LOCAL
 - Local Administrator User Name: Administrator
- Server Selection:**
 - Server 1: WZP22020L9E / HXAF240C-M5SX
 - Server 2: WZP220216WY / HXAF240C-M5SX
 - Server 4: WZP22020L9B / HXAF240C-M5SX
- UCSM Configuration:**
 - VLAN Name: hx-inband-mgmt
 - VLAN ID: 613
 - VLAN Name: hx-storage-data

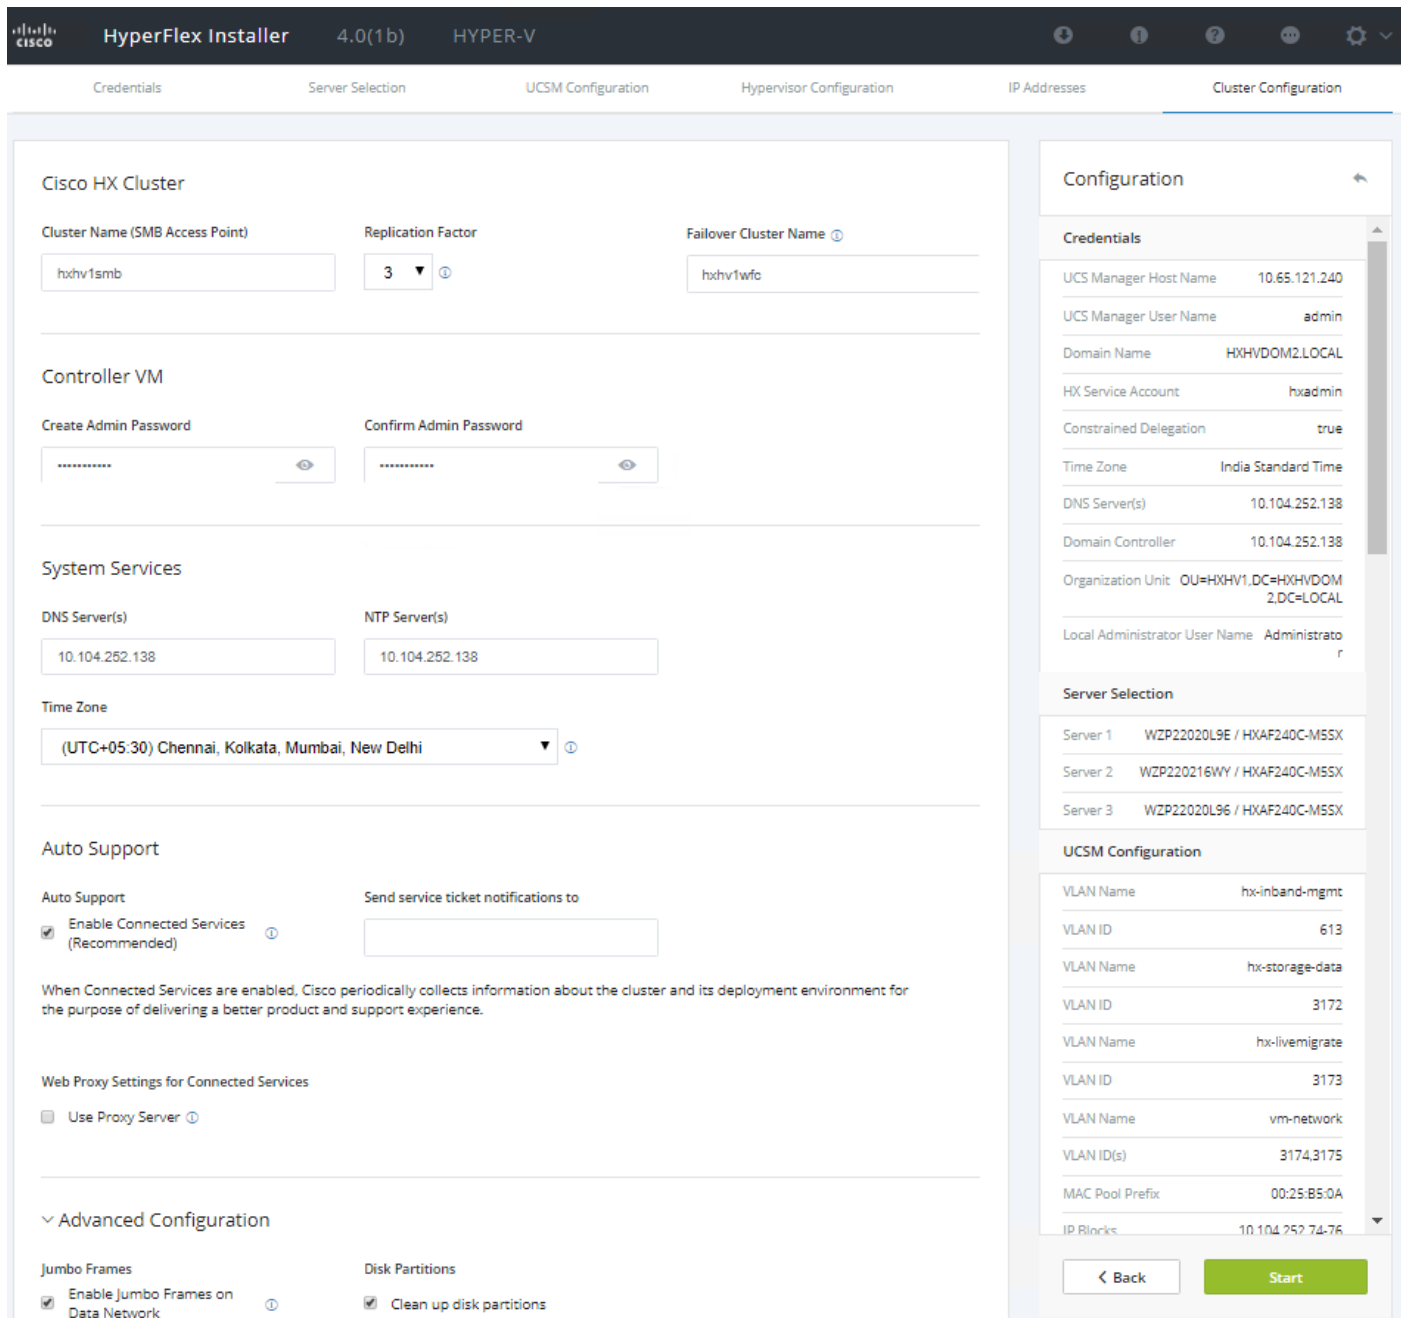
Buttons: < Back, Continue

11. On the Cluster Configuration page:

- a. Cisco HX Cluster - Enter the Cluster Name (SMB Access Point). Select a Replication Factor from the drop-down list and enter a Windows Failover Cluster Name that was created in the pre-installation phase.
- b. Controller virtual machine - Enter the Password that will be assigned to the Controller virtual machines.
- c. System Services - Enter the AD DNS server IP address, and make sure to use the Active Directory domain name for NTP Server for Controller virtual machines to synchronize time with the Active Directory.
- d. Time Zone - Select a time zone from the drop-down list.
- e. Auto Support - Enable Connected Services to enable management via Cisco Intersight and enter the email address to receive service ticket alerts, then scroll down.

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

12. Click Start.



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



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

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

Below the progress bar, a section titled "UCSM Configuration in Progress" is visible. The main area of the interface is divided into two columns. The left column displays validation results under a "Validations" dropdown menu. The right column shows configuration details for Credentials, Server Selection, and UCSM Configuration.

Validations - Overall (Succeeded):

- Cluster Management IP Not Used
- Cluster Data IP Not Used
- Verify DNS Servers
- Verify NTP Servers
- Verify SMTP Server

UCSM Validation:

- Hardware: Validating cluster has servers with same number of data disks
- Hardware: Validating cluster has servers with consistent coprocessor card presence
- Hardware: Validating homogeneity of all selected servers as either data at rest encryption capable or non-capable

UCS Manager:

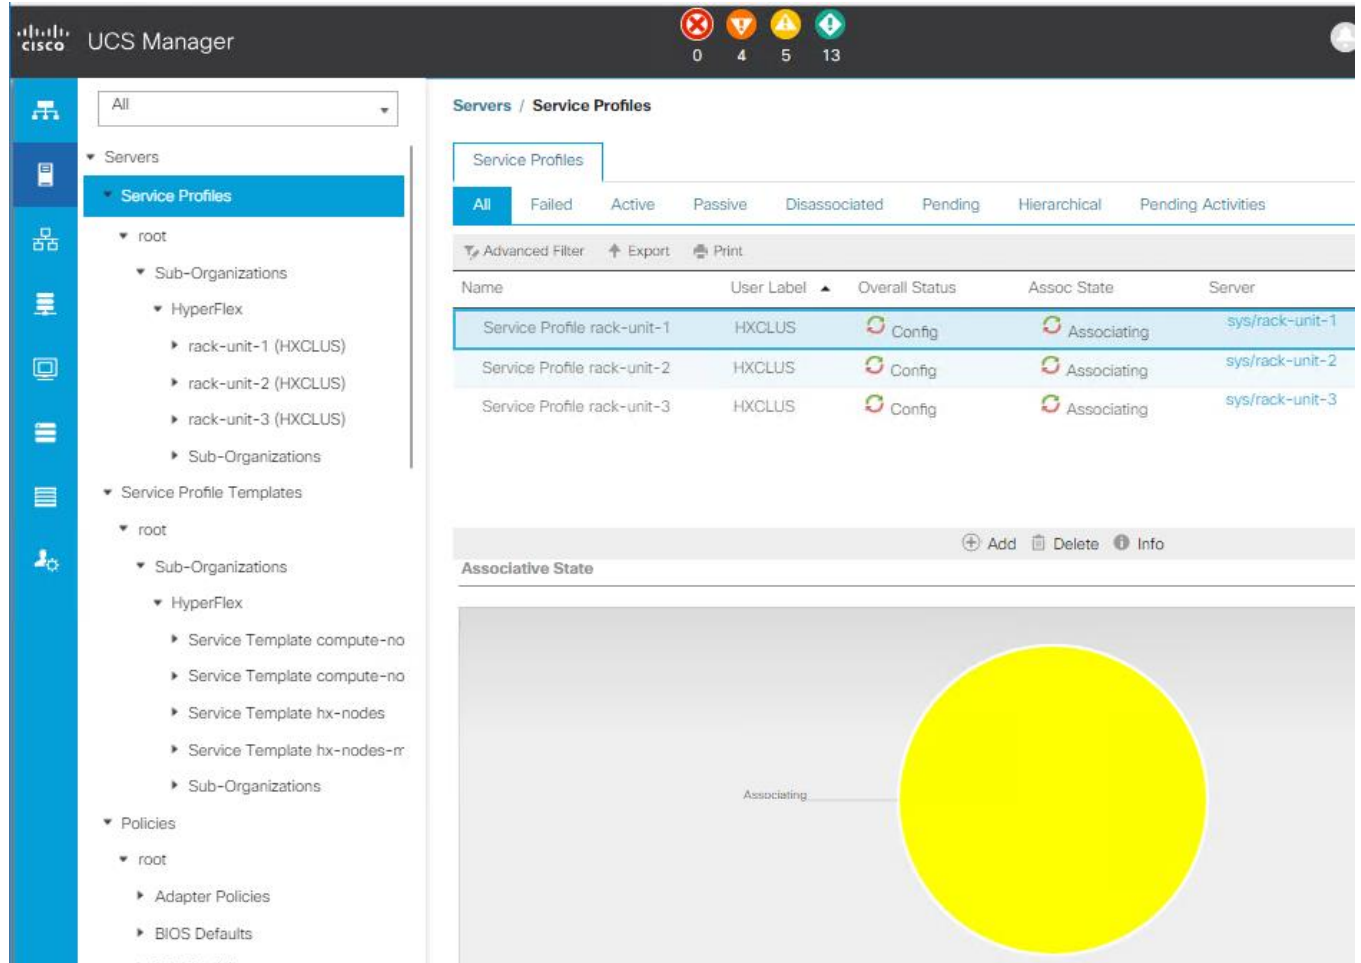
- Firmware: Validating HyperFlex version '3.5(2a)' is compatible with UCS M5 server bundle version '4.0(1b)'
- QoS: Validating QoS class parameter(s) change for system class: 'platinum'
- QoS: Validating QoS class parameter(s) change for system class: 'gold'

Configuration Details:

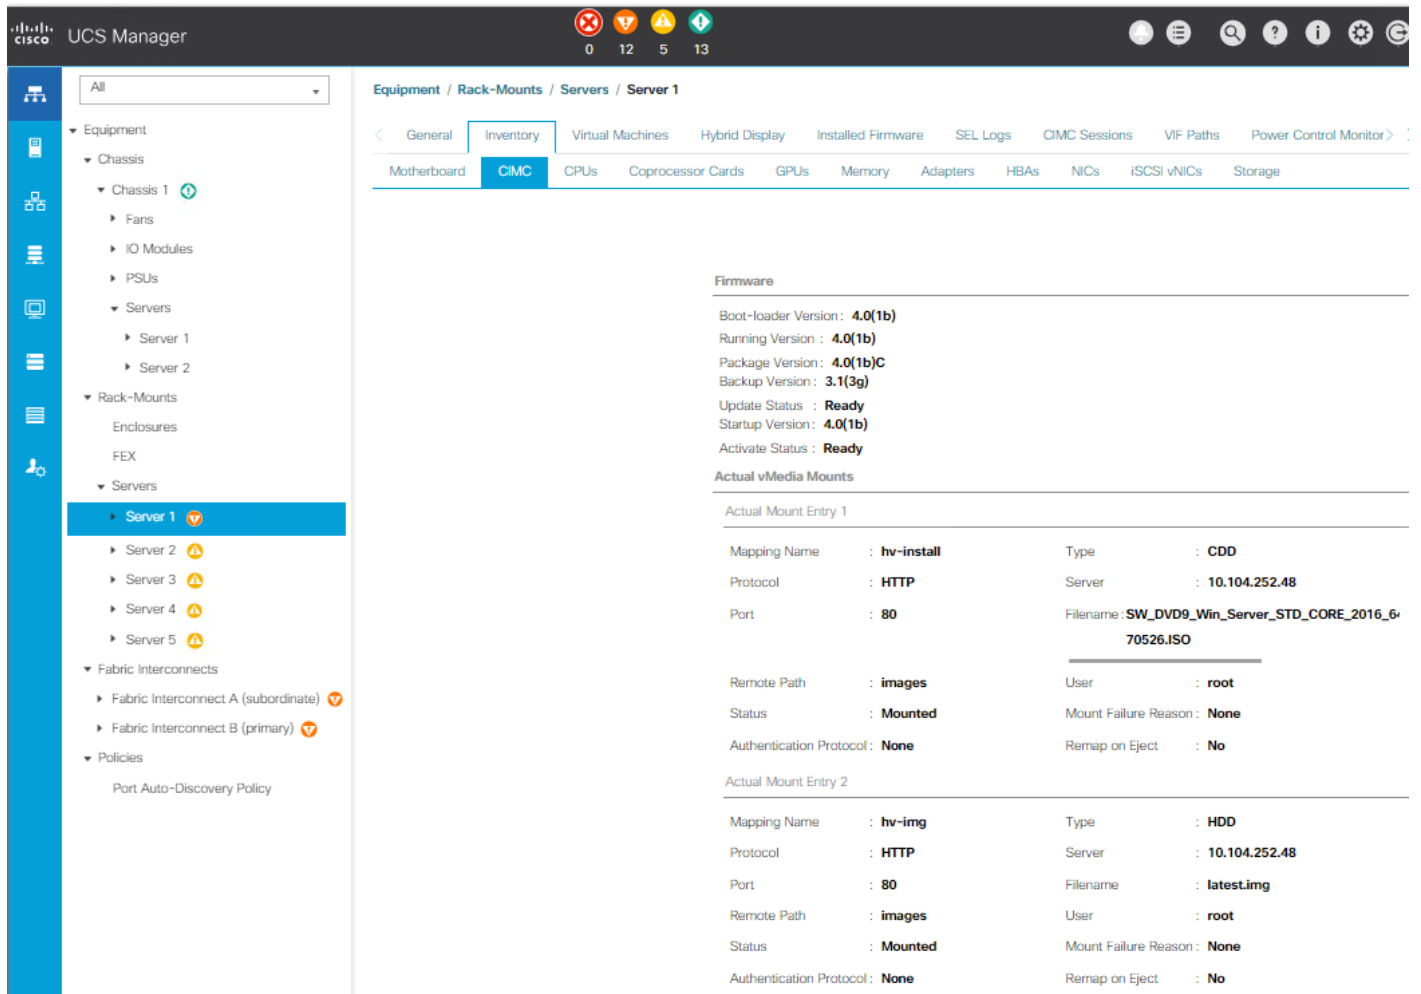
- Credentials:**
 - UCS Manager Host Name: 10.65.121.240
 - UCS Manager User Name: admin
 - Domain Name: HXHVDOM1.LOCAL
 - HX Service Account: hxadmin
 - Constrained Delegation: true
 - Time Zone: India Standard Time
 - DNS Server(s): 10.104.252.44
 - Domain Controller: 10.104.252.44
 - Organization Unit: OU=HyperFlex.DC=hxhvdom1.DC=local
 - Local Administrator User Name: Administrator
- Server Selection:**
 - Server 2: WZP220216WY / HXAF240C-M55X
 - Server 3: WZP22020L96 / HXAF240C-M55X
 - Server 1: WZP22020L9E / HXAF240C-M55X
- UCSM Configuration:**
 - VLAN Name: hx-inband-mgmt, VLAN ID: 613
 - VLAN Name: hx-storage-data, VLAN ID: 3172
 - VLAN Name: hx-livemigrate, VLAN ID: 3173
 - VLAN Name: vm-network, VLAN ID(s): 3174

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

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



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



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

Progress

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

Hypervisor Configuration in Progress

Hypervisor Configuration

Hypervisor Configuration - Overall In Progress

- ✓ Login to UCS API
- ✓ Setting up configuration for OS installation
- ✓ Creating temporary boot policy for OS install on MS
- ✓ Inventorying physical servers
- ✓ Logout from UCS API
- ✓ CONFIGURATION COMPLETED SUCCESSFULLY
- ⌚ Waiting for all servers to install and acquire IP address...

rack-unit-1 In Progress

- ✓ Setting up node for OS installation
- ✓ Powering off server
- ✓ Changing boot policy
- ✓ Mounting remote vMedia image
- ✓ Powering on server
- ✓ Waiting for BIOS post completion
- ✓ Waiting for BIOS boot message
- ⌚ Waiting for server to install and acquire IP address...

rack-unit-2

- ✓ Setting up node for OS installation

Configuration

Credentials

UCS Manager Host Name	10.65.121.240
UCS Manager User Name	admin
Domain Name	HXHVDOM1.LOCAL
HX Service Account	hxadmin
Constrained Delegation	true
Time Zone	India Standard Time
DNS Server(s)	10.104.252.44
Domain Controller	10.104.252.44
Organization Unit	OU=HyperFlex,DC=hxhvdom1,DC=local
Local Administrator User Name	Administrator

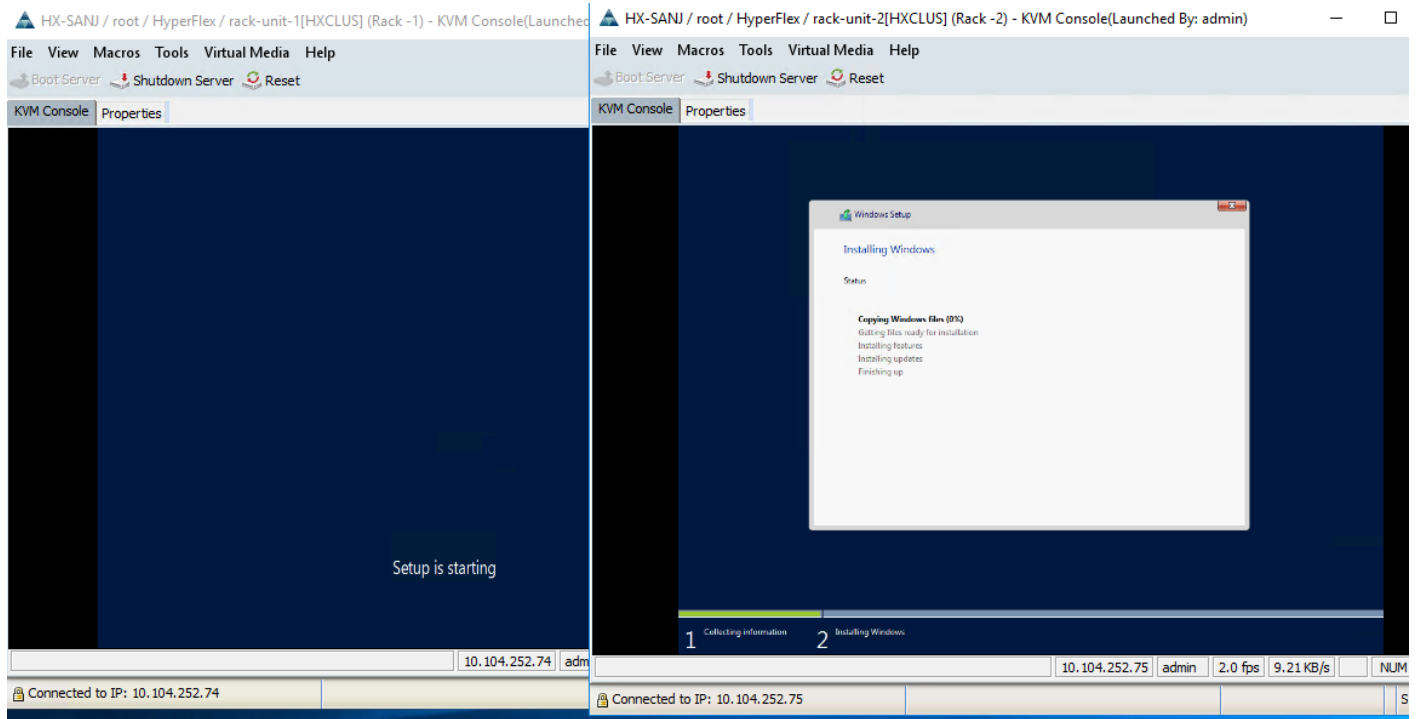
Server Selection

Server 2	WZP220216WY / HXAF240C-M5SX
Server 3	WZP22020L96 / HXAF240C-M5SX
Server 1	WZP22020L9E / HXAF240C-M5SX

UCSM Configuration

VLAN Name	hx-inband-mgmt
VLAN ID	613
VLAN Name	hx-storage-data
VLAN ID	3172
VLAN Name	hx-livemigrate
VLAN ID	3173
VLAN Name	vm-network
VLAN ID(s)	3174

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



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

Progress

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

Deploy in Progress

Deploy Validation

Deploy Validation - Overall Succeeded

- Cluster Management IP resolveable
- Nodes Compatible check
- Storage Controller Management IP List Name Resolution Check
- Storage Controller Data IP List Name Resolution Check
- Hypervisor Management IP List Name Resolution Check
- Hypervisor Data IP List Name Resolution Check
- Hypervisor host check
- Hypervisor max cluster size check
- Data IP's specified check
- Data IP subnet specified check
- Data Network IP's in the same subnet
- Management IP's specified check
- Management IP subnet specified check
- Management Network IP's in the same subnet
- NTP reachability
- DNS reachability

hxxv1.HXHVDOM1.LOCAL Succeeded

- HyperV authentication and reachability check

Configuration

Credentials

UCS Manager Host Name	10.65.121.240
UCS Manager User Name	admin
Domain Name	HXHVDOM1.LOCAL
HX Service Account	hxadmin
Constrained Delegation	true
Time Zone	India Standard Time
DNS Server(s)	10.104.252.44
Domain Controller	10.104.252.44
Organization Unit	OU=HyperFlex.DC=hxhvdom1.DC=local
Local Administrator User Name	Administrator

Server Selection

Server 2	WZP220216WY / HXAF240C-M55X
Server 3	WZP22020L96 / HXAF240C-M55X
Server 1	WZP22020L9E / HXAF240C-M55X

UCSM Configuration

VLAN Name	hx-inband-mgmt
VLAN ID	613
VLAN Name	hx-storage-data
VLAN ID	3172
VLAN Name	hx-livemigrate
VLAN ID	3173
VLAN Name	vm-network
VLAN ID(s)	3174

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

The screenshot shows the 'Summary' tab of the HyperFlex Installer. At the top, it displays 'HyperFlex Installer 4.0(1b) HYPER-V'. The cluster name is 'hxxhv1smb', which is marked as 'ONLINE' and 'HEALTHY'. Below this, there are two columns of configuration details:

Version	4.0.1b-33133	Domain Name	HXXHVDOM2.LOCAL
Cluster Management IP Address	hxxhv1cip.HXXHVDOM2.LOCAL	Failover cluster Name	hxxhv1wfc
Cluster Data IP Address	192.168.11.135	DNS Server(s)	10.104.252.138
Replication Factor	Three copies	NTP Server(s)	10.104.252.138
Available Capacity	8.0 TB		

Below the configuration details is a 'Servers' table:

Model	Serial Number	Management Hypervisor	Management Storage Controller	Data Network Hypervisor	Data Network Storage Controller
HXXAF240C-M55X	WZP22020L9B	10.104.252.130	10.104.252.134	192.168.11.130	192.168.11.134
HXXAF240C-M55X	WZP22020L9E	10.104.252.127	10.104.252.131	192.168.11.127	192.168.11.131
HXXAF240C-M55X	WZP220216WY	10.104.252.128	10.104.252.132	192.168.11.128	192.168.11.132

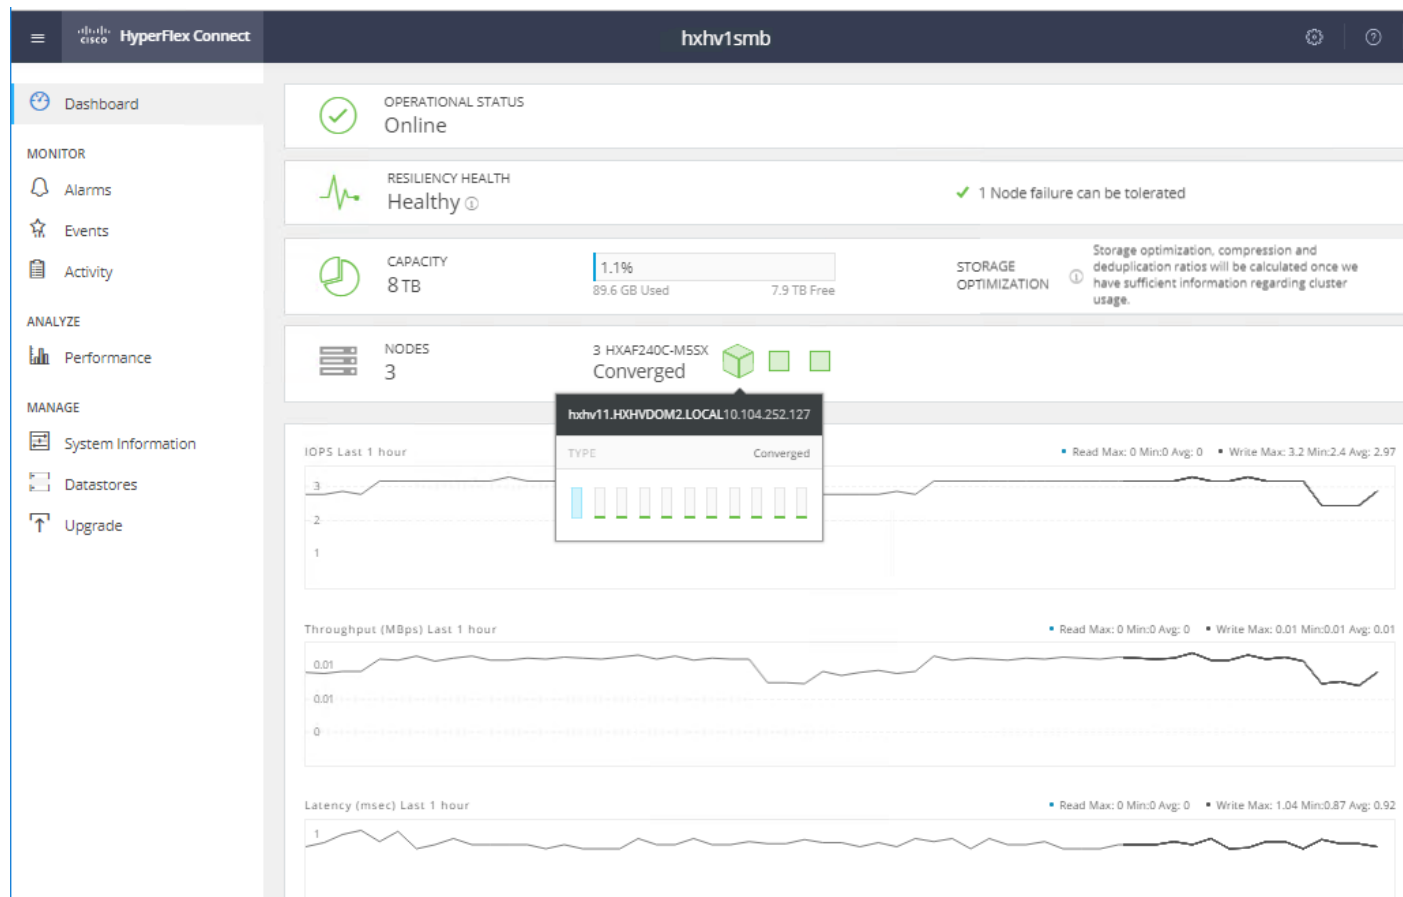
At the bottom right, there are two buttons: 'Back to Workflow Selection' and 'Launch HyperFlex Connect'.

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

This screenshot is similar to the previous one, but with a red box highlighting the 'Export Configuration' button in the top right corner of the application window. The button is located in the top right toolbar, next to a downward-pointing arrow icon.

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

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



Post Installation Tasks

Create Datastores

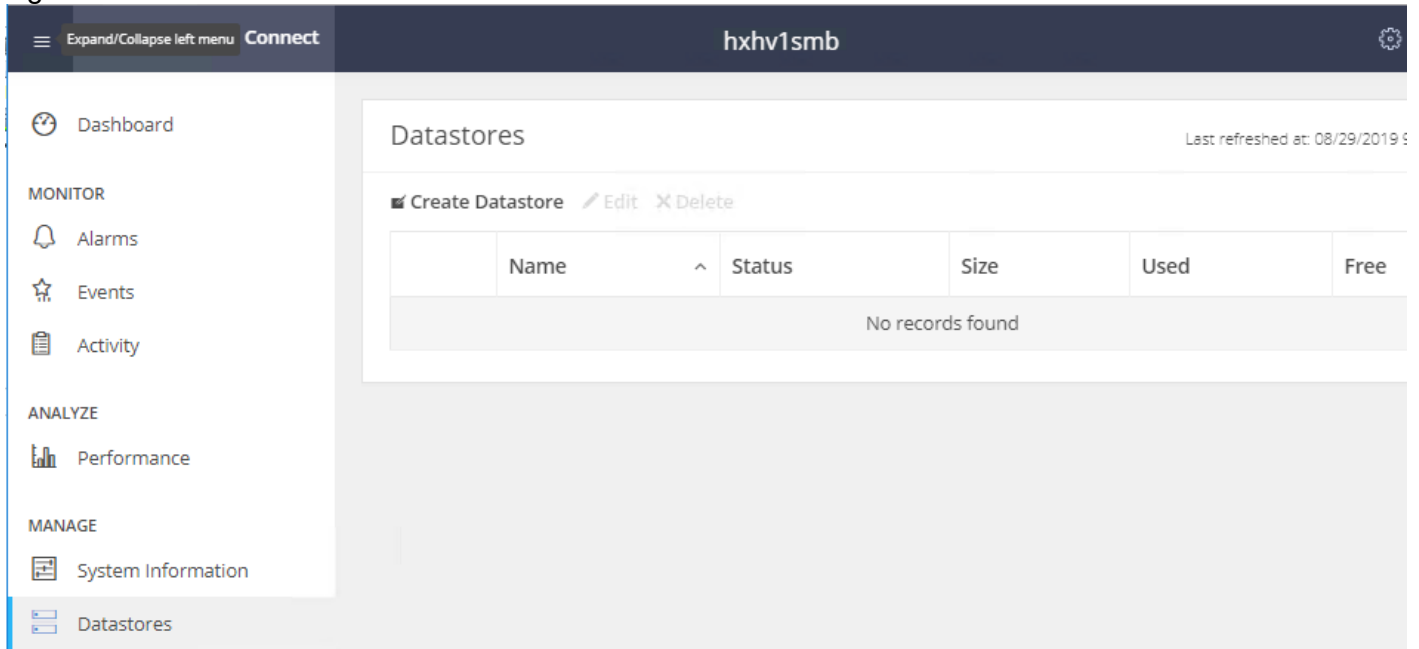
Create a datastore for storing the virtual machines. This task can be completed by using the HyperFlex Connect HTML management webpage. The Datastores created using HX Connect creates a SMB share which the HyperFlex Hyper-V nodes can use it to store virtual machine files. To configure a new datastore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



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

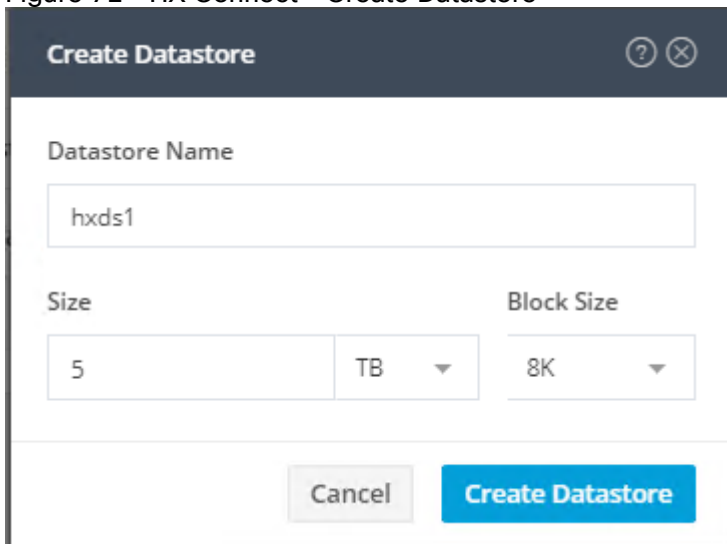
1. Use a web browser to open the HX cluster IP management URL.
2. Enter the credentials.
3. Click Login.
4. Click Datastores in the left pane and click Create Datastore.

Figure 71 HX Connect - Datastores



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

Figure 72 HX Connect - Create Datastore



6. Click Create Datastore.

Figure 73 HX Connect – Datastore Status

The screenshot shows the HyperFlex Connect web interface. At the top, the title bar reads 'HyperFlex Connect' and 'hxhv1smb'. A green notification banner at the top center says 'Created datastore hxds1'. Below this, the 'Datastores' section is visible, showing a table with the following data:

	Name	Status	Size	Used	Free
<input type="checkbox"/>	hxds1	Normal	5 TB	0 B	5 TB

The interface also shows a navigation menu on the left with sections for MONITOR (Alarms, Events, Activity), ANALYZE (Performance), and MANAGE (System Information, Datastores). The bottom of the table indicates 'Showing 1 - 1 of 1'.

Constrained Delegation (Optional)

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

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

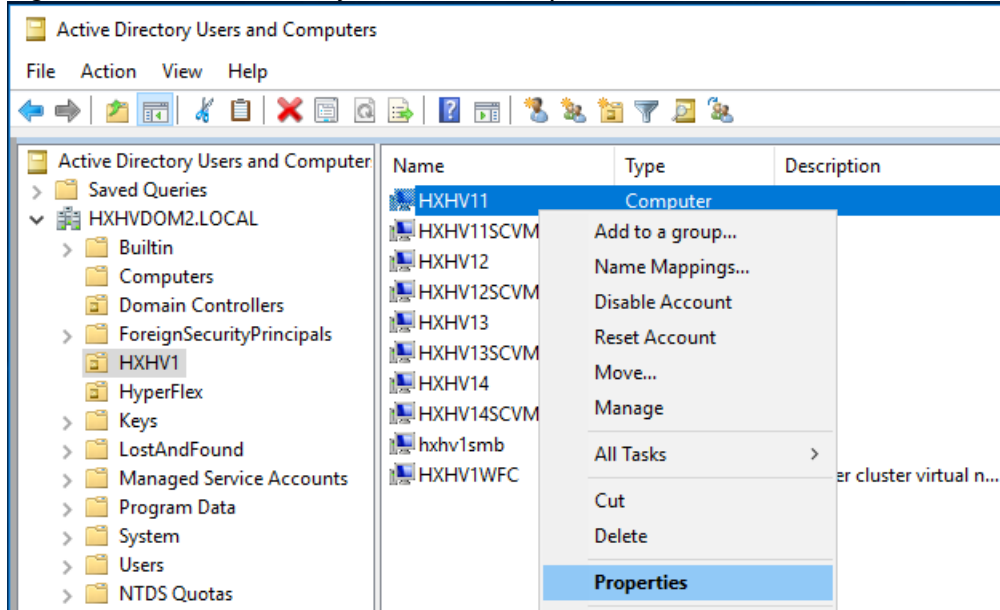


This step must be performed only if Constrained Delegation was not configured during initial installation. It is recommended that you perform this procedure using the HX Installer and not as part of post-installation.

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

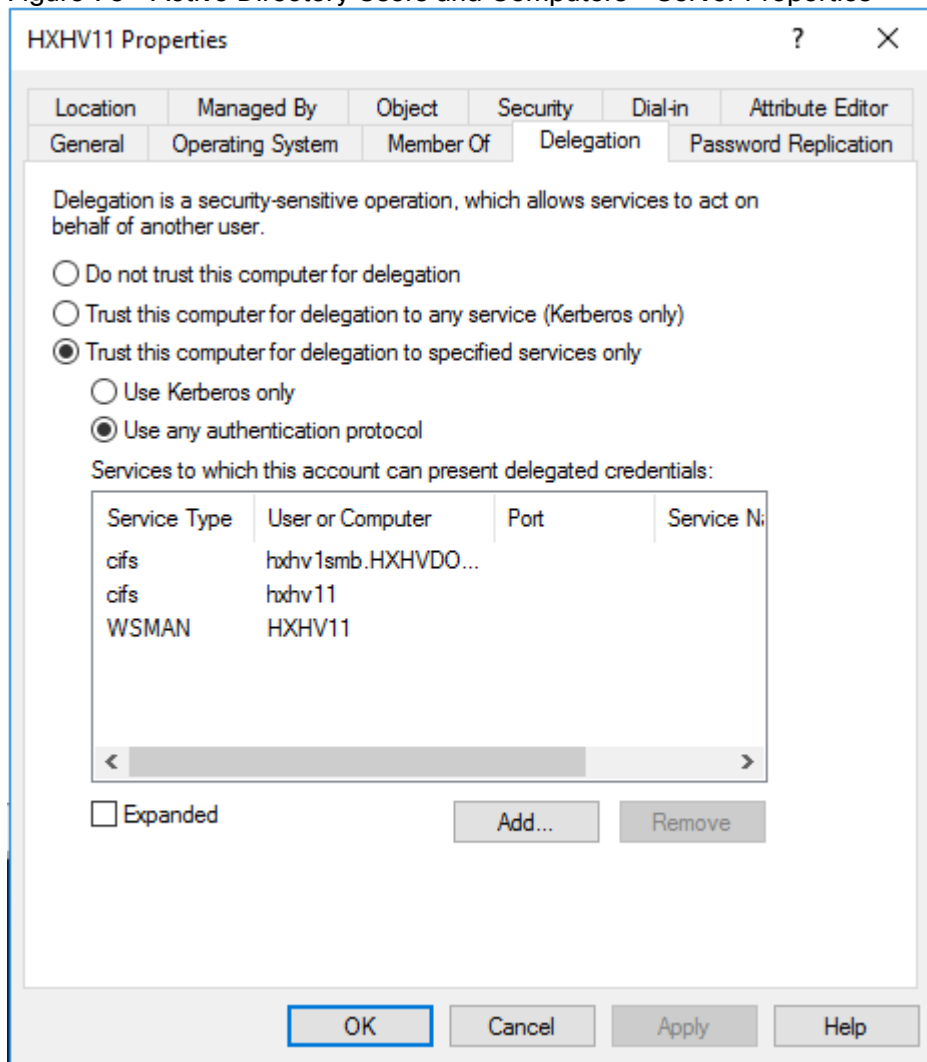
1. Open the Active Directory Users and Computers snap-in. (From Server Manager, select the server if it is not selected, click Tools >> Active Directory Users and Computers).
2. From the navigation pane in Active Directory Users and Computers, select the domain and double-click the Computers folder.
3. From the Computers folder, right-click the computer account of the source server and then click Properties.

Figure 74 Active Directory Users and Computers



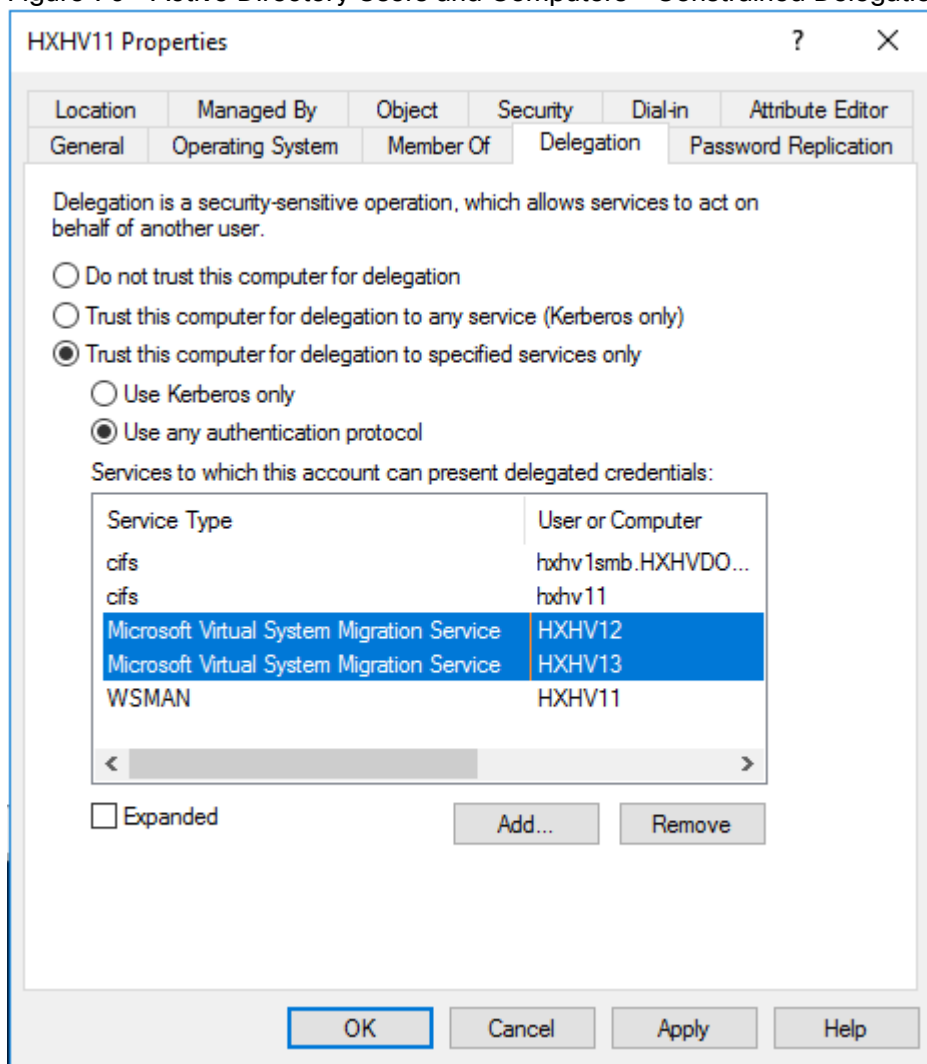
4. From the Properties tab, click the Delegation tab.
5. On the delegation tab, select Trust this computer for delegation to the specified services only and then select Use any authentication protocol.

Figure 75 Active Directory Users and Computers - Server Properties



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

Figure 76 Active Directory Users and Computers - Constrained Delegation



11. On the Delegation tab of the Properties dialog box, verify that the services you selected in the previous step are listed as the services to which the destination computer can present delegated credentials. Click OK.
12. From the Computers folder, select the computer account of the destination server and repeat the process. In the Select Users or Computers dialog box, be sure to specify the name of the source server.

Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces

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

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

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

Figure 77 PowerShell - Get VMSwitch

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Get-VMSwitch -name * | select PSComputerName,Name} | Format-Table -AutoSize
```

Name	PSComputerName	RunspaceId
vswitch-hx-inband-mgmt	hxhv12	5902f1e8-7a62-4535-9648-0bc0c9216bf5
vswitch-hx-storage-data	hxhv12	5902f1e8-7a62-4535-9648-0bc0c9216bf5
vswitch-hx-vm-network	hxhv12	5902f1e8-7a62-4535-9648-0bc0c9216bf5
vswitch-hx-livemigration	hxhv12	5902f1e8-7a62-4535-9648-0bc0c9216bf5
vswitch-hx-vm-network	hxhv14	ddc2de20-d3c7-4c07-9382-535dd69f72f9
vswitch-hx-storage-data	hxhv14	ddc2de20-d3c7-4c07-9382-535dd69f72f9
vswitch-hx-livemigration	hxhv14	ddc2de20-d3c7-4c07-9382-535dd69f72f9
vswitch-hx-inband-mgmt	hxhv14	ddc2de20-d3c7-4c07-9382-535dd69f72f9
vswitch-hx-vm-network	hxhv11	bc454e5a-58fe-45a5-a247-40465649b814
vswitch-hx-inband-mgmt	hxhv11	bc454e5a-58fe-45a5-a247-40465649b814
vswitch-hx-storage-data	hxhv11	bc454e5a-58fe-45a5-a247-40465649b814
vswitch-hx-livemigration	hxhv11	bc454e5a-58fe-45a5-a247-40465649b814

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

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

Figure 78 PowerShell - Remove VMSwitch

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhv11,hxhv12,hxhv14 -ScriptBlock {Get-VMSwitch -name * | select PSComputerName,Name} | Format-Table -AutoSize
```

Name	PSComputerName	RunspaceId
vswitch-hx-inband-mgmt	hxhv12	c2002602-d1fc-4dc7-8099-5517acb3b244
vswitch-hx-storage-data	hxhv12	c2002602-d1fc-4dc7-8099-5517acb3b244
vswitch-hx-vm-network	hxhv12	c2002602-d1fc-4dc7-8099-5517acb3b244
vswitch-hx-vm-network	hxhv11	79a70bcc-2cee-4ea2-b38e-cea2329cc0fa
vswitch-hx-inband-mgmt	hxhv11	79a70bcc-2cee-4ea2-b38e-cea2329cc0fa
vswitch-hx-storage-data	hxhv11	79a70bcc-2cee-4ea2-b38e-cea2329cc0fa
vswitch-hx-vm-network	hxhv14	3ca68630-3192-486d-9087-7d748d86e4ba
vswitch-hx-storage-data	hxhv14	3ca68630-3192-486d-9087-7d748d86e4ba
vswitch-hx-inband-mgmt	hxhv14	3ca68630-3192-486d-9087-7d748d86e4ba

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

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

Figure 79 PowerShell - Assign Static IP

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhv11 -ScriptBlock {New-NetIPAddress -ifAlias "team-hx-livemigration" -IPAddress 192.168.73.127 -PrefixLength 24}
```

```
ifIndex : 6
PSComputerName : hxhv11
RunspaceId : 11e201e8-0654-416e-a32e-6207c56d38d0
Caption :
Description :
ElementName :
InstanceID :
CommunicationStatus :
DetailedStatus :
HealthState :
InstallDate :
Name : ;C<8;@B8A=8; <A55@55;55;
OperatingStatus :
OperationalStatus :
PrimaryStatus :
Status :
StatusDescriptions :
AvailableRequestedStates :
EnabledDefault : 2
EnabledState :
OtherEnabledState :
RequestedState : 12
TimeOfLastStateChange :
TransitioningToState : 12
CreationClassName :
SystemCreationClassName :
SystemName :
NameFormat :
OtherTypeDescription :
ProtocolIFType : 4096
ProtocolType :
Address :
AddressOrigin : 0
AddressType :
IPv4Address : 192.168.73.127
IPv6Address :
```

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

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

- Repeat steps 3 and 4 to assign static IP addresses to the live migration and VM networks on all the Hyper-V hosts
- Verify the jumbo frame settings on the live migration network adapters using the below command:

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

Figure 80 PowerShell – Verify jumbo frame settings

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

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

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

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

Figure 81 PowerShell – Configure jumbo frame settings

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

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

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

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

Ping statistics for 192.168.73.128:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms

```
PS C:\Users\administrator.HXHVDOM2> ping -l 8972 -f 192.168.73.130
```

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

Ping statistics for 192.168.73.130:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 0ms, Maximum = 0ms, Average = 0ms

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

Rename the Cluster Network in Windows Failover Cluster - Optional

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

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

Figure 82 PowerShell – Get Cluster Network

```
[hvh11]: PS C:\Users\administrator.HXHVDOM2\Documents> Get-ClusterNetwork
```

Name	State	Metric	Role
Cluster Network 1	Up	69600	ClusterAndClient
Cluster Network 2	Up	29600	Cluster
Cluster Network 3	Up	29601	Cluster

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

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

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

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

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

Figure 83 PowerShell – Rename the Cluster Network

```
[hvh11]: PS C:\Users\administrator.HXHVDOM2\Documents> (Get-ClusterNetwork | Where-Object {$_.Address -eq "10.104.252.0"}).Name = "hx-inband-mgmt"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.11.0"}).Name = "hx-storage-data"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "192.168.73.0"}).Name = "LiveMigration"
(Get-ClusterNetwork | Where-Object {$_.Address -eq "172.18.0.0"}).Name = "vm-network"
```

3. Verify now the cluster network using the “*Get-ClusterNetwork*” command:

Figure 84 PowerShell – Verify Cluster Network

```
[hvh11]: PS C:\Users\administrator.HXHVDOM2\Documents> Get-ClusterNetwork
```

Name	State	Metric	Role
hx-inband-mgmt	Up	69600	ClusterAndClient
hx-storage-data	Up	29601	Cluster
LiveMigration	Up	29602	Cluster
vm-network	Up	28800	Cluster

Configure the Windows Failover Cluster Network Roles

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

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

```
(Get-ClusterNetwork -Name "hx-inband-mgmt").Role = 3
(Get-ClusterNetwork -Name "hx-storage-data").Role = 0
(Get-ClusterNetwork -Name "LiveMigration").Role = 1
```

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

Figure 85 PowerShell – Configure Cluster Network Roles

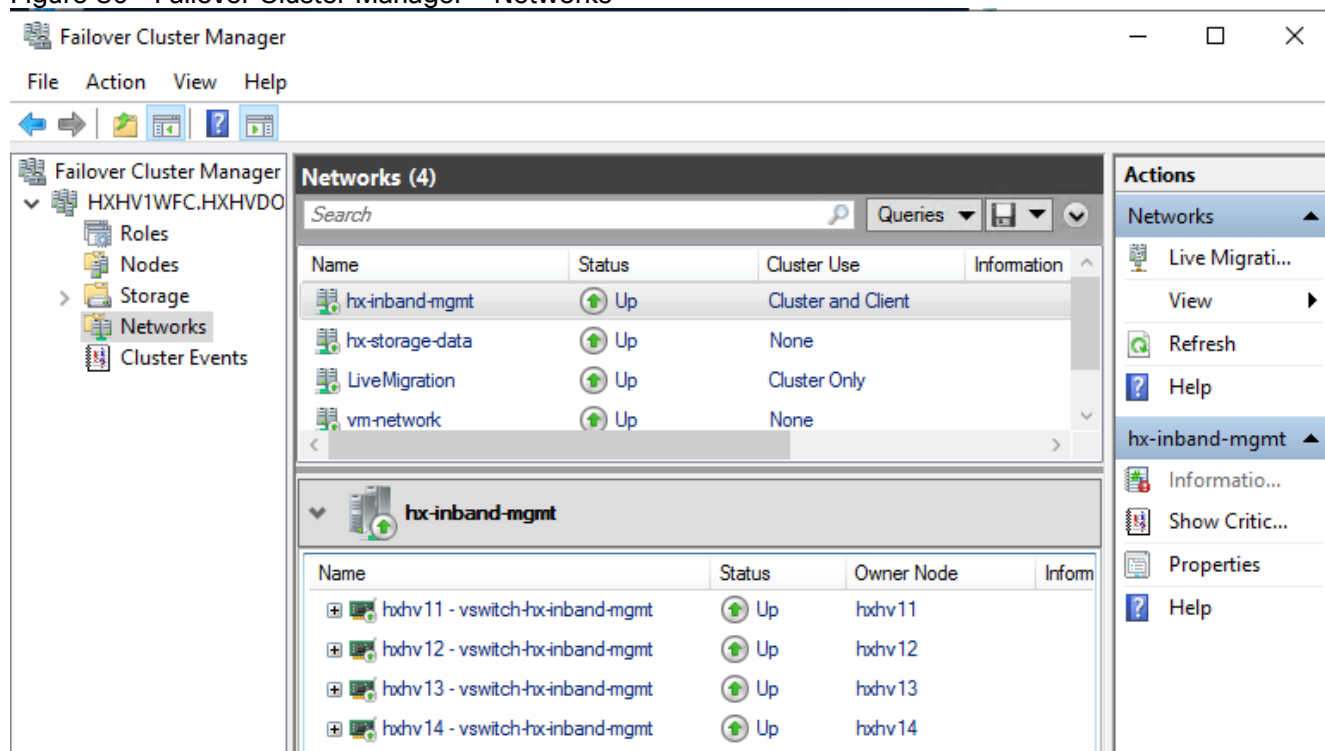
```
[hvh11]: PS C:\Users\administrator.HXHVDM2\Documents> (Get-ClusterNetwork -Name "hx-inband-mgmt").Role = 3
(Get-ClusterNetwork -Name "hx-storage-data").Role = 0
(Get-ClusterNetwork -Name "LiveMigration").Role = 1
(Get-ClusterNetwork -Name "vm-network").Role = 0
```

Role = 0 to disable cluster communication

Role = 1 to enable only cluster communication

Role = 3 to enable both cluster & client communication

Figure 86 Failover Cluster Manager – Networks



Configure the Windows Failover Cluster Network for Live Migration

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

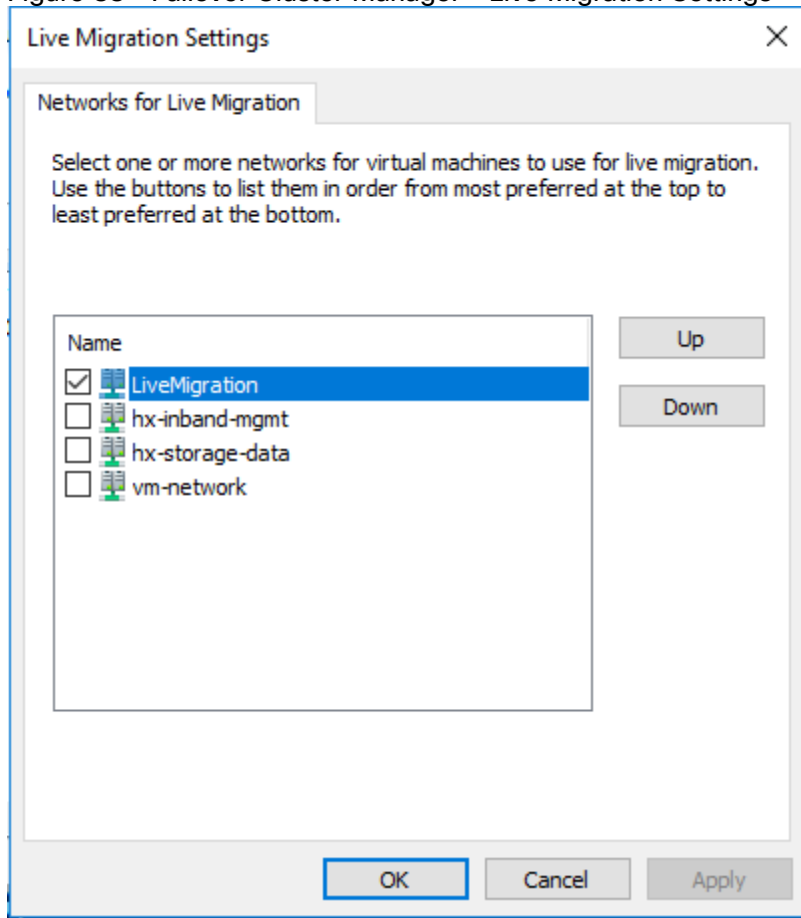
1. Run the PowerShell command shown below to configure the cluster network for live migration traffic:

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

Figure 87 PowerShell – Configure Live Migration Network

```
[hvh11]: PS C:\Users\administrator.HXHVDM2\Documents> Get-ClusterResourceType -Name "Virtual Machine" | Set-ClusterParameter -Name MigrationExcludeNetworks -Value ([String]::Join(";", (Get-ClusterNetwork | Where-Object {$_.Name -ne "LiveMigration"}).ID))
```


Figure 88 Failover Cluster Manager - Live Migration Settings




Create Folders on the HX Datastore

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

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

```
mkdir \\hxnv1smb.hxnvdom.local\hnds1\<folder Name>
```

Figure 89 Create a Folder on SMB Share



2. Create folders for different purposes and requirements.

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

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

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

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

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

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

Invoke-Command -ComputerName $hosts -ScriptBlock {SET-VMHOST -virtualharddiskpath
"\hxhv1smb.hxhvd0m2.local\hxds1\Hyper-V\VHDs" -virtualmachinepath
"\hxhv2smb.hxhvd0m2.local\HXDS1\Hyper-V\"}
```

Figure 90 PowerShell – Configure Virtual Machine Files Store Location

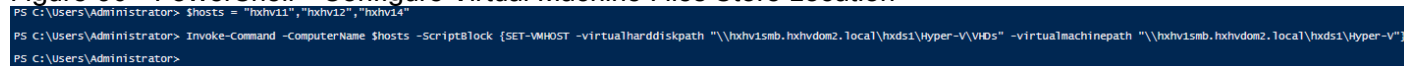


Figure 91 PowerShell – Configure Virtual Machine Files Store Location

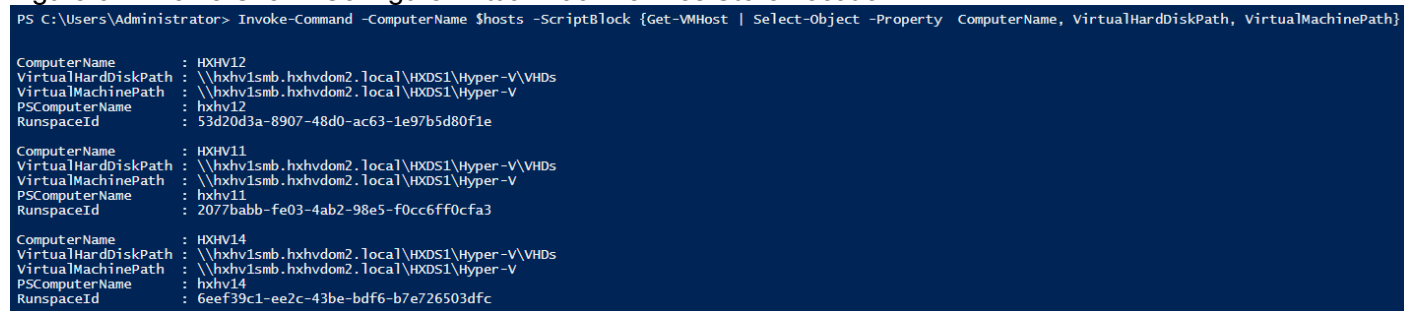



Figure 92 Hyper-V Settings – VHD Store Location



Validate the Windows Failover Cluster Configuration

It is a good practice to validate the Windows failover cluster by running the Validate a Configuration Wizard from the Failover Cluster Manager, or the Test-Cluster Windows PowerShell cmdlet and fix any errors or warnings reported in the results page. [Figure 93](#) shows the command to run the cluster validation.

Figure 93 Failover Cluster - Validate

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

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

Configure Quorum in Windows Server Failover Cluster

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

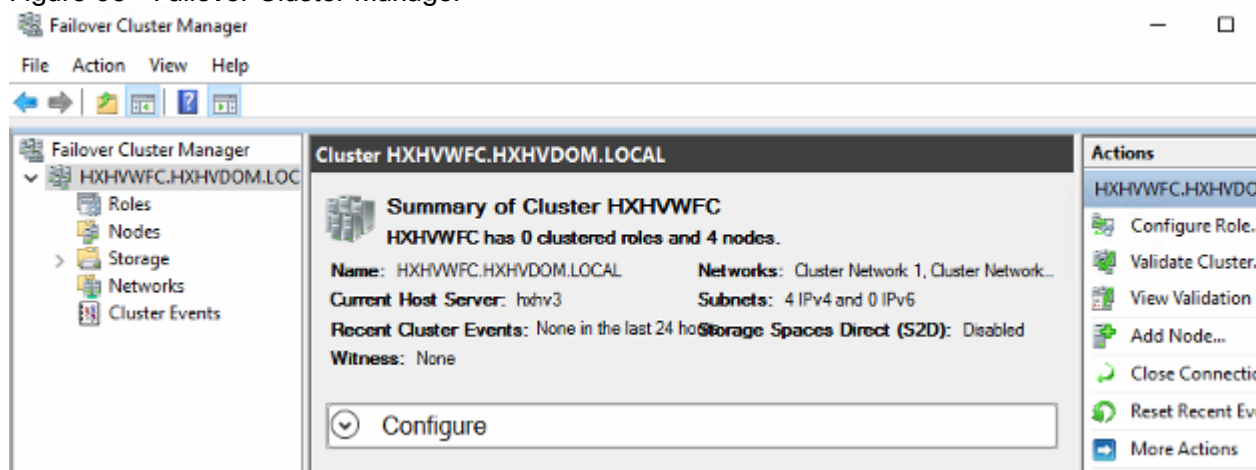
Review the information about quorum resources using the “Get-ClusterQuorum” PowerShell cmdlet or from the Summary page of failover cluster manager as shown below:

Figure 94 PowerShell - Get Cluster Quorum

```
PS C:\Users\Administrator> Get-ClusterQuorum -Cluster hxhv1wfc.hxhvd0m2.local

Cluster              QuorumResource
-----
HXHV1WFC
```

Figure 95 Failover Cluster Manager



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

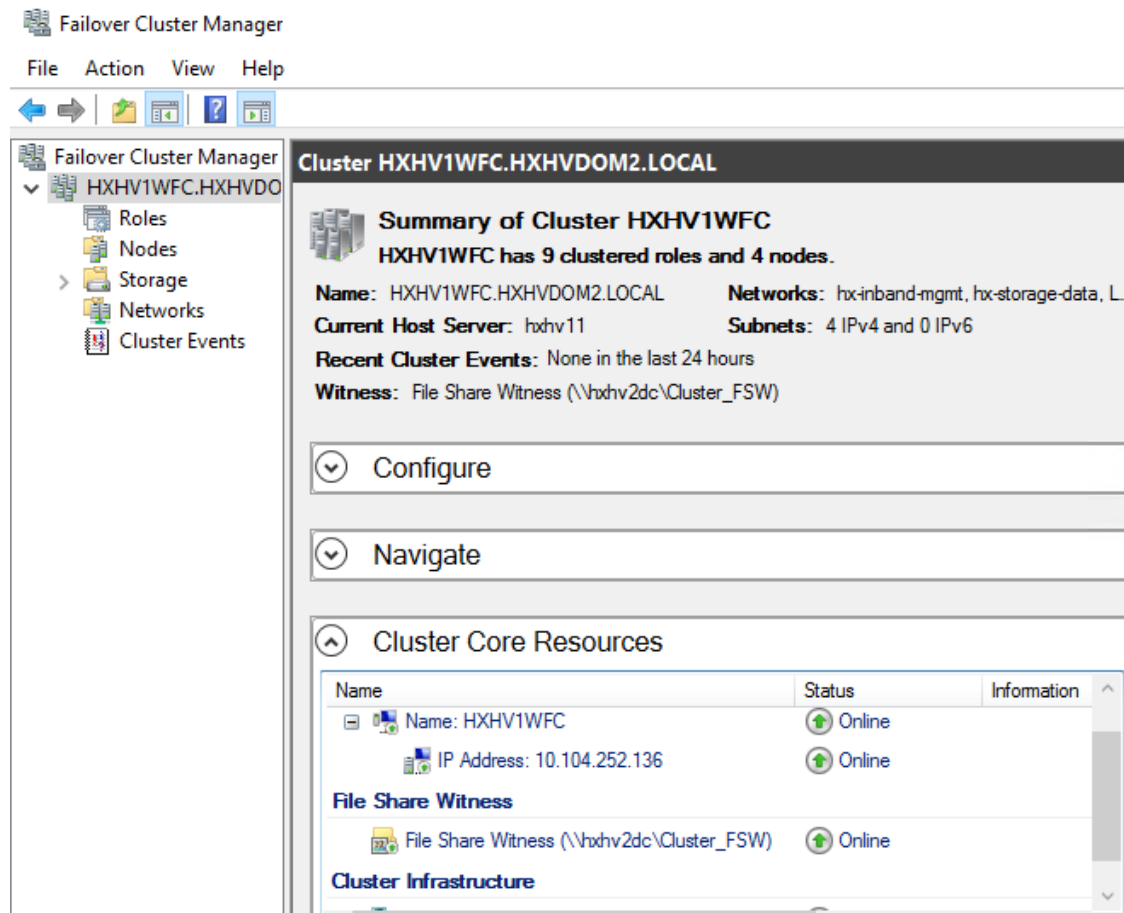
```
Set-ClusterQuorum -NodeAndFileShareMajority "\\fileserver\fsw"
```

Figure 96 PowerShell - Configure Cluster Quorum

```
PS C:\Users\Administrator> Set-ClusterQuorum -Cluster hxhv1wfc.hxhvd0m2.local -NodeAndFileShareMajority \\hxhv2dc\Cluster_FSW

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

Figure 97 Failover Cluster Manager – Quorum Configuration



It is recommended to place the file share witness on a server outside of the cluster. Windows Server 2019 supports using a USB drive connected to a network (router) device as a file share witness for failover clustering.

Initial Tasks and Testing

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

Ready Clones

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

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

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

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

2. Log into a HX Hyper-V node or a remote management station with RSAT tools.
3. Open PowerShell and run the downloaded script as shown below:

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

Figure 98 PowerShell – Create Ready Clones

```
PS C:\Users\administrator.HXHVDOM2>
PS C:\Users\administrator.HXHVDOM2> C:\HxClone-HyperV-v4.0.1b-33133.ps1 -VmName RCVM1 -ClonePrefix c14 -CloneCount 1 -AddToCluster $true

Directory: \\hxhv2smb.hxhvd0m2.local\hxds1\Hyper-V\Virtual Hard Disks

Mode                LastWriteTime         Length Name
----                -
d-----          9/11/2019   7:16 PM          tmp1417411279
\\hxhv2smb.hxhvd0m2.local\hxds1\Hyper-V\Virtual Hard Disks\tmp1417411279\RCVM1\Virtual Machines\9b535cbb-c0a8-4b77-9142-284525fb3033.vmcx

Directory: \\hxhv2smb.hxhvd0m2.local\hxds1

Mode                LastWriteTime         Length Name
----                -
d-----          9/11/2019   7:16 PM          c141
\\hxhv2smb.hxhvd0m2.local\hxds1\c141

Name                : c141
OwnerNode           : hxhv21
State                : Offline

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

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

Figure 99 Failover Cluster Manager - Roles View

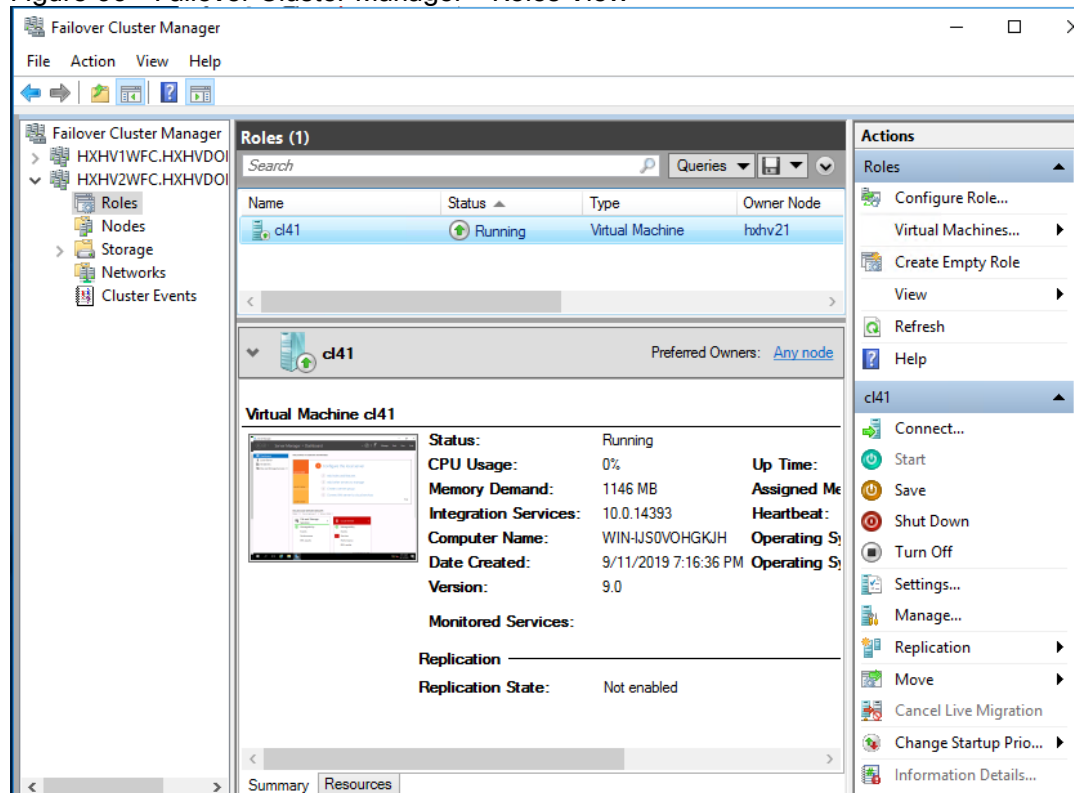


Table 45 Script Parameters

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

Auto-Support and Notifications

Auto-Support should be enabled for all clusters during the initial HyperFlex installation. Auto-Support enables Call Home to automatically send support information to Cisco TAC, and notifications of tickets to the email address specified. If the settings need to be modified, they can be changed in the HyperFlex Connect HTML management webpage.

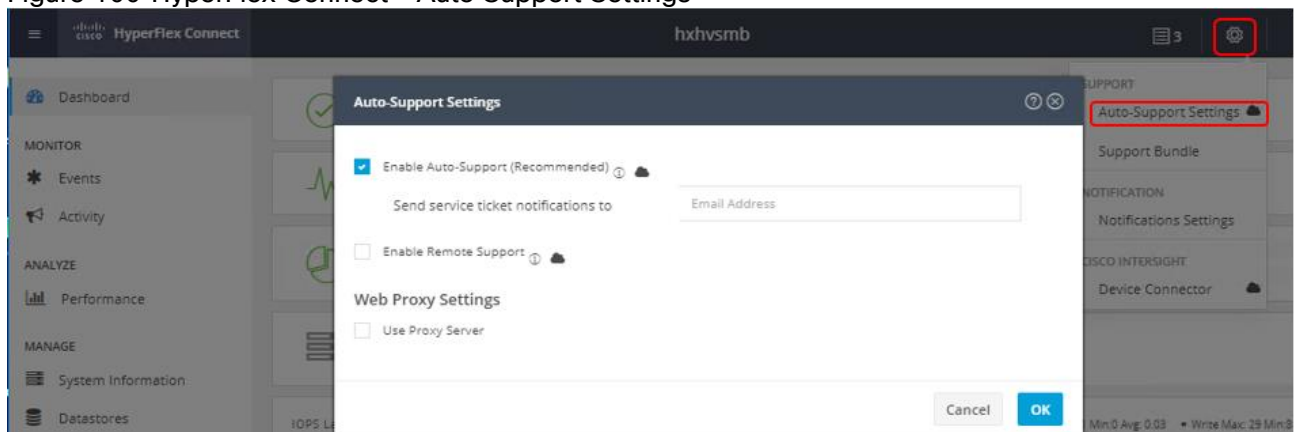
A list of events that automatically open a support ticket with Cisco TAC are as follows:

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

To change Auto-Support settings, follow these steps:

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

Figure 100 HyperFlex Connect – Auto Support Settings





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

To enable direct email notifications, follow these steps:

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

Figure 101 HyperFlex Connect – Notification Settings

The screenshot shows the 'Notifications Settings' dialog box. It contains the following fields and values:

- Send email notifications for alarms
- Mail Server Address:
- From Address:
- Recipient List (Comma separated):

Buttons: Cancel, OK

Smart Licensing

HyperFlex 2.5 and later utilizes Cisco Smart Licensing, which communicates with a Cisco Smart Account to validate and check out HyperFlex licenses to the nodes, from the pool of available licenses in the account. At the beginning, Smart Licensing is enabled but the HX storage cluster is unregistered and in a 90-day evaluation period or EVAL MODE. For the HX storage cluster to start reporting license consumption, it must be registered with the Cisco Smart Software Manager (SSM) through a valid Cisco Smart Account. Before beginning, verify that you have a Cisco Smart account, and valid HyperFlex licenses are available to be checked out by your HX cluster.

To create a Smart Account, go to [Create Smart Accounts](#).

To activate and configure smart licensing, follow these steps:

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

```
stcli license show status
```


Figure 102 Storage Controller Virtual Machine – View License

```

root@hxhvl1scvm:~# stcli license show status

Smart Licensing is ENABLED

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

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

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

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

Utility:
  Status: DISABLED

Transport:
  Type: TransportCallHome
root@hxhvl1scvm:~#

```



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

2. Navigate to Cisco Software Central (<https://software.cisco.com/>) and log in to your Smart Account.
3. From Cisco Smart Software Manager, generate a registration token.
4. In the License pane, click Smart Software Licensing to open Cisco Smart Software Manager.
5. Click Inventory.
6. From the virtual account where you want to register your HX storage cluster, click General, and then click New Token.
7. In the Create Registration Token dialog box, add a short Description for the token, enter the number of days you want the token to be active and available to use on other products, and check Allow export-controlled functionality on the products registered with this token.
8. Click Create Token.
9. From the New ID Token row, click the Actions drop-down list, and click Copy.

10. Log into a controller virtual machine.
11. Register your HX storage cluster, where `idtoken-string` is the New ID Token from Cisco Smart Software Manager.

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

12. Confirm that your HX storage cluster is registered.

```
# stcli license show summary
```

HyperFlex Cluster Expansion

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

Expansion with Converged Nodes

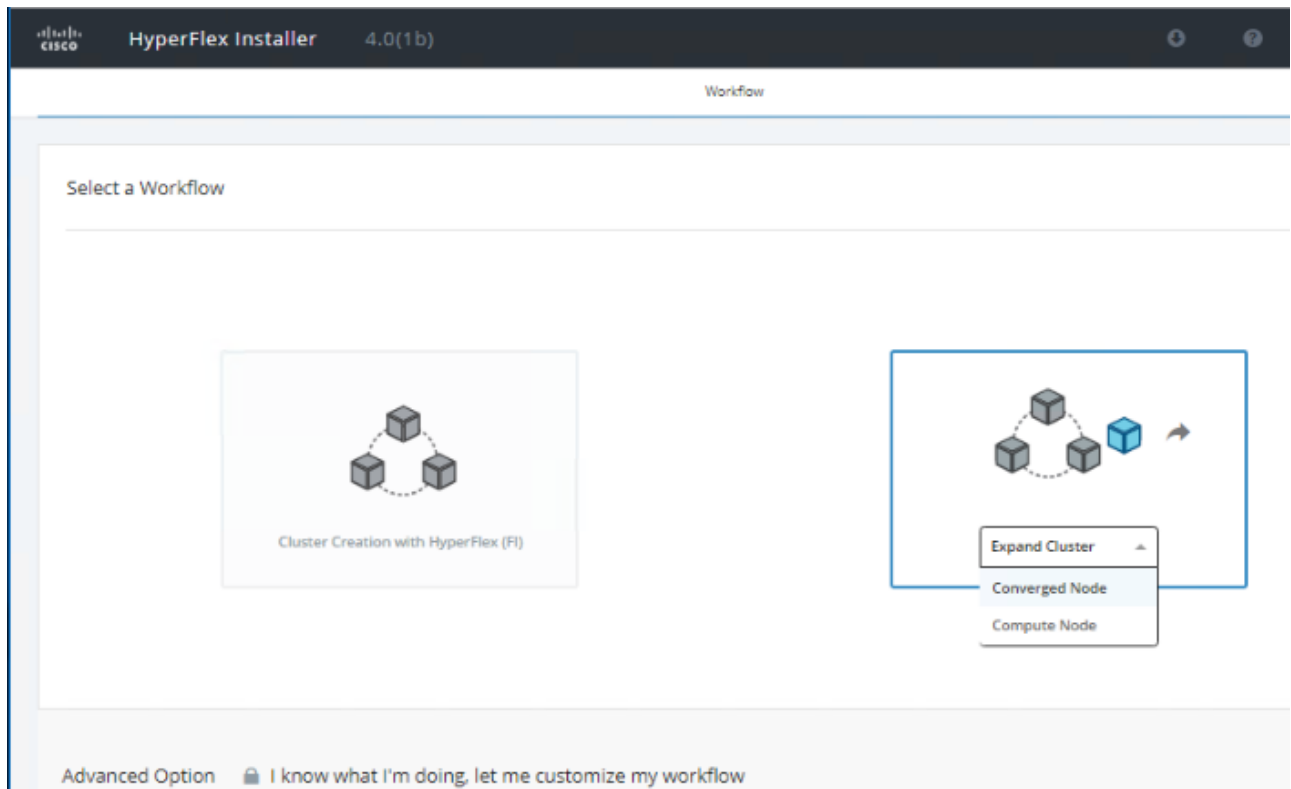
The HX installer has a wizard for Cluster Expansion with Converged Nodes. This procedure is very similar to the initial HyperFlex cluster setup. The following process assumes a new Cisco HX node has been ordered, therefore it is pre-configured from the factory with the proper hardware and firmware installed.

Prerequisites

Refer to [IP Addressing](#), [Prepopulate AD DNS with Records](#), and [Cabling](#) in the “Installation” section before continuing with the following steps.

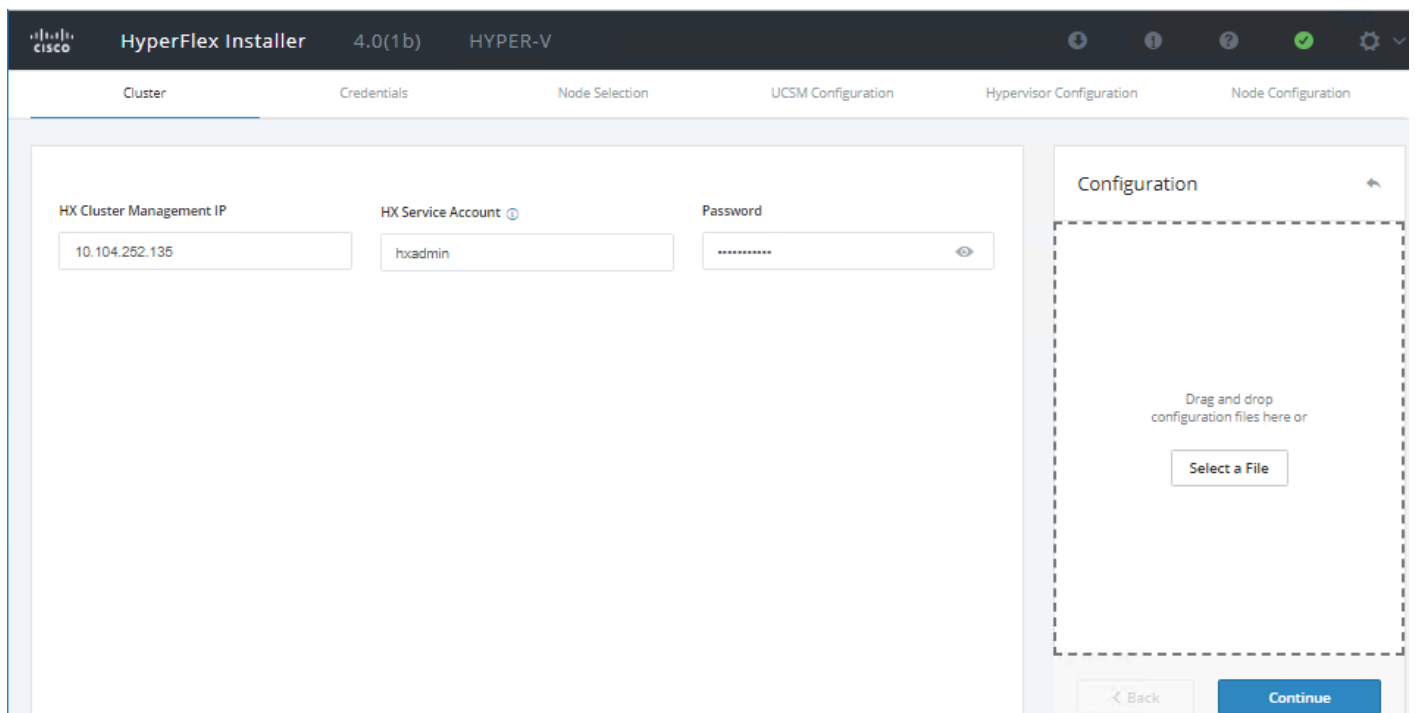
To add converged storage nodes to an existing HyperFlex cluster, follow these steps:

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

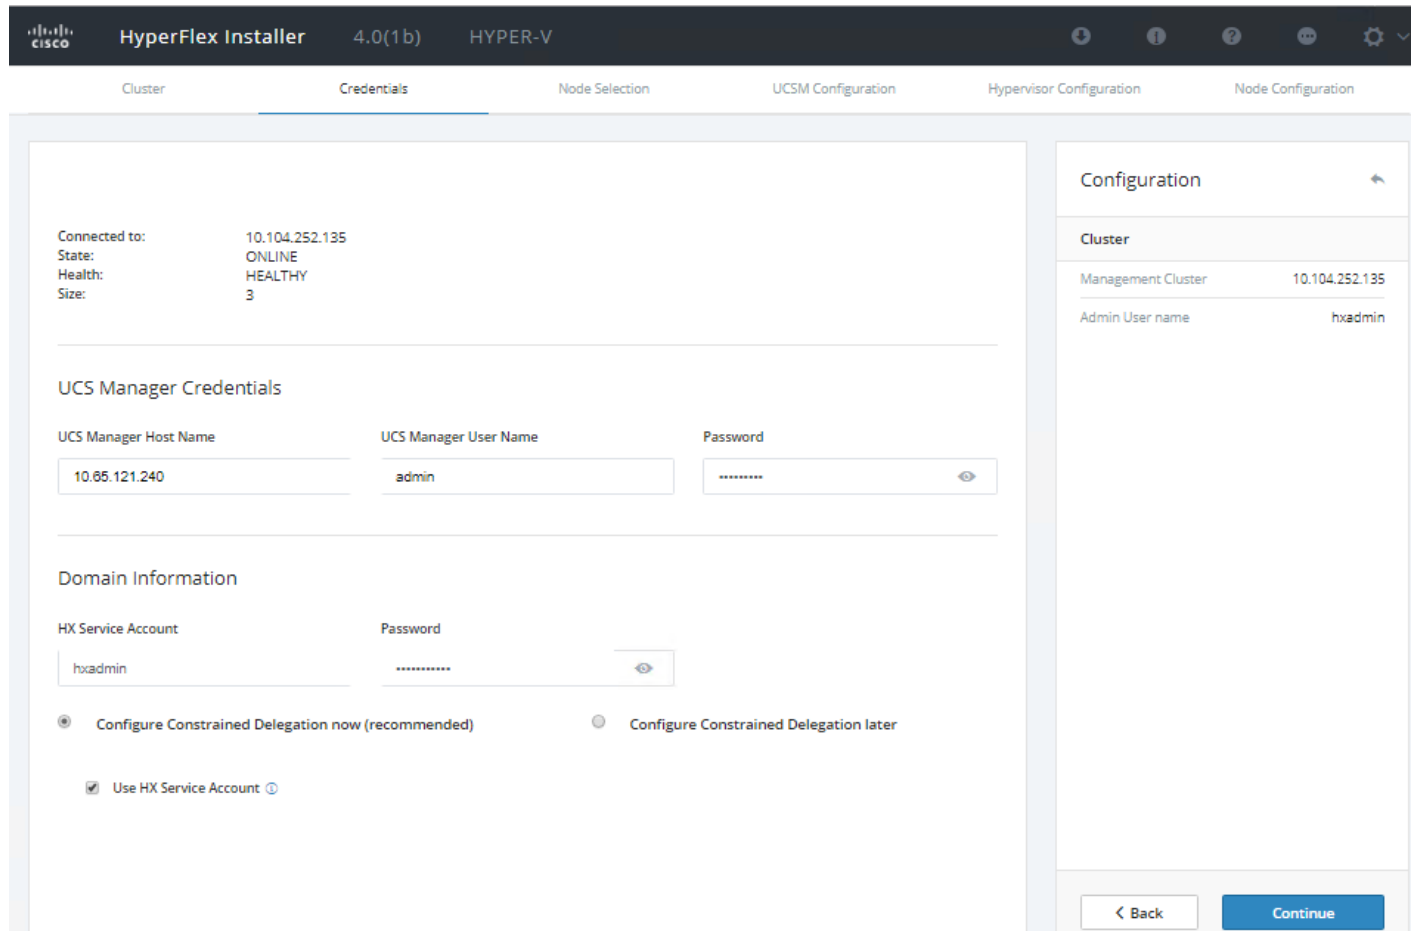


2. On the Cluster Page:

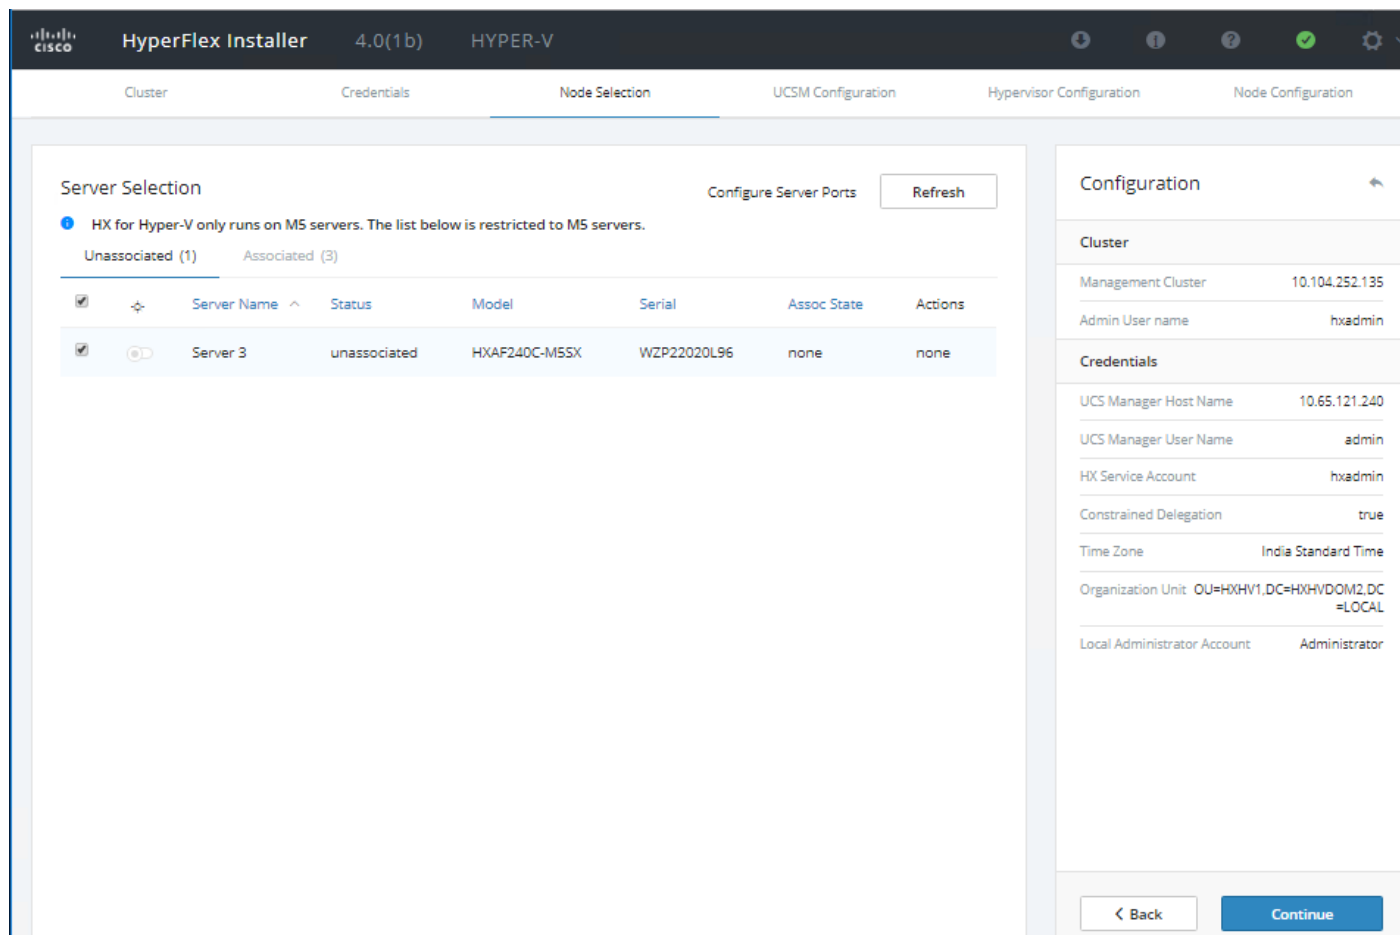
- a. Enter the HX Cluster Management IP address, Cluster Admin User name and password. You can select the option to see the passwords in clear text.
- b. Optionally, you can import a JSON file that has the configuration information, except for the appropriate passwords.



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



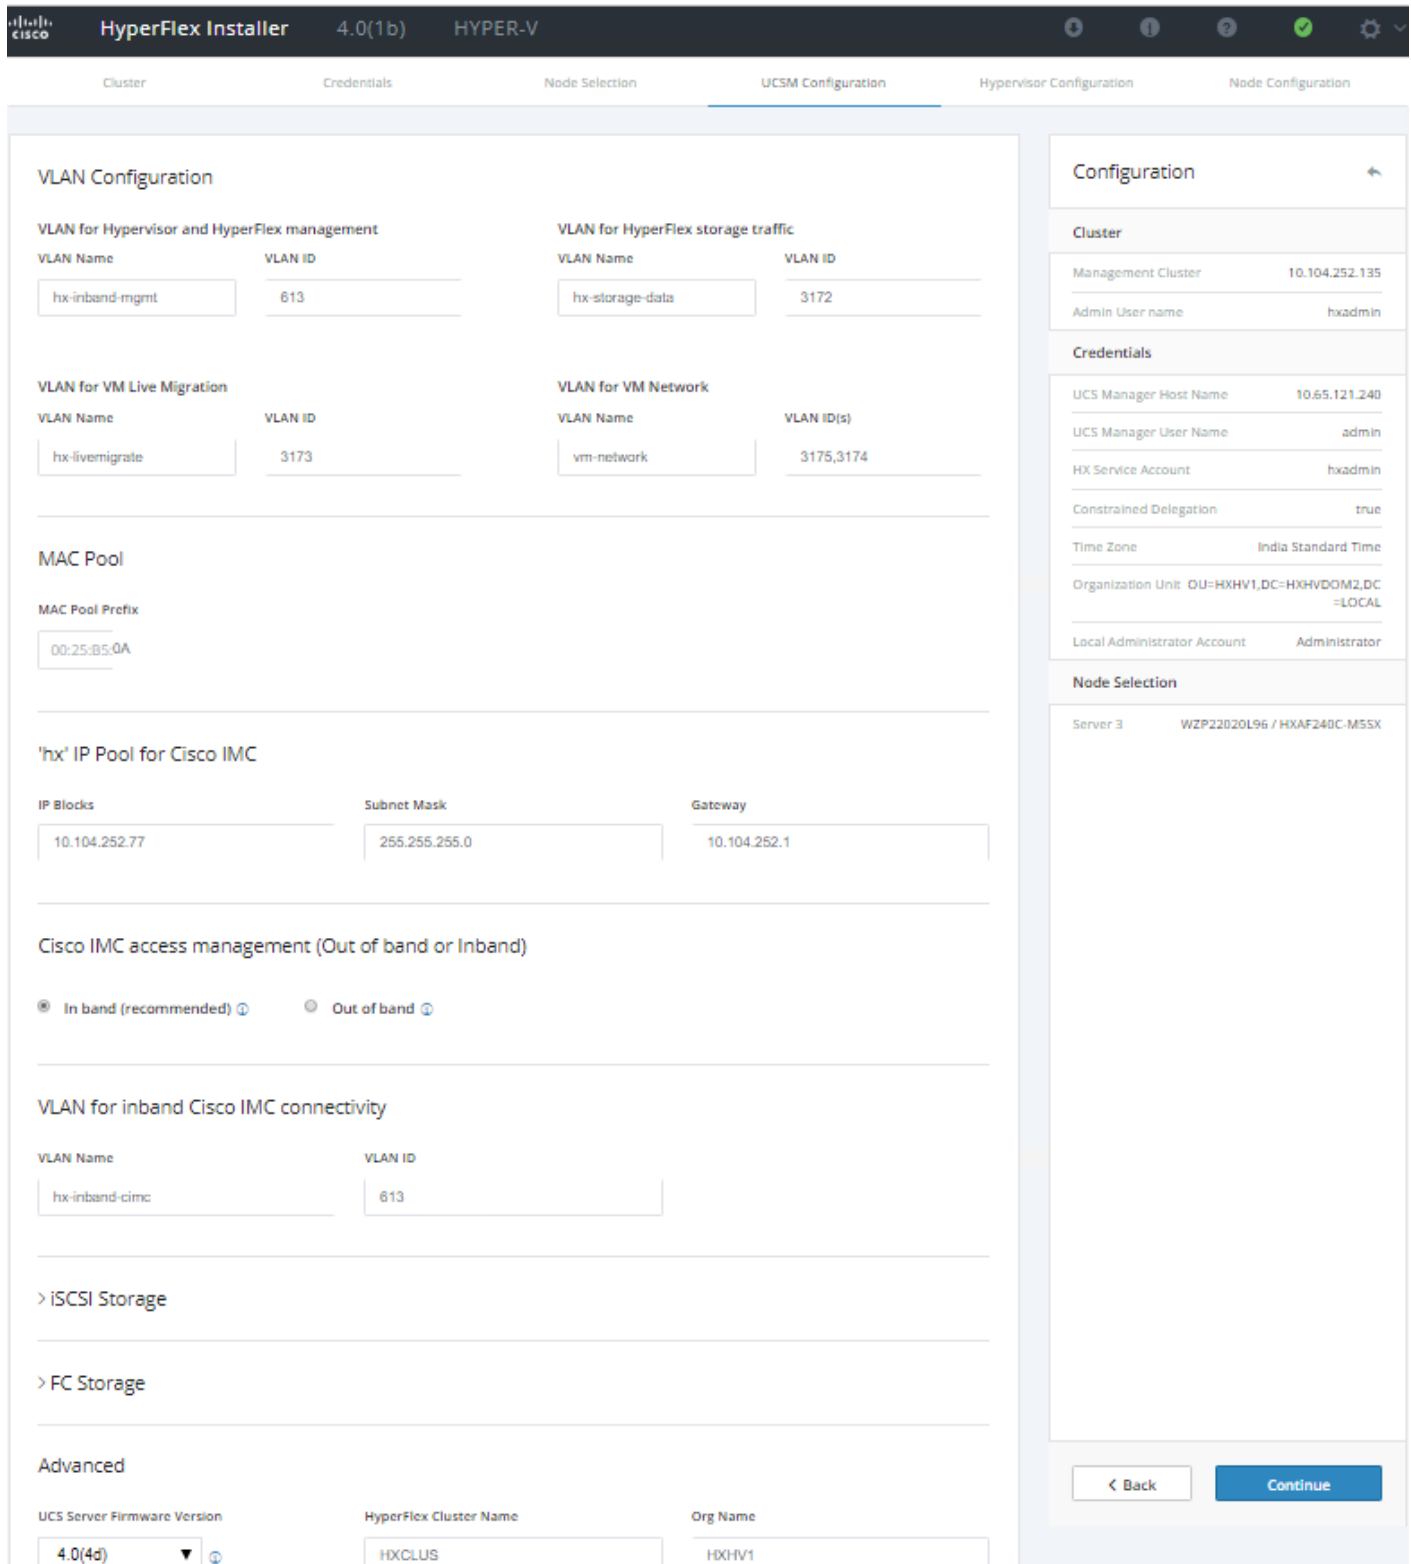
5. On the Node Selection page:
 - a. Select the Unassociated HX servers you want to add to the existing HX cluster.



6. Click Continue.

7. On the Cisco UCS Manager Configuration page:

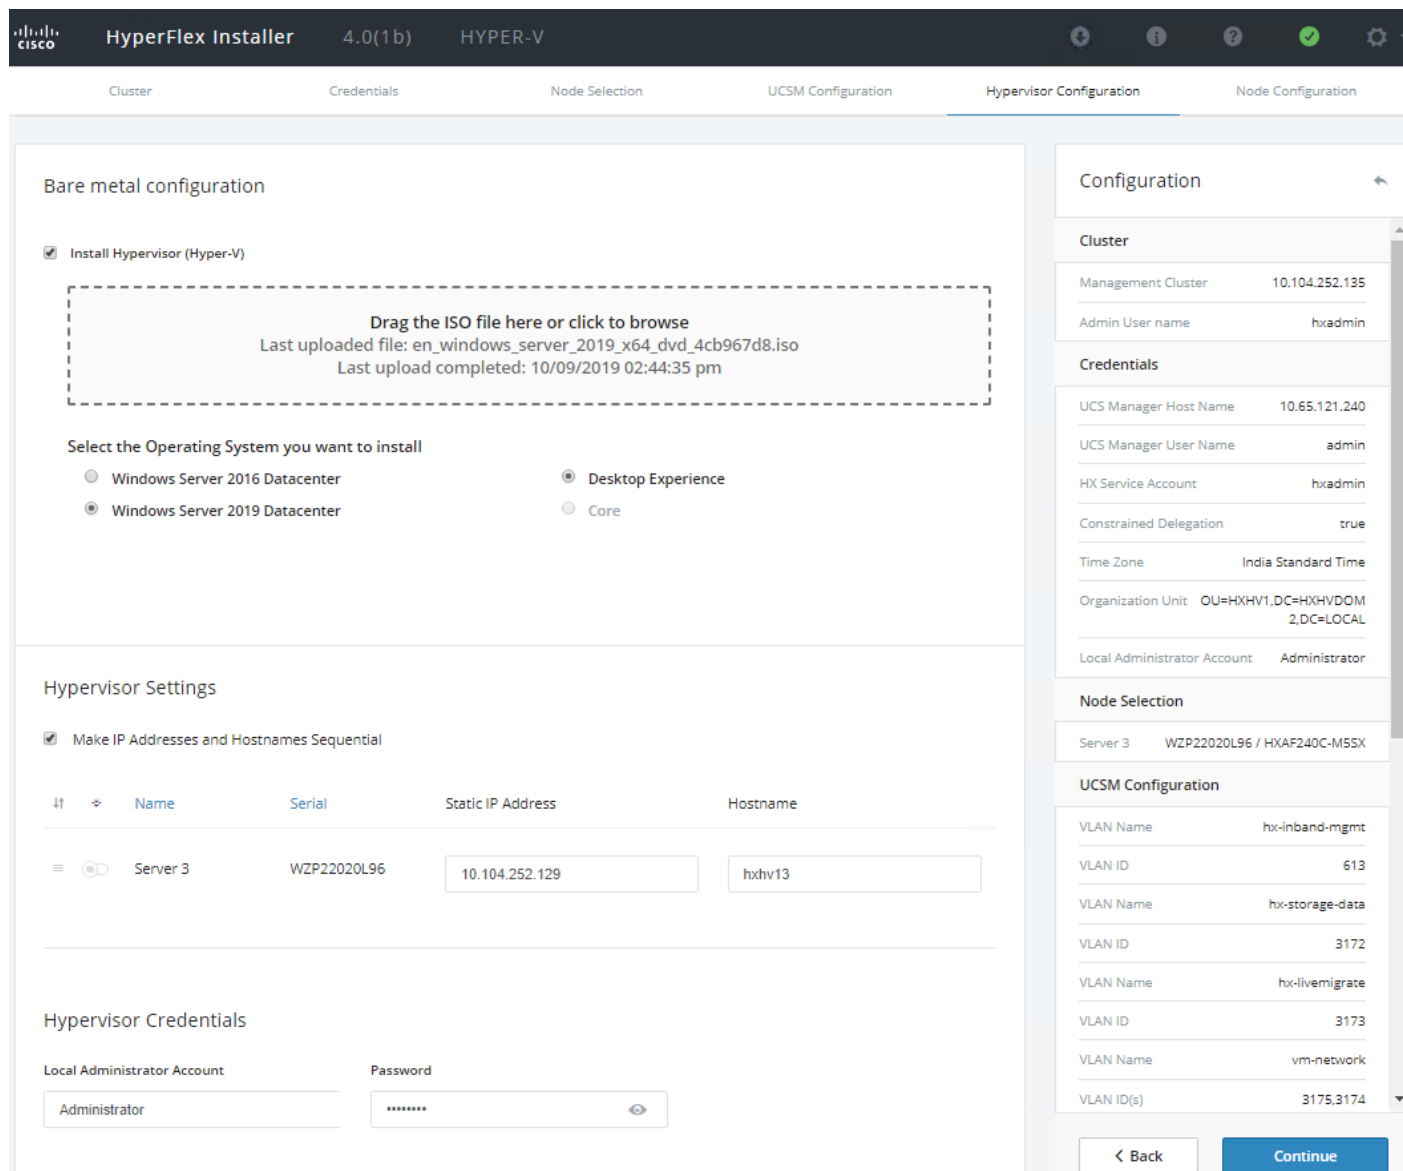
- a. VLAN Configuration – Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed here.
- c. MAC Pool – Enter the MAC Pool prefix, only enter the 4th byte value, for example: 00:25:B5:0A.
- d. 'hx' IP Pool for Cisco IMC – Enter the IP address range, subnet mask and gateway to be used by the CIMC interfaces of the servers in this HX cluster.
- e. Cisco IMC access Management (Out-of-band or inband) – Select the recommended 'in band' option for faster installation of hypervisor OS on all the hx nodes.
- f. The Out-Of-Band network needs to be on the same subnet as the Cisco UCS Manager. You can add multiple blocks of addresses as a comma separated line.
- g. VLAN for Inband Cisco IMC Connectivity – Enter a VLAN name and ID.
- h. iSCSI/FC Storage (optional) – iSCSI Storage and FC Storage are used for adding external storage to the HyperFlex cluster. Not defined for this setup.
- i. Advanced – If multiple firmware packages exist on the Fabric Interconnect, choose the version to be installed on the servers that will comprise this cluster. Enter a unique Org name for the HyperFlex Cluster.



8. Click Continue.

9. On the Hypervisor Configuration page:

- a. Bare Metal Configuration - If Windows Server 2016 is not installed on the nodes, select “Install Hypervisor (Hyper-V)”, drag or click Browse to upload OS media file in the box and select the radio button to choose OS you want to install.
- b. Hypervisor Settings - Enter IP addresses and hostnames for the Hypervisors that were created in the pre-installation section phase. The IP addresses will be assigned through Serial over Lan (SoL) through Cisco UCS Manager to the Hyper-V host systems as their management IP addresses.
- c. Hypervisor Credentials - this field is pre-populated and can't be edited.

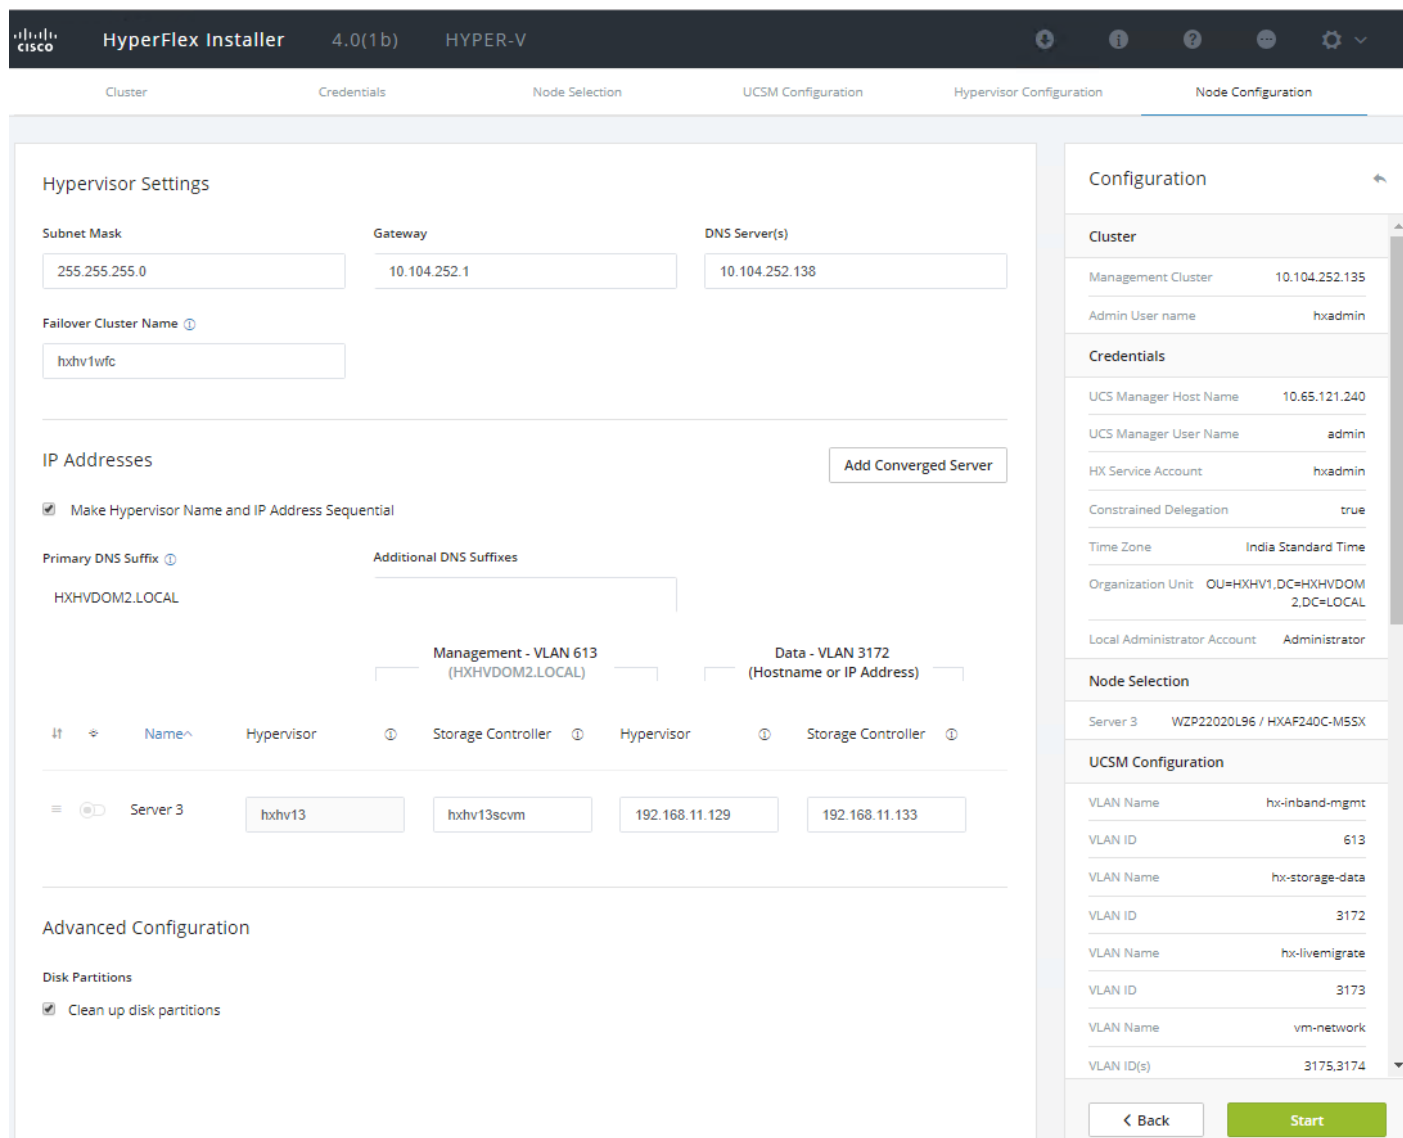


10. Click Continue.

11. On the Node Configuration Page:

- a. Hypervisor Settings - Here enter the Subnet Mask, Gateway, and DNS Server IP addresses.
- b. Failover Cluster Name - Enter the Windows failover cluster name.

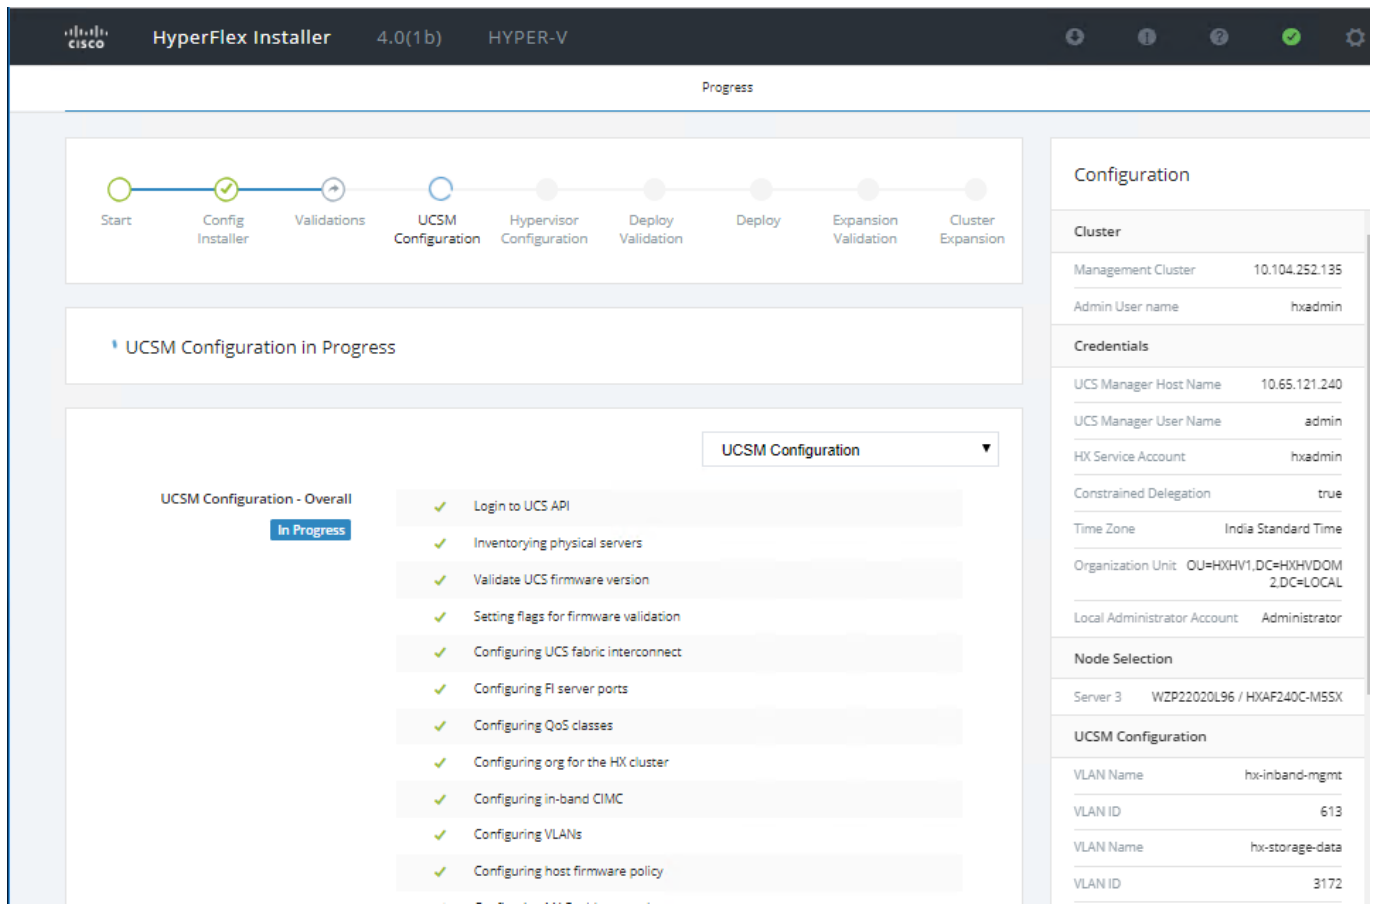
- c. IP Addresses - If you leave the checkbox, Make IP Addresses and Hostnames Sequential as checked, then the installer will automatically fill the rest of the servers sequentially.
- d. Assign the additional IP addresses for the Management and Data networks as well as the cluster IP addresses.
- e. (Optional) At this step you can manually add more servers for expansion if these servers already have service profiles associated and the hypervisor is ready, by clicking on Add Converged Server and then entering the IP addresses for the storage controller management and data networks.
- f. Advanced Configuration - Select Clean up disk partitions in case if the server has been used previously.



12. Click Start.

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

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



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

Cluster Name hxxhv1 smb **ONLINE** **HEALTHY**

Version	4.0.1b-33133	Domain Name	HXHVDOM2.LOCAL
Cluster Management IP Address	hxxhv1.cip.HXHVDOM2.LOCAL	Failover cluster Name	hxxhv1.wfc
Cluster Data IP Address	192.168.11.135	DNS Server(s)	10.104.252.138
Replication Factor	Three copies	NTP Server(s)	10.104.252.138
Available Capacity	10.7 TB		

Servers

Model	Serial Number	Management Hypervisor	Management Storage Controller	Data Network Hypervisor	Data Network Storage Controller
HXAF240C-M5SX	WZP22020L96	10.104.252.129	10.104.252.133	192.168.11.129	192.168.11.133
HXAF240C-M5SX	WZP22020L9B	10.104.252.130	10.104.252.134	192.168.11.130	192.168.11.134
HXAF240C-M5SX	WZP22020L9E	10.104.252.127	10.104.252.131	192.168.11.127	192.168.11.131
HXAF240C-M5SX	WZP220216WY	10.104.252.128	10.104.252.132	192.168.11.128	192.168.11.132

[Back to Workflow Selection](#)
[Launch HyperFlex Connect](#)

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

HyperFlex Connect | hxxhv1smb

Filter | Filter listed tasks

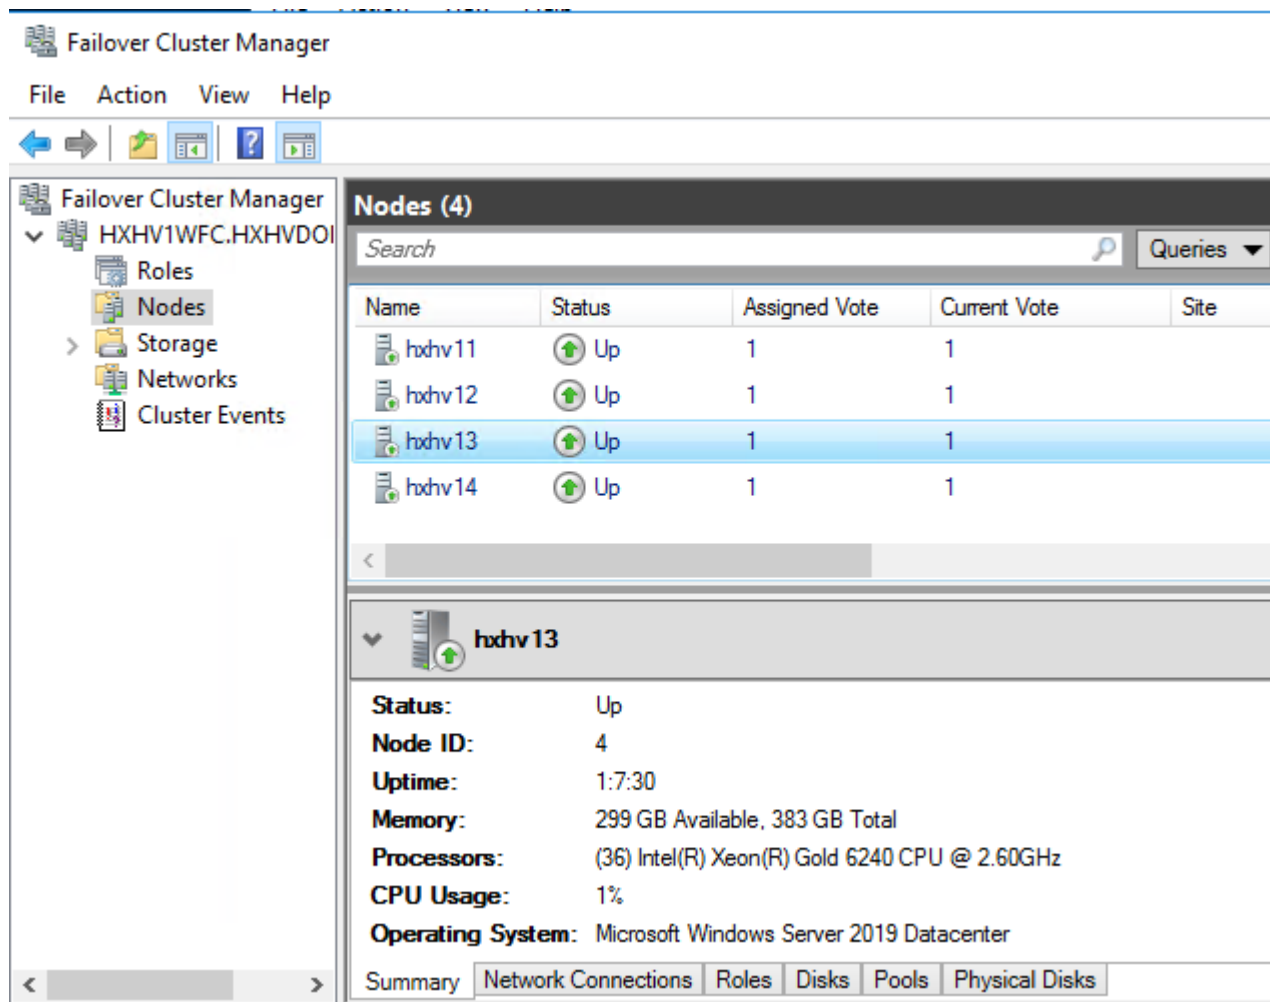
Activity | Monitor progress of recent tasks on the HX storage cluster. | Last refreshed at: 09/10/2019 6:17:33 PM

^ Collapse All

- Expand Hx Cluster
Status: Success
09/10/2019 5:44:07 PM
 - 192.168.11.133
 - ✓ Add to ZK Ensemble: Add to ZK ensemble DAD5F723-20E4-CA4B-A44F-AE14F03C48A9
 - ✓ HxCluster node join: HxCluster node join
 - 192.168.11.133
 - ✓ Initializing hxMgmtService: Initializing hxMgmtService
- Validating if we can Expand a Cluster.
Status: Success
09/10/2019 5:43:26 PM
 - 192.168.11.133
 - ✓ Validate Controller Nodes: Validate Controller Nodes
 - ✓ Check ZooKeeper Ensemble: Check ZooKeeper Ensemble
 - ✓ Check ControllerVM Version: Check ControllerVM Version
- Create Hx Cluster
Status: Success
08/29/2019 3:18:28 AM
 - hxxhv1smb
 - ✓ ZK ensemble: HxCluster ZK ensemble
 - ✓ Persist Encrypted HyperV Credentials: Persist Encrypted HyperV Credentials
 - ✓ Init Management Service: HxCluster Init Management Service
 - ✓ Storage HxCluster: Storage HxCluster
 - ✓ System Datastore: System Datastore
 - ✓ Firewall Configuration: Firewall Configuration
 - hxxhv11scvm.HXHVDO...
 - ✓ HxCluster node join: HxCluster node join
 - hxxhv12scvm.HXHVDO...
 - ✓ HxCluster node join: HxCluster node join
 - hxxhv14scvm.HXHVDO...
 - ✓ HxCluster node join: HxCluster node join

The screenshot displays the Cisco HyperFlex Connect interface for a system named 'hxhv1smb'. The interface is divided into a left-hand navigation menu and a main content area. The navigation menu includes sections for 'MONITOR' (Alarms, Events, Activity), 'ANALYZE' (Performance), and 'MANAGE' (System Information, Datastores, Upgrade). The main content area shows a 'System Overview' for 'hxhv1smb', which is currently 'ONLINE'. Key system details include an uptime of 12 days, 15 hours, 0 minutes, and 42 seconds, running on a Microsoft Windows Server 2019 Datacenter Hypervisor with HXDP Version 4.0.1b-33133. Capacity information shows a total of 10.71 TB and an available capacity of 10.61 TB, with a data replication factor of 3. DNS and NTP servers are listed as 10.104.252.138, and SSH access is enabled. Below this, four nodes are listed: hxhv11, hxhv12, hxhv13, and hxhv14. Each node is 'Online' and has 11 disks (1 Caching, 10 Persistent). The nodes are connected to HXAF240C-M5SX storage. Each node's status is accompanied by a bar chart showing disk health and an 'Hypervisor Status' section indicating 'Online' and the hypervisor address (10.104.252.127, 10.104.252.128, and 10.104.252.129 respectively).

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



17. For more information, refer to section [Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces](#).

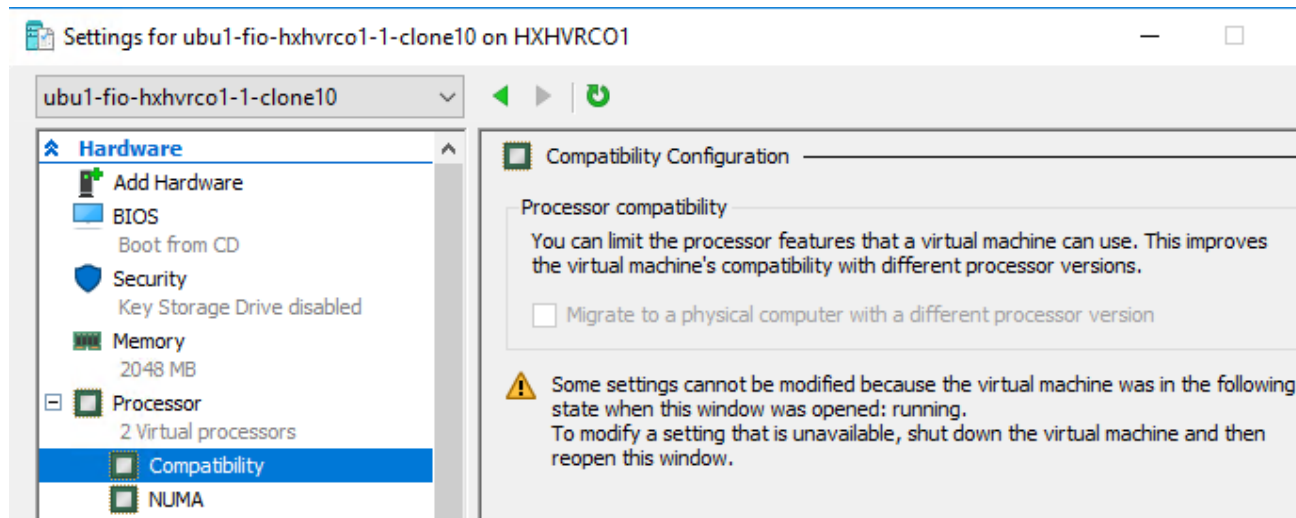
Expansion with Compute-Only Nodes

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

- The number of compute-only nodes cannot exceed the number of HyperFlex converged nodes within a single HyperFlex cluster.
- The Cisco UCS infrastructure firmware revision, which provides the firmware for Cisco UCS Manager and the Fabric Interconnects, must be maintained at the minimum version required for the HyperFlex converged nodes, or higher, at all times.
- The version of Microsoft Hyper-V installed on the compute-only nodes must be compatible with the Cisco HyperFlex version in use, and it must match the version installed on the HyperFlex converged nodes.
- While the CPU models and memory capacities between the compute-only nodes and the HyperFlex converged nodes do not have to match, configuring the nodes to have similar capacities is recommended.
- Care must be taken that the addition of the compute-only nodes will not significantly impact the HyperFlex cluster by creating additional load, or by consuming too much space. Pay close attention to the space

consumption and performance requirements of any net-new virtual machines that will run on the additional compute-only nodes and note the current cluster performance and space utilization. If no new virtual machines will be created, then the current cluster performance will not be impacted.

- Mixing different models of compute-only nodes is allowed within the same cluster. Example: using Cisco UCS B220 M4, B200 M5 and C220 M5 servers as compute-only nodes is allowed.
- Connectivity between compute-only nodes and the HyperFlex cluster must be within the same Cisco UCS domain and must be 10 GbE or 40 GbE. Connecting compute-only nodes from a different Cisco UCS domain is not allowed and connecting standalone rack-mount servers from outside of the Cisco UCS domain is not allowed.
- Blade servers installed in the Cisco UCS 5108 Blade Chassis can connect through 10 GbE or 40 GbE chassis links, using the Cisco UCS 2204XP, 2208XP, or 2304 model Fabric Extenders. The Fabric Extenders, also called I/O Modules (IOMs), are typically installed in pairs, and connect the 5108 chassis to the Fabric Interconnects, which provide all the networking and management for the blades. Care must be taken not to oversubscribe and saturate the chassis links.
- Mixing CPU generations will require configuring Processor Compatibility in order to allow Live Migration to work between the compute-only nodes and the converged nodes. Enabling Processor Compatibility typically requires all virtual machines to be powered off. If it is known ahead of time that Processor Compatibility will be needed, then it is easier to enable it on the Hyper-V Manager prior to installing HyperFlex as shown below:

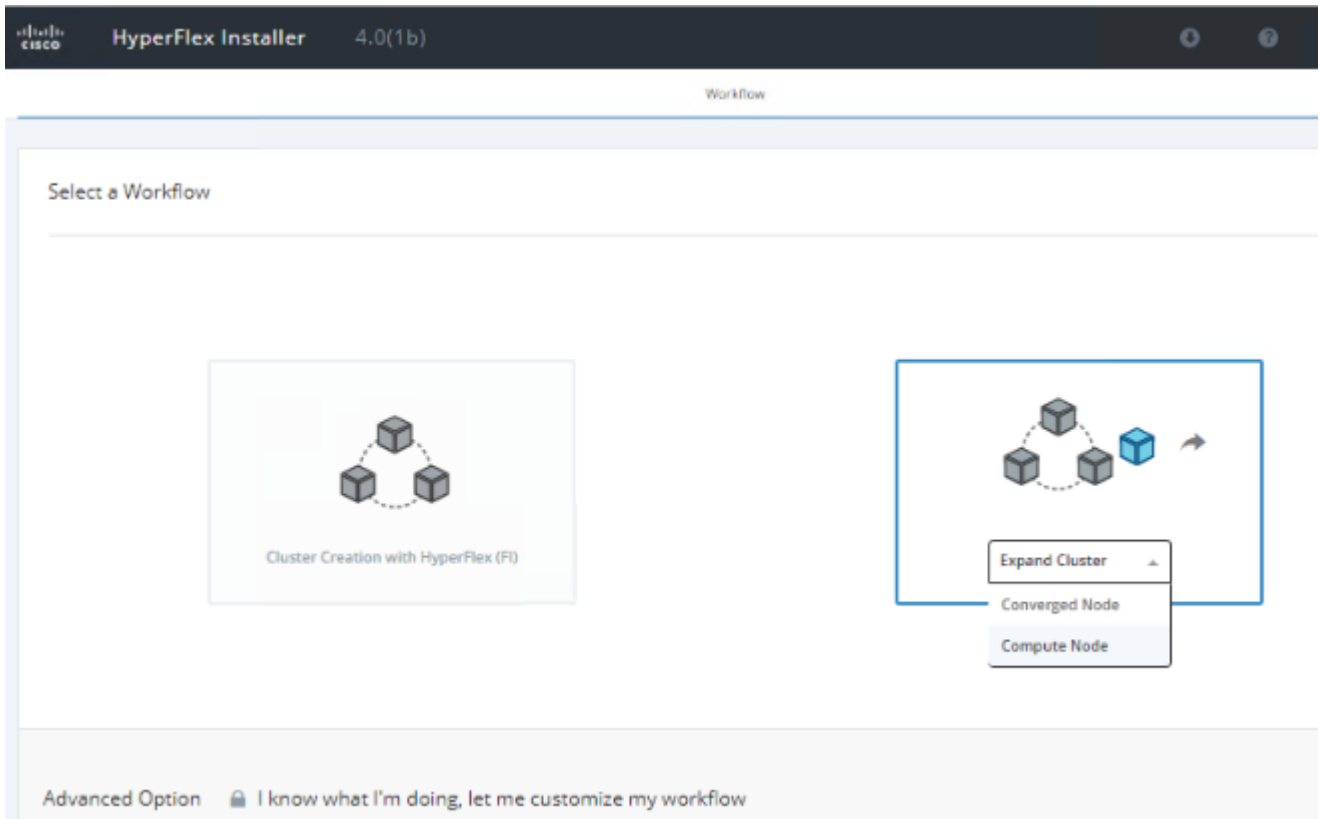


- Compute-only nodes can be configured to boot from SAN or local disks. No other internal storage should be present in a compute-only node. Manual configuration of the boot policy will be necessary if booting from any device other than M.2 drive.

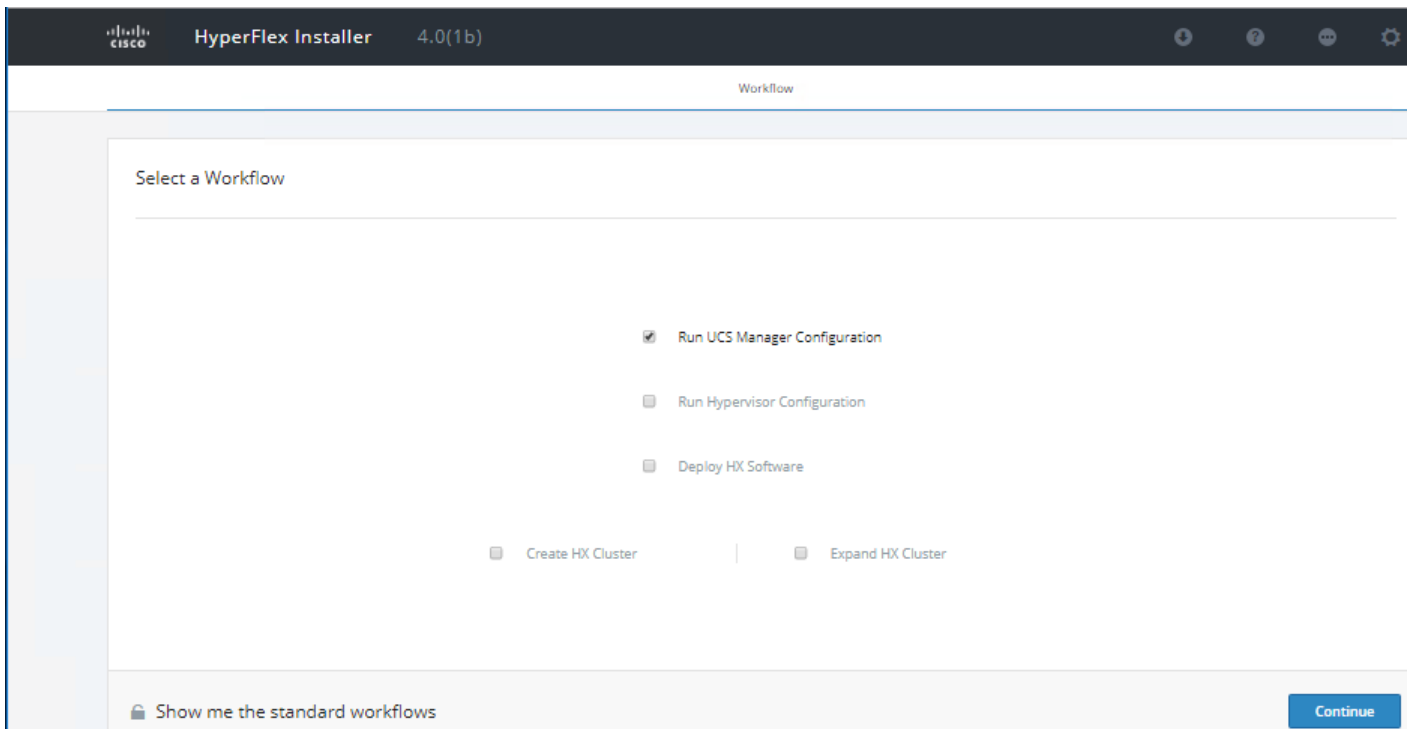
Prerequisites

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

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



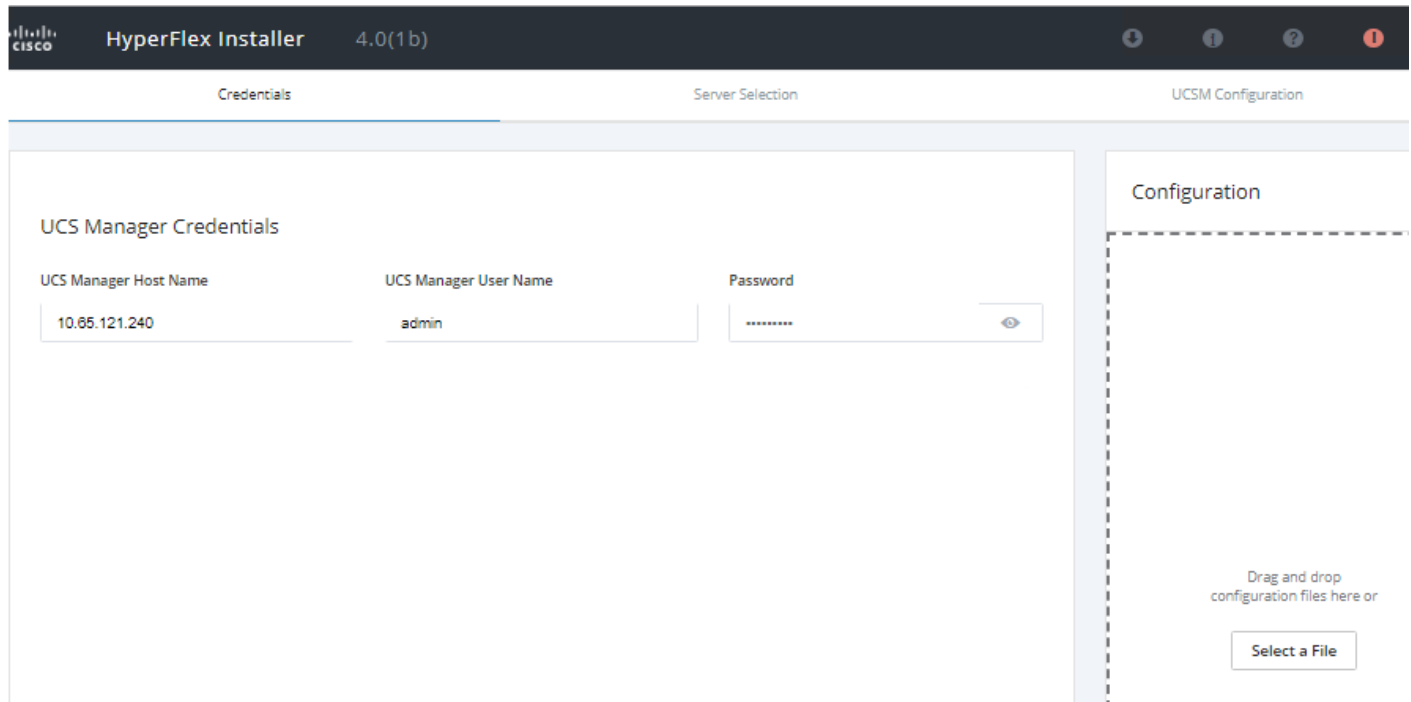
2. Select "Run UCS Manager Configuration" as shown below:



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

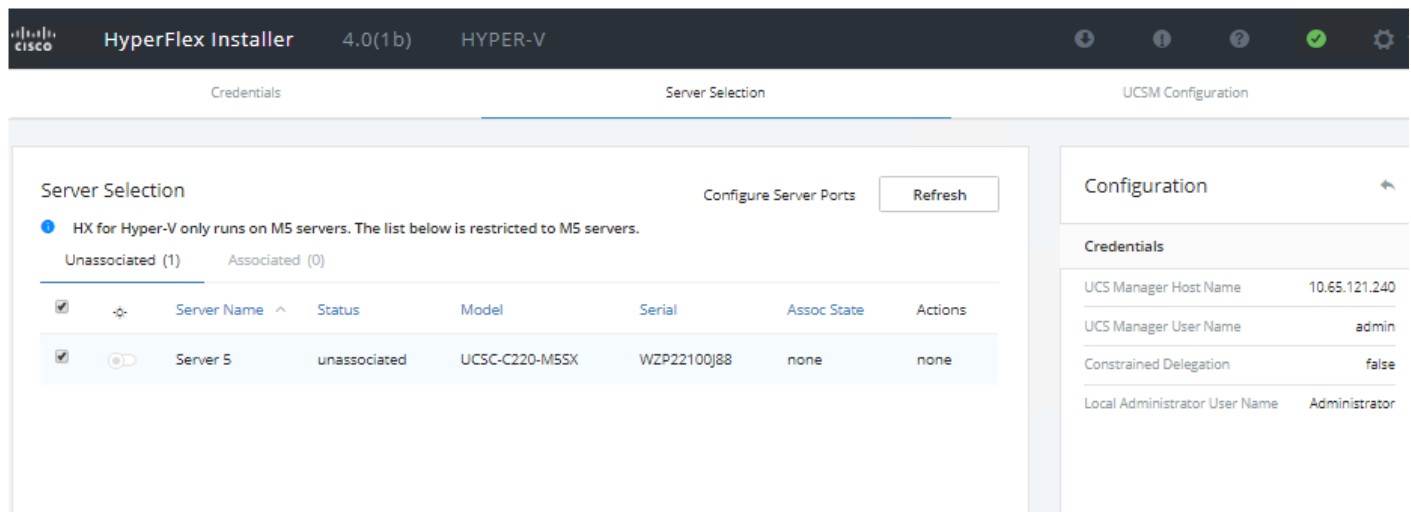
4. On the Credentials Page:

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



5. On the Server Selection page:

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



6. Click Continue.

7. On the Cisco UCSM Configuration page:

- VLAN Configuration – Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed.
- MAC Pool – Enter the MAC Pool prefix, only enter the 4th byte value, for example: 00:25:B5:0A.

- c. 'hx' IP Pool for Cisco IMC - Enter the IP address range, subnet mask and gateway to be used by the CIMC interfaces of the servers in this HX cluster.
- d. Cisco IMC access Management (Out of band or inband) - Select the recommended 'in band' option for faster installation of hypervisor OS on all the hx nodes.
- e. The Out-Of-Band network needs to be on the same subnet as the Cisco UCS Manager. You can add multiple blocks of addresses as a comma separated line.
- f. VLAN for Inband Cisco IMC Connectivity - Enter a VLAN name and ID.
- g. iSCSI/FC Storage (optional) - iSCSI Storage and FC Storage are used for adding external storage to the HyperFlex cluster. Not defined for this setup.
- h. Advanced - If multiple firmware packages exist on the Fabric Interconnect, choose the version to be installed on the servers that will comprise this cluster. Enter a unique Org name for the HyperFlex Cluster.

The screenshot displays the HyperFlex Installer 4.0(1b) HYPER-V interface during the UCSM Configuration step. The main configuration area is split into several sections:

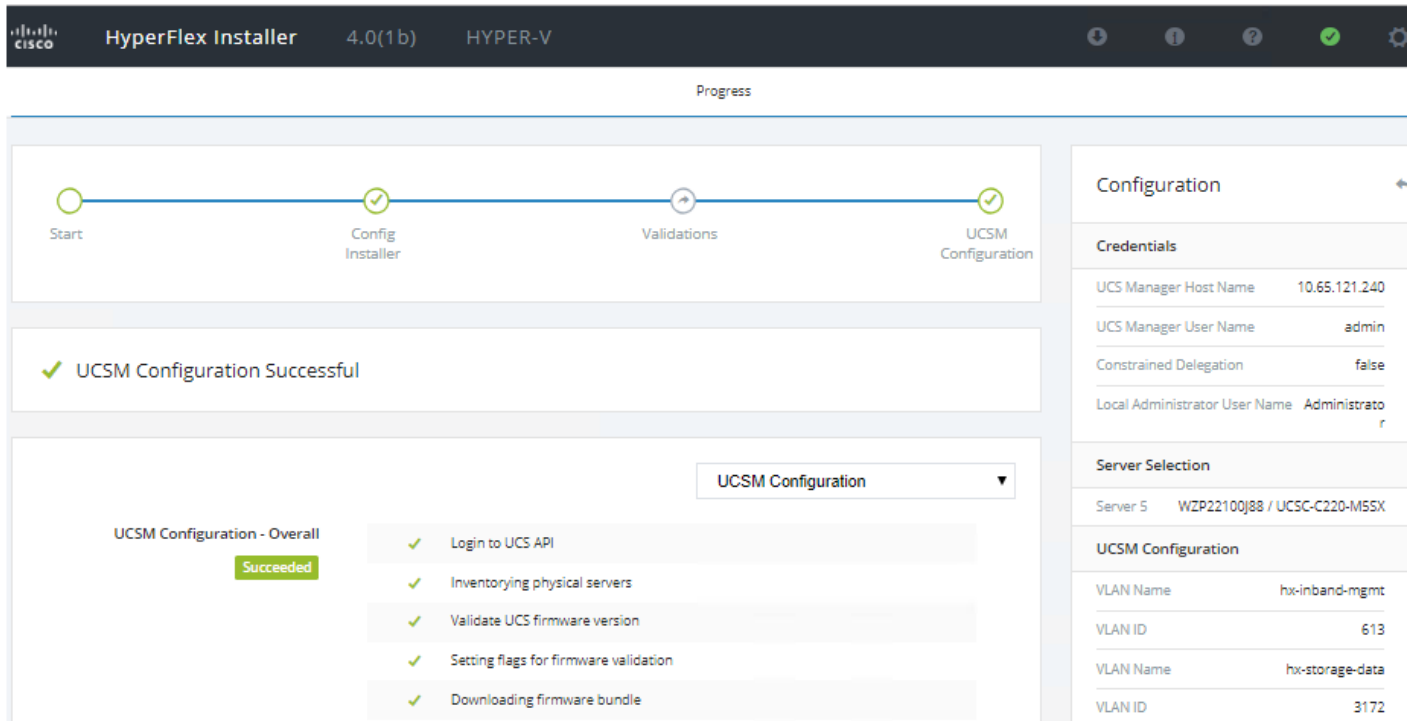
- VLAN Configuration:**
 - VLAN for Hypervisor and HyperFlex management:** VLAN Name: hx-inband-mgmt, VLAN ID: 613
 - VLAN for HyperFlex storage traffic:** VLAN Name: hx-storage-data, VLAN ID: 3172
 - VLAN for VM Live Migration:** VLAN Name: hx-livemigrate, VLAN ID: 3173
 - VLAN for VM Network:** VLAN Name: vm-network, VLAN ID(s): 3174,3175
- MAC Pool:** MAC Pool Prefix: 00:25:B5:0A
- 'hx' IP Pool for Cisco IMC:** IP Blocks: 10.104.252.81, Subnet Mask: 255.255.255.0, Gateway: 10.104.252.1
- Cisco IMC access management (Out of band or Inband):** Radio buttons for In band (recommended) and Out of band.
- VLAN for inband Cisco IMC connectivity:** VLAN Name: hx-inband-cimc, VLAN ID: 613
- > iSCSI Storage**
- > FC Storage**
- Advanced:** UCS Server Firmware Version: 4.0(4d), HyperFlex Cluster Name: HXCLUS, Org Name: HXHV1

The right-hand panel, titled **Configuration**, provides a summary of the settings:

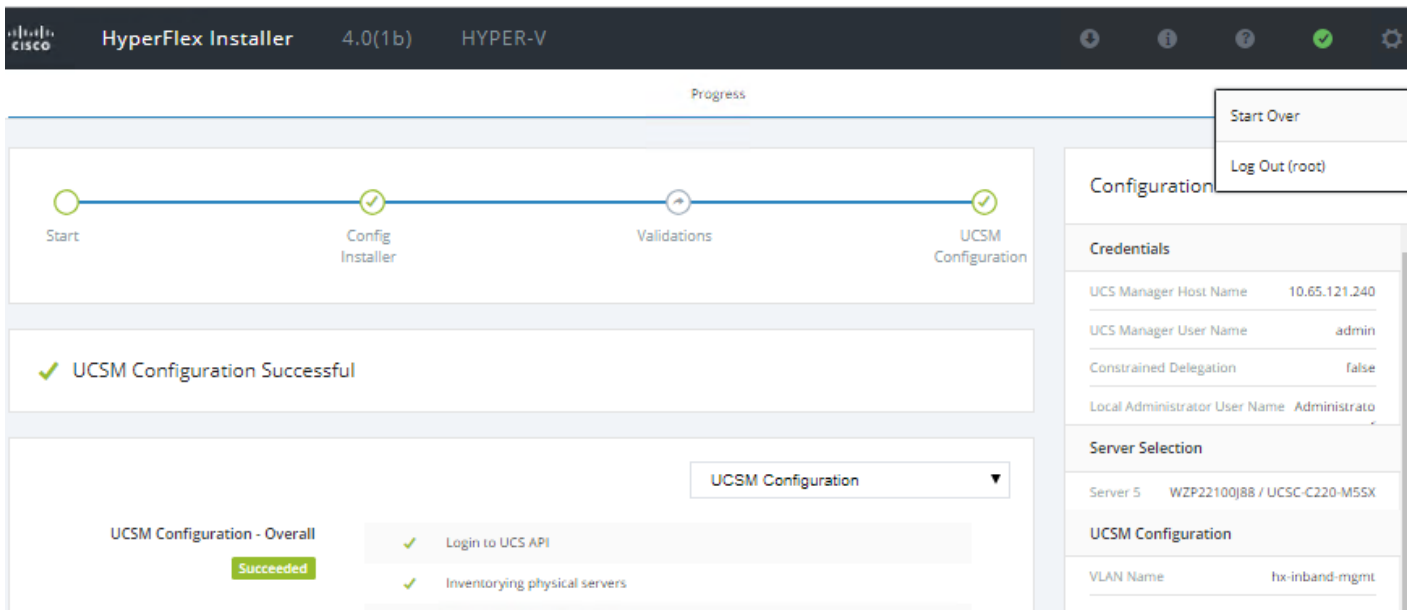
- Credentials:** UCS Manager Host Name: 10.65.121.240, UCS Manager User Name: admin, Constrained Delegation: false, Local Administrator User Name: Administrator
- Server Selection:** Server 5: WZP22100J88 / UCSC-C220-M55X
- UCSM Configuration:** A list of parameters including VLAN Name, VLAN ID, IP Blocks, Subnet Mask, Gateway, and various storage-related settings (VLAN A/B Name, FC Storage, WWxN Pool, VSAN A/B Name).

At the bottom of the right panel are **Back** and **Continue** buttons.

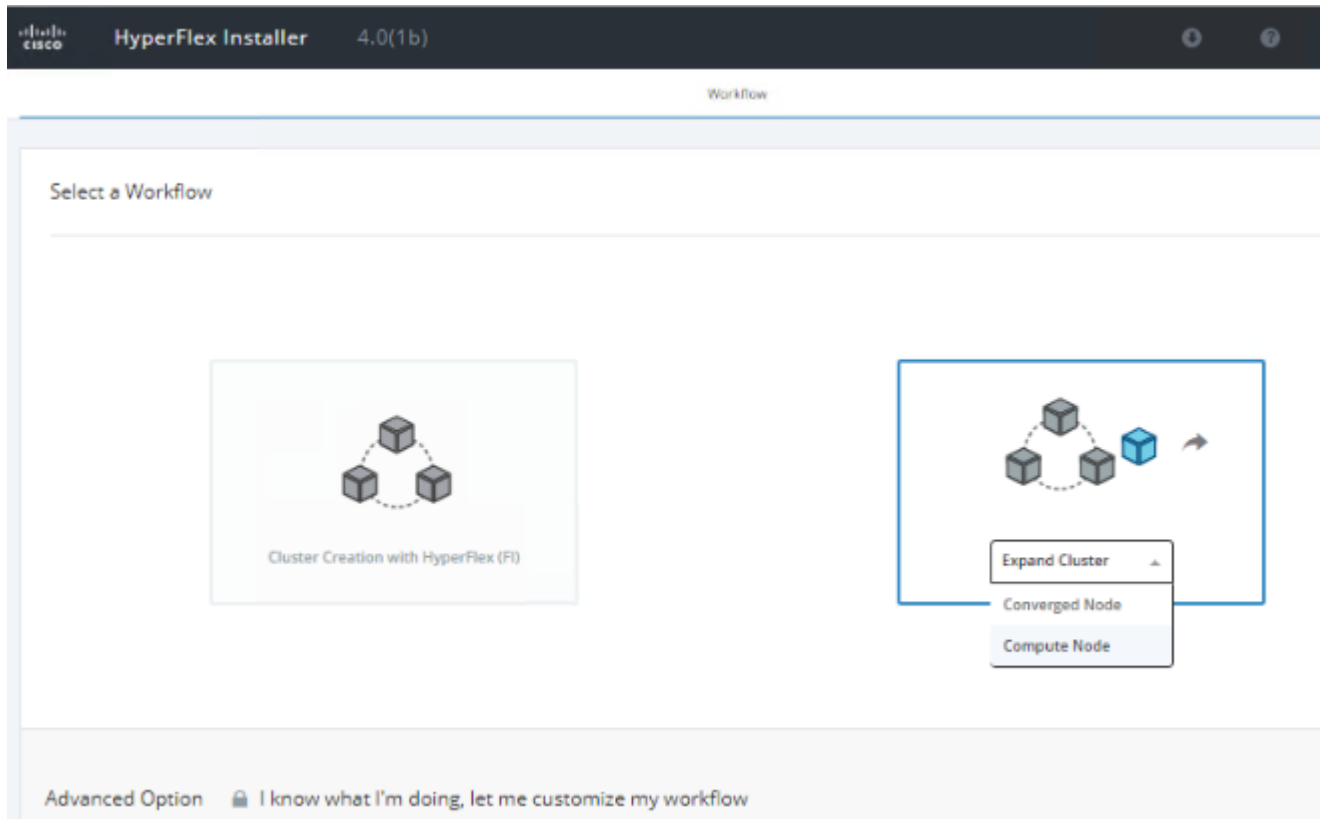
8. Click Start.
9. Validation of the configuration will now start. After validation, the installer will create the compute-only node service profiles and associate them with the selected servers.



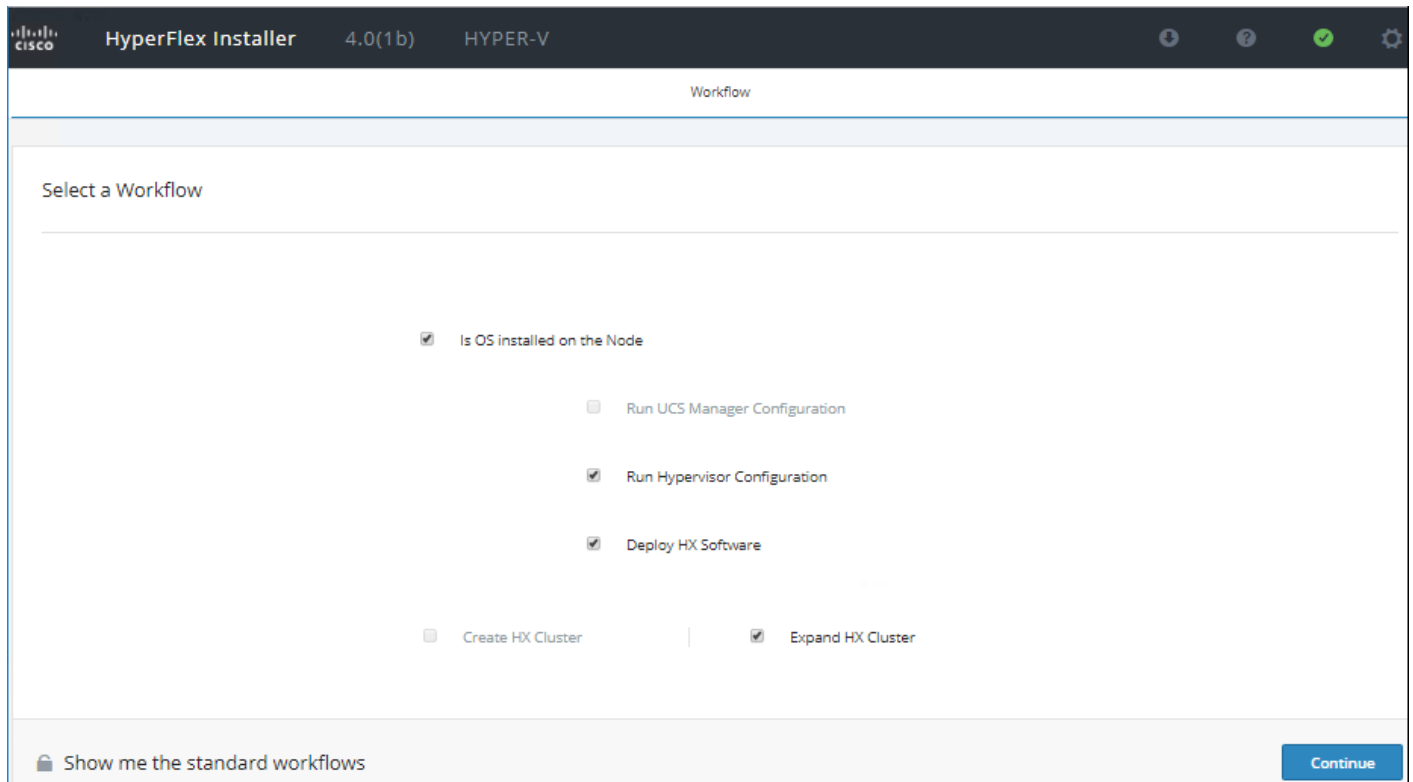
10. If the Hyper-V hypervisor has not been previously installed on the compute-only nodes, refer to the Appendix, section B: [Install Microsoft Windows Server 2016](#).
11. If the hypervisor is already installed, then continue with the next steps in this section.
12. You might need to “start over” since the previous workflow was finished. Click the gear icon in the top right corner and select Start Over.



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



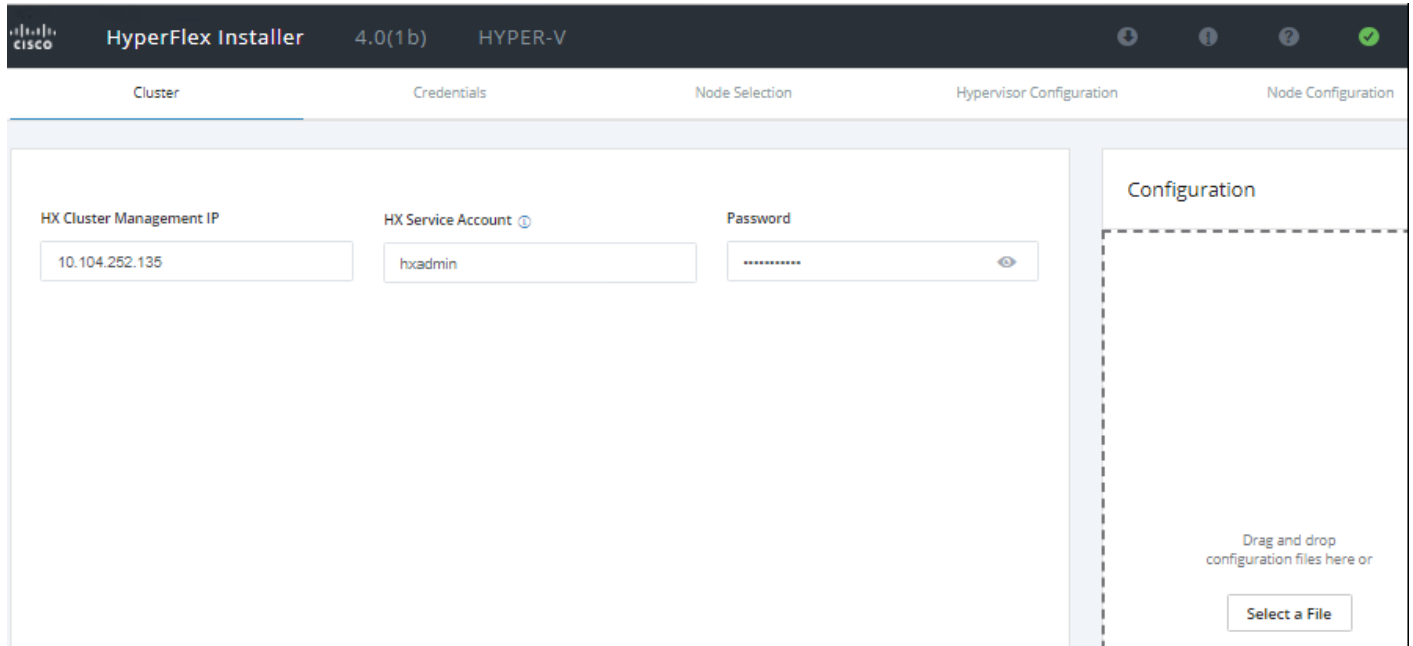
14. On the Select a Workflow page, select "Is OS installed on the node". With this selection, Run Hypervisor Configuration, Deploy HX Software and Expand Cluster are also automatically selected.



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

16. On the Cluster Page:

- a. Enter the HX Cluster Management IP address, Cluster Admin User name and password. You can select the option to see the passwords in clear text.
- b. Optionally, you can import a JSON file that has the configuration information, except for the appropriate passwords.



17. Click Continue.

18. On the Credentials Page:

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

HyperFlex Installer 4.0(1b) HYPERS-V

Cluster Credentials Node Selection Hypervisor Configuration Node Configuration

Connected to: 10.104.252.135
State: ONLINE
Health: HEALTHY
Size: 4

UCS Manager Credentials

UCS Manager Host Name: 10.65.121.240
UCS Manager User Name: admin
Password:

Domain Information

HX Service Account: hxadmin
Password:

Configure Constrained Delegation now (recommended) Configure Constrained Delegation later

Use HX Service Account

Configuration

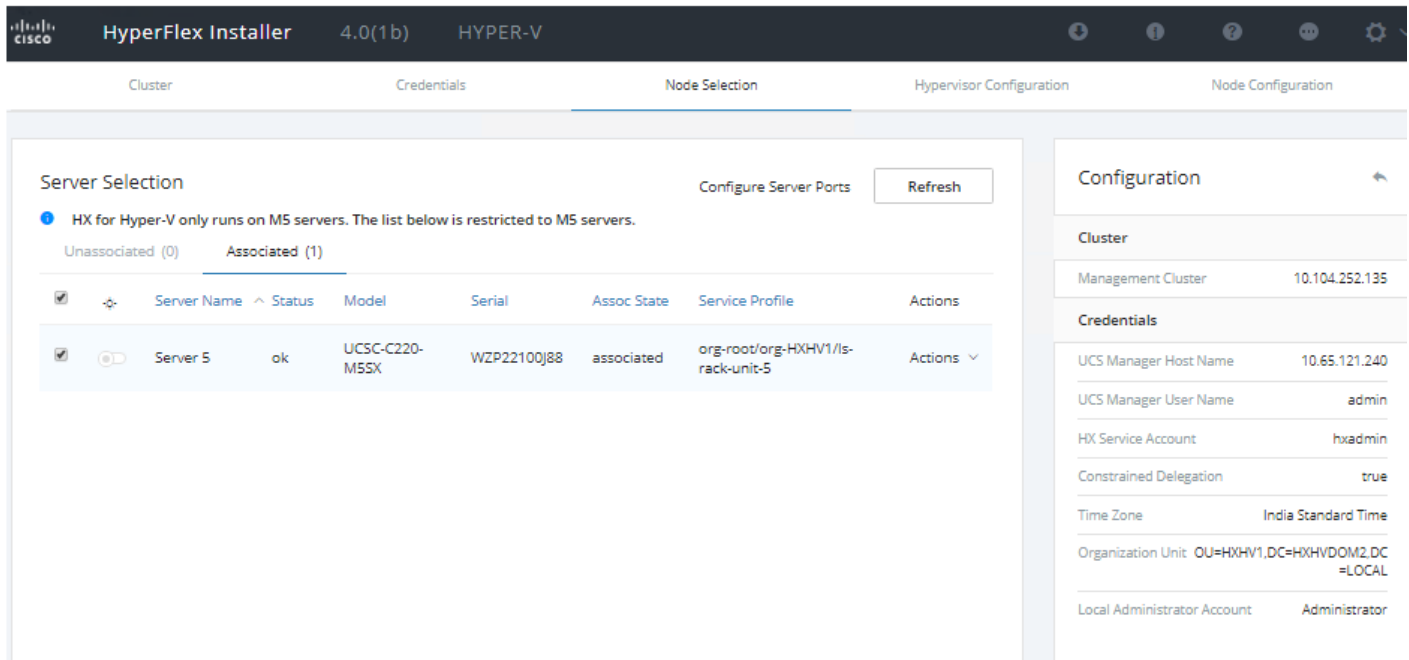
Cluster

Management Cluster 10.104.252.135

< Back Continue

19. On the Node Selection page:

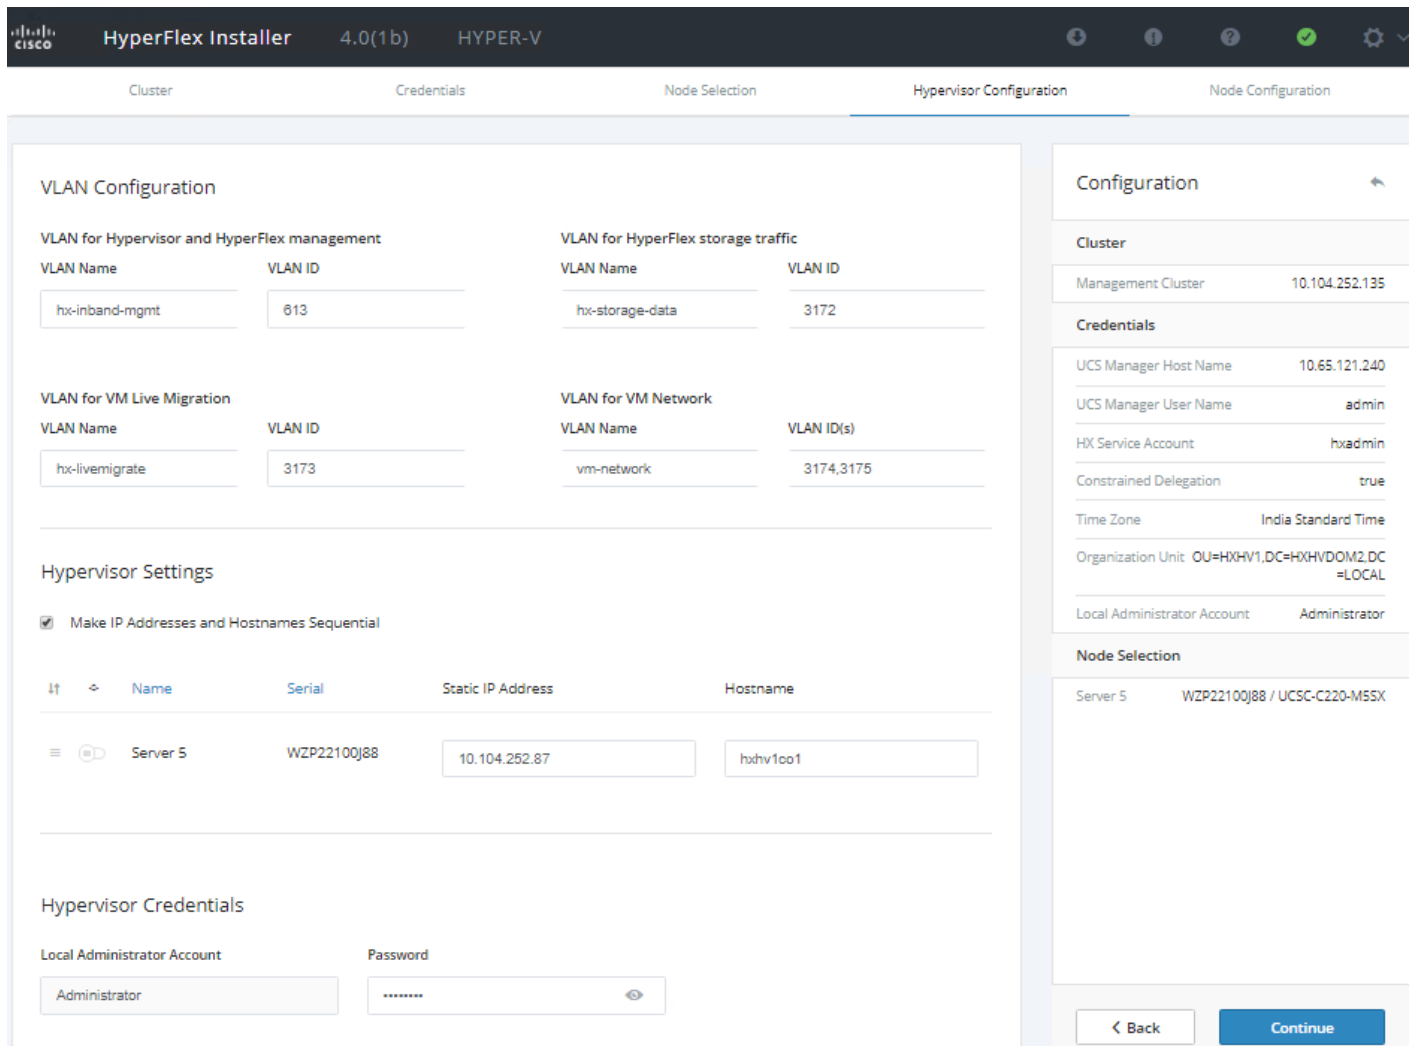
- Select the unassociated HX servers you want to add to the existing HX cluster.



20. Click Continue.

21. On the Hypervisor Configuration page:

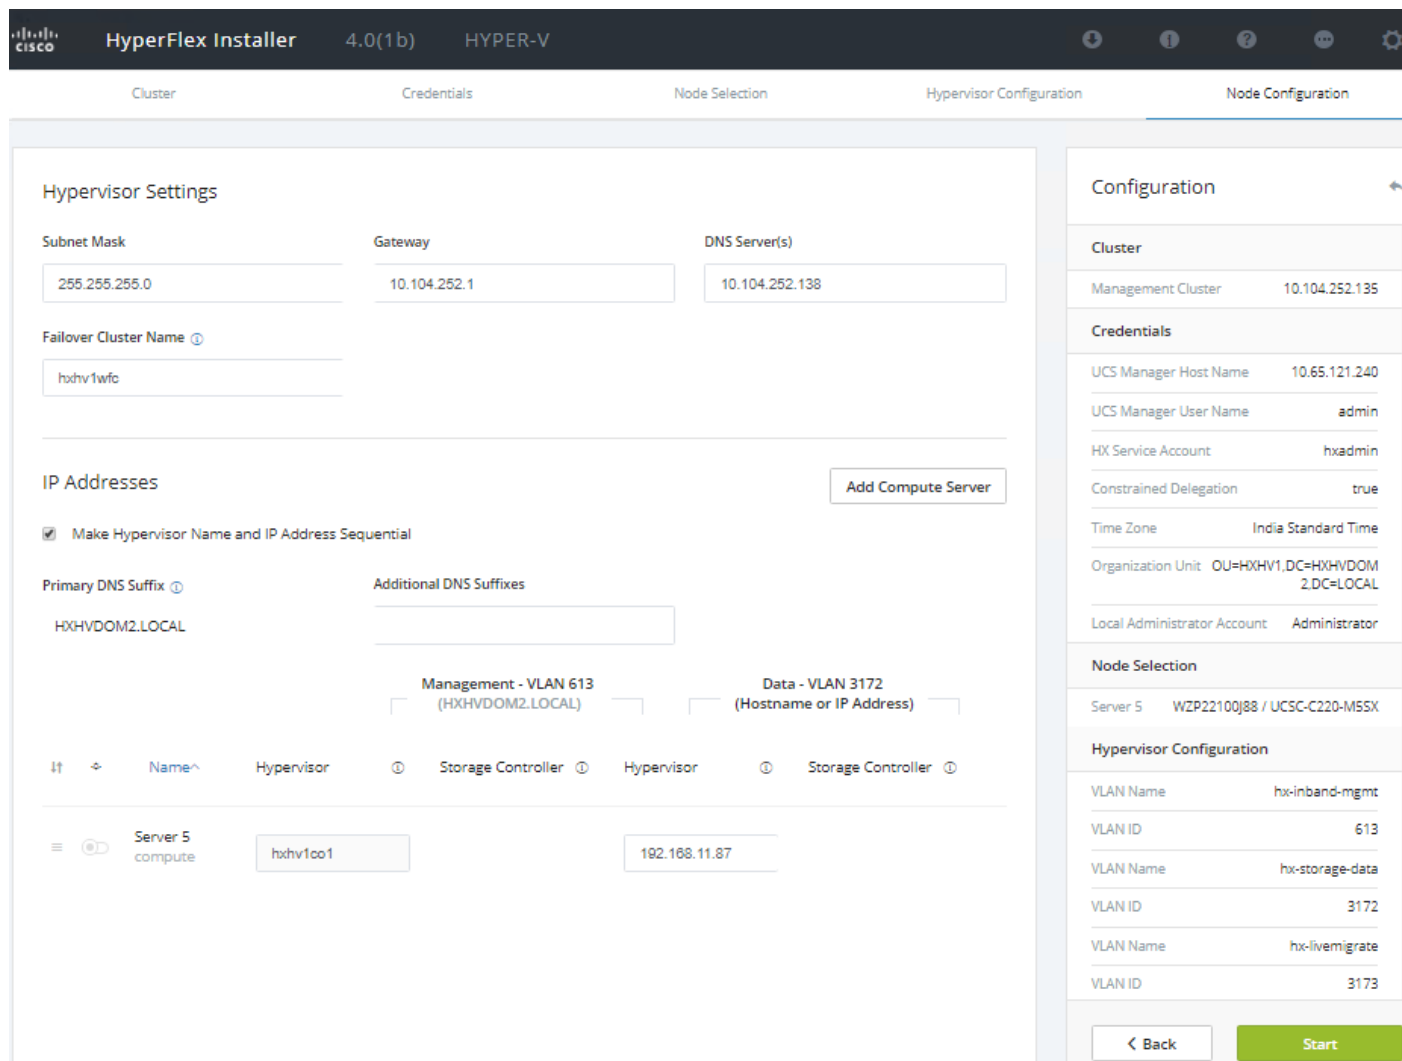
- a. VLAN Configuration - Enter the VLAN names and VLAN IDs that are to be created in Cisco UCS, multiple comma-separated VLAN IDs for different guest virtual machine networks are allowed here.
- b. Hypervisor Settings - Enter IP addresses and hostnames for the Hypervisors that were created in the pre-installation section phase. The IP addresses will be assigned via Serial over Lan (SoL) through Cisco UCS Manager to the Hyper-V host systems as their management IP addresses.
- c. Hypervisor Credentials - this field is pre-populated and can't be edited.



22. Click Continue.

23. On the Node Configuration Page:

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



24. Click Start.

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

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

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

Progress

Start Config Installer Hypervisor Configuration Deploy Validation Deploy Expansion Validation Cluster Expansion

Cluster Expansion in Progress

Deploy

Deploy - Overall
Succeeded

hxxhvc02.HXHVDM1.LOCAL
Succeeded

- ✓ HyperV Host Add
- ✓ Getting Ready to Copy Required Powershell, Task and XML files to Windows Host
- ✓ Validate HyperV Hosts
- ✓ Main Playbook [HyperV Host Configuration]
- ✓ Main Playbook [Storage Controller VMs Configuration]
- ✓ Generate Host Data Facts Related JSON Files
- ✓ Generate Controller/VM Data Facts Related JSON Files

Configuration

Cluster	
Management Cluster	10.104.252.135
Credentials	
UCS Manager Host Name	10.65.121.240
UCS Manager User Name	admin
HX Service Account	hxadmin
Constrained Delegation	true
Time Zone	Pacific Standard Time
Organization Unit	OU=HyperFlex,DC=hxxhvc02,DC=local
Local Administrator Account	Administrator
Node Selection	
Server 1/1	FCH2141JBKY / UCSB-B200-M5
Server 5	WZP2208115W / UCSC-C220-M55X
Hypervisor Configuration	
VLAN Name	hx-inband-mgmt
VLAN ID	613
VLAN Name	hx-storage-data

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

The screenshot shows the HyperFlex Installer Summary page. At the top, it displays 'HyperFlex Installer 4.0(1b) HYPER-V'. The page is divided into 'Progress' and 'Summary' tabs. The cluster name is 'hxhv1smb', with 'ONLINE' and 'HEALTHY' status indicators. Below this, a table lists various configuration parameters. At the bottom, there is a table of servers with columns for Model, Serial Number, Management Hypervisor, Management Storage Controller, Data Network Hypervisor, and Data Network Storage Controller. Two buttons are visible at the bottom right: 'Back to Workflow Selection' and 'Launch HyperFlex Connect'.

Parameter	Value	Parameter	Value
Version	4.0.1b-33133	Domain Name	HXHVDOM2.LOCAL
Cluster Management IP Address	hxhv1cip.HXHVDOM2.LOCAL	Failover cluster Name	hxhv1wfc
Cluster Data IP Address	192.168.11.135	DNS Server(s)	10.104.252.138
Replication Factor	Three copies	NTP Server(s)	10.104.252.138
Available Capacity	10.7 TB		

Model	Serial Number	Management Hypervisor	Management Storage Controller	Data Network Hypervisor	Data Network Storage Controller
HXAF240C-M55X	WZP22020L96	10.104.252.129	10.104.252.133	192.168.11.129	192.168.11.133
HXAF240C-M55X	WZP22020L9E	10.104.252.127	10.104.252.131	192.168.11.127	192.168.11.131
HXAF240C-M55X	WZP22020L9B	10.104.252.130	10.104.252.134	192.168.11.130	192.168.11.134
UC5C-C220-M55X	WZP22100J88	10.104.252.87		192.168.11.87	
HXAF240C-M55X	WZP220216WY	10.104.252.128	10.104.252.132	192.168.11.128	192.168.11.132

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

HyperFlex Connect hxhv1smb

OPERATIONAL STATUS
Online

RESILIENCY HEALTH
Healthy 1 Node failure can be tolerated

CAPACITY
10.7TB
156.7 GB Used | 10.6 TB Free

STORAGE OPTIMIZATION
88.6%

Compression: 29%
Deduplication: 84%

NODES
5
4 HXAF240C-M5SX Converged

1 UCSC-C220-M5SX Compute

IOPS Last 1 hour
Read Max: 56.6 Min: 37.2 Avg: [unclear]

hxbv1co1.HXHVDOM2.LOCAL:10.104.252.87

Connect hxhv1smb

Activity Monitor progress of recent tasks on the HX storage cluster. Last refreshed at: 09/29/2019 5:53:34 PM

Filter: Filter listed tasks

^ Collapse All

- Expand HyperV Cluster - COMPUTE
Status: Success
09/29/2019 4:48:46 PM
 - Initialize: hxhv1co1.HXHVDOM2.LOCAL compute initialization succeeded.
 - Whitelist: hxhv1co1.HXHVDOM2.LOCAL white listing for compute 192.168.11.87 succeeded.
 - Placement: hxhv1co1.HXHVDOM2.LOCAL placement succeeded, 192.168.11.87 compute is assigned to 192.168.11.131 controller vm.
 - Redirect IOvisor: hxhv1co1.HXHVDOM2.LOCAL compute IOvisor redirection succeeded.
 - ZK Update: hxhv1co1.HXHVDOM2.LOCAL compute ZK update succeeded.
- Check Expand HyperV Cluster
Status: Success
09/29/2019 4:48:21 PM
 - Duplicate IP Check: Duplicate IP check is successful
 - Resolve Cluster Data IP: Cluster data IP is successfully resolved



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

Expand/Collapse left menu Connect
hxhv1smb ⚙ ?

System Overview
Nodes
Disks

Last refreshed at: 09/29/2019 5:54:30 PM ↻

✔

hxhv1smb

ONLINE

Actions ▼

Uptime	31 days, 14 hours, 33 minutes, 41 sec onds	Hypervisor	Microsoft Windows Server 2019 Datacenter	Total Capacity	10.71 TB	DNS Server(s)	10.104.252.138
		HXDP Version	4.0.1b-33133	Available Capacity	10.56 TB	NTP Server(s)	10.104.252.138
				Data Replication Factor	3	Controller Access over SSH	Enabled

hxhv11.HXHVDOM2.LOCAL
HXAF240C-M55X
11 Disks (1 Caching, 10 Persistent)

Online ✔

HXDP Version 4.0(1b)

Type Hyper Converged

Hypervisor Status Online
Hypervisor Address 10.104.252.127

hxhv12.HXHVDOM2.LOCAL
HXAF240C-M55X
11 Disks (1 Caching, 10 Persistent)

Online ✔

HXDP Version 4.0(1b)

Type Hyper Converged

Hypervisor Status Online
Hypervisor Address 10.104.252.128

hxhv13.HXHVDOM2.LOCAL
HXAF240C-M55X
11 Disks (1 Caching, 10 Persistent)

Online ✔

HXDP Version 4.0(1b)

Type Hyper Converged

Hypervisor Status Online
Hypervisor Address 10.104.252.129

hxhv14.HXHVDOM2.LOCAL
HXAF240C-M55X
11 Disks (1 Caching, 10 Persistent)

Online ✔

HXDP Version 4.0(1b)

Type Hyper Converged

Hypervisor Status Online
Hypervisor Address 10.104.252.130

hxhv1co1.HXHVDOM2.LOCAL
UCSC-C220-M55X

Online ✔

Version Microsoft Windows Server 2019 Datacenter

Type Compute

Dashboard

MONITOR

Alarms

Events

Activity

ANALYZE

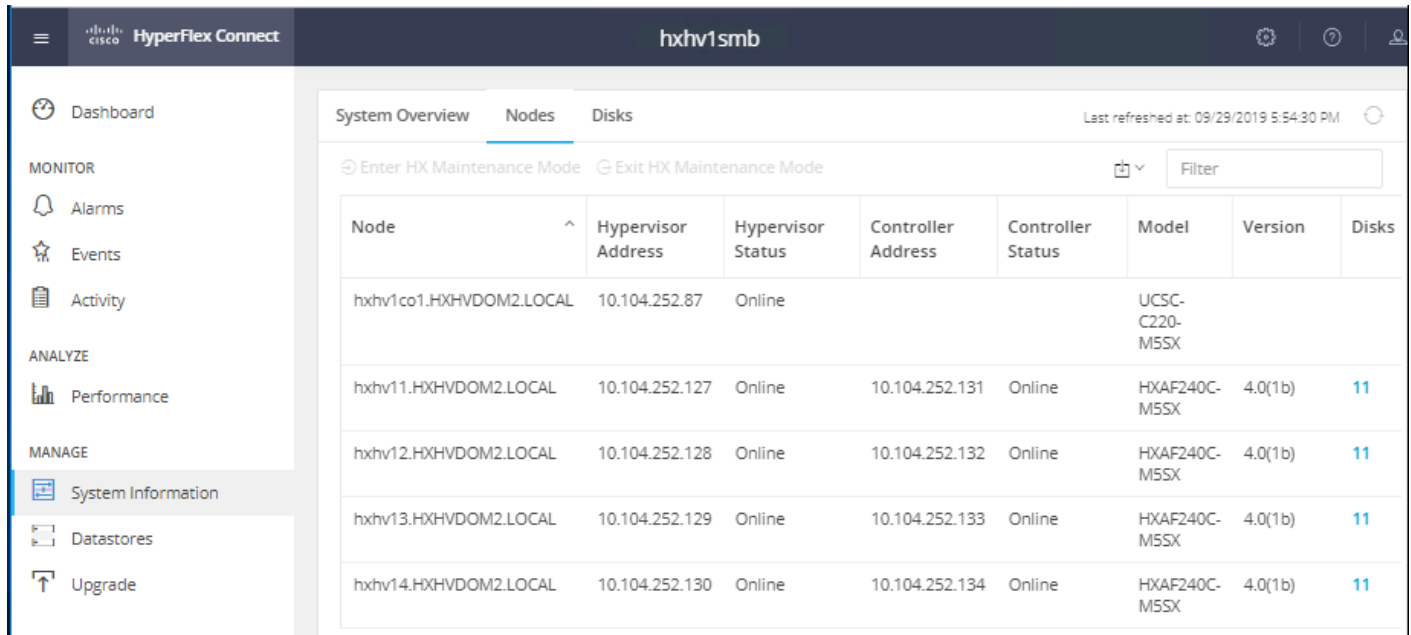
Performance

MANAGE

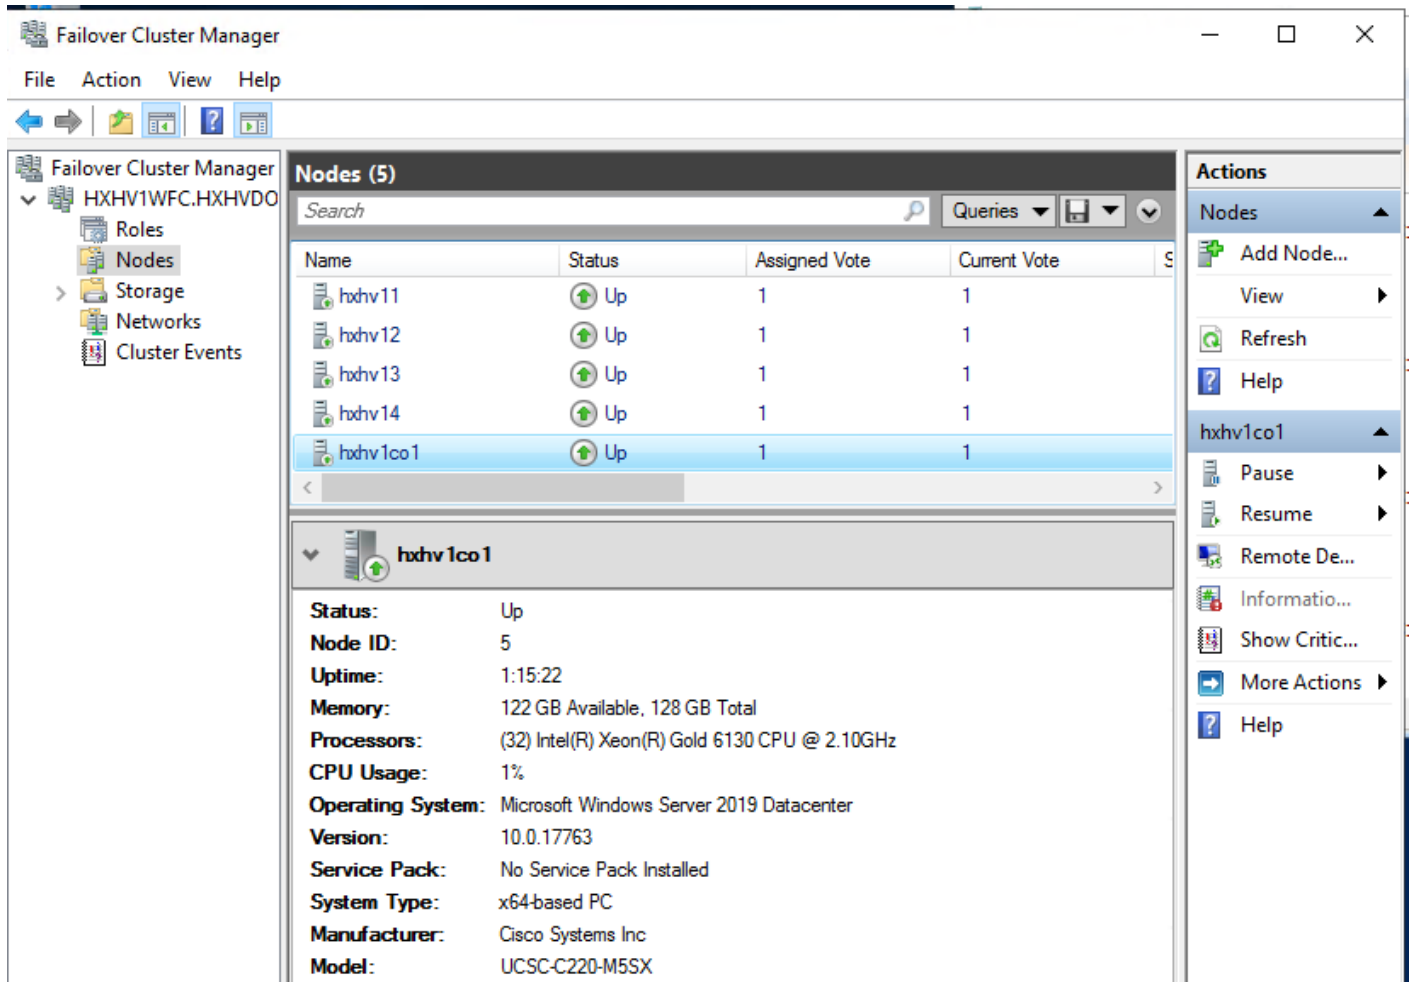
System Information

Datastores

Upgrade



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



31. For more information, refer to section [Assign IP Addresses to Live Migration and Virtual Machine Network Interfaces](#).
32. After assigning IP addresses, verify the jumbo frame settings on the live migration network of newly expanded compute node using the below PowerShell cmdlet:

a. From local host, use the following command:

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

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

```
Invoke-Command -ComputerName hxhvlco1 -ScriptBlock {Get-NetAdapterAdvancedProperty -name hv-livemigrate* | Where-Object {$_.DisplayName -Match "Jumbo*"}}
```

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhvlco1 -ScriptBlock {Get-NetAdapterAdvancedProperty -name hv-livemigrate* | Where-Object {$_.DisplayName -Match "Jumbo*"}}
```

Name	DisplayName	DisplayValue	RegistryKeyword	RegistryValue	PSComputerName
hv-livemigrate-a	Jumbo Packet	Bytes 1514	*JumboPacket	{1514}	hxhvlco1
hv-livemigrate-b	Jumbo Packet	Bytes 1514	*JumboPacket	{1514}	hxhvlco1

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

a. From local host, use the following command:

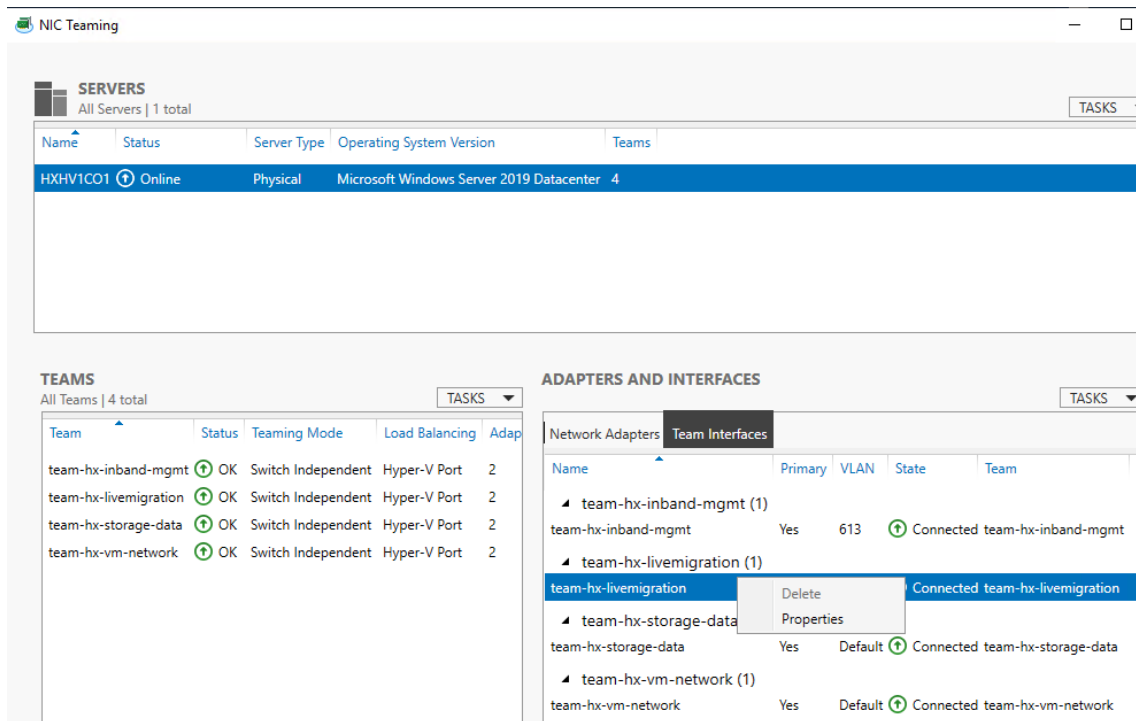
```
Set-NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword "*JumboPacket" -RegistryValue 9014
```

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

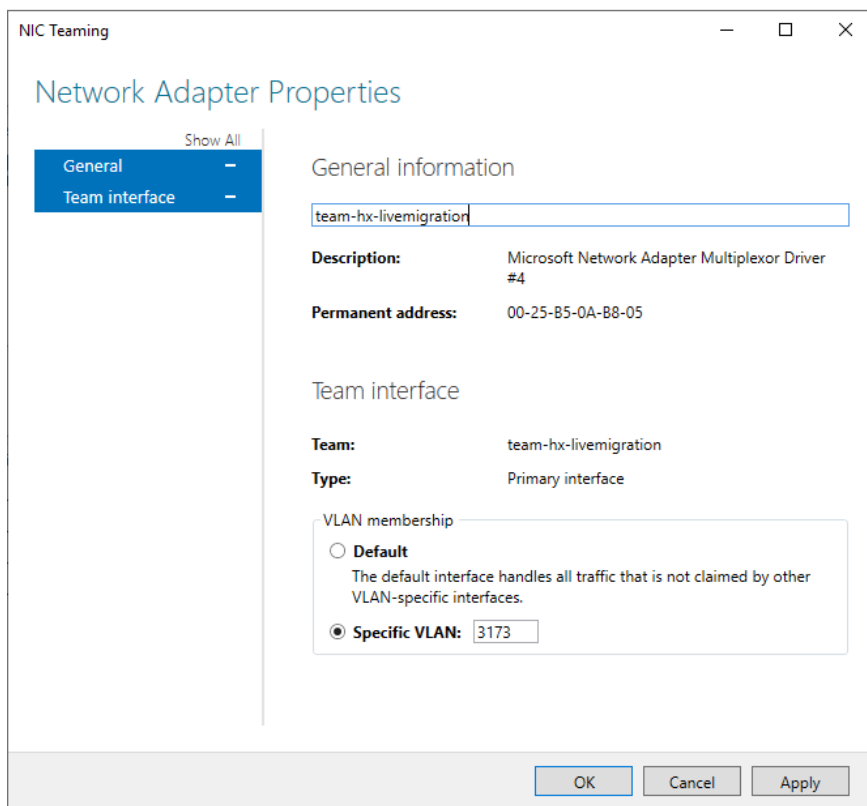
```
Invoke-Command -ComputerName hxhvlco1 -ScriptBlock {Set-NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword "*JumboPacket" -RegistryValue 9014}
```

```
PS C:\Users\Administrator> Invoke-Command -ComputerName hxhvlco1 -ScriptBlock {Set-NetAdapterAdvancedProperty -Name "hv-livemigrate*" -RegistryKeyword "*JumboPacket" -RegistryValue 9014}
```

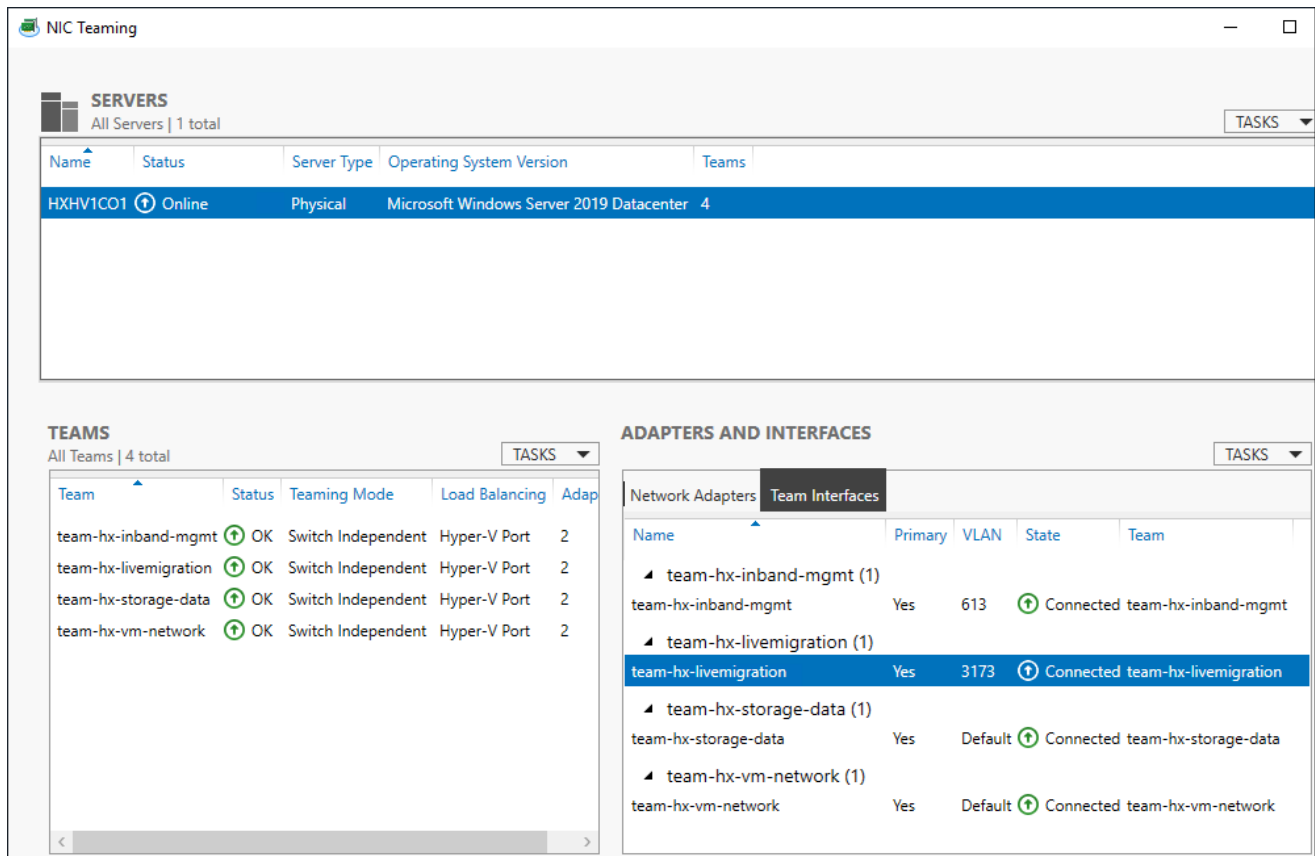
34. On the newly added Compute node, open Server Manager > Local Server and click Enabled next to the NIC Teaming.
35. On the NIC Teaming window, right-click "team-hx-livemigration" and click properties under the "Adapter and Interfaces" section as shown below:



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



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



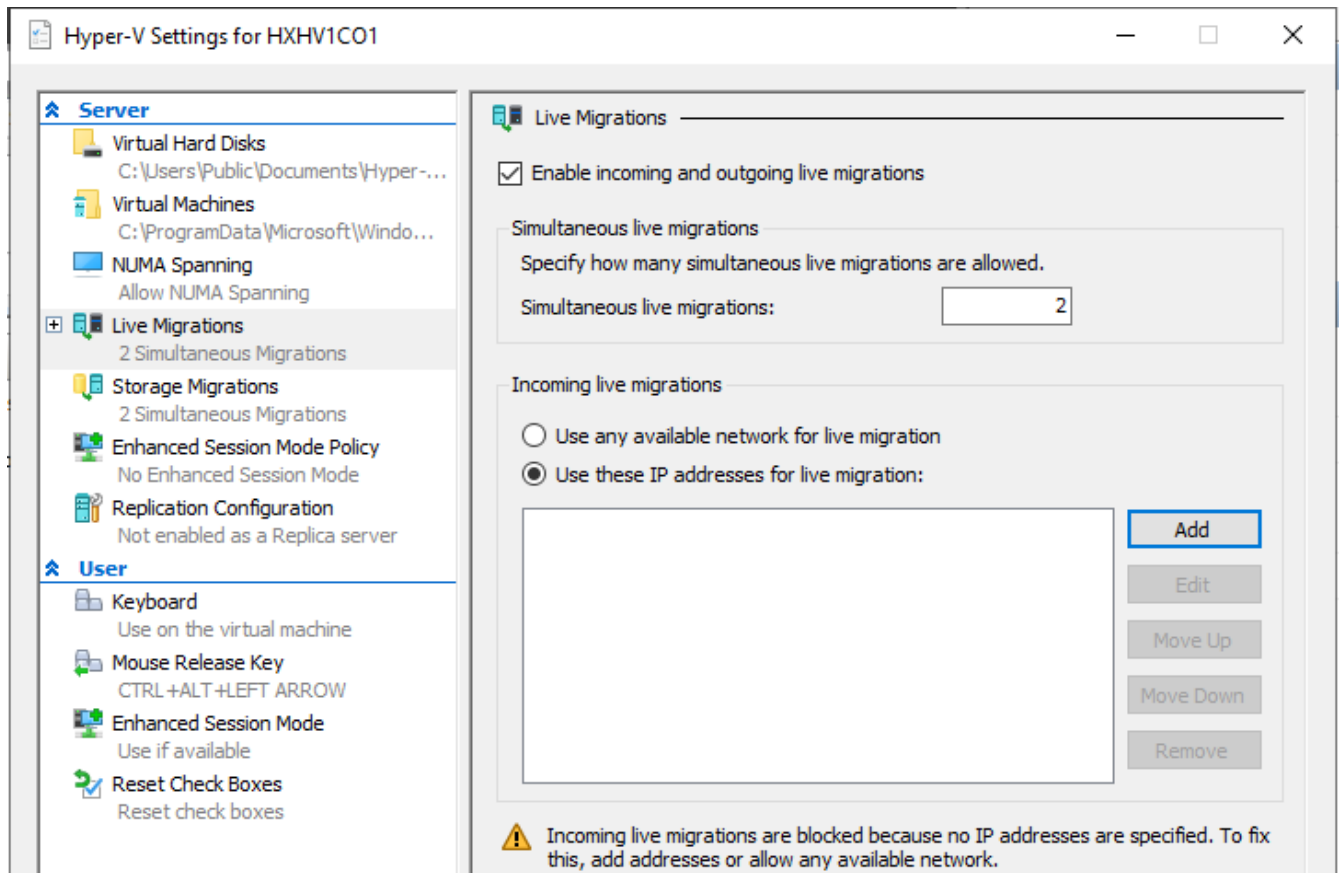
38. Test/validate by pinging the LM network interface of another hx node from the newly added compute node if jumbo frame is working as shown below:

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


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

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

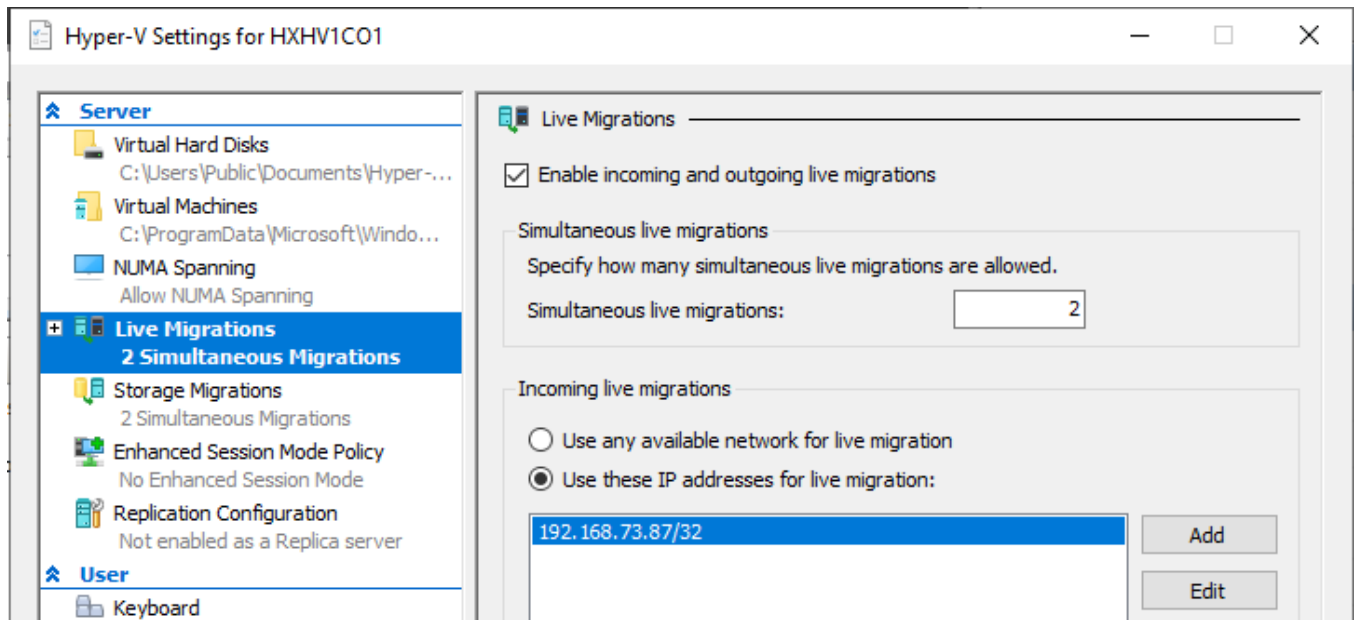
39. Open Hyper-V Manager on the newly added compute node, select the server and in the Actions pane, click Hyper-V Settings > Live Migrations.



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



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



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

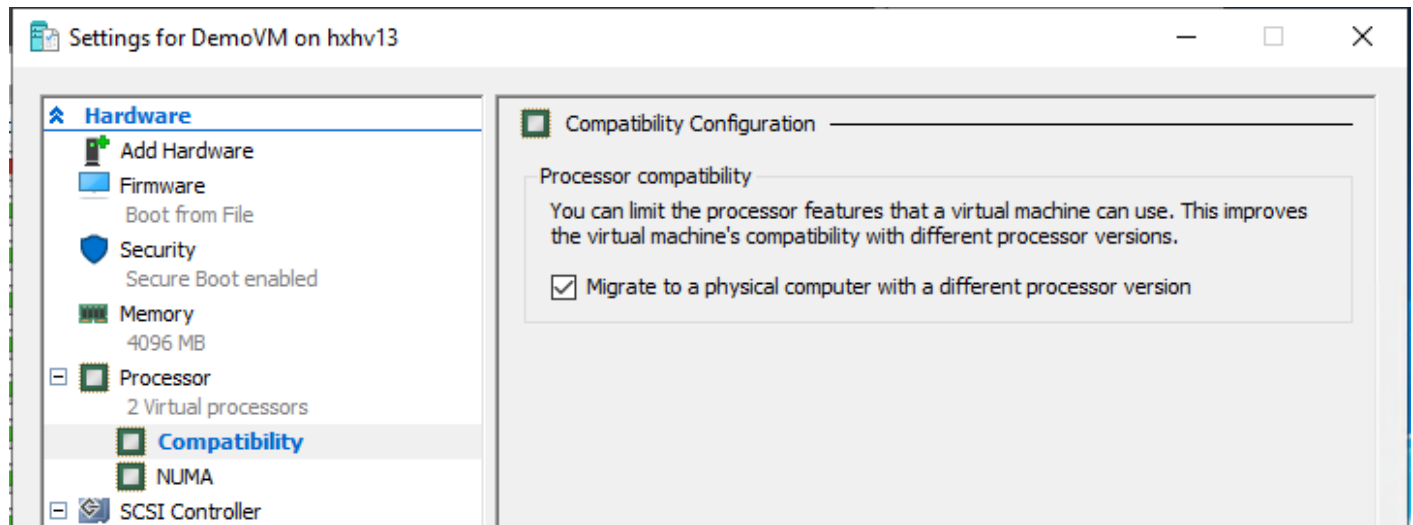
```
PS C:\Users\administrator.HXHVDOM2> Get-VM
```

Name	State	CPUUsage(%)	MemoryAssigned(M)	Uptime	Status	Version
DemoVM	Running	0	4096	00:00:35.9470000	Operating normally	9.0
StCt1VM	Running	0	73728	6.09:23:45.3590000	Operating normally	9.0
VM01	Running	0	2048	00:33:16.8980000	Operating normally	9.0
VM05	Running	0	8192	6.09:17:27.6200000	Operating normally	9.0
VM06	Running	0	8192	6.09:17:14.1760000	Operating normally	9.0

```
PS C:\Users\administrator.HXHVDOM2> Measure-Command { Move-VM -Name DemoVM -DestinationHost HXHV1C01 }
```

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

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



Management

HyperFlex Connect

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

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

1. Using a web browser, open the HyperFlex cluster's management IP address via HTTPS, for example, <https://10.29.149.230>
2. Enter a local credential, such as admin, and the password.
3. Click Login.

The Dashboard view will display after a successful login.

Figure 103 Cisco HyperFlex Connect – Login Page


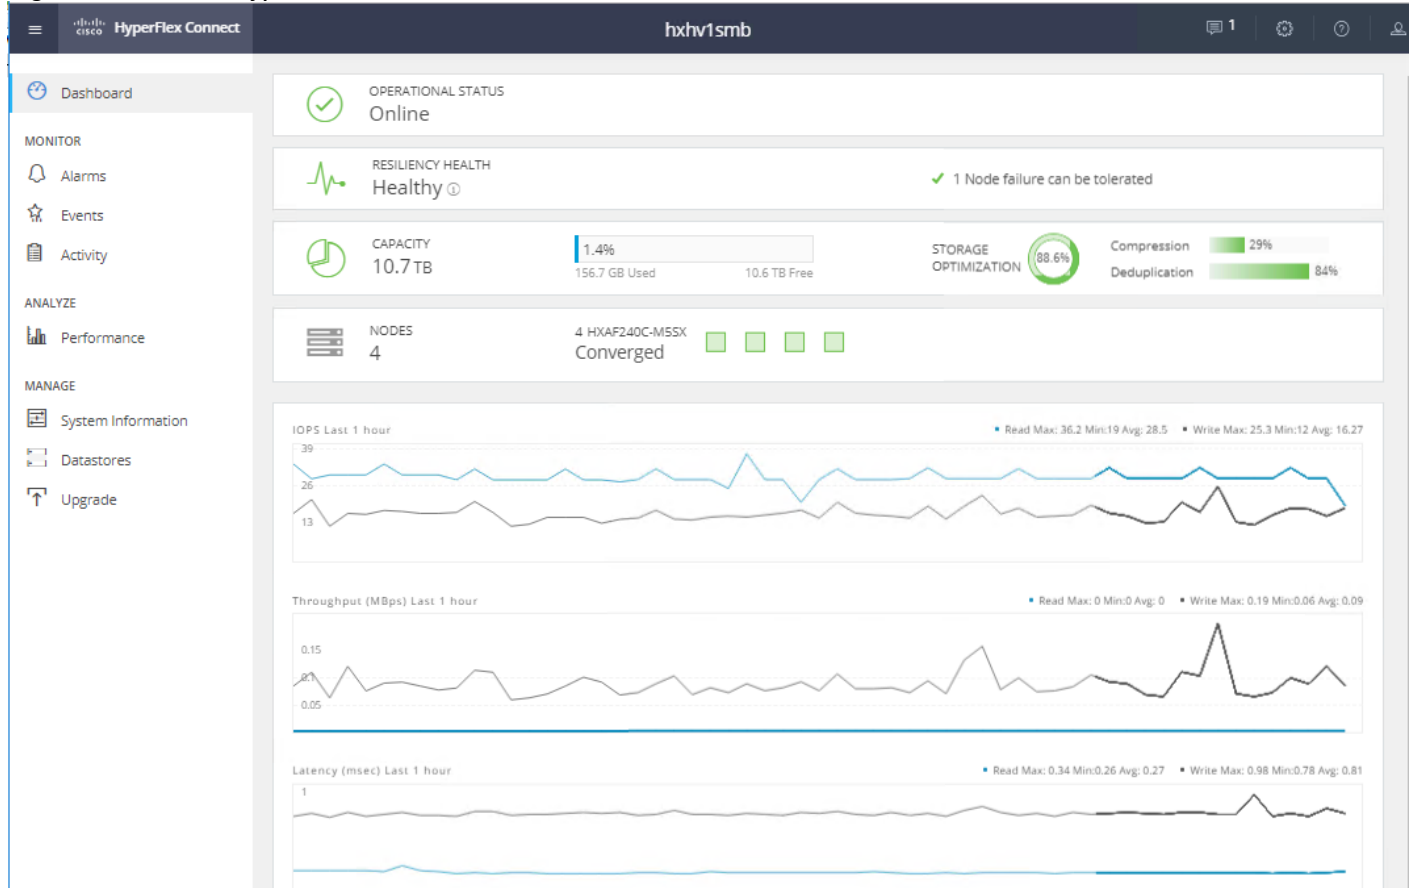


Figure 104 Cisco HyperFlex Connect – Dashboard



Dashboard

From the Dashboard view, the following elements are presented:

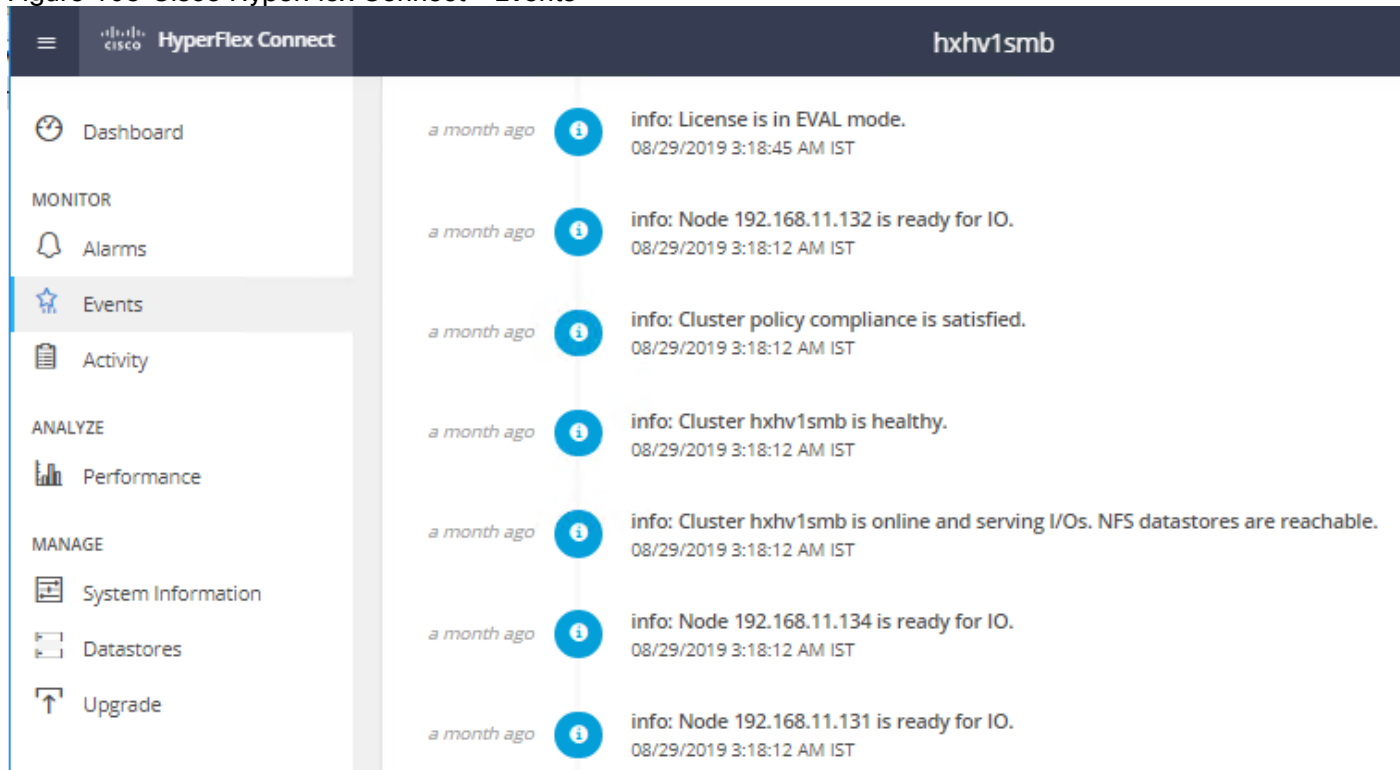
- Cluster operational status, overall cluster health, and the cluster's current node failure tolerance.
- Cluster storage capacity used and free space, compression and deduplication savings, and overall cluster storage optimization statistics.
- Cluster size and individual node health.
- Cluster IOPs,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

Monitor

HyperFlex Connect provides for additional monitoring capabilities, including:

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

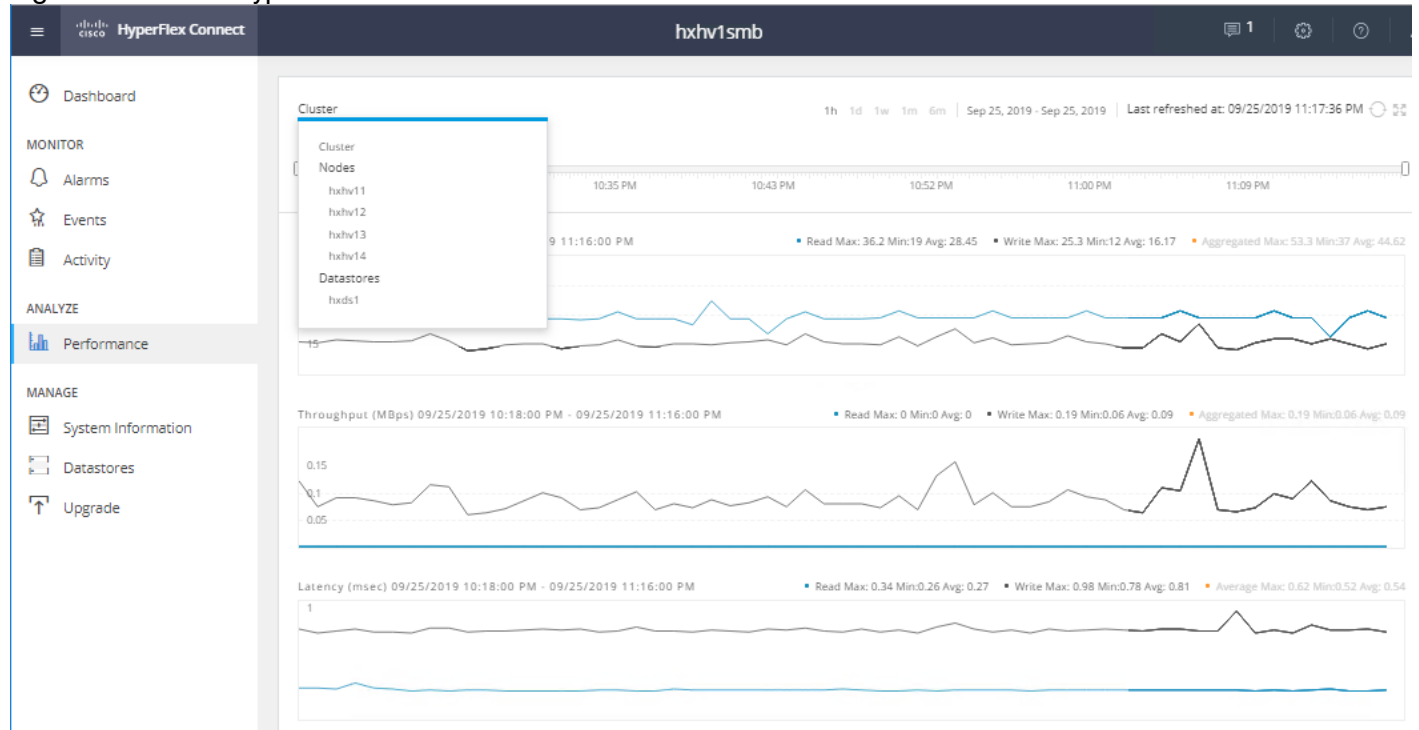
Figure 105 Cisco HyperFlex Connect - Events



Analyze

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

Figure 106 Cisco HyperFlex Connect – Performance

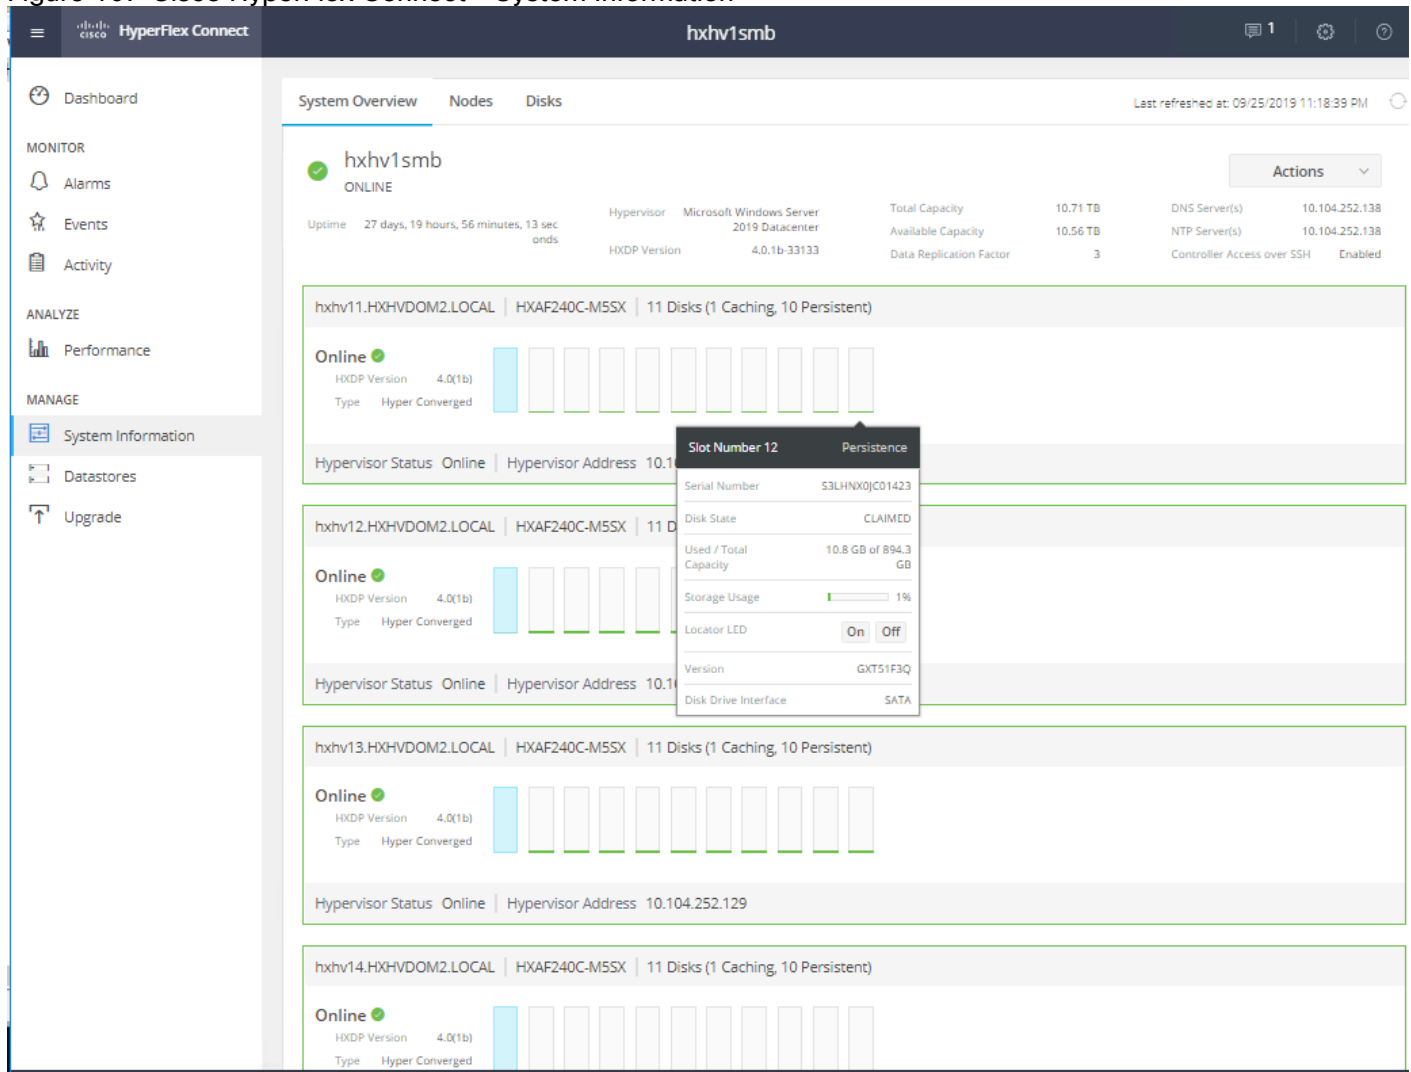


Manage

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

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

Figure 107 Cisco HyperFlex Connect – System Information



Configure Remote Management Stations for Managing VMs in HX Cluster

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

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

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

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

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

```
HyperFlex StorageController 4.0(1b)
Last login: Wed Sep 25 18:34:47 2019 from 10.104.252.51
root@hxhv11scvm:~# cat /etc/springpath/storfs.cfg | grep crmZKEnsemble
crmZKEnsemble=192.168.11.133:2181,192.168.11.134:2181,192.168.11.132:2181,192.168.11.131:2181
root@hxhv11scvm:~# █
```

- To allow access to the SMB share on remote management station, execute the following command without any parameters as shown below:

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

```
root@hxhv11scvm:~# python /opt/springpath/storfs-hyperv/FixScvmmAccess.py
Enter Ip address of SCVMM: 10.104.252.77
PING 10.104.252.77 (10.104.252.77) 56(84) bytes of data.
64 bytes from 10.104.252.77: icmp_seq=1 ttl=128 time=0.230 ms

--- 10.104.252.77 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.230/0.230/0.230/0.000 ms
root@hxhv11scvm:~# █
```

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

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

```
PS C:\Users\administrator.HXHVDOM2> test-path \\hxhv1smb.hxhvdome2.local\hxds1
False
PS C:\Users\administrator.HXHVDOM2> █
```

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

```

PS C:\Users\administrator.HXHVDOM2> Get-Content C:\Windows\System32\drivers\etc\hosts
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com           # source server
#       38.25.63.10      x.acme.com              # x client host
#
# localhost name resolution is handled within DNS itself.
#       127.0.0.1        localhost
#       ::1             localhost
10.104.252.135          hxhv1smb.hxhvdome2.local
PS C:\Users\administrator.HXHVDOM2>


```

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

```

PS C:\Users\administrator.HXHVDOM2> test-path \\hxhv1smb.hxhvdome2.local\hxds1
True
PS C:\Users\administrator.HXHVDOM2>

```



Microsoft Hyper-V Manager

Hyper-V Manager is a free GUI-based administration tool for managing Hyper-V hosts and virtual machines, both locally and remotely. By default, the management tool is installed locally on Windows Servers where the Hyper-V role is enabled. To manage Hyper-V hosts from remote Windows client/server, Remote Server Administration Tools (RSAT) feature needs to be installed.

To manage Hyper-V hosts from a remote management station, follow these steps:

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

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

2. Complete the steps described in the above section "Configuring Remote Management Stations for Managing VMs in HX Cluster."

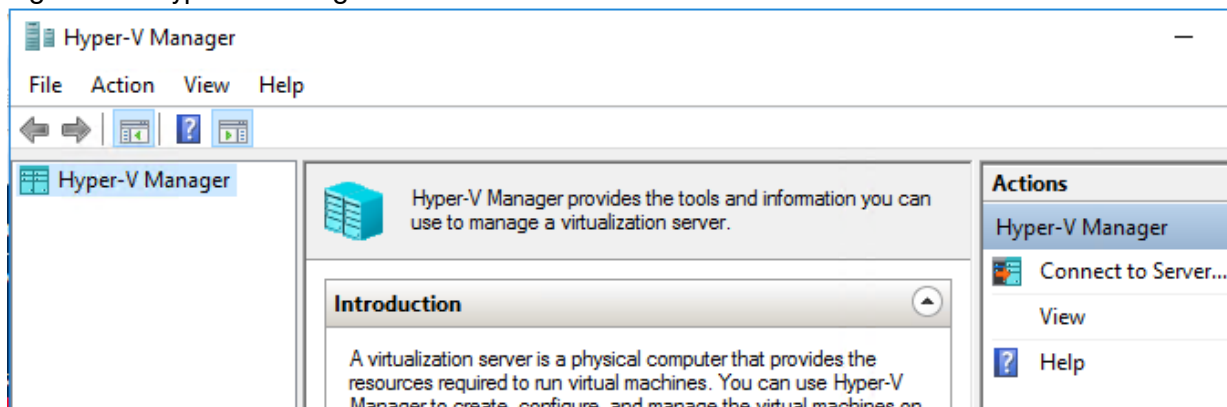
Basic management functions like changing the default store location for virtual machine files and creating virtual machines are described in the following sections. For more information about managing Hyper-V using Hyper-V manager refer to the Microsoft website.

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

To change the default location, follow these steps:

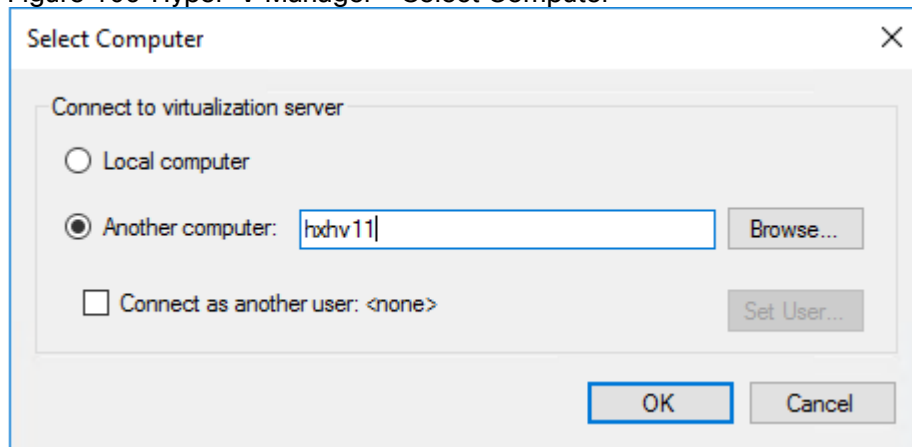
1. Open the Server Manager dashboard and click Tools. Click Hyper-V Manager. The Hyper-V Manager console appears.

Figure 108 Hyper-V Manager - Connect to Server



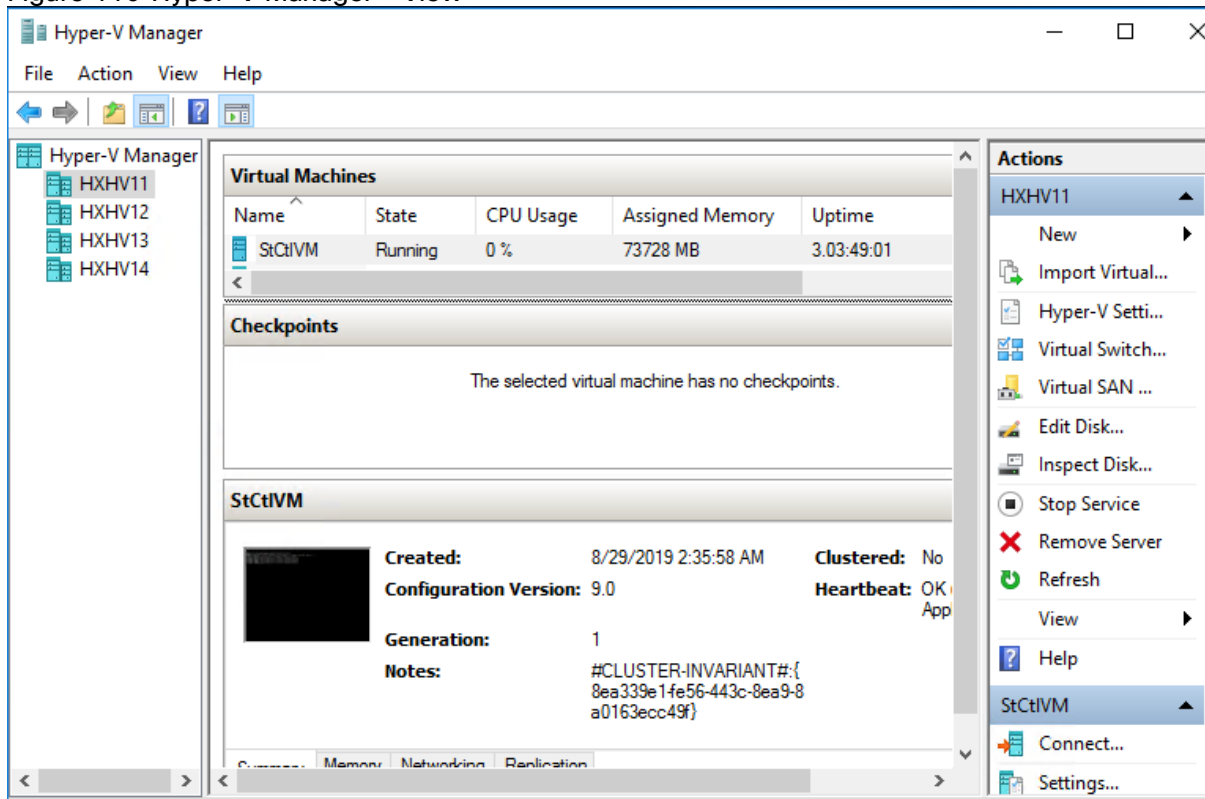
2. In the left pane, select Hyper-V Manager and click Connect to Server....

Figure 109 Hyper-V Manager - Select Computer



3. In the Select Computer dialog box, select Another computer and type in the name of the Hyper-V node (for example, HXHV1) that belongs to the Hyper-V cluster. Click OK.
4. Repeat steps 1-3 for each Hyper-V node in the HyperFlex cluster.

Figure 110 Hyper-V Manager – View



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

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

Figure 111 Hyper-V Manager – Virtual Machine VHD Location

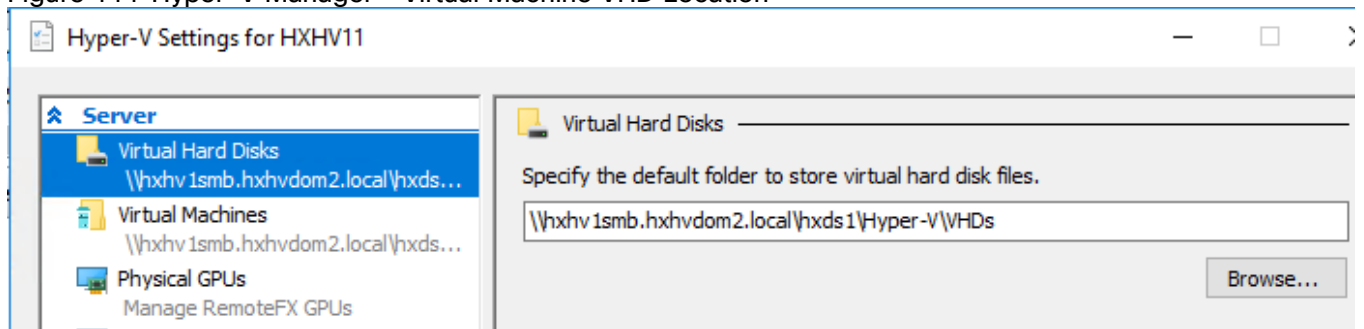
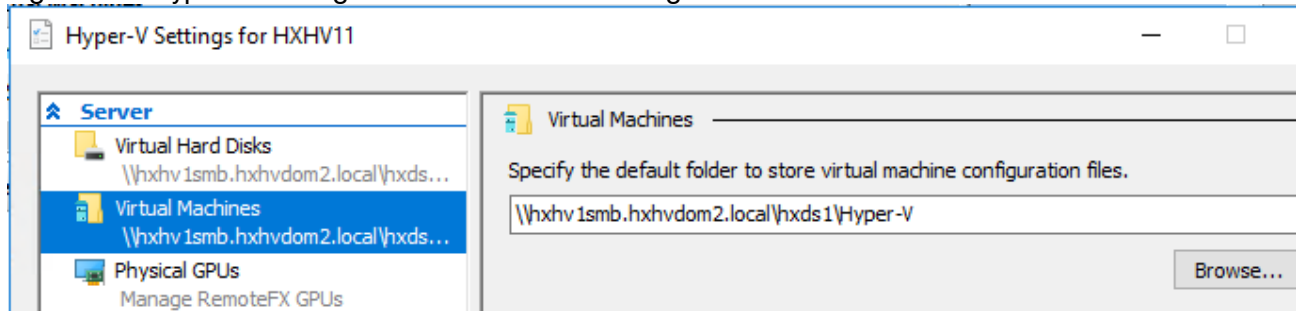


Figure 112 Hyper-V Manager – Virtual Machine Configuration Files Location

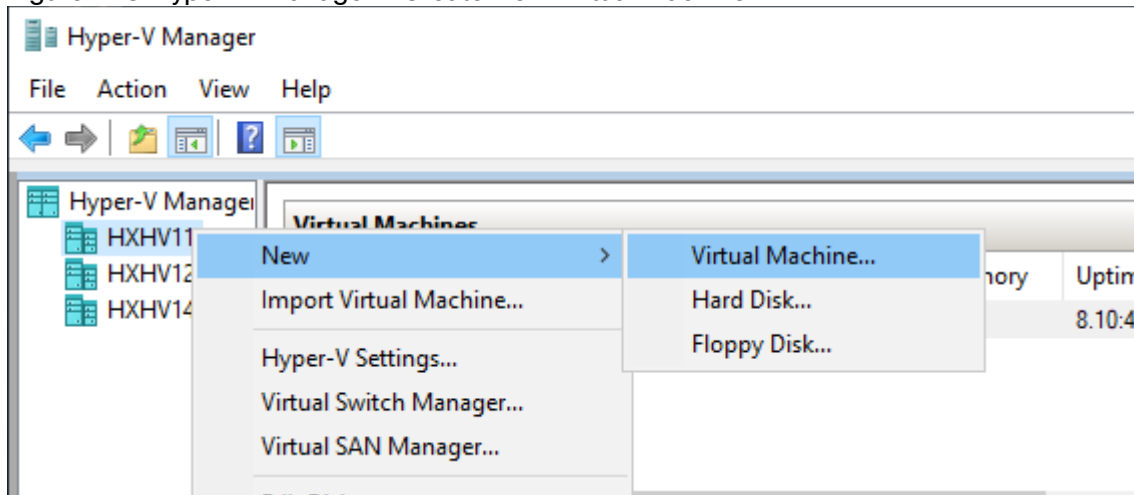


Create Virtual Machines using Hyper-V Manager

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

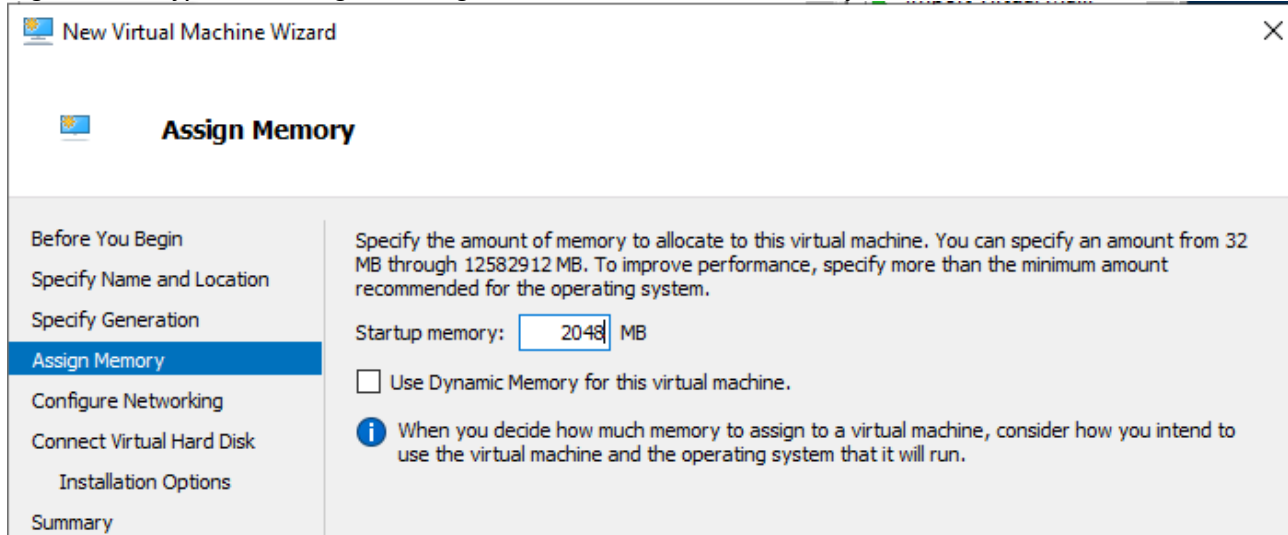
1. Open Hyper-V Manager from the Server Manager > Tools.
2. Select a Hyper-V server, and right-click and select New > Create a virtual machine. The Hyper-V Manager New Virtual Machine wizard displays.

Figure 113 Hyper-V Manager – Create New Virtual Machine



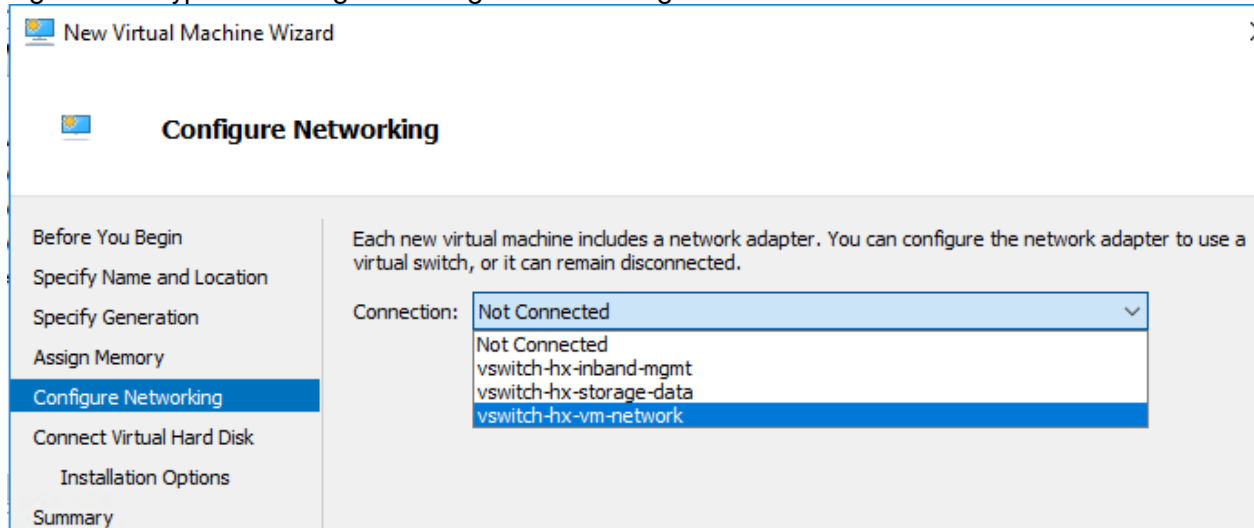
3. In the Before you Begin page, click Next.
4. In the Specify Name and Location page, enter a name for the virtual machine configuration file. The location for the virtual machine click Next.
5. In the Specify Generation page, choose either Generation 1 or Generation 2.
6. In the Assign Memory page, set the start memory value (For example - 2048 MB). Click Next.

Figure 114 Hyper-V Manager – Assign New Virtual Machine Memory



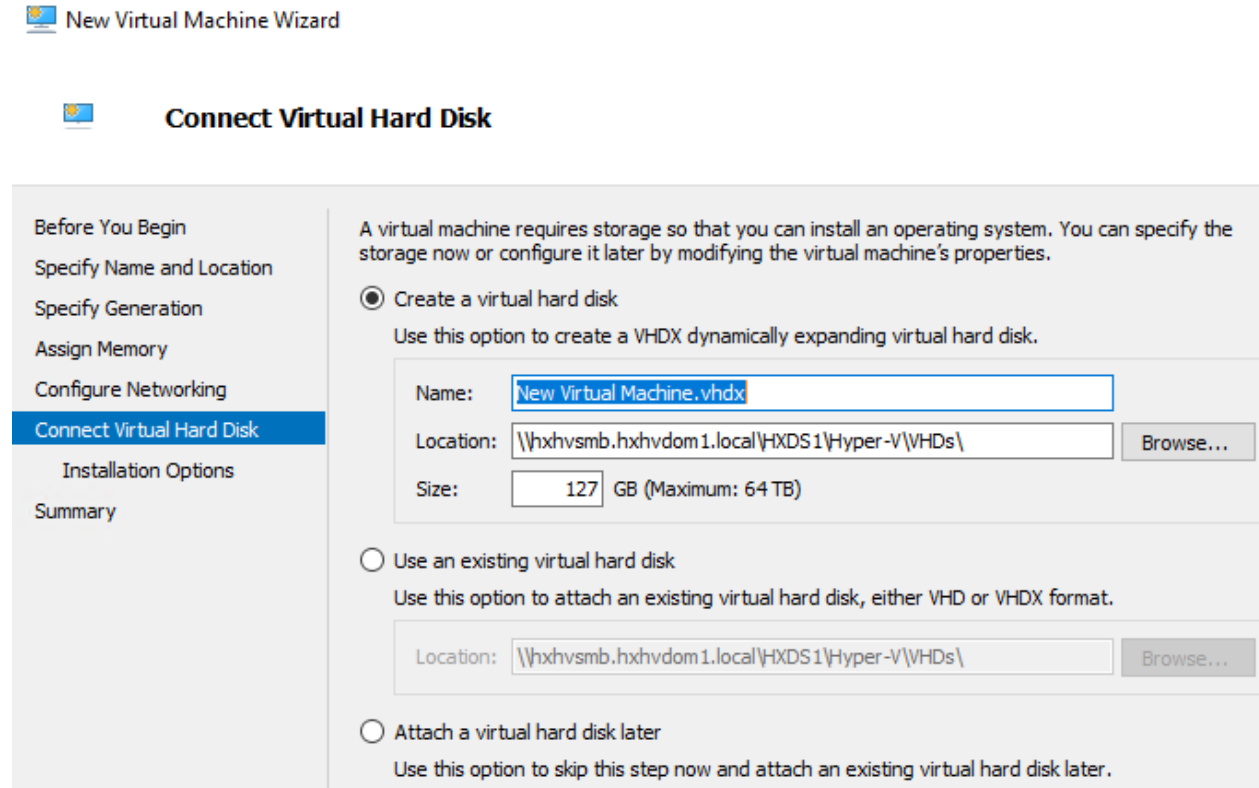
7. In the Configure Networking page, select a network connection for the virtual machine to use from a list of existing virtual switches.

Figure 115 Hyper-V Manager – Configure Networking



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

Figure 116 Hyper-V Manager – Connect VHD to New Virtual Machine



9. In the Installation Options, you can leave the default option Install an operating system later selected. Click Next.
10. In the Summary page, verify that the list of options displayed are correct. Click Finish.
11. In Hyper-V Manager, right-click the virtual machine to perform various operations like Connect, Edit Settings, Start/Stop, and so on.



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

Windows Failover Cluster Manager

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



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

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

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

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

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

2. Complete the steps described in the above section "Configuring Remote Management Stations for Managing VMs in HX Cluster"

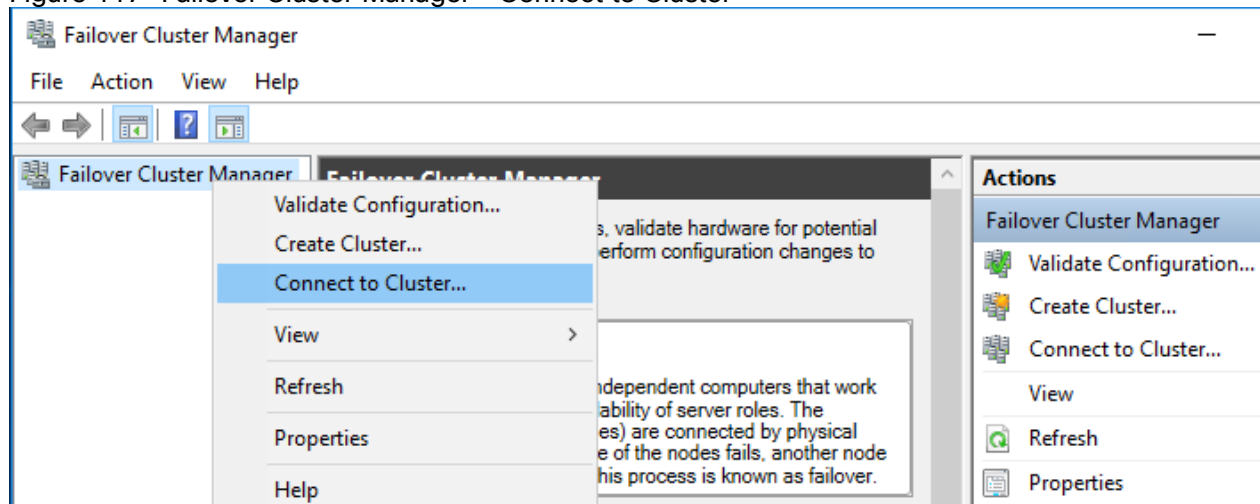
To create the clustered virtual machine role using the Failover Cluster Manager, follow these steps:



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

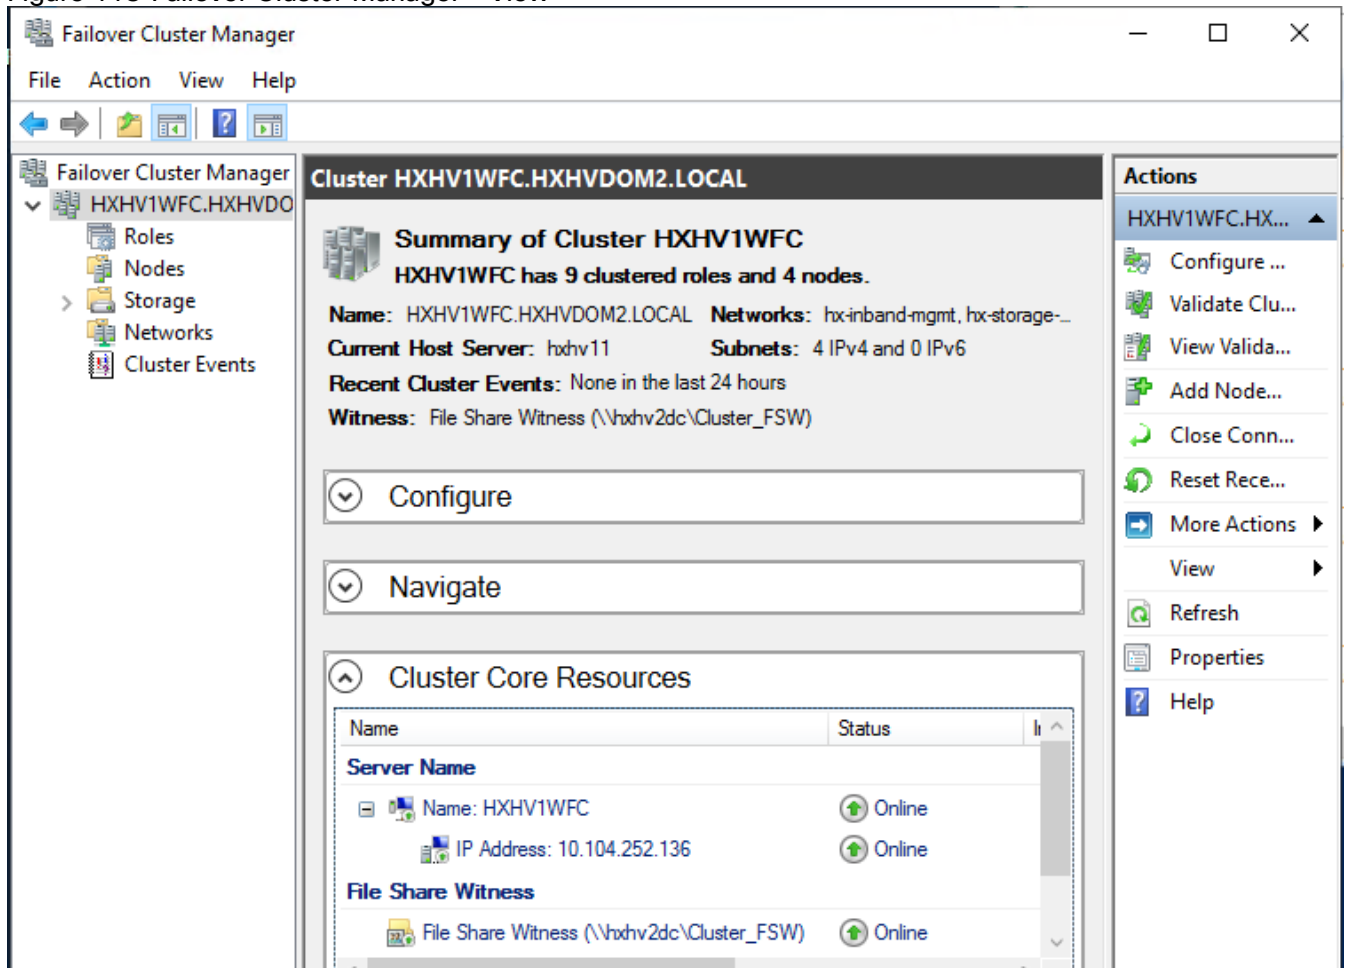
1. In the Failover Cluster Manager console, under the Actions pane, click Connect to Cluster...

Figure 117 Failover Cluster Manager – Connect to Cluster



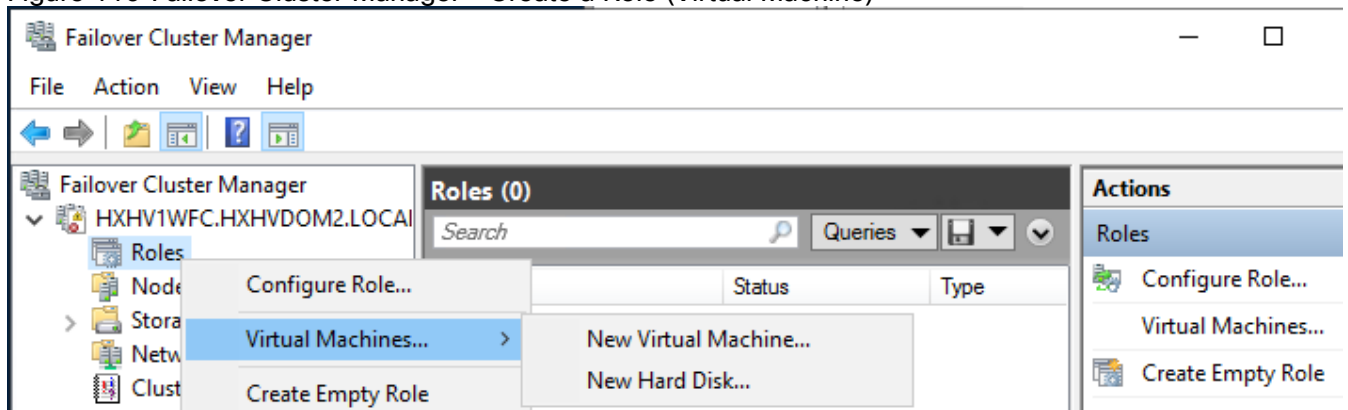
2. In the Select Cluster dialog box, click Browse to navigate to the Hyper-V cluster name. Click OK.

Figure 118 Failover Cluster Manager - View




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

Figure 119 Failover Cluster Manager - Create a Role (Virtual Machine)



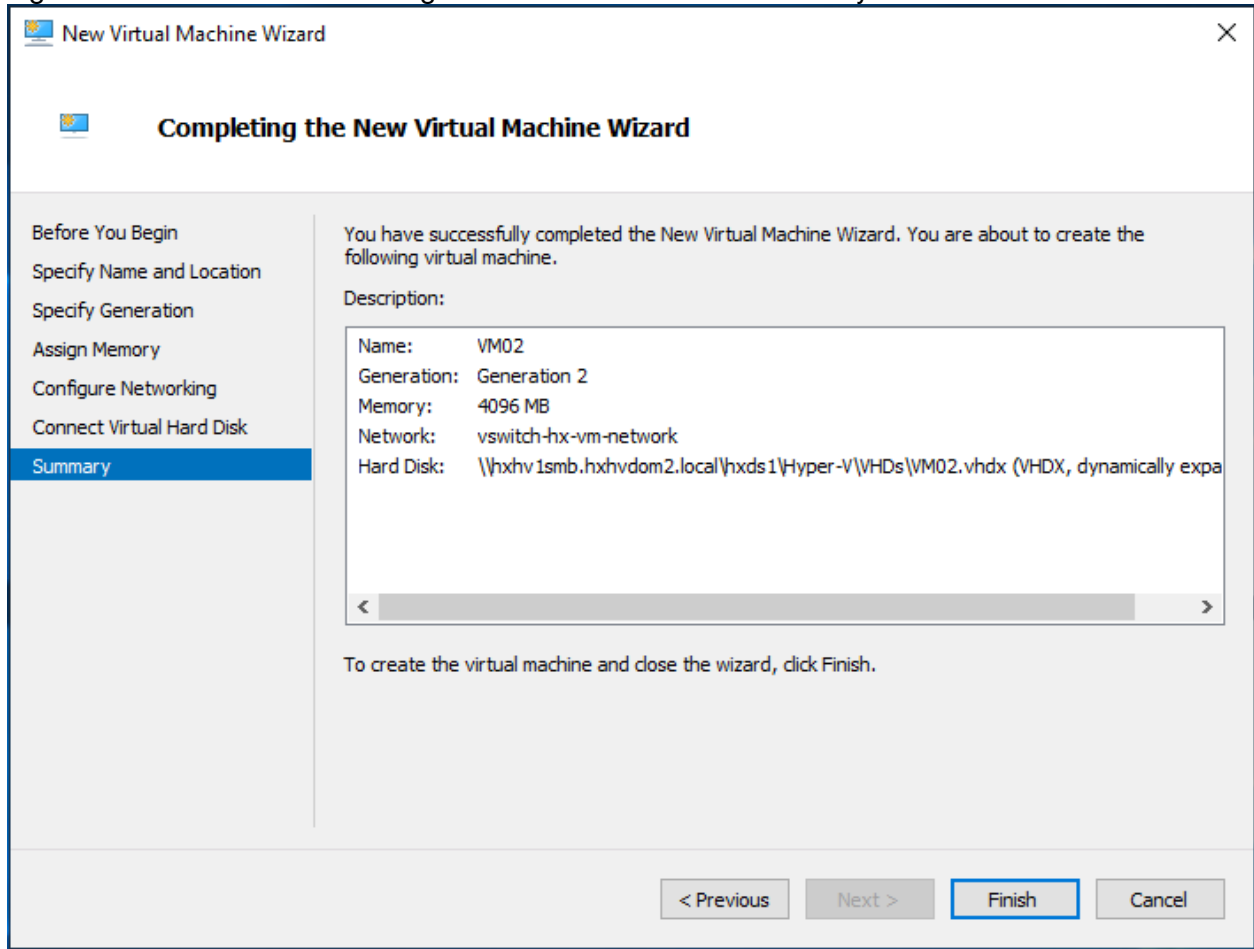
4. In the New Virtual Machine dialog box, search and select the Hyper-V node where you wish to create new virtual machines. Click OK. The New Virtual Machine wizard appears.

Figure 120 Failover Cluster Manager - New Virtual Machine Target Host



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

Figure 121 Failover Cluster Manager – New Virtual Machine Summary



12. After clicking Finish on the Summary page, the VM is configured for high availability as shown in the following screenshots. Next power ON the VM and install the operating system.

Figure 122 Failover Cluster Manager - High Availability Wizard

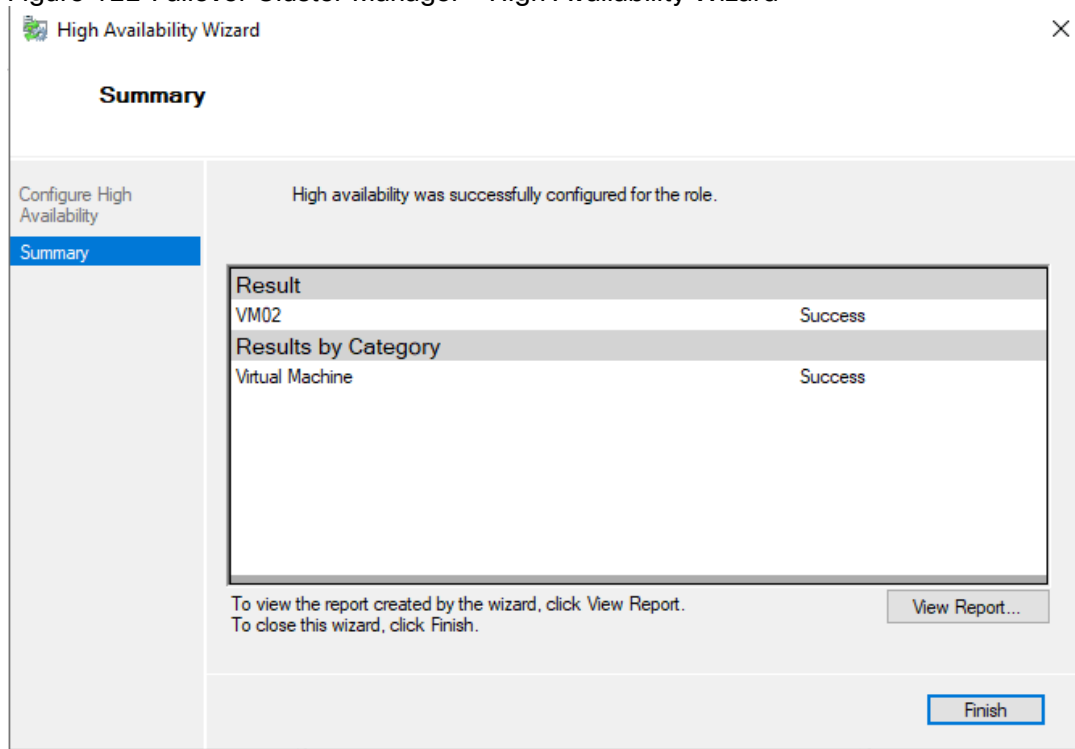
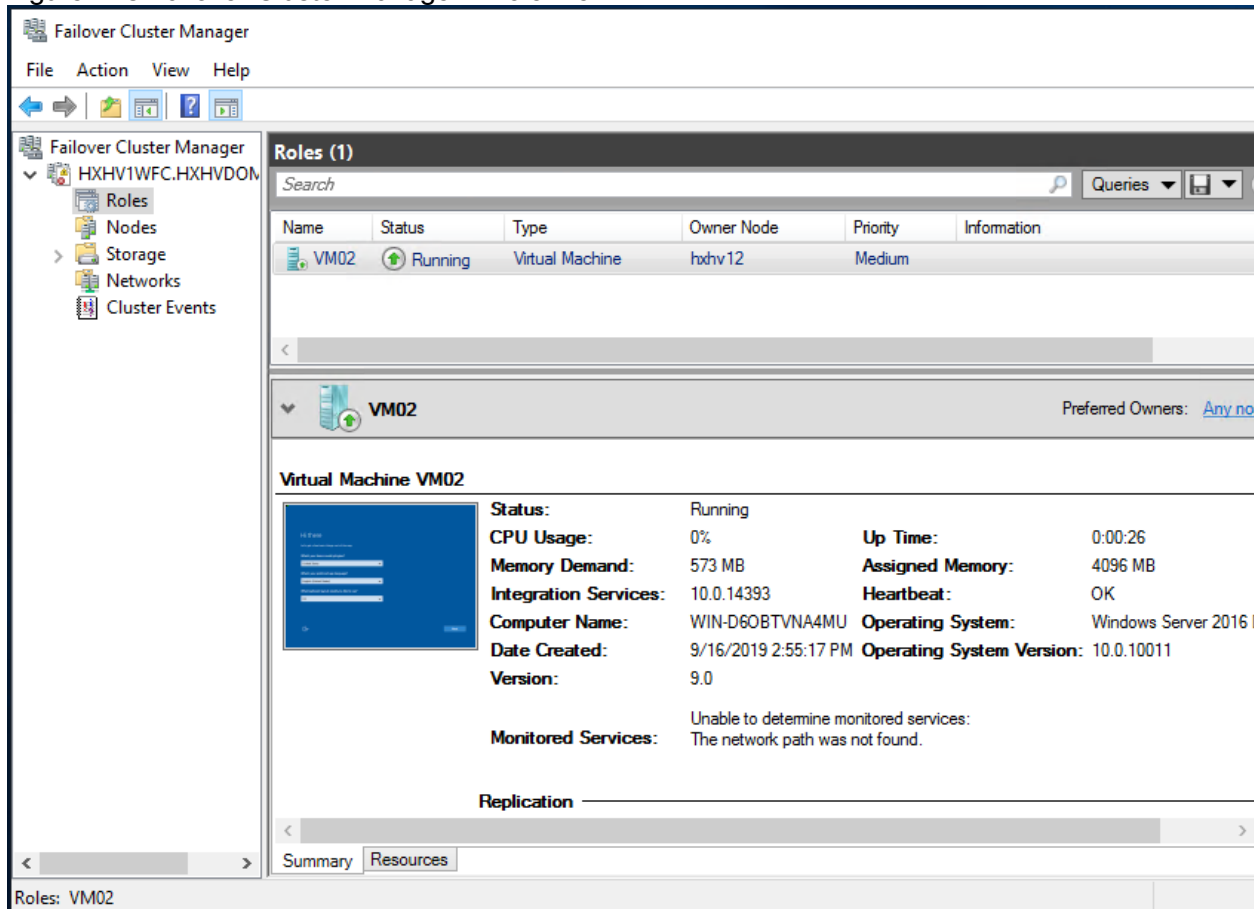


Figure 123 Failover Cluster Manager - Role View



Microsoft System Center Virtual Machine Manager 2019

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

To manage the Hyper-V Failover Cluster in HyperFlex System using SCVMM, follow these steps:

1. Install System Center VMM 2019 to manage Windows Servers 2016/2019 in VMM fabric. For more details on system requirement, refer below link

<https://docs.microsoft.com/en-us/system-center/vmm/system-requirements?view=sc-vmm-2019>

2. Complete the steps described in the above section "Configuring Remote Management Stations for Managing VMs in HX Cluster"



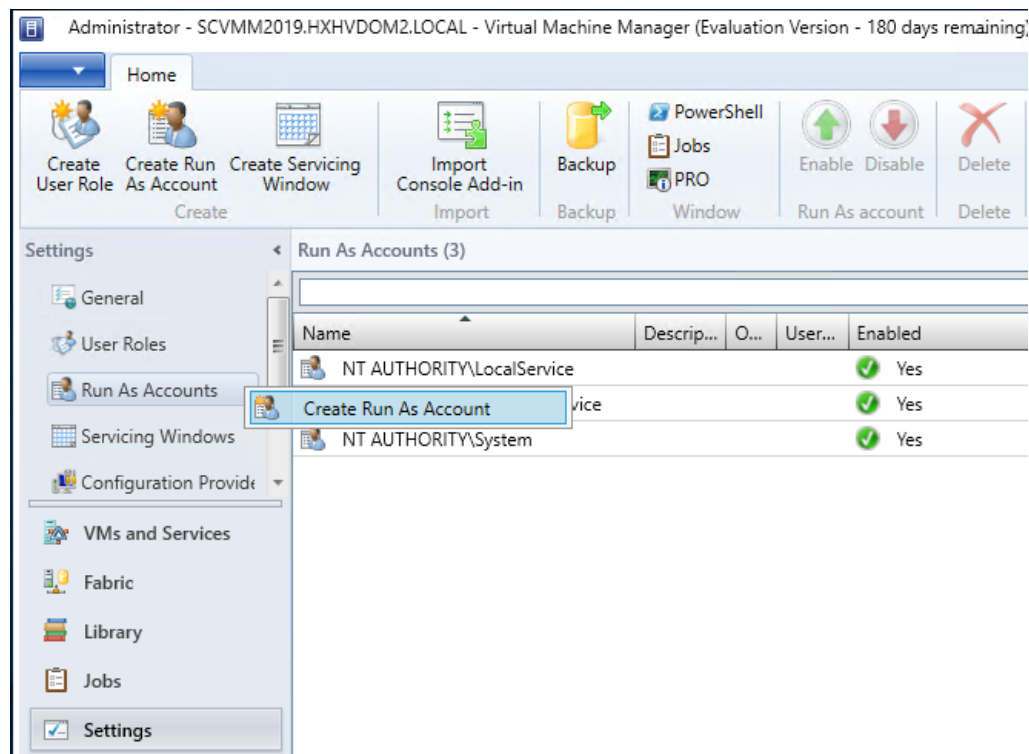
Installing Microsoft SCVMM is beyond the scope of this document. The following steps cover the basic procedure to add the HyperFlex Hyper-V Cluster to SCVMM and configure storage for managing.

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

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

To create a Run As account, follow these steps:

1. Click Settings and in Create click Create Run As Account.



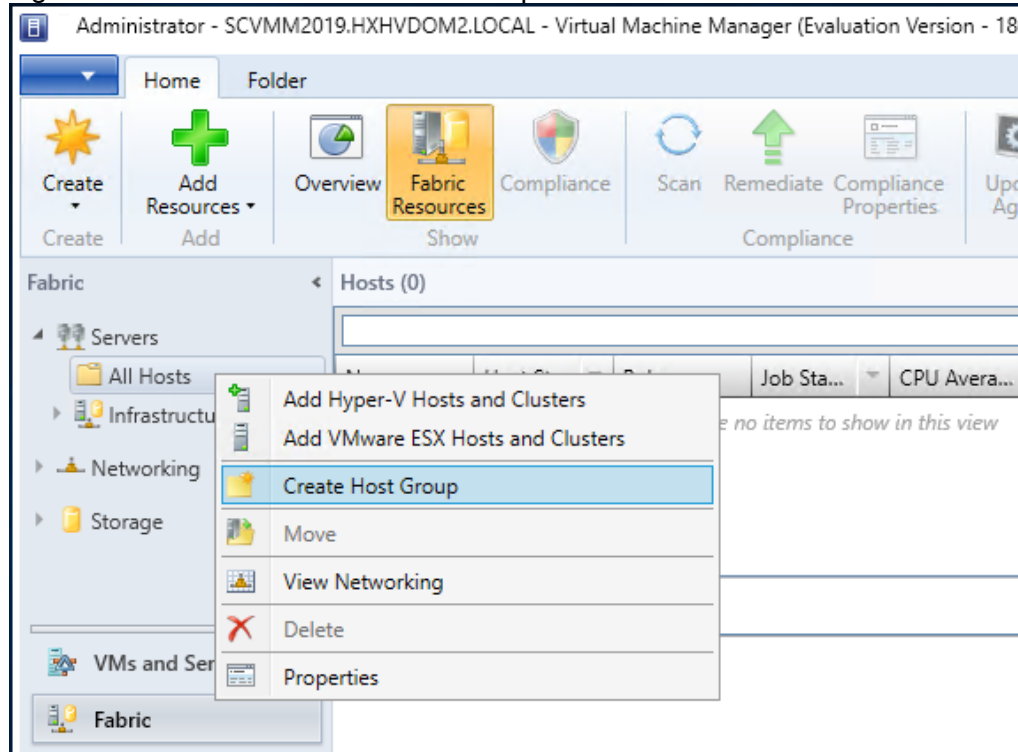
2. In Create Run As Account specify name and optional description to identify the credentials in VMM.
3. In User name and Password specify the credentials. The credentials can be a valid Active Directory user or group account, or local credentials.
4. Clear Validate domain credentials if it is not required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

Manage Servers and Clusters

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

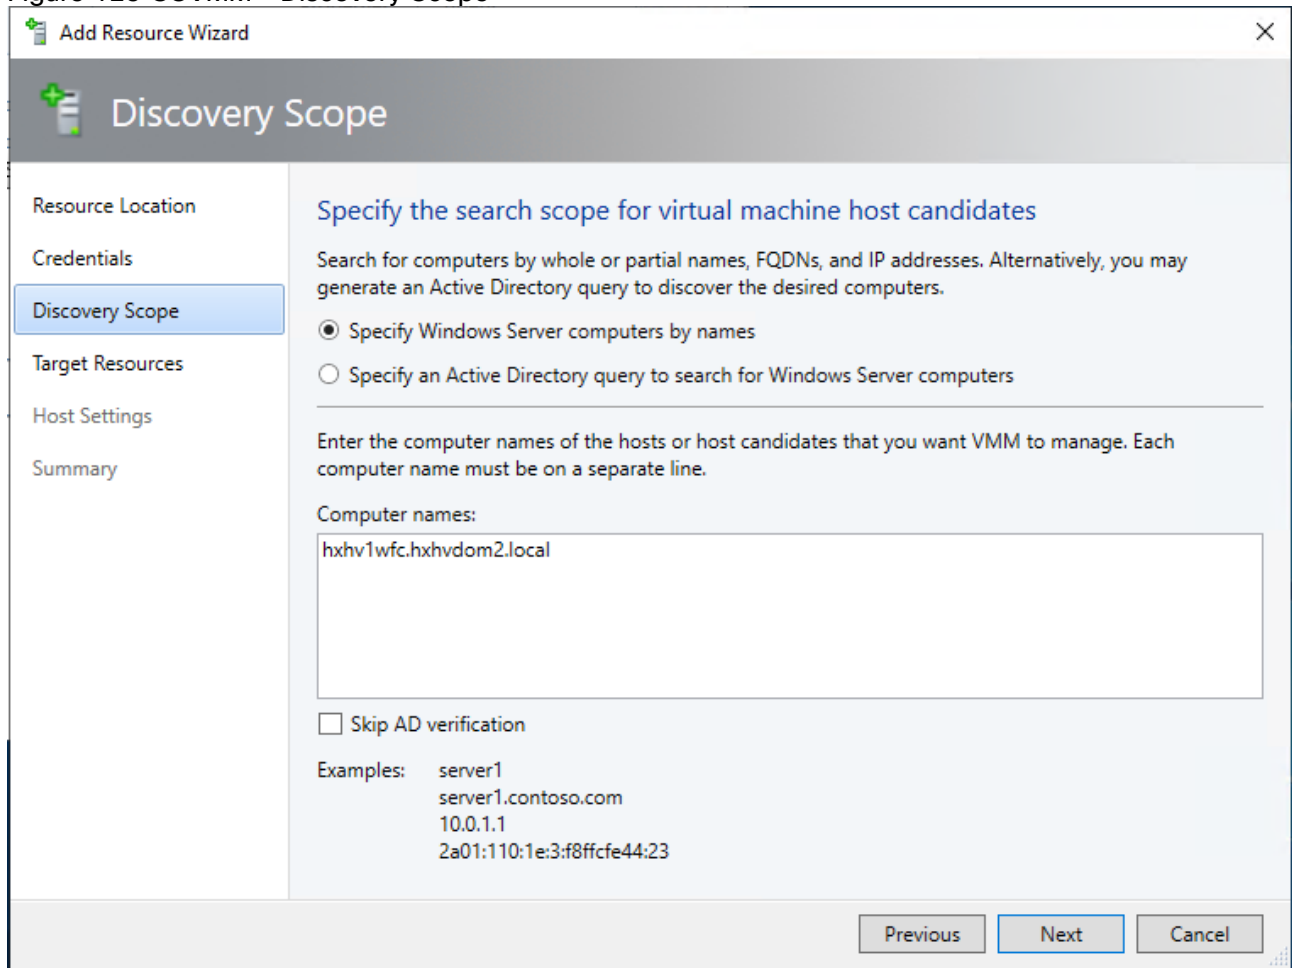
1. Open the SCVMM administrator Console and click Fabric > Servers > All Hosts.
2. Right-click All Hosts and Create a folder for the HyperFlex Hyper-V Cluster.
3. Right-click the newly created folder and click Add Hyper-V Hosts and Clusters.

Figure 124 SCVMM - Create a Host Group



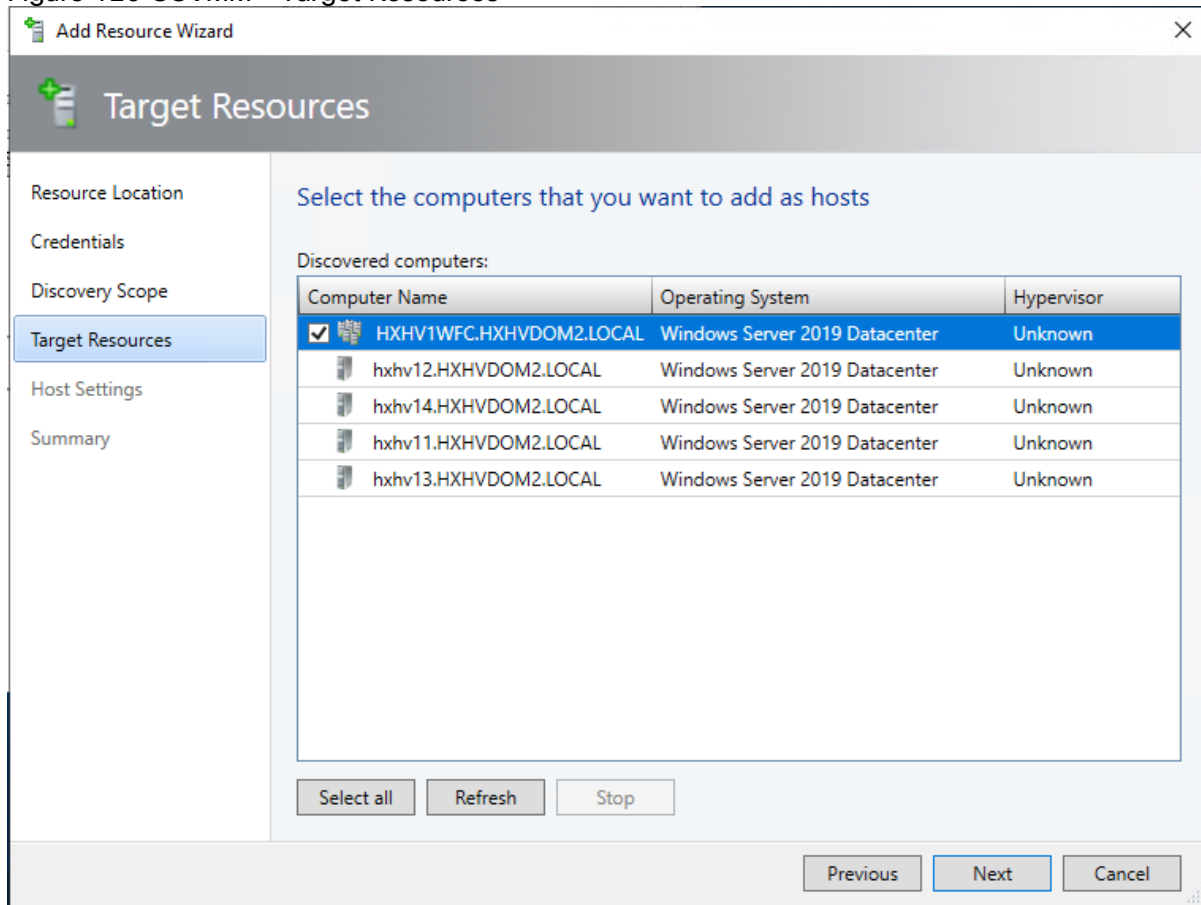
4. In the Credentials section, select Use an existing Run As account and select the account created in the previous section.
5. In the Discovery scope, enter the FQDN of HyperFlex Hyper-V Cluster as shown below.

Figure 125 SCVMM – Discovery Scope



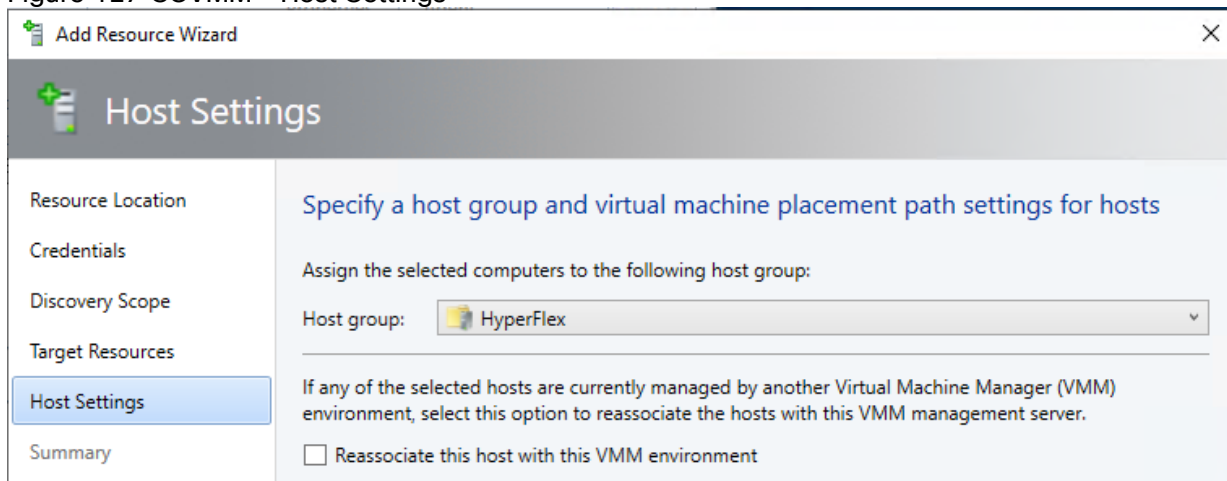
6. In the Target Resources page, select all the discovered hosts and click Next.

Figure 126 SCVMM – Target Resources



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

Figure 127 SCVMM – Host Settings



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

9. A job window pops-up showing the progress of adding virtual machine hosts.

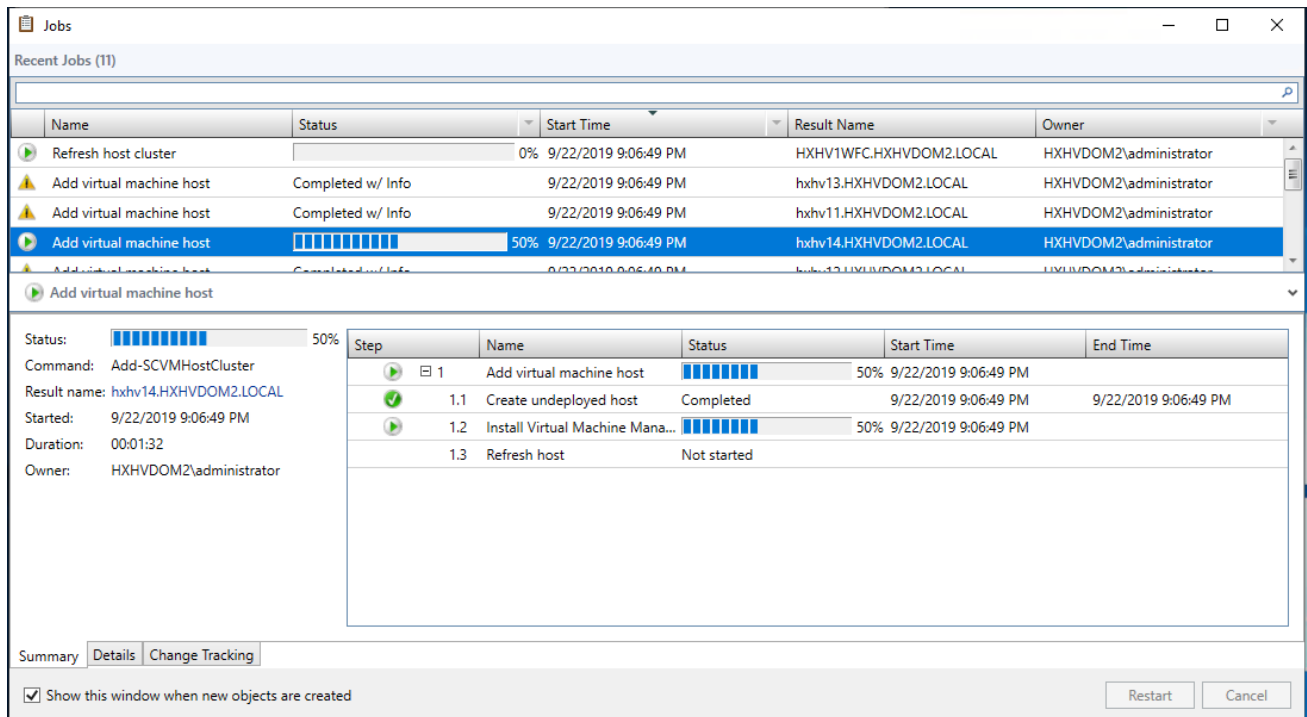
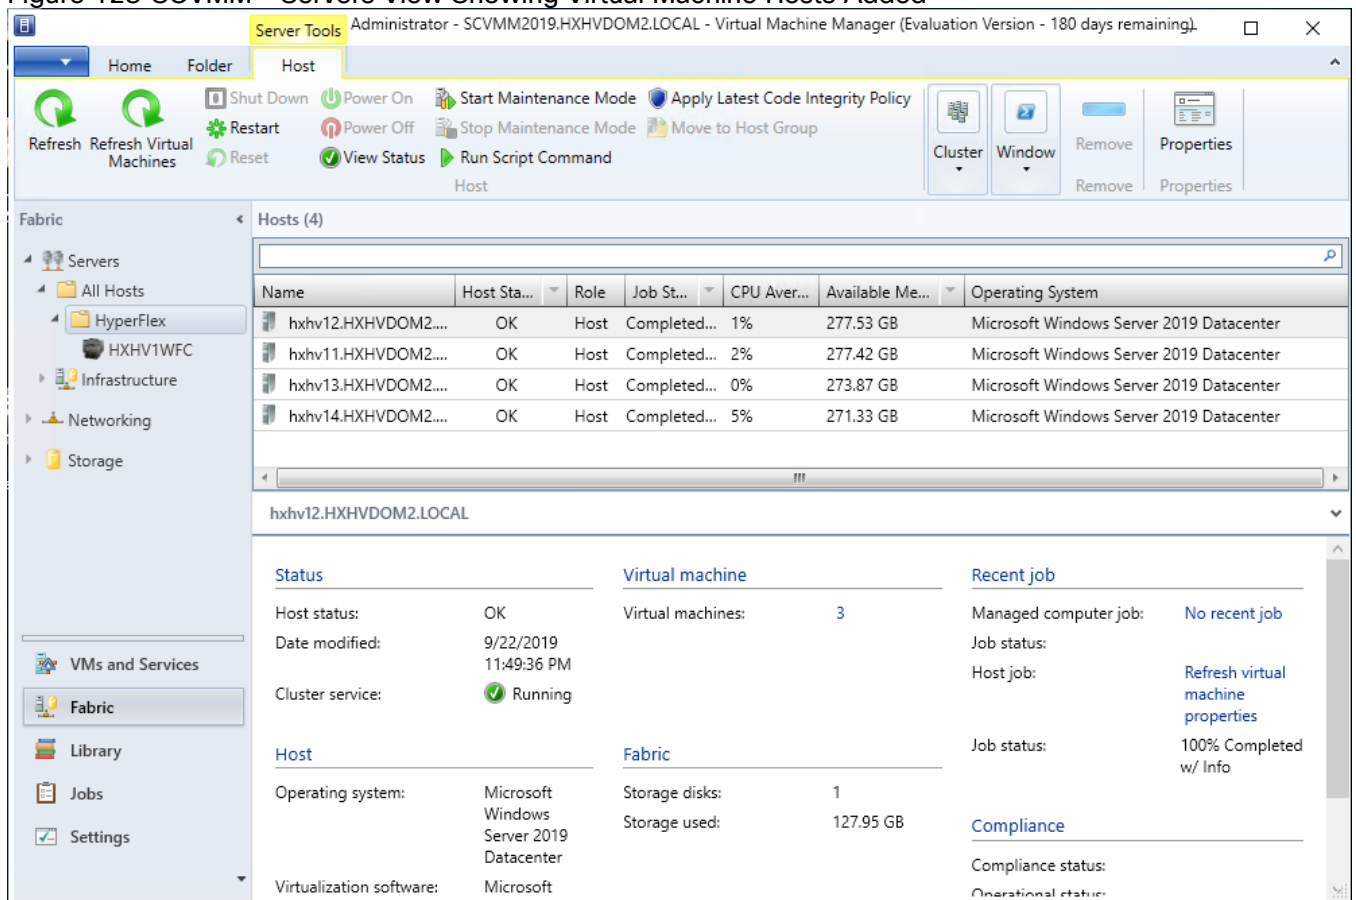


Figure 128 SCVMM - Servers View Showing Virtual Machine Hosts Added



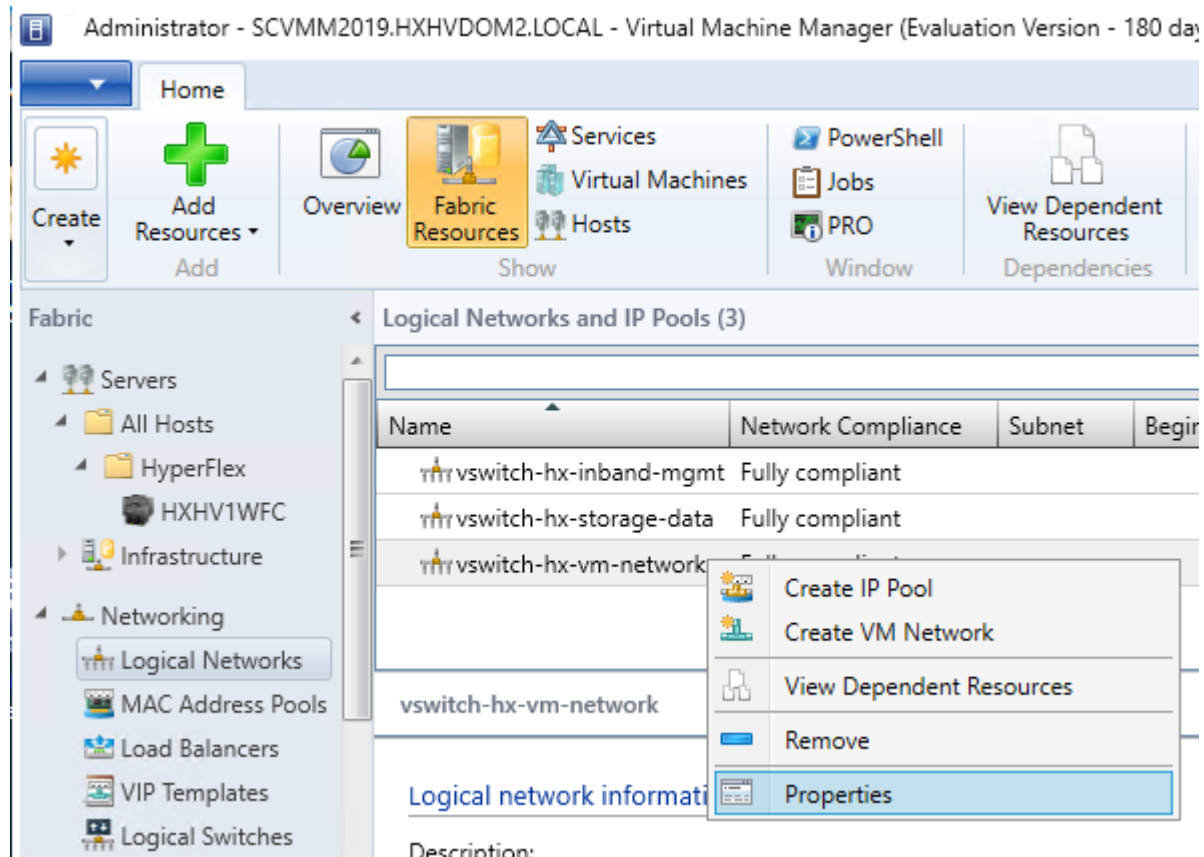
Networking

HyperFlex installer configures the cluster networking end-to-end. Network teams and Hyper-V virtual switches are created for the Management, Storage, VM Network and Live Migration networks specified during the installer.

To create the Network Sites and add them to the logical networks created by the installer, follow these steps:

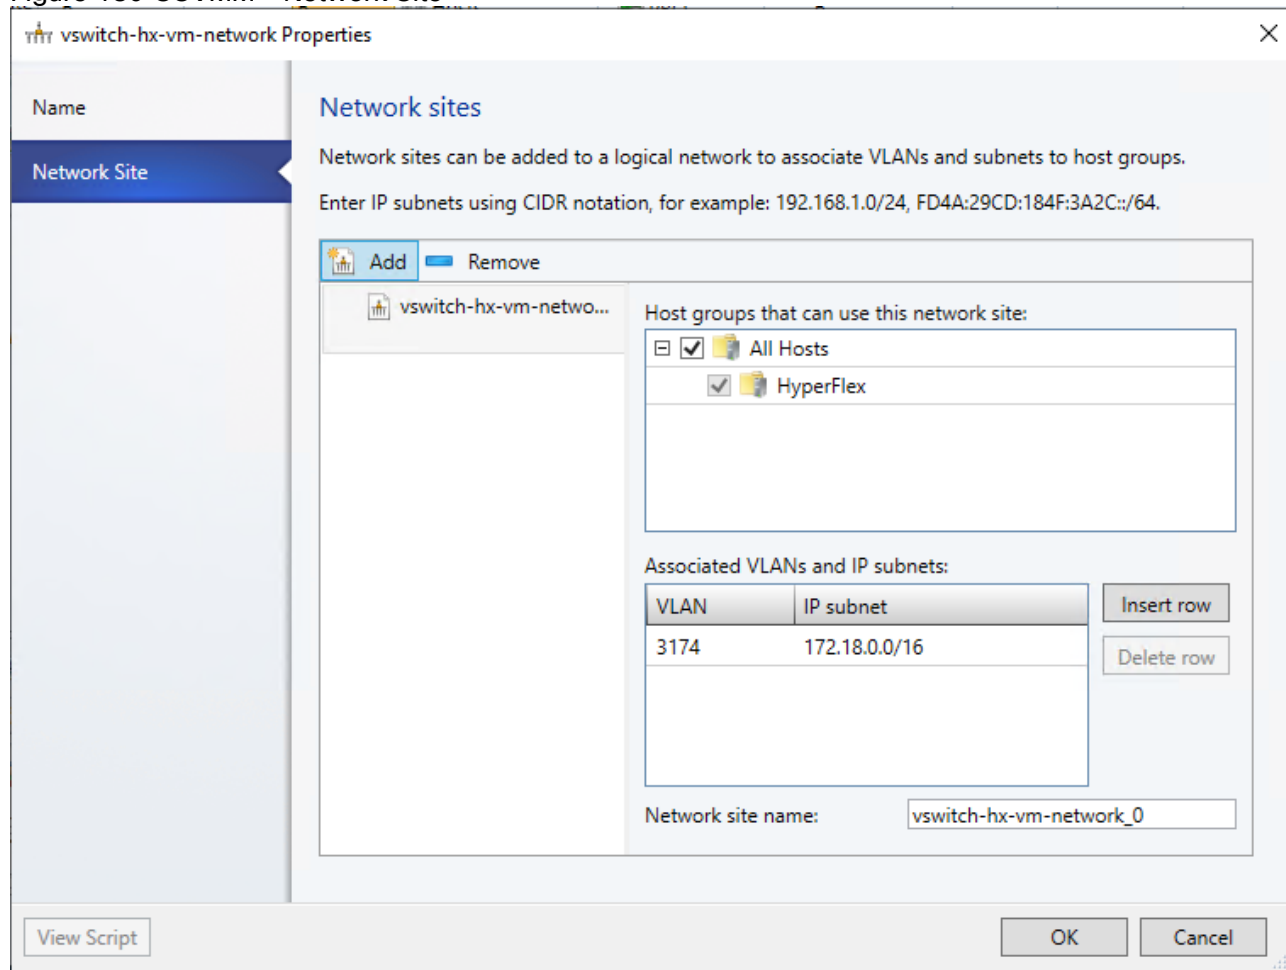
1. Under Fabric -> Networking -> Logical Networks, find the Logical Network created by the installer.
2. Right-click and select 'Properties' of the logical network.

Figure 129 SCVMM - Logical Network



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

Figure 130 SCVMM - Network Site



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

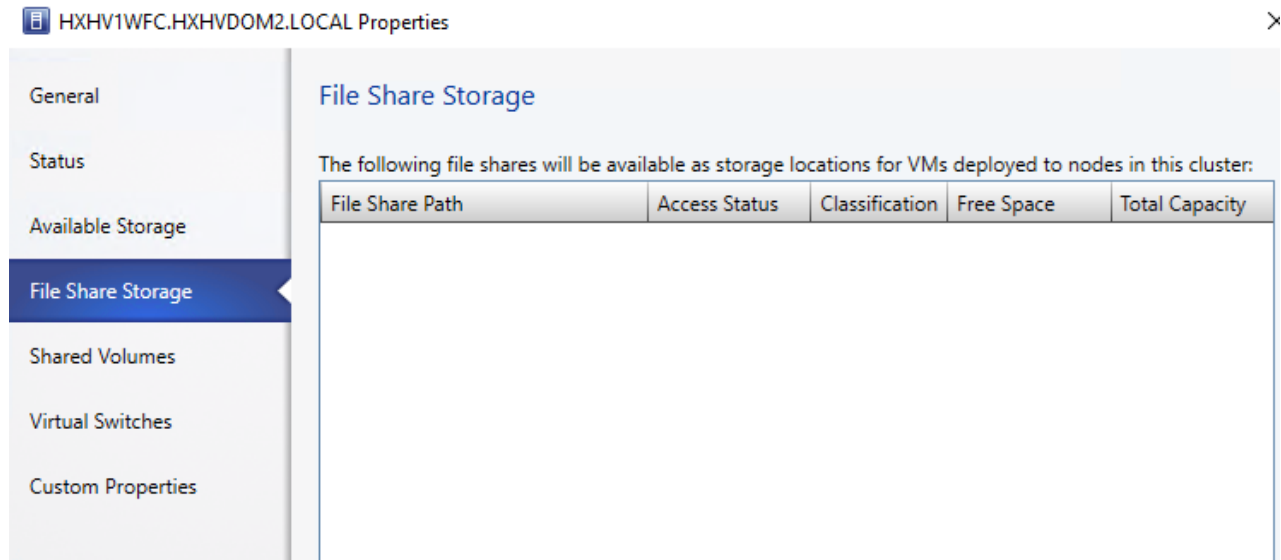
Storage

Datastores created in the HyperFlex Connect, presents a SMB share to be used by the HyperFlex Hyper-V nodes to place Virtual Machine files on it. The naming convention is '\\<hxClustername>\<DatastoreName>'.

To add the HX Datastores (SMB Share path) to the Hyper-V Cluster nodes, follow these steps:

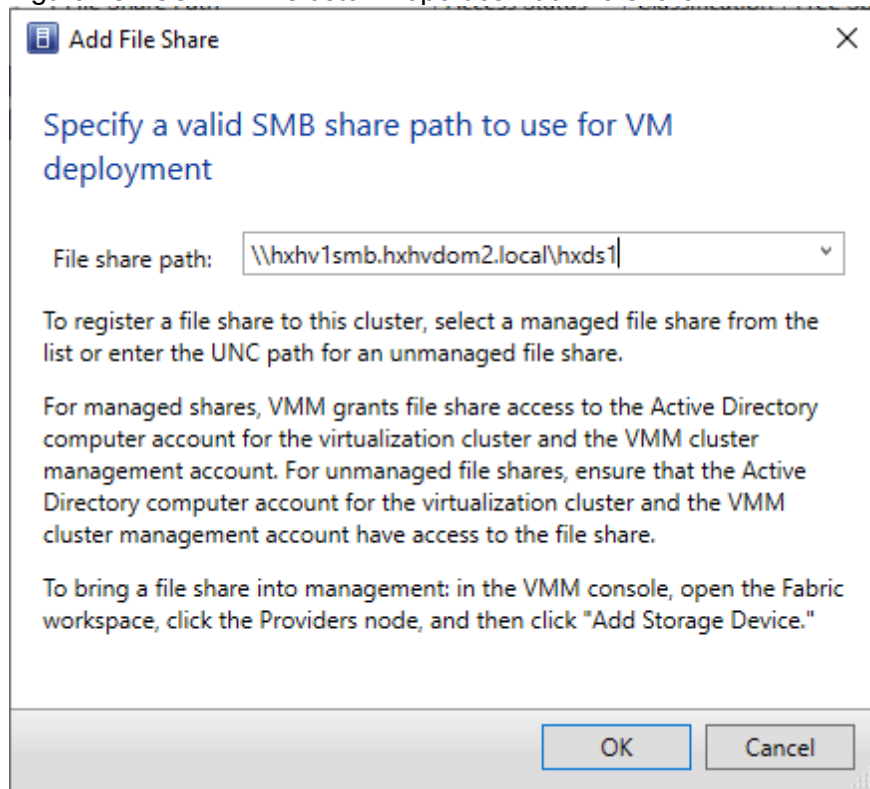
- Right-click the Cluster 'HXHWFC', select Properties and click 'File Share Storage'.

Figure 131 SCVMM – Cluster Properties File Share Storage




2. Click Add to specify the UNC path for the datastore and enter the File share path.

Figure 132 SCVMM – Cluster Properties Add File Share



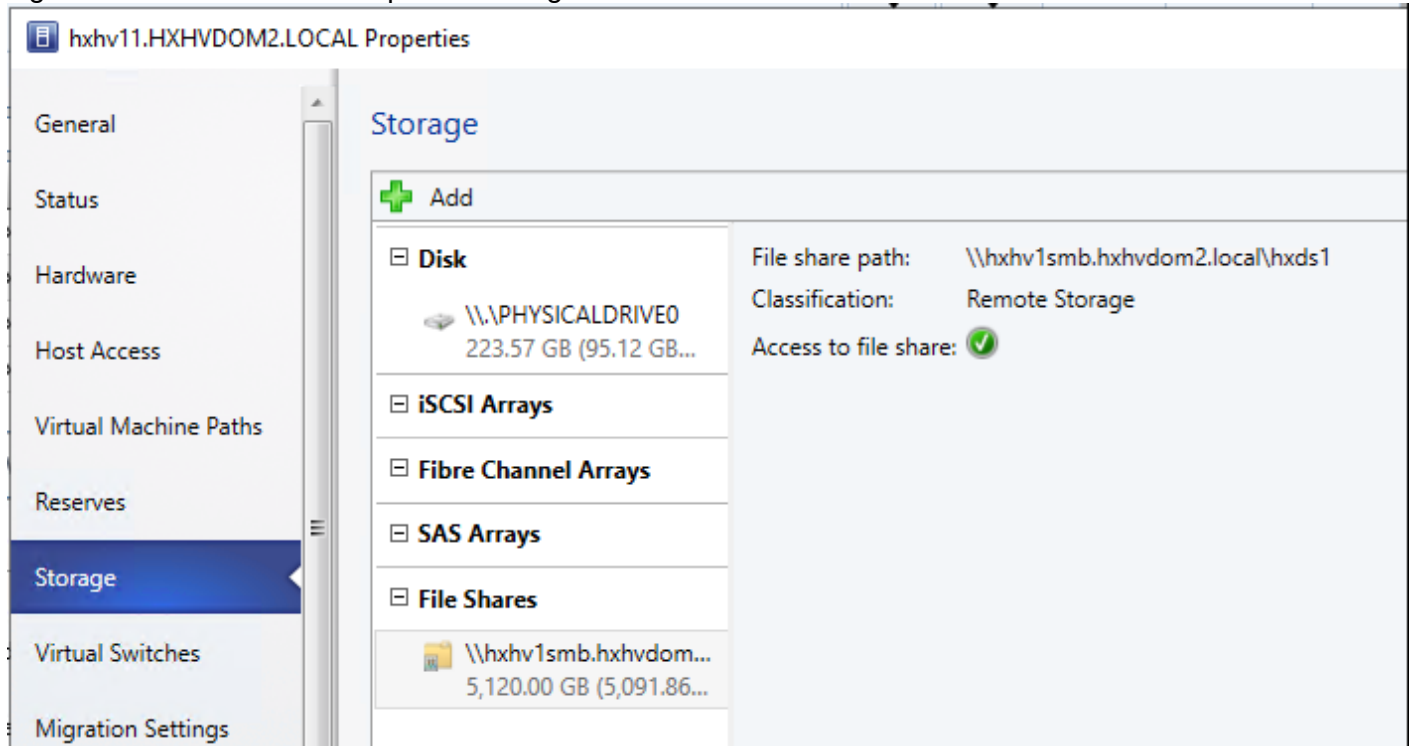
3. Click OK.
4. Repeat steps 2 and 3 to add other datastores if you want to deploy new virtual machines from SCVMM.
5. Right-click and click the host cluster and the access status of the file share path turns green as shown below:

Figure 133 SCVMM – Cluster Properties File Share Storage Status



- Verify the above configuration is applied to all Hyper-V nodes in the cluster by checking their properties by clicking Storage as shown below.

Figure 134 SCVMM – Host Properties Storage



Create a Virtual Machine using SCVMM

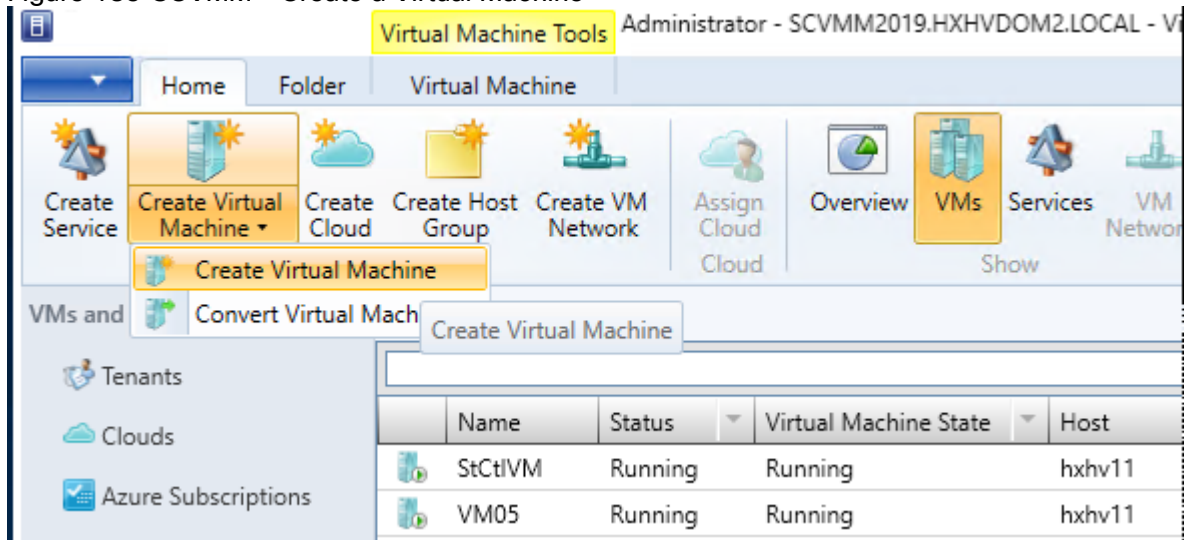
To create and deploy virtual machines in the System Center – Virtual Machine Manager (VMM) fabric, from an existing virtual hard disk (VHDX) that has been generalized using Sysprep and copied to the VMM library, follow these steps:



Using SCVMM, you have the option to either create highly available or non-highly available virtual machines.

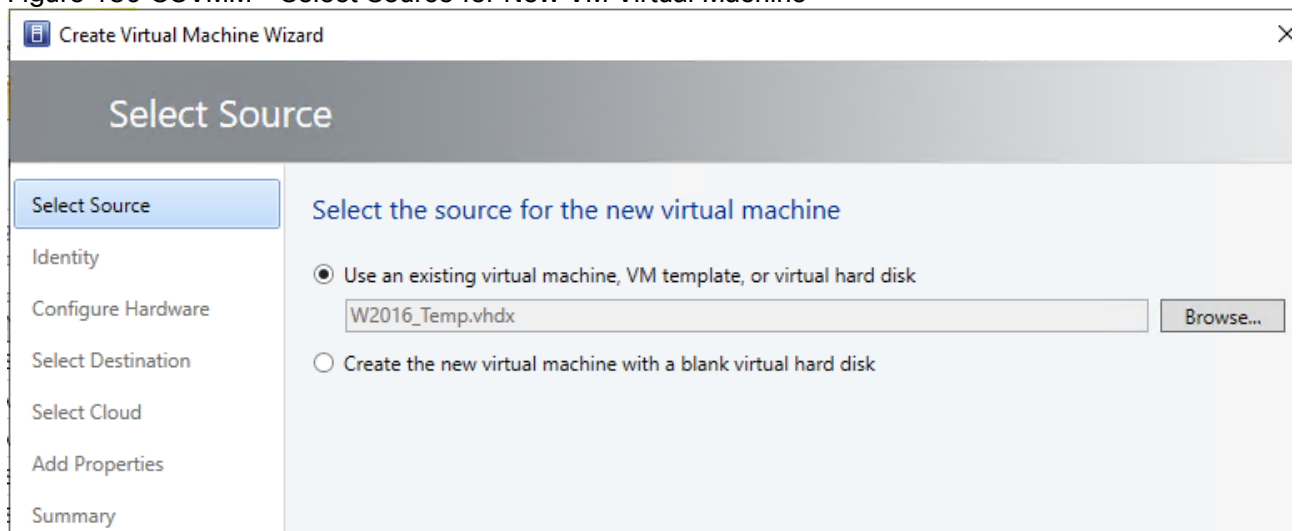
- Click Virtual Machines and Services > Create Virtual Machine > Create Virtual Machine.

Figure 135 SCVMM - Create a Virtual Machine



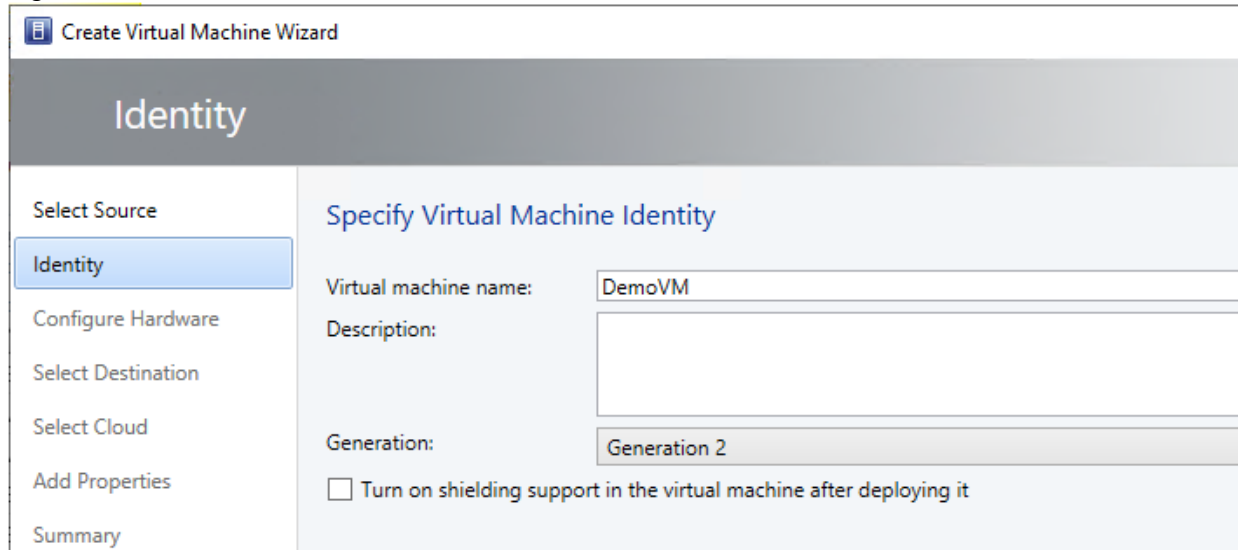
- In Create Virtual Machine Wizard > Select Source, click Use an existing virtual machine, virtual machine template, or virtual hard disk > Browse. Select an existing VHD.

Figure 136 SCVMM - Select Source for New VM Virtual Machine



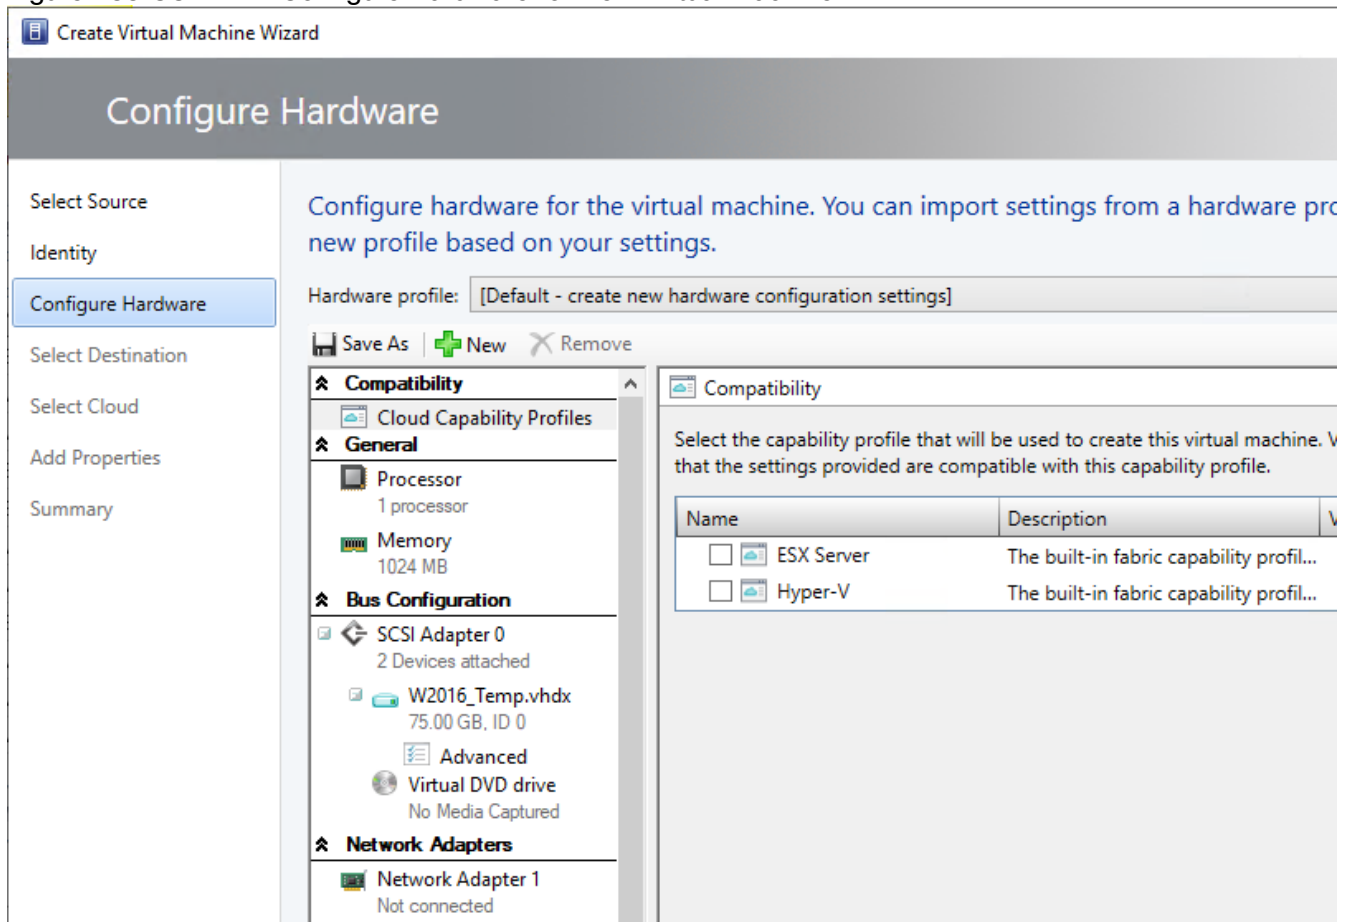
- In Identity, specify the virtual machine name and an optional description. If the VHD you choose is in the .vhdx format, in the Generation box, select Generation 1 or Generation 2. Click Next.

Figure 137 SCVMM - Create New Virtual Machine Name and Generation



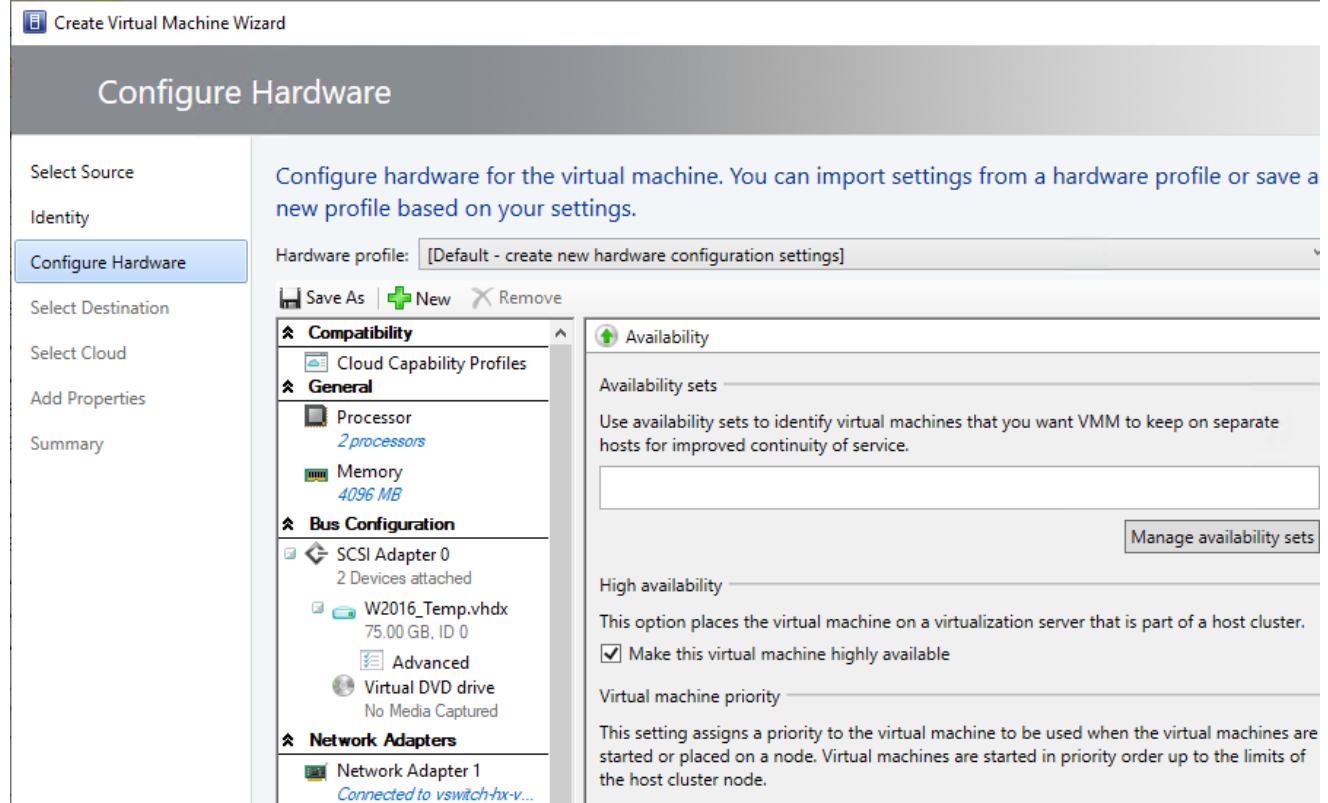
- In Configure Hardware, either select the profile that you want to use from the Hardware profile list or configure the hardware settings manually. Click Next.

Figure 138 SCVMM - Configure Hardware for New Virtual Machine



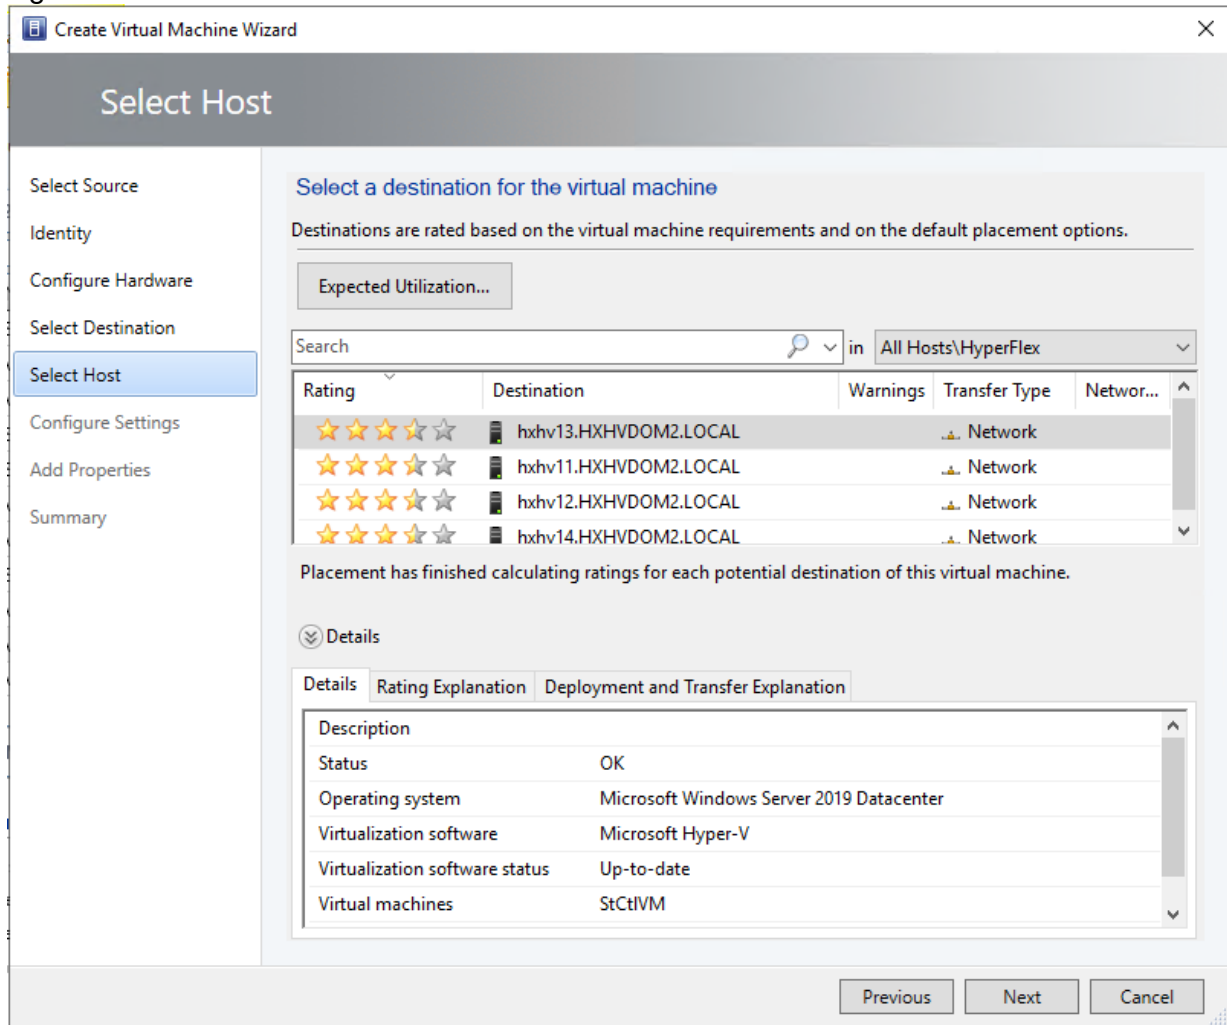
- In the Configure Hardware > Advanced > Availability section, you have the option to make the virtual machine highly available by selecting the 'Make the virtual machine highly available'.

Figure 139 SCVMM - Configure Hardware for Highly Available Virtual Machine



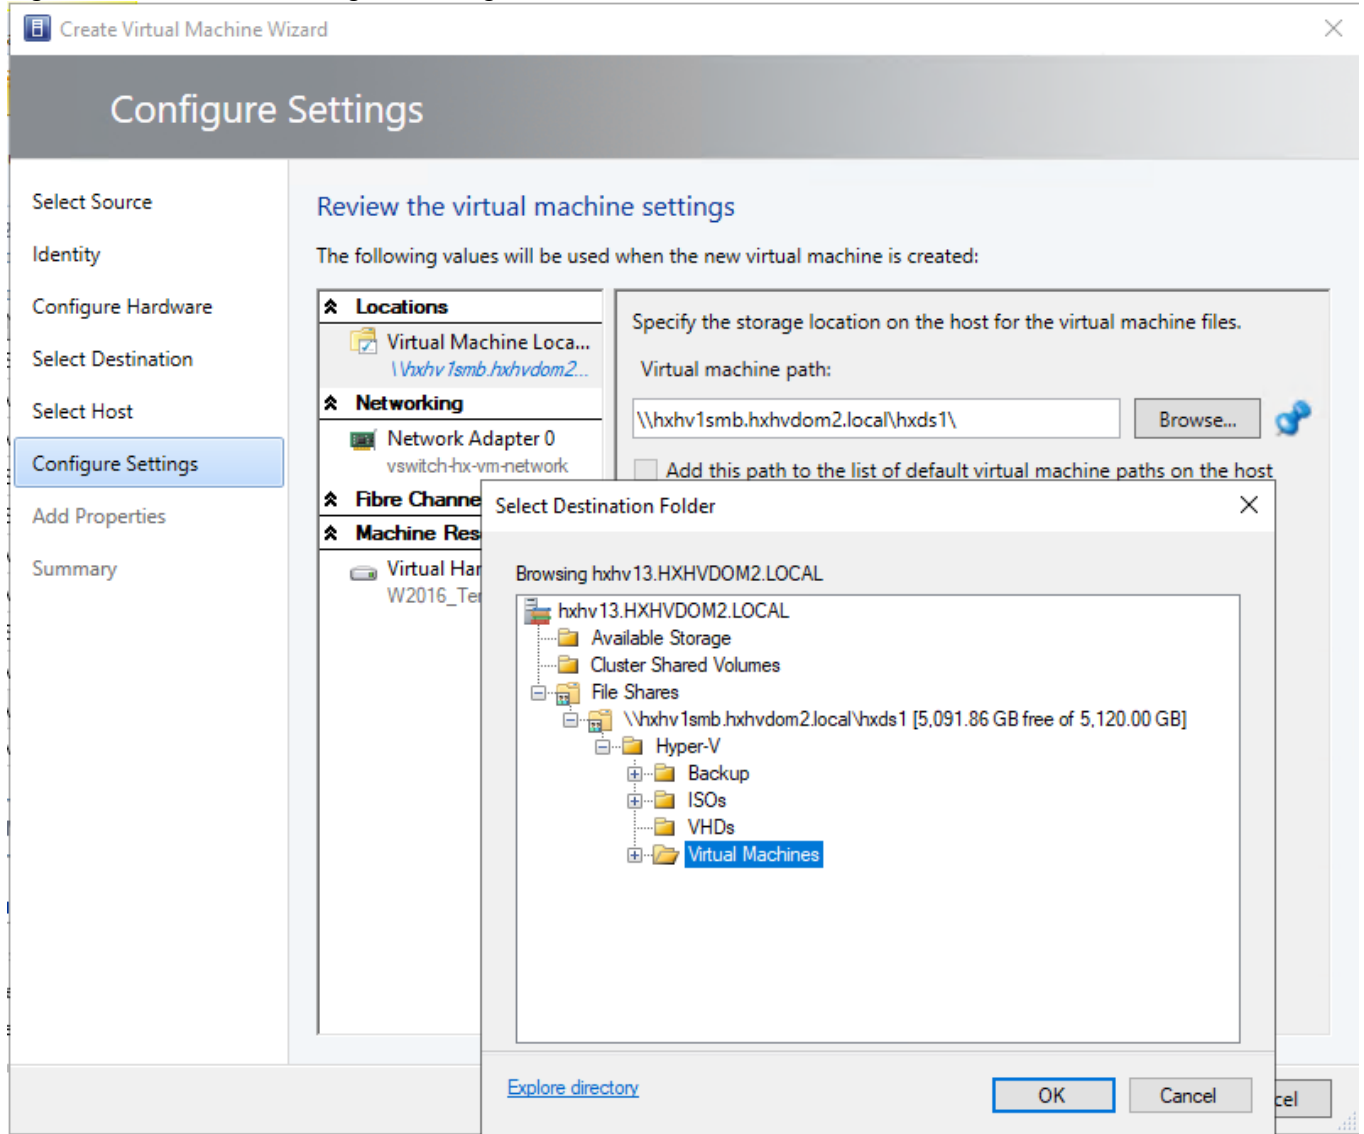
6. In Select Destination, place the virtual machine on a host by selecting the destination folder. Click Next.
7. In Select Host, select a Hyper-V host for the virtual machine or leave with default selection. Click Next.

Figure 140 SCVMM - Select Host for New Virtual Machine



8. In Configure Settings, under locations browse to the HX Datastore SMB share mapped in previous sections and select a vSwitch for the Network Adapter under Networking. Click Next.


Figure 141 SCVMM - Configure Settings for New Virtual Machine



9. In Add Properties section, select appropriate actions and click Next.

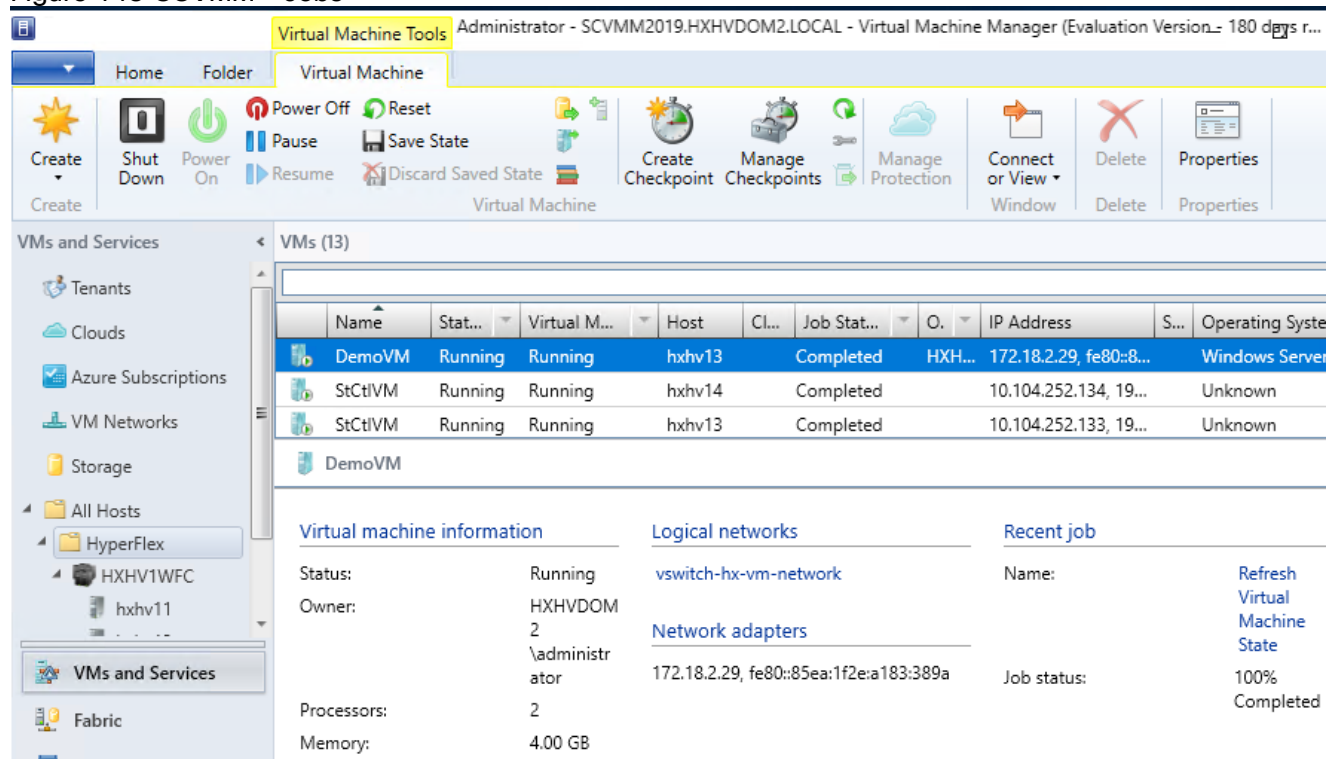
10. In the Summary page, review and confirm the settings and click Create.

Figure 142 SCVMM - Jobs



11. In the VMs and Services page, right-click the newly created VM and power it ON.

Figure 143 SCVMM - Jobs



PowerShell

Windows PowerShell is a Windows command-line shell designed especially for system administrators. Windows PowerShell includes an interactive prompt and a scripting environment that can be used independently or in combination. It comes installed by default in every Windows, starting with Windows 7 SP1 and Windows Server 2008 R2 SP1. Using PowerShell, the HyperFlex Hyper-V cluster environment can be managed locally or remotely from a Windows management host running PowerShell (latest version recommended).

The figure below shows an example to create a virtual machine on a HX Hyper-V node from a remote management station.

Figure 144 PowerShell – Create a New Virtual Machine

```
PS C:\Windows\System32\WindowsPowerShell\v1.0> Enter-PSSession -ComputerName hxhvv2
[hxhvv2]: PS C:\Users\Administrator.HXHVDOM\Documents> New-VM -Name VM02 -Generation 2 -MemoryStartupBytes 4GB -VHDPath "\\hxhvsmb.hxhvdom.local\HXD51\Hyper-V\VHDs
Name State CPUUsage(%) MemoryAssigned(M) Uptime Status Version
-----
VM02 off 0 0 00:00:00 Operating normally 8.0

[hxhvv2]: PS C:\Users\Administrator.HXHVDOM\Documents> Set-VM -Name VM02 -ProcessorCount 2
[hxhvv2]: PS C:\Users\Administrator.HXHVDOM\Documents> Set-VMNetworkAdapterVlan -VMName VM02 -Access -VlanId 3174
[hxhvv2]: PS C:\Users\Administrator.HXHVDOM\Documents> Start-VM -Name VM02
[hxhvv2]: PS C:\Users\Administrator.HXHVDOM\Documents> Get-VM
Name State CPUUsage(%) MemoryAssigned(M) Uptime Status Version
-----
STCCLVM Running 0 49152 6:04:46:36.5160000 Operating normally 8.0
VM02 Running 0 4096 00:00:18.3190000 Operating normally 8.0
```

Microsoft Windows Admin Center (WAC)

Microsoft has recently launched a management tool called "Windows Admin Center". It is a locally deployed, browser-based app for managing servers, clusters, hyper-converged infrastructure, and Windows 10 PCs. It comes at no additional cost beyond Windows and is ready to use in production. Windows Admin Center is the modern evolution of "in-box" management tools, like Server Manager and MMC. It complements System Center - it's not a replacement.



Microsoft Windows Admin Center is a management tool launched recently and it is evolving. Cisco has not extensively tested WAC to manage HyperFlex Hyper-V Cluster.

For more information about Windows Admin Center, refer to [What is Windows Admin Center?](#)

To download Windows Admin Center, refer to [Hello, Windows Admin Center!](#)

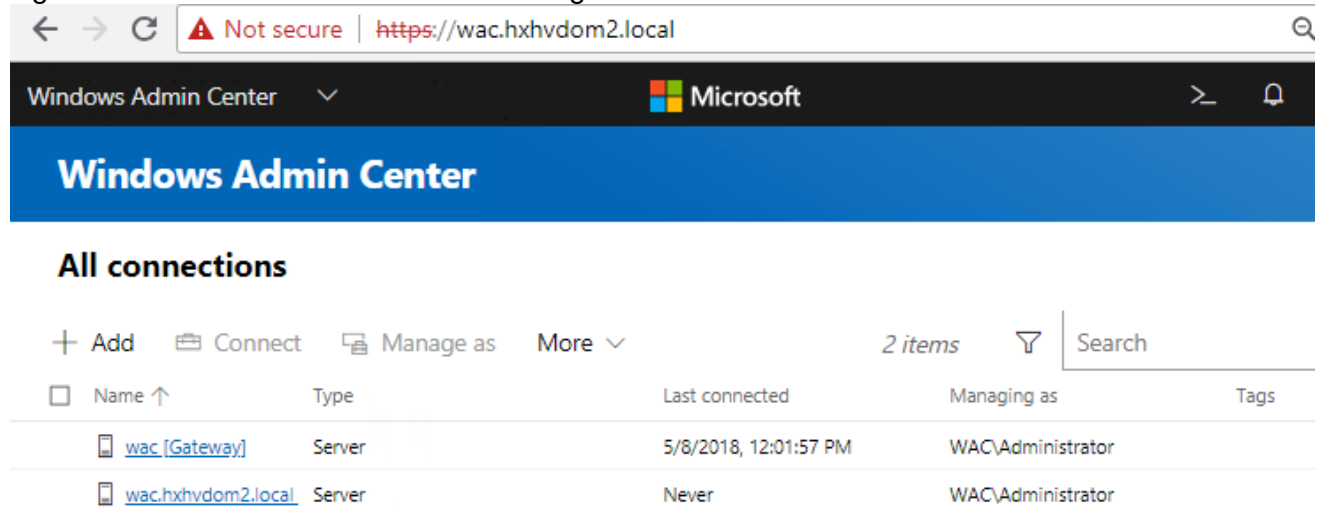
To install Windows Admin Center, refer to [Install Windows Admin Center](#)

Connect to Managed Nodes and Clusters

After you have completed the installation of Windows Admin Center, you can add servers or clusters to manage from the main overview page. To add a single server or a cluster as a managed node, follow these steps:

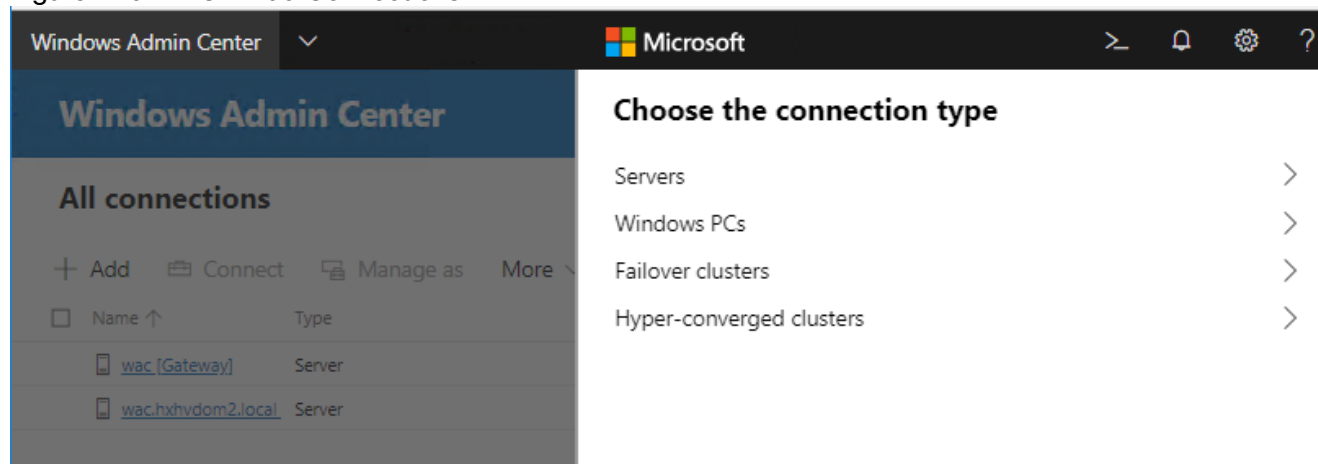
1. Open a browser and launch the Windows Admin Center.
2. Click + Add under All Connections.

Figure 145 Windows Admin Center Home Page



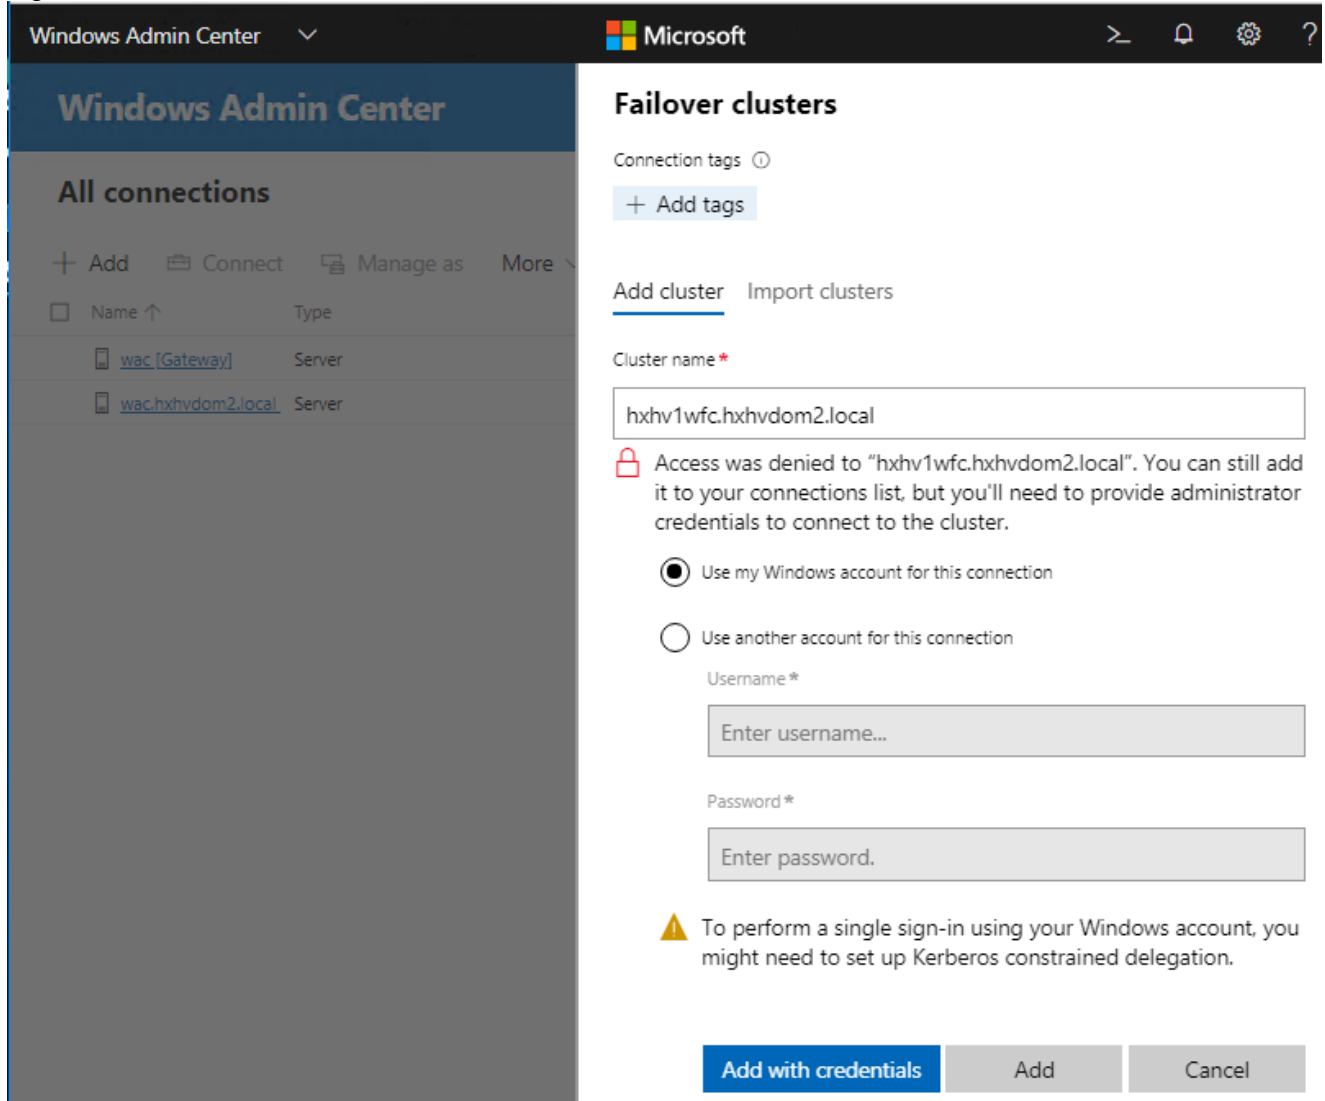
3. Choose to add a Server, Failover Cluster connection.

Figure 146 WAC – Add Connections



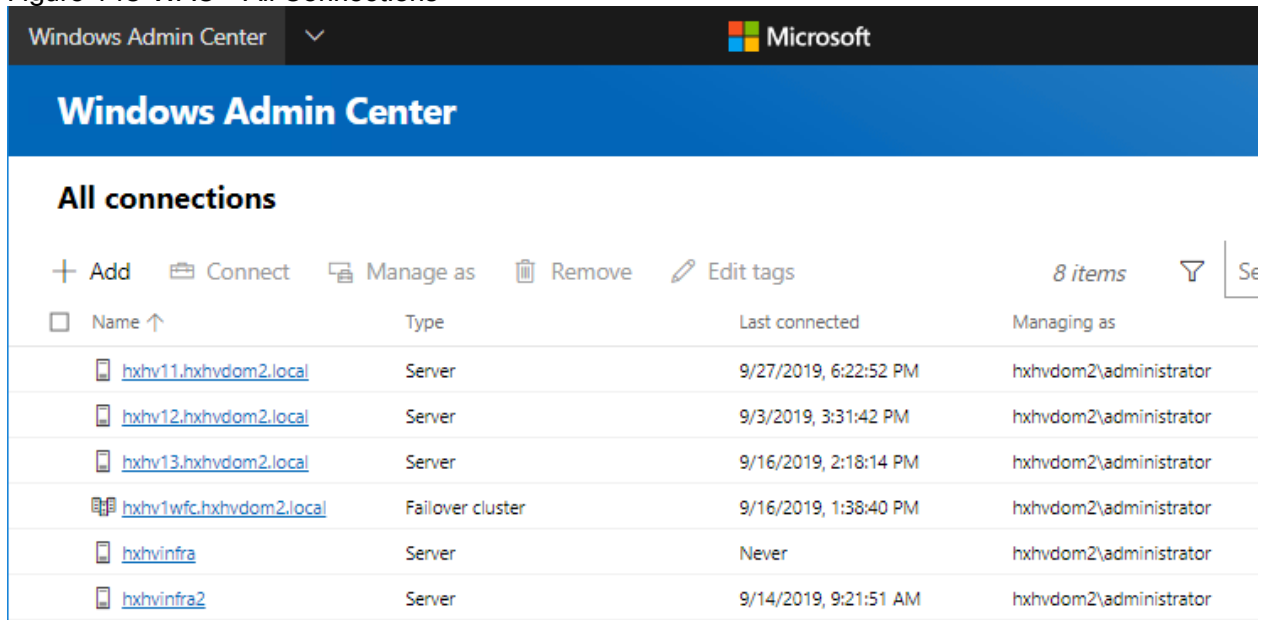
4. Type the name of the server or cluster to manage and click Submit. The server or cluster will be added to your connection list on the overview page. In this example, Hyper-V cluster is added to manage.

Figure 147 WAC - Add Failover Cluster Connections



5. Authenticate with a managed node using 'Single Sign-on' or 'Manage As' by entering the credentials.
6. Select a server/cluster and click Edit Tags to organize your connections. This will help to filter your connection lists.

Figure 148 WAC – All Connections

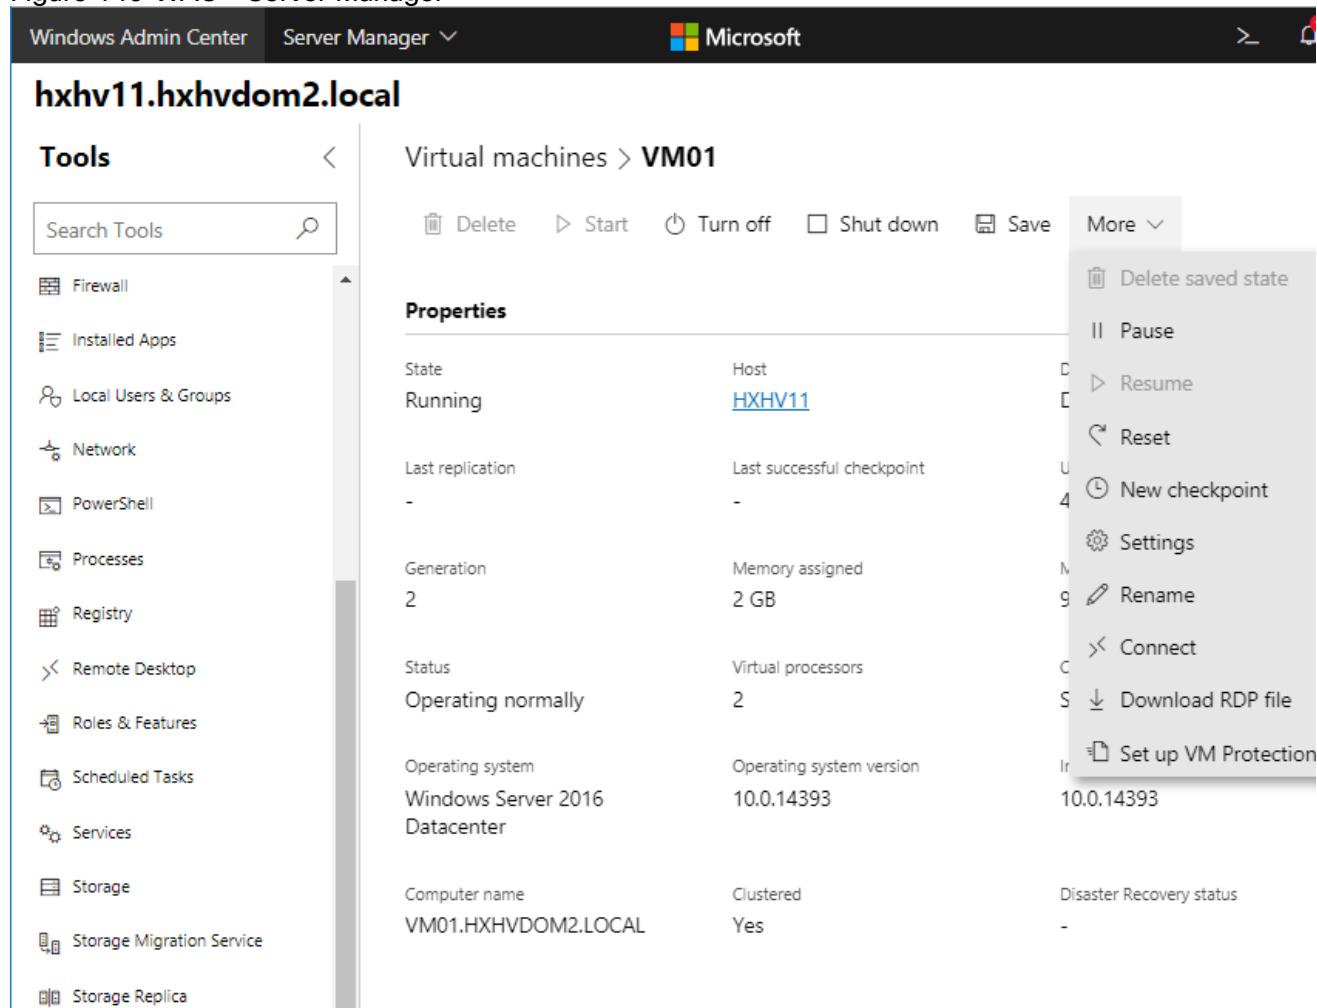


Manage Servers with WAC

To add and manage individual servers running Windows Server 2012 or later to Windows Admin Center with a comprehensive set of tools including Devices, Events, Processes, Roles and Features, Updates, Virtual Machines and more, follow these steps:

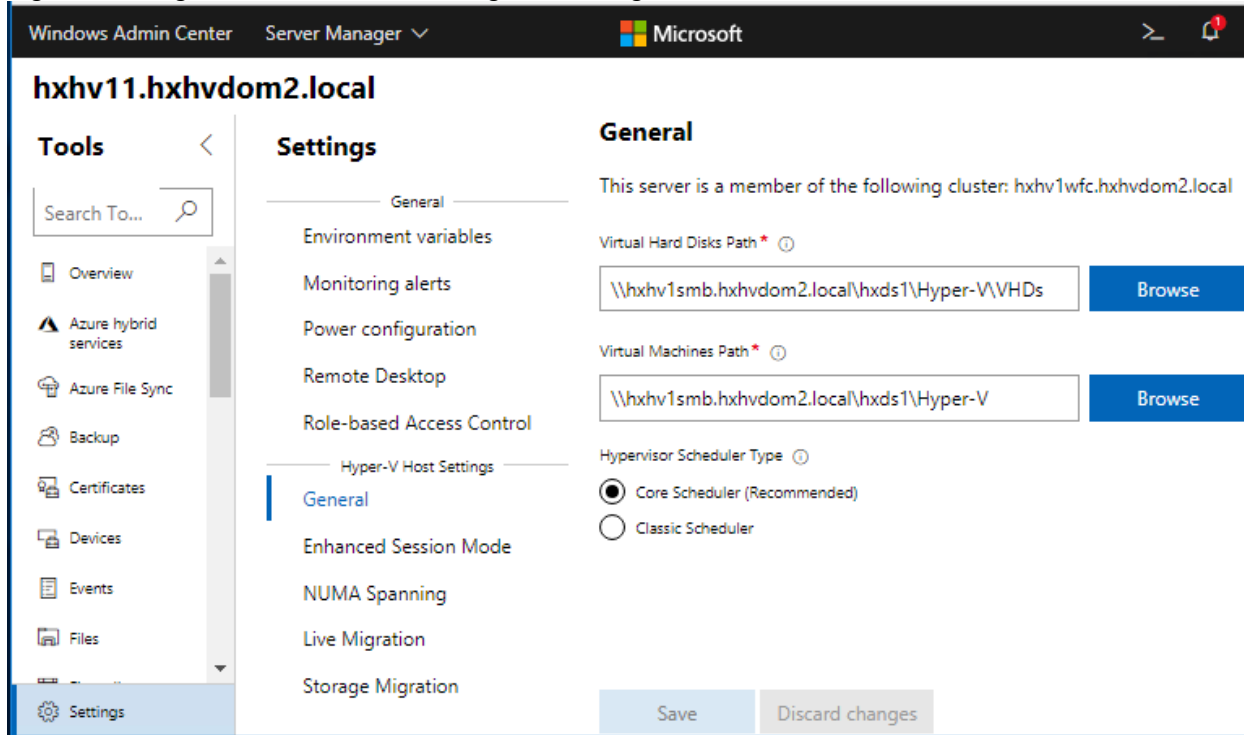
1. Launch WAC from a browser.
2. Click a server under All Connections. The figure below shows the tools available to manage servers.

Figure 149 WAC – Server Manager



3. Click Settings under the Tools pane and change the virtual machine files default location as shown below:

Figure 150 Figure WAC - Server Manager - Settings



Manage a Failover Cluster with WAC

To manage a failover cluster with WAC, follow these steps:

1. Add Failover clusters to view and manage cluster resources, storage, network, nodes, roles, virtual machines and virtual switches.
2. Launch WAC from a browser.
3. Select and click the cluster under All Connections. The figure below shows the tools available to view and manage Cluster.

Figure 151 WAC – Failover Cluster Manager

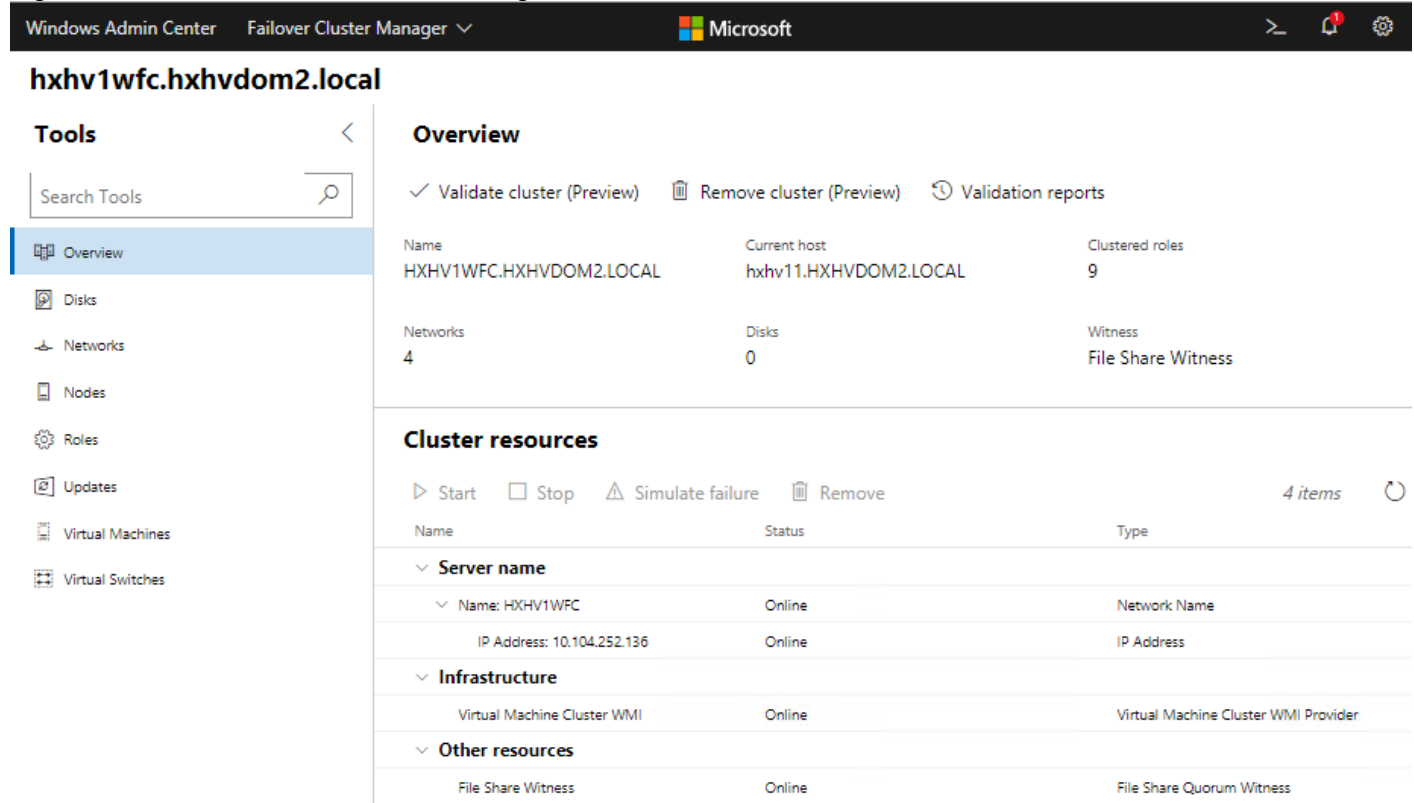
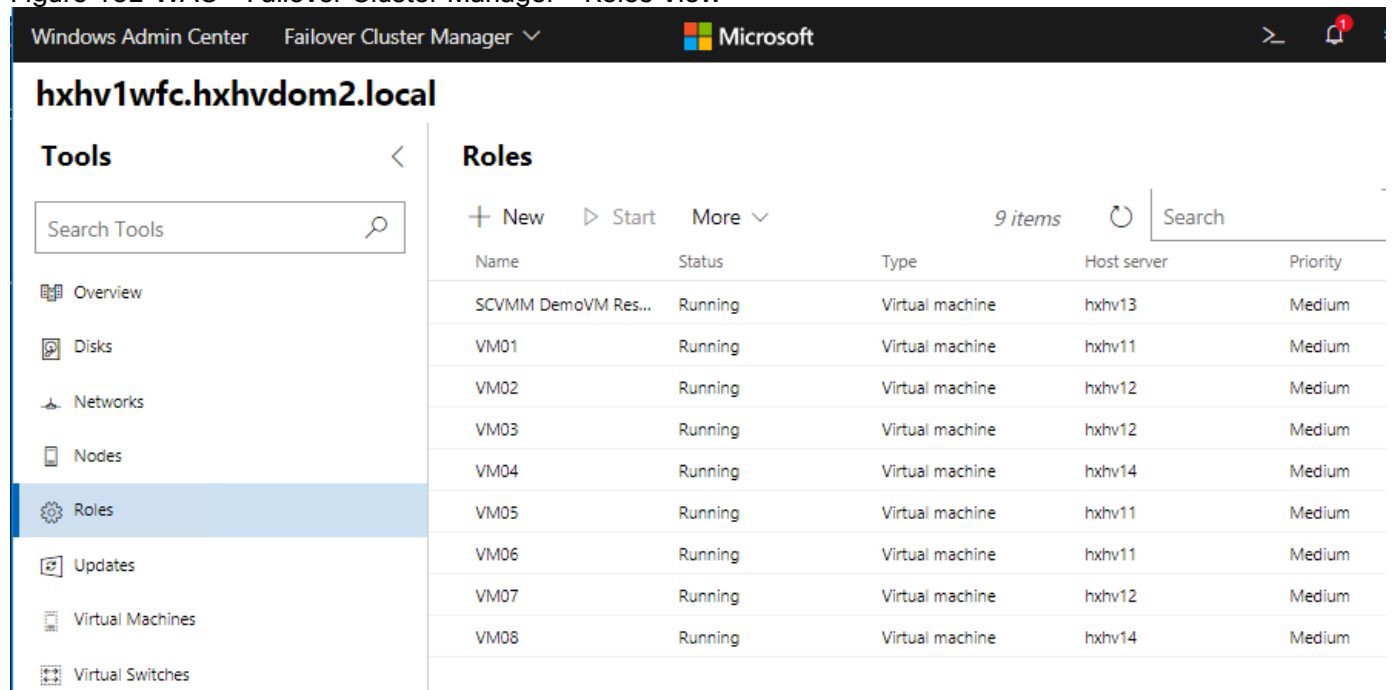


Figure 152 WAC – Failover Cluster Manager – Roles View



- Click Virtual Machines under the Tools pane and create a new Highly Available Virtual Machine as shown below:

The screenshot displays the Windows Admin Center interface for configuring a new virtual machine. The main window title is 'hxhv1wfc.hxhvd0m2.local'. On the left, a 'Tools' sidebar lists various management options, with 'Virtual Machines' selected. The central pane shows the 'Virtual machines' section with a 'New' button and a table that currently contains no records. The right-hand pane is the 'New virtual machine' configuration form, which is pre-filled with the following details:

- Name:** TestVM1
- Generation:** Generation 2 (Recommended)
- Host:** hxhv12.hxhvd0m2.local (Recommended)
- Path:** \\hxhv1smb.hxhvd0m2.local\hxd1\H: (with a 'Browse' button)
- Virtual processors:** Count is set to 2.
- Memory:** Set to 4 GB. There are checkboxes for 'Enable nested virtualization' and 'Use dynamic memory' (with sub-fields for Minimum and Maximum RAM).
- Network:** Network adapter is set to vswitch-hx-vm-network.
- Storage:** This section is currently empty.

At the bottom right of the configuration pane, there are 'Create' and 'Cancel' buttons. The browser address bar at the bottom left shows 'https://wac.hxhvd0m2.local'.



Refer to the Appendix section [C: Live Migration of Virtual Machines Between a Standalone Hyper-V Host and HyperFlex Hyper-V Host](#) to configure the Live Migration of virtual machines between a Standalone Hyper-V host and HyperFlex Hyper-V hosts.

Appendix

A: Cluster Capacity Calculations

HyperFlex HX Data Platform cluster capacity is calculated as follows:

$$\left(\left(\frac{\langle \text{capacity disk size in GB} \rangle \times 10^9}{1024^3} \right) \times \langle \text{number of capacity disks per node} \rangle \times \langle \text{number of HyperFlex nodes} \rangle \times 0.92 \right) / \text{replication factor}$$

Divide the result by 1024 to get a value in TiB

The replication factor value is 3 if the HX cluster is set to RF=3, and the value is 2 if the HX cluster is set to RF=2.

The 0.92 multiplier accounts for an 8% reservation set aside on each disk by the HX Data Platform software for various internal filesystem functions.

Calculation example:

$\langle \text{capacity disk size in GB} \rangle = 960$

$\langle \text{number of capacity disks per node} \rangle = 8$ for an HXAF220c-M5SX model server

$\langle \text{number of HyperFlex nodes} \rangle = 8$

replication factor = 3

Result: $\left(\left(\frac{960 \times 10^9}{1024^3} \right) \times 8 \times 8 \times 0.92 \right) / 3 = 17547.6074$

$17547.6074 / 1024 = 17.14$ TiB

B: Install Microsoft Windows Server 2016/2019

The Windows ISO and HyperFlex Driver image files must be placed on a shared location (such as HX Installer) that is reachable from Cisco UCS Manager and the Out-of-band subnet. Use the following steps to download and host the HyperFlex Driver Image and Windows ISO in a shared location within the installer VM.

To install Windows Server 2016/2019 and apply Cisco HyperFlex driver image on all HX nodes, follow these high-level steps.

- [Configure Cisco UCS Manager using HX Installer](#)
- [Configure Cisco UCS vMedia and Boot Policies](#)
- [Install Microsoft Windows Server 2019 OS](#)
- [Undo vMedia and Boot Policy Changes](#)

Configure Cisco UCS Manager using HX Installer

To complete this step, refer to section [HyperFlex Installation](#).

Configure Cisco UCS vMedia and Boot Policies

By using a Cisco UCS vMedia policy, the Windows Server 2016 media ISO file and [Cisco HyperFlex Driver image](#) can be mounted to all of the HX servers automatically. The existing vMedia policy, named “HyperFlex” must be modified to mount this file, and the boot policy must be modified temporarily to boot from the remotely mounted vMedia file. Once these two tasks are completed, the servers can be rebooted, and they will automatically boot from the When mounted vMedia file, installing and configuring Windows Server 2016 on the HX nodes.

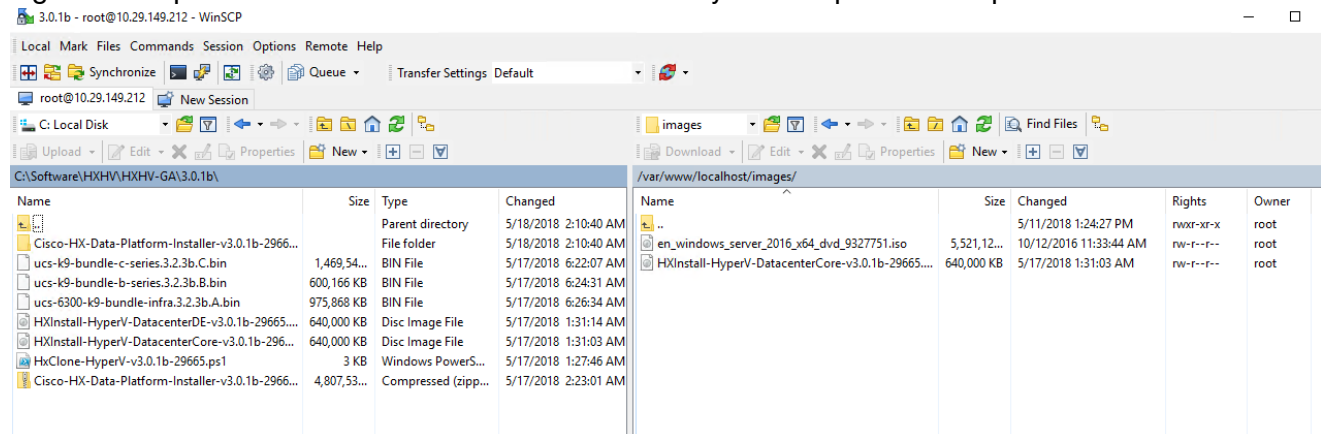


WARNING! While vMedia policies are very efficient for installing multiple servers, using vMedia policies as described could lead to an accidental reinstall of Windows on any existing server that is rebooted with this policy. Please be certain that the servers being rebooted while the policy is in effect are the servers you wish to reinstall. Even though the custom ISO will not continue without a secondary confirmation, extreme caution is recommended. This procedure needs to be carefully monitored and the boot policy should be changed back to original settings immediately after the intended servers are rebooted, and the Windows installation begins. Using this policy is only recommended for new installs or rebuilds. Alternatively, you can manually select the boot device using the KVM console during boot, and pressing F6, instead of making the vMedia device the default boot selection.

To configure the Cisco UCS vMedia and Boot Policies, follow these steps:

1. Connect to your HX Installer VM and browse to the folder that contains the `/var/www/localhost/images/`.
2. Copy the HyperFlex Driver Image “latest.img” (`/opt/springpath/packages/latest.img`) to the images folder (`/var/www/localhost/images/`).
3. Copy *Windows Server 2016 iso* image to the HXDP installer’s images folder (`/var/www/localhost/images/`)

Figure 153 Upload Windows ISO and Cisco Driver and System Preparation Script



For Windows Server 2019, follow these additional steps to prepare a suitable Driver image for automated OS install.

1. Copy the HyperFlex Driver Image. For example, run the following command:

```
rsync -avzP /opt/springpath/packages/latest.img
/var/www/localhost/images/install.img
```

```

root@HyperFlex-Installer:~# rsync -avzP /opt/springpath/packages/latest.img /var/www/localhost/images/install.img
sending incremental file list
latest.img
 655,360,000 100% 26.52MB/s 0:00:23 (xfr#1, to-chk=0/1)

sent 583,007,126 bytes received 35 bytes 23,796,210.65 bytes/sec
total size is 655,360,000 speedup is 1.12
root@HyperFlex-Installer:~# _

```

2. Mount the HyperFlex Driver Image. For example, run the following command:

```

mkdir -p /mnt/install-img && mount -o loop,rw
/var/www/localhost/images/install.img /mnt/install-img

```

3. Copy the answer file specific to Windows Server 2019. For example, run the following command:

```

cp
/opt/springpath/packages/FactoryUnattendXML/WindowsServer2019/Autounattend.xml.ro
/mnt/install-img/Autounattend.xml

```

4. Unmount the HyperFlex Driver Image. For example, run the following command:

```


umount /mnt/install-img

```

```

root@HyperFlex-Installer:~# cp /opt/springpath/packages/FactoryUnattendXML/WindowsServer2019/Autounattend.xml.ro /mnt/install-img/Autounattend.xml
root@HyperFlex-Installer:~# umount /mnt/install-img/
root@HyperFlex-Installer:~#

```

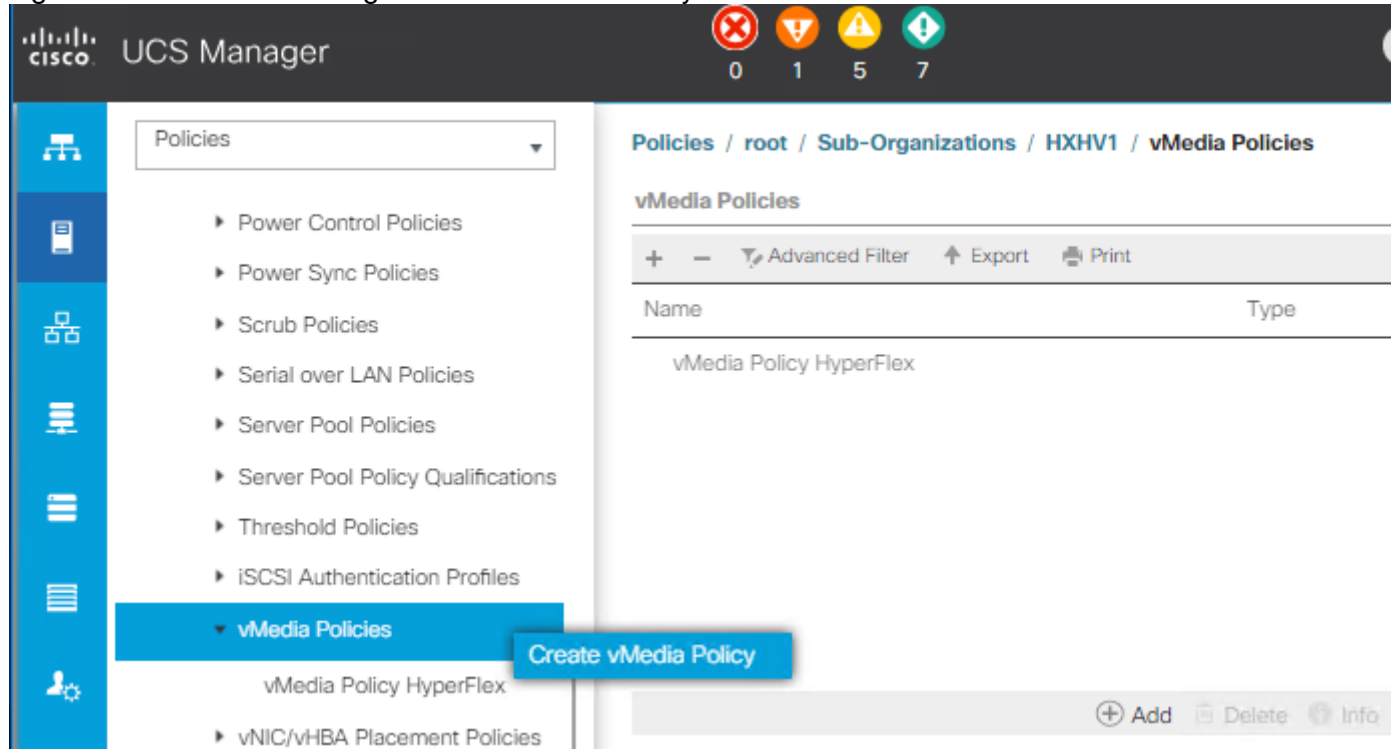


Make sure network connectivity exists between the file share and all server management IPs.

5. Configure the vMedia and Boot policies using Cisco UCS Manager to mount the above images.
6. Launch Cisco UCS Manager by accessing the Cisco UCS Manager IP address in a browser of your choice.
7. Click Launch UCS Manager and log in with administrator username and the password you used at the beginning of the installation.

8. In the left navigation pane, click Servers.
9. Expand Servers > Policies > root > Sub-Organizations > hx-cluster_name>vMedia Policies to view the list of vMedia Policies.

Figure 154 Cisco UCS Manager – Create vMedia Policy



10. Double-click vMedia Policy HyperFlex.
11. In the properties for vMedia Policy HyperFlex, click Create vMedia Mount to add the mount points.
12. In the Create vMedia Mount dialog box, complete the following fields in Table 46

Table 46 Create vMedia Mount Details

Field Name	Action	Example Value
Name	Name for the mount point.	Windows -ISO
Description	Can be used for more information.	
Device Type	Type of image that you want to mount	CDD
Protocol	The protocol used for accessing the share where the ISO files are located.	HTTP
Hostname/IP Address	IP address or FQDN of the server hosting the images.	10.104.252.48

Field Name	Action	Example Value
Image Name Variable	This value is not used in HyperFlex installation.	None
Remote File	The filename of the ISO file that you want to mount.	
Remote Path	The path on the remote server to where the file resides	
Username	If you use CIFS or NFS a username might be necessary	
Password	If you use CIFS or NFS a password might be necessary	

Figure 155 Cisco UCS Manager - Create vMedia Mount CDD

Create vMedia Mount [?] [X]

Name : W2019-ISO

Description :

Device Type : CDD HDD

Protocol : NFS CIFS HTTP HTTPS

Hostname/IP Address : 10.104.252.48

Image Name Variable : None Service Profile Name

Remote File : en_windows_server_2019_x64_dvd_4cb967d8.iso

Remote Path : /images/

Username :

Password :

Remap on Eject :

OK Cancel

13. Click Save Changes and click OK.
14. Click OK. When you click OK, you are returned to the vMedia policy and will see the information that you submitted.
15. Repeat steps 5 and 6 but change the type to HDD and the filename to the Cisco HyperFlex driver image.

Figure 156 Cisco UCS Manager - Create vMedia Mount HDD

Create vMedia Mount ? X

Name :

Description :

Device Type : CDD HDD

Protocol : NFS CIFS HTTP HTTPS

Hostname/IP Address :

Image Name Variable : None Service Profile Name

Remote File :

Remote Path :

Username :

Password :

Remap on Eject :

OK
Cancel

16. On completion, the following screen displays:

Figure 157 Cisco UCS Manager - Create vMedia Policy

Modify vMedia Policy ?

vMedia Policy:

Create vMedia Policy

Name : **HX-401b**

Description :

Retry on Mount Failure: **Yes**

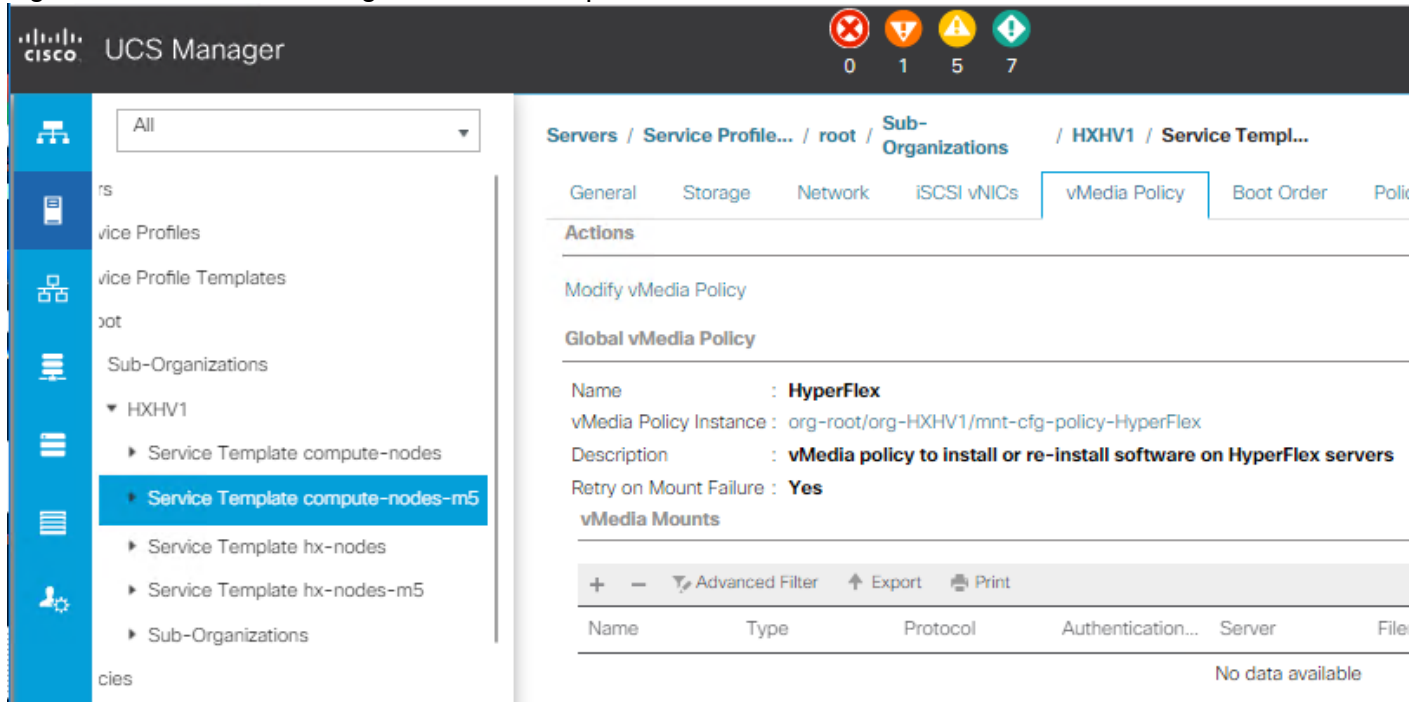
vMedia Mounts

+ - ▼ Advanced Filter ↑ Export 🖨 Print ⚙

Name	Type	Protocol	Authen...	Server	Filename	Re... ▲	User	Remap...
HX-Driver	HDD	HTTP	None	10.104.252.48	latest.img	/image...		No
W2019-ISO	CDD	Unkno...	None	10.104.252.48	en_windows_server_...	/image...		No

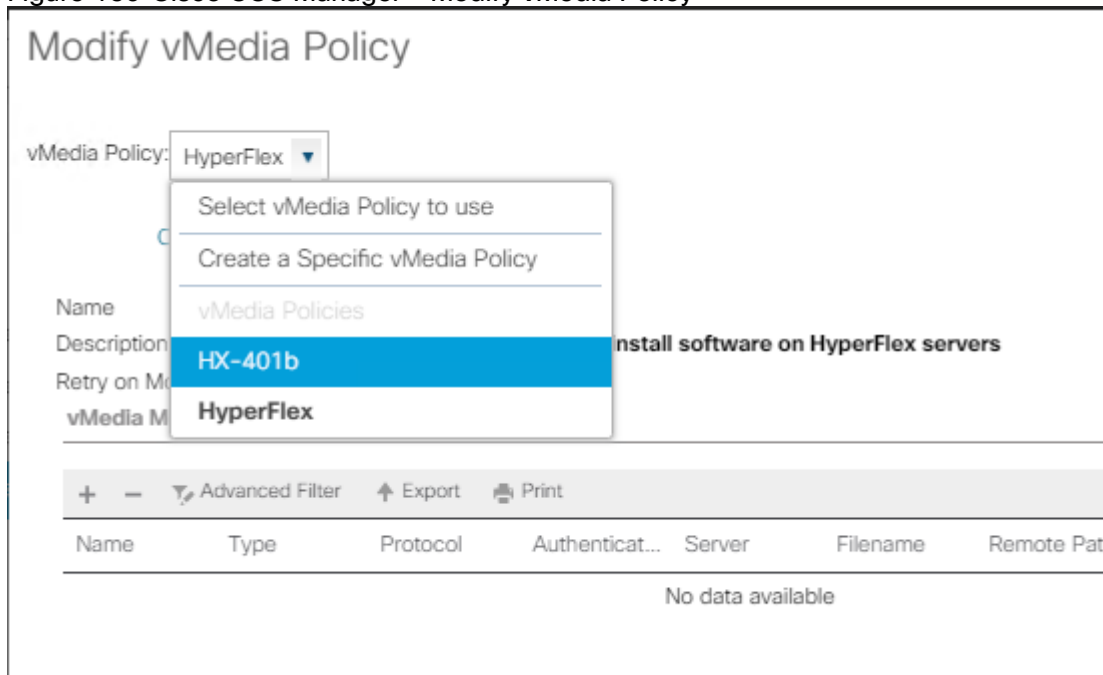
17. In the left navigation pane, select Servers > Service Profile Templates > root > Sub-Organizations > hx-cluster_name > Service Template hx-nodes_name (example: compute-nodes-m5).

Figure 158 Cisco UCS Manager – Service Template



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

Figure 159 Cisco UCS Manager – Modify vMedia Policy



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

19. Select Servers > Policies > root > Sub-Organizations > hx-cluster_name > Boot Policies Boot Policy HyperFlex-m5.
20. In the configuration pane, click CIMC Mounted vMedia. Click Add CIMC Mounted CD/DVD to add this to the boot order.
21. Select the CIMC Mounted CD/DVD entry in the list and move it to the top of the boot order by clicking Move Up.

Figure 160 Cisco UCS Manager - Boot Order

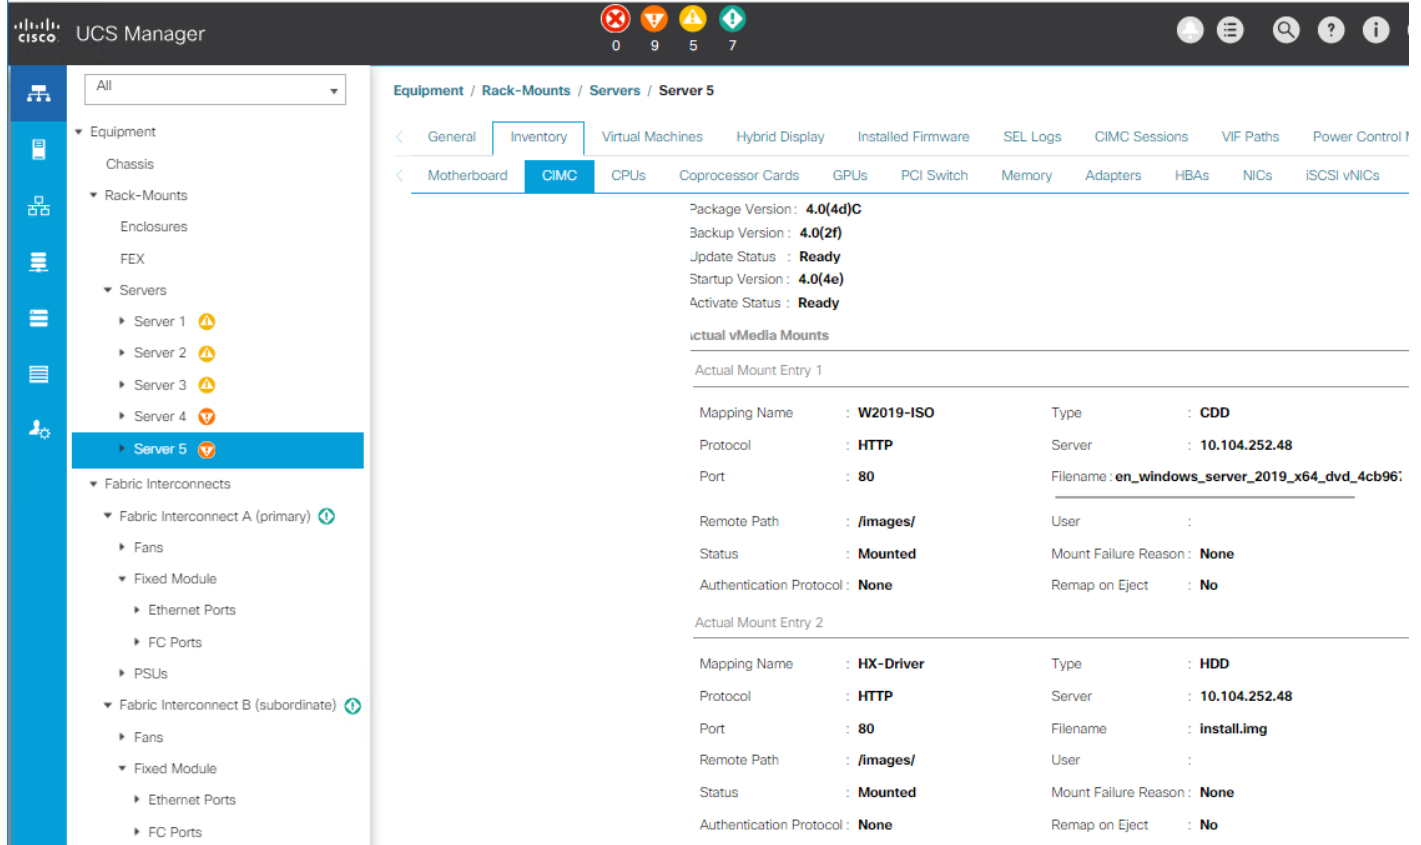
Name	Ord...	vNIC/v...	Type	LUN N...	WWN	Slot N...	Boot N...	Boot P...	Descri...
CD/DVD	1								
Embedded Disk	2								
CIMC Mounted CD/...	3								

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

To verify the images are mounted correctly, follow these steps:

1. On the Equipment tab, select one of the servers.
2. Click Inventory > CIMC, scroll down and make sure for the mount entry #1(OS image) and mount entry #2 (Cisco HyperFlex driver image) the status is Mounted and there are no failures.

Figure 161 Cisco UCS Manager – Validate vMedia Mount



Install Microsoft Windows Server 2019 OS

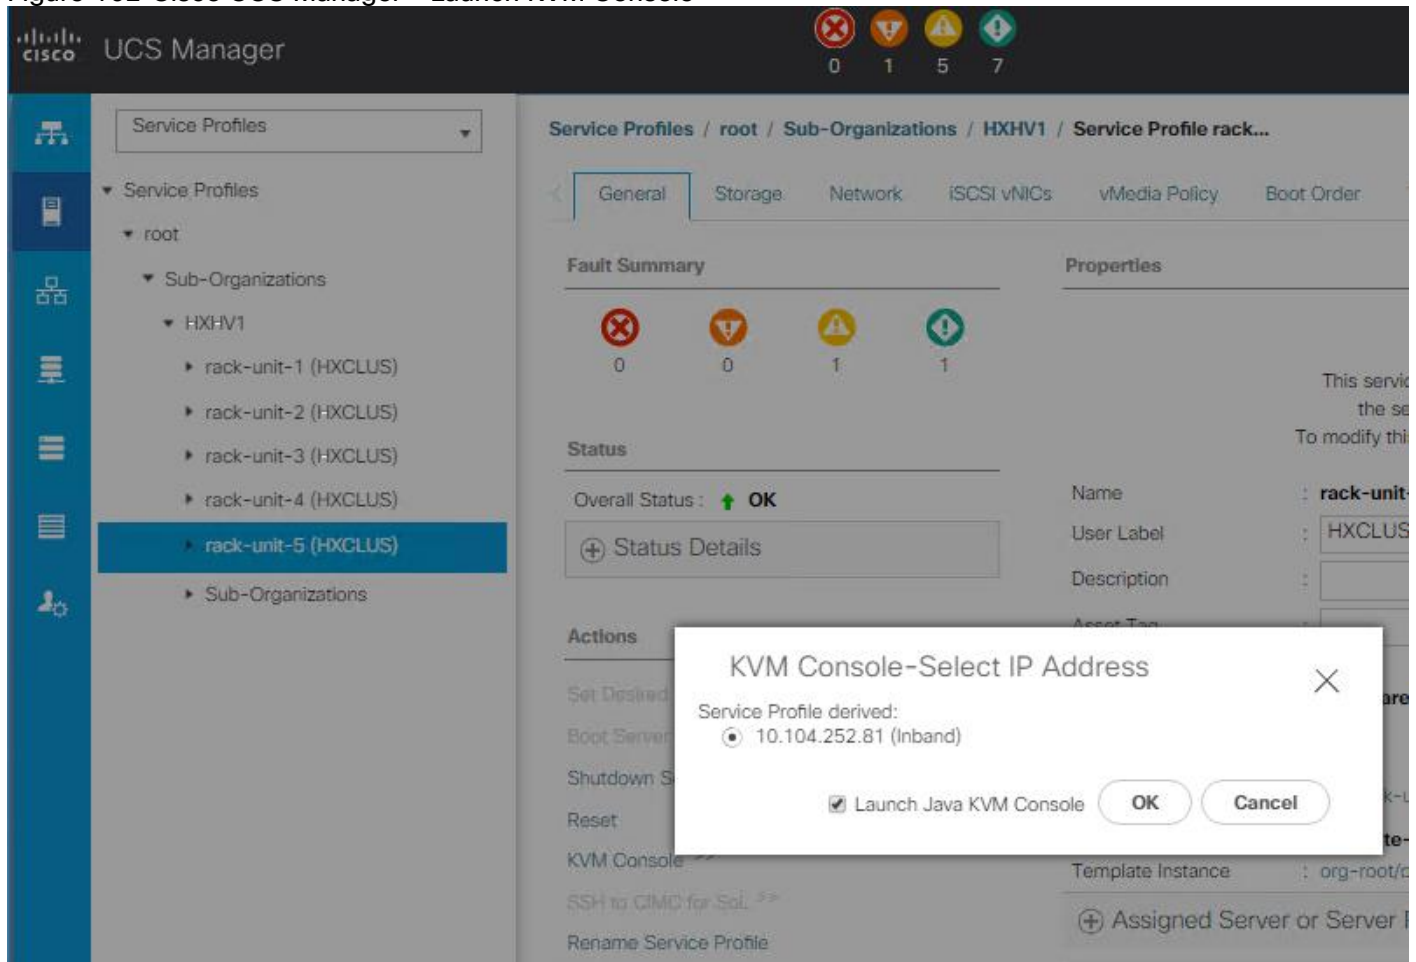
To install Microsoft Windows Server 2016 OS, follow these steps:

1. In the menu bar, click Servers and choose the first HyperFlex service profile.
2. Click the General tab and choose Actions > KVM Console.




The KVM console will try to open in a new browser. Be aware of any pop-up blockers. Allow the pop-ups and re-open the KVM.

Figure 162 Cisco UCS Manager – Launch KVM Console



3. Reboot the server. In the KVM console choose Server Actions and click Reset.

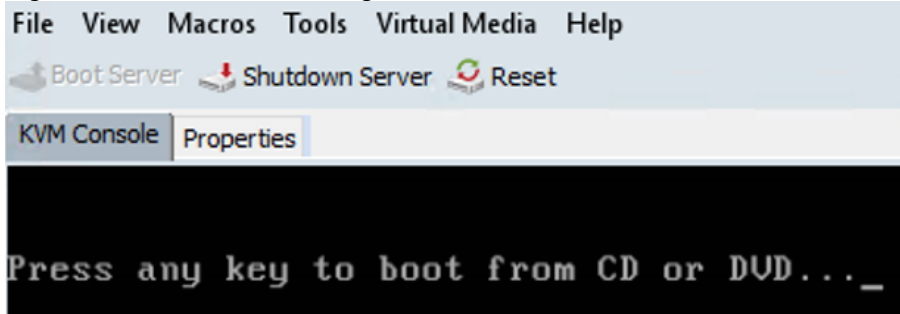
Figure 163 Cisco UCS Manager – Server KVM Console




4. Choose Power Cycle.

5. When the server is rebooting, remember to press any key to start the Windows installation.


Figure 164 Cisco UCS Manager - KVM Console Server Boot



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


6. When the Windows installation is complete, you will see some tasks running as shown in the below and the host will reboot a few times. Allow some time for the system preparation to complete.

Figure 165 System Preparation



7. The installation is complete when a clean command prompt with no activity is displayed as shown below.

Figure 166 Windows Server Command Prompt



8. Repeat steps 1-7 on all the HX nodes in the cluster and verify the below task is running as shown in below. The 'HXInstallbootstraplauncherTask' in running state is an indication of successful installation Windows OS and system preparation.

Figure 167 Validate Windows Server Installation Completion

```

File View Macros Tools Virtual Media Help
Boot Server Shutdown Server Reset
KVM Console Properties
Administrator: C:\Windows\system32\cmd.exe - powershell
PS C:\Users\Administrator> Get-ScheduledTask -TaskName hx*

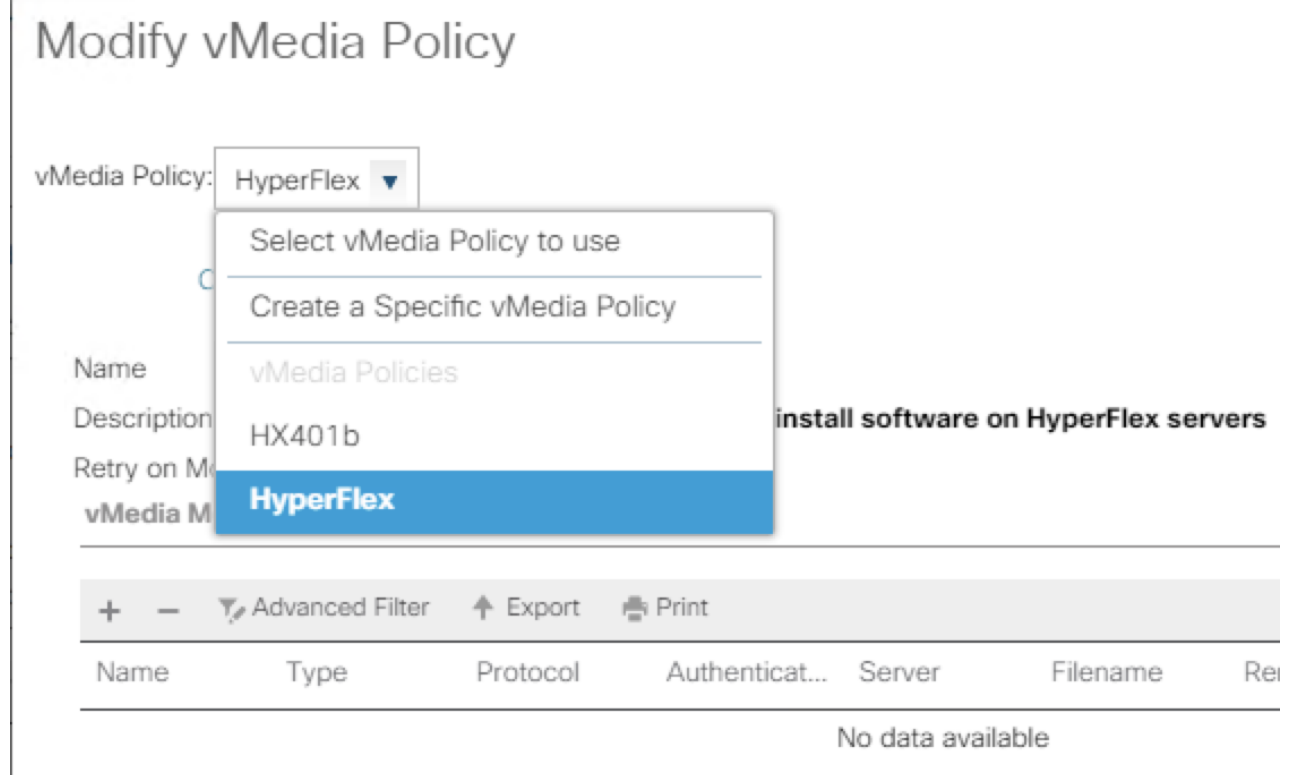
TaskPath TaskName State
-----
\ HXInstallbootstraplauncherTask Running
  
```

Undo vMedia and Boot Policy Changes

To undo vMedia and boot policy changes, follow these steps:

1. When all the servers have booted from the remote vMedia file and begun their installation process, the changes to the boot policy need to be quickly undone, to prevent the servers from going into a boot loop, constantly booting from the installation ISO file. To revert the vMedia policy settings, follow these steps: In Cisco UCS Manager select Servers > Service Profile Templates > root > Sub-Organizations > hx-cluster_name > Service Template hx-nodes-m5. Then, click on Modify vMedia Policy.
2. Under the vMedia Policy drop-down selection, choose "HyperFlex" policy as shown in the figure below.

Figure 168 Cisco UCS Manager – Undo vMedia and Boot Policy Changes



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

Undo vMedia and Boot Policy Changes (for Compute-only Nodes)

To undo vMedia and boot policy changes, follow these steps:

1. Reset the vMedia policy back to the default HyperFlex policy:
 - a. Update the vMedia policy for compute nodes. Go to Servers > Service Profile Templates > root > Sub-Organizations > hx-cluster_name > Service Template compute-nodes, or compute-nodes-m5. Click Modify vMedia Policy.
 - b. Under the vMedia Policy drop-down selection, choose "HyperFlex" policy.
2. Restore the boot order to the one before installation:
 - a. In the Navigation pane, click the Servers tab.
 - b. Expand Servers > Policies > root > > Boot Policies > hx-compute, or hx-compute-m5.
 - c. In the Boot Order configuration pane, use the Move Down button to move CIMC Mounted CD/DVD option to the bottom of the list.

C: Live Migration of Virtual Machines Between a Standalone Hyper-V Host and HyperFlex Hyper-V Host

This section covers the steps to configure and perform live migration of virtual machines between two non-clustered Hyper-V hosts. This is sometimes called as 'shared-nothing' live migration. Basically, this Hyper-V feature allows migration of workloads running on existing standalone Hyper-V hosts or Hyper-V cluster to the newly deployed HyperFlex Cluster

Prerequisites

Before you perform the steps in this document, make sure that your environment meets the following prerequisites:

- The source and destination computers either belong to the same Active Directory domain or belong to domains that trust each other.
- The user account has the appropriate permission to perform the various steps:
 - Configure HX storage access to the standalone Hyper-V host
 - To configure and perform the live migrations, the account must be a member of the local Hyper-V Administrators group or the Administrators group on both the source and destination computers.
- A computer running Windows with the Hyper-V management tools installed. For instructions, see <http://technet.microsoft.com/library/hh857623.aspx>

To manage the live migration tasks with remote management tools, configure constrained delegation and select Kerberos as the authentication protocol. Otherwise, you must sign on to the source computer to perform a live migration, and CredSSP is used to authenticate the live migration.

- hxxvinfra2. hxxhvd0m2.local – is a standalone hyper-v host in AD domain “hxxhvd0m2.local”
- hxxhv11. hxxhvd0m2.local – is a Hyper-V host in a HyperFlex cluster belonging to the same above AD domain “hxxhvd0m2.local”

Configure HX Storage Access to the Standalone Hyper-V Host

To configure HX storage access to the standalone Hyper-V host, follow these steps:

1. Note down the IP address of the standalone Hyper-V host
2. Log into one of the StCtIVM using SSH and execute the below command to open the port.

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

3. Enter the IP address of the standalone Hyper-V host at the “Enter Ip address of SCVMM” prompt as shown below. (This step is to allow access to the SMB share on HX Datastore)

Figure 169 Storage Access script to open port

```

root@hxhv1scvm:~# python /opt/springpath/storfs-hyperv/FixScvmmAccess.py
Enter Ip address of SCVMM: 10.104.252.52
PING 10.104.252.52 (10.104.252.52) 56(84) bytes of data.
64 bytes from 10.104.252.52: icmp_seq=1 ttl=128 time=0.577 ms

--- 10.104.252.52 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.577/0.577/0.577/0.000 ms
root@hxhv1scvm:~# █

```

4. On the standalone Hyper-V host, edit the host file and add an entry with IP address of HX SMB Access Point mapping to its FQDN. The following PowerShell command will add the above entry to the host file:

```

Add-Content -Path C:\Windows\System32\drivers\etc\hosts -Value
" `r`n10.104.252.135`thxhv1smb.hxhvd0m2.local" -Force

```

Figure 170 Edit Hosts file to add an entry

```

PS C:\Users\administrator.HXHVDOM2> Add-Content -Path C:\Windows\System32\drivers\etc\hosts -Value " `r`n10.104.252.135`thxhv1smb.hxhvd0m2.local" -Force

```



`r` and `n` used after the value parameter are PowerShell notation of carriage return and new line and the IP address and hostname are separated by `t` which is the PowerShell notation for a tab character.

Figure 171 Source computer - Hosts file entry Mapping

```

PS C:\Users\administrator.HXHVDOM2> Get-Content -Path C:\Windows\System32\drivers\etc\hosts
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com          # source server
#       38.25.63.10      x.acme.com              # x client host
#
# localhost name resolution is handled within DNS itself.
#       127.0.0.1        localhost
#       ::1              localhost
10.104.252.135 hxhv1smb.hxhvd0m2.local
PS C:\Users\administrator.HXHVDOM2> █

```

5. Verify now if the HX Cluster SMB share is accessible now using the below PowerShell command or browse the SMB Share

```

test-path \\hxhv1smb.hxhvd0m2.local\hxds1

```

Figure 172 Command to Verify Access to SMB Share

```

PS C:\Users\administrator.HXHVDOM2> test-path \\hxhv1smb.hxhvd0m2.local\hxds1
True

```

Figure 173 Command to Verify Access to SMB Share

```
PS C:\Users\administrator.HXHVDOM2> dir \\hxhv1smb.hxhvdome2.local\hxds1

Directory: \\hxhv1smb.hxhvdome2.local\hxds1

Mode                LastWriteTime         Length Name
----                -
d-----          10/17/2019   9:27 PM         Hyper-V
```

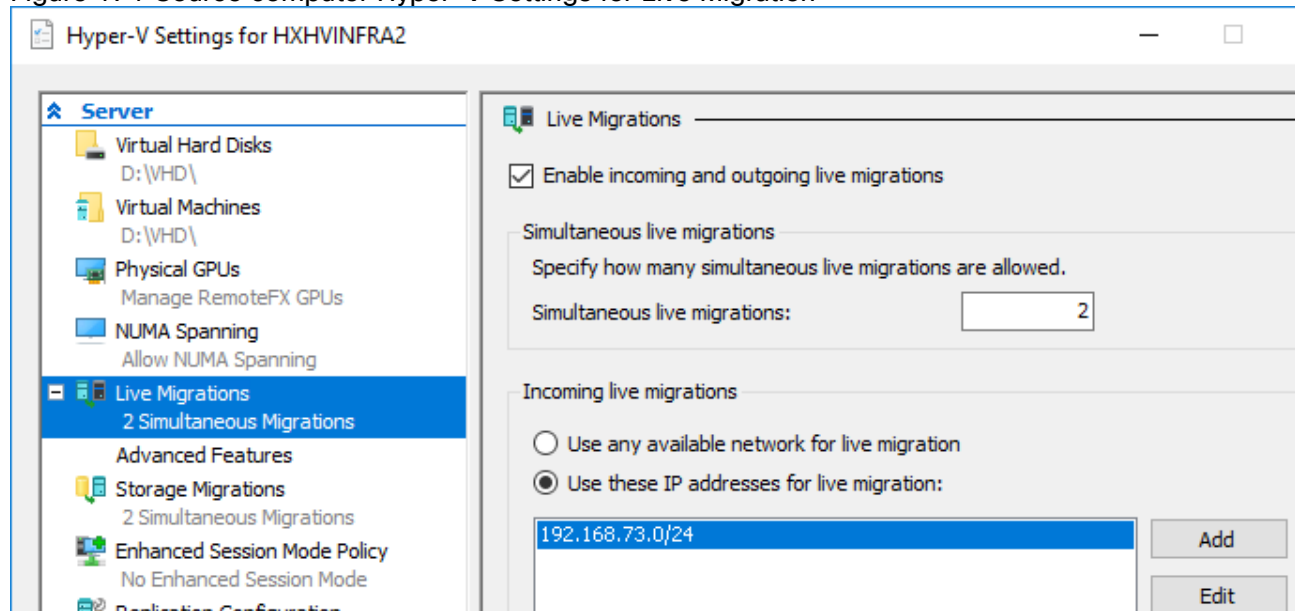
Configure the Source and Destination Computers for Live Migration

This section covers the steps to enable live migrations and configure the source and destination servers to send and receive live migrations. As a security best practice, it is recommended that you select specific networks to use for live migration traffic.

To configure the source computer (standalone Hyper-V hosts) for live migration, follow these steps:

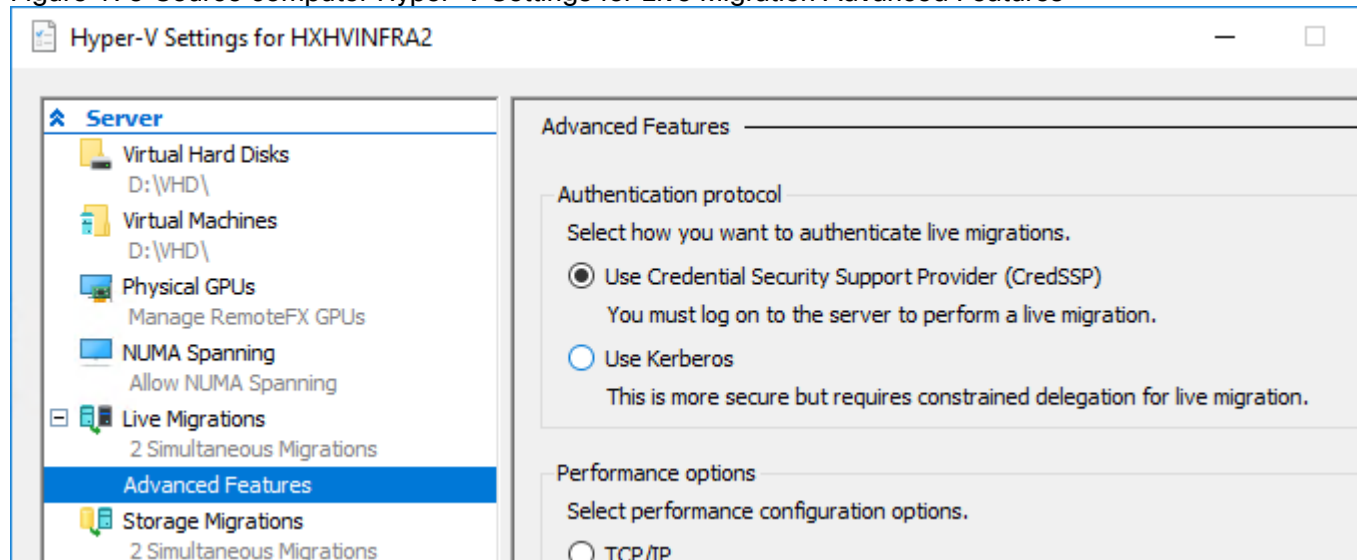
1. Open Hyper-V Manager
2. In the navigation pane, select the source servers in navigation pane
3. In the Action pane, click Hyper-V Settings.
4. In Hyper-V Settings dialog box, click Live Migrations.
5. In the Live Migrations pane, check Enable incoming and outgoing live migrations.
6. Under Simultaneous live migrations, specify a different number if you don't want to use the default of 2.
7. Under Incoming live migrations, if you want to use specific network connections to accept live migration traffic, click Add to type the IP address information. Otherwise, click Use any available network for live migration. Click OK.

Figure 174 Source computer Hyper-V Settings for Live Migration



8. If you have configured constrained delegation in the earlier section, expand Live Migrations and then select Advanced Features.
9. In the Advanced Features pane, under Authentication protocol, select CredSSP.
10. Click OK.

Figure 175 Source computer Hyper-V Settings for Live Migration Advanced Features



CredSSP authentication protocol requires user must logon to the server to perform a live migration. Kerberos requires constrained delegation for live migration.

Configure Performance Options for Live Migration – Optional

This section covers the steps to configure the live migration performance options. The default performance options can reduce overhead on the network and CPU usage in addition to the reducing the amount of time for a live migration. To configure the appropriate live migration performance options based on your environment and requirements, follow these steps:

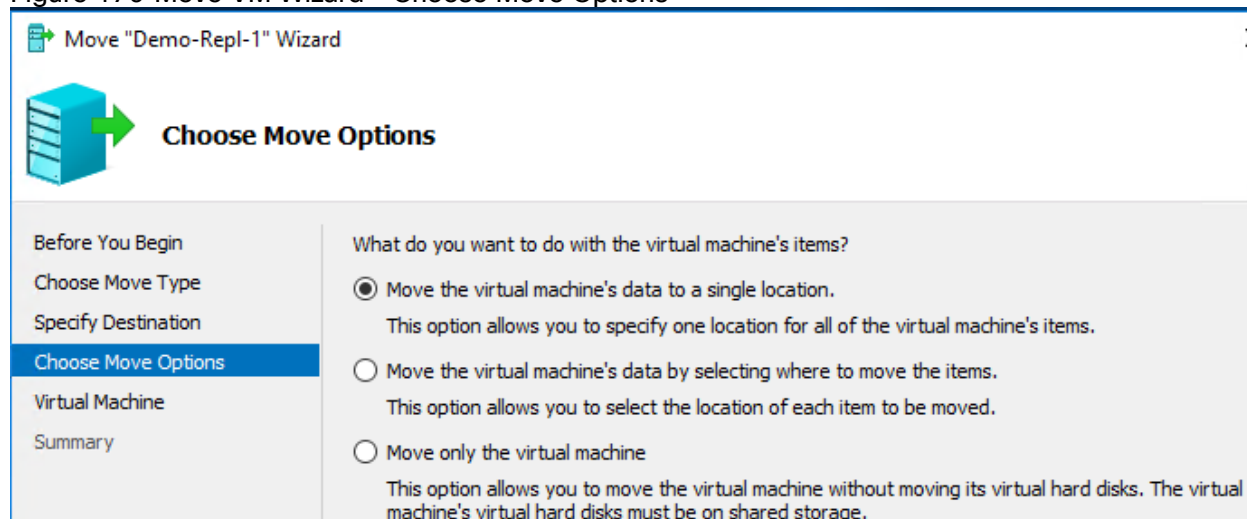
1. From Hyper-V Manager, select one of the servers you want to configure for live migrations.
2. In the Action pane, click Hyper-V Settings.
3. In the Hyper-V Settings dialog box, expand Live Migrations and then select Advanced Features.
4. In the Advanced Features pane, under Performance options, select the appropriate option for your environment, and then click OK.

Move a Running Virtual Machine

To move a running virtual machine, follow the below steps:

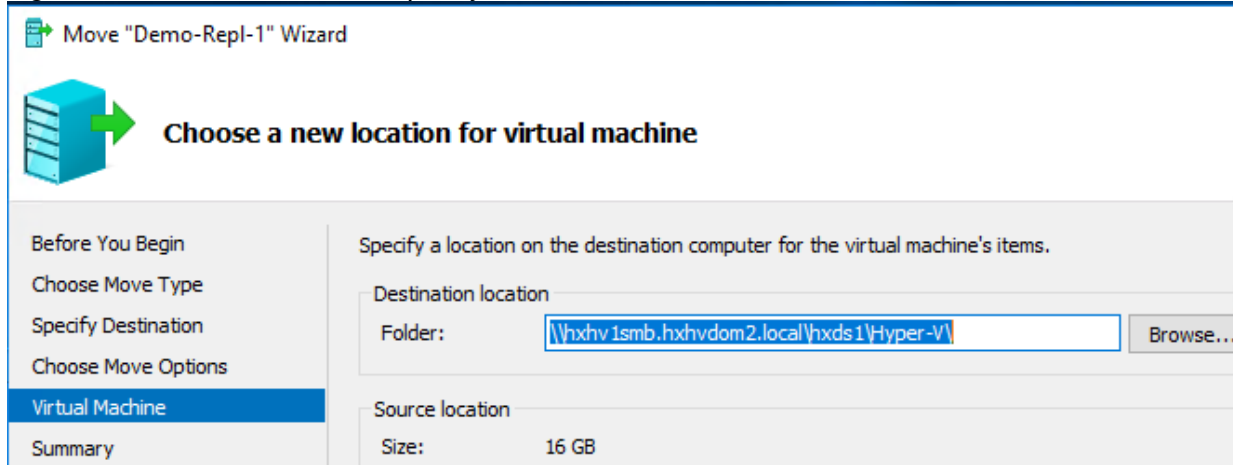
1. From the remote host with RSAT tools installed, open Hyper-V Manager.
2. From Hyper-V Manager, click the name of the source server.
3. From the Virtual Machines section of Hyper-V Manager, right-click the virtual machine and then click Move.
4. On the Choose Move Type page of the Move Wizard, choose Move the virtual machine.
5. On the Specify Destination page, type the name or browse to the destination computer.
6. On the Choose Move Options page, select 'Move the virtual machine's data to a single location' and click Next

Figure 176 Move VM Wizard – Choose Move Options



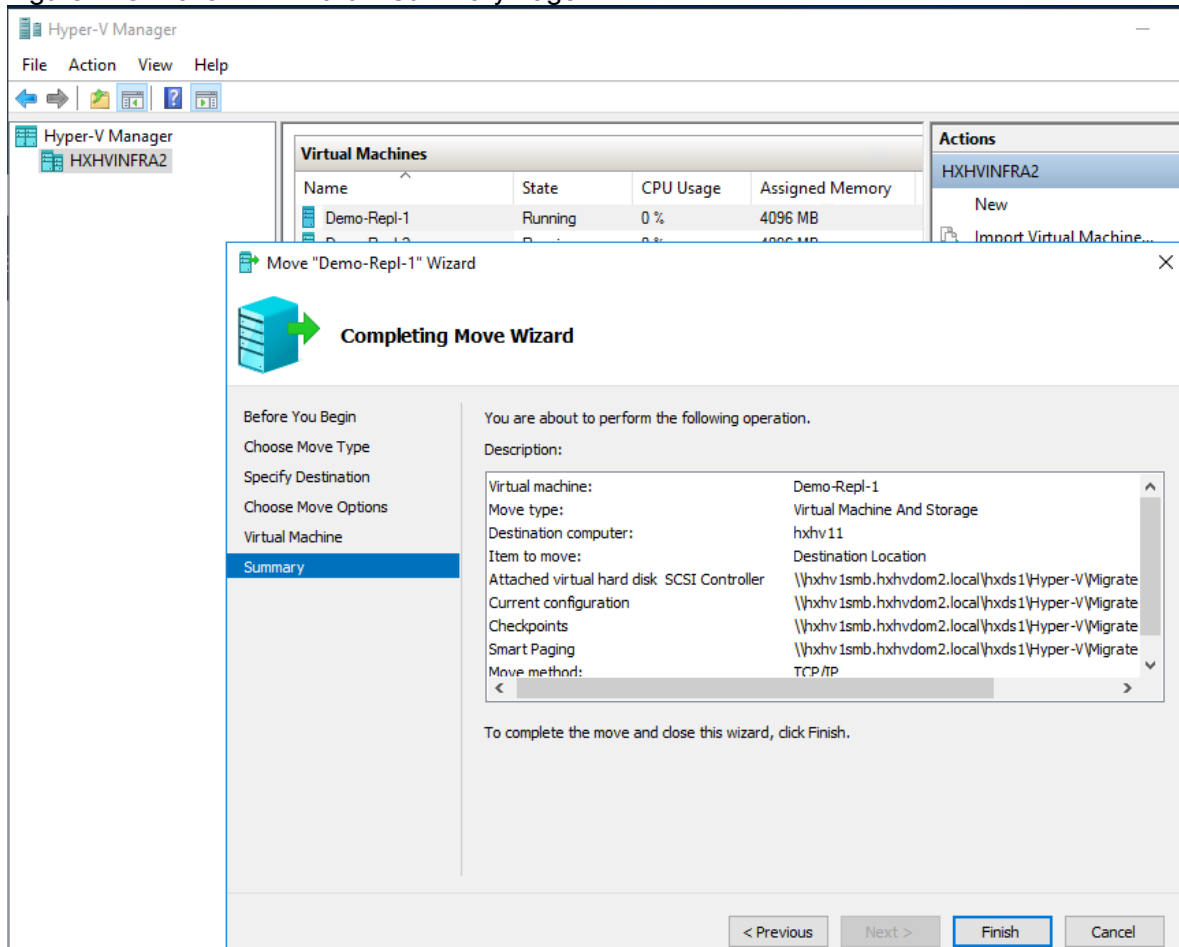
7. On the Choose a New Location for Virtual Machine, browse and select the destination folder location (Hyper-Flex SMB access path).

Figure 177 Move VM Wizard – Specify Destination Location

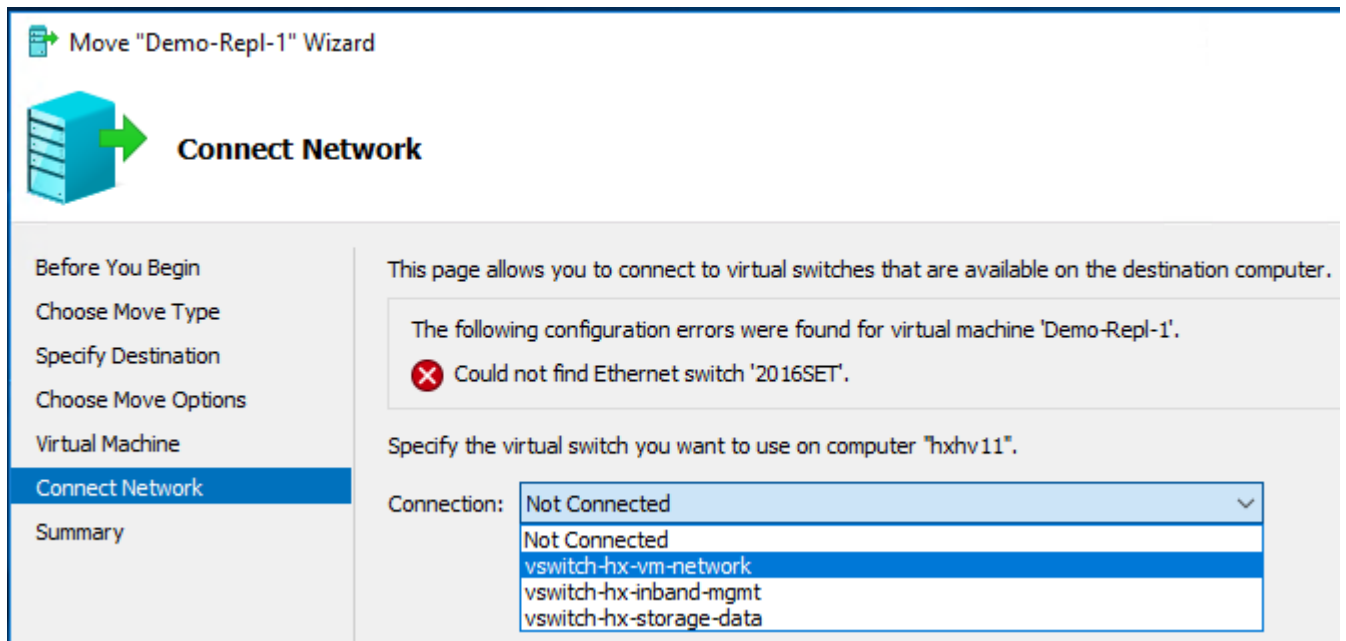


8. On the Summary page, review your choices and then click Finish.

Figure 178 Move VM Wizard – Summary Page

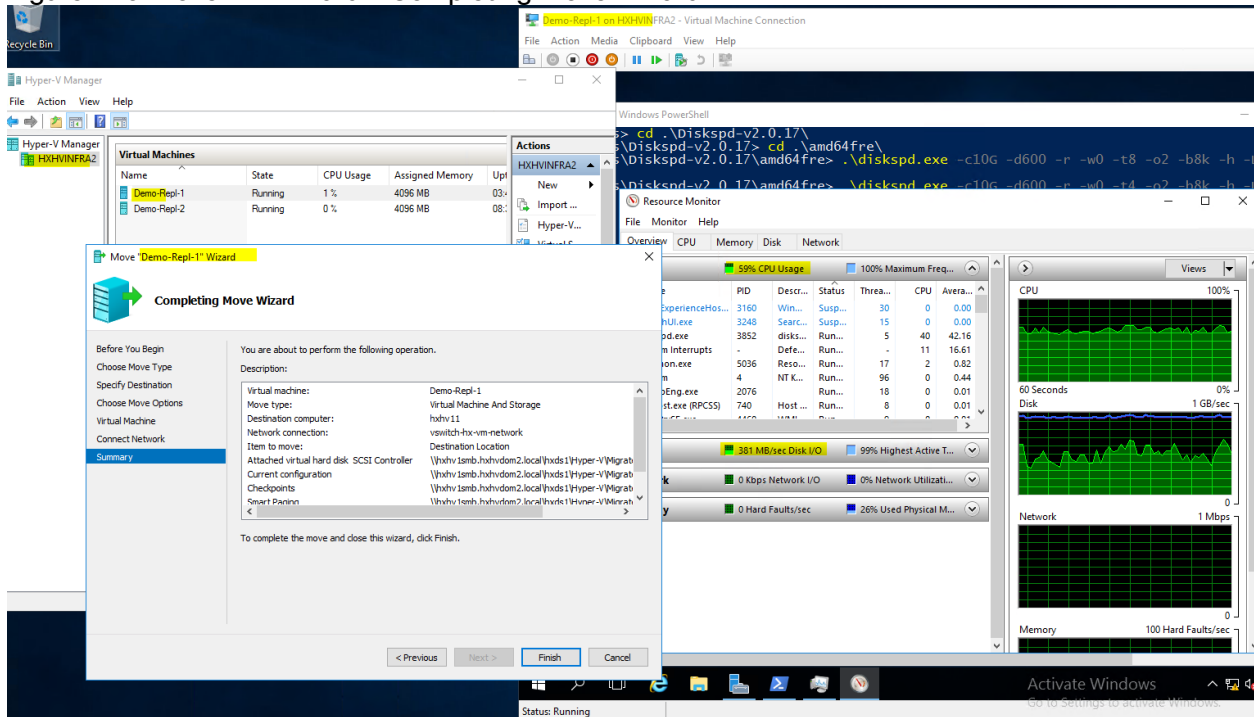


9. If you encounter an error during performing the move, It could be because the vSwitch to which the VM is attached on the source computer is not found on the destination host. Click the drop-down menu and select an appropriate vSwitch listed that is available on the destination host and click Finish.



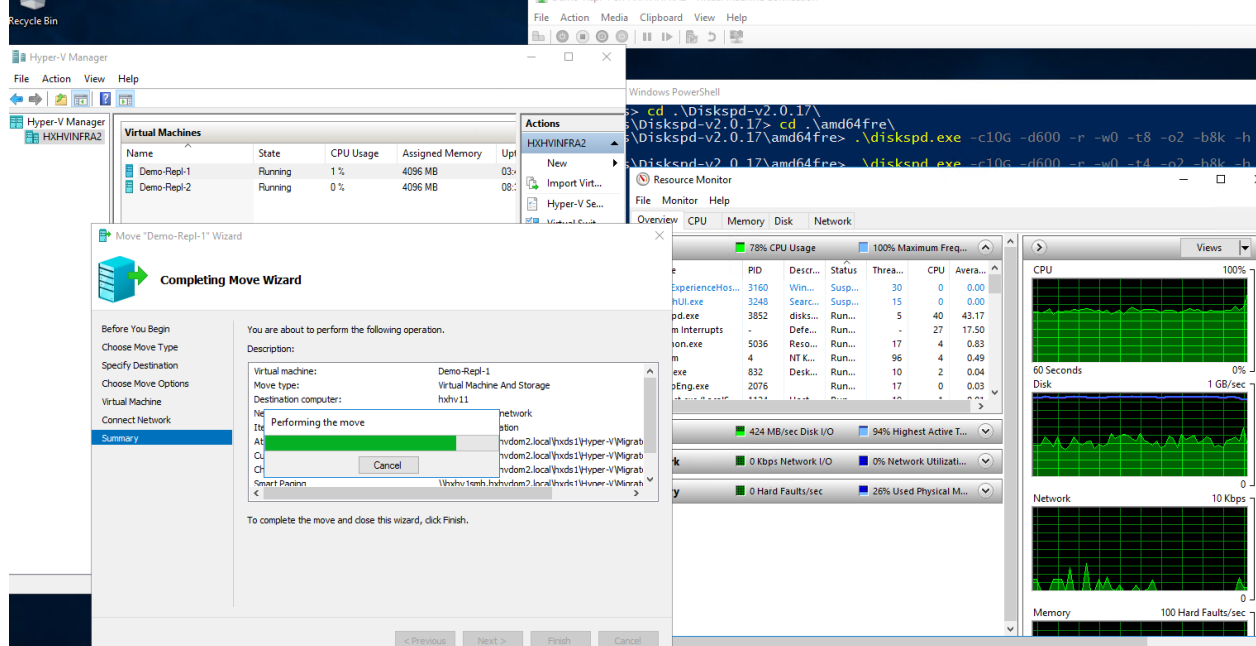
10. Click Finish.

Figure 179 Move VM Wizard – Completing Move Wizard




Live migration begins.

Figure 180 Move VM Wizard – Live Migration in Progress



11. After the VM migration is complete. The VM “Demo-Rep1-1” can now be seen on the destination HX Hyper-V host.

Figure 181 VM Live Migrated on Destination Host



D: Delegate HX Service Account with Least Privileges for Administrative Tasks

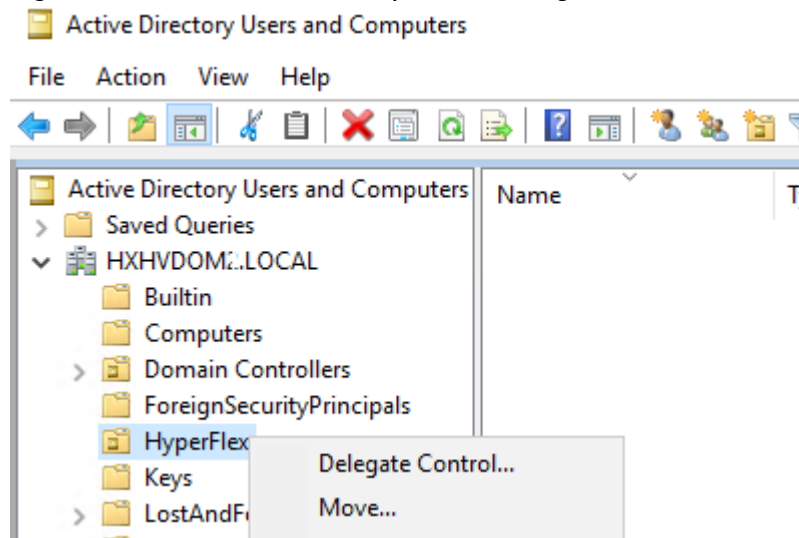
A service account in the AD domain with appropriate privileges is required to deploy and manage HyperFlex clusters successfully. You can either grant full domain admin privileges or grant absolute minimum permissions necessary to the service account for deploying and managing HyperFlex Hyper-V clusters. From a security perspective, it is highly recommended to grant service accounts with least privileges security for any administrative tasks.

In the HyperFlex Installation section, the service account was granted with full domain admin privileges for HyperFlex deployment. This section explains the steps to delegate the service account for HyperFlex with least security privileges necessary for administrative tasks.

To delegate HX Service Account with least privileges for administrative tasks, follow these steps:


1. Log into the Active Directory server as a domain administrator.
2. Open Active Directory Users & Computers and navigate to the Organizational Unit (OU) created in the above steps.
3. Right-click the OU and click Delegate Control.

Figure 182 AD Users and Computers – Delegate Control to OU



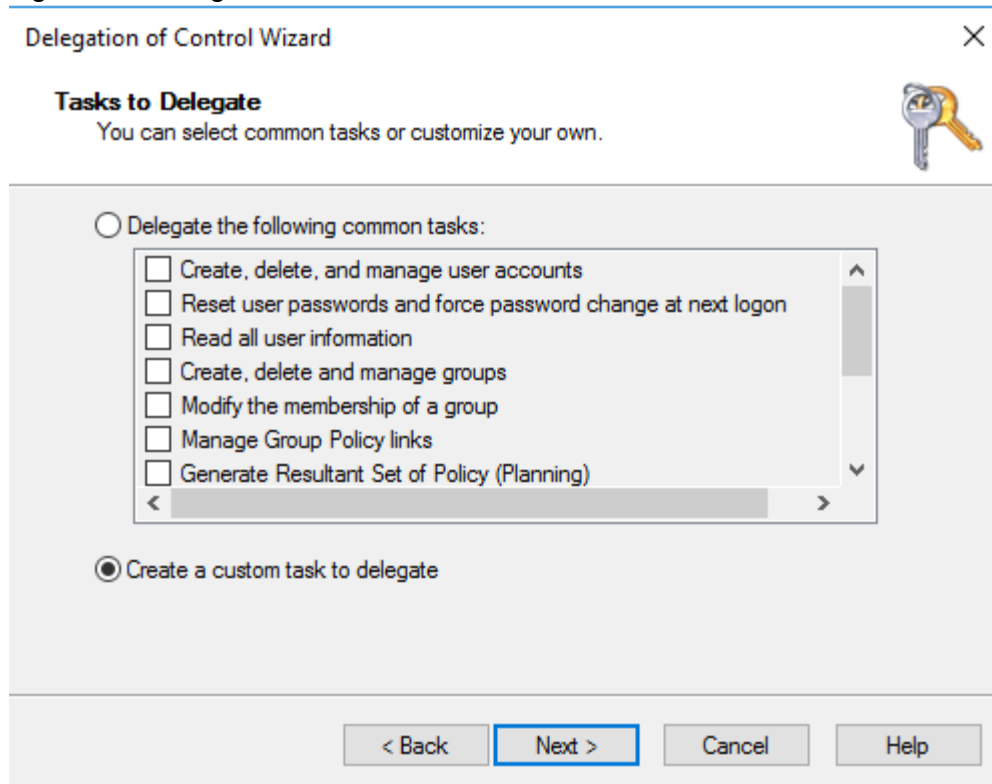
4. Click Next on the Delegation of Control Wizard
5. In Users or Groups window of Delegation of Control Wizard, click Add.
6. In Select Users, Computers, or Groups window, enter the HX service account (hxadmin) and click Next.

Figure 183 AD Users and Computers – Select Users, Computers, or Groups



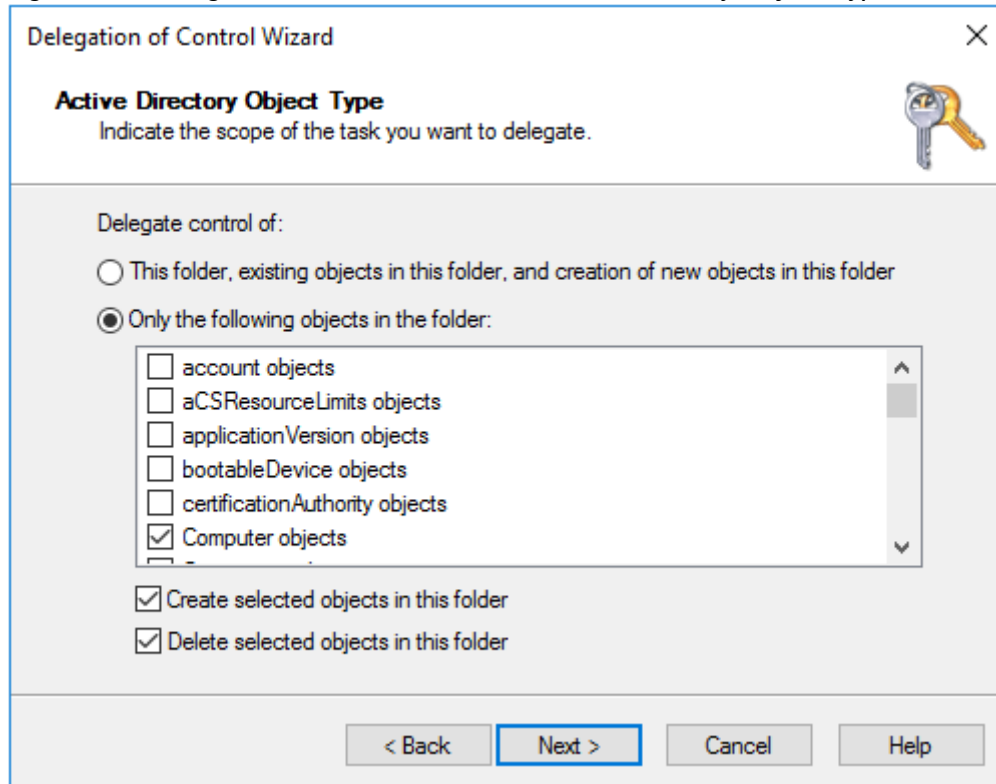
7. In Tasks to Delegate, select Create a custom task to delegate. Click Next.

Figure 184 Delegation of Control Wizard



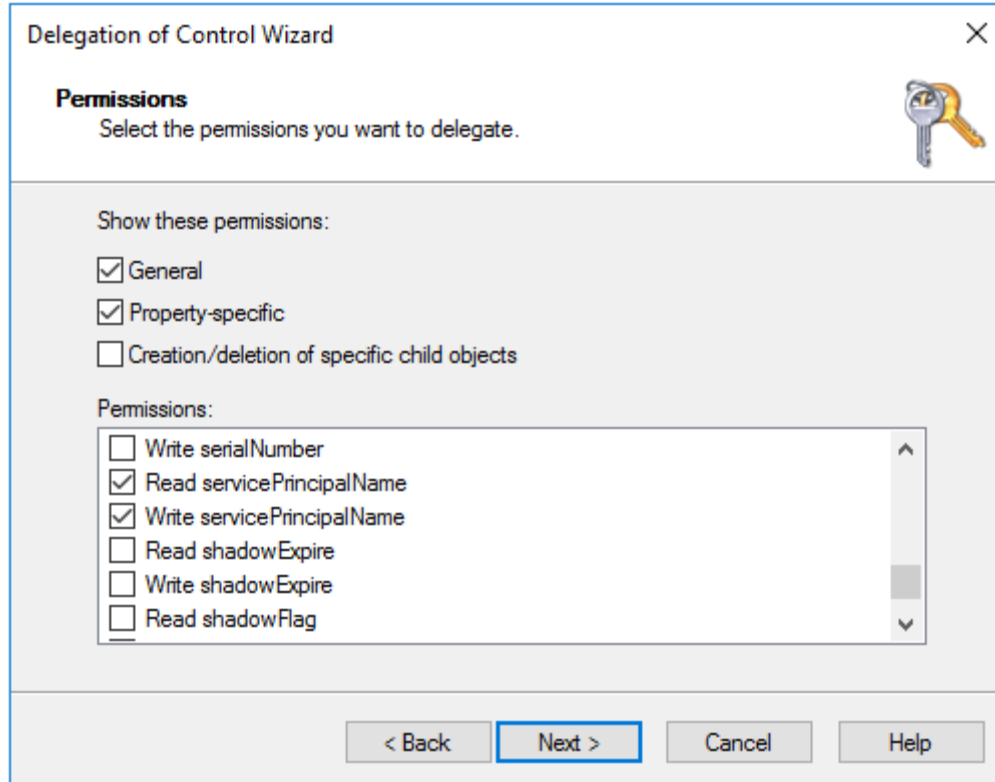
8. In Active Directory Object Type, complete the following and then click Next:
- Select Only the following objects in the folder and check Computer Objects.
 - Check the boxes:
 - Create selected objects in the folder.
 - Delete selected objects in this folder.

Figure 185 Delegation of Control Wizard – Active Directory Object Type



9. In Permissions, check the following and then click Next:
 - a. In Show these permissions, select General and Property-specific.
 - b. In Permissions, select the following permissions to ensure AD includes the appropriate permissions.
 - i. Read servicePrincipalName.
 - ii. Write servicePrincipalName.

Figure 186 Delegation of Control Wizard – Permissions



- iii. Click Finish.
- iv. AD replication might take a few minutes to complete user delegation on all domain controllers.

E: Common Resiliency and High Availability Scenario for HX Hyper-V

1. If a Node Failure or OS/Windows Crash - VMs are failed over to the surviving node.
2. If a Drive Failure occurs - This is invisible to the OS for HX Disks as the HX File system will auto repair degraded data.
3. If the OS boot disk fails - VMs are failed over to the surviving node.
4. If the Controller VM or Controller VM's disk fails - SMB traffic redirection through DFS and SMB client.
5. If there is a Network Failure - VMs will be isolated for up to 4 minutes on isolated node. After that VMs will be failed over and node quarantined.

Hyper-V Failover Clustering Resiliency

1. Virtual Machine Load Balancing (move VMs on memory pressure and cpu utilization).
2. High Availability increased with Transient Error reactions.
3. Set a Resiliency level and period.
4. Isolated node(s) don't immediately failover VMs (transient network error).

5. Unhealthy nodes quarantined (VMs failed over).
6. Host Resource Protection (throttle noisy neighbor VMs).

F: Antivirus Best practices for Hyper-V and Windows Server Failover Cluster

Refer to the Microsoft link below to configure the Windows Defender Antivirus exclusions on Windows Server and make sure you include the recommended files and folders in the exclusion list not on default paths as suggested in the document:

<https://docs.microsoft.com/en-us/windows/security/threat-protection/windows-defender-antivirus/configure-server-exclusions-windows-defender-antivirus>

The following Microsoft URL also provides recommended antivirus exclusions for Hyper-V hosts:

<https://support.microsoft.com/en-us/help/3105657/recommended-antivirus-exclusions-for-hyper-v-hosts>

About the Authors

Sanjeev Naldurgkar, Technical Marketing Engineer, Cisco Systems, Inc.

Sanjeev has been with Cisco for eight years focusing on delivering customer-driven solutions on Microsoft Hyper-V and VMware vSphere. He has over 18 years of experience in the IT Infrastructure, Server virtualization, and Cloud Computing. He holds a bachelor's degree in Electronics and Communications Engineering, and leading industry certifications from Microsoft and VMware.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the author would like to thank:

- Babu Mahadevan V, Cisco Systems, Inc.
- Jeffrey Nichols, Cisco Systems, Inc.
- Nathan Tran, Cisco Systems, Inc.
- Charles Back, Cisco Systems, Inc.