

Cisco HyperFlex Edge 4.0 with Cisco Intersight

Deployment Guide for HyperFlex Edge Systems with Cisco Intersight Cloud Management Platform and VMware ESXi

Last Updated: October 11, 2019



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. (o8ogR)

© 2019 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary.....	6
Solution Overview.....	7
Introduction.....	7
Audience.....	8
Purpose of this Document.....	8
What's New in this Release?.....	8
Solution Summary.....	8
Technology Overview.....	10
Cisco HyperFlex HX-Series Nodes.....	10
Cisco HyperFlex HXAF220c-M5SX All-Flash Node.....	10
Cisco HyperFlex HX220c-M5SX Hybrid Node.....	11
Cisco HyperFlex Edge HXAF-E-220M5SX All-Flash Node.....	11
Cisco HyperFlex Edge HX-E-220M5SX Hybrid Node.....	12
All-Flash Versus Hybrid.....	12
Cisco VIC 1457 MLOM Interface Cards.....	13
Cisco HyperFlex Data Platform Software.....	14
Cisco HyperFlex Connect HTML5 Management Web Page.....	15
Cisco HyperFlex HX Data Platform Administration Plug-in.....	16
Cisco HyperFlex HX Data Platform Controller.....	16
Data Operations and Distribution.....	17
Cisco Intersight - Cloud Based Management Platform.....	19
HyperFlex Installation with Cisco Intersight.....	20
Invisible Cloud Witness for HyperFlex Edge 2-node cluster.....	22
Intersight Virtual Appliance.....	23
Solution Design.....	25
Requirements.....	25
Physical Components.....	25
Software Components.....	26
Licensing.....	27
Considerations.....	29
Hypervisors.....	29
vCenter Server.....	29
Version Control.....	29
Scale.....	29
Capacity.....	30
Data Protection.....	32
Network Design.....	32
HyperFlex Networking.....	32
VLANs and Subnets.....	34
Jumbo Frames.....	35

QoS System Classes	35
Physical Topology	35
Topology Overview	35
1GE Single-Switch Connectivity for HyperFlex Edge 3- or 4-Node Cluster	37
1GE Dual-Switch Connectivity for HyperFlex Edge 3- or 4-Node Cluster	38
1GE Single-Switch Connectivity for HyperFlex Edge 2-Node Cluster	39
1GE Dual-Switch Connectivity for HyperFlex Edge 2-Node Cluster.....	40
10GE Single-Switch Connectivity for HyperFlex Edge Clusters	41
10GE Dual-Switch Connectivity for HyperFlex Edge Clusters.....	43
Logical Topology	44
Topology Overview	44
1GE Single-Switch Logical Networking for HyperFlex Edge 3- or 4-Node Cluster.....	45
1GE Dual-Switch Logical Networking for HyperFlex Edge 3- or 4-Node Cluster	46
1GE Single-Switch Logical Networking for HyperFlex Edge 2-Node Cluster	47
1GE Dual-Switch Logical Networking for HyperFlex Edge 2-Node Cluster.....	47
10GE Single-Switch Logical Networking for HyperFlex Edge Clusters	49
10GE Dual-Switch Logical Networking for HyperFlex Edge Clusters.....	49
Deployment of Hardware and Software	51
Deployment Options	51
Solution Architecture.....	52
HXDP 4.0(1a) Validation	52
HXDP 4.0(1b) Validation	53
Preinstallation Checklist.....	54
Network Topology	54
CIMC Mode	54
vCenter Configuration	55
Network Services	55
VLANs.....	57
IP Addressing	58
Usernames and Passwords	58
Physical Installation.....	59
Cabling.....	59
Intersight Connectivity	59
Install Cisco HyperFlex Cluster	59
Cisco Intersight Account.....	60
Cisco UCS IMC Configuration	60
Claim Devices in Intersight	65
HyperFlex Edge Installation – Single Site	68
HyperFlex Edge Installation – Multiple Sites	81
Post-install Configuration.....	98
Check ESXi Configuration.....	101
HyperFlex Licensing	103

Unclaim HyperFlex Edge Cluster in Intersight	104
Reclaim HyperFlex Edge Cluster in Intersight.....	105
Upgrade Cisco HyperFlex Edge Cluster	107
ESXi Hypervisor Installation	112
Validation.....	119
Initial Functionality Validation.....	119
Verify Cluster Status.....	119
Create Datastores	120
Create Virtual Machines	121
Snapshots	121
Ready Clones	123
vMotion	124
Verify Redundancy.....	124
Test - HyperFlex Node Failover.....	125
Test – Network Uplink Failover.....	127
Test - Capacity Drive Failure	129
Test - Caching Drive Failure	133
Bill of Materials	135
Summary	138
For More Information.....	138
Appendix A: HyperFlex Cluster Capacity Calculations.....	139
Appendix B: HyperFlex Sizer	140
Appendix C: Example Cisco Nexus 9348GC-FXP Switch Configuration	145
Appendix D: Example Cisco Catalyst 9300-48P Switch Configuration	147
Appendix E: Example Multi-site Deployments using Intersight API.....	149
Appendix F: Example Script for ESXi Post-Install Configuration	152
About the Authors.....	153
Acknowledgements	153

Executive Summary

Cisco HyperFlex™ systems provide an all-purpose virtualized server platform, with hypervisor hosts, network connectivity, and virtual server storage across a set of Cisco HyperFlex HX-Series x86 rack-mount servers. The platform combines the converged computing and networking capabilities provided by the Cisco Unified Computing System™ (Cisco UCS®) with next-generation hyperconverged storage software to uniquely provide the computing resources, network connectivity, storage, and hypervisor platform needed to run an entire virtual environment, all contained in a single uniform system.

Cisco HyperFlex Edge System is a new flavor of Cisco HyperFlex system which is optimized for remote sites, branch offices, and edge environments. As a smaller form factor of Cisco hyperconverged solution, Cisco HyperFlex Edge offers the full power of our next generation hyperconverged platform without the need for connecting to Cisco UCS Fabric Interconnects. Cisco HyperFlex Edge systems support a configuration of two, three or four HyperFlex HX220c converged nodes, and allows the ability to scale up capacity by hot-adding additional capacity drives to the nodes.

Traditionally, the configuration, deployment, management, and monitoring datacenter solutions are done with existing separate tools for Cisco UCS, HyperFlex or VMware. The management of a Cisco HyperFlex cluster can be done through a VMware vSphere web client plug-in or through a HTML5 based native HyperFlex Connect management tool. Since Cisco HyperFlex version 2.6, the support of Cisco Intersight cloud-based remote monitoring and management has raised the management of Cisco UCS and HyperFlex systems to a new level. With Cisco Intersight the deployment and management can be done from the cloud, providing a low-cost, easy-to-deploy, remote management feature set for Cisco HyperFlex systems to your edge environments without requiring experienced IT at the deployment site.

The Cisco Validated Design (CVD) program consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. This CVD presents a validated Remote Office and Branch Office (ROBO) solution in a VMware ESXi virtual environment, optimized with the Cisco HyperFlex platform that can be deployed and managed from Cisco Intersight cloud platform.

Solution Overview

Introduction

Deploying IT infrastructure at the remote sites in a fast, efficient and consistent way is always challenging for many customers, especially for the scenarios where fewer or nearly zero IT professionals are present at the site. At the same time the requirements for IT resources in edge environments keep increasing, while evolving technologies such as image processing, AI/ML, autonomously driven vehicles, and intelligent manufacturing are growing rapidly. These changes result in many new challenges for customers who are transitioning their applications and workloads. The choice of hardware and software layers for the distributed environments becomes very important, because the efficiency of the infrastructure affects the efficiency of the applications, the speed of data collection and processing, storage performance, and resource management.

Virtualization is an ideal solution for many of these challenges. It is a technology that allows for the sharing and easy expansion of underlying hardware resources by multiple workloads. This approach leads to higher utilization of IT resources while providing necessary fault tolerance. Hyperconvergence is an evolving technology that leverages many benefits of virtualization. Hyperconverged infrastructures bring the simplification of deployment, centralized management, as well as increased agility, thereby reducing the amount operational costs of daily IT operations.

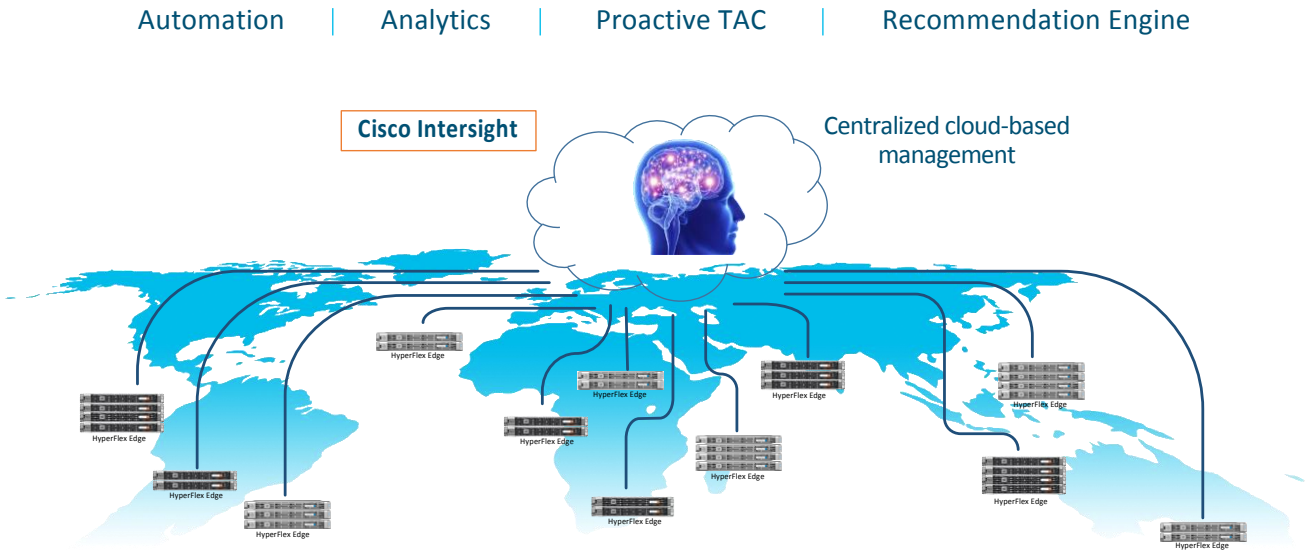
As a proven industry leading hyperconvergence platform, a Cisco HyperFlex system delivers many enterprise-class features, such as:

- A fully distributed log-structured file system that supports thin provisioning
- High performance and low latency flash-friendly architecture
- In-line data optimization with deduplication and compression
- Fast and space-efficient clones through metadata operations
- Native replication of virtual machine snapshots
- Cloud based centralized deployments, witnessing, and ongoing management

Cisco HyperFlex systems let you unlock the full potential of hyperconvergence and adapt IT to the needs of many workloads. With Cisco HyperFlex systems, customers have the flexibility to support different types of workloads without comprising their performance requirements. Cisco Intersight cloud based central management provides additional advantages for deploying and managing Cisco HyperFlex Edge infrastructures at multiple sites in parallel.

Cloud-based management platforms provide unified access to applications and infrastructure monitoring, configuration and orchestration, therefore reducing the complexity of IT management by simplifying and unifying the deployment and management of many edge devices. Cisco Intersight is Cisco's new system management platform for Cisco UCS servers and Cisco HyperFlex systems that delivers intuitive computing through cloud-powered intelligence. This cloud-based platform offers a centralized management that enables IT organizations to analyze, simplify, and automate their environments in ways that were not possible with traditional tools. This capability empowers organizations to achieve significant savings in Total Cost of Ownership (TCO) and to deliver applications faster, so they can support new business initiatives. To access the Cisco Intersight platform, just go to the website: <https://intersight.com/>.

Figure 1 Solution Overview



Audience

The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering and customers who want to deploy Cisco HyperFlex systems for ROBO or edge environments. The readers of this document are expected to have the necessary understanding of VMWare vSphere virtual architecture, Cisco UCS and HyperFlex servers, and Cisco Intersight cloud management platform. External references are provided where applicable and it is recommended that the reader be familiar with these documents.

Purpose of this Document

This document describes how to deploy a Cisco HyperFlex Edge cluster to a single site or to multiple sites from the cloud using Cisco Intersight. It includes design guidance explaining the architecture and topologies, capacity and scalability, and a bill of materials. This document presents a tested and validated solution and provides insight into operational best practices.

What's New in this Release?

The initial version of this CVD was published on September 11, 2019 with the validation of the solution with HXDP version 4.0(1a).

This new release extends the validation of the solution with HXDP version 4.0(1b). It includes support for the second generation Intel® Xeon® Scalable processor (formerly Cascade Lake).

Solution Summary

A HyperFlex Edge cluster is composed of two, three or four Cisco UCS HX-Series Converged Nodes (with disk storage). Cluster node expansion is not supported but is planned for a future software release with 10GE topologies. Each node is equipped with at least one high-performance SSD drive for data caching and rapid acknowledgment of write requests. Each node is also equipped with additional disks, up to the platform's physical limit, for long term storage capacity.

The following are the components of a Cisco HyperFlex Edge system:

- Cisco HyperFlex HX-Series Rack-Mount Servers, the supported models are listed below:
 - Cisco HyperFlex HX220c-M5SX Rack-Mount Servers (2-, 3-, or 4-node configuration)
 - Cisco HyperFlex HXAF220c-M5SX All-Flash Rack-Mount Servers (2-, 3-, or 4-node configuration)
 - Cisco HyperFlex HX220c-M4S Rack-Mount Servers (3-node configuration only)
 - Cisco HyperFlex HXAF220c-M4S All-Flash Rack-Mount Servers (3-node configuration only)
- Cisco HyperFlex Data Platform Software
- VMware vSphere ESXi Hypervisor
- VMware vCenter Server (end-user supplied)

In this solution Cisco Intersight is the management platform that performs the deployment and administration of the HyperFlex Edge systems across multiple sites.



Note: Cisco UCS M4 servers are discontinued and are only supported for 3-node HyperFlex Edge clusters with existing purchase. Therefore, the configuration and deployment of HyperFlex M4 edge cluster is not explained in this CVD.

Technology Overview

Cisco HyperFlex systems are built on the Cisco UCS platform which can be deployed quickly and are highly flexible and efficient, reducing risk for the customer. Cisco HyperFlex delivers the simplicity, agility, scalability, and pay-as-you-grow economics of the cloud with the benefits of multisite, distributed computing at global scale.

Cisco HyperFlex Edge is a new version of the Cisco HyperFlex system that is optimized for remote sites, branch offices, and Edge environments. A smaller form factor of the Cisco hyperconverged solution, Cisco HyperFlex Edge offers the full power of a next generation hyperconverged platform without the need to connect to Cisco UCS Fabric Interconnects.

Cisco HyperFlex HX-Series Nodes

The Cisco HyperFlex Edge systems are built with Cisco UCS rack mount servers without the requirement of using Cisco UCS fabric interconnects. To integrate management of the HyperFlex Edge nodes with Cisco Intersight cloud management platform, the Cisco Integrated Management Controller (CIMC) service needs to run within the servers. The Cisco Integrated Management Controller (CIMC) is a baseboard management software that provides embedded server management for Cisco UCS C-Series and HX-Series Rack Servers. It can be configured to operate in Dedicated Mode or Shared Mode. The Dedicated Mode uses the dedicated management port on the server motherboard. In the Shared Mode any LOM port or VIC adapter card port can be used to access the CIMC.

A Cisco HyperFlex Edge cluster requires a minimum of two HX-Series edge nodes. The HX-Series edge nodes combine the CPU and RAM resources for hosting guest virtual machines with a shared pool of the physical storage resources used by the HX Data Platform software. HX-Series hybrid edge nodes use a combination of solid-state disks (SSDs) for caching and hard-disk drives (HDDs) for the capacity layer. HX-Series all-flash edge nodes use SAS SSD for the caching layer and SATA SSDs for the capacity layer.

Cisco HyperFlex HXAF220c-M5SX All-Flash Node

The HXAF220c-M5SX servers extend the capabilities of Cisco's HyperFlex portfolio in a 1U form factor with the addition of the Intel® Xeon® Processor Scalable Family, 24 DIMM slots with configuration options ranging from 128GB up to 3TB of DRAM, and an all flash footprint of cache and capacity drives for highly available, high performance storage.

Figure 2 HXAF220c-M5SX All-Flash Node

Front View with Bezel Attached



Front View with Bezel Removed



Rear View (no VIC or PCIe adapters installed)



Cisco HyperFlex HX220c-M5SX Hybrid Node

The HX220c M5 servers extend the capabilities of Cisco's HyperFlex portfolio in a 1U form factor with the addition of the Intel® Xeon® Processor Scalable Family, 24 DIMM slots with configuration options ranging from 128GB up to 3TB of DRAM, and an all flash footprint of cache drives and hard disk capacity drives for highly available, high performance storage.

Figure 3 HX220c-M5SX Node

Front View with Bezel Attached



Front View with Bezel Removed



Rear View (no VIC or PCIe adapters installed)



Note: Neither NVMe SSDs nor SEDs are supported on the HyperFlex Edge systems.

The HX220 hybrid or all-flash nodes used for building HyperFlex Edge systems are the same nodes used for building non-Edge standard HyperFlex clusters. The HyperFlex Edge nodes can be ordered from a specifically created edge bundle - HX-E-M5S-HXDP, in which two separate product identifiers are created: HXAF-E-220M5SX for all-flash server and HX-E-220M5SX for hybrid server. This allows different choices of hardware components for the edge nodes such as CPU, disks and network adapters. Customers have the flexibility of configuring HX Edge nodes with a single CPU with SKU HX-CPU-4114 and above. Lower bin CPU SKUs such as HX-CPU-3106, HX-CPU-4108 or HX-CPU-4110 are only supported in dual CPU configured HyperFlex Edge systems. The customers also have the flexibility of configuring HyperFlex Edge nodes with minimum three capacity disks.

Cisco HyperFlex Edge HXAF-E-220M5SX 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 1.6 TB or 400GB SAS SSD write-log drive, a minimum of three, and up to eight 960 GB or 3.8 terabyte (TB) SATA SSD drives for storage capacity.

Figure 4 HXAF-E-220M5SX All-Flash Node

Front View with Bezel Attached



Front View with Bezel Removed



Rear View (no VIC or PCIe adapters installed)



Cisco HyperFlex Edge HX-E-220M5SX Hybrid Node

This small footprint Cisco HyperFlex hybrid model contains a minimum of three, and up to eight 2.4 terabyte (TB), 1.8TB or 1.2 TB SAS hard disk drives (HDD) that contribute to cluster storage capacity, a 240 GB SSD housekeeping drive, a 480 GB SATA SSD or 800 GB SAS SSD caching drive, and a 240 GB M.2 form factor SSD that acts as the boot drive.

Figure 5 HX-E-220M5SX Node

Front View with Bezel Attached



Front View with Bezel Removed



Rear View (no VIC or PCIe adapters installed)



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 nodes in the cluster.
- 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.

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 VIC 1457 MLOM Interface Cards


A Cisco innovation, the Cisco UCS Virtual Interface Card (VIC) 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 Express (PCIe) virtual devices, either virtual Network Interface Cards (vNICs) or virtual Host Bus Adapters (vHBAs), with a high rate of I/O Operations Per Second (IOPS), support for lossless Ethernet, and 10/25/40/100-Gbps connection to servers. 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 Cisco VIC 1400 series is designed exclusively for the M5 generation of Cisco UCS and HyperFlex servers and is recommended for the deployment of the 10GE HyperFlex Edge cluster.

The Cisco UCS VIC 1457 is a quad-port Small Form-Factor Pluggable (SFP28) 10/25-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 or C-Series Rack Servers.

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).

Figure 6 Cisco VIC 1457 mLOM Card



Cisco HyperFlex Data Platform Software

The Cisco HyperFlex delivers a new generation of flexible, scalable, enterprise-class hyperconverged solutions. The solution also delivers storage efficiency features such as thin provisioning, data deduplication, and compression for greater capacity and enterprise-class performance. Additional operational efficiency is facilitated through features such as cloning and snapshots.

The complete end-to-end hyperconverged solution provides the following benefits to customers:

- **Simplicity:** The solution is designed to be deployed and managed easily and quickly through familiar tools and methods. No separate management console is required for the Cisco HyperFlex solution.
- **Centralized hardware management:** The cluster hardware is managed in a consistent manner by Cisco Intersight. Cisco Intersight also provides a single console for solution management, including firmware management. Cisco HyperFlex HX Data Platform clusters are managed through a plug-in to VMware vCenter, or through HyperFlex Connect, a native HTML5 UI.
- **High availability:** Component redundancy is built into most levels at the node. Cluster-level tolerance of node and network failures is implemented as well.
- **Enterprise-class storage features:** Complementing the other management efficiencies are features such as thin provisioning, data deduplication, compression, cloning, and snapshots to address concerns related to overprovisioning of storage.
- **Flexibility with a "pay-as-you-grow" model:** Customers can purchase the exact amount of computing and storage they need and expand one node at a time up to the supported cluster node limit.
- **Agility to support different workloads:** Support for both hybrid and all-flash models allows customers to choose the right platform configuration for capacity-sensitive applications or performance-sensitive applications according to budget requirements.

The Cisco HyperFlex HX Data Platform is a purpose-built, high-performance, distributed file system with a wide array of enterprise-class data management services. The data platform's innovations redefine distributed storage technology,

exceeding the boundaries of first-generation hyperconverged infrastructures. The data platform has all the features that you would expect of an enterprise shared storage system, eliminating the need to configure and maintain complex Fibre Channel storage networks and devices. The platform simplifies operations and helps ensure data availability. Enterprise-class storage features include the following:

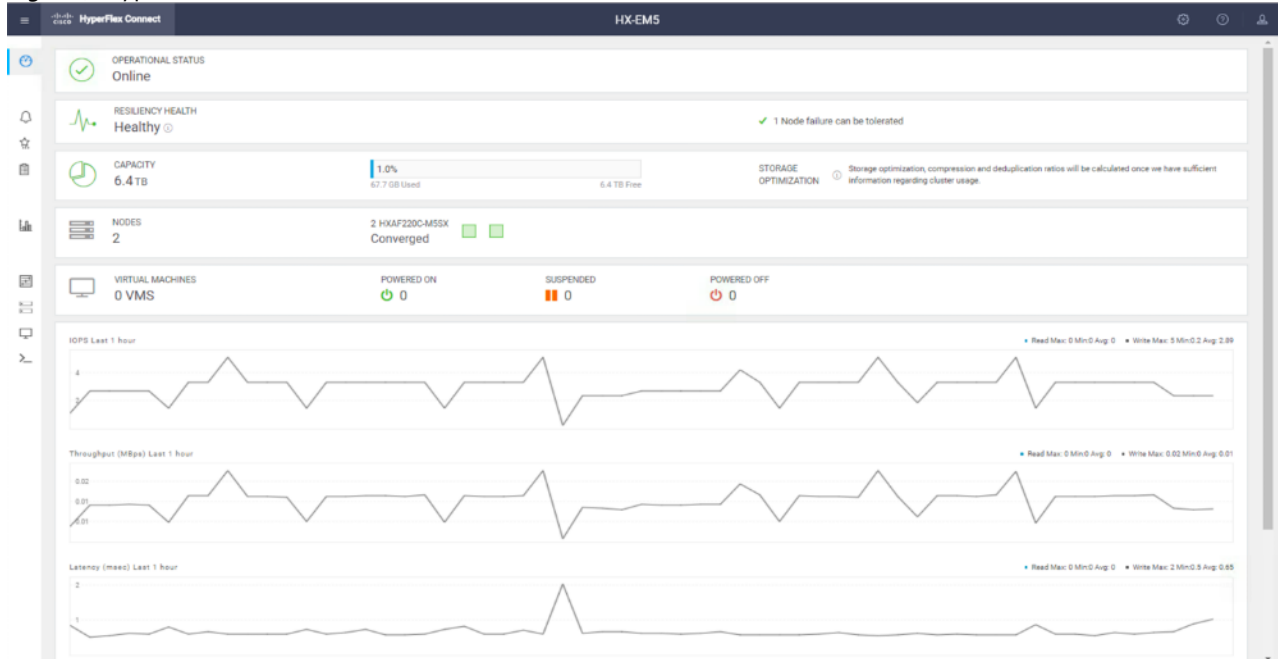
- Replication of all written data across the cluster so that data availability is not affected if single or multiple components fail (depending on the failure scenario).
- Deduplication is always on, helping reduce storage requirements in which multiple operating system instances in client virtual machines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

The HX Data Platform can be administered through a VMware vSphere web client plug-in or through the HTML5-based native Cisco HyperFlex Connect management tool. Additionally, since the HX Data Platform Release 2.6, Cisco HyperFlex systems can also be managed remotely by the Cisco Intersight™ cloud-based management platform. Through the centralized point of control for the cluster, administrators can create datastores, monitor the data platform health, and manage resource use.

Cisco HyperFlex Connect HTML5 Management Web Page

A 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>>.

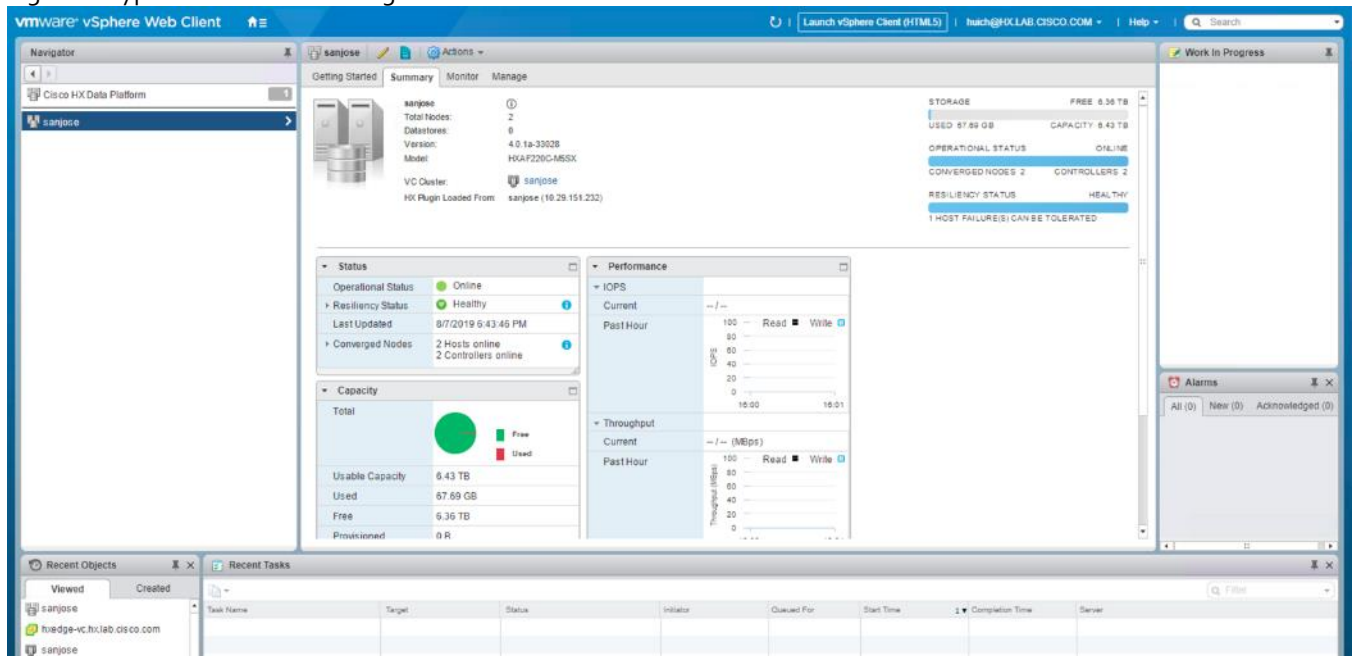
Figure 7 HyperFlex Connect GUI



Cisco HyperFlex HX Data Platform Administration Plug-in

The Cisco HyperFlex HX Data Platform is also administered secondarily through a VMware vSphere web client plug-in.

Figure 8 HyperFlex Web Client Plug-in



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. The Storage Controller Virtual Machine (SCVM) uses the VMDirectPath I/O feature to provide PCI passthrough

control of the physical server's SAS disk controller. This method gives the controller VM full control of the physical disk resources, utilizing the SSD drives as a read/write caching layer, and the HDDs or SDDs as a capacity layer for distributed storage. The controller integrates the data platform into the VMware vSphere cluster through the use of three preinstalled VMware ESXi vSphere Installation Bundles (VIBs) on each node:

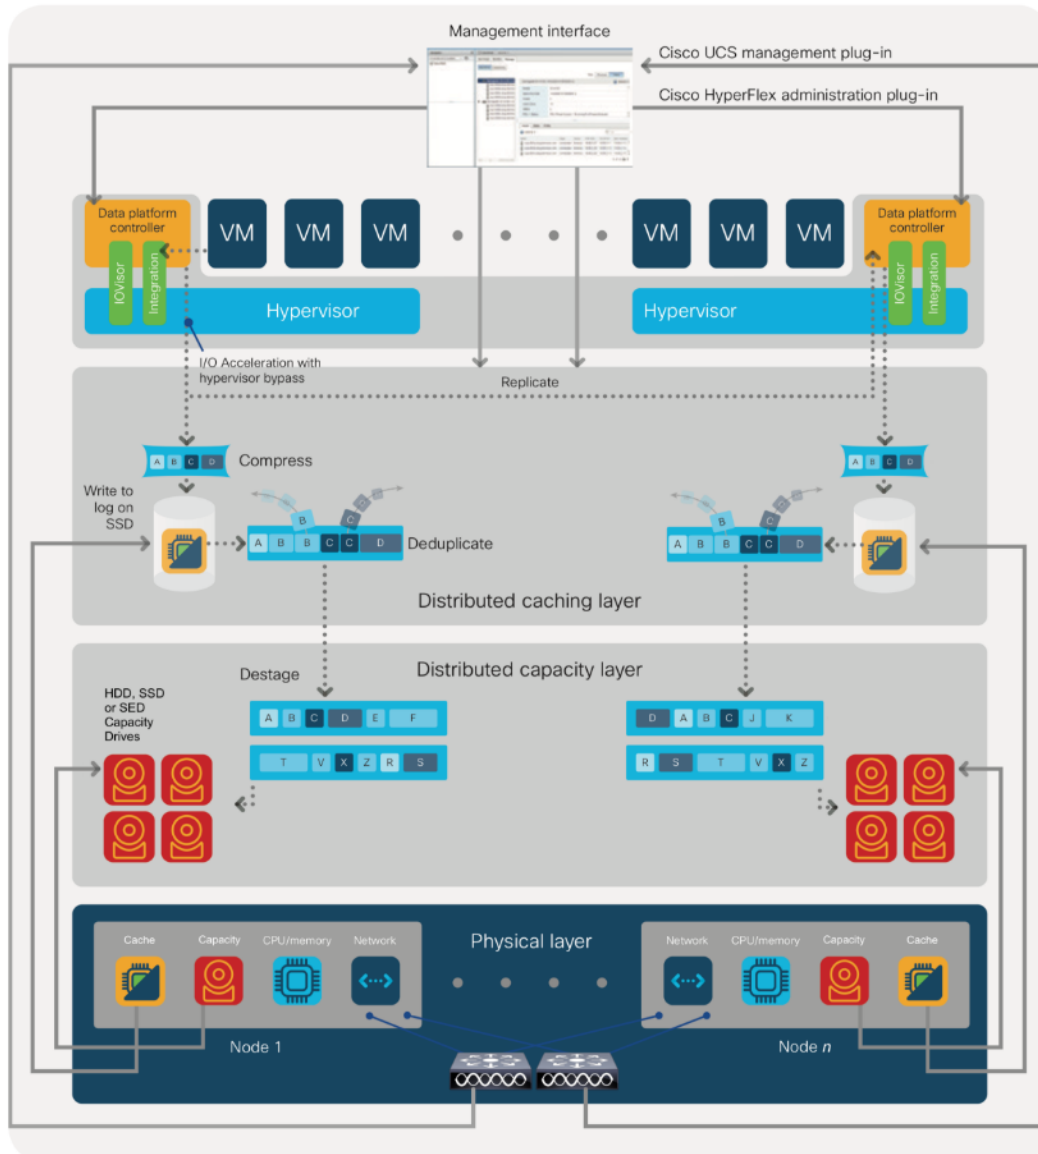
- **IO Visor:** This VIB provides a network file system (NFS) mount point so that the ESXi hypervisor can access the virtual disks that are attached to individual virtual machines. From the hypervisor's perspective, it is simply attached to a network file system. The IO Visor intercepts guest VM IO traffic,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

Data Operations and Distribution

The HX Data Platform controllers handle all read and write requests from the guest virtual machines to the virtual machine disks (VMDKs) stored in the distributed data stores in the cluster. The data platform distributes the data across multiple nodes of the cluster and across multiple capacity disks in each node according to the replication-level policy selected during cluster setup. The replication-level policy is defined by the replication factor (RF) parameter. When RF = 3, a total of three copies of the blocks are written and distributed to separate locations for every I/O write committed to the storage layer; when RF = 2, a total of two copies of the blocks are written and distributed. The HyperFlex Edge clusters only support RF=2.

Figure 9 shows the movement of data in the HX Data Platform.

Figure 9 Cisco HyperFlex HX Data Platform Data Movement



For each write operation, the data is intercepted by the IO Visor module on the node on which the virtual machine is running, a primary node is determined for that particular operation through a hashing algorithm, and the data is then sent to the primary node. The primary node compresses the data in real time and 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. Because the virtual disk contents have been divided and spread out through the hashing algorithm, the result of this method is that all write operations are spread across all nodes, avoiding problems related to data locality and helping prevent “noisy” virtual machines from consuming all the I/O capacity of a single node. The write operation will not be acknowledged until all the desired 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 read operations are requested on data that has recently been written.

The 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. Based on policies accounting for I/O load and access patterns, those write-cache segments are locked, and new write operations 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 a hybrid system or to the SSD capacity layer of the nodes for an 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 SSDs 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 for 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 a minimal amount of time. Deduplication, compression, and destaging take place with no delays or I/O penalties for the guest virtual machines making requests to read or write data, which benefits both the HDD and SSD configurations.

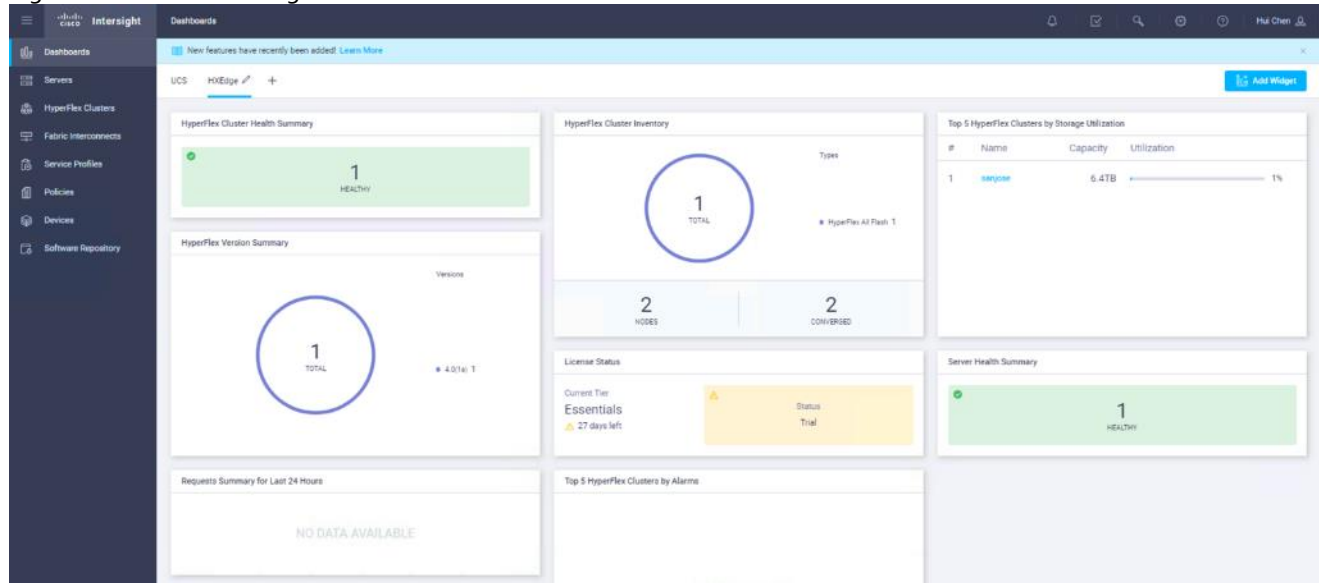
For data read operations, data may be read from multiple locations. For data that was very recently written, the data is likely to still exist in the write log of the local platform controller memory or in the write log of the local caching-layer disk. If local write logs do not contain the data, the distributed file system 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 file system will retrieve the requested data from the distributed capacity layer. As requests for read operations are made to the distributed file system 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. 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.

Cisco Intersight - Cloud Based Management Platform

Cisco Intersight (<https://intersight.com>) is an API driven, cloud-based system management platform. It is designed to help organizations to achieve their IT management and operations with a higher level of automation, simplicity, and operational efficiency. It is a new generation of global management tool for the Cisco Unified Computing System (UCS) and Cisco HyperFlex systems and provides a holistic and unified approach to managing the customers' distributed and virtualized environments. Cisco Intersight simplifies the installation, monitoring, troubleshooting, upgrade and support for your infrastructure with the following benefits:

- **Cloud Based Management:** The ability to manage Cisco UCS and HyperFlex from the cloud provides the customers the speed, simplicity and easy scaling in the management of their infrastructure whether in the datacenters or remote and branch office locations.
- **Automation:** Unified API in Cisco UCS and Cisco HyperFlex systems enables policy driven configuration and management of the infrastructure and it makes Intersight itself and the devices connected to it fully programmable and DevOps friendly.
- **Analytics and Telemetry:** Intersight monitors the health and relationships of all the physical and virtual infrastructure components. It also collects telemetry and configuration information for developing the intelligence of the platform in the way in accordance with Cisco information security requirements.
- **Connected TAC:** Solid integration with Cisco TAC enables more efficient and proactive technical support. Intersight provides enhanced operations automation by expediting sending files to speed troubleshooting.
- **Recommendation Engine:** Driven by analytics and machine learning, Intersight recommendation engine provides actionable intelligence for IT operations management from daily increasing knowledge base and practical insights learned in the entire system.
- **Management as A Service:** Cisco Intersight provides management as a service and is designed to be infinitely scale and easy to implement. It relieves users of the burden of maintaining systems management software and hardware.

Figure 10 Cisco Intersight



Cisco Intersight enables the support of monitoring and deploying Cisco HyperFlex clusters. The Cisco Intersight website and framework can be upgraded with new and enhanced feature sets 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.

The minimum supported web browser versions to run Cisco Intersight are:

- Google Chrome: 62.0.3202.94
- Firefox: 57.0.1
- Microsoft Edge: 40.15063.674.0
- Safari: 10.1.1

The minimum supported firmware versions for Cisco HyperFlex Edge systems to be deployed by Intersight are:

- Cisco IMC Software: 3.1(2d)
- Cisco HyperFlex Data Platform (HXDP): 2.6.1b
- Device Connector: 1.0.4-10 or later

HyperFlex Installation with Cisco Intersight

Cisco Intersight provides comprehensive lifecycle management for the HyperFlex systems, including remote cloud-based installation. Since HXDP version 4.0, support has been extended from 3-node only Edge clusters to include 2, 3 and 4-node Edge clusters to run HyperFlex in different environments. Also, since HXDP version 4.0, the Cisco Intersight Invisible Cloud Witness service is available for supporting 2-node Cisco HyperFlex Edge cluster deployments. The Cisco HX-series rack-mount servers are connected to Cisco Intersight via the network, so that Cisco Intersight can manage and configure the nodes.

Cisco Intersight provides an installation wizard to install, configure, and deploy Cisco HyperFlex Edge clusters. The wizard constructs a pre-configuration definition of the cluster called a HyperFlex Cluster Profile. HyperFlex Cluster Profiles are built on policies in which administrators define sets of rules and operating characteristics, such as the node identity, interfaces, and network connectivity. Every active node in the HyperFlex cluster must be associated with a HyperFlex Cluster Profile. After gathering the node configuration settings to build the HyperFlex Cluster Profile, the installation wizard will validate and deploy the HyperFlex Cluster Profile to your Cisco HX-series nodes, thereby creating a Cisco HyperFlex Edge cluster. You can clone a successfully deployed HyperFlex Cluster Profile, and then use that copy as the basis to create a new cluster.

HyperFlex Policies in Cisco Intersight provide different configurations including Auto Support, security, network configuration and more. A policy that has been configured can be assigned to any number of servers in order to provide a configuration baseline.

Table 2 lists all HyperFlex policies that are required to define your HyperFlex Edge Cluster Profile.

Table 1 Intersight Policies for HyperFlex Edge Cluster Profile

Policies	Purpose	Fields of Configuration	Notes
Security Policy	Configures ESXi and Controller VM password for the HyperFlex cluster.	Hypervisor Admin _____ Hypervisor password _____ Controller VM Admin Password _____	This policy presents an option to update the Hypervisor password in Intersight.
System Configuration Policy (DNS, NTP, and Timezone)	Configures DNS, NTP, and Timezone on all servers.	Timezone _____ DNS Suffix _____ DNS Server _____ NTP Server _____	DNS and NTP servers should reside outside of the HyperFlex storage cluster.
vCenter Policy	An optional policy for registering the cluster to vCenter during installation of the HyperFlex cluster.	vCenter Server _____ vCenter Username _____ vCenter Password _____ vCenter Datacenter Name _____	
Storage Configuration Policy	Configures the options for VDI Optimization (for hybrid HyperFlex systems) and cleanup of Disk Partitions.	<input type="checkbox"/> VDI Optimization <input type="checkbox"/> Cleanup of Disk Partitions	
Auto Support Policy	The option to enable Auto Support. If enabled, notifications are sent to designated email addresses or email aliases that you want to receive the notifications.	<input type="checkbox"/> Auto-Support Send Notification to _____	Auto Support is the alert notification service provided through HXDP. Auto Support is configured by configuring the SMTP mail server and adding email recipients.
Node Configuration Policy (IP & Hostname)	Configures the management IP pool for the hypervisors and the controller VMs.	Hostname Prefix _____ Management Starting IP _____	Designate which the HX edge nodes are going to associate to this HX Cluster Profile.

Policies	Purpose	Fields of Configuration	Notes
		Management Ending IP _____ Management Subnet Mask _____ Management Subnet Gateway _____ Controller VM Starting IP _____ Controller VM Ending IP _____ Controller VM Subnet Mask _____ Controller VM Subnet Gateway _____	
Network Configuration Policy	Configures the Uplink Speed, MAC Prefix, VLAN, and Jumbo Frames for the management network in the Edge clusters.	Uplink Speed _____ MAC Prefix Starting Address _____ MAC Prefix Ending Address _____ Management VLAN ID _____ <input type="checkbox"/> Jumbo Frames	
Proxy Setting Policy	The option to specify the HTTP proxy settings for HX installation process and the HyperFlex Storage Controller VMs.	Proxy Hostname _____ Port _____ Username _____ Password _____	This policy is required when the internet access of your servers is secured by an HTTP proxy.
HyperFlex Storage Network Policy	Configures the HX storage network VLAN ID.	Storage network VLAN ID _____	This policy cannot be saved as the VLAN ID is required to be input uniquely every time.

Invisible Cloud Witness for HyperFlex Edge 2-node cluster

The Cisco HyperFlex Data Platform software runs a fully distributed clustered file system across all the nodes in the cluster. The clustered file system utilizes a quorum mechanism that is used to guarantee data consistency and availability across different nodes. The quorum mechanism works well for three-node and larger clusters that can tolerate the failure of one or more nodes and still be able to obtain a majority consensus and continue operations.

However, fault tolerance and file system consistency become more challenging when only two nodes are deployed at a customer's remote-office or branch-office (ROBO) location. In this scenario, if one of the two nodes failed, a quorum can no longer be established using a node majority algorithm alone. In the unlikely event that the communication pathway between the two node is disrupted, a "split brain" condition may occur if both nodes continue to process data without obtaining a quorum. The opposite outcome — the loss of availability — is also possible. To avoid these scenarios the hyperconverged two-node architectures normally require an additional component, sometimes referred to as a witness or arbiter, that can vote if a failure occurs within the cluster. The witness breaks the impasse that occurs when a node fails, or when the two nodes of the cluster can no longer communicate, and therefore cannot agree with one another about the condition of the cluster. This additional witness node must be provisioned on the existing infrastructure and connected to the remote cluster over the customer's network.

With the innovative Cisco Intersight cloud platform, a new invisible witness architecture has been developed for Cisco HyperFlex Edge deployments. This architecture allows Cisco HyperFlex Edge two-node clusters be deployed without the user having to install and configure any witness nodes, which are required to achieve a quorum to maintain cluster consistency and ensure high availability. This innovative process eliminates the cost and complexity of deploying and maintaining a dedicated witness server or multiple servers. With Intersight there is no need for witness upgrades or security patches as this burden is all seamlessly handled by Cisco.


The Invisible Cloud Witness is automatically activated and configured from Cisco Intersight when a 2-node Cisco HyperFlex Edge cluster is deployed. The Invisible Cloud Witness does not require any configuration input, instead, after the physical servers are securely claimed into the Cisco Intersight, all components are configured transparently during a 2-node Cisco HyperFlex Edge cluster installation. The Invisible Cloud Witness communicates with the Cisco Intersight service through an embedded device connector inside the Cisco HyperFlex controller virtual machines. For the two-node HyperFlex Edge clusters, Cisco Intersight performs arbitration in the scenario that a node fails or the communication between nodes fails. Arbitration and avoidance of a split-brain scenario are accomplished by using an internal election protocol that grants the right to continue operating to a single winner. The Cisco Intersight maintains the election state for each remote cluster and will essentially vote with one of the surviving nodes to help ensure continued cluster operation in the event of a failure. Only when a failure is detected, the surviving nodes reach out to the Cisco Intersight platform for permission to continue I/O operations. The witness service then grants privileges to only one of the nodes to continue operation. The two Cisco HyperFlex clustered nodes also periodically perform heartbeat tests over the local network. If reachability is interrupted because node or link failure, the process of reaching out to the Cisco Intersight platform will begin.

For more details on how the Invisible Cloud Witness it is built and how it operates, please refer to the [Cisco white paper Cisco Intersight Invisible Cloud Witness for Cisco HyperFlex Edge](#).

Intersight Virtual Appliance

The Cisco Intersight Virtual Appliance delivers the management features of Intersight for Cisco UCS and HyperFlex into the on-premise environment. It is deployed from a VMware OVA that enables the additional control to specify what data is sent back to Cisco with a single point of egress within the enterprises network. The virtual appliance form factor enables additional data locality, security, or compliance needs that are not completely met by connecting directly to intersight.com in the cloud. However, The Cisco Intersight Virtual Appliance is not intended for an environment with no external connectivity, the Cisco Intersight virtual appliance requires an internet connection back to Cisco and the cloud-based Intersight services for updates and to deliver some of the product features. Communication back to Cisco can be redirected via a proxy server if direct connectivity is not available or allowed by policy. Updates to the virtual appliance are automated and applied during a user specified recurring maintenance window. This connection also facilitates the streamlining of Cisco TAC services for Cisco UCS and HyperFlex systems, with features like automated support log collection.

Figure 11 Cisco Intersight Virtual Appliance



Cisco Intersight Virtual Appliance OVA can be downloaded from Cisco website and can be deployed as a virtual machine in your existing environment. Cisco Intersight Virtual Appliance uses a subscription-based license delivered via Cisco Smart Licensing. After the installation of the appliance OVA is completed, you must connect the appliance to Cisco Intersight, and register the license as part of the initial setup process.

Solution Design

Requirements

The following sections detail the physical hardware, software revisions, and firmware versions required for the solution, which deploys the Cisco HyperFlex Edge system with Cisco Intersight cloud management platform.

Physical Components

Table 2 lists the hardware components required to install a single cluster of the Cisco HyperFlex Edge system.

Table 2 HyperFlex Edge System Components

Component	Hardware Required
HX-Series Servers	Two to Four Cisco HyperFlex HXAF-E-220M5SX All-Flash rack servers, or Two to Four Cisco HyperFlex HX-E-220M5SX Hybrid rack servers

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

HXAF-E-220M5SX Spec Sheet:

<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hxaf-e-220m5sx-edge.pdf>

HX-E-220M5SX Spec Sheet:

<https://www.cisco.com/c/dam/en/us/products/collateral/hyperconverged-infrastructure/hyperflex-hx-series/hx-e-220m5sx-edge-specsheet.pdf>

For the complete server specifications and additional information, see [Cisco's Hyperconverged Infrastructure](#) site.

Table 3 lists the required hardware components and disk options for the Cisco HXAF-E-220M5SX All-Flash rack servers, which are required for creating the HyperFlex Edge cluster:

Table 3 HXAF-E-220M5SX Edge Server Options

HXAF-E-220M5SX options	Hardware Required
Processors	Choose single or a matching pair of Intel® Xeon® Processor Scalable Family or 2 nd Generation Intel® Xeon® Processor Scalable Family CPUs
Memory	128 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 or 2933 MHz 1.2v modules
Disk Controller	Cisco 12Gbps Modular SAS HBA
SSDs	One 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One 400 GB 2.5 Inch Enterprise Performance 12G SAS SSD, or one 1.6TB 2.5in Enterprise Performance 12G SAS SSD Three to eight 3.8 TB 2.5 Inch Enterprise Value 6G SATA SSDs, or three to eight 960 GB 2.5 Inch Enterprise Value 6G SATA SSDs
Network	Onboard LOM NIC ports (1G), or optional Intel i350-T4 Quad-Port PCIe NIC Card (1GE), or optional Cisco UCS VIC1457 VIC MLOM (10GE)

HXAF-E-220M5SX 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

Table 4 lists the hardware component options for the HX-E-220M5SX server model:

Table 4 HX-E-220M5SX Edge Server Options

HX-E-220M5SX Options	Hardware Required
Processors	Choose single or a matching pair of Intel® Xeon® Processor Scalable Family or 2 nd Generation Intel® Xeon® Processor Scalable Family CPUs
Memory	128 GB to 3 TB of total memory using 16 GB, 32 GB, 64 GB, or 128 GB DDR4 2666 or 2933 MHz 1.2v modules
Disk Controller	Cisco 12Gbps Modular SAS HBA
SSDs	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	Three to eight 2.4 TB, 1.8 TB or 1.2 TB SAS 12Gbps 10K rpm SFF HDD
Network	Onboard LOM NIC ports (1GE), or optional Intel i350-T4 Quad-Port PCIe NIC Card (1GE), or optional Cisco UCS VIC1457 VIC MLOM (10GE)
Boot Device	One 240 GB M.2 form factor SATA SSD
microSD Card	One 32GB microSD card for local host utilities storage



Note: To deploy 10GE HyperFlex Edge clusters, Cisco VIC card that supports 10GE connectivity is required on each server. M5 servers require VIC 1387 and two QSAs, or VIC1457 to support 10GE connection. For 2-node and 4-node 10GE edge clusters only VIC 1457 is supported. VIC 1457 is highly recommended for M5 servers and is the only VIC card that is explained in this document.



Note: HX Edge nodes can be configured with a single CPU of more than eight cores. Lower bin CPU SKUs below HX-CPU-4114 or HX-CPU-I4210 such as HX-CPU-3106, HX-CPU-4108, HX-CPU-4110 or HX-CPU-I4208 are only supported in dual CPU configured HyperFlex Edge systems.

Software Components

Table 5 lists the software components and the versions required for deploying the HyperFlex Edge system with Cisco HyperFlex Data Platform (HXDP) software release 4.0:

Table 5 Software Components

Component	Software Required – Minimum Version
Hypervisor	VMware ESXi 6.0 Update 3, or 6.5 Update 2, or 6.7 Update 2 ESXi 6.7 U2 is recommended (CISCO Custom Image for ESXi 6.7 Update 2: HX-ESXi-6.7U2-13473784-Cisco-Custom-6.7.2.2-install-only.iso) Note: Use of a published Cisco custom ESXi ISO installer file is required when installing/reinstalling ESXi or upgrading to a newer version prior to installing HyperFlex. An offline bundle file is also provided to upgrade ESXi on running

Component	Software Required – Minimum Version
	<p>clusters.</p> <p>Note: If using VIC 1457 in 10GE topology, ESXi 6.0 is not supported.</p> <p>Note: VMware vSphere Standard, Essentials Plus, ROBO, Enterprise or Enterprise Plus licensing is required from VMware.</p>
Management Server	<p>VMware vCenter Server for Windows or vCenter Server Appliance 6.0 U3c or later.</p> <p>For information about the interoperability of your ESXi version and vCenter Server, go to: http://www.vmware.com/resources/compatibility/sim/interop_matrix.php</p> <p>Note: Using ESXi 6.5 on the HyperFlex nodes also requires using vCenter Server 6.5 and ESXi 6.7 requires vCenter Server 6.7.</p>
Cisco HyperFlex Data Platform	Cisco HXDP Software version 4.0(1a)
Cisco Integrated Management Controller (CIMC)	CIMC version 4.0(1a)
Cisco UCS Firmware	<p>Cisco UCS firmware 4.0(1a) or later.</p> <p>Note: BIOS version 4.0.1c0 or later.</p>



Note: To support the 2nd Generation Intel® Xeon® scalable family CPUs, minimum HXDP version 4.0(1b) and UCS firmware version 4.0(4f) are required.

Licensing

Cisco HyperFlex 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 time consuming and error prone manual 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.

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 on the Cisco Smart Software Manager satellite server, go to: <https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager-satellite.html>

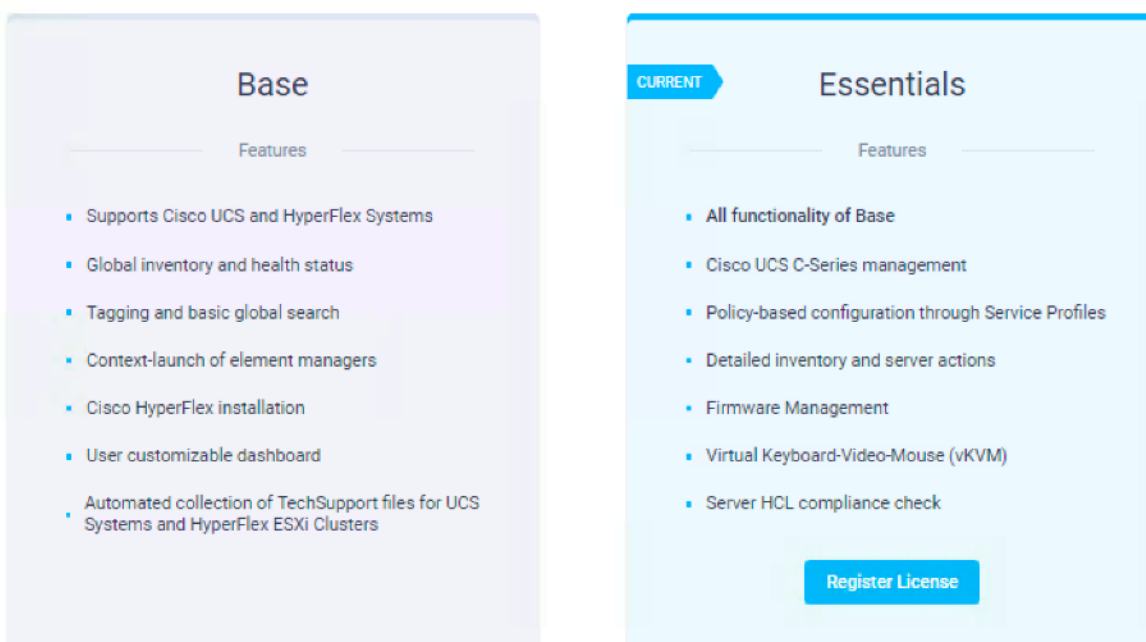
Beginning with Cisco HyperFlex 3.0, licensing of the system requires one license per node from one of three different licensing editions; Edge licenses, Standard licenses, or Enterprise licenses. The features available with a Cisco HyperFlex Edge license include:

- HyperFlex Edge deployments at single site or multiple sites
- Fixed scale HX220c SFF model servers (2-, 3-, 4-node configuration)

- Hybrid or All-Flash servers
- 1 GbE or 10 GbE Ethernet without use of Cisco UCS Fabric Interconnects
- In line Compression and Deduplication
- HyperFlex native snapshots
- Rapid Clones
- Asynchronous Replication for Disaster Recovery
- Management via vCenter plugin, HyperFlex Connect, or Cisco Intersight
- REST API

Cisco Intersight Licensing

Two licensing editions are available for Cisco Intersight customers: Base edition and Essentials edition.



The Base edition is available at no additional cost to customers and automatically included with every Cisco UCS and Cisco HyperFlex system purchase. Cisco Intersight Base edition provides access to a portal that delivers centralized monitoring and basic inventory of managed systems, organizational capabilities including tagging and search, and the capability to launch native endpoint management interfaces including Cisco UCS Manager and HyperFlex Connect. Starting release 4.0(1a), Cisco Intersight provides comprehensive lifecycle management and includes remote cloud-based installation, and invisible witnessing of HyperFlex Edge clusters in 1GE and 10GE networking topology options. This release extends support to 2-Node Edge clusters to run HyperFlex in environments requiring a small footprint and 4-Node Edge clusters to enable scaling-up HyperFlex Edge clusters.

The Essentials edition provides extra value in addition to what has been provided in the Base edition. It includes features to enable the customers to centralize configuration management through a unified policy engine, determine compliance with the Cisco UCS Hardware Compatibility List (HCL), and initiate firmware updates. The Essentials edition provides a single interface for monitoring, management, and operations, with the capability to launch the virtual keyboard, video, and mouse (vKVM) console directly from Cisco Intersight. Cisco Intersight Essentials edition can be ordered in one-year, three-

year, and five-year subscription periods. You can start a free 90-day evaluation trial of the Essentials edition by signing up from the licensing page in Intersight.

HyperFlex cloud deployment, HyperFlex dashboards, day two monitoring, centralized alarms, HyperFlex cluster list view, vKVM launch (HyperFlex PID servers only), and cross launch to HyperFlex Connect are all included in the Base edition. If the only requirements are to install and upgrade the HyperFlex Data Platform software, those activities are supported using the Basic edition license. For full stack HyperFlex upgrades, including chassis firmware and ESXi hypervisor updates, the HCL feature in the Essentials edition can be purchased and used to audit compliance between hardware, OS, driver, and component firmware versions.

For the latest licensing information, refer to the [Licensing Requirements](#) in the Cisco Intersight Help Center.

Considerations

Before installing the HyperFlex Edge system, the following factors are important to be considered regarding the functionality and usable capacity required to build your cluster.

Hypervisors

The standard Cisco HyperFlex cluster can be built in VMware ESXi, Microsoft Hyper-V and Kubernetes Container environment; however, Cisco HyperFlex Edge clusters are only supported using the VMware ESXi hypervisor.

vCenter Server

It is highly recommended that the VMware vCenter server which will manage the new Cisco HyperFlex cluster, has been installed and operational prior to the installation of the Cisco HyperFlex HX Data Platform software. However, using a new VMware vCenter server that will be deployed within the Cisco HyperFlex cluster you plan to install is also supported using a different method. The following best practice guidance applies to installations of HyperFlex clusters:

- Do not modify the default TCP port settings of the vCenter installation. Using non-standard ports can lead to failures during the installation.
- It is recommended to build the vCenter server on a physical server or in a virtual environment outside of the HyperFlex cluster. Building the vCenter server as a virtual machine inside the HyperFlex cluster environment is highly discouraged. There is a tech note for multiple methods of deployment if no external vCenter server is already available:
http://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HyperFlex_HX_DataPlatformSoftware/TechNotes/Nested_vcenter_on_hyperflex.html



Note: This document does not explain the installation and configuration of VMware vCenter Server for Windows, or the vCenter Server Appliance.

Version Control

The software revisions listed in Table 5 are the only valid and supported configuration at the time of the publishing of this validated design. Special care must be taken not to alter the revision of the hypervisor, vCenter server, Cisco HX platform software, or the Cisco UCS firmware without first consulting the appropriate release notes and compatibility matrixes to ensure that the system is not being modified into an unsupported configuration.

Scale

A HyperFlex Edge system has a cluster size of 2, 3 or 4 converged HX220 hybrid or all-flash nodes. Unlike a standard HyperFlex cluster in the datacenter, Edge clusters do not support Cisco UCS compute-only nodes and the cluster cannot be expanded with additional converged nodes in the current release.

Capacity

Overall usable HyperFlex 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. In addition, configuring a cluster as a stretched cluster across two sites modifies the data distribution method, which reduces capacity in favor of data availability. Caching disk sizes are not calculated as part of 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 120×10^9 bytes of usable addressable capacity, or 120 billion bytes. However, many operating systems and filesystems report their space based on standard computer binary exponentiation, or calculation by powers of 2, also called binary prefix. In this example, 2^{10} or 1024 bytes make up a kilobyte, 2^{10} kilobytes make up a megabyte, 2^{10} megabytes make up a gigabyte, and 2^{10} gigabytes make up a terabyte. As the values increase, the disparity between the two systems of measurement and notation get worse, at the terabyte level, the deviation between a decimal prefix value and a binary prefix value is nearly 10 percent.

The International System of Units (SI) defines values and decimal prefix by powers of 10 as are listed in Table 6.

Table 6 SI Unit Values (Decimal Prefix)

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

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

Table 7 IEC Unit Values (Binary Prefix)

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

For the purpose of this document, the decimal prefix numbers are used only for raw disk capacity as listed by the respective manufacturers. For all calculations where raw or usable capacities are shown from the perspective of the HyperFlex Data Platform software, 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 Cisco Intersight, the HyperFlex vCenter Web Plugin and HyperFlex Connect GUI when viewing cluster capacity, allocation and consumption.

Appendix A: HyperFlex Cluster Capacity Calculations describes how to calculate the usable capacity of the HyperFlex Cluster. The HyperFlex tool to help with sizing is listed in **Appendix B: HyperFlex Sizer**.

Table 8 lists a set of HyperFlex M5 edge cluster usable capacity values, using binary prefix, for an array of cluster configurations. The binary prefix numbers are used for all calculations where raw or usable capacities are shown from the perspective of the HyperFlex software, filesystems or operating systems, which is different from the decimal prefix numbers used by manufacturers for raw disk capacity. This is done primarily to show a consistent set of values as seen by the end user from within the HyperFlex vCenter Web Plugin and HyperFlex Connect GUI when viewing cluster capacity, allocation and consumption, and also within most operating systems. These values are useful for determining the appropriate size of HyperFlex cluster to initially purchase, and how much capacity can be gained by adding capacity disks.

Table 8 HyperFlex Edge Cluster Usable Capacities

HX-Series Server Model	Node Quantity	Capacity Disk Size (each)	Capacity Disk Quantity (per node)	Cluster Usable Capacity at RF=2	
HXAF-E-220M5SX	2	3.8 TB	3	9.7 TiB	
			8	25.7 TiB	
		960 GB	3	2.4 TiB	
			8	6.4 TiB	
	3	3.8 TB	3	14.5 TiB	
			8	38.6 TiB	
		960 GB	3	3.6 TiB	
			8	9.6 TiB	
	4	3.8 TB	3	19.3 TiB	
			8	51.5 TiB	
		960 GB	3	4.8 TiB	
			8	12.8 TiB	
HX-E-220M5SX	2	2.4 TB	3	6.0 TiB	
			8	16.0 TiB	
		1.8 TB	3	4.5 TiB	
			8	12.0 TiB	
		1.2 TB	3	3.0 TiB	
			8	8.0 TiB	
		3	2.4 TB	3	9.0 TiB
				8	24.2 TiB
	1.8 TB		3	6.8 TiB	
			8	18.2 TiB	
	1.2 TB		3	4.5 TiB	
			8	12.1 TiB	
	4	2.4 TB	3	12.0 TiB	
			8	32.2 TiB	
		1.8 TB	3	9.0 TiB	
			8	24.2 TiB	

HX-Series Server Model	Node Quantity	Capacity Disk Size (each)	Capacity Disk Quantity (per node)	Cluster Usable Capacity at RF=2
		1.2 TB	3	6.0 TiB
			8	16.1 TiB

Data Protection

Cisco HyperFlex supports data protection features including snapshot-based VM level replication between two HyperFlex clusters. Replication can be used to migrate or recover a single VM in the secondary HyperFlex cluster, groups of VMs can be coordinated and recovered, or all VMs can be recovered as part of a disaster recovery scenario. In order to start using replication, two HyperFlex clusters must be installed and have network connectivity between them. It is possible to replicate between hybrid and all-flash clusters. To avoid complications with duplicate VM IDs, it is recommended that the two replicating HyperFlex clusters be managed by two different VMware vCenter servers.

To enable replication, the replication networking must first be configured in HyperFlex Connect. Once the networking configuration work is completed for both clusters that will replicate to each other, a partnership, or pairing between the two clusters is established. After this replication pair is established, VMs can be protected individually, or they can be placed into protection groups, which are created to protect multiple VMs with the same replication settings. Two paired HyperFlex clusters can replicate VMs in both directions, therefore the replication status of all VMs and Protection Groups, incoming and outgoing, are presented in the replication views of both clusters. Virtual machines can be configured for protection, i.e. replication, individually, or be placed into a Protection Group. The protection settings that can be configured on an individual VM are the same as the settings that are configured for a protection group. In most cases, it is easier to configure multiple Protection Groups, each with the settings that are required, and then add VMs to those groups.



Note: HyperFlex native replication is not available for use with 2-node Edge clusters as of the time of the publication of this document.

Cisco HyperFlex Edge system can also be protected with the third-party data protection software, which is not explained in this document.

Network Design

HyperFlex Networking

The Cisco HyperFlex system has communication pathways that fall into four defined zones:

- **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 HyperFlex 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:
 - Cisco HX-series rack-mount server CIMC management interfaces.
 - ESXi host management interfaces.
 - Storage Controller Virtual Machine management interfaces.
 - A roaming HyperFlex cluster management interface.
 - Storage Controller Virtual Machine replication interfaces.

- A roaming HyperFlex cluster replication interface.
- **VM Zone:** This zone comprises the connections needed to service network IO to the guest virtual machines that will run inside the HyperFlex hyperconverged system. This zone typically contains multiple VLANs, that are trunked to the Cisco HX-series servers from the upstream switches 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 HyperFlex system, throughout the LAN/WAN.
- **Storage Zone:** This zone comprises the connections used by the Cisco HX Data Platform software, ESXi hosts, and the storage controller virtual machines to service the HyperFlex 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, in a dual switch configuration this traffic all traverses a single network switch, however there are hardware failure scenarios where this traffic would need to traverse the network from one upstream switch to the other. For that reason, the VLAN used for HyperFlex storage traffic must be able to traverse from one switch to the other in a dual switch configuration, reaching switch A from switch B, and vice-versa. This zone can be configured to carry jumbo frame traffic for higher performance, and in such a scenario jumbo frames must be enabled on all the interfaces. In this zone are multiple components:
 - A VMkernel interface used for storage traffic on each ESXi host in the HyperFlex cluster.
 - Storage Controller Virtual Machine storage interfaces.
 - A roaming HyperFlex cluster storage interface.
- **VMotion Zone:** This zone comprises the connections used by the ESXi hosts to enable vMotion of the guest virtual machines from host to host. During normal operation, in a dual switch configuration this traffic all traverses a single network switch, however there are hardware failure scenarios where this traffic would need to traverse the network from one upstream switch to the other. For that reason, the VLAN used for HyperFlex VMotion traffic must be able to traverse from one switch to the other in a dual switch configuration, reaching switch A from switch B, and vice-versa. This zone can be configured to carry jumbo frame traffic for higher performance, and in such a scenario jumbo frames must be enabled on all the interfaces.

This method of zoning provides separate networks for different type of traffic with special purpose and avoids multiplexing all traffic onto a single network. As a result, the HyperFlex installer creates four pre-defined virtual networks on all HyperFlex hosts at the ESXi hypervisor level. Four different virtual switches are created by the HyperFlex installer, which are each serviced by a virtual NIC (vmnic) created from the physical adapters or from the virtual adapters hosted on the Cisco VIC card. The vSwitches created are:

- **vswitch-hx-inband-mgmt:** This is the default vSwitch which is renamed by the ESXi kickstart file as part of the automated installation. The default VMkernel port, vmk0, is configured in the standard Management Network port group. A second port group is created for the Storage Platform Controller VMs to connect to with their individual management interfaces. A third port group is created for cluster to cluster VM snapshot replication traffic.
- **vswitch-hx-storage-data:** This vSwitch is created as part of the automated installation. A VMkernel port, vmk1, is configured in the Storage Hypervisor Data Network port group, which is the interface used for connectivity to the HyperFlex Datastores via NFS. Enabling Jumbo frames on the uplinks for this switch is highly recommended. A second port group is created for the Storage Platform Controller VMs to connect to with their individual storage interfaces.
- **vswitch-hx-vm-network:** This vSwitch is created as part of the automated installation. One port group needs to be created for each guest VM network. By default, none guest VM port group is created during HyperFlex Edge cluster installation. You need to manually create the port groups for your VM networks or create them using the post_install script.

- vmotion:** This vSwitch is created as part of the automated installation. The IP addresses of the VMkernel ports (vmk2) can be configured manually or during the post_install script execution. Enabling Jumbo frames on the uplinks for this switch is highly recommended.

The HyperFlex installer will create four vSwitches described as the above for 10GE HyperFlex Edge system, each using two uplinks which are serviced from the virtual adapters hosted on the Cisco VIC 1457 card.

For 1GE HyperFlex Edge cluster, the HyperFlex installer will create four vSwitches but only the first two vSwitches are utilized: vswitch-hx-inband-mgmt and vswitch-hx-storage-data. Depending on the variable topologies (single-switch or dual-switch; 2-node, 3-node or 4-node), each of these two vSwitches might use one or two uplinks that are serviced by virtual NICs (vmnic) created from LOM ports or from the physical adapters.

The tables in the [Logical Topology](#) section give more details into the ESXi virtual networking design as built by the HyperFlex Edge installer within Cisco Intersight by default.

VLANs and Subnets

Without Fabric Interconnects Cisco HyperFlex Edge servers are administrated directly from Cisco Integrated Management Controller (CIMC) Software. The Cisco Integrated Management Controller (CIMC) is a baseboard management controller that provides embedded server management for Cisco HyperFlex rack servers and the configuration of CIMC fits into the management zone. By registering a device connector embedded in CIMC, HyperFlex Edge servers can be easily and remotely managed with Cisco Intersight. The CIMC VLAN and the HyperFlex management VLAN are required to have outbound Internet access to reach Cisco Intersight.

Based on the design of HyperFlex networking plus CIMC communication required for server management, multiple VLANs are recommended to use for the HyperFlex Edge system configuration:

VLAN Function	VLAN ID	Purpose	Notes
CIMC Traffic	Customer supplied	Cisco Integrated Management Controller interfaces for UCS standalone servers	This VLAN must be routable and must have access to Intersight.
ESXi and HyperFlex Management Traffic	Customer supplied	ESXi host management interfaces HX Storage Controller VM management interfaces HX Storage Cluster roaming management interface	This VLAN must be routable and must have access to Intersight. It can be same or different from the CIMC VLAN.
HyperFlex Storage Traffic	Customer supplied	ESXi host storage VMkernel interfaces HX Storage Controller storage network interfaces HX Storage Cluster roaming storage interface	This VLAN is used for storage traffic. Only L2 connectivity is required. It may not be combined with the management VLAN and must be dedicated for storage traffic.
vMotion Traffic	Customer supplied	ESXi host vMotion VMkernel interfaces	This VLAN can be the same as the management VLAN but it is recommended to segregate this traffic onto a unique VLAN.

VLAN Function	VLAN ID	Purpose	Notes
Guest VM Traffic	Customer supplied	Networks for Guest Virtual Machines	There can be multiple VLAN port groups created for different applications.

Jumbo Frames

The Cisco UCS best practice is to use standard ethernet frames MTU 1500 for the external management network, and to enable Jumbo Frames MTU 9000 for any Storage facing networks. 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. Therefore, 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-vmotion VLAN is also set to use jumbo frames. This guideline, however, is not valid for HyperFlex Edge clusters. By default, Jumbo frames are not required for the edge system. All traffic traversing the VLANs and subnets on the HyperFlex Edge system is configured by default to use standard ethernet frames (MTU 1500). Jumbo frames can be enabled for data traffic and vMotion traffic in the HyperFlex Edge cluster if all types of workloads in the Virtual Machine zone allow it.

While enabling the Jumbo frames in the edge system, it means that the uplinks for Cisco HyperFlex servers must be configured to pass jumbo frames. Failure to configure the uplink switches to allow jumbo frames can lead to service interruptions during some failure scenarios. So that Cisco recommends that this configuration only be used in environments where the uplink switches are capable of passing jumbo frames, and that jumbo frames be enabled through the whole pathways.

QoS System Classes

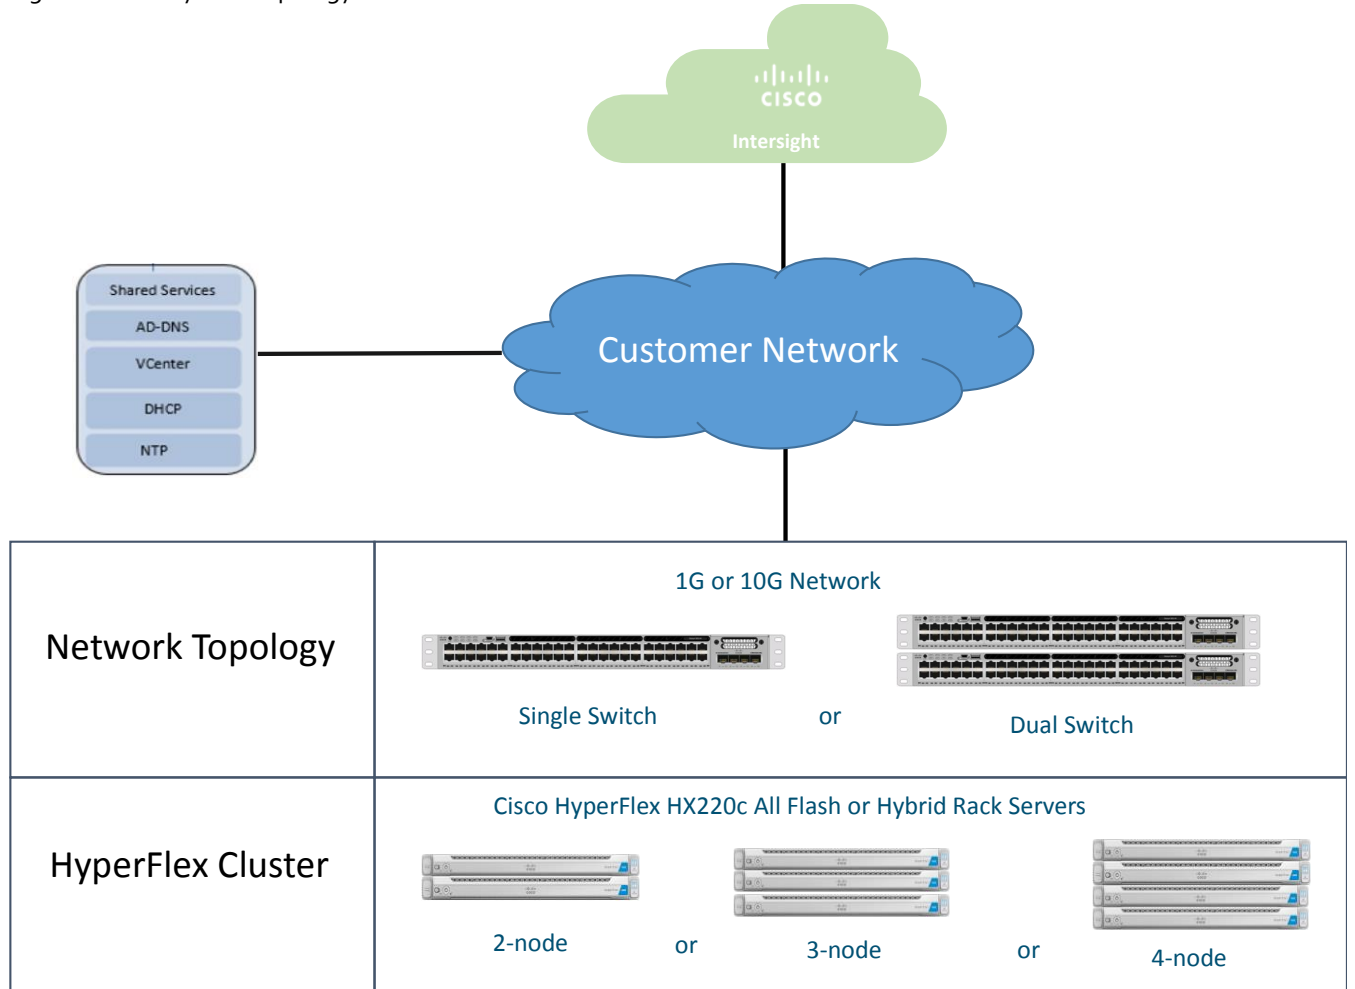
Currently the HyperFlex Edge system does not support QoS services.

Physical Topology

Topology Overview

A Cisco HyperFlex Edge cluster is built using Cisco HX-Series rack-mount servers without connecting them to Cisco UCS Fabric Interconnects. Upstream network connections, also referred to as "northbound" network connections, are made directly from the servers to the customer chosen data center top-of-rack (ToR) switches at the time of installation.

Figure 12 Physical Topology Overview



In traditional Cisco UCS deployments, the Cisco UCS Fabric Interconnects provide the management and configuration services for the connected rack-mount servers and blade servers attached to them. Without Fabric Interconnects, Cisco Intersight communicates with the servers through a device connector that is embedded in the server’s firmware. The device connector can be enabled from Cisco Integrated Management Controller (CIMC) Software. Once connected to Cisco Intersight, the Cisco Integrated Management Controller (CIMC) gives Intersight full control of the server, allowing for complete configuration and management via the cloud. The CIMC of each server can be configured to operate in one of two modes; Dedicated Mode or Shared Mode. Dedicated Mode uses a dedicated 1GbE management LAN-on-motherboard (LOM) port on the rear of the server. In Shared Mode, any LOM port or VIC adapter card port can be used to access the CIMC instead of using the dedicated management LOM port. Internally, shared modes redirect the CIMC communication to ingress/egress using the configured LOM or Cisco VIC interfaces. This method reduces cabling and port count requirements in the system, however it does require a different configuration of the CIMC itself, and also in the upstream connected switches. The Shared modes include Shared LOM, Shared LOM 10GE, Cisco VIC Card, and Shared LOM Extended modes.

If there are plenty of network ports on the upstream switch or available on an out-of-band management switch, the dedicated CIMC management port can be used. However, dedicated mode consumes more switch ports, requires extra cabling, additional power usage, and additional configuration. Depending on the availability of the networking resources in the user environment, Shared LOM, Cisco VIC or Shared LOM Extended mode can simplify the configuration and also reduces switch ports and cabling. In shared LOM mode, all host ports must be able to pass the same management subnet by allowing the appropriate VLAN.

The topology of Cisco HyperFlex Edge system is straightforward. It is composed of the fixed amount of two, three or four HX-Series 220c rack-mount servers (hybrid or all-flush) per cluster. Depending on the speed of the customer network,

different network adapters are required to be used on the Cisco HyperFlex Edge servers. There are two types of connection that are supported – 1G or 10G Ethernet. For the 1GE connection, the HX-Series converged servers are connected directly to the customer’s network via onboard LOM ports or with the optional Intel i350-T4 Quad-Port PCIe NIC Card. For the 10GE connection, the HX-Series converged servers are connected directly to the customer’s network via Cisco VIC 1457 Quad-Port unified virtual interface card.

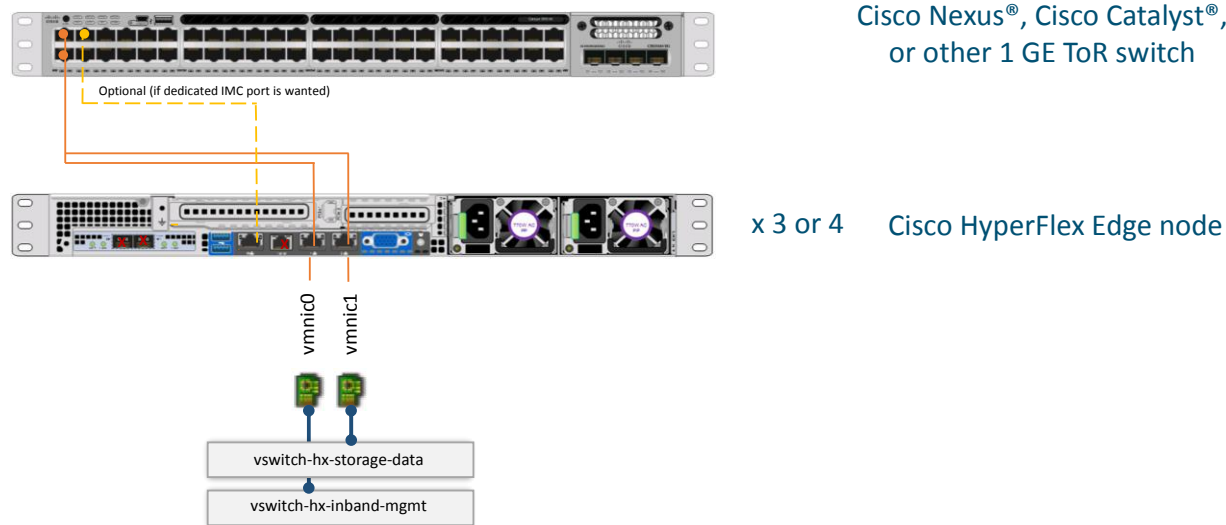
Taking into the consideration of the availability of single switch or dual switches at the customer site, the following six physical topologies between the Cisco HX-series servers and upstream network are supported and will be depicted in this document. The choice depends on the requirements and the availability of hardware or networking resources.

- 1GE single-switch connectivity for HyperFlex Edge 3- or 4-node cluster
- 1GE dual-switch connectivity for HyperFlex Edge 3- or 4-node cluster
- 1GE single-switch connectivity for HyperFlex Edge 2-node cluster
- 1GE dual-switch connectivity for HyperFlex Edge 2-node cluster
- 10GE single-switch connectivity for HyperFlex Edge clusters
- 10GE dual-switch connectivity for HyperFlex Edge clusters

1GE Single-Switch Connectivity for HyperFlex Edge 3- or 4-Node Cluster

Single switch configurations provide a simple topology that requires only a single switch, and minimum two 1GbE ports per server. Both link and switch redundancy are not provided in this basic topology. Access ports or trunk ports are the two supported network port configurations.

Figure 13 1GE Single-Switch Topology for HyperFlex Edge 3- or 4-Node Cluster



The single switch topology requires at least two separate networks: Management network (includes guest VM network and vMotion traffic) and Data network (for storage traffic).

Two 1GbE ports on each server are required with one of these two ports dedicated for HyperFlex storage traffic:

- Port 1 — management (ESXi and CIMC), vMotion traffic, and VM guest traffic
- Port 2 — HyperFlex storage traffic

If desired, use an additional dedicated CIMC port connecting to the same switch or to an out-of-band management switch.

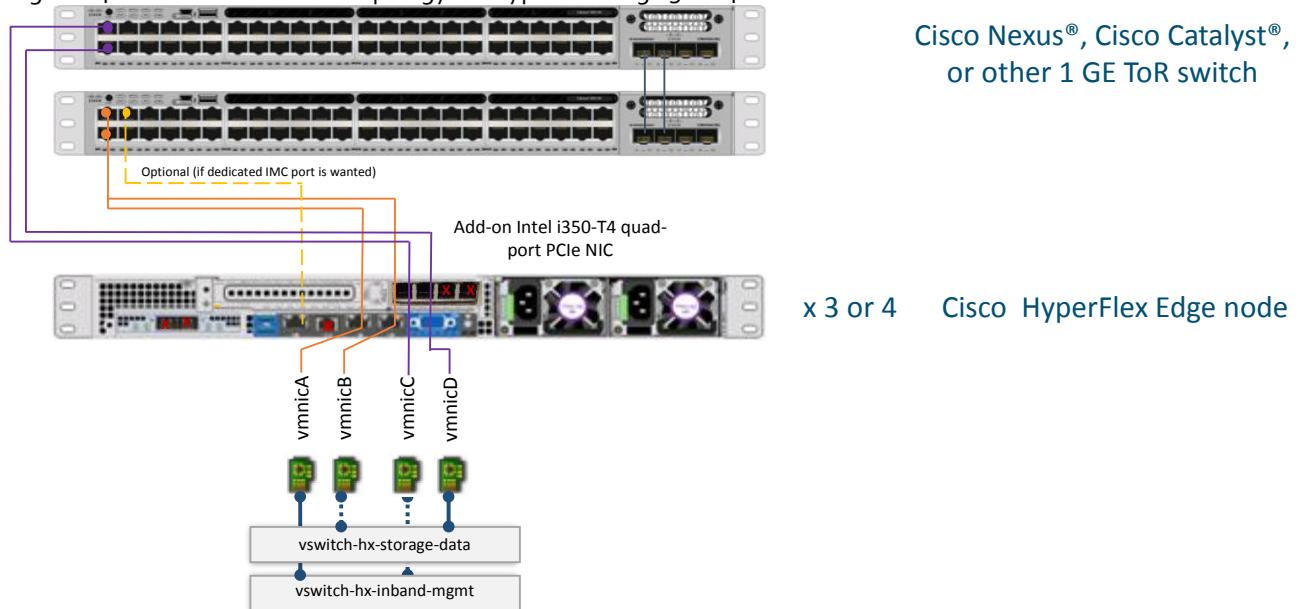
The requirements of the upstream network:

- A managed switch with VLAN capability
- Minimum 6 physical 1GbE ports for three HyperFlex nodes or 8 ports for four HyperFlex nodes provided that inband CIMC is configured on the 1GbE LOM ports in Shared Mode. Users may optionally configure out of band CIMC using a dedicated port on the LOM card, which will require an additional 1GbE switch port per server.
- Jumbo frames are not required
- Spanning tree portfast (access ports) or portfast trunk (trunk ports) must be enabled for all network ports connected to HyperFlex servers for uninterrupted CIMC access.

1GE Dual-Switch Connectivity for HyperFlex Edge 3- or 4-Node Cluster

Dual switch configurations provide a slightly more complex topology with full redundancy that protects against switch failure, link and port failure, and LOM or PCIe NIC HW failures. This topology requires two switches that may be standalone or stacked, using four 1GbE ports per server via the addition of an Intel i350 PCIe NIC card per server. Trunk ports are the only supported network port configuration.

Figure 14 1GE Dual-Switch Topology for HyperFlex Edge 3- or 4-Node Cluster



The dual switch topology also requires two separate networks: Management network (includes guest VM network) and Data network (for storage and vMotion traffic). Four 1GbE ports on each server are required, two ports from the onboard LOM card and two ports from a PCIe add-in NIC card. One LOM and one PCIe port serve management and VM guest traffic in a redundant configuration while second LOM and second PCIe port serve storage data and vMotion traffic in a redundant and load balanced configuration.

- Onboard LOM Port 1 — management (ESXi, HyperFlex controller, and CIMC) and VM guest traffic
- Onboard LOM Port 2 — vMotion traffic (and storage standby)
- Intel i350 PCIe NIC Port 1 — HyperFlex storage traffic (and vMotion standby)
- Intel i350 PCIe NIC Port 2 — VM guest traffic (and management standby)



Note: Do not connect more than two ports on the Intel i350 PCIe NIC prior to cluster installation. After cluster deployment, using the additional two ports is optional and allowed.

The requirements of the upstream network:

- Two managed switches with VLAN capability
- Minimum 12 physical 1GbE ports for three HyperFlex nodes or 16 ports for four HyperFlex nodes provided that inband CIMC is configured on the 1GbE LOM ports in Shared Mode. Users may optionally configure out of band CIMC using dedicated port on the LOM card, which will require an additional 1GbE switch port per server.
- Jumbo frames are not required
- Portfast trunk should be configured for uninterrupted CIMC access

1GE Single-Switch Connectivity for HyperFlex Edge 2-Node Cluster

2-node HyperFlex Edge systems are supported from HXDP version 4.0 onward, and it has a new network design. This design provides a fully redundant topology that protects against switch, link and port failures when using dual or stacked switches. With the single switch configuration, the design provides improved redundancy for the cluster but without protection against switch failure.

Cisco IMC Connectivity for the 2-node 1 Gigabit Ethernet (GE) topology requires the use of the dedicated 1GE Cisco IMC management port. Other operating modes, including shared LOM mode, are not available due to the use of direct connect cables in this topology.

Both the single-switch and dual-switch topologies require two separate networks: Management network (includes guest VM network) and Data network (for storage and vMotion traffic). Two 1GbE ports from the add-in Intel i350 PCIe NIC card and two 10GbE Lan-on-motherboard (LOM) ports on each server are required. The direct 10GE connectivity between two ports from the onboard LOM card between the two nodes, without the need for a 10GE capable switch, provides high speed and full redundancy for the storage data and vMotion traffic. Two ports from a PCIe add-in NIC card serve management and VM guest traffic in a redundant configuration.

- Onboard 10GbE LOM Port 1 — HyperFlex storage traffic (and vMotion standby)
- Onboard 10GbE LOM Port 2 — vMotion traffic (and storage standby)
- Intel i350 PCIe NIC Port 1 — Management (ESXi, HyperFlex controller, and CIMC) and VM guest traffic
- Intel i350 PCIe NIC Port 2 — VM guest traffic (and management standby)



Note: Do not connect more than two ports on the Intel i350 PCIe NIC prior to cluster installation. After the cluster deployment, using the additional two ports is optional and allowed.

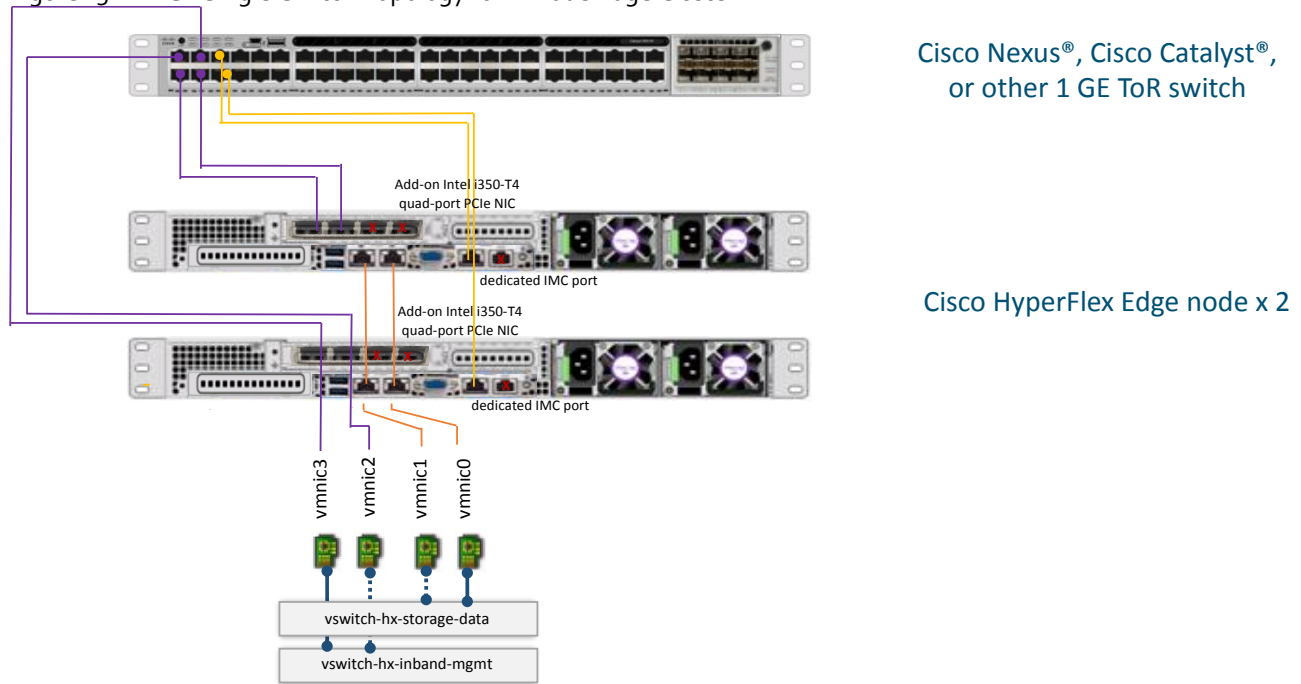
The following requirements for the HyperFlex Edge nodes are common to both topologies and must be met before starting deployment:

- Dedicated 1GE Cisco IMC management port per server is required
- Intel i350 Quad Port PCIe NIC Card installed in a PCIe slot in each server is required
- Cisco VIC is not used in this 1GE topology
- Two 10GE Direct Connect LAN-on-Motherboard (LOM) connections are required

The single switch configuration provides a simple topology that requires only a single switch, such that switch redundancy is not provided. Here are the requirements of the upstream network:

- A managed switch with VLAN capability
- Six physical 1GE switch ports are required as inband CIMC is configured on one 1GE LOM port in Dedicated Mode on each server.
- Switchports connected to the dedicated Cisco IMC management ports are configured in Access Mode on the appropriate VLAN.
- Switchports connected to the Intel i350 ports are configured in Trunk Mode with the appropriate VLANs allowed to pass. Portfast trunk should be configured for these ports.
- Jumbo frames are not required to be configured on the ToR switch.

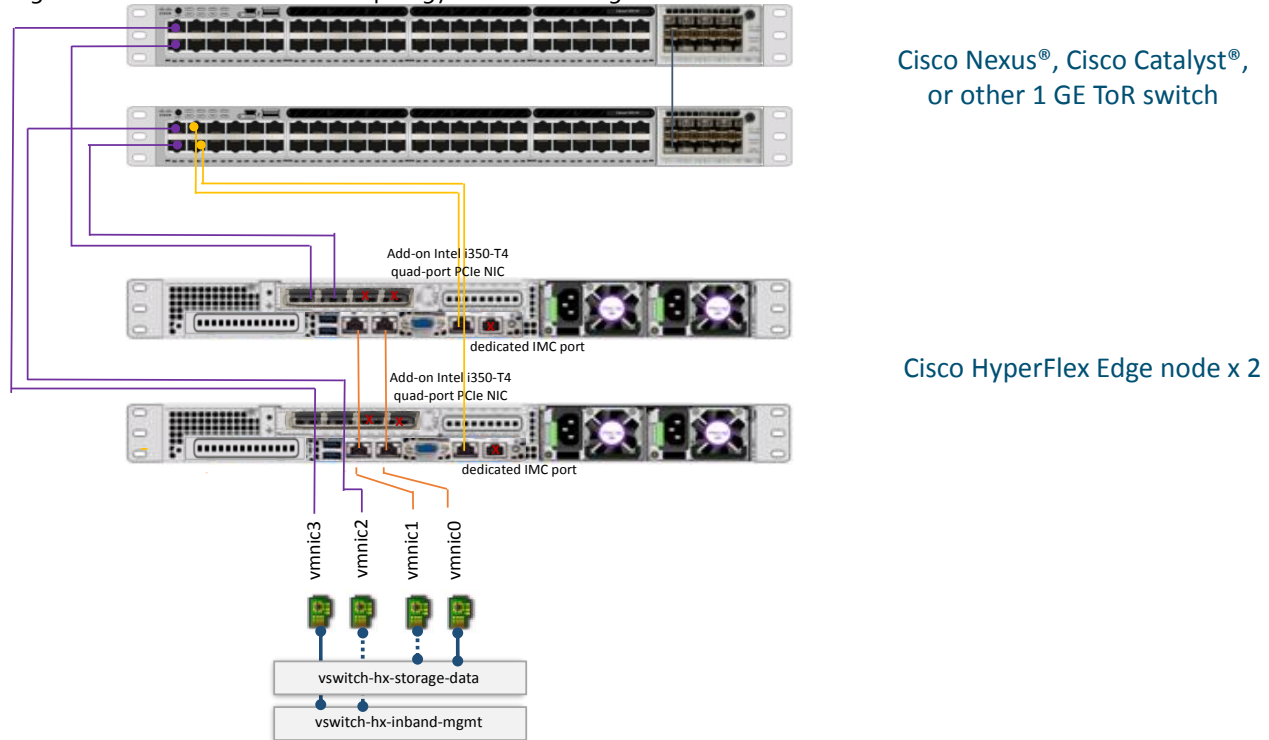
Figure 15 1GE Single-Switch Topology for 2-Node Edge Cluster



1GE Dual-Switch Connectivity for HyperFlex Edge 2-Node Cluster

The dual switch configuration provides a topology that requires two Ethernet ToR switches, such that switch redundancy is also provided in addition to link and port redundancy. The requirements of the upstream network are the same as the single-switch topology except that two managed switches with VLAN capability are required in this topology. The connection to the two ports from the Intel i350 PCIe NIC card on each server goes to one switch port on each of the two ToR switches.

Figure 16 1GE Dual-Switch Topology for 2-Node Edge Cluster



Note: Cisco highly recommends the 10GE topology for higher performance and future node expansion capabilities. 1GE topologies will not be supported for future expansion.

10GE Single-Switch Connectivity for HyperFlex Edge Clusters

10GE HyperFlex Edge system have been supported since HXDP version 3.5 and requires installing Cisco Virtual Interface Cards (VICs) in each HyperFlex server. This design provides a fully redundant topology that protects against switch, link and port failures when using dual or stacked switches. With the single switch configuration, the topology provides no protection against switch failure.

There are two options for Cisco IMC Connectivity supported for the 10GE HyperFlex Edge system.

- Use of a dedicated 1GE Cisco IMC management port: this option requires additional switch ports and cables; however, it avoids network contention and ensures always on, out of band access to each physical server.
- Use of shared LOM extended mode (EXT). In this mode, single wire management is used, and Cisco IMC traffic is multiplexed onto the 10GE VIC connections. When operating in this mode, multiple streams of traffic are shared on the same physical link. The uninterrupted reachability to CIMC is not guaranteed by QoS that is not enforced for Cisco HyperFlex Edge systems. This deployment option is not recommended and is not explained in this document.

Both the single-switch and dual-switch topologies for 10GE HyperFlex clusters require the same four separate networks as the standard HyperFlex clusters: Management traffic network, Data traffic network, vMotion network, and guest VM network.

Any two 10GE ports from Cisco VIC 1457 adapter on each server are required. Four virtual NIC interfaces are created from each port and are assigned to the four virtual networks in a HyperFlex system serving different types of the traffic.

- Cisco VIC 1457 Port 1
 - VMNIC2 – Management (ESXi, HyperFlex controller, and CIMC) traffic

- VMNIC4 – Storage Traffic (Standby)
- VMNIC6 – VM guest traffic
- VMNIC8 – vMotion traffic
- Cisco VIC 1457 Port 2
 - VMNIC3 – Management (ESXi, HyperFlex controller, and CIMC) traffic (Standby)
 - VMNIC5 – Storage Traffic
 - VMNIC7 – VM guest traffic
 - VMNIC9 – vMotion traffic (Standby)



Note: Do not connect additional 10GE ports on the Cisco VIC 1457 prior to cluster installation. After cluster deployment, using the additional two 10GE ports for guest VM traffic is optional and is allowed.

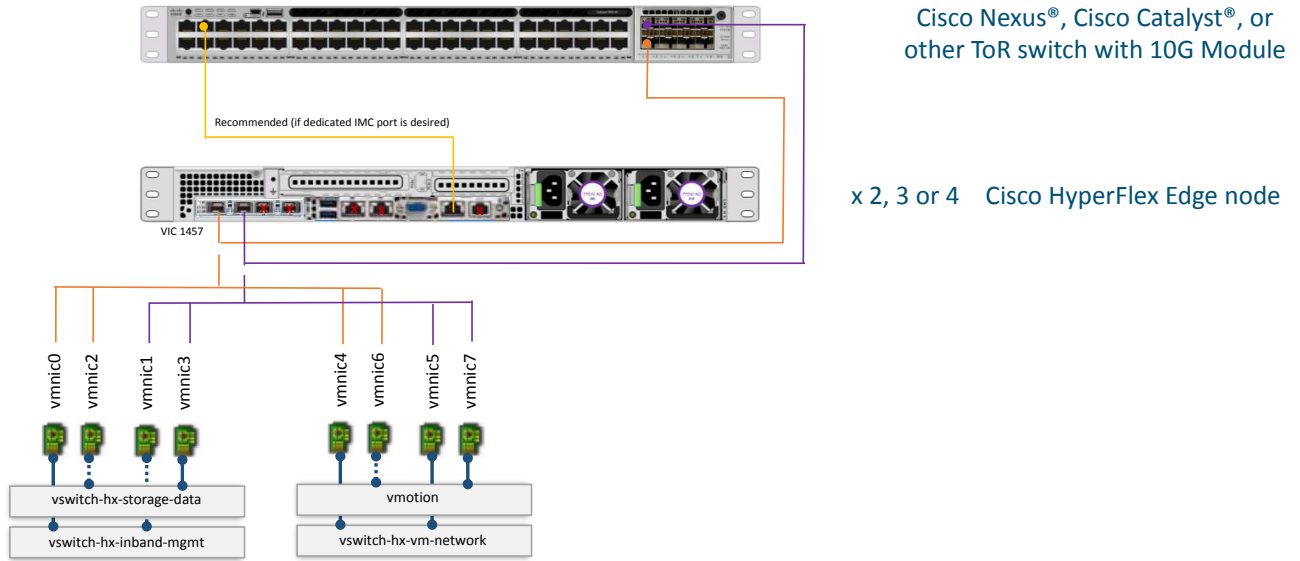
The following requirements for the HyperFlex Edge nodes are common to both topologies and must be met before starting deployment:

- Cisco VIC 1457 (installed in the MLOM slot in each server) (recommended)
- Dedicated 1 Gigabit Ethernet (GbE) Cisco IMC management port per server (recommended)
- Only 10GbE speed is supported (not 1GbE, 25GbE, or 40GbE) in this topology

Single switch configuration provides a simple topology that requires only a single switch, such that switch redundancy is not provided. Here are the requirements of the upstream network:

- A managed switch with VLAN capability
- One 1GE ToR switch port for dedicated Cisco IMC management port on each server.
- Two physical 10GE switch ports are required to connect to two VIC 1457 ports on each server.
- Switchports connected to the dedicated Cisco IMC management ports are configured in Access Mode on the appropriate VLAN.
- Switchports connected to the VIC 1457 ports are configured in Trunk Mode with the appropriate VLANs allowed to pass. Portfast trunk should be configured for these ports.
- Jumbo frames are not required to be configured on the ToR switch.

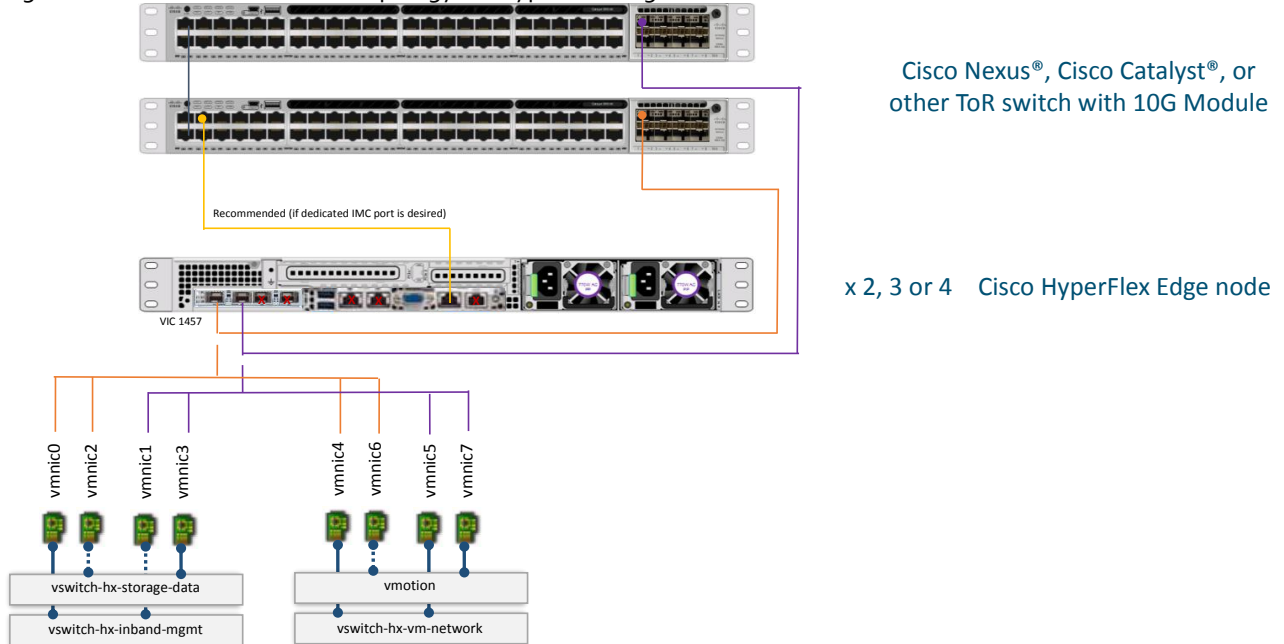
Figure 17 10GE Single-switch topology for HyperFlex Edge clusters



10GE Dual-Switch Connectivity for HyperFlex Edge Clusters

Dual-switch configuration provides a topology that requires two Ethernet ToR switches such that the switch redundancy is also provided in addition to the link and port redundancy. The requirements of the upstream network are the same as the single-switch topology except that two managed switches with VLAN capability are required in this topology. The connection to any two ports from the Cisco VIC 1457 adapter on each server goes to one 10GE switch port on each of the two ToR switches.

Figure 18 10GE Dual-Switch Topology for HyperFlex Edge Clusters



Note: The sample configuration presented in this document is based on the 10GE dual switch topology.

Logical Topology

Topology Overview

As described in the HyperFlex network design section, four vSwitches are required for creating a HyperFlex standard system:

- vswitch-hx-inband-mgmt — ESXi management (vmko), storage controller management network
- vswitch-hx-storage-data — ESXi storage interface (vmk1), HX storage controller data network
- vmotion — vMotion interface (vmk2)
- vswitch-hx-vm-network — VM guest portgroups

Depending on the network topology the customer chooses the vSwitches required for a HyperFlex Edge system will be different. With the multiple virtual interfaces provided by the innovative Cisco VIC technology, the installer will create four vSwitches described as the above on the 10GE HyperFlex Edge system. For 1GE HyperFlex Edge system, only the first two vSwitches are utilized: vswitch-hx-inband-mgmt and vswitch-hx-storage-data.

For the single switch configuration for the 1GE 3-node or 4-node HyperFlex Edge clusters, one LOM port is assigned to vswitch-hx-inband-mgmt, and another LOM port is assigned to vswitch-hx-storage-data. The data storage network requires a dedicated port in this configuration; therefore vswitch-hx-storage-data carries the storage data traffic and vswitch-hx-inband-mgmt carries the traffic for management, vMotion and guest virtual machines. There is no uplink redundancy for vSwitches in this configuration.

For the dual switch configuration of the 1GE 3-node or 4-node HyperFlex Edge clusters, redundancy occurs at the vSwitch level; one LOM port and one PCIe NIC port are assigned to vswitch-hx-inband-mgmt while the second LOM port and the second PCIe NIC port are assigned to vswitch-hx-storage-data. One difference from the single switch configuration is the placement of vmotion network. To better leverage the port resources, the vmotion VMkernel port is on the vswitch-hx-storage-data and uses the opposite failover order of the storage data network. In this way the uplinks for storage and the uplinks for vmotion are set as active/standby but each service will use opposite active links. The vswitch-hx-inband-mgmt is set to active/standby with the same failover order. By default, all network services will use the active path and failover to the standby path only when needed during a failure scenario. But the failover order for guest VM portgroups may be overridden as needed and to achieve better load balancing.

For the 1GE 2-node HyperFlex Edge cluster, redundancy occurs at the vSwitch level when two LOM ports are assigned to vswitch-hx-inband-mgmt while two PCIe NIC ports are assigned to vswitch-hx-storage-data. Both vSwitches are set for active/standby across the two uplinks. It is important to note that upstream network redundancy will only be possible in a dual-switch configuration, because in a single switch configuration all uplinks would be down during a switch failure. All services by default consume a single uplink port and failover when needed. Failover order for guest VM portgroups may be overridden as needed and to achieve better load balancing. In addition, the vmotion VMkernel port on the vswitch-hx-storage-data uses the opposite failover order of the storage data network. In this way the uplinks for storage and the uplinks for vmotion are set as active/standby but each service will use opposite active links. The vswitch-hx-inband-mgmt is set to active/standby with the same failover order for all port groups. By default, all network services will use the active path and failover to the standby path only when needed during a failure scenario.

For 1GE network topology, since only two vSwitches are used and assigned with physical vmnics, you can manually delete vswitch-hx-vm-network and vmotion vSwitches. Alternatively, run the post_install script and it will automatically remove vswitch-hx-vm-network and vmotion vSwitches from all clustered HyperFlex Edge nodes. The post_install script will also set the appropriate failure order automatically for the vmotion VMkernel port in the dual switch configuration.

For the 10GE HyperFlex Edge clusters, the redundancy occurs at the vSwitch level when eight virtual interfaces (VMNIC) are created from two Cisco VIC 1457 ports and assigned to the four vSwitches: vswitch-hx-inband-mgmt, vswitch-hx-storage-data, vswitch-hx-vm-network and vmotion. vswitch-hx-vm-network is set for active/active across the two uplinks while the other three vSwitches are set for active/standby across the two uplinks. All services by default consume a single uplink port

and failover when needed except for guest VM portgroups. Failover order for guest VM portgroups may be overridden as needed and to achieve better load balancing.

The following six logical topologies are explained in this document:

- 1GE single-switch logical networking for HyperFlex Edge 3- or 4-node cluster
- 1GE dual-switch logical networking for HyperFlex Edge 3- or 4-node cluster
- 1GE single-switch logical networking for HyperFlex Edge 2-node cluster
- 1GE dual-switch logical networking for HyperFlex Edge 2-node cluster
- 10GE single-switch logical networking for HyperFlex Edge clusters
- 10GE dual-switch logical networking for HyperFlex Edge clusters

1GE Single-Switch Logical Networking for HyperFlex Edge 3- or 4-Node Cluster

Figure 19 illustrates a logical view of network configuration for the 1GE single switch topology for HyperFlex Edge 3- or 4- node cluster.

Figure 19 1GE Single-Switch Logical Topology for HyperFlex Edge 3- or 4-Node Clusters

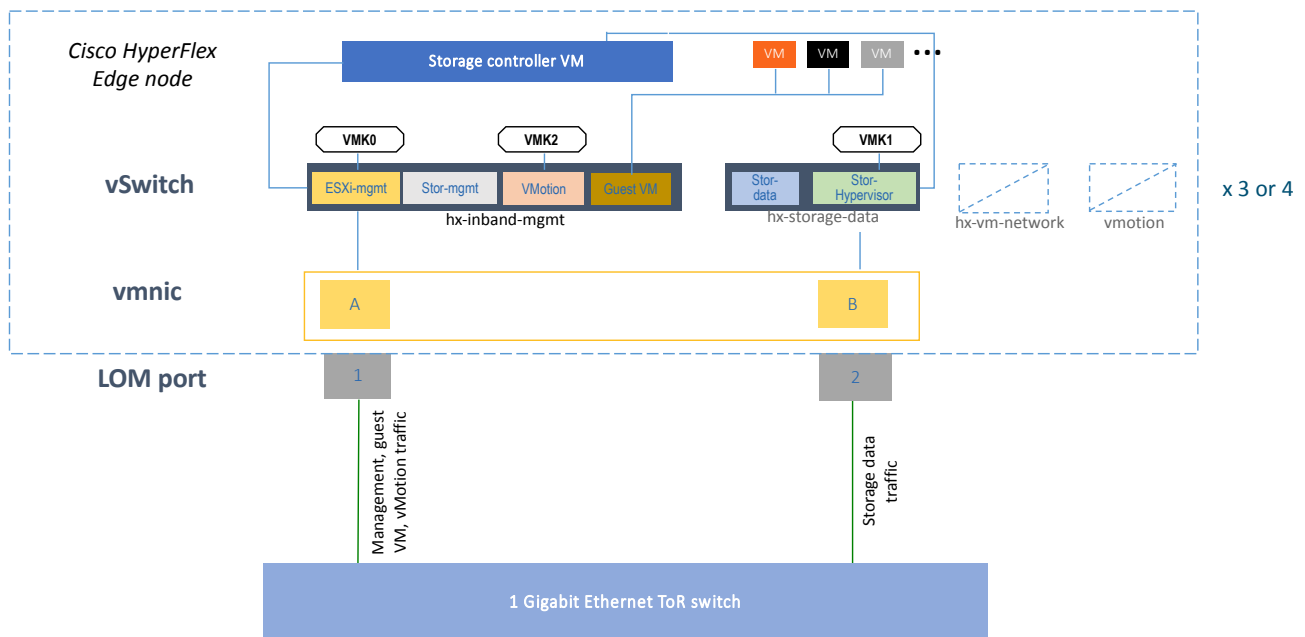


Table 9 gives the details into the ESXi virtual networks on the HyperFlex Edge 3- or 4-node cluster with 1GE single switch topology.

Table 9 Virtual Switches for the HX Edge 3- or 4-Node Cluster with 1GE Single-Switch Topology

Virtual Switch	Port Groups	VLAN IDs	Active vmnic(s)	Passive vmnic(s)	Jumbo
vswitch-hx-inband-mgmt	Management Network Storage Controller Management Network	<<hx-inband-mgmt>>	vmnico	NA	no

Virtual Switch	Port Groups	VLAN IDs	Active vmnic(s)	Passive vmnic(s)	Jumbo
	Storage Controller Replication Network	<<hx-inband-repl>>			
	vm-network-<<VLAN ID>>	<<vm-network>>			
	vmotion-<<VLAN ID>>	<<hx-vmotion>>			
vswitch-hx-storage-data	Storage Controller Data Network Storage Hypervisor Data Network	<<hx-storage-data>>	vmnic1	NA	no

1GE Dual-Switch Logical Networking for HyperFlex Edge 3- or 4-Node Cluster

Figure 20 illustrates a logical view of network configuration for the 1GE dual switch topology for HyperFlex Edge 3- or 4- node cluster.

Figure 20 1GE Dual-Switch Logical Topology for HyperFlex Edge 3- or 4-Node Clusters

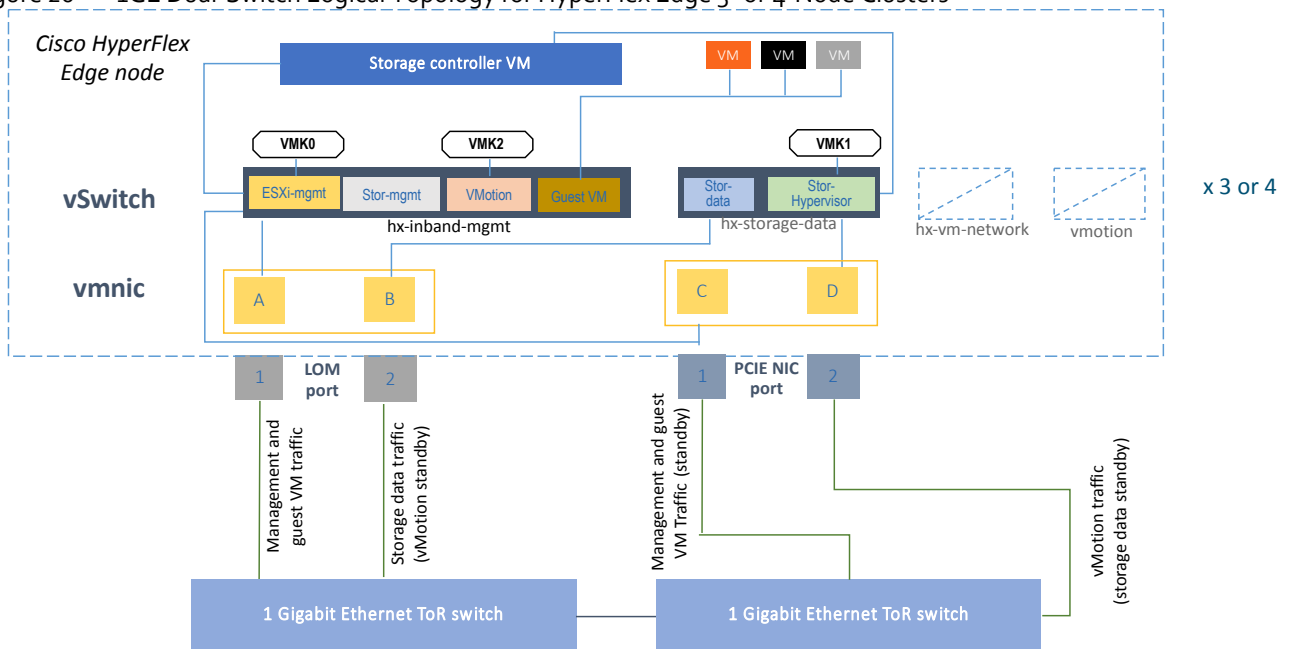


Table 10 lists the details into the ESXi virtual networks on the HyperFlex Edge 3- or 4-node cluster with 1GE dual switch topology.

Table 10 Virtual Switches for the HX Edge 3- or 4-Node Cluster with 1GE Dual-Switch Topology

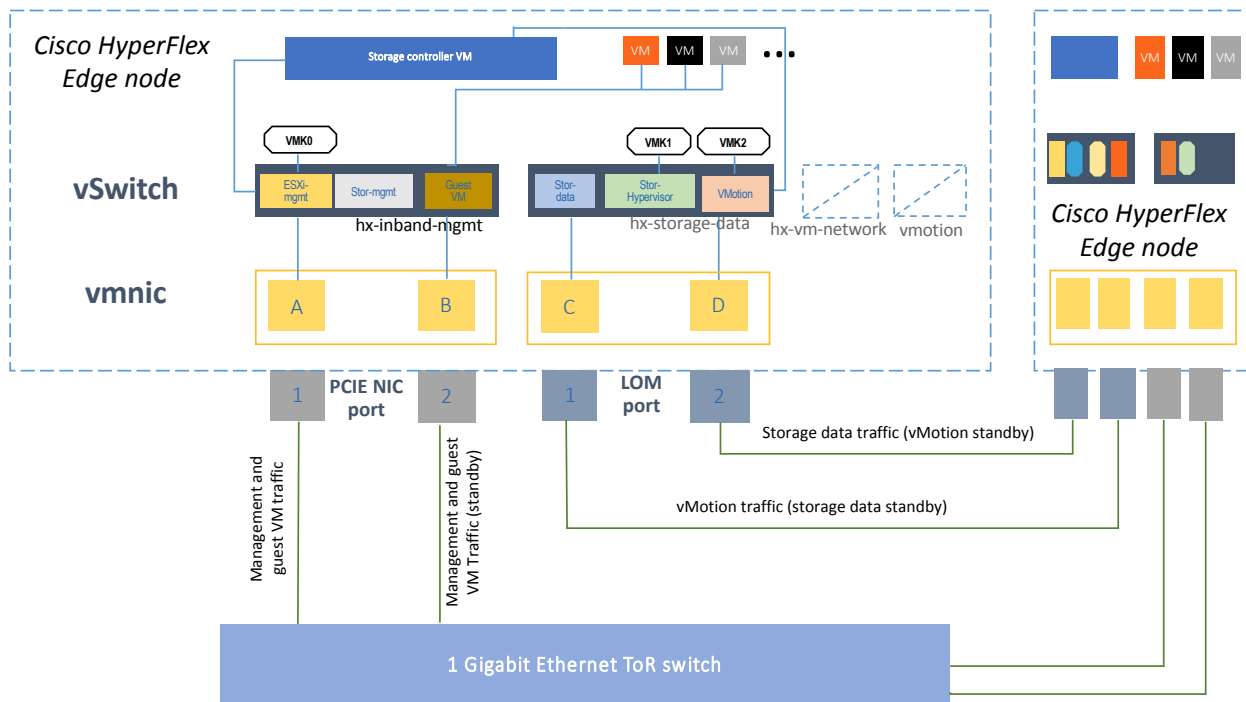
Virtual Switch	Port Groups	VLAN IDs	Active vmnic(s)	Passive vmnic(s)	Jumbo
vswitch-hx-inband-mgmt	Management Network Storage Controller Management Network	<<hx-inband-mgmt>>	vmnico	vmnic2	no

Virtual Switch	Port Groups	VLAN IDs	Active vmnic(s)	Passive vmnic(s)	Jumbo
	Storage Controller Replication Network	<<hx-inband-repl>>			
	vm-network-<<VLAN ID>>	<<vm-network>>	vmnic0 vmnic2		
vswitch-hx-storage-data	Storage Controller Data Network	<<hx-storage-data>>	vmnic3	vmnic1	no
	Storage Hypervisor Data Network				
	vmotion-<<VLAN ID>>	<<hx-vmotion>>	vmnic1	vmnic3	

1GE Single-Switch Logical Networking for HyperFlex Edge 2-Node Cluster

Figure 21 illustrates a logical view of network configuration for the 1GE single switch topology for HyperFlex Edge 2-node cluster.

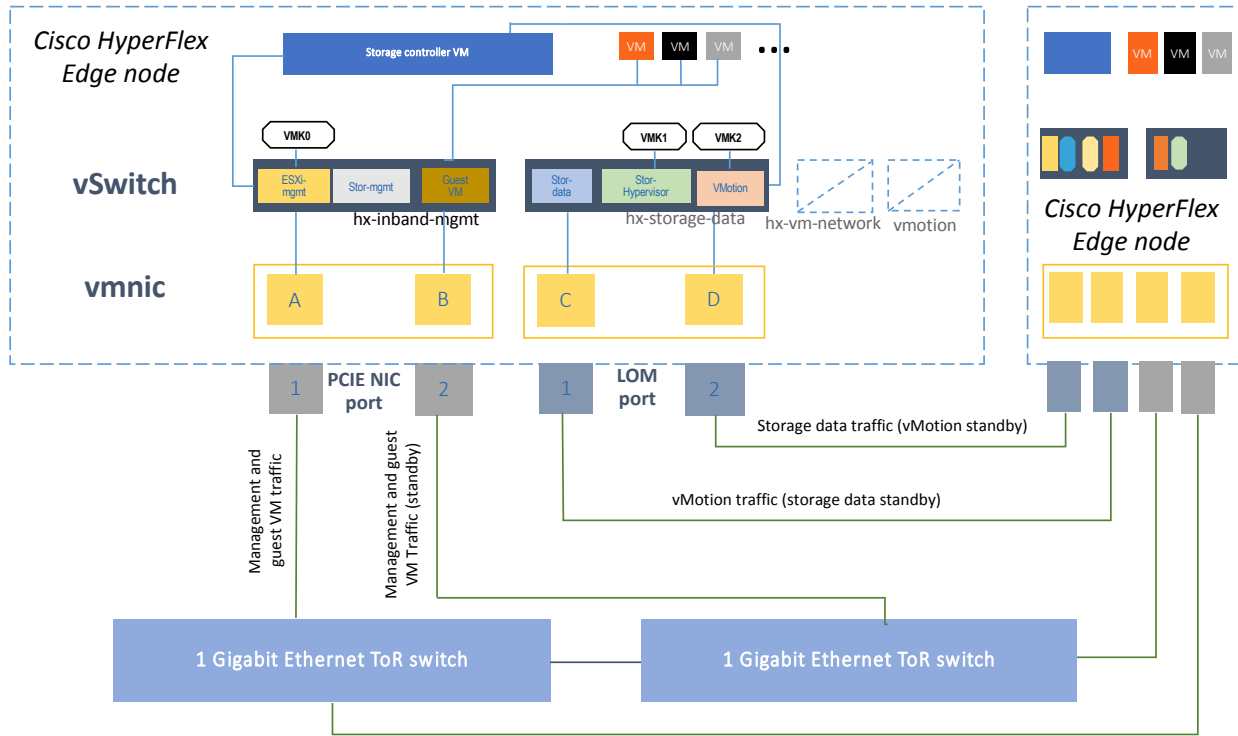
Figure 21 1GE Single-Switch Logical Topology for HyperFlex Edge 2-Node Cluster



1GE Dual-Switch Logical Networking for HyperFlex Edge 2-Node Cluster

Figure 22 illustrates a logical view of network configuration for the 1GE dual switch topology for HyperFlex Edge 2-node cluster.

Figure 22 1GE Dual-Switch Logical Topology for HyperFlex Edge 2-Node Cluster



The same ESXi virtual networking design is applied to the 1GE HyperFlex Edge 2-node cluster, in single switch or dual switch topology. Table 11 lists the details into the ESXi virtual networks.

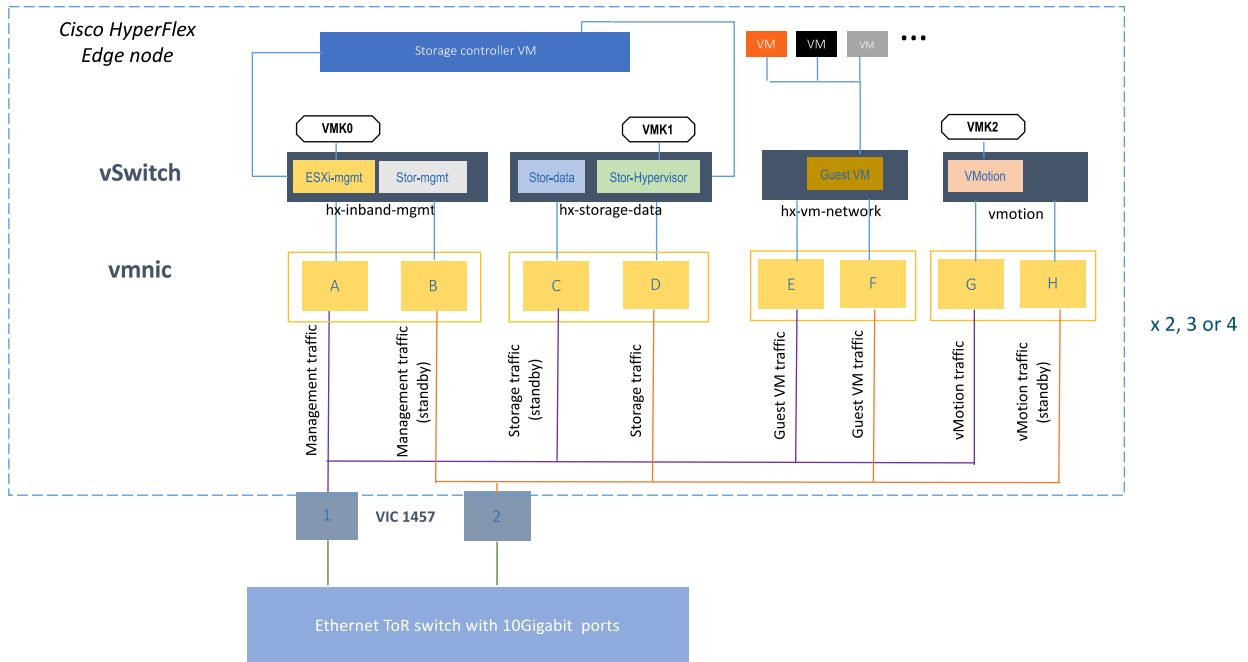
Table 11 Virtual switches for the 1GE HX Edge 2-node cluster

Virtual Switch	Port Groups	VLAN IDs	Active vmnic(s)	Passive vmnic(s)	Jumbo
vswitch-hx-inband-mgmt	Management Network	<<hx-inband-mgmt>>	vmnic2	vmnic3	no
	Storage Controller Management Network				
	Storage Controller Replication Network	<<hx-inband-repl>>			
	vm-network-<<VLAN ID>>	<<vm-network>>	vmnic2 vmnic3		
vswitch-hx-storage-data	Storage Controller Data Network	<<hx-storage-data>>	vmnic0	vmnic1	no
	Storage Hypervisor Data Network				
	vmotion-<<VLAN ID>>	<<hx-vmotion>>	vmnic1	vmnic0	

10GE Single-Switch Logical Networking for HyperFlex Edge Clusters

Figure 23 illustrates a logical view of network configuration for the 10GE single switch topology for HyperFlex Edge clusters.

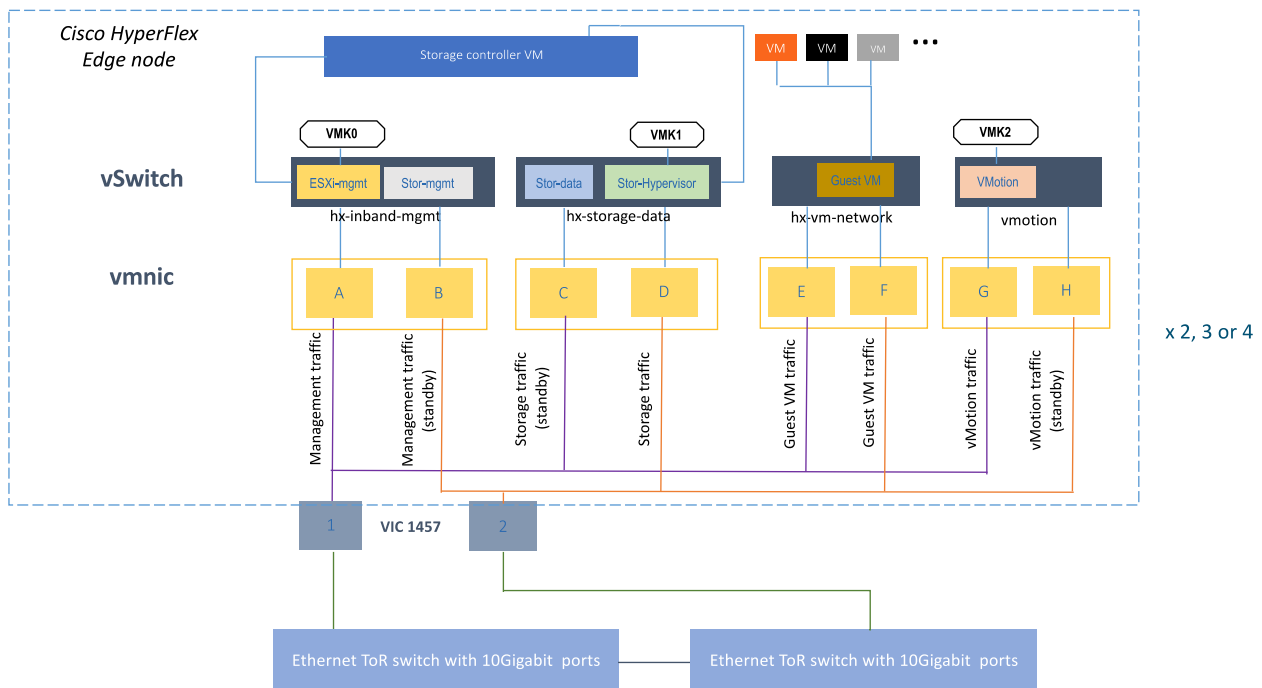
Figure 23 10GE Single-Switch Logical Topology for HyperFlex Edge Clusters



10GE Dual-Switch Logical Networking for HyperFlex Edge Clusters

Figure 24 illustrates a logical view of network configuration for the 10GE dual switch topology for HyperFlex Edge clusters.

Figure 24 10GE Dual-Switch Logical Topology for HyperFlex Edge Clusters



The same ESXi virtual networking design is applied to the 10GE HyperFlex Edge clusters of the size of 2-node, 3-node or 4-node, in single switch or dual switch topology. Table 12 lists the details into the ESXi virtual networks.

Table 12 Virtual Switches for the 10GE HX Edge Clusters

Virtual Switch	Port Groups	VLAN IDs	Active vmnic(s)	Passive vmnic(s)	Jumbo
vswitch-hx-inband-mgmt	Management Network Storage Controller Management Network	<<hx-inband-mgmt>>	vmnic2	vmnic3	no
	Storage Controller Replication Network	<<hx-inband-repl>>			
vswitch-hx-storage-data	Storage Controller Data Network Storage Hypervisor Data Network	<<hx-storage-data>>	vmnic5	vmnic4	no
	vm-network-<<VLAN ID>>	<<vm-network>>	vmnic6 vmnic7		no
vmotion	vmotion-<<VLAN ID>>	<<hx-vmotion>>	vmnic8	vmnic9	no

Deployment of Hardware and Software

This section provides the guidelines to deploy the solution. The steps required to deploy HyperFlex Edge clusters from Cisco Intersight are described in detail here.

Deployment Options

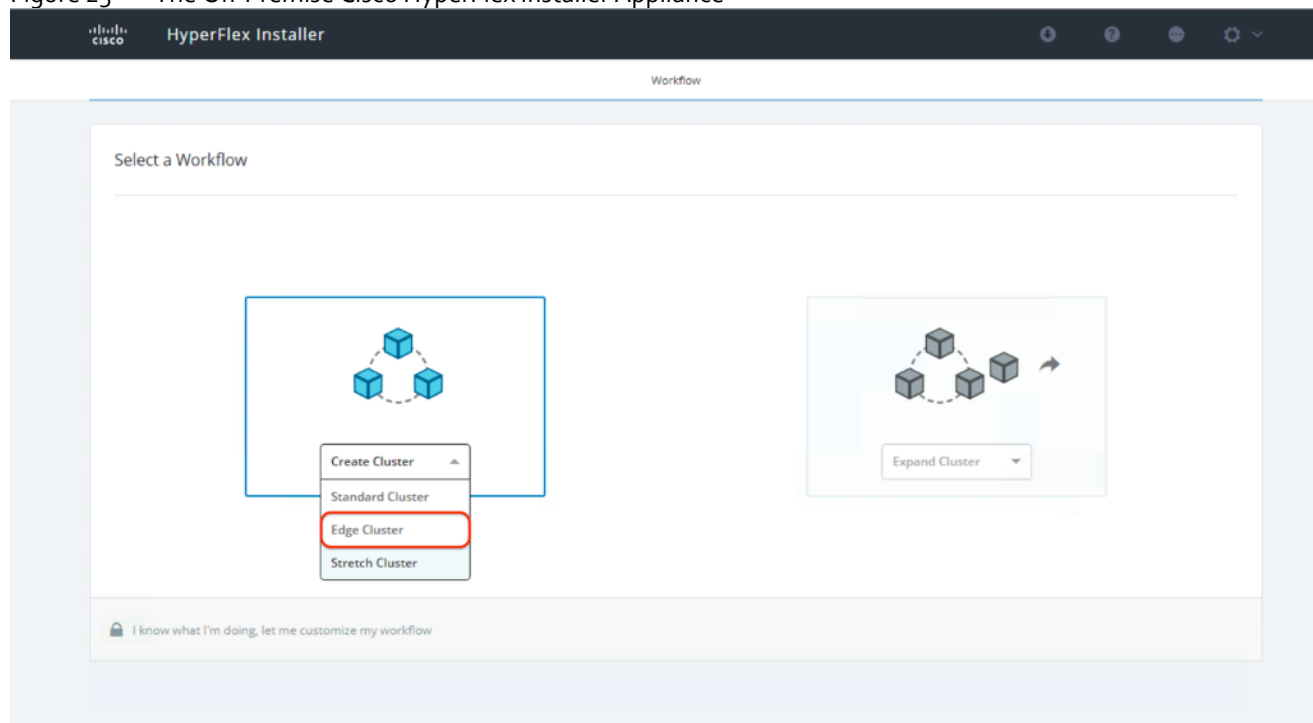
A Cisco HyperFlex Edge system can be deployed using:

- The on-premise Cisco HyperFlex Installer appliance, or
- Cisco Intersight from the cloud

The on-premise Cisco HyperFlex Installer appliance is available as an OVA file that can be downloaded from the Cisco website. After the HX Installer VM has been installed in your existing ESXi environment, it can be accessed via a web browser on the local computer by navigating to the IP address of the HX Installer VM.

The HyperFlex installer will guide you through the process of setting up your cluster. There are some specific workflows to support the installation for Standard clusters, Edge clusters or Stretched clusters. For Edge clusters, the installer will assign IP addresses to the HX-series servers that come from the factory with ESXi hypervisor software preinstalled. The installer will deploy the HyperFlex controller VMs and software on the nodes, add the nodes to the vCenter cluster, then finally create the HyperFlex cluster and distributed filesystem. All of these processes can be completed via a single workflow from the HyperFlex Installer webpage.

Figure 25 The On-Premise Cisco HyperFlex Installer Appliance



HyperFlex Edge clusters can also be deployed rapidly via Cisco Intersight. The cloud installer constructs a pre-configuration definition of your cluster, called an HyperFlex Cluster Profile. This definition is a logical representation of the HX nodes in your HyperFlex Edge cluster. Each HyperFlex node provisioned in Cisco Intersight is specified in a HyperFlex Cluster profile.

The Invisible Cloud Witness service is required for 2-node HyperFlex Edge deployment and is only available from Cisco Intersight.

The on-premise HyperFlex installer option supports both 1GE and 10GE network topologies but can only be used for Edge deployments for three- and four-node clusters. This methodology requires downloading and installation of the appliance along with local network access.

The Intersight cloud option can be used for Edge deployments for all cluster sizes including two-, three- and four-node clusters. This methodology supports all the 1GE and 10GE network topologies listed in the above topology sections.

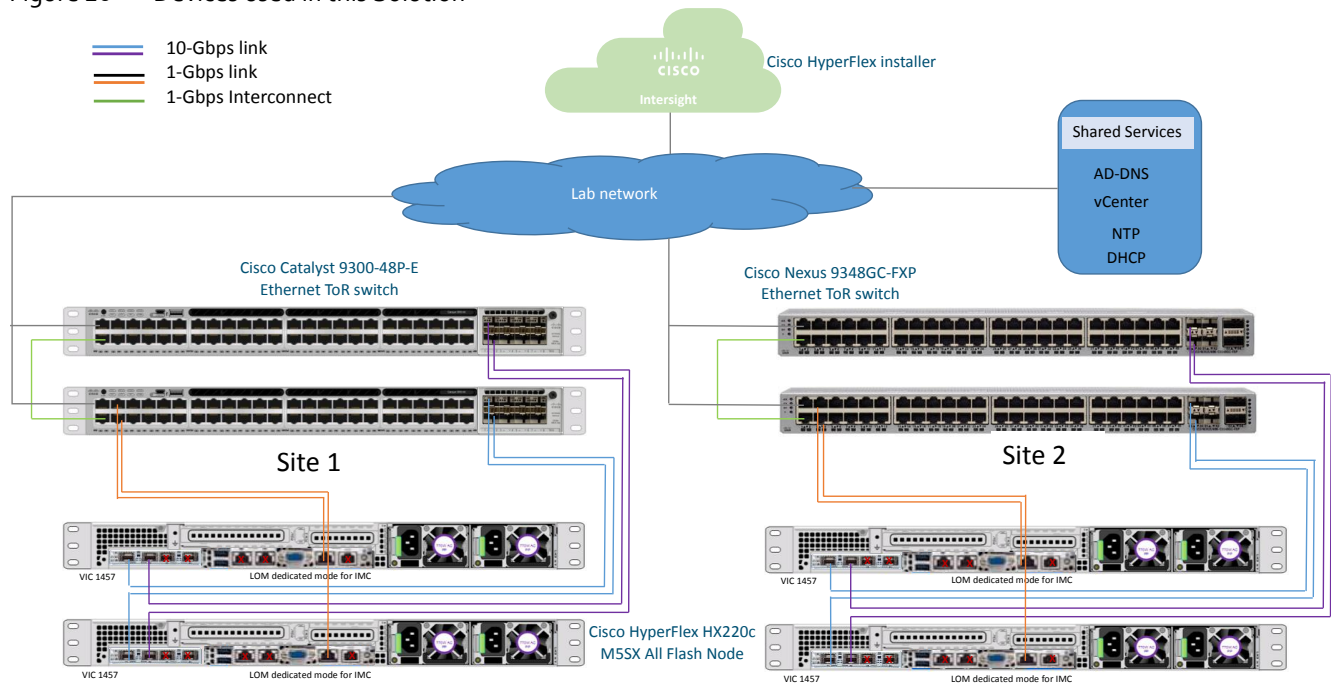


Note: Use of the on-premises HyperFlex installer is not supported for two-node HyperFlex Edge cluster. Alternatively, two-node HyperFlex Edge cluster can be deployed from the on-premises Intersight virtual appliance. The procedures are the same as being deployed from the cloud version of Cisco Intersight.

Solution Architecture

Figure 26 shows the topology used in the deployment and validation of this solution. It demonstrates the cabling of the HyperFlex Edge servers in the 10GE Dual Switch topology.

Figure 26 Devices used in this Solution



In this solution, four (4) Cisco HXAF-E-220M5SX All-Flash rack servers are used to build two 2-node HyperFlex Edge clusters at different sites.



Note: Use of Cisco Intersight is required for deploying 2-node HyperFlex Edge cluster and is the only deployment option that is explained in this CVD.

HXDP 4.0(1a) Validation

Table 13 lists the required hardware components and disk options for the Cisco HXAF-E-220M5SX All-Flash rack servers, which are required for creating the HyperFlex Edge clusters in this solution.

Table 13 HXAF-E-220M5SX Server Components 4.0(1a)

HXAF-E-220M5SX options	Hardware Required
Processors	Two (2) Intel Xeon Silver 4116 12-core CPUs
Memory	Twelve (12) 32 GB D 2666 MHz RDIMMs are chosen (total 384 GB)
Disk Controller	Cisco 12Gbps Modular SAS HBA
SSDs	One (1) 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One (1) 400GB 2.5in Enterprise Performance 12G SAS SSD (10Xendurance) Eight (8) 960 GB SSDs are chosen.
Network	Cisco UCS VIC1457 VIC MLOM (10GE)
Boot Device	One 240 GB M.2 form factor SATA SSD
microSD Card	One 32GB microSD card for local host utilities storage

Table 14 lists the software components and the versions as tested and validated in this document:

Table 14 Software Component Versions 4.0(1a)

Component	Software Version
Cisco HyperFlex Data Platform Software	4.0(1a)
Cisco UCS Firmware	4.0(2c) BIOS 4.0.2a
Cisco IMC Software	4.0(2c)
VMWare vSphere ESXi Hypervisor	6.5.0 Build 8935087 (factory installed), and 6.7.0 Build 13473784
VMWare vSphere vCenter Appliance	6.7.0, Build 13010631

HXDP 4.0(1b) Validation

Table 15 lists the required hardware components and disk options for the Cisco HXAF-E-220M5SX All-Flash rack servers, which are required for creating the HyperFlex Edge clusters in this solution.

Table 15 HXAF-E-220M5SX Server Components 4.0(1b)

HXAF-E-220M5SX options	Hardware Required
Processors	Two (2) 2 nd Gen Intel Xeon Silver 4214 12-core CPUs for cluster 1 Two (2) 2 nd Gen Intel Xeon Gold 5218 12-core CPUs for cluster 2
Memory	Twelve (12) 32 GB 2933 MHz RDIMMs are chosen (total 384 GB)
Disk Controller	Cisco 12Gbps Modular SAS HBA
SSDs	One (1) 240 GB 2.5 Inch Enterprise Value 6G SATA SSD One (1) 400GB 2.5in Enterprise Performance 12G SAS SSD (10Xendurance) Eight (8) 960 GB SSDs are chosen.

HXAF-E-220M5SX options	Hardware Required
Network	Cisco UCS VIC1457 VIC MLOM (10GE)
Boot Device	One 240 GB M.2 form factor SATA SSD
microSD Card	One 32GB microSD card for local host utilities storage

Table 16 lists the software components and the versions as tested and validated in this document:

Table 16 Software Component Versions 4.0(1b)

Component	Software Version
Cisco HyperFlex Data Platform Software	4.0(1b)
Cisco UCS Firmware	4.0(4f) BIOS 4.0.4g
Cisco IMC Software	4.0(4f)
VMWare vSphere ESXi Hypervisor	6.7.0 Build 13473784
VMWare vSphere vCenter Appliance	6.7.0, Build 13010631

The configuration examples of Nexus 9348GC-FXP and Catalyst 9300-48P ToR Switches are in [Appendix C: Example Cisco Nexus 9348GC-FXP Switch Configuration](#) and [Appendix D: Example Cisco Catalyst 9300-48P Switch Configuration](#).

Preinstallation Checklist

Prior to beginning the installation activities, it is important to gather the information contained in the following sections.

Network Topology

To get started, select the size of your HyperFlex Edge cluster and select the network topology according to the available network resources in your IT environment.

Table 17 Network Topology

Item	Selection
Size	<input type="checkbox"/> 2-node <input type="checkbox"/> 3-node <input type="checkbox"/> 4-node
Network Speed	<input type="checkbox"/> 1GE <input type="checkbox"/> 10GE
Uplink	<input type="checkbox"/> Single-switch <input type="checkbox"/> Double-switch
Jumbo Frames	<input type="checkbox"/> Support Jumbo Frames

CIMC Mode

Choose one of the Cisco IMC Connectivity options – Shared or Dedicated Mode.

Table 18 CIMC Mode

Item	Configuration
CIMC Mode	<input type="checkbox"/> Shared Mode <input type="checkbox"/> Dedicated Mode

vCenter Configuration

Virtual or physical vCenter that runs on an external server and is local to the site. VMWare vCenter server or appliance is required for operating the HyperFlex Edge system properly. However, installation of a HyperFlex Edge cluster may be initially performed without a vCenter. But the cluster must be registered to a vCenter server before running production workloads. Deployment with an external virtual or physical vCenter that is local to the site is highly desired.

Table 19 vCenter Configuration

Item	Configuration
vCenter Option	<input type="checkbox"/> Single vCenter (local site) <input type="checkbox"/> Centralized vCenter (multiple sites) <input type="checkbox"/> Nested vCenter
vCenter FQDN	<vCenter name>
vCenter IP Address	<vCenter IP>
Data Center Name	<datacenter name>
Administrator Username	<admin-user>
Administrator Password	<admin-password>

Network Services

DHCP versus Static IP

By default, the HyperFlex installation will assign a static IP address to the management interface of the ESXi servers. Using Dynamic Host Configuration Protocol (DHCP) for automatic IP address assignment is not recommended for HX/ESXi management, storage, or vMotion networks. DHCP or static assignment may be used for the IMC network configuration to ease configuration.

DNS

DNS servers are highly recommended to be configured for querying Fully Qualified Domain Names (FQDN) in the HyperFlex and ESXi Management group. DNS forward and reverse lookup records need to be created prior to beginning the installation. At a minimum, it is highly recommended to create A records and reverse PTR records for the ESXi hypervisor hosts' management interfaces. In addition, all device connectors on the HyperFlex nodes and all HyperFlex Controller VMs must properly resolve public domains via DNS and permit outbound initiated HTTPS connections on port 443 so that the latest HyperFlex software packages can be auto downloaded and the HyperFlex Edge cluster can be fully deployed from Intersight.

The following tables will assist with gathering the required DNS information for the installation, by listing the information required, and an example configuration:

Table 20 DNS Server Information

Item	Value
DNS Server #1	
DNS Server #2	
DNS Domain	
SMTP Server Name	
HX Domain Name	
HX Edge Server #1 Name	
HX Edge Server #2 Name	
HX Edge Server #3 Name	
HX Edge Server #4 Name	

NTP

Consistent time clock synchronization is required across the components of the HyperFlex system, provided by reliable NTP servers, accessible to the Cisco UCS CIMC Management network group, and the HyperFlex and ESXi Management group. NTP is used by Cisco IMC, vCenter, the ESXi hypervisor hosts, and the HyperFlex Storage Platform Controller VMs. The use of public NTP servers is highly discouraged; instead, a reliable internal NTP server should be used.

The following tables will assist with gathering the required NTP information for the installation by listing the information required, and an example configuration:

Table 21 NTP Server Information

Item	Value
NTP Server #1	
NTP Server #2	
Timezone	

Auto-Support Service

It is recommended to enable Connected Services on the HyperFlex cluster so that the emails for the auto support notifications will be sent out to the appropriate contacts.

Table 22 Auto-support Service

Item	Configuration
------	---------------

Item	Configuration
Option	<input type="checkbox"/> Enable Auto Support
Email for service request notifications	<service email address>

Proxy Server

Use of a proxy server is optional if direct connectivity to Intersight is not available. When using a proxy, the device connectors in each server must be configured to use the proxy in order to claim the servers into an Intersight account. In addition, the proxy information must be provided in the HyperFlex Cluster Profile to ensure the HyperFlex Data Platform can be successfully downloaded. Use of username and password is optional.

Table 23 Proxy Server

Item	Configuration
Proxy Required	<input type="checkbox"/> Yes <input type="checkbox"/> No
Proxy Hostname	<Proxy name>
Proxy IP Address	<Proxy IP>
Proxy Port	<Proxy port>
Username	<user>
Password	<password>

VLANs

Prior to the Edge installation, the required VLAN IDs need to be documented and created in the upstream network if necessary. At a minimum there are two VLANs that need to be trunked to the upstream Ethernet switches: a VLAN for the HyperFlex and ESXi Management group, and a VLAN for the HyperFlex Storage group. Following Cisco’s best practices, an additional two VLANs can be added: a VLAN for the VMotion group, and at least one VLAN for the guest VM traffic. For the 10GE topology, all these four VLANs need to be trunked to the upstream Ethernet switches. If HyperFlex Replication is to be used, another VLAN must be created and trunked for the replication traffic. The VLAN for VMotion can be the same as the management VLAN but it is not recommended. CIMC can be in a separate VLAN but usually it is fine if it shares the same VLAN as the HyperFlex and ESXi Management group. During the installation, only the VLAN ID for Management network and the VLAN ID for Storage network must be supplied at the step to create the HyperFlex Cluster Profile.

The following tables will assist with gathering the required VLAN information for the installation by listing the information required, and an example configuration:

Table 24 VLAN Information

VLAN Name	VLAN ID
<<hx-inband-mgmt>>	
<<hx-inband-repl>> (if applicable)	
<<hx-storage-data>>	
<<hx-vm-data>>	
<<vmotion>>	

IP Addressing

To install the HX Data Platform from Cisco Intersight, several IP addresses need to be allocated. The UCS CIMC IP addresses, and HXDP management IP addresses must allow communication with Intersight. Additional IP addresses for the Cisco HyperFlex Edge system need to be allocated from the appropriate subnets and VLANs to be used. IP addresses that are used by the system include the Cisco UCS Integrated Management Controller, HyperFlex and ESXi Management, HyperFlex Storage, and vMotion. The Cisco UCS IMC IP addresses and HyperFlex and ESXi Management IP addresses can come from the same subnet, or be separate, as long as the HyperFlex management subnet permits outbound access to Intersight.

Use the following table to gather and input the required IP addresses for the installation of a HyperFlex Edge cluster.

Table 25 HyperFlex Edge Cluster IP Addressing

Address Group:	UCS Management	HyperFlex and ESXi Management		HyperFlex Storage		VMotion
VLAN ID:						
Subnet:				169.254.1.0		
Subnet Mask:				255.255.255.0		
Gateway:						
Device	UCS CIMC Addresses	ESXi Management Interfaces	Storage Controller VM Management Interfaces	ESXi Hypervisor Storage VMkernel Interfaces	Storage Controller VM Storage Interfaces	VMotion VMkernel Interfaces
HX Cluster					169.254.1.20	
HX Node #1				169.254.1.11	169.254.1.21	
HX Node #2				169.254.1.12	169.254.1.22	
HX Node #3	(If applicable)					
HX Node #4	(If applicable)					



Note: Table cells highlighted in yellow have auto-assigned IP addresses and don't require user configuration.

Usernames and Passwords

Several usernames and passwords need to be defined or known as part of the HyperFlex installation process: Cisco Intersight account, UCS IMC Administrator, ESXi Administrator, HyperFlex Administrator and vCenter Administrator. The following table will assist with gathering the required username and password information by listing the information required and an example configuration:

Table 26 Usernames and Passwords

Account	Username	Password
Cisco Intersight account	<cisco_account_username>	<cisco_account_pw>
CIMC Administrator	admin	<<cimc_admin_pw>>
ESXi Administrator	root	<<esxi_root_pw>>
HyperFlex Administrator	root	<<hx_admin_pw>>
vCenter Administrator	<<vcenter_administrator>>	<<vcenter_admin_pw>>

Physical Installation

Install the HX-Series rack-mount servers according to their corresponding hardware installation guides listed below.

HX220c M5 Server:

https://www.cisco.com/c/en/us/td/docs/hyperconverged_systems/HX_series/HX220c_M5/HX220c_M5.html

Cabling

The physical layout of the HyperFlex system was previously described in the section covering Physical Topology. The HX-series rack-mount servers need to be cabled properly to the customer's upstream networking switches before beginning the installation activities. Depending on the availability of the upstream switches or switch ports, you can choose single switch or dual switch topology, along with the CIMC NIC policy (shared LOM or dedicated). Once these decisions are made, the nodes can be cabled appropriately. Make sure no unnecessary ports on the Quad-port Intel i350 PCIe NIC card or on the Cisco VIC 1457 card are connected prior to cluster installation.

Intersight Connectivity

The following are the prerequisites for Intersight connectivity:

- Before installing the HyperFlex cluster on a set of HyperFlex servers, make sure that the device connector on the corresponding Cisco IMC instance is properly configured to connect to Cisco Intersight and claimed.
- All device connectors must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of the HX Installer supports the use of an HTTP proxy.
- All controller VM management interfaces must properly resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. The current version of HX Installer supports the use of an HTTP proxy if direct Internet connectivity is unavailable.
- IP connectivity (L2 or L3) is required from the CIMC management IP on each server to all of the following: ESXi management interfaces, HyperFlex controller VM management interfaces, and vCenter server. Any firewalls in this path should be configured to allow the necessary ports as outlined in the Hyperflex Hardening Guide.
- Starting with HXDP release 3.5(2a), the Intersight installer does not require a factory installed controller VM to be present on the HyperFlex servers.
- When redeploying HyperFlex on the same servers, new controller VMs must be downloaded from Intersight into all ESXi hosts. This requires each ESXi host to be able to resolve `svc.intersight.com` and allow outbound initiated HTTPS connections on port 443. Use of a proxy server for controller VM downloads is supported and can be configured in the HyperFlex Cluster Profile if desired.
- In addition, on post-cluster deployment the new HyperFlex cluster is automatically claimed in Intersight for ongoing management.

Install Cisco HyperFlex Cluster

Cisco HyperFlex systems are ordered with a factory pre-installed configuration. This factory integration work will deliver the HyperFlex servers with the proper firmware revisions pre-set, a copy of the VMware ESXi hypervisor software pre-installed, and some components of the Cisco HyperFlex software already pre-staged. Once on site, the final steps to be performed are reduced and simplified due to the already completed factory work. As outlined in the previous section, installation of the Cisco HyperFlex system can be done via a deployable HyperFlex installer virtual machine from an OVA file. Another option to install Cisco HyperFlex system is using the Cisco Intersight cloud management platform, wherein the HyperFlex installer function is delivered from the cloud with no need for the users to have their own installer virtual machine.

Use the following procedures to deploy and configure a 2-node Cisco HyperFlex Edge systems from Cisco Intersight cloud platform. The procedures describe how to deploy and run an HX Data Platform configuration where an external vCenter appliance has already been installed and available on an existing ESXi host. Although using an external vCenter appliance is used as an example for this solution, embedded VMware vSphere vCenter is also supported via a separate procedure.

Cisco Intersight Account

Prior to beginning the installation activities, a Cisco Intersight account must be set up. Additionally, you should gather all the configurations settings contained in the pre-install [checklist](#).

A Cisco Intersight account is required for this solution. To create your Intersight account you must have a valid Cisco ID first. If you do not have a Cisco ID yet and need to create an account, follow these steps:

1. Visit <https://intersight.com> from your workstation.
2. Click Sign In with Cisco ID.
3. On the Cisco Login page, you have the option to Log into an Existing Account or click Register Now to create a new account.

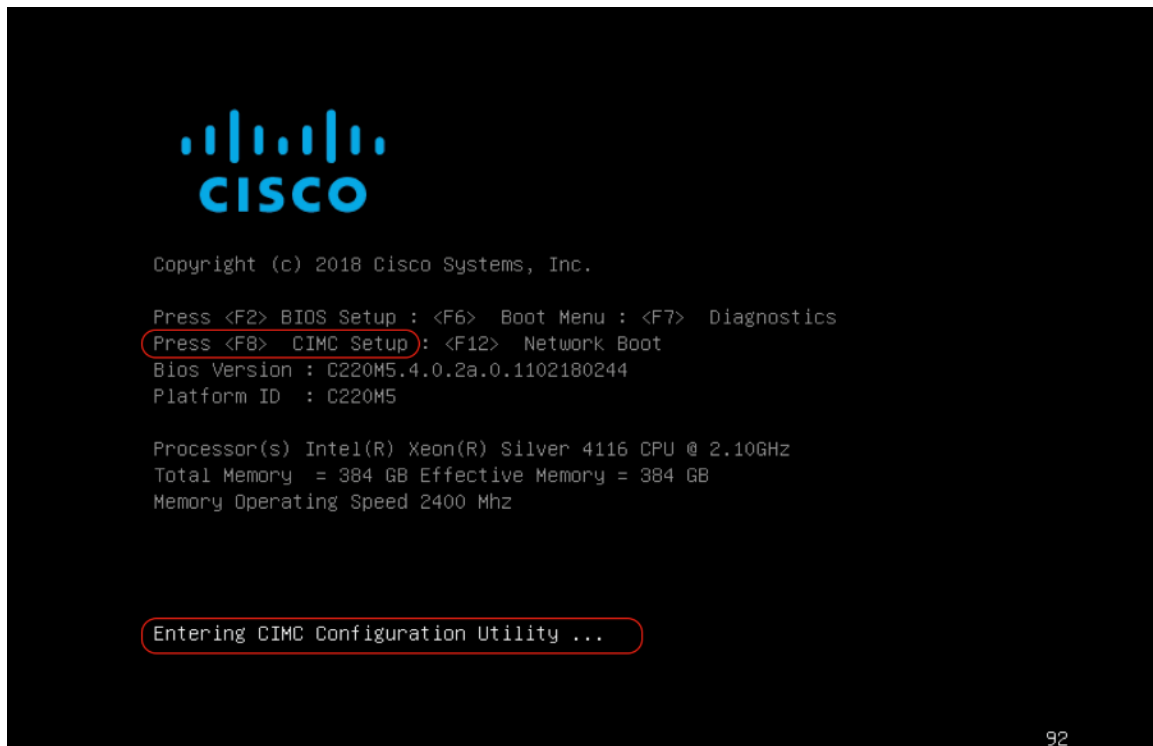
4. Click Register Now and provide the requested information to create a cisco.com account.
5. Once a valid account is created, you can use it to login to Cisco Intersight.

Cisco UCS IMC Configuration


Cisco Intersight cloud-based management platform provides remote management of HyperFlex Edge servers through a device connector that is embedded in the servers' firmware and can be enabled from Cisco Integrated Management Controller (CIMC) Software. The UCS device connector provides access to UCS server management functionality built within the server. By registering the device connector with Intersight, UCS servers or HyperFlex clusters will be claimed as serviceable devices and will be easily and remotely manageable. Cisco Intersight needs access to the CIMC and also the HyperFlex management network.

In order to enable the device connector on the HyperFlex nodes for Cisco Intersight to communicate with, the Cisco Integrated Management Controller (CIMC) interface needs to be configured on each server. There are two options to set an IP address for the CIMC, DHCP or Static. If you already have a DHCP server set up in your environment, DHCP assignment will be the simplest approach; otherwise you can assign the CIMC a static IP address. To configure Static, follow these steps:

1. Attach a Cisco KVM dongle to the KVM connector port in the front of the new HyperFlex server and connect it to a monitor and USB keyboard.
2. Power on the server.
3. Wait for some minutes until the screen with Cisco logo is displayed.
4. Press the F8 button to enter Cisco IMC configuration utility when prompted for boot options.

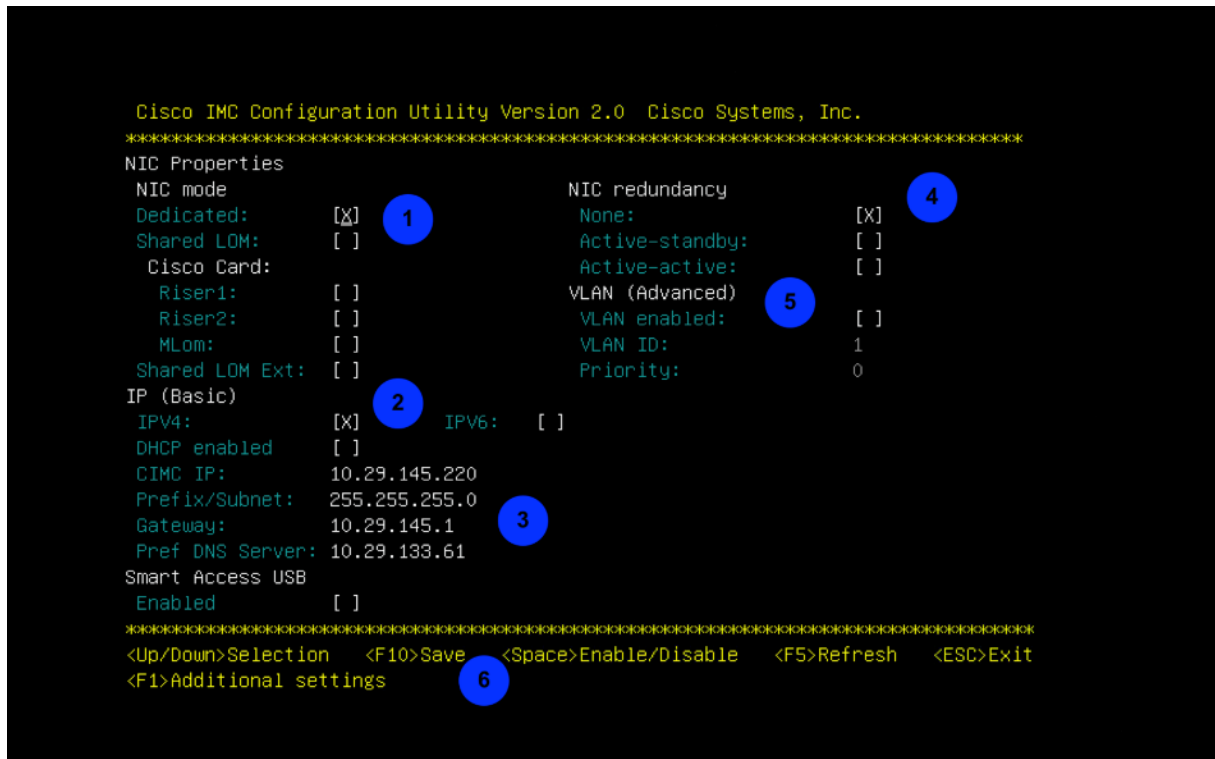


5. For the first time configuration, you will be prompted for a new CIMC password. Input the default password of "password" as the current CIMC password and set a new CIMC password. This configuration will use Cisco123 as the new password.

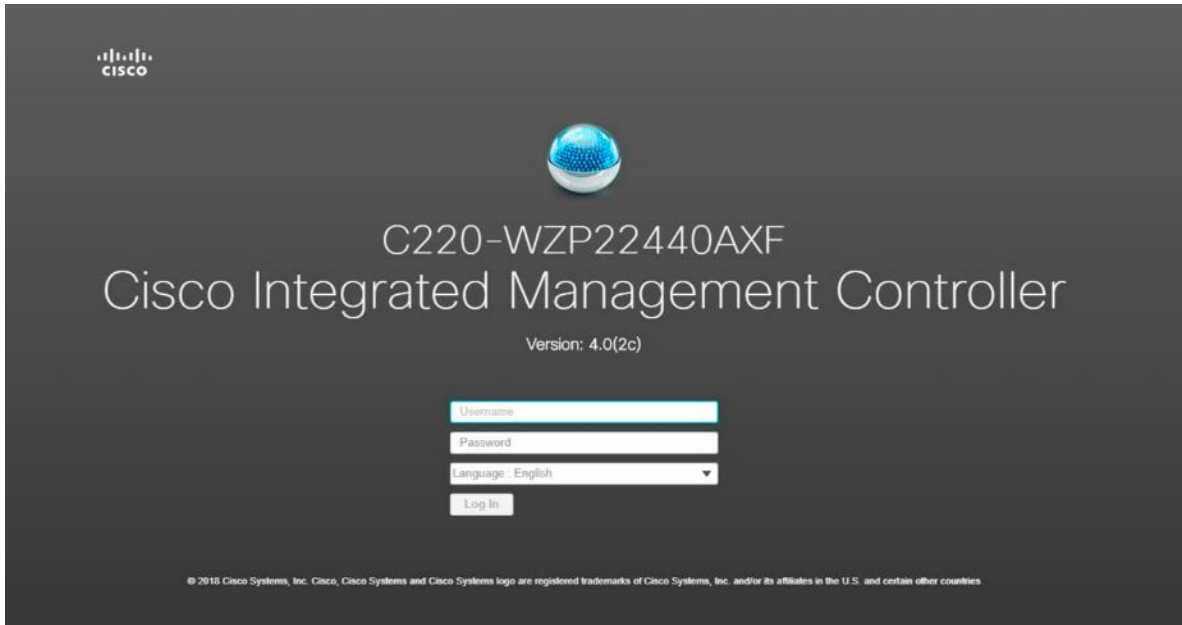


6. On the Cisco IMC configuration utility window, choose the desired NIC Mode. Select Dedicated if you want to access CIMC through the dedicated management port. Select other modes, e.g. Shared LOM Ext mode, if you want to access CIMC through shared LOM or adapter port. In this solution Dedicated mode is chosen.
7. Uncheck DHCP enabled, check IPV4 to set a static IP address.
8. Input IPV4 settings for CIMC IP, Prefix/Subnet, and Gateway.
9. Select NIC redundancy. In this solution None is selected as Dedicated Mode was chosen.

- 10. If needed check VLAN enabled, input VLAN ID for CIMC management. In this solution this field is left blank.
- 11. Hit F10 to save the configuration and then hit ESC to exit the utility. The server will reload and might boot up into the ESXi operating system preloaded from factory. Do not worry if the server does not fully boot into an OS at this point. The installation workflow will modify the boot order automatically at the correct stage of deployment, so do not make manual configuration changes in the CIMC utility that deviate from the steps below.

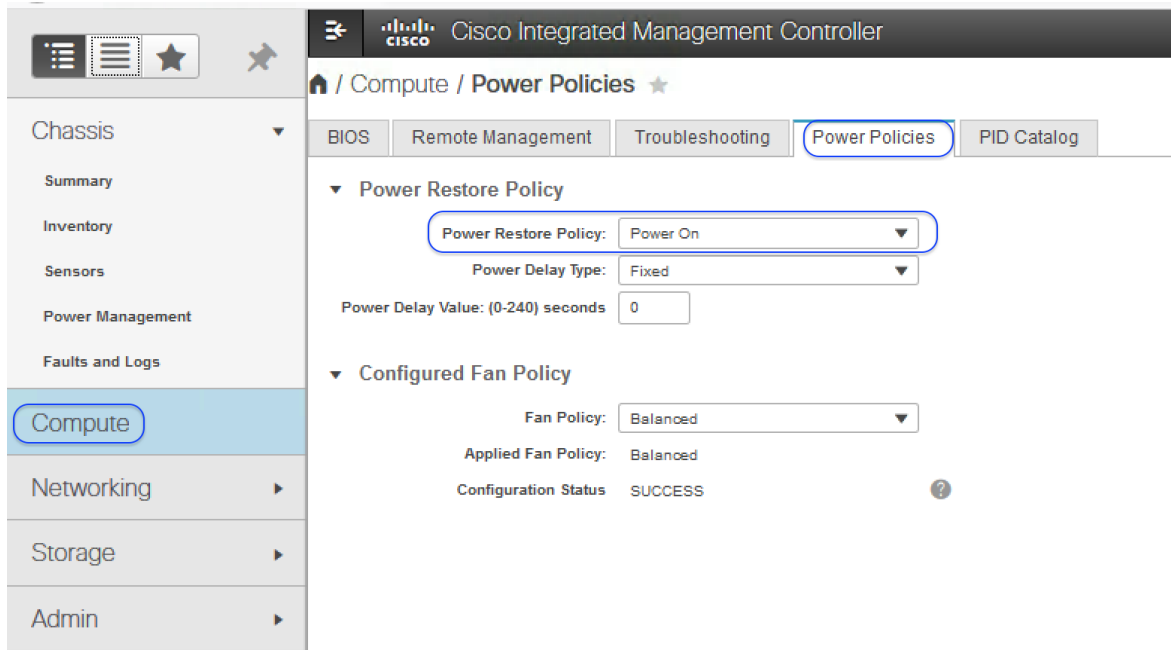


- 12. Ping the CIMC IP address that was just set to validate the configuration.
- 13. Open a web browser with **https://<CIMC-IP-Address>** to CIMC management GUI. Login with admin and the configured password.

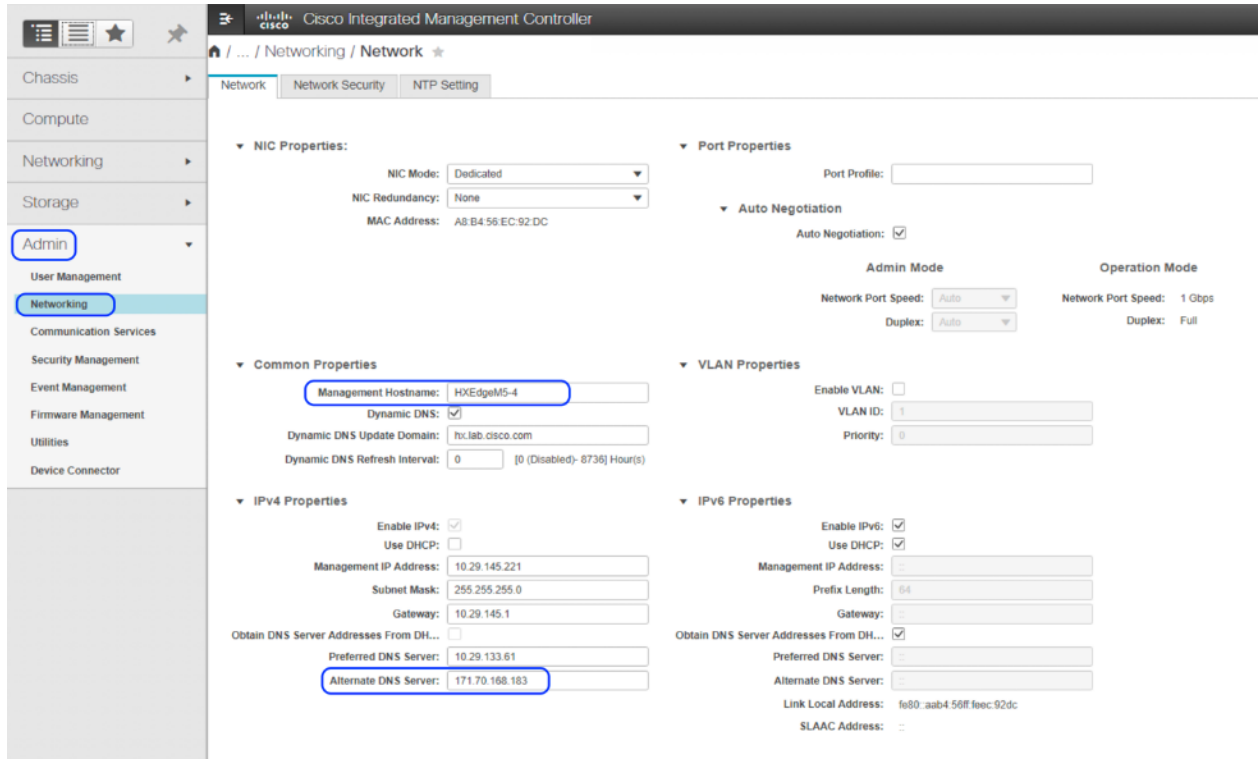


- On the Summary page, check the BIOS running version and IMC Firmware running version to ensure these versions meet the requirements for HyperFlex and Intersight deployment. The factory shipped systems should already have the recommended firmware preloaded. If any firmware upgrade is needed, go to Cisco website to download the latest [Host Upgrade Utility \(HUU\)](#) for this model of server, and use that utility to upgrade the firmware to a compatible version. The HUU is used by booting the server from the HUU ISO file via the CIMC remote KVM.

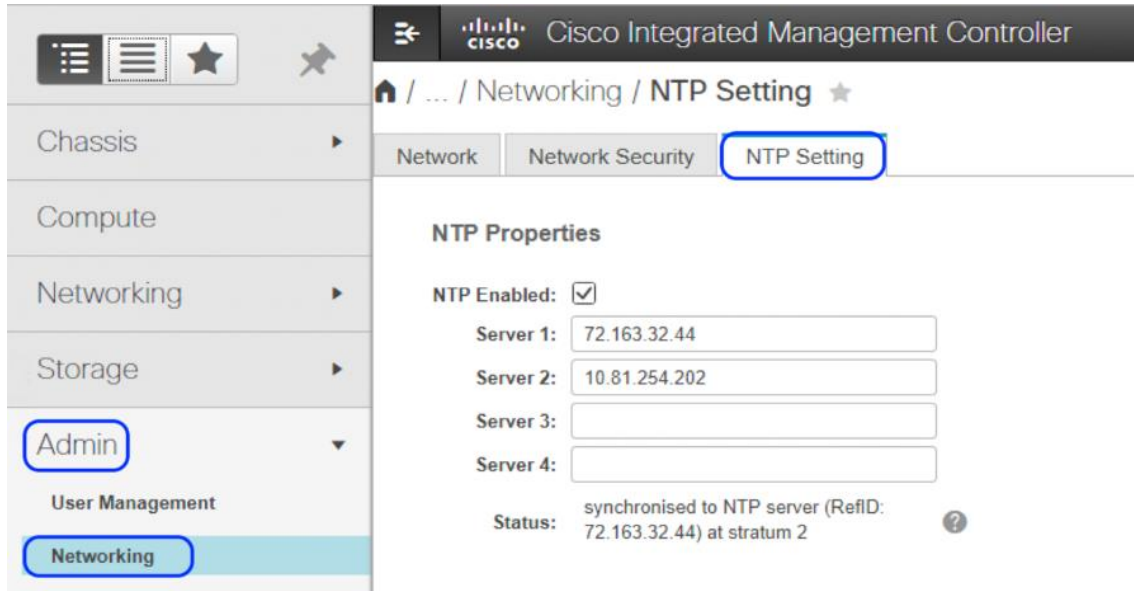
- Click Select Timezone to set the right time zone for the server.
- From the navigation pane choose Compute, then from the Power Policy tab change CIMC Power Restore Policy from default "Power Off" to the desired action.



17. From the navigation pane expand Admin, then click Networking, under Network tab, review your CIMC's IP settings, add an alternative DNS server here if desired. You can change the hostname here as well.



18. Click NTP Setting tab, then add your NTP server information.

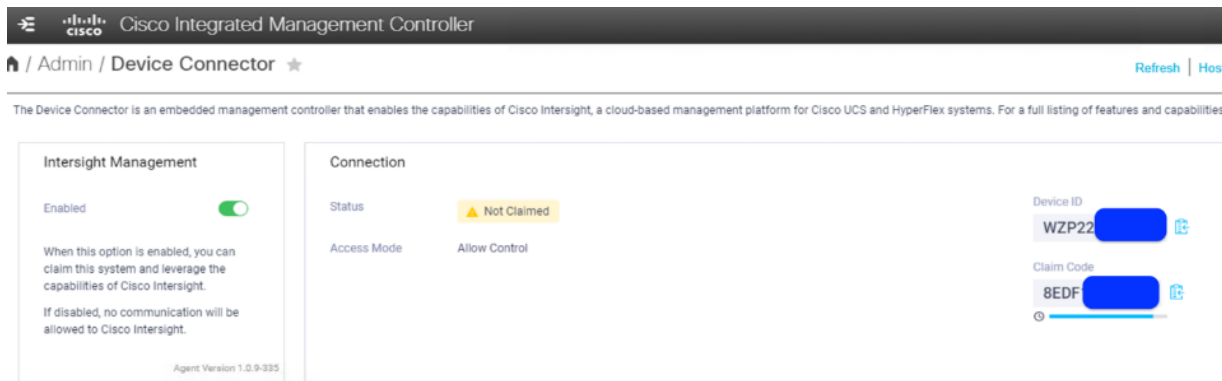


19. After making the changes, click Save Changes to save and apply the changed values. Change of hostname will create a new certificate for the web browser requiring you to re-login to the web GUI. It might take several minutes to commit the changes.
20. Repeat steps 1-19 for all the HyperFlex Edge servers.

Claim Devices in Intersight

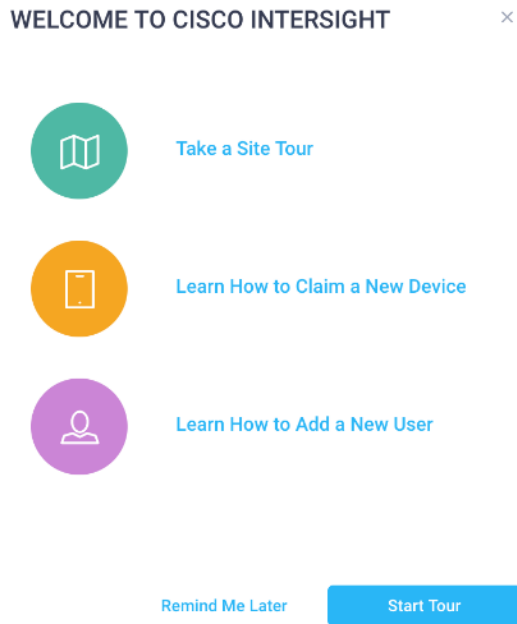
To enable the device connector on the HyperFlex servers and let Intersight claim them for cloud management, follow these steps:

1. Log into CIMC web management GUI with **https://<CIMC-IP-Address>**.
2. From the navigation pane expand Admin, then click Device Connector, turn on the Device Connector to enable Intersight Management. This enables the CIMC to establish a connection to Cisco Intersight.
3. Wait until the connection succeeds and a Claim Code has been generated for this device. Note that the device is Not Claimed yet. Take note of the Device ID (server S/N) and claim code as they will be needed in a subsequent step.

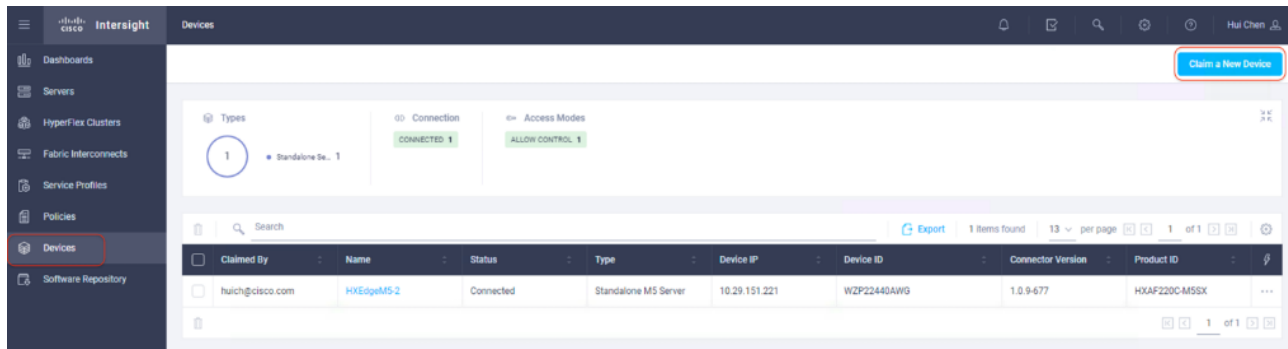


4. Repeat Steps 1-3 for all the HyperFlex Edge servers to gather all Device ID's and claim codes. Note: claim codes do have an expiration time as indicated by the colored bar under the claim code. Be sure to claim the servers before the codes expire.

- Open a web browser and navigate to the Cisco Intersight Cloud Management platform <https://intersight.com/>.
- Login with your Cisco ID and password. If this is the first time using Intersight, it is recommended you take a site tour to be guided through some main features.



- To Claim a new device, from the left-hand Navigation pane, click Devices, in the Device window, choose Claim a New Device at the right top corner.



- Input the Device ID and Claim Code obtained from UCS IMC management GUI. Use copy and paste for accuracy. Click Claim.

CLAIM A NEW DEVICE

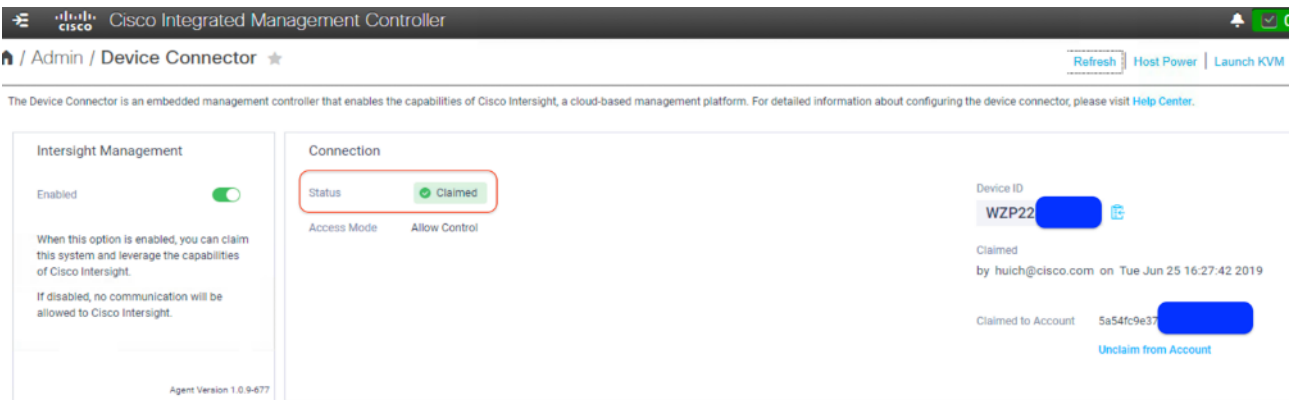
To claim your device, you must have the Device ID and Claim Code.

Device ID *

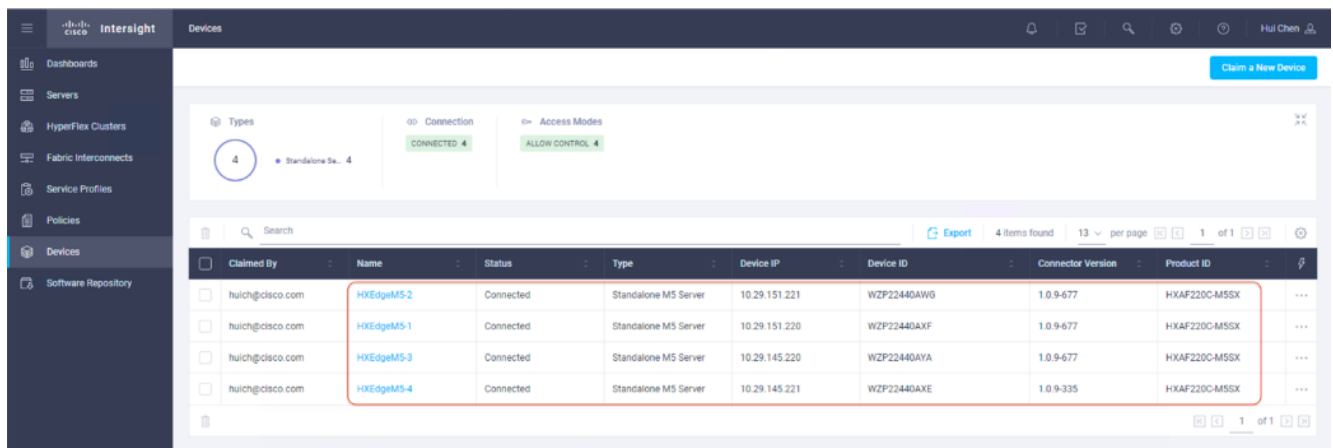
Claim Code *

Cancel Claim

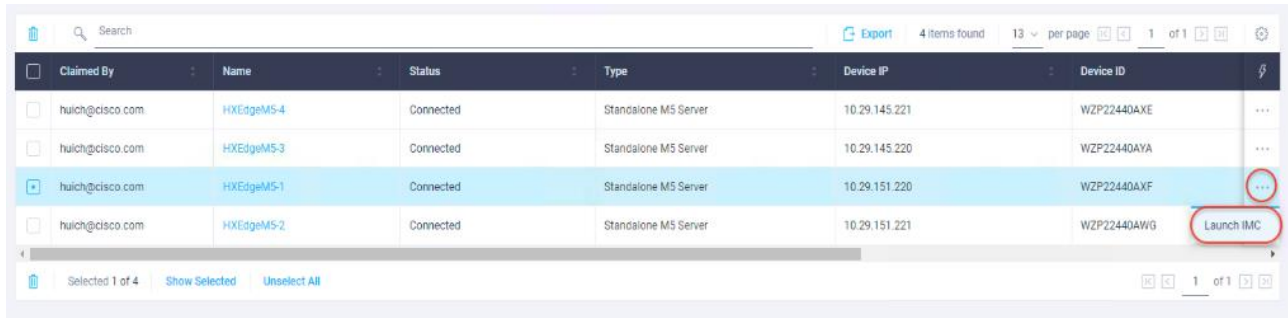
9. Wait until the device is claimed successfully.
10. Go to the server’s CIMC page to double-check the status. The Device Connector page now shows this device is claimed.



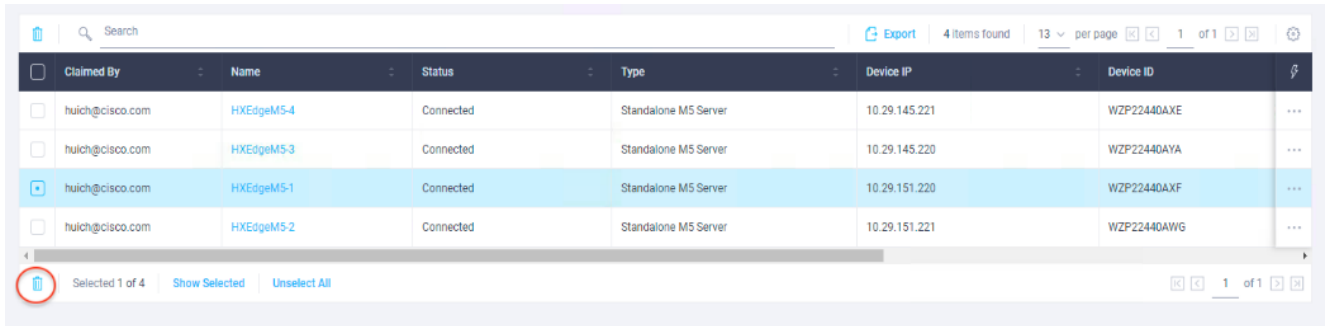
11. Repeat Steps 7-10 to claim all the HyperFlex servers in Intersight.
12. In the Device window, these HyperFlex nodes should now display as connected devices.



13. You now have the option to launch the UCS CIMC management GUI directly from Intersight.



- To Un-claim (remove) a device from Intersight, select the device(s), click the trash icon, then click the pop-up confirmation page and click Remove to confirm the removal.

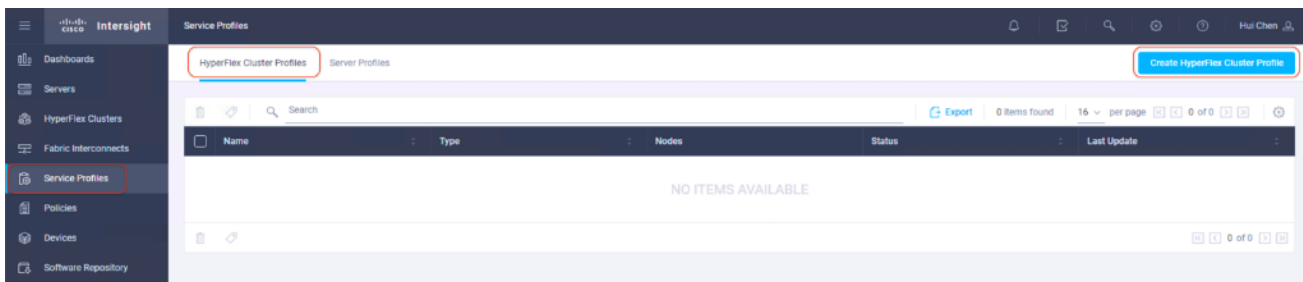


HyperFlex Edge Installation – Single Site

Cisco Intersight provides an installation wizard to install, configure, and deploy Cisco HyperFlex Edge (HX Edge) clusters. The wizard constructs a pre-configuration definition of an HX Edge cluster called an HX Cluster Profile. The cluster profile is policy driven with administrator-defined sets of rules and operating characteristics such as the node identity, interfaces, and vCenter connectivity. Every active node in the HX Edge cluster must be associated with an HX Cluster Profile. After the user inputs all configuration settings, the installation wizard will validate and deploy the HX Cluster Profile on the HyperFlex Edge nodes. You can clone a successfully deployed HX Cluster Profile, and then use that copy as the basis to easily create many more new clusters at the same site or at the remote site.

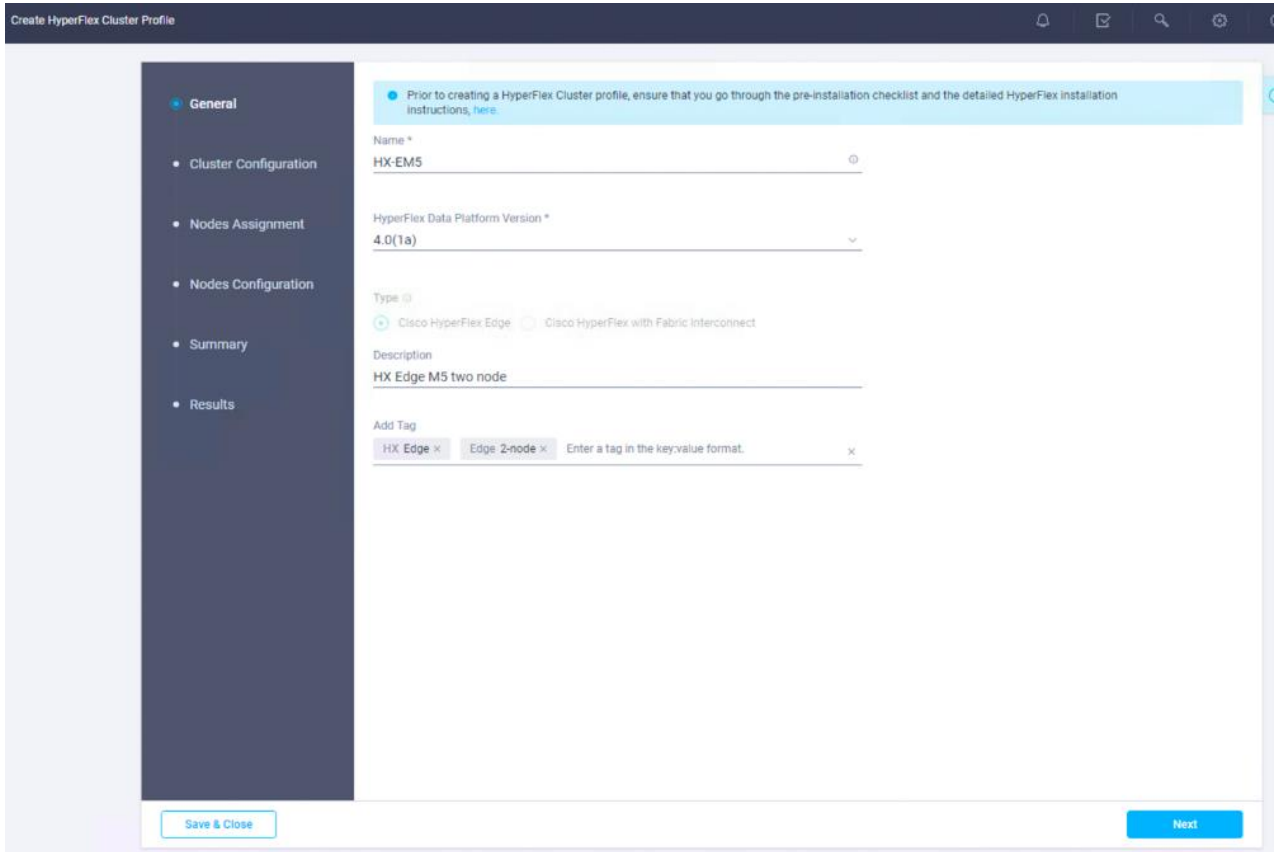
To install and configure a HyperFlex Edge cluster with Intersight, follow these steps:

- Log into Cisco Intersight Cloud Management platform <https://intersight.com/> with your Cisco ID and password.
- From the left-hand navigation pane, choose Service Profiles. On the Service Profiles page, click HyperFlex Cluster Profile tab then choose Create HyperFlex Cluster Profile.

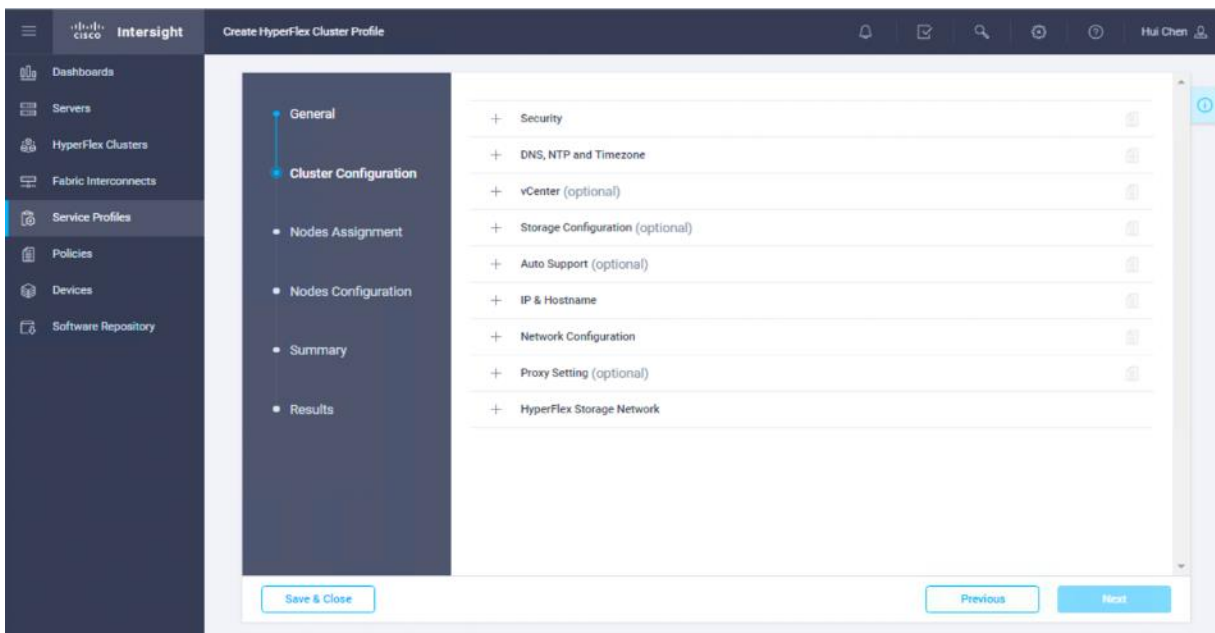


- The HyperFlex Cluster Profile installation wizard is displayed. On the General page, enter a cluster name under Name. This cluster name must be unique and will be used as the HXDP cluster name, vCenter cluster name, and Intersight

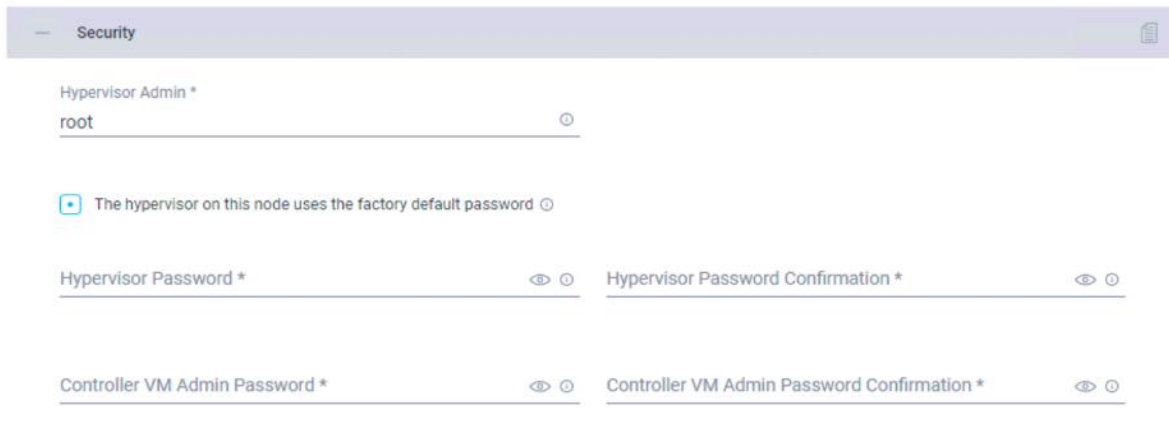
cluster name. Select the appropriate HXDP version. Under Type, select Cisco HyperFlex Edge. Add the necessary description and tags for this cluster for good reference.



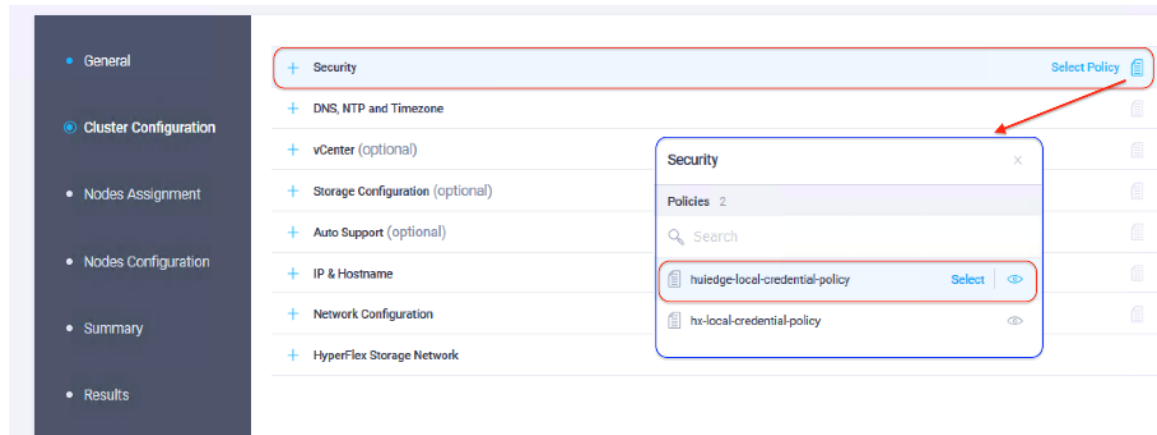
4. Click Next on the Cluster Configuration page.



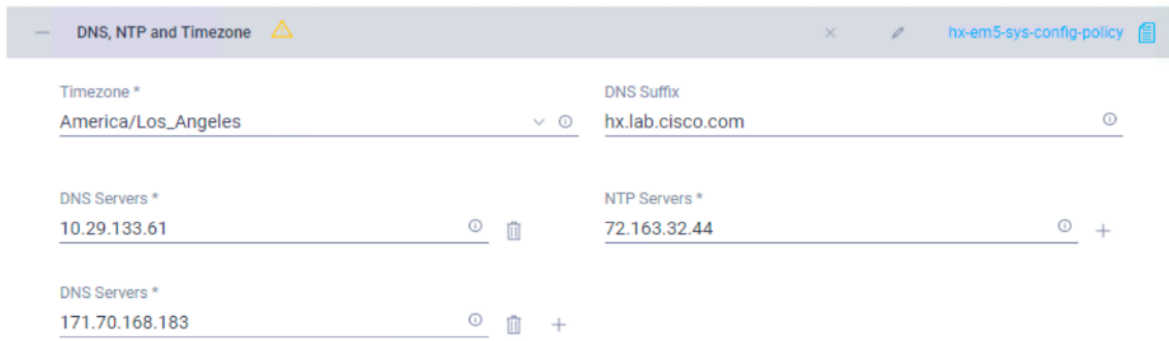
- Click + to expand Security configuration. Enter root as the Hypervisor administrative user. Click the checkbox if the hypervisor on this node uses the factory default password. Input a user supplied new password for the Hypervisor and a user supplied password for the HX controller VM. Confirm that password in the second text field.



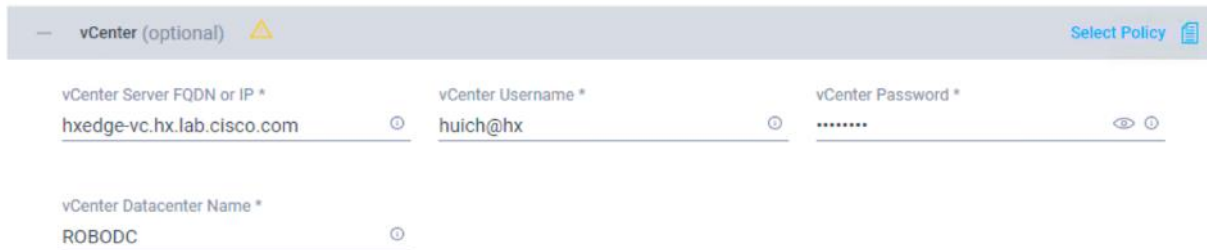
- Once you close the security configuration, the settings are automatically saved to a policy named <HX-Cluster-Name>-local-credential-policy. This policy is reusable and can be selected when you create your next HX Cluster Profile.
- (Optional) To choose an existing policy for your cluster profile, at the policy line, click Select Policy icon, to choose the desired policy from the available policy list and click Select.



- Click + to expand DNS, NTP and Timezone configuration. Choose a time zone from the drop-down list and enter the DNS server and NTP server information. Click + to input secondary DNS or NTP servers.

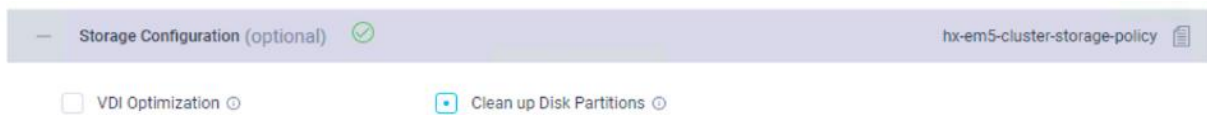


9. Once you close the DNS, NTP, and Timezone configuration, the settings are automatically saved to a reusable policy named <HX-Cluster-Name>-sys-config-policy.
10. Click + to expand vCenter configuration. Enter the vCenter server FQDN or IP address, the administrative username and password. Enter the Datacenter name in vCenter hosting the HyperFlex Edge cluster. If deploying in a nested vCenter configuration, you may leave the vCenter fields blank and register the cluster to vCenter once it has been deployed on the cluster. For more information, see the [How to Deploy vCenter on the HX Data Platform](#) guide.



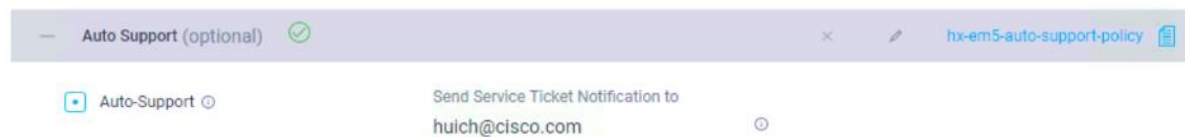
The screenshot shows a configuration panel titled "vCenter (optional)" with a warning icon. It contains three input fields: "vCenter Server FQDN or IP *" with the value "hxedge-vc.hx.lab.cisco.com", "vCenter Username *" with the value "huich@hx", and "vCenter Password *" with masked characters. Below these is a "vCenter Datacenter Name *" field with the value "ROBODC". A "Select Policy" button is visible in the top right corner.

11. Once you close the vCenter configuration, the settings are automatically saved to a reusable policy named <HX-Cluster-Name>-vcenter-config-policy.
12. Click + to expand Storage configuration. Select Clean Up Disk Partitions if performing a reinstallation on top of an existing deployment. If deploying a VDI environment on a hybrid HX cluster, check the box to enable filesystem optimizations.



The screenshot shows a configuration panel titled "Storage Configuration (optional)" with a checkmark icon. It features two checkboxes: "VDI Optimization" (unchecked) and "Clean up Disk Partitions" (checked). A policy name "hx-em5-cluster-storage-policy" is displayed in the top right.

13. Once you close the Storage configuration, the settings are automatically saved to a reusable policy named <HX-Cluster-Name>-cluster-storage-policy.
14. Click + to expand Auto Support configuration. Check the box to enable Auto-Support. Enter your email address for the service ticket notification.



The screenshot shows a configuration panel titled "Auto Support (optional)" with a checkmark icon. It includes a checked "Auto-Support" checkbox and a text input field "Send Service Ticket Notification to" containing the email address "huich@cisco.com". A policy name "hx-em5-auto-support-policy" is shown in the top right.

15. Once you close the Auto Support configuration, the settings are automatically saved to a reusable policy named <HX-Cluster-Name>-auto-support-policy.
16. Click + to expand IP & Hostname configuration. Enter a Hostname Prefix. In a later step, hostnames will be sequentially assigned to hosts using this prefix. Enter a starting IP address, an ending IP address, the subnet mask and gateway for the management IP pool. IPs from this range will be automatically assigned to hosts during the node configuration step. If only the management network IPs are entered, the same range will be used for both ESXi management and HX Controller VM management IPs. If you desire to use a second, non-contiguous range of IPs for the HX Controller VMs, you may optionally enter starting and ending IP addresses, subnet mask and gateway for the HX Controller VM management IP pool. Note these two IP ranges must fall within the same IP subnet and VLAN.

IP & Hostname hx-em5-node-config-policy

Hostname Prefix *
HX-E

Management Network Starting IP * Management Network Ending IP *
10.29.151.224 10.29.151.225

Management Network Subnet Mask * Management Network Gateway *
255.255.255.0 10.29.151.1

Controller VM Management Network Starting IP Controller VM Management Network Ending IP
10.29.151.228 10.29.151.229

Controller VM Management Network Subnet Mask Controller VM Management Network Gateway
255.255.255.0 10.29.151.1

17. Once you close the IP & Hostname configuration, the settings are automatically saved to a reusable named <HX-Cluster-Name>-node-config-policy.
18. Click + to expand Network configuration. Select the Uplink Speed for the upstream network topology. There are two options to choose from for HyperFlex Edge, 1GE or 10GE. Choose 10GE for this deployment.

Uplink Speed

10G (HyperFlex Edge)

1G (HyperFlex Edge)

10G (HyperFlex Edge)

19. Enter MAC Prefix Starting and Ending addresses. Enter the VLAN ID for the management network. Leave the Jumbo Frames checkbox unchecked for HyperFlex Edge deployments.

Network Configuration hx-em5-cluster-network-policy

Uplink Speed
10G (HyperFlex Edge)

MAC Prefix Starting Address * MAC Prefix Ending Address *
00:25:B5:E0 00:25:B5:E0

Management Network VLAN ID * Jumbo Frames
101

20. Once you close the Network configuration, the settings are automatically saved to a reusable policy named <HX-Cluster-Name>-cluster-network-policy.
21. Leave the optional Proxy Setting blank.

Proxy Setting (optional)

Hostname * Port * Username

Password

- 22. Click + to expand HyperFlex Storage Network configuration. Enter the VLAN ID for the data storage network. It is required to use a unique storage VLAN per cluster if multiple clusters are to be deployed within the same network. To avoid the conflict this policy is not saved for reuse.

HyperFlex Storage Network

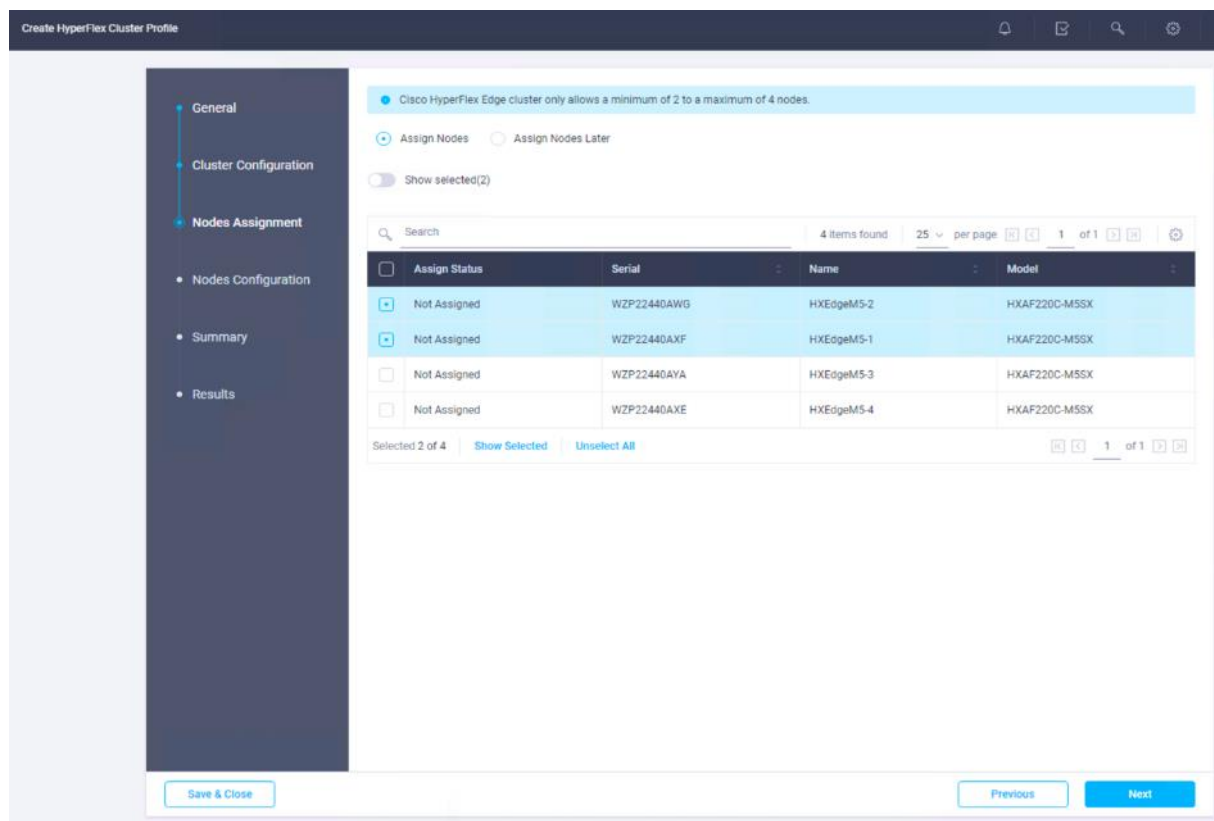
Storage Network VLAN ID *

102

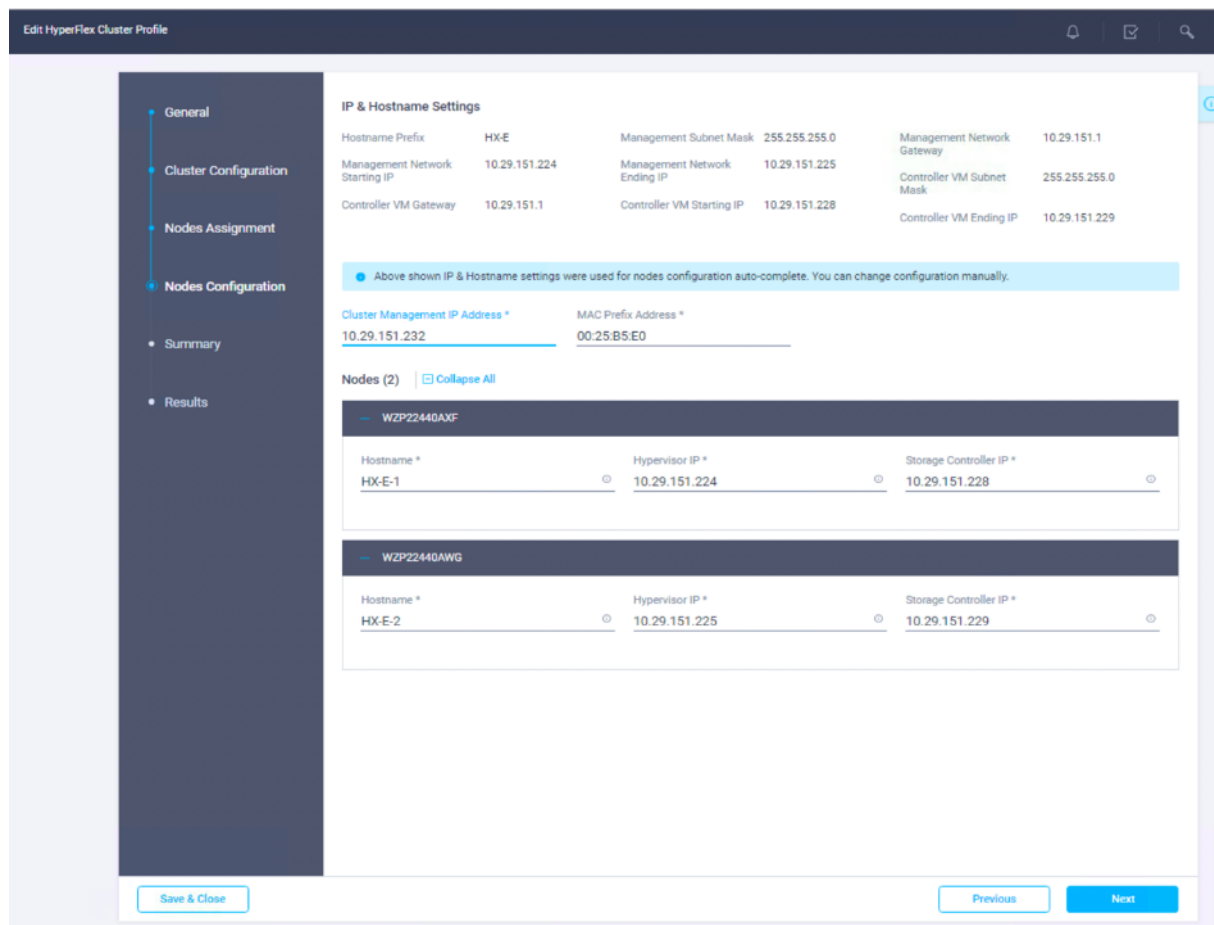
- 23. Once all the policies are configured the saved policies will be listed on this page.

+ Security	hx-em5-local-credential-policy
+ DNS, NTP and Timezone	hx-em5-sys-config-policy
+ vCenter (optional)	hx-em5-vcenter-config-policy
+ Storage Configuration (optional)	hx-em5-cluster-storage-policy
+ Auto Support (optional)	hx-em5-auto-support-policy
+ IP & Hostname	hx-em5-node-config-policy
+ Network Configuration	hx-em5-cluster-network-policy
+ Proxy Setting (optional)	
+ HyperFlex Storage Network	

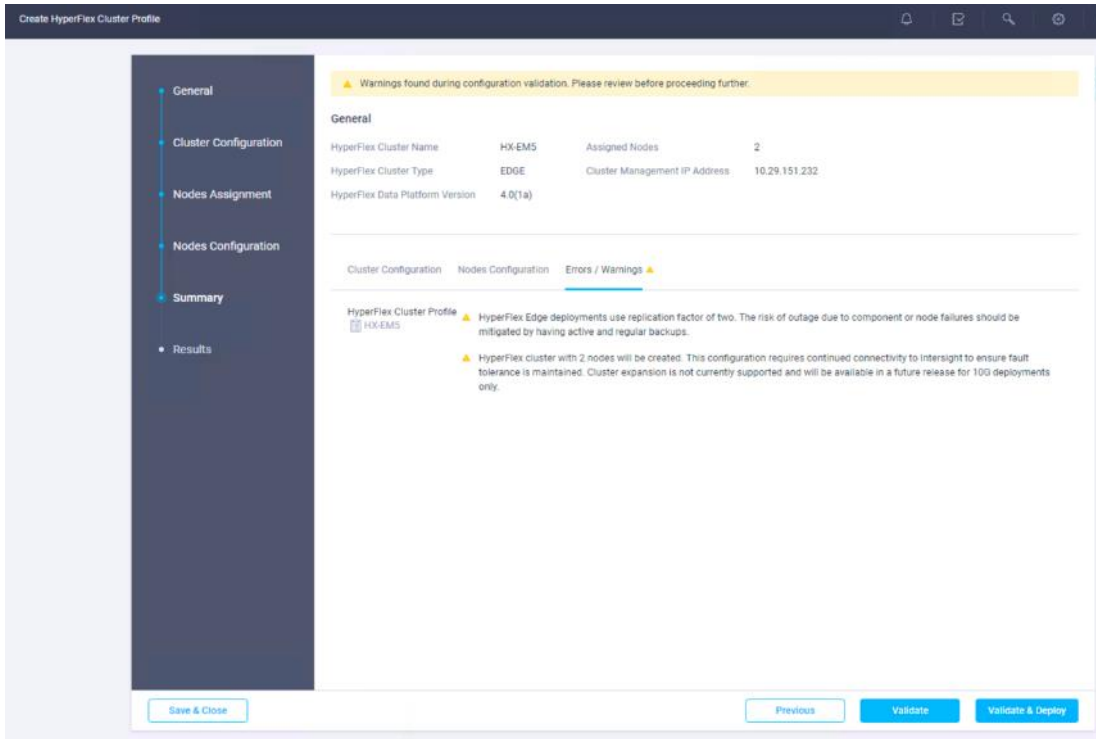
- 24. Click Next to proceed to the Nodes Assignment page. Select two available HyperFlex Edge nodes. You also have the option to Assign Nodes Later. This allows you to save the HX cluster profile now, for example, in a situation where the hardware is not yet installed but you want to perform these configuration steps early and return when the hardware is available to assign once claimed in Intersight.



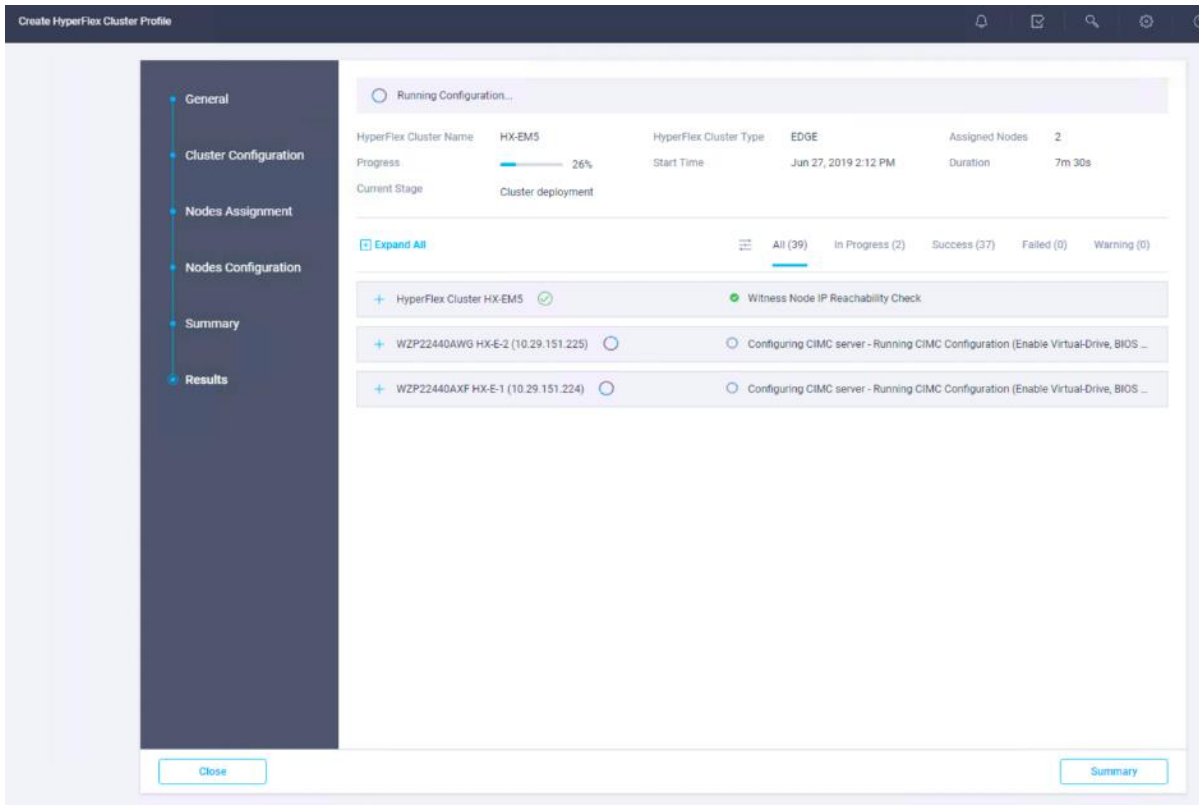
- Click Next to navigate to the Nodes Configuration page. Check the node configuration for both HyperFlex Edge nodes. You may freely modify the hostname of automatic IP address assignment if desired. Enter the cluster management IP address within the same IP management subnet.



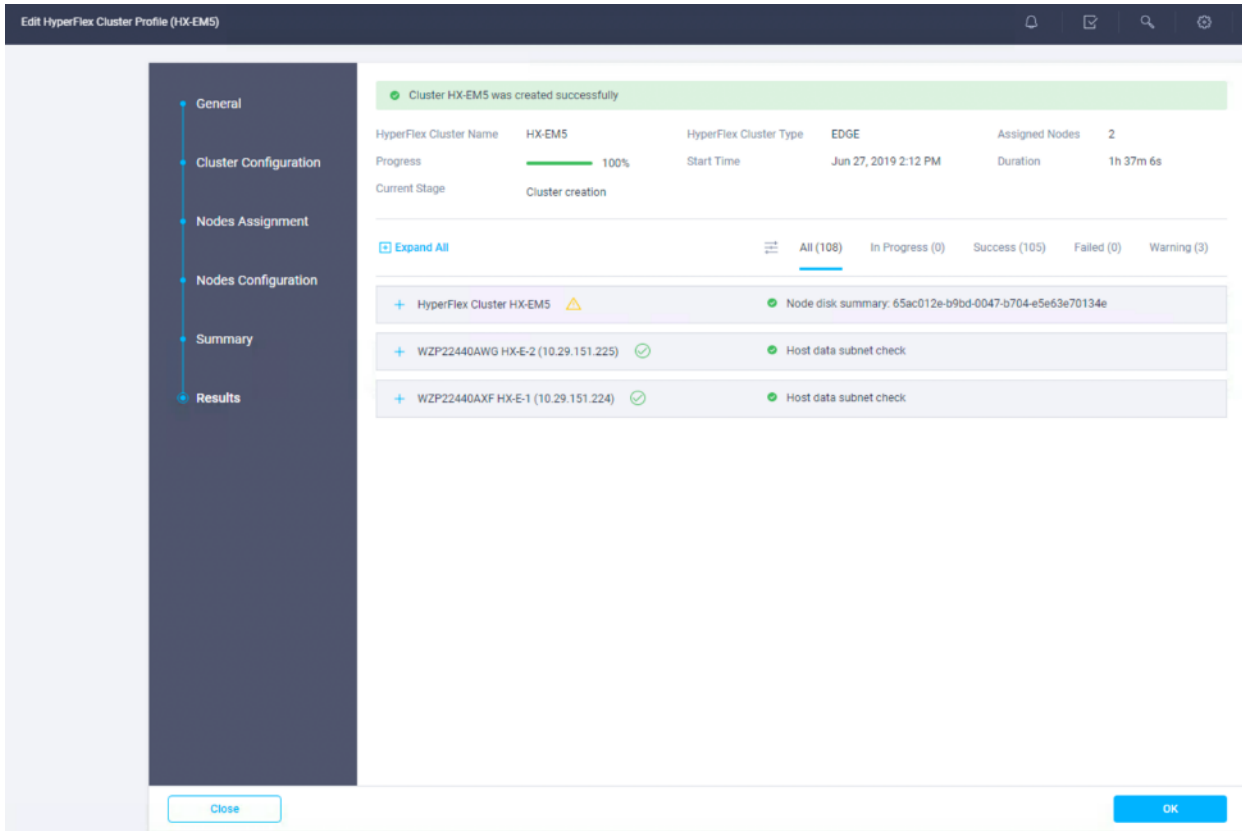
26. Click Next to the Summary page. Review the Cluster Configuration and Nodes Configuration. Check if there are any errors. Ignore the warnings about RF=2 and Intersight continued connectivity.




- Click Validate to validate the HyperFlex Edge cluster configuration only without starting deployment. This will start a series of hardware, software, and environmental checks that will take a few minutes to complete. Alternatively, click Validate & Deploy to complete validation and deployment together. This document follows the path of performing Validate & Deploy at one step.



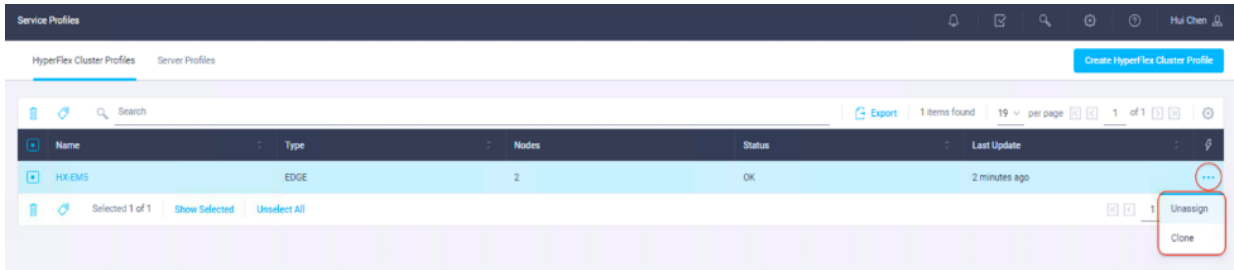
28. Optionally, you can click Save and Close to complete deployment later. Installation time will vary based on network bandwidth, but typically takes about one hour. You can remain on the results page to watch cluster deployment progress in real time. Alternatively, you may click Close to send the task into the background and navigate elsewhere with in Intersight. To return to this results view, navigate back to the Service Profiles > HX Cluster Profile list view and select the cluster name.



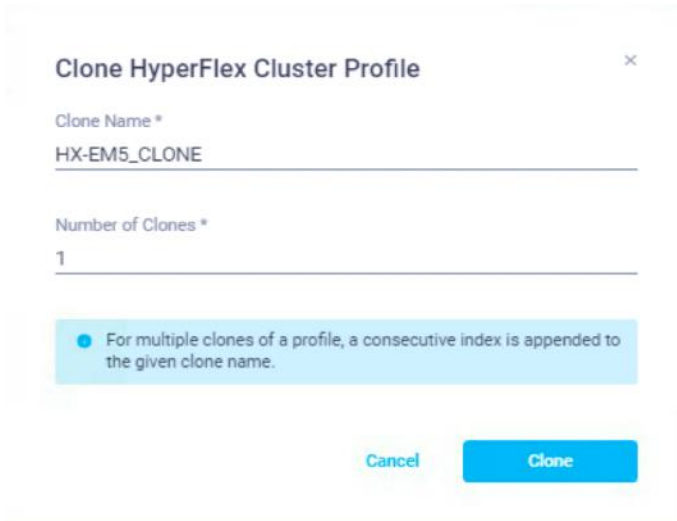
29. Once deployment has completed successfully, click OK
30. Once back on the Service Profiles > HX Cluster Profile page, find the newly deployed HX cluster profile with a status of OK.



31. You can clone & modify the HX Cluster Profile to quickly create many new cluster profiles. To clone a Cluster Profile, select the profile for cloning, under the Settings column, click ..., and then click Clone.



32. In the pop-up Clone window, enter the name for the cloned cluster and choose the number of clones. Click Clone.

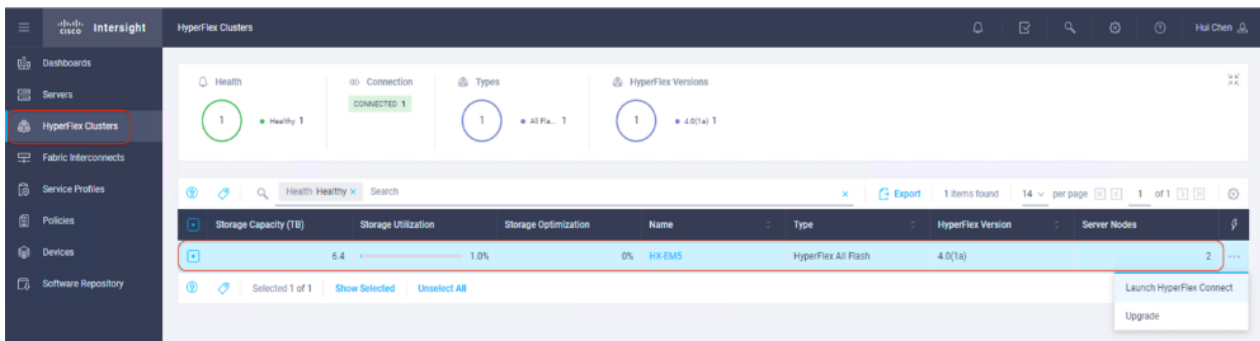


- 33. Before deploying the new copy of the cloned Cluster Profile, you must review the configuration, and make any necessary changes to the policies, or create the new policies for your new needs.
- 34. To un-assign the servers from the selected Cluster Profile, under the Settings column, click ..., and then click Unassign and Unassign again.

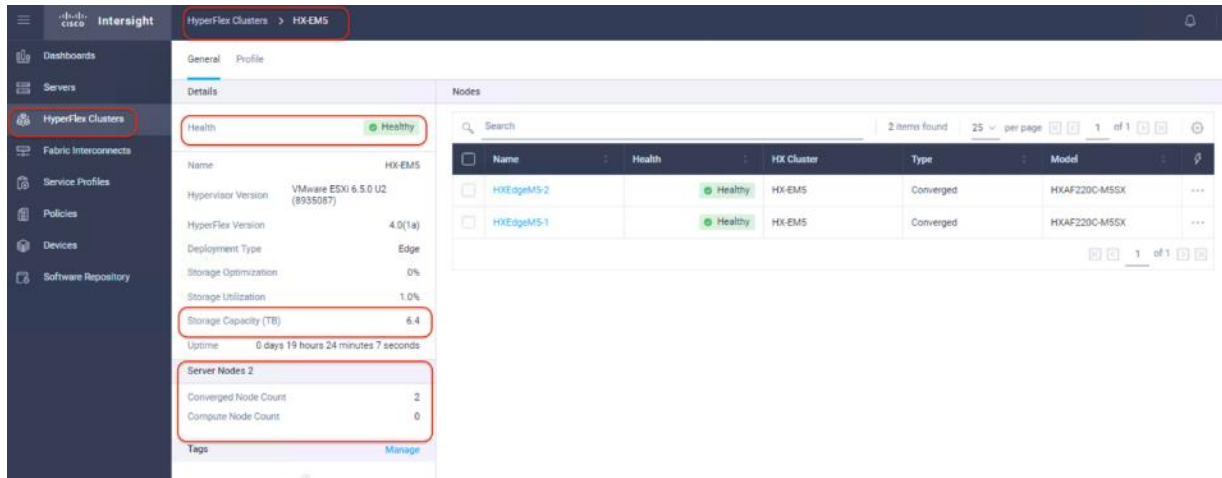


Note: Unassigning will remove all assigned servers from the Cluster Profile. This action is specific to Intersight and will not result in a factory reset or affect the operation of the running cluster. To fully clean up this cluster from Intersight, please un-claim the HyperFlex cluster and associated servers. Contact Cisco TAC for RMA support or cluster re-installation procedures.

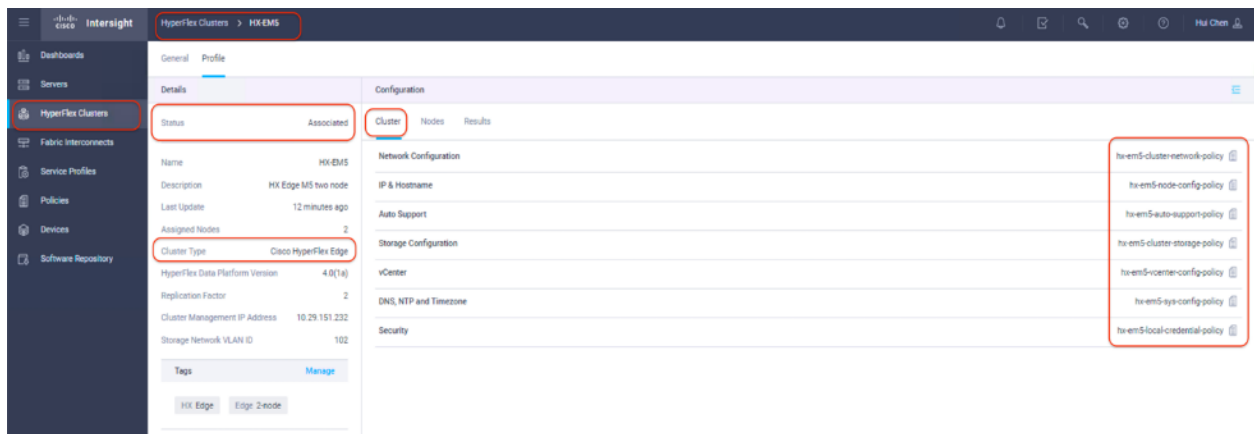
35. Verify the HyperFlex Edge cluster information in Intersight. In the navigation pane, go to the HyperFlex Clusters page, then click the hyperlinked cluster name.



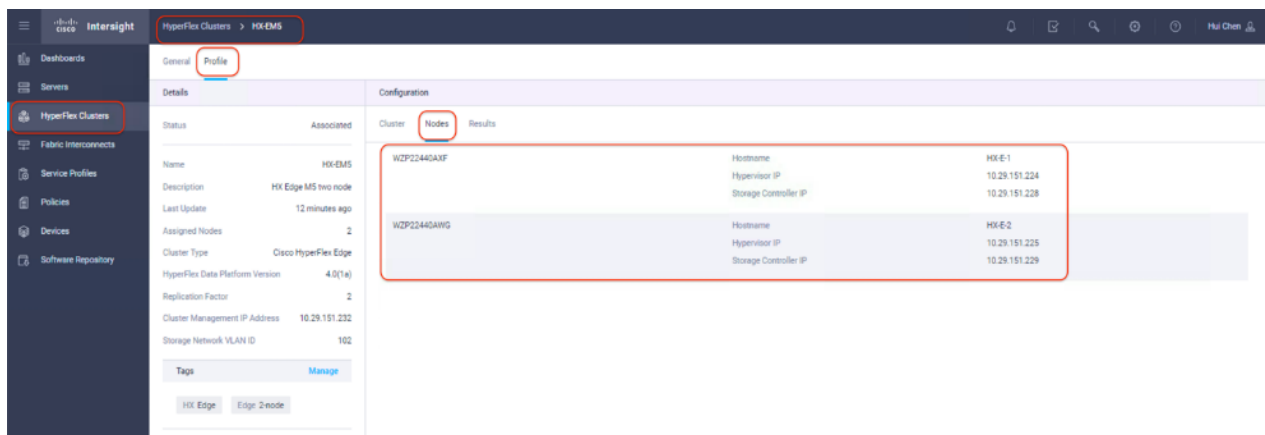
- On the HyperFlex Cluster page, under General tab, review the summary information about the new cluster including the health state, storage capacity, node count, and so on. Check if there are any critical alarms requiring your attention.



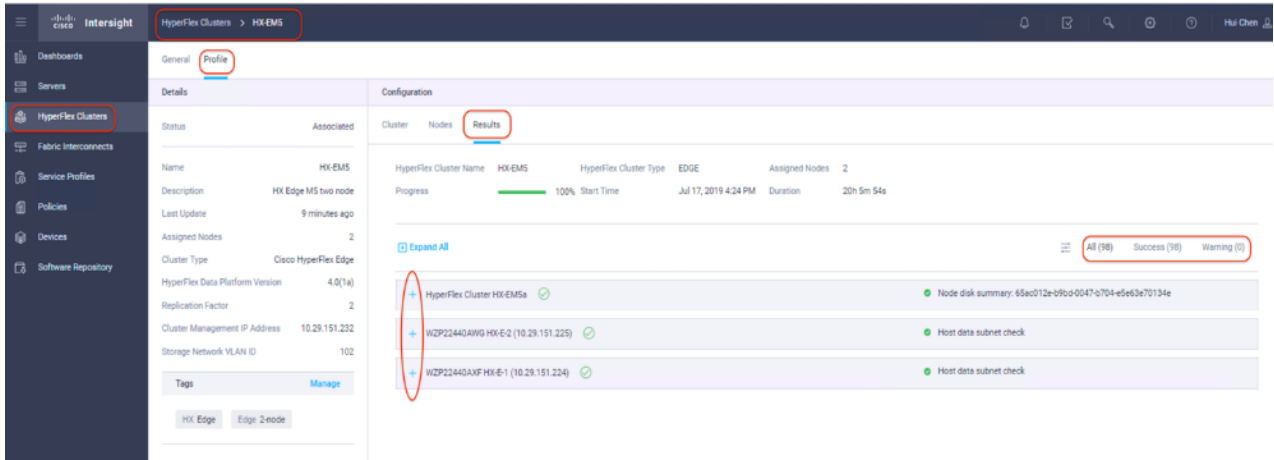
- On the same page, click the Profile tab, allowing you to review the newly created HX Cluster Profile information. Under Cluster Configuration, the details of the policies you just configured can be reviewed.



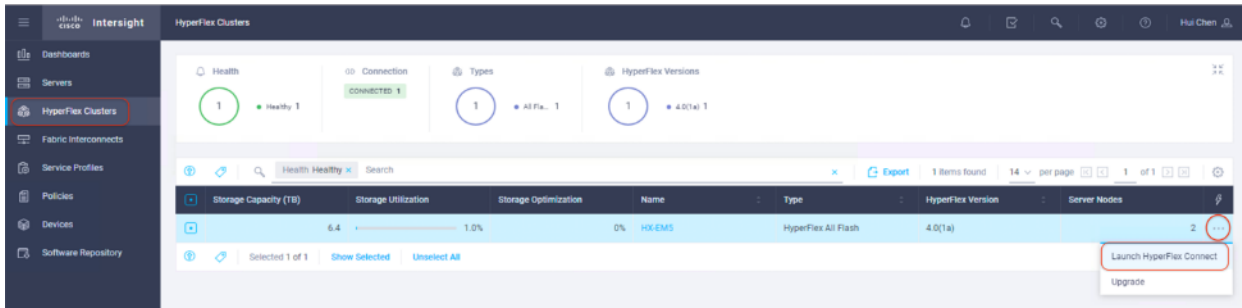
- Review the Nodes Configuration.



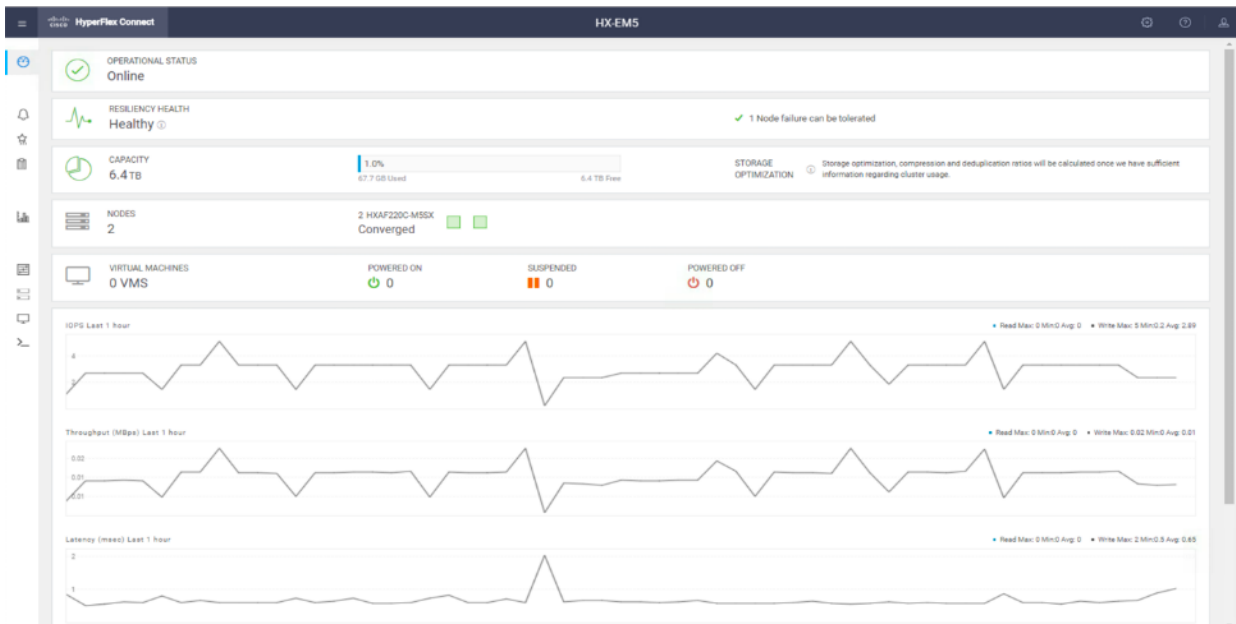
- And review the details of the installation process under Results. Click + at the top of each panel to expand to the detailed information.



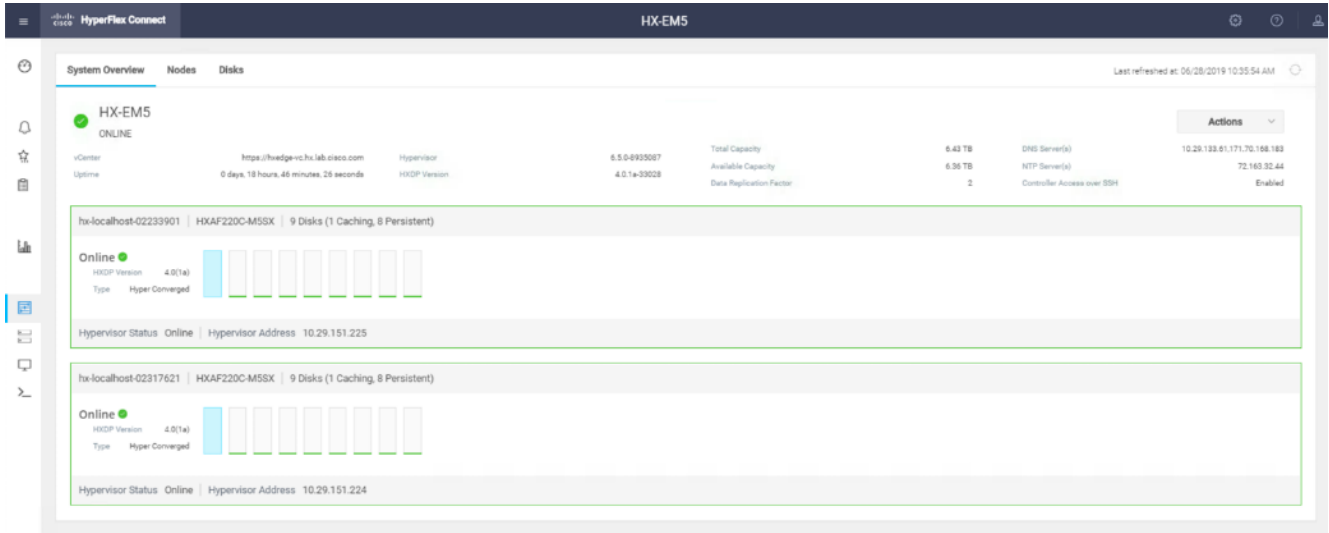
40. You can review a more detailed status of the new HyperFlex Edge cluster in the HyperFlex Connect management GUI. HyperFlex Connect can be directly launched from Intersight. To find the cross-launch feature, first navigate back to the cluster list view by selecting HyperFlex Clusters on the left-hand navigation pane.



41. Then find the desired cluster from the list, under the Settings column, click ..., and then click Launch HyperFlex Connect. The native HyperFlex Connect UI will launch in a new browser tab. All HyperFlex functions are available through cross launch.



- In the HyperFlex Connect management console, click System Information, verify the status of HX Edge converged nodes.

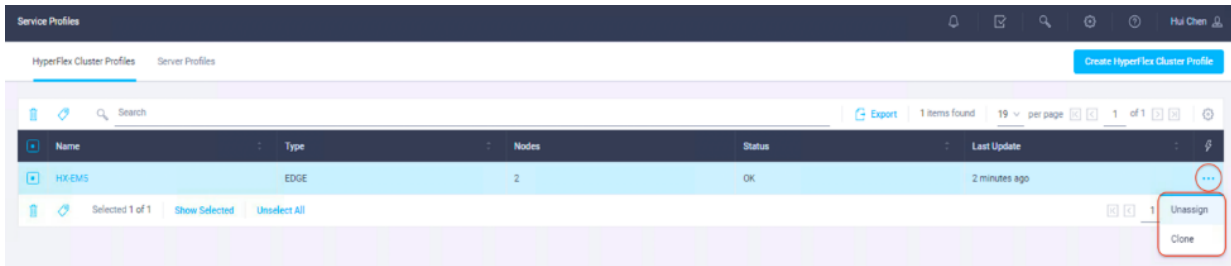


HyperFlex Edge Installation – Multiple Sites

Clone HyperFlex Cluster Profile

To clone the HyperFlex Cluster Profile, follow these steps:

- You can clone and modify the HyperFlex Cluster Profile to quickly deploy another HyperFlex Edge cluster at a different site.



- In the pop-up Clone window, enter the name for the cloned cluster and choose the number of clones. Click Clone.

Clone HyperFlex Cluster Profile

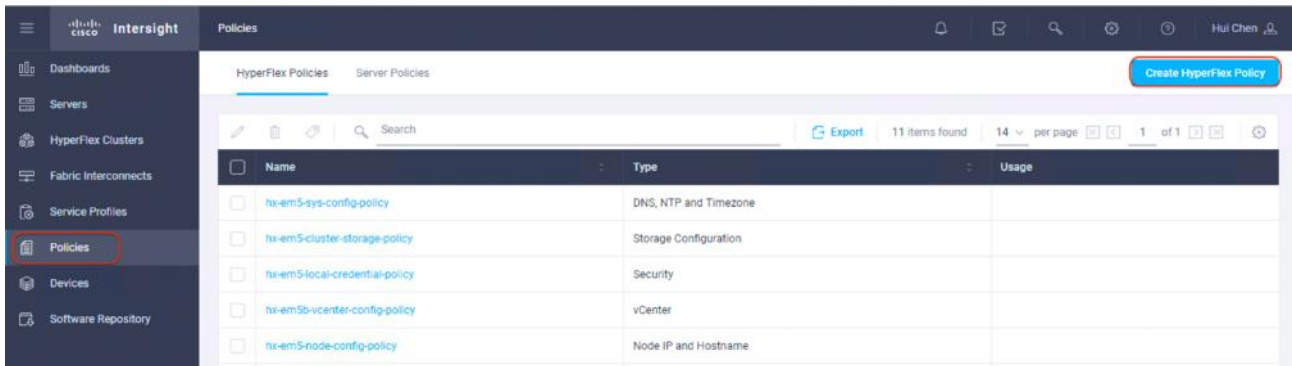
Clone Name *
HX-EM5b

Number of Clones *
1


For multiple clones of a profile, a consecutive index is appended to the given clone name.

Cancel Clone

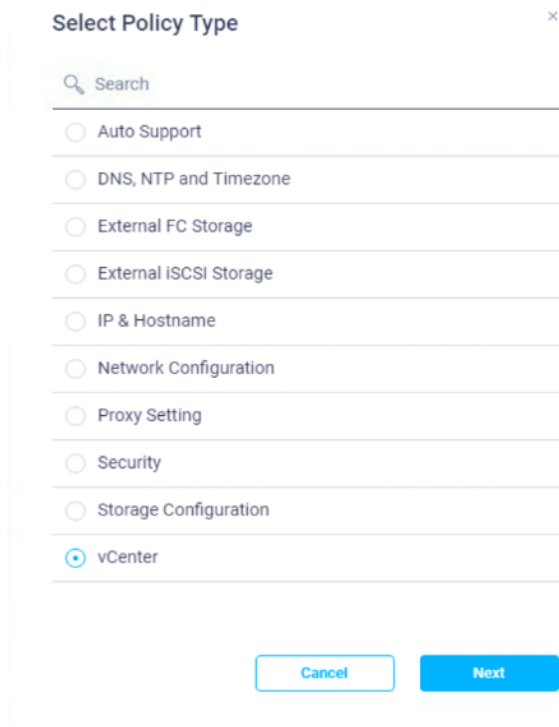
- Before deploying the new copy of the cloned Cluster Profile, you must review the configuration, and make any necessary changes to the policies. If the policies are in use, they cannot be changed so you have to create the new policies for your needs.
- From the left-hand navigation pane of Cisco Intersight, choose Policies. On the Policies page, click Create HyperFlex Policy tab to create a new HyperFlex policy.



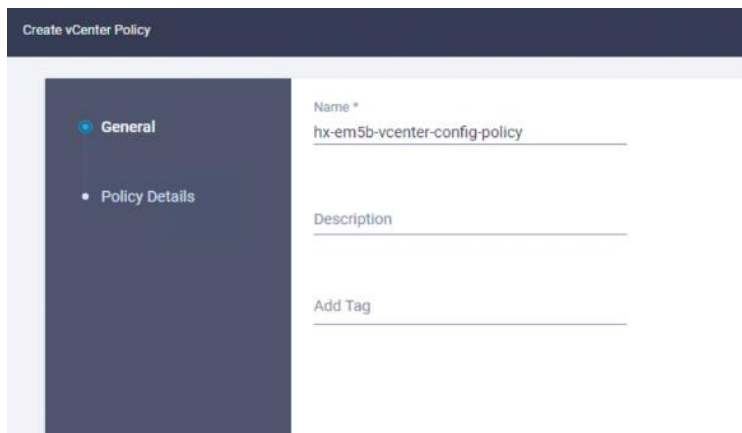
- From the left-hand navigation pane of Cisco Intersight, choose Policies. On the Policies page, click Create HyperFlex Policy tab to create a new HyperFlex policy.



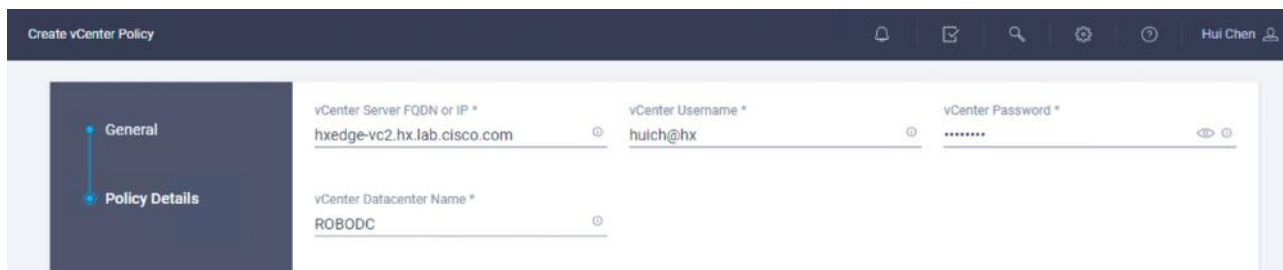
- Choose vCenter as Policy Type, click Next.



7. Enter a policy name, click Next. It is optional to add Description and Tag here.

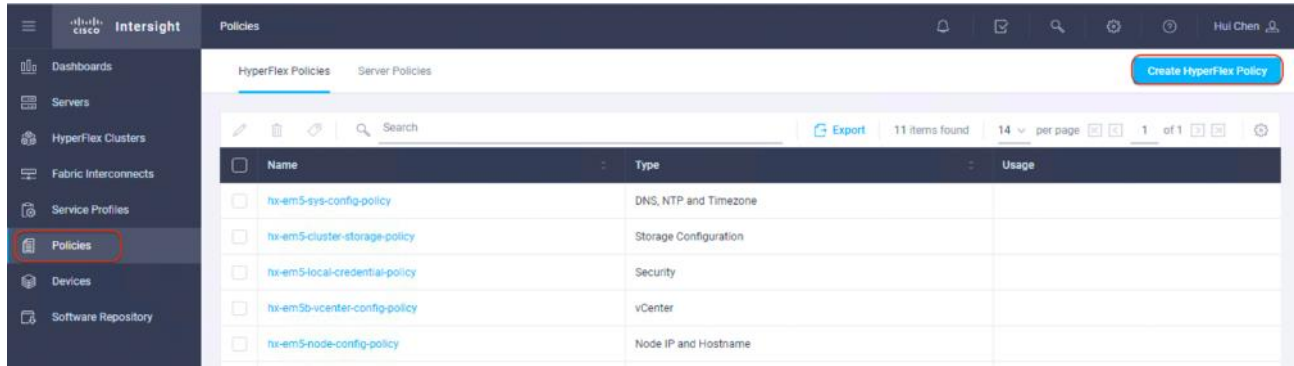


8. Enter the vCenter server FQDN or IP address, the administrative username and password. Enter the Datacenter name in vCenter hosting the HyperFlex Edge cluster.

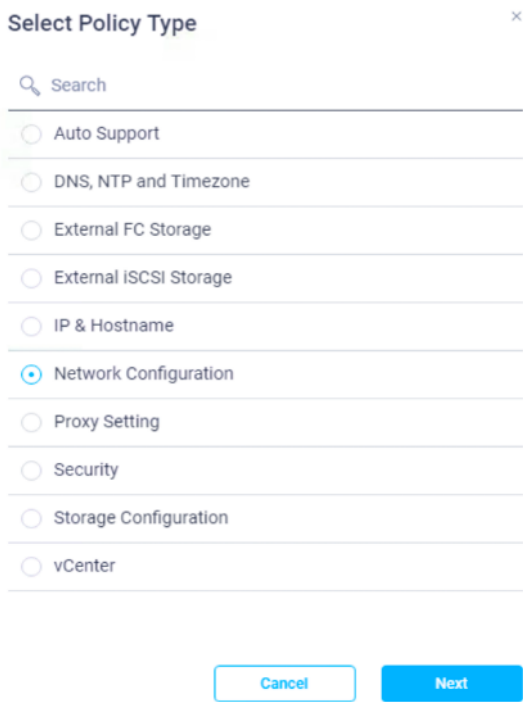


9. Click Create to continue. This will create a new vCenter policy, e.g. hx-em5b-vcenter-config-policy.

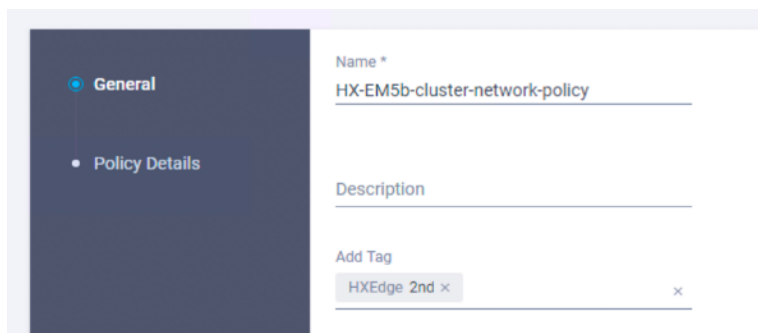
- From the left-hand navigation pane of Cisco Intersight, choose Policies. On the Policies page, click Create HyperFlex Policy tab to create a new HyperFlex policy.



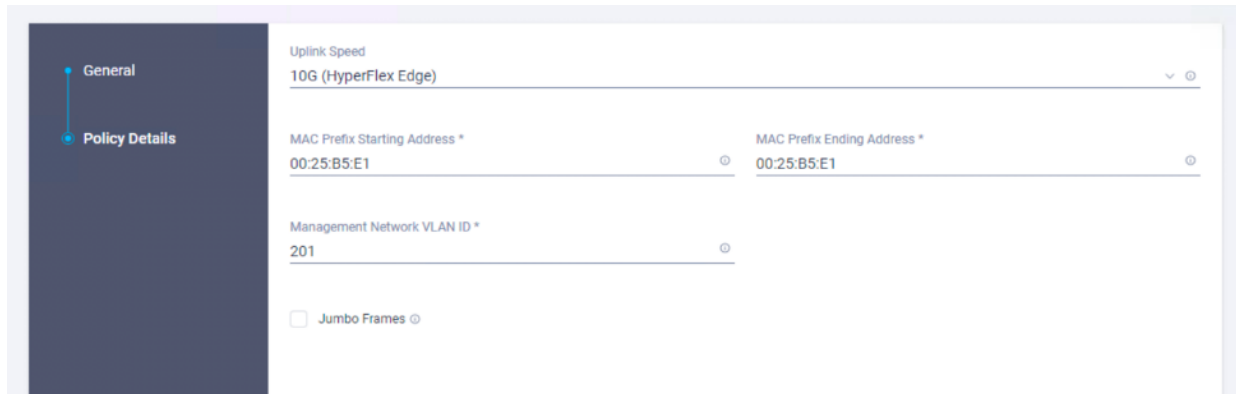
- Choose Network Configuration as Policy Type and click Next.



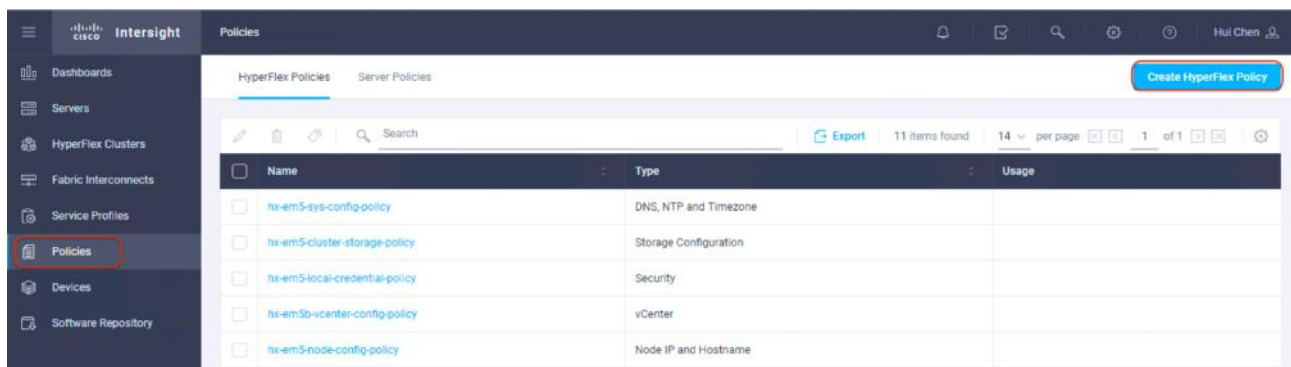
- Enter a policy name and click Next. It is optional to add a Description and Tag.



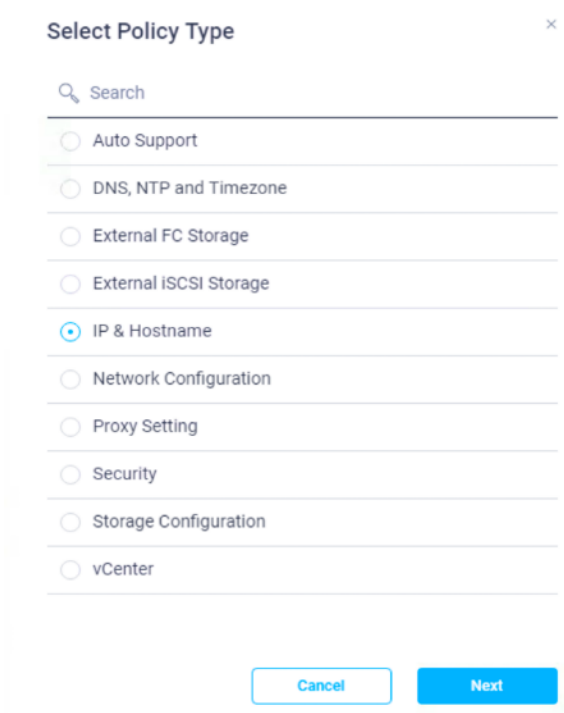
13. Select the Uplink Speed for the upstream network topology. There are two options of selection for HyperFlex Edge, 1GE or 10GE. Choose 10GE for this deployment.
14. Enter the MAC Prefix Starting and Ending addresses. Enter the VLAN ID for the management network. Leave the Jumbo Frames checkbox unchecked for HyperFlex Edge deployments.



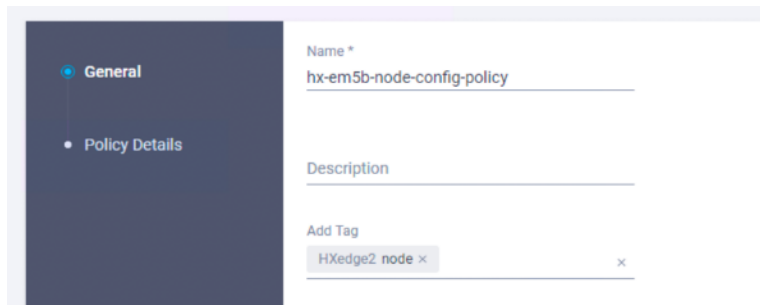
15. Click Create to continue. This will create a new cluster network policy, e.g. HX-EM5b-cluster-network-policy.
16. From the left-hand navigation pane of Cisco Intersight, choose Policies. On the Policies page, click Create HyperFlex Policy tab to create a new HyperFlex policy.



17. Choose IP & Hostname as Policy Type, click Next.



18. Enter a policy name and click Next. It is optional to add Description and Tag here.



19. Enter a Hostname Prefix. In a later step, hostnames will be sequentially assigned to hosts using this prefix. Enter the starting IP address, ending IP address, subnet mask and gateway for the management IP pool. IPs from this range will be automatically assigned to hosts on the node configuration step. Enter starting and ending IP addresses, netmask and gateway for the HX Controller VM management IP pool. Note these two IP ranges must fall within the same IP subnet and VLAN.

20. Click Create to continue. This will create a new cluster network policy, e.g. HX-em5b-node-config-policy.

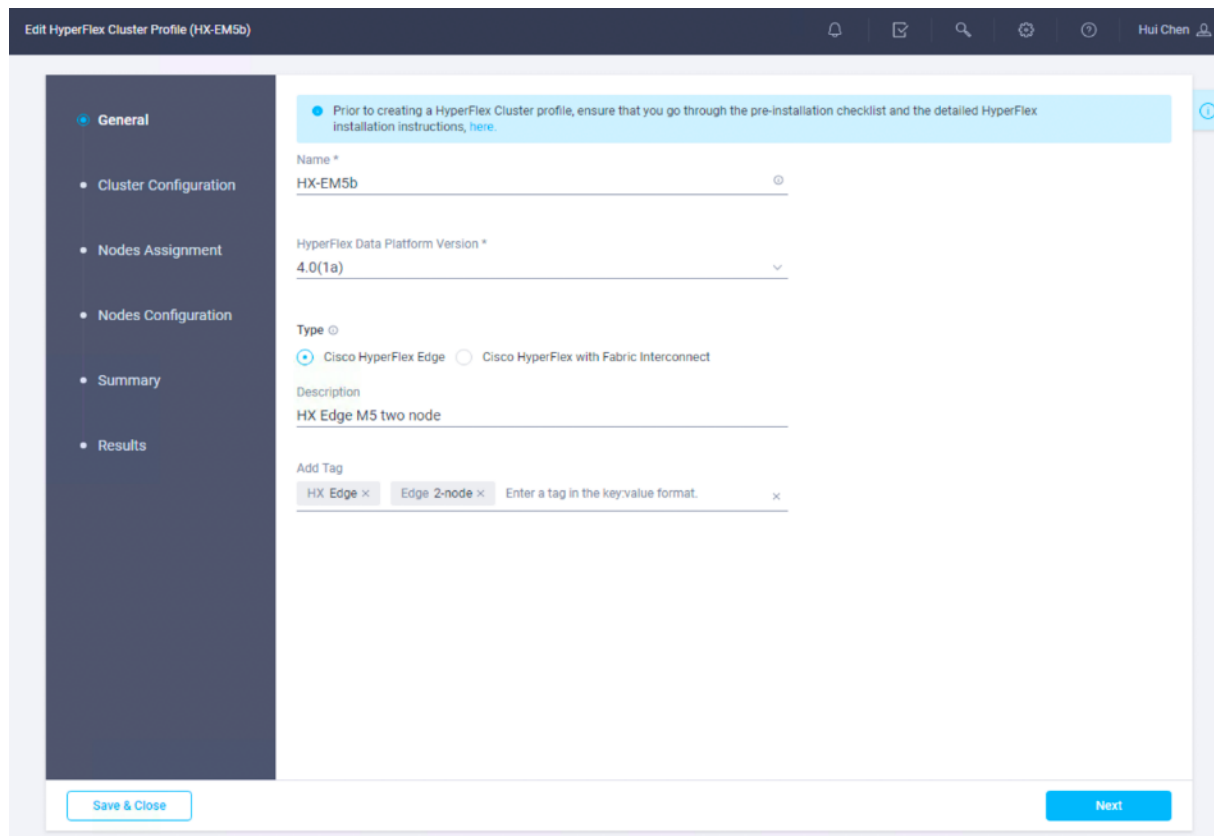


Note: Cloning of HyperFlex policies is not supported with Cisco Intersight in the current release. The support will be available in the future release.

21. From the left-hand navigation pane, choose Service Profiles. On the Service Profiles page, select the cloned (at Step 2) HyperFlex Cluster Profile. Click on the name, for example, HX-EM5b.

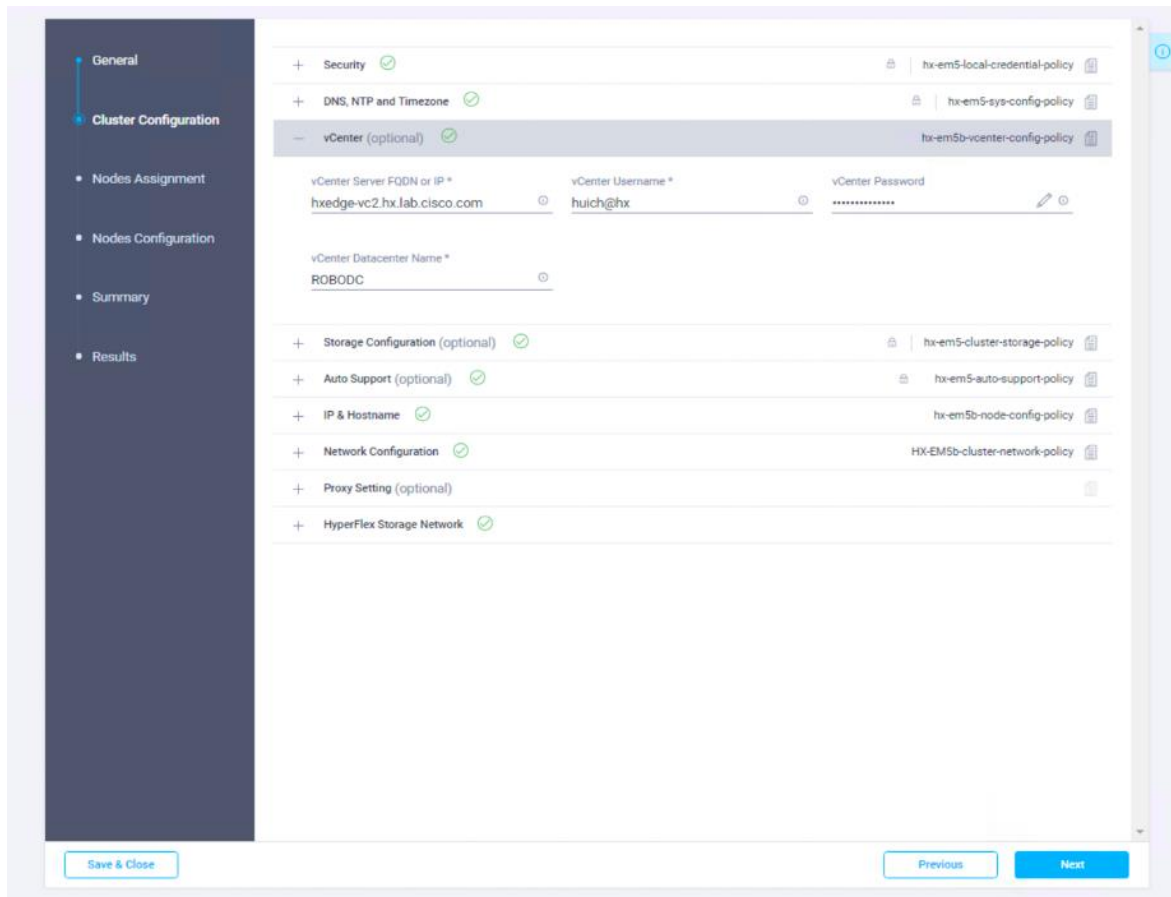
Name	Type	Nodes	Status	Last Update
HX-EM5a	EDGE	0	Not Assigned	Jul 26, 2019 1:52 PM
HX-EM5b	EDGE	0	Not Assigned	Jul 24, 2019 9:45 AM
HX-EM5	EDGE	0	Not Assigned	Jul 16, 2019 4:20 PM

22. The Edit HyperFlex Cluster Profile wizard is displayed. On the General page, you can keep the input intact. Optionally you can modify Description and Tags if needed.

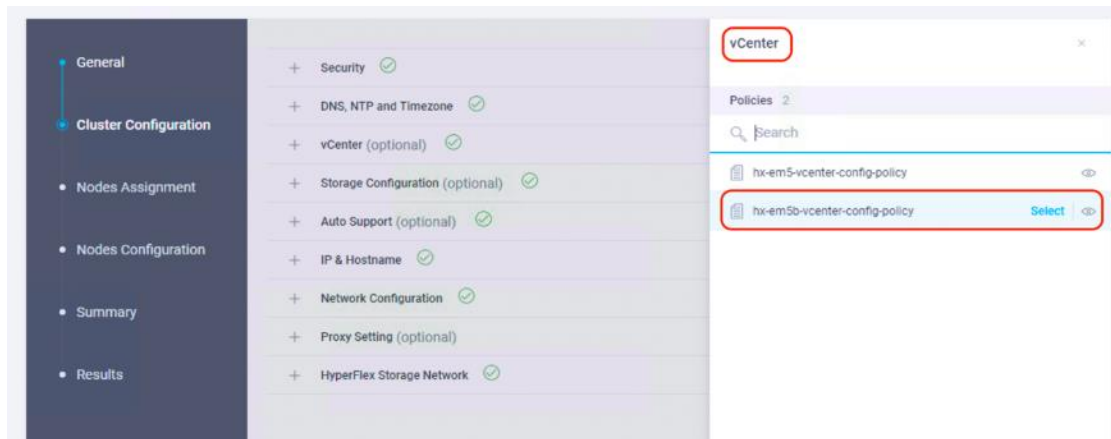


23. Click Next to the Cluster Configuration page.

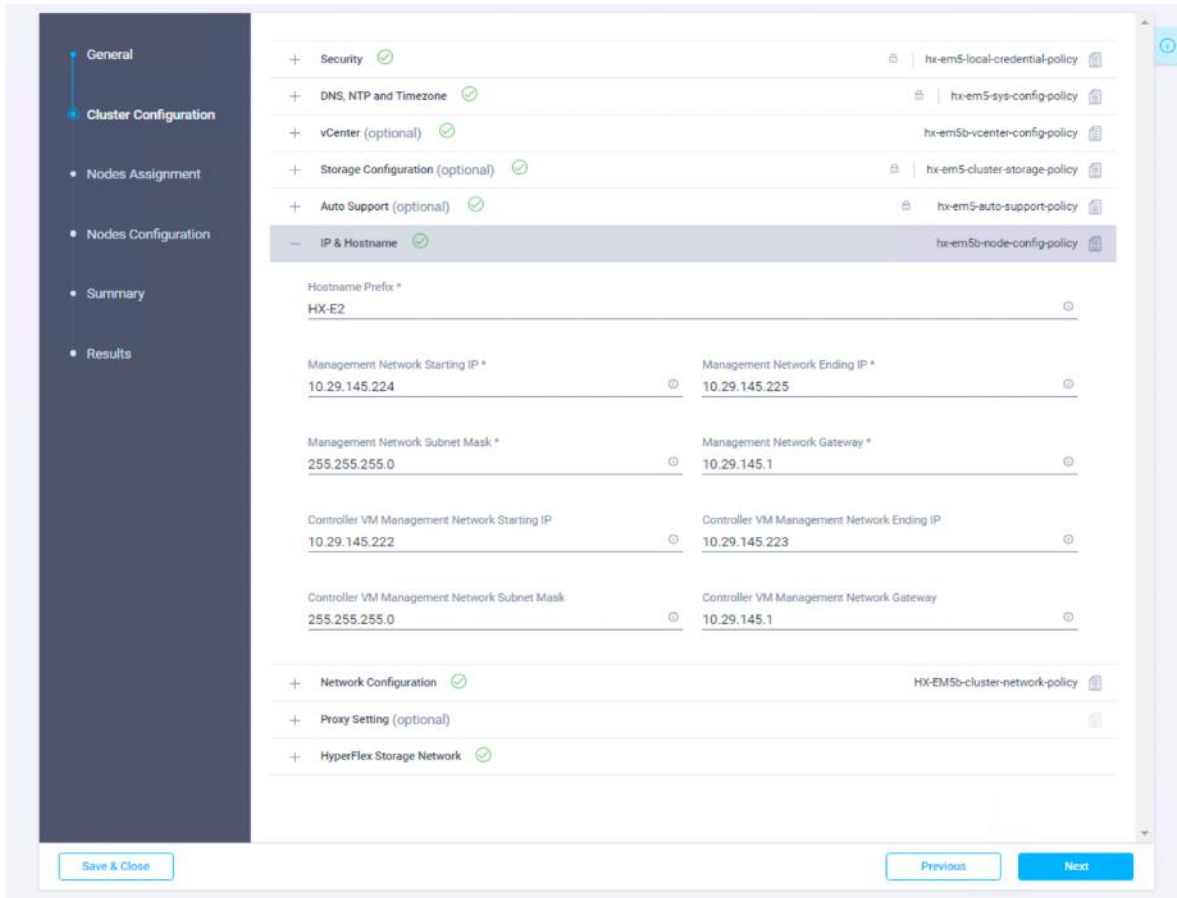
24. Choose the newly created vCenter policy for the new cluster, for example, hx-em5b-vcenter-config-policy.



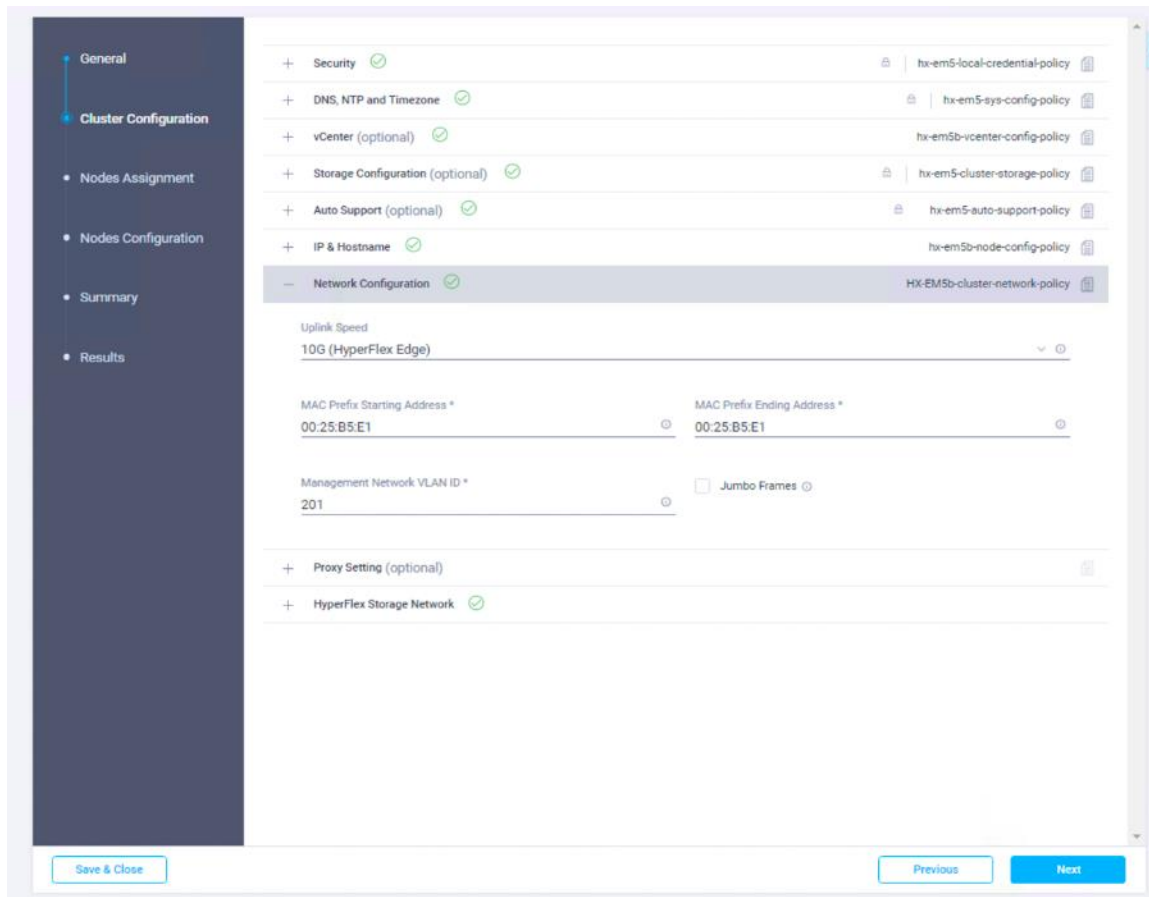
25. To choose an existing policy for your cluster profile, at the policy line, click Select Policy icon to choose the desired policy from the available policy list and click Select.



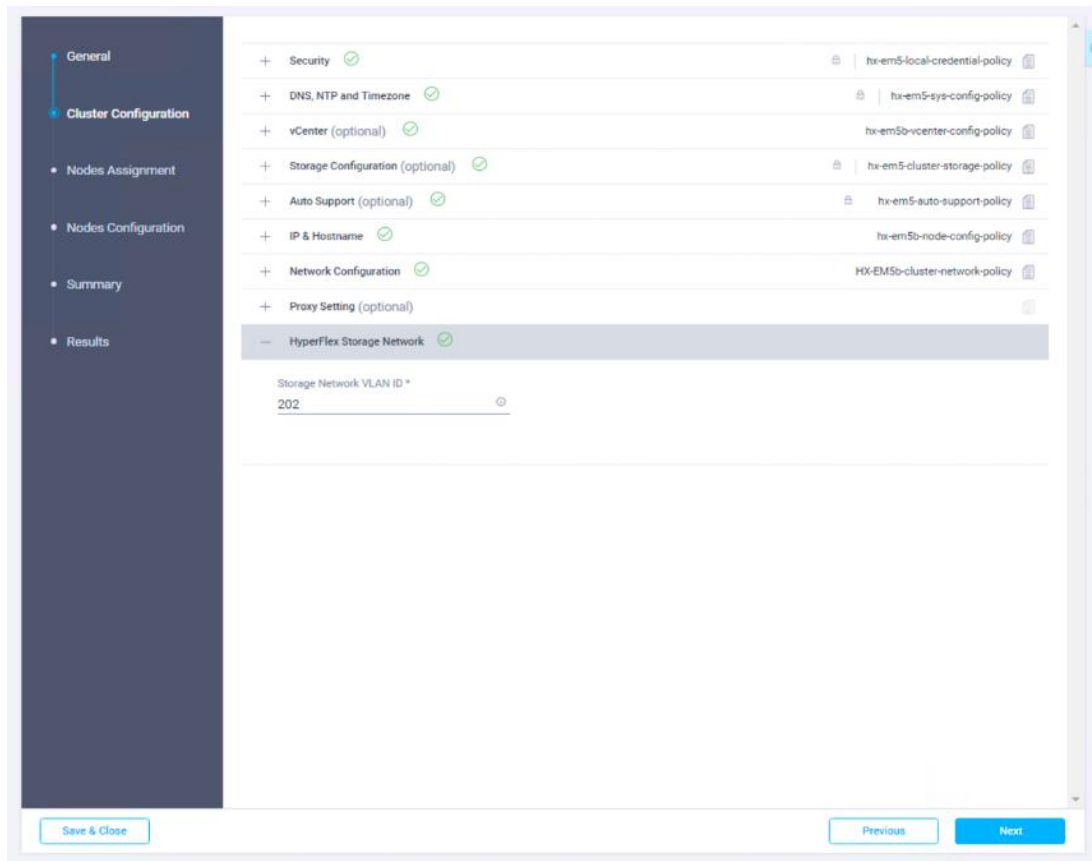
26. Choose the newly created IP & Hostname policy for the new cluster, for example; hx-em5b-node-config-policy.



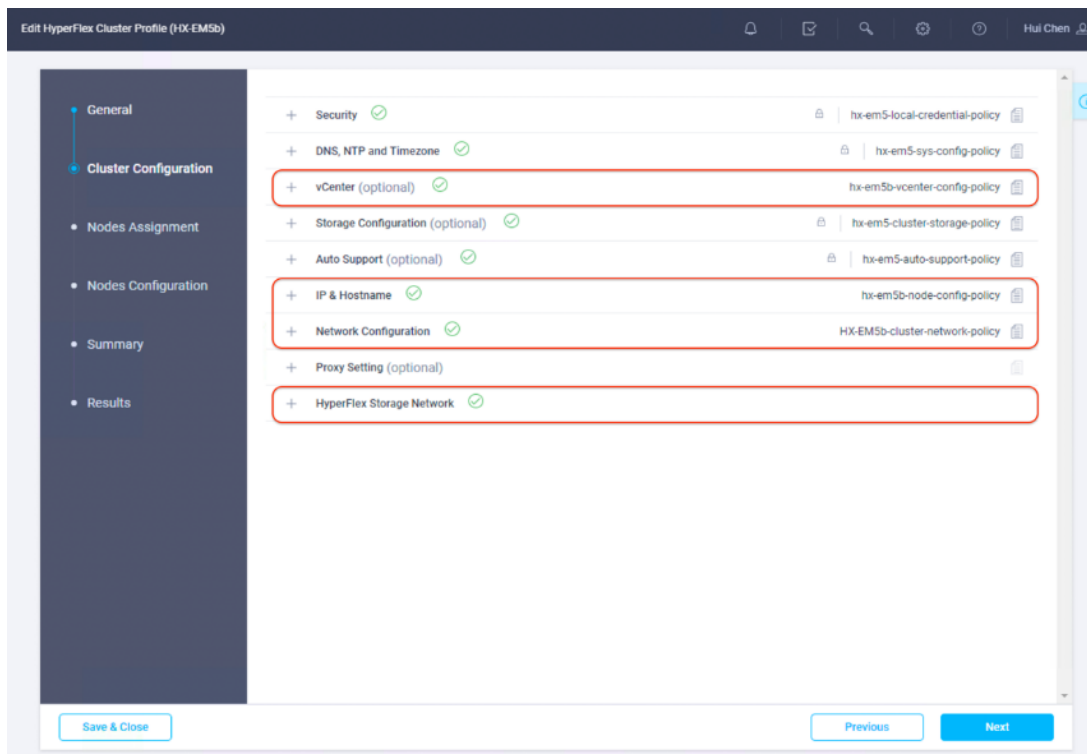
27. Choose the newly created Network policy for the new cluster, for example; HX-EM5b-cluster-network-policy.



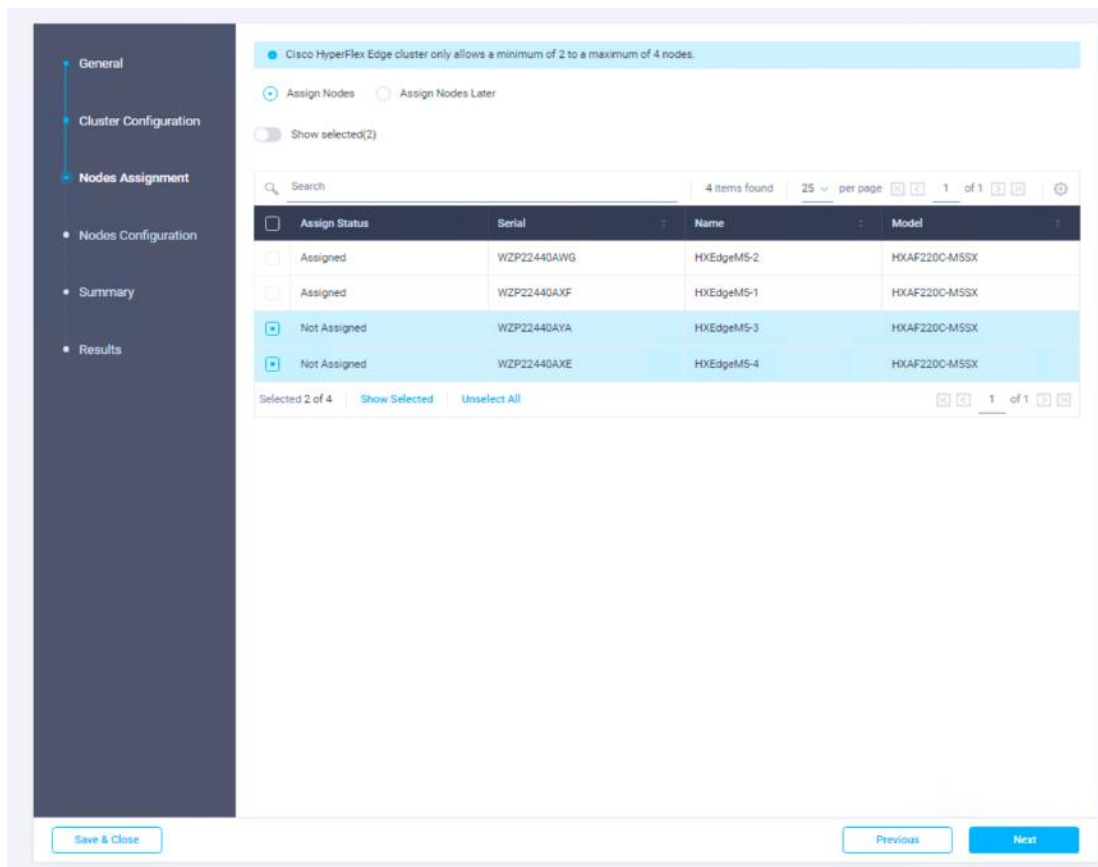
28. Click + to expand HyperFlex Storage Network configuration. Enter the VLAN ID for the data storage network. It is required to use a unique storage VLAN per cluster.



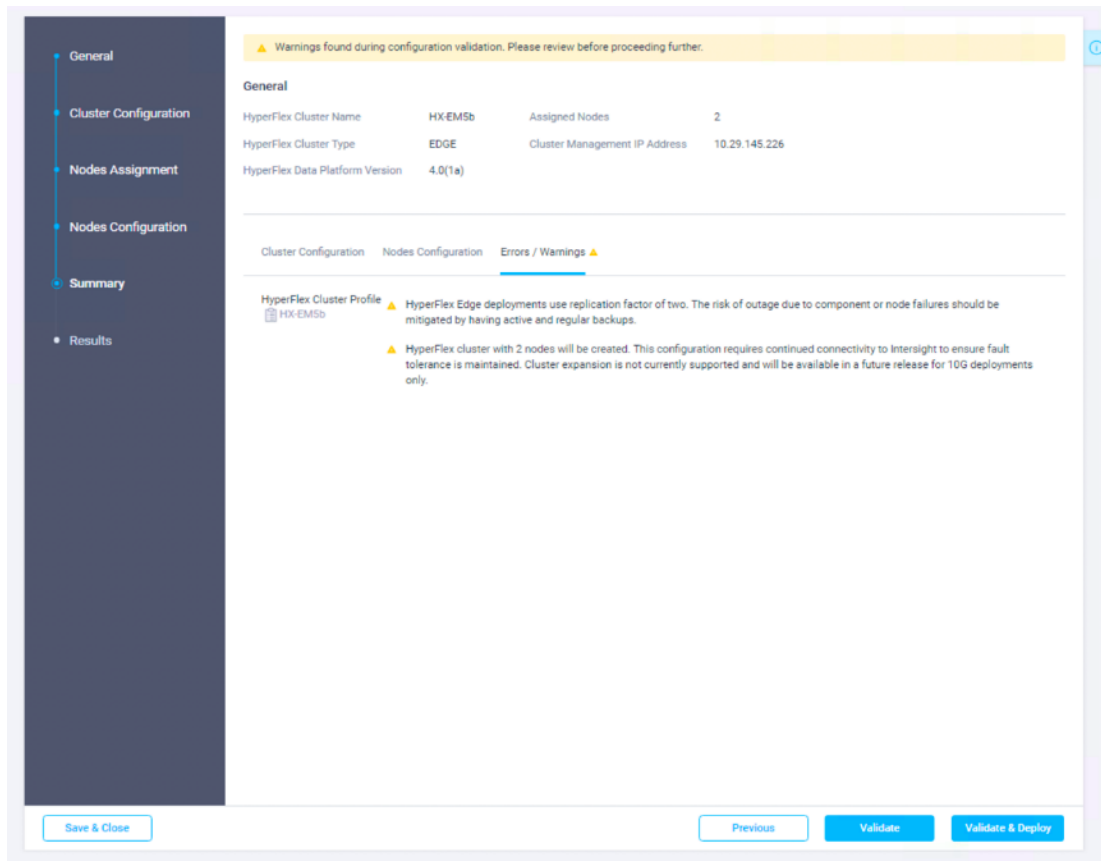
29. Now all the policies are configured and are listed as shown in the screenshot. The policies are mixed with the reusable policies derived from the cloned HyperFlex Cluster Profile, and also the newly created policies. Depending on the resources available for being shared by multiple sites, more new policies might need to be created.



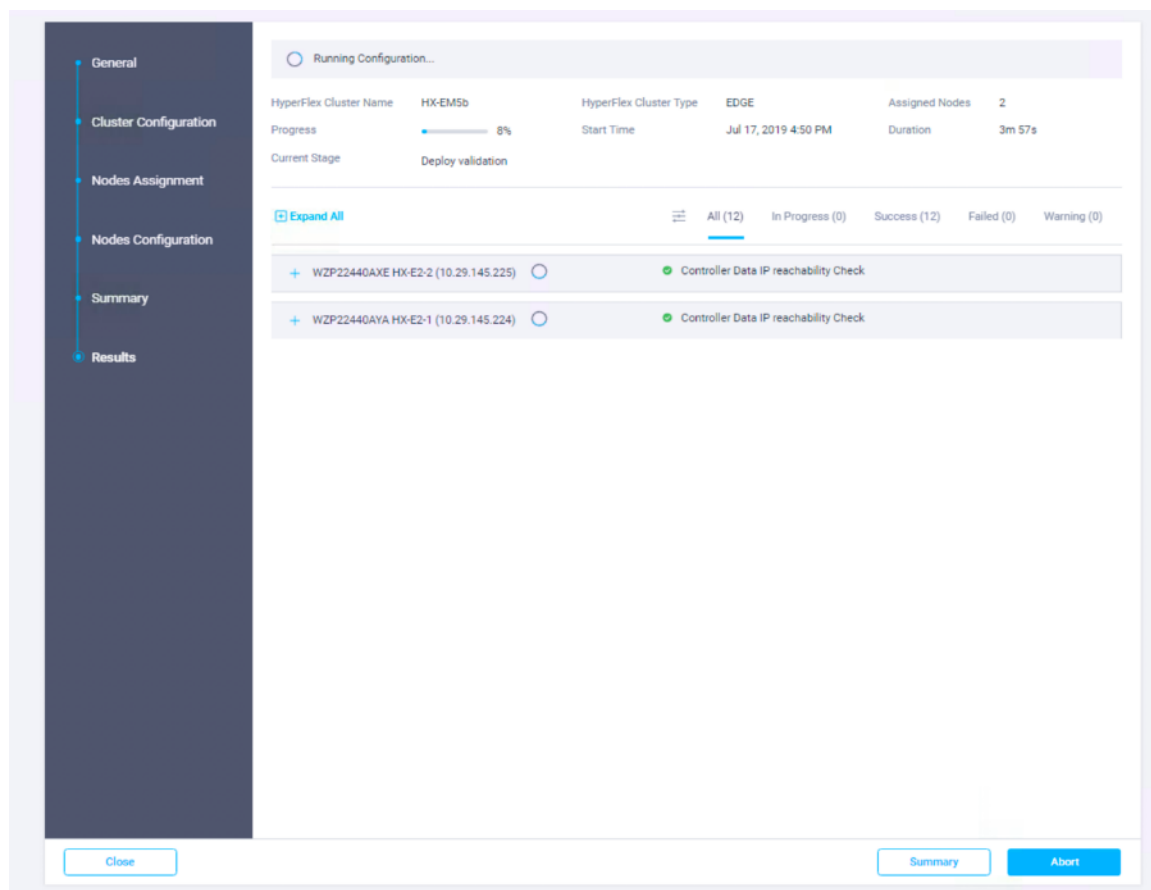
30. Click Next to the Nodes Assignment page. Select two available HyperFlex Edge nodes. You also have the option to Assign Nodes Later to save off the HyperFlex cluster profile now and return when the hardware is available to assign once claimed in Intersight.



31. Click Next to navigate to the Nodes Configuration page. Check the node configuration for both HyperFlex Edge nodes. You may freely modify the hostname of automatic IP address assignment if desired. Enter the cluster management IP address within the same IP management subnet.
32. Click Next to proceed to the Summary page. Review the Cluster Configuration and Nodes Configuration. Check if there are any errors. Ignore the warnings about RF=2 and Intersight continued connectivity.



33. Click Validate & Deploy to complete validation and deployment of the second HyperFlex Edge cluster together.



34. Wait until the deployment has completed successfully, click OK.

Intersight API

Cisco Intersight provides a cloud-based RESTful API to manage Cisco UCS and HyperFlex systems across multiple Data Centers. The Intersight API is a programmatic interface that uses the REST architecture to provide access to the Intersight Management Information Model. API requests may be read-only query with no side-effects or produce modifications of the resources. The response may confirm that some alteration has been made to the resource and it may provide hypertext links to related resources or collections of resources. The Intersight API accepts and returns messages that are encapsulated through JavaScript Object Notation (JSON) documents and uses HTTP over TLS as the transport protocol. REST API is a feature that requires Intersight Essentials license.

Cisco Intersight provides the benefits of cloud-based management that customers have come to appreciate with Software-as-a-Service (SaaS) products. For example, the Intersight API is automatically updated when new features are deployed to the cloud, providing programmatic access to new IT infrastructure capabilities.

The Intersight API is based on the OpenAPI standard, which defines a programming language-agnostic interface description for describing, producing, consuming, and visualizing RESTful Web services. The OpenAPI specification for Intersight allows both humans and computers to discover and understand the capabilities of the service without requiring access to source code, additional documentation, or inspection of network traffic.

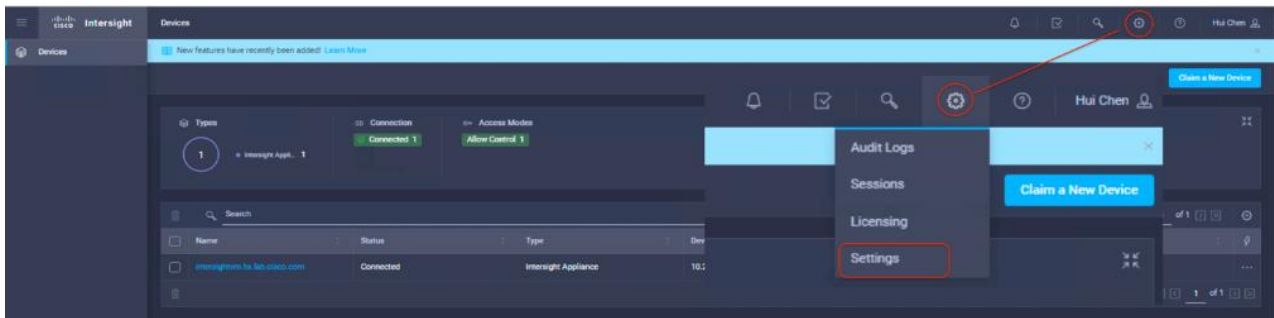
Intersight provides a downloadable Software Development Kit (SDK) for the Python programming language and PowerShell scripting. You can generate SDKs for other programming languages using the open-source OpenAPI tools. The link below provides access to download the SDKs and other Resources:

<https://www.intersight.com/apidocs/downloads/>


With Intersight API, you can easily automate many management workflows including automatically generating multiple HyperFlex cluster profiles in Intersight from a centralized data file. An example for multi-site HyperFlex Edge deployments using Intersight Python SDK is shown in Appendix E. The open source Python codes and instructions are available to download from the GitHub web site.

You must generate the required API Keys before you run the program calling the Intersight API, and then proceed to download the Intersight SDK. An API key is used to register your application with Cisco Intersight. As a Cisco Intersight user, to generate and manage API keys, follow these steps:

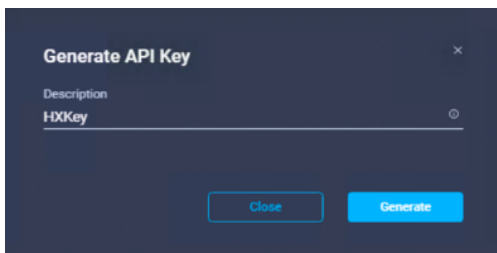
1. From Cisco Intersight, select Settings and click the Settings menu.



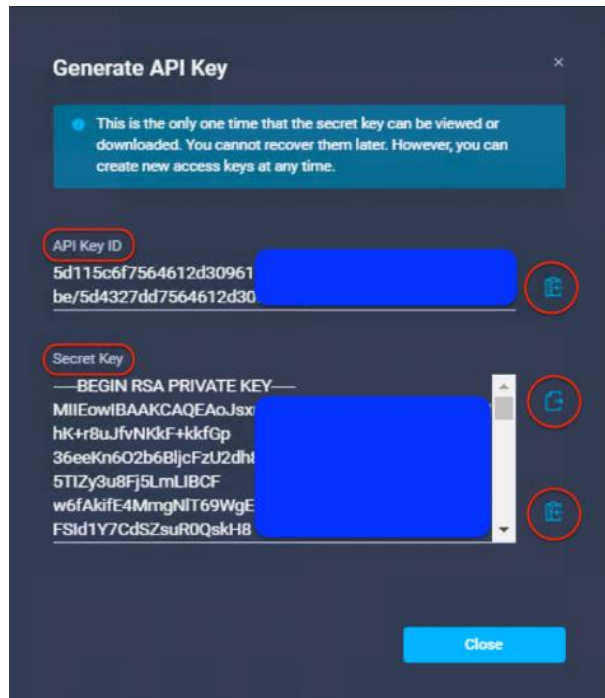
2. In the General page, click API -> API Keys -> Generate API Key.



3. In the Generate API Key screen, enter the Description for the API Key and click Generate.



4. The API Key ID and RSA Private Key are displayed.



5. Save the private Key ID and Secret Key information in files and save them to a location accessible from your scripts. You can use Copy and Paste method to obtain the key information in an accurate way. There is a Save button next to the Secret key. By clicking on that you can quickly save the Secret Key into a file.

Post-install Configuration

Prior to putting a new HyperFlex Edge cluster into production, a few post-install tasks must be completed. To automate the post installation procedures and verify the HyperFlex cluster configuration, a `post_install` script has been provided on the HyperFlex Controller VMs. These steps can also be performed manually or with a Power CLI script in vCenter if preferred (a sample script is shown in Appendix F). The following procedure will use the `post_install` script. To run this script, follow these steps:

1. SSH to the cluster management IP address and login using `<root>` username and the controller VM password provided during installation. Verify the cluster is online and healthy using `"stcli cluster info"` or `"stcli cluster storage-summary."`

```

root@SpringpathController6K0AYKDTI4:~# stcli cluster storage-summary
address: 169.254.1.20
name: HX-EM5
state: online
uptime: 5 days 0 hours 8 minutes 20 seconds
activeNodes: 2 of 2
compressionSavings: 82.54%
deduplicationSavings: 0.0%
freeCapacity: 6.3T
healingInfo:
  inProgress: False
resiliencyInfo:
  messages:
    Storage cluster is healthy.
  state: 1
  nodeFailuresTolerable: 1
  cachingDeviceFailuresTolerable: 1
  persistentDeviceFailuresTolerable: 1
  zoneResInfoList: None
spaceStatus: normal
totalCapacity: 6.4T
totalSavings: 82.54%
usedCapacity: 81.6G
zkHealth: online
arbitrationServiceState: online
clusterAccessPolicy: lenient
dataReplicationCompliance: compliant
dataReplicationFactor: 2

```

2. Run the following command in the shell and press enter:

```
/usr/share/springpath/storfs-misc/hx-scripts/post_install.py
```

```

root@SpringpathControllerYP403DRZ9C:~# /usr/share/springpath/storfs-misc/hx-scripts/post_install.py
Select post_install workflow-

1. New/Existing Cluster
2. Expanded Cluster
3. Generate Certificate

Note: Workflow No.3 is mandatory to have unique SSL certificate in the cluster.
By Generating this certificate, it will replace your current certificate.
If you're performing cluster expansion, then this option is not required.

Selection: 1
Logging in to controller localhost
HX CVM admin password:
Getting ESX hosts from HX cluster...
vCenter URL: 10.29.151.36
Enter vCenter username (user@domain): huich@hx
vCenter Password:
Found datacenter ROBODC
Found cluster HX-EM5

post_install to be run for the following hosts:
hx-e-1.hx.lab.cisco.com
hx-e-2.hx.lab.cisco.com

.....

Enter ESX root password:
HX Edge configuration detected
Uplink speed is detected as: 10G
Uplink count is detected as: 2

Enter vSphere license key? (y/n) █

```

3. Select the first post_install workflow type – New/Existing Cluster.

4. Enter the HX Storage Controller VM root password for the HX cluster (use the one entered during the HX Cluster installation).
5. Enter the vCenter user name and password.
6. Enter ESXi host root password (use the one entered during the HX Cluster installation).
7. You must license the vSphere hosts through the script or complete this task in vCenter before continuing. Failure to apply a license will result in an error when enabling HA or DRS in subsequent steps. Enter "n" if you have already registered the license information in vCenter.

```
Enter vSphere license key? (y/n) n
```

8. Enter "y" to enable HA/DRS.

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

9. Enter "y" to disable the ESXi hosts' SSH warning. SSH running in ESXi is required in HXDP 2.6.

```
Disable SSH warning? (y/n) y
```

10. Add the vMotion VMkernel interfaces to each node by entering "y". Input the netmask, the vMotion VLAN ID, and the vMotion IP addresses for each of the hosts as prompted. vMotion will be configured per best practices depending on the choice of single or dual switch configuration. In the case of single switch configuration, a traffic shaper will be automatically applied to the VMkernel port to ensure vMotion does not consume all bandwidth available on the shared uplink port.

```
Add vMotion interfaces? (y/n) y
Netmask for vMotion: 255.255.255.0
VLAN ID: (0-4096) 103
vMotion IP for hx-e-1.hx.lab.cisco.com: 192.168.151.11
Adding vMotion-103 to hx-e-1.hx.lab.cisco.com
Adding vmkernel to hx-e-1.hx.lab.cisco.com
vMotion IP for hx-e-2.hx.lab.cisco.com: 192.168.151.12
Adding vMotion-103 to hx-e-2.hx.lab.cisco.com
Adding vmkernel to hx-e-2.hx.lab.cisco.com
```

11. Add VM network portgroups for guest VM traffic. Enter "n" to skip this step and create the portgroups manually in vCenter. This step will add identical network configuration to all nodes in the cluster. By default, the Intersight HX Edge installer will not create any guest VM port groups.

```
Add VM network VLANs? (y/n) y
Port Group Name to add (VLAN ID will be appended to the name): vmNet
VLAN ID: (0-4096) 104
Adding vmNet-104 to hx-e-1.hx.lab.cisco.com
Adding vmNet-104 to hx-e-2.hx.lab.cisco.com
Add additional VM network VLANs? (y/n) n
```

12. Enter "y" to run the health check on the cluster.

```

Run health check? (y/n) y
Validating cluster health and configuration...

Cluster Summary:
Version - 4.0.1a-33028
Model - HXAF220C-M5SX
Health - HEALTHY
ASUP enabled - False

```

13. A summary of the cluster will be displayed upon completion of the script. Make sure the cluster is healthy.

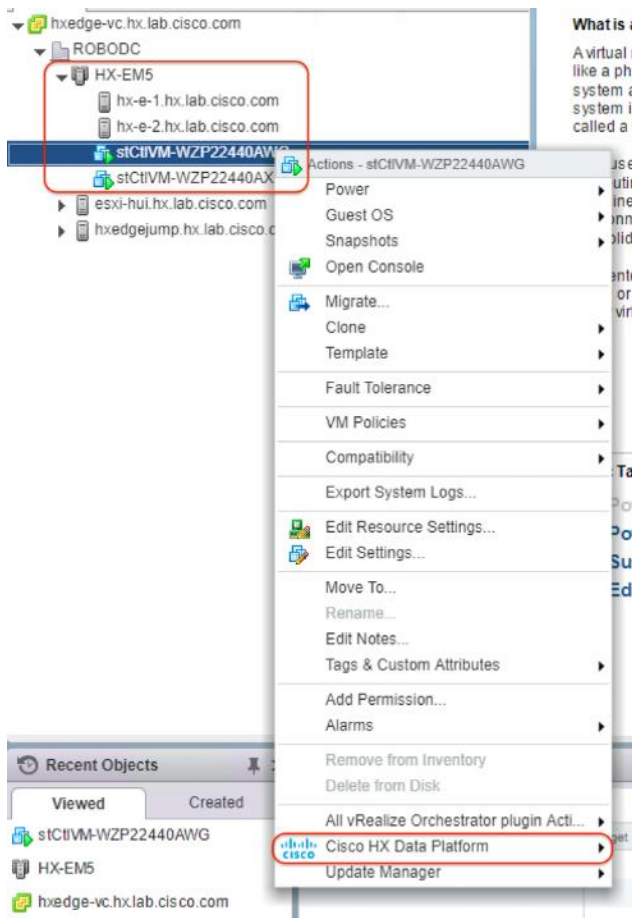


Note: It is recommended to enable a syslog destination for permanent storage of the ESXi host logs. The configuration can be done manually in the vCenter or through a Power CLI script.

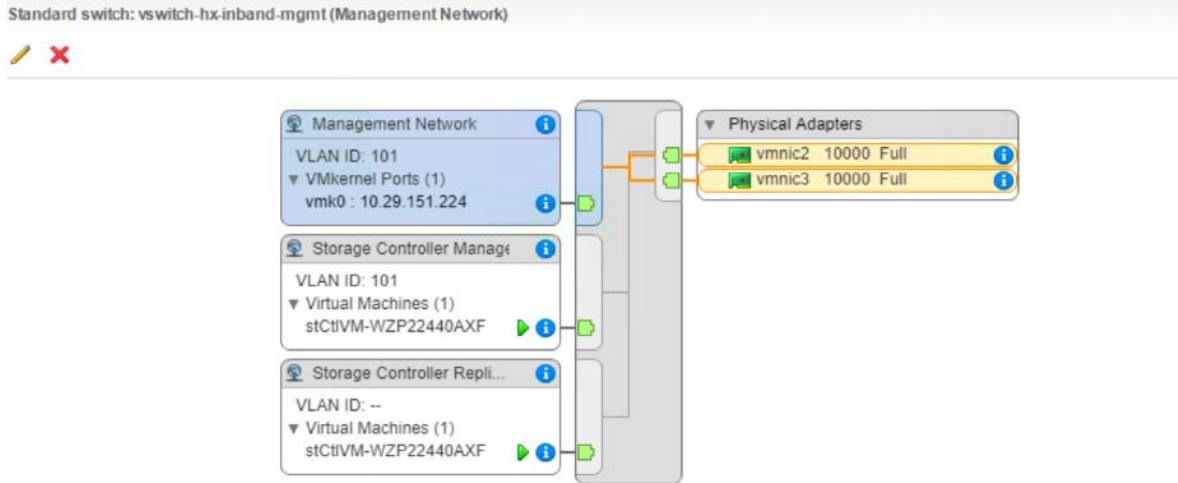
Check ESXi Configuration

After completing the post-install script, log in to the vSphere Web Client to check the cluster status and to verify ESXi host configuration in the vCenter server. To check the ESXi configuration, follow these steps:

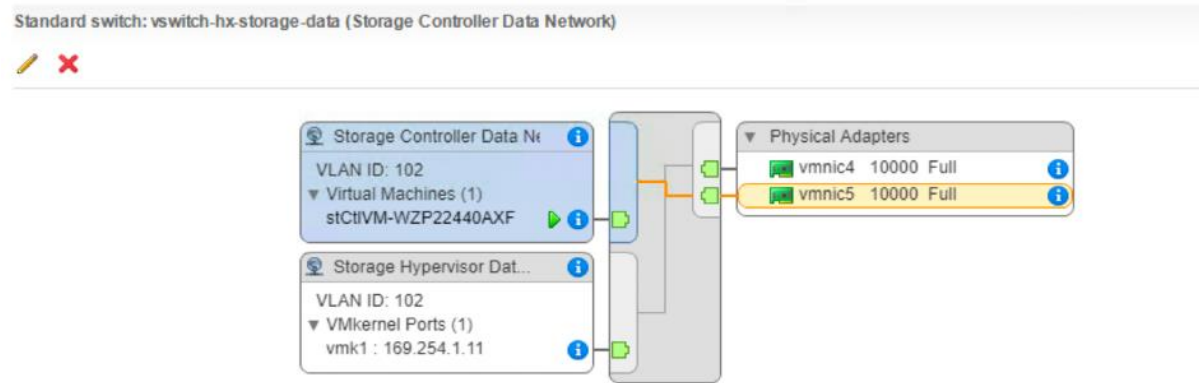
1. Log into vSphere Web Client and verify that the HX Data Platform plug-in appears in the extension list.



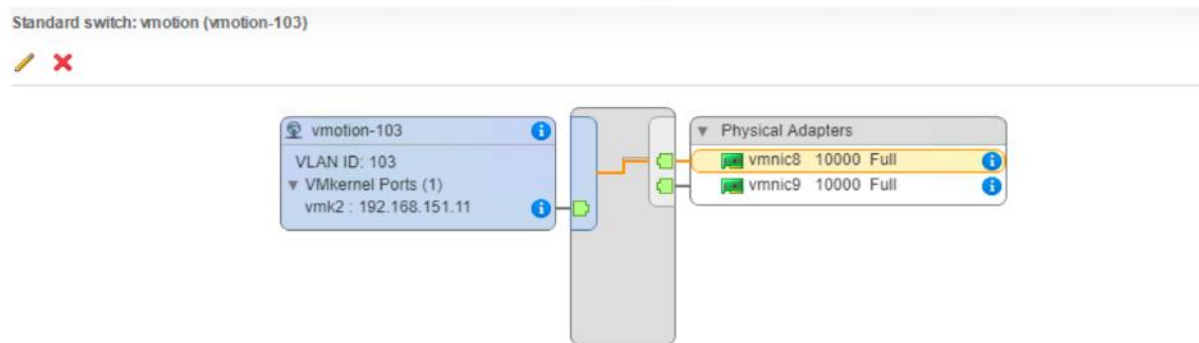
2. Verify that the appropriate VMkernel ports and VM network portgroups are created properly for each host in vCenter.
3. On the Management vswitch:



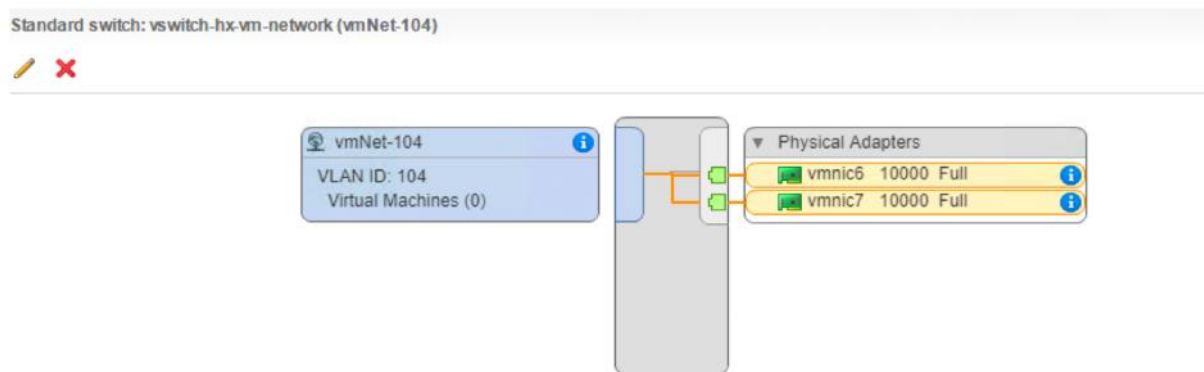
4. On the Storage Data vswitch:



5. On the vMotion vswitch:



6. On the guest VM-network vswitch:



HyperFlex 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 HyperFlex storage cluster is unregistered and in a 90-day evaluation period or EVAL MODE. For the HyperFlex 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 HyperFlex cluster.

To create a Smart Account see; **Cisco Software Central > Request a Smart Account**

<https://webapps.cisco.com/software/company/smartaccounts/home?route=module/accountcreation> .

To activate and configure smart licensing, follow these steps:

1. Log into a controller VM. Confirm that your HyperFlex storage cluster is in Smart Licensing mode.

```
# stcli license show status
```

```
Smart Licensing is ENABLED
```

```
Registration:
```

```
Status: UNREGISTERED
```

```
Export-Controlled Functionality: Not Allowed
```

```
License Authorization:
```

```
Status: EVAL MODE
```

```
Evaluation Period Remaining: 79 days, 8 hr, 52 min, 57 sec
```

```
Last Communication Attempt: NONE
```

The 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 into 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 HyperFlex 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 VM.
11. Register your HyperFlex 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 HyperFlex storage cluster is registered.

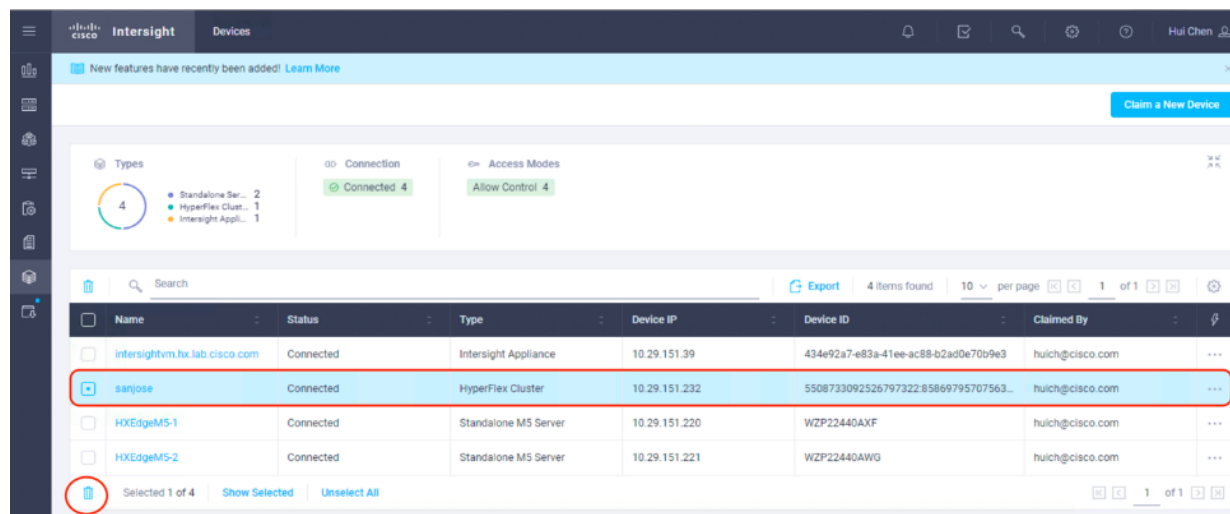
```
# stcli license show summary
```

The cluster is now ready. You may run any other preproduction tests that you wish to run at this point.

Unclaim HyperFlex Edge Cluster in Intersight

To unclaim the HyperFlex Edge cluster, follow these steps:

1. To unclaim the HyperFlex Edge cluster from Cisco Intersight, go to devices, choose the HyperFlex Edge cluster you want to remove. Click Delete.



2. Click Delete again to confirm the deletion.

Remove Device From Intersight

Device "sanjose" will be removed from Intersight

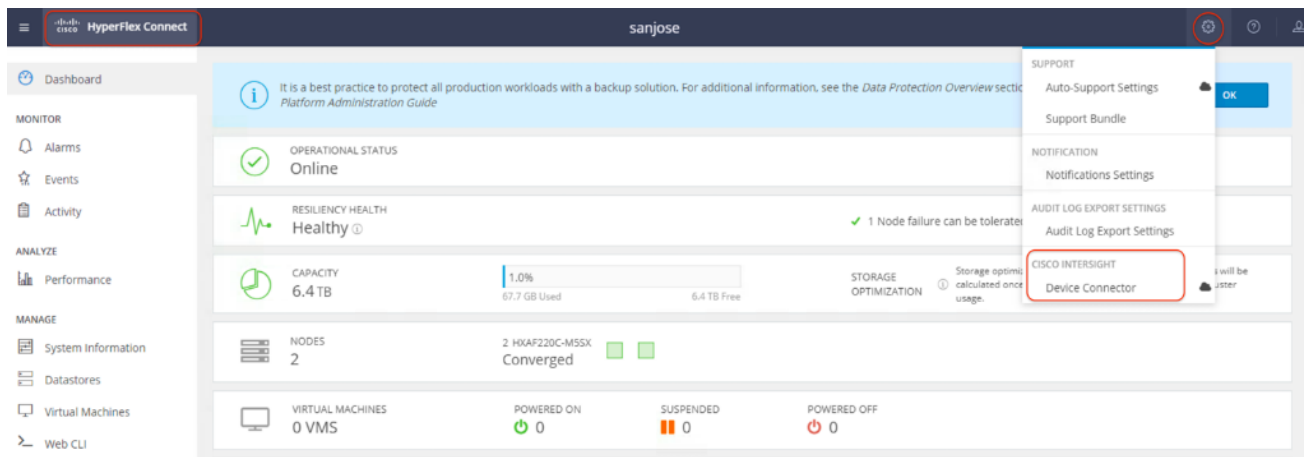


3. Verify the edge cluster is removed from Cisco Intersight.

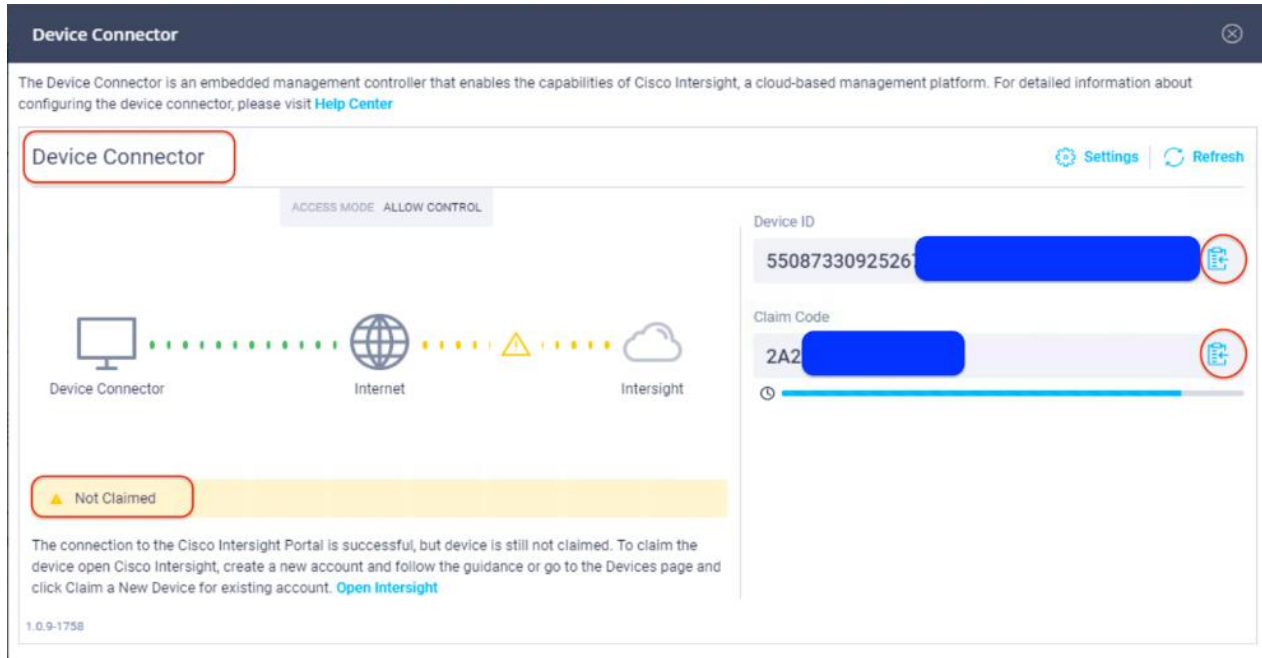
Reclaim HyperFlex Edge Cluster in Intersight

To reclaim the HyperFlex Edge cluster into Cisco Intersight, follow these steps:

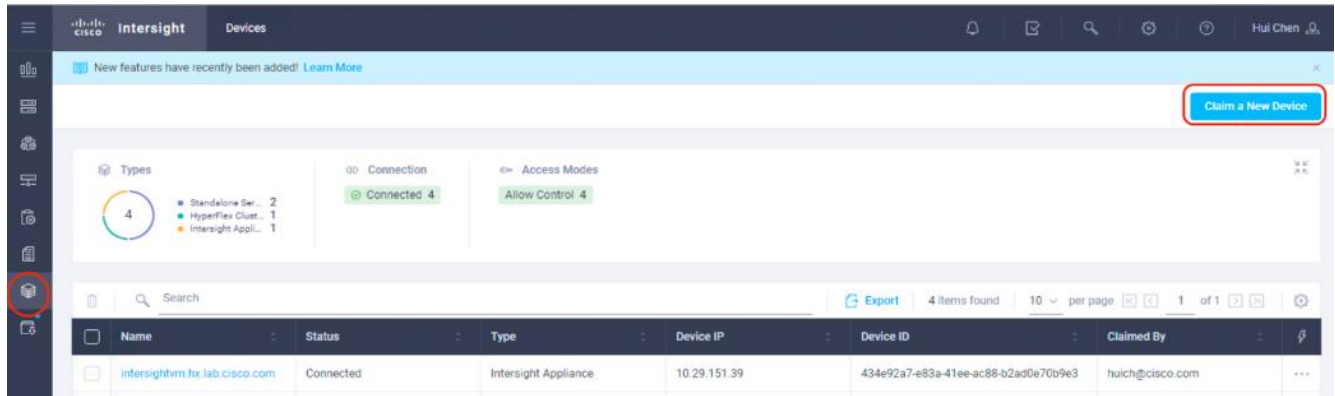
1. In the HyperFlex Connect management console, choose Settings, click Cisco Intersight Device Connector going to the Device Connector page.



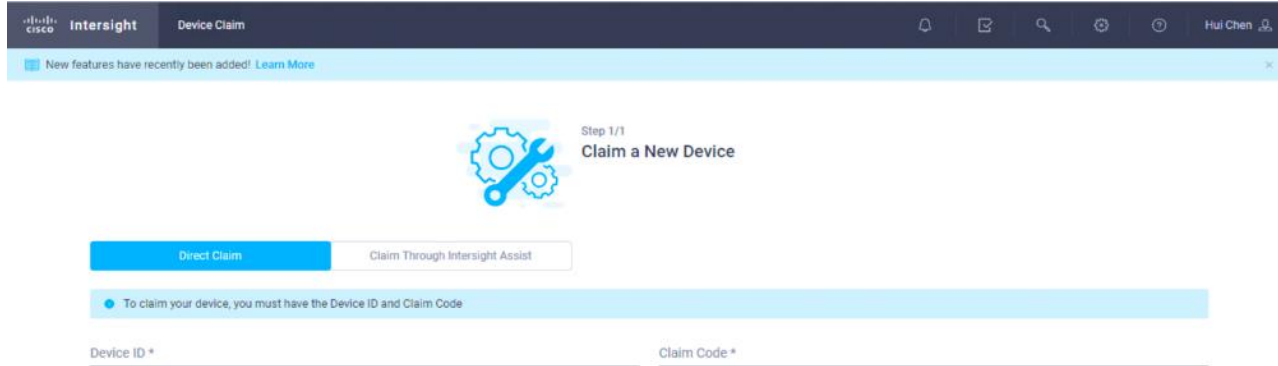
2. The Device ID and Claim Code are provided in this page. One easy way to input these values is clicking on the Copy icon at the right side of the values. Then those values will be placed in Clipboard and can be pasted anywhere.



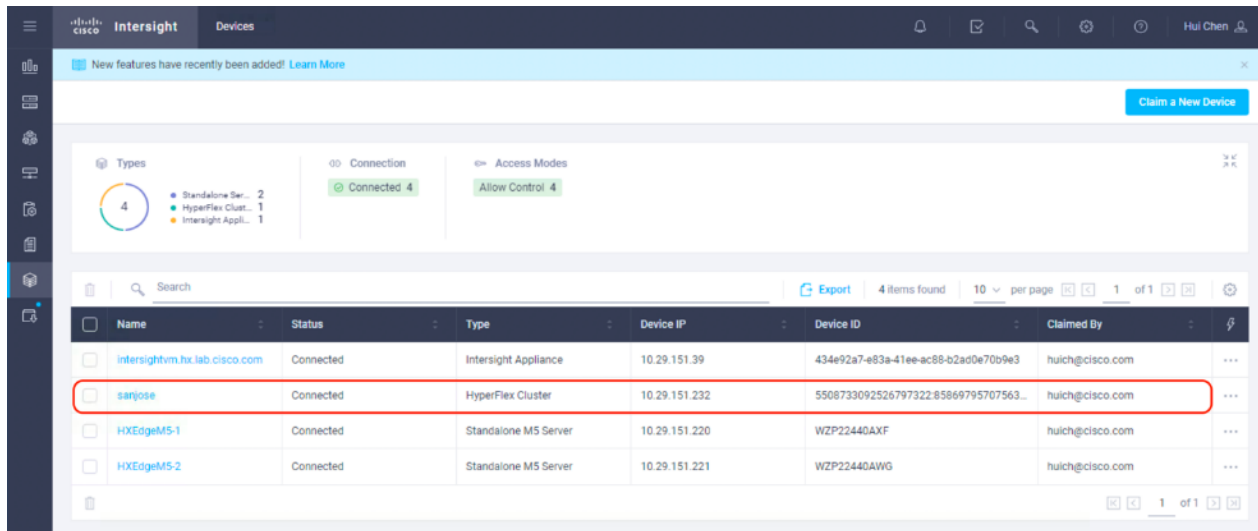
- Log into Cisco Intersight with your Cisco account. From the left-hand Navigation pane, click Devices, in the Device window, choose Claim a New Device at the right top corner.



- Enter the Device ID and Claim Code obtained from Step 2. Use copy and paste for accuracy. Click Claim.



- Wait until the edge cluster is claimed successfully.



Upgrade Cisco HyperFlex Edge Cluster

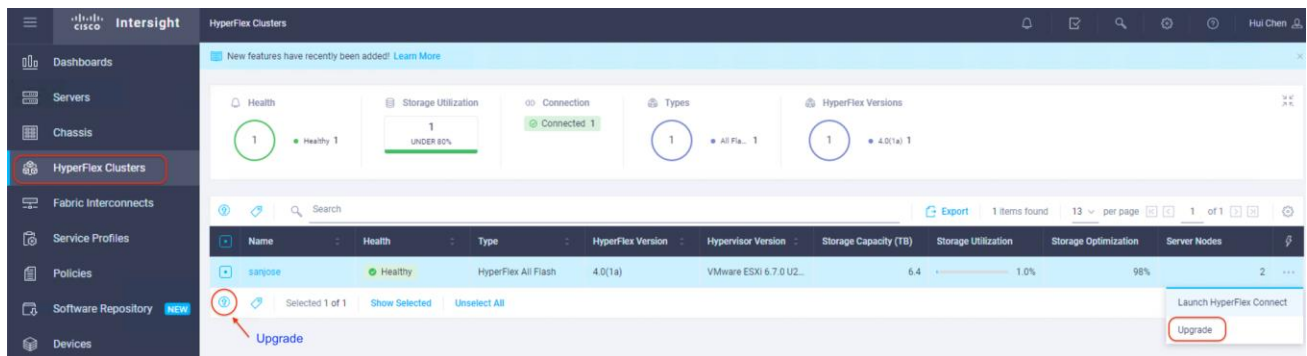
Cisco Intersight delivers Cloud-managed HyperFlex Edge cluster upgrade including support for multi-site orchestrated remote upgrades of the HyperFlex Data Platform. It allows HyperFlex Edge clusters deployed via Intersight to perform orchestrated upgrades across one or many sites in parallel.

As of the time of the publication of this document, upgrading Cisco HyperFlex Edge clusters using Cisco Intersight is supported with the following limitations:

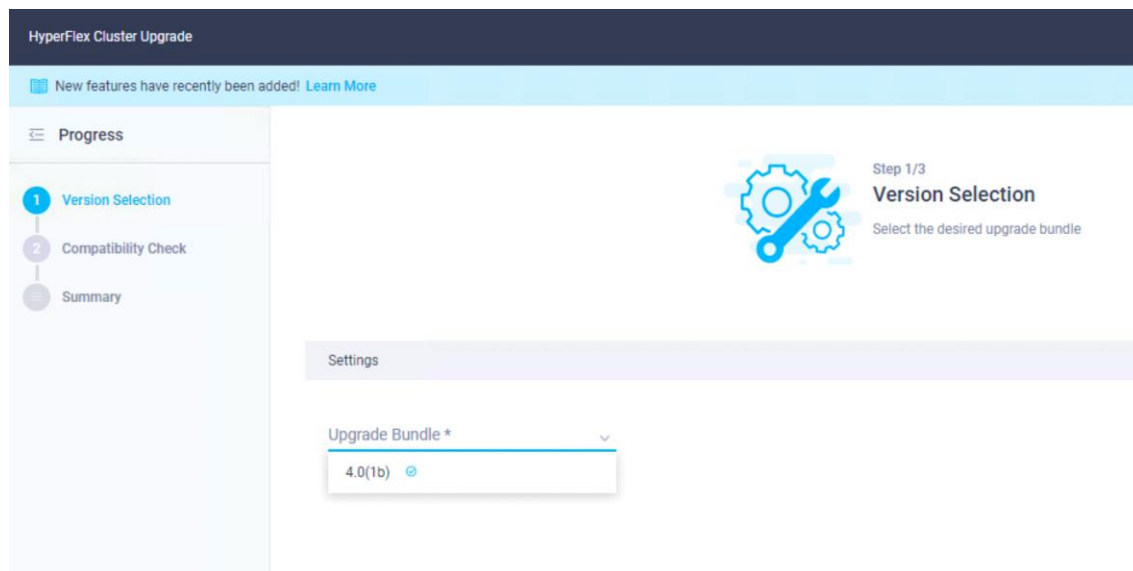
- Upgrade is supported only for the Edge clusters deployed on Cisco UCS HX220c M5 servers.
- The HyperFlex Edge clusters were created through Cisco Intersight.
- The HXDP version on the cluster is 4.0(1a) or later.
- Only the HXDP software will be upgraded.

To upgrade the HyperFlex Edge cluster, for example from HXDP version 4.0(1a) to 4.0(1b), from Cisco Intersight, follow these steps:

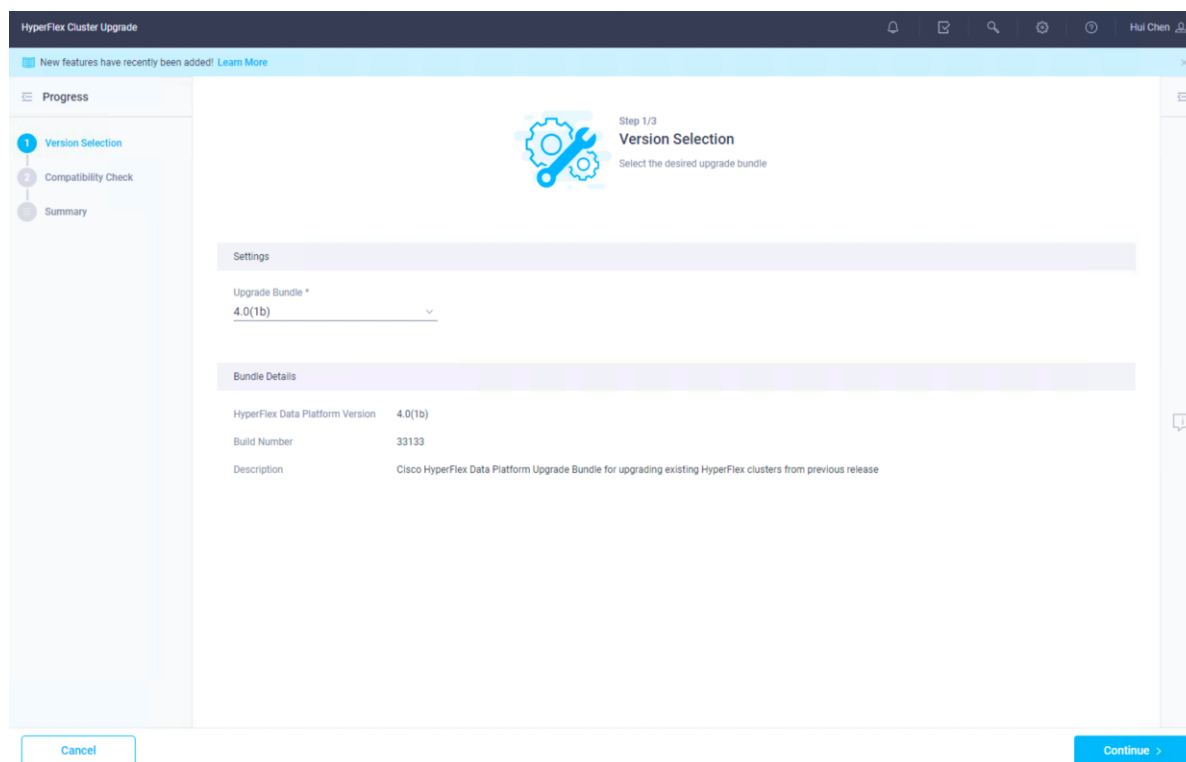
1. From Cisco Intersight, go to HyperFlex Clusters.
2. From the list of the clusters, choose the HyperFlex Edge cluster you want to upgrade. Click Upgrade icon. You can choose multiple clusters at this step for multi-site upgrade. To upgrade a single cluster, you can also right click on "...” for that cluster, then choose Upgrade.



3. Select the Upgrade Bundle from the list, for example 4.0(1b). Click Continue.



4. Verify the bundle details. Click Continue.




5. Wait for the running of Pre-upgrade validations to be completed.

Running Pre-Upgrade Validations

- Checking Cluster Profiles
- Initiating Cluster Pre-upgrade validations
- Waiting for Cluster Pre-upgrade Validations to Complete
- Verifying Configuration Results
- Retrieving Configuration Result Details

- Verify the upgrade requirements are met and the compatibility check succeeds. Choose the cluster from the list, click Continue. Multi-site choice is supported.



Step 2/3

Compatibility Check

Ensure all clusters meet the upgrade requirements. If any cluster does not meet the pre-upgrade requirements, delete the cluster and continue with the upgrade. Deleting a HyperFlex cluster from this page does not remove the cluster from Intersight but removes it from the upgrade workflow.

✔ All clusters meet upgrade requirements

	Name	Hypervisor Version	HyperFlex Version
<input type="checkbox"/>	sanjose	VMware ESXi 6.7.0 U2 (13473784)	4.0(1a)

1 items found
18 per page
1 of 1

incol
Continue >

- Read the warning message, verify the VM settings for your environment meet the requirements. Click Upgrade to continue.

Upgrade (1 HyperFlex Cluster)

Upgrade starts immediately on all selected clusters. When the data path must be upgraded, some clusters may require the evacuation of VMs. For two node clusters, the migration of VMs will be attempted automatically. For three node or larger clusters, enable and set DRS to fully automatic mode. If DRS is not available, manually migrate the VMs when prompted, to ensure that the upgrade can continue.



- From the request list, double click the cluster name to check the upgrade process.

The screenshot shows a management console interface for a 'Two Node Hyperflex Cluster Upgrade'. The top navigation bar includes 'Requests > Two Node Hyperflex Cluster Upgrade' and a user profile 'Hui Chen'. A notification banner states 'New features have recently been added! Learn More'. The main content is divided into 'Details' and 'Execution Flow'.

Details:

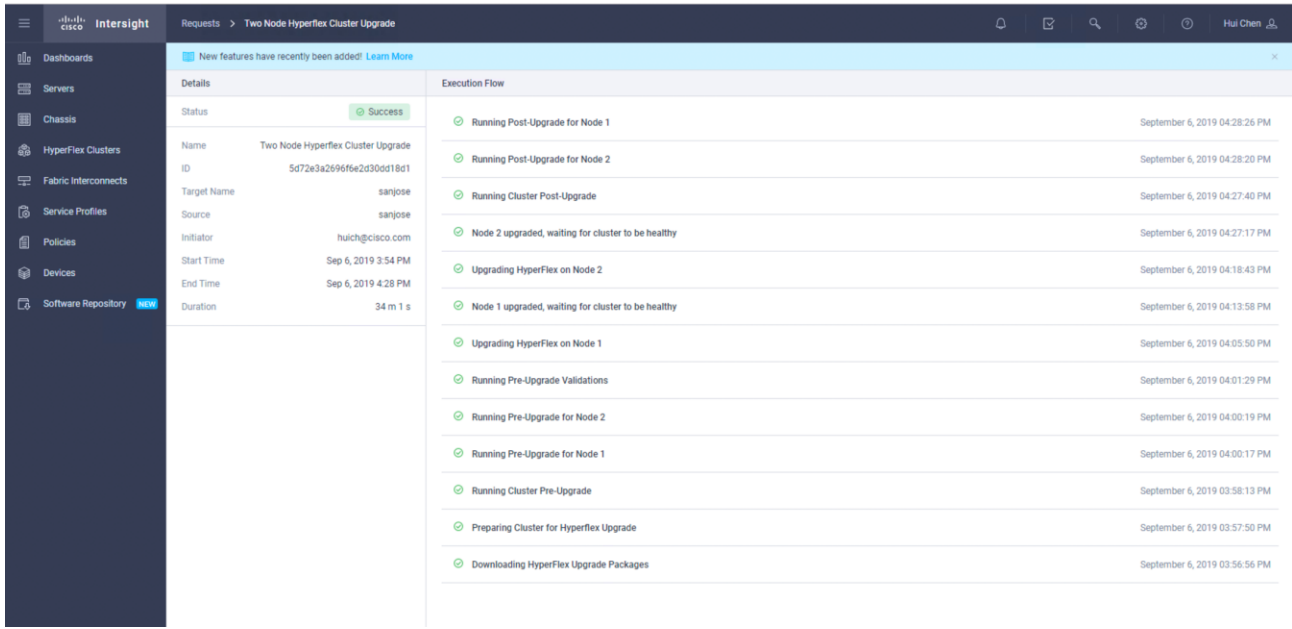
- Status: In Progress
- Name: Two Node Hyperflex Cluster Upgrade
- ID: 5d72e3a2696f6e2d30dd18d1
- Target Name: sanjose
- Source: sanjose
- Initiator: huich@cisco.com
- Start Time: Sep 6, 2019 3:54 PM
- End Time: -
- Duration: 29 m 6 s

Execution Flow:

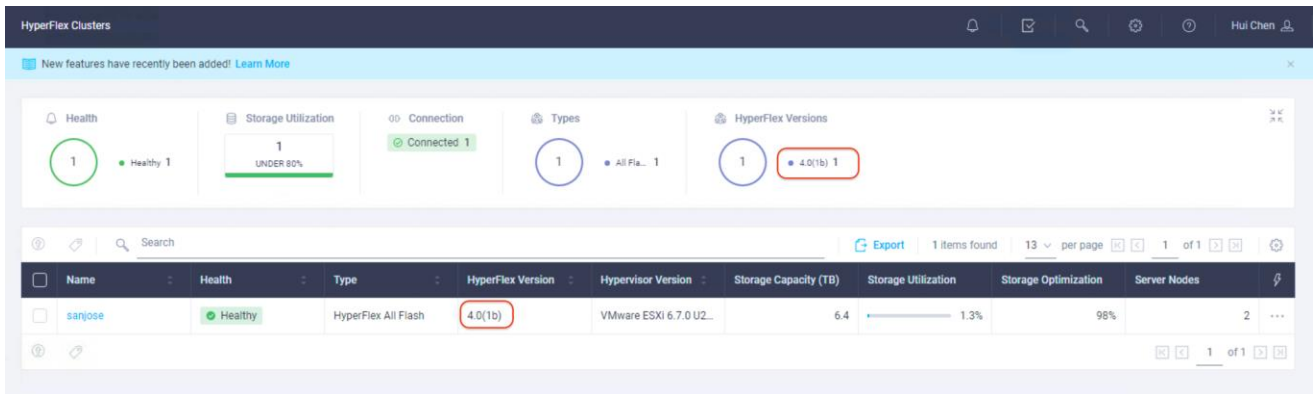
A progress bar at the top of the Execution Flow section shows 69% completion. The flow consists of the following steps:

- Node 2 upgraded, waiting for cluster to be healthy
- Upgrading HyperFlex on Node 2 (September 6, 2019 04:18:43 PM)
- Node 1 upgraded, waiting for cluster to be healthy (September 6, 2019 04:13:58 PM)
- Upgrading HyperFlex on Node 1 (September 6, 2019 04:05:50 PM)
- Running Pre-Upgrade Validations (September 6, 2019 04:01:29 PM)
- Running Pre-Upgrade for Node 2 (September 6, 2019 04:00:19 PM)
- Running Pre-Upgrade for Node 1 (September 6, 2019 04:00:17 PM)
- Running Cluster Pre-Upgrade (September 6, 2019 03:58:13 PM)
- Preparing Cluster for Hyperflex Upgrade (September 6, 2019 03:57:50 PM)
- Downloading HyperFlex Upgrade Packages (September 6, 2019 03:56:56 PM)

- Wait until the upgrade process is completed successfully.



- Go to HyperFlex Clusters. From the list of the clusters, verify the status of the HyperFlex Edge cluster and that the cluster has been upgraded to the desired version.



- Launch HyperFlex Connect management console for this cluster, check the system information. verify the status of the cluster and verify that the cluster has been upgraded to the desired version.

The screenshot shows the Cisco HyperFlex Connect dashboard for a cluster named 'sanjose'. The cluster is in an 'ONLINE' state. Key metrics include: vCenter (https://hvedge-vc.hx.lab.cisco.com), Uptime (31 days, 17 hours, 1 minutes, 5 seconds), Hypervisor (6.7.0-13473784), and HXDP Version (4.0.1b-33133). The dashboard also displays two nodes: 'hx-localhost-55803731' and 'hx-localhost-55807221', both with HXDP Version 4.0(1b) and Type 'Hyper Converged'. The interface includes a sidebar with navigation options like Dashboard, MONITOR, ANALYZE, and MANAGE.



Note: With the current release, upgrading Cisco HyperFlex Edge clusters using Cisco Intersight only upgrades the HXDP software. It is recommended, if necessary, that you upgrade the IMC firmware on the servers one by one using the Cisco Host Upgrade Utility (HUU) tool.

ESXi Hypervisor Installation

HyperFlex nodes come from the factory with a copy of the ESXi hypervisor pre-installed, however there are scenarios where it may be necessary to redeploy the HyperFlex cluster or reinstall ESXi on an HyperFlex node.

The following is a high-level example of a HyperFlex Edge rebuild procedure:

1. Unclaim the HyperFlex Edge cluster from Cisco Intersight.
2. Clean-up the existing environment by:
 - a. Delete the existing HX virtual machines and HX datastores.
 - b. Remove the HX cluster in vCenter.
 - c. Remove vCenter MOB entries for the HX extension.
3. Do a fresh ESXi re-installation on all the HyperFlex Edge nodes.
4. In Cisco Intersight, re-associate the HyperFlex Cluster Profile with the edge nodes and proceed with redeployment.

The HyperFlex system requires a Cisco custom ESXi ISO file to be used, which has Cisco hardware specific drivers pre-installed, and customized settings configured to ease the installation process. The Cisco custom ESXi ISO file is available to download at cisco.com. The HyperFlex custom ISO based on ESXi 6.7 Update 2 ISO release with the filename: ***HX-ESXi-6.7U2-13473784-Cisco-Custom-6.7.2.2-install-only.iso*** is available on the Cisco website and is used to replace the factory preinstalled ESXi 6.5 U2 image in this document:

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

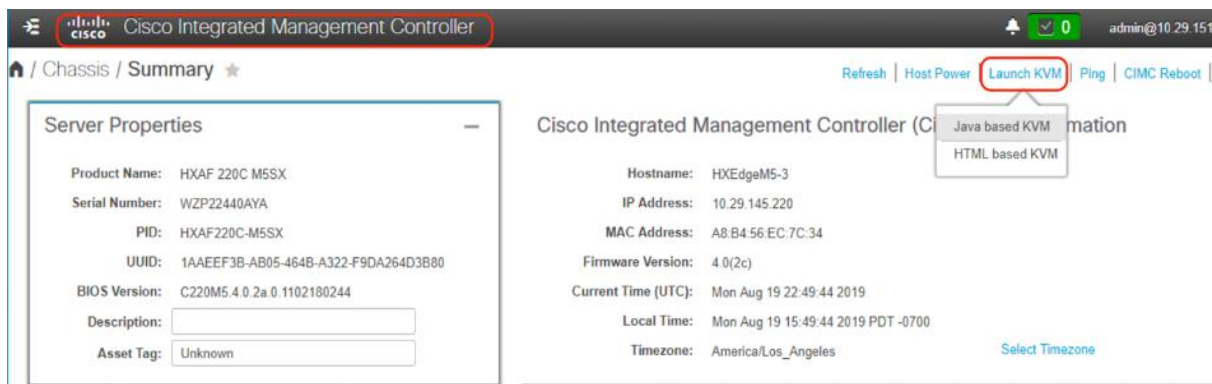
The custom Cisco HyperFlex ESXi ISO will automatically perform the following tasks with no user interaction required:

- Accept the End User License Agreement

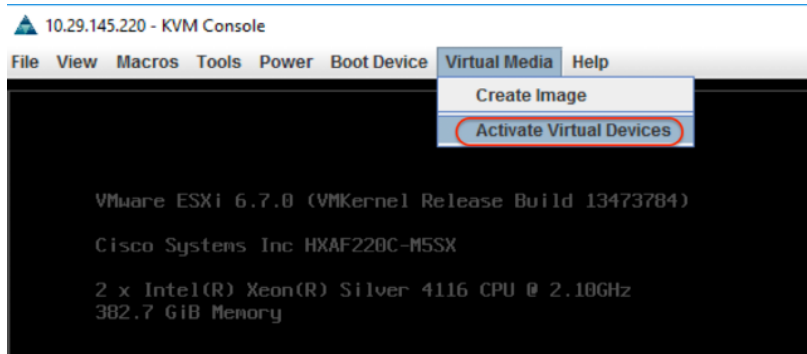
- Configure the root password to: Cisco123
- Install ESXi to the internal mirrored Cisco FlexFlash SD cards, or the internal M.2 SSD
- Set the default management network to use the correct vmnic, and obtain an IP address via DHCP
- Enable SSH access to the ESXi host
- Enable the ESXi shell
- Enable serial port com1 console access to facilitate Serial over LAN access to the host
- Configure the ESXi configuration to always use the current hardware MAC address of the network interfaces, even if they change
- Rename the default vSwitch to vswitch-hx-inband-mgmt

To re-install ESXi hosts, it is necessary to open a remote KVM console session to each server being worked on. To open the KVM console and reboot the servers, follow these steps:

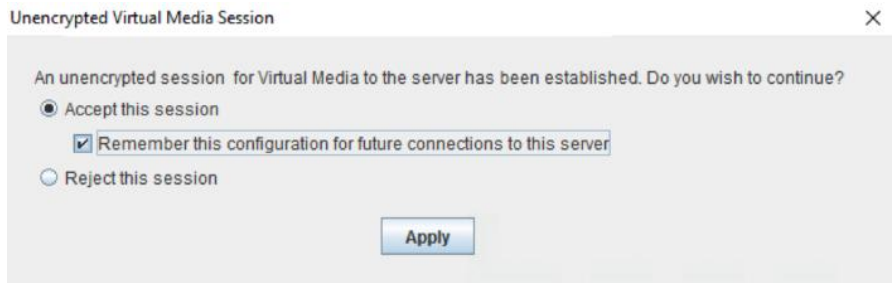
1. Download the latest HyperFlex custom ISO file from Cisco website to the local machine.
2. Launch KVM from CIMC management console for every server.



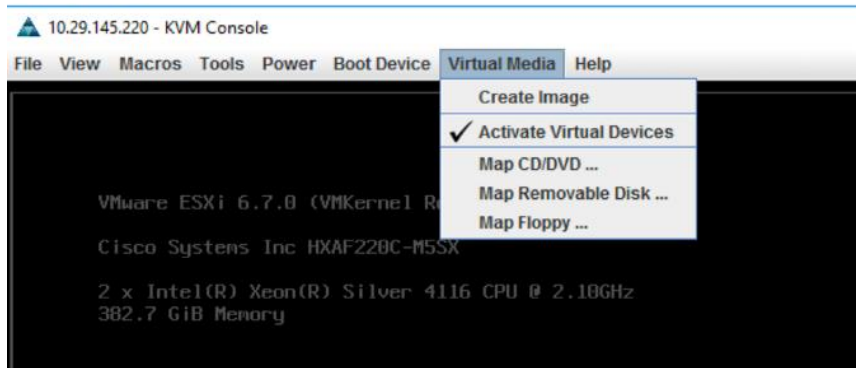
3. In the KVM Console, click Virtual Media, then Activate Virtual Devices.



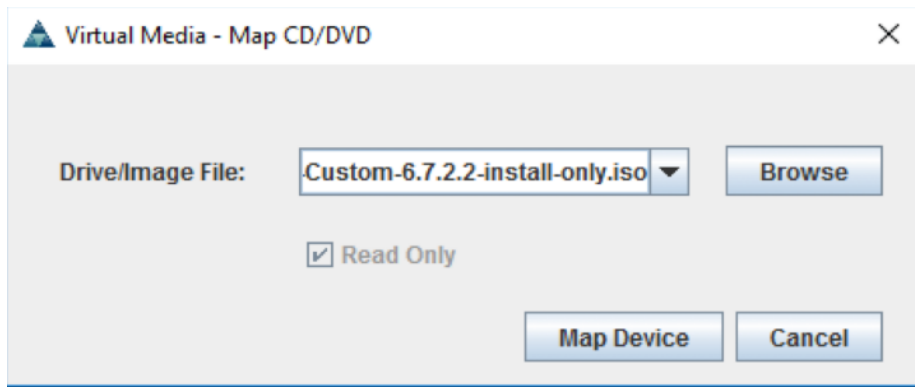
4. Accept the unencrypted session for Virtual Media and apply the response.



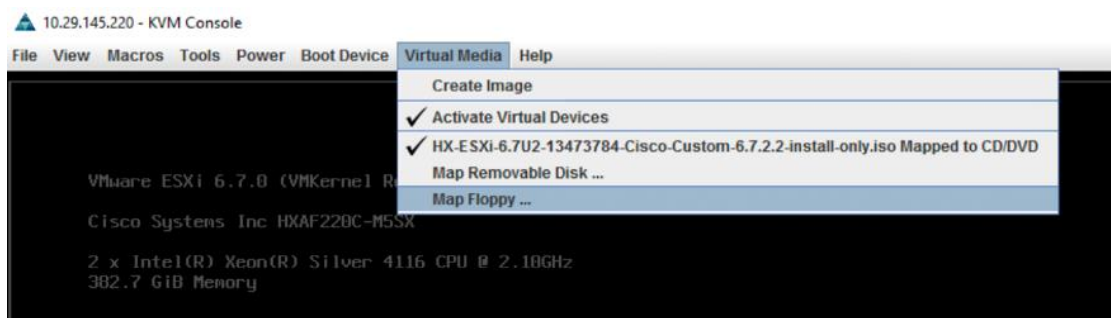
5. Once the session of Activate Virtual Devices is completed, click Map CD/DVD.



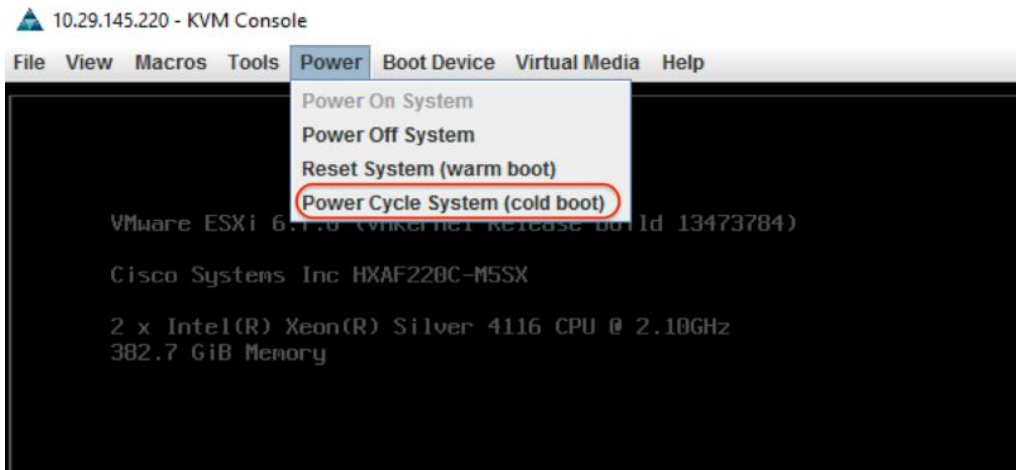
6. Browse to the folder where the downloaded HyperFlex custom ISO image is saved, select that file. Click Map Device to continue.



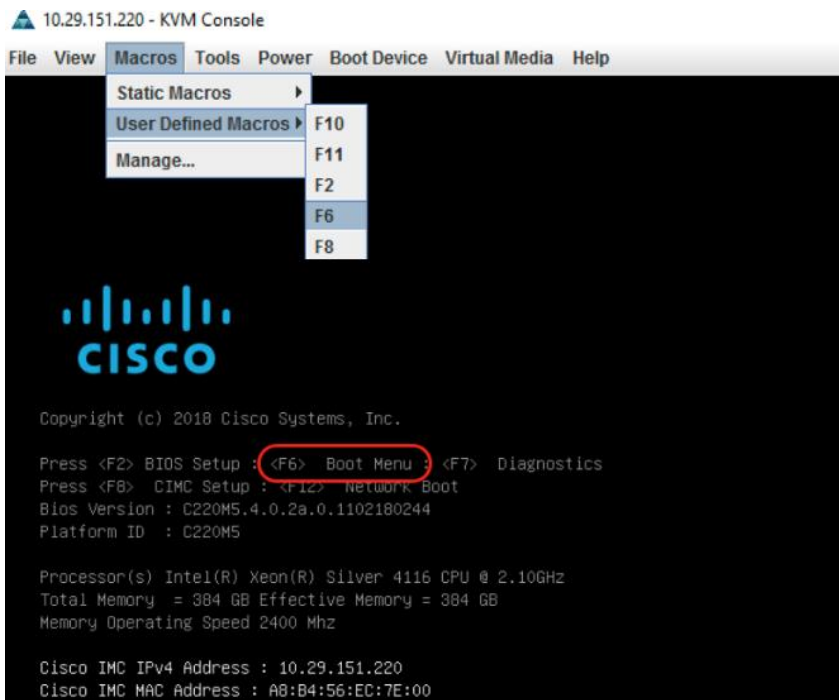
7. Click Virtual Media again to verify that the appropriate ISO file is mapped as virtual device.



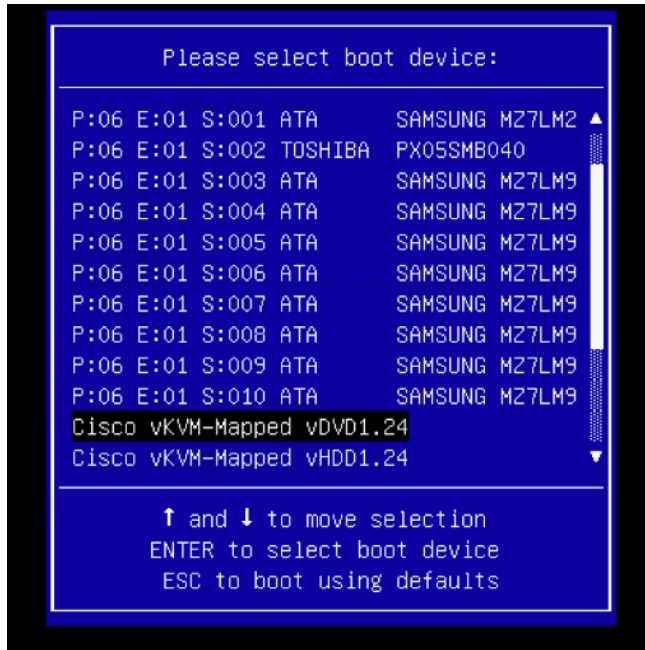
8. Click Power, choose Power Cycle System to reboot the server.



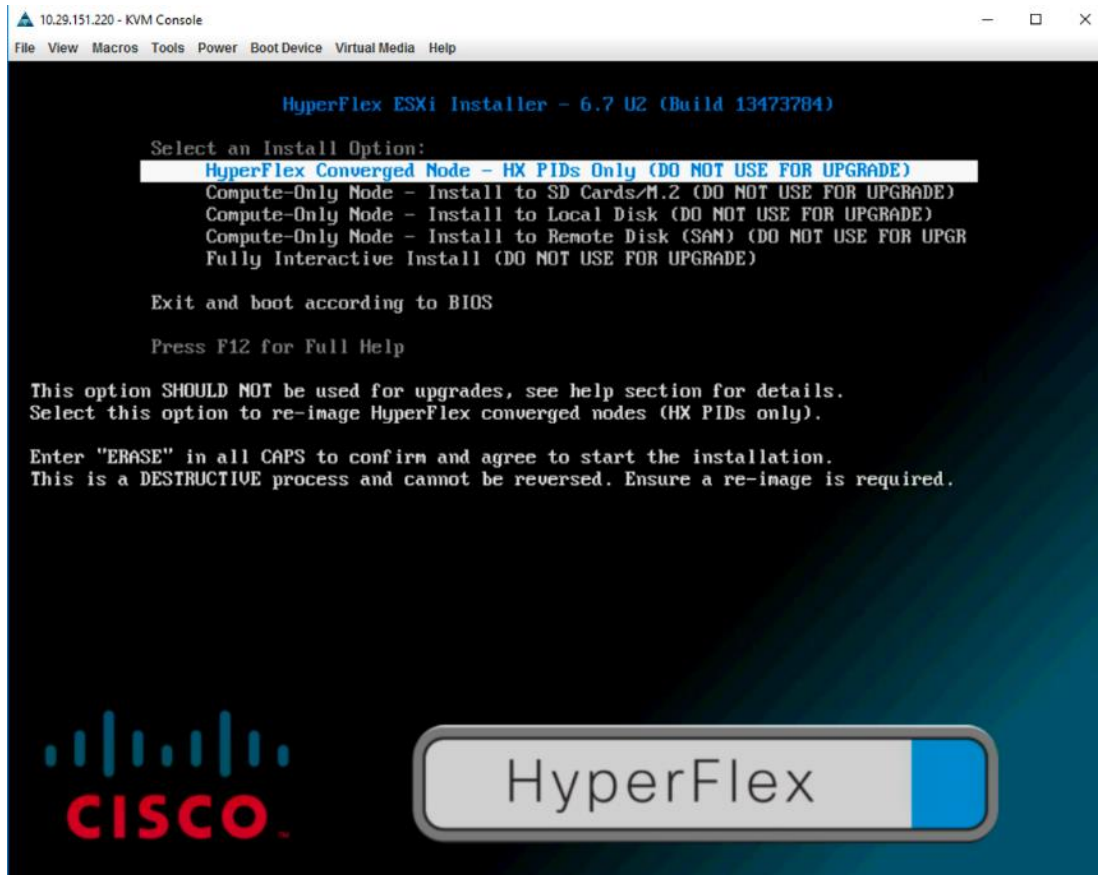
9. Click Yes to continue.
10. The server you are monitoring in the KVM console window will now immediately reboot. During booting, send User Defined Macros F6 to enter the Boot Menu.



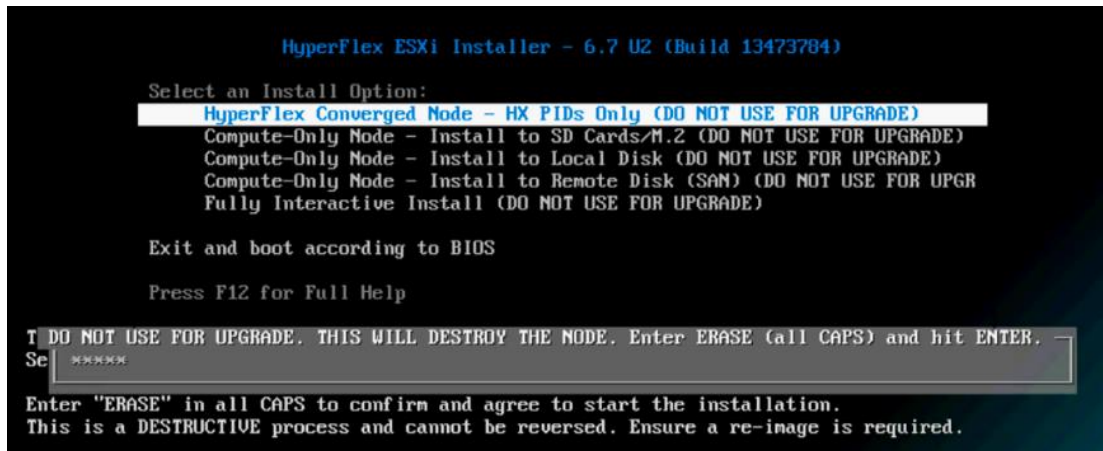
11. Select Cisco vKVM-Mapped vDVD device as boot device.




- Now the server boots from the installation ISO file. You will see a customized Cisco boot menu. In the Cisco customized boot menu, select "HyperFlex Converged Node – HX PIDs Only" and press Enter.



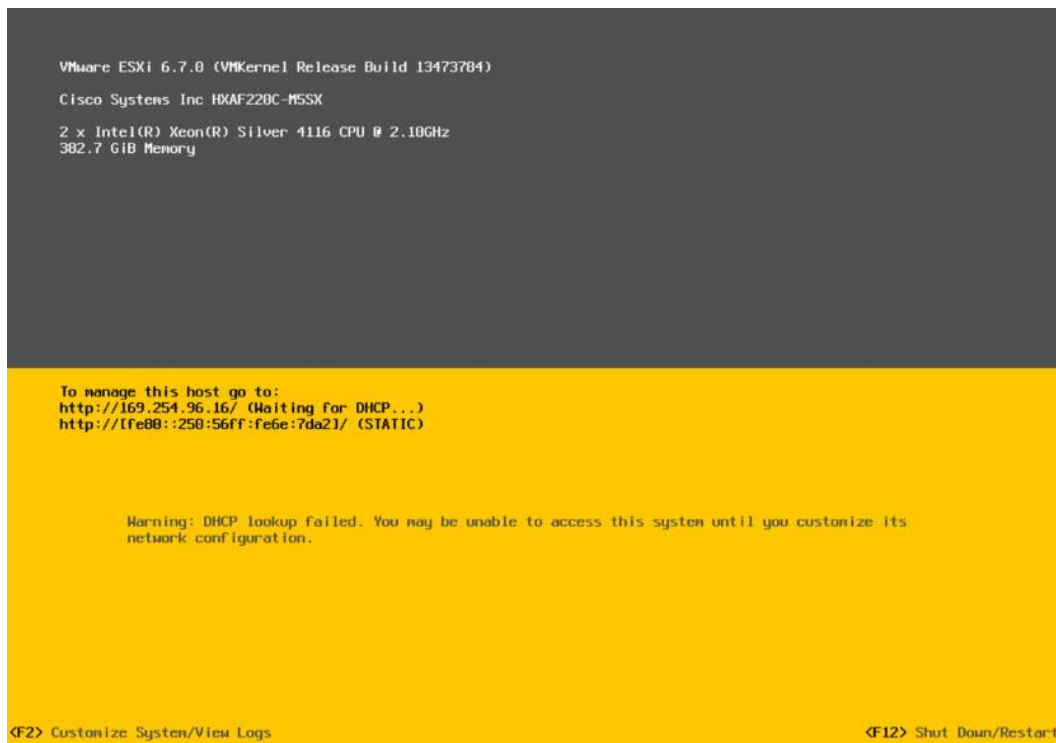
- Enter "ERASE" then press Enter.



- 14. The ESXi installer will continue the installation process automatically; you might receive an error messages but these messages can be safely ignored.



- 15. When the process is complete, the standard ESXi console screen displays:



16. Repeat Steps 3-15 for all additional HyperFlex Edge servers.

Validation

This section provides a list of items that should be reviewed and validated after the HyperFlex Edge system has been deployed. The goal of this section is to verify the configuration and functionality of the solution and ensure that the configuration supports high availability requirements.

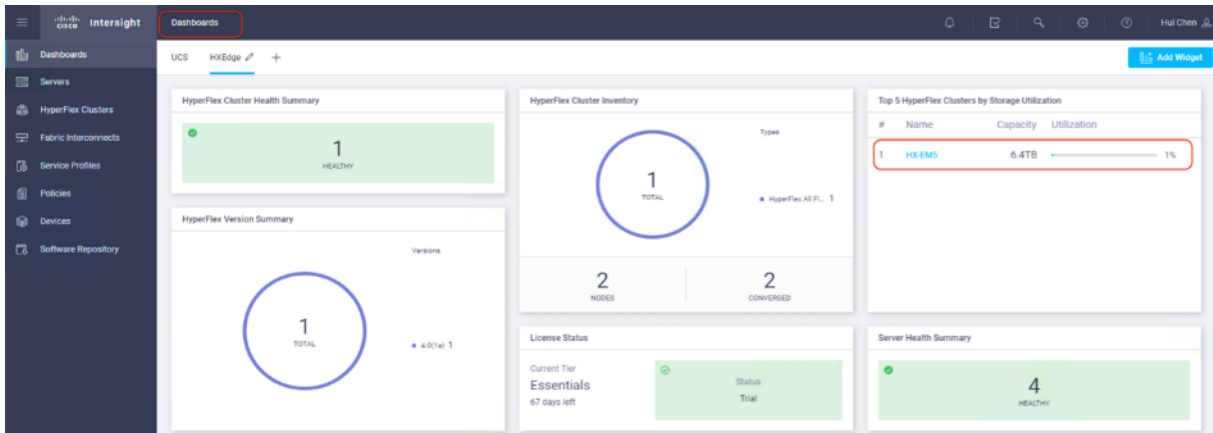
Initial Functionality Validation

The following tests are critical to functionality of the solution and should be verified before deploying for production.

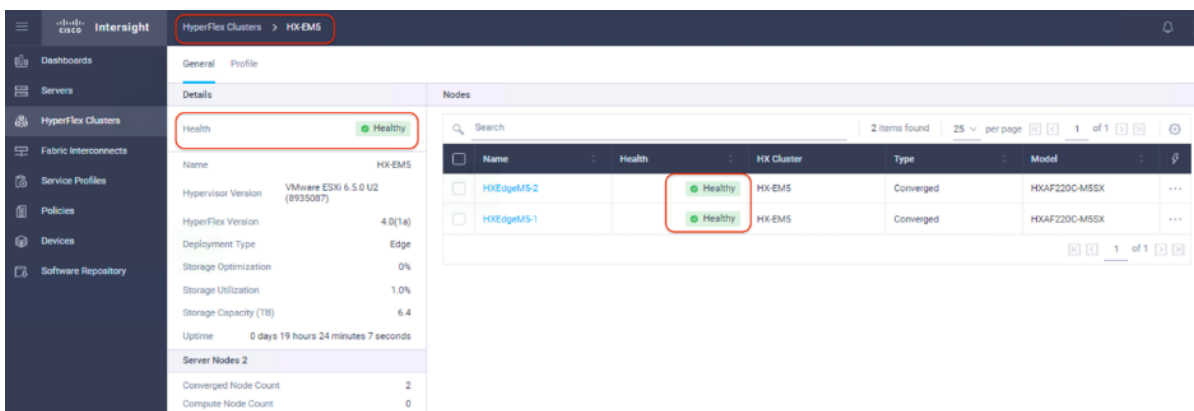
Verify Cluster Status

To verify the cluster status, follow these steps:

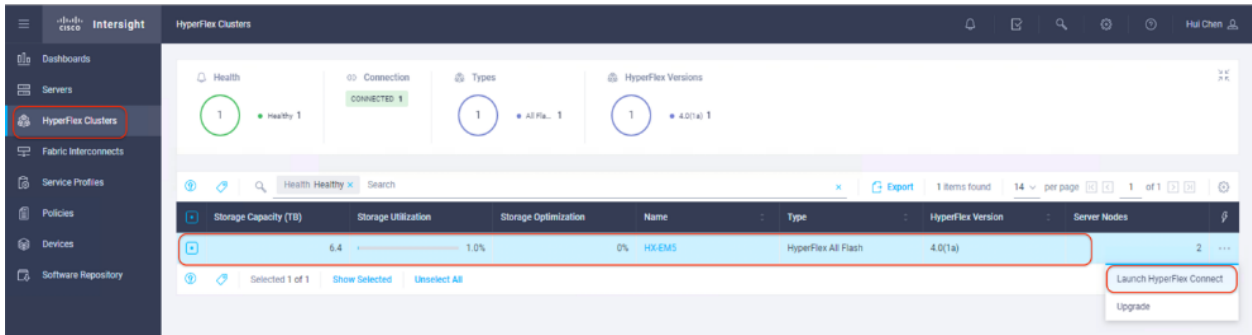
1. From Cisco Intersight Dashboard, check the HyperFlex Cluster Health Summary, ensure the cluster is healthy.



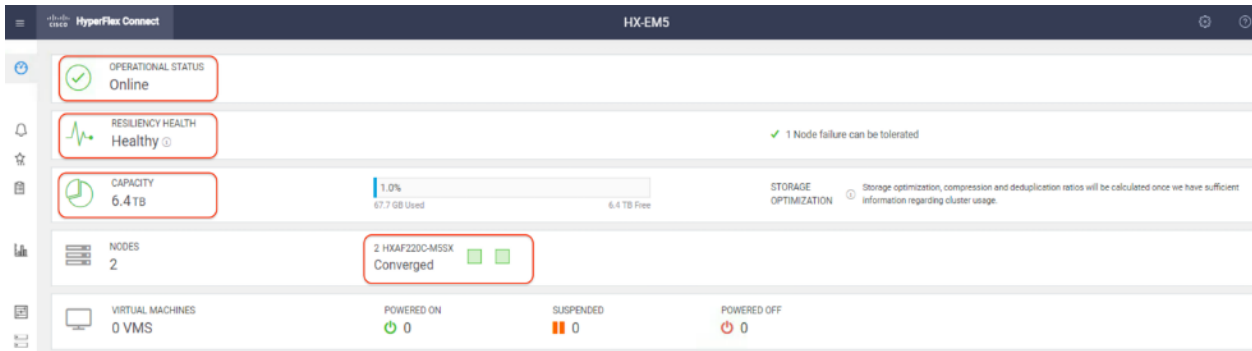
2. Click the name of the cluster for validation.
3. On the page for Cluster Details, verify the cluster health and the nodes health.



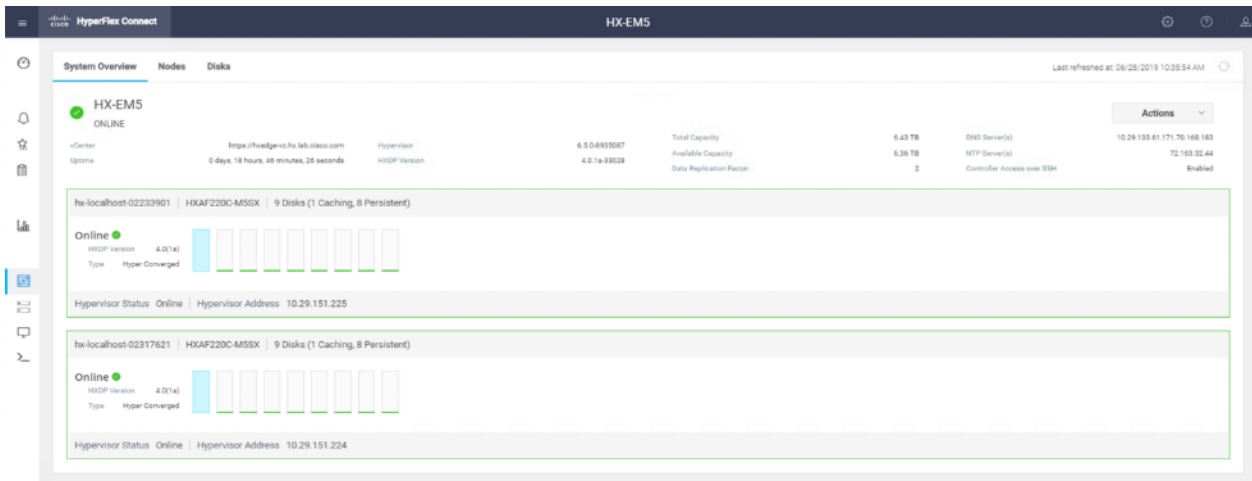
4. Click HyperFlex Clusters on the left pane, go back to the cluster list.
5. Select the cluster, click ... then click Launch HyperFlex Connect to Open HyperFlex Connect management console.



6. On HyperFlex Connect Dashboard, verify the operational status and resiliency health of the cluster, and also the capacity of the cluster.



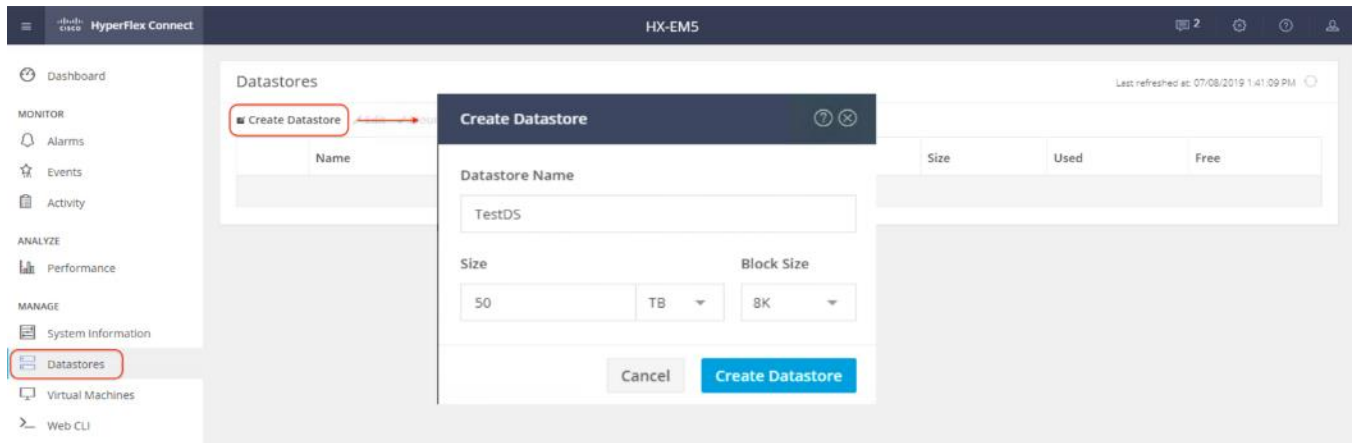
7. Click System Information on the left pane, verify the status of the system, the nodes and the disks.




Create Datastores

To configure a new datastore via the HyperFlex Connect management console, follow these steps:

1. Click Datastores and then click Create Datastore.
2. In the popup screen, enter Datastore Name (for example, TestDS), Size, and Block Size, then click Create Datastore.

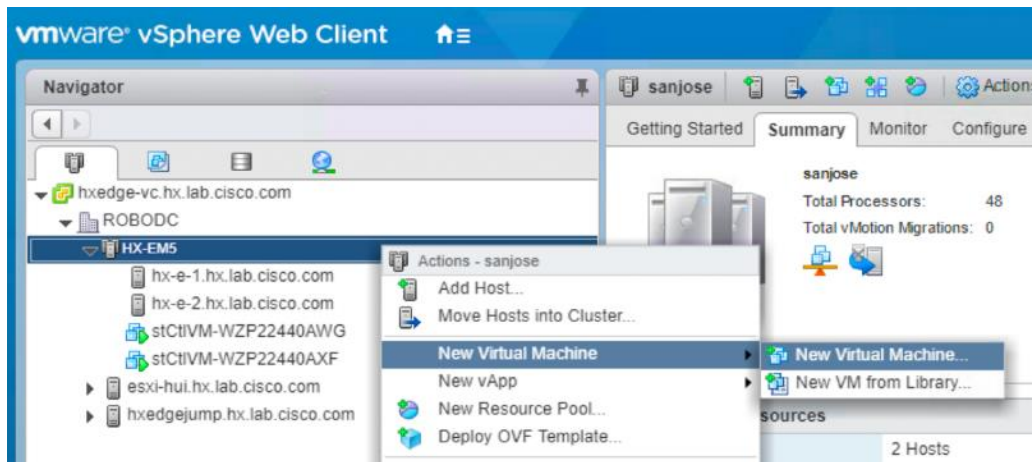


3. After the data store has been created successfully, check the data store status. Make sure that the data store shows as Normal and Mounted.



Create Virtual Machines

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.

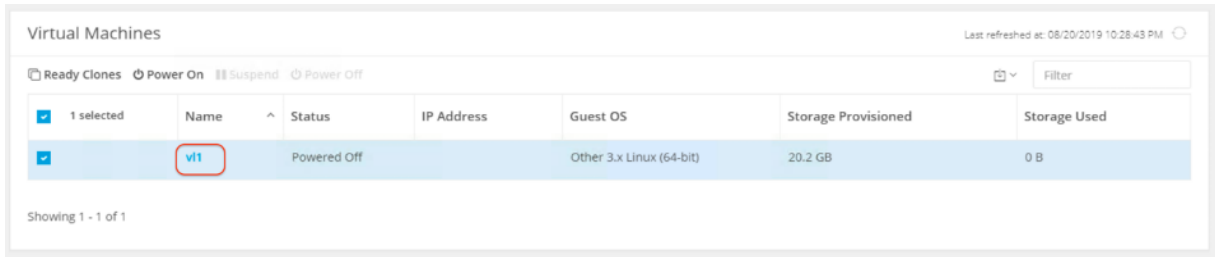


Snapshots

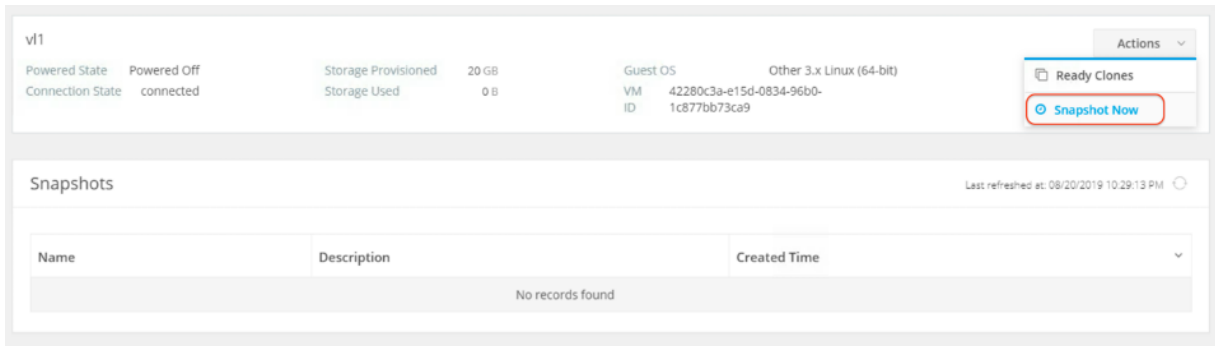
Take a snapshot of the new virtual machine prior to powering it on.

To take an instant snapshot of a VM, follow these steps:

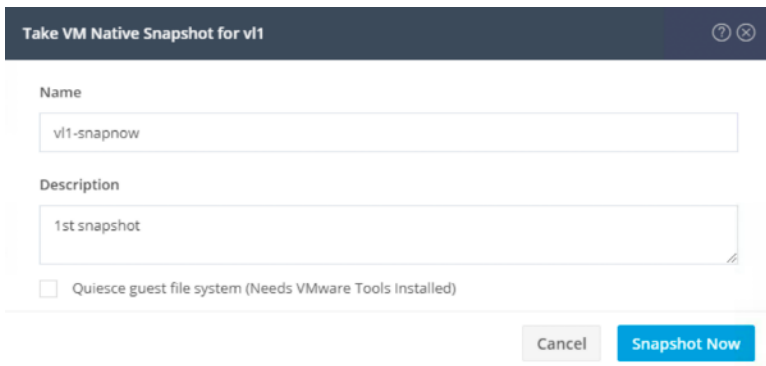
1. In the HyperFlex Connect webpage, click the Virtual Machines menu, then click the name of the VM to snapshot.



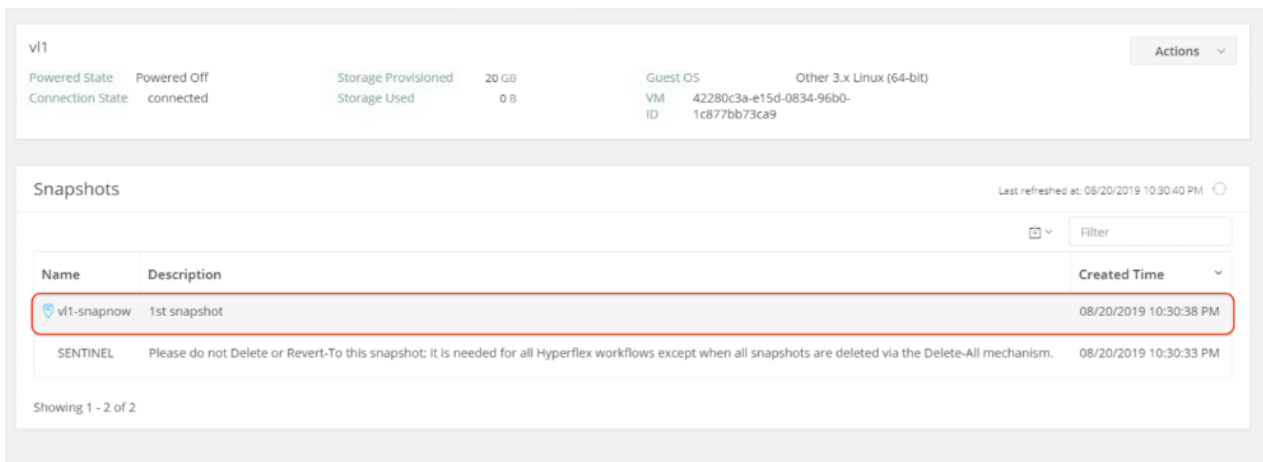
2. Click the Actions drop-down list, then select Snapshot Now.



3. Input the snapshot name, a description if desired, and choose whether to quiesce the VM, then click Snapshot Now.



4. Verify that the snapshot of the VM is created successfully.

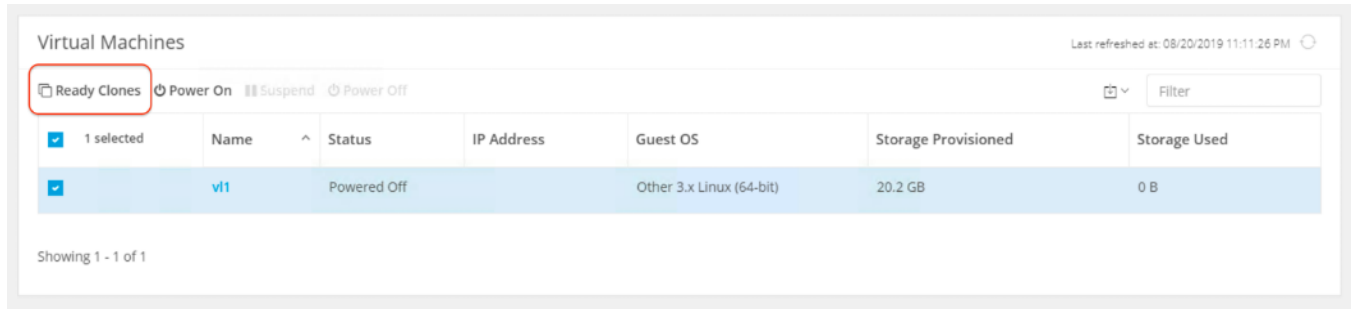


Ready Clones

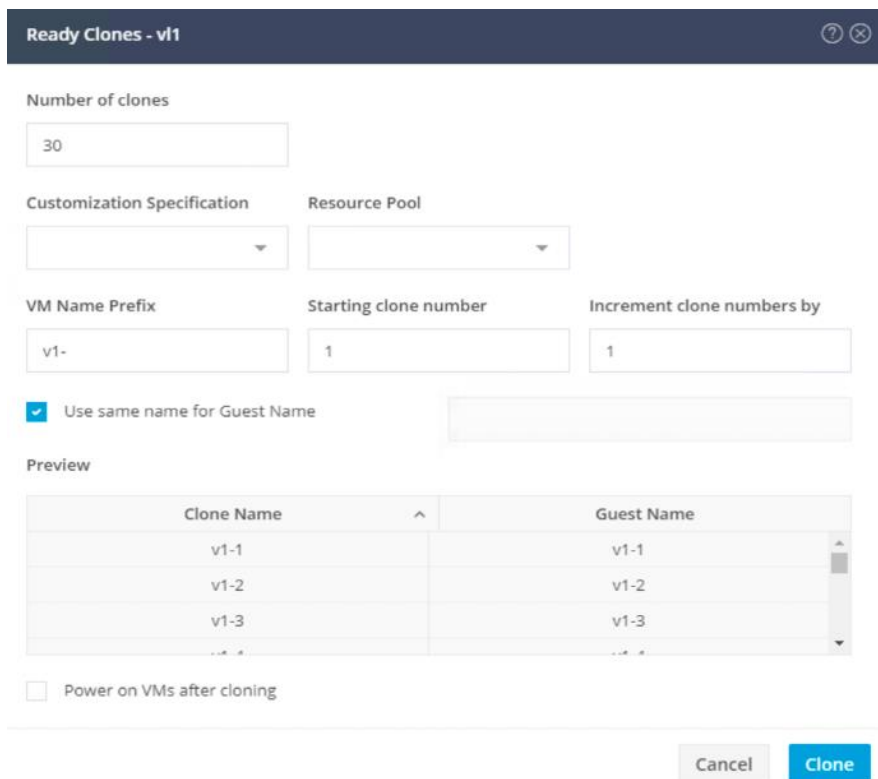
Create a few clones of our test virtual machine.

To create the Ready Clones, follow these steps:

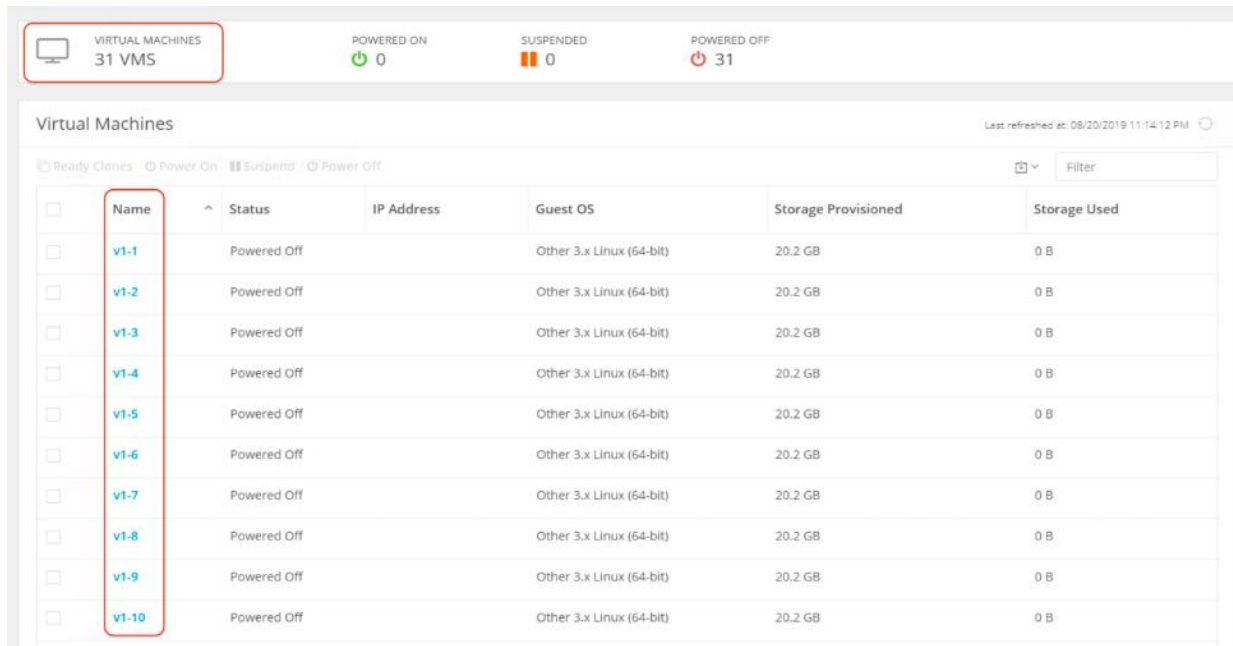
1. In the HyperFlex Connect webpage, click the Virtual Machines menu, click the checkbox to select the VM to clone, then click Ready Clones.



2. Input the Number of clones to create, a customization specification if needed, a resource pool if needed, and a naming prefix, then click Clone to start the operation. The clones will be created in seconds.

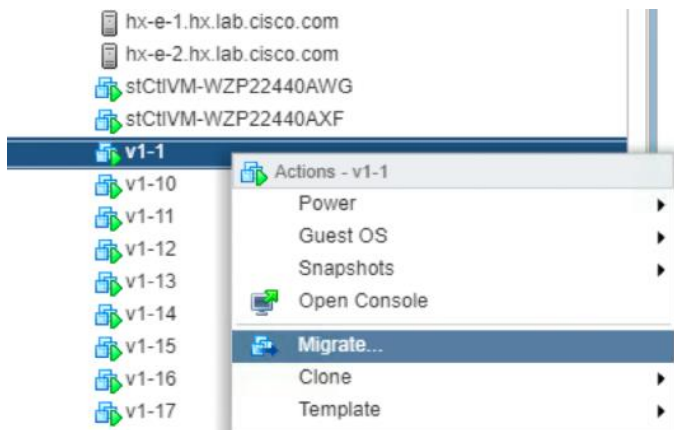


3. In the HyperFlex Connect Virtual Machines screen, verify the expected number of VMs are cloned.



vMotion

1. Power on the virtual machines.
2. Perform the virtual machine migration (vMotion) of one test virtual machine to a different host on the cluster.



3. During the vMotion of the virtual machine, make sure the test virtual machine can perform a continuous ping to default gateway and to check if the network connectivity is maintained during and after the migration.

Verify Redundancy

The following redundancy testing can be performed to verify the robustness of the system. With the Invisible Cloud Witness provided by Cisco Intersight since HXDP 4.0 release, the need for an external witness node is eliminated for 2-node HyperFlex Edge clusters.

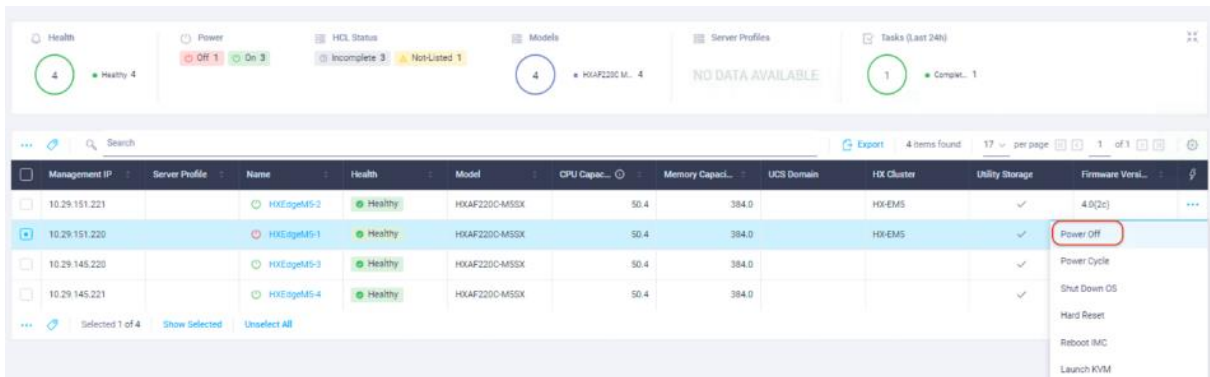
Cisco HyperFlex Invisible Cloud Witness is designed to maintain file system consistency by preventing a “split brain” condition if the nodes in the ROBO cluster lose direct communication with each other and is designed to tolerate failures of the following components:

- WAN/Internet Link
- LAN Connectivity between nodes (may be direct connect)
- Node Availability

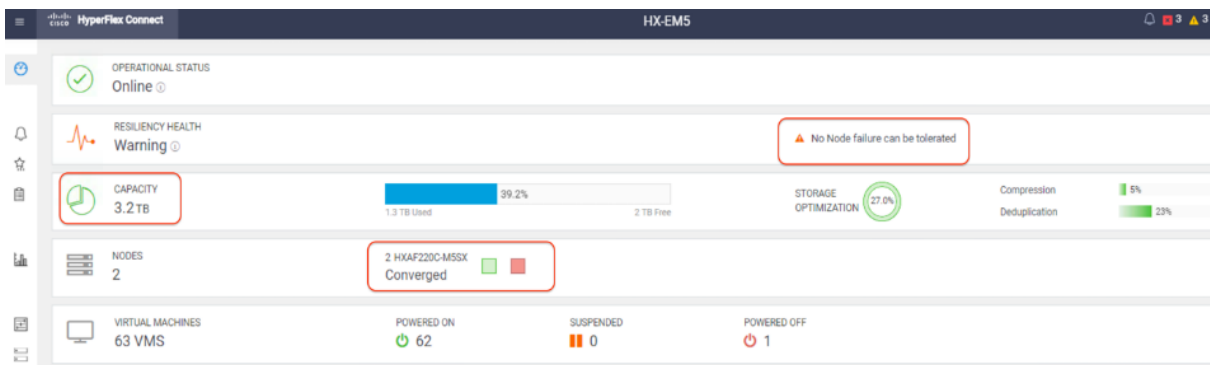
These scenarios need to be tested. In addition, removing or adding capacity disks on the HyperFlex Edge nodes should not cause the system downtime and also needs to be tested.

Test - HyperFlex Node Failover

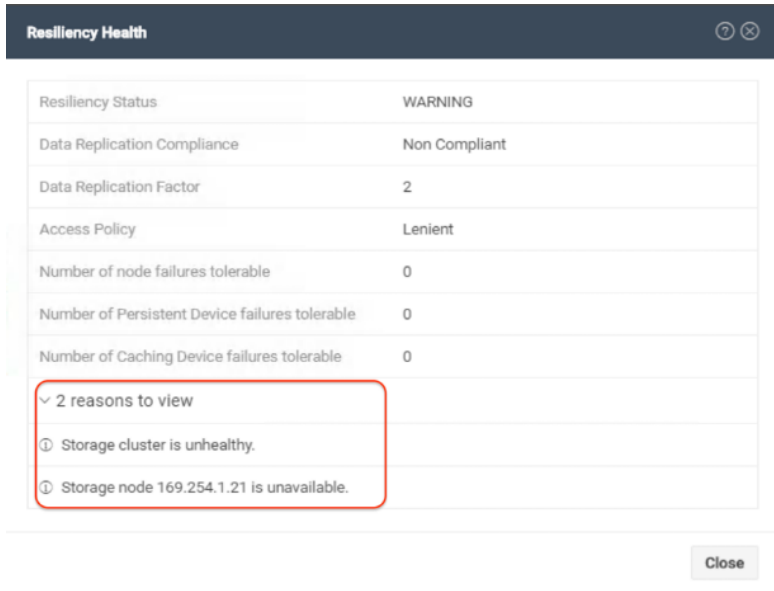
1. Check and verify that the cluster is healthy.
2. Shutdown one of the two HyperFlex Edge nodes, which can be done from Cisco Intersight.



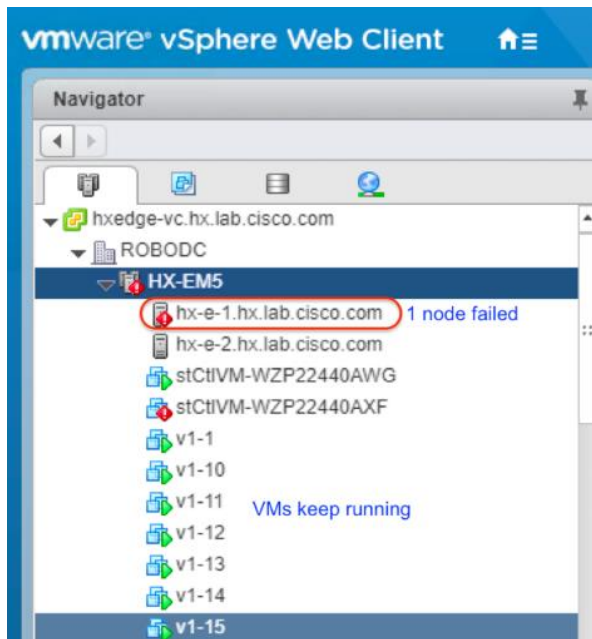
3. Check the cluster status from HyperFlex Connect management console. In HyperFlex Connect, one node became red, and the cluster capacity has been reduced to half.



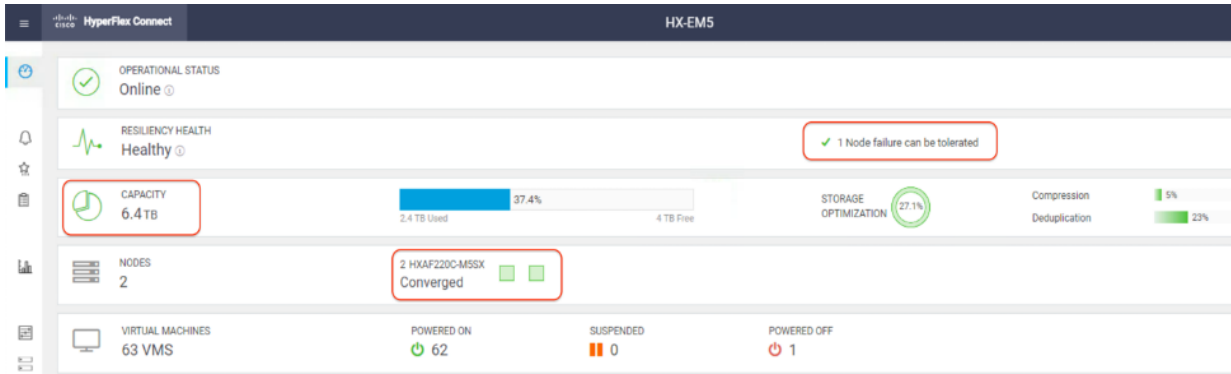
4. Click the Warning Information sign under Resiliency Health, it shows one node has failed so the cluster is unhealthy.



- Verify that VMs on the surviving node should not experience any downtime, and VMs from the failed node should automatically be restarted on the surviving node by the vCenter High Availability process after a few minutes.



- Recover the failed node. Wait until the cluster is fully recovered. Verify the cluster and nodes status is healthy.



7. vMotion some VMs back to the recovered node, verify that those VMs run properly.

Test – Network Uplink Failover

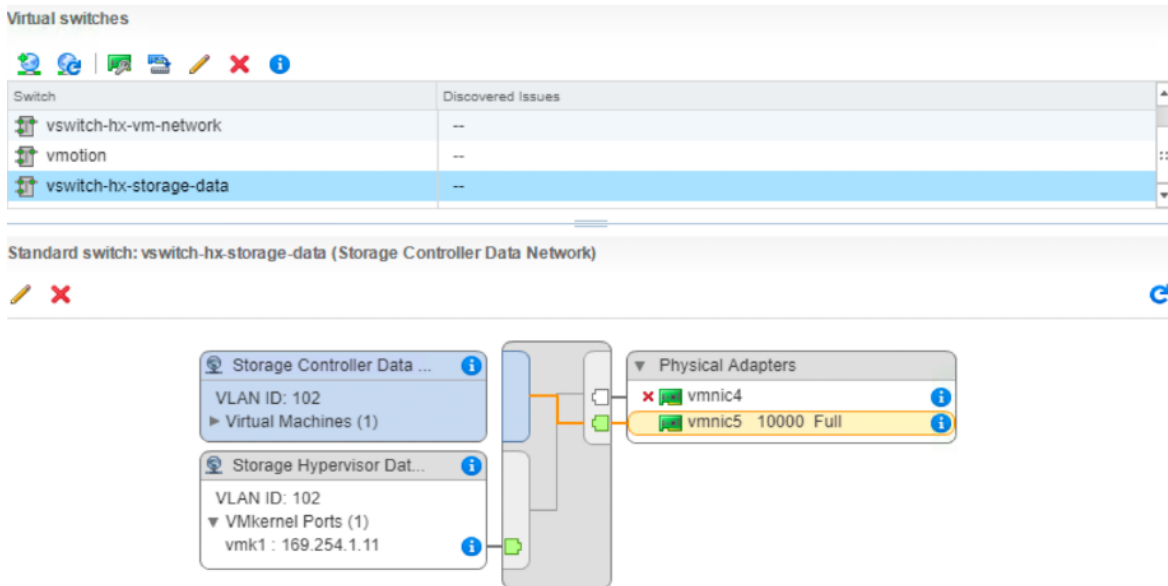
1. Check and verify that the cluster is healthy.
2. Fail one of the network uplinks from the Cisco VIC 1457 card via CLI.

```

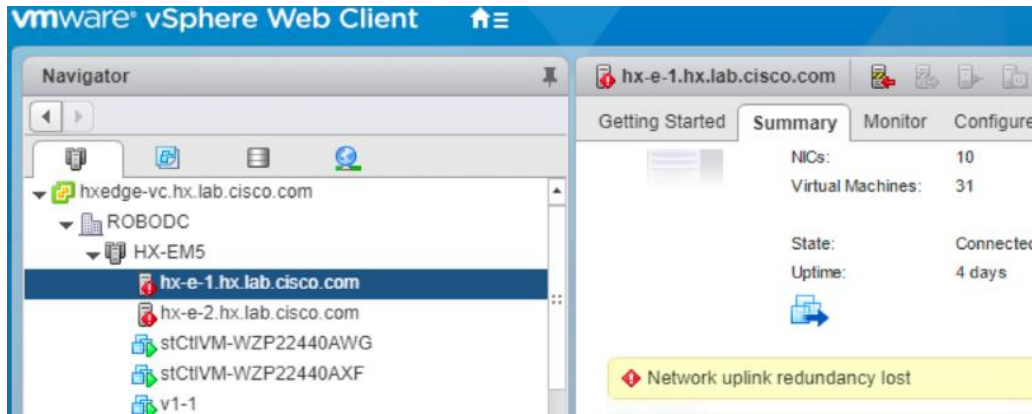
HXEDGE-N9Ka(config-if)# show interface brief | grep up
mgmt0  --          up          10.29.151.30          100      1500

Eth1/1    101    eth  access up      none          1000(D)  --
Eth1/5    1       eth  trunk up      none          1000(D)  --
Eth1/49   1       eth  trunk up      none          10G(D)   --
Eth1/50   1       eth  trunk up      none          10G(D)   --
HXEDGE-N9Ka(config-if)# interface e1/49-50
HXEDGE-N9Ka(config-if-range)# shutdown
HXEDGE-N9Ka(config-if-range)# no shutdown
HXEDGE-N9Ka(config-if-range)#
    
```

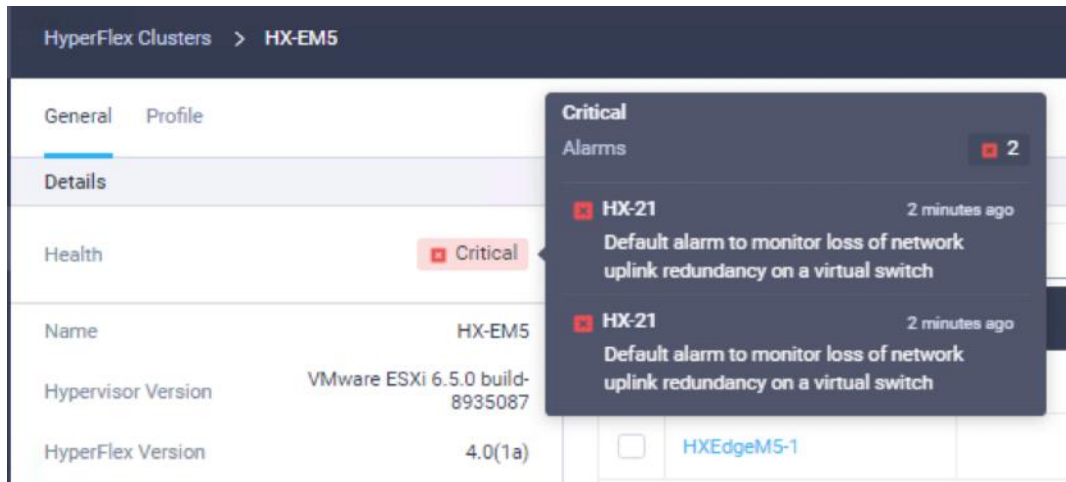
3. In vCenter check that one network uplink on the ESXi host has failed.



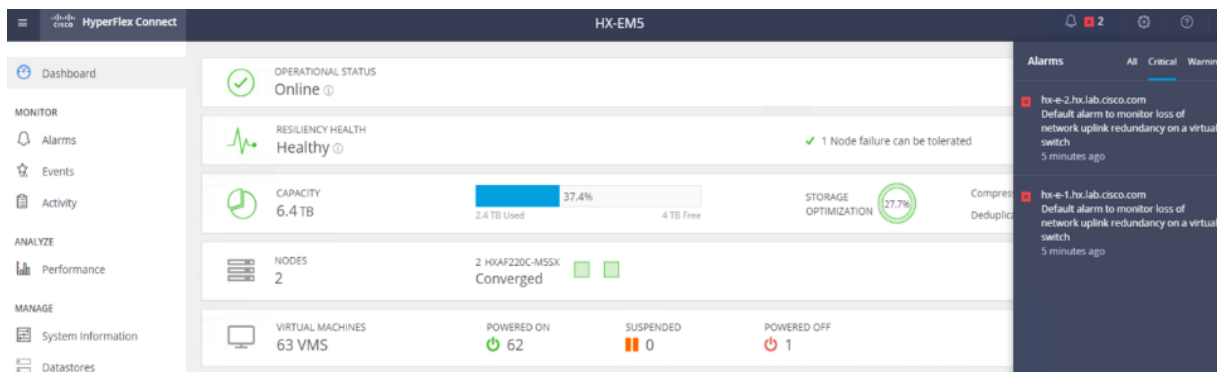
4. Verify an alarm "Network uplink redundancy lost" is displayed in vCenter.



5. The cluster status in Cisco Intersight became Critical with error of Loss of Network Uplink Redundancy.



6. Critical alarms show up in HyperFlex Connect GUI too.



7. Enable the VIC ports.

```

HXEDGE-N9Ka(config-if)# show interface brief | grep up
mgmt0  --  up  10.29.151.30  100  1500

Eth1/1  101  eth  access  up  none  1000(D)  --
Eth1/5  1  eth  trunk  up  none  1000(D)  --
Eth1/49 1  eth  trunk  up  none  10G(D)  --
Eth1/50 1  eth  trunk  up  none  10G(D)  --
HXEDGE-N9Ka(config-if)# interface e1/49-50
HXEDGE-N9Ka(config-if-range)# shutdown
HXEDGE-N9Ka(config-if-range)# no shutdown
HXEDGE-N9Ka(config-if-range)#
    
```

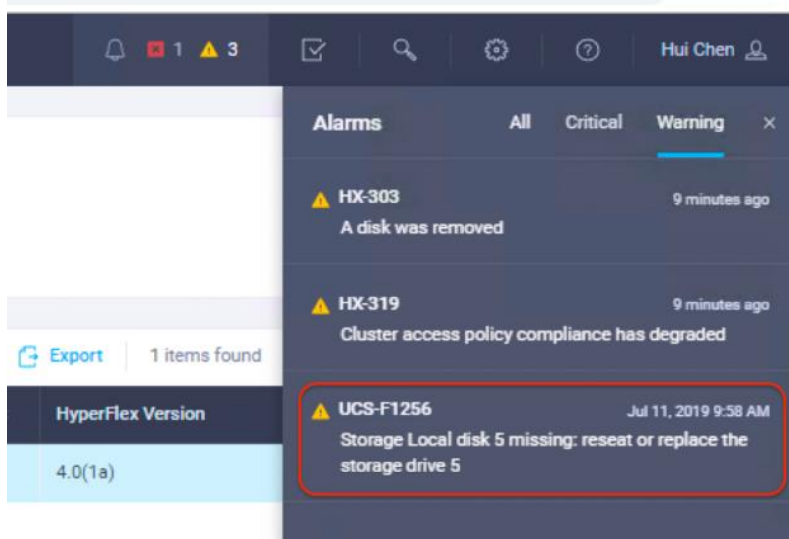

- Verify that all the alarms have been cleared.



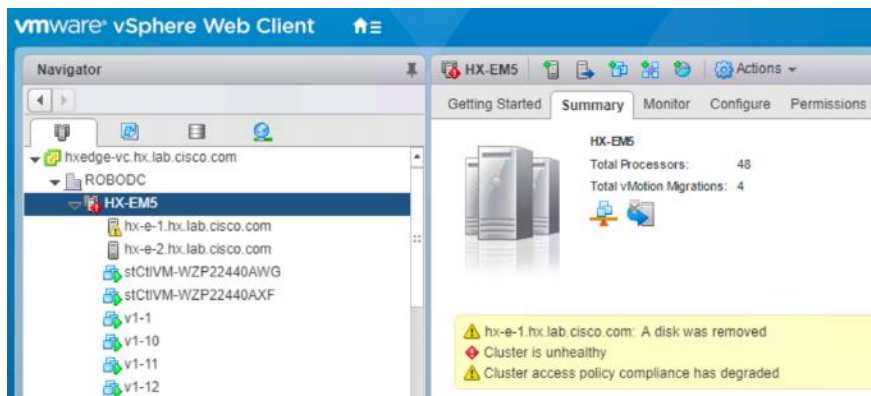
Note: One port from each server is disabled and then re-enabled for this testing, which can also simulate the scenario of loss of one switch in the dual-switch 10GE topology. The switch failover testing should have the same results.

Test - Capacity Drive Failure

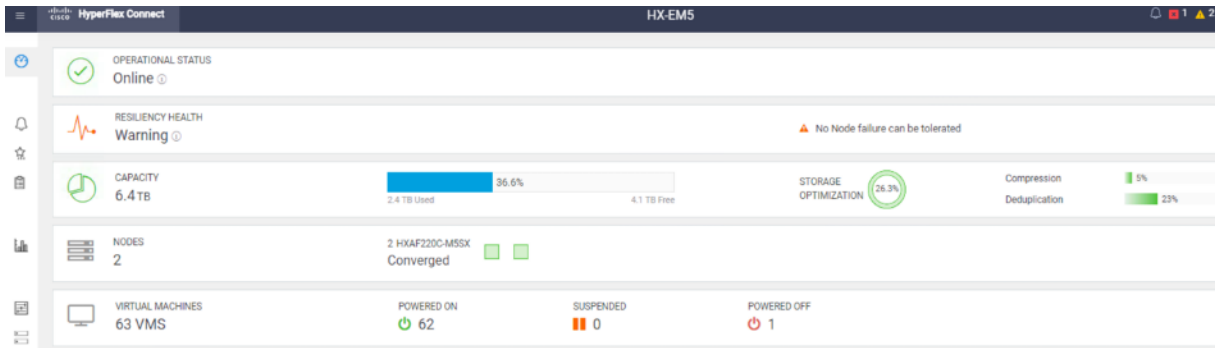
- Check and verify that the cluster is healthy.
- Pull out a capacity drive on node 1 of the cluster.
- View the alarms displayed in Cisco Intersight.



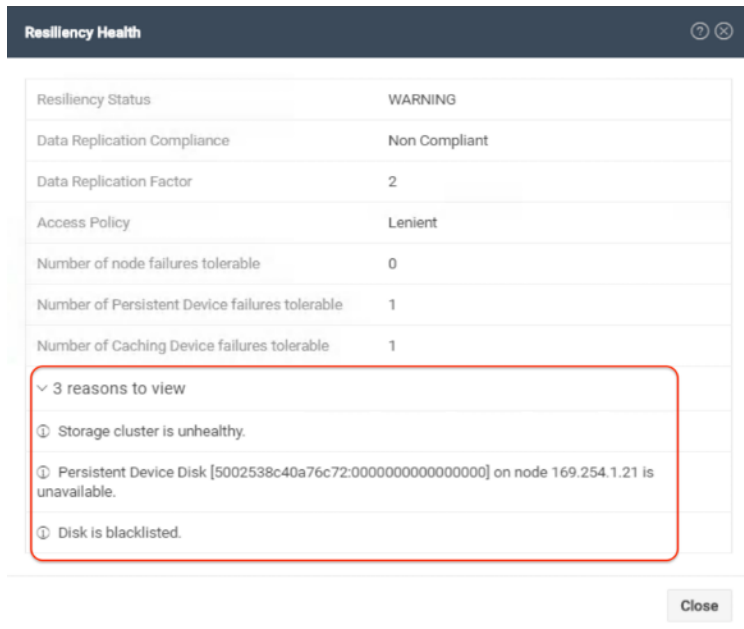
- View the alarms displayed in vCenter as well.



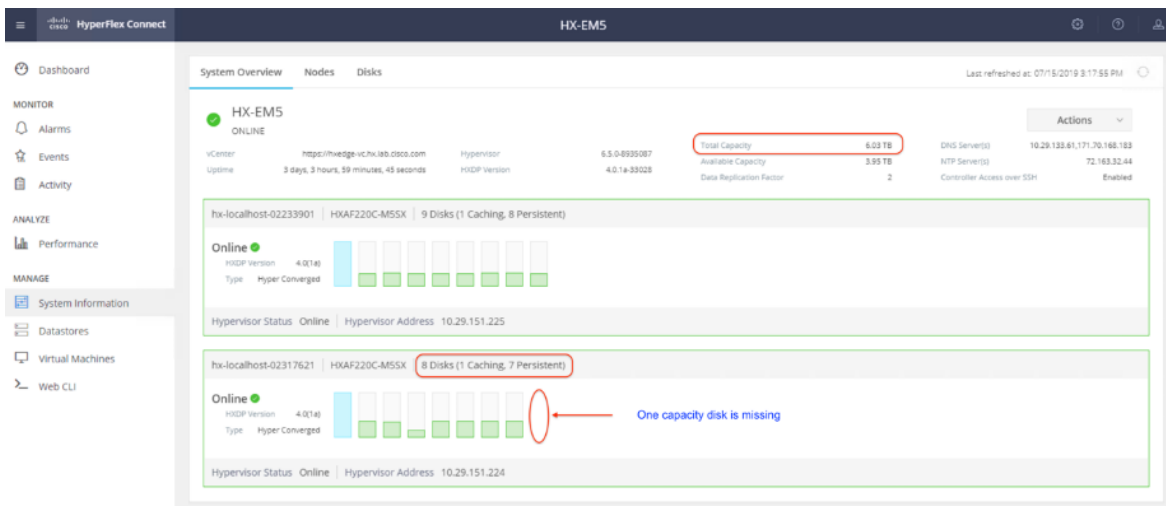
- Check the cluster status from the HyperFlex Connect management console.



- Click the Warning Information sign under Resiliency Health, note that one capacity disk is backlisted so the cluster is unhealthy.



- Wait until the system is auto healed. Check the disk status and verify that the cluster total capacity has been decreased.



- Pull out another capacity disk on the second node now. Check the Auto Healing process.

Command	stcli cluster storage-summary
---------	-------------------------------

Only direct commands are supported. To run interactive commands, login to an HX Controller VM command line.

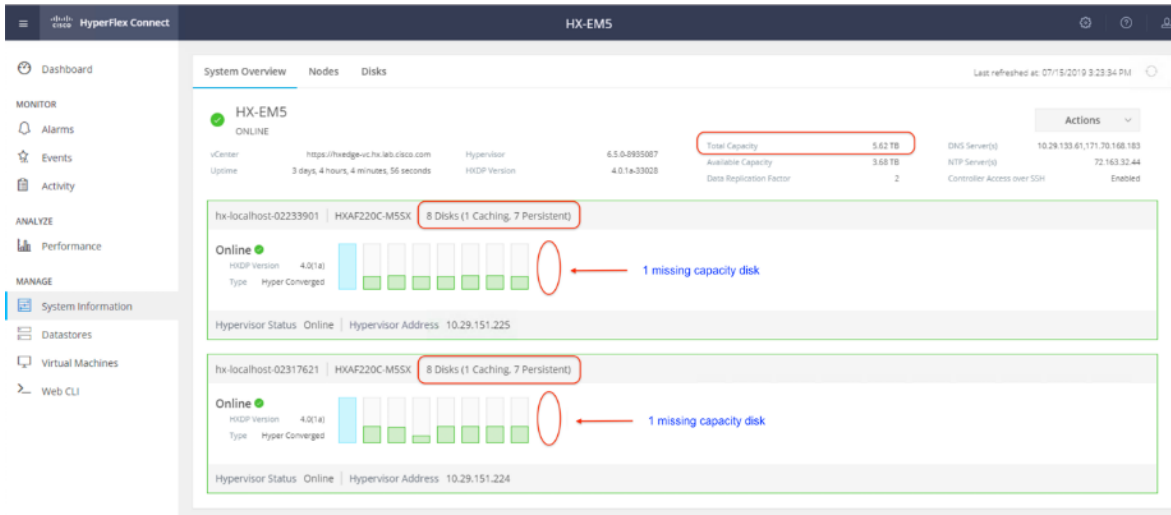
```
stcli cluster storage-summary
```

Output

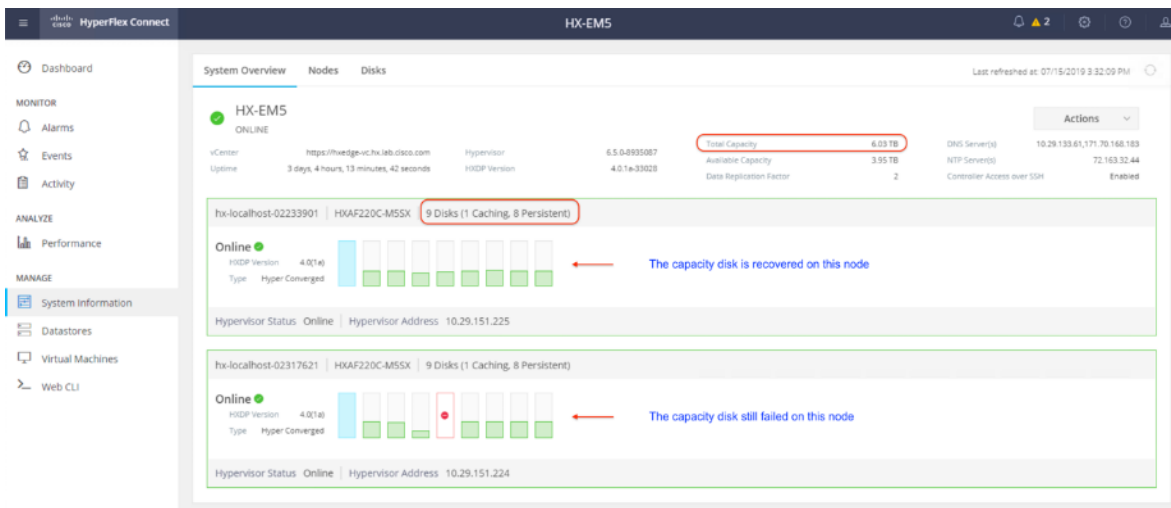
```
address: 169.254.1.20
name: HX-EM5
state: online
uptime: 3 days 4 hours 8 minutes 1 seconds
activeNodes: 2 of 2
compressionSavings: 4.77%
deduplicationSavings: 23.84%
freeCapacity: 3.6T
healingInfo:
  messages:
    Auto healing in progress, 50% completed.
  inProgress: True
  percentComplete: 50
  estimatedCompletionTimeInSeconds: 121
resiliencyInfo:
  messages:
    Storage cluster is unhealthy.
    -----
    Persistent Device Disk [5002538c40a76f2e:0000000000000000] on node 169.254.1.22 is unavailable.
    -----
    Disk is blacklisted.

state: 2
nodeFailuresTolerable: 0
cachingDeviceFailuresTolerable: 1
persistentDeviceFailuresTolerable: 0
zoneResInfoList: None
spaceStatus: normal
totalCapacity: 5.6T
totalSavings: 27.47%
usedCapacity: 2.0T
zkHealth: online
arbitrationServiceState: online
clusterAccessPolicy: lenient
dataReplicationCompliance: non_compliant
```

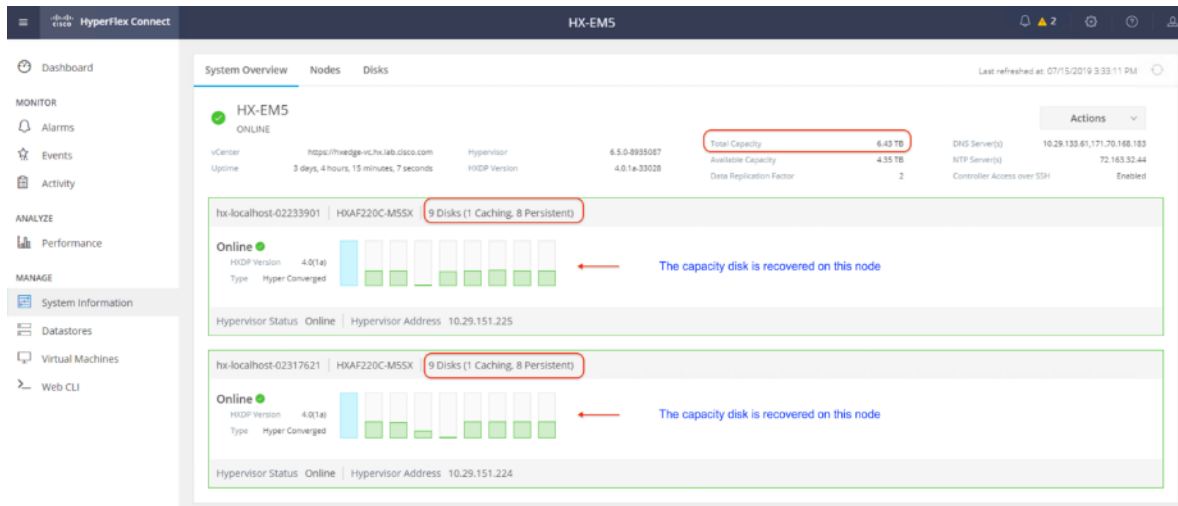
- Wait until the system is auto healed. Check the disk status and verify that the cluster total capacity has been decreased again.



- Put back the capacity disk on the second node. Wait until the system is auto healed. Check the disk status and verify that the cluster total capacity has been increased.

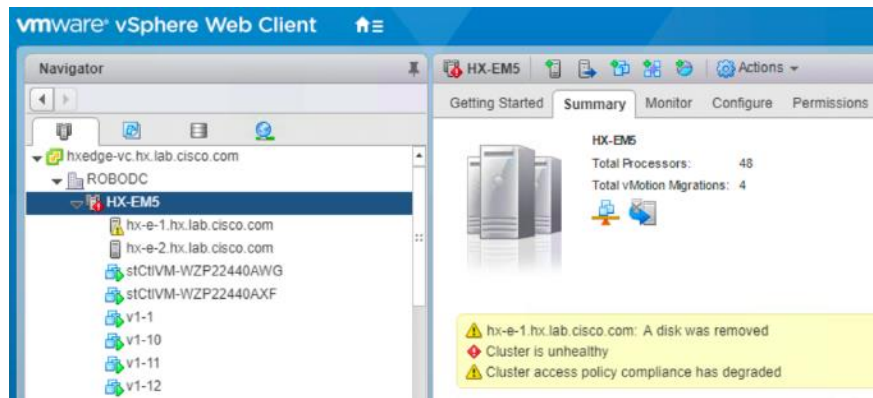


- Put back the capacity disk on another node now. Wait until the system is auto healed. Check the disk status and verify that the cluster total capacity has been restored back to the original number.



Test - Caching Drive Failure

1. Check and verify that the cluster is healthy.
2. Pull out a caching drive on node 1 of the cluster.
3. View alarms displayed in vCenter.



4. Check the Auto Healing process from CLI.

HX-EM5

Command stcli cluster storage-summary

① Only direct commands are supported through HX Connect. To run interactive commands, login to an HX Controller VM command line.

Output

```

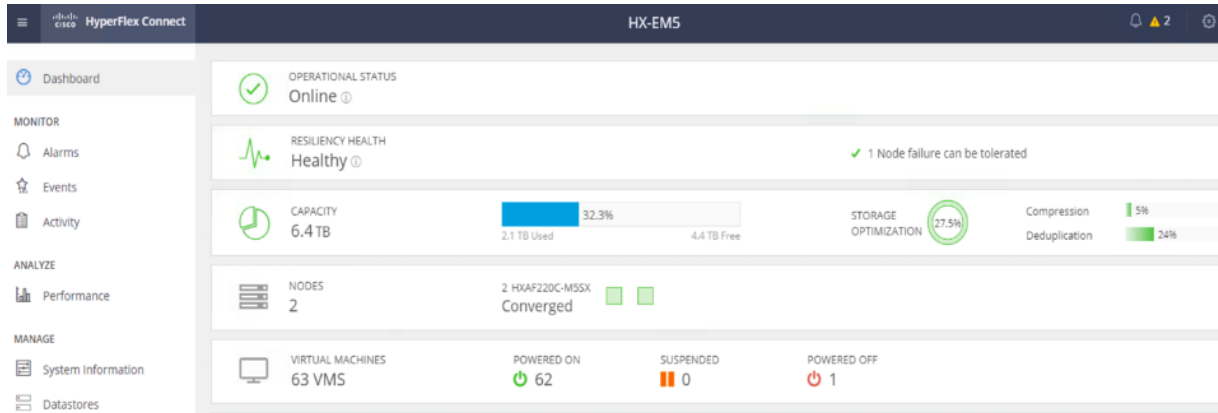
address: 169.254.1.20
name: HX-EM5
state: online
uptime: 4 days 3 hours 34 minutes 3 seconds
activeNodes: 2 of 2
compressionSavings: 4.81%
deduplicationSavings: 24.58%
freeCapacity: 3.6T
healingInfo:
  messages:
    Auto healing in progress, 75% completed.
  inProgress: True
  percentComplete: 75
  estimatedCompletionTimeInSeconds: 0
resiliencyInfo:
  messages:
    -----
    Storage cluster is unhealthy.
    Caching Device Disk [58ce38ee202e9b11:0000000000000000] on node 169.254.1.21 is unavailable.
    Disk is blacklisted.
    -----
  state: 2
  nodeFailuresTolerable: 0
  cachingDeviceFailuresTolerable: 0
  persistentDeviceFailuresTolerable: 1
  zoneResInfoList: None
spaceStatus: normal
totalCapacity: 6.4T
totalSavings: 28.21%
usedCapacity: 2.9T
zkHealth: online
arbitrationServiceState: online
clusterAccessPolicy: lenient
dataReplicationCompliance: non_compliant
dataReplicationFactor: 2
        
```

5. Wait until the system is auto healed. Check the cluster and caching disk status.



Note: Loss of a caching disk is similar to loss of a node. The system cannot bear two caching disk failure simultaneously.

6. Insert the caching disk. Wait until the system is auto healed. Verify that the cluster is back to the healthy state.



Bill of Materials

This section provides the BOM of the Cisco devices that are ordered to build the HyperFlex 2-node Edge Clusters to validate this solution.

Line Number	Item Name	Item Description	Total Line Ordered Quantity	Product Family	Notes
1	HX-E-M5S-HXDP	Cisco HyperFlex M5 Edge Hybrid & All Flash (w/o FI)	1	UCSHX	HX Edge Bundle
1.1	HXAF-E-220M5SX	Cisco HyperFlex All Flash Edge 220 M5 system	4	UCSHX	Node Configuration
1.1.1	HX-MR-X32G2RS-H	32GB DDR4-2666-MHz RDIMM/PC4-21300/dual rank/x4/1.2v	48	UCSHX	
1.1.2	HX-SD960G61X-EV	960GB 2.5 inch Enterprise Value 6G SATA SSD	32	UCSHX	
1.1.3	HX-SD400G12TX-EP	400GB 2.5in Enterprise Performance 12G SAS SSD (10Xendurance)	4	UCSHX	

Line Number	Item Name	Item Description	Total Line Ordered Quantity	Product Family	Notes
1.1.4	HX-SD240G61X-EV	240GB 2.5 inch Enterprise Value 6G SATA SSD	4	UCSHX	
1.1.5	HX-M2-240GB	240GB SATA M.2	4	UCSHX	
1.1.6	HX-PSU1-1050W	Cisco UCS 1050W AC Power Supply for Rack Server	8	UCSHX	
1.1.7	CAB-C13-C14-AC	Power cord, C13 to C14 (recessed receptacle), 10A	8	DSBUOTH	
1.1.8	HX-MSD-32G	32GB Micro SD Card for UCS M5 servers	4	UCSHX	
1.1.9	HX-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	4	UCSHX	
1.1.10	UCSC-HS-C220M5	Heat sink for UCS C220 M5 rack servers 150W CPUs & below	8	UCSC	
1.1.11	UCS-MSTOR-M2	Mini Storage carrier for M.2 SATA/NVME (holds up to 2)	4	UCS	
1.1.12	HX-E-220C-BZL-M5	HX220C M5 Edge Security Bezel	4	UCSHX	
1.1.13	HX-SAS-M5	Cisco 12G Modular SAS HBA (max 16 drives)	4	UCSHX	

Line Number	Item Name	Item Description	Total Line Ordered Quantity	Product Family	Notes
1.1.14	HX-VSP-6-5-FND-D	Factory Installed -vSphere SW 6.5 End user to provide License	4	UCSHX	
1.1.15	HX-VSP-6-5-FND-DL	Factory Installed - VMware vSphere 6.5 SW Download	4	UCSHX	
1.1.16	HX-CPU-4116	2.1 GHz 4116/85W 12C/16.50MB Cache/DDR4 2400MHz	8	UCSHX	
1.2	HXDP-E001-1YR=	Cisco HyperFlex Data Platform Edge Edition 1 Year Subscription	4	UCSHX	License
1.2.0.1	HXDPE001-1YR	HyperFlex Data Platform Edge Edition 1 Year Subscription	4	UCSHX	
2	HX-MLOM-C25Q-04=	Cisco UCS VIC 1457 Quad Port 10/25G SFP28 CNA MLOM	4	HX Compute	Cisco VIC 1457

Summary

Cisco HyperFlex Edge Systems combine software-defined computing in the form of Cisco UCS servers and software-defined storage with the powerful Cisco HyperFlex HX Data Platform software. This unique combination of hardware and software being engineered together as a single solution, reduces risk and offers simplicity for remote location, ROBO and edge environments.

Cisco Intersight cloud management platform is designed to deploy, monitor, manage and provide additional value-added services with strengthened security to Cisco UCS and HyperFlex products. The cloud-powered intelligence from Intersight allows all features to be enabled and performed remotely without requiring IT staff to be physical present near the hardware. Cisco Intersight also brings in the innovative Invisible Cloud Witness service for HyperFlex Edge two-node clusters so that the user does not need to manually install and configure any witness nodes that are required to create a quorum to maintain cluster consistency and ensure high availability. This process eliminates the cost and complexity of deploying and maintaining a dedicated witness server or multiple servers.

As described in this document, HyperFlex Edge Systems can be remotely deployed and managed from the cloud with Cisco Intersight. This is a big step forward towards Cisco's strategy to build cloud managed, on-premise, intelligent infrastructures. With Cisco Intersight, deploying HyperFlex Edge systems can be performed through fully automated policies and zero touch installations, thereby eliminating configuration errors with guaranteed consistency. Deployment through the cloud enables parallel installations from one to many sites simultaneously. Cisco Intersight combined with HyperFlex will completely transform how we manage IT infrastructure in the future.

For More Information

For additional information, see the following:

- Cisco HyperFlex products, services, and solutions: <https://www.cisco.com/go/hyperflex>
- Cisco Intersight Cloud Management Platform: <https://www.cisco.com/go/intersight>
- Cisco Intersight Help Center: <https://www.intersight.com/help/resources>

Appendix A: HyperFlex Cluster Capacity Calculations

A HyperFlex HX Data Platform cluster capacity is calculated as follows:

$$\frac{((\langle \text{capacity disk size in GB} \rangle \times 10^9) / 1024^3) \times \langle \text{number of capacity disks per node} \rangle \times \langle \text{number of HyperFlex nodes} \rangle \times 0.92}{\text{replication factor}}$$

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:

$\langle \text{capacity disk size in GB} \rangle = 960$ GB disks

$\langle \text{number of capacity disks per node} \rangle = 8$ for an HXAF-E-220M5SX model server

$\langle \text{number of HyperFlex nodes} \rangle = 2$

replication factor = 2

Result: $\frac{((960 \times 10^9) / 1024^3) \times 8 \times 2 \times 0.92}{2} = 6580.3528$

$6580.3528 / 1024 = 6.42$ TiB

Appendix B: HyperFlex Sizer

HyperFlex sizer is a cloud based end-to-end tool that can help the customers and partners find out how many Cisco HyperFlex nodes are needed and how the nodes can be configured to meet their needs for the compute resources, storage capacity and performance requirements in the datacenter. This cloud application can be accessed from anywhere from the Cisco website (CCO login required):

<https://hyperflexsizer.cloudapps.cisco.com>

Figure 27 HyperFlex Sizer



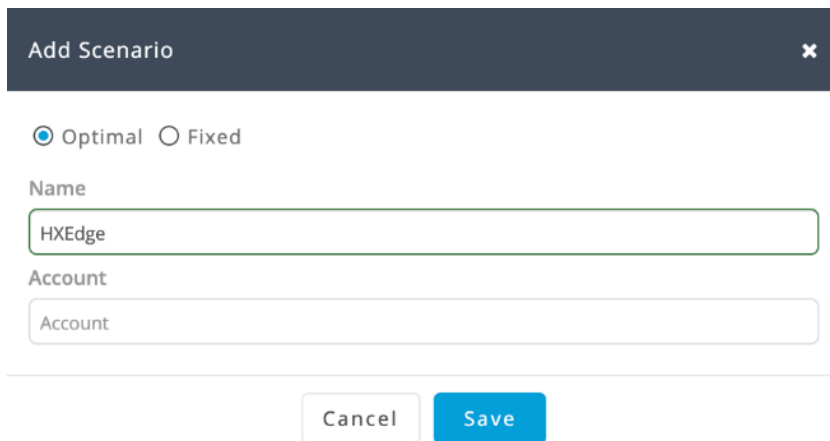
Note: The HyperFlex Sizer tool is designed to provide general guidance in evaluating the optimum solution for using selected Cisco products. The tool is not intended as a substitute for your own judgment or for that of your professional advisors.

The sizing guidance of the HyperFlex system is calculated according to the information of workloads collected from the users. The HyperFlex Sizer supports the following workloads:

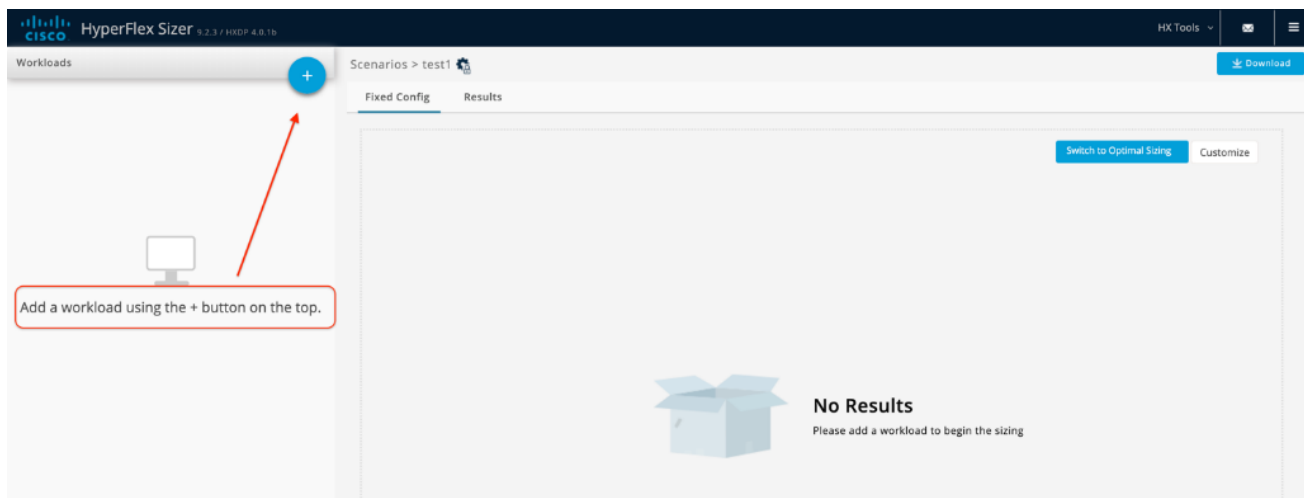
- Virtual Desktop Infrastructure (VDI)
- General Server Virtualized Environment (VSI)
- Microsoft SQL database
- Oracle
- Microsoft Exchange Server
- Compute and Capacity Sizer (RAW)
- HX Edge (ROBO)

To calculate a cluster size for the pre-defined HX Edge workload, follow these steps:

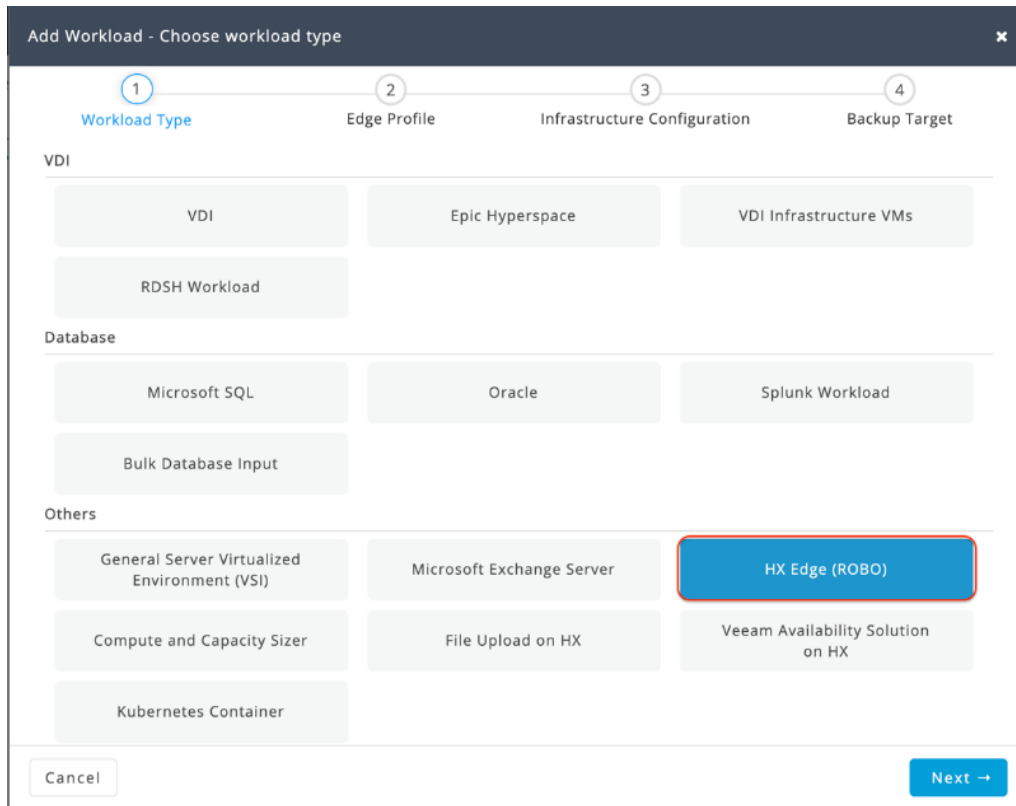
1. Using a web browser, navigate to the Cisco HyperFlex Sizer: <https://hyperflexsizer.cloudapps.cisco.com>.
2. Enter the username and password, login with your Cisco account.
3. click the Add Scenario button to create a new scenario.
4. Enter the name for the new Scenario, then click Save.



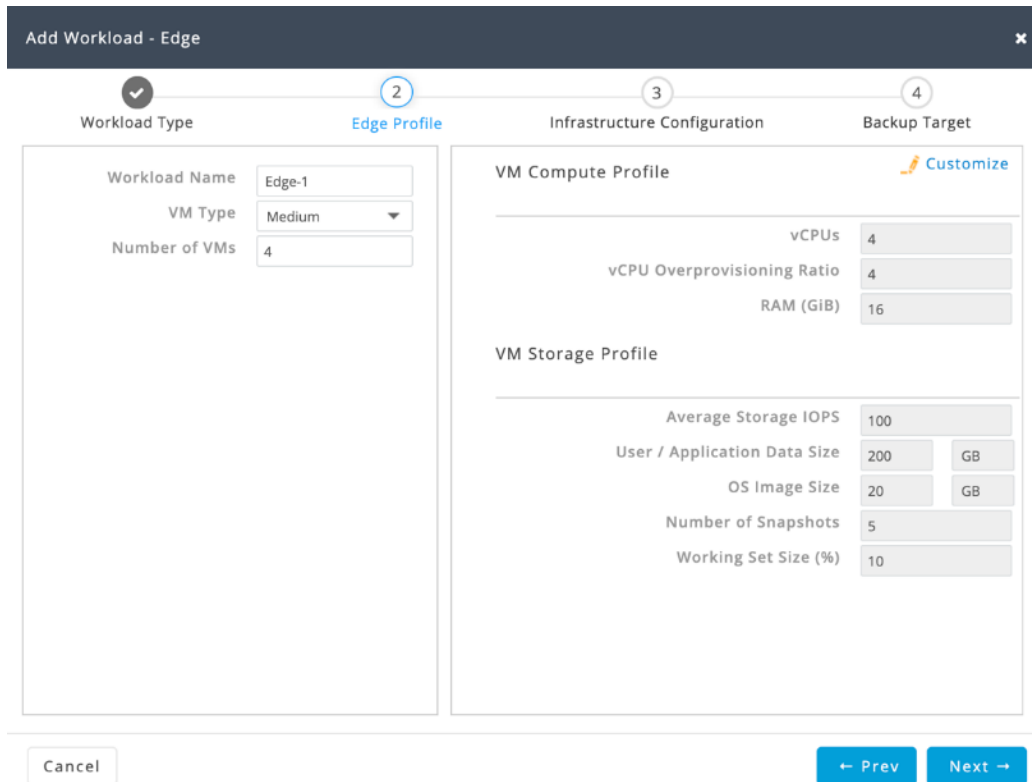
5. When the Workloads page displays, click + to add a new workload.



6. Choose HX Edge (ROBO) as the new workload type, then click Next.



7. In the Edge Profile screen, choose the VM Type and define the number of VMs in your environment. Customize the VM Compute Profile and VM Storage Profile appropriately if needed. Click Next to continue.



- In the Infrastructure Configuration screen, choose the network topology for the HyperFlex Edge cluster, set the values for Replication Factor and Performance Headroom. Define the Compression and Deduplication savings you expect. Click Next to continue.

The screenshot shows the 'Add Workload - Edge' configuration window with four steps: Workload Type, Edge Profile, Infrastructure Configuration (active), and Backup Target. The Infrastructure Configuration step includes the following fields and values:

- Data Replication Factor:** RF 2
- Performance Headroom (nodes):** 1
- NIC Details:** 10G
- Compression Savings (%):** 20
- Deduplication Savings (%):** 10
- Net Savings:** 28.0% (1.39 : 1)

Informational messages are present:

- "Edge workload is only supported with RF2 and N+0 / N+1 configuration"
- "These compression savings can be increased by using the HyperFlex Acceleration Engine."
- Example for compression: "Eg: 100.0 GB effective capacity, with 20% compression savings will require 80.0 GB of usable space"
- Example for deduplication: "Eg: 80.0 GB of capacity post compression, with 10% dedupe will require 72.0 GB usable space"

Navigation buttons at the bottom include 'Cancel', '← Prev', and 'Next →'.

- Ignore the choice of Backup strategy for now. Click Save.

Add Workload - Edge
✕

✓ Workload Type
✓ Edge Profile
✓ Infrastructure Configuration
4

Backup Target

We recommend having a backup strategy for your production environment. In addition to protection from other failures – having an active backup strategy and regular backups does mitigate risk of outage due to component and node failures.

For further documentation, the links to the specific backup solutions for HX are provided below.

Veeam Availability on HX

Cohesity

Veeam

Commvault

Cancel

← Prev

Save

- Based on the workload information you input, the sizer will suggest you with an optimized HyperFlex Edge cluster with the node configuration provided.

The screenshot shows the HyperFlex Sizer interface with the following details:

- Workloads:** Edge-1 (Medium), Cluster 1
- Scenarios > hxedge:**
 - Lowest Cost | All-Flash
 - Threshold: Standard (selected)
 - Node Choice: HyperFlex & Compute (selected)
 - Hypervisor: ESXi (selected)
- Aggregate Summary:**
 - 1 Workloads, 1 Clusters
 - 1+1 (FT) Nodes, 2 Rack Units
- Utilization-Cluster 1:**
 - CPU: 27% (green), 54% (orange)
 - RAM: 40% (green), 80% (orange)
 - Storage Capacity: 21% (green), 21% (orange)
 - Storage IOPS: 3% (green), 6% (orange)
- Node Results Table:**

Cluster	Settings	Part	Type	Description	Count
Cluster 1	RF 2 N+1	HX-E-220M55X	CTO	1x Intel Xeon Silver 4114 Processor, 10 cores, 2.20 GHz 128 [8x16] GiB DDR4 RAM 3x 1.2TB, 2.5" HDD 1x 480GB SATA 1 RU 10G/40G modular LAN	2

Appendix C: Example Cisco Nexus 9348GC-FXP Switch Configuration

```

!Command: show running-config
!Running configuration last done at: Thu Jul 18 22:32:00 2019
!Time: Thu Aug  8 17:19:24 2019

version 9.2(2) Bios:version 05.33
switchname HXEDGE-N9Ka
vdc HXEDGE-N9Ka id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8

feature sftp-server

no password strength-check
username admin password 5 xxxxxx role network-admin
username huich password 5 xxxxxx role network-admin
ip domain-lookup
ip domain-name hx.lab.cisco.com
ip name-server 10.29.x.x
system default switchport
copp profile strict
snmp-server user admin network-admin auth md5 0x804a62fe353f07743c01527714b9559c
  priv 0x804a62fe353f07743c01527714b9559c localizedkey
snmp-server user huich network-admin auth md5 0x804a62fe353f07743c01527714b9559c
  priv 0x804a62fe353f07743c01527714b9559c localizedkey
rmon event 1 description FATAL(1) owner PMON@FATAL
rmon event 2 description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 171.68.x.x use-vrf default

vlan 1,101-104
vlan 101
  name HX-MGMT
vlan 102
  name HX-Data
vlan 103
  name HX-VMotion
vlan 104
  name HX-VMnet

vrf context management
  ip route 0.0.0.0/0 10.29.x.x

interface Ethernet1/1
  description accessLink-lab151
  switchport access vlan 101
  spanning-tree port type network

interface Ethernet1/2

interface Ethernet1/3

interface Ethernet1/4

```

```
interface Ethernet1/5
  description peer-link
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type network
  mtu 9216

interface Ethernet1/6

.....

interface Ethernet1/47

interface Ethernet1/48

interface Ethernet1/49
  description HXEdgeM5-1-Port2
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  mtu 9216

interface Ethernet1/50
  description HXEdgeM5-2-Port2
  switchport mode trunk
  switchport trunk allowed vlan 101-104
  spanning-tree port type edge trunk
  mtu 9216

interface Ethernet1/51

interface Ethernet1/52

interface Ethernet1/53

interface Ethernet1/54

interface mgmt0
  vrf member management
  ip address 10.29.x.x/24
line console
line vty
boot nxos bootflash:/nxos.9.2.2.bin
no feature signature-verification
no system default switchport shutdown
```

Appendix D: Example Cisco Catalyst 9300-48P Switch Configuration

```

Current configuration : 7537 bytes
!
! Last configuration change at 22:48:38 UTC Thu Jul 18 2019 by admin
!
version 16.6
no service pad
service timestamps debug datetime msec
service timestamps log datetime msec
no platform punt-keepalive disable-kernel-core
!
hostname HXEDGE-C9Ka
!
!
vrf definition Mgmt-vrf
!
  address-family ipv4
  exit-address-family
!
  address-family ipv6
  exit-address-family
!
enable secret 5 xxxxxx
enable password xxxxxx
!
no aaa new-model
switch 1 provision c9300-48p
!

.....

!
interface GigabitEthernet0/0
vrf forwarding Mgmt-vrf
ip address 10.29.x.x 255.255.255.0
speed 1000
negotiation auto
!
interface GigabitEthernet1/0/1
description accessLink-lab145
switchport access vlan 201
!
interface GigabitEthernet1/0/2
description peer-link
switchport trunk allowed vlan 201-204
switchport mode trunk
!
interface GigabitEthernet1/0/3
!

.....

!
interface GigabitEthernet1/0/48
!
interface GigabitEthernet1/1/1
!
interface GigabitEthernet1/1/2
!

```

```

interface GigabitEthernet1/1/3
!
interface GigabitEthernet1/1/4
!
interface TenGigabitEthernet1/1/1
description HXEdgeM5-3-Port2
switchport trunk allowed vlan 201-204
switchport mode trunk
spanning-tree portfast
!
interface TenGigabitEthernet1/1/2
description HXEdgeM5-4-Port2
switchport trunk allowed vlan 201-204
switchport mode trunk
spanning-tree portfast
!
interface TenGigabitEthernet1/1/3
!
.....

!
interface TenGigabitEthernet1/1/8
!
interface FortyGigabitEthernet1/1/1
!
interface FortyGigabitEthernet1/1/2
!
interface Vlan1
no ip address
shutdown
!
ip forward-protocol nd
ip http server
ip http authentication local
ip http secure-server
ip route 0.0.0.0 0.0.0.0 10.29.151.1
!
!
!
!
control-plane
service-policy input system-cpp-policy
!
!
line con 0
stopbits 1
line vty 0 4
password xxxxxx
login
line vty 5 15
password xxxxxx
login
!
!
mac address-table notification mac-move
!
!
!
!
!
End

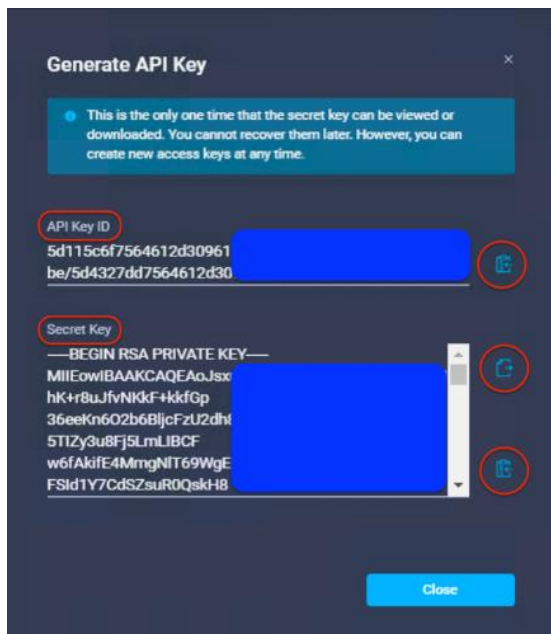
```

Appendix E: Example Multi-site Deployments using Intersight API

An open repository called **hx_intersight_deploy** in GitHub holds some pre-production Python codes that can do HyperFlex Edge deployments at scale via Intersight API. The codes are programmed with Intersight Python SDK and are available here: https://github.com/michzimm/hx_intersight_deploy. What **hx_intersight_deploy** do is that it automatically generates multiple HyperFlex cluster profiles in Intersight from data provided in an excel spreadsheet. The codes were tested and worked but certainly not with the quality for production yet. Extra efforts might need, from yourself, to solidify the codes to meet your own requirements.

To use this GitHub **hx_intersight_deploy** package to simultaneously deploy HyperFlex Edge clusters at multiple sites, follow these steps:

1. Install Python 2.7 on your work station.
2. Download Intersight Python SDK, install the package following the instructions: <https://intersight.com/apidocs/downloads/>.
3. Download the "hx_intersight_deploy" repository from GitHub to your work station: https://github.com/michzimm/hx_intersight_deploy.
4. Set up your Python environment with the required versions by using the provided file of "requirements.txt": `pip install -r requirements.txt`
5. Log into <https://intersight.com> using your Cisco Intersight account.
6. Go to Settings, click API -> API Keys -> Generate API Key, to create the unique API Key ID and Secret Key.



7. Copy the API Key ID. Edit the file with the name of "auth", paste the API Key ID as the value for "api_key_id" in the file.
8. Save the Secret Key to a text file on the work station. Edit the file "auth" and add the full path to this file as the value for "api_private_key_file" in the file.

Example "auth" file:

```
{
  "api_base_uri": "https://intersight.com/api/v1",
  "api_key_id": "<api_key_id>",
  "api_private_key_file": "<path_to_private_key_file">
}
```

Intersight API Key ID

full path to Secret Key text file

9. Edit the file with the name of "input.xlsx", put the definition of your HyperFlex cluster profiles and the values for different HyperFlex policies in the file.
10. The file "input.xlsx" is an excel spreadsheet file used to pass your configuration data to the Intersight API. Each row in the spreadsheet represents one HyperFlex Edge cluster profile which will be created. Follow the format of the sample file, and instructions at the GitHub page how to define the policies and cluster profiles for multiple clusters. You can define multiple cluster profiles with the same or variable policies in one single file for multi-site deployments.
11. From your work station, run the command: `hx_intersight_deploy.py -a auth -f ./input.xlsx`.

```
[root@rh73hui hx_intersight_deploy-master]# ./hx_intersight_deploy.py -a auth -f input.xlsx

Please select from the following options:
 1. Claim HyperFlex nodes in Intersight
 2. Create HyperFlex Cluster Profiles in Intersight
 3. Assign claimed HyperFlex nodes to HyperFlex Cluster Profiles in Intersight
 4. Perform all of the above

Enter the number for your selection: █
```

12. There are three things that `hx_intersight_deploy.py` can do:
 - Claim HyperFlex nodes in Intersight
 - Create HyperFlex Cluster Profiles in Intersight
 - Assign claimed HyperFlex nodes to HyperFlex Cluster Profiles in Intersight
13. You can choose option 1, 2, or 3 to complete these things separately one by one.
14. Or you can perform all of the above with the choice of option 4. Then you will be asked for the new passwords for ESXi Hypervisor and the HyperFlex cluster, and the admin password for your vCenter and CIMC for the HyperFlex servers. After that the HyperFlex Cluster Profiles for multi-site will be created.

```
[root@rh73hui hx_intersight_deploy-master]# ./hx_intersight_deploy.py -a auth -f input.xlsx

Please select from the following options:
 1. Claim HyperFlex nodes in Intersight
 2. Create HyperFlex Cluster Profiles in Intersight
 3. Assign claimed HyperFlex nodes to HyperFlex Cluster Profiles in Intersight
 4. Perform all of the above

Enter the number for your selection: 4

collecting required passwords...
Please enter the new ESXi hypervisor password:
Please confirm the new ESXi hypervisor password:
Do you want the HyperFlex cluster to use the same password? (yes/no): yes
Is the vCenter password the same? (yes/no): no
Please enter the vCenter password:
Please enter the CIMC password for the HyperFlex nodes:

-----
| HyperFlex Cluster Profiles | HyperFlex Nodes |
|-----|-----|
| SanJose=created           | 10.29.151.220=already claimed:assigned, 10.29.151.221=already claimed:assigned |
| Oakland=created          | 10.29.145.221=already claimed:assigned, 10.29.145.220=already claimed:assigned |
|-----|-----|
```

15. Log back into Cisco Intersight, you still need to complete the last step of Validate and Deploy for each cluster profile in Intersight to create the HyperFlex Edge clusters at different sites.

Appendix F: Example Script for ESXi Post-Install Configuration

```
# Configure_ESX_post_install.ps1
# Description: Configures ESXi options and settings after HyperFlex installation.
# Usage: Modify the variables to specify the ESXi root password, the servers to be
# configured, the guest VLAN ID, and the IP addresses used for the vMotion vmkernel
# interfaces.
#
Set-PowerCLIConfiguration -InvalidCertificateAction Ignore -Confirm:$false | Out-Null

$domainname ="hx.lab.cisco.com"

$array="hx3edge-1.hx.lab.cisco.com","hx3edge-2.hx.lab.cisco.com","hx3edge-3.hx.lab.cisco.com"

$ip=11

Foreach ($i in $array)
{

Connect-VIServer -server $i -user root -password xxxxxx
$vh=Get-VMHost -Name $i

#disable shell warning
$vh | Set-VMHostAdvancedConfiguration UserVars.SuppressShellWarning 1

#Configuring default DNS suffix
#Write-Host "Configuring DNS and Domain Name on $vmhost" -ForegroundColor Green
#Get-VMHostNetwork -VMHost $vmhost | Set-VMHostNetwork -DomainName $domainname -hostName $myar-
ray2 -DNSAddress $dnsip -Confirm:$false
Get-VMHostNetwork -VMHost $vh | Set-VMHostNetwork -SearchDomain $domainname -Confirm:$false

#configure syslog traffic to send to vCenter or syslog server
Set-VMHostSysLogServer -SysLogServer '10.29.x.x:514' -VMHost $vh

#add port groups to VM-Network
$vswo = Get-VirtualSwitch -VMHost $vh -Name "vswitch-hx-inband-mgmt"
write-host $vswo
$vpgo = Get-VirtualPortGroup -VirtualSwitch $vswo -Name "VM Network"
Set-VirtualPortGroup -VirtualPortGroup $vpgo -VlanID 104

$vmip="169.254.2."+$ip
New-VirtualPortGroup -VirtualSwitch $vswo -Name "VMotion" -VlanID 103
New-VMHostNetworkAdapter -VMHost $vh -VirtualSwitch $vswo -PortGroup "VMotion" -VMotionEnabled
$true -IP $vmip -SubnetMask 255.255.255.0 -Confirm:$false

$ip=$ip+1

Disconnect-VIServer -server $server -Confirm:$False
}
```


About the Authors

Hui Chen, Technical Marketing Engineer, Cisco UCS Data Center Engineering Group, Cisco Systems, Inc.

Hui is a network and storage veteran with over 15 years of experience on the unified computing, Fibre Channel-based storage area networking, the LAN/SAN convergence systems; and how to build end-to-end, from the application, server, networking to storage, solutions in the data center. Currently he focuses on Cisco's Software Defined Storage (SDS) and Hyperconverged Infrastructure (HCI) solutions. Hui is also a seasoned CCIE.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the author would like to thank:

- Jonathan Gorlin, Product Manager, Cisco Systems Inc.
- Brian Everitt, Technical Marketing Engineer, Cisco Systems Inc.
- Michael Zimmerman, Technical Marketing Manager, Cisco Systems Inc.