



Cisco HyperFlex Systems Installation Guide for VMware ESXi, Release 4.5

First Published: 2021-01-06

Last Modified: 2023-05-05

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT 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 THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned 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. (1721R)

© 2021–2023 Cisco Systems, Inc. All rights reserved.



CONTENTS

[Full Cisco Trademarks with Software License](#) ?

PREFACE

[Communications, Services, Bias-free Language, and Additional Information](#) ix

CHAPTER 1

[New and Changed Information](#) 1

[New and Changed Information](#) 1

CHAPTER 2

[Overview](#) 3

[Cisco HyperFlex HX-Series System](#) 3

[Cisco HyperFlex HX-Series System Components](#) 3

[Cisco HyperFlex HX-Series System Configuration Options](#) 5

[Cisco HyperFlex HX-Series System Management Components](#) 8

[Cisco HyperFlex Connect User Interface and Online Help](#) 9

[Dashboard Page](#) 11

[Operational Status Dialog Box](#) 12

[Resiliency Health Dialog Box](#) 13

CHAPTER 3

[Installation Prerequisites](#) 15

[Supported Versions and System Requirements for Cisco HXDP Release 4.5\(x\)](#) 15

[Required Hardware Cables](#) 16

[Host Requirements](#) 16

[Disk Requirements](#) 17

[Port Requirements](#) 19

[HyperFlex External Connections](#) 20

[Fabric Interconnect Uplink Provisioning](#) 22

[Network Settings](#) 25

- VLAN and vSwitch Requirements 26
- Cisco UCS Requirements 27
- Hypervisor Requirements 28
- Storage Cluster Requirements 29
- vCenter Configuration Requirements 30
- System Services Requirements 30
- CPU Resource Reservation for Controller VMs 32
- Memory Usage for Controller VMs - NVMe Drives 33
- Auto Support Requirements 33
- Single Sign On Requirements 34

CHAPTER 4 Install Cisco HyperFlex Systems Servers 35

- Rack Cisco HyperFlex Nodes 35
- Setting Up the Fabric Interconnects 36
 - Configuring the Primary Fabric Interconnect Using Cisco UCS Manager GUI 37
 - Configuring the Secondary Fabric Interconnect Using Cisco UCS Manager GUI 38
 - Configure the Primary Fabric Interconnect Using CLI 39
 - Configure the Subordinate Fabric Interconnect Using CLI 41
 - Verify Console Setup 42
- Connecting HX-Series Servers to Cisco UCS Fabric Interconnects 43
- Overview 43
 - Connecting Converged Nodes to the Fabric Interconnect 44
 - Physical Connectivity Illustrations for Direct Connect Mode Cluster Setup 45

CHAPTER 5 Configure Cisco HyperFlex Systems 47

- Installation Workflow 47
- Deploy HX Data Platform Installer OVA Using vSphere Web Client 48
- Deploy the HX Data Platform Installer OVA with a Static IP Address 50
- Configure Syslog 51
- Configure and Deploy Your HyperFlex Cluster 52
 - Associate HyperFlex Servers 52
 - Configure UCS Manager 53
 - Configure Hypervisor 57
 - Configure IP Addresses 58

Configure Your HyperFlex Cluster	59
Installation of HyperFlex Nodes with GPUs	62
HX Data Platform Installer Navigation Aid Buttons	63
Warnings and Error Messages	63

CHAPTER 6

Configure Licensing with HyperFlex Data Platform	65
Smart Licensing and HyperFlex	65
License Compliance and Feature Functionality	68
License Management for Connected Environments	69
Registering a Cluster with Smart Licensing	69
Registering a Cluster with Smart Software Licensing through HX Connect	69
Registering a Cluster with Smart Software Licensing through a Controller VM	71
Deregistering a Cluster from Smart Licensing	72
Renewing Smart Licensing Authorization	73
License Management for Disconnected Environments	73
Smart Licensing and Smart Software Manager Satellite	74
Specific License Reservation and HyperFlex	74
Installing Specific License Reservation (SLR) Licenses	75
Canceling Specific License Reservation (SLR) Licenses	83
Returning Specific License Reservation (SLR) Licenses	84
Troubleshooting Specific License Reservation (SLR)	87
Facilitating Controller VM Root Access for Air-Gapped Clusters	88

CHAPTER 7

Configure HyperFlex Hardware Acceleration Cards	91
Overview of HyperFlex Hardware Acceleration Cards	91
Install HyperFlex Hardware Acceleration Cards	91
Deploy HX Data Platform Installer OVA Using vSphere Web Client	92
Deploy the HX Data Platform Installer OVA with a Static IP Address	94
Configure and Deploy Your HyperFlex Cluster	96
Enter Credentials	96
Associate HyperFlex Servers	98
Configure UCS Manager	99
Configure Hypervisor	102
Configure IP Addresses	104

Configure Your HyperFlex Cluster	105
Verify Installation of HyperFlex Hardware Acceleration Cards	108
Troubleshoot HyperFlex Hardware Acceleration Cards	108
Additional Information on HyperFlex Hardware Acceleration Cards	108

CHAPTER 8**Post Cluster Configuration Tasks 109**

Post Cluster Configuration Guidelines	109
Enabling PCI Passthrough for a Network Device on a Host	109
Run Post-Installation Script	110
Changing ESXi Host Root Password	113
Changing Storage Controller Password	114
Cisco HyperFlex HTML Plugin for VMware vCenter	114
Add Datastores in the Storage Cluster	115
Set HA Heartbeat	115
Auto Support and Smart Call Home for HyperFlex	115
Configuring Auto Support Using HX Connect	117
Configuring Notification Settings Using CLI	118
Configuring Smart Call Home for Data Collection	118
Replacing Self-Signed with CA-Signed Certificate	121
Replication Pairing	122
Adding Private VLAN	122
About Private VLANs	122
Configuring a Private VLAN on a VM Network without Existing VMs	123
Configuring Private VLAN on ESX Hosts	123
Configuring a Private VLAN on a VM Network with Existing VMs	124
Deleting VMNICs on the vSphere Standard Switch	124
Creating vSphere Distributed Switch	124
Creating Private VLANs on vSphere Distributed Switch	125
Set Private VLAN in Distributed Port Group	125
Distributed Virtual Switches and Cisco Nexus 1000v	126
Hosting vCenter on the HX Data Platform	127
Deploying AMD GPUs	127

CHAPTER 9**Setting Up Multiple HX Clusters 129**

Setting Up Multiple Clusters 129

CHAPTER 10

Expand Cisco HyperFlex System Clusters 131

Cluster Expansion Guidelines 131

ESXi Installation Guidelines 132

Prerequisites When Expanding M5/M6 Clusters 132

Mixed Cluster Expansion Guidelines - Cisco HX Release 5.5(x) 133

Steps During Mixed Cluster Expansion 133

Prerequisites for Adding a Converged Node 134

Preparing a Converged Node 135

Adding a Converged Node to an Existing Cluster 135

Prerequisites for Adding a Compute-Only Node 140

Preparing a Compute-Only Node 141

Verify the HX Data Platform Installer 142

Apply an HX Profile on a Compute-only Node Using UCS Manager 142

Install VMware ESXi on Compute Nodes 142

Adding a Compute-Only Node to an Existing Cluster 143

Resolving Failure of Cluster Expansion 148

Logical Availability Zones 148

CHAPTER 11

Set Up Clusters with Mixed CPUs 153

Overview 153

Prerequisites for Using Mixed CPUs 153

CPU Compatibility with EVC Mode 154

Enable Enhanced vMotion Compatibility (EVC) on an Existing Cluster 154

Adding Newer Generation Servers to Uniform Clusters 155

Adding Mixed or Older Generation Servers to Existing Clusters 155

CHAPTER 12

Cisco HyperFlex Systems Customized Installation Method 157

Overview 157

Installation and Configuration Workflow for Non Pre-Configured Cisco HyperFlex Systems 157

Installing VMware ESXi 158

Configure vMedia and Boot Policies Through Cisco UCS Manager 159

Opening a Remote KVM Console 160

Rebooting Servers 160
 Undo vMedia and Boot Policy Changes 160

APPENDIX A

Lockdown Mode 163

Enable or Disable Lockdown Mode 163
 Enable or Disable Lockdown Mode from the DCUI: 163
 Enable or Disable Lockdown Mode from the vSphere Web Client: 163
 Troubleshoot Lockdown Mode 164
 When at least one host is in Lockdown mode 164
 When the host is in Lockdown mode while the upgrade in progress: 164



Communications, Services, Bias-free Language, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.



CHAPTER 1

New and Changed Information

- [New and Changed Information](#), on page 1

New and Changed Information

The following table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to this guide or of all new features in this release.

Feature	Description	Date Added	Where Documented
Cisco HyperFlex Systems Installation Guide for VMware ESXi, Release 4.5	First release of the 4.5 guide.	January 5, 2021	This guide



CHAPTER 2

Overview

This chapter provides an overview of the components in Cisco HyperFlex Systems:

- [Cisco HyperFlex HX-Series System, on page 3](#)
- [Cisco HyperFlex HX-Series System Components, on page 3](#)
- [Cisco HyperFlex HX-Series System Configuration Options, on page 5](#)
- [Cisco HyperFlex HX-Series System Management Components, on page 8](#)
- [Cisco HyperFlex Connect User Interface and Online Help, on page 9](#)

Cisco HyperFlex HX-Series System

Cisco HyperFlex HX-Series System provides a fully contained, virtual server platform that combines all three layers of compute, storage, and network with the powerful Cisco HX Data Platform software tool resulting in a single point of connectivity for simplified management. Cisco HyperFlex HX-Series System is a modular system designed to scale out by adding HX nodes under a single UCS management domain. The hyperconverged system provides a unified pool of resources based on your workload needs.

Cisco HyperFlex HX-Series System Components

- **Cisco HX-Series Server**—You can use any of the following servers to configure the Cisco HyperFlex System:
 - Converged nodes—**All Flash**: Cisco HyperFlex HXAF240c M5, HXAF220c M5, HXAF240c M4, and HXAF220c M4.
 - Converged nodes—**Hybrid**: Cisco HyperFlex HX240c M5, HX220c M5, HX240c M4, and HX220c M4.
 - Compute-only—Cisco B200 M3/M4, B260 M4, B420 M4, B460 M4, B480 M5, C240 M3/M4, C220 M3/M4, C480 M5, C460 M4, B200 M5, C220 M5, and C240 M5.
- **Cisco HX Data Platform** —The HX Data Platform consists of the following components:
 - **Cisco HX Data Platform Installer**: Download this installer to a server connected to the storage cluster. The HX Data Platform Installer configures the service profiles and policies within Cisco UCS Manager, deploys the controller VMs, installs the software, creates the storage cluster, and updates the VMware vCenter plug-in.

- **Storage Controller VM:** Using the HX Data Platform Installer, installs the storage controller VM on each converged node in the managed storage cluster.
- **Cisco HX Data Platform Plug-in:** This integrated VMware vSphere interface monitors and manages the storage in your storage cluster.

- **Cisco UCS Fabric Interconnects (FI)**

Fabric Interconnects provide both network connectivity and management capabilities to any attached Cisco HX-Series Server.

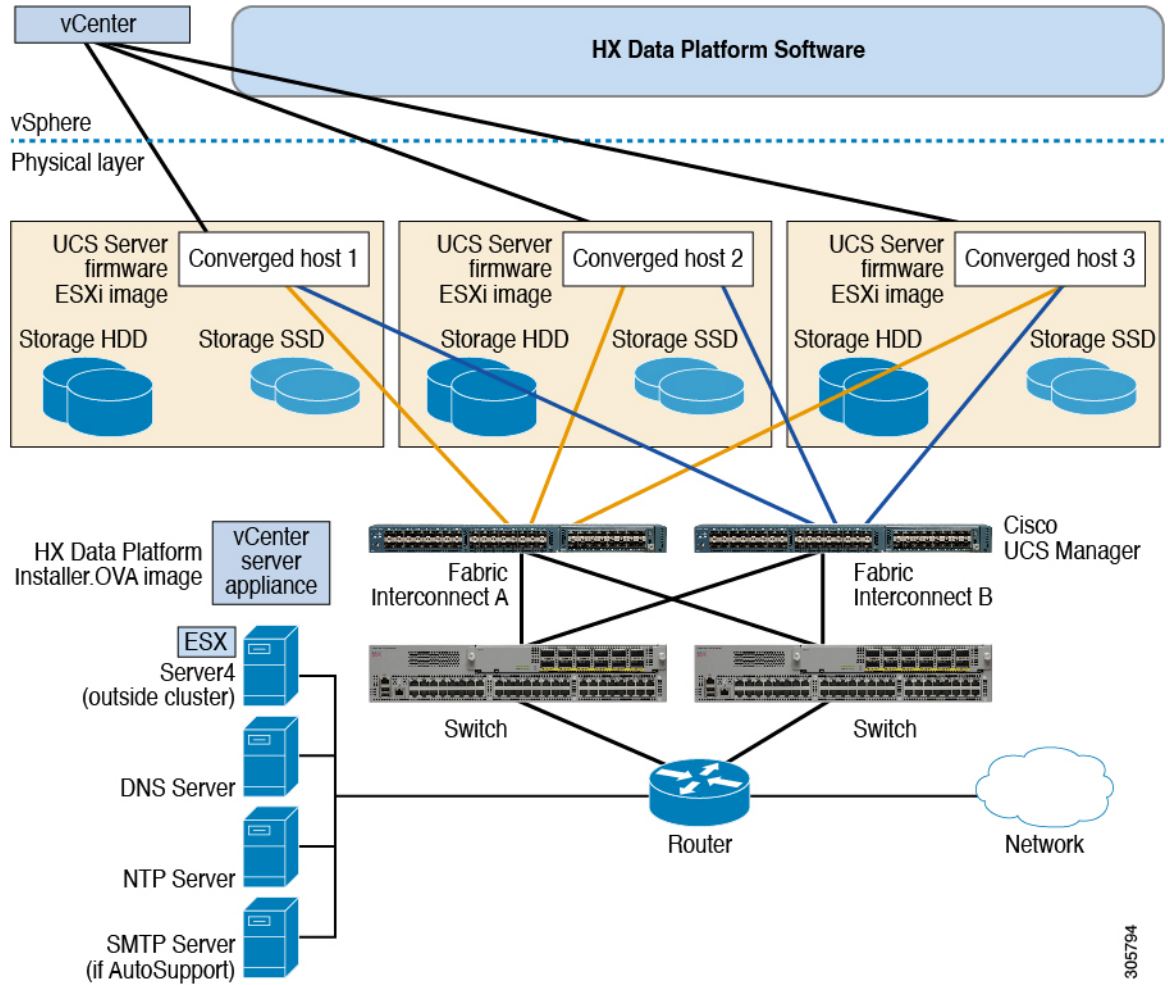
FI that were purchased and deployed as part of the Cisco HyperFlex System are also referred to as an **HX FI Domain** in this document. The following Fabric Interconnects are supported:

- Cisco UCS 6200 Series Fabric Interconnects
- Cisco UCS 6300 Series Fabric Interconnects
- Cisco UCS 6400 Series Fabric Interconnects

- **Cisco Nexus Switches**

Cisco Nexus switches deliver high-density, configurable ports for flexible access deployment and migration.

Figure 1: Cisco HyperFlex HX-Series System Component Details



Cisco HyperFlex HX-Series System Configuration Options


The Cisco HyperFlex HX-Series System offers flexible and scalable options to expand storage and compute capabilities in your environment. To add more storage capabilities to your Cisco HyperFlex System, you simply add a Cisco HyperFlex Server.



Note An **HX Cluster** is a group of HX-Series Servers. Each HX-Series Server in the cluster is referred to as an HX node or a Host.

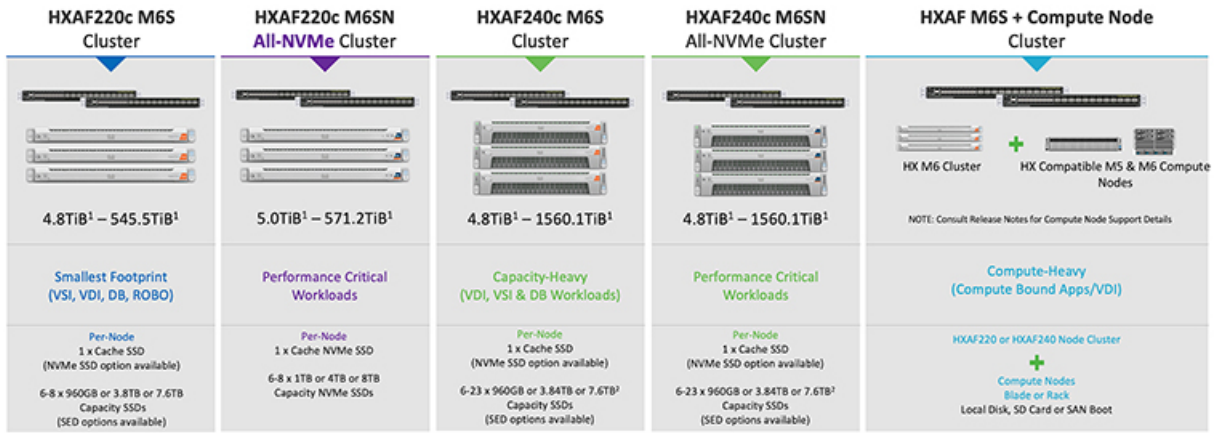
You can configure a HX Cluster in many ways, the following images provide general configuration examples. For the latest compatibility and scalability details consult the [Cisco HX Data Platform Compatibility and Scalability Details - 5.5\(x\) Releases](#) chapter in the [Cisco HyperFlex Recommended Software Release and Requirements Guide](#):

Figure 2: Cisco HyperFlex Hybrid M6 Configurations



¹Usable capacity w/ RFB before compression and deduplication

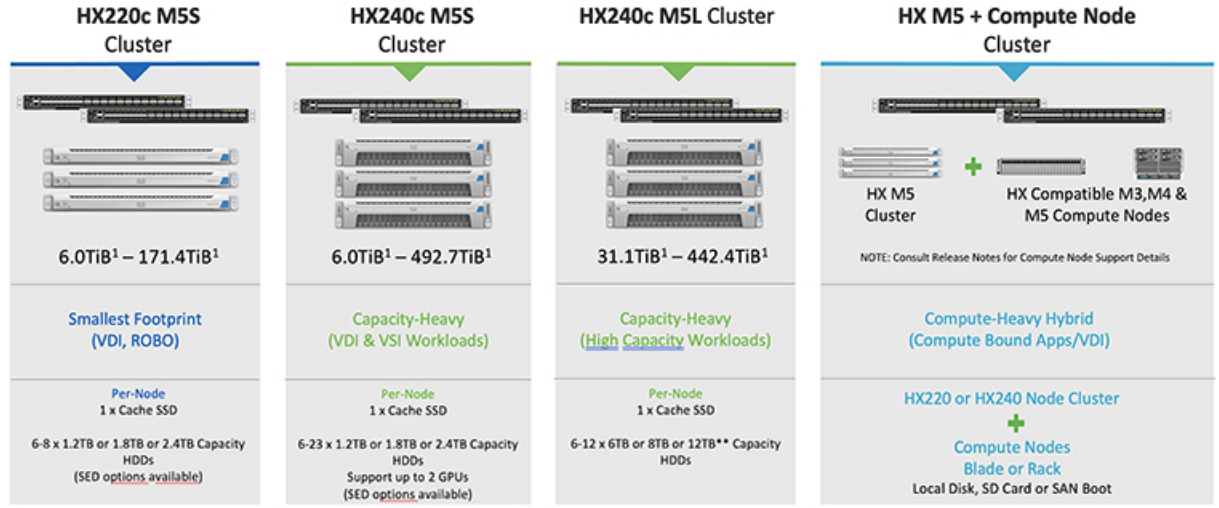
Figure 3: Cisco HyperFlex Hybrid M6 Configurations



¹Usable capacity w/ RFB before compression and deduplication


²Max Converged node limit is 16 when using more than 12 x 7.6TB drives per node

Figure 4: Cisco HyperFlex Hybrid M5 Configurations



¹Usable capacity w/ Rf3 before compression and deduplication

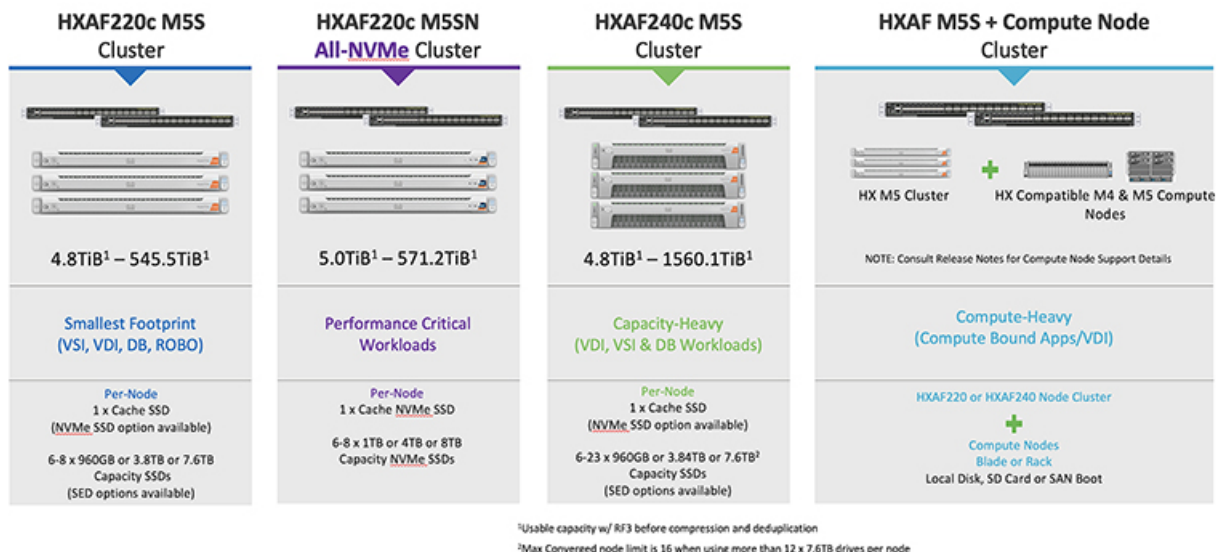
Figure 5: Cisco HyperFlex All Flash M6 Configurations



¹Usable capacity w/ Rf3 before compression and deduplication

²Max Converged node limit is 16 when using more than 12 x 7.6TB drives per node

Figure 6: Cisco HyperFlex All Flash M5 Configurations



Cisco HyperFlex HX-Series System Management Components

The Cisco HyperFlex HX-Series System is managed using the following Cisco software components:

Cisco UCS Manager

Cisco UCS Manager is embedded software that resides on a pair of Fabric Interconnects providing complete configuration and management capabilities for Cisco HX-Series Server. The most common way to access UCS Manager is to use a web browser to open the GUI. UCS Manager supports role-based access control.

The configuration information is replicated between two Cisco UCS Fabric Interconnects (FI) providing a high-availability solution. If one FI becomes unavailable, the other takes over.

A key benefit of UCS Manager is the concept of Stateless Computing. Each node in an HX Cluster has no set configuration. MAC addresses, UUIDs, firmware, and BIOS settings, for example, are all configured on UCS Manager in a Service Profile and applied uniformly to all the HX-Series servers. This enables consistent configuration and ease of reuse. A new Service Profile can be applied within a matter of minutes.

Cisco HX Data Platform

Cisco HX Data Platform is a hyperconverged software appliance that transforms Cisco servers into a single pool of compute and storage resources. It eliminates the need for network storage and tightly integrates with VMware vSphere and its existing management application to provide a seamless data management experience. In addition, native compression and deduplication reduce storage space occupied by the VMs.

HX Data Platform is installed on a virtualized platform, such as vSphere. It manages the storage for your virtual machines, applications, and data. During installation, you specify the Cisco HyperFlex HX Cluster name, and Cisco HX Data Platform creates a hyperconverged storage cluster on each of the nodes. As your storage needs increase and you add nodes to the HX Cluster, Cisco HX Data Platform balances the storage across the additional resources.

VMware vCenter Management

Cisco HyperFlex System has VMware vCenter-based management. The vCenter Server is a data center management server application developed to monitor virtualized environments. The HX Data Platform is also accessed from the preconfigured vCenter Server to perform all storage tasks. vCenter supports key shared storage features like VMware vMotion, DRS, HA, and vSphere replication. More scalable, native HX Data Platform snapshots and clones replace VMware snapshots and cloning capability.

You must have a vCenter installed on a separate server to access HX Data Platform. vCenter is accessed through the vSphere Client, which is installed on the administrator's laptop or PC.

Cisco HyperFlex Connect User Interface and Online Help



Cisco HyperFlex Connect (HX Connect) provides a user interface to Cisco HyperFlex. It is divided into two main sections, a Navigation pane on the left and a Work pane on the right.



Important To perform most actions in HX Connect, you must have administrative privileges.

Table 1: Header Icons

Icon	Name	Description
	Menu	Toggles between the full-size Navigation pane and the icon-only, hover-over Navigation pane.
	Messages	Displays a list of user initiated actions; for example, datastore created, disk removed. Use Clear All to remove all of the messages and hide the Messages icon.
	Settings	Accesses Support , Notification , and Cloud Management settings. You can also access the Support Bundle page.
	Alarms	Displays an alarm count of your current errors or warnings. If there are both errors and warnings, the count shows the number of errors. For more detailed alarm information, see the Alarms page.
	Help	Opens the context-sensitive HX Connect Online Help file.

Icon	Name	Description
	User	Accesses your configurations, such as timeout settings, and log out. User Settings is visible only to administrators.
	Information	Accesses more detailed data about that element.

To access the online help for:

- A particular page in the user interface, click **Help** in the header.
- A dialog box, click **Help** in that dialog box.
- A wizard, click **Help** in that wizard.

Table Header Common Fields

Several tables in HX Connect provide one or more of the following three fields that affect the content displayed in the table.

UI Element	Essential Information
Refresh field and icon	The table automatically refreshes for dynamic updates to the HX Cluster. The timestamp indicates the last time the table was refreshed. Click the circular icon to refresh the content now.
Filter field	Display in the table only list items that match the entered filter text. The items listed in the current page of the table below are automatically filtered. Nested tables are not filtered. Type in the selection text in the Filter field. To empty the Filter field, click the x . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the filter.
Export menu	Save a copy of the current page of table data. The table content is downloaded to the local machine in the selected file type. If the listed items are filtered, the filtered subset list is exported. Click the down arrow to select an export file type. The file type options are: <code>csv</code> , <code>xls</code> , and <code>doc</code> . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the export.

Dashboard Page



Important If you are a read-only user, you may not see all of the options available in the Help. To perform most actions in HyperFlex (HX) Connect, you must have administrative privileges.

Displays a status summary of your HX storage cluster. This is the first page that you see when you log into Cisco HyperFlex Connect.

UI Element	Essential Information
Operational Status section	<p>Provides the functional status of the HX storage cluster and application performance.</p> <p>Click Information (i) to access the HX storage cluster name and status data.</p>
Cluster License Status section	<p>Displays the following link when you log into the HX storage cluster for the first time or till the HX storage cluster license is registered:</p> <p>Cluster License not registered link—Appears when the HX storage cluster is not registered. To register a cluster license, click this link and provide product instance registration token in the Smart Software Licensing Product Registration screen. For more information on how to get a product instance registration token, refer the Registering a Cluster with Smart Licensing section in the Cisco HyperFlex Systems Installation Guide for VMware ESXi.</p> <p>Beginning with HXDP Release 5.0(2a), HX Connect users with expired or insufficient licenses will be unable to access certain features or have limited feature functionality, for more information see License Compliance and Feature Functionality.</p>
Resiliency Health section	<p>Provides the data health status and ability of the HX storage cluster to tolerate failures.</p> <p>Click Information (i) to access the resiliency status, and replication and failure data.</p>
Capacity section	<p>Displays a breakdown of the total storage versus how much storage is used or free.</p> <p>Also displays the storage optimization, compression-savings, and deduplication percentages based on the data stored in the cluster.</p>
Nodes section	<p>Displays the number of nodes in the HX storage cluster, and the division of converged versus compute nodes. Hovering over a node icon displays that node's name, IP address, node type, and an interactive display of disks with access to capacity, usage, serial number, and disk type data.</p>
Performance section	<p>Displays an HX storage cluster performance snapshot for a configurable amount of time, showing IOPS, throughput, and latency data.</p> <p>For full details, see Performance Page.</p>

UI Element	Essential Information
Cluster Time field	System date and time for the cluster.

Table Header Common Fields

Several tables in HX Connect provide one or more of the following three fields that affect the content displayed in the table.

UI Element	Essential Information
Refresh field and icon	The table automatically refreshes for dynamic updates to the HX Cluster. The timestamp indicates the last time the table was refreshed. Click the circular icon to refresh the content now.
Filter field	Display in the table only list items that match the entered filter text. The items listed in the current page of the table below are automatically filtered. Nested tables are not filtered. Type in the selection text in the Filter field. To empty the Filter field, click the x . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the filter.
Export menu	Save a copy of the current page of table data. The table content is downloaded to the local machine in the selected file type. If the listed items are filtered, the filtered subset list is exported. Click the down arrow to select an export file type. The file type options are: <code>cvs</code> , <code>xls</code> , and <code>doc</code> . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the export.

Operational Status Dialog Box

Provides the functional status of the HX storage cluster and application performance.

UI Element	Essential Information
Cluster Name field	Name of this HX storage cluster.
Cluster Status field	<ul style="list-style-type: none"> • Online—Cluster is ready. • Offline—Cluster is not ready. • Read Only—Cluster cannot accept write transactions, but can continue to display static cluster information. • Out of space—Either the entire cluster is out of space or one or more disks are out of space. In both cases, the cluster cannot accept write transactions, but can continue to display static cluster information.

UI Element	Essential Information
Data-at-rest encryption capable field	<ul style="list-style-type: none"> • Available • Not supported <p>Alternatively, Yes and No can be used.</p>
Reason to view drop-down list	Displays the number of messages to explain what is contributing to the current status.

Click **Close**.

Resiliency Health Dialog Box

Provides the data health status and ability of the HX storage cluster to tolerate failures.

Name	Description
Resiliency Status field	<ul style="list-style-type: none"> • Healthy—Cluster is healthy with respect to data and availability. • Warning—Either data or cluster availability is being adversely affected. • Unknown—Transitional state while the cluster is coming online. <p>Color coding and icons are used to indicate various status states. Click an icon to display additional information.</p>
Data Replication Compliance field	<ul style="list-style-type: none"> • Compliant
Data Replication Factor field	Displays the number of redundant data replicas across the HX storage cluster.
Number of node failures tolerable field	Displays the number of node disruptions the HX storage cluster can handle.
Number of Persistent Device failures tolerable field	Displays the number of persistent device disruptions the HX storage cluster can handle.
Number of Caching Device failures tolerable field	Displays the number of cache device disruptions the HX storage cluster can handle.
Reason to view drop-down list	Displays the number of messages to explain what is contributing to the current status.

Click **Close**.



CHAPTER 3

Installation Prerequisites

- [Supported Versions and System Requirements for Cisco HXDP Release 4.5\(x\)](#), on page 15
- [Required Hardware Cables](#), on page 16
- [Host Requirements](#), on page 16
- [Disk Requirements](#), on page 17
- [Port Requirements](#), on page 19
- [HyperFlex External Connections](#), on page 20
- [Fabric Interconnect Uplink Provisioning](#), on page 22
- [Network Settings](#), on page 25
- [VLAN and vSwitch Requirements](#), on page 26
- [Cisco UCS Requirements](#), on page 27
- [Hypervisor Requirements](#), on page 28
- [Storage Cluster Requirements](#), on page 29
- [vCenter Configuration Requirements](#), on page 30
- [System Services Requirements](#), on page 30
- [CPU Resource Reservation for Controller VMs](#), on page 32
- [Memory Usage for Controller VMs - NVMe Drives](#), on page 33
- [Auto Support Requirements](#), on page 33
- [Single Sign On Requirements](#), on page 34

Supported Versions and System Requirements for Cisco HXDP Release 4.5(x)

Cisco HX Data Platform requires specific software and hardware versions, and networking settings for successful installation.

Table 2: Supported Versions and System Requirements for Cisco HXDP Release 4.5(x)

Requirement	Link to Details
Confirm the component firmware on the server meets the minimum versions supported.	Recommended FI/Server Firmware
List of recommended browsers.	Browser Recommendations

Required Hardware Cables

- **6200/6400/6500 series FI:** Use at least two 10-Gb Small Form-Factor Pluggable (SFP) cables per server.
- **6300 series FI:** Use at least two 40-GbE QSFP cables per server.
- Ensure that the Fabric Interconnect console cable (CAB-CONSOLE-RJ45) has an RJ-45 connector on one end and a DB9 connector on the other. This cable is used to connect into the RS-232 console connection on a laptop.

- Ensure that the standard power cords have an IEC C13 connector on the end that plugs into the power supplies. Make sure that the optional jumper power cords have an IEC C13 connector on the end that plugs into the power supplies and an IEC C14 connector on the end that plugs into an IEC C13 outlet receptacle.

For further details, see the [Cisco UCS 6300 Series Fabric Interconnect Hardware Guide](#).

- The KVM cable provides a connection for the Cisco HX-Series Servers into the system. It has a DB9 serial connector, a VGA connector for a monitor, and dual USB 2.0 ports for a keyboard and mouse. With this cable, you can create a direct connection to the operating system and the BIOS running on the system.



Note This same KVM cable is used for both UCS rack mount and blade servers.

For further details on cables and ordering information for M series servers, see the respective [Cisco HyperFlex HX-Series Models](#) and [Cisco UCS B200 Blade Server Installation and Service Note](#).

Host Requirements

A Cisco HyperFlex cluster contains a minimum of three converged HyperFlex nodes. There is an option of adding compute-only nodes to provide additional compute power if there is no need for extra storage. Each server in a HyperFlex cluster is also referred as a HyperFlex node. Make sure that each node has the following settings installed and configured before you deploy the storage cluster.

For further information, refer to the [Cisco HX240c/220c HyperFlex Node Installation Guides](#).

Ensure that the following host requirements are met.

- Use the same VLAN IDs for all the servers (node or hosts) in the cluster.
- Use the same administrator login credentials for all the ESXi servers across the storage cluster.



Note The root user is created with the same password as the admin user during cluster creation. It is important to track the root user password because future changes to the admin password do not automatically update the root password .

- Keep SSH enabled on all ESXi hosts.

- Configure DNS and NTP on all servers.
- Install and configure VMware vSphere.
- **VIC and NIC Support:** For details, see the [Cisco HyperFlex Systems—Networking Topologies](#) document.

Disk Requirements

The disk requirements vary between converged nodes and compute-only nodes. To increase the available CPU and memory capacity, you can expand the existing cluster with compute-only nodes as needed. These compute-only nodes provide no increase to storage performance or storage capacity.

Alternatively, adding converged nodes increase storage performance and storage capacity alongside CPU and memory resources.

Servers with only Solid-State Disks (SSDs) are All-Flash servers. Servers with both SSDs and Hard Disk Drives (HDDs) are hybrid servers.

The following applies to all the disks in a HyperFlex cluster:

- All the disks in the storage cluster must have the same amount of storage capacity. All the nodes in the storage cluster must have the same number of disks.
- All **SSDs** must support TRIM and have TRIM enabled.
- All **HDDs** can be either SATA or SAS type. All SAS disks in the storage cluster must be in a pass-through mode.
- Disk partitions must be removed from SSDs and HDDs. Disks with partitions are ignored and not added to your HX storage cluster.
- Moving operational disks between servers within same cluster or moving them into expansion nodes within the same active cluster is not supported.
- Optionally, you can remove or backup existing data on disks. All existing data on a provided disk is overwritten.



Note New factory servers are shipped with appropriate disk partition settings. Do not remove disk partitions from new factory servers.

- Only the disks ordered directly from Cisco are supported.
- On servers with Self Encrypting Drives (SED), both the cache and persistent storage (capacity) drives must be SED capable. These servers support Data at Rest Encryption (DARE).
- In the event you see an error about unsupported drives or catalog upgrade, see the [Compatibility Catalog](#).

In addition to the disks listed in the table below, all M4 converged nodes have 2 x 64-GB SD FlexFlash cards in a mirrored configuration with ESX installed. All M5 converged nodes have M.2 SATA SSD with ESXi installed.



-
- Note** Do not mix storage disks type or storage size on a server or across the storage cluster. Mixing storage disk types is not supported.
- When replacing cache or persistent disks, always use the same type and size as the original disk.
 - Do not mix any of the persistent drives. Use all HDD or SSD and the same size drives in a server.
 - Do not mix hybrid and All-Flash cache drive types. Use the hybrid cache device on hybrid servers and All-Flash cache devices on All-Flash servers.
 - Do not mix encrypted and non-encrypted drive types. Use SED hybrid or SED All-Flash drives. On SED servers, both the cache and persistent drives must be SED type.
 - All nodes must use same size and quantity of SSDs. Do not mix SSD types.
-

Please refer to the corresponding server model spec sheet for details of drives capacities and number of drives supported on the different servers.

For information on compatible PIDs when performing an expansion of existing cluster, please refer to the [Cisco HyperFlex Drive Compatibility](#) document.

Compute-Only Nodes

The following table lists the supported compute-only node configurations for compute-only functions. Storage on compute-only nodes is not included in the cache or capacity of storage clusters.



-
- Note** When adding compute nodes to your HyperFlex cluster, the compute-only service profile template automatically configures it for booting from an SD card. If you are using another form of boot media, update the local disk configuration policy. See the *Cisco UCS Manager Server Management Guide* for server-related policies.
-

Supported Compute-Only Node Servers	Supported Methods for Booting ESXi
<ul style="list-style-type: none"> • Cisco B200 M4/M5 • B260 M4 • B420 M4 • B460 M4 • C240 M4/M5 • C220 M4/M5 • C460 M4 • C480 M5 • B480 M5 	<p>Choose any method.</p> <p>Important Ensure that only one form of boot media is exposed to the server for ESXi installation. Post install, you may add in additional local or remote disks.</p> <p>USB boot is not supported for HX Compute-only nodes.</p> <ul style="list-style-type: none"> • SD Cards in a mirrored configuration with ESXi installed. • Local drive HDD or SSD. • SAN boot. • M.2 SATA SSD Drive. <p>Note HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform release 4.5(1a) and later.</p>

Port Requirements

If your network is behind a firewall, in addition to the standard port requirements, VMware recommends ports for VMware ESXi and VMware vCenter.

- CIP-M is for the cluster management IP.
- SCVM is the management IP for the controller VM.
- ESXi is the management IP for the hypervisor.

The comprehensive list of ports required for component communication for the HyperFlex solution is located in Appendix A of the [HX Data Platform Security Hardening Guide](#)



Tip If you do not have standard configurations and need different port settings, refer to Appendix A of the [HX Data Platform Security Hardening Guide](#) for customizing your environment.

HyperFlex External Connections

External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
Intersight Device Connector	Supported HX systems are connected to Cisco Intersight through a device connector that is embedded in the management controller of each system.	HTTPS Port Number: 443 1.0.5-2084 or later (Auto-upgraded by Cisco Intersight)	

External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
			<p>All device connectors must properly resolve <code>svc.intersight.com</code> and allow outbound-initiated HTTPS connections on port 443. The current HX Installer supports the use of an HTTP proxy.</p> <p>The IP addresses of ESXi management must be reachable from Cisco UCS Manager over all the ports that are listed as being needed from installer to ESXi management, to ensure deployment of ESXi management from Cisco Intersight.</p> <p>Note Outbound HTTPS connections on port 443 initiated by ESXi hosts can be blocked by the default ESXi firewall. The ESXi firewall can be temporarily disabled to allow this connectivity.</p> <p>To disable the ESXi firewall, use the <code>esxcli network firewall set --enabled=false</code> command and after the installation has completed use the <code>esxcli network firewall set --enabled=false</code> command to re-enable the firewall.</p>

External Connection	Description	IP Address/ FQDN/ Ports/Version	Essential Information
			For more information, see the Network Connectivity Requirements section of the Intersight Help Center.
Auto Support	Auto Support (ASUP) is the alert notification service provided through HX Data Platform.	SMTP Port Number: 25	Enabling Auto Support is strongly recommended because it provides historical hardware counters that are valuable in diagnosing future hardware issues, such as a drive failure for a node.

Fabric Interconnect Uplink Provisioning

Prior to setting up the HyperFlex cluster, plan the upstream bandwidth capacity for optimal network traffic management. This ensures that the flow is in steady state, even if there is a component failure or a partial network outage.

By default, the *hx-vm-network* vSwitch is configured as **active/active**. All other vSwitches are configured as **active/standby**.

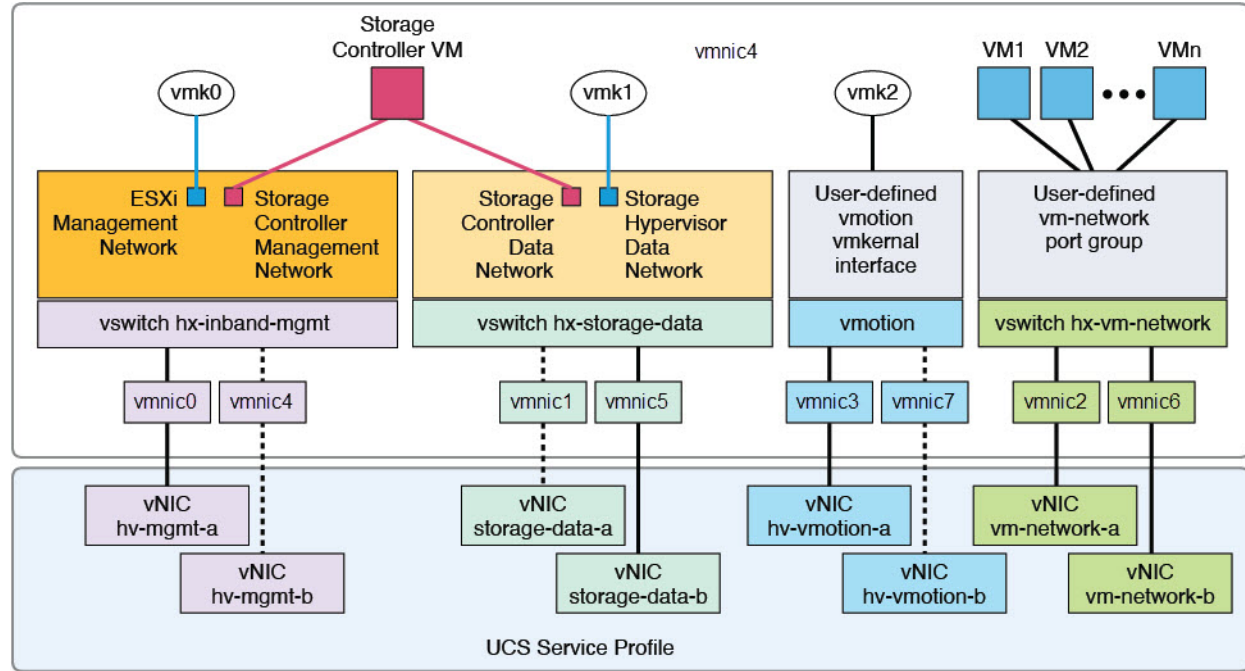


Note All VLANs (including storage-data and vmotion) must be configured upstream. It is imperative that storage-data is configured upstream to prevent fail-back timing issues which require temporary upstream connectivity.



Note For clusters running Catalyst switches upstream to the FI's, set the best effort Quality of Service (QOS) MTU to 9216 (located in LAN > LAN Cloud > QoS System Class), otherwise failover will fail.


Figure 7: HyperFlex Data Platform Connectivity for a Single Host



- Note: 1. Dotted lines represent a “standby” link.
 2. All “a” vNICs connect to FI-A.
 3. All “b” vNICs connect to FI-B.
 4. MTU of 9000 is needed for storage-data and vmotion networks.
 5. All VLANs by default are tagged on the FI so frames are passed untagged to each vswitch.
 6. The vm network port groups are automatically created in 1.8 installer with vlan suffix.


Set the default vSwitch NIC teaming policy and failover policy to **yes** to ensure that all management, vMotion, and storage traffic are locally forwarded to the fabric interconnects to keep the flow in steady state. When vNIC-a fails, ESXi computes the load balancing and all the virtual ports are repinned to vNIC-b. When vNIC-a comes back online, repinning does apply and virtual ports are rebalanced across vNIC-a and vNIC-b. This reduces the latency and bandwidth utilization upstream of the Cisco UCS fabric interconnects.

Figure 8: Traffic Flow in Steady State



In case one or more server links fail, for instance, if Host 1 loses connectivity to Fabric A while Host 2 loses connectivity to Fabric B, the traffic must go through the upstream switches. Therefore, the uplink network bandwidth usage increases, and you must add more uplinks.

Figure 9: Traffic Flow During Link Failure





Note When you have uplinks from a fabric interconnect to two different upstream switches, you encounter a condition called **Disjoint Layer 2** (DJL2) on the FI. This is known to happen on the FI on End Host Mode and if the DJL2 is not configured properly.

To deploy the DJL2 properly, refer to the [Cisco UCS 6300 Series Fabric Interconnect Hardware Guide—Deploy Layer 2 Disjoint Networks Upstream in End Host Mode](#) white paper.

Network Settings



Important All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.



Note You cannot use the storage-data VLAN for User VM traffic.

Best Practices

- Must use different subnets and VLANs for each network.
- Directly attach each host to a Cisco UCS fabric interconnect using a 10-Gbps cable.
- Do not use VLAN 1 which is the default VLAN as it can cause networking issues, especially if Disjoint Layer 2 configuration is used.
- Installer sets the VLANs as *non-native* by default. Ensure to configure the upstream switches to accommodate the non-native VLANs.
- Uplinks from the UCS Fabric Interconnects to all top of rack switch ports need to have Port-fast, spanning-tree port type edge trunk, or similar spanning tree configuration that immediately put ports into forwarding mode. Configure spanning tree in **edge trunk** or **portfast edge** mode depending on the vendor and model of the switch. This extra configuration ensures that when links flap or change state, they do not transition through unnecessary spanning tree states and incur an extra delay before traffic forwarding begins. Failure to properly configure FI uplinks in **portfast edge** mode may result in network and cluster outages during failure scenarios and during infrastructure upgrades that leverage the highly available network design native to HyperFlex.

Each ESXi host needs the following networks.

- **Management traffic network**—From the vCenter, handles the hypervisor (ESXi server) management, and storage cluster management.
- **Data traffic network**— Handles the hypervisor and storage data traffic and is required to be a unique VLAN per Hyperflex Cluster.
- **vMotion network**
- **VM network**

There are four vSwitches, each carrying a different network.

- **vswitch-hx-inband-mgmt**—Used for ESXi management, storage controller management, and Replication. These two vSwitches are further divided in two port groups with assigned static IP addresses to handle traffic between the storage cluster and the ESXi host.
- **vswitch-hx-storage-data**—Used for ESXi storage data and HX Data Platform ISCSI. These two vSwitches are further divided in two port groups with assigned static IP addresses to handle traffic between the storage cluster and the ESXi host. These two vSwitches are further divided in two port groups with assigned static IP addresses to handle traffic between the storage cluster and the ESXi host.
- **vswitch-hx-vmotion**—Used for VM and storage vMotion.

This vSwitch, has one port group for management, defined through vSphere that connects to all the hosts in the vCenter cluster.

- **vswitch-hx-vm-network**—Used for VM data traffic.

You can add or remove VLANs on the corresponding vNIC templates in Cisco UCS Manager. See [Managing VLANs in Cisco UCS Manager](#) and [Managing vNIC templates in Cisco UCS Manager](#) for the detailed steps. To create port groups on the vSwitch, refer to [Adding Virtual Port Groups to VMware Standard vSwitch](#).



- Note**
1. The Cisco HX Data Platform Installer automatically creates the vSwitches.
 2. The following services in vSphere must be enabled after the HyperFlex storage cluster is created.
 - DRS (Optional, if licensed)
 - vMotion
 - High Availability

VLAN and vSwitch Requirements

- Provide at least three VLAN IDs.
- All VLANs must be configured on the fabric interconnects during the installation.
- All VLANs (including storage-data and vmotion) must be configured upstream. It is imperative that storage-data is configured upstream to prevent fail-back timing issues which require temporary upstream connectivity.

VLAN Type	Description
Note	Must use different subnets and VLANs for each of the following networks.
VLAN ESXi and HyperFlex Management Traffic	VLAN Name: <user-defined> (for example "hx-inband-mgmt") VLAN ID: <user-defined>

VLAN Type	Description
VLAN HyperFlex Storage Data	VLAN Name: <user-defined> (for example, "hx-storage-data") VLAN ID: <user-defined>
VLAN VM vMotion	VLAN Name: <user-defined> (for example, "hx-vmotion") VLAN ID: <user-defined>
VLAN VM Network	VLAN VM Network: <user-defined> (for example, "hx-vm-network"). This is required to be a unique VLAN HyperFlex Cluster. VLAN ID: <user-defined>

The VLAN tagging with External Switch VLAN Tagging (EST) and vSwitch settings are applied using UCS Manager profiles. The HX Data Platform Installer, simplifies this process.


Note

- Do not use VLAN 1 which is the default VLAN as it can cause networking issues, especially if Disjoint Layer 2 configuration is used. Use a different VLAN other than VLAN 1.
Installer sets the VLANs as *non-native* by default. Configure the upstream switches to accommodate the non-native VLANs.
- Inband Management is not supported on VLAN 2 or VLAN 3.

Cisco UCS Requirements

Provide the listed content for the UCS Fabric Interconnect and UCS Manager when prompted.

Cisco UCS Fabric Interconnect Requirements

UI Element	Essential Information
Uplink Switch Model	Provide the switch type and connection type (SFP + Twin Ax or Optic).
Fabric Interconnect Cluster IP address	<IP address>.
FI-A IP Address	<IP address>.
FI-B IP Address	<IP address>.
MAC Address Pool	Check 00:00:00 MAC address pool.
IP Blocks	KVM IP pool. A minimum of 4 IP addresses.
Subnet mask	For example, 255.255.0.0.
Default Gateway	For example, 10.193.0.1.

Cisco UCS Manager Requirements

UI Element	Essential Information
UCS Manager Host Name	<i>Hostname or IP address.</i>
User Name	<i><admin username></i>
Password	<i><admin username></i>

Hypervisor Requirements

Enter the IP address from the range of addresses that are available to the ESXi servers on the storage management network or storage data network through vCenter. Provide static IP addresses for all network addresses.

**Note**

- Data and Management networks must be on different subnets.
- Data network IP addresses cannot be changed after the storage cluster is created. Contact Cisco TAC for assistance changing Management Network IPs.
- Though, not required by itself, if you are specifying DNS names, enable IP addresses forward and reverse DNS lookup.
- The installer IP address must be reachable from the management subnet used by the hypervisor and the storage controller VMs. The installer appliance must run on the ESXi host or on a VMware workstation that is not a part of the cluster to be installed.

Management Network IP Addresses		Data Network IP Addresses	
Hypervisor	Storage Controller	Hypervisor	Storage Controller
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>	<i><IP Address ></i>
<i>VLAN Tag</i>	<i>VLAN_ID</i>	<i>VLAN Tag</i>	<i>VLAN_ID</i>
<i>Subnet Mask</i>		<i>Subnet Mask</i>	
<i>Default Gateway</i>		<i>Default Gateway</i>	
Installer Appliance IP Addresses			
<i><IP Address ></i>		<i><IP Address ></i>	

Storage Cluster Requirements

Storage cluster is a component of the Cisco HX Data Platform which reduces storage complexity by providing a single datastore that is easily provisioned in the vSphere Web Client. Data is fully distributed across disks in all the servers that are in the storage cluster, to leverage controller resources and provide high availability.

A storage cluster is independent of the associated vCenter cluster. You can create a storage cluster using ESXi hosts that are in the vCenter cluster.

To define the storage cluster, provide the following parameters.

Field	Description
Name	Enter a name for the storage cluster.
Management IP Address	<p>This provides the storage management network, access on each ESXi host.</p> <ul style="list-style-type: none"> • The IP address must be on the same subnet as the Management IP addresses for the nodes. • Do not allow cluster management IPs to share the last octet with another cluster on the same subnet. • These IP addresses are in addition to the four IP addresses we assign to each node in the Hypervisor section.
Storage Cluster Data IP Address	<p>This provides the storage data network and storage controller VM network, access on each ESXi host.</p> <p>The same IP address must be applied to all ESXi nodes in the cluster.</p>
Data Replication Factor	<p>Data Replication Factor defines the number of redundant replicas of your data across the storage cluster.</p> <p>This is set during HX Data Platform installation and cannot be changed.</p> <p>Choose a Data Replication Factor. The choices are:</p> <ul style="list-style-type: none"> • Data Replication Factor 3—A replication factor of three is highly recommended for all environments except HyperFlex Edge. A replication factor of two has a lower level of availability and resiliency. The risk of outage due to component or node failures should be mitigated by having active and regular backups. <p>Attention This is the recommended option.</p> <ul style="list-style-type: none"> • Data Replication Factor 2—Keep two redundant replicas of the data. This consumes less storage resources, but reduces your data protection in the event of simultaneous node or disk failure. <p>If nodes or disks in the storage cluster fail, the cluster's ability to function is affected. If more than one node fails or one node and disk(s) on a different node fail, it is called a simultaneous failure.</p>

vCenter Configuration Requirements

Provide administrator level account and password for vCenter. Ensure that you have an existing vCenter server. Ensure that the following vSphere services are operational.

- Enable Dynamic Resource Scheduler (DRS) [Optional, enable if licensed].
- Enable vMotion.
- Enable High availability (HA) [Required to define failover capacity and for expanding the datastore heartbeat].
- User VMs must be version 9 or later [Required to use HX Data Platform, Native Snapshots, and ReadyClones].

Field	Description
vCenter Server	Enter your current vCenter server web address. For example, <i>http://<IP address></i> .
User Name	Enter <i><admin username></i> .
Password	Enter <i><admin password></i> .
Datacenter Name Note An existing datacenter object can be used. If the datacenter doesn't exist in vCenter, it will be created.	Enter the required name for the vCenter datacenter.
Cluster Name	Enter the required name for the vCenter cluster. The cluster must contain a minimum of three ESXi servers.

System Services Requirements

Before installing Cisco HX Data Platform, ensure that the following network connections and services are operational.

- DNS server



Caution DNS servers should reside outside of the HX storage cluster. Nested DNS servers can cause a cluster to not start after entire cluster is shutdown, such as during DC power loss.

- NTP server



Caution NTP servers should reside outside of the HX storage cluster. Nested NTP servers can cause a cluster to not start after entire cluster is shutdown, such as during DC power loss.



- Note**
- Before configuring the storage cluster, manually verify that the NTP server is working and providing a reliable source for the time.
 - Use the same NTP server for all nodes (both converged and compute) and all storage controller VMs.
 - The NTP server must be stable, continuous (for the lifetime of the cluster), and reachable through a static IP address.
 - If you are using Active Directory as an NTP server, please make sure that the NTP server is setup according to Microsoft best practices. For more information, see [Windows Time Service Tools and Settings](#). Please note that if the NTP server is not set correctly, time sync may not work, and you may need to fix the time sync on the client-side. For more information, see [Synchronizing ESXi/ESX time with a Microsoft Domain Controller](#).

- Time Zone

Field	Essential Information
DNS Server(s)	<p data-bbox="699 1119 846 1146"><IP address></p> <p data-bbox="699 1163 1523 1224">DNS server address is required if you are using hostnames while installing the HyperFlex Data Platform.</p> <p data-bbox="699 1245 1523 1451"> Note <ul style="list-style-type: none"> • If you do not have a DNS server, do not enter a hostname under System Services in the Cluster Configuration page of the HX Data Platform Installer. Use only IP addresses. • To provide more than one <i>DNS servers address</i>, separate the address with a comma. Check carefully to ensure that DNS server addresses are entered correctly. </p>

Field	Essential Information
NTP Server(s) (A reliable NTP server is required)	<p><IP address></p> <p>NTP server is used for clock synchronization between:</p> <ul style="list-style-type: none"> • Storage controller VM • ESXi hosts • vCenter server <p>Important Static IP address for an NTP server is required to ensure clock synchronization between the storage controller VM, ESXi hosts, and vCenter server.</p> <p>During installation, this information is propagated to all the storage controller VMs and corresponding hosts. The servers are automatically synchronized on storage cluster startup.</p>
Time Zone	<p><your time zone></p> <p>Select a time zone for the storage controller VMs. It is used to determine when to take scheduled snapshots.</p> <p>Note All the VMs must be in the same time zone.</p>

CPU Resource Reservation for Controller VMs

As the storage controller VMs provide critical functionality for the HyperFlex Data Platform, the HX Data Platform Installer configures CPU resource reservations for the controller VMs. This reservation guarantees that the controller VMs have the minimum required CPU resources. This is useful in situations where the physical CPU resources of the ESXi hypervisor host are heavily consumed by the guest VMs. The following table details the CPU resource reservation for storage controller VMs.

Product ID	Number of VM CPU	Shares	Reservation	Limit
HXAF220c-M5SN (All NVMe 220)	12	Low	10,800 MHz	Unlimited
With HX Boost Mode enabled: HXAF220c-M5SN/HXAF220c-M5N	16	Low	10,800 MHz	Unlimited
With HX Boost Mode enabled: HXAF220c-M4/M5/M6 HXAF240c-M4/M5SX/M6	12	Low	10,800 MHz	Unlimited
All Other Models	8	Low	10,800 MHz	Unlimited



Note Enabling HX Boost Mode does not change the controller VM CPU reservation. It merely changes the upper limit on how much CPU the controller VM can consume.

Memory Usage for Controller VMs - NVMe Drives

The following table details the memory usage for storage controller VMs for NVMe drives.

Table 3: Memory Usage for Controller VMs - NVMe Drives

Server Model	Drive Size		
	1 / 4 TB	8 TB	15 TB (HX Release 5.0(2a) and later)
M5 C220 All NVMe	72 GB	84 GB	No Support
M6 C220 All NVMe	72 GB	84 GB	96 GB
M6 240C All NVMe	72 GB	84 GB	110 GB

- 15 TB persistent drives on All Flash configurations is not supported.
- Support for 16 TB LFF drives is not enabled.

Auto Support Requirements

Auto Support (ASUP) is the alert notification service provided through HX Data Platform. If you enable Auto Support, notifications are sent from HX Data Platform to designated email addresses or email aliases that you want to receive the notifications.

To configure Auto Support, you need the following information:

Auto Support	
Enable Auto Support check box	Check this box during HX storage cluster creation.
Mail Server	<p><IP address></p> <p>SMTP mail server must be configured in your network to enable Auto Support. Used for handling email sent from all the storage controller VM IP addresses.</p> <p>Note Only unauthenticated SMTP is supported for ASUP.</p>
Mail Sender	<p><username@domain.com></p> <p>Email address to use for sending Auto Support notifications.</p>

Auto Support	
ASUP Recipient	List of email addresses or email aliases to receive Auto Support notifications.



Note Enabling Auto Support is strongly recommended because it provides historical hardware counters that are valuable in diagnosing future hardware issues, such as drive failure for a node.

Single Sign On Requirements

The SSO URL is provided by vCenter. If it is not directly reachable from the controller VM, then configure the location explicitly using **Installer Advanced Settings**.

Single Sign On (SSO)	
SSO Server URL	SSO URL can be found in vCenter at vCenter Server > Manage > Advanced Settings , key <code>config.vpxd.sso.sts.uri</code>



CHAPTER 4

Install Cisco HyperFlex Systems Servers

This chapter describes how to install the physical components for setting up a HyperFlex cluster:

- [Rack Cisco HyperFlex Nodes, on page 35](#)
- [Setting Up the Fabric Interconnects, on page 36](#)
- [Connecting HX-Series Servers to Cisco UCS Fabric Interconnects, on page 43](#)
- [Overview, on page 43](#)

Rack Cisco HyperFlex Nodes

For details on the HyperFlex cluster and node limits, see **Cisco HX Data Platform Storage Cluster Specifications** in the latest release of the [Release Notes for Cisco HX Data Platform](#).

For UCS C-Series integration guidelines, see the [Cisco UCS C-Series Server Integration with Cisco UCS Manager Configuration Guide](#) for your release.

For details on the installation of Cisco HyperFlex nodes, refer to respective links from the following table:

Type of Node To Be Installed	Reference
Converged Nodes	
HyperFlex HX245c M5/M6 Nodes	Cisco HyperFlex HX245c Node Installation Guides
HyperFlex HX240c M5/M6 Nodes	Cisco HyperFlex HX240c Node Installation Guides
HyperFlex HX225c M5/M6 Nodes	Cisco HyperFlex HX225c Node Installation Guides
HyperFlex HX220c M5/M6 Nodes	Cisco HyperFlex HX220c Node Installation Guides
Compute-only Nodes	
Cisco UCS B200 M5 Nodes	Cisco UCS B200 M3/M4/M5 Blade Server Installation and Service Note
Cisco UCS B480 M5 Nodes	Cisco UCS B480 M5 Blade Server Installation and Service Note
Cisco UCS C240 M5/M6 Rack Nodes	Cisco UCS C240 Server Installation and Service Guide

Type of Node To Be Installed	Reference
Cisco UCS C220 M5/M6 Rack Nodes	Cisco UCS C220 Server Installation and Service Guide
Cisco UCS C480 M5 Nodes	Cisco UCS C480 M5 Server Installation and Service Guide

Setting Up the Fabric Interconnects

Configure a redundant pair of fabric interconnects for high availability as follows:

1. Connect the two fabric interconnects directly using Ethernet cables between the L1 and L2 high availability ports.
2. Connect Port L1 on fabric interconnect A to port L1 on fabric interconnect B, and Port L2 on fabric interconnect A to port L2 on fabric interconnect B.

This allows both the fabric interconnects to continuously monitor the status of each other.

Verify and obtain the following information before connecting the fabric interconnects.

Item	Description
Verify the physical connections of the fabric interconnects.	<ul style="list-style-type: none"> • Console port for the first fabric interconnect must be physically connected to a computer or console server. • Management Ethernet port (mgmt0) must be connected to an external hub, switch, or router. • L1 ports on both the fabric interconnects must be directly connected to each other. • L2 ports on both the fabric interconnects must be directly connected to each other.
Verify console port parameters on the computer terminal.	<ul style="list-style-type: none"> • 9600 baud • 8 data bits • No parity • 1 stop bit

Item	Description
Obtain information for initial setup.	Collect the following information for initial setup: <ul style="list-style-type: none"> • System name • Password for admin account • Three static IP addresses • Subnet mask for three static IP addresses • Default gateway IP address • DNS server IP address • Domain name for the system

Both fabric interconnects must go through the same setup process. Set up the primary fabric interconnect and enable for cluster configuration. When you use the same process to set up the secondary fabric interconnect, it detects the first fabric interconnect as a peer.

Configuring the Primary Fabric Interconnect Using Cisco UCS Manager GUI

Specify the following three IP addresses in the same subnet before you begin the configuration.

- Management Port IP address for the primary fabric interconnect, FI A.
- Management Port IP address for the secondary fabric interconnect, FI B.
- IP address of the HyperFlex Cluster.

Configure the Primary Fabric Interconnect using the Cisco UCS Manager GUI as follows:

-
- Step 1** Connect to the console port. See [Cisco 6300 Series Fabric Interconnect Hardware Installation Guide](#) for more details.
- Step 2** Power on the fabric interconnect. You will see the *Power On* self-test message as the fabric interconnect boots.
- Step 3** At the installation method prompt, enter *gui*.
- Step 4** If the system cannot access the DHCP server, you will be prompted to enter the following information:
- IPv4 address for the management port on the fabric interconnect.
 - IPv4 subnet mask for the management port on the fabric interconnect.
 - IPv4 for the default gateway assigned to the fabric interconnect.
- Important** All IP addresses must be IPv4. HyperFlex does not support IPv6 addresses.
- Step 5** Copy the web link from the prompt into a web browser and navigate to the Cisco UCS Manager launch page.
- Step 6** Select **Express Setup**.
- Step 7** Select **Initial Setup** and click **Submit**.
- Step 8** In the **Cluster and Fabric Setup** area, complete the following fields:

Name	Description
Enable Cluster option	Select the enable cluster option.
Fabric Setup option	Select <i>Fabric A</i> .
Cluster IP Address field	Enter the IPv4 address that Cisco UCS Manager will use.

Step 9 In the **System Setup** area, complete the following fields:

Field	Description
System Name field	The name assigned to the Cisco UCS domain.
Admin Password field	The password used for the admin account on the fabric interconnect. Choose a strong password that meets the guidelines for Cisco UCS Manager passwords. This password cannot be blank.
Confirm Admin Password field	The password used for the admin account on the fabric interconnect.
Mgmt IP Address field	The static IP address for the management port on the fabric interconnect.
Mgmt IP Netmask field	The IP subnet mask for the management port on the fabric interconnect.
Default Gateway field	The IP address for the default gateway assigned to the management port on the fabric interconnect.
DNS Server IP field	The IP address for the DNS server assigned to the management port on the fabric interconnect.
Domain name field	The name of the domain in which the fabric interconnect resides.

Step 10 Click **Submit**.
A page displays the results of your setup operations.

Configuring the Secondary Fabric Interconnect Using Cisco UCS Manager GUI

Make sure that the console port of the secondary fabric interconnect is physically connected to a computer or a console server. Ensure that you know the password for the admin account on the primary fabric interconnect that you configured earlier.

Step 1 Connect to the console port. See [Cisco 6300 Series Fabric Interconnect Hardware Installation Guide](#) for more details.

Step 2 Power on the fabric interconnect. You will see the *Power On* self-test message as the fabric interconnect boots.

Step 3 At the installation method prompt, enter *gui*.

Step 4 If the system cannot access the DHCP server, you will be prompted to enter the following information:

- IPv4 address for the management port on the fabric interconnect.
- IPv4 subnet mask for the management port on the fabric interconnect.
- IPv4 address for the default gateway assigned to the fabric interconnect.

Note Both the fabric interconnects must be assigned the same management interface address type during setup.

Step 5 Copy the web link from the prompt into a web browser and go to the Cisco UCS Manager GUI launch page.

Step 6 Copy the web link from the prompt into a web browser and navigate to the Cisco UCS Manager launch page.

Step 7 Select **Express Setup**.

Step 8 Select **Initial Setup** and click **Submit**.

The fabric interconnect should detect the configuration information for the first fabric interconnect.

Step 9 In the **Cluster and Fabric Setup** area, complete the following fields:

Name	Description
Enable Cluster option	Select the enable cluster option.
Fabric Setup option	Select <i>Fabric B</i> .

Step 10 In the **System Setup** area, enter the password for the Admin account into the **Admin Password of Master** field. The **Manager Initial Setup** Area is displayed.

Step 11 In the **Manager Initial Setup** area, the field that is displayed depends on whether you configured the first fabric interconnect with an IPv4 management address. Complete the field that is appropriate for your configuration as follows:

Field	Description
Peer FI is IPv4 Cluster enabled. Please provide local FI Mgmt0 IPv4 address field	Enter an IPv4 address for the Mgmt0 interface on the local fabric interconnect.

Step 12 Click **Submit**.

A page displays the results of your setup operations.

Configure the Primary Fabric Interconnect Using CLI

Step 1 Connect to the console port.

Step 2 Power on the fabric interconnect.

You will see the power-on self-test messages as the fabric interconnect boots.

Step 3 When the unconfigured system boots, it prompts you for the setup method to be used. Enter **console** to continue the initial setup using the console CLI.

Step 4 Enter **setup** to continue as an initial system setup.

Step 5 Enter **y** to confirm that you want to continue the initial setup.

- Step 6** Enter the password for the admin account.
- Step 7** To confirm, re-enter the password for the admin account.
- Step 8** Enter **yes** to continue the initial setup for a cluster configuration.
- Step 9** Enter the fabric interconnect fabric (either **A** or **B**).
- Step 10** Enter the system name.
- Step 11** Enter the IPv4 address for the management port of the fabric interconnect.
You will be prompted to enter an IPv4 subnet mask.
- Step 12** Enter the IPv4 subnet mask, then press **Enter**.
You are prompted for an IPv4 address for the default gateway, depending on the address type you entered for the management port of the fabric interconnect.
- Step 13** Enter the IPv4 address of the default gateway.
- Step 14** Enter **yes** if you want to specify the IP address for the DNS server, or **no** if you do not.
- Step 15** (Optional) Enter the IPv4 address for the DNS server.
The address type must be the same as the address type of the management port of the fabric interconnect.
- Step 16** Enter **yes** if you want to specify the default domain name, or **no** if you do not.
- Step 17** (Optional) Enter the default domain name.
- Step 18** Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the Setup wizard again to change some of the settings.
If you choose to go through the Setup wizard again, it provides the values you previously entered, and the values appear in brackets. To accept previously entered values, press **Enter**.

Example

The following example sets up the first fabric interconnect for a cluster configuration using the console and IPv4 management addresses:

```

Enter the installation method (console/gui)? console
Enter the setup mode (restore from backup or initial setup) [restore/setup]? setup
You have chosen to setup a new switch. Continue? (y/n): y
Enter the password for "admin": adminpassword%958
Confirm the password for "admin": adminpassword%958
Do you want to create a new cluster on this switch (select 'no' for standalone setup
or if you want this switch to be added to an existing cluster)? (yes/no) [n]: yes
Enter the switch fabric (A/B): A
Enter the system name: foo
Mgmt0 IPv4 address: 192.168.10.10
Mgmt0 IPv4 netmask: 255.255.255.0
IPv4 address of the default gateway: 192.168.10.1
Virtual IPv4 address: 192.168.10.12
Configure the DNS Server IPv4 address? (yes/no) [n]: yes
  DNS IPv4 address: 20.10.20.10
Configure the default domain name? (yes/no) [n]: yes
  Default domain name: domainname.com
Join centralized management environment (UCS Central)? (yes/no) [n]: no
Following configurations will be applied:
  Switch Fabric=A

```

```

System Name=foo
Management IP Address=192.168.10.10
Management IP Netmask=255.255.255.0
Default Gateway=192.168.10.1
Cluster Enabled=yes
Virtual Ip Address=192.168.10.12
DNS Server=20.10.20.10
Domain Name=domainname.com
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

Configure the Subordinate Fabric Interconnect Using CLI

This procedure describes setting up the second fabric interconnect using the IPv4 address for the management port.



Note When adding a new Fabric Interconnect to an existing High Availability cluster, for example, during a new install or when replacing a Fabric Interconnect, the new device will not be able to log into the cluster as long as the authentication method is set to remote. To successfully add a new Fabric Interconnect to the cluster, the authentication method must be temporarily set to local and the local admin credentials of the primary Fabric Interconnect must be used.

- Step 1** Connect to the console port.
- Step 2** Power up the fabric interconnect.
You will see the power-on self-test messages as the fabric interconnect boots.
- Step 3** When the unconfigured system boots, it prompts you for the setup method to be used. Enter **console** to continue the initial setup using the console CLI.

Note The fabric interconnect should detect the peer fabric interconnect in the cluster. If it does not, check the physical connections between the L1 and L2 ports, and verify that the peer fabric interconnect has been enabled for a cluster configuration.

- Step 4** Enter **y** to add the subordinate fabric interconnect to the cluster.
- Step 5** Enter the admin password of the peer fabric interconnect.
- Step 6** Enter the IP address for the management port on the subordinate fabric interconnect.
- Step 7** Review the setup summary and enter **yes** to save and apply the settings, or enter **no** to go through the Setup wizard again to change some of the settings.

If you choose to go through the Setup wizard again, it provides the values you previously entered, and the values appear in brackets. To accept previously entered values, press **Enter**.

Example

The following example sets up the second fabric interconnect for a cluster configuration using the console and the IPv4 address of the peer:

```

Enter the installation method (console/gui)? console
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect

```

```

will be added to the cluster. Continue (y/n) ? y
Enter the admin password of the peer Fabric Interconnect: adminpassword%958
Peer Fabric interconnect Mgmt0 IPv4 Address: 192.168.10.11
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

```

Verify Console Setup

You can verify that both fabric interconnect configurations are complete by logging in to the fabric interconnect through SSH.

Use the following commands to verify the cluster status using Cisco UCS Manager CLI.

Command	Purpose	Sample Output
show cluster state	Displays the operational state and leadership role for both fabric interconnects in a high availability cluster.	The following example displays that both fabric interconnects are in the Up state, HA is in the Ready state, fabric interconnect A has the primary role, and fabric interconnect B has the subordinate role. <pre> UCS-A# show cluster state Cluster Id: 0x4432f72a371511de-0xb97c000de1blada4 A: UP, PRIMARY B: UP, SUBORDINATE HA READY </pre>

Command	Purpose	Sample Output
<code>show cluster extended-state</code>	Displays extended details about the cluster state and typically used when troubleshooting issues.	<p>The following example shows how to view the extended state of a cluster.</p> <pre>UCSC# show cluster extended-state 0x2e95deacbd0f11e2- 0x8ff35147e84f3de2Start time: Thu May 16 06:54:22 2013Last election time: Thu May 16 16:29:28 2015System Management Viewing the Cluster State A: UP, PRIMARY B: UP, SUBORDINATE A: memb state UP, lead state PRIMARY, mgmt services state: UP B: memb state UP, lead state SUBORDINATE, mgmt services state: UP heartbeat state PRIMARY_OK HA READY Detailed state of the device selected for HA quorum data: Device 1007, serial: a66b4c20-8692-11df-bd63-1b72ef3ac801, state: active Device 1010, serial: 00e3e6d0-8693-11df-9e10-0f4428357744, state: active Device 1012, serial: 1d8922c8-8693-11df-9133-89fa154e3fa1, state: active</pre>

Connecting HX-Series Servers to Cisco UCS Fabric Interconnects

Overview

The Cisco HX220c and HX240c Servers connect directly to the fabric interconnects. The direct connection enables Cisco UCS Manager to manage the HX-Series servers using a single cable for both management traffic and data traffic.



Note After connecting the server with the fabric interconnect, when the server is discovered, update the C-Series software bundle available for Cisco UCS Manager using the UCS Manager configuration form.

When you use direct connect mode, all Cisco UCS managed adapters must be connected to the server ports on the fabric interconnects. Make sure that the HX servers have the recommended firmware as listed in the *Requirements Chapter*. If not, use Cisco UCS Manager to update the firmware.



Note Note the following UCS configuration limits:

- General Cisco UCS configuration limits: see the [Cisco UCS 6200, 6332, 6324 and 6400 Configuration Limits for Cisco UCS Manager](#).

Connecting Converged Nodes to the Fabric Interconnect

This topic describes how to physically add converged nodes for creating a HX cluster or adding to an existing HX cluster.

Before you begin



Important

- Set the CIMC server to factory default settings before integrating with Cisco UCS Manager.
- Do not connect dedicated CIMC ports to the network for integrated nodes. Doing so causes the server to not be discovered in Cisco UCS Manager. If the server is not discovered, reset CIMC to factory settings for each server.
- If there is no foreseeable future need to connect FC Storage, only use ports 1-16.
- Cisco UCS FI 6200/6300/6400 and 6400 only support configuring ports 1-6 as FC ports. If there is a future need to connect FC Storage, convert ports 1-6 to FC.



Note The conversion may disrupt the HX deployment.

- Before you connect the CIMC server, make sure a Cisco VIC 1227 is installed in the PCIe slot 2 of an HXc240, or Riser 1 slot 1 for an HXc220 for integration with Cisco UCS Manager. If the card is not installed in the correct slot, you cannot enable direct connect management for the server.
- Complete the physical cabling of servers to the fabric interconnects, and configure the ports as server ports.

Step 1 Install the HX server in the rack. See [Rack Cisco HyperFlex Nodes, on page 35](#) for more details.

Step 2 Configure the server ports on the fabric interconnect.

- Connect a 10-Gb SFP+ cable from one port on the server to fabric interconnect A. You can use any port on fabric interconnect A, but the port must be enabled for server traffic.

Connect one cable from the VIC to the fabric interconnect for one card. Do not connect both ports to the same fabric interconnect.

- b) Configure that port on FI-A as a server port. For the detailed steps, refer to the *Configuring Port Modes for a 6248 Fabric Interconnect* section of the [Cisco UCS Manager Network Management Guide](#).
- c) Connect 10-Gb SFP+ cable from the other port on the server to FI B. You can use any port on FI B, but the port must be enabled for server traffic.

Note Do not mix SFP+ types on an uplink. If you do, you will get `Discovery Failed` errors.

- d) Configure that port on FI-B as a server port. For the detailed steps, refer to the *Configuring Port Modes for a 6248 Fabric Interconnect* section of the [Cisco UCS Manager Network Management Guide](#).

Step 3

Attach a power cord to each power supply in your node, and to a grounded AC power outlet. During initial boot up, wait for approximately two minutes to let the node boot in standby power.

- Note**
- a. When powered up, the server is discovered by the fabric interconnects. You can monitor node discovery in UCS Manager.
 - b. Verify the node's power status by looking at the node **Power Status LED** on the front panel. A node is in the standby power mode when the LED is amber.


Step 4

Repeat steps one through four to connect the remaining HX-Series servers to the HyperFlex cluster.

Physical Connectivity Illustrations for Direct Connect Mode Cluster Setup


The following images shows a sample of direct connect mode physical connectivity for C-Series Rack-Mount Server with Cisco UCS Domain, Cisco UCS Manager, release 3.1 and later. This image shows the cabling configuration for Cisco UCS Manager integration with a C-Series Rack-Mount Server. The paths shown in gold carry both management traffic and data traffic.

Figure 10: Direct Connect Cabling Configuration



305210

Figure 11: Direct Connect Cabling Configuration with Cisco VIC 1455



1	Cisco UCS 6454 Fabric Interconnect or Cisco UCS 6200, or 6300 Series FI (Fabric A)	3	C-Series Rack-Mount Server
2	Cisco UCS 6454 Fabric Interconnect or Cisco UCS 6200, or 6300 Series FI (Fabric B)	4	Cisco UCS VIC in supported PCIe slot

XGb represents a 40 Gigabit Ethernet connection or a 10 Gigabit Ethernet connection. For the 10 Gigabit Ethernet, the following cables are used:

- 4x10 Breakout Small Form-Factor Pluggable (SFP) cables
- 4x10 Active Optical (OAC) cables
- 10G Small Form-Factor Pluggable (SFP) cable that uses the Qualified Security Assessor (QSA) module



CHAPTER 5

Configure Cisco HyperFlex Systems

This chapter describes how to configure the components of the Cisco HyperFlex Systems:

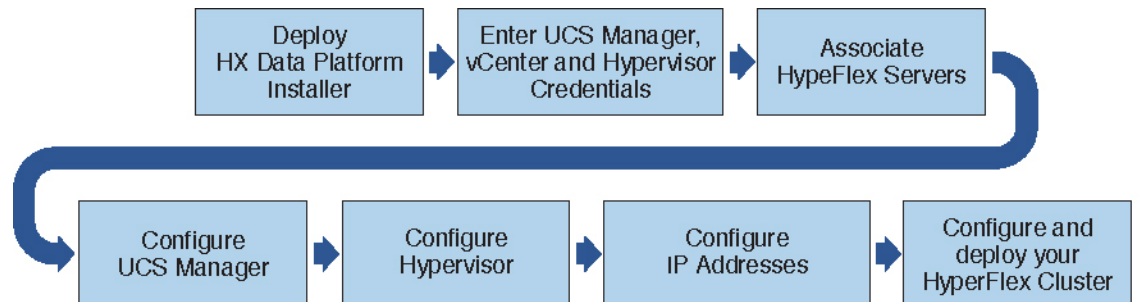
- [Installation Workflow](#), on page 47
- [Deploy HX Data Platform Installer OVA Using vSphere Web Client](#), on page 48
- [Deploy the HX Data Platform Installer OVA with a Static IP Address](#), on page 50
- [Configure Syslog](#), on page 51
- [Configure and Deploy Your HyperFlex Cluster](#), on page 52
- [Installation of HyperFlex Nodes with GPUs](#), on page 62
- [HX Data Platform Installer Navigation Aid Buttons](#), on page 63
- [Warnings and Error Messages](#), on page 63

Installation Workflow



Note If the HyperFlex cluster nodes were part of any other HyperFlex cluster before (or not factory shipped), follow the node cleanup procedure before starting the cluster deployment. For more information, see [HyperFlex Customer Cleanup Guides for FI and Edge](#).

The following installation workflow summarizes the steps involved in creating a Standard Cluster, using the HX Data Platform Installer.



Follow this workflow during installation:

1. Deploy the HX Data Platform Installer OVA using the vSphere Web Client. If your hypervisor wizard defaults to DHCP for assigning IP addresses to new VMs, deploy the HX Data Platform Installer OVA

with a static IP address. See [Deploy HX Data Platform Installer OVA Using vSphere Web Client, on page 48](#) or [Deploy the HX Data Platform Installer OVA with a Static IP Address, on page 50](#) for more information.

2. Configure syslog to send all logging information to a centralized syslog repository. See [Configure Syslog, on page 51](#) for more information.
3. Enter UCS Manager, vCenter, and Hypervisor credentials.
4. Configure server ports and associate HyperFlex servers. See [Associate HyperFlex Servers, on page 52](#) for more information.
5. Configure VLAN, MAC Pool, 'hx-ext-mgmt' IPPool for Out-of-Band CIMC, iSCSi Storage, and FC Storage. See [Configure UCS Manager, on page 53](#) for more information.
6. Configure the Hypervisor. See [Configure Hypervisor, on page 57](#) for more information.
7. Configure the IP addresses. See [Configure IP Addresses, on page 58](#) for more information.
8. Configure and deploy the HyperFlex cluster. See [Configure Your HyperFlex Cluster, on page 59](#) for more information.

Deploy HX Data Platform Installer OVA Using vSphere Web Client

In addition to installing the HX Data Platform on an ESXi host, you may also deploy the HX Data Platform Installer on either VMware Workstation, VMware Fusion, or Virtual Box.



Note

- Connect to vCenter to deploy the OVA file and provide the IP address properties. Deploying directly from an ESXi host will not allow you to set the values correctly.
- Do not deploy the HX Data Platform Installer to an ESXi server that is going to be a node in the Cisco HX Storage Cluster.

Step 1

Locate and download the HX Data Platform Installer OVA from [Download Software](#). Download the HX Data Platform Installer to a node that is on the storage management network, which will be used for the HX Data Platform storage cluster.

Example:

```
Cisco-HX-Data-Platform-Installer-v4.5.1a-26363.ova
```

Step 2

Deploy the HX Data Platform Installer using VMware hypervisor, to create a HX Data Platform Installer virtual machine.

Note Use a release of the virtualization platform that supports virtual hardware version 10.0 or greater.

vSphere is a system requirement. You can use either vSphere thick client, vSphere thin client, or vSphere Web Client. To deploy the HX Data Platform Installer, you can also use VMware Workstation, VMware Fusion, or VirtualBox.

- a) Open a virtual machine hypervisor, such as vSphere, VirtualBox, Workstation, or Fusion.

b) Select the node where you want to deploy the HX Data Platform Installer.

Important Ensure that you provide user credentials while deploying the HX Installer OVA using vSphere Web Client.

- Using vSphere thick Client—Expand **Inventory list > Host > File > Deploy OVA**.
- Using vSphere Web Client—Expand **vCenter Inventory list > Hosts > Host > Deploy OVA**.

Step 3

Select where the HX Data Platform Installer is located. Accept the defaults, and select the appropriate network.

Step 4

Enter a static IP address for use by the HX Data Platform Installer VM.

Note

- Static IP Address is necessary even if DHCP is configured for the network. You need the static IP address to run the HX Data Platform Installer, to install the HX Data Platform, and to create the HX Data Platform storage cluster.
- If your hypervisor wizard defaults to DHCP for assigning IP addresses to new VMs, then complete the steps in [Deploy the HX Data Platform Installer OVA with a Static IP Address, on page 50](#), to install the HX Data Platform Installer VM with a static IP address. DNS must be reachable from the Installer VM.

Field	Description
Hostname	The hostname for this VM. Leave blank to try to reverse lookup the IP address.
Default Gateway	The default gateway address for this VM. Leave blank if DHCP is desired.
DNS	The domain name servers for this VM (comma separated). Leave blank if DHCP is desired.
IP Address	The IP address for this interface. Leave blank if DHCP is desired.
Netmask	The netmask or prefix for this interface. Leave blank if DHCP is desired.
Root Password	The root user password. This field is a required field.

Step 5

Click **Next**. Verify if the options listed are correct and select **Power on after deployment**.

To power on the HX Data Platform Installer manually, navigate to the virtual machine list and power on the installer VM.

Note

The preferred settings for the HX Data Platform Installer virtual machine is 3 vCPU and 4 GB of memory. Reducing these settings can result in 100% CPU usage and spikes for the host.

Step 6

Click **Finish**. Wait for the HX Data Platform Installer VM to be added to the vSphere infrastructure.

Step 7 Open the HX Data Platform Installer virtual machine console.

The initial console display lists the HX Data Platform Installer virtual machine IP address.

```
Data Platform Installer.
*****
You can start the installation by visiting
the following URL:
http://192.168.10.210
*****
Cisco-HX-Data-Platform-Installer login:
```

Step 8 Use the URL to log in to the HX Data Platform Installer.

```
Example:
http://192.168.10.210
```

Step 9 Accept the self-signed certificate.

Step 10 Log in using the username **root** and the password you provided as part of the OVA deployment.

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

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

Step 1 Install the VMware OVF Tool 4.1 or later on a node that is on the storage management network that will be used for the HX Data Platform storage cluster. See [OVF Tool Documentation](#) for more details.

Step 2 Locate and download HX Data Platform installer OVA from [Download Software](#) on the node where VMware OVF was installed.

Step 3 Deploy the downloaded HX Data Platform installer OVA, using the `ovftool` command. For example:

```
root@server:/tmp/test_ova# ovftool --noSSLVerify --diskMode=thin
--acceptAllEulas=true --powerOn --skipManifestCheck --X:injectOvfEnv
--datastore=qa-048-ssdl --name=rfsi_static_test1 --network='VM Network'
--prop:hx.3gateway.Cisco_HX_Installer_Appliance=10.64.8.1
--prop:hx.4DNS.Cisco_HX_Installer_Appliance=10.64.1.8
--prop:hx.5domain.Cisco_HX_Installer_Appliance=cisco
--prop:hx.6NTP.Cisco_HX_Installer_Appliance=10.64.8.5
--prop:hx.1ip0.Cisco_HX_Installer_Appliance=10.64.8.36
--prop:hx.2netmask0.Cisco_HX_Installer_Appliance=255.255.248.0
--prop:hx.7root_password.Cisco_HX_Installer_Appliance=mypassword
/opt/ovf/rfsi_test/Cisco-HX-Data-Platform-Installer-v1.7.1-14786.ova
vi://root:password@esx_server
```

The command deploys the HX Data Platform installer, powers on the HX Data Platform installer VM, and configures the provided static IP address. A sample of processing response:

```
Opening OVA source:
/opt/ovf/rfsi_test/Cisco-HX-Data-Platform-Installer-v1.7.1-14786.ova
Opening VI target: vi://root@esx_server:443/
Deploying to VI: vi://root@esx_server:443/
Transfer Completed
Powering on VM: rfsi_static_test
```

Task Completed
Completed successfully

DNS must be reachable from the Installer VM. The required command options for the static IP address to be configured successfully are:

Command	Description
powerOn	To power on the HX Data Platform installer VM after it is deployed.
X:injectOvfEnv	To insert the static IP properties onto the HX Data Platform installer VM.
prop:hx.3gateway.Cisco_HX_Installer_Appliance=10.64.8.1	Specify the appropriate gateway IP address.
prop:hx.4DNS.Cisco_HX_Installer_Appliance=10.64.1.8	Specify the appropriate DNS IP address.
prop:hx.5domain.Cisco_HX_Installer_Appliance=cisco	Specify the appropriate domain.
prop:hx.6NTP.Cisco_HX_Installer_Appliance=10.64.8.5	Specify the appropriate NTP IP address.
prop:hx.1ip0.Cisco_HX_Installer_Appliance=10.64.8.36	Specify the appropriate installer static IP address.
prop:hx.2netmask0.Cisco_HX_Installer_Appliance=255.255.248.0	Specify the appropriate netmask address.
prop:hx.7root_password.Cisco_HX_Installer_Appliance=mypassword	Specify the root user password.
/opt/ovf/rfsi_test/Cisco-HX-Data-Platform-Installer-v1.7.1-14786.ova	The source address of the HX Data Platform installer OVA.
vi://root:password@esx_server	The destination ESX server where the HX Data Platform installer VM is installed. Include the appropriate ESX server root login credentials.

Configure Syslog

It is best practice to send all logging information to a centralized syslog repository.



Attention

In general, configuring audit log export using syslog is recommended if long term retention of audit log is required. Specifically, for HX220c nodes and compute-only nodes booting from SD card, configuring syslog is required for persistent logging. If you do not configure the syslog server, audit logs are overwritten because of the log rotation policy.



Note You cannot select an NFS datastore as a destination for the persistent scratch location on ESXi. If you select the HX datastore for the persistent scratch location, it will be removed after the ESXi host reloads.

For all M5 and M6 servers, M.2 boot SSD is automatically selected for use as scratch. This is configured out of the box on any new install.

For HX240M4 (non-SED), Intel SSD is used for persistent logs/scratch (same applies on 220M5/240M5, but on a different local SSD).

For HX220M4 and HX240M4 (SED), there is no location to store the scratch partition. So, the only option is to use syslog for persistent logging over the network.

Step 1 Verify that the syslog server is up and running and TCP/UDP ports are open to receive logs from ESXi servers.

Step 2 SSH to the ESXi shell and execute the following commands.

```
a) esxcli system syslog config set --loghost='udp://remote-syslog-server-ip'
b) esxcli system syslog reload
c) esxcli network firewall ruleset set -r syslog -e true
d) esxcli network firewall refresh
```

Step 3 Repeat steps 1 and 2 for all ESXi hosts in the cluster.

Step 4 At the remote syslog server, verify if the logs are being received in the designated directory.

Configure and Deploy Your HyperFlex Cluster

Associate HyperFlex Servers

On the **Server Selection** page, the **Configuration** pane on the right displays a detailed list of the **Credentials** used. The **Server Selection** page displays a list of unassociated HX servers under the **Unassociated** tab, and the list of discovered servers under the **Associated** tab.

Field	Description
Locator LED	Turn on to locate a server.
Server Name	Name assigned to the server.
Status	<ul style="list-style-type: none"> Inaccessible—
Model	Displays the server model.
Serial	Displays the serial number of the server.
Assoc State	<ul style="list-style-type: none"> Associated Unassociated

Field	Description
Service Profile [Only for Associated Servers]	Service profile assigned to the server. Note Editing the HyperFlex Service Profile templates is not recommended.
Actions	<ul style="list-style-type: none"> • Launch KVM Console—Choose this option to launch the KVM Console directly from the HX Data Platform Installer. • Disassociate Server—Choose this option to remove a service profile from that server.

Before you begin

Ensure that you completed entering UCS Manager, vCenter, and Hypervisor credentials.

Step 1 Click the **Configure Server Ports** button to discover any new HX nodes. In the **Configure Server Ports** dialog box, list all ports to be configured as server ports. Click **Configure**.

Note Typically, the server ports are configured in Cisco UCS Manager before you start the configuration.

Step 2 Select the servers under the **Unassociated** tab to include in the HyperFlex cluster.

If HX servers do not appear in this list, check Cisco UCS Manager and ensure that they have been discovered.

Note If there are no unassociated servers, the following error message is displayed:

No unassociated servers found. Login to UCS Manager and ensure server ports are enabled.

Step 3 Click **Continue** to configure UCS Manager. See [Configure UCS Manager, on page 53](#).

Configure UCS Manager

On the **UCSM Configuration** page, you can configure VLAN, MAC Pool, 'hx-ext-mgmt' IP Pool for CIMC, iSCSi Storage, and FC Storage.

Before you begin

Associate servers to the HyperFlex cluster. See [Associate HyperFlex Servers, on page 52](#).

Step 1 In the **VLAN Configuration** section, complete the following fields:

Note Use separate subnet and VLANs for each of the following networks.

Field	Description
VLAN for Hypervisor and HyperFlex management	

Field	Description
VLAN Name	hx-inband-mgmt
VLAN ID	Default—3091
VLAN for HyperFlex storage traffic	
VLAN Name	hx-storage-data
VLAN ID	No default value.
VLAN for VM vMotion	
VLAN Name	hx-vmotion
VLAN ID	Default—3093
VLAN for VM Network	
VLAN Name	vm-network
VLAN ID(s)	Default—3094 A comma-separated list of guest VLANs.

Step 2 In the **MAC Pool** section, configure **MAC Pool Prefix** by adding in two more hex characters (0-F).

Note Select a prefix that is not used with any other MAC address pool across all UCS domains.

Example:
00:25:B5:A0:

Step 3 In the '**hx-ext-mgmt**' **IP Pool for CIMC** section, complete the following fields:

Field	Description
IP Blocks	The range of management IP addresses assigned to the CIMC for each HyperFlex server. The IP addresses are specified as a range, and multiple blocks of IPs may be specified as a comma-separated list. Ensure you have at least one unique IP per server in the cluster. When selecting to use out-of-band, this range must fall into the same IP subnet used on the mgmt0 interfaces on the Fabric Interconnects. For example, <i>10.193.211.124-127, 10.193.211.158-163.</i>
Subnet Mask	Specify the subnet mask for the IP range provided above. For example, <i>255.255.0.0.</i>
Gateway	Enter the Gateway IP address. For example, <i>10.193.0.1.</i>

The management IP addresses used to access the CIMC on a server can be either:

- **Out of band:** The CIMC management traffic traverses the Fabric Interconnect through the limited bandwidth management interface, mgmt0, on the Fabric Interconnect. This option is most commonly used and shares the same VLAN as the Fabric Interconnect management VLAN.
- **In-band:** The CIMC management traffic traverses the Fabric Interconnect through the uplink ports of the Fabric Interconnect. The bandwidth available for management traffic in this case would be equivalent to the Fabric Interconnect uplink bandwidth. If you are using the In-band option, the Cisco HyperFlex installer will create a dedicated VLAN for the CIMC management communication. This option is useful when large files such as a Windows Server installation ISO must be mounted to the CIMC for OS installation. This option is only available in the HyperFlex installer VM and is not available for deployments through Intersight.

Step 4 Select either **Out of band** or **In-band** based on the type of connection you want to use for CIMC management access. If you select In-band, provide the VLAN ID for the management VLAN. Make sure to create the CIMC management VLAN in the upstream switch for seamless connectivity.

Step 5 If you want to add external storage, configure **iSCSI Storage** by completing the following fields:

Field	Description
Enable iSCSI Storage check box	Select to configure iSCSI storage.
VLAN A Name	Name of the VLAN associated with the iSCSI vNIC, on the primary Fabric Interconnect (FI-A).
VLAN A ID	ID of the VLAN associated with the iSCSI vNIC, on the primary Fabric Interconnect (FI-A).
VLAN B Name	Name of the VLAN associated with the iSCSI vNIC, on the subordinate Fabric Interconnect (FI-B).
VLAN B ID	ID of the VLAN associated with the iSCSI vNIC, on the subordinate Fabric Interconnect (FI-A).

Step 6 If you want to add external storage, configure **FC Storage** by completing the following fields:

Field	Description
Enable FC Storage check box	Select to enable FC Storage.
WWxN Pool	A WWN pool that contains both WW node names and WW port names. For each Fabric Interconnect, a WWxN pool is created for WWPN and WWNN.
VSAN A Name	The name of the VSAN for the primary Fabric Interconnect (FI-A). Default— <code>hx-ext-storage-fc-a</code> .

Field	Description
VSAN A ID	The unique identifier assigned to the network for the primary Fabric Interconnect (FI-A). Caution Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.
VSAN B Name	The name of the VSAN for the subordinate Fabric Interconnect (FI-B). Default— <code>hx-ext-storage-fc-b</code> .
VSAN B ID	The unique identifier assigned to the network for the subordinate Fabric Interconnect (FI-B). Caution Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.

Step 7 In the **Advanced** section, do the following:

Field	Description
UCS Server Firmware Release	Select the UCS firmware release to associate with the HX servers from the drop-down list. The UCS firmware release must match the UCSM release. See the latest Cisco HX Data Platform Release Notes for more details. For example, <code>3.2(1d)</code> .
HyperFlex Cluster Name	Specify a user-defined name. The HyperFlex cluster name is applied to a group of HX Servers in a given cluster. The HyperFlex cluster name adds a label to service profiles for easier identification.
Org Name	Specify a unique <i>Org Name</i> to ensure isolation of the HyperFlex environment from the rest of the UCS domain.

Step 8 Click **Continue** to configure the Hypervisor. See [Configure Hypervisor, on page 57](#).

Configure Hypervisor



Note Review the VLAN, MAC pool, and IP address pool information on the **Hypervisor Configuration** page, in the **Configuration** pane. These VLAN IDs may be changed by your environment. By default, the HX Data Platform Installer sets the VLANs as non-native. You must configure the upstream switches to accommodate the non-native VLANs by appropriately applying a trunk configuration.



Attention You can skip configuring Hypervisor in case of a reinstall, if ESXi networking has been completed.

Before you begin

Configure VLAN, MAC Pool, and 'hx-ext-mgmt' IP Pool for Out-of-Band CIMC. If you are adding external storage, configure iSCSI Storage and FC Storage. Select the UCS Server Firmware Version and assign a name for the HyperFlex cluster. See [Configure UCS Manager, on page 53](#).

Step 1

In the **Configure Common Hypervisor Settings** section, complete the following fields:

Field	Description
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, <i>255.255.0.0</i> .
Gateway	IP address of gateway. For example, <i>10.193.0.1</i> .
DNS Server(s)	IP address for the DNS Server. Note <ul style="list-style-type: none"> • If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform Installer. Use only static IP addresses and hostnames for all ESXi hosts. • If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.

Step 2

On the **Hypervisor Settings** section, select **Make IP Addresses and Hostnames Sequential** to make the IP addresses sequential. Complete the following fields:

Note You can rearrange the servers using drag and drop.

Field	Description
Name	Name assigned to the server.
Locator LED	Turn on to locate a server.
Serial	Displays the serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.
Hostname	Do not leave the hostname fields empty.

Step 3 Click **Continue** to configure IP Addresses. See [Configure IP Addresses, on page 58](#).

Configure IP Addresses

Before you begin

Ensure that you completed configuring Hypervisor on the **Hypervisor Configuration** page. See [Configure Hypervisor, on page 57](#).

Step 1 On the **IP Addresses** page, select **Make IP Addresses Sequential** to make the IP Addresses sequential.

Step 2 When you enter IP addresses in the first row for Hypervisor, Storage Controller (Management) and Hypervisor, Storage Controller (Data) columns, the HX Data Platform Installer incrementally autofills the node information for the remaining nodes. The minimum number of nodes in the storage cluster is three. If you have more nodes, use the **Add** button to provide the address information.

Note Compute-only nodes can be added only after the storage cluster is created.

For each HX node, enter the Hypervisor, Storage Controller, Management, and Data IP addresses. For the IP addresses, specify if the network belongs to the Data Network or the Management Network.

Field	Description
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and the storage cluster.
Management Storage Controller	Enter the static IP address that handles the storage controller VM management network connection between the storage controller VM and the storage cluster.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	Enter the static IP address that handles the storage controller VM data network connection between the storage controller VM and the storage cluster.

Step 3 The IP address provided here are applied to one node in the storage cluster. In the event the node becomes unavailable the affected IP address is moved to another node in the storage cluster. All nodes must have a port configured to accept these IP addresses.

Provide the following IP addresses:

Field	Description
Management Cluster Data IP Address	Enter the management network IP address for the HX Data Platform storage cluster.
Data Cluster Data IP Address	Enter the IP address of data network for the HX Data Platform storage cluster.
Management Subnet Mask	Enter the subnet information for your VLAN and vSwitches. Provide the management network value. For example, 255.255.255.0.
Data Subnet Mask	Provide the network value for the data network. For example, 255.255.255.0.
Management Gateway	Provide the network value for your management network. For example, 10.193.0.1.
Data Gateway	Provide the network value for your data network. For example, 10.193.0.1.

Step 4 Click **Continue** to configure the HyperFlex cluster. See [Configure Your HyperFlex Cluster, on page 59](#).

Configure Your HyperFlex Cluster

On the **Cluster Configuration** page, for the Cisco HX Storage Cluster complete the following fields to begin deploying the HyperFlex cluster.

Before you begin

Ensure that you completed configuring IP addresses on the **IP Addresses** page. See [Configure IP Addresses, on page 58](#).

Step 1 In the **Cisco HX Cluster** section, complete the following fields:

Field	Description
Cluster Name	Specify a name for the HX Data Platform storage cluster.

Field	Description
Replication Factor	<p>Specify the number of redundant replicas of your data across the storage cluster. Set the replication factor to either 2 or 3 redundant replicas.</p> <ul style="list-style-type: none"> • For hybrid servers (servers that contain SSD and HDDs), the default value is 3. • For flash servers (servers that contain only SSDs), select either 2 or 3. • A replication factor of three is highly recommended for all environments except HyperFlex Edge. A replication factor of two has a lower level of availability and resiliency. The risk of outage due to component or node failures should be mitigated by having active and regular backups.

Step 2 In the **Controller VM** section, create a new password for the Administrative User of the HyperFlex cluster.

A default administrator username and password is applied to the controller VMs. The VMs are installed on all converged and compute-only nodes.

- Important**
- You cannot change the name of the controller VM or the controller VM's datastore.
 - Use the same password for all controller VMs. The use of different passwords is not supported.
 - Provide a complex password that includes 1 uppercase character, 1 lowercase character, 1 digit, 1 special character, and a minimum of 10 characters in total.
 - You can provide a user-defined password for the controller VMs and for the HX cluster to be created. For password character and format limitations, see the section on Guidelines for HX Data Platform Special Characters in the *Cisco HX Data Platform Management Guide*.

Step 3 In the **vCenter Configuration** section, complete the following fields:

Field	Description
vCenter Datacenter Name	Enter the vCenter datacenter name for the Cisco HyperFlex cluster.
vCenter Cluster Name	Enter the vCenter cluster name.

Step 4 In the **System Services** section, complete the following fields:

DNS Server(s)	A comma-separated list of IP addresses of each DNS server.
NTP Server(s)	<p>A comma-separated list of IP addresses of each NTP server.</p> <p>Note All hosts must use the same NTP server, for clock synchronization between services running on the storage controller VMs and ESXi hosts.</p>
DNS Domain Name	DNS FQDN or IP address.
Time Zone	The local time zone for the controller VM, to determine when to take scheduled snapshots. Scheduled native snapshot actions are based on this setting.

Step 5 In the **Connected Services** section, select **Enable Connected Services** to enable Auto Support and Intersight Management.

Field	Description
Enable Connected Services (Recommended)	Enables Auto Support and Intersight management. Log on to HX Connect to configure these services or selectively turn them On or Off .
Send service ticket notifications to	Email address where SR notifications are sent when triggered by Auto Support.

Step 6 In the **Advanced Configuration** section, do the following:

Field	Description
Jumbo frames Enable Jumbo Frames	<p>Check to set the MTU size for the storage data network on the host vSwitches and vNICs, and each storage controller VM.</p> <p>The default value is 9000.</p> <p>Note To set your MTU size to a value other than 9000, contact Cisco TAC.</p>
Disk Partitions Clean up Disk Partitions	<p>Check to remove all existing data and partitions from all nodes added to the storage cluster for manually prepared servers. Select this option to delete existing data and partitions. You must backup any data that should be retained.</p> <p>Attention Do not select this option for factory prepared systems. The disk partitions on factory prepared systems are properly configured.</p>
Virtual Desktop (VDI)	<p>Check for VDI only environments.</p> <p>Note To change the VDI settings after the storage cluster is created, shut down or move the resources, make the changes (described in the steps below), then restart the cluster.</p> <p>The HyperFlex cluster by default is configured to be performance tuned for VSI workloads.</p> <p>You may change this performance customization by performing the following steps on your HyperFlex Data Platform cluster. To change the HyperFlex cluster from VDI to VSI workloads (and vice versa):</p> <p>WARNING: A maintenance window is required as this will cause data to be unavailable while the cluster is offline.</p> <ol style="list-style-type: none"> Shut down the cluster (<code>hxcli cluster shutdown</code>). Edit the <code>storfs.cfg</code> in all the controller VMs to modify the <code>workloadType</code> to <code>Vsi</code> or <code>Vdi</code>. Start the cluster (<code>hxcli cluster start</code>) to enable the tune changes after the cluster is created.

Field	Description
(Optional) vCenter Single-Sign-On Server	<p>This information is only required if the SSO URL is not reachable.</p> <p>Note Do not use this field. It is used for legacy deployments.</p> <p>You can locate the SSO URL in vCenter by navigating to vCenter Server > Manage > Advanced Settings > key config.vpxd.sso.sts.uri.</p>

Step 7 Click **Start** to begin deploying the HyperFlex cluster. The **Progress** page displays the progress of various configuration tasks.

Caution Do not skip validation warnings.
See the **Warnings** section for more details.

What to do next

- Some validation errors require you to go back and re-enter a parameter (for example, an invalid ESXi password, incorrect NTP server, bad SSO server, or other incorrect input). Click **Re-enter Values** to return to the **Cluster Configuration** page and resolve the issue.
- When complete, the HyperFlex servers are installed and configured. The deployed cluster status shows as **Online** and **Healthy**.
- Click **Launch HyperFlex Connect** to create datastores and manage your cluster.

Installation of HyperFlex Nodes with GPUs

A specific BIOS policy change is required when installing HyperFlex nodes with GPUs. All supported GPU cards require enablement of BIOS setting that allows greater than 4 GB of Memory Mapped I/O (MMIO). For more information, see [Requirement for All Supported GPUs](#).

Installing GPU After the HyperFlex Cluster Is Created

If the GPUs are installed after a cluster is created, then the service profile associated with the servers must be modified to have the BIOS policy setting enabled.

Enable the BIOS Setting as detailed in [Cisco UCS Manager Controlled Server](#). Set Memory Mapped I/O above 4 GB config to **Enabled** as specified in *step 3*.

Installing GPU Before the HyperFlex Cluster Is Created

If the GPU card is installed before the cluster is created, then during cluster creation, select the *Advanced workflow*.

1. On the HX Data Platform Installer page, select **I know what I'm doing, let me customize my workflow**.
2. Check **Run UCS Manager Configuration** and click **Continue**.

This creates the necessary service profiles for the HyperFlex nodes.

3. Enable the BIOS Setting as detailed in [Cisco UCS Manager Controlled Server](#). Set Memory Mapped I/O above 4 GB config to **Enabled** as specified in *step 3*.
4. Go back to the *Advanced workflow* on the HX Data Platform Installer page to continue with **Run ESX Configuration, Deploy HX Software, and Create HX Cluster** to complete cluster creation.

HX Data Platform Installer Navigation Aid Buttons

- **Export Configuration**—Click the down arrow icon to download a JSON configuration file.
- **Workflow Info**—Hover over the information icon to view the current workflow. For HyperFlex cluster creation, the workflow info is *Create Workflow = Esx*.
- **Tech Support**—Click the question mark icon to view details related to the HyperFlex Data Platform software version. Click **Create New Bundle** to create a Tech Support Bundle for Cisco TAC.
- **Save Changes**—Click the circle icon to save changes made to the HyperFlex cluster configuration parameters.
- **Settings**—Click the gear icon to **Start Over** or **Log Out**.

Warnings and Error Messages

- UCSM configuration and Hypervisor configuration succeeded, but deployment or cluster creation fails—Click **Settings Icon > Start Over**. Select **I know what I'm doing, let me customize my workflow** to start the cluster configuration from the point where the failure occurred.
- IP Address screen shows as blank when you go back to re-enter values—Add the IP addresses manually. Click **Add Server** for the number of servers in your cluster and re-input all of the IP addresses on this page.
- Server reachability issues are seen observed when DNS is not properly configured on the Installer VM (SSO Error)—Edit the **SSO** field manually and either substitute IP address in place of FQDN or troubleshoot and remediate the DNS configuration.
- Ensure that a matching Cisco UCS Manager release to Cisco HyperFlex release is selected when creating another cluster—If a matching release is not selected, manually enter the correct release.

For the current compatibility matrix, refer to the *Software Versions* table in the latest [Release Notes for Cisco HX Data Platform](#).



CHAPTER 6

Configure Licensing with HyperFlex Data Platform

- [Smart Licensing and HyperFlex](#), on page 65
- [License Compliance and Feature Functionality](#), on page 68
- [License Management for Connected Environments](#), on page 69
- [License Management for Disconnected Environments](#), on page 73

Smart Licensing and HyperFlex

Overview

Cisco Smart Software Licensing (Smart Licensing) is an intelligent software license management system that automates time-consuming, manual licensing tasks, such as procuring, deploying, and managing licenses across your entire organization. It provides visibility into your license ownership and consumption so you know what you own and how you are using it.

Smart Licensing introduces company-wide license pooling. Server-based licenses, or smart licenses, are not node-locked to devices so they can be used on any compatible device owned by your company. Using virtual accounts to organize licenses and product instances for your company into logical entities, for example, by business unit, product type, or IT group, makes it easy to transfer devices and licenses across virtual accounts.

The Smart Licensing feature integrates with Cisco HyperFlex and is automatically enabled as soon as you create an HX storage cluster. For your HX storage cluster to start reporting license consumption, you must register it with Cisco Smart Software Manager (SSM) through your Cisco Smart Account. A Smart Account is a cloud-based repository that provides full visibility and access control to Cisco software licenses and product instances across your company. Registration is valid for one year.

Registration enables HyperFlex to be identified to a Smart Account and allows license usage to be reported to Cisco Smart Software Manager or Smart Software Manager satellite. After registration, HyperFlex reports license usage to Cisco Smart Software Manager or Smart Software Manager satellite with the current license status. See the following License Status section for details.



Note You will need to open ports 80 and 443 to `tools.cisco.com` for all HyperFlex management IPs in order for this to work.

After you register your HX storage cluster, communication messages are signed by a certificate used to identify HyperFlex to Cisco Smart Software Manager or Smart Software Manager satellite. HyperFlex automatically sends the following requests:

- A renew registration request every six months. In case the automatic registration renewal does not occur, use the `stcli license renew id` command to manually renew.
- A renew authorization request every 30 days is required by Smart Licensing. In case the automatic authorization renewal does not occur, use the `stcli license renew auth` command to manually renew. You need to manually renew Smart Licensing authorization only if connectivity is not available when renewal is attempted or your renewal time is outside of your connectivity window.
- A renew authorization request is also sent to Cisco Smart Software Manager or Smart Software Manager satellite whenever license consumption changes. This authorization is valid for 90 days. If HyperFlex doesn't contact Cisco Smart Software Manager or Smart Software Manager satellite for 90 days to renew its authorization, licenses consumed by HyperFlex are reclaimed and put back in the pool.

License Status

Registration Status	Description	Verify Status	System Functionality
Evaluation Mode	Smart Licensing is enabled but your HX storage cluster is not registered to Cisco Smart Software Manager or Smart Software Manager satellite and in a 90-day evaluation period.	To verify the status or check the remaining time left in the evaluation period, run #stcli license show all Result: Mode = Eval & Remaining period (Number of Days:Hours:Minutes)	No impact on features or functionality.
Evaluation Expired	Smart Licensing is enabled but your HX storage cluster is not registered to Cisco Smart Software Manager or Smart Software Manager satellite. Your license is in the Initial Unidentified state and not considered out of compliance.	To verify the status, run #stcli license show all Result: Mode = Evaluation Expired	No impact on features or functionality. <ul style="list-style-type: none"> • Generates syslog message. • Generates an <i>Evaluation Expired</i> alarm in the HX Connect UI.
In Compliance	Smart Licensing is enabled and your HX storage cluster is registered to Cisco Smart Software Manager or Smart Software Manager satellite. You are consuming <i>less</i> licenses than you own.	—	—

Registration Status	Description	Verify Status	System Functionality
HyperFlex Release 5.0(2a) and later Out of Compliance	-	-	See the License Compliance and Feature Functionality , on page 68 section for the impact to features and functionality.
HyperFlex Release 5.0(1b) and earlier Out of Compliance	<p>You are consuming <i>more</i> licenses than you own.</p> <p>Important Cisco will never interfere with or shutdown a customer network when a device is out of compliance.</p> <ul style="list-style-type: none"> • Out of Compliance at Initial Registration state—Smart Licensing is enabled and your HX storage cluster is registered to Cisco Smart Software Manager or Smart Software Manager Satellite but after initial registration you do not have enough licenses. • Out of Compliance after Initial state or after being in-compliance for some period— Smart Licensing is enabled and your HX storage cluster is registered to Cisco Smart Software Manager or Smart Software Manager Satellite but you no longer have enough licenses. 	<p>To verify the status, run <code>#stcli license show all</code>.</p> <p>Result: Mode = Out of Compliance</p>	<p>No impact on features or functionality.</p> <ul style="list-style-type: none"> • Generates syslog message. • Generates an <i>Out of Compliance</i> alarm in the HX Connect UI at cluster level. <p>Note The Out of Compliance state breaches the intellectual property EULA and the license must be purchased/renewed in order to continue receiving support.</p>

Registration Status	Description	Verify Status	System Functionality
Authorization Expired	Smart Licensing is enabled and your HX storage cluster is registered to Cisco Smart Software Manager or Smart Software Manager satellite but has not communicated to Cisco Smart Software Manager or Smart Software Manager satellite for more than 90 days.	To verify the status, run <code>#stcli license show status</code> . Result: Mode = Authorization Expired	No impact on features or functionality. <ul style="list-style-type: none"> Generates syslog message. No event or alarm on HX Connect. Cisco Smart Software Manager portal displays flags and notifications.
Export Control Flag Set to "Not Allowed"	Smart Licensing is enabled and your HX storage cluster is registered to Cisco Smart Software Manager or Smart Software Manager satellite but cannot register to use Export Control.	—	Behavior is mostly controlled by the Cisco Smart Software Manager server. Note This status is applicable only if the HX storage cluster contains restricted functionality.
ID Certificate Expired	Smart Licensing is enabled and your HX storage cluster is registered to Cisco Smart Software Manager or Smart Software Manager satellite but your ID certificate has not renewed for more than six months. Your license is in the Subsequent Unidentified state and is considered out of compliance.	To verify the status, run <code>#stcli license show status</code> Result: Mode: ID Certificate Expired To clear all conditions and return to <i>In Compliance status</i> , run the following command: <code>#stcli license renew <auth>/<id></code>	<ul style="list-style-type: none"> Generates syslog message. No event or alarm on HX Connect. Cisco Smart Software Manager portal displays flags and notifications.

License Compliance and Feature Functionality

Beginning with Cisco HXDP Release 5.0(2a), full feature functionality and configuration changes require a valid Cisco HyperFlex Software License. HX Connect users with expired or insufficient licenses at the end of the evaluation or the grace period after the license compliance date, view a prominent countdown banner that alerts the user to the license compliance need and provides a link to the license renewal page until the license expiration is remedied.

In the event a license passes both the license expiration date and the grace period countdown, the current configurations will operate as expected with limited information. Renewing the license allows a user to resume full feature functionality, and make configuration changes. For details and examples of the banners, see the [License Compliance and Feature Functionality](#) section of the Cisco HyperFlex Systems Ordering and Licensing Guide.

To review the Cisco End User Agreement (Cisco EULA), see https://www.cisco.com/c/en/us/about/legal/cloud-and-software/end_user_license_agreement.html

License Management for Connected Environments

To manage your licenses for connected environments, proceed as follows:

Registering a Cluster with Smart Licensing

Smart Licensing automatically integrates with your HX storage cluster and is enabled by default. You need not install Smart Licensing. Your HX storage cluster is unregistered with Smart Licensing and in a 90-day EVAL MODE. Within the 90 days, you need to register your HX storage cluster to use full functionality.



Attention Before registering the HyperFlex cluster with Smart Software Manager satellite, ensure that proxy is not configured. If proxy is configured, remove proxy before registering the cluster with Smart Software Manager satellite.

Before you begin

- Smart Licensing was introduced in Cisco HX Release 2.5. It is recommended that you confirm your cluster is running a HX 4.0 release and later.
- Before you can begin using Smart Licensing, you need to have a Cisco Smart Account. You can create (or select) a Smart Account while placing an order, or create a Smart Account outside of placing an order and add new or existing licenses over time.

To create a Smart Account, see **Cisco Software Central > Request a Smart Account** (<https://webapps.cisco.com/software/company/smartaccounts/home?route=module/accountcreation>).

You can register the HX storage cluster with Cisco Smart Software Manager (SSM) in one of the following ways:

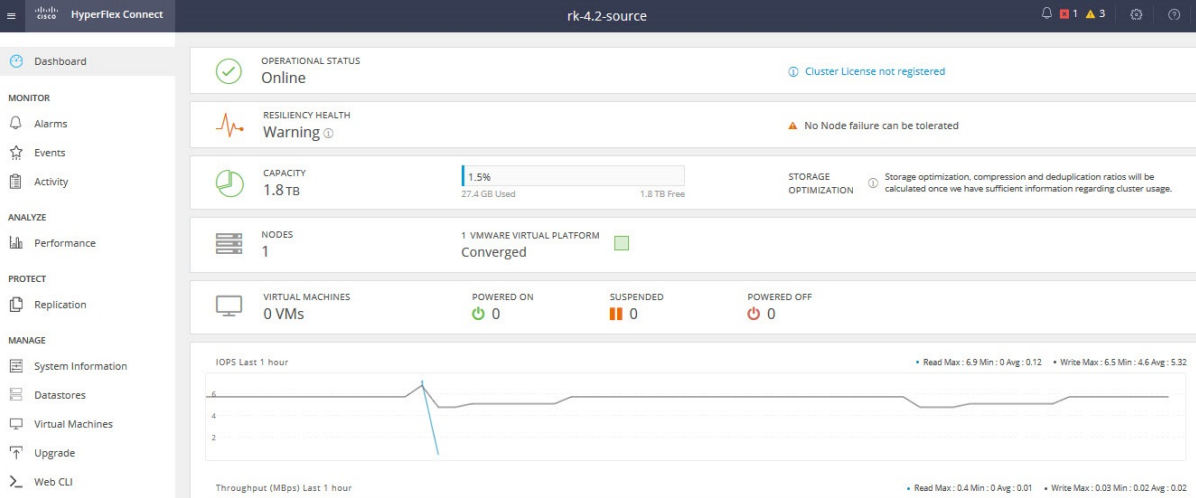
Registering a Cluster with Smart Software Licensing through HX Connect

Cisco recommends registration of cluster with Smart Software Licensing through HX Connect.

Before you begin

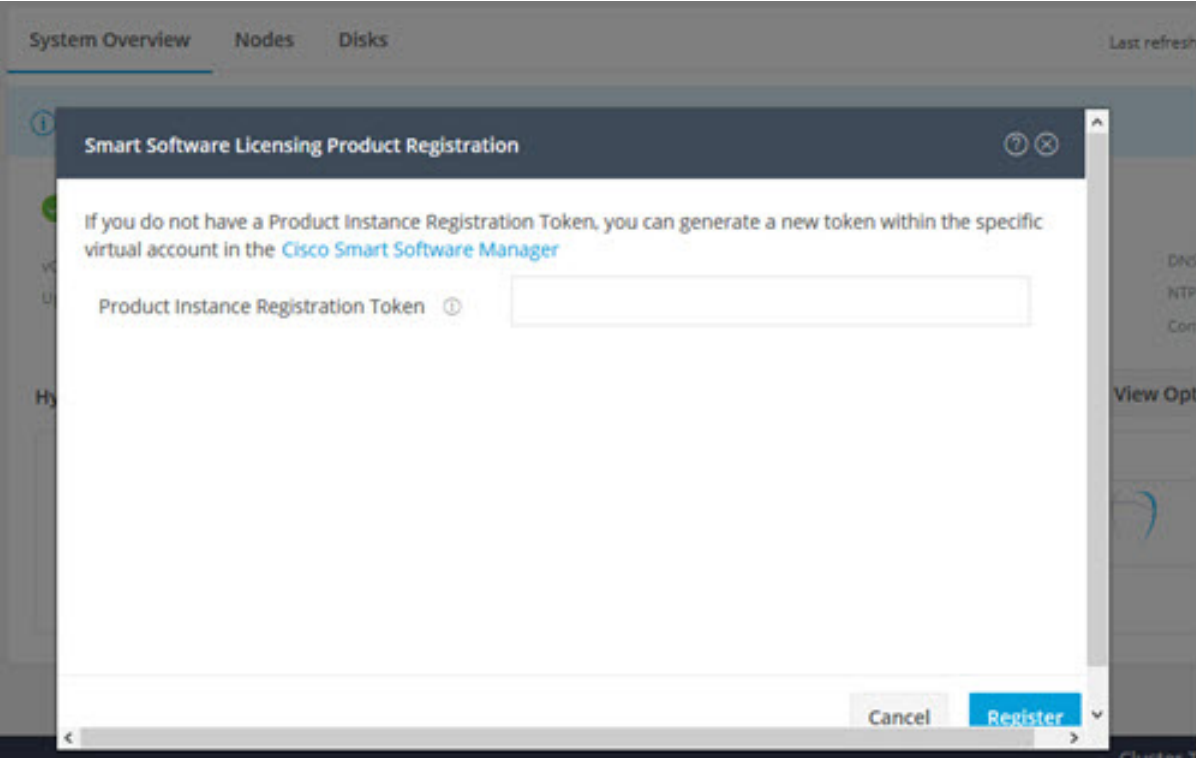
- You must have a product instance registration token. If you do not have a token, you can create a token in Cisco Smart Software manager. For more information on how to create a product instance registration token, see [Creating a Registration Token, on page 71](#).

- Step 1** Log into HX Connect.
- Step 2** In the **Dashboard** page, click **Cluster License not registered**.



Alternatively, you can do the registration by clicking the **Register Now** link in the **System Information** page.

- Step 3** In the **Smart Software Licensing Product Registration** dialog box, enter the registration token in the **Product Instance Registration Token** field.



For more information on how to create a product instance registration token, see [Creating a Registration Token](#), on page 71.

Step 4 Click **Register**.

On successful registration, the license type and license status appears in the **System Information** page.

- **License Type**—Displays Evaluation, Edge, or Data Centre as the HX storage cluster license type.
- **License Status**—Displays one of the following as the HX storage cluster license status:
 - In compliance
 - License expires in <n> days. Cluster not registered - Register Now. (This status appears only for Evaluation type license)
 - License expired. Cluster not registered - Register Now. (This status appears only for Evaluation type license)
 - Out of compliance - Insufficient license
 - Authentication expired—This status appears when HX is unable to communicate with Cisco Smart Software Manager or Smart Software Manager satellite for more than 90 days.

Creating a Registration Token

Registration Token is used to register and consume a product for smart licensing. You must create a token to register the product and add the product instance to a specified virtual account.

Step 1 Log into the software manager depending on which release you are using.

Option	Description
Cisco Smart Software Manager	Navigate to Cisco Software Central (https://software.cisco.com/) and log in to your Smart Account. In the License pane, click Smart Software Licensing . Click Inventory .
Smart Software Manager satellite	Navigate to <code>https://<IP address of the satellite>:8443</code> , and log in to the satellite using admin credentials.

Step 2 From the virtual account where you want to register your HX storage cluster, click **General**, and then click **New Token**.**Step 3** In the **Create Registration Token** dialog box, do the following, and then click **Create Token**:

- 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. Maximum = 365 days
- Check **Allow export-controlled functionality on the products registered with this token**.

Step 4 From the **New ID Token** row, click the **Actions** drop-down list, and click **Copy**.**Registering a Cluster with Smart Software Licensing through a Controller VM**

This section covers the alternate method for registering a cluster with Smart Software Licensing.

Step 1 Log into a controller VM.

Step 2 Confirm that your HX storage cluster is in Smart Licensing mode.

```
# stcli license show status
```

Feedback should 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). The Smart Licensing evaluation period starts when the HX storage cluster begins using the licensing feature and is not renewable. When the evaluation period expires, the Smart Agent sends a notification.

Step 3 Register your HX storage cluster, where *idtoken-string* is the **New ID Token** from Cisco Smart Software Manager or Smart Software Manager satellite.

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

Step 4 Confirm that your HX storage cluster is registered.

```
# stcli license show summary
```

Alternatively, you can confirm that your HX storage cluster is registered in **Cisco Smart Software Manager > Inventory > Product Instances**.

Example:

```
root@SpringpathController80IW1HJOKW:~# stcli license show summary
```

```
Smart Licensing is ENABLED
```

```
Registration:
```

```
Status: REGISTERED
Smart Account: Corp X HyperFlex License
Virtual Account: west-region
Last Renewal Attempt: None
Next Renewal Attempt: Aug 1 17:47:06 2017 PDT
```

```
License Authorization:
```

```
Status: AUTHORIZED
Last Communication Attempt: SUCCEEDED
Next Communication Attempt: Mar 4 16:47:11 2017 PST
```

```
License Usage:
```

```
License                               Entitlement Tag
Count  Status
```

```
-----
Cisco Vendor String XYZ
regid.2016-11.com.cisco.HX-SP-DP-S001,1.0_1c06ca12-18f2-47bd-bcea-518ab1fd4520 3      InCompliance
```

Deregistering a Cluster from Smart Licensing

You deregister an HX storage cluster when you want to release the licenses back into the pool to be used by another HX storage cluster or to remove Cisco Smart Software Manager registration; for example, if you want to decommission a cluster. After the HX storage cluster is deregistered, HyperFlex runs in evaluation mode, if any evaluation period remains. If not, HyperFlex is in evaluation expired status. Deregistering a cluster that is in an expired evaluation state, does not impact the cluster production data.

The Smart Agent contacts the licensing cloud and unregisters itself. All Smart Licensing entitlements and certificates on the platform are removed. All certificates and registration information on the trusted store are removed. The Smart Agent can unregister even if it is unable to communicate with Cisco to deregister. If you need to use Smart Licensing again, reregister your HX storage cluster. See [Registering a Cluster with Smart Software Licensing through a Controller VM, on page 71](#).

Before you begin

- Verify that your HX storage cluster is registered with Smart Licensing using the following command:

```
# stcli license show status
```

- Step 1** Log into a controller VM.
- Step 2** Deregister your HX storage cluster from Smart Licensing.
- ```
stcli license deregister
```
- Step 3** Confirm that your HX storage cluster is deregistered.
- ```
# stcli license show summary
```
-

Renewing Smart Licensing Authorization

Before you begin

- Verify that your HX storage cluster is registered with Smart Licensing using the following command:

```
# stcli license show status
```

- Step 1** Log into a controller VM.
- Step 2** Renew the Smart Licensing authorization using the following commands:
- ```
stcli license renew id
stcli license renew auth
```
- Step 3** Confirm that the HX storage cluster is renewed and authorized.
- ```
# stcli license show summary
```
-

License Management for Disconnected Environments

To manage your license for disconnected environments, proceed as follows:

Smart Licensing and Smart Software Manager Satellite

If you would prefer not to, or are not permitted to manage your installed base using an Internet connection, you can install Smart Software Manager Satellite on-premises to manage your licenses locally using a subset of Cisco Smart Software Manager. [Download Smart Software Manager satellite](#).

To configure Smart Software Manager satellite for HyperFlex, from the HX Data Platform CLI, run the following command:

```
stcli services sch set --portal-url
http://<satellite-host>/Transportgateway/services/DeviceRequestHandler --email
<user-email-address>
```

After Smart Software Manager satellite registers with Cisco Smart Software Manager and is fully functional, it needs to synchronize with Cisco Smart Software Manager every 30 days. There are two options for syncing:

- On-demand or scheduled synchronization when your network is connected.
- Manual synchronization by downloading and then uploading the license file.



Note Once the HX cluster is configured with a Smart Satellite server, the token can no longer be generated from the Smart Licensing registration from the Smart portal, as it needs to be generated from the Smart Satellite Server UI.



Note Smart Software Manager Satellite on-premises is not supported with Hyper-V.

Specific License Reservation and HyperFlex

Cisco Specific License Reservation (SLR) is a new software license management system that allows customers to use their devices in a disconnected mode, specifically for environments without external network connectivity (air-gapped). SLR also provides the following additional benefits:

- Automates time-consuming licensing tasks
- Allows tracking of the status of license and software usage trends
- Simplifies core purchasing, management, and reporting functions

SLR allows customers to reserve licenses from their virtual account, tie them to a device's UDI, and then use their device with these licenses in a disconnected mode. It also allows customers to operate normally without ongoing communication with either the Cisco Smart Software Manager (CSSM) or a Smart Software Satellite (on-site collector).

HyperFlex SLR enabled PIDs

Only the following HyperFlex PIDs are eligible for use under the SLR mode:

Table 4: Cisco HyperFlex Data Platform (HXDP) Software SKUs for Disconnected and Air-gapped Deployments

HXDP SKU	Description
Specific License Registration	
HXDP-S-SLR	Cisco HyperFlex Data Platform Standard Edition Specific License Registration Subscription
HXDP-P-SLR	Cisco HyperFlex Data Platform Enterprise Edition Specific License Registration Subscription
HXDP-E-SLR	Cisco HyperFlex Data Platform Edge Edition Specific License Registration Subscription

Installing Specific License Reservation (SLR) Licenses

This procedure explains how to install an SLR license, return an SLR license (for re-use by CSSM), or cancel an SLR license request.

The SLR installation process is very similar to the regular Smart License installation process. Since there is no communication between the cluster and the Smart account, a manual process must be used to carry out the same conversation that previously existed with the HTTP connection.

These conversations start from the Cisco hardware cluster where request codes are generated. The request code includes some basic cluster identification information. The request then gets carried to the Smart account in the Cisco portal to request an authorization code based on the request code. Once Cisco gets the authorization code, which includes both license identification and entitlement information, the authorization code is carried back to the cluster and installation can begin. Once the installation is complete, the license is fully activated on the cluster.

The enable/disable command is for getting to reservation mode. This is the default mode in the current registration feature. You have to specifically enable reservation mode to configure all of the reservations inside of the command. If you already have a license registered with CSSM, you should unregister it so that it can be recycled. Once you get the authorization code back from CSSM, you can use the `reservation install` command to install the reservation code. At some point, if you want to destroy the cluster or just return the license back to CSSM so that it can be re-used, you can use the `reservation return` command to generate a return code that can be deposited back.

The following steps describe how to install, return and cancel your SLR license.

Step 1

Enter the command `stcli license reservation enable` on the HX node to enable the reservation mode.

Entering this command switches the configuration mode into reservation mode. No license status is changed.

On the left side of the following screen, you can view the status for a typical cluster with a regular registration. In reservation mode, you can see the difference in that status. For a typical device, if you look at the reservation status, it shows as registered. The license authorization status will be authorized. If you have an individual license, it will tell you which license is in compliance.

On the right side, you can see the system is in the unregistered state and using the evaluation license.

Installing Specific License Reservation (SLR) Licenses

```

root@SpringpathControllerGVSTXU995I:~# stcli license show all
Smart Licensing Status
=====
Smart Licensing is ENABLED

Registration:
Status: REGISTERED
Smart Account: BU Production Test
Virtual Account: HXDP-Lic-Production-Test
Export-Controlled Functionality: Allowed
Initial Registration: SUCCEEDED on Jun 14 15:49:41 2017 PDT
Last Renewal Attempt: SUCCEEDED on Jun 14 15:49:41 2017 PDT
Next Renewal Attempt: Dec 11 14:49:41 2017 PST
Registration Expires: Jun 14 15:43:40 2018 PDT

License Authorization:
Status: AUTHORIZED on Jun 14 15:49:48 2017 PDT
Last Communication Attempt: SUCCEEDED on Jun 14 15:49:48 2017 PDT
Next Communication Attempt: Jul 14 15:49:48 2017 PDT
Communication Deadline: Sep 12 15:43:48 2017 PDT

Evaluation Period:
Evaluation Mode: Not In Use
Evaluation Period Remaining: 89 days, 12 hr, 40 min, 5 sec

License Usage
=====
License Authorization Status: AUTHORIZED as of Jun 14 15:49:48 2017 PDT

Cisco SP HyperFlex HX Data Platform SW v2.0 (regid.2016-11.com.cisco.HX-SP-DP-S001.1.0_10c6
Description: Cisco SP HyperFlex HX Data Platform SW v2.0
Count: 3
Version: 1.0
Status: InCompliance

Product Information
=====
UDI: PID:HX240C-M4SX,SN:9140506151354670828,VID: 5510173717264294049

Agent Version
=====
Smart Agent for Licensing: 1.3.5

```

- Login to control VM console
- `stcli license reservation enable`

```

root@SpringpathController2SAPEP8VJ9:~# stcli license reservation enable
root@SpringpathController2SAPEP8VJ9:~# stcli license show status
Smart licensing is ENABLED
License Reservation is ENABLED

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

License Authorization:
Status: EVAL MODE
Evaluation Period Remaining: 89 days, 23 hr, 54 min, 59 sec
Last Communication Attempt: NONE

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

Utility:
Status: DISABLED

Transport:
Type: TransportCallHome

```

Step 2 Enter the command `stcli license reservation request` to make a reservation request.

The license request code is shown in the blue box in the following screen.

```

root@SpringpathController2SAPEP8VJ9:~# stcli license reservation request
CB-PHX240C-M4SX,S1743837435069904050,V7822371211685355448-B6jnU5MNT-D4

root@SpringpathController2SAPEP8VJ9:~# stcli license show status
Smart Licensing is ENABLED
License Reservation is ENABLED

Registration:
Status: RESERVATION IN PROGRESS
Reservation process started on: Thu Aug 30 15:04:25 PDT 2018
Export-Controlled Functionality: Not Allowed

License Authorization:
Status: Evaluation Mode (84 days, 16 hr, 54 min, 14 sec remaining)

Utility:
Status: DISABLED

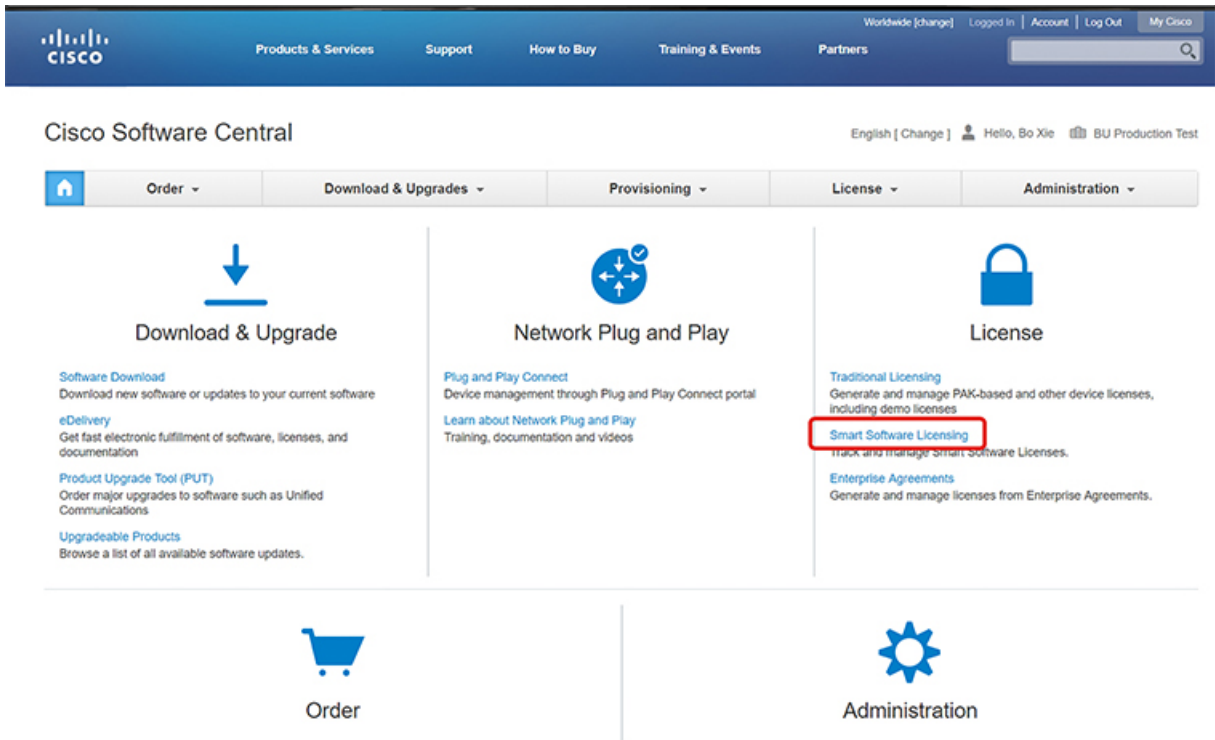
Transport:
Type: TransportCallHome
root@SpringpathController2SAPEP8VJ9:~# stcli license show reservation
Smart Licensing is ENABLED
License Reservation is ENABLED
RESERVATION IN PROGRESS
Request Code:CE-PHX240C-M4SX,S1743837435069904050,V7822371211685355448-B6jnU5MNT-B3
Last Return Code:CABuDC-HKhr56-uRDagz-poweUT-16GFAX-TKF

```

Once you start a reservation request, you can see that the registration status is `RESERVATION IN PROGRESS` (as shown in the red box). Once you have the Request Code, you can go to CSSM to convert it into an authorization code.


Step 3 Log into CSSM (located at <https://software.cisco.com>).

Step 4 In the **License** section, click on the **Smart Software Licensing** link.



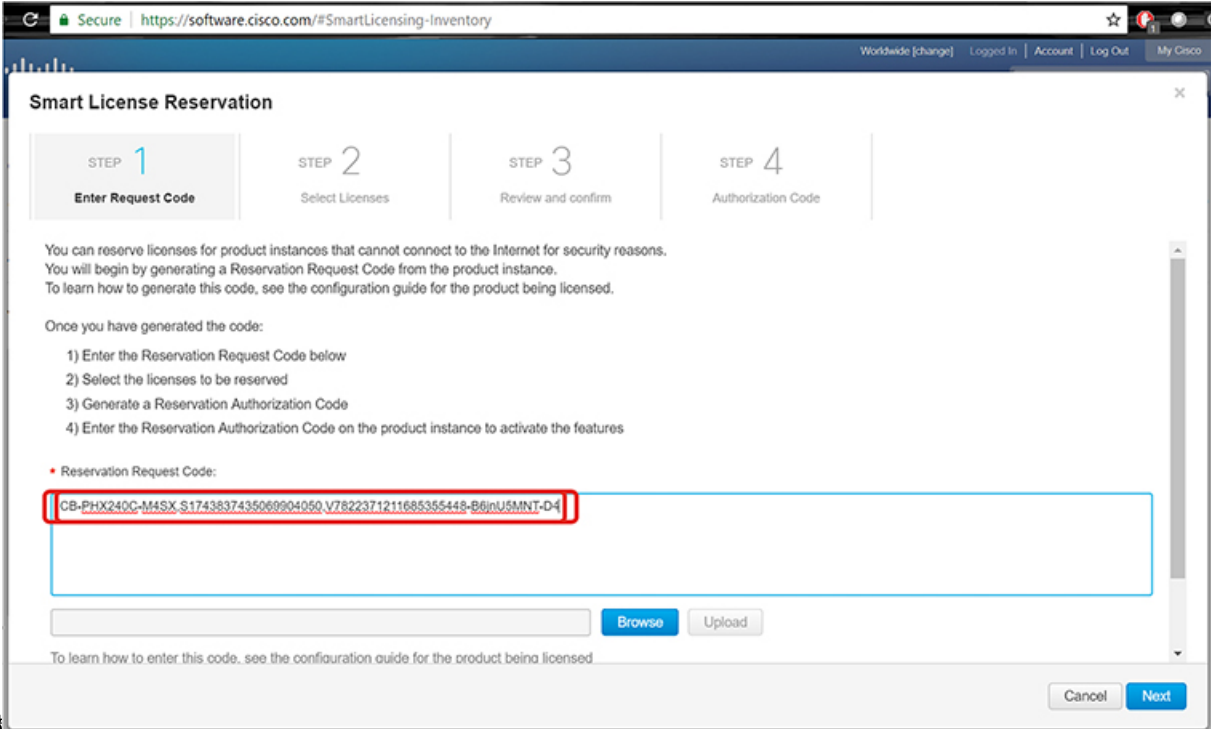
This takes you to the **Smart Software Licensing** page.

Step 5 Under the **Licenses** tab, click **License Reservation**.

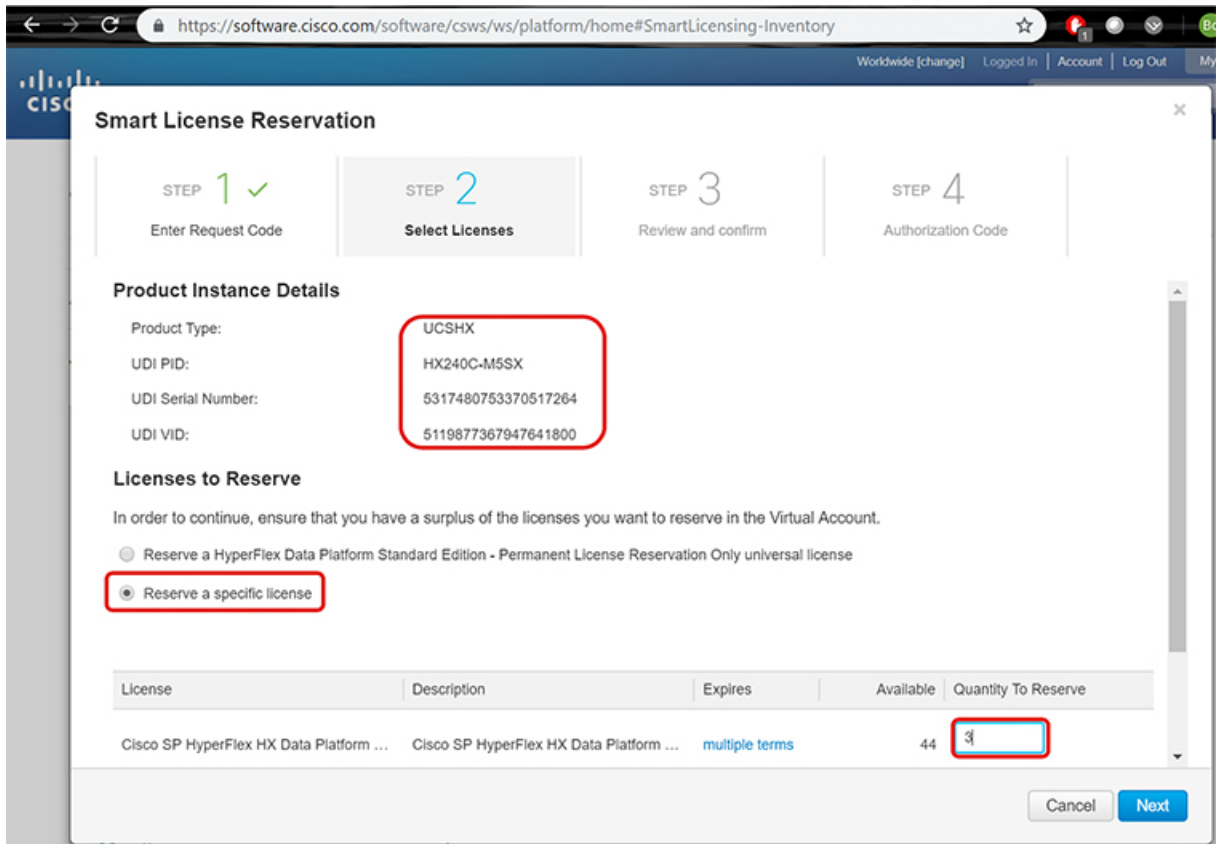


Step 6 Complete the four-step SLR process to Enter your Request Code, Select Licenses, Review and confirm the Authorization Code, then download the Authorization Code.

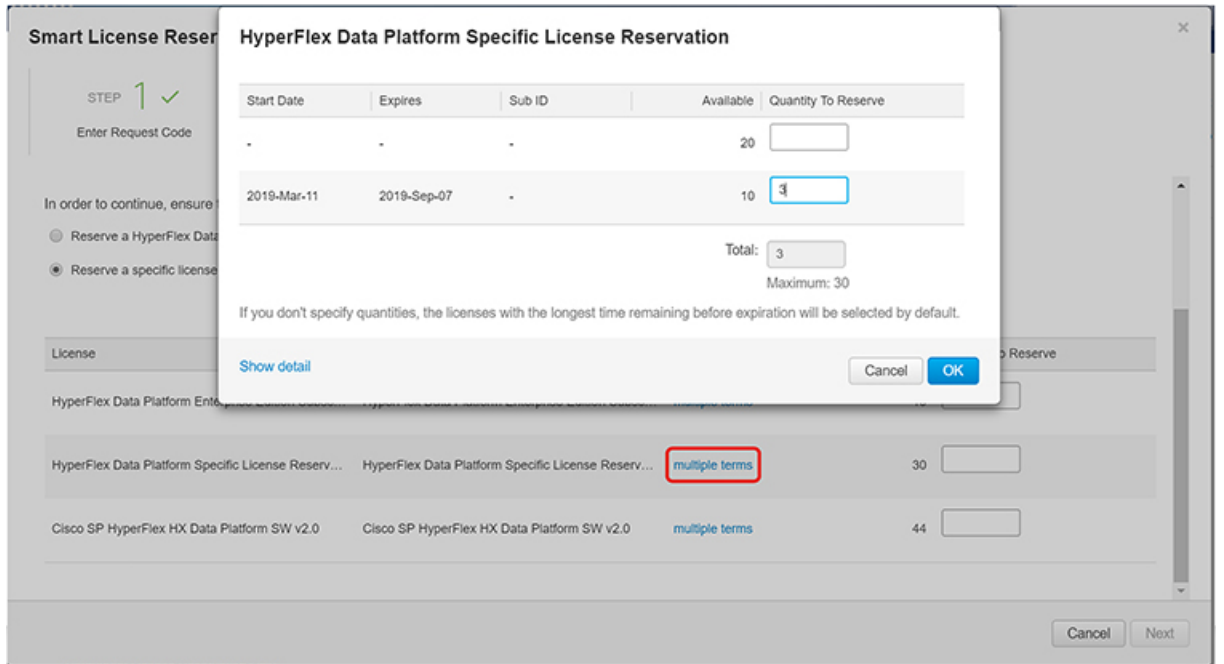
a. Enter the Request Code – Enter the Reservation Request Code that was just generated on the




b. Select Licenses – This screen tells you what the license is for and provides Product Instance Details such as the Product Type, UDI PID, UDI Serial Number and UDI VID. Verify the information provided, then select the checkbox to Reserve a specific license.



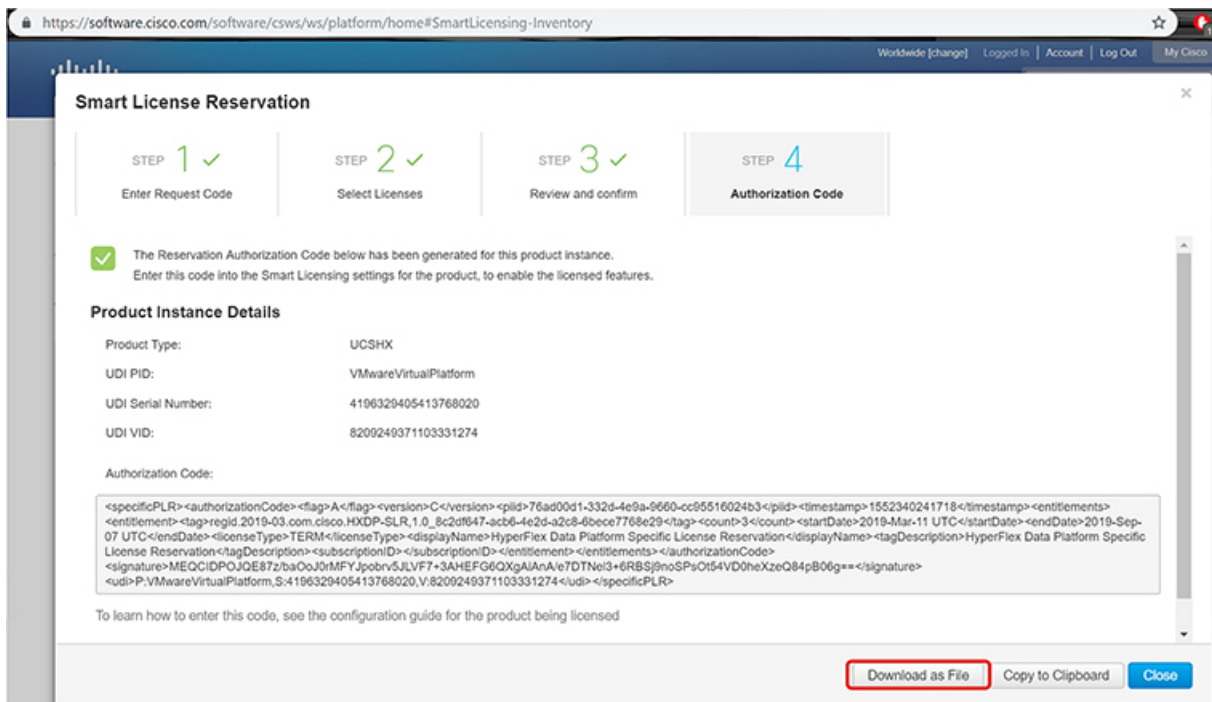
For the Specific License Reservation you have selected, enter the **Quantity to Reserve**.



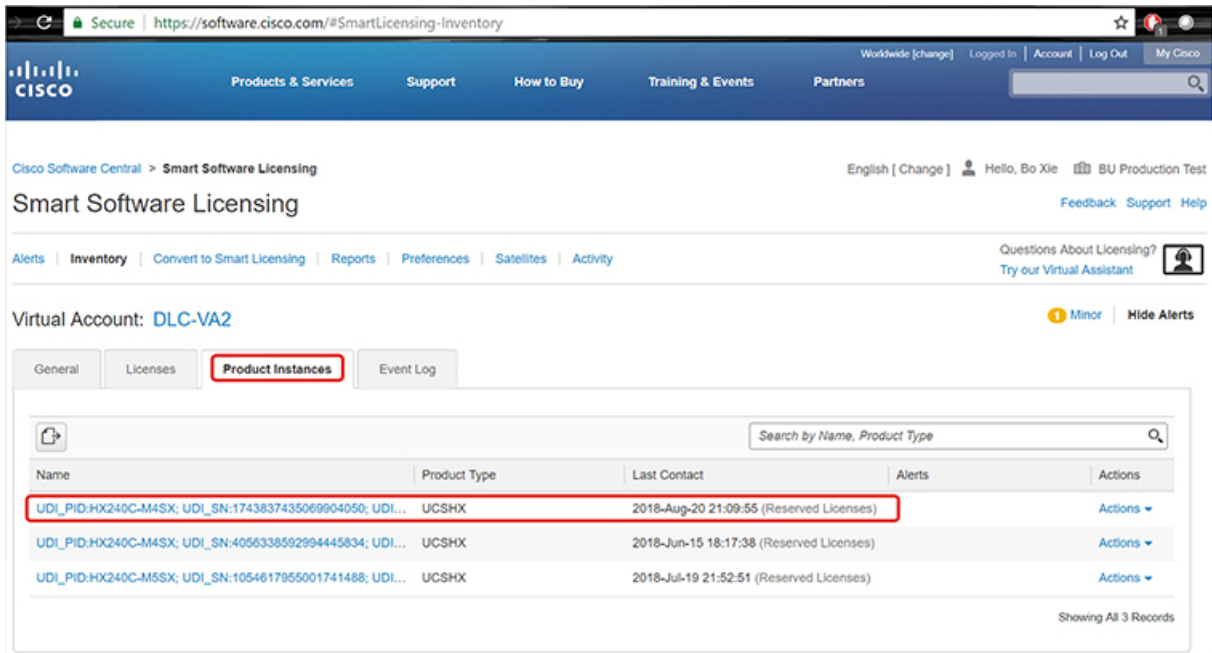
- c. **Review and Confirm** – Review and confirm the product instance details and license to reserve, then click the **Generate Authorization Code** button at the bottom of the screen.



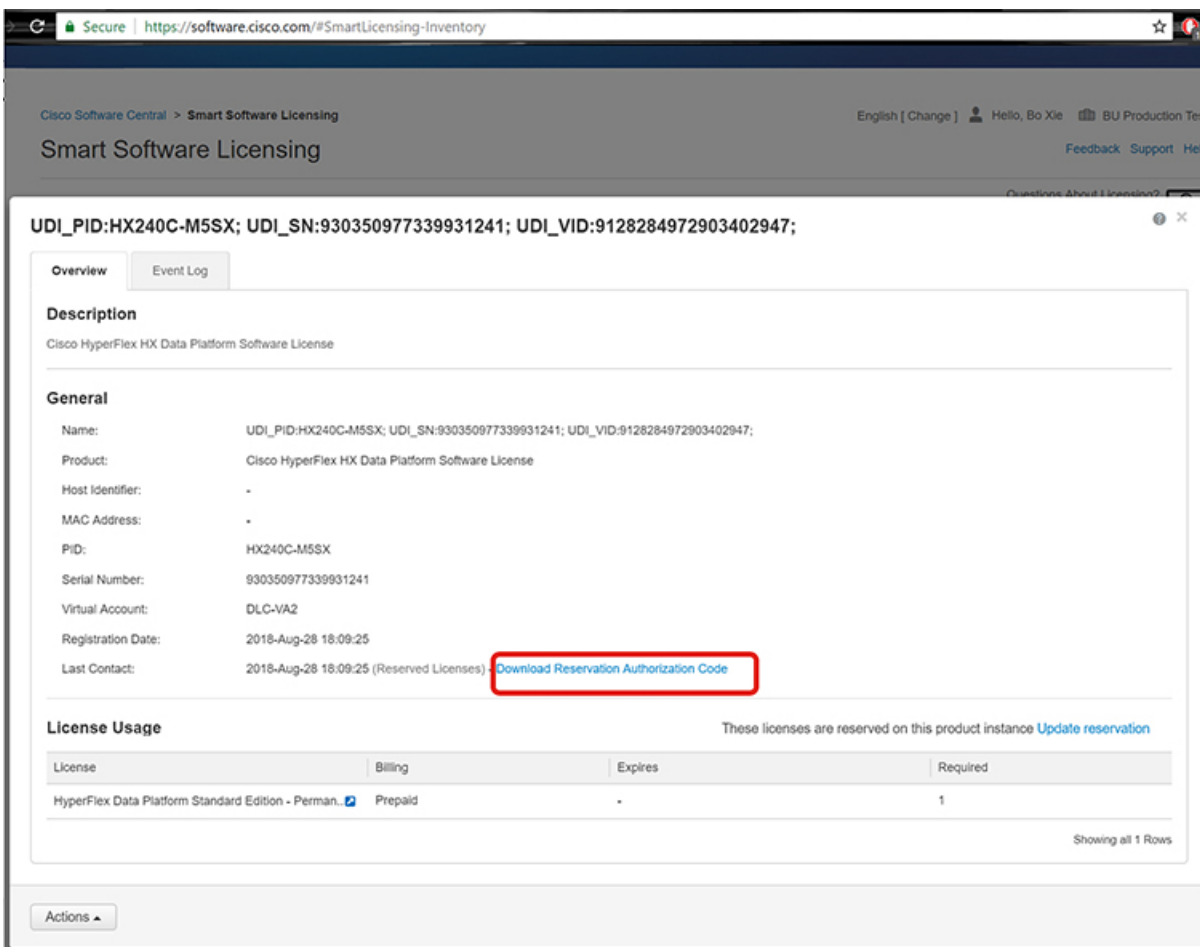
- d. **Authorization Code** – View the authorization code, which you can use on the device side, then click the **Download as File** button at the bottom of the screen.



- Step 7** Go to the **Product Instances** tab to view the license you just reserved, which is located in the row corresponding to the class and serial number of your reservation. Click the link in this row (shown in the red box).



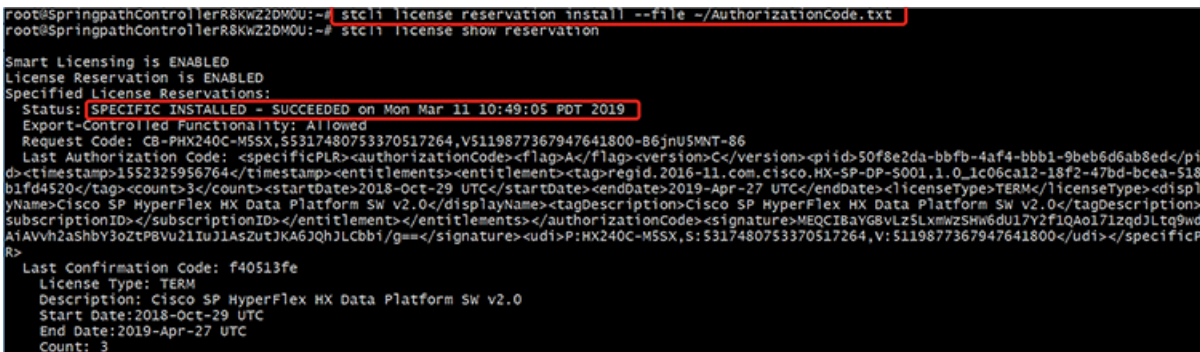
A dialog box appears showing a description of the license:



From this page, you can view general details of the license. You can also download the reservation authorization code (highlighted in red above) in case you lose it in any given instance. You can always return to this page to retrieve it again.

You can then go back to the **Licenses** tab to view the current license consumption.

Step 8 Enter the command `stcli license reservation install <enter authorization code>` on the HX node followed by your authorization code.



Once the reservation is successful, you can view the status, which shows as `REGISTERED - SPECIFIC LICENSE RESERVATION`. The authorization also shows that it is `AUTHORIZED - RESERVED`.

```

root@slcvm3:~# stcli license show all

Smart Licensing Status
=====
Smart Licensing is ENABLED
License Reservation is ENABLED

Registration:
  Status: REGISTERED - SPECIFIC LICENSE RESERVATION
  Export-Controlled Functionality: Allowed
  Initial Registration: SUCCEEDED on Thu Feb 21 09:23:52 PST 2019
  Last Renewal Attempt: None

License Authorization:
  Status: AUTHORIZED - RESERVED on Thu Feb 21 09:23:52 PST 2019

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

Utility:
  Status: DISABLED

Transport:
  Type: TransportCallHome

Evaluation Period:
  Evaluation Mode: Not In Use
  Evaluation Period Remaining: 87 days, 4 hr, 20 min, 45 sec

License Usage
=====
License Authorization Status:
  Status: AUTHORIZED - RESERVED on Mon Mar 04 14:37:18 PST 2019
  Last Communication Attempt: SUCCEEDED on Mar 4 14:37:18 2019 PST
  Next Communication Attempt: NONE

HyperFlex Data Platform Enterprise Edition Subscription (regid.2017-12.com.cisco.HXDP-P001,1.0_454a6b54-8b33-45bd-9d11-b1567c4a741e)
Description:
  Count: 3
  Version: 1.0
  Status: ReservedInCompliance
  Export status: NOI_RESTRICTED
  Feature Name: null
  Feature Description: null
Reservation:
  Reservation Status: SPECIFIC INSTALLED
  Total Reserved Count: 4

Product Information
=====
UDI: PID:HX240C-MSSX,SN:2432415792187588918,VID: 6984912678611873514

Agent Version
=====
Smart Agent for Licensing: 2.1.3

```


You can also enter the command `stcli license show reservation` on the HX node to show the SLR reservation. The response shows that SLR has been installed.

Canceling Specific License Reservation (SLR) Licenses

This procedure explains how to cancel an SLR license request.

- Step 1** To cancel a reservation request that you started (before going to CSSM to get an authorization code), use the `stcli license reservation cancel` command.
- Step 2** Verify that the reservation request was canceled by using the `stcli license show reservation` command. After entering this command, you should see that the status has returned to unregistered.

Returning Specific License Reservation (SLR) Licenses



Returning Specific License Reservation (SLR) Licenses

Now that the license on the cluster is fully active, at a later time, you can destroy the cluster and return the license to CSSM so it can be re-used for another cluster. The following steps describe how to return your SLR license.

Step 1 Enter the command `stcli license reservation return`. A return code is then generated that you can use in CSSM. If you look at the status, the license reverts back to being an unregistered evaluation license just as it was before you registered it.

```

root@SpringpathController2SAPEP8VJ9:~# stcli license reservation return
CABeUN-BvP26i-yju9Pc-Tw59i1-cNTFmt-MRq
root@SpringpathController2SAPEP8VJ9:~# stcli license show reservation

Smart Licensing is ENABLED
License Reservation is ENABLED
  Last Return Code:CABeUN-BvP26i-yju9Pc-Tw59i1-cNTFmt-MRq
root@SpringpathController2SAPEP8VJ9:~# stcli license show status

Smart Licensing is ENABLED
License Reservation is ENABLED

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

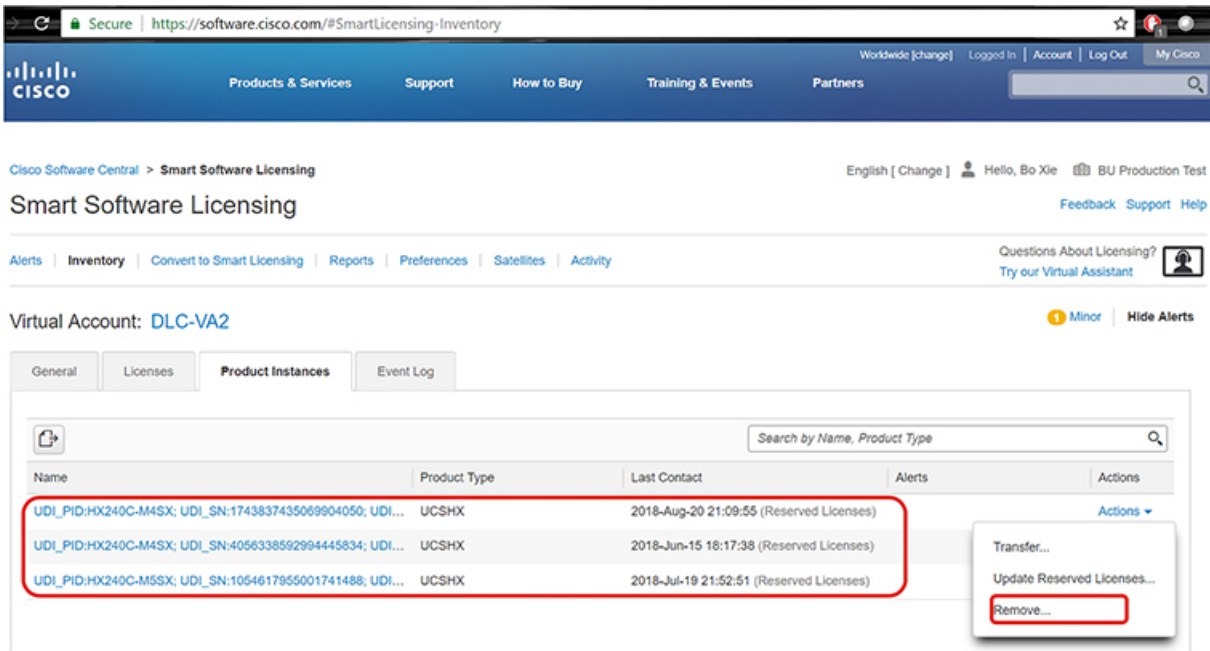
License Authorization:
  Status: EVAL MODE
  Evaluation Period Remaining: 89 days, 23 hr, 32 min, 3 sec
  Last Communication Attempt: SUCCEEDED on Aug 20 14:12:06 2018 PDT
  Next Communication Attempt: NONE

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

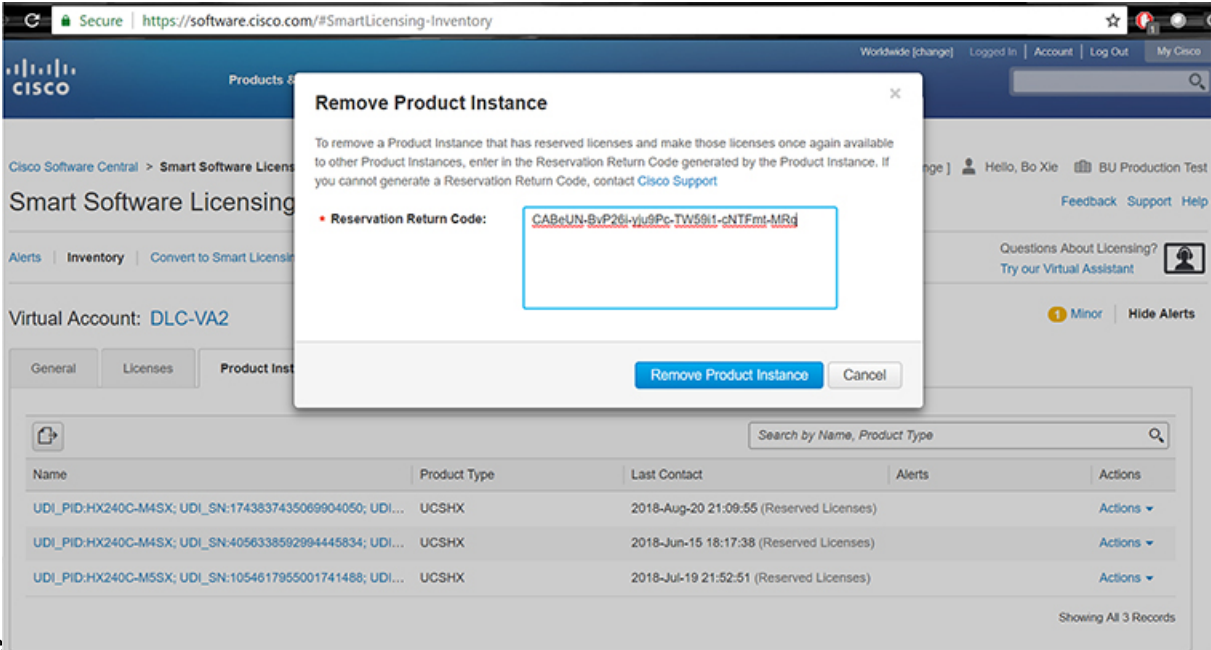
Utility:
  Status: DISABLED

Transport:
  Type: TransportCallHome
    
```

Step 2 Go back to CSSM to return the license back to the pool. Go back to the **Product Instances** tab, use the **Actions** menu, then click **Remove**.




The **Remove Product Instance** dialog box appears, which allows you to enter the return code. Enter the return code, then click **Remove Product**



Instance

In the **Product Instances** tab, you can see that the SLR license you previously registered is now gone. You can see that there are now only two in use, whereas there used to be three. At this point the license has been successfully returned.



Troubleshooting Specific License Reservation (SLR)

This section describes common error messages that may appear as you configure and use Specific License Reservation (SLR). It also provides recommendations on how to troubleshoot if applicable.

Two common error messages that may appear are as follows

- During configuration, if you issue a reservation request command before you enable the reservation mode, an error message appears indicating that "License Reservation is not enabled". Or, if you request to cancel an operation that you did not request, a message appears indicating that "No reservation process is pending". The following image shows these errors:

- Error you see from command line configuration output.

- Making reservation request before reservation is enabled. Issue "`stcli license reservation enable`" first

```
root@SpringpathController2SAPEP8VJ9:~# stcli license reservation request
Internal error processing slRequestReservation: 'com.cisco.nesla.agent.SmartAgentException: License Reservation is not enabled.'
```

- Making reservation cancellation when there is no pending request to cancel

```
root@SpringpathController2SAPEP8VJ9:~# stcli license reservation cancel
Internal error processing slCancelReservation: 'com.cisco.nesla.agent.SmartAgentException: No reservation process is pending.'
```

- During run-time when the license status changes, initial registration fails with a communication send error. This message typically appears when entering the `show status` command. The library does not propagate the error to the top but you can find the actual reason by using the logs.

- Smart Licensing Agent only runs on the node with the `mgmtip`.

- Do "`stcli license show status`" or "`stcli license show tech-support`" for check the status

```
root@SpringpathController2SAPEP8VJ9:~# stcli license show status
Smart Licensing is ENABLED

Registration:
Status: UNREGISTERED - REGISTRATION FAILED
Initial Registration: FAILED
Failure Reason: Communication send error.
Export-Controlled Functionality: Not Allowed

License Authorization:
Status: EVAL MODE
Evaluation Period Remaining: 84 days, 17 hr, 48 min, 14 sec
Last Communication Attempt: NONE

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

Utility:
Status: DISABLED

Transport:
Type: TransportCallHome
```

Step 1

Issue a `grep` command and view the SL column to identify errors and all other SL-related messages in the log. For example, the following image shows that a proxy is being used and that the proxy connection failed. This gives a good indication that the proxy configuration for your license server is not correct.

- `grep` for "ERROR\|SL:" in `/var/log/springpath/stNodeMgr.log`

```

2018-08-30-21:10:37.833 [ ] [Thread-6067] DEBUG c.s.s.stNodeMgr.StNodeMgrImpl5 - SL: getRegInfo mode!: Set(HX240C-M4SX), serials: Set(FCH2025V301, FCH2025V3HP, FCH2025V3FK)
2018-08-30-21:10:37.895 [ ] [Thread-6067] ERROR event_msg_sender_log - exception
2018-08-30-21:10:37.895 [ ] [Thread-6067] ERROR event_msg_sender_log - GCH Internal Set HTTPS Proxy [proxy-1.cisco.com : 0]connection Failed, Please check it.
2018-08-30-21:10:37.896 [ ] [Thread-6067] ERROR event_msg_sender_log - proxy check exception
2018-08-30-21:10:37.896 [ ] [Thread-6067] ERROR send_client_msg_log - Send SL data to URL [https://tools.cisco.com/its/service/oddce/services/DOCEService] : [fail]:GCH Internal Set HTTPS Proxy [proxy-1.cisco.com : 0]connection Failed, Please check it.
2018-08-30-21:10:37.896 [ ] [Thread-6067] ERROR c.c.n.p.EmbeddedGCHCommunication - ResultEntity.getError_msg(): Cannot send out SL Message.GCH Internal Set HTTPS Proxy [proxy-1.cisco.com : 0]connection Failed, Please check it.
2018-08-30-21:10:37.901 [ ] [Thread-6067] ERROR c.c.n.a.impl.AsyncRequestProcessor - failed to send request / process response: SmartAgentMessageReg
2018-08-30-21:10:37.901 [ ] [Thread-6067] ERROR c.c.n.a.impl.AgentKeyStoreManager - saving to keystore failed
2018-08-30-21:10:37.901 [ ] [Thread-6067] ERROR c.c.n.a.impl.AgentKeyStoreManager - saving to keystore failed
2018-08-30-21:10:37.901 [ ] [Thread-6067] ERROR c.c.n.a.impl.AsyncRequestProcessor - scheduled RegisterRetryJob
2018-08-30-21:10:37.901 [ ] [Thread-6067] INFO c.s.s.stNodeMgr.StNodeMgrImpl5 - SL: --> received global notification....
2018-08-30-21:10:37.901 [ ] [Thread-6067] INFO c.s.s.stNodeMgr.StNodeMgrImpl5 - SL: notification type: NotifyRegisterFailed
2018-08-30-21:10:37.901 [ ] [Thread-6067] INFO c.s.s.stNodeMgr.StNodeMgrImpl5 - SL: enforce mode: NotApplicable
    
```

- "`stcli services sch show`" reveals the proxy server setting error

- Step 2** To resolve this, use the `stcli services sch show` command to see the proxy set-up, correct any errors, then try again.
- Step 3** You can also issue a `grep` command on "CISCO-SMART-LIC" to check the Smart License syslog messages which are generated during the transition.

- `grep "CISCO-SMART-LIC" /var/log/syslog`

```

Aug 20 23:30:42.373 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% Smart Agent is initialized
Aug 20 23:30:42.376 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% Smart Agent is enabled
Aug 20 23:30:42.551 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% Smart Agent is initialized
Aug 20 23:31:34.420 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% Smart Agent DeRegistration with CSSM failed: Agent is not registered.
Aug 20 23:32:23.359 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% CD-PHX240C-M4SX,S1743837435069904050,V7822371211685355448-B6jnuSMT-BD License Reservation process must be completed with the 'license smart reservation install' command. Reservation started on PHX240C-M4SX,S1743837435069904050,V7822371211685355448
Aug 23 16:09:07.265 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% CE-PHX240C-M4SX,S1743837435069904050,V7822371211685355448-B6jnuSMT-B3 License Reservation process must be completed with the 'license smart reservation install' command. Reservation started on PHX240C-M4SX,S1743837435069904050,V7822371211685355448
Aug 23 16:17:32.621 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% CAABYP-doDLDA-qs2XGw-uZwHeY-XXwZf7-7n23 License Reservation Authorization code installed
Aug 23 16:17:32.622 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% Usage of export controlled features is true
Aug 23 16:17:32.793 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% Smart Agent Registration with Cisco licensing cloud successful
Aug 23 16:17:32.800 SpringpathController2SAPE8V39 root: %CISCO-SMART-LIC% All entitlements are authorized
    
```

Facilitating Controller VM Root Access for Air-Gapped Clusters

Starting with Cisco HX Release 5.0(2b), for HX nodes licensed with Subscription License Reservations or Permanent License Reservations, you now have the option to enable a persistent advance shell for troubleshooting, after a one-time Consent Token authentication with TAC. This enables access to root on-demand thereafter without needing to go through Consent Token (CT) challenge-response workflows with TAC repeatedly.



Note To enable this feature, you must have a license reservation with either SLR or PLR.



Note Once you enable this feature, you cannot invalidate the Consent Token.



Note Once a cluster is enabled with persistent root shell, all expanded nodes are enabled with persistent root automatically.

To facilitate Controller VM Root Access for Air-Gapped Clusters, proceed as follows:

-
- Step 1** From `hxshell`, issue an `su root` command. A warning message appears indicating the following:
- ```
WARNING: By accepting this support session, you give your consent and hereby authorize Cisco to have privileged access to the supported Cisco device for the purpose of providing technical support. At the conclusion of this session you must exit root shell from all the open ssh sessions of all the controller vms of the cluster and invalidate the consent token in order to terminate Cisco's access and close the privileged access portal. You are hereby advised that failure to do so may create a vulnerability in your product.
```
- Step 2** Click `y` to Accept. A message appears indicating that "Consent token is needed to access root shell !! with 4 options:"
1. Generate Challenge For root Shell Access
  2. Accept Response
  3. Exit
  4. Generate Challenge for Persistent Root Shell Access
- Step 3** Enter CLI Option 4. A Warning message appears indicating:
- ```
WARNING: Installation of this consent token will enable persistent root shell access on this device. You are hereby advised that this action may create a vulnerability in your product. Accept (N/y):
```
- Step 4** Click `y` to Accept.
- Step 5** Copy the Challenge String provided.
- Step 6** Contact TAC and provide the Challenge String and obtain the Consent Token.
- Step 7** If starting a new SSH session, re-enter `su root`. Enter CLI Option 2. A message appears indicating a Starting background timer of 30 mins. Please input the response when you are ready
- Step 8** Enter the Consent Token.
- When done, a message appears indicating the status of the Response Signature and Response processed. Next, a message appears requesting "Do you want to sync the consent token to other controller VMs in the cluster (y/n)?"
- Step 9** Enter "y" to sync the consent token with all other nodes in the cluster.
- Note** Sync the consent token with all other nodes in the cluster. You do not need to complete these procedures on multiple nodes.

EXAMPLE:

```

hxshell:~$ su root
Password:
WARNING: By accepting this support session, you give your consent and hereby authorize Cisco to have privileged access to
Cisco device for the purpose of providing technical support. At the conclusion of this session you must exit root shell
pen ssh sessions of all the controller vms of the cluster and invalidate the consent token in order to terminate Cisco's
se the privileged access portal. You are hereby advised that failure to do so may create a vulnerability in your product
Accept(Y/n): y
Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
4. Generate Challenge for Persistent Root Shell Access
Enter CLI Option:
4
WARNING: Installation of this consent token will enable persistent root shell access on this device. You are hereby adv
ction may create a vulnerability in your product. Accept (N/y):y
Generating Challenge.....
Challenge String (Please copy everything between the asterisk lines exclusively):
*****BEGIN TOKEN*****
6+3MygAAAEBAQAAAAABAgAEAAAAAwMACOIomz15anMJBAAQeK2gdNTjjV/DtGjnkS5QeQUABAAAAAGAAALeXBlcmZsZXgHAAxTeXBlcmZsZXhfQ1QIAAL
1ZDlLZDc0N2JiNmVmYzQzOGMwMmQzMzRjNmM0MTZhNw==
*****END TOKEN*****
Consent token is needed to access root shell !!
1. Generate Challenge For root Shell Access
2. Accept Response
3. Exit
4. Generate Challenge for Persistent Root Shell Access
Enter CLI Option:
2
Starting background timer of 30 mins
Please input the response when you are ready .....
1TPXIQAAAQEBAAQAAAAABAgAEAAAAAwMBYmxxMk1TczVrZWVUVVSeKJBVTdXTmhYRZta10xSudiVesvk2kxSFhXeTBjK2LZdmorakFBbGhSbWFXcTRWejg
4NUMrWUZ1MG01MU5MYllwCuxzYTLYeW1ZOVN6Mk9xRvd6dXJ5MFZMSmVJBFJ6ald0TwhITmMNCnp1TFLwckVQaHp3ckRLVklrait1b2cxeldpNkxscGtyNH
NaZkFWM0pob3F6WGRjaE93U0sNCnM3RVJscUdqVTAY0WkzRWNhRVMSdFdLZU5IUFR4UnVLRndhZHL1SL1JzeTVteXhTS3BtWULzMO5RTHZPTFLCZXkNCnrD
21lWmhvcTBMbG9aRzNoZmVwS3M0L2ZzT3BKs3FCSUcxwEXwbzNVVkh1KohL0XVCNw0NcjN4aVM2bE1wTmFQa1g2NXUyb3U4T3c9PQ=
Response Signature Verified successfully !
Response processed successfully.
Do you want to sync the consent token to other controller VMs in the cluster(y/n)? : y
Please re-enter root password:
Token synchronized successfully !!
Providing persistent root shell!!
bash-4.2#

```



CHAPTER

7

Configure HyperFlex Hardware Acceleration Cards

This chapter describes how to configure Cisco HyperFlex Hardware Acceleration Cards.

- [Overview of HyperFlex Hardware Acceleration Cards, on page 91](#)
- [Install HyperFlex Hardware Acceleration Cards, on page 91](#)
- [Deploy HX Data Platform Installer OVA Using vSphere Web Client, on page 92](#)
- [Deploy the HX Data Platform Installer OVA with a Static IP Address, on page 94](#)
- [Configure and Deploy Your HyperFlex Cluster, on page 96](#)
- [Verify Installation of HyperFlex Hardware Acceleration Cards, on page 108](#)
- [Troubleshoot HyperFlex Hardware Acceleration Cards, on page 108](#)
- [Additional Information on HyperFlex Hardware Acceleration Cards, on page 108](#)

Overview of HyperFlex Hardware Acceleration Cards

This chapter provides details for installation, post-installation, and troubleshooting of HyperFlex Hardware Acceleration Cards (PID: HX-PCIE-OFFLOAD-1) on HyperFlex nodes and for configuring your initial cluster. These cards provide improved performance and compression efficiency for most storage workloads.



Note An HXDP-P Enterprise license is required for installing and configuring HyperFlex Hardware Acceleration Cards.

For a description of PCIe slots and riser cards (and other related information), please refer to the [Cisco HX240c M5 HyperFlex Node \(Hybrid and All-Flash Models\) Installation Guide](#).

Install HyperFlex Hardware Acceleration Cards

Before you begin

Before beginning the HyperFlex Hardware Acceleration Cards installation process, note the following:

- Installation is only supported on NEW deployments.

- Installation is only supported on the following HX 240 M5/M6 servers:
 - HXAF240C-M5SX Cisco HyperFlex HX240c M5 All Flash
 - HXAF240C-M6S Cisco HyperFlex HX240c M6 All Flash
 - HX240C-M5SX Cisco HyperFlex HX240c M5
 - HX240C-M6SX Cisco HyperFlex HX240c M6
 - HX240C-M5L HyperFlex HX240c M5 LFF
 - HX240C-M6S HyperFlex HX240c M6 LFF
- Installation is not supported in Hyper-V and is not supported for stretched clusters.
- All nodes in the HX cluster must contain the HyperFlex Hardware Acceleration Cards.



Note If any nodes do not contain an acceleration card during validation, the installer fails to proceed, and an error message appears.

- The cluster type must be All Flash/Hybrid ESXi.
- Installation is only supported on HX 240 M5/M6 servers.
- Hardware Acceleration Cards do not work with Cisco HX Data Replication.
- Starting with HX Release 4.0(2b), Hardware Offload option with Stretched cluster configurations is supported.

Step 1 Install the PCIe cards. This is a Cisco-assisted installation.

Step 2 Configure your cluster using the HX installer. For more information, see:

- [Deploying HX Data Platform Installer OVA Using vSphere Web Client](#)
- [Configure and Deploy your HX Cluster](#)

Deploy HX Data Platform Installer OVA Using vSphere Web Client

In addition to installing the HX Data Platform on an ESXi host, you may also deploy the HX Data Platform Installer on either VMware Workstation, VMware Fusion, or Virtual Box.

**Note**

- Connect to vCenter to deploy the OVA file and provide the IP address properties. Deploying directly from an ESXi host will not allow you to set the values correctly.
- Do not deploy the HX Data Platform Installer to an ESXi server that is going to be a node in the Cisco HX Storage Cluster.

Step 1

Locate and download the HX Data Platform Installer OVA from [Download Software](#). Download the HX Data Platform Installer to a node that is on the storage management network, which will be used for the HX Data Platform storage cluster.

Example:

```
Cisco-HX-Data-Platform-Installer-v4.5.1a-26363.ova
```

Step 2

Deploy the HX Data Platform Installer using VMware hypervisor, to create a HX Data Platform Installer virtual machine.

Note Use a release of the virtualization platform that supports virtual hardware version 10.0 or greater.

vSphere is a system requirement. You can use either vSphere thick client, vSphere thin client, or vSphere Web Client. To deploy the HX Data Platform Installer, you can also use VMware Workstation, VMware Fusion, or VirtualBox.

- Open a virtual machine hypervisor, such as vSphere, VirtualBox, Workstation, or Fusion.
- Select the node where you want to deploy the HX Data Platform Installer.

Important Ensure that you provide user credentials while deploying the HX Installer OVA using vSphere Web Client.

- Using vSphere thick Client—Expand **Inventory list > Host > File > Deploy OVA**.
- Using vSphere Web Client—Expand **vCenter Inventory list > Hosts > Host > Deploy OVA**.

Step 3

Select where the HX Data Platform Installer is located. Accept the defaults, and select the appropriate network.

Step 4

Enter a static IP address for use by the HX Data Platform Installer VM.

Note

- Static IP Address is necessary even if DHCP is configured for the network. You need the static IP address to run the HX Data Platform Installer, to install the HX Data Platform, and to create the HX Data Platform storage cluster.
- If your hypervisor wizard defaults to DHCP for assigning IP addresses to new VMs, then complete the steps in [Deploy the HX Data Platform Installer OVA with a Static IP Address, on page 50](#), to install the HX Data Platform Installer VM with a static IP address. DNS must be reachable from the Installer VM.

Field	Description
Hostname	The hostname for this VM. Leave blank to try to reverse lookup the IP address.
Default Gateway	The default gateway address for this VM. Leave blank if DHCP is desired.

Field	Description
DNS	The domain name servers for this VM (comma separated). Leave blank if DHCP is desired.
IP Address	The IP address for this interface. Leave blank if DHCP is desired.
Netmask	The netmask or prefix for this interface. Leave blank if DHCP is desired.
Root Password	The root user password. This field is a required field.

Step 5 Click **Next**. Verify if the options listed are correct and select **Power on after deployment**.

To power on the HX Data Platform Installer manually, navigate to the virtual machine list and power on the installer VM.

Note The preferred settings for the HX Data Platform Installer virtual machine is 3 vCPU and 4 GB of memory. Reducing these settings can result in 100% CPU usage and spikes for the host.

Step 6 Click **Finish**. Wait for the HX Data Platform Installer VM to be added to the vSphere infrastructure.

Step 7 Open the HX Data Platform Installer virtual machine console.

The initial console display lists the HX Data Platform Installer virtual machine IP address.

```
Data Platform Installer.
*****
You can start the installation by visiting
the following URL:
http://192.168.10.210
*****
Cisco-HX-Data-Platform-Installer login:
```

Step 8 Use the URL to log in to the HX Data Platform Installer.

```
Example:
http://192.168.10.210
```

Step 9 Accept the self-signed certificate.

Step 10 Log in using the username **root** and the password you provided as part of the OVA deployment.

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

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

- Step 1** Install the VMware OVF Tool 4.1 or later on a node that is on the storage management network that will be used for the HX Data Platform storage cluster. See [OVF Tool Documentation](#) for more details.
- Step 2** Locate and download HX Data Platform installer OVA from [Download Software](#) on the node where VMware OVF was installed.
- Step 3** Deploy the downloaded HX Data Platform installer OVA, using the `ovftool` command. For example:

```
root@server:/tmp/test_ova# ovftool --noSSLVerify --diskMode=thin
--acceptAllEulas=true --powerOn --skipManifestCheck --X:injectOvfEnv
--datastore=qa-048-ssd1 --name=rfsi_static_test1 --network='VM Network'
--prop:hx.3gateway.Cisco_HX_Installer_Appliance=10.64.8.1
--prop:hx.4DNS.Cisco_HX_Installer_Appliance=10.64.1.8
--prop:hx.5domain.Cisco_HX_Installer_Appliance=cisco
--prop:hx.6NTP.Cisco_HX_Installer_Appliance=10.64.8.5
--prop:hx.1ip0.Cisco_HX_Installer_Appliance=10.64.8.36
--prop:hx.2netmask0.Cisco_HX_Installer_Appliance=255.255.248.0
--prop:hx.7root_password.Cisco_HX_Installer_Appliance=mypassword
/opt/ovf/rfsi_test/Cisco-HX-Data-Platform-Installer-v1.7.1-14786.ova
vi://root:password@esx_server
```

The command deploys the HX Data Platform installer, powers on the HX Data Platform installer VM, and configures the provided static IP address. A sample of processing response:

```
Opening OVA source:
/opt/ovf/rfsi_test/Cisco-HX-Data-Platform-Installer-v1.7.1-14786.ova
Opening VI target: vi://root@esx_server:443/
Deploying to VI: vi://root@esx_server:443/
Transfer Completed
Powering on VM: rfsi_static_test
Task Completed
Completed successfully
```

DNS must be reachable from the Installer VM. The required command options for the static IP address to be configured successfully are:

Command	Description
<code>powerOn</code>	To power on the HX Data Platform installer VM after it is deployed.
<code>X:injectOvfEnv</code>	To insert the static IP properties onto the HX Data Platform installer VM.
<code>prop:hx.3gateway.Cisco_HX_Installer_Appliance=10.64.8.1</code>	Specify the appropriate gateway IP address.
<code>prop:hx.4DNS.Cisco_HX_Installer_Appliance=10.64.1.8</code>	Specify the appropriate DNS IP address.
<code>prop:hx.5domain.Cisco_HX_Installer_Appliance=cisco</code>	Specify the appropriate domain.
<code>prop:hx.6NTP.Cisco_HX_Installer_Appliance=10.64.8.5</code>	Specify the appropriate NTP IP address.
<code>prop:hx.1ip0.Cisco_HX_Installer_Appliance=10.64.8.36</code>	Specify the appropriate installer static IP address.
<code>prop:hx.2netmask0.Cisco_HX_Installer_Appliance=255.255.248.0</code>	Specify the appropriate netmask address.
<code>prop:hx.7root_password.Cisco_HX_Installer_Appliance=mypassword</code>	Specify the root user password.
<code>/opt/ovf/rfsi_test/Cisco-HX-Data-Platform-Installer-v1.7.1-14786.ova</code>	The source address of the HX Data Platform installer OVA.

Command	Description
vi://root:password@esx_server	The destination ESX server where the HX Data Platform installer VM is installed. Include the appropriate ESX server root login credentials.

Configure and Deploy Your HyperFlex Cluster

Enter Credentials

On the **Credentials** page, you can choose to import the required configuration data from a JSON file or input data into the required fields manually.



Note For first-time installation of a HyperFlex cluster, contact your Cisco representative to procure the factory preinstallation JSON file.



Note The root user is created with the same password as the admin user during cluster creation. It is important to track the root user password because future changes to the admin password do not automatically update the root password .

To perform cluster creation, by importing the configuration data from a *JSON configuration* file, do the following:

1. Click **Select a file** and choose your *JSON file* to load the configuration data. Select **Use Configuration**.
2. An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Step 1 In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. On the HX Data Platform Installer login page, verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.

Step 2 In the login page, enter the following credentials:

Username: root

Password (Default): Cisco123

Attention Systems ship with a default password of Cisco123 that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Step 3 Select the **I accept the terms and conditions** check box, and click **Login**.

Step 4 On the **Select a Workflow** page, from the **Create Cluster** drop-down list, select **Standard Cluster**.

Step 5 On the **Credentials** page, enter the following configuration data:

Field	Description
UCS Manager Host Name	Enter the UCS Manager FQDN or IP address. For example, <i>10.193.211.120</i> .
UCS Manager User Name	Enter the administrative username.
Password	Enter the administrative password.

vCenter Credentials

Field	Description
vCenter Server	Enter the vCenter server FQDN or IP address. For example, <i>10.193.211.120</i> . Note <ul style="list-style-type: none"> • A vCenter server is required before the cluster can be made operational. • The vCenter address and credentials must have root level administrator permissions to the vCenter. • vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.
User Name	Enter the administrative username. For example, <i>administrator@vsphere.local</i> .
Admin Password	Enter the administrative password.

Step 6 On the **Hypervisor Configuration** page, enter the following configuration data:

Hypervisor Credentials

Field	Description
Admin User Name	Enter the administrative username. The username is root for factory nodes.

Field	Description
New Password	<p>Important You are required to change the factory default password for the hypervisor.</p> <p>Create a new password for the hypervisor using the following guidelines:</p> <ul style="list-style-type: none"> • Must be 6 to 80 characters in length. • Must have 1 uppercase character, 1 lowercase character, 1 digit, and 1 special character. • If password starts with uppercase character, then 2 uppercase characters are required. • If password ends with a digit, then 2 digits are required.
Confirm New Password	Re-enter the new password for the hypervisor.

Step 7 Click **Continue** to begin associating HyperFlex Servers. See [Associate HyperFlex Servers, on page 52](#).

Associate HyperFlex Servers

On the **Server Selection** page, the **Configuration** pane on the right displays a detailed list of the **Credentials** used. The **Server Selection** page displays a list of unassociated HX servers under the **Unassociated** tab, and the list of discovered servers under the **Associated** tab.

Field	Description
Locator LED	Turn on to locate a server.
Server Name	Name assigned to the server.
Status	<ul style="list-style-type: none"> • Inaccessible—
Model	Displays the server model.
Serial	Displays the serial number of the server.
Assoc State	<ul style="list-style-type: none"> • Associated • Unassociated
Service Profile [Only for Associated Servers]	<p>Service profile assigned to the server.</p> <p>Note Editing the HyperFlex Service Profile templates is not recommended.</p>

Field	Description
Actions	<ul style="list-style-type: none"> • Launch KVM Console—Choose this option to launch the KVM Console directly from the HX Data Platform Installer. • Disassociate Server—Choose this option to remove a service profile from that server.

Before you begin

Ensure that you completed entering UCS Manager, vCenter, and Hypervisor credentials.

Step 1 Click the **Configure Server Ports** button to discover any new HX nodes. In the **Configure Server Ports** dialog box, list all ports to be configured as server ports. Click **Configure**.

Note Typically, the server ports are configured in Cisco UCS Manager before you start the configuration.

Step 2 Select the servers under the **Unassociated** tab to include in the HyperFlex cluster.

If HX servers do not appear in this list, check Cisco UCS Manager and ensure that they have been discovered.

Note If there are no unassociated servers, the following error message is displayed:

No unassociated servers found. Login to UCS Manager and ensure server ports are enabled.

Step 3 Click **Continue** to configure UCS Manager. See [Configure UCS Manager, on page 53](#).

Configure UCS Manager

On the **UCSM Configuration** page, you can configure VLAN, MAC Pool, 'hx-ext-mgmt' IP Pool for CIMC, iSCSI Storage, and FC Storage.

Before you begin

Associate servers to the HyperFlex cluster. See [Associate HyperFlex Servers, on page 52](#).

Step 1 In the **VLAN Configuration** section, complete the following fields:

Note Use separate subnet and VLANs for each of the following networks.

Field	Description
VLAN for Hypervisor and HyperFlex management	
VLAN Name	hx-inband-mgmt
VLAN ID	Default—3091
VLAN for HyperFlex storage traffic	

Field	Description
VLAN Name	hx-storage-data
VLAN ID	No default value.
VLAN for VM vMotion	
VLAN Name	hx-vmotion
VLAN ID	Default—3093
VLAN for VM Network	
VLAN Name	vm-network
VLAN ID(s)	Default—3094 A comma-separated list of guest VLANs.

Step 2 In the **MAC Pool** section, configure **MAC Pool Prefix** by adding in two more hex characters (0-F).

Note Select a prefix that is not used with any other MAC address pool across all UCS domains.

Example:
00:25:B5:A0:

Step 3 In the '**hx-ext-mgmt**' **IP Pool for CIMC** section, complete the following fields:

Field	Description
IP Blocks	The range of management IP addresses assigned to the CIMC for each HyperFlex server. The IP addresses are specified as a range, and multiple blocks of IPs may be specified as a comma-separated list. Ensure you have at least one unique IP per server in the cluster. When selecting to use out-of-band, this range must fall into the same IP subnet used on the mgmt0 interfaces on the Fabric Interconnects. For example, <i>10.193.211.124-127, 10.193.211.158-163.</i>
Subnet Mask	Specify the subnet mask for the IP range provided above. For example, <i>255.255.0.0.</i>
Gateway	Enter the Gateway IP address. For example, <i>10.193.0.1.</i>

The management IP addresses used to access the CIMC on a server can be either:

- **Out of band:** The CIMC management traffic traverses the Fabric Interconnect through the limited bandwidth management interface, mgmt0, on the Fabric Interconnect. This option is most commonly used and shares the same VLAN as the Fabric Interconnect management VLAN.

- **In-band:** The CIMC management traffic traverses the Fabric Interconnect through the uplink ports of the Fabric Interconnect. The bandwidth available for management traffic in this case would be equivalent to the Fabric Interconnect uplink bandwidth. If you are using the In-band option, the Cisco HyperFlex installer will create a dedicated VLAN for the CIMC management communication. This option is useful when large files such as a Windows Server installation ISO must be mounted to the CIMC for OS installation. This option is only available in the HyperFlex installer VM and is not available for deployments through Intersight.

Step 4 Select either **Out of band** or **In-band** based on the type of connection you want to use for CIMC management access. If you select In-band, provide the VLAN ID for the management VLAN. Make sure to create the CIMC management VLAN in the upstream switch for seamless connectivity.

Step 5 If you want to add external storage, configure **iSCSI Storage** by completing the following fields:

Field	Description
Enable iSCSI Storage check box	Select to configure iSCSI storage.
VLAN A Name	Name of the VLAN associated with the iSCSI vNIC, on the primary Fabric Interconnect (FI-A).
VLAN A ID	ID of the VLAN associated with the iSCSI vNIC, on the primary Fabric Interconnect (FI-A).
VLAN B Name	Name of the VLAN associated with the iSCSI vNIC, on the subordinate Fabric Interconnect (FI-B).
VLAN B ID	ID of the VLAN associated with the iSCSI vNIC, on the subordinate Fabric Interconnect (FI-A).

Step 6 If you want to add external storage, configure **FC Storage** by completing the following fields:

Field	Description
Enable FC Storage check box	Select to enable FC Storage.
WWxN Pool	A WWN pool that contains both WW node names and WW port names. For each Fabric Interconnect, a WWxN pool is created for WWPN and WWNN.
VSAN A Name	The name of the VSAN for the primary Fabric Interconnect (FI-A). Default—hx-ext-storage-fc-a.
VSAN A ID	The unique identifier assigned to the network for the primary Fabric Interconnect (FI-A). Caution Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.

Field	Description
VSAN B Name	The name of the VSAN for the subordinate Fabric Interconnect (FI-B). Default— <code>hx-ext-storage-fc-b</code> .
VSAN B ID	The unique identifier assigned to the network for the subordinate Fabric Interconnect (FI-B). Caution Do not enter VSAN IDs that are currently used on the UCS or HyperFlex system. If you enter an existing VSAN ID in the installer which utilizes UCS zoning, zoning will be disabled in your existing environment for that VSAN ID.

Step 7 In the **Advanced** section, do the following:

Field	Description
UCS Server Firmware Release	Select the UCS firmware release to associate with the HX servers from the drop-down list. The UCS firmware release must match the UCSM release. See the latest Cisco HX Data Platform Release Notes for more details. For example, <code>3.2(1d)</code> .
HyperFlex Cluster Name	Specify a user-defined name. The HyperFlex cluster name is applied to a group of HX Servers in a given cluster. The HyperFlex cluster name adds a label to service profiles for easier identification.
Org Name	Specify a unique <i>Org Name</i> to ensure isolation of the HyperFlex environment from the rest of the UCS domain.

Step 8 Click **Continue** to configure the Hypervisor. See [Configure Hypervisor, on page 57](#).

Configure Hypervisor



Note Review the VLAN, MAC pool, and IP address pool information on the **Hypervisor Configuration** page, in the **Configuration** pane. These VLAN IDs may be changed by your environment. By default, the HX Data Platform Installer sets the VLANs as non-native. You must configure the upstream switches to accommodate the non-native VLANs by appropriately applying a trunk configuration.



Attention You can skip configuring Hypervisor in case of a reinstall, if ESXi networking has been completed.

Before you begin

Configure VLAN, MAC Pool, and 'hx-ext-mgmt' IP Pool for Out-of-Band CIMC. If you are adding external storage, configure iSCSI Storage and FC Storage. Select the UCS Server Firmware Version and assign a name for the HyperFlex cluster. See [Configure UCS Manager, on page 53](#).

Step 1

In the **Configure Common Hypervisor Settings** section, complete the following fields:

Field	Description
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, <i>255.255.0.0</i> .
Gateway	IP address of gateway. For example, <i>10.193.0.1</i> .
DNS Server(s)	IP address for the DNS Server. Note <ul style="list-style-type: none"> • If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform Installer. Use only static IP addresses and hostnames for all ESXi hosts. • If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.

Step 2

On the **Hypervisor Settings** section, select **Make IP Addresses and Hostnames Sequential** to make the IP addresses sequential. Complete the following fields:

Note You can rearrange the servers using drag and drop.

Field	Description
Name	Name assigned to the server.
Locator LED	Turn on to locate a server.
Serial	Displays the serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.
Hostname	Do not leave the hostname fields empty.

Step 3

Click **Continue** to configure IP Addresses. See [Configure IP Addresses, on page 58](#).

Configure IP Addresses

Before you begin

Ensure that you completed configuring Hypervisor on the **Hypervisor Configuration** page. See [Configure Hypervisor, on page 57](#).

Step 1 On the **IP Addresses** page, select **Make IP Addresses Sequential** to make the IP Addresses sequential.

Step 2 When you enter IP addresses in the first row for Hypervisor, Storage Controller (Management) and Hypervisor, Storage Controller (Data) columns, the HX Data Platform Installer incrementally autofills the node information for the remaining nodes. The minimum number of nodes in the storage cluster is three. If you have more nodes, use the **Add** button to provide the address information.

Note Compute-only nodes can be added only after the storage cluster is created.

For each HX node, enter the Hypervisor, Storage Controller, Management, and Data IP addresses. For the IP addresses, specify if the network belongs to the Data Network or the Management Network.

Field	Description
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and the storage cluster.
Management Storage Controller	Enter the static IP address that handles the storage controller VM management network connection between the storage controller VM and the storage cluster.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	Enter the static IP address that handles the storage controller VM data network connection between the storage controller VM and the storage cluster.

Step 3 The IP address provided here are applied to one node in the storage cluster. In the event the node becomes unavailable the affected IP address is moved to another node in the storage cluster. All nodes must have a port configured to accept these IP addresses.

Provide the following IP addresses:

Field	Description
Management Cluster Data IP Address	Enter the management network IP address for the HX Data Platform storage cluster.
Data Cluster Data IP Address	Enter the IP address of data network for the HX Data Platform storage cluster.

Field	Description
Management Subnet Mask	Enter the subnet information for your VLAN and vSwitches. Provide the management network value. For example, 255.255.255.0.
Data Subnet Mask	Provide the network value for the data network. For example, 255.255.255.0.
Management Gateway	Provide the network value for your management network. For example, 10.193.0.1.
Data Gateway	Provide the network value for your data network. For example, 10.193.0.1.

Step 4 Click **Continue** to configure the HyperFlex cluster. See [Configure Your HyperFlex Cluster, on page 59](#).

Configure Your HyperFlex Cluster

On the **Cluster Configuration** page, for the Cisco HX Storage Cluster complete the following fields to begin deploying the HyperFlex cluster.

Before you begin

Ensure that you completed configuring IP addresses on the **IP Addresses** page. See [Configure IP Addresses, on page 58](#).

Step 1 In the **Cisco HX Cluster** section, complete the following fields:

Field	Description
Cluster Name	Specify a name for the HX Data Platform storage cluster.
Replication Factor	Specify the number of redundant replicas of your data across the storage cluster. Set the replication factor to either 2 or 3 redundant replicas. <ul style="list-style-type: none"> For hybrid servers (servers that contain SSD and HDDs), the default value is 3. For flash servers (servers that contain only SSDs), select either 2 or 3. A replication factor of three is highly recommended for all environments except HyperFlex Edge. A replication factor of two has a lower level of availability and resiliency. The risk of outage due to component or node failures should be mitigated by having active and regular backups.

Step 2 In the **Controller VM** section, create a new password for the Administrative User of the HyperFlex cluster.

A default administrator username and password is applied to the controller VMs. The VMs are installed on all converged and compute-only nodes.

- Important**
- You cannot change the name of the controller VM or the controller VM's datastore.
 - Use the same password for all controller VMs. The use of different passwords is not supported.
 - Provide a complex password that includes 1 uppercase character, 1 lowercase character, 1 digit, 1 special character, and a minimum of 10 characters in total.
 - You can provide a user-defined password for the controller VMs and for the HX cluster to be created. For password character and format limitations, see the section on Guidelines for HX Data Platform Special Characters in the *Cisco HX Data Platform Management Guide*.

Step 3 In the **vCenter Configuration** section, complete the following fields:

Field	Description
vCenter Datacenter Name	Enter the vCenter datacenter name for the Cisco HyperFlex cluster.
vCenter Cluster Name	Enter the vCenter cluster name.

Step 4 In the **System Services** section, complete the following fields:

DNS Server(s)	A comma-separated list of IP addresses of each DNS server.
NTP Server(s)	A comma-separated list of IP addresses of each NTP server. Note All hosts must use the same NTP server, for clock synchronization between services running on the storage controller VMs and ESXi hosts.
DNS Domain Name	DNS FQDN or IP address.
Time Zone	The local time zone for the controller VM, to determine when to take scheduled snapshots. Scheduled native snapshot actions are based on this setting.

Step 5 In the **Connected Services** section, select **Enable Connected Services** to enable Auto Support and Intersight Management.

Field	Description
Enable Connected Services (Recommended)	Enables Auto Support and Intersight management. Log on to HX Connect to configure these services or selectively turn them On or Off .
Send service ticket notifications to	Email address where SR notifications are sent when triggered by Auto Support.

Step 6 In the **Advanced Configuration** section, do the following:

Field	Description
Jumbo frames Enable Jumbo Frames	Check to set the MTU size for the storage data network on the host vSwitches and vNICs, and each storage controller VM. The default value is 9000. Note To set your MTU size to a value other than 9000, contact Cisco TAC.

Field	Description
Disk Partitions Clean up Disk Partitions	<p>Check to remove all existing data and partitions from all nodes added to the storage cluster for manually prepared servers. Select this option to delete existing data and partitions. You must backup any data that should be retained.</p> <p>Attention Do not select this option for factory prepared systems. The disk partitions on factory prepared systems are properly configured.</p>
Virtual Desktop (VDI)	<p>Check for VDI only environments.</p> <p>Note To change the VDI settings after the storage cluster is created, shut down or move the resources, make the changes (described in the steps below), then restart the cluster.</p> <p>The HyperFlex cluster by default is configured to be performance tuned for VSI workloads.</p> <p>You may change this performance customization by performing the following steps on your HyperFlex Data Platform cluster. To change the HyperFlex cluster from VDI to VSI workloads (and vice versa):</p> <p>WARNING: A maintenance window is required as this will cause data to be unavailable while the cluster is offline.</p> <ol style="list-style-type: none"> Shut down the cluster (<code>hxcli cluster shutdown</code>). Edit the <code>storfs.cfg</code> in all the controller VMs to modify the <code>workloadType</code> to <code>Vsi</code> or <code>Vdi</code>. Start the cluster (<code>hxcli cluster start</code>) to enable the tune changes after the cluster is created.
(Optional) vCenter Single-Sign-On Server	<p>This information is only required if the SSO URL is not reachable.</p> <p>Note Do not use this field. It is used for legacy deployments.</p> <p>You can locate the SSO URL in vCenter by navigating to vCenter Server > Manage > Advanced Settings > key config.vpxd.sso.sts.uri.</p>

Step 7 Click **Start** to begin deploying the HyperFlex cluster. The **Progress** page displays the progress of various configuration tasks.

Caution Do not skip validation warnings.
See the **Warnings** section for more details.

What to do next

- Some validation errors require you to go back and re-enter a parameter (for example, an invalid ESXi password, incorrect NTP server, bad SSO server, or other incorrect input). Click **Re-enter Values** to return to the **Cluster Configuration** page and resolve the issue.

- When complete, the HyperFlex servers are installed and configured. The deployed cluster status shows as **Online** and **Healthy**.
- Click **Launch HyperFlex Connect** to create datastores and manage your cluster.

Verify Installation of HyperFlex Hardware Acceleration Cards

You can verify that you installed the HyperFlex Hardware Acceleration Card successfully as follows:

Step 1 Log into the controller VM.

Step 2 Find the following tunes file: `/opt/springpath/config/offload.tunes`.

Note The tunes file can only be seen by System Administrators and root users. If these users are able to access this file, the installation is successful. If you do not have System Administrator or root user permissions, you can verify that the installation is successful if there are no error messages or events in the UI.

Troubleshoot HyperFlex Hardware Acceleration Cards

Troubleshoot post-installation HyperFlex Hardware Acceleration Card related issues as follows:

Symptom	Workaround
Although the cluster is operational, you may note generic alerts in the vCenter and HX Connect UI.	Contact your Cisco representative for assistance.
The Cluster reports a pairing error when NR pairing is attempted.	NR pairing is not allowed if either cluster is on release prior to 4.5(1a), and has been enabled with HX Hardware Acceleration Card. NR pairing with HX Hardware Acceleration card is only supported if both the clusters in the pair have HX Hardware Acceleration Cards.

Additional Information on HyperFlex Hardware Acceleration Cards

Additional notes on HyperFlex Hardware Acceleration Cards include the following:

- Compression gain value can be seen in HX connect UI dashboard.
- Enhance performance for the 8K read workload using the following commands:
 - `root@ucs984scvm:~# echo 3 > /sys/module/fdma/parameters/decompress_min_t`
 - `root@ucs984scvm:~# echo 3 > /sys/module/fdma/parameters/decompress_max_t`



CHAPTER 8

Post Cluster Configuration Tasks

- [Post Cluster Configuration Guidelines](#), on page 109
- [Enabling PCI Passthrough for a Network Device on a Host](#), on page 109
- [Run Post-Installation Script](#), on page 110
- [Changing ESXi Host Root Password](#), on page 113
- [Changing Storage Controller Password](#), on page 114
- [Cisco HyperFlex HTML Plugin for VMware vCenter](#), on page 114
- [Add Datastores in the Storage Cluster](#), on page 115
- [Set HA Heartbeat](#), on page 115
- [Auto Support and Smart Call Home for HyperFlex](#), on page 115
- [Replacing Self-Signed with CA-Signed Certificate](#), on page 121
- [Replication Pairing](#), on page 122
- [Adding Private VLAN](#), on page 122
- [Distributed Virtual Switches and Cisco Nexus 1000v](#), on page 126
- [Hosting vCenter on the HX Data Platform](#), on page 127
- [Deploying AMD GPUs](#), on page 127

Post Cluster Configuration Guidelines



Important

- Keep SSH enabled on all ESXi hosts. This is required for the following Cisco HyperFlex post cluster configuration operations.
 - Do not change these pre-configured values without approval from Cisco.
-

Enabling PCI Passthrough for a Network Device on a Host

Passthrough devices provide the means to more efficiently use resources and improve performance in your environment. Enabling PCI passthrough allows a VM to use a host device as if the device were directly attached to the VM.



Caution Never setup up HXDP cluster critical devices for PCI passthrough.

The following procedure describes how to configure a network device (such as NVIDIA GPUs) for PCI passthrough on an ESXi host.

- Step 1** In vSphere Client, browse to the ESXi host in the Navigator panel.
- Step 2** Enter HX maintenance mode on the node that has the GPUs installed. To enter maintenance mode, right click on the node > **Cisco HX Maintenance Mode** > **Enter HX Maintenance Mode**
- Step 3** In a new browser window, login directly to the ESXi node.
- Step 4** Click **Manage**.
- Step 5** Under the **Hardware** tab, click **PCI Devices**. A list of available passthrough devices appears.
- Step 6** Select PCI device you want to enable for passthrough. Click **Toggle passthrough**.
- Step 7** Reboot the host to make the PCI device available for use.
- Step 8** When the reboot completes, ensure that the node is not in maintenance mode.
- Step 9** Log into vCenter Server.
- Step 10** Locate the VM, right click and select elect **Edit Settings**.
- Step 11** From the **New device** drop-down, select **PCI Device**, and click **Add**.
- Step 12** Click the passthrough device to use (example: NVIDIA GPU) and click **OK**.
- Step 13** Log into the ESXi host and open the virtual machine configuration file (.vmx) in a text editor.

```
cd /vmfs/volumes/[datastore_name]/[vm_name]
vi [vmname].vmx
```

- Step 14** Add the following lines, save, and exit the text editor.

```
# pciPassthru.64bitMMIOSizeGB = "64"
# Firmware = "efi"
# pciPassthru.use64bitMMIO = "TRUE"
```

Run Post-Installation Script

To complete the post-installation tasks, you can run the post-installation script.



Important

- Ensure that you run *hx_post_install* and confirm network operation immediately following the deployment of the HyperFlex System.
-

1. Use an SSH client to connect to cluster virtual IP using admin login.
2. Type `hx_post_install` and hit `Enter`.
3. Set the post-install script parameters as specified in the following table:



Note If you run into any post-install script issues, set the post-install script parameters manually.

Parameter	Description
Enable HA/DRS on cluster?	Enables vSphere High Availability (HA) feature per best practice.
Disable SSH warning?	Suppresses the SSH and shell warnings in the vCenter.
Add vMotion interfaces	Configure vMotion interfaces per best practice. Requires <i>IP address</i> and <i>VLAN ID</i> input.
Add VM network VLANs	Add additional guest VLANs to Cisco UCS Manager and within ESXi on all cluster hosts.

4. Correct network errors reported, if any.

Sample Post-Install Script: Option 1. New/Existing Cluster

```
admin@SpringpathController:~$ hx_post_install

Select hx_post_install workflow-

1. New/Existing Cluster
2. Expanded Cluster (for non-edge clusters)
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 HX-01-cmip.example.com
HX CVM admin password:
Getting ESX hosts from HX cluster...
vCenter URL: 192.168.202.35
Enter vCenter username (user@domain): administrator@vsphere.local
vCenter Password:
Found datacenter HX-Clusters
Found cluster HX-01

post_install to be run for the following hosts:
HX-01-esxi-01.example.com
HX-01-esxi-02.example.com
HX-01-esxi-03.example.com

Enter ESX root password:

Enter vSphere license key? (y/n) n

Enable HA/DRS on cluster? (y/n) y
Successfully completed configuring cluster HA.

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

```

Add vmotion interfaces? (y/n) y
Netmask for vMotion: 255.255.254.0
VLAN ID: (0-4096) 208
vMotion MTU is set to use jumbo frames (9000 bytes). Do you want to change to 1500 bytes?
(y/n) y
vMotion IP for HX-01-esxi-01.example.com: 192.168.208.17
Adding vmotion-208 to HX-01-esxi-01.example.com
Adding vmkernel to HX-01-esxi-01.example.com
vMotion IP for HX-01-esxi-02.example.com: 192.168.208.18
Adding vmotion-208 to HX-01-esxi-02.example.com
Adding vmkernel to HX-01-esxi-02.example.com
vMotion IP for HX-01-esxi-03.example.com: 192.168.208.19
Adding vmotion-208 to HX-01-esxi-03.example.com
Adding vmkernel to HX-01-esxi-03.example.com

Add VM network VLANs? (y/n) y
Attempting to find UCSM IP
Found UCSM 10.75.61.254, logging with username admin. Org is HX-Cluster
UCSM Password:
Port Group Name to add (VLAN ID will be appended to the name): USERS
VLAN ID: (0-4096) 1219
Adding VLAN 1219 to FI
Adding VLAN 1219 to vm-network-a VNIC template
Adding USERS-1219 to HX-01-esxi-01.example.com
Adding USERS-1219 to HX-01-esxi-02.example.com
Adding USERS-1219 to HX-01-esxi-03.example.com
Add additional VM network VLANs? (y/n) n

Run health check? (y/n) y

Validating cluster health and configuration...

Cluster Summary:
Version - 4.5.1a-39020
Model - HXAF220C-M5SX
Health - HEALTHY
ASUP enabled - False
admin@SpringpathController:~$

```

Sample Post-Install Script: Option 3. Generate Certificate

```

admin@SpringpathController:~$ hx_post_install

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: 3
Certificate generation workflow selected

Logging in to controller 10.20.1.64
HX CVM admin password:
Getting ESX hosts from HX cluster...

Select Certificate Generation Workflow-

1. With vCenter
2. Without vCenter

```

```
Selection: 1
vCenter URL: 10.33.16.40
Enter vCenter username (user@domain): administrator@vsphere.local
vCenter Password:
Starting certificate generation and re-registration.
Trying to retrieve vCenterDatacenter information ....
Trying to retrieve vCenterCluster information ....
Certificate generated successfully.
Cluster re-registration in progress ....
Cluster re-registered successfully.
admin@SpringpathController:~$
```

Sample Network Errors

```
Host: esx-hx-5.cpoc-rtp.cisco.com
No errors found
```

```
Host: esx-hx-6.cpoc-rtp.clsco.com
No errors found
```

```
Host: esx-hx-1.cpoc-rtp.cisco.com
No errors found
```

```
Host: esx-hx-2.cpoc-rtp.cisco.com
No errors found
```

```
controller VM clocks:
stctlVM-FCH1946V34Y - 2016-09-16 22:34:04
stctlVM-FCH1946V23M - 2016-09-16 22:34:04
stctlVM-FCH1951V2TT - 2016-09-16 22:34:04
stctlVM-FCH2004VINS - 2016-09-16 22:34:04
```

```
Cluster:
Version - 1.8.1a-19499
Model - HX220C-M4S
Health - HEALTHY
Access policy - LENIENT
ASUP enabled - False
SMTP server - smtp.cisco.com
```

Changing ESXi Host Root Password

You can change the default ESXi password for the following scenarios:

- During creation of a standard and stretch cluster (supports only converged nodes)
- During expansion of a standard cluster (supports both converged or compute node expansion)
- During Edge cluster creation



Note In the above cases, the ESXi root password is secured as soon as installation is complete. In the event a subsequent password change is required, the procedure outlined below may be used after installation to manually change the root password.

As the ESXi comes up with the factory default password, you should change the password for security reasons. To change the default ESXi root password post-installation, do the following.



Note If you have forgotten the ESXi root password, for password recovery please contact Cisco TAC.

Step 1 Log into the ESXi host service control using SSH.

Step 2 Acquire root privileges.

```
su -
```

Step 3 Enter the current root password.

Step 4 Change the root password.

```
passwd root
```

Step 5 Enter the new password, and press **Enter**. Enter the password a second time for confirmation.

Note If the password entered the second time does not match, you must start over.

Changing Storage Controller Password

To reset the HyperFlex storage controller password post-installation, do the following.

Step 1 Log into a storage controller VM.

Step 2 Change the Cisco HyperFlex storage controller password.

```
# hxcli security password set
```

This command applies the change to all the controller VMs in the storage cluster.

Note If you add new compute nodes and try to reset the cluster password using the `hxcli security password set` command, the converged nodes get updated, but the compute nodes may still have the default password.

Step 3 Type the `new password`.

Step 4 Press **Enter**.

Cisco HyperFlex HTML Plugin for VMware vCenter

The Cisco HyperFlex vCenter Plugin is integrated with the vSphere Web Client and supports all of the HX Data Platform post-installation management and monitoring functions. For the complete installation and usage information. See the *Cisco HyperFlex HTML Plugin for VMware vCenter* chapter of the [Cisco HyperFlex Data Platform Administration Guide](#) for your release.

Add Datastores in the Storage Cluster

A new HyperFlex cluster has no default datastores configured for virtual machine storage, so the datastores must be created using VMware vSphere Web Client.



Note A minimum of two datastores is recommended for high availability.

-
- Step 1** From the vSphere Web Client Navigator, **Global Inventory Lists** expand **Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores**.
 - Step 2** Click the **Create Datastore** icon.
 - Step 3** Enter a **Name** for the datastore. The vSphere Web Client enforces a 42 character limit for the datastore name. Assign each datastore a unique name.
 - Step 4** Specify the **Size** for the datastore. Choose **GB** or **TB** from the drop-down list. Click **OK**.
 - Step 5** Click the **Refresh** button to display your new datastore.
 - Step 6** Click the **Hosts** tab to see the **Mount Status** of the new datastore.
-

Set HA Heartbeat

Under the vSphere HA settings, ensure that you set the Datastore for Heartbeating option to allow selecting any datastore from the list of available datastores.

-
- Step 1** Login to vSphere.
 - Step 2** Verify DRS is enabled.
From vSphere **Home > Hosts and Clusters > cluster > ConfigureServices**. Click **vSphere DRS**.
 - Step 3** Select the **Edit** button. Click **vSphere HA**. Click **Edit**.
 - Step 4** Select **Turn on vSphere HA** if it is not selected.
 - Step 5** Expand **Admission Control > Define Failover capacity by > Cluster resource percentage** from the drop-down menu. You may use the default value or enable **Override calculated failover capacity** and enter a percentage.
 - Step 6** Expand **Heartbeat Datastores** and select **Use datastore only from the specified list**. Select which datastores to include.
 - Step 7** Click **OK**.
-

Auto Support and Smart Call Home for HyperFlex

You can configure the HX storage cluster to send automated email notifications regarding documented events. You can use the data collected in the notifications to help troubleshoot issues in your HX storage cluster.



Note Auto Support (ASUP) and Smart Call Home (SCH) support the use of a proxy server. You can enable the use of a proxy server and configure proxy settings for both using HX Connect.

Auto Support (ASUP)

Auto Support is the alert notification service provided through HX Data Platform. If you enable Auto Support, notifications are sent from HX Data Platform to designated email addresses or email aliases that you want to receive the notifications. Typically, Auto Support is configured during HX storage cluster creation by configuring the SMTP mail server and adding email recipients.



Note Only unauthenticated SMTP is supported for ASUP.

If the **Enable Auto Support** check box was not selected during configuration, Auto Support can be enabled post-cluster creation using the following methods:

Post-Cluster ASUP Configuration Method	Associated Topic
HX Connect user interface	Configuring Auto Support Using HX Connect, on page 117
Command Line Interface (CLI)	Configuring Notification Settings Using CLI, on page 118
REST APIs	Cisco HyperFlex Support REST APIs on Cisco DevNet .

Auto Support can also be used to connect your HX storage cluster to monitoring tools.

Smart Call Home (SCH)

Smart Call Home is an automated support capability that monitors your HX storage clusters and then flags issues and initiates resolution before your business operations are affected. This results in higher network availability and increased operational efficiency.

Call Home is a product feature embedded in the operating system of Cisco devices that detects and notifies the user of a variety of fault conditions and critical system events. Smart Call Home adds automation and convenience features to enhance basic Call Home functionality. After Smart Call Home is enabled, Call Home messages/alerts are sent to Smart Call Home.

Smart Call Home is included with many Cisco service contracts and includes:

- Automated, around-the-clock device monitoring, proactive diagnostics, real-time email alerts, service ticket notifications, and remediation recommendations.
- Proactive messaging sent to your designated contacts by capturing and processing Call Home diagnostics and inventory alarms. These email messages contain links to the Smart Call Home portal and the TAC case if one was automatically created.
- Expedited support from the Cisco Technical Assistance Center (TAC). With Smart Call Home, if an alert is critical enough, a TAC case is automatically generated and routed to the appropriate support team through `https`, with debug and other CLI output attached.
- Customized status reports and performance analysis.

- Web-based access to all Call Home messages, diagnostics, and recommendations for remediation in one place; TAC case status; and up-to-date inventory and configuration information for all Call Home devices.

To ensure automatic communication among your HX storage cluster, you, and Support, see [Configuring Smart Call Home for Data Collection](#), on page 118.

Configuring Auto Support Using HX Connect

Typically, Auto Support (ASUP) is configured during HX storage cluster creation. If it was not, you can enable it post cluster creation using the HX Connect user interface.

Step 1 Log into HX Connect.

Step 2 In the banner, click **Edit settings (gear icon)** > **Auto Support Settings** and fill in the following fields.

UI Element	Essential Information
Enable Auto Support (Recommended) check box	Configures Call home for this HX storage cluster by enabling: <ul style="list-style-type: none"> • Data delivery to Cisco TAC for analysis. • Notifications from Support as part of proactive support.
Send service ticket notifications to field	Enter the email address that you want to receive the notifications.
Terms and Conditions check box	End user usage agreement. The check box must be checked to use the Auto-Support feature.
Use Proxy Server check box	<ul style="list-style-type: none"> • Web Proxy Server url • Port • Username • Password

Step 3 Click **OK**.

Step 4 In the banner, click **Edit settings (gear icon)** > **Notifications Settings** and fill in the following fields.

UI Element	Essential Information
Send email notifications for alarms check box	If checked, fill in the following fields: <ul style="list-style-type: none"> • Mail Server Address • From Address—Type the email address used to identify your HX storage cluster in the Support Service Tickets, and as the sender for Auto Support Notifications. Support information is not sent to this email address. • Recipient List (Comma separated)

Step 5 Click **OK**.

Configuring Notification Settings Using CLI

Use the following procedure to configure and verify that you are set up to receive alarm notifications from your HX storage cluster.



Note Only unauthenticated SMTP is supported for ASUP.

Step 1 Log into a storage controller VM in your HX storage cluster using `ssh`.

Step 2 Configure the SMTP mail server, then verify the configuration.

Email address used by the SMTP mail server to send email notifications to designated recipients.

Syntax: `stcli services smtp set [-h] --smtp SMTPSERVER --fromaddress FROMADDRESS`

Example:

```
# stcli services smtp set --smtp mailhost.eng.mycompany.com --fromaddress smtpnotice@mycompany.com
# stcli services smtp show
```

Step 3 Enable ASUP notifications.

```
# hxcli services asup enable
```

Step 4 Add recipient email addresses, then verify the configuration.

List of email addresses or email aliases to receive email notifications. Separate multiple emails with a space.

Syntax: `hxcli services asup recipients add --recipients RECIPIENTS`

Example:

```
# hxcli services asup recipients add --recipients user1@mycompany.com user2@mycompany.com
# hxcli services asup show
```

Step 5 From the controller VM that owns the eth1:0 IP address for the HX storage cluster, send a test ASUP notification to your email.

```
# sendasup -t
```

To determine the node that owns the eth1:0 IP address, log into each storage controller VM in your HX storage cluster using `ssh` and run the `ifconfig` command. Running the `sendasup` command from any other node does not return any output and tests are not received by recipients.

Step 6 Configure your email server to allow email to be sent from the IP address of all the storage controller VMs.

Configuring Smart Call Home for Data Collection

Data collection is enabled by default but, you can opt-out (disable) during installation. You can also enable data collection post cluster creation. During an upgrade, Smart Call Home enablement is determined by your

legacy configuration. For example, if `hxcli services asup show` as enabled, Smart Call Home is enabled on upgrade.

Data collection about your HX storage cluster is forwarded to Cisco TAC through `https`. If you have a firewall installed, configuring a proxy server for Smart Call Home is completed after cluster creation.



Note Smart Call Home does not support the use of a proxy server in deployments where outgoing connections from an HX cluster require to go through a proxy server.

Using Smart Call Home requires the following:

- A Cisco.com ID associated with a corresponding Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service contract for your company.
- Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service for the device to be registered.

Step 1 Log into a storage controller VM in your HX storage cluster.

Step 2 Register your HX storage cluster with Support.

Registering your HX storage cluster adds identification to the collected data and automatically enables Smart Call Home. To register your HX storage cluster, you need to specify an email address. After registration, this email address receives support notifications whenever there is an issue and a TAC service request is generated.

Syntax:

```
stcli services sch set [-h] --email EMAILADDRESS
```

Example:

```
# stcli services sch set --email name@company.com
```

Step 3 Verify data flow from your HX storage cluster to Support is operational.

Operational data flow ensures that pertinent information is readily available to help Support troubleshoot any issues that might arise.

Note Contact TAC to verify connectivity.

```
# asupcli [--all] ping
```

`--all` option runs the commands on all the nodes in the HX cluster.

Step 4 (Optional) Configure a proxy server to enable Smart Call Home access through port 443.

If your HX storage cluster is behind a firewall, after cluster creation, you must configure the Smart Call Home proxy server. Support collects data at the url: `https://diag.hyperflex.io:443` endpoint.

a. Clear any existing registration email and proxy settings.

```
# stcli services sch clear
```

b. Set the proxy and registration email.

Syntax:

```
stcli services sch set [-h] --email EMAILADDRESS [--proxy-url PROXYURL]
[--proxy-port PROXYPORT] [--proxy-user PROXYUSER] [--portal-url PORTALURL]
[--enable-proxy ENABLEPROXY]
```

Option	Required or Optional	Description
--email EMAILADDRESS	Required.	Add an email address for someone to receive email from Cisco support. Recommendation is to use a distribution list or alias.
--enable-proxy ENABLEPROXY	Optional.	Explicitly enable or disable use of proxy.
--portal-url PORTALURL	Optional.	Specify an alternative Smart Call Home portal URL, if applicable.
--proxy-url PROXYURL	Optional.	Specify the HTTP or HTTPS proxy URL, if applicable.
--proxy-port PROXYPORT	Optional.	Specify the HTTP or HTTPS proxy port, if applicable.
--proxy-user PROXYUSER	Optional.	Specify the HTTP or HTTPS proxy user, if applicable. Specify the HTTP or HTTPS proxy password, when prompted.

Example:

```
# stcli services sch set
--email name@company.com
--proxy-url www.company.com
--proxy-port 443
--proxy-user admin
--proxy-password adminpassword
```

- c. Ping to verify the proxy server is working and data can flow from your HX storage cluster to the Support location.

Note Contact TAC to verify connectivity.

```
# asupcli [--all] ping
```

--all option runs the command on all the nodes in the HX cluster.

Step 5 Verify Smart Call Home is enabled.

When Smart Call Home configuration is `set`, it is automatically enabled.

```
# stcli services sch show
```

Step 6 Enable Auto Support (ASUP) notifications.

Typically, Auto Support (ASUP) is configured during HX storage cluster creation. If it was not, you can enable it post cluster creation using HX Connect or CLI.

If Smart Call Home is disabled, enable it manually.

```
# stcli services sch enable
```

Replacing Self-Signed with CA-Signed Certificate



Note For Releases 5.0(1x) and earlier, root level access to Controller VM is required to run the following certificate replacement script. Please contact TAC to complete the certificate replacement process. For Releases 5.0(2a) and later, you must access the **diag** user shell and complete the CAPTCHA test. For a description of the process, see the [Diag User Overview](#) in the [Cisco HyperFlex Data Platform Administration Guide, Release 5.0](#).

Import CA certificate is automated through shell script. Generate CSR (certificate signing request) from any CVM, preferably from the CIP node. Only one CSR is required for the cluster as each CVM must be installed with the same certificate. When generating the CSR, you should enter the hostname assigned to the management CIP as the Common Name of the Subject's Distinguished Name.

For example:

```
You are about to be asked to enter information that will be incorporated
into your certificate request.
What you are about to enter is what is called a Distinguished Name or a DN.
There are quite a few fields but you can leave some blank
For some fields there will be a default value,
If you enter '.', the field will be left blank.
-----
Country Name (2 letter code) [AU]:US
State or Province Name (full name) [Some-State]:CA
Locality Name (eg, city) []:San Jose
Organization Name (eg, company) [Internet Widgits Pty Ltd]:Cisco
Organizational Unit Name (eg, section) []:HyperFlex
Common Name (e.g. server FQDN or YOUR name) []:<hostname-cluster-management-IP>
Email Address []:support@cisco.com
```

After you get the CA certificate, import the certificate using the automated script. The script will update the certificate on that CVM only.



Note For cluster expansion, the script has to be run again on the expanded node CVM with the same certification and key files to import the certificate.

After accessing the **diag** shell, take the following steps:

Step 1

Script Location in CVM: /usr/share/springpath/storfs-misc/hx-scripts/

```
diag/usr/share/springpath/storfs-misc/hx-scripts/certificate_import_input.certificate_import_input.sh

run stcli cluster reregister
```

Step 2

In the Controller VM (Pointing to CIP), execute this commands to generate the CSR request.

```
openssl req -nodes -newkey rsa:2048 -keyout /etc/ssl/private/<Host Name of the CVM>.key -out
/etc/ssl/certs/<Host Name of the CVM>.csr
cat /etc/ssl/certs/<host name mapped to the management CIP>.csr - Copy the request to any notepad.

Send the request to CA to generate the certificate
```

Step 3 Once you receive the certificate from CA (.crt files), copy the certificate and key to each CVM.

Step 4 On each CVM, use this script to import the certificate: `./certificate_import_input.sh`.

```
root@SpringpathControllerVUFSTDS58L:/usr/share/springpath/storfs-misc/hx-scripts#
./certificate_import_input.sh
```

Step 5 Enter the path for the key: `/etc/ssl/private/<Host Name of the CVM>.key`.

Step 6 Enter the path for the certificate in certificate format: `<Path to the CA .crt file>`

Note After providing all the inputs, it takes some time to finish the import process.

Step 7 From the CVM pointing to CIP, run `stcli reregister` command to reregister the cluster to vCenter. It is mandatory to reregister the cluster once the certificate is imported.

Replication Pairing

Creating a replication cluster pair is a pre-requisite for setting up VMs for replication. The replication network and at least one datastore must be configured prior to creating the replication pair.

By pairing cluster 1 with cluster 2, you are specifying that all VMs on cluster 1 that are explicitly set up for replication can replicate to cluster 2, and that all VMs on cluster 2 that are explicitly set up for replication can replicate to cluster 1.

By pairing a datastore A on cluster 1 with a datastore B on cluster 2, you are specifying that for any VM on cluster 1 that is set up for replication, if it has files in datastore A, those files will be replicated to datastore B on cluster 2. Similarly, for any VM on cluster 2 that is set up for replication, if it has files in datastore B, those files will be replicated to datastore A on cluster 1.

Pairing is strictly 1-to-1. A cluster can be paired with no more than one other cluster. A datastore on a paired cluster, can be paired with no more than one datastore on the other cluster.

For the detailed procedure on creating, editing, and deleting replication pairs, see the [Cisco HyperFlex Systems Administration Guide](#).

Adding Private VLAN

About Private VLANs

A private VLAN partitions the Layer 2 broadcast domain of a VLAN into subdomains, allowing you to isolate the ports on the switch from each other. A subdomain consists of a primary VLAN and one or more secondary VLANs. A private VLAN domain has only one primary VLAN. Each port in a private VLAN domain is a member of the primary VLAN, and the primary VLAN is the entire private VLAN domain.

Understanding Private VLAN Ports

Table 5: Types of Private VLAN Ports

VLAN Port	Description
Promiscuous Primary VLAN	Belongs to the primary VLAN. Can communicate with all interfaces that belong to those secondary VLANs that are associated to the promiscuous port and associated with the primary VLAN. Those interfaces include the community and isolated host ports. All packets from the secondary VLANs go through this VLAN.
Isolated Secondary VLAN	Host port that belongs to an isolated secondary VLAN. This port has complete isolation from other ports within the same private VLAN domain, except that it can communicate with associated promiscuous ports.
Community Secondary VLAN	Host port that belongs to a community secondary VLAN. Community ports communicate with other ports in the same community VLAN and with associated promiscuous ports.

Following HX deployment, a VM network uses a regular VLAN by default. To use a Private VLAN for the VM network, see the following sections:

- [Configuring a Private VLAN on a VM Network without Existing VMs, on page 123.](#)
- [Configuring a Private VLAN on a VM Network with Existing VMs, on page 124.](#)

Configuring a Private VLAN on a VM Network without Existing VMs

-
- Step 1** To configure a private VLAN on Cisco UCS Manager, see the [Cisco UCS Manager Network Management Guide](#).
 - Step 2** To configure a private VLAN on the upstream switch, see the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#).
 - Step 3** To configure a private VLAN on ESX hosts, see [Configuring Private VLAN on ESX Hosts, on page 123](#).
-

Configuring Private VLAN on ESX Hosts

To configure private VLANs on the ESX hosts do the following:

-
- Step 1** Delete VMNICs on the vSphere Standard Switches from the VMware vSphere Client.
 - Step 2** Create new vSphere Distributed Switch with the VMNICs deleted from the previous step.
 - Step 3** Create promiscuous, isolated, and community VLAN.
-

Configuring a Private VLAN on a VM Network with Existing VMs

- Step 1** To configure a private VLAN on Cisco UCS Manager, see the [Cisco UCS Manager Network Management Guide](#).
- Step 2** To configure a private VLAN on the upstream switch, see the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#).
- Step 3** To configure a private VLAN on ESX hosts, see [Configuring Private VLAN on ESX Hosts, on page 123](#)
- Step 4** Migrate VMs from vSphere standard switch to the newly created vSphere distributed switch.
- Right-click the vCenter Virtual Machine and click **Migrate Virtual Machine Networking**.
 - Choose **source network** and **destination network** from the drop-down list.
 - Click **Next**.
 - Select the **Virtual Machines** that you want to migrate.
 - Click **Finish**.
- Step 5** Change network connection of the network adapter on the VMs to private VLAN.
- Right-click the vCenter Virtual Machine and click **Edit Settings**.
 - Under the **Hardware** tab, select the network adapter you want to modify.
 - Select the **Network Connection** you want to use from the **Network Label** drop-down list.
 - Click **OK**.
-

Deleting VMNICs on the vSphere Standard Switch

- Step 1** Log on to VMware vSphere Client.
- Step 2** Select **Home > Hosts and Clusters**.
- Step 3** Select the ESX host from which you want to delete the VMNIC.
- Step 4** Open the **Configuration** tab.
- Step 5** Click **Networking**.
- Step 6** Select the **switch** you wish to remove a VMNIC from.
- Step 7** Click the **Manage the physical adapters connected to the selected switch** button.
- Step 8** Select the **vmnic** you want to delete and click **Remove**.
- Step 9** Confirm your selection by clicking **Yes**.
- Step 10** Click **Close**.
-

Creating vSphere Distributed Switch

- Step 1** Log on to the VMware vSphere Client.
- Step 2** Select **Home > Networking**.
- Step 3** Right click on the cluster **Distributed Switch > New Distributed Switch**.
- Step 4** In the **Name and Location** dialog box, enter a name for the distributed switch.

- Step 5** In the **Select Version** dialog box, select the distributed switch version that correlates to your version and configuration requirements.
- Step 6** Click **Next**.
- Step 7** In the **Edit Settings** dialog box, specify the following:
- Number of uplink ports
 - **Enable** Network I/O Control.
 - **Create a default port group** should be checked.
 - Enter the default port group **name** in the **Port Group Name** box.
- Step 8** Click **Next**.
- Step 9** Review the settings in the **Ready to complete** dialog box.
- Step 10** Click **Finish**.
-

Creating Private VLANs on vSphere Distributed Switch

- Step 1** From the VMware vSphere Client, select **Inventory > Networking**.
- Step 2** Right-click on the dvSwitch.
- Step 3** Click **Edit Settings**.
- Step 4** Select the **Private VLAN** tab.
- Step 5** On the **Primary private VLAN ID** tab, enter a **private VLAN ID**.
- Step 6** On the **Secondary private VLAN ID** tab, enter a **private VLAN ID**.
- Step 7** Select the type of VLAN from the **Type** drop-down list. Valid values include:
- **Isolated**
 - **Community**
 - **Promiscuous** (Default)
- Step 8** Click **OK**.
-

Set Private VLAN in Distributed Port Group

Before you begin

Create Private VLAN on the vSphere Distribute Switch.

- Step 1** Right click **dvPortGroup** under **dvSwitch**, and click **Edit Settings**.
- Step 2** Click **Policies > VLAN**.
- Step 3** Select **Private VLAN**, from the **VLAN type** drop-down list.

Step 4 From the **Private VLAN Entry** drop-down list, select the type of private VLAN. It can be one of the following:

- **Isolated**
- **Community**

Note Community private VLAN is recommended.
Promiscuous ports are not supported

Step 5 Click **OK**.

Distributed Virtual Switches and Cisco Nexus 1000v

Considerations when Deploying Distributed Switches



Note

- Using Distributed Virtual Switches (DVS) or Cisco Nexus 1000v (N1Kv) is an optional and not a required step.
 - DVS for your vMotion network is available only if your environment has Enterprise Plus License for vSphere.
 - You can use only one of the two switches at a given time.
 - There may be a potential conflict between the Quality of Service (QoS) policy for HyperFlex and Nexus 1000v. Make sure that the QoS classes for N1Kv are set as per the HyperFlex policy. See *Creating a QoS Policy*, in the [Network and Storage Management Guide](#).
 - If you choose to deploy N1Kv switch, apply the settings as described, so that the traffic between the HyperFlex hosts flows locally on the FIs in a steady state. If not configured accurately, it could lead to a situation where most traffic will go through the upstream switches leading to latency. In order to avoid that scenario, ensure that the Storage Controller, Management Network, and vMotion port groups are configured with active/standby and failover enabled.
1. Set the **link status** for the **Network Control Policy** using UCS Manager. For details, see the "Configuring Network Control Policy" section in the [Cisco UCS Manager GUI Configuration Guide](#).
 2. Set the vSwitch properties in vCenter.
 - a. Set the **Network Failure Detection** to **Link Status only**.
 - b. Set **Failback** to **Yes**. For details, see the "Configuring the VM-FEX for VMware" section in the [Cisco UCS Manager VM-FEX for VMware Configuration guide](#)

Distributed switches ensure that each node is using the same configuration. It helps prioritize traffic and allows other network streams to utilize available bandwidth when no vMotion traffic is active.

The HyperFlex (HX) Data Platform can use Distributed Virtual Switch (DVS) Networks for non-HyperFlex dependent networks.

These non-HX dependent networks include:

- VMware vMotion networks
- VMware applications networks

The HX Data Platform has dependency that the following networks use standard vSwitches.

- vswitch-hx-inband-mgmt: Storage Controller Management Network
- vswitch-hx-inband-mgmt: Management Network
- vswitch-hx-storage-data: Storage Hypervisor Data Network
- vswitch-hx-storage-data: Storage Controller Data Network

During HX Data Platform installation, all the networks are configured with standard vSwitch networks. After the storage cluster is configured, the non-HX dependent networks can be migrated to DVS networks. For example:

- vswitch-hx-vm-network: VM Network
- vmotion: vmotion pg

For further details on how to migrate the vMotion Network to Distributed Virtual Switches, please see the *Migrating vMotion Networks to Distributed Virtual Switches (DVS) or Cisco Nexus 1000v (N1Kv)* in the [Network and Storage Management Guide](#).

Hosting vCenter on the HX Data Platform

Deployment of vCenter on the HyperFlex cluster is supported with some constraints. See the [How to Deploy vCenter on the HX Data Platform](#) TechNote for more details.

Deploying AMD GPUs

AMD FirePro S7150 series GPUs are supported in HX240c M5/M6 nodes. These graphic accelerators enable highly secure, high performance, and cost-effective VDI deployments. Follow the steps below to deploy AMD GPUs in HyperFlex.

Step	Action	Step Instructions
1	For the service profiles attached to the servers modify the BIOS policy.	Requirement For All Supported GPUs: Memory-Mapped I/O Greater than 4 GB
2	Install the GPU card in the servers.	GPU Card Installation
3	Power on the servers, and ensure that the GPUs are visible in the Cisco UCS Manager inventory for the servers.	—

Step	Action	Step Instructions
4	Install the vSphere Installation Bundle (VIB) for the AMD GPU card and reboot.	Download the inventory list from Cisco Software Downloads that includes the latest driver ISO for C-series standalone firmware / software version bundle 3.1(3) for AMD on VMware ESXi .
5	Create a Win10 VM on the cluster with the VM configuration.	Specifying Eligible Virtual Machines
6	On each ESXi hosts run the <code>MxGPU.sh</code> script to configure the GPUs and to create virtual functions from the GPU.	Using the MxGPU Setup Script
7	Assign the virtual functions (VFs) created in the previous step to the Win10 VMs.	—



CHAPTER 9

Setting Up Multiple HX Clusters

- [Setting Up Multiple Clusters, on page 129](#)

Setting Up Multiple Clusters

Multiple HyperFlex clusters may coexist under the same UCS domain (pair of fabric interconnects). The following guidelines should be followed to ensure smooth ongoing operations for all equipment connected to the domain.



Note Having HX clusters with two different HX releases on the same UCSM domain is a supported configuration as long as the UCSM infrastructure release supports the needed server firmware bundles. HXDP releases are mapped to UCSM server firmware bundles per the release notes. The release of UCSM infrastructure is independent.

Recommendations

- The number of compute-only nodes must be less than or equal to the number of converged nodes when using the standard HXDP license. The enterprise HXDP licenses allows a maximum 2:1 ratio of compute to converged nodes.
- To create a new HyperFlex cluster, ensure that all the requirements listed in Chapter 2 are met. In addition, follow the process as outlined in Chapter 4.
- All nodes in an HX cluster should reference the same policies and service profiles.



Note Editing the HyperFlex Service Profile templates is not recommended.

- Assign a unique name for each HX cluster.
- Each HX cluster will be created in a unique sub-org as part of installation. Do not modify this hierarchy as it ensures unique policies are created per cluster.
- Each cluster should use a unique storage data VLAN to keep all storage traffic isolated. Reuse of this VLAN across multiple clusters is highly discouraged.

- Cisco requires unique storage data VLAN for each Fabric Interconnect attached cluster, including Stretch Clusters. Use of a shared VLAN for storage data in such deployments is not supported, as it can lead to cluster outages.

For other deployment types, it is strongly recommended to use a unique storage data VLAN for each cluster to keep all storage traffic isolated. Ensure proper network isolation if you are using the same storage data VLAN for multiple clusters. Using the same storage data VLAN on multiple clusters without proper network isolation is not supported.

- When reusing VLANs (e.g. management and guest traffic VLANs), create a new unique VLAN name for each VLAN even if it already exists in UCSM. This will ensure no disruption to other clusters and servers in that domain.
- Ensure that you select the compatible Cisco UCS Manager and Cisco HyperFlex releases.

For the latest compatibility matrix, refer to the *Software Versions* table in the latest [Release Notes for Cisco HX Data Platform](#).

- Ensure that you clear browser cache before setting up a second cluster on a different vCenter using the same Cisco HX Data Platform Installer. This avoids issues such as caching the older cluster's IP address which could result in deployment failure.



Note You may need to add more uplink ports depending on your traffic.



Note It is possible to use the same VLANs across multiple clusters connected to the same pair of Fabric Interconnects, for example - Management, vMotion and VM guest VLANs. This is possible as long as you do not overlap IPs. It is, however, recommended to keep the HX storage VLANs different per cluster to ensure storage traffic is secure and isolated. If deciding to reuse the storage VLAN against best practices, be extremely vigilant to avoid duplicate IPs. A duplicate IP can disrupt existing storage traffic on another cluster.

Step 1 Log into the HX Data Platform Installer.

Step 2 Follow the **Create Cluster** workflow for a standard cluster to create additional clusters. See [Configure Cisco HyperFlex Systems, on page 47](#) for more details.



CHAPTER 10

Expand Cisco HyperFlex System Clusters

- [Cluster Expansion Guidelines, on page 131](#)
- [Prerequisites When Expanding M5/M6 Clusters, on page 132](#)
- [Mixed Cluster Expansion Guidelines - Cisco HX Release 5.5\(x\) , on page 133](#)
- [Steps During Mixed Cluster Expansion, on page 133](#)
- [Prerequisites for Adding a Converged Node, on page 134](#)
- [Preparing a Converged Node, on page 135](#)
- [Adding a Converged Node to an Existing Cluster, on page 135](#)
- [Prerequisites for Adding a Compute-Only Node, on page 140](#)
- [Preparing a Compute-Only Node, on page 141](#)
- [Adding a Compute-Only Node to an Existing Cluster, on page 143](#)
- [Resolving Failure of Cluster Expansion, on page 148](#)
- [Logical Availability Zones, on page 148](#)

Cluster Expansion Guidelines

Please review these guidelines before expanding your cluster.



Note If you have LAZ configured (enabled by default for clusters of size 8 or more), please review [Logical Availability Zones, on page 148](#) prior to moving ahead with expansion.

- If you have replication configured, put replication in pause mode before performing upgrade, expansion or cluster maintenance. After the upgrade, expansion or cluster maintenance is completed, then resume replication. Perform the pause and resume on any cluster that has replication configured to or from this local cluster.
- If you are using RESTful APIs to perform cluster expansion, sometimes the task may take longer than expected.
- ESXi installation is supported on SD cards for M4 converged nodes and M.2 SATA SSD for M5/M6 converged nodes. For compute-only nodes, ESXi installation is supported for SD Cards, SAN boot, front SSD/HDD, or single M.2 SSD (using UCS-MSTOR-M2 controller). Installing ESXi on USB Flash is not supported for compute-only nodes



Note HW RAID M.2 (UCS-M2-HWRAID and HX-M2-HWRAID) is a supported boot configuration starting with HX Data Platform release 4.5(1a) and later.

- You must click on the discovered cluster to proceed with expanding a standard ESX cluster. Not doing so results in errors.
- Use only Admin credentials for the Controller VM during expansion workflow. Using any other credentials other than Admin may cause the expansion to fail.
- In the event you see an error about unsupported drives or catalog upgrade, see the [Compatibility Catalog](#).
- Starting with HX Release 5.0(1b) and later, you can expand ESXi based 10/25 GbE HyperFlex Edge clusters with 2 nodes via Intersight.

Please refer to the Intersight documentation for all requirements: [Cluster Expansion Requirements](#).

- Starting with HX Release 5.0(2b) you can not add new nodes with 375G WL cache drives to an existing cluster with nodes that have 1.6TB cache drives.
- Moving operational disks between servers within same cluster or moving them into expansion nodes within the same active cluster is not supported.

ESXi Installation Guidelines

1. Modify boot policy for compute node.

To modify the template and boot policy for HyperFlex Stretched Cluster compute only node on M5 server:

 - a. Clone the template.
 - b. Uncheck the Flex flash from local boot policy, if the compute M5 node does not have flash cards.
 - c. Add the SAN boot with proper WWPN to the boot order.
2. Start the DPI expansion workflow.
3. When prompted, install ESXi using an ISO image.
4. Return to the DPI expansion workflow and complete the ESXi installation workflow.



Note If the Hypervisor configuration fails with the SOL logging failure message, access the installer CLI through SSH with root and default password and configure the ESXi hypervisor. Then, run the advanced installer and check the **HX Storage Software** and **Expand Cluster** check boxes to proceed with the ESXi installation process.

Prerequisites When Expanding M5/M6 Clusters

Prior to beginning cluster expansion in M5/M6 clusters, perform the following tasks:

- **Hypercheck Health Check Utility**— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and will help ensure a seamless upgrade experience. For more information, see the [Hyperflex Health & Pre-Upgrade Check Tool](#) TechNote for full instructions on how to install and run Hypercheck.
- Upgrade the HX cluster and UCS Manager to the appropriate recommended release for your deployment. For more information, see the [Cisco HyperFlex Recommended Software Release and Requirements Guide](#).
- Download and deploy the matching HX Data Platform Installer (release should be same as cluster) release to run the expansion workflow.

Mixed Cluster Expansion Guidelines - Cisco HX Release 5.5(x)

General Guidelines:

- HX240c M6 is not able to use the additional slots if combined in a cluster with M5 nodes.
- All servers must match the form factor (220/240), type (Hybrid/AF/NVME), security capability (Non-SED only) and disk configuration (QTY, capacity, and non-SED) across the cluster.

Mixed Cluster Expansion Options: Supported

- Expanding existing M5 clusters with M6 converged nodes is supported.
- Expanding existing mixed M5/M6 cluster with M5 or M6 converged nodes is supported.
- Only expansion workflow is supported to create a mixed cluster (Initial cluster creation with mixed M5/M6 servers is not supported).
- Adding any supported compute-only nodes is permitted with all M5, M6 and mixed M5/M6 clusters using the HX Data Platform 5.0 or later Installer. Some example combinations are listed here, many other combinations are possible.

Mixed Cluster Expansion Options: Not Supported

- Expanding existing M6 cluster with M5 converged nodes is NOT supported.
- Mixing Intel and AMD M6 is not supported.
- HX Edge does not support mixed M5 and M6 clusters.
- Initial cluster creation with mixed M5/M6 servers is not supported.

Steps During Mixed Cluster Expansion

- During the validation steps, before expansion begins, an EVC check is performed. Follow the displayed guidance to manually enable EVC mode on the existing cluster at this time.



Caution Failure to enable EVC at the time of the warning will require a complete shutdown of the storage cluster and all associated VMs at a later point in time. Do not skip this warning.

- Perform the EVC mode configuration in vCenter and then retry the validation.
- Cluster expansion will then validate a second time and then continue with the expansion.

Prerequisites for Adding a Converged Node

A converged node can be added to a HyperFlex cluster after cluster creation. The storage on a converged node is automatically added to the cluster's storage capacity.

Before you start adding a converged node to an existing storage cluster, make sure that the following prerequisites are met.

- Ensure that the storage cluster state is healthy.
- Ensure that the new node meets the system requirements listed under **Installation Prerequisites**, including network and disk requirements.
- Ensure that the new node uses the same configuration as the other nodes in the storage cluster. This includes VLAN IDs and switch types (whether vSwitches), VLAN tagging with External Switch VLAN Tagging (EST), VLAN tagging with Virtual Switch Tagging (VST), or Virtual Distributed Switch.



Note If the storage cluster is in an out of space condition, when you add a new node, the system automatically rebalances the storage cluster. This is in addition to the rebalancing that is performed every 24 hours.

- Ensure that the node you add is of the same model (HX220 or HX240) type (Hybrid, All Flash or NVME), and disk configuration (SED or non-SED). In addition, ensure that the number of capacity disks matches the existing cluster nodes.
- To add a node that has a different CPU family from what is already in use in the HyperFlex cluster, enable EVC. For more details, see the *Setting up Clusters with Mixed CPUs* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Ensure that the software version on the node matches the Cisco HX Data Platform release, the ESXi version, and the vCenter version. To identify the software version, go to the Storage Cluster Summary tab in vCenter and check the HX Data Platform release in the top section. Upgrade if necessary.



Note If you upgraded the cluster, you must download and install a new installer VM, that matches the current release of HXDP running on the cluster.

- Ensure that the new node has at least one valid DNS and NTP server configured.

- If you are using SSO or Auto Support, ensure that the node is configured for SSO and SMTP services.
- Allow ICMP for ping between the HX Data Platform Installer and the existing cluster management IP address.

Preparing a Converged Node

Step 1 Connect the converged node to the hardware and the network of the existing storage cluster.

Step 2 Ensure that the HX node is a node prepared at factory.

Note Do not reuse a removed converged node or its disks in the original cluster.

Adding a Converged Node to an Existing Cluster



Note If you are using RESTful APIs to perform cluster expansion, the task may take longer than expected.

Step 1 Launch the Cisco HX Data Platform Installer.

- a) In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. The Cisco HX Data Platform Installer login page appears. Verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.
- b) In the login page, enter the following credentials:

Username: `root`

Password (Default): `Cisco123`

Note Systems ship with a default password of `Cisco123` that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

- c) Read the EULA, check the **I accept the terms and conditions** checkbox, and click **Login**.

Step 2 On the **Workflow** page, select **Cluster Expansion**.

Step 3 On the **Credentials** page, complete the following fields.

To perform cluster creation, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

Note For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.

- a. Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
- b. An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Field	Description
UCS Manager Credentials	
UCS Manager Host Name	UCS Manager FQDN or IP address. For example, <i>10.193.211.120</i> .
User Name	<admin> username.
Password	<admin> password.
vCenter Credentials	
vCenter Server	vCenter server FQDN or IP address. For example, <i>10.193.211.120</i> . Note <ul style="list-style-type: none"> • A vCenter server is required before the cluster can be made operational. • The vCenter address and credentials must have root level administrator permissions to the vCenter. • vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.
User Name	<admin> username. For example, <i>administrator@vsphere.local</i> .
Admin Password	<root> password.
Hypervisor Credentials	
Admin User Name	<admin> username. This is root for factory nodes.
Admin Password	<root> password. Default password is <code>Cisco123</code> for factory nodes. Note Systems ship with a default password of <code>Cisco123</code> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Step 4 Click **Continue**. A **Cluster Expand Configuration** page is displayed. Select the *HX Cluster* that you want to expand. If the HX cluster to be expanded is not found, or if loading the cluster takes time, enter the IP of the Cluster Management Address in the **Management IP Address** field.

Step 5 The **Server Selection** page displays a list of unassociated HX servers under the **Unassociated** tab, and the list of discovered servers under the **Associated** tab. Select the servers under the **Unassociated** tab to include in the HyperFlex cluster.

If HX servers do not appear in this list, check Cisco UCS Manager and ensure that they have been discovered.

For each server you can use the **Actions** drop-down list to set the following:

- **Launch KVM Console**—Choose this option to launch the KVM Console directly from the HX Data Platform Installer.
- **Disassociate Server**—Choose this option to remove a service profile from that server.

Note If there are no unassociated servers, the following error message is displayed:

```
No unassociated servers found. Please login to UCS Manager and ensure server ports are enabled.
```

The **Configure Server Ports** button allows you to discover any new HX nodes. Typically, the server ports are configured in Cisco UCS Manager before you start the configuration.

Step 6 Click **Continue**. The **UCSM Configuration** page appears.

Note If you imported a JSON file at the beginning, the **Credentials** page should be populated with the required configuration data from the preexisting HX cluster. This information must match your existing cluster configuration.

Step 7 Click **Continue**. The **Hypervisor Configuration** page appears. Complete the following fields:

Attention You can skip the completion of the fields described in this step in case of a reinstall, and if ESXi networking has been completed.

Field	Description
Configure Common Hypervisor Settings	
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, <i>255.255.0.0</i> .
Gateway	IP address of gateway. For example, <i>10.193.0.1</i> .
DNS Server(s)	IP address for the DNS Server. If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform installer. Use only static IP addresses and hostnames for all ESXi hosts. Note If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.
Hypervisor Settings Ensure to select Make IP Addresses and Hostnames Sequential , to make the IP addresses sequential. Note You can rearrange the servers using drag and drop.	
Name	Server name.
Serial	Serial number of the server.

Field	Description
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.
Hostname	Do not leave the hostname fields empty.

Step 8 Click **Continue**. The **IP Addresses** page appears. You can add more compute or converged servers, by clicking **Add Compute Server** or **Add Converged Server**.

Ensure to select **Make IP Addresses Sequential**, to make the IP addresses sequential. For the IP addresses, specify if the network should belong to Data Network or Management Network.

For each HX node, complete the following fields for Hypervisor Management and Data IP addresses.

Field	Description
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and the storage cluster.
Management Storage Controller	Enter the static IP address that handles the HX Data Platform storage controller VM management network connection between the storage controller VM and the storage cluster.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	Enter the static IP address that handles the HX Data Platform storage controller VM data network connection between the storage controller VM and the storage cluster.

When you enter IP addresses in the first row for Hypervisor (Management), Storage Controller VM (Management), Hypervisor (Data), and Storage Controller VM (Data) columns, the HX Data Platform Installer applies an incremental auto-fill to the node information for the rest of the nodes. The minimum number of nodes in the storage cluster is three. If you have more nodes, use the **Add** button to provide the address information.

Note Compute-only nodes can be added only after the storage cluster is created.

Controller VM Password	<p>A default administrator username and password are applied to the controller VMs. The VMs are installed on all converged and compute-only nodes.</p> <p>Important</p> <ul style="list-style-type: none"> You cannot change the name of the controller VM or the controller VM's datastore. Use the same password for all controller VMs. The use of different passwords is not supported. Provide a complex password that includes 1 uppercase character, 1 digit, 1 special character, and a minimum of 10 characters in total. You can provide a user-defined password for the controller VMs and for the HX cluster to be created. For password character and format limitations, see the section on Guidelines for HX Data Platform Special Characters in the <i>Cisco HX Data Platform Management Guide</i>.
-------------------------------	--

Field	Description
Advanced Configuration	
Jumbo frames Enable Jumbo Frames checkbox	Check to set the MTU size for the storage data network on the host vSwitches and vNICs, and each storage controller VM. The default value is 9000. Note To set your MTU size to a value other than 9000, contact Cisco TAC.
Disk Partitions Clean up Disk Partitions checkbox	Check to remove all existing data and partitions from all nodes added to the storage cluster. You must backup any data that should be retained. Important Do not select this option for factory prepared systems. The disk partitions on factory prepared systems are properly configured. For manually prepared servers, select this option to delete existing data and partitions.

Step 9 Click **Start**. A **Progress** page displays the progress of various configuration tasks.

- Note** If the vCenter cluster has EVC enabled, the deploy process fails with a message: `The host needs to be manually added to vCenter`. To successfully perform the deploy action, do the following:
- Log into the ESXi host to be added in vSphere Client.
 - Power off the controller VM.
 - Add the host to the vCenter cluster in vSphere Client.
 - In the HX Data Platform Installer, click **Retry Deploy**.

Step 10 When cluster expansion is complete, click **Launch HyperFlex Connect** to start managing your storage cluster.

- Note** When you add a node to an existing storage cluster, the cluster continues to have the same HA resiliency as the original storage cluster until auto-rebalancing takes place at the scheduled time.
- Rebalancing is typically scheduled during a 24-hour period, either 2 hours after a node fails or if the storage cluster is out of space.

Step 11 Create the required VM Network port groups and vMotion vmkernel interfaces using HyperFlex `hx_post_install` script or manually to match the other nodes in the cluster.

- a) SSH to HyperFlex cluster management IP.
- b) Log in as the admin user.
- c) Run the `hx_post_install` command.
- d) Follow the on-screen instructions, starting with vMotion and VM network creation. The other configuration steps are optional.

Step 12 After the new nodes are added to the storage cluster the High Availability (HA) services are reset so that HA can recognize the added nodes.

- a) Log into vCenter.

- b) In the vSphere Web Client, navigate to the Host: **Home > vCenter > Inventory Lists > Hosts and Clusters > vCenter > Server > Datacenter > Cluster > Host**
- c) Select the new node.
- d) Right-click and select **Reconfigure for vSphere HA**.

Prerequisites for Adding a Compute-Only Node

You can add a compute-only node to a HyperFlex cluster after cluster creation. It is added to provide extra compute resources. The Cisco UCS server does not need to have any caching or persistent drives as they do not contribute any storage capacity to the cluster.

Before you start adding a compute-only node, make sure that the following prerequisites are met.

- Ensure that the storage cluster state is healthy.
- Ensure that the new node meets the compute-only system requirements listed in *Installation Prerequisites*, including network and disk requirements.
- Install ESXi hypervisor after service profile association.
- Ensure that the new node uses the same configuration as the other nodes in the storage cluster. This includes VLAN IDs and switch types (whether vSwitches), VLAN tagging with External Switch VLAN Tagging (EST), VLAN tagging with Virtual Switch Tagging (VST), or Virtual Distributed Switch.
- Enable EVC if the new node to be added has a different CPU family than what is already used in the HX cluster. For more details, see the *Setting up Clusters with Mixed CPUs* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.
- Ensure that the software release on the node matches the Cisco HX Data Platform release, the ESXi release and the vCenter release. To identify the software release, go to the **Storage Cluster Summary** tab in vCenter and check the *HX Data Platform version* in the top section. Upgrade if necessary.
- Ensure that the new node has at least one valid DNS and NTP server configured.
- If you are using SSO or Auto Support, ensure that the node is configured for SSO and SMTP services.
- Compute-only nodes are deployed with automatic detection and configuration of disk and boot policies based on the boot hardware.

Starting with HX Data Platform release 4.5(1a) and later, compute-only nodes are deployed with automatic detection and configuration of disk and boot policies based on the inventoried boot hardware. Users cannot directly select the UCSM policies. Instead, the boot device is automatically determined based on the first acceptable boot media discovered in the server. The tables below show the priority order for M5/M6 generation servers. Reading from top to bottom, the first entry that is a match based on the inventoried hardware are selected automatically during cluster expansion. For example, when expanding with a B200 compute node with a single M.2 boot SSD, the second rule in the table below is a match and used for SPT association.

If the server is booted using a mechanism not listed (such a SAN boot), the catch-all policy of **anyld** is selected and administrators may subsequently modify the UCSM policies and profiles as needed to boot the server.

Table 6: Priority for M6

Priority for M6			
Priority	SPT Name	Boot Device	Number of disks
1	compute-nodes-m6-m2r1	M6 - M.2 - 2 Disks	2
2	compute-nodes-m6-m2sd	M6 - M.2 - 1 Disk	1
3	compute-nodes-m6-ldr1	MegaRAID Controller	2
4	compute-nodes-m6-anyld	M6 - Generic	Any

Table 7: Priority for M5

Priority for M5			
Priority	SPT Name	Boot Device	Number of disks
1	compute-nodes-m5-m2r1	M.2 Raid	2
2	compute-nodes-m5-m2pch	PCH/Non-RAID M.2	1
3	compute-nodes-m5-sd	FlexFlash	2
4	compute-nodes-m5-ldr1	MegaRAID	2
5	compute-nodes-m5-sd	FlexFlash	1
6	compute-nodes-m5-anyld	Any other config	Any

Preparing a Compute-Only Node

Step 1 Ensure that the server is a supported HX server and meets the requirements. For more details, see the *Host Requirements* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*.

Step 2 Log into Cisco UCS Manager.

- a) Open a browser and enter the Cisco UCS Manager address for the fabric interconnect of the storage cluster network.
- b) Click the **Launch UCS Manager** button.
- c) If prompted, download, install, and accept Java.
- d) Log in with administrator credentials.

Username: admin

Password: <admin password>

Step 3 Locate the server to ensure that the server has been added to the same FI domain as the storage cluster and is an approved compute-only model. Check the latest [Release Notes for Cisco HX Data Platform](#) for a full list of compatible Compute-only nodes.

Verify the HX Data Platform Installer

-
- Step 1** Verify that the HX Data Platform installer is installed on a node that can communicate with all the nodes in the storage cluster and compute nodes that are being added to the storage cluster.
- Step 2** If the HX Data Platform installer is not installed, see Deploy the HX Data Platform Installer.
-

Apply an HX Profile on a Compute-only Node Using UCS Manager

In Cisco UCS Manager Cisco UCS Manager the network policies are grouped into an HX profile. The HX installer handles automatic service profile association for compute-only nodes. Manual association is not required.

Once install begins, you should monitor compute-only node service profile association in UCS Manager. Wait until the server is fully associated before continuing on to install ESXi.

Install VMware ESXi on Compute Nodes



-
- Important** Install VMware ESXi on each compute-only node.
- Install a Cisco HyperFlex Data Platform supported release of ESXi. See the [Cisco HyperFlex Data Platform Release Notes](#) for a list of supported ESXi versions.
- If the compute only node already has ESXi installed, it must be re-imaged with the Cisco HX Custom image.
-

Before you begin

Ensure the required hardware and network settings are met. For more details, see the *Installation Prerequisites* section in the *Cisco HyperFlex Systems Installation Guide for VMware ESXi*. Ensure the service profiles in the previous step have finished associating.

-
- Step 1** Download the *HX Custom Image for ESXi* from the Cisco.com download site for Cisco HyperFlex. See [Download Software](#).
- Select a networked location that can be accessed through Cisco UCS Manager.
- Step 2** Log into Cisco UCS Manager.
- Step 3** Log into the KVM console of the server through Cisco UCS Manager.
- In the Navigation Pane, click **Servers > Service Profiles > Sub-Organizations > hx-cluster**.
 - Right click the *hx-cluster* and choose **KVM Console**.
- Step 4** Copy the *HX-Vmware.iso* image to the KVM path for the compute server.

Example:

HX-ESXi-7.0U3-20328353-Cisco-Custom-7.3.0.10-install-only.iso

Step 5 From the KVM console session, select **Virtual Media > Map CD/DVD** and mount the *HX Custom Image for ESXi* image. If you do not see the **Map CD/DVD** option, first activate virtual devices.

a) Select **Virtual Media > Activate Virtual Devices**.

This opens in a pop-up window.

b) Click **Accept the session > Apply**.

Step 6 From the **Map CD/DVD** option, map to the location of the *HX-Vmware.iso* file.

a) Select the *HX-Vmware.iso* file.

b) Select **Map Device**.

There is a check mark indicating that the file is on a mapped location, once the process is complete. The mapped file's full name includes the ESXi build ID.

Step 7 Reset the compute server.

a) Click the **Reset** button on the KVM console. Click **OK** to confirm.

b) Select **Power Cycle**. Click **OK**.

Step 8 Change the boot path to point to the *HX-Vmware.iso* file.

a) Press **F6**.

b) From the **Enter boot selection** menu, use the arrow keys to highlight the *Cisco vKVM-Mapped vDVD1.22* option.

c) Press **Enter** to select.

This launches the ESXi installer bootloader. Select one of the three compute-only node options based on desired boot type: SD Card, Local Disk, or Remote Disk. Type in **yes** (all lowercase) to confirm selection. The rest of the installation is automated. ESXi will reboot several times. It is normal to see warnings that automatically dismiss after a short wait period. Wait for the *ESXi DCUI* to fully appear, signaling the end of installation.

Step 9 Repeat steps 3 to 8 for each Cisco HyperFlex server.

Step 10 Once ESXi is fully installed, click **continue**. Then click **Retry Hypervisor Configuration** to complete the rest of the cluster expansion.

Adding a Compute-Only Node to an Existing Cluster

To add a HyperFlex compute-only node to an existing HyperFlex system cluster, complete the following steps.



Note If you are using RESTful APIs to perform cluster expansion, sometimes the task may take longer than expected.



Note After you add a compute-only node to an existing cluster, you must manually configure the vmk2 interface for vmotion.

Step 1

Launch the Cisco HX Data Platform Installer.

- a) In your web browser, enter the IP address or the node name for the HX Data Platform Installer VM. Click **Accept** or **Continue** to bypass any SSL certificate errors. The Cisco HX Data Platform Installer login page appears. Verify the HX Data Platform Installer **Build ID** in the lower right corner of the login screen.
- b) In the login page, enter the following credentials:

Username: root

Password (Default): Cisco123

Note Systems ship with a default password of Cisco123 that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

- c) Read the EULA, check the **I accept the terms and conditions** checkbox, and click **Login**.

Step 2

On the **Workflow** page, select **Cluster Expansion**.

Step 3

On the **Credentials** page, complete the following fields.

To perform cluster expansion, you can import a *JSON configuration* file with the required configuration data. The following two steps are optional if importing a JSON file, otherwise you can input data into the required fields manually.

- Note**
- a. Click **Select a file** and choose your *JSON file* to load the configuration. Select **Use Configuration**.
 - b. An **Overwrite Imported Values** dialog box displays if your imported values for Cisco UCS Manager are different. Select **Use Discovered Values**.

Field	Description
UCS Manager Credentials	
UCS Manager Host Name	UCS Manager FQDN or IP address. For example, 10.193.211.120.
User Name	<admin> username.
Password	<admin> password.
vCenter Credentials	
vCenter Server	vCenter server FQDN or IP address. For example, 10.193.211.120. Note <ul style="list-style-type: none"> • A vCenter server is required before the cluster can be made operational. • The vCenter address and credentials must have root level administrator permissions to the vCenter. • vCenter server input is optional if you are building a nested vCenter. See the Nested vCenter TechNote for more details.

Field	Description
User Name	<admin> username. For example, <i>administrator@vsphere.local</i> .
Admin Password	<root> password.
Hypervisor Credentials	
Admin User Name	<admin> username. This is root for factory nodes.
Admin Password	<root> password. Default password is <code>Cisco123</code> for factory nodes. Note Systems ship with a default password of <code>Cisco123</code> that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

Step 4 Click **Continue**. A **Cluster Expand Configuration** page is displayed. Select the *HX Cluster* that you want to expand. If the HX cluster to be expanded is not found, or if loading the cluster takes time, enter the IP of the Cluster Management Address in the **Management IP Address field**.

Step 5 **(M6 Servers Only)** Click **Continue**. A **Server Selection** page is displayed. On the **Server Selection** page, the **Associated** tab lists all the HX servers that are already connected. Do not select them, on the **Unassociated** tab, select the servers you wish to add to the cluster.

Step 6 Click **Continue**. The **Hypervisor Configuration** page appears. Complete the following fields:

Attention You can skip the completion of the fields described in this step in case of a reinstall, and if ESXi networking has been completed.

Field	Description
Configure Common Hypervisor Settings	
Subnet Mask	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, <i>255.255.0.0</i> .
Gateway	IP address of gateway. For example, <i>10.193.0.1</i> .
DNS Server(s)	IP address for the DNS Server. If you do not have a DNS server, do not enter a hostname in any of the fields on the Cluster Configuration page of the HX Data Platform installer. Use only static IP addresses and hostnames for all ESXi hosts. Note If you are providing more than one DNS server, check carefully to ensure that both DNS servers are correctly entered, separated by a comma.

Field	Description
Hypervisor Settings	
Ensure to select Make IP Addresses and Hostnames Sequential , to make the IP addresses sequential.	
Note You can rearrange the servers using drag and drop.	
Name	Server name.
Serial	Serial number of the server.
Static IP Address	Input static IP addresses and hostnames for all ESXi hosts.
Hostname	Do not leave the hostname fields empty.

Step 7

Click **Continue**. An **IP Addresses** page is displayed. Click **Add Compute-only Node** to add a new node. If you are adding more than one compute-only node, select **Make IP Addresses Sequential**.

Field	Information
Management Hypervisor	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and storage cluster.
Management Storage Controller	None.
Data Hypervisor	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller	None.
Controller VM	Enter the default Admin username and password that were applied to controller VMs when they were installed on the existing HX Cluster. Note The name of the controller VM cannot be changed. Use the existing cluster password.

Step 8

Click **Start**. A **Progress** page displays the progress of various configuration tasks.

Note By default no user intervention is required if you are booting from FlexFlash (SD Card). However, if you are setting up your compute-only node to boot from a local disk, complete the following steps in Cisco UCS Manager :

- a. Click the service profile created by the HX Data Platform Installer .
For example, *blade-1(HX_Cluster_Name)*.
- b. On the **General** tab, click **Unbind from the Template**.
- c. In the working pane, click the **Storage** tab. Click the **Local Disk Configuration Policy** sub tab.
- d. In the Actions area, select **Change Local Disk Configuration Policy > Create Local Disk Configuration Policy**.

- e. Under **Create Local Disk Configuration Policy**, enter a name for the policy, and keep the rest as default. Click **Ok**.
- f. In the **Change Local Disk Configuration Policy** Actions area, select the newly created local disk configuration policy from the drop-down list. Click **Ok**.
- g. Now, go back to the HX Data Platform Installer UI and click **Continue**, and then click **Retry UCSM Configuration**.

Compute Node Expansion - ESXi Installation Required


ESXi must be installed on all nodes being added at this point using the HX ESXi ISO on [cisco.com](https://www.cisco.com)

Using an existing installation of ESXi will cause installation to fail. Other ESXi ISOs other than the one posted on Cisco are not supported.

Once ESXi is installed, select Continue and then Retry to continue installation.
Full instructions can be found below.

If ESXi is already installed using the HX ESXi ISO wait for it to boot and then select Continue and Retry to continue installation.

 Instructions

 Launch UCS Manager

Continue

Note If the vCenter cluster has EVC enabled, the deploy process fails, The host needs to be manually added to vCenter. To successfully perform the deploy action, do the following:

- a) Log into the ESXi host to be added in vSphere Client.
- b) Power off the controller VM.
- c) Add the host to the vCenter cluster in vSphere Web Client.
- d) In the HX installer, click **Retry Deploy**.

Step 9 When installation is complete, start managing your storage cluster by clicking **Launch HyperFlex Connect**.

Step 10 After the new nodes are added to the storage cluster, HA services are reset so that HA is able to recognize the added nodes.

- a) Log on to VMware vSphere Client.
- b) Select **Home > Hosts and Clusters > Datacenter > Cluster > Host**.
- c) Select the new node.
- d) Right-click and select **Reconfigure for vSphere HA**.

Step 11 After adding compute-only nodes to an existing cluster, you must manually configure the vmk2 interface for vmotion.

Resolving Failure of Cluster Expansion

If you receive an error dialog box and the storage cluster expansion doesn't complete, proceed with the resolution options listed below:

-
- Step 1** **Edit Configuration** - Returns you to the Cluster Configuration page. You fix the issues listed in the validation page.
- Step 2** **Start Over** - Allows you to reverse the settings you applied by clearing progress table entries and you are returned to the Cluster Configuration page to restart a new deployment. See Technical Assistance Center (TAC).
- Step 3** **Continue** - Adds the node to the storage cluster in spite of the failure generating errors. See Technical Assistance Center (TAC).
- Note** Select the Continue button only if you understand the failures and are willing to accept the possibility of unpredictable behavior.

For more information about cleaning up a node for the purposes of redeploying HyperFlex, see the [HyperFlex Customer Cleanup Guides for FI and Edge](#).

Logical Availability Zones

The Logical Availability Zones (LAZ) feature groups cluster storage nodes in fixed number pools of nodes which enable higher resiliency. The number of zones that can be set automatically or selected manually based on cluster parameters, such as replication factor and cluster size. LAZ is enabled by default on HyperFlex clusters with 8 or more storage nodes. The feature remains enabled through the life cycle of the cluster unless explicitly disabled either at install time or post installation.

Advantages of Logical Availability Zones

Reducing the failure of large clusters in a distributed system is the primary advantage of enabling LAZ on install. In any distributed storage system, when the number of resources in the cluster grow, so does the failure risk. Multiple simultaneous failures could result in permanent data unavailability.

LAZ helps reduce risk of multiple simultaneous component and node failures from causing a catastrophic failure. It does this by grouping resources based on some basic constraints, you can improve the availability from 20% up to 70% in comparison to the same cluster without LAZ. The amount of improvement depends on the cluster replication factor (RF) as well as the number of zones configured. In principle, a cluster with fewer zones and a higher replication factor provides optimal results. Additionally, LAZ saves time by performing maintenance tasks on multiple resources grouped in the same zone, an option not possible in clusters without LAZ.

It is recommended that LAZ be enabled during the HyperFlex cluster installation. Enabling LAZ during install provides optimal cluster performance and data availability. With the guidance of support, LAZ can be enabled or disabled at a later time using the command line interface (CLI). Review the LAZ guidelines before disabling.

Specifying the Number of Zones and Optimizing Balance

The number of zones is set automatically by default and recommended. When you let the installer decide the number of zones, the number of zones is decided based on the number of nodes in the cluster.

To maintain the most balanced consumption of space and data distribution, it is recommended that the number of nodes in a cluster are whole multiple of number of zones, which is either 3, 4, or 5. For example, 8 nodes would evenly divide into 4 zones of 2 servers each, and 9 nodes would divide evenly into 3 zones of 3 servers each. Eleven nodes would create an unbalanced number of nodes across the zones, leading to unbalanced space consumption on the nodes. Users with a need may manually specify 3, 4 or 5 zones.

LAZ Guidelines and Considerations

- HyperFlex clusters determine which nodes participate in each zone. This configuration cannot be modified.
- When changing the number of resources, add or remove an equal number of resources from each configured zone.
- **Cluster Expansion:** Perform expansions in the same increment number of nodes as zones in order to maintain a balanced zone. A balanced zone is when the number of nodes and zones added during install or expansion (or a permanent failure of nodes from zone(s) occurs) are equal. For example, a cluster with 12 nodes and 4 zones is a balanced zone. In this case, it is recommended to add 4 nodes during expansion.
- **Imbalanced Zones:** Zones may become imbalanced when the number of nodes and zones added during install or expansion (or permanent failure of nodes from zone(s)) are not equal. Imbalanced zones can lead to non-optimal performance and are **not** recommended. For example, a cluster with 11 nodes and 4 zones will have 3 nodes per zone except the last zone. In this case, you need to add 1 node to make it balanced. The new node is added automatically to the last zone.
- **Disabling and Re-enabling LAZ:** You can disable and enable LAZ dynamically. It is not recommended to disable and re-enable LAZ in the same cluster with a different number of zones. Doing so could result in an excessive amount of movement and reorganization of data across the cluster - to comply with existing data distribution rules if LAZ is turned on in a cluster already containing data. This can result in the cluster becoming no longer zone compliant for example, if the cluster usage is already greater than 25%.

Viewing LAZ Status and Connections

- To view LAZ information from the HX Connect dashboard, log into HX Connect and use the **System information** and **HyperFlex Connect > Dashboard** menu.
- You can also view LAZ details through CLI by running the `stcli cluster get-zone` command. The following is sample output from the `stcli cluster get-zone` command:

```
stcli cluster get-zone

zones:
-----
pNodes:
-----
state: ready
name: 10.10.18.61
-----
state: ready
name: 10.10.18.59
-----
zoneId: 0000000057eebaab:00000000000000003
numNodes: 2
-----
pNodes:
-----
state: ready
```

```

name: 10.10.18.64
-----
state: ready
name: 10.10.18.65
-----
zoneId: 0000000057eebaab:0000000000000001
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.60
-----
state: ready
name: 10.10.18.63
-----
zoneId: 0000000057eebaab:0000000000000004
numNodes: 2
-----
pNodes:
-----
state: ready
name: 10.10.18.58
-----
state: ready
name: 10.10.18.62
-----
zoneId: 0000000057eebaab:0000000000000002
numNodes: 2
-----
isClusterZoneCompliant: True
zoneType: logical
isZoneEnabled: True
numZones: 4
AboutCluster Time : 08/22/2019 2:31:39 PM PDT

```

LAZ Related Commands

The following STCLI commands are used for LAZ operations. For more information on CLI commands, see the [Cisco HyperFlex Data Platform CLI Guide](#).

Please be advised to wait at least 10 seconds between successive invocations of LAZ disable and LAZ enable operations in that order.

Command	Description
stcli cluster get-zone	Gets the zone details. This option is used to check if the zone is enabled.
stcli cluster set-zone --zone 0	Enables or Disables zones.

Command	Description
<pre>stcli cluster set-zone --zone 1 stcli rebalance start</pre>	<p>(Recommended) Enables and creates zones (default number of zones)</p> <p>Important You must execute the rebalance start command after you enable and create zones.</p> <p>A cluster created without zoning enabled, will become zone compliant only after enabling zoning and successful completion of rebalance.</p> <p>Warning Rebalance is a critical background service. Disabling the service may lead to unexpected behavior including loss of cluster resiliency. Support for this command is limited to Cisco Tech support only. General use is not supported.</p> <p>Triggering rebalance activity may involve large scale data movements across several nodes in the cluster which may decrease the IO performance in the cluster.</p>
<pre>stcli cluster set-zone --zone 1 --numzones <integer-value> stcli rebalance start</pre>	<p>Enables zones and creates a specific number of zones.</p> <p>Important The number of zones can only be 3, 4, or 5.</p> <p>Important You must execute the rebalance start command after you enable and create zones.</p> <p>Warning Rebalance is a critical background service. Disabling the service may lead to unexpected behavior including loss of cluster resiliency. Support for this command is limited to Cisco Tech support only. General use is not supported.</p>



CHAPTER 11

Set Up Clusters with Mixed CPUs

This chapter describes how to add HX nodes with different Intel CPU versions on the same FI.

- [Overview, on page 153](#)
- [Prerequisites for Using Mixed CPUs, on page 153](#)
- [CPU Compatibility with EVC Mode, on page 154](#)
- [Enable Enhanced vMotion Compatibility \(EVC\) on an Existing Cluster, on page 154](#)

Overview

HyperFlex supports Intel v3 and Intel v4 CPUs on the same fabric interconnect. Enable VMware Enhanced vMotion Compatibility (EVC) to migrate virtual machines between Intel v3 and Intel v4 CPUs. After EVC is enabled, all hosts in the HyperFlex cluster are configured to give them a baseline with features of the lower model of CPU. Identical CPU features are exposed to virtual machines regardless of which host they are running on, so virtual machines can be migrated between hosts in a HyperFlex cluster. This ensures CPU compatibility for vMotion even if the hosts have different underlying hardware.

Prerequisites for Using Mixed CPUs

- You can only use CPUs from a single vendor in an EVC-enabled cluster. You cannot add a host from a different vendor into an EVC-enabled cluster.
- When enabling EVC for a cluster with Intel processors of the Xeon E3 or Xeon E5 family that have different revision numbers (v2, v3, or v4), an EVC baseline is required.
- Enable advanced virtualization CPU features in the BIOS if they are available. Otherwise, it can cause problems in enabling EVC, as EVC compatibility checks may fail to detect the features that are expected to be present for a particular CPU.
- Migration of virtual machines using vMotion may fail, even if they are within an EVC cluster under the following scenarios:
 - When a host is not connected to a vCenter Server system.
 - When a host is not configured for vMotion.
 - If the virtual machine does not reside on storage shared by the source and destination hosts.

CPU Compatibility with EVC Mode

To determine the Enhanced vMotion Compatibility (EVC) modes compatible with your CPU, search the [VMware Compatibility Guide](#). Search for the server model or CPU family, and click the entry in the CPU Series column to display the compatible EVC modes.

Finding the Current EVC Mode for a Host

Each EVC mode corresponds closely to the features available in processors with the same name.

Using vSphere Web Client

1. From the vSphere Web Client Navigator, select **Hosts and Clusters** > **HX Cluster** > **Summary**. The **Summary** tab indicates whether EVC is enabled, and displays the current EVC mode for the host.
2. Click the blue icon next to the EVC mode to display a list of all the supported EVC modes for the host.

Using VMware Shared Utilities Tool

VMware provides a free CPU identification utility that displays the compatible EVC modes, in addition to other CPU features. You can download this utility and boot the host from the ISO image from [Shared Utilities](#).

Enable Enhanced vMotion Compatibility (EVC) on an Existing Cluster

Enable EVC to ensure that migration with vMotion is possible between hosts in the cluster. EVC mode is required when mixing different CPU families within the same HyperFlex cluster. Once EVC mode is enabled, only hosts that meet the minimum requirements of the configured EVC mode may be added to the cluster. EVC mode can be enabled, even during cluster expansion without disruption.



Note

- EVC is disabled by default. It can be enabled in the cluster settings under *VMware EVC*.
- This is not an HX Data Platform constraint but rather a VMware limitation. For further details, refer to the VMware KB article [EVC and CPU Compatibility FAQ \(1005764\)](#).



Note

- If you are enabling EVC mode, verify that EVC mode supports Advanced Encryption Standard New Instructions (AES-NI).
- This is not an HX Data Platform constraint but rather a VMware limitation. For further details, refer to the VMware KB article [EVC and CPU Compatibility FAQ \(1005764\)](#).

There are two paths to consider, uniform clusters that you want to add newer generation servers to, and existing clusters with mixed generations of servers.

Adding Newer Generation Servers to Uniform Clusters

If the cluster is currently uniform and you want to add newer generation servers to the cluster, you can enable EVC online and without disruption by selecting the current generation EVC mode in VC. Then proceed normally with expansion (converged or compute only). It is imperative the EVC mode is set before any expansion is attempted.

To enable EVC mode on a uniform cluster prior to performing a cluster expansion, perform the following steps:

-
- Step 1** Enable Enhanced vMotion Compatibility (EVC) in the HX cluster.
- From the vSphere Web Client Navigator, select **Hosts and Cluster > Datacenter > HX Cluster**.
 - Select the cluster for which you want to enable EVC. In the Work pane, click the **Manage** or **Configure** tab. Select **VMware EVC**.
 - Click the **Edit** button and select the desired *EVC mode*. Click **Ok**.
- Step 2** Proceed with compute only or converged node expansion using the HyperFlex installer.
-

Adding Mixed or Older Generation Servers to Existing Clusters

The cluster already has mixed generations of servers OR you wish to add older generation servers to an existing cluster (compute only nodes).



Note These steps should be followed if EVC mode was not enabled during the cluster expansion workflow with newer nodes.

To add an older generation server to an existing cluster perform the following steps:

-
- Step 1** Ensure that the HyperFlex cluster is healthy, and all the nodes are online before you start.
- **Using vSphere Web Client**

From the vSphere Web Client Navigator, select **Home > Global Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > HX Cluster > Summary**.

Example response:

```
Operational Status: Online
Resiliency Status: Healthy
```
 - **Using the Controller VM**

In the controller VM, run the command `# hxcli cluster info`.

Example response:

```
healthstate: healthy
state: online
```
- Step 2** Power off all non-storage controller virtual machines.

- Step 3** Log into one storage controller VM and execute the command `hxcli cluster shutdown`. Wait until the execution is complete.
- Step 4** Shut down all the storage controller VMs.
- From the vSphere Web Client Navigator, select **VMs and Templates** > **vCenter server** > **Datacenter** > **Discovered virtual machine** > **Virtual Machines** > **controller_vm**.
 - Right-click the **controller_vm** or from the **Actions** menu, select **Power** > **Shut Down Guest OS**.
- Step 5** Put each HX host in maintenance mode.
- From the vSphere Web Client Navigator, select **Hosts and Cluster** > **Datacenter** > **HX Cluster** > **node**.
 - Right-click the node and select **Maintenance Mode** > **Enter Maintenance Mode**.
- Note** Do not use the Cisco HX Maintenance Mode menu for this operation.
- Step 6** Enable Enhanced vMotion Compatibility (EVC) in the HX cluster.
- From the vSphere Web Client Navigator, select **Hosts and Cluster** > **Datacenter** > **HX Cluster**.
 - Select the cluster for which you want to enable EVC. In the Work pane, click the **Manage** or **Configure** tab. Select **VMware EVC**.
 - Click the **Edit** button and select the desired *EVC mode*. Click **Ok**.
- Step 7** Exit maintenance mode.
- From the vSphere Web Client Navigator, select **Hosts and Cluster** > **Datacenter** > **HX Cluster** > **node**.
 - Right-click the node and select **Maintenance Mode** > **Exit Maintenance Mode**.
- Note** Do not use the Cisco HX Maintenance Mode menu for this operation.
- Step 8** The controller VMs should automatically power on after the host exits maintenance mode. If the controller VMs do not power on automatically, do the following:
- From the vSphere Web Client Navigator, select **VMs and Templates** > **vCenter server** > **Datacenter** > **ESXi Agents** > **Virtual Machines** > **controller_vm**.
 - Right-click the **controller_vm** or from the **Actions** menu. Select **Power** > **Power On or Power** > **Power ON**.
- Step 9** Make sure all controller VMs are fully booted. Then, log in to one of the controller VMs, and run the command `hxcli cluster start`.
- Step 10** Make sure that all datastores are mounted from vCenter HX plug-in and ensure that the cluster is healthy.
- Step 11** Start the user VMs.
- Step 12** Proceed with compute only expansion using the HyperFlex installer.
-



CHAPTER 12

Cisco HyperFlex Systems Customized Installation Method

- [Overview, on page 157](#)
- [Installation and Configuration Workflow for Non Pre-Configured Cisco HyperFlex Systems, on page 157](#)
- [Installing VMware ESXi, on page 158](#)

Overview

This chapter describes the process for manually preparing the HyperFlex servers, before moving to the installation steps. This customized installation method can be used during the following scenarios:

- Adding a compute-only node to the HyperFlex cluster.
- Redeploying your Cisco HyperFlex System.

The tasks to perform are essentially the same as for deploying pre-configured HyperFlex Systems. The new task specific to non pre-configured systems is, installing VMware ESXi.

Installation and Configuration Workflow for Non Pre-Configured Cisco HyperFlex Systems

Before you begin

Review the installation and configuration requirements for Cisco HyperFlex Systems. See *Installation Prerequisites* for more details.

Step 1 Clean up the existing environment by deleting the cluster in vCenter. Remove the vCenter MOB entries, UCS Manager service profiles, and VLANs in Cisco UCS.

Step 2 Download the **Cisco HyperFlex Data Platform Installer OVA** file from [Download Software](#).

Example:

```
Cisco-HX-Data-Platform-Installer-v5.0.2a-41731-esx.ova
```

- Step 3** Launch the HX Data Platform Installer and sign-in.
- Select **Customized Workflow**.
 - Select **Run UCS Manager configuration** to configure UCS service profiles. Follow the steps as described in Configuring Cisco UCS Manager and HX Data Platform section of the Cisco HyperFlex Systems Installation Guide for VMware ESXi.
- Step 4** Perform a fresh ESXi installation using the vMedia method.
- Note** Using Dynamic Host Configuration Protocol (DHCP) for automatic IP address assignment is not recommended. By default, the HX Data Platform Installer assigns static IP addresses to the ESXi servers. If you are using DHCP, configure the networking in the ESXi manually with the proper VLANs.
- Step 5** Launch the HX Data Platform Installer again.
- Select **Customized Workflow**.
 - Select **Run ESX Configuration, Deploy HX Software, and Create HX Cluster**.
- Ensure that you select **Delete Existing Partitions** in the wizard.

Installing VMware ESXi

The servers must have a supported version of VMware ESXi. To attain the best HX snapshot performance and functionality, Cisco recommends installing ESXi 7.0 U2 or later. See the latest [Release notes for Cisco HX Data Platform](#) for a list of supported ESXi versions.



Important Install ESXi on each HX server.

Download the VMware ESXi image from the Cisco HyperFlex Data Platform [Download Software](#) page. Select a networked location that can be accessed through Cisco UCS Manager.

The HX custom ISO is based on the Cisco custom ESXi release.

For example,

```
HX-ESXi-7.0U3-21930508-Cisco-Custom-7.3.0.16-install-only.iso
```

What to do next

- Configure vMedia and Boot Policies through Cisco UCS Manager.
- Open a remote KVM console.
- Reboot the servers to begin installation.
- Undo vMedia and Boot Policy changes through Cisco UCS Manager.

Configure vMedia and Boot Policies Through Cisco UCS Manager

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

Before you begin

Log into HX Data Platform Installer. Run the **Create Cluster** workflow for standard cluster. up to Cisco UCS Manager configuration.



Note Create a cluster with name as *Temporary*, so it is easier to identify it during service profile disassociation from the server.

-
- Step 1** In Cisco UCS Manager, click the **Servers** tab in the Navigation Pane.
 - Step 2** Expand **Servers > Policies > root > Sub-Organizations > hx-cluster > vMedia Policies**
 - Step 3** Click **vMedia Policy HyperFlex**.
 - Step 4** In the configuration pane, click **Create vMedia Mount**.
 - Step 5** Enter a name for the mount, for example: **ESX**.
 - Step 6** Select the **CDD option**.
 - Step 7** Select the **HTTP** as the protocol.
 - Step 8** Enter the **IP Address** of the HyperFlex installer VM, for example: **192 . 168 . 10 . 210**.
 - Step 9** Select **None** as the Image Variable Name.
 - Step 10** Enter **HX-ESXi-7.0U3-21930508-Cisco-Custom-7.3.0.16-install-only.iso** as the Remote File.
 - Step 11** Enter **/images/** as the **Remote Path**.
 - Step 12** Click **Save Changes**, and click **OK**.
 - Step 13** In the Configuration Pane, select **Servers > Service Profile Templates > root > Sub-Organizations > hx-cluster > Service Template hx-nodes**
 - Step 14** Click the **vMedia Policy** tab.
 - Step 15** Click **Modify vMedia Policy**.
 - Step 16** Choose the **HyperFlex vMedia Policy** from the drop-down selection, and click **OK** twice.
 - Step 17** Select **Servers > Policies > root > Sub-Organizations > hx-cluster > Boot Policy HyperFlex**.
 - Step 18** In the Navigation Pane, expand the section titled **CIMC Mounted vMedia**.
 - Step 19** Click the entry labeled **Add CIMC Mounted CD/DVD**.
 - Step 20** Select the **CIMC Mounted CD/DVD** entry in the Boot Order list, and click the **Move Up** button until the CIMC Mounted CD/DVD entry is listed first.
 - Step 21** Click **Save Changes**, and click **OK**.
-

What to do next

Delete the Sub-Organization *Temporary*.

Opening a Remote KVM Console

To monitor the progress of one or more servers, it is advisable to open a remote KVM console session to watch the installation.

To open the KVM console, complete the following steps:

-
- Step 1** In Cisco UCS Manager, click **Servers** in the Navigation pane.
 - Step 2** Expand **Servers** > **Service Profiles** > **Root** > **Sub-Organizations** > **hx-cluster** > **rack-unit-number**.
 - Step 3** In the Work pane, click the **General** tab.
 - Step 4** In the Actions area, click **KVM Console**.
 - Step 5** Click **Continue** to any security alerts that appear. The remote **KVM Console** window appears shortly and shows the server's local console output.
 - Step 6** Repeat Steps 2-4 for any additional servers whose **KVM Console** you wish to monitor during the installation.
-

Rebooting Servers

Reboot the servers to begin the installation after modifying the vMedia policy, boot policy, and service profile template.

To reboot servers, complete the following steps:

Before you begin

Open a remote KVM Console sessions to monitor the progress of rebooting servers.

-
- Step 1** In Cisco UCS Manager, click **Servers** in the Navigation pane.
 - Step 2** Expand **Equipment** > **Rack Mounts** > **Servers**.
 - Step 3** In the Work pane, click the first server to be rebooted, then **shift+click** the last server to be rebooted, selecting them all.
 - Step 4** **Right-click** the mouse, and click **Reset**.
 - Step 5** Click **OK**.
 - Step 6** Select **Power Cycle** and click **OK**.

Now, the servers that you are monitoring in the KVM console windows immediately reboot, then boot from the remote vMedia mount, and install the Cisco customized ESXi ISO. If there are any error messages, they can be safely ignored.

Undo vMedia and Boot Policy Changes

To prevent the servers from going into a boot loop, constantly booting from the installation ISO file, undo the changes to the boot policy.

Before you begin

Ensure that all the servers have booted from the remote vMedia file and have begun their installation process.

-
- Step 1** In Cisco UCS Manager, click **Servers** in the Navigation pane.
- Step 2** Expand **Servers > Policies > Root > Sub-Organizations > *hx-cluster_name* > Boot Policies > Boot Policy HyperFlex**.
- Step 3** In the Work pane, click the **General** tab.
- Step 4** In the Actions area, click **CIMC Mounted CD/DVD**.
- Step 5** Select the **CIMC Mounted CD/DVD** entry in the **Boot Order** list, and click **Delete**.
- Step 6** Click **Save Changes**, and click **OK**.
-

What to do next

New Node

If you are adding a new node, never used before in the cluster, expand the HX cluster. See the Cluster Expansion Guidelines section for more details.

Reinstalling an Existing Node

If this node was part of the cluster in the past and was reimaged to fix something, contact Cisco TAC for guidance.



APPENDIX **A**

Lockdown Mode

Overview

This section provides an introduction to Lockdown mode, which is used to increase the security of an ESXi host by limiting the access allowed to the host. When this mode is enabled, the ESXi host can only be accessed through the vCenter server or the Direct Console User Interface (DCUI). Enabling Lockdown mode affects which users are authorized to access host services.



Note When enabling lockdown mode, you must add the hxuser account to each ESXi host exception user list.

- [Enable or Disable Lockdown Mode, on page 163](#)
- [Troubleshoot Lockdown Mode, on page 164](#)

Enable or Disable Lockdown Mode

This section describes how to enable or disable Lockdown mode either from the DCUI or from the vSphere web client.

Enable or Disable Lockdown Mode from the DCUI:

- Step 1** Log directly in to the ESXi host.
 - Step 2** Open the Direct Console User Interface (DCUI) on the host.
 - Step 3** Press **F2** for Initial Setup.
 - Step 4** Press **Enter** to toggle the Configure Lockdown Mode setting.
 - Step 5** Browse to the host in the vSphere Web Client inventory.
-

Enable or Disable Lockdown Mode from the vSphere Web Client:

- Step 1** Browse to the host in the vSphere Web Client inventory.

- Step 2** Click the **Manage** tab and click Settings.
 - Step 3** Under System, select **Security Profile**.
 - Step 4** In the Lockdown Mode panel, click **Edit**.
 - Step 5** Click **Exception Users** and select +Add user to add hxuser (all lowercase).
 - Step 6** Click **Lockdown Mode** and select one of the Lockdown mode options.
-

Troubleshoot Lockdown Mode

If you receive an error dialog box and/or the software upgrade fails in Lockdown mode, proceed with the following resolution options based the following scenarios:

When at least one host is in Lockdown mode

If adding hosts to vCenter in the deploy phase fails, and you receive the error message **Failed to add hosts to vCenter**:

- Step 1** Check the host Lockdown mode in pre-upgrade validation.
 - Step 2** Detect the situation, throw an error, and terminate the cluster upgrade.
 - Step 3** Disable Lockdown mode and try the upgrade again.
-

When the host is in Lockdown mode while the upgrade in progress:

If adding hosts to vCenter in the deploy phase fails, and you receive the error message **Failed to add hosts to vCenter**:

- Step 1** Check the host Lockdown mode before host upgrade.
 - Step 2** Detect the situation and error out and failed the upgrade.
 - Step 3** Disable Lockdown mode and try the upgrade again.
-