



Cisco HyperFlex Systems Stretch Cluster Guide, Release 6.0

First Published: 2024-07-18

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)

© 2024 Cisco Systems, Inc. All rights reserved.



CONTENTS

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

CHAPTER 1	Introduction	1
	Introduction	1

CHAPTER 2	Preinstallation Checklist	3
	Common Site Interlink and Witness Requirements	3
	Invisible Cloud Witness Preinstallation Checklist for All New Installs	8
	Stretch Witness VM	8
	Using the Intersight Arbitrator Swichover	9
	Changing the Witness VM Password	12
	Witness VM Network IP Addressing	13

CHAPTER 3	Guidelines and Limitations	15
	Guidelines	15
	Limitations	16

CHAPTER 4	Expanding a Stretch Cluster	17
	Cluster Expansion Guidelines	17
	Configuring the Sites for Expand Cluster	17
	Enter Credentials	18
	Associate HyperFlex Servers	19
	Configure Hypervisor	20
	Cluster Expansion Workflow	22
	Enter Credentials	22
	Associate HyperFlex Servers	25

Configure Nodes 26

CHAPTER 5

Replication 29

Configuring Replication 29

Configuring Replication VLAN in Cisco UCS Manager 29

Creating Replication VLAN 29

Associating Replication VLAN to a HyperFlex Cluster 32

CHAPTER 6

Managing HyperFlex Stretch Clusters 35

Overview 35

Monitoring the Health of a Stretch Cluster 35

Viewing System Information 37

Creating a Datastore 42

Shutting a Stretch Cluster Site Down Gracefully 42

Changing the Preferred Site 43

CHAPTER 7

Troubleshooting 45

Viewing Resiliency Status in HX Connect 45

Troubleshooting the Network Configuration 47

Troubleshooting for Site-to-Site Failover 48

CHAPTER 8

Post Installation 49

Confirm That the Installation is Complete 49

Running Post Install Script 51



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

Introduction

- [Introduction, on page 1](#)

Introduction

HyperFlex Stretch Cluster enables you to deploy an Active-Active disaster avoidance solution for mission critical workloads requiring high uptime (near zero Recovery Time Objective) and no data loss (zero Recovery Point Objective). A stretch cluster is a single cluster with geographically distributed nodes. Both sides of the cluster act as primary for certain user VMs. The data for these VMs is replicated synchronously on the other site. Stretch clusters enable you to access the entire cluster even if one of the sites were to completely go down. Typically these sites are connected with a low latency, dedicated, high-speed link between them.

Beginning with HyperFlex 5.5(1a) there are two Stretch Cluster models. Witness VM and Invisible Cloud Witness. The Invisible Cloud Witness also introduces the Preferred Site feature, designed to serve requests in the event that the connectivity between sites is down.

Stretch Cluster Invisible Cloud Witness Model

Stretch Clusters installed using HXDP 5.5(1a) will auto-configure an Invisible Cloud Witness for site arbitration, without the need for of an external witness service in a third site. Invisible Cloud Witness automatically runs the latest Arbitrator version, user maintenance of this component is not required.

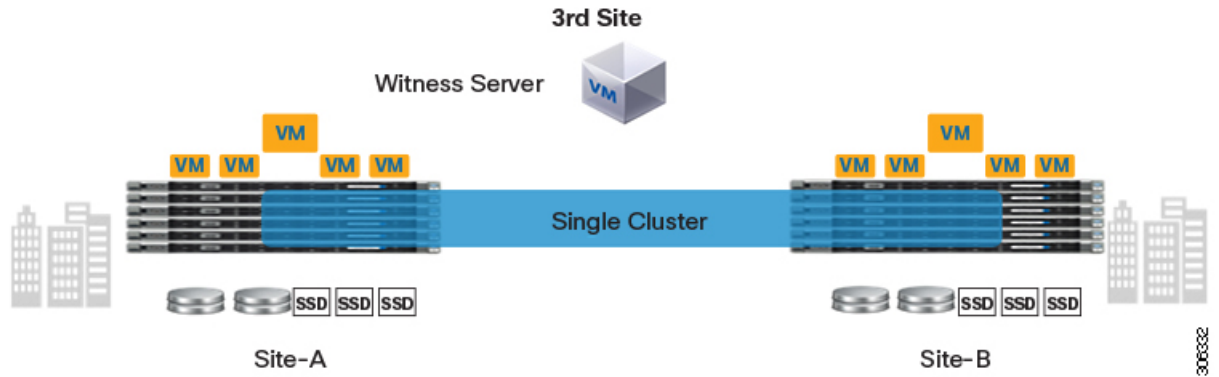


Stretch Cluster Witness VM Model

If you have an existing Stretch Cluster (HXDP Release 5.0(x) and earlier), you will continue to use the VM based witness as described in the [Cisco HyperFlex Systems Stretch Cluster Guide, Release 5.0](#) and the [Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 5.0](#).



Attention This model only supports existing Stretch clusters deployed on HX Release 5.0(x) and earlier. All new Stretch Clusters will use Invisible Cloud Witness Model.



This chapter focuses on new Invisible Cloud Witness stretch cluster installations, provides information about monitoring stretch clusters, and discusses guidelines to follow when deploying Invisible Cloud Witness stretch clusters.



CHAPTER 2

Preinstallation Checklist

- [Common Site Interlink and Witness Requirements, on page 3](#)
- [Invisible Cloud Witness Preinstallation Checklist for All New Installs, on page 8](#)
- [Stretch Witness VM, on page 8](#)
- [Using the Intersight Arbitrator Swichover , on page 9](#)
- [Changing the Witness VM Password, on page 12](#)
- [Witness VM Network IP Addressing, on page 13](#)

Common Site Interlink and Witness Requirements


The following describe general requirements, and the HXDP release they apply to.

Network Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
10 Gbps dedicated, 5-ms Round-Trip Time (RTT) latency between the two active sites is required.	✓	✓
100 Mbps, 200 ms-RTT worst case latency for 16 kilobyte packet sizes between the active sites and witness site is required.	-	✓
Existing fabric interconnects are supported, provided the fabric interconnects support M5 or M6 servers.	✓	✓
User VMs should be capable of vMotioning to any site, without impacting external network connectivity to these VMs.	✓	✓
Similar to regular HX Clusters, two separate IP subnets are required—Both over Stretch L2. One subnet for data traffic and one for management traffic, with the management subnet reachable from vCenter and the witness node.	✓	✓

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
<p>FI facing ports need to have Port-fast, spanning-tree port type edge trunk, or similar spanning tree configuration that immediately put ports into forwarding mode.</p>	✓	✓
<p>QoS</p> <p>If the HyperFlex cluster is a stretch cluster, you should enable QoS end-to-end between the sites. The QoS policies in the Cisco UCS and upstream network should also be aligned so that HyperFlex traffic can receive consistent QoS end-to-end.</p> <p>If the HyperFlex cluster is a stretch cluster that connects to an ACI Multi-Pod fabric, you should enable QoS policies in the ACI fabric and extend it across the IPN as well.</p> <p>The QoS policies in the Cisco UCS and ACI fabrics should also be aligned so that HyperFlex traffic can receive consistent QoS end-to-end.</p> <p>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.</p>	✓	✓

Network Topology



Witness VM Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
Configure and enable NTP on all servers.	-	✓
An independent third witness site is required.	-	✓
A separate witness is required for each cluster.	-	✓
Both the main sites must have connectivity to the third witness site with a minimum bandwidth of 100 MBPS, 200 ms-RTT worst case latency for 16 kilobyte packet sizes.	-	✓
Site must have the capability to deploy and run Open Virtualization Format (OVF) image.	-	✓

Network Latency Requirements for the Witness VM

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
The HyperFlex Stretch Cluster solution requires that the Witness VM be located in a third site to ensure that a storage site failure does not affect the Witness VM.	-	✓
The solution can support a Witness bandwidth as low as 100 Mbps, 200 ms-RTT worst case latency for 16 kilobyte packet sizes.	-	✓
Latency to the witness impacts site failure times, and it is recommended that for larger clusters with significant load and data, to have RTT times in the order of 10ms or lower.	-	✓

Witness VM Fabric Interconnect and Node Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
Symmetric configuration is required across both sites.	✓	✓
There must be a minimum of two converged nodes on each site.	✓	✓
A maximum of 16 converged nodes on each site is supported. Ensure that both sites have the same number of converged nodes.	✓	✓
There must be a redundant fabric interconnect configuration on each site.	✓	✓
Converged nodes have to be M5 or M6 nodes.	✓	✓
Ensure that the Fabric Interconnect pair is of the same model in the same domain.	✓	✓
Compute-only nodes are supported.	✓	✓

VMware Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
VMware Enterprise Plus edition with HyperFlex Stretch Cluster is highly recommended to ensure proper failover behavior and guarantee high performance during normal operations. While it is possible to run Stretch Cluster without VMware Enterprise Plus edition, the advanced DRS features such as site affinity will not be available, negating some of the intended operational functions of Stretch Cluster.	✓	✓
Use a single vCenter for both sites.	✓	✓
The vCenter can be a VM running at the same site as the witness.	-	✓
Nested vCenter is not supported in Stretch Cluster.	✓	✓
The vCenter must be configured independently for High Availability, as required.	✓	✓

UCS Manager Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
Two separate, manually synchronized fabric interconnect domains are required.	✓	✓

VLAN Requirements

Requirement	HXDP Release 5.5(1a) and Later (Invisible Cloud Witness)	HXDP Release 5.0(x) and Earlier (Witness VM)
IP addresses for nodes on both sites are required.	✓	✓
Stretch VLANs across both sites are required.	✓	✓

Invisible Cloud Witness Preinstallation Checklist for All New Installs

This checklist applies to all new Invisible Cloud Witness Stretch Cluster installs using HXDP 5.5(1a) and later.



Remember All new Stretch Cluster installs auto-configure an Invisible Cloud Witness for site arbitration. Fresh installs using the Witness VM (HXDP 5.0(x) and earlier) is not supported.

- Using the Intersight Arbitrator instead of witness setup for switchover is supported.
- Network connectivity for the install powered by Installer VM.
- Intersight deployment of stretch-cluster is NOT supported.
- Intersight access
- Auxiliary ZooKeeper (AUX ZK IP): This must be in the same data network.
- Preferred site: The site designated to serve requests in the event that the connectivity between sites is down.



Note If you have an existing Stretch Cluster (HXDP Release 5.0(x) and earlier), you will continue to use upgrade the VM based witness as described in the [Cisco HyperFlex Systems Stretch Cluster Guide, Release 5.0](#) and the [Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 5.0](#).

Stretch Witness VM


Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

Table 1: Port Requirements

Port Number	Service/Protocol	Source	Port Destinations	Essential Information
2181 2888 3888	(Zookeeper lifecycle)/TCP	Witness	Each CVM Node	Bidirectional, management addresses
8180	Exhibitor (Zookeeper lifecycle)/TCP	Witness	Each CVM Node	Bidirectional, management addresses
80	HTTP/TCP	Witness	Each CVM Node	Potential future requirement
443	HTTPS/TCP	Witness	Each CVM Node	Potential future requirement

Using the Intersight Arbitrator Swichover

Beginning with HXDP Release 6.0(x), you can use the Intersight Arbitrator instead of witness setup for swichover



- A HyperFlex Witness VM node is mandatory in a stretch cluster environment to achieve quorum in case of total failure in any of the sites or when the network link between the sites encounters a failure.
- In case of accidental deletion or loss of Witness VM, to replace the Witness VM please contact Cisco TAC.
- The Witness VM requires a Static IP address that cannot be changed without cluster redeployment. If DHCP server is used to define the network configuration, the IP address needs to be dedicated to Witness VM.

Complete the following tasks to deploy Stretched Cluster Witness using Intersight's Arbitrator for swichover using hxcli.

SUMMARY STEPS

1. Run the `hxcli` command for swichover from Controller vm.
2. To check the current witness related details of the stretch cluster use the command `hxcli cluster witness`.
3. To locate the initiate swichover parameters definitions, use the command `hxcli cluster swichover -h`.
4. To initiate swichover, use the command `hxcli cluster swichover`.

- Review your Stretch Cluster Infrastructure using the command `hxcli cluster witness`.

DETAILED STEPS

- Step 1** Run the `hxcli` command for swichover from Controller vm.

```

hxshell:~$ hxcli cluster
Displays the list of commands available in the cluster namespace.
Usage :
hxcli cluster [flags]
hxcli cluster [command]
Aliases :
cluster, cl
Available Commands:
about Cluster version, model and other metadata details
detail Short summary of cluster details
health Health and resiliency information of the cluster
info Detailed information about cluster
shutdown Shuts down the Hyperflex cluster
start Start the Hyperflex cluster
stats Cluster capacity and space savings
Cluster Operational Status and Resiliency Health
Stretch Cluster witness infrasturcture swichover process
Stretch Cluster witness infrastructure status
Flags:
-h, -help help for cluster
status
swichover
witness

```

- Step 2** To check the current witness related details of the stretch cluster use the command `hxcli cluster witness`. The following example shows the output of command before swichover is done on a Stretch cluster. Note that the Stretch Cluster has `ZK_WITNESS` and the witness ip is displayed to user.

```

hxshell:~$ hxcli cluster witness
Stretch Cluster witness type : ZK_WITNESS Stretch Cluster witness ip : 10.198.148.26
Stretch Cluster witness state : ONLINE hxshell:~$ |

```

- Step 3** To locate the initiate swichover parameters definitions, use the command `hxcli cluster swichover -h`.

```

hxshell:~$ hxcli cluster swichover -h
Stretch Cluster witness infrasturcture swichover process

Usage:
  hxcli cluster swichover [flags]

Aliases:
  swichover, so

Flags:
  -l, --arb-ip string      Stretch Cluster witness local arbitrator ip
  -n, --arb-name string    Stretch Cluster witness local arbitrator user name
  -p, --arb-pwd string     Stretch Cluster witness local arbitrator password
  -t, --arb-type string    Stretch Cluster witness arbitrator type [intersight | local]
  -a, --aux-ip string     Stretch Cluster witness aux ip
  -h, --help               help for swichover
  -s, --site string       Stretch Cluster witness preferred site

```

Common Error Messages: Input Arguments Checking

- a) **Aux IP Expectations:** Aux ip provided if not valid, unused, and present in same subnet of the data network, is reported as invalid input with the following error message:

```
Internal Server Error (500): ("cause":"Aux Ip provided should be valid, unused, and present in same subnet as the data network ", "errorCode":400, "message":"!)
```

- b) **Site Input Expectations:** Site provided should be one of your stretch cluster sites. If not, it is reported as invalid parameter with the following error message:

```
Internal Server Error (500): ("cause":"Field name: site has illegal value. Should be one of: SiteA.SiteB", "errorCode":400, "message":"failure: Witness vm to Arbitrator swichover process failed.")
```

- c) **Arbitrator Type Expectations:** Arbitrator type should be local. If not it is reported as an invalid parameter with the following error message:

```
Internal Server Error (500): ("cause":"Field name: arb type has illegal value. Should be one of: intersight local", "errorCode":400, "message":"failure: Wittness vm to Arbitrator swichover process failed.
```

Step 4 To initiate swichover , use the command `hxcli cluster swichover`.

The following example shows a successful swichover.

```
hxshell:~$ hxcli cluster swichover --aux-ip 172.16.60.102 --site SiteA --arb-type intersight

Stretch cluster swichover from Witness vm to Arbitrator has been initiated..
Swichover process might take 5 - 6 mins or more based on certain conditions: /

Awaiting ensemble changes to update on all nodes
Swichover process might take 5 - 6 mins or more based on certain conditions: / Task Id : 750975cb-abd6-433a-aa5d-1c419e4ee4a2
Description : Notify swichover ensemble changes to all
              : nodes
Status      : Succeeded
Name       : Swichover ensemble update workflow
Category   : STRETCH_CLUSTER_SWICHOVER
Submitted At : 2024-04-05 01:06:28.804 -0700 PDT
Swichover process might take 5 - 6 mins or more based on certain conditions: -
Initiating conversion of standalone nodes to observers

Stretch Cluster witness swichover to arbitrator success

Stretch Cluster preferred site has been updated successfully
hxshell:~$
```

Note Post swichover, standalone nodes are converted to observers. To verify conversion, use the command `zkstatus`.

Step 5 Review your Stretch Cluster Infrastructure using the command `hxcli cluster witness`.

Example output:

```
hxshell:~$ hxcli cluster witness
Stretch Cluster witness type : INTERSIGHT
Stretch Cluster witness ip : INTERSIGHT
Stretch Cluster witness state : ONLINE
```

The screenshot displays the Cisco HyperFlex Connect interface for the simulator_witness node. At the top, there is a notification: "Cluster not registered with Cisco Licensing. Register Now". Below this, the node status is shown as "ONLINE". A table provides details about the node's configuration and license status:

License Type	Evaluation	Actions	
License Status	License expires in 89 days. Cluster not registered with Cisco Licensing.		
Hypervisor	7.0.3-20842708	Total Capacity	1.84 TiB
HXDP Version	6.0.1a-44239	Available Capacity	1.8 TiB
Witness	Online (INTERSIGHT)	Data Replication Factor	2+ 2
		DNS Server(s)	10.107.128.8
		NTP Server(s)	10.107.129.14
		Controller Access over SSH	Enabled

Below the table, the node is identified as "SiteA | **PREFERRED SITE** | Hyperconverged Nodes". The interface also shows "Disk View Options" and a "Disk View Legend".

The bottom portion of the screenshot shows the "Activity" log for the simulator_witness node. The "Activity" tab is selected in the left sidebar. The log contains the following entries:

- Notify switchover ensemble changes to all nodes
Status: Success
04/05/2024 1:36:28 PM
- Switch a stretch cluster from using Witness VM to a local/intersight arbitrator
Status: Success
04/05/2024 1:36:23 PM

Below the activity log, there are several task cards with green checkmarks indicating successful completion:

- Validate inputs and cluster condition: Validate inputs and cluster condition
- Remove Witness VM and Reconfigure Storage Cluster: Remove Witness VM and Reconfigure Storage Cluster
- Configuring new services to provide Arbitration Capability: Configuring new services to provide Arbitration Capability

What to do next

After successfully deploying the witness node, you can proceed to installing your Stretch cluster. When prompted enter the IP address of this witness node on the **IP Address** page when creating a HyperFlex Stretch cluster. The witness node is automatically used during configuration of the stretch cluster.

Changing the Witness VM Password

After successful deployment of witness VM, you must change the default password.

Before you begin

Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

Download and deploy the witness VM.

Step 1 Log into the witness VM using SSH.

```
$ ssh root@<IP address of witness VM>
```

Step 2 Enter the default password.

Step 3 Enter the `passwd` command and change the password.

```
$ passwd
Changing password for user admin.
```

Step 4 Log out from the witness VM.

What to do next

Log into the witness VM with the new password.

Witness VM Network IP Addressing

Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

IP addresses for HyperFlex Stretch Cluster need to be allocated from the appropriate subnets and VLANs to be used.



Important

- Required only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.
- Ensure that the Data and Management Networks are on different subnets for a successful installation.
- Ensure that the IP addressing is in the same subnet for the same components (Management Network IP addresses, Data Network IP addresses).

Table 2: Example: Stretch Cluster IP Addressing

Storage Cluster Management IP address	10.10.10.128	Storage Cluster Data IP address	192.168.10.160
Subnet mask IP address	255.255.255.0	Subnet mask IP address	255.255.255.0
Default gateway IP address	10.10.10.1	Default gateway IP address	192.168.10.1

Table 3: Example: Stretch Cluster IP Addressing for Site A

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
ESXi Hostname*	Hypervisor Management Network	Storage Controller Management Network	Hypervisor Data Network (Not Required for Cisco Intersight)	Storage Controller Data Network (Not Required for Cisco Intersight)
Server 1:	10.10.10.2	10.10.10.32	192.168.10.2	192.168.10.32

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
Server 2:	10.10.10.3	10.10.10.33	192.168.10.3	192.168.10.33
Server 3:	10.10.10.4	10.10.10.34	192.168.10.4	192.168.10.3
Server 4:	10.10.10.5	10.10.10.35	192.168.10.5	192.168.10.35
Server 5:	10.10.10.6	10.10.10.36	192.168.10.6	192.168.10.36

Table 4: Example: Stretch Cluster IP Addressing for Site B

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
ESXi Hostname*	Hypervisor Management Network	Storage Controller Management Network	Hypervisor Data Network (Not Required for Cisco Intersight)	Storage Controller Data Network (Not Required for Cisco Intersight)
Server 1:	10.10.10.64	10.10.10.96	192.168.10.64	192.168.10.96
Server 2:	10.10.10.65	10.10.10.97	192.168.10.65	192.168.10.97
Server 3:	10.10.10.66	10.10.10.98	192.168.10.66	192.168.10.98
Server 4:	10.10.10.67	10.10.10.99	192.168.10.67	192.168.10.99
Server 5:	10.10.10.68	10.10.10.100	192.168.10.68	192.168.10.100



CHAPTER 3

Guidelines and Limitations

- [Guidelines, on page 15](#)
- [Limitations, on page 16](#)

Guidelines

Consider the following guidelines when creating a HyperFlex Stretch Cluster:

- HXDP Enterprise Edition licensing is required to use the HyperFlex Stretch Cluster feature.
- vSphere Enterprise Plus licensing is required for full featured HyperFlex Stretch Cluster functionalities like VM load balancing and VM migration.
- Ensure that DRS is enabled.
- Supported Stretch Cluster scale:
 - There must be a minimum of 2 nodes on each site. The minimum overall cluster size across both sites is 4.
 - Small Form Factor (SFF). The maximum cluster size across both sites is 64. The maximum converged node count is 16 per site. The compute to converged node ratio can be 2:1, while the maximum limit of nodes per site must not exceed 32. For example, you can have 11 converged and 21 compute nodes per site.
 - Large Form Factor (LFF). The maximum cluster size is 48. The maximum converged node count is 8 per site. The compute to converged node ratio can be 2:1, while the maximum limit of nodes per site must not exceed 24. For example, you can have 8 converged and 16 compute nodes per site.
- There must be a redundant fabric interconnect configuration on each site.
- There must be symmetric cluster configuration across both sites. The number of nodes and the model of HX nodes should be the same on both sites.
- VMs are placed correctly within site affinity only if VMware HA and DRS are enabled before VM creation. Otherwise, the correct placement of VM for proper affinity is not guaranteed.
- HyperFlex Native Replication is supported between Stretch Clusters, and between Stretch Clusters and standard clusters.
- It is a best practice to have two datastores one per site, with the respective affinity.

- **Guest Traffic:** Due to the nature of the Cisco VIC carving up multiple vNICs from the same physical port, it is not possible for guest VM traffic configured on vswitch-hx-vm-network to communicate L2 to interfaces or services running on the same host. It is recommended to either a) use a separate VLAN and perform L3 routing or b) ensure any guest VMs that need access to management interfaces be placed on the vswitch-hx-inband-mgmt vSwitch. In general, guest VMs should not be put on any of the HyperFlex configured vSwitches except for the vm-network vSwitch. An example use case would be if you need to run vCenter on one of the nodes and it requires connectivity to manage the ESXi host it is running on. In this case, use one of the recommendations above to ensure uninterrupted connectivity.

Limitations

Consider the following limitations when creating a HyperFlex Stretch Cluster:

- Self Encrypting Drives (SEDs) are not supported. However, VM based third-party software encryption is supported.
- Overlay networking and L3 protocols are not supported. L2 adjacency is required for the data and management networks.
- Stretch Cluster is not supported on Hyper-V platform.
- iSCSI is not supported in Stretch cluster.
- Node removal is not supported.
- Online rolling upgrades are supported only for the HX Data Platform. Cisco UCS Manager upgrades and VMware ESXi upgrades must be performed manually on one node at a time or performed offline.
- Upgrade from standalone cluster to Stretch Cluster configuration is not supported.
- Stretch Cluster is supported only on M5 or M6 nodes. M5/M6 is supported when reached by expansion.



CHAPTER 4

Expanding a Stretch Cluster

- [Cluster Expansion Guidelines, on page 17](#)
- [Configuring the Sites for Expand Cluster, on page 17](#)
- [Cluster Expansion Workflow, on page 22](#)
- [Enter Credentials, on page 22](#)
- [Associate HyperFlex Servers, on page 25](#)
- [Configure Nodes, on page 26](#)

Cluster Expansion Guidelines

- Stretch cluster expansion supports both converged nodes and compute-only nodes.
- When adding a converged node, ensure that the configuration is symmetric across both sites. For instance if Site 1 is expanded with two nodes, Site 2 must also be expanded with two converged nodes.
- When adding compute nodes, ensure that you do not exceed the supported node count.

Configuring the Sites for Expand Cluster

Expanding a Stretch Cluster includes various steps that are involved in adding a Compute-only or a Converged node to a Stretch Cluster. See [Cluster Expansion Workflow, on page 22](#) for more details.

Before executing the cluster expansion workflow, re-create the sites in the installer as they are deployed using this procedure.

Before you begin

- Complete the prerequisites for your install. or [Invisible Cloud Witness Preinstallation Checklist for All New Installs](#)
- Ensure that you have a functional, symmetric, and running Stretch Cluster.

Step 1 Log into the HX Data Platform Installer.

Step 2 On the **Workflow** page, from the **Expand Cluster** drop-down list, select **Stretch Cluster**. Click **Continue**.

- Step 3** On the **Cluster** page, enter hostname and credentials of management cluster.
- Step 4** On the **Credentials** page, enter UCS Manager and Hypervisor credentials. See [Enter Credentials](#) for more details.
- Step 5** On the **Server Selection** page, configure the server ports and associate the new HX expansion nodes with the site. See [Associate HyperFlex Servers](#) for more details.
- Step 6** On the **Hypervisor Configuration** page, configure the subnet mask, gateway, and complete the hypervisor settings. See [Configure Hypervisor](#) for more details.
- Step 7** Click **Start** to begin site configuration for expand 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

Using the HX Data Platform Installer, proceed to create your expanded HyperFlex Stretch Cluster. See [Cluster Expansion Workflow, on page 22](#) for more details.

Enter Credentials

Enter UCS Manager credentials and hypervisor credentials for both the sites across the Stretch Cluster.

Before you begin

- Complete the prerequisites for your install. or [Invisible Cloud Witness Preinstallation Checklist for All New Installs](#)
- Download and deploy the Witness VM¹.

- Step 1** Log in to the HX Data Platform Installer with root user credentials. For first time log in, you will be prompted to change the default password set in the factory.
- In a browser, enter the URL for the VM where HX Data Platform Installer was installed.
 - Enter the following login credentials.
 - Username: **root**
 - Password: **Cisco123**

c) Read the EULA, select the **I accept the terms and conditions** checkbox, and click **Login**.

- Step 2** In the **Change factory default password** screen, complete the following fields, and click **Change Password & Login**.

Field	Description
New password	Enter a new hypervisor password.
Re-enter new password	Re-enter the new hypervisor password.

¹ Applies to Witness VM only.

Step 3 On the **Workflow** page, from the **Expand Cluster** drop-down list, select **Stretch Cluster**. Click **Continue**.

Step 4 On the **Credentials** page, select **Configure Site**.

Step 5 Enter the following credentials for UCS Manager:

UCS Manager Credentials for this site

Field	Description
UCS Manager Host Name field	Enter the UCS Manager FQDN or IP address. For example, <i>10.193.211.120</i> .
UCS Manager User Name field	Enter the administrative level username. For example, <i><admin></i> .
Password field	Enter the administrative level password. For example, <i><root></i> .
Site Name field	Specify a unique site name.
Org Name field	Specify a unique org name, to ensure isolation of the HyperFlex environment from the rest of the UCS domain.

Step 6 Enter the following credentials for the Hypervisor:

Hypervisor Credentials

Field	Description
User Name field	Enter the administrative username. For example, <i><admin></i> username. The username is root for factory nodes.
Password	Enter the administrative level password.

Step 7 Click **Continue** to begin associating HyperFlex Servers. See [Associate HyperFlex Servers, on page 19](#).

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 column	Turn on to locate a server.
Server Name column	Name that is assigned to the server.

Field	Description
Status column	<ul style="list-style-type: none"> • Inaccessible • Ok
Model column	Displays the server model.
Serial column	Displays the serial number of the server.
Service Profile column [Only for Associated Servers]	Service profile that is assigned to the server.
Actions drop-down list	<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 on the **Credentials** page. See [Enter Credentials, on page 18](#).

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 Hypervisor. See [Configure Hypervisor, on page 20](#).

Configure Hypervisor

Step 1 In the **Configure Common Hypervisor Settings** section, complete the following fields:

Field	Description
Subnet Mask field	Set the subnet mask to the appropriate level to limit and control IP addresses. For example, <i>255.255.0.0</i> .

Field	Description
Gateway field	IP address of gateway. For example, <i>10.193.0.1</i> .
DNS Server(s) field	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.
DNS Domain Name field	Enter the DNS FQDN or IP address.

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

Step 3 In the **Advanced Configuration** section, select **Clean up disk partitions** to remove all existing data and partitions from all nodes added to the storage cluster.

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.
- Ensure that you backup any data which must be retained.

Step 4 Click **Start** to begin site configuration for expand 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

Run the Cluster Expansion workflow from the HX Data Platform Installer to create your HyperFlex Stretch Cluster expansion. See [Cluster Expansion Workflow, on page 22](#) for more details.

Cluster Expansion Workflow

The following cluster expansion workflow summarizes the steps that are involved in adding a Compute-only or a Converged node to a Stretch Cluster, using the HX Data Platform installer.

Step	Description	Reference
1.	Enter UCS Manager Credentials for Site A and Site B, vCenter credentials, and the Hypervisor Credentials.	Enter Credentials, on page 22
2.	Configure the server ports and associate HyperFlex servers.	Associate HyperFlex Servers, on page 25
3.	Configure hypervisor, IP addresses, and start the cluster expansion process.	Configure Nodes, on page 26

Enter Credentials

Before you begin

- Complete the cluster expansion prerequisites.
- Run the *Configure Site* workflow once for each site. Refer to the [Invisible Cloud Witness Preinstallation Checklist for All New Installs, on page 8](#) for more details.

Step 1 Log into the Cisco HX Data Platform Installer.

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

Username: root

Password (Default): cisco123

Important 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.

- Read the EULA, check the **I accept the terms and conditions** check box, and click **Login**.

Step 2 On the **Workflow** page, from the **Expand Cluster** drop-down list, select **Stretch Cluster**. Click **Continue**.

Step 3 On the **Cluster** page, enter the following credentials:

To perform Stretch 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 For a first-time installation, contact your Cisco representative to procure the factory preinstallation JSON file.

Field	Description
Cluster Management FQDN/IP field	Enter the HyperFlex cluster FQDN or IP address. For example, <i>10.193.211.120</i> .
User Name field	Enter the administrative level username. For example, <i><admin></i> .
Password field	Enter the administrative level password.

Click **Continue**.

Step 4 On the **Credentials Page**, select **Expand Stretch Cluster**.

Step 5 Enter the following credentials for UCS Manager:

UCS Manager Credentials for Site 1

Field	Description
UCS Manager Hostname field	UCS Manager FQDN or IP address for site 1. For example, <i>10.193.211.120</i> .
User Name field	Enter the administrative level username. For example, <i><admin></i> username.
Password field	Enter the administrative level password.
Site Name field	Specify a unique site name.
Org Name field	Specify a unique org name, to ensure isolation of the HyperFlex environment from the rest of the UCS domain.

UCS Manager Credentials for Site 2

Field	Description
UCS Manager Host Name field	Enter the UCS Manager FQDN or IP address for site 2. For example, <i>10.193.211.120</i> .
User Name field	Enter the administrative level username. For example, <i><admin></i> username.
Password field	Enter the administrative level password.
Site Name field	Specify a unique site name.

Field	Description
Org Name field	Specify a unique org name, to ensure isolation of the HyperFlex environment from the rest of the UCS domain.

Step 6 Enter the following credentials for the vCenter.

vCenter Credentials

Field	Description
vCenter Server field	<p>Enter the vCenter server FQDN or IP address.</p> <p>For example, <i>10.193.211.120</i>.</p> <p>Note</p> <ul style="list-style-type: none"> vCenter Server input is optional if building a nested vCenter. See the Nested vCenter TechNote for more details. 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.
User Name field	<p>Enter the administrative username.</p> <p>For example, <i>administrator@vsphere.local</i>.</p>
Admin Password field	<p>Enter the administrative level password.</p> <p>Enter the <i><root></i> password.</p>

Step 7 Enter the following credentials for the Hypervisor:

Hypervisor Credentials

Field	Description
Admin User Name field	<p>Enter the administrative username.</p> <p>For example, <i><admin></i> username.</p> <p>The username is root for factory nodes.</p>
The hypervisor on this node uses the factory default password check box	<p>Select the check box, if you are changing the factory default password for the hypervisor.</p> <p>To enter the current hypervisor password, uncheck this check box. You can now enter the hypervisor password in the Enter current hypervisor password field.</p>
New Password	<p>Create a new password for the hypervisor.</p> <p>Important You are required to change the factory default password.</p>
Confirm New Password	Reenter the new password for the hypervisor.

What to do next

Click **Continue** to begin associating HyperFlex servers.

Associate HyperFlex Servers

On the **Node 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 column	Turn on to locate a server.
Server Name column	Name that is assigned to the server.
Site column	Name of the site where the server is physically located.
Status column	<ul style="list-style-type: none"> • Inaccessible • Ok
Model column	Displays the server model.
Serial column	Displays the serial number of the server.
Service Profile column [Only for Associated Servers]	Service profile that is assigned to the server.
Actions column	<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 credentials for both sites and vCenter credentials, and Hypervisor credentials on the **Credentials** page. See [Enter Credentials, on page 22](#).

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.

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 HyperFlex nodes. See [Configure Nodes, on page 26](#).

Configure Nodes

On the **IP Addresses** page, complete the following:

Before you begin

Associate servers on the HyperFlex cluster. See [Associate HyperFlex Servers, on page 25](#).

Step 1 In the **IP Addresses** section, select **Make IP Addresses Sequential** to make the IP addresses sequential. 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.

You can add more compute-only or converged servers, by clicking **Add Compute Server** or **Add Converged Server**.

- Note**
- When adding a converged node, ensure that the configuration is symmetric across both sites. You can add as many compute nodes as needed. There is no restriction.
 - 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
Locator LED column	Turn on to locate a server.
Name column	Displays the name assigned to the server.
Site column	Displays the site where the server is physically located.
Management Hypervisor field	Enter the static IP address that handles the Hypervisor management network connection between the ESXi host and the storage cluster.
Management Storage Controller field	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 field	Enter the static IP address that handles the Hypervisor data network connection between the ESXi host and the storage cluster.
Data Storage Controller field	Enter the static IP address that handles the storage controller VM data network connection between the storage controller VM and the storage cluster.

Step 2 In the **Advanced Configuration** section, select **Clean up disk partitions** to remove all existing data and partitions from all nodes added to the storage cluster.

- 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.
 - Ensure that you backup any data which must be retained.

Step 3 Click **Start** to begin adding HyperFlex nodes to the Stretch cluster.

Note When adding nodes to the Stretch cluster, add only the nodes already associated with that cluster.



CHAPTER 5

Replication

- [Configuring Replication, on page 29](#)
- [Configuring Replication VLAN in Cisco UCS Manager, on page 29](#)

Configuring Replication

To configure replication on a HyperFlex Controller VM in a Stretch Cluster, you can complete the necessary VLAN configuration on the fabric interconnects using UCS Manager manually. Alternately, you can create a new VLAN in HX Connect. Ensure to use the same VLAN configuration on the UCSM fabric interconnects, on both the sites in a Stretch Cluster.



Attention For configuring the replication network in a Stretch Cluster deployment, see the [Cisco HyperFlex Administration Guide](#) for your release.

Configuring Replication VLAN in Cisco UCS Manager

To configure the replication VLAN using UCS Manager, complete the following steps in Site-A and Site-B:

1. Create a replication VLAN.
2. Associate the replication VLAN to the HX service profile of the HyperFlex cluster.

Creating Replication VLAN



Attention It is best practice to delete unnecessary and stale VLANs. If the VLANs already exist on your fabric interconnect, you do not have to recreate it.

Create the following replication VLANs, if you do not see them in your HyperFlex cluster. Strictly follow the naming conventions that are specified in the table.

VLAN ID	Specify Name	Name Displayed in UCSM
<VLAN ID>	<hx-inband-repl-VLAN ID>	hx-inband-repl-VLAN ID (VLAN ID)

Step 1 Log into Cisco UCS Manager. In the Navigation pane, click **LAN**.

Step 2 On the **LAN** tab, right-click **VLANS** and click **Create VLANs**.

The screenshot shows the Cisco UCS Manager interface. The navigation pane on the left has 'LAN' selected. The main content area shows a list of VLANs, with 'Create VLANs' highlighted in a red box. The right pane shows details for a selected VLAN, including a table of VLANs and a fault summary.

Name	ID
VLAN 710 (710)	710
VLAN 711 (711)	711
VLAN default (1)	1
VLAN hx-inband-mgmt (210)	210
VLAN hx-inband-repl-810..	810
VLAN hx-inband-repl-910..	910

Details

General | Org Permissions | VLAN

Fault Summary

0 0 0

Actions

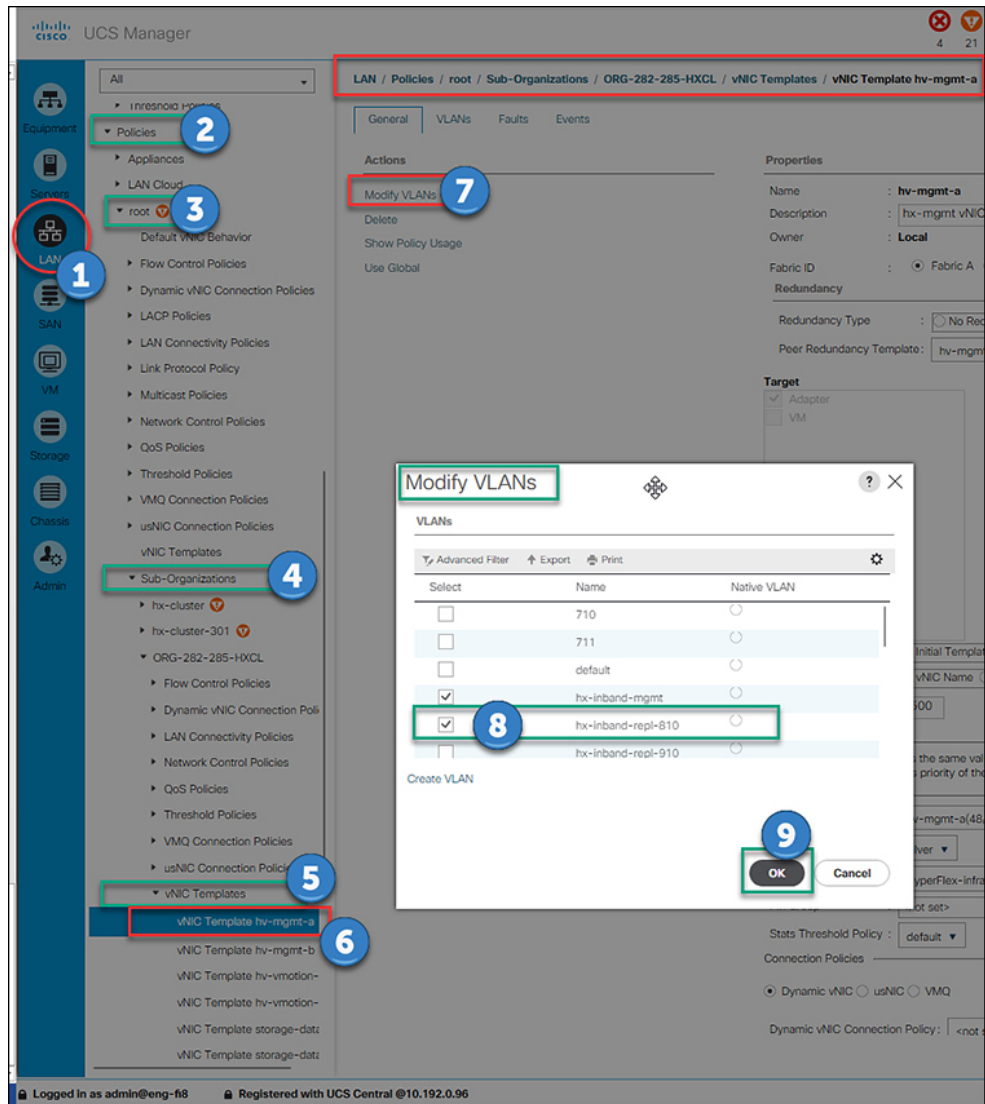
Step 3 In the **Create VLANs** dialog box, complete the following fields for all the VLANs that must be supported on the fabric interconnect:

UI Element	Essential Information
VLAN Name/Prefix field	Enter the VLAN name prefix. For example, hx-inband-repl-810.
Multicast Policy Name field	<not set>
Common/Global check box Fabric A check box Fabric B check box Both Fabric Configured Differently check box	Choose the fabric configuration option.
VLAN IDs field	Enter the VLAN ID, to create a global VLAN that maps to the same VLAN IDs in all available fabrics. For example, 810.
Sharing Type check box	Choose the sharing type: None Primary Isolated Community

Step 4 Click **OK**.

Associating Replication VLAN to a HyperFlex Cluster

To associate the replication VLANs to the HX service profile of the HyperFlex cluster, complete the following steps.



Before you begin

Create VLANs for replication in UCS Manager.

- Step 1** Log into Cisco UCS Manager. In the Navigation pane, click **LAN**.
- Step 2** Navigate to **Policies > root > Sub-Organizations > hx-cluster > vNIC Templates > vNIC Template hv-mgmt**.
- Step 3** Select **Modify VLANs** on the work pane.
- Step 4** Select the VLANs to associate with the HyperFlex cluster.

Step 5 Click **OK**.



CHAPTER 6

Managing HyperFlex Stretch Clusters

- [Overview, on page 35](#)
- [Monitoring the Health of a Stretch Cluster, on page 35](#)
- [Viewing System Information, on page 37](#)
- [Creating a Datastore, on page 42](#)
- [Shutting a Stretch Cluster Site Down Gracefully, on page 42](#)
- [Changing the Preferred Site, on page 43](#)

Overview

You can view do the following for the sites and the witness VM or Invisible Cloud Witness node in a stretch cluster using HX Connect:

- View the Operational Status and Resiliency Status of the HyperFlex cluster on both the sites and the witness VM or Invisible Cloud Witness node.
- View the Functional status of the sites and the witness VM or Invisible Cloud Witness node, the IP address of the witness VM, and the status of the Invisible Cloud Witness node. View HX storage cluster system-related information, including node and disk data.
- Enter/Exit HX maintenance mode.
- Associate a datastore with one of the sites in a stretch cluster, when creating the datastore.

Monitoring the Health of a Stretch Cluster

You can view the Operational Status and Resiliency Status of the HyperFlex cluster on both the sites and the witness VM or Invisible Cloud Witness node in HX Connect on the Dashboard page.



Important

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



Step 1 Log into HX Connect.

- a) Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
- b) Enter the administrative username and password.
- c) Click **Login**.

Step 2 On the left navigation pane, click **Dashboard**.

Step 3 On the **Dashboard** you can view the following details for the HyperFlex Stretch Cluster:

Displays a status summary of your HX storage cluster for the sites across the Stretch cluster.

UI Element	Essential Information
Operational Status section	<p>Provides the functional status and application performance of the HX storage cluster on Site A and Site B.</p> <p>Provides the functional status and application performance of the witness VM or Invisible Cloud Witness node.</p> <p>Click Information () to access the HX storage cluster name and status data.</p>
Resiliency Health section	<p>Provides the data health status and the ability of the HX storage cluster on Site A and Site B to tolerate failures.</p> <p>Click Information () to access the resiliency status, replication and failure data. This also provides information about data replication compliance, caching device failures tolerable, and device failures tolerable on each node in both the sites.</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 and the division of converged versus compute nodes across Site A and Site B in the Stretch Cluster.</p> <p>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>
Cluster Time field	System date and time for the cluster.

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 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 out 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: <i>cvs</i> , <i>xls</i> , and <i>doc</i> . To export content from other pages in the table, scroll to the bottom, click through the page numbers, and apply the export.

Viewing System Information

On the **System Information** page, you can view HX storage cluster system-related information, including node and disk data. You can also Enter or Exit Maintenance Mode for the sites.

Step 1 Log into HX Connect.

- Enter the HX Storage Cluster management IP address in a browser. Navigate to *https://<storage-cluster-management-ip>*.
- Enter the administrative username and password.
- Click **Login**.

Step 2 On the left navigation pane, select **System Information**.

Step 3 Under the **System Overview** tab you can view the following information for both the sites and the witness VM or Invisible Cloud Witness node:

HX Storage Cluster Configuration Data tab

Displays the basic configuration information the HX storage cluster on the Stretch cluster sites.

Field	Description
HX storage cluster field	Name of this storage cluster.

Field	Description
Cluster License Status section	<p>Displays the Register Now link when you log into the HX storage cluster for the first time or till the HX storage cluster license is registered:</p> <p>Register Now link—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>Note To register a cluster license, you can also choose Register Cluster from the Actions drop-down field.</p>
License section	<ul style="list-style-type: none"> • License Type—Displays Evaluation, Edge, Standard, or Enterprise as the HX storage cluster license type. • License Status—Displays one of the following as the HX storage cluster license status: <ul style="list-style-type: none"> • In compliance • License expires in <n> days. Cluster not registered - Register Now. (This status appears for Evaluation type license) • License expired. Cluster not registered - Register Now. (This status appears for Evaluation type license) • Out of compliance - License expired on <date> • Out of compliance - Insufficient license • Certificate expired—This status appears when ID certificate is not renewed for more than six months. • 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. <p>Note To refresh license certificate or renew license authorization, choose the respective options from the Actions drop-down field.</p>
HX storage cluster status field	<p>Provides functional status of the HX storage cluster in Site A and Site B:</p> <ul style="list-style-type: none"> • Online—Cluster is ready. • Offline—Cluster is not ready. • Read Only—Cluster is out of space. • Unknown—Transitional state while the cluster is coming online.
vCenter link	<p>Secure URL to the VMware vSphere associated with this HX storage cluster. Click the link to remotely access the vSphere Web Client.</p>
Hypervisor field	<p>Hypervisor version installed on this HX storage cluster.</p>

Field	Description
HXDP Version field	Installer package version installed on this HX storage cluster.
Data Replication Factor field	Number of the redundant data replicas stored on this HX storage cluster.
Uptime field	Length of time this HX storage cluster has been online.
Total Capacity field	Overall storage size of this cluster.
Available Capacity field	Amount of free storage in this cluster.
DNS Server(s) field	IP address for the DNS server(s) for this HX storage cluster.
NTP Server(s) field	IP address for the NTP server(s) for this HX storage cluster.
Witness IP Address field	Witness VM: Provides the IP address Invisible Cloud Witness: Shows the status and (INTERSIGHT)

Step 4 Use **Actions** list to access the controller VM using SSH as an administrator and perform actions in the following table.

Note Actions to enable or disable SSH can only be performed by **domain** users, and not local users. Domain users are users in VC (ESXi).

UI Element	Essential Information
Disable Controller Access over SSH	Secure Shell (SSH) is disabled by default.
Register Now	Register your license.
Re-register vCenter	Re-register your license via vCenter
Check Secure Boot Status	Verify your Secure Boot Status
Change Preferred Site	The site designated to serve requests in the event that the connectivity between sites is down. In the Change Preferred Site window, use the radio button to select the preferred site, and click Update to confirm your selection.

Step 5 Under the **Nodes** tab, you can view the following information:

Displays data about individual nodes in this HX storage cluster. To see this information in tabular format, go to the **Nodes** page.

UI Element	Essential Information
Node field	Name of a node on this cluster.
Hypervisor Address field	IP address for the management network for this HX storage cluster.

UI Element	Essential Information
Hypervisor Status field	<ul style="list-style-type: none"> • Online • Offline • In Maintenance • In Progress
Controller Address field	IP address of the controller VM on this HX storage cluster.
Controller Status field	Status of the controller VM on this HX storage cluster. <ul style="list-style-type: none"> • Online • Offline • In Maintenance • Healthy • Warning
Model field	Physical hardware model number of this node.
Version	Installer package version installed on this node.
Disks field	Number of caching versus persistent disks in this node.
Command Line Mode field	Valid values are: Root shell or Admin shell .
Site	The site the node belongs to.

Step 6 Under the **Disks** tab, you can view the following information:

For nodes with disks, an interactive display of disks is included with the following pop-up data:

Table 5: Caching Disks

UI Element	Essential Information
Slot Number field	Location of the drive.
Storage Usage field	Percentage of disk storage used.
Status field	<ul style="list-style-type: none"> • Claimed • Available • Ignored • Blocked • Ok to Remove • Unknown

UI Element	Essential Information
Locator LED	Activates a physical light on the host to help locate a disk; options are On and Off .
Capacity field	Total disk size.
Serial Number field	Physical serial number of this disk.
Version	Disk version
Disk Drive Interface	Describes how the disk is connected. SATA is one example.

Table 6: Persistent Disks

UI Element	Essential Information
Status field	<ul style="list-style-type: none"> • Claimed • Available • Ignored • Blocked • Ok to Remove • Unknown
Locator LED action	Activates a physical light on the host to help locate a disk; options are On and Off .
Used / Total Capacity field	Amount of the disk used versus the total disk size.
Serial Number field	Physical serial number of this disk.
Storage Usage field	Percentage of disk storage used.
Version	Disk version
Disk Drive Interface	Describes how the disk is connected. SATA is one example.
Slot Number field	s Location of the drive.
Type field	Describes the disk type. Values include: Rotational or Solid State.
Usage field	Describes how the disk is being used: Persistent, System, or Cache.

Creating a Datastore

To associate a datastore with one of the sites in a Stretch cluster, do the following:

-
- Step 1** Log into HX Connect.
- Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
 - Enter the administrative username and password.
 - Click **Login**.
- Step 2** On the left navigation page, click **Datastores**.
- Step 3** In the work pane, click **Create Datastore**.
- Step 4** Enter a datastore name and capacity.

UI Element	Essential Information
Datastore Name field	Enter a unique datastore name for this HX Storage Cluster.
Size field	Enter the quantity of the datastore. Select the unit of measure. Options are: GB and TB . Ensure it is sufficient to support the virtual machines in this HX Storage Cluster.
Block Size	Select a block size. <ul style="list-style-type: none"> 8K—Default 4K
Site Affinity	Select the radio button to associate the datastore with the site.

- Step 5** Click **Create Datastore** to complete the action.
- HX Data Platform creates a datastore and mounts it on every node in this HX Storage Cluster.
-

Shutting a Stretch Cluster Site Down Gracefully

There may be a need to shut a Stretch Cluster site down (for example, when relocating equipment, routine maintenance, site move, site work, failure testing etc.). If you have removed the VMware EAM dependency from your cluster (default in HX 4.0(2b) and later releases), then you can use HX maintenance mode in vCenter or through HX Connect to shutdown the control VMs on each node in a site. When this occurs, the site will shut down gracefully and the guest VMs will failover to the surviving site. You can then also power down ESXi on the nodes if you need to. If EAM is enabled on your CVMs, and you require more information on shutting down a site, see Appendix A in [Operating Cisco HyperFlex HX Data Platform Stretch Clusters](#).

Changing the Preferred Site

The preferred site is designated to serve requests in the event that the connectivity between sites is down. You can change the preferred site selection after install from HX Connect by performing the following steps:

-
- Step 1** Log into HX Connect
- Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
 - Enter the administrative username and password.
 - Click **Login**.
- Step 2** On the left navigation pane, select **System Information**.
- The preferred site is identified with a badge near the Site Name.
- Step 3** Under the **System Overview** tab you can view the following information for both the sites and the witness node:**HX Storage Cluster Configuration Data** tab and select **Actions** from the list.
- Step 4** In the **Change Preferred Site** window, use the radio button to select the preferred site
- Step 5** Click **Update** to confirm your selection.
-




CHAPTER 7

Troubleshooting

- [Viewing Resiliency Status in HX Connect, on page 45](#)
- [Troubleshooting the Network Configuration, on page 47](#)
- [Troubleshooting for Site-to-Site Failover, on page 48](#)

Viewing Resiliency Status in HX Connect

The Dashboard page in HX Connect displays the status summary of your HX storage cluster. Click **Information** () to access the resiliency status. Resiliency status is the data resiliency health status and the ability of the storage cluster to tolerate failures.

To access the resiliency status in HX Connect:

1. Log into HX Connect.
 - a. Enter the HX Storage Cluster management IP address in a browser. Navigate to `https://<storage-cluster-management-ip>`.
 - b. Enter the administrative username and password. Click **Login**.
2. On the left navigation pane, click **Dashboard**.
3. On the **Dashboard** page, you can view the resiliency status for the HyperFlex Stretch Cluster:

Color coding and icons are used to indicate various status states. Click an icon to display additional information, such as reason messages that explain what is contributing to the current state. The various Resiliency Status states are:

- **Healthy**—The cluster is healthy with respect to data and availability.
- **Warning**—Either the data or the cluster availability is being adversely affected.
- **Unknown**—A transitional state while the cluster is coming online.



Important

The disk failures that are mentioned below are for the converged nodes only (the witness VM or Invisible Cloud Witness node and the compute nodes do not have disks hosting user data).

Independent and Non-Simultaneous Failure Scenarios

Failure Scenario	Expected Behavior	Resiliency Status in HX Connect
One node failure	VMs will failover to the remaining nodes on the same site as long as those nodes can accommodate the VM resources.	Warning—Cluster shows status of unhealthy until the cluster recovers.
All node failure on any one site	User VMs failover to other site. Note For the recovery operation to succeed after a site failure, the witness VM or Invisible Cloud Witness node must be online.	Warning—HX Connect shows the site failure details. Cluster shows status of unhealthy until the cluster recovers.
One disk failure	Cluster recovers after failover.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Two disk failure on a single site (one disk on each node) simultaneously	Same as one disk failure.	Warning—Cluster shows status of unhealthy until the cluster recovers. Same as one disk failure.
Witness VM or Invisible Cloud Witness failure	Cluster remains online.	Witness VM :No visible indication. Invisible Cloud Witness: Online/Offline/Unknown is shown on the System Information Page
vCenter failure (platform impact)	Cluster remains online.	No visible indication.
Network isolation between a given site and a witness VM or Invisible Cloud Witness	Cluster remains online.	No visible indication.
Network isolation between sites	All user VMs from one site will fail over to the other site. The VMs from the other site will continue to run. Note For the recovery operation to succeed after a site failure, the witness VM or Invisible Cloud Witness node must be online.	Warning—HX Connect shows the site failure details. Cluster shows status of unhealthy until the cluster recovers.

Failure Scenarios with witness VM or Invisible Cloud Witness Failure


Failure Scenario	Expected Behavior	Visible Indications in HX Connect
Disk failure	Cluster remains online.	Warning—Cluster shows status of unhealthy until the cluster recovers.

Failure Scenario	Expected Behavior	Visible Indications in HX Connect
Single node failure	Cluster remains online. Cluster becomes unhealthy.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Link failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness VM or Invisible Cloud Witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Either fabric interconnect pair failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness VM or Invisible Cloud Witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Switch failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness VM or Invisible Cloud Witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.
Site power failure	Cluster halts resulting in All Path Down on both sites, until either the second failure is healed or the witness VM or Invisible Cloud Witness is restored.	Warning—Cluster shows status of unhealthy until the cluster recovers.

Troubleshooting the Network Configuration

Do the following to verify if the network configuration prerequisites are met:

- Ensure that the nodes physically residing on different sites are pingable, once the stretch VLANs are created.
- Ensure that the Round-Trip Time (RTT) and bandwidth requirements are met across all the sites. You can use iPerf and Ping or any Cisco supported tool to measure the RTT and bandwidth.



Troubleshooting for Site-to-Site Failover

Do the following during site-to-site failover (if needed):

- During site-to-site failure, depending on the time it takes to failover, VMs may see a transient All Paths Down (APD) event from the ESX layer. This APD is expected when the witness VM or Invisible Cloud Witness connection is slow.
- In some cases, you may need to manually restart the user VMs if they have timed-out during site-to-site failover.

Do not try to resolve the APD by rebooting the nodes. Depending on the size of the cluster and amount of data, the failover can take 20-30 minutes or more if the witness VM or Invisible Cloud Witness transaction latencies are high.



CHAPTER 8

Post Installation

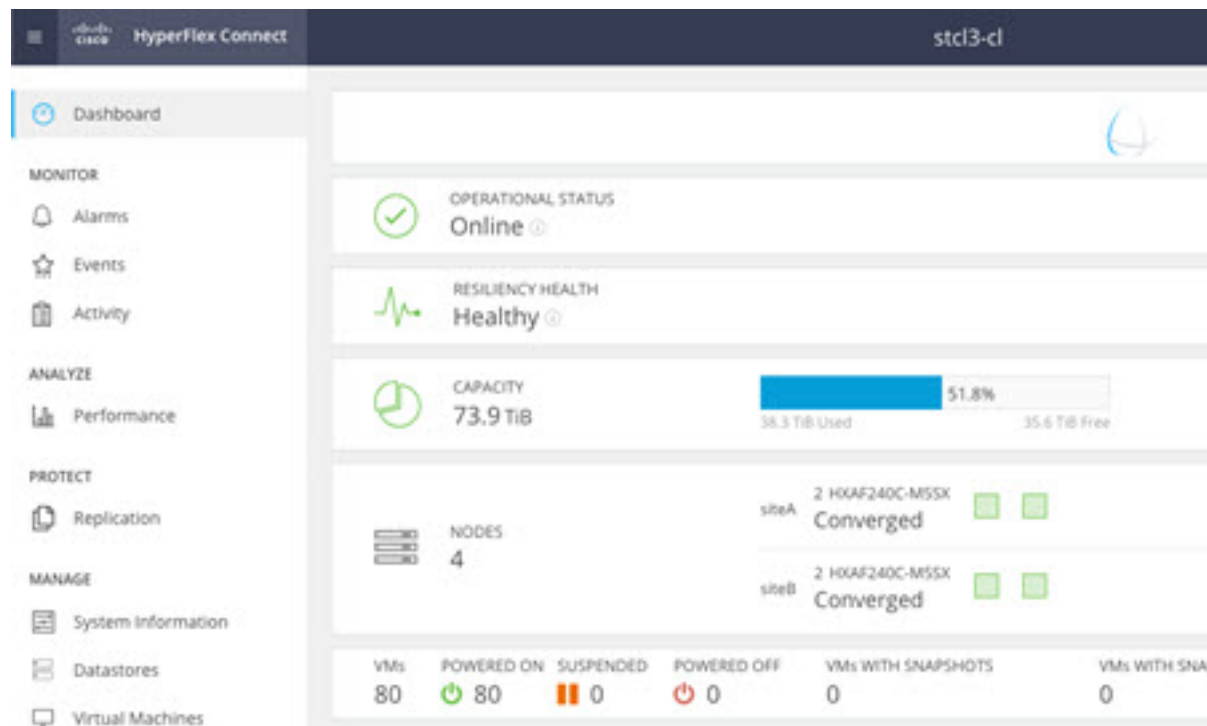
- [Confirm That the Installation is Complete](#), on page 49
- [Running Post Install Script](#), on page 51

Confirm That the Installation is Complete

Post Installation Checklist

After a successful installation. It is recommended that you validate the following in HX Connect:

- Dashboard Panel



- Operational Status Online
- Resiliency Health: Healthy

- Nodes: Correct number of nodes, site names, and nodes per site
- System Information Panel

The screenshot displays the Cisco HyperFlex Connect interface for a cluster named 'stcl3-cl'. The left sidebar shows navigation options: Dashboard, MONITOR (Alarms, Events, Activity), ANALYZE (Performance), PROTECT (Replication), and MANAGE (System Information, Datastores, Virtual Machines, Upgrade, Web CLI). The main content area is titled 'System Overview' and includes a warning: 'Cluster not registered with Cisco Licensing. Register Now'. Below this, cluster details are shown: 'stcl3-cl ONLINE', License Type: License Status, Evaluation: License expires in 86 days, Cluster not registered with Cisco Licensing, vCenter: https://mbodge-1.nc7.eng.storvisor.com, Hypervisor: 7.0.3-21424296, vSphere: 8 days, 13 hours, 55 minutes, 51 seconds, HADP version: 5.5.1a-43230, Witness: Online (INTERSIGHT). The 'Hyperconverged Nodes' section is divided into 'Site A (PREFERRED SITE)' and 'Site B'. Site A nodes include ucs798 and HXAF240C-M55X, both with Online Hypervisor and HyperFlex Controller. Site B node is ucs801, also Online. Disk usage is shown as 24 in use for each node.

Node	Hypervisor	HyperFlex Controller	Disk Overview (24 in use)
ucs798	Online	Online	1 2 3 4 5 6 7 8
HXAF240C-M55X	10.64.64.45 7.0.3-21424296	10.64.64.86 5.5.1a-43230	1 2 3 4 5 6 7 8
ucs799	Online	Online	1 2 3 4 5 6 7 8
HXAF240C-M55X	10.64.64.46 7.0.3-21424296	10.64.64.87 5.5.1a-43230	1 2 3 4 5 6 7 8
ucs801	Online	Online	1 2 3 4 5 6 7 8

- Cluster name displayed is correct
- Cluster status is shown as Online
- Witness: Online (INTERSIGHT) or (Witness VM IP)
- Preferred Site badge on configured site.
- Nodes are shown in expected sites
- Hypervisor and Hyperflex Controllers are Online

Verifying the AuxZK IP

To verify the AuxZK IP, run the command `stcli cluster info` command.


```
clusterAccessPolicy: lenient
auxZkIp: 10.64.72.99
zoneType: physical
size: 4
clusterType: stretch_cluster
```

Running Post Install Script

After the installation of a Stretch Cluster using the HX Data Platform Installer, run the post installation script to finalize the configuration and set the vMotion network up. You can also run this script at a future time if needed.

1. Log into a Cluster IP (CIP) through an SSH server using admin login.
2. Run the `hx_post_install` script.
3. Follow the prompts and enter the required information.

