Cisco HyperFlex Systems Stretched Cluster Guide, Release 5.0

First Published: 2021-11-22

Last Modified: 2023-04-07

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

CONTENTS

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.

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

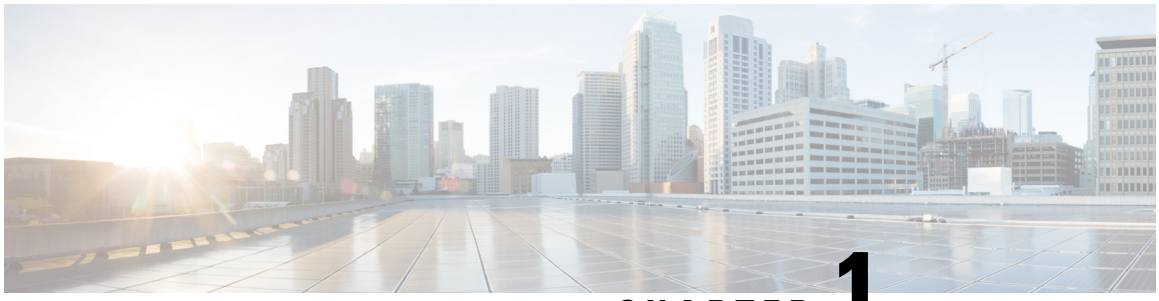

CHAPTER 1

Introduction

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.

Stretch Cluster Witness VM Model

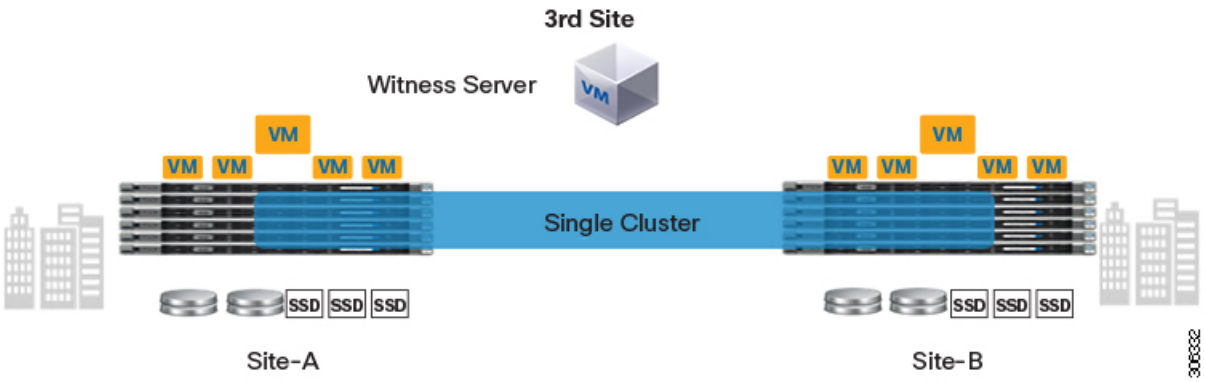

This document describes the deployment process for stretch clusters, gives information about monitoring stretch clusters, and discusses guidelines to follow when deploying stretch clusters.

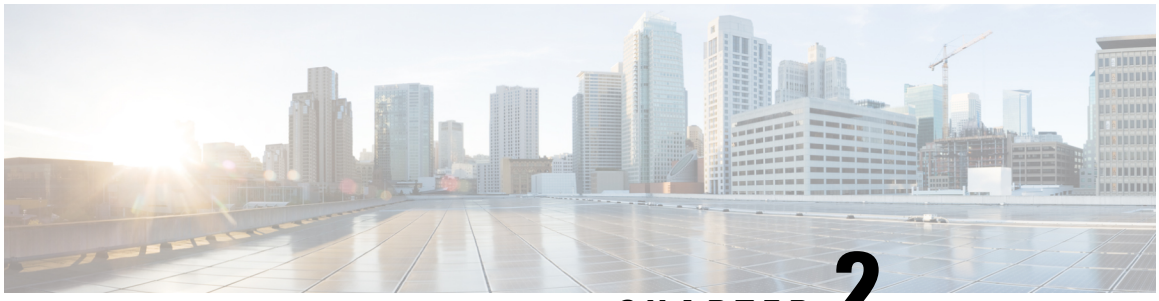

CHAPTER 2

Preinstallation Checklist

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.

- Network connectivity for the install powered by Installer VM.

Note 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.

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 Stretched L2. One subnet for data traffic and one for management traffic, with the management subnet reachable from vCenter and the witness node.	✓	✓
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.	✓	✓
QoS If the HyperFlex cluster is a stretched 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. If the HyperFlex cluster is a stretched 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. 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. 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.	✓	✓

Network Topology

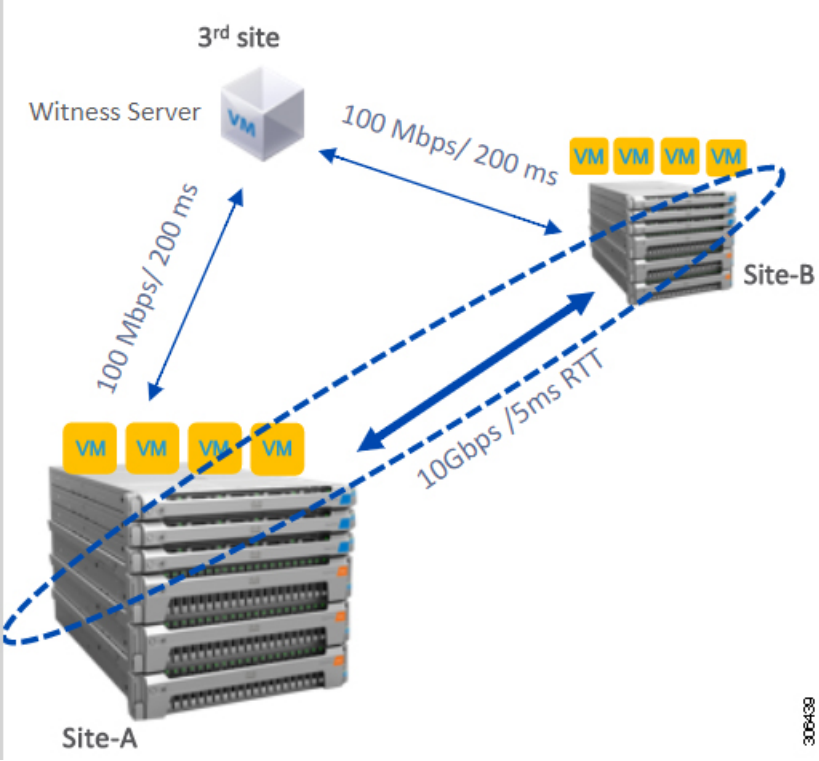

Witness VM Requirements

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 Stretched 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 Stretched Cluster is highly recommended to ensure proper failover behavior and guarantee high performance during normal operations. While it is possible to run Stretched 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 Stretched 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 Stretched 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.	✓	✓
Stretched VLANs across both sites are required.	✓	✓

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

Deploying the Witness VM Node

Important

- 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.

Note If the IP address is requested from a DHCP server please verify that the clustering service is running after deploying the Witness VM, perform the following steps:

- Check for the clustering service by running the following command after every reboot of the Witness VM:

  ```
  service exhibitor status
  ```

- If the output does not show the service as running, then restart exhibitor using the following command after ensuring that the witness VM does have an IP address:

  ```
  service exhibitor restart
  ```

The following procedure details the steps to follow, for deploying a witness VM virtual node on a physical ESXi host.

Before you begin

This section is for use only with HyperFlex Stretch Cluster Installed with HXDP Release 5.0(x) or earlier.

Attention

- HyperFlex witness VM node version 1.1.3 is supported in Cisco HXDP Release 4.5(2a) and later. To review the recommended version for your specific release, see the *HX Data Platform Software Versions for HyperFlex Witness Node* sections of the *Cisco HyperFlex Software Requirements and Recommendations* guide.

- Ensure that the virtual network on this ESXi host is reachable from both the stretch cluster sites.
- Download the HyperFlex witness VM node on to your desktop or host that runs vSphere Web Client from Download Software.

```
Example:
HyperFlex-Witness-1.0.2.ova
```

- High Availability is optional for the witness VM node.

Step 1 Log into vSphere Web Client. Choose the ESXi server where the witness VM should be deployed. Right-click the ESXi host and select **Deploy OVF Template**.

Step 2 Browse and select the *HyperFlex-Witness.ova* file. Click **Next**.

Step 3 Specify a unique name for the witness VM node in the **Virtual Machine Name** field. Select a location for the virtual machine from the drop-down list. Click **Next**.

Step 4 From the **Select a compute resource** drop-down list, choose the ESXi host where you want to deploy the witness VM node. Click **Next**.

Step 5 In the **Review details** pane, verify the template details. Click **Next**.

Step 6 In the **Select Storage** pane, do the following:

Field	Description
Select virtual disk format drop-down list	• Thick Provision Lazy Zeroed • Thick Provision Eager Zeroed • Thin Provision
VM Storage Policy drop-down list	Datastore Default

Select the datastore where the virtual machine will reside. Ensure that this datastore has at least 40 GB of available free space. Click **Next**.

Step 7 In the **Select Networks** pane, select a *Destination Network* port group, where the witness VM has to connect. Click **Next**.

Step 8 On the **Customize Template** page, complete the fields that are relevant for your configuration. If no values are entered, the VM uses DHCP server provided network configuration parameters.

Field	Description
Static IP Address field	The IP address for Witness VM. If the DHCP server is used to define the network configuration, the IP address needs to be dedicated to Witness VM. Leave blank if DHCP is desired.
Netmask field	The netmask or prefix for this interface. Leave blank if DHCP is desired.
Default Gateway field	The default gateway address for this VM. Leave blank if DHCP is desired.
DNS field	The domain name servers for this VM (comma separated). Leave blank if DHCP is desired.
NTP field	NTP servers for this VM (comma separated) to sync time. Leave blank if DHCP is desired.

Click **Next**.

Step 9 On the **Ready to complete** page, verify all the details entered. Click **Finish**.

Step 10 Repeat this process for each cluster.

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
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

Management Network IP Addresses (must be routable)			Data Network IP Addresses (does not have to be routable)	
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

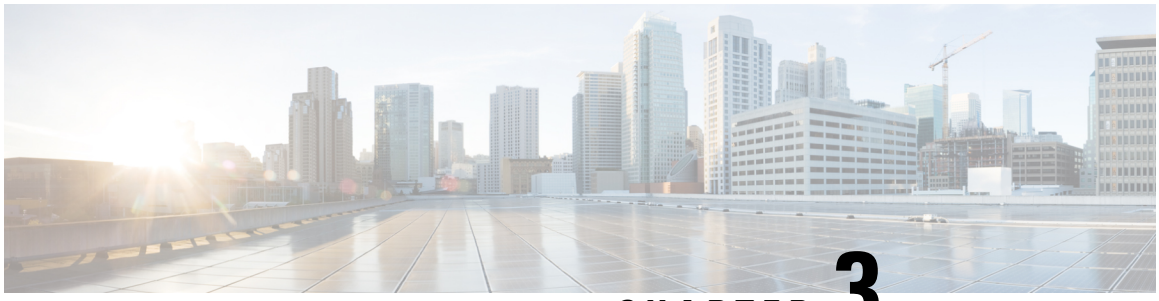

CHAPTER 3

Guidelines and Limitations

Guidelines

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.
- 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. M4/M5, M4/M6, M5/M6 is supported when reached by expansion.
- Shared Witness VM is not supported for Stretch Cluster deployments.

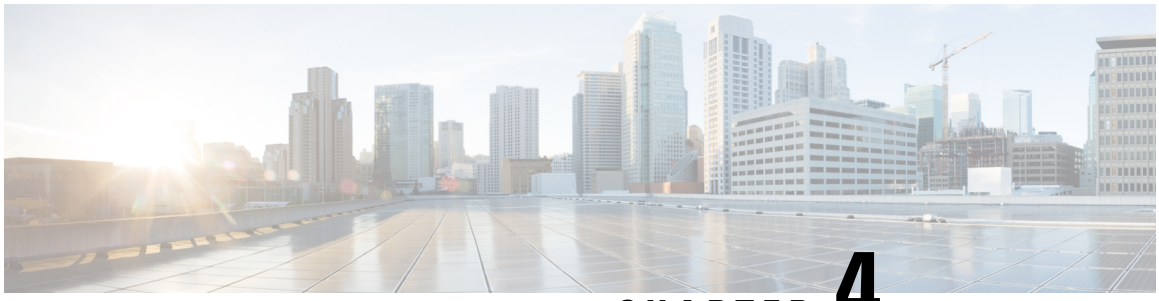

CHAPTER 4

Installation

Installation Overview

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

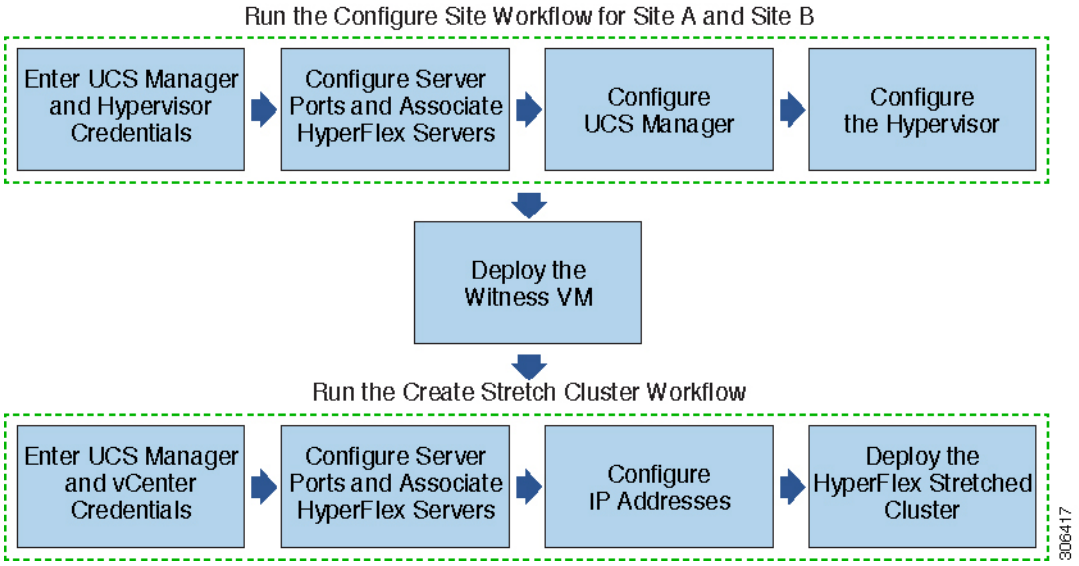

Follow this workflow during installation:

Step	Description	Reference
Create the Stretch Cluster sites—Run the Configure Site workflow for Site A and Site B.	Log into the HX Data Platform Installer. Enter UCS Manager credentials and Hypervisor credentials for both the sites.	Enter Credentials, on page 17
	Configure the server ports and associate HyperFlex servers.	Associate HyperFlex Servers, on page 19
	Configure VLAN, MAC Pool, 'hx-ext-mgmt' IP Pool for Out-of-Band CIMC, inband CIMC, iSCSi Storage, and FC Storage.	Configure UCS Manager, on page 20
	Configure the Hypervisor.	Configure Hypervisor: HX Release 5.0(x) and Earlier, on page 23
Download and deploy the Witness VM.	**Note** A witness VM is mandatory in a HyperFlex Stretch Cluster environment.	Deploying the Witness VM Node, on page 8
Create your HyperFlex Stretch Cluster—Run the Create Stretch Cluster workflow.	Enter UCS Manager Credentials for Site A and Site B, and vCenter credentials.	Enter Credentials, on page 25
	Configure the server ports and associate HyperFlex servers.	Associate HyperFlex Servers, on page 27
	Configure IP addresses.	Configure IP Addresses, on page 28
	Deploy the HyperFlex Stretch Cluster.	Deploy a HyperFlex Stretched Cluster, on page 29

Create the Stretch Cluster Sites

Before you begin

- Complete the prerequisites. Refer to the Common Site Interlink and Witness Requirements, on page 4 for more details.
- Download and deploy the Witness VM for each cluster.

Use the following procedure to install and configure a Cisco HyperFlex Stretch Cluster using the HX Data Platform Installer.

Step 1 Log into the HX Data Platform Installer.

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

Step 3 On the **Credentials** page, enter UCS Manager and Hypervisor credentials. See Enter Credentials, on page 17 for more details.

Step 4 On the **Server Selection** page, configure the server ports and associate HyperFlex servers. See Associate HyperFlex Servers, on page 19 for more details.

Step 5 On the **UCSM Configuration** page, configure VLAN, MAC Pool, 'hx-ext-mgmt' IP Pool for Out-of-Band CIMC, iSCSi Storage, and FC Storage. See Configure UCS Manager, on page 20 for more details.

Step 6 On the **Hypervisor Configuration** page, configure the subnet mask, gateway, and complete the hypervisor settings. See Configure Hypervisor: HX Release 5.0(x) and Earlier, on page 23 for more details.

Step 7 Click **Configure Site** to begin cluster creation. 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 HyperFlex stretch cluster. See Create Your HyperFlex Stretch Cluster, on page 24 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 installation prerequisites. Refer to the Common Site Interlink and Witness Requirements, on page 4 for more details.
- Download and deploy the Witness VM.

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.

a) In a browser, enter the URL for the VM where HX Data Platform Installer was installed.

b) 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.

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

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

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**.

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, *10.193.211.120*.
UCS Manager User Name field	Enter the administrative level username. For example, <*admin*>.
Password field	Enter the administrative level password. For example, <*root*>.
Site Name field	Specify a unique site name.

Step 6 Enter the following credentials for the Hypervisor:

Hypervisor Credentials

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

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.
Status column	• 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	• **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 17.

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 20 .

Configure UCS Manager

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

Before you begin

Associate servers on the HyperFlex cluster. See Associate HyperFlex Servers, on page 19.

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 field	hx-inband-mgmt
VLAN ID field	Default—3091
VLAN for HyperFlex storage traffic	
VLAN Name field	hx-storage-data
VLAN ID field	Default—3092
VLAN for VM vMotion	
VLAN Name field	hx-vmotion
VLAN ID field	Default—3093
VLAN for VM Network	
VLAN Name field	vm-network
VLAN ID(s) field	Default—3094 A comma-separated list of guest VLANs.

Step 2 In the **MAC Pool** section, configure the **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 Out-of-Band CIMC** section, complete the following fields:

Field	Description
IP Blocks field	The range of IP addresses designated for the HyperFlex nodes. The IP addresses can be a comma-separated list of values for the guest VLANs. For example, *10.193.211.124-127, 10.193.211.158-163*.
Subnet Mask field	Set the subnet to the appropriate level to limit and control IP addresses. For example, *255.255.0.0*.
Gateway field	Enter the IP address. For example, *10.193.0.1*.

Step 4 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 field	Name of the VLAN associated with the iSCSI vNIC, on the primary fabric interconnect (FI-A).
VLAN A ID field	ID of the VLAN associated with the iSCSI vNIC, on the primary fabric interconnect (FI-A).
VLAN B Name field	Name of the VLAN associated with the iSCSI vNIC, on the subordinate fabric interconnect (FI-B).
VLAN B ID field	ID of the VLAN associated with the iSCSI vNIC, on the subordinate fabric interconnect (FI-A).

Step 5 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 field	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 field	The name of the VSAN for the primary fabric interconnect (FI-A). Default—`hx-ext-storage-fc-a`.
VSAN A ID field	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 field	The name of the VSAN for the subordinate fabric interconnect (FI-B). Default—`hx-ext-storage-fc-b`.
VSAN B ID field	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 6 In the **Advanced** section, do the following:

Field	Description
UCS Server Firmware Version drop-down list	Select the UCS server firmware version to associate with the HX servers from the drop-down list. The UCS firmware version must match the UCSM version. See the latest Cisco HX Data Platform Release Notes for more details. For example, *3.2(1d)*.
HyperFlex Cluster Name field	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 field	Specify a unique *Org Name* to ensure isolation of the HyperFlex environment from the rest of the UCS domain.

Step 7 Click **Continue** to configure the Hypervisor. See Configure Hypervisor: HX Release 5.0(x) and Earlier, on page 23.

Configure Hypervisor: HX Release 5.5 and later

Starting with HXDP release 5.5(1a) all new stretch cluster installations auto-configure Intersight Arbitrator via HyperFlex Classic installer without the need of an external witness service in a third site. (Stretch Cluster deployments from Intersight is not supported).

If your installation is before 5.5(1a) continue to the Configure Hypervisor: HX Release 5.0(x) and Earlier, on page 23section of this guide.

Configure Hypervisor: HX Release 5.0(x) and Earlier

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. 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 Stretch Cluster. See Configure UCS Manager, on page 20.

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, *255.255.0.0*.
Gateway field	IP address of gateway. For example, *10.193.0.1*.
DNS Server(s) field	IP address for the DNS Server. **Note** • 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 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 Click **Configure Site** to begin cluster creation. 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 Create Stretch Cluster workflow from the HX Data Platform Installer to create your HyperFlex Stretch Cluster. See Create Your HyperFlex Stretch Cluster, on page 24 for more details.

Create Your HyperFlex Stretch Cluster

Use the following procedure to create a Cisco HyperFlex Stretch Cluster using the Cisco HX Data Platform Installer.

Before you begin

- Complete the prerequisites. Refer to the Common Site Interlink and Witness Requirements, on page 4 for more details.
- Ensure that you run the *Configure Site* workflow for both sites.
- Download and deploy the Witness VM.

Step 1 Log into the Cisco HX Data Platform Installer.

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

Step 3 On the **Credentials** page, select **Create Stretch Cluster**. Enter UCS Manager credentials for site 1 and site 2. Enter vCenter credentials. Click **Continue**. See Enter Credentials, on page 25 for more details.

Step 4 On the **Server Selection** page, configure the server ports and associate HyperFlex servers. See Associate HyperFlex Servers, on page 27 for more details.

Step 5 On the **IP Addresses** page, configure the IP addresses to be applied to the storage cluster. See Configure IP Addresses, on page 28 for more details.

Step 6 On the **Cluster Configuration** page, configure and deploy the HyperFlex Stretch cluster. See Deploy a HyperFlex Stretched Cluster, on page 29 for more details.

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

What to do next

Click **Launch HyperFlex Connect** to manage your HyperFlex Stretch Cluster.

Enter Credentials

Enter UCS manager credentials for site 1 and site 2, and enter credentials for the vCenter.

Before you begin

- Complete the prerequisites. Refer to the Common Site Interlink and Witness Requirements, on page 4 for more details.
- Download and deploy the Witness VM.
- Complete configuring the HyperFlex Stretch Cluster. See Create the Stretch Cluster Sites, on page 16 for more details.

Step 1 Log into the Cisco HX Data Platform Installer.

a) 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 page, verify the Cisco 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`

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.

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

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

Step 3 On the **Credentials** page, select **Create Stretch Cluster**.

Step 4 Enter the following UCS Manager credentials for Site 1 and Site 2.

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**.

UCS Manager Credentials for Site 1

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

Step 5 Enter the following credentials for the vCenter.

vCenter Credentials

Field	Description
vCenter Server field	Enter the vCenter server FQDN or IP address. For example, *10.193.211.120*. **Note** • 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.

Field	Description
User Name field	Enter the administrative username. For example, *administrator@vsphere.local*.
Admin Password field	Enter the administrative level password. Enter the <*root*> password.

Step 6 Click **Continue** to begin associating HyperFlex servers. See Associate HyperFlex Servers, on page 27.

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.
Status column	• 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	• **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 25.

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 IP addresses. See Configure IP Addresses, on page 28.

Configure IP Addresses

Before you begin

Associate servers on the HyperFlex cluster. See Associate HyperFlex Servers, on page 27.

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.

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

Field	Description
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 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 field	Enter the management network IP address for the HX Data Platform storage cluster.

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

Step 4 Click **Continue** to deploy the HyperFlex Stretch cluster. See Deploy a HyperFlex Stretched Cluster, on page 29 for more details.

Deploy a HyperFlex Stretched 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 28.

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

Field	Description
Cluster Name field	Specify a name for the HX Data Platform storage cluster.
Replication Factor drop-down list	Specify the number of redundant replicas of your data across the storage cluster. Set the replication factor to 4, for 2 copies in each site.

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 field	Enter the vCenter datacenter name for the Cisco HyperFlex cluster.
vCenter Cluster Name field	Enter the vCenter cluster name.

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

Field	Description
DNS Server(s) field	Enter a comma-separated list of IP addresses of each DNS server.
NTP Server(s) field	Enter 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 field	Enter the DNS FQDN or IP address.
Time Zone drop-down list	Select 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 **Auto Support** section, select **Enable Connected Services** to enable Auto Support and Cisco Intersight management.

Field	Description
Enable Connected Services (Recommended) check box	Select to enable Auto Support and Cisco Intersight management. Log on to HX Connect to configure these services or selectively turn them **On** or **Off**.
Send service ticket notifications to field	Enter the email address where SR notifications should be sent when triggered by Auto Support.

Step 6 In the **Advanced Networking** section, complete the following fields:

Note When repurposing existing HyperFlex hardware for stretched cluster, the VM Network port group may not have configured with the correct vSwitch. You must manually configure the VM network on the correct vSwitch.

Field	Description
Management VLAN Tag - Site 1 field	Enter the correct VLAN ID if using trunk ports. Enter 0 if using access ports. **Note** Use the same Management VLAN Tag for site 1 and site 2.
Management VLAN Tag - Site 2 field	Enter the correct VLAN ID if using trunk ports. Enter 0 if using access ports. **Note** Use the same Management VLAN Tag for site 1 and site 2.
Management vSwitch field	Default is vswitch-hx-inband-mgmt. **Note** Do not modify the name of the vSwitch.
Data VLAN Tag - Site 1 field	Enter the correct VLAN ID if using trunk ports. Enter 0 if using access ports. **Note** Use the same Data VLAN Tag for site 1 and site 2.
Data VLAN Tag - Site 2 field	Enter the correct VLAN ID if using trunk ports. Enter 0 if using access ports. **Note** Use the same Data VLAN Tag for site 1 and site 2.
Data vSwitch field	Default is vswitch-hx-storage-data.

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

Field	Description
Jumbo frames **Enable Jumbo Frames** check box	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** At the time of installation, you can select the MTU size as either 1500 or 9000 bytes. **Note** If you want to change the MTU size from 9000 to 1500 after installation, you must change the MTU size in Cisco UCS, Data vSwitch, VMkernel, and Controller VM Ethernet1 interface.

Field	Description
Disk Partitions **Clean up Disk Partitions** check box	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. **Attention** Do not select this option for factory prepared systems. The disk partitions on factory prepared systems are properly configured.
Virtual Desktop (VDI) check box	Check for VDI only environments. **Note** To change the VDI settings after the storage cluster is created, shutdown or move the resources, make changes, and restart the cluster.

Step 8 Click **Start** to begin deploying the HyperFlex stretched 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 manage your HyperFlex stretched cluster.

Configuring VMware vCenter High Availability Settings for Stretch Cluster

The HyperFlex Stretch Cluster works as designed using the default installation values. This section provides the recommended settings for configuring VMware vSphere high availability for Stretch Cluster during custom configuration.

Note Use the default value for the fields for which the settimgs are not captured in this section.

Before you begin

Create the HyperFlex Stretch Cluster sites.

Step 1 Log on to vCenter with the vSphere Web Client.

Step 2 Select the Stretch cluster and click **Configure**.

Step 3 Choose **Configure** > **vSphere Availability** and click **Edit**.

Step 4 Check the **Turn on vSphere HA** check box.

Step 5 Click **Failures and Responses**.

Step 6 In the **Failure conditions and responses** screen, complete the following fields:

- **Enable Host Monitoring**—Check this check box to enable the host monitoring.
- **Host Failure Response**—From the drop-down list, choose **Restart VMs**.
- **Response for Host Isolation**—From the drop-down list, choose **Power off and Restart VMs**.
- **Datastore with PDL**—From the drop-down list, choose **Power off and Restart VMs**.
- **Datastore with APD**—From the drop-down list, choose **Power off and Restart VMs (conservative)**.
- **VM Monitoring**—From the drop-down list, choose the required option for VM monitoring. By default, **Disabled** is displayed.

Step 7 Click **Admission Control** and set to **Disable**.

Step 8 Click **Heartbeat Datastores**. In the **Heartbeat Datastores** screen, click the **Use datastores only from the specified list** radio button and choose the HyperFlex datastores.

Step 9 Click **Advanced Options**, and add the following options:

Option	Value
das.usedefaultisolationaddress	Leave entries for the das.isolationaddresses blank, in which case VMware defaults to the management network gateway address. If you prefer to enter values, then use the Management Network gateway for das.isolationaddress0. You can use any other existing IPs for additional isolation addresses, however it is a best practice that these are not cluster hosts, including the cluster CIP, or FI VIPs for additional isolation addresses. If you enter addresses, be sure to set das.usedefaultisolationaddress to **False**.
das.isolationaddress0	The IP address for Management Network Gateway.
das.isolationaddress1	The existing IP address that is outside the cluster. Do not use FI VIPs, Cluster IP (CIP), or cluster host IP.

Step 10 Click **OK**.

Using Intersight Private Virtual Appliance

This workflow intended for users who are unable to use the Intersight cloud (for example air-gaped/dark site network).

Step 1 Install new Stretch Cluster cluster. For additional information see, Installation Overview, on page 15

Step 2 If this is a air-gaped/dark site install (no internet) **HxConnect** > **System Information** > **Witness (INTERSIGHT)** should show offline status.

Step 3 Claim the cluster via the appliance.

Example:

netstat -tanp |grep <PVA IP>:443

tcp 0 0 10.64.65.21:40150 <PVA IP>:443 ESTABLISHED 9280/hxdp

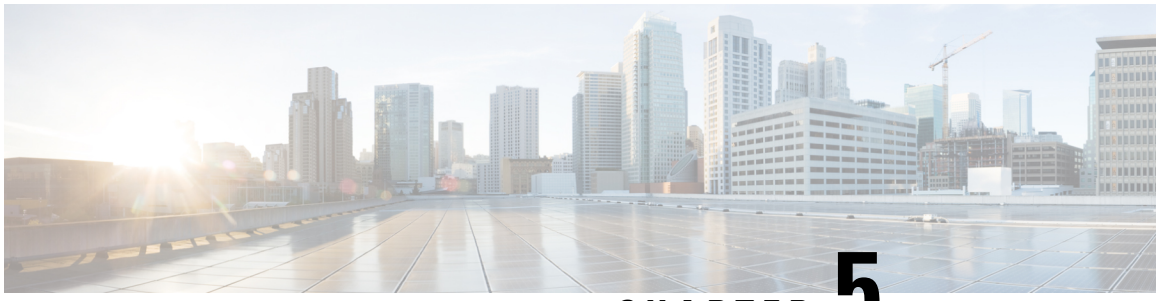

CHAPTER 5

Expanding a Stretch Cluster

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 40 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. See Common Site Interlink and Witness Requirements, on page 4 for more details.
- 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 40 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. See Common Site Interlink and Witness Requirements, on page 4 for more details.
- Download and deploy the Witness VM[1].

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.

a) In a browser, enter the URL for the VM where HX Data Platform Installer was installed.

b) 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, *10.193.211.120*.
UCS Manager User Name field	Enter the administrative level username. For example, <*admin*>.
Password field	Enter the administrative level password. For example, <*root*>.
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, <*admin*> 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 37.

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	• 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	• **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 36.

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 38.

Configure Hypervisor

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

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, *255.255.0.0*.
Gateway field	IP address of gateway. For example, *10.193.0.1*.
DNS Server(s) field	IP address for the DNS Server. **Note** • 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 40 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 40
2.	Configure the server ports and associate HyperFlex servers.	Associate HyperFlex Servers, on page 43
3.	Configure hypervisor, IP addresses, and start the cluster expansion process.	Configure Nodes, on page 44

Enter Credentials

Before you begin

- Complete the cluster expansion prerequisites.
- Run the *Configure Site* workflow once for each site. Refer to the Common Site Interlink and Witness Requirements, on page 4 for more details.

Step 1 Log into the Cisco HX Data Platform Installer.

a) 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.

b) 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.

c) 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, *10.193.211.120*.
User Name field	Enter the administrative level username. For example, <*admin*>.
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, *10.193.211.120*.
User Name field	Enter the administrative level username. For example, <*admin*> 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, *10.193.211.120*.

Field	Description
User Name field	Enter the administrative level username. For example, <*admin*> 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.

Step 6 Enter the following credentials for the vCenter.

vCenter Credentials

Field	Description
vCenter Server field	Enter the vCenter server FQDN or IP address. For example, *10.193.211.120*. **Note** • 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	Enter the administrative username. For example, *administrator@vsphere.local*.
Admin Password field	Enter the administrative level password. Enter the <*root*> password.

Step 7 Enter the following credentials for the Hypervisor:

Hypervisor Credentials

Field	Description
Admin User Name field	Enter the administrative username. For example, <*admin*> username. The username is **root** for factory nodes.
The hypervisor on this node uses the factory default password check box	Select the check box, if you are changing the factory default password for the hypervisor. To enter the current hypervisor password, uncheck this check box. You can now enter the hypervisor password in the **Enter current hypervisor password** field.

Field	Description
New Password	Create a new password for the hypervisor. **Important** You are required to change the factory default password.
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	• 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	• **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 40.

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 44.

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 43.

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.

Field	Description
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.

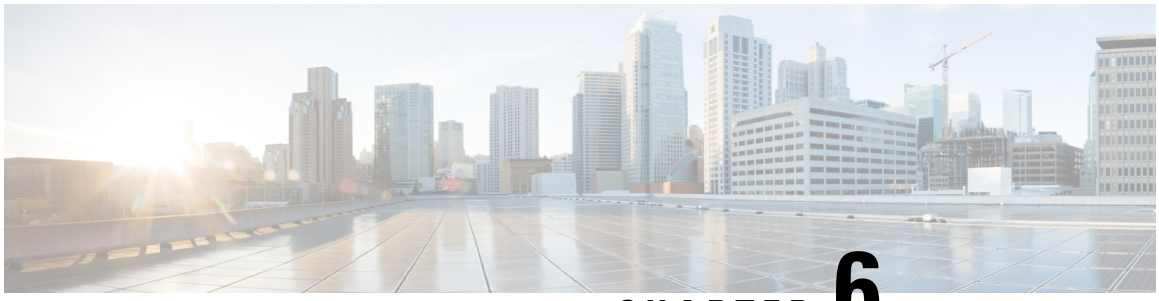

CHAPTER 6

Replication

Configuring Replication

To configure replication on a HyperFlex Controller VM in a Stretch Cluster, you complete the necessary VLAN configuration on the fabric interconnects using UCS Manager manually. 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**.

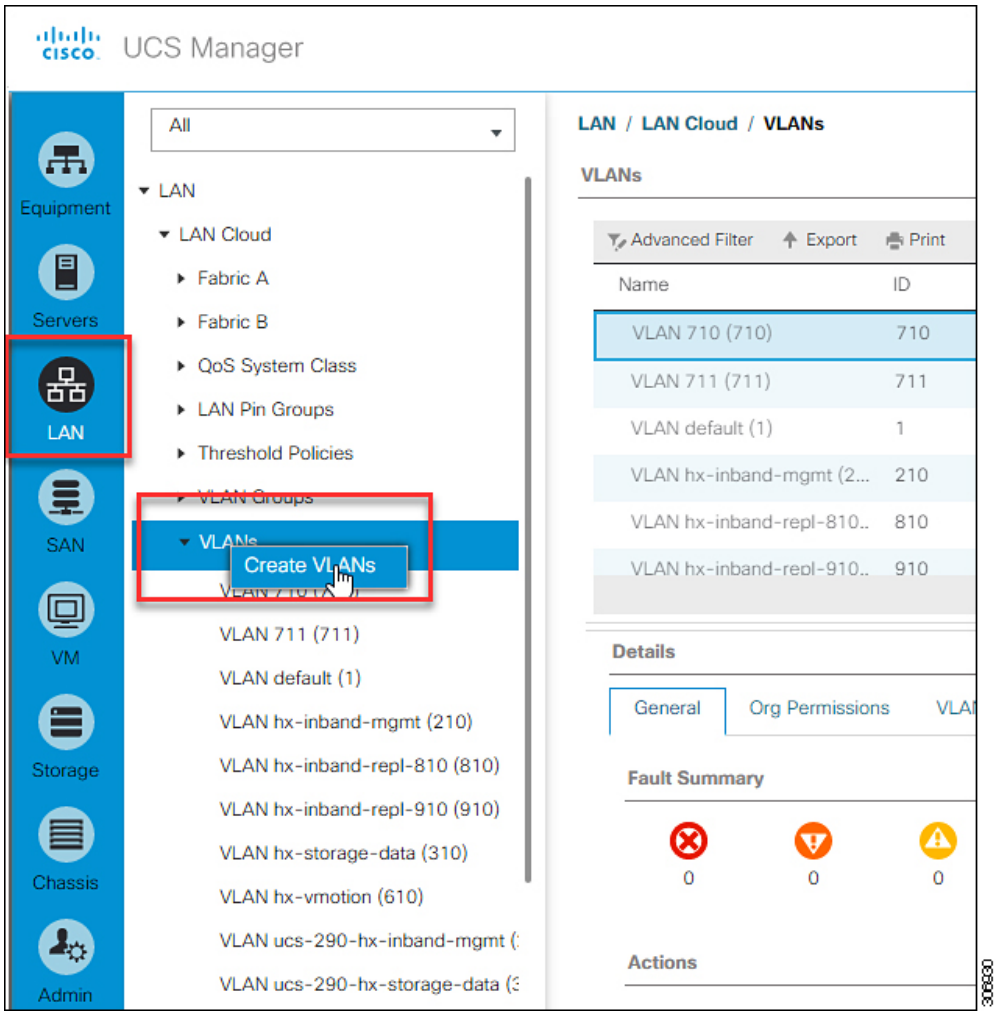

Step 3 In the **Create VLANs** dialog box, complete the following fields for all the VLANs that must be supported on the fabric interconnect:

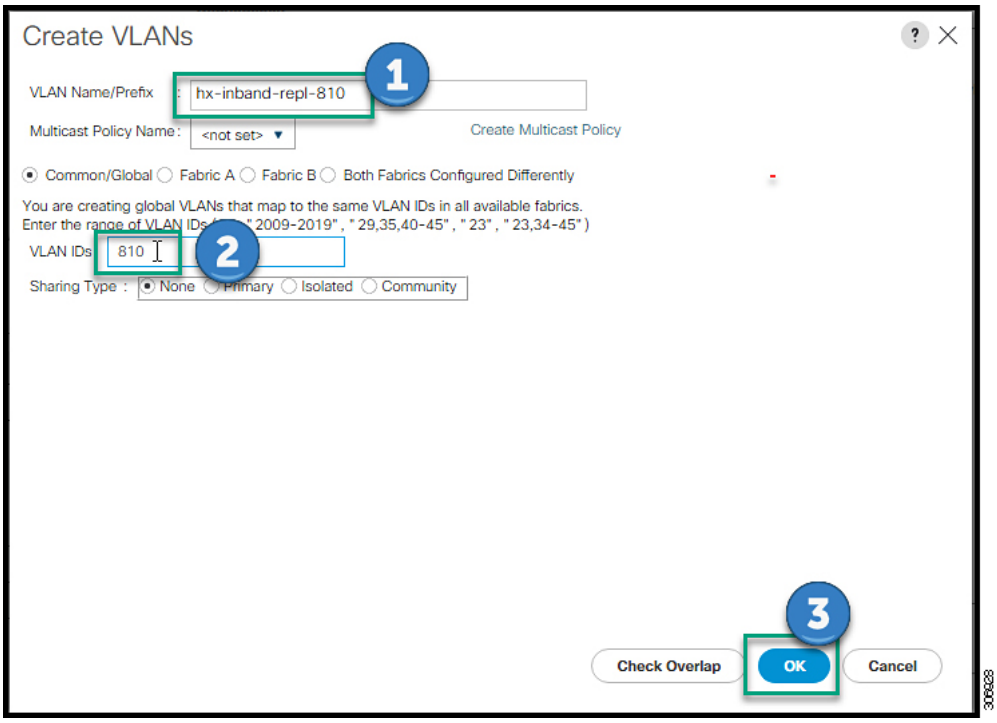

UI Element	Essential Information
VLAN Name/Prefix field	Enter the VLAN name prefix. For example, `hx-inband-repl-810`.
Multicast Policy Name field	*<not set>*
Common/Global check box **Fabric A** check box **Fabric B** check box **Both Fabric Configured Differently** check box	Choose the fabric configuration 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.

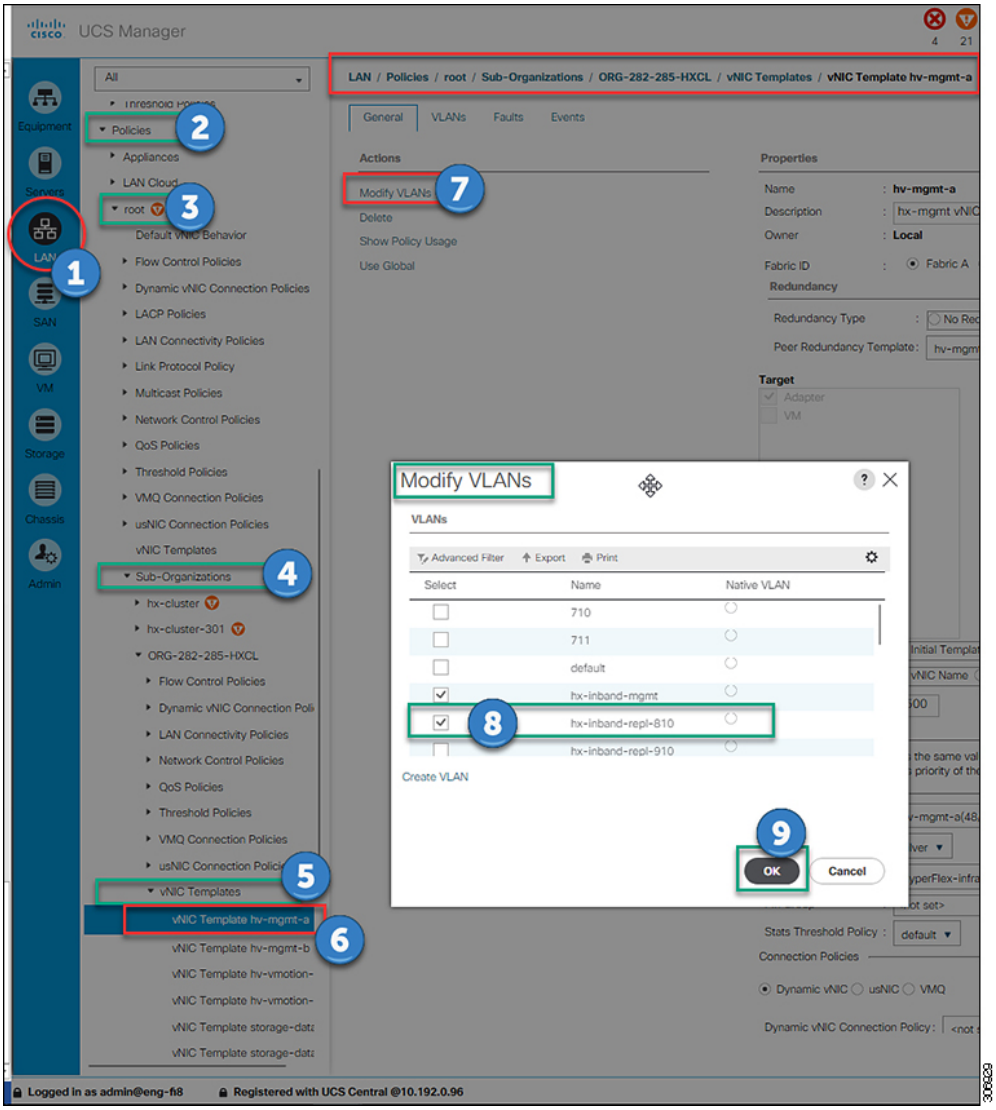

Before you begin

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**.

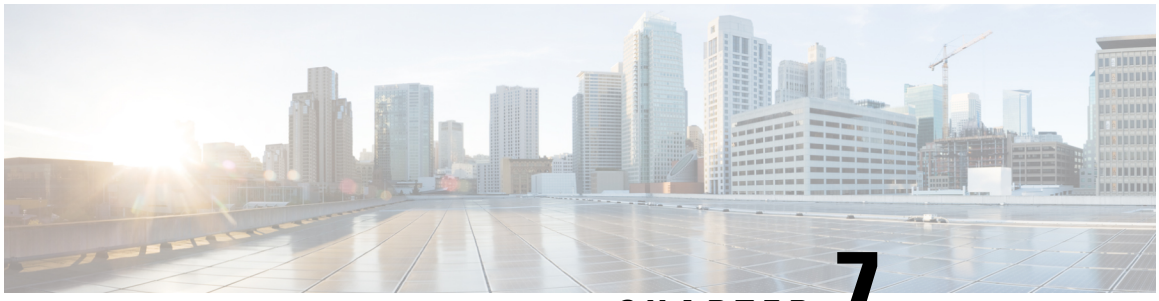

CHAPTER 7

Managing HyperFlex Stretch Clusters

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	Provides the functional status and application performance of the HX storage cluster on Site A and Site B. Provides the functional status and application performance of the witness VM or Invisible Cloud Witness node. Click **Information** ( ) to access the HX storage cluster name and status data.
Resiliency Health section	Provides the data health status and the ability of the HX storage cluster on Site A and Site B to tolerate failures. 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.
Capacity section	Displays a breakdown of the total storage versus how much storage is used or free. Also displays the storage optimization, compression-savings, and deduplication percentages based on the data stored in the cluster.
Nodes section	Displays the number of nodes and the division of converged versus compute nodes across Site A and Site B in the Stretch Cluster. 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.
Performance section	Displays an HX storage cluster performance snapshot for a configurable amount of time, showing IOPS, throughput, and latency data. For full details, see **Performance Page**.
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: `cvs`, `xls`, and `doc`. 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.

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, 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	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: **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. **Note** To register a cluster license, you can also choose **Register Cluster** from the **Actions** drop-down field.

License section	• **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:
	  • 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.
	Note To refresh license certificate or renew license authorization, choose the respective options from the **Actions** drop-down field.

HX storage cluster status field	Provides functional status of the HX storage cluster in Site A and Site B: • **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	Secure URL to the VMware vSphere associated with this HX storage cluster. Click the link to remotely access the vSphere Web Client.
Hypervisor field	Hypervisor version installed on this HX storage cluster.

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	• **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. • **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	• **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	• **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: Peristent, 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.

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 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. • 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

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, 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.

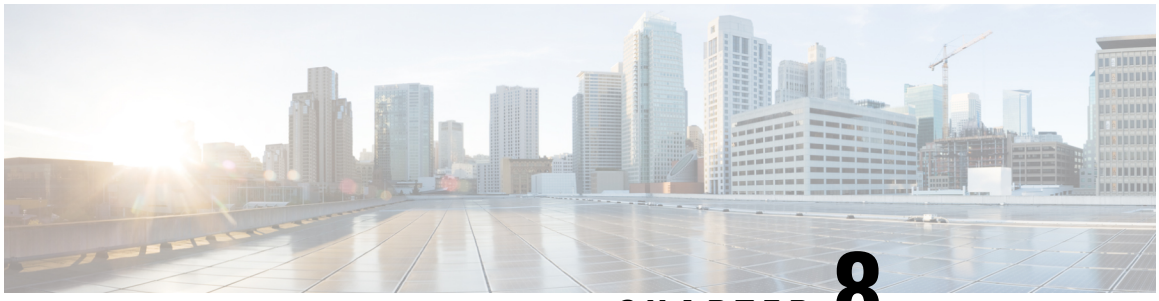

CHAPTER 8

Troubleshooting

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.
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.

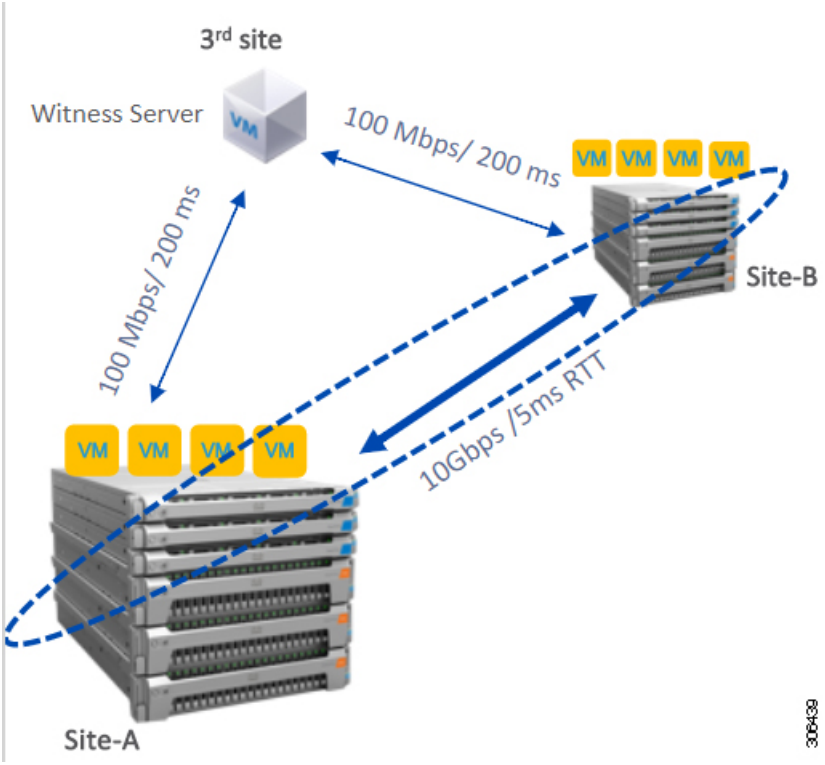

Increasing Datastore Capacity May Not Show as Free Space

Issue:

Increasing the datastore capacity may not show as free space in the HX Connect UI and stcli.

Recommended Solution:

Resize the datastore depending on the actual space that is used by each VM.

Note Use the maximum cluster capacity to resize the datastore. It reflects the actual used spaced of the datastore, which can be used as a hint to resize the datastore based on your requirements.

1. From HX Connect, select **Datastores**.
2. Select a *datastore*. Click **Edit**.
3. In the **Edit Datastore** dialog box, to resize the datastore, change the **Size** as needed.
4. Click **Edit Datastore**.

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.

Redeploy HyperFlex Stretch Cluster Witness VM

In some circumstances, it may be required to re-deploy the Stretch cluster witness VM.

Some examples include:

- Witness VM is deleted
- Witness VM is corrupted

Note This process should not be used in an attempt to change IP of the Witness VM.

To redeploy the Stretch Cluster Witness VM:

Step 1 Download the appropriate version of Stretch cluster witness .ova.

Step 2 Verify the cluster is healthy.

Step 3 Deploy a new witness VM with the same IP address as the previous one.

Step 4 Log into the new witness VM and stop exhibitor.

Example:

```
root@Cisco-HX-Witness-Appliance:~# service exhibitor stop
```

Step 5 Locate an SCVM in the cluster which is part of the zookeeper ensemble.

Note There should be two nodes per site that are part of the ensemble, and they show as either **leader** or **follower** in the example output.

Example:

```
root@SCVM-1:~# echo srvr | nc localhost 2181
Zookeeper version: 3.4.6--1, built on 07/26/2017 20:05 GMT
Latency min/avg/max: 0/0/3043
Received: 531158372
Sent: 531194128
```

```
Connections: 9
Outstanding: 0
Zxid: 0x350002e75c
Mode: follower   <<<<<<<<<<<<
Node count: 3804
```

Step 6 Copy the exhibitor.properties file from the SCVM to the witness VM under the /tmp location.

Example:

```
root@SCVM-1:~# scp /usr/share/exibitor/exhibitor.properties
root@<Witness-VM-IP>:/tmp/exhibitor.properties
```

Step 7 Edit the file **on the witness VM** and locate the line starting with **com.netflix.exhibitor.servers-spec**.

Example:

```
root@Cisco-HX-Witness-Appliance:~# vi /tmp/exhibitor.properties#

Auto-generated by Exhibitor
#Mon Oct 29 15:51:29 PDT 2018
com.netflix.exhibitor-rolling-hostnames=
com.netflix.exhibitor-rolling.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling.servers-spec=0\:192.168.4.159,1\:10.9.47.50
,2\:10.9.47.49,3\:10.9.47.45,4\:10.9.47.44
com.netflix.exhibitor.java-environment=ZOO_LOG_DIR\=/var/log/zookeeper\
nZOO_LOG4J_PROP\="INFO,ROLLINGFILE"\nZOOPIDFILE\=/tmp/zookeeper/
zookeeper_server.pid
com.netflix.exhibitor.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling-hostnames-index=0
com.netflix.exhibitor-rolling.java-environment=ZOO_LOG_DIR\=/var/log/
zookeeper\nZOO_LOG4J_PROP\="INFO,ROLLINGFILE"\nZOOPIDFILE\=/tmp/zookeeper/
zookeeper_server.pid
com.netflix.exhibitor-rolling.observer-threshold=0
com.netflix.exhibitor.servers-spec=0\:192.168.4.159,1\:10.9.47.50,2\
:10.9.47.49,3\:10.9.47.45,4\:10.9.47.44
com.netflix.exhibitor.cleanup-period-ms=0
com.netflix.exhibitor.auto-manage-instances-fixed-ensemble-size=0
com.netflix.exhibitor.zookeeper-install-directory=/usr/share/zookeeper
com.netflix.exhibitor.check-ms=30000
```

Note In the above example, the 10.9.47.x IP's are the storage data subnet while the 192.168.4.x IP's are the management subnet.

Step 8 With all of the data IP addresses of SCVM's in the zookeeper cluster, and the management IP address of the witness VM displayed, replace all of the data IP addresses with the management IP addresses for each corrseponding controller VM.

Example:

```
root@Cisco-HX-Witness-Appliance:~# vi /tmp/exhibitor.properties#

Auto-generated by Exhibitor
#Mon Oct 29 15:51:29 PDT 2018
com.netflix.exhibitor-rolling-hostnames=
com.netflix.exhibitor-rolling.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling.servers-spec=0\:192.168.4.159,1\:10.9.47.50,2\:10.9.47.49,3\:10.9.47.45,4\:10.9.47.44
com.netflix.exhibitor.java-environment=ZOO_LOG_DIR\=/var/log/zookeeper\nZOO_LOG4J_PROP\="INFO,ROLLINGFILE"\nZOOPIDFILE\=/tmp/zookeeper/zookeeper_server.pid
com.netflix.exhibitor.zookeeper-data-directory=/var/zookeeper
com.netflix.exhibitor-rolling-hostnames-index=0
com.netflix.exhibitor-rolling.java-environment=ZOO_LOG_DIR\=/var/log/zookeeper\nZOO_LOG4J_PROP\="INFO,ROLLINGFILE"\nZOOPIDFILE\=/tmp/zookeeper/zookeeper_server.pid
com.netflix.exhibitor-rolling.observer-threshold=0
com.netflix.exhibitor.servers-spec=0\:192.168.4.159,1\:192.168.4.50,2\:192.168.4.49,3\:192.168.4.45,4\:192.168.4.44
com.netflix.exhibitor.cleanup-period-ms=0
com.netflix.exhibitor.auto-manage-instances-fixed-ensemble-size=0
```

```
com.netflix.exhibitor.zookeeper-install-directory=/usr/share/zookeeper
com.netflix.exhibitor.check-ms=30000
```

Step 9 Save changes to the file and copy the file from /tmp to /usr/share/exhibitor/.

Example:

```
root@Cisco-HX-Witness-Appliance:~# cp /tmp/exhibitor.properties
/usr/share/exhibitor/exhibitor.properties
```

Step 10 Wait for the witness VM to sync with the cluster (this is typically quick but can take up to an hour). You can verify the Witness VM is online from HX Connect under **System Information**.

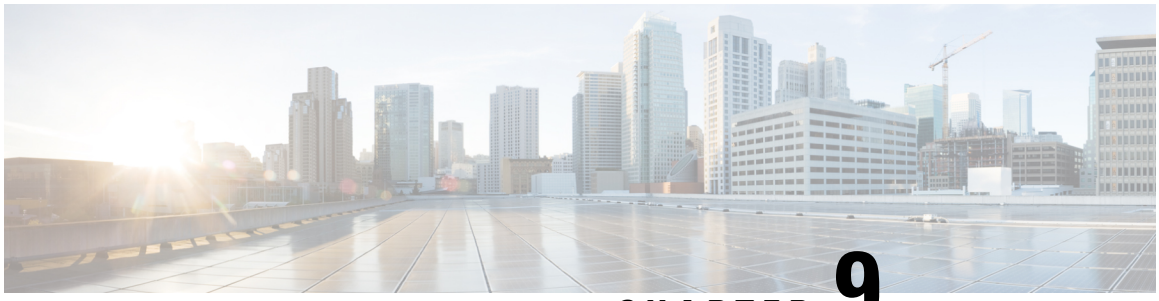

CHAPTER 9

Post Installation

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

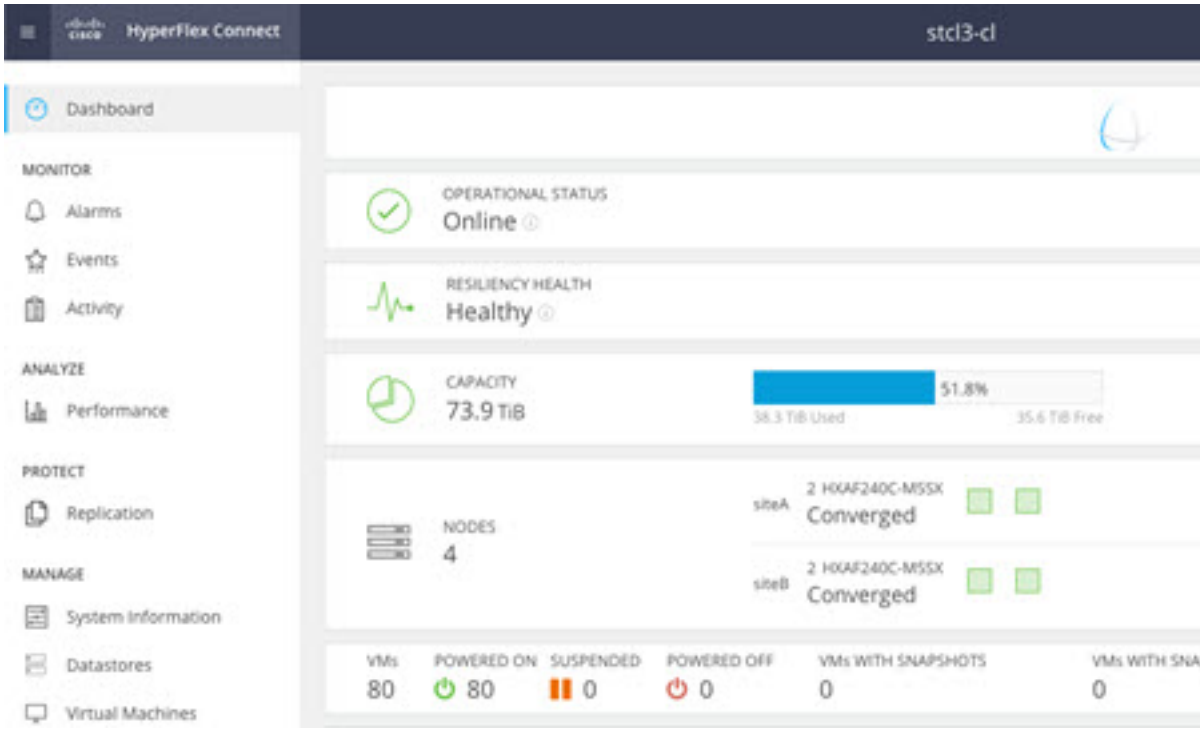

- Operational Status Online

- Resiliency Health: Healthy
- Nodes: Correct number of nodes, site names, and nodes per site

- System Information Panel

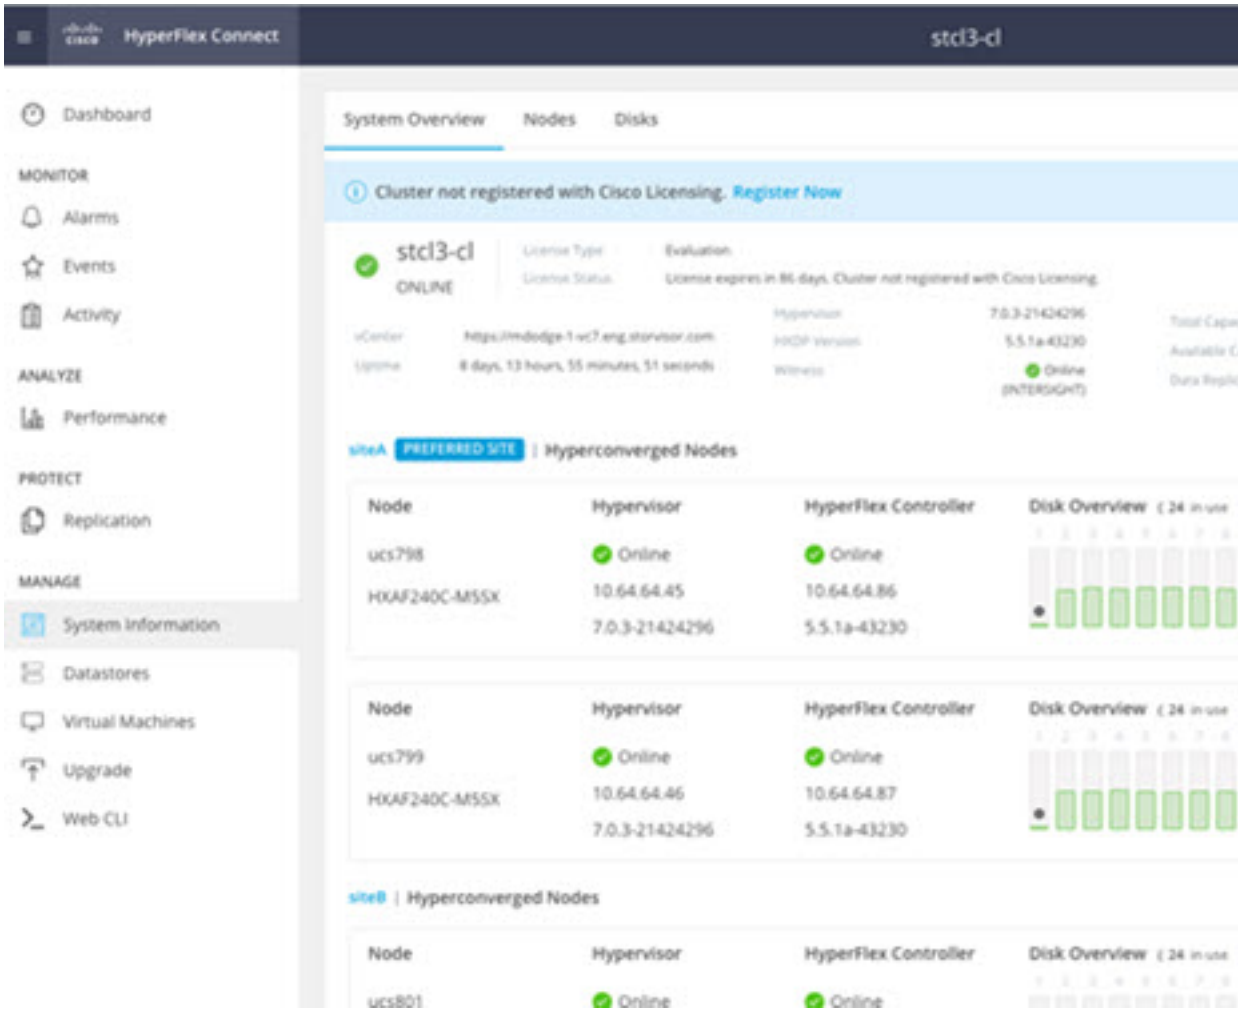

- Cluster name displayed is correct
- Cluster status is shown as Online
- Witness: Online (INTERSIGHT) or (WItness VM IP)
- Preferred Site badge on configured site.
- Nodes are shown in expected sites
- Hypervisor and Hyperflex Controllers are Online

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.

Verifying Which Site a Node Belongs To

Use the command, `stcli cluster get-zone`, to check which node belongs to which site.

Example:

```
admin@ucs-stctlvm-230-1:~# stcli cluster get-zonezones:
----------------------------------------
pNodes:
----------------------------------------
state: ready
name: 10.104.49.115
----------------------------------------
state: ready
name: 10.104.49.116
----------------------------------------
zoneId: 7f2bf7811475cacc:44dd22fa3eadfd4d
numNodes: 2
----------------------------------------
pNodes:
----------------------------------------
state: ready
name: 10.104.49.113
----------------------------------------
state: ready
name: 10.104.49.114
----------------------------------------
zoneId: 422fe637cab59ec5:4b49875b5641bf8a
numNodes: 2
----------------------------------------
isClusterZoneCompliant: True
zoneType: 2
isZoneEnabled: True
numZones: 2
```