



Post Cluster Configuration Tasks

- [Post Cluster Configuration Guidelines](#), on page 1
- [Enabling PCI Passthrough for a Network Device on a Host](#), on page 1
- [Run Post-Installation Script](#), on page 2
- [Changing ESXi Host Root Password](#), on page 4
- [Changing Storage Controller Password](#), on page 5
- [Access the HX Data Platform Plugin through vSphere](#), on page 6
- [Add Datastores in the Storage Cluster](#), on page 6
- [Set HA Heartbeat](#), on page 6
- [Auto Support and Smart Call Home for HyperFlex](#), on page 7
- [Replacing Self-Signed with CA-Signed Certificate](#), on page 12
- [Replication Pairing](#), on page 14
- [Adding Private VLAN](#), on page 14
- [Distributed Virtual Switches and Cisco Nexus 1000v](#), on page 18
- [Hosting vCenter on the HX Data Platform](#), on page 19
- [Deploying AMD GPUs](#), on page 19

Post Cluster Configuration Guidelines



Important

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

Enabling PCI Passthrough for a Network Device on a Host

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

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

Procedure

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

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

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

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

Run Post-Installation Script

To complete the post-installation tasks, you can run a post-installation script on the Installer VM.



Important

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

1. Use an SSH client to connect to the shell on the installer VM.
2. Log in with installer VM *root* credentials.
3. Type `post_install` and hit `Enter`.
4. Set the post-install script parameters as specified in the following table:



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

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

5. Correct network errors reported, if any.

Sample Post-Install Script

```
root@Cisco-HX-Data-Platform-Installer:~# post_install
Select post_install workflow-
```

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

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

```
Selection: 3
Certificate generation workflow selected
```

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

```
Select Certificate Generation Workflow-
```

1. With vCenter
2. Without vCenter

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

Sample Network Errors

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

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

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

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

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

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

Changing ESXi Host Root Password

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

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



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

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



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

Procedure

Step 1 Log in to the ESXi host service control using SSH.

Step 2 Acquire root privileges.

```
su -
```

Step 3 Enter the current root password.

Step 4 Change the root password.

```
passwd root
```

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

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

Changing Storage Controller Password

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

Procedure

Step 1 Log in to a storage controller VM.

Step 2 Change the Cisco HyperFlex storage controller password.

```
# stcli security password set
```

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

Note If you add new compute nodes and try to reset the cluster password using the `stcli security password set` command, the converged nodes get updated, but the compute nodes may still have the default password. To change the compute node password, use the following procedure.

To change the password on compute nodes:

- a. Vmotion all the user VMs off the ESXi hosts.
- b. Launch the storage controller VM console from vCenter and log in as the root user.
- c. Run the `passwd` command to change the password.
- d. Log out and re-login to confirm that the password changed successfully.
- e. Run the `stcli node add -f` command to add the node back into the cluster.

Step 3 Type in the new password.

Step 4 Press **Enter**.

Access the HX Data Platform Plugin through vSphere

To manage your storage cluster through a GUI, launch the vSphere Web Client. You access your storage cluster through the vSphere Web Client and HX Data Platform plug-in.

Procedure

-
- Step 1** From the HX Data Platform installer, after installation is completed, on the Summary page, click **Launch vSphere Web Client**.
 - Step 2** Display the login page, click **Login to vSphere Web Client** and enter your vSphere credentials.
 - Step 3** View the HX Data Platform plug-in.
From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform**.
-

Add Datastores in the Storage Cluster

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



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

Procedure

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

Set HA Heartbeat

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

Procedure

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

Auto Support and Smart Call Home for HyperFlex

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

Auto Support (ASUP)

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



Note Only unauthenticated SMTP is supported for ASUP.

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

Post-Cluster ASUP Configuration Method	Associated Topic
HX Connect user interface	Configuring Auto Support Using HX Connect, on page 8
Command Line Interface (CLI)	Configuring Notification Settings Using CLI, on page 9
REST APIs	Cisco HyperFlex Support REST APIs on Cisco DevNet .

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

Smart Call Home (SCH)

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

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

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

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

To ensure automatic communication among your HX storage cluster, you, and Support, see [Configuring Smart Call Home for Data Collection](#), on page 10.

Configuring Auto Support Using HX Connect

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

Procedure

- Step 1** Log in to HX Connect.
- Step 2** In the banner, click **Edit settings (gear icon)** > **Auto Support Settings** and fill in the following fields.

UI Element	Essential Information
Enable Auto Support (Recommended) check box	Configures Call home for this HX storage cluster by enabling: <ul style="list-style-type: none"> • Data delivery to Cisco TAC for analysis. • Notifications from Support as part of proactive support.
Send service ticket notifications to field	Enter the email address that you want to receive the notifications.

UI Element	Essential Information
Terms and Conditions check box	End user usage agreement. The check box must be checked to use the Auto-Support feature.
Use Proxy Server check box	<ul style="list-style-type: none"> • Web Proxy Server url • Port • Username • Password

Step 3 Click **OK**.

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

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

Step 5 Click **OK**.

Configuring Notification Settings Using CLI

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



Note Only unauthenticated SMTP is supported for ASUP.

Procedure

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

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

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

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

Example:

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

Step 3 Enable ASUP notifications.

```
# stcli services asup enable
```

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

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

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

Example:

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

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

```
# sendasup -t
```

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

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

Configuring Smart Call Home for Data Collection

Data collection is enabled by default but, during installation, you can opt-out (disable). You can also enable data collection post cluster creation. During an upgrade, Smart Call Home is set up based on your legacy configuration. For example, if `stcli services asup show` is enabled, Smart Call Home is enabled on upgrade.

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



Note In HyperFlex Data Platform release 2.5(1.a), Smart Call Home Service Request (SR) generation does not use a proxy server.

Using Smart Call Home requires the following:

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

Procedure

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

Step 2 Register your HX storage cluster with Support.

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

Note Upon configuring Smart Call Home in Hyperflex, an email will be sent to the configured address containing a link to complete registration. If this step is not completed, the device will remain in an inactive state and an automatic Service Request will not be opened.

Syntax:

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

Example:

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

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

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

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

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

If you upgraded your HX storage cluster from HyperFlex 1.7.1 to 2.1.1b, also run the following command:

```
# asupcli [--all] post --type alert
```

Contact Support if you receive the following error:

```
root@ucs-stctlvm-554-1:/tmp# asupcli post --type alert
/bin/sh: 1: ansible: not found
Failed to post - not enough arguments for format string
root@ucs-stctlvm-554-1:/tmp#
```

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

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

- a. Clear any existing registration email and proxy settings.

```
# stcli services sch clear
```

- b. Set the proxy and registration email.

Syntax:

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

Syntax Description

Option	Required or Optional	Description
--email EMAILADDRESS	Required.	Add an email address for someone to receive email from Cisco support. Recommendation is to use a distribution list or alias.

Option	Required or Optional	Description
--enable-proxy ENABLEPROXY	Optional.	Explicitly enable or disable use of proxy.
--portal-url PORTALURL	Optional.	Specify an alternative Smart Call Home portal URL, if applicable.
--proxy-url PROXYURL	Optional.	Specify the HTTP proxy URL, if applicable.
--proxy-port PROXYPORT	Optional.	Specify the HTTP proxy port, if applicable.
--proxy-user PROXYUSER	Optional.	Specify the HTTP proxy user, if applicable. Specify the HTTP proxy password, when prompted.

Example:

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

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

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

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

Step 5 Verify Smart Call Home is enabled.

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

```
# stcli services sch show
```

If Smart Call Home is disabled, enable it manually.

```
# stcli services sch enable
```

Step 6 Enable Auto Support (ASUP) notifications.

Typically, Auto Support (ASUP) is configured during HX storage cluster creation. If it was not, you can enable it post cluster creation using HX Connect or CLI. For more information, see [Auto Support and Smart Call Home for HyperFlex](#).

Replacing Self-Signed with CA-Signed Certificate

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

For example:

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

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



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

Procedure

- Step 1** Script Location in CVM: /usr/share/springpath/storfs-misc/hx-scripts/
 certificate_import_input.sh
 run stcli cluster reregister
- Step 2** In the Controller VM (Pointing to CIP), execute this commands to generate the CSR request.
 openssl req -nodes -newkey rsa:2048 -keyout /etc/ssl/private/<Host Name of the CVM>.key
 -out /etc/ssl/certs/<Host Name of the CVM>.csr
 cat /etc/ssl/certs/<host name mapped to the management CIP>.csr - Copy the request to
 any notepad.
 Send the request to CA to generate the certificate
- Step 3** Once you receive the certificate from CA (.crt files), copy the certificate and key to each CVM.
- Step 4** On each CVM, use this script to import the certificate: ./certificate_import_input.sh
 root@SpringpathControllerVUFSTDS58L:/usr/share/springpath/storfs-misc/hx-scripts#
 ./certificate_import_input.sh
- Step 5** Enter the path for the key: /etc/ssl/private/<Host Name of the CVM>.key.
- Step 6** Enter the path for the certificate in certificate format: <Path to the CA .crt file>
- Note** After providing all the inputs, it takes some time to finish the import process.
- Step 7** From the CVM pointing to CIP, run **stcli reregister** commmand to reregister the cluster to vCenter. It is mandatory to reregister the cluster once the certificate is imported.
-

Replication Pairing

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

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

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

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

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

Adding Private VLAN

About Private VLANs

A private VLAN partitions the Layer 2 broadcast domain of a VLAN into subdomains, allowing you to isolate the ports on the switch from each other. A subdomain consists of a primary VLAN and one or more secondary VLANs. A private VLAN domain has only one primary VLAN. Each port in a private VLAN domain is a member of the primary VLAN, and the primary VLAN is the entire private VLAN domain.

Understanding Private VLAN Ports

Table 1: Types of Private VLAN Ports

VLAN Port	Description
Promiscuous Primary VLAN	Belongs to the primary VLAN. Can communicate with all interfaces that belong to those secondary VLANs that are associated to the promiscuous port and associated with the primary VLAN. Those interfaces include the community and isolated host ports. All packets from the secondary VLANs go through this VLAN.
Isolated Secondary VLAN	Host port that belongs to an isolated secondary VLAN. This port has complete isolation from other ports within the same private VLAN domain, except that it can communicate with associated promiscuous ports.
Community Secondary VLAN	Host port that belongs to a community secondary VLAN. Community ports communicate with other ports in the same community VLAN and with associated promiscuous ports.

Following HX deployment, a VM network uses a regular VLAN by default. To use a Private VLAN for the VM network, see the following sections:

- [Configuring a Private VLAN on a VM Network without Existing VMs, on page 15.](#)
- [Configuring a Private VLAN on a VM Network with Existing VMs, on page 15.](#)

Configuring a Private VLAN on a VM Network without Existing VMs

Procedure

- Step 1** To configure a private VLAN on Cisco UCS Manager, see the [Cisco UCS Manager Network Management Guide](#).
- Step 2** To configure a private VLAN on the upstream switch, see the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#).
- Step 3** To configure a private VLAN on ESX hosts, see [Configuring Private VLAN on ESX Hosts, on page 15](#).
-

Configuring Private VLAN on ESX Hosts

To configure private VLANs on the ESX hosts do the following:

Procedure

- Step 1** Delete VMNICs on the vSphere Standard Switches from the VMware vSphere Client.
- Step 2** Create new vSphere Distributed Switch with the VMNICs deleted from the previous step.
- Step 3** Create promiscuous, isolated, and community VLAN.
-

Configuring a Private VLAN on a VM Network with Existing VMs

Procedure

- Step 1** To configure a private VLAN on Cisco UCS Manager, see the [Cisco UCS Manager Network Management Guide](#).
- Step 2** To configure a private VLAN on the upstream switch, see the [Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide](#).
- Step 3** To configure a private VLAN on ESX hosts, see [Configuring Private VLAN on ESX Hosts, on page 15](#)
- Step 4** Migrate VMs from vSphere standard switch to the newly created vSphere distributed switch.
- a) Right-click the vCenter Virtual Machine and click **Migrate Virtual Machine Networking**.
 - b) Choose **source network** and **destination network** from the drop-down list.
 - c) Click **Next**.

- d) Select the **Virtual Machines** that you want to migrate.
- e) Click **Finish**.

- Step 5** Change network connection of the network adapter on the VMs to private VLAN.
- a) Right-click the vCenter Virtual Machine and click **Edit Settings**.
 - b) Under the **Hardware** tab, select the network adapter you want to modify.
 - c) Select the **Network Connection** you want to use from the **Network Label** drop-down list.
 - d) Click **OK**.
-

Deleting VMNICs on the vSphere Standard Switch

Procedure

- Step 1** Log on to VMware vSphere Client.
 - Step 2** Select **Home > Hosts and Clusters**.
 - Step 3** Select the ESX host from which you want to delete the VMNIC.
 - Step 4** Open the **Configuration** tab.
 - Step 5** Click **Networking**.
 - Step 6** Select the **switch** you wish to remove a VMNIC from.
 - Step 7** Click the **Manage the physical adapters connected to the selected switch** button.
 - Step 8** Select the **vmnic** you want to delete and click **Remove**.
 - Step 9** Confirm your selection by clicking **Yes**.
 - Step 10** Click **Close**.
-

Creating vSphere Distributed Switch

Procedure

- Step 1** Log on to the VMware vSphere Client.
- Step 2** Select **Home > Networking**.
- Step 3** Right click on the cluster **Distributed Switch > New Distributed Switch**.
- Step 4** In the **Name and Location** dialog box, enter a name for the distributed switch.
- Step 5** In the **Select Version** dialog box, select the distributed switch version that correlates to your version and configuration requirements.
- Step 6** Click **Next**.
- Step 7** In the **Edit Settings** dialog box, specify the following:
 - Number of uplink ports
 - **Enable Network I/O Control**.
 - **Create a default port group** should be checked.

- Enter the default port group **name** in the **Port Group Name** box.

- Step 8** Click **Next**.
- Step 9** Review the settings in the **Ready to complete** dialog box.
- Step 10** Click **Finish**.
-

Creating Private VLANs on vSphere Distributed Switch

Procedure

- Step 1** From the VMware vSphere Client, select **Inventory > Networking**.
- Step 2** Right-click on the dvSwitch.
- Step 3** Click **Edit Settings**.
- Step 4** Select the **Private VLAN** tab.
- Step 5** On the **Primary private VLAN ID** tab, enter a **private VLAN ID**.
- Step 6** On the **Secondary private VLAN ID** tab, enter a **private VLAN ID**.
- Step 7** Select the type of VLAN from the **Type** drop-down list. Valid values include:
- **Isolated**
 - **Community**
 - **Promiscuous (Default)**
- Step 8** Click **OK**.
-

Set Private VLAN in Distributed Port Group

Before you begin

Create Private VLAN on the vSphere Distributed Switch.

Procedure

- Step 1** Right click **dvPortGroup** under **dvSwitch**, and click **Edit Settings**.
- Step 2** Click **Policies > VLAN**.
- Step 3** Select **Private VLAN**, from the **VLAN type** drop-down list.
- Step 4** From the **Private VLAN Entry** drop-down list, select the type of private VLAN. It can be one of the following:
- **Isolated**
 - **Community**

Note Community private VLAN is recommended.
Promiscuous ports are not supported

Step 5 Click OK.

Distributed Virtual Switches and Cisco Nexus 1000v

Considerations when Deploying Distributed Switches



Note

- Using Distributed Virtual Switches (DVS) or Cisco Nexus 1000v (N1Kv) is an optional and not a required step.
 - DVS for your vMotion network is available only if your environment has Enterprise Plus License for vSphere.
 - You can use only one of the two switches at a given time.
 - There may be a potential conflict between the Quality of Service (QoS) policy for HyperFlex and Nexus 1000v. Make sure that the QoS classes for N1Kv are set as per the HyperFlex policy. See *Creating a QoS Policy*, in the [Network and Storage Management Guide](#).
 - If you choose to deploy N1Kv switch, apply the settings as described, so that the traffic between the HyperFlex hosts flows locally on the FIs in a steady state. If not configured accurately, it could lead to a situation where most traffic will go through the upstream switches leading to latency. In order to avoid that scenario, ensure that the Storage Controller, Management Network, and vMotion port groups are configured with active/standby and failover enabled.
1. Set the **link status** for the **Network Control Policy** using UCS Manager. For details, see the "Configuring Network Control Policy" section in the [Cisco UCS Manager GUI Configuration Guide](#).
 2. Set the vSwitch properties in vCenter.
 - a. Set the **Network Failure Detection** to **Link Status only**.
 - b. Set **Failback** to **Yes**. For details, see the "Configuring the VM-FEX for VMware" section in the [Cisco UCS Manager VM-FEX for VMware Configuration guide](#)

Distributed switches ensure that each node is using the same configuration. It helps prioritize traffic and allows other network streams to utilize available bandwidth when no vMotion traffic is active.

The HyperFlex (HX) Data Platform can use Distributed Virtual Switch (DVS) Networks for non-HyperFlex dependent networks.

These non-HX dependent networks include:

- VMware vMotion networks
- VMware applications networks

The HX Data Platform has dependency that the following networks use standard vSwitches.

- vswitch-hx-inband-mgmt: Storage Controller Management Network
- vswitch-hx-inband-mgmt: Management Network
- vswitch-hx-storage-data: Storage Hypervisor Data Network
- vswitch-hx-storage-data: Storage Controller Data Network

During HX Data Platform installation, all the networks are configured with standard vSwitch networks. After the storage cluster is configured, the non-HX dependent networks can be migrated to DVS networks. For example:

- vswitch-hx-vm-network: VM Network
- vmotion: vmotion pg

For further details on how to migrate the vMotion Network to Distributed Virtual Switches, please see the *Migrating vMotion Networks to Distributed Virtual Switches (DVS) or Cisco Nexus 1000v (N1Kv)* in the [Network and Storage Management Guide](#).

Hosting vCenter on the HX Data Platform

Deployment of vCenter on the HyperFlex cluster is supported with some constraints. See the [How to Deploy vCenter on the HX Data Platform](#) TechNote for more details.

Deploying AMD GPUs

AMD FirePro S7150 series GPUs are supported in HX240c M5 nodes. These graphic accelerators enable highly secure, high performance, and cost-effective VDI deployments. Follow the steps below to deploy AMD GPUs in HyperFlex.

Step	Action	Step Instructions
1	For the service profiles attached to the servers modify the BIOS policy.	Requirement For All Supported GPUs: Memory-Mapped I/O Greater than 4 GB
2	Install the GPU card in the servers.	GPU Card Installation
3	Power on the servers, and ensure that the GPUs are visible in the Cisco UCS Manager inventory for the servers.	—
4	Install the vSphere Installation Bundle (VIB) for the AMD GPU card and reboot.	Download the inventory list from Cisco Software Downloads that includes the latest driver ISO for C-series standalone firmware / software version bundle 3.1(3) for AMD on VMware ESXi .
5	Create a Win10 VM on the cluster with the VM configuration.	Specifying Eligible Virtual Machines

Step	Action	Step Instructions
6	On each ESXi hosts run the <code>MxGPU.sh</code> script to configure the GPUs and to create virtual functions from the GPU.	Using the MxGPU Setup Script
7	Assign the virtual functions (VFs) created in the previous step to the Win10 VMs.	—