



Administering Clusters on vSphere

You can create, upgrade, modify, or delete vSphere on-prem Kubernetes clusters using the Cisco Container Platform web interface.

Cisco Container Platform supports v2 and v3 clusters on vSphere. The v2 clusters use a single master node for its Control Plane, whereas the v3 clusters can use 1 or 3 master nodes for its control plane. The multi-master approach of v3 clusters is the preferred cluster type as this approach ensures high availability for the Control Plane.



Note The UI differences between v2 and v3 clusters are called out in the cluster creation task.

This chapter contains the following topics:

- [Creating Clusters on vSphere, on page 1](#)
- [Configuring Add-ons for Clusters on vSphere, on page 4](#)
- [Deleting Add-ons for v3 Clusters, on page 5](#)
- [Upgrading Clusters on vSphere, on page 6](#)
- [Scaling Clusters on vSphere, on page 6](#)
- [Configuring Node Pools, on page 6](#)
- [Deleting Clusters on vSphere, on page 8](#)

Creating Clusters on vSphere

Before you begin

Ensure that your subnets do not overlap. The list of subnets to consider are:

- 172.17.0.0/16: Docker Bridge uses this default subnet for networking. You can change the Default Docker Bridge IP address during cluster deployment.
- 10.96.0.0/12: This subnet is defined as `--service-cluster-ip-range` in each guest cluster. Kubernetes allocates and assigns IP addresses from this subnet for Kubernetes services. You cannot change the Kubernetes cluster IP address range.
- Routable CIDR subnet: The routable CIDR subnet for node and load balancer services. The subnet is defined during cluster deployment.

- Pod subnet: This subnet applies to the ACI-CNI and Calico options. The IP addresses for the pods are assigned from this subnet. The subnet is defined either in the ACI-CNI profile for the ACI-CNI option or during cluster deployment for the Calico CNI option.
- Service subnet: This subnet applies to the ACI-CNI option. Cisco Container Platform assigns each Kubernetes node an IP address from this subnet. It is used by ACI as PBR target for Kubernetes services of type load balancer. By default, this subnet is part of the same ACI VRF as the pod subnet and the Routable CIDR subnet.
- ACI infrastructure network: This subnet applies to the ACI-CNI option.

Step 1 In the left pane, click **Clusters**, and then click the **vSphere** tab.

Step 2 Click **NEW CLUSTER**.

Step 3 In the **Basic Information** screen:

- From the **INFRASTRUCTURE PROVIDER** drop-down list, choose the provider related to your Kubernetes cluster. For more information, see [Adding vSphere Provider Profile](#).
- In the **KUBERNETES CLUSTER NAME** field, enter a name for your Kubernetes tenant cluster.
- In the **DESCRIPTION** field, enter a description for your cluster.
- In the **KUBERNETES VERSION** drop-down list, choose the version of Kubernetes that you want to use for creating the cluster.
- If you are using ACI, specify the ACI profile.
For more information, see [Adding ACI Profile](#).
- Click **NEXT**.

Step 4 In the **Provider Settings** screen:

- From the **DATA CENTER** drop-down list, choose the data center that you want to use.
- From the **CLUSTERS** drop-down list, choose a cluster.

Note Ensure that DRS and HA are enabled on the cluster that you choose. For more information on enabling DRS and HA on clusters, see *Cisco Container Platform Installation Guide*.

- From the **DATASTORE** drop-down list, choose a datastore.

Note Ensure that the datastore is accessible to the hosts in the cluster.

- From the **VM TEMPLATE** drop-down list, choose a VM template.
- From the **NETWORK** drop-down list, choose a network.

Note

- Ensure that you select a subnet with an adequate number of free IP addresses. For more information, see [Managing Networks](#). The selected network must have access to vCenter.

- For v2 clusters that use HyperFlex systems:

- The selected network must have access to the HyperFlex Connect server to support HyperFlex Storage Provisioners.
- For HyperFlex Local Network, select **k8-priv-iscsivm-network** to enable HyperFlex Storage Provisioners.

- From the **RESOURCE POOL** drop-down list, choose a resource pool.

g) Click **NEXT**.

Step 5 In the **Node Configuration** screen:

a. From the **GPU TYPE** drop-down list, choose a GPU type.

Note GPU Configuration applies only if you have GPUs in your HyperFlex cluster.

b. For v3 clusters, under **MASTER**, choose the number of master nodes, and their VCPU and memory configurations.

Note You may skip this step for v2 clusters. You can configure the number of master nodes only for v3 clusters.

c. Under **WORKER**, choose the number of worker nodes, and their VCPU and memory configurations.

d. In the **SSH USER** field, enter the ssh user name.

e. In the **SSH KEY** field, enter the SSH public key that you want to use for creating the cluster.

Note Ensure that you use the Ed25519 or ECDSA format for the public key. As RSA and DSA are less secure formats, Cisco prevents the use of these formats.

f. In the **ROUTABLE CIDR** field, enter the IP addresses for the pod subnet in the CIDR notation.

For more information on the routable CIDR, see [Tenant Cluster with ACI Deployment](#).

g. From the **SUBNET** drop-down list, choose the subnet that you want to use for this cluster.

h. In the **POD CIDR** field, enter the IP addresses for the pod subnet in the CIDR notation.

i. In the **DOCKER HTTP PROXY** field, enter a proxy for the docker.

j. In the **DOCKER HTTPS PROXY** field, enter an https proxy for the docker.

k. In the **DOCKER BRIDGE IP** field, enter a valid CIDR to override the default Docker bridge.

Note If you want to install the HX-CSI addon, ensure that you set the CIDR network prefix of the **DOCKER BRIDGE IP** field to /24.

l. Under **DOCKER NO PROXY**, click **ADD NO PROXY**, and then specify a comma-separated list of hosts that you want to exclude from proxying.

m. In the **VM USERNAME** field, enter the VM username that you want to use as the login for the VM.

n. Under **NTP POOLS**, click **ADD POOL** to add a pool.

o. Under **NTP SERVERS**, click **ADD SERVER** to add an NTP server.

p. Under **ROOT CA REGISTRIES**, click **ADD REGISTRY** to add a root CA certificate to allow tenant clusters to securely connect to additional services.

q. Under **INSECURE REGISTRIES**, click **ADD REGISTRY** to add docker registries created with unsigned certificates.

r. For v2 clusters, the Istio add-on is deprecated.

s. Click **NEXT**.

Step 6 For v2 clusters, to integrate Harbor with Cisco Container Platform:

Note Harbor is currently not available for v3 clusters.

a) In the **Harbor Registry** screen, click the toggle button to enable Harbor.

b) In the **PASSWORD** field, enter a password for the Harbor server administrator.

- c) In the **REGISTRY** field, enter the size of the registry in gigabits.
- d) Click **NEXT**.

Step 7 In the **Summary** screen, verify the configuration, and then click **FINISH**.

The cluster deployment takes a few minutes to complete. The newly created cluster is displayed on the **Clusters** screen.

For more information on deploying applications on clusters, see [Deploying Applications on Kubernetes Clusters](#).

Configuring Add-ons for Clusters on vSphere



Note This section applies to v3 clusters.

In v3 clusters, the monitoring, logging, Istio, Harbor, and Kubernetes dashboard functions are available as configurable add-ons.

In v2 clusters, the monitoring, logging, Harbor, and Kubernetes dashboard add-ons are installed by default. The Istio add-on has been deprecated.

Step 1 In the left pane, click **Clusters**, and then click the **vSphere** tab.

Step 2 From the **VERSIONS** drop-down list, choose **VERSION 3** to view the v3 clusters.

Step 3 Choose the cluster for which you want to configure add-ons.

Step 4 Click the **ADD-ONS** tab.
The **Installed Add-ons** page appears.

Step 5 Click **INSTALL ADD-ON**.
The **Install Add-on** page appears.

Step 6 In the **Select an Add-on** area, click one of the following add-ons:

- Monitoring: For monitoring clusters
- Logging: For logging
- Dashboard: For deploying and managing the applications that are deployed on the clusters
- Kubeflow: For deploying machine learning (ML) workloads
- HyperFlex Storage (CSI): For deploying HyperFlex storage
- Istio Operator: For deploying the Istio operator service, which is required for running Istio
- Istio: For deploying the Istio services, which requires the Istio Operator to be running beforehand
- Harbor Operator: For deploying the Harbor operator service, which is required for running Harbor

Note The default registry size of a Harbor instance is 20Gi. You can modify the default size using the **REGISTRY SIZE** field in the **Configure the Add-on** area. Customizing the Chartmuseum size using the Cisco Container Platform web interface is not currently supported. As a workaround, see [Customizing Chartmuseum Size of Harbor Instance, on page 5](#).

- Harbor: For deploying the Harbor service, which requires the Harbor Operator to be running beforehand

Step 7 Click **INSTALL**.

Customizing Chartmuseum Size of Harbor Instance



Note This section applies to v3 clusters.

The default chartmuseum size of a Harbor instance is 5Gi. Customizing the Chartmuseum size using the Cisco Container Platform web interface is not currently supported. As a workaround, you can use the following steps:

Step 1 Install the Harbor operator add-on as described in [Configuring Add-ons for Clusters on vSphere, on page 4](#).

Step 2 SSH into the master of the tenant cluster.

Step 3 Customize the Chartmuseum size.

For example, to set the size of the chartmuseum to 40Gi, run the following command:

```
helm install -n harbor-cr /opt/ccp/charts/ccp-harbor-cr.tgz --set chartmuseumSize=40Gi
```

Deleting Add-ons for v3 Clusters



Note This section applies to v3 clusters.

In v3 clusters, the monitoring, logging, Istio, Harbor, and Kubernetes dashboard functions are removable through the Cisco Container Platform web interface.

In v2 clusters, you cannot delete these add-ons through the Cisco Container Platform web interface.

Step 1 In the left pane, click **Clusters**, and then click the **vSphere** tab.

Step 2 From the **VERSIONS** drop-down list, choose **VERSION 3** to view the v3 clusters.

Step 3 Choose the cluster for which you want to delete add-ons.

Step 4 Click the **ADD-ONS** tab.
The **Installed Add-ons** page appears.

Step 5 From the drop-down list displayed under the **ACTIONS** column, click **Delete** for the add-on that you want to delete.

Step 6 Click **Close**.

Upgrading Clusters on vSphere

Before you begin

Ensure that you have imported the latest tenant cluster OVA to the vSphere environment.

Ensure that an adequate number of free IP addresses are available. For more information, see [Managing Networks](#).

For more information on importing the tenant cluster OVA, see the *Cisco Container Platform Installation Guide*.

-
- Step 1** In the left pane, click **Clusters**, and then click the **vSphere** tab.
- Step 2** From the drop-down list displayed under the **ACTIONS** column, choose **Upgrade** for the cluster that you want to upgrade.
- Step 3** In the **Upgrade Cluster** dialog box, choose a Kubernetes version and a new template for the VM, and then click **Submit**. It may take a few minutes for the Kubernetes cluster upgrade to complete.
-

Scaling Clusters on vSphere

You can scale clusters by adding or removing worker nodes to them based on the demands of the workloads you want to run. You can add worker nodes in a default or custom node pool.

For more information on adding worker node pools, see [Configuring Node Pools, on page 6](#).

Configuring Node Pools

Node pools allow the creation of worker nodes with varying configurations. Nodes belonging to a single node pool have identical characteristics.

In the Cisco Container Platform vSphere implementation, a node pool has the following properties:

- vcpus
- memory
- template
- [labels](#)
- [taints](#)

Labels and taints are optional parameters. All nodes that belong to a nodepool are tagged with labels and they are tainted. Taints are key-value pairs, which are associated with an *effect*.

The following table describes the available *effects*.

Effect	Description
NoSchedule	Ensures that the pods that do not contain this taint are not scheduled on the node.
PreferNoSchedule	Ensures that Kubernetes avoids scheduling pods that do not contain this taint on the node.
NoExecute	Ensures that a pod is removed from the node if it is already running on the node, and is not scheduled on the node if it is not yet running on the node.

During cluster creation, each cluster is assigned a default node pool. Cisco Container Platform supports the ability for different master and worker configurations. Upon cluster creation, the master node is created in the default-master-pool and the worker nodes are created in the default-pool.

Cisco Container Platform supports the ability to create multiple node pools and customize each pool characteristics such as vCCPUs, memory, labels, and taints.

Adding Node Pools

Cisco Container Platform allows you to add custom node pools to an existing cluster.

-
- Step 1** Click the cluster for which you want to add a node pool.
The **Cluster Details** page displays the node pools of the cluster that you have selected.
 - Step 2** In the right pane, click **ADD NODE POOL**.
The **Add Node Pool** page appears.
 - Step 3** Under **POOL NAME**, enter a name for the node pool.
 - Step 4** Ensure that an adequate number of free IP addresses is available in the subnet that you have selected during tenant cluster creation. For more information, see [Managing Networks](#).
 - Step 5** Under **Kubernetes Labels**, enter the key-value pair of the label.
You can click the **Delete** icon to delete a label and the **+LABEL** icon to add a label.
 - Step 6** Under **Kubernetes Taints**, enter the key-value pair and the effect you want to set for the label.
You can click the **Delete** icon to delete a taint and the **+TAINT** icon to add a taint.
 - Step 7** Click **ADD**.
The **Cluster Details** page displays the node pools. You can point the mouse over the **Labels** and **Taints** to view a summary of the labels and taints that are assigned to a pool.
-

Modifying Node Pools

Cisco Container Platform allows you to modify the worker node pools.

-
- Step 1** Click the cluster that contains the node pool that you want to modify.
The **Cluster Details** dialog box appears displaying the node pools of the cluster that you have chosen.

- Step 2** From the drop-down list next to the name of the node pool, click **Edit**. The **Update Node Pool** page appears.
- Step 3** Ensure that an adequate number of free IP addresses is available in the subnet that you have selected during tenant cluster creation. For more information, see [Managing Networks](#).
- Step 4** Under **Kubernetes Labels**, modify the key-value pair of the label.
- Step 5** Under **Kubernetes Taints**, modify the key-value pair and the effect you want to set for the label.
- Step 6** Click **UPDATE**.
-

Deleting Node Pools

Cisco Container Platform allows you to delete the worker node pools. You cannot delete the default master pool.

- Step 1** Click the cluster that contains the node pool that you want to delete. The **Cluster Details** page displays the node pools of the cluster that you have chosen.
- Step 2** From the drop-down list next to the worker pool that you want to delete, choose **Delete**. The worker pool is deleted from the **Cluster Details** page.
-

Deleting Clusters on vSphere

Before you begin

Ensure that the cluster you want to delete is not currently in use, as deleting a cluster removes the containers and data associated with it.

- Step 1** In the left pane, click **Clusters**, and then click the **vSphere** tab.
- Step 2** From the drop-down list displayed under the **ACTIONS** column, choose **Delete** for the cluster that you want to delete.
- Step 3** Click **DELETE** in the confirmation dialog box.
-