



## **Cisco Container Platform 4.1.0 User Guide**

**First Published:** 2019-07-25

**Last Modified:** 2019-08-13

### **Americas Headquarters**

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

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

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

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

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

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

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019 Cisco Systems, Inc. All rights reserved.



## CONTENTS

---

### CHAPTER 1

#### **Cisco Container Platform 1**

- Administrator Workflow 1
- User Workflow 2
- Accessing Cisco Container Platform Web Interface 3
- Setting Up Cisco Container Platform 3

---

### CHAPTER 2

#### **Configuring Cisco Smart Software Licensing 5**

- Connected Model 5
- Disconnected Model 6
- License Usage and Compliance 6
- Workflow in a Connected Model 6
  - Configuring Transport Settings 7
  - Registering Using Registration Token 7
  - Renewing Authorization 8
  - Reregistering Cisco Container Platform License 9
  - Deregistering Registration 9
- Workflow in a Disconnected Model 10
  - Registering Using Specific License Reservations 10
  - Updating Reserved Licenses 11
  - Deregistering Registration by Returning Reserved Licenses 12

---

### CHAPTER 3

#### **Managing Cisco Container Platform Infrastructure Configuration 13**

- Managing Users and RBAC 13
  - Configuring Local Users 13
  - Changing Login Passphrase 14
  - Recovering Login Passphrase for Local Admin 14

Configuring AD Servers	14
Configuring AD Groups	15
Managing Provider Profile	15
Adding Provider Profile	15
Adding vSphere Provider Profile	16
Adding Amazon Provider Profile	16
Adding Azure Provider Profile	16
Modifying Provider Profile	18
Deleting Provider Profile	18
Managing ACI Profile	18
Adding ACI Profile	18
Modifying ACI Profile	19
Deleting ACI Profile	19
Managing Networks	19
Modifying Networks	20
Adding Subnets	20
Modifying Subnets	21
Adding VIP Pool	21
Modifying VIP Pool	21
<hr/>	
<b>CHAPTER 4</b>	<b>Administering vSphere On-prem Clusters 23</b>
	Creating Kubernetes Clusters on vSphere On-prem Clusters 23
	Upgrading vSphere Clusters 25
	Scaling vSphere Clusters 25
	Configuring Node Pools 25
	Adding Node Pools 26
	Modifying Node Pools 26
	Deleting Node Pools 27
	Deleting vSphere Clusters 27
<hr/>	
<b>CHAPTER 5</b>	<b>Administering Amazon Web Services (AWS) EKS Clusters 29</b>
	Prerequisites for Configuring AWS EKS Clusters 29
	Amazon Resource Requirements 29
	Adding AMI Files to your Amazon Account 30

Creating AWS Roles	30
AWS Account Policy Requirements	31
Provider permissions	31
Amazon IAM Authentication	32
Enabling Common Identity	32
Configuring Control Plane Proxy for EKS Access	33
Creating AWS EKS Clusters	33
Scaling AWS EKS Clusters	34
Deleting AWS EKS Clusters	35

---

**CHAPTER 6**
**Administering Azure Kubernetes Service (AKS) Clusters 37**

Prerequisites for Configuring AKS Clusters	37
Azure User Account Requirements	37
Creating Resource Groups	37
Creating AKS Clusters	38
Deleting AKS Clusters	39

---

**CHAPTER 7**
**Managing Kubernetes Clusters 41**

Setting Up Kubernetes Dashboard to Access Clusters	41
Setting Up Kubernetes Dashboard for vSphere On-prem Clusters	41
Setting Up Kubernetes Dashboard for On-prem AWS IAM Enabled	42
Setting Up Kubernetes Dashboard for AWS EKS Clusters	42
Setting Up Kubernetes Dashboard for AKS Clusters	42
Monitoring Health of Cluster Deployments	43
Example of Monitoring Multiple Prometheus Instances	44
Monitoring Logs from Cluster Deployments	44
Viewing EFK Logs Using Kibana (Tenant Cluster)	45
Viewing EFK Logs Using Kibana (Control Plane Cluster)	45
Forwarding Logs to External Elasticsearch Server	46

---

**CHAPTER 8**
**Services and Networking 49**

Load Balancing Kubernetes Services using NGINX	49
L7 Ingress	49
L4 Ingress	51

Configuring L4 Load Balancing 52

Ingress CA 52

Network Policies 55

Load Balancer Services 55

---

**CHAPTER 9 Istio Service Mesh 57**

Introduction to Istio Service meshes 57

Configuring Istio Service meshes 58

Monitoring Service meshes 59

---

**CHAPTER 10 Harbor Registry 63**

Using Harbor Registry in Tenant Clusters 63

---

**CHAPTER 11 Deploying Applications on Kubernetes Clusters 65**

Workflow of Deploying Applications 65

Downloading Kubeconfig File 65

Sample Scenarios 66

Deploying a Pod with Persistent Volume 66

Deploying Cafe Application with Ingress 68

---

**APPENDIX A User Privileges on vSphere 71**

User Privileges on vSphere 71

Erase User Data 75

---

**APPENDIX B Accessibility Features in Cisco Container Platform 77**



# CHAPTER 1

## Cisco Container Platform

Cisco Container Platform is a turnkey, production grade, extensible platform to deploy and manage multiple Kubernetes clusters. It runs on 100% upstream Kubernetes. Cisco Container Platform offers seamless container networking, enterprise-grade persistent storage, built-in production-grade security, integrated logging, monitoring and load balancing.

Cisco Container Platform provides authentication and authorization, security, high availability, networking, load balancing, and operational capabilities to effectively operate and manage Kubernetes clusters. Cisco Container Platform also provides a validated configuration of Kubernetes and can integrate with underlying infrastructure components such as [Cisco HyperFlex](#) and [Cisco ACI](#). The infrastructure provider for Cisco Container Platform is Hyperflex.

Using the Cisco Container Platform web interface, you can create Kubernetes clusters on which you can deploy containerized applications. The clusters are created on the infrastructure provider platform.

The two user personas in Cisco Container Platform are as follows:

- The **Administrator** persona, which is associated with the **Administrator** role.
- The **User** persona, which is associated with the **User** role.

This chapter contains the following topics:

- [Administrator Workflow, on page 1](#)
- [User Workflow, on page 2](#)
- [Accessing Cisco Container Platform Web Interface, on page 3](#)
- [Setting Up Cisco Container Platform, on page 3](#)

## Administrator Workflow

The following table lists the workflow for Cisco Container Platform administrators.

Task	Related Section
Access the Cisco Container Platform web interface with <i>Administrator</i> credentials.	<a href="#">Accessing Cisco Container Platform Web Interface, on page 3</a>
Set up the Cisco Container Platform infrastructure configuration.	<a href="#">Setting Up Cisco Container Platform, on page 3</a>

Task	Related Section
Configure Cisco Smart Software Licensing for your Cisco Container Platform instance.	<a href="#">Configuring Cisco Smart Software Licensing, on page 5</a>
Manage the Cisco Container Platform infrastructure configurations using which clusters are created.	<a href="#">Managing Cisco Container Platform Infrastructure Configuration, on page 13</a>
Create Kubernetes clusters.	<a href="#">Creating Kubernetes Clusters on vSphere On-prem Clusters, on page 23</a> <a href="#">Creating AWS EKS Clusters, on page 33</a>
Add users, assign appropriate roles, and associate the new users to the Kubernetes clusters that you have created.	<a href="#">Managing Users and RBAC, on page 13</a>
Monitor Kubernetes clusters.	<a href="#">Monitoring Health of Cluster Deployments, on page 43</a> <a href="#">Monitoring Logs from Cluster Deployments, on page 44</a>
Manage Kubernetes cluster using the Kubernetes Dashboard.	<a href="#">Managing Kubernetes Clusters, on page 41</a>
Manage the lifecycle of Kubernetes clusters by scaling or upgrading the clusters.	<a href="#">Scaling vSphere Clusters, on page 25</a> <a href="#">Upgrading vSphere Clusters, on page 25</a> <a href="#">Scaling AWS EKS Clusters, on page 34</a>

## User Workflow

The following table lists the workflow for developers assigned with the *User* role.

Task	Related Section
Access the Cisco Container Platform web interface with user credentials.	<a href="#">Accessing Cisco Container Platform Web Interface, on page 3</a>
Monitor Kubernetes clusters that are assigned to the user.	<a href="#">Monitoring Health of Cluster Deployments, on page 43</a> <a href="#">Monitoring Logs from Cluster Deployments, on page 44</a>
Manage the assigned Kubernetes clusters using the Kubernetes Dashboard or CLI.	<a href="#">Managing Kubernetes Clusters, on page 41</a>
Deploy applications on the assigned Kubernetes clusters.	<a href="#">Deploying Applications on Kubernetes Clusters, on page 65</a>



# Accessing Cisco Container Platform Web Interface

## Before you begin

Ensure that you have configured the prerequisites for integrating ACI with Cisco Container Platform.

For more information, refer to the following documents:

- *ACI Integration Requirements* section of the *Cisco Container Platform Installation Guide*
- [Planning](#) and [Prerequisites](#) section of the Cisco ACI and Kubernetes Integration page

Ensure that you have powered on the installer VM on vCenter. The URL of the installer appears on the vCenter **Web console**.

---

**Step 1** Obtain the URL to access the Cisco Container Platform web interface from the vCenter **Web console**.

**Step 2** Access the URL using your web browser.

`https://<Cisco Container Platform IP Address>`

**Note** We recommend that you use the Chrome, Safari, or Firefox browser to access the URL.

**Step 3** Log in to the web interface as an *admin* user using the passphrase given during the Cisco Container Platform installation.

---

## Setting Up Cisco Container Platform



---

**Note** This topic is applicable only for an ACI environment. In a non-ACI environment, the IP address range of the default VIP pool must be expanded to include the additional VIPs for tenant clusters. For more information, see [Managing Networks, on page 19](#).

---

When you log in to Cisco Container Platform for the first time, you need to configure the Cisco Container Platform initial setup using the **Cisco Container Platform Setup** wizard.

---

**Step 1** On the **Welcome** page, click **START THE SETUP**.

**Step 2** In the **ACI Credentials** screen, specify information such as IP address, username, and passphrase of the APIC instance, click **CONNECT**, and then click **NEXT**.

**Step 3** In the **ACI Configuration** screen, perform these steps:

- In the **NAMESERVERS** field, enter the IP address of all the DNS servers that the ACI fabric can access.
- From the **VMM DOMAIN** drop-down list, choose the Virtual Machine Manager Domain (VMM) that you want to use.
- In the **INFRASTRUCTURE VLAN ID** field, enter the VLAN number for layer 2 networking.
- From the **VRF** drop-down list, choose the Virtual Routing and Forwarding (VRF) IP address.
- From the **L3OUT POLICY NAME** drop-down list, choose the ACI object for allowing external internet connectivity.

- f) From the **L3OUT NETWORK NAME** drop-down list, choose the external network that is reachable through the L3OUT object.
- g) From the **AAEP NAME** drop-down list, choose an Attachable Access Entity Profile (AAEP) name to associate the VMM domain with an AAEP.
- h) In the **STARTING SUBNET FOR PODS** field, enter the starting IP address for the IP pool that is used to allocate IP addresses to the pods.
- i) In the **STARTING SUBNET FOR SERVICE** field, enter the starting IP address for the IP pool that is used to allocate IP addresses to the service VLAN.
- j) In the **CONTROL PLANE CONTRACT NAME** field, enter the name of the contract that is provided by the Control Plane endpoint group to allow traffic from the Control Plane cluster to the tenant cluster.
- k) In the **NODE VLAN START ID** field, enter the starting VLAN ID that is used to allocate VLAN to the node.
- l) In the **NODE VLAN END ID** field, enter the ending VLAN ID that is used to allocate VLAN to the node.
- m) In the **OPFLEX MULTICAST RANGE** field, enter a range for the Opflex multicast.
- n) Click **CONNECT**.

**Step 4**

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

For more information on adding, modifying, or deleting an ACI profile, see [Managing ACI Profile, on page 18](#).

---



## CHAPTER 2

# Configuring Cisco Smart Software Licensing

You need to configure Cisco Smart Software Licensing on [Cisco Smart Software Manager \(Cisco SSM\)](#) to easily procure, deploy, and manage licenses for your Cisco Container Platform instance. The number of licenses required depends on the number of VMs necessary for your deployment scenario.

Cisco SSM enables you to manage your Cisco Smart Software Licenses from one centralized website. With Cisco SSM, you can organize and view your licenses in groups called virtual accounts. You can also use Cisco SSM to transfer the licenses between virtual accounts as needed.

You can access Cisco SSM from the Cisco Software Central homepage, under the Smart Licensing area.

Cisco Container Platform is initially available for a 90-day evaluation period after which, you need to register the product.

Based on the level of security required for your environment, you can use one of the following licensing models:

- [Connected Model, on page 5](#)
- [Disconnected Model, on page 6](#)

You need to choose the licensing model when you register Cisco Container Platform.

This chapter contains the following topics:

- [Connected Model, on page 5](#)
- [Disconnected Model, on page 6](#)
- [License Usage and Compliance, on page 6](#)
- [Workflow in a Connected Model, on page 6](#)
- [Workflow in a Disconnected Model, on page 10](#)

## Connected Model

In a connected deployment model, the license usage information is directly sent over the Internet or through an HTTP proxy server to Cisco SSM.

For a higher degree of security, you can opt to use a partially connected deployment model, where the license usage information is sent from Cisco Container Platform to a locally installed VM-based satellite server (Cisco SSM satellite). Cisco SSM satellite synchronizes with Cisco SSM on a daily basis.



**Note** Ensure that you use Cisco SSM Satellite version 5.0 or later. For more information, see [Installing and Configuring Cisco SSM Satellite](#).

See also [Workflow in a Connected Model](#), on page 6.

## Disconnected Model

For the highest degree of security, you can opt to use Specific License Reservations in a fully disconnected model. In this model, you do not need to have access to the Cisco SSM or Cisco SSM satellite. All license changes are processed manually.

See also [Workflow in a Disconnected Model](#), on page 10.

## License Usage and Compliance

Once you register Cisco Container Platform with Cisco SSM, you will receive the **Cisco Container Platform License with Support** license.

Cisco SSM totals the license requirements for all your Cisco Container Platform instances and compares the total license usage to the number of licenses purchased. After the data synchronization, your Cisco Container Platform instance displays one of the following status indicators:

- **Authorized**, when the number of licenses purchased is sufficient
- **Out of Compliance**, when the number of licenses is insufficient

See also [Updating Reserved Licenses](#), on page 11.

- **Authorization Expired**, when the product has not communicated with Cisco SSM for a period of 90 days.

## Workflow in a Connected Model

The following table describes the workflow of Cisco Smart Software Licensing in a connected model.

Task	Related Section
Configure the transport settings using which Cisco Container Platform connects to Cisco SSM or Cisco SSM satellite	<a href="#">Configuring Transport Settings</a> , on page 7
Register the Cisco Container Platform instance	<a href="#">Registering Using Registration Token</a> , on page 7
Manage licenses	<a href="#">Renewing Authorization</a> , on page 8 <a href="#">Reregistering Cisco Container Platform License</a> , on page 9 <a href="#">Deregistering Registration</a> , on page 9

See also [Connected Model, on page 5](#).

## Configuring Transport Settings

By default, Cisco Container Platform directly communicates with the Cisco SSM. You can modify the mode of communication by configuring the transport settings.

### Before you begin

Ensure that you have obtained the registration token for the Cisco Container Platform instance.

- 
- Step 1** Log in to the Cisco Container Platform web interface.
- Step 2** From the left pane, click **Licensing**.  
If you are running Cisco Container Platform in the Evaluation mode, a license notification is displayed on the **Smart Software Licensing** pane.
- Step 3** If a license notification is displayed, click the **edit the Smart Call Home Transport Settings** link.  
Alternatively, click the **Licensing Status** tab, and then click the **View/Edit** link that appears under **Transport Settings**.
- Step 4** In the **Transport Settings** dialog box, perform one of these steps:
- To configure Cisco Container Platform to send the license usage information to Cisco SSM using the Internet:
    1. Click the **DIRECT** radio button.
    2. Configure a DNS on Cisco Container Platform to resolve *tools.cisco.com*.


This is the default setting.
  - To configure Cisco Container Platform to send the license usage information to Cisco SSM using the Cisco SSM satellite:
    1. Click the **TRANSPORT GATEWAY** radio button.
    2. Enter the URL of the Cisco SSM satellite.
  - To configure Cisco Container Platform to send the license usage information to Cisco SSM using a proxy server. For example, an off-the-shelf proxy, such as Cisco Transport Gateway or Apache:
    1. Click the **HTTP/HTTPS PROXY** radio button.
    2. Enter the IP address and port number of the proxy server.
- Step 5** Click **SAVE**.
- 

## Registering Using Registration Token

You need to register your Cisco Container Platform instance with Cisco SSM or Cisco SSM satellite before the 90-day evaluation period expires.

The following figure shows the workflow for registering Cisco Container Platform using a registration token.

Figure 1: Workflow of Registering Using Registration Token

**Before you begin**

Ensure that you have configured the [Configuring Transport Settings](#).

- 
- Step 1** Perform these steps on [Cisco SSM](#) or Cisco SSM satellite to generate a registration token:
- Go to **Inventory** > **<Choose your virtual account>** > **General**, and then click **New Token**.
  - If you want to enable higher levels of encryption for the products registered using the registration token, check the **Allow export-controlled functionality on the products registered with this token** check box.
 

**Note** This option is available only if you are compliant with the Export-Controlled functionality.
  - Download or copy the token.
- Step 2** Perform these steps in the Cisco Container Platform web interface to register the registration token and complete the license registration process:
- From the left pane, click **Licensing**.
  - In the license notification, click **Register**.  
The **Smart Software Licensing Product Registration** dialog box appears.
  - In the **Product Instance Registration Token** field, enter, paste, or upload the registration token that you generated in Step 1.
  - Click **REGISTER** to complete the registration process.
- If registering the token fails, you can [Reregistering Cisco Container Platform License](#) the Cisco Container Platform instance using a new token.
- 

## Renewing Authorization

By default, the authorization is automatically renewed every 30 days. However, Cisco Container Platform allows a user to manually initiate the authorization renew in case the automatic renewal process fails. The authorization expires if Cisco Container Platform is not connected to Cisco SSM or Cisco SSM satellite for 90 days and the licenses consumed by Cisco Container Platform are reclaimed and put back to the license pool.

**Before you begin**

Ensure that the Cisco Container Platform instance is registered with Cisco SSM or Cisco SSM satellite.

- 
- Step 1** Log in to the Cisco Container Platform web interface.
  - Step 2** From the left pane, click **Licensing**.
  - Step 3** From the **ACTIONS** drop-down list, choose **Renew Authorization Now**.
  - Step 4** Click **OK** in the **Renew Authorization** dialog box.  
Cisco Container Platform synchronizes with Cisco SSM or Cisco SSM satellite to check the license authorization status and Cisco SSM or Cisco SSM satellite reports back the status to Cisco Container Platform, on a daily basis.
- 

## Reregistering Cisco Container Platform License

You can reregister Cisco Container Platform with Cisco SSM or Cisco SSM satellite by deregistering it and registering it again, or by using a register force option.

### Before you begin

Ensure that you have obtained a new registration token from Cisco SSM or Cisco SSM satellite.

- 
- Step 1** Log in to the Cisco Container Platform web interface.
  - Step 2** From the left pane, click **Licensing**.
  - Step 3** From the **ACTIONS** drop-down list, choose **Reregister**.
  - Step 4** In the **Product Instance Registration Token** field of the **Smart Software Licensing Product Reregistration** dialog box, enter the registration token that you generated using Cisco SSM or Cisco SSM satellite.
  - Step 5** Click **REGISTER** to complete the registration process.  
Cisco Container Platform sends a request to Cisco SSM or Cisco SSM satellite to check the registration status and Cisco SSM or Cisco SSM satellite reports back the status to Cisco Container Platform, on a daily basis.
- 

## Deregistering Registration

You can deregister the Cisco Container Platform instance from Cisco SSM or Cisco SSM satellite to release all the licenses from the current Virtual account and the licenses are available for use by other products in the virtual account. Deregistering disconnects Cisco Container Platform from Cisco SSM or Cisco SSM satellite.

### Before you begin

Ensure that the Cisco Container Platform instance is registered with Cisco SSM or Cisco SSM satellite.

- 
- Step 1** Log in to the Cisco Container Platform web interface.
  - Step 2** From the left pane, click **Licensing**.
  - Step 3** From the **ACTIONS** drop-down list, choose **Deregister**.
  - Step 4** Click **DEREGISTER** in the confirmation dialog box.  
Cisco Container Platform sends a request to Cisco SSM or Cisco SSM satellite to check the deregistration status and Cisco SSM or Cisco SSM satellite reports back the status to Cisco Container Platform, on a daily basis.
-

## Workflow in a Disconnected Model

The following table describes the workflow of Cisco Smart Software Licensing in a disconnected model.


Task	Related Section
Register the Cisco Container Platform instance	<a href="#">Registering Using Specific License Reservations, on page 10</a>
Manage licenses	<a href="#">Updating Reserved Licenses, on page 11</a> <a href="#">Deregistering Registration by Returning Reserved Licenses, on page 12</a>

See also [Disconnected Model, on page 6](#).

## Registering Using Specific License Reservations

The following figure shows the workflow for registering Cisco Container Platform using Specific License Reservations (SLRs).

**Figure 2: Workflow of Registering Using SLRs**



### Before you begin

Ensure that your smart account is authorized for License Reservation.

**Step 1** Perform these steps in the Cisco Container Platform web interface:

- From the left pane, click **Licensing**.
- In the license notification, click **Register**.  
The **Smart Software Licensing Product Registration** dialog box appears.



- c) Click the **start here** link to initiate **Specific License Reservation**.
- d) In the **Ensure Smart Account Can Use License Reservation** dialog box, check the **YES, MY SMART ACCOUNT IS LICENSE RESERVATION ENABLED** check box, and then click **CONTINUE**.
- e) In the **Smart License Reservation** dialog box, click **GENERATE REQUEST CODE**.
- f) Copy and save the **License Reservation Request Code**.

**Step 2** Perform these steps on Cisco SSM:

- a) Go to **Inventory > <Choose your virtual account> > Licenses**, and then click **License Reservation**.

**Note** If you cannot find the **License Reservation** button, contact Cisco support to enable this feature.

The **Smart License Reservation** dialog box appears.

- b) In the **Reservation Request Code** field, enter, paste, or upload the **License Reservation Request Code** from Step 1, and then click **Next**.
- c) In the **Licenses to Reserve** area, check the **Reserve a specific license** check box.
- d) Under the **Quantity To Reserve** column, type the number of licenses to reserve for your product instance, and then click **Next**.
- e) Click **Generate Authorization Code**.
- f) Copy or download the authorization code.

**Step 3** Perform these steps in the Cisco Container Platform web interface:

- a) From the left pane, click **Licensing**.
- b) In the license notification, click **ENTER RESERVATION AUTHORIZATION CODE**.  
The **Enter Reservation Authorization Code** dialog box appears.
- c) In the **RESERVATION AUTHORIZATION CODE** field, enter, paste, or upload the authorization code from Step 2.
- d) Click **RESERVE**.  
The number of licenses reserved is displayed in **Licensing > Smart Licensing Usage**. When your license usage is out of compliance, you can [Updating Reserved Licenses](#).

---

See also [License Usage and Compliance](#), on page 6.

## Updating Reserved Licenses

If your license usage is out of compliance, or you want to increase the number of licenses for your Cisco Container Platform instance, you can update the number of reserved licenses.

**Step 1** Perform these steps in Cisco SSM:

- a) Go to **Inventory > <Choose your virtual account> > Product Instances**.
- b) From the **ACTIONS** drop-down list of your Cisco Container Platform instance, choose **Update reserved licenses**.  
The **Update License Reservation** dialog box appears.
- c) In the **Licenses to Reserve** area, check the **Reserve a specific license** check box.
- d) Under the **Quantity To Reserve** column, modify the number of reserved licenses for your product instance, and then click **Next**.
- e) Click **Generate Authorization Code**.
- f) Copy or download the authorization code.

**Step 2** Perform these steps in the Cisco Container Platform web interface:

- a) From the left pane, click **Licensing > Smart Licensing Usage**.
- b) From **ACTIONS** drop-down list, click **Update Reservation**.  
The **Enter Reservation Authorization Code** dialog box appears.
- c) In the **RESERVATION AUTHORIZATION CODE** field, enter, paste, or upload the reservation code from Step 1.
- d) Click **RESERVE**.  
A license reservation confirmation code is generated.

**Step 3** Perform these steps in Cisco SSM:

- a) In the **Update License Reservation** dialog box, click **Enter Confirmation Code**.
- b) In the **Reservation Confirmation Code** field, enter, paste, or upload the confirmation code that you received in Step 2.
- c) Click **OK**.  
The number of reserved licenses for your product instance is updated on Cisco SSM.

See also [License Usage and Compliance](#), on page 6.

## Deregistering Registration by Returning Reserved Licenses

You can deregister your Cisco Container Platform instance from Cisco SSM by returning the reserved licenses. Consequently, these licenses are available for other product instances.

**Step 1** Perform these steps in the Cisco Container Platform web interface:

- a) From the left pane, click **Licensing**.  
The **Smart Software Licensing** page appears.
- b) From **ACTIONS** drop-down list, click **Return reserved licenses**.  
The **Confirm Return Licenses** dialog box appears.
- c) Click **GENERATE RESERVATION RETURN CODE**.  
If you cannot generate a **Reservation Return Code**, contact Cisco support.

**Caution** The licenses remain reserved in Cisco SSM and cannot be used for other product instances until they are released on Cisco SSM.

**Step 2** Follow these steps in the Cisco SSM to release the licenses:

- a) Go to **Inventory > <Choose your virtual account> > Product Instances**.
- b) From the **ACTIONS** drop-down list of your Cisco Container Platform instance, choose **Remove**.  
The **Remove Product Instance** dialog box appears.
- c) In the **Reservation Return Code** field, enter, paste, or upload the reservation return code from Step 1.
- d) Click **Remove Product Instance**.

The reserved licenses are released, and you can use them for other product instances. Your Cisco Container Platform instance goes back to the evaluation mode.



## CHAPTER 3

# Managing Cisco Container Platform Infrastructure Configuration

---

This chapter contains the following topics:

- [Managing Users and RBAC, on page 13](#)
- [Managing Provider Profile, on page 15](#)
- [Managing ACI Profile, on page 18](#)
- [Managing Networks, on page 19](#)

## Managing Users and RBAC

Cisco Container Platform provides Role-based Access Control (RBAC) through built-in static roles, namely the *Administrator* and *User* roles. Role-based access allows you to use local accounts and LDAP for authentication and authorization.

## Configuring Local Users

Cisco Container Platform allows you to manage local users. An administrator can add a user, and assign an appropriate role and cluster(s) to the user.



---

**Caution** Use of local authentication is not recommended and is considered less secure for production data.

---

### Before you begin

Ensure that you have configured LDAP Server for authentication of Cisco Container Platform users.

For more information, see [Configuring AD Servers, on page 14](#).

- 
- Step 1** From the left pane, click **User Management**, and then click the **Users** tab.
- Step 2** Click **NEW USER**.
- Step 3** Specify information such as first name, last name, username, passphrase, and role for the user.
- Step 4** Click **SUBMIT**.

The new user is displayed on the **User Management** page.

**Note** You can edit or delete a user by using the options available under the **ACTIONS** column.

## Changing Login Passphrase

**Step 1** From the left pane, click **User Management**, and then click the **Users** tab.

**Step 2** From the drop-down list displayed under the **ACTIONS** column, choose **Edit** corresponding to your name.

**Note** Administrators can change passphrase and role for other users as well.

**Step 3** Change the passphrase and role assigned as necessary, and click **SUBMIT**.

## Recovering Login Passphrase for Local Admin

**Step 1** Perform one of the following steps:

- a) If you have SSH access to the Cisco Container Platform Control Plane nodes, log in to a Cisco Container Platform Control Plane node.
- b) If you have the Kubeconfig file, save it in the `$HOME/.kube` directory. You can specify other kubeconfig files by [setting the KUBECONFIG environment variable](#) or by setting the `--kubeconfig` flag.

**Step 2** List the available pods.

```
kubectl get pods
```

**Step 3** Search for the pod that has the following format:

```
kaas-corc-xxxxxxxx-xxxx
```

**Step 4** Reset the login passphrase for the admin user.

```
kubectl exec kaas-corc-7df5d76f87-55n7b ./password_reset
Password reset for 'admin' user : <50-char-long-random-string>
```

The local admin passphrase is reset to a 50-character random string. You can choose to continue using this passphrase, or reset the passphrase by [Changing Login Passphrase](#).

## Configuring AD Servers

LDAP authentication is performed using a service account that can access the LDAP database and query for user accounts. You will need to configure the AD server and service account in Cisco Container Platform.

**Step 1** From the left pane, click **User Management**, click the **Active Directory** tab, and then click **EDIT**.

**Step 2** In the **SERVER IP ADDRESS** field, type the IP address of the AD server.

- Step 3** In the **PORT** field, type the port number for the AD server.
- Step 4** For improved security, we recommend that you check **STARTTLS**.
- Step 5** In the **BASE DN** field, specify the domain name of the AD server for all the accounts that you have.
- Step 6** In the **ACCOUNT USERNAME** field, specify the service account name that is used for accessing the LDAP server.
- Step 7** In the **PASSPHRASE** field, type the passphrase of the AD account.
- Step 8** Click **SUBMIT**.
- 

## Configuring AD Groups

Cisco Container Platform allows you to manage users using AD groups. An administrator can add users to AD groups, and then assign appropriate roles and clusters to the groups.

### Before you begin

Ensure that you have configured the AD server that you want to use.

For more information on configuring AD servers, see [Configuring AD Servers, on page 14](#).

---

- Step 1** From the left pane, click **User Management**, and then click the **Groups** tab.
- Step 2** Click **ADD GROUP**.
- Step 3** Specify information such as the name of the AD group and the role you want to assign to the group.
- Note** If the AD group is associated with the *Administrator* role, by default, access is provided to all clusters. But, if the AD group is associated with the *User* role, you need to assign a cluster.
- Step 4** From the **CLUSTERS** drop-down list, choose the names of the cluster that you want to assign to the AD group.
- Step 5** Click **SUBMIT**.
- 

## Managing Provider Profile

Cisco Container Platform enables you to define the provider profile on which clusters can be created.

You can configure multiple provider profiles in an instance of Cisco Container Platform and use the same provider profile for multiple clusters.

## Adding Provider Profile

After your Cisco Container Platform control plane is available, log in to the Cisco Container Platform web interface, and then add the required provider profiles.

This section contains the following topics:

## Adding vSphere Provider Profile

### Before you begin

Cisco Container Platform interacts with vSphere through the user that you configure when you add a provider profile. Hence, you need to ensure that this user has the necessary privileges.

For more information on the vSphere user privileges, see [User Privileges on vSphere, on page 71](#).

- 
- Step 1** From the left pane, click **Infrastructure Providers**.  
The **Infrastructure Providers** screen appears.
- Step 2** Click **NEW PROVIDER** and enter information such as name, description, address, port, username and passphrase of the provider profile.
- Step 3** Click **ADD**.  
The vSphere provider profile that you added is displayed on the **Infrastructure Providers > vSphere** screen.
- 

## Adding Amazon Provider Profile

- 
- Step 1** From the left pane, click **Infrastructure Provider**.  
The **Infrastructure Providers** screen appears.
- Step 2** Click **NEW PROVIDER** and enter the following information:
- In the **PROVIDER NAME** field, enter a name for the related Amazon account.
  - In the **ACCESS KEY ID** field, enter the key ID for the related Amazon account.
  - In the **SECRET ACCESS KEY** field, enter the access key for the related Amazon account.
  - Click **ADD**.

**Note** The access key and secret must not be from your AWS root user account.

The Amazon provider profile that you added is displayed on the **Infrastructure Providers > AWS** screen.

For more information on administering AWS EKS clusters, see [Administering Amazon Web Services \(AWS\) EKS Clusters, on page 29](#).

---

## Adding Azure Provider Profile

### Before you begin

Ensure that you have created a service principal in your Azure account and noted down the values of the `id`, `appId`, `password`, and `tenant` parameters. For more information, see [Creating Service Principals, on page 17](#).

- 
- Step 1** From the left pane, click **Infrastructure Provider**.  
The **Infrastructure Providers** screen appears.
- Step 2** Click the **NEW PROVIDER** and enter the following information:

- a) In the **NAME** field, enter a name for your Azure account.
- b) In the **CLIENT ID** field, enter the value of the `appID` parameter from [Creating Service Principals, on page 17](#).
- c) In the **CLIENT SECRET** field, enter the value of the `password` parameter from [Creating Service Principals, on page 17](#).
- d) In the **TENANT ID** field, enter the value of the `tenant` parameter from [Creating Service Principals, on page 17](#).
- e) In the **SUBSCRIPTION ID** field, enter the value of the `id` parameter from [Creating Service Principals, on page 17](#).
- f) Click **ADD**.

---

The Azure provider profile is displayed on the **Infrastructure Providers > Azure** screen.

For more information on administering Azure Kubernetes Service (AKS) clusters, see [Administering Azure Kubernetes Service \(AKS\) Clusters, on page 37](#).

## Creating Service Principals

---

**Step 1** Login to the [Azure Portal](#).

**Step 2** [Install the Azure CLI](#).

**Step 3** Follow these steps to configure the Azure CLI to use the Azure account that you want use with Cisco Container Platform.

- a) Log in to the Azure CLI.

```
az login
```

The URL to the device login page and an authentication code is displayed.

- b) Use a browser to access the device login page, enter the code that you have received, and then click **Continue**.
- c) Choose your Azure account.

**Step 4** From the command output on the Azure CLI, note down the value of the `id` parameter. This value is required while [Adding Azure Provider Profile](#) to Cisco Container Platform.

For example:

```
{
  "cloudName": "AzureCloud",
  "id": "aaaaaaaa-bbbb-1111-cc22-ddddd3333444ddd",
  "isDefault": true,
  "name": "Microsoft Azure Enterprise",
  "state": "Enabled",
  "tenantId": "xxxxx-yyyy-999z-9090-uuuuu999uuuu",
  "user": {
    "name": "user@org.com",
    "type": "user"
  }
}
```

**Step 5** Create a service principal using the Azure CLI.

```
az ad sp create-for-rbac -n myserviceprincipal
```

Where, `myServicePrincipal` is the name of the service principal. You may give any name for your service principle.

**Step 6** From the command output on the Azure CLI, note down the values of the `appID`, `password`, and `tenant` parameters. These values are required while [Adding Azure Provider Profile](#) to Cisco Container Platform.

For example:

```
{
  "appID": "qqqqqqqq-a1a1-2b2b-9z9z-www1111vvvv",
```

```

    "displayName": "myserviceprincipal",
    "name": "http://myserviceprincipal",
    "password": "mmmmmm-n0n0-plp1-q3q3-uuuu0000vvvv",
    "tenant": "xxxxx-yyyy-999z-9090-uuuuu999uuuu"
  }

```

---

## Modifying Provider Profile

---

- Step 1** From the left pane, click **Infrastructure Providers**. The **Infrastructure Providers** screen appears.
- Step 2** Click the **vSphere**, **AWS**, or **Azure** tab as necessary.
- Step 3** From the drop-down list displayed under the **ACTIONS** column, choose **Edit** for the provider profile that you want to modify.
- Step 4** Change the provider details as necessary and click **SUBMIT**.
- 

## Deleting Provider Profile

---

- Step 1** From the left pane, click **Infrastructure Providers**.
- Step 2** Click the **vSphere**, **AWS**, or **Azure** tab as necessary.
- Step 3** From the drop-down list displayed under the **ACTIONS** column, choose **Delete** corresponding to the provider profile that you want to delete.
- Step 4** Click **DELETE** in the confirmation dialog box.
- 

## Managing ACI Profile

Cisco Container Platform enables you to define ACI profiles using which tenant clusters can be created.

You can define multiple ACI profiles and use the same profile for multiple clusters.

## Adding ACI Profile

---

- Step 1** From the left pane, click **ACI Profiles**.
- Step 2** Click **Add New ACI Profile** and perform these steps:
- a) Specify information such as profile name, IP address, username, and passphrase of the ACI instance.
 

**Note** If there is more than one host, use a comma-separated host list in the **APIC IP ADDRESSES** field.
  - b) In the **NAMESERVERS** field, enter the IP address of all the DNS servers that the ACI fabric can access.



- c) From the **VMM DOMAIN** drop-down list, choose the Virtual Machine Manager Domain (VMMD) that you want to use.
- d) In the **INFRASTRUCTURE VLAN ID** field, enter the VLAN number for layer 2 networking.
- e) From the **VRF** drop-down list, choose the Virtual Routing and Forwarding (VRF) IP address.
- f) From the **L3OUT POLICY NAME** drop-down list, choose the ACI object for allowing external internet connectivity.
- g) From the **L3OUT NETWORK NAME** drop-down list, choose the external network that is reachable through the L3OUT object.
- h) From the **AAEP NAME** drop-down list, choose an Attachable Access Entity Profile (AAEP) name to associate the VMM domain with an AAEP.
- i) In the **STARTING SUBNET FOR PODS** field, enter the starting IP address for the IP pool that is used to allocate IP addresses to the pods.
- j) In the **STARTING SUBNET FOR SERVICE** field, enter the starting IP address for the IP pool that is used to allocate IP addresses to the service VLAN.
- k) In the **CONTROL PLANE CONTRACT NAME** field, enter the name of the contract that is provided by the Control Plane endpoint group to allow traffic from the Control Plane cluster to the tenant cluster.
- l) In the **NODE VLAN START ID** field, enter the starting VLAN ID that is used to allocate VLAN to the node.
- m) In the **NODE VLAN END ID** field, enter the ending VLAN ID that is used to allocate VLAN to the node.
- n) In the **OPFLEX MULTICAST RANGE** field, enter a range for the Opflex multicast.

**Step 3** Click **SUBMIT**.

---

## Modifying ACI Profile

---

**Step 1** From the left pane, click **ACI Configuration**.

**Step 2** From the drop-down list displayed under the **ACTIONS** column, choose **Edit** for the ACI profile that you want to modify.

**Step 3** Change the ACI profile details as necessary and click **SUBMIT**.

---

## Deleting ACI Profile

---

**Step 1** From the left pane, click **ACI Configuration**.

**Step 2** From the drop-down list displayed under the **ACTIONS** column, choose **Delete** for the ACI profile that you want to delete.

**Step 3** Click **DELETE** in the confirmation dialog box.

---

## Managing Networks



**Note** This section is applicable only for a non-ACI environment.

---

Based on the information that you provided during installation, Cisco Container Platform creates a network, subnet, and an IP pool. Cisco Container Platform requires a minimum of six IP addresses. After installation, you can add or modify the IP pool range, subnet, or network by using the Cisco Container Platform web interface. The IP address pools define the IP address ranges that are managed by Cisco Container Platform.




---

**Note** You must ensure that the range of IP addresses in the VIP pools is outside of the IP addresses that are assigned by DHCP.

---

The IP addresses that are managed by Cisco Container Platform are used for the following purposes:

- A VIP for the Cisco Container Platform Kubernetes Master
- A VIP for the external Ingress access of Cisco Container Platform
- Static Interface IP addresses for master and worker nodes in each tenant cluster
- A VIP for the Kubernetes master of each tenant cluster
- A VIP for the external NGINX Ingress Controller of each tenant cluster
- VIPs for any LoadBalancer type Kubernetes Service of a tenant cluster

To create tenant clusters, you need to configure a subnet during cluster creation. The total number of free IP addresses across all the pools for that subnet must be at least:

$3 + (\text{Number of tenant worker nodes})$

## Modifying Networks

---

- Step 1** From the left pane, click **Networks**.  
The **Networks** page displays the default network.
- Step 2** From the drop-down list displayed under the **ACTIONS** column, choose **Edit** for the network that you want to modify. Alternatively, click the **SUBNETS** tab or the **POOLS** tab, and then click **EDIT** from the right pane to view the **Edit** dialog box.
- Step 3** Modify the network name as necessary and click **SUBMIT**.
- 

## Adding Subnets

If you want to allocate VIP from a different subnet CIDR you need to add the subnet.

---

- Step 1** From the left pane, click **Networks**, and then click the network to which you want to add a subnet.
- Step 2** From the right pane, click **NEW SUBNET**.
- Step 3** Enter a name and CIDR for the subnet.
- Step 4** Enter a gateway IP address that you want to use.  
A gateway IP address allows a cluster to access other networks.

- Step 5** Enter the IP address of the necessary DNS nameserver.  
You can click **+NAMESERVER** to enter IP addresses of additional nameservers.
- Step 6** Click **SUBMIT**.
- 

## Modifying Subnets

---

- Step 1** From the left pane, click **Networks**, and then click the network that contains the subnet you want to modify.
- Step 2** Click the **SUBNETS** tab.
- Step 3** From the drop-down list displayed under the **ACTIONS** column, choose **Edit** for the subnet that you want to modify.
- Step 4** Modify the subnet name, CIDR, gateway IP or list of nameservers as necessary.
- Step 5** Click **SUBMIT**.
- 

## Adding VIP Pool

---

- Step 1** From the left pane, click **Networks**, and then click the network to which you want to add a VIP pool.
- Step 2** From the right pane, click **NEW POOL**.
- Step 3** Specify a name, subnet and IP address range for the VIP pool.
- Step 4** Click **SUBMIT**.
- 

## Modifying VIP Pool

---

- Step 1** From the left pane, click **Networks**, and then click the network that contains the VIP pool you want to modify.
- Step 2** Click the **POOLS** tab.
- Step 3** From the drop-down list displayed under the **ACTIONS** column, choose **Edit** for the VIP pool that you want to modify.
- Step 4** Change the pool name and the IP address as necessary, and then click **SUBMIT**.
-





## CHAPTER 4

# Administering vSphere On-prem Clusters

---

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

This chapter contains the following topics:

- [Creating Kubernetes Clusters on vSphere On-prem Clusters, on page 23](#)
- [Upgrading vSphere Clusters, on page 25](#)
- [Scaling vSphere Clusters, on page 25](#)
- [Configuring Node Pools, on page 25](#)
- [Deleting vSphere Clusters, on page 27](#)

## Creating Kubernetes Clusters on vSphere On-prem Clusters

---

- Step 1** From the left pane, click **Clusters**, and then click the **vSphere** tab.
- Step 2** Click **NEW CLUSTER**.
- Step 3** In the **Basic Information** screen, specify the following information:
- a) From the **INFRASTRUCTURE PROVIDER** drop-down list, choose the provider related to your Kubernetes cluster. For more information, see [Adding vSphere Provider Profile , on page 16](#).
  - b) In the **KUBERNETES CLUSTER NAME** field, enter a name for your Kubernetes tenant cluster.
  - c) In the **VERSION** drop-down list, choose the version of Kubernetes that you want to use for creating the cluster.
  - d) In the **CNI** field, enter the Container Network Interface (CNI) that you want to use.
  - e) In the **DESCRIPTION** field, add a description.
  - f) If you are using ACI, specify the ACI profile, see [Adding ACI Profile, on page 18](#).
  - g) Click **NEXT**.
- Step 4** In the **Provider Settings** screen, specify the data center, cluster, resource pool, storage class, network, HyperFlex local network, datastore, and VM template that you want to use, and then click **NEXT**.

- Note**
- Ensure that DRS and HA are enabled on the cluster that you choose in this step. For more information on enabling DRS and HA on clusters, refer to the *Cisco Container Platform Installation Guide*.
  - Ensure that the datastore that you choose in this step is accessible to the hosts in the cluster.
  - For Network, select a subnet with an adequate number of free IP addresses. For more information, see [Managing Networks, on page 19](#). The selected network must have access to vCenter.
  - The following steps are only required on HyperFlex systems:
    - The selected network must have access to the HypexFlex Connect server to support HyperFlex Storage Provisioners.
    - For HyperFlex Local Network, select **k8-priv-iscsivm-network** to enable HyperFlex Storage Provisioners.

**Step 5** In the **GPU Configuration** screen, specify the type and number of GPUs that you want to use for worker nodes in the cluster.

**Note** GPU Configuration is applicable only if you have GPUs in your HyperFlex cluster.

**Step 6** In the **Node Configuration** screen, specify the following information, and then click **NEXT**:

- The number of worker and master nodes, and their VCPU and memory configurations.
- If you have GPUs in your HyperFlex cluster, you can select the type of GPU you want to use for cluster creation. This is an optional field.
- The SSH public key that you want to use for creating the cluster. Ensure that you use the Ed25519 or ECDSA format for the public key.

**Note:** As RSA and DSA are less secure formats, Cisco prevents the use of these formats.

- The VM username that you want to use as the login for the VM.
- The subnet that you want to use for this cluster.
- The number of load balancer IP addresses for this cluster.  
For more information, see [Load Balancer Services, on page 55](#).
- The IP addresses in CIDR notation that you want to use as the pod subnet.
- Whether or not you want to enable Istio
- A root CA certificate to allow tenant clusters to securely connect to additional services

**Step 7** In the **Harbor Registry** screen, specify if you want to enable Harbor. If no, click **NEXT**. If yes, you must specify the following information, and then click **NEXT**:

- a) Ensure the switch to enable Harbor is activated
- b) A password for Harbor server admin
- c) The immutable registry size in gigabits

**Step 8** In the **Summary** screen, verify the configuration, and then click **FINISH**.

The cluster deployment takes 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](#), on page 65.

---

## Upgrading vSphere Clusters

### Before you begin

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

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

---

- Step 1** From 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**.
- 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 vSphere Clusters

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 25.

## 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** From 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, on page 19](#).
  - 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, on page 19](#).
  - 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 vSphere Clusters

### 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** From 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.
-





## CHAPTER 5

# Administering Amazon Web Services (AWS) EKS Clusters

Integrating Cisco Container Platform with Amazon Web Services (AWS) allows you to deploy and run containerized applications across both Cisco-based on-prem environments and the AWS cloud.

This chapter contains the following topics:

- [Prerequisites for Configuring AWS EKS Clusters, on page 29](#)
- [Amazon IAM Authentication, on page 32](#)
- [Configuring Control Plane Proxy for EKS Access, on page 33](#)
- [Creating AWS EKS Clusters, on page 33](#)
- [Scaling AWS EKS Clusters, on page 34](#)
- [Deleting AWS EKS Clusters, on page 35](#)

## Prerequisites for Configuring AWS EKS Clusters

The prerequisites for configuring AWS EKS clusters are as follows:

### Amazon Resource Requirements

The following table describes the default limits for the Amazon resources that you may need to increase depending on your Cisco Container Platform deployment requirements.



**Note** To increase the limits for a specific resource, you need to contact [Amazon support](#).

Amazon Resource	Default Limit	Description
Network Address Translation (NAT) gateway for each AWS account	14	Each EKS cluster uses three NAT gateways. With the default setting, you are limited to four clusters.
Amazon Virtual Private Cloud (Amazon VPC) for each AWS account	3	Each tenant cluster requires a separate Amazon VPC.

Amazon Resource	Default Limit	Description
Amazon Elastic Container Service for Kubernetes (Amazon EKS) cluster for each AWS account	3	<b>Note</b> Changes to the Amazon EKS cluster limit are updated only on Thursdays.
Elastic IP address for each region	5	Each EKS cluster uses three elastic IP addresses. For more information, see <a href="#">Amazon VPC Limits</a> .
Internet gateway for each region	5	Each EKS cluster uses one internet gateway.

## Adding AMI Files to your Amazon Account


Cisco Container Platform generates a specific AMI (Amazon Machine Image) file with each product release. The AMI file ensures that compatible packages are available for successful tenant cluster creation.

To make the AMI file available to your Amazon account, you must [submit a support case](#) that includes your 12 digit Amazon account ID. You will be notified when the AMI is available within your Amazon account.

## Creating AWS Roles

- 
- Step 1** Log in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
- Step 2** From the navigation pane of the IAM console, click **Roles**, and then click **Create role**.
- Step 3** Under **Select type of trusted entity**, click **Another AWS account**.
- Step 4** In the **Account ID** field, enter your **AWS Account ID**, and then click **Next**.  
The AWS account number must be a trusted entity so that Cisco Container Platform can use the Role ARN during EKS cluster creation.
- Step 5** Skip the screen to choose permission policies and permission boundary and click **Next**.
- Step 6** Add metadata to the role by attaching tags of your choice as key–value pairs and click **Next**.
- Step 7** In the **Role name** field, enter the name for the role as `k8s-ccp-user` or any other name of your choice.
- Step 8** In the **Description** field, enter a description of your choice and click **Create role**.
- Step 9** After the role is created, navigate to the created role and verify the following details of the role:
- Click the **Permissions** tab to verify that permissions are not set.
  - Click the **Trust Relationships** tab to verify that a trust relationship exists for the AWS account that you entered during creation of the Role ARN.

Figure 3: AWS Management Console-Trust Relationships Tab



## AWS Account Policy Requirements

### Provider permissions

If the AWS provider account is not a root account then you must ensure that the account has the permissions needed to create the EKS and EC2 resources.

The minimum permissions needed are included in the [Sample aws-provider-policy.json File](#). You can create and import this file to configure the necessary permissions.

### Sample aws-provider-policy.json File

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "cloudformation:*",
        "elasticloadbalancing:*",
        "autoscaling:*",
        "ec2:*",
        "eks:*",
        "ecr:*",
        "ecs:*",
        "s3:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:List*"
      ],
      "Resource": "*"
    }
  ]
}
```

```

        "iam:Get*",
        "iam:PassRole",
        "iam:AddRoleToInstanceProfile",
        "iam:RemoveRoleFromInstanceProfile",
        "iam:CreateRole",
        "iam:CreateInstanceProfile",
        "iam>DeleteInstanceProfile",
        "iam>DeleteRole",
        "iam>DeleteRolePolicy",
        "iam:AttachRolePolicy",
        "iam:DetachRolePolicy",
        "iam:PutRolePolicy",
        "iam:*AccessKey*",
        "iam:*MFA*"
    ],
    "Resource": "*"
}
]
}

```

## Amazon IAM Authentication

By default, the AWS IAM identity is used to authenticate EKS clusters to establish connection. Cisco Container Platform uses [AWS IAM Authenticator](#) to authenticate on-prem cluster using the AWS IAM identity. This authentication provides a consistent, unified identity scheme across both on-premise and AWS EKS clusters.

The AWS IAM Authenticator fulfills both a client and server function. On the client side, the authenticator generates, tokenizes and transmits a pre-signed URL to the server-side for identity validation. The client is a Go binary, installed on your workstation, which is transparently invoked by kubectl each time you interact with your Kubernetes cluster. The server-side is a containerized instance of AWS IAM Authenticator running as a DaemonSet on the Kubernetes master nodes. This interacts with the AWS Secure Token Service (STS) to perform identity validation. Cisco Container Platform takes care of the initial server-side configuration and provides a preconfigured `kubeconfig` file for admin users to download.




---

**Note** You need to ensure that the AWS IAM Authenticator is available within your `$PATH` while using kubectl to interact with the clusters.

---

## Enabling Common Identity

Within the Cisco Container Platform web interface, users are able to select a common identity scheme for clusters. After the clusters are provisioned, you can apply a shared RBAC policy.




---

**Note** The use of IAM Authentication is implicitly enabled for EKS clusters. Cisco Container Platform can map a user supplied IAM role to the EKS cluster and configuring IAM auth for on-premises clusters.

---

# Configuring Control Plane Proxy for EKS Access

If your Control Plane VMs need proxy configuration to access the internet, specifically AWS API endpoints, you need to configure Cisco Container Platform application deployments with the proxy information.

**Step 1** SSH to the Control Plane cluster master VM.

**Step 2** Run the following commands to specify the proxy information:

**Note** Replace `<Proxy_IP_address>` with the IP address of your proxy server.

```
kubectl patch deployment kaas-api --patch
'{"spec":{"template":{"spec":{"containers":[{"name":"api","containers":[{"name":"proxy","image":"<Proxy_IP_address>:5000/cisco-ccp-eks-api","imagePullPolicy":"Always"}]}]}}}}'

kubectl patch deployment kaas-ccp-eks-operator --patch
'{"spec":{"template":{"spec":{"containers":[{"name":"operator","containers":[{"name":"proxy","image":"<Proxy_IP_address>:5000/cisco-ccp-eks-operator","imagePullPolicy":"Always"}]}]}}}}'

kubectl patch daemonset aws-iam-authenticator -n kube-system --patch
'{"spec":{"template":{"spec":{"containers":[{"name":"authenticator","containers":[{"name":"proxy","image":"<Proxy_IP_address>:5000/cisco-ccp-eks-iam-authenticator","imagePullPolicy":"Always"}]}]}}}}'
```

## Creating AWS EKS Clusters

### Before you begin

- Ensure that you have added your Amazon provider profile. For more information, see [Adding Amazon Provider Profile, on page 16](#).
- Ensure that you have added the required AMI files to your account. For more information, see [Adding AMI Files to your Amazon Account, on page 30](#).
- Ensure that you have created an AWS IAM Role for the Cisco Container Platform usage to create AWS EKS Clusters. For more information, see [Creating AWS Roles, on page 30](#).

**Step 1** From the left pane, click **Clusters**, and then click the **AWS** tab.

**Step 2** Click **NEW CLUSTER**.

**Step 3** In the **Basic Information** screen, enter the following information:

- From the **INFRASTRUCTURE PROVIDER** drop-down list, choose the provider related to the appropriate Amazon account.
- From the **AWS REGION** drop-down list, choose an appropriate AWS region.

**Note** Not all regions support EKS. Ensure that you select a supported region. Currently, Cisco Container Platform supports the **ap-northeast-1**, **ap-northeast-2**, **ap-southeast-1**, **ap-southeast-2**, **eu-central-1**, **eu-north-1**, **eu-west-1**, **eu-west-2**, **eu-west-3**, **us-east-1**, **us-east-2**, and **us-west-2** regions.

- In the **KUBERNETES CLUSTER NAME** field, enter a name for your cluster.
- Click **NEXT**.

- Step 4** In the **Node Configuration** screen, specify the following information:
- From the **INSTANCE TYPE** drop-down list, choose an [instance type](#) for your cluster.
  - From the **MACHINE IMAGE** drop-down list, choose an appropriate Cisco Container Platform Amazon Machine Image (AMI) file.  
To add AMI files to your Amazon account, see [Adding AMI Files to your Amazon Account, on page 30](#).
  - In the **WORKER COUNT** field, enter an appropriate number of worker nodes.
  - In the **SSH PUBLIC KEY** drop-down field, choose an appropriate authentication key.  
This field is optional. It is needed if you want to ssh to the worker nodes for troubleshooting purposes. Ensure that you use the Ed25519 or ECDSA format for the public key.  
**Note:** As RSA and DSA are less secure formats, Cisco prevents the use of these formats.
  - In the **IAM ACCESS ROLE ARN** field, enter the Amazon Resource Name (ARN) information.  
**Note** By default, the AWS credentials specified at the time of Amazon EKS cluster creation, that is the credentials configured in the Infrastructure Provider, are mapped to the `Kubernetes cluster-admin ClusterRole`. A default `ClusterRoleBinding` binds the credentials to the `system:masters` group, thereby granting super-user access to the holders of the IAM identity. The **IAM ACCESS ROLE ARN** field allows you to specify the ARN of an additional AWS IAM role or IAM user who is also granted administrative control of the cluster.
  - Click **NEXT**.
- Step 5** In the **VPC Configuration** screen, specify the following information:
- In the **SUBNET CIDR** field, enter a value of the overall subnet CIDR for your cluster.
  - In the **PUBLIC SUBNET CIDR** field, enter values for your cluster on separate lines.
  - In the **PRIVATE SUBNET CIDR** field, enter values for your cluster on separate lines.
- Step 6** In the **Summary** screen, review the cluster information and then click **FINISH**.  
Cluster creation can take up to 20 minutes. You can monitor the cluster creation status on the **Clusters** screen.
- Note** If you receive the Could not get token: AccessDenied error message, it indicates that the AWS account is not a trusted entity for the Role ARN.
- For information on adding your AWS account as a trusted entity, see [Creating AWS Roles, on page 30](#).

## Scaling AWS EKS Clusters

You can scale EKS clusters by adding or removing worker nodes to them based on the demands of the workloads you want to run.

- Step 1** From the right pane, click **EDIT**.  
The **Edit Cluster** dialog box appears.
- Step 2** From the **INSTANCE TYPE** drop-down list, choose an [instance type](#) for your cluster.
- Step 3** From the **MACHINE IMAGE** drop-down list, choose an appropriate Cisco Container Platform Amazon Machine Image (AMI) file.  
To add AMI files to your Amazon account, see [Adding AMI Files to your Amazon Account, on page 30](#).
- Step 4** In the **WORKER COUNT** field, change the number of work nodes as necessary.



**Step 5** Click **UPDATE**.

---

## Deleting AWS EKS Clusters

### Before you begin

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

---

**Step 1** From the left pane, click **Clusters**, and then click the **EKS Clusters** 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.

Upon deleting an AWS EKS cluster, it takes about 10 minutes for the cluster resources to be released.

---





## CHAPTER 6

# Administering Azure Kubernetes Service (AKS) Clusters

---

Integrating Cisco Container Platform with Azure Kubernetes Service (AKS) allows you to deploy and run containerized applications on the Azure cloud.

This chapter contains the following topics:

- [Prerequisites for Configuring AKS Clusters, on page 37](#)
- [Creating AKS Clusters, on page 38](#)
- [Deleting AKS Clusters, on page 39](#)

## Prerequisites for Configuring AKS Clusters

The prerequisites for configuring AKS clusters are as follows:

### Azure User Account Requirements

The following roles are necessary for your Azure user account that you want to use with Cisco Container Platform:

- An **App administrator** within the directory
- An **Owner** within the Azure subscription

You can contact your Azure administrator to set up these roles for your Azure account.

### Creating Resource Groups

A resource group is a logical grouping of the resources that are required to deploy Cisco Container Platform on Azure. They allow you to manage your resources efficiently.

- 
- Step 1** Log in to the [Microsoft Azure portal](#).
  - Step 2** From the left navigation pane, click **All services > Resource groups**. The **Resource groups** pane is displayed.
  - Step 3** Click **Add** to create a new resource group.

**Step 4** In the **Create an empty resource group** pane, specify the following information:

- a) In the **Resource Group name** field, enter a unique name.
- b) From the **Subscription** drop-down list, choose a subscription type.
- c) From the **Resource group location** drop-down list, choose the region in which you want to create your AKS cluster.

**Note** Not all regions support AKS. Ensure that you select a supported region. Currently, Cisco Container Platform supports only the **eastus** and **westus** regions.

For more information on a resource group, refer to [Resource Group Overview](#).

---

## Creating AKS Clusters

### Before you begin

- Ensure that you have added your Azure provider profile.  
For more information, see [Adding Azure Provider Profile, on page 16](#).
  - Ensure that your Azure account that is used with Cisco Container Platform has the necessary roles.  
For more information, see [Azure User Account Requirements, on page 37](#).
  - Ensure that you have created the necessary resource groups in your Azure account.  
For more information, see [Creating Resource Groups, on page 37](#).
- 

**Step 1** From the left pane, click **Clusters**, and then click the **Azure** tab.

**Step 2** Click **NEW CLUSTER**.

**Step 3** In the **Basic Information** screen, specify the following information:

- a) From the **INFRASTRUCTURE PROVIDER** drop-down list, choose the provider related to the appropriate Azure account.
- b) From the **AZURE REGION** drop-down list, choose an appropriate AKS region.

**Note** Not all regions support AKS. Ensure that you select a supported region. Currently, Cisco Container Platform supports only the **eastus** and **westus** regions.

- c) In the **KUBERNETES CLUSTER NAME** field, enter a name for your cluster.
- d) In the **KUBERNETES VERSION** field, choose the kubernetes version for your cluster.

**Note** Not all Kubernetes versions are supported in all Azure regions. Ensure that you select an available Kubernetes version.

- e) Click **NEXT**.

**Step 4** In the **Cluster settings** screen, specify the following information:

- a) In the **NODE POOL NAME** field, enter an alphanumeric name for the primary node pool of your cluster.
- b) In the **RESOURCE GROUP NAME** field, enter the name of the resource group that you have created in your Azure account.

For more information on creating a resource group, see [Creating Resource Groups, on page 37](#).

- c) Click **NEXT**.

**Step 5** In the **Node Configuration** screen, specify the following information:

- a) From the **WORKER INSTANCE TYPE** drop-down list, choose an [instance type](#) for your cluster.
- b) In the **WORKER COUNT** field, enter an appropriate number of worker nodes.
- c) In the **POD CIDR** field, enter the desired CIDR value.

**Note** This is an optional field. So you can leave the field empty or leave the default as is.

- d) In the **SERVICE CIDR** field, enter the desired CIDR value.

**Note** This is an optional field. So you can leave the field empty or leave the default as is.

- e) Click **NEXT**.

**Step 6** In the **Summary** screen, review the cluster information, and then click **FINISH**.

---

Cluster creation can take up to 20 minutes. You can monitor cluster creation status on the **Clusters** screen.

## Deleting AKS Clusters

### Before you begin

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

---

**Step 1** From the left pane, click **Clusters**, and then click the **Azure** 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.

---

Upon deleting an AKS cluster, it takes about 15 minutes for the cluster resources to be released.





## CHAPTER 7

# Managing Kubernetes Clusters

---

The Cisco Container Platform web interface allows you to manage Kubernetes clusters by using the **Kubernetes Dashboard**. Once you set up the **Kubernetes Dashboard**, you can deploy applications on the authorized Kubernetes clusters, and manage the application and the cluster itself.

This chapter contains the following topics:

- [Setting Up Kubernetes Dashboard to Access Clusters, on page 41](#)
- [Monitoring Health of Cluster Deployments, on page 43](#)
- [Monitoring Logs from Cluster Deployments, on page 44](#)

## Setting Up Kubernetes Dashboard to Access Clusters

The steps to set up the Kubernetes dashboard differ based on the method used for cluster creation.

This section contains the following topics:

### Setting Up Kubernetes Dashboard for vSphere On-prem Clusters

For a vSphere on-prem cluster, you can access the Kubernetes dashboard using the `Kubeconfig` file or the Kubernetes default token. The steps to use the default token are same as those provided in [Setting up Kubernetes Dashboard for On-prem AWS IAM Enabled Clusters].

Follow these steps to set up the Kubernetes dashboard using the `Kubeconfig` file.

- 
- Step 1** To download the `Kubeconfig` file that provides you access to the Kubernetes cluster, perform these steps on the Cisco Container Platform web interface:
- a) From the left pane, click **Clusters**.
  - b) From the drop-down list displayed under the **ACTIONS** column, choose **Download Token** to get the `Kubeconfig` file of the vSphere cluster.
- The `Kubeconfig` file is downloaded to your local system.
- Step 2** To set up the Kubernetes dashboard access, perform these steps in the Kubernetes dashboard:
- a) Click the **Kubeconfig** radio button.
  - b) Select the `Kubeconfig` file that you got in Step 1-b.
-

## Setting Up Kubernetes Dashboard for On-prem AWS IAM Enabled

For an on-prem AWS cluster that has IAM enabled, you can access the Kubernetes dashboard only by using the Kubernetes default token.

- 
- Step 1** To download the `Kubeconfig` file that provides you access to the Kubernetes cluster, perform these steps on the Cisco Container Platform web interface:
- From the left pane, click **Clusters**.
  - From the drop-down list displayed under the **ACTIONS** column, choose **Download Token** to get the `Kubeconfig` file of the on-prem AWS cluster.  
The `Kubeconfig` file is downloaded to your local system.
- Step 2** To get the Kubernetes default token, perform these steps in the `kubectl` utility:
- List the Kubernetes secrets in the `kube-system` namespace.  

```
kubectl get secrets -n kube-system
```
  - Search for the secret that has the following format:  

```
default-token-XXXXX
```
  - Get the default token.  

```
kubectl describe secret default-token-XXXXX -n kube-system
```
- Step 3** To set up the Kubernetes dashboard access, perform these steps in the Kubernetes dashboard:
- Click the **Token** radio button.
  - In the **Enter token** field, enter the Kubernetes default token that you got in Step 2-c.
- 

## Setting Up Kubernetes Dashboard for AWS EKS Clusters

- 
- Step 1** To download the `Kubeconfig` file that provides you access to the AWS EKS cluster, perform these steps on the Cisco Container Platform web interface:
- From the left pane, click **Clusters**.
  - From the drop-down list displayed under the **ACTIONS** column, choose **Download Token** to get the `Kubeconfig` file of the AWS EKS cluster.  
The `Kubeconfig` file is downloaded to your local system.
- Step 2** To set up the Kubernetes dashboard access, follow the steps provided on the [Dashboard Tutorial](#) page.
- 

## Setting Up Kubernetes Dashboard for AKS Clusters

- 
- Step 1** To download the `Kubeconfig` file that provides you access to the AKS cluster, perform these steps on the Cisco Container Platform web interface:
- From the left pane, click **Clusters** and then click the **Azure** tab.



- b) For the cluster whose dashboard you would like to access, from the drop-down list displayed under the **ACTIONS** column, choose **Download Kubeconfig** to get the `kubeconfig` file of the AKS cluster.  
The `kubeconfig` file is downloaded to your local system.

**Step 2** To set up Kubernetes dashboard access, follow these steps:

- a) Create the necessary cluster role binding.

```
kubectl --kubeconfig=[location_of_kubeconfig_downloaded_from_ccp_dashboard] create
clusterrolebinding kubernetes-dashboard --clusterrole=cluster-admin
--serviceaccount=kube-system:kubernetes-dashboard
```

- b) Connect to the Kubernetes dashboard using the Azure CLI.

```
az aks browse --resource-group [name_of_your_resource_group] --name [name_of_your_aks_cluster]
```

The Kubernetes dashboard is displayed in a new browser tab.

For more information, refer to the [Dashboard Tutorial](#) page.



**Note** In the [Dashboard Tutorial](#) page, you must follow the steps in the RBAC enabled cluster section because currently, in Cisco Container Platform, RBAC is enabled by default for all AKS clusters.

## Monitoring Health of Cluster Deployments

It is recommended to continuously monitor the health of your cluster deployment to improve the probability of early detection of failures and avoid any significant impact from a cluster failure.

Cisco Container Platform is deployed with Prometheus and Grafana, which are configured to start monitoring and logging services automatically when a Kubernetes cluster is created.

[Prometheus](#) is an open-source systems monitoring and alerting toolkit and [Grafana](#) is an open source metric analytics and visualization suite.

Prometheus collects the data from the cluster deployment, and Grafana provides a general purpose dashboard for displaying the collected data. Grafana offers a highly customizable and user-friendly dashboard for monitoring purposes.



**Note** A user with *Administrator* role can view all the cluster deployments, but a user with *User* role can view only those clusters for which the user has permission to view.

The following example shows a use case in which Grafana is available in the `ccp` namespace.

**Step 1** Access the Kubernetes cluster master node using `ssh`.

```
ssh -l <username> <IP address of master node>
```

Once you create a Kubernetes cluster, it may take a few minutes for the necessary services to start. If `ssh` to a cluster fails, we recommend that you try again after a few minutes.

**Step 2** Follow these steps to access Grafana.

**Note** On the Control Plane, Grafana is installed in the default namespace, therefore you must skip `-n ccp` in the commands.

```
export INGRESS_IP=$(kubectl get svc nginx-ingress-controller -n ccp
-o=jsonpath='{.spec.loadBalancerIP}')
export GRAFANA_USER=$(kubectl get secret ccp-monitor-grafana -n ccp -o=jsonpath='{.data.admin-user}'
| base64 --decode)
export GRAFANA_PASSWORD=$(kubectl get secret ccp-monitor-grafana -n ccp
-o=jsonpath='{.data.admin-password}' | base64 --decode)
echo "Login to Grafana at http://{INGRESS_IP}/grafana as user ${GRAFANA_USER} with password
${GRAFANA_PASSWORD}" && unset GRAFANA_PASSWORD GRAFANA_USER
```

**Note** It is important to either change or retain the original login credentials since the secret that was used to initialize the Grafana login may be lost or changed with future upgrades.

**Step 3** Add Prometheus as the data source and configure the Grafana dashboard to monitor the health of your cluster deployments.

## Example of Monitoring Multiple Prometheus Instances

To monitor multiple Prometheus instances you must expose Prometheus as an [Ingress resource](#) so that you can access it from a Grafana instance that is running in a different cluster.



**Note** The following example is valid only if Harbor is not installed.

```
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    kubernetes.io/ingress.class: nginx
    nginx.ingress.kubernetes.io/add-base-url: "true"
    nginx.ingress.kubernetes.io/rewrite-target: /
  name: ccp-monitor-prometheus-server
  namespace: ccp
spec:
  rules:
  - http:
      paths:
      - backend:
          serviceName: ccp-monitor-prometheus-server
          servicePort: 9090
        path: /
```

After Prometheus is accessible externally from the cluster, you can add it as a new datasource in Grafana.

## Monitoring Logs from Cluster Deployments

The Elasticsearch, Fluentd, and Kibana (EFK) stack enables you to collect and monitor log data from containerized applications for troubleshooting or compliance purposes. These components are automatically installed when you install Cisco Container Platform.

Fluentd is an open source data collector. It works at the backend to collect and forward the log data to Elasticsearch.

Kibana is an open source analytics and visualization platform designed to work with Elasticsearch. It allows you to create rich visualizations and dashboards with the aggregated data.



**Note** A user with the *Administrator* role can view all logs, but a user with *User* role can view logs for only those clusters for which the user has permission to view.

This section contains the following topics:

## Viewing EFK Logs Using Kibana (Tenant Cluster)

### Before you begin

Ensure that you have installed the `kubectl` utility.

- 
- Step 1** Download the Kubeconfig file of the cluster whose logs you want to view, see [Downloading Kubeconfig File, on page 65](#).
- Step 2** Copy the contents of the downloaded Kubeconfig file to:
- Your local host `~/ .kube/config`
  - A local file and export `KUBECONFIG=<Downloaded Kubeconfig file>`
- Step 3** Create a port-forward using `kubectl` to access Kibana from outside a cluster.
- a) Determine the pod.
- ```
kubectl -n ccp get pods
```
- Example
- ```
ccp-efk-kibana-6d7c97575c-9qxbf
```
- b) Open a port-forward.
- Example
- ```
kubectl port-forward -n ccp ccp-efk-kibana-6d7c97575c-9qxbf 5601:5601
```
- Step 4** Access the Kibana UI and view the data from the target tenant cluster using a web browser.  
`http://localhost:5601/app/kibana`  
For more information on customizing the Kibana UI, refer to the [latest Kibana documentation](#).
- 

## Viewing EFK Logs Using Kibana (Control Plane Cluster)

### Before you begin

Ensure that you have installed the `kubectl` utility.

- 
- Step 1** Access the Kubernetes cluster master node using `ssh`.

```
ssh ccpuser@control plane master node
sudo cat /etc/kubernetes/admin.conf
```

**Step 2** Copy the contents of the downloaded Kubeconfig file to:

- Your local host `~/ .kube/config`
- A local file and export `KUBECONFIG=<Full path of the Kubeconfig local file>`

For more information on setting Kubeconfig, see [Configure Access to Multiple Clusters](#).

**Step 3** Create a port-forward using kubectl to access Kibana from outside a cluster.

a) Determine the pod.

```
kubectl get pods
```

Example

```
ccp-efk-kibana-6d7c97575c-9qxbf
```

b) Open a port-forward.

Example

```
kubectl port-forward ccp-efk-kibana-6d7c97575c-9qxbf 5601:5601
```

**Step 4** Access the Kibana UI and view the data from the target tenant cluster using a web browser.

```
http://localhost:5601/app/kibana
```

For more information on customizing the Kibana UI, refer to the [latest Kibana documentation](#).

## Forwarding Logs to External Elasticsearch Server

Use the following Curl commands to configure forwarding of logs to an external Elasticsearch server:

**Step 1** Open a terminal that has a curl client installed.

**Step 2** Configure Cisco Container Platform login credentials.

```
export MGMT_HOST=https://<Cisco Container Platform IP address>:<Port>
export CCP_USER=<Username>
export CCP_PASSPHRASE=<Passphrase>
```

**Step 3** Login to Cisco Container Platform and save the session cookie for future requests into the cookies.txt local file.

```
curl -k -j -c cookies.txt -X POST -H "Content-Type:application/x-www-form-urlencoded" -d
"username=$CCP_USER&password=$CCP_PASSWORD" $MGMT_HOST/2/system/login/
```

**Step 4** Get the list of cluster names.

```
curl -s -k -b cookies.txt -H "Content-Type: application/json"
$MGMT_HOST/2/clusters/ | jq -r '.[].name'
```

**Step 5** Set the CLUSTER\_NAME environment variable to the cluster that you are working on.

```
export CLUSTER_NAME="<A cluster name from Step 2>"
```

**Step 6** Configure the cluster UUID.

```
export CLUSTER_UUID=$(curl -s -k -b cookies.txt -H "Content-Type: application/json"
$MGMT_HOST/2/clusters/$CLUSTER_NAME | jq -r '.uuid')
```

**Step 7** Configure the Elasticsearch server IP address and port number.

```
export EFK_SERVER=<IP address of Elasticsearch server>
export EFK_PORT=<Port number of Elasticsearch server>
```

**Step 8** Install the helm chart to configure the custom Elasticsearch server.

```
curl -s -k -b cookies.txt -X POST --header 'Content-Type: application/json' --header 'Accept:
application/json' -d '{"chart_url": "/opt/ccp/charts/ccp-agent.tgz", "name": "ccpagent", "options":
"cp-efk.localLogForwarding.enabled=false,cp-efk.localLogForwarding.elasticsearchHost='$EFK_SERVER',cp-efk.localLogForwarding.elasticsearchPort='$EFK_PORT'"}'
$MGMT_HOST/2/clusters/$CLUSTER_UUID/helmcharts
```

---





## CHAPTER 8

# Services and Networking

This chapter contains the following topics:

- [Load Balancing Kubernetes Services using NGINX, on page 49](#)
- [L7 Ingress, on page 49](#)
- [L4 Ingress, on page 51](#)
- [Ingress CA , on page 52](#)
- [Network Policies, on page 55](#)
- [Load Balancer Services, on page 55](#)

## Load Balancing Kubernetes Services using NGINX

Cisco Container Platform uses NGINX to offer advanced layer 7 load balancing solutions. NGINX can handle a large number of requests and at the same time, it can be run on Kubernetes containers.

The NGINX load balancer is automatically provisioned as part of Kubernetes cluster creation. Each Kubernetes cluster is provisioned with a single L7 NGINX load balancer. You can access the load balancer using its virtual IP address, which can be found by running the command `kubectl get svc -n ccp`.

To use the NGINX load balancer, you must create an Ingress resource. Ingress is a Kubernetes object that allows you to define HTTP load balancing rules to allow inbound connections to reach the cluster services. You can configure Ingress to create external URLs for services, load balance traffic, terminate SSL, offer name-based virtual hosting, and so on.

## L7 Ingress

Cisco Container Platform supports the following types of L7 Ingresses:

- **Simple fanout**

It enables you to access the website using http.

**Example**

```
cafe.test.com -> 10.1.1.1 -> /tea tea-svc:80 /coffee coffee-svc:80
```

For this type of Ingress, you need to create a yaml file that defines the Ingress rules.

**Sample yaml file**

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: cafe-ingress
spec:
  rules:
  - host: cafe.test.com
    http:
      paths:
      - path: /
        backend:
          serviceName: tea-svc
          servicePort: 80
      - path: /
        backend:
          serviceName: tea-svc
          servicePort: 80

```

- **Simple fanout with SSL termination**

It enables you to access the website using https.

**Example**

```

https://cafe.test.com  ->  10.1.1.1  ->  /tea      tea-svc:80
                        /coffee   coffee-svc:80

```

For this type of Ingress, you need to create the following yaml files:

- A yaml file that defines the Secret

**Sample yaml file**

```

apiVersion: v1
kind: Secret
metadata:
  name: cafe-secret
type: Opaque
data:
  tls.crt: base64 encoded cert
  tls.key: base64 encoded key

```

- A yaml file that defines the Ingress rules

**Sample yaml file**

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  name: cafe-ingress
spec:
  tls:
  - hosts:
    - cafe.test.com
      secretName: cafe-secret
  rules:
  - host: cafe.example.com
    http:
      paths:
      - path: /
        backend:
          serviceName: tea-svc
          servicePort: 80
      - path: /
        backend:

```



```

serviceName: coffee-svc
servicePort: 80

```

- **Name based virtual hosting**

It enables you to access the website using multiple host names.

**Example**

```

tea.test.com    --|          |-> tea.test.com    s1:80
                |10.1.1.1 |
coffee.test.com --|          |-> coffee.test.com s2:80

```

For this type of Ingress, you need to create a yaml file that defines the Ingress rules.

**Sample yaml file**

```

apiVersion: extensions/v1beta1
kind: Ingress
metadata
name: cafe-ingress
spec:
  rules:
    -host: tea.test.com
      http:
        paths:
          -path:/
            backend:
              serviceName: tea-svc
              servicePort: 80
    -host: coffee.test.com
      http:
        paths:
          -path:/
            backend:
              serviceName: coffee-svc
              servicePort: 80

```



**Note** You can download the yaml files that are shown in this topic from the following link:

<https://github.com/nginxinc/kubernetes-ingress/tree/master/examples/complete-example>

For more information on a sample scenario of implementing Ingress, see [Deploying Cafe Application with Ingress, on page 68](#).

## L4 Ingress

NGINX supports L4 TCP and UDP Ingress load balancing. It uses the NGINX helm chart that contains the TCP or UDP service mappings, instead of the Ingress resources as in the case of L7 support.

## Configuring L4 Load Balancing




---

**Note** NGINX supports either TCP or UDP L4 load balancing, but not both simultaneously.

---

**Step 1** Access the Kubernetes cluster master node using ssh.

```
ssh -l <username> <IP address of master node>
```

**Note** Once you create a Kubernetes cluster, it may take a few minutes for the necessary services to start. If ssh to a cluster fails, we recommend that you try again after a few minutes.

**Step 2** Get the current helm configuration values.

```
helm get values --all nginx-ingress > l4.yaml
```

**Step 3** Edit the l4.yaml file.

You can search for *tcp* or *udp* in the l4.yaml file, and then add your L4 services.

The following example shows adding the tcp-test-svc TCP service that uses port 3333.

```
tcp:
  "9000": default/tcp-test-svc:3333
```

The following example shows adding the udp-test-svc UDP service that uses port 5005.

```
udp:
  "9001": default/udp-test-svc:5005
```

**Step 4** Update the NGINX helm chart with the L4 service mappings.

```
helm upgrade --install nginx-ingress /opt/ccp/charts/nginx-ingress.tgz -f l4.yaml
```

**Note** You need to restart the NGINX Ingress controller pods for the new configuration to take effect.

**Step 5** Verify that ingress has successfully mapped the port.

```
kubectl get services -o wide -w nginx-ingress-controller
```

---

## Ingress CA


Cisco Container Platform by default creates an L7 Ingress service in order to support [Monitoring Health of Cluster Deployments](#), [Monitoring Logs from Cluster Deployments](#), and [Setting Up Kubernetes Dashboard for vSphere On-prem Clusters](#). All of these services are exposed with TLS enabled, and the certificate authority (CA) that is used to sign the Ingress controller server certificate is self-signed and per cluster based.

In order to reach the services without triggering SSL warning, you can either add the CA as part of your application that needs to interact with services behind Cisco Container Platform ingress (preferred), or add the CA to your system trusted CA list. The following section describes how to obtain the CA certificate.

**Step 1** Log in to the Kubernetes dashboard from browser as described in [Setting Up Kubernetes Dashboard for vSphere On-prem Clusters](#) section, download the kubeconfig file, and then use it to login to the Kubernetes dashboard.

**Step 2** From the right pane, click the dropdown box under **Namespace**, click the **ccp** namespace.

*Figure 4: Kubernetes Dashboard*



**Step 3** Click the **Secrets** tab.  
The **Secrets** pane appears.



```
curl --cacert ./ca.crt -I https://10.10.99.185/dashboard
HTTP/1.1 200 OK
Server: nginx/1.13.12
Date: Mon, 30 Jul 2018 19:08:11 GMT
Content-Type: text/html; charset=utf-8
Connection: keep-alive
Vary: Accept-Encoding
Accept-Ranges: bytes
Cache-Control: no-store
Strict-Transport-Security: max-age=15724800; includeSubDomains
```

---

## Network Policies

Cisco Container Platform supports [Kubernetes NetworkPolicies](#). The NetworkPolicies are independent of the underlying container network plugin.

## Load Balancer Services

Cisco Container Platform supports load balancer services on tenant clusters.

While creating a tenant cluster, you need to choose the number of load balancer IP addresses that you want to allocate for a tenant cluster from a VIP pool that you want to use.



**Note** The cluster creation operation fails if the number of requested load balancer IP addresses is more than the available IP addresses in the pool.

For more information, see [Creating Kubernetes Clusters on vSphere On-prem Clusters, on page 23](#).

Once load balancer IP addresses are allocated for a tenant cluster, externally reachable load balancer IP addresses are automatically provisioned for the load balancer services.

The following code provides an example of creating a service of type **LoadBalancer**.

```
apiVersion: v1
kind: Service
metadata:
  name: frontend
  labels:
    app: guestbook
    tier: frontend
type: LoadBalancer
```

You can update the number of available load balancer IP addresses from the **Edit Cluster** screen. You need to be aware of the number of used addresses in order to update the number of allocated load balancer IP addresses.

For example:

Suppose the current tenant is allocated with five load balancer IP addresses. If there are three load balanced services running, you cannot reduce the number of load balancer IP addresses to three or less as there are services using those IP addresses already.



---

**Note** When you delete a tenant cluster, the allocated load balancer IP addresses are recycled to the VIP pool.

---



## CHAPTER 9

# Istio Service Mesh

---

This chapter contains the following topics:

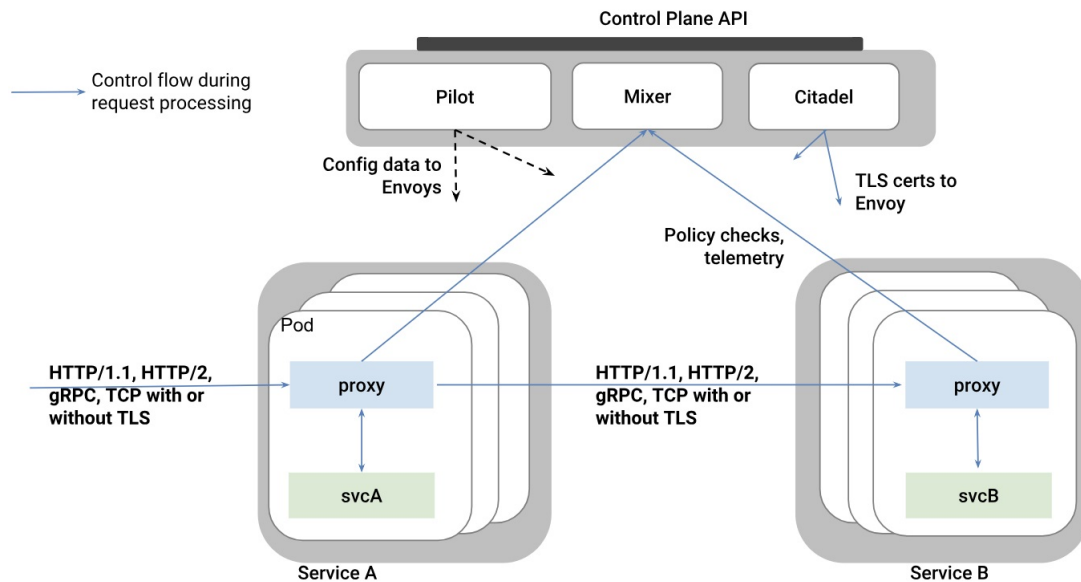
- [Introduction to Istio Service meshes, on page 57](#)
- [Configuring Istio Service meshes, on page 58](#)
- [Monitoring Service meshes, on page 59](#)

## Introduction to Istio Service meshes

Cisco Container Platform includes support for Istio service meshes. An Istio service mesh is logically split into a Data Plane and a Control Plane. The Data Plane includes a set of intelligent proxies (Envoy) and the Control Plane provides a reliable Istio framework. The term Istio is sometimes also used as a synonym to refer to the entire service mesh stack that includes the Control Plane and the Data Plane components.

The service mesh technology allows you to construct North-South and East-West L4 and L7 application traffic meshes. It provides containerized applications a language-independent framework that removes several common tasks related to L4 and L7 application networking from the actual application code. The common tasks include L4 and L7 service routing and load balancing, support for polyglot environments in a language-independent manner and advanced telemetry. The service mesh technology enhances operational capabilities such as monitoring, security, load balancing and troubleshooting for the applications. You can deploy a service mesh in a multi-cloud topology allowing these functions to operate with applications that run across multiple independent cloud deployments.

The following figure shows the high-level architecture of an Istio service mesh.



### Istio Architecture

In Cisco Container Platform, the components of Istio and Envoy are supported in the upstream Istio community. The Control and Data Plane components of the solution, such as Pilot, Mixer, Citadel and the Data Plane Envoy proxy for both North-South and East-West load balancing, are supported on Cisco Container Platform.

For more information on these technologies, refer to the upstream community documentation pages for [Istio](#) and [Envoy](#).



**Note** Currently, the Istio service mesh feature is marked as a Tech Preview feature and uses the Istio community version v1.0. You need to contact your service representative for support on the version of Cisco Container Platform you have deployed.

## Configuring Istio Service meshes

An Istio service mesh is a configurable feature on Cisco Container Platform. You can configure a separate instance of the service mesh stack on each tenant cluster. Support for Istio must be configured at the time of creating a tenant Kubernetes cluster. You can perform this configuration using APIs or the Cisco Container Platform web interface.

Each instance of the Istio service mesh uses an IP address from the Virtual IP address pool that is associated with the tenant cluster. Consequently, you need to ensure that there is sufficient number of IP addresses free and available in the VIP pool before enabling Istio. Typically, at least three IP addresses are required, one each for the Kubernetes API, Kubernetes Ingress, and Istio Ingress gateway. This number may change in future when additional features require more virtual IP addresses.



For more information on the required number of virtual IP addresses for a given software version of Cisco Container Platform, refer to the Virtual IP address section.

The following figure shows the **Node Configuration** screen, using which you can enable the Istio service mesh on a tenant cluster of the Cisco Container Platform.

In the current version of Cisco Container Platform, you can use a boolean flag to enable an Istio service mesh in a tenant Kubernetes cluster of Cisco Container Platform. If you enable the flag, a predetermined configuration of an Istio-based service mesh with Envoy as the Data Plane is configured in the tenant Kubernetes cluster. An internal instance of a service load balancer is automatically configured and a virtual IP address is automatically allocated for the Ingress gateway function of Istio.

## Monitoring Service meshes

On Cisco Container Platform, the Istio Control Plane is deployed in a special **istio-system** namespace of a tenant Kubernetes cluster. This is similar to how other add-on services such as Prometheus based monitoring or NGINX based Kubernetes ingress are provided. In a production deployment, a tenant Kubernetes cluster administrator grants read-write access to your development namespaces but not to the namespaces of system add-on services such as Istio, thereby protecting the Control Plane of such services from getting over-written accidentally or maliciously by your application containers.

The following is a checklist of monitoring and troubleshooting steps when using Istio on Cisco Container Platform:

1. If Istio fails to be enabled on your tenant Kubernetes cluster, in addition to the usual troubleshooting steps for Cisco Container Platform, also ensure that there is a sufficient number of virtual IP addresses available in the pool configured for this Kubernetes tenant cluster. In the current version of Cisco Container Platform,

at least three IP addresses need to be free and available for a tenant Kubernetes cluster that has Istio enabled.

2. Confirm that all pods are running in the `istio-system` namespace of the tenant Kubernetes cluster. The following figure shows a sample CLI output indicating that all Istio control pods are running correctly in a tenant Kubernetes cluster. If one or more pods continuously fails to run, use `kubectl describe pod <name_of_pod>` to troubleshoot the issue.

```
ccpuser@vhosakot-istio14-master5ebb31962c:~$ kubectl get pods -n istio-system
NAME                                READY   STATUS    RESTARTS   AGE
grafana-5b977b576f-2r5gs            1/1     Running   0           20h
istio-citadel-5ff4f56f56-1k6wz      1/1     Running   0           20h
istio-egressgateway-6567bc7ffb-84tj8 1/1     Running   0           20h
istio-ingressgateway-5dfb78f45b-c6jxc 1/1     Running   0           20h
istio-mixer-post-install-w56cx       0/1     Completed 0           20h
istio-pilot-6ddc9b5b49-hl5nd        2/2     Running   0           20h
istio-policy-f67cb98b5-n2q2m        2/2     Running   0           20h
istio-sidecar-injector-5545db64bf-tttc9 1/1     Running   0           20h
istio-statsd-prom-bridge-949999c4c-82spd 1/1     Running   0           20h
istio-telemetry-667d4c6765-2s9hj    2/2     Running   0           20h
istio-tracing-754cdfd695-2wd45      1/1     Running   0           20h
prometheus-86cb6dd77c-4cj77         1/1     Running   0           20h
servicegraph-ccd4d4859-sgcwc        1/1     Running   0           20h
```

3. Confirm that all Istio services are running in the `istio-system` namespace of the tenant Kubernetes cluster.

The following figure shows a CLI output with the Istio services up and running.

```
ccpuser@vhosakot-istio14-master5ebb31962c:~$ kubectl get svc -n istio-system
NAME                                TYPE           CLUSTER-IP      EXTERNAL-IP      PORT(S)
grafana                             ClusterIP      10.98.223.200   <none>           3000/TCP
istio-citadel                       ClusterIP      10.97.93.126    <none>           8060/TCP,9093/TCP
istio-egressgateway                 ClusterIP      10.108.19.80    <none>           80/TCP,443/TCP
istio-ingressgateway                 LoadBalancer  10.111.228.87   10.10.99.148     80:31380/TCP,443:31390/TCP,31400:31400/TCP
istio-pilot                         ClusterIP      10.104.249.174  <none>           15003/TCP,15005/TCP,15007/TCP,15010/TCP,15011/TCP,8080/TCP,9091/TCP
istio-policy                         ClusterIP      10.108.75.85    <none>           9091/TCP,15004/TCP,9093/TCP
istio-sidecar-injector               ClusterIP      10.109.55.202   <none>           443/TCP
istio-statsd-prom-bridge             ClusterIP      10.107.183.156  <none>           9102/TCP,9125/UDP
istio-telemetry                     ClusterIP      10.110.209.16   <none>           9091/TCP,15004/TCP,9093/TCP,42422/TCP
prometheus                          ClusterIP      10.101.6.183    <none>           9090/TCP
servicegraph                        ClusterIP      10.105.53.151   <none>           8088/TCP
tracing                             LoadBalancer  10.101.62.116   <pending>        80:31960/TCP
zipkin                               ClusterIP      10.99.116.160   <none>           9411/TCP
ccpuser@vhosakot-istio14-master5ebb31962c:~$
```

4. Confirm that the Ingress gateway service has an external IP address allocated and that this IP address is one of the previously available IP addresses in the virtual IP address pool associated with this tenant Kubernetes cluster. An example of this CLI output is shown in the preceding figure.
5. Deploy the [bookinfo example application](#) provided in the Istio upstream community web site.
6. The `istioctl` CLI utility is not deployed in the current version of the Cisco Container Platform. Most of the Istio functionality is now available through the `kubectl` CLI, but if you want to use `istioctl`, run these steps to deploy `istioctl` on a tenant Kubernetes cluster of the Cisco Container Platform:

```
export ISTIO_VERSION=1.0
curl -L https://git.io/getLatestIstio | sh -
chmod +x istio-${ISTIO_VERSION}/bin/istioctl
sudo mv istio-${ISTIO_VERSION}/bin/istioctl /usr/local/bin/
istioctl version
```

For more information and operational guidelines, refer to the [Istio upstream documentation](#).





## CHAPTER 10

# Harbor Registry

---

Using a Harbor registry, you can host container images in a local, private Docker registry. Harbor is an extension of the basic Docker registry that implements access controls, identity management, and a graphical interface. Using imagePullSecrets, Kubernetes resources can connect to a Harbor Registry to retrieve container images on other systems.

This chapter contains the following topic:

- [Using Harbor Registry in Tenant Clusters, on page 63](#)

## Using Harbor Registry in Tenant Clusters

Follow these steps to create a new tenant cluster with access to the Harbor registry:

---

**Step 1** Obtain the Ingress Root CA Certificate from the Kubernetes UI in one of the following ways:

- Use the steps in [Ingress CA , on page 52](#).
- Run the following command on the tenant cluster where Harbor registry is installed.

```
kubectl get secrets -n ccp ccp-ingress-tls-ca -o jsonpath='{.data.tls\.crt}' | base64 --decode
```

You can view the Harbor endpoint at `https://<LOAD_BALANCER_IP>:443` of the cluster where it is installed.

**Step 2** Create a new tenant cluster.

For more information, see [Creating Kubernetes Clusters on vSphere On-prem Clusters, on page 23](#).

**Step 3** In the **Node Configuration** screen, copy and paste the Root CA certificate obtained in Step 1.

Adding CA certificates to the Root CA is the only supported method of enabling secure registries in Cisco Container Platform tenant clusters.

**Note** Do not enable Harbor in the **Harbor Registry** screen.

**Step 4** After tenant cluster creation, SSH to one of the VMs in the cluster and login to the Harbor registry with the password you provided during the installation of Harbor.

```
docker login -u admin -p *****  
https://<LOAD_BALANCER_IP>:443
```





# CHAPTER 11

## Deploying Applications on Kubernetes Clusters

Once you have created Kubernetes cluster using the Cisco Container Platform web interface, you can deploy containerized applications on top of it.

This chapter contains the following topics:

- [Workflow of Deploying Applications, on page 65](#)
- [Downloading Kubeconfig File, on page 65](#)
- [Sample Scenarios, on page 66](#)

### Workflow of Deploying Applications

| Task                                                                                                                 | Related Section                                                                      |
|----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Create Kubernetes clusters using the Cisco Container Platform web interface.                                         | <a href="#">Creating Kubernetes Clusters on vSphere On-prem Clusters, on page 23</a> |
| Download the kubeconfig file that contains the cluster information and the certificates required to access clusters. | <a href="#">Downloading Kubeconfig File, on page 65</a>                              |
| Use the kubectl utility to deploy the application and test the scenario.                                             | <a href="#">Sample Scenarios, on page 66</a>                                         |

### Downloading Kubeconfig File

You must download the cluster environment to access the Kubernetes clusters using command line tools such as `kubectl` or using APIs.

**Step 1** From the left pane, click **Clusters**.

**Step 2** Click the **Download** icon corresponding to the cluster environment that you want to download.

The `kubeconfig` file that contains the cluster information and the certificates required to access clusters is downloaded to your local system.

# Sample Scenarios

This topic contains a few sample scenarios of deploying applications.

## Deploying a Pod with Persistent Volume

Tenant clusters are deployed with a default storage class named **standard**, and a default storage class provider named **vSphere provider**.

If you select a HyperFlex local network during cluster creation, HyperFlex storage class and storage class provisioner are created by default. In Cisco container Platform 4.0+, when deployed with Hyperflex 4.0+, the following two HyperFlex provisioners are supported:

- **hyperflex**, the HyperFlex FlexVolume provisioner available with HyperFlex 3.5+
- **hyperflex-csi**, the HyperFlex Container Storage Interface(CSI) provisioner available with HyperFlex 4.0+

### Step 1 Configure a tenant Kubernetes cluster.

```
export KUBECONFIG=<Path to kubeconfig file>
```

### Step 2 Verify if the storage cluster is created.

```
kubectl describe storageclass standard
```

```
Name:                standard
IsDefaultClass:      Yes
Provisioner:         kubernetes.io/vsphere-volume
Parameters:          diskformat=thin
ReclaimPolicy:       Delete
VolumeBindingMode:   Immediate
```

On HyperFlex 4.0+, if you have selected a HyperFlex local network, additional storage classes are displayed when you run the following command:

```
kubectl get sc
```

| NAME               | PROVISIONER                  | AGE |
|--------------------|------------------------------|-----|
| hyperflex          | hyperflex.io/hxvolume        | 22h |
| hyperflex-csi      | csi-hxcsi                    | 22h |
| standard (default) | kubernetes.io/vsphere-volume | 22h |

### Step 3 Create the persistent volume claim to request for storage.

```
cat <<EOF > pvc.yaml
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: pv-claim
spec:
  storageClassName: standard
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 3Gi
EOF
```



**Note** The `storageClassName` field is optional. For HyperFlex 4.0+, you must use **hyperflex-csi** as the storage class.

```
kubectl create -f pvc.yaml
persistentvolumeclaim "pv-claim" created
```

**Note** The HyperFlex storage class supports the `ReadWriteOnce` or `ReadOnlyMany` access modes and the vSphere storage class supports the `ReadWriteOnce` access mode.

#### Step 4 Verify if the persistent volume claim (pvc) is created.

```
kubectl describe pvc pv-claim
Name:          pv-claim
Namespace:    default
StorageClass: standard
Status:       Bound
Volume:       hx-default-pv-claim-5c4e8978-cdd2-11e8-9a07-005056b8fd7b
Labels:
Annotations:  pv.kubernetes.io/bind-completed=yes
              pv.kubernetes.io/bound-by-controller=yes
Finalizers:   [kubernetes.io/pvc-protection]
Capacity:    3Gi
Access Modes: RWO,ROX
Events:       \
```

Persistent Volume is automatically created and is bounded to this pvc.

**Note** When **VSPHERE** is used as the default storage class, a VMDK file is created inside the **kubevols** folder in the datastore which is specified during the creation of the tenant Kubernetes cluster.

#### Step 5 Create a pod that uses persistent volume claim with storage class.

```
cat <<EOF > pvc-pod.yaml
kind: Pod
apiVersion: v1
metadata:
  name: pvc-pod
spec:
  volumes:
    - name: pvc-storage
      persistentVolumeClaim:
        claimName: pv-claim
  containers:
    - name: pvc-container
      image: nginx
      ports:
        - containerPort: 80
          name: "http-server"
      volumeMounts:
        - mountPath: "/usr/share/nginx/html"
          name: pvc-storage
EOF
```

```
kubectl create -f pvc-pod.yaml
pod "pvc-pod" created
```

#### Step 6 Verify if the pod is up and running.

```
kubectl get pod pvc-pod

NAME        READY   STATUS    RESTARTS   AGE
pvc-pod    1/1     Running   0           16s
```

When **VSPHERE** is used as the default storage class, you can access vCenter and view the dynamically provisioned VMDKs of the pod.

## Deploying Cafe Application with Ingress

This scenario describes deploying and configuring the *Cafe application* with Ingress rules to manage incoming HTTP requests. It uses a **Simple fanout with SSL termination Ingress**.

For more information on Ingress, see [Load Balancing Kubernetes Services using NGINX, on page 49](#).

**Step 1** Go to the following URL:

<https://github.com/nginxinc/kubernetes-ingress/tree/master/examples/complete-example>

**Step 2** Download the following yaml files:

- tea-rc.yaml
- tea-svc.yaml
- coffee-rc.yaml
- coffee-svc.yaml
- cafe-secret.yaml
- cafe-ingress.yaml

**Step 3** Open the **kubectl** utility.

**Step 4** Obtain the IP address of the L7 NGINX load balancer that Cisco Container Platform automatically installs:

```
kubectl get pods --all-namespaces -l app=nginx-ingress -o wide
```

| NAMESPACE    | NAME                                            | READY | STATUS  | RESTARTS | AGE | IP           | NODE                    |
|--------------|-------------------------------------------------|-------|---------|----------|-----|--------------|-------------------------|
| ingressnginx | nginx-<br>ingresscontroller<br>-66974b775-jnmp1 | 1/1   | Running | 0        | 3d  | 10.10.45.235 | test-clusterwc5729f9ce2 |

**Step 5** Deploy the Cafe application.

a) Create the coffee and the tea services and replication controllers:

```
kubectl create -f tea-rc.yaml<br>
kubectl create -f tea-svc.yaml<br>
kubectl create -f coffee-rc.yaml<br>
kubectl create -f coffee-svc.yaml
```

**Step 6** Configure load balancing.

a) Create a Secret with an SSL certificate and a key:

```
kubectl create -f cafe-secret.yaml
```

b) Create an Ingress Resource:

```
kubectl create -f cafe-ingress.yaml
```

**Step 7** Verify that the Cafe application is deployed.

```
kubectl get pods -o wide
```

| NAMESPACE       | READY | STATUS  | RESTARTS | AGE | IP              | NODE                     |
|-----------------|-------|---------|----------|-----|-----------------|--------------------------|
| coffee-rc-jb9sx | 1/1   | Running | 0        | 3d  | 192.168.151.134 | test-cluster-wb3d42afeff |
| coffee-rc-tjwgj | 1/1   | Running | 0        | 3d  | 192.168.44.133  | test-cluster-wc5729f9ce2 |
| tea-rc-6qvmv    | 1/1   | Running | 0        | 3d  | 192.168.44.132  | test-cluster-wc5729f9ce2 |
| tea-rc-ms46j    | 1/1   | Running | 0        | 3d  | 192.168.151.132 | test-cluster-wb3d42afeff |
| tea-rc-tnftv    | 1/1   | Running | 0        | 3d  | 192.168.151.133 | test-cluster-wb3d42afeff |

### Step 8 Verify if the coffee and tea services are deployed.

```
kubectl get svc
```

| NAME       | TYPE      | CLUSTER-IP    | EXTERNAL-IP | PORT(S) | AGE |
|------------|-----------|---------------|-------------|---------|-----|
| coffee-svc | ClusterIP | 10.105.139.1  | 80/TCP      | 3d      |     |
| kubernetes | ClusterIP | 10.96.0.1     | 443/TCP     | 4d      |     |
| tea-svc    | ClusterIP | 10.109.34.129 | 80/TCP      | 3d      |     |

### Step 9 Verify if the Ingress is deployed.

```
kubectl describe ing
```

```
Name: cafe-ingress
Namespace: default
Address:
Default backend: default-http-backend:80 (<none>)
TLS: cafe-secret terminates cafe.example.com
Rules:
```

| Host             | Path    | Backends               |
|------------------|---------|------------------------|
| cafe.example.com | /tea    | tea-svc:80 (<none>)    |
|                  | /coffee | coffee-svc:80 (<none>) |

```
Annotations:
Events: <none>
```

### Step 10 Test the application.

- Access the load balancer IP address 10.10.45.235, which is obtained in Step2.
- Test if the Ingress controller is load balancing as expected.

```
curl --resolve cafe.example.com:443:10.10.45.235 https://cafe.example.com/coffee --insecure
<!DOCTYPE html>
...
<p><span>Server &nbsp; address:</span> <span>192.168.151.134:80</span></p>
...
curl --resolve cafe.example.com:443:10.10.45.235 https://cafe.example.com/coffee --insecure
<!DOCTYPE html>
...
<p><span>Server &nbsp; address:</span> <span>192.168.44.133:80</span></p>
...

```





# APPENDIX **A**

## User Privileges on vSphere

This appendix contains the following topic:

- [User Privileges on vSphere, on page 71](#)
- [Erase User Data, on page 75](#)

## User Privileges on vSphere

The following table provides the minimal set of privileges that are required by the vSphere user to execute the relevant operations in vCenter.

| Roles              | Privileges                                                                                                                                                                                                                                                                                                | Entities                  | Propagate to Children |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------|-----------------------|
| manage-k8s-node-vm | Resource.AssignVMToPool<br>System.Anonymous<br>System.Read<br>System.View<br>VirtualMachine.Config.AddExistingDisk<br>VirtualMachine.Config.AddNewDisk<br>VirtualMachine.Config.AddRemoveDevice<br>VirtualMachine.Config.RemoveDisk<br>VirtualMachine.Inventory.Create<br>VirtualMachine.Inventory.Delete | Cluster, Hosts, VM folder | Yes                   |
| manage-k8s-volumes | Datastore.AllocateSpace<br>Datastore.FileManagement<br>System.Anonymous<br>System.Read<br>System.View                                                                                                                                                                                                     | Datastore                 | No                    |

| <b>Roles</b>                         | <b>Privileges</b>                                                     | <b>Entities</b>                                                        | <b>Propagate to Children</b> |
|--------------------------------------|-----------------------------------------------------------------------|------------------------------------------------------------------------|------------------------------|
| k8s-system-read-and-spbmprofile-view | StorageProfile.View<br>System.Anonymous<br>System.Read<br>System.View | vCenter                                                                | No                           |
| ReadOnly                             | System.Anonymous<br>System.Read<br>System.View                        | Datacenter,<br>Datastore<br>cluster,<br>Datastore<br>storage<br>folder | Yes                          |
| ccp-register-extension               | Extension.Register<br>Extension.Unregister<br>Extension.Update        | vCenter                                                                | No                           |

| Roles     | Privileges | Entities                                                             | Propagate to Children |
|-----------|------------|----------------------------------------------------------------------|-----------------------|
| CCP_Admin |            | Cluster,<br>Hosts,<br>Vcenter,<br>Datastore,<br>Datastore<br>cluster | Yes                   |

| Roles | Privileges                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Entities | Propagate to Children |
|-------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|-----------------------|
|       | Extension.Register<br>Extension.Unregister<br>Extension.Update<br>Resource.AssignVMToPool<br>Network.Assign<br>StorageProfile.View<br>System.Anonymous<br>System.Read<br>System.View<br>VirtualMachine.Config.AddExistingDisk<br>VirtualMachine.Config.AddNewDisk<br>VirtualMachine.Config.AddRemoveDevice<br>VirtualMachine.Config.RemoveDisk<br>VirtualMachine.Config.CPUCount<br>VirtualMachine.Config.AdvancedConfig<br>VirtualMachine.Config.Resource<br>VirtualMachine.Config.ManagedBy<br>VirtualMachine.Config.DiskExtend<br>VirtualMachine.Config.Memory<br>VirtualMachine.Config.Settings<br>VirtualMachine.Config.RawDevice<br>VirtualMachine.Inventory.Create<br>VirtualMachine.Inventory.Remove<br>VirtualMachine.Provisioning.Clone<br>VirtualMachine.Provisioning.CreateTemplateFromVM<br>VirtualMachine.Provisioning.DeployTemplate<br>VApp.Import<br>VApp.PowerOn<br>VApp.PowerOff<br>VApp.Suspend<br>VApp.ResourceConfig<br>VApp.InstanceConfig |          |                       |



| Roles | Privileges                                     | Entities | Propagate to Children |
|-------|------------------------------------------------|----------|-----------------------|
|       | VApp.ApplicationConfig<br>VApp.ManagedByConfig |          |                       |

For more information on adding a provider profile, see [Adding vSphere Provider Profile](#), on page 16.

## Erase User Data

You need to erase user data and return a cluster to a clean state when its physical media is replaced or removed. When working with **Virtual Volumes**, deleting or overwriting a file is not adequate for completely erasing user data. File systems do not overwrite the disk blocks that contain data. This means that deletion of a VM or datastore does not erase user data. In order to securely erase user data, you need to erase the physical storage underlying the datastore.

For more information on securely erasing user data from a cluster, refer to the latest documentation from your storage vendor.





## APPENDIX **B**

# Accessibility Features in Cisco Container Platform

---

The list of accessibility features in Cisco Container Platform is available on the [Voluntary Product Accessibility Template \(VPAT\)](#) page under the [Cloud](#) section. For further assistance, you can contact [accessibility@cisco.com](mailto:accessibility@cisco.com).

All product documents are accessible except for images, graphics, and some charts. If you would like to receive the product documentation in audio format, braille, or large print, contact [accessibility@cisco.com](mailto:accessibility@cisco.com).

