



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:

- [Accessing Kubernetes Clusters, on page 1](#)
- [Monitoring Health of Cluster Deployments, on page 5](#)
- [Monitoring Logs from Cluster Deployments, on page 6](#)
- [Renewing Kubernetes API Certificates, on page 11](#)

Accessing Kubernetes Clusters

The steps to access the Kubernetes clusters differ based on the method used for cluster creation.

This section contains the following topics:

Accessing Kubernetes Clusters on vSphere

You can access clusters on vSphere using the Kubernetes dashboard and 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:

- a) In 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 get the Kubernetes default token, perform these steps in the `kubectl` utility:

- a) SSH into the master of the tenant cluster.
- b) List the Kubernetes secrets in the `kube-system` namespace.

```
kubectl get secrets -n kube-system
```

- c) Search for the secret that has the following format:

```
default-token-XXXXX
```

d) Get the default token in one of the following ways:

- `kubectl describe secret default-token-XXXXX -n kube-system`
- `kubectl get secret default-token-XXXXX -n kube-system -o jsonpath='{.data.token}' | base64 -d`

Step 3 To set up the Kubernetes dashboard access, perform these steps:

- a) In the left pane of the Cisco Container Platform web interface, click **Clusters**.
- b) From the drop-down list displayed under the **ACTIONS** column, choose **Kubernetes Dashboard**. The Kubernetes Dashboard login screen is displayed.
- c) Log in to the Kubernetes Dashboard in one of the following ways:
 - Using the **TOKEN**: Click the **TOKEN** radio button and enter the token from Step 2-c.
 - Using the **kubeconfig** file: Click the **Kubeconfig** radio button and select the Kubeconfig file from Step 1-b.

Accessing Kubernetes Clusters for On-prem AWS IAM Enabled Clusters

For an on-prem AWS cluster that has IAM enabled, you can access clusters using the Kubernetes dashboard and 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:

- a) In 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 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:

- a) List the Kubernetes secrets in the `kube-system` namespace.

```
kubectl get secrets -n kube-system
```

- b) Search for the secret that has the following format:

```
default-token-XXXXX
```

- c) 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:

- a) Click the **Token** radio button.
- b) In the **Enter token** field, enter the Kubernetes default token from Step 2-c.

Accessing Kubernetes Clusters on AWS EKS

You can use the `Kubeconfig` file along with the `kubectl` command line tool to access the clusters on AWS EKS.

-
- 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:
- In 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.
-

Accessing Kubernetes Clusters on AKS

You can use the `Kubeconfig` file along with the `kubectl` command line tool to access the clusters on AKS.

-
- 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:
- In the left pane, click **Clusters** and then click the **Azure** tab.
 - 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 environment.
- Step 2** To set up Kubernetes dashboard access, follow these steps:
- 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
```
 - 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, see [Dashboard Tutorial](#).



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 clusters on AKS.

Accessing Kubernetes Clusters on GKE

Before you begin

- Install Google Cloud SDK on your computer. For more information, see [Installing Google Cloud SDK](#).



Note You can choose a suitable installation option depending on the OS installed on your computer.

- - After the SDK is installed, set the environment variable `G_CLOUD_SDK_PATH` to `GOOGLE_CLOUD_SDK_INSTALLATION_DIR/bin`.
- - Setup the GCP credentials.

For more information, see [Configuring cluster access for kubectl](#).

You can configure your computer to access clusters on GKE in one of the following ways:

Accessing Clusters on GKE using GKE Dashboard

We recommend that you use the GKE dashboard that is available on the [GCP console](#) to view, inspect, manage, and delete resources in your clusters.

For more information, see [GKE Dashboard](#).



Note The open-source Kubernetes dashboard is deprecated for clusters on GKE.

Accessing Clusters on GKE using Kubeconfig and Kubectl

You can use the `kubeconfig` file along with the `kubectl` command line tool to access the clusters on GKE.

Generate the `kubeconfig` file in one of the following ways:

- Use GCP to generate the `kubeconfig` file in your environment.

When you create a cluster using GCP, an entry is automatically added to the `kubeconfig` in the `$HOME/.kube/config` file.

For more information, see [Setting up kubeconfig on your machine](#).

- Alternatively, use Cisco Container Platform to generate the `kubeconfig` file:

- In the left pane, click **Clusters**, and then click the **GKE** tab.
- For the cluster whose dashboard you want to access, from the drop-down list displayed under the **ACTIONS** column, choose **Download Kubeconfig** to get the `kubeconfig` file of the cluster.

The `kubeconfig` file is downloaded to your environment.

- Edit the `kubeconfig` file to replace `$G_CLOUD_SDK_PATH` with the local path where `google-cloud-sdk` is installed.
 - A separate `kubeconfig` file is generated to allow access to the GKE 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](#).

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 Perform one of the following steps to access Kibana from outside a cluster.

- 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
```

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

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

- Run a bash script to create a Kibana login and access Kibana through Ingress from outside a cluster.

a. Create ingress resource yaml for kibana.

```
cat > kibana_ingress.yaml <<EOF
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    kubernetes.io/ingress.class: nginx
    nginx.ingress.kubernetes.io/add-base-url: "false"
    nginx.ingress.kubernetes.io/auth-realm: Kibana
    nginx.ingress.kubernetes.io/auth-secret: ccp-kibana-basic-auth
    nginx.ingress.kubernetes.io/auth-type: basic
    nginx.ingress.kubernetes.io/rewrite-target: /
  name: ccp-kibana
  namespace: ccp
spec:
  rules:
  - http:
      paths:
      - path: /kibana/?(.*)
        backend:
          serviceName: ccp-efk-kibana
          servicePort: 5601
EOF
```

b. Create a bash script.

```
cat > ccp_kibana_ingress_setup.sh << EOF

#!/bin/bash

# Kibana username
KIBANA_USERNAME=<USER>
KIBANA_PASSWORD=<PASSWORD>

# Configure Kibana server.basePath param
kubectl -n ccp get cm ccp-efk-kibana -o json |
jq '.["data"]["kibana.yml"] += "server.basePath: \"/kibana\\"'\n' |
```

```
kubectl replace -f -

# Generating password and creating secret
KIBANA_SECRET=`htpasswd -n -b $KIBANA_USERNAME $KIBANA_PASSWORD | head -n 1`
kubectl -n ccp create secret generic ccp-kibana-basic-auth --from-literal=auth=$KIBANA_SECRET

# Creating ingress
kubectl apply -f kibana_ingress.yaml

EOF
```

c. Run the script.

```
sudo apt-get install htpasswd
chmod 700 ccp_kibana_ingress_setup.sh
./ccp_kibana_ingress_setup.sh
```

d. Restart the Kibana pod.

```
kubectl -n ccp delete $(kubectl -n ccp get pods -o name | grep ccp-efk-kibana)
```

e. Access the Kibana UI and view the data using a web browser.

```
http://<INGRESS IP>/kibana
```

For more information on customizing the Kibana UI, see 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 Perform one of the following steps to access Kibana from outside a cluster:

- 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**

```
kubect1 port-forward ccp-efk-kibana-6d7c97575c-9qxbf 5601:5601
```

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

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

- Run a bash script to create a Kibana login and access Kibana through Ingress from outside a cluster.

a. Create ingress resource yaml for kibana.

```
cat > kibana_ingress.yaml <<EOF
apiVersion: extensions/v1beta1
kind: Ingress
metadata:
  annotations:
    kubernetes.io/ingress.class: nginx
    nginx.ingress.kubernetes.io/add-base-url: "false"
    nginx.ingress.kubernetes.io/auth-realm: Kibana
    nginx.ingress.kubernetes.io/auth-secret: ccp-kibana-basic-auth
    nginx.ingress.kubernetes.io/auth-type: basic
    nginx.ingress.kubernetes.io/rewrite-target: /
  name: ccp-kibana
  namespace: ccp
spec:
  rules:
  - http:
    paths:
    - path: /kibana/(?.* )
      backend:
        serviceName: ccp-efk-kibana
        servicePort: 5601
EOF
```

b. Create a bash script.

```
cat > ccp_kibana_ingress_setup.sh << EOF

#!/bin/bash

# Kibana username
KIBANA_USERNAME=<USER>
KIBANA_PASSWORD=<PASSWORD>

# Configure Kibana server.basePath param
kubect1 -n ccp get cm ccp-efk-kibana -o json |
jq '["data"]["kibana.yml"] += "server.basePath: \"/kibana\"\n" ' |
kubect1 replace -f -

# Generating password and creating secret
KIBANA_SECRET=`htpasswd -n -b $KIBANA_USERNAME $KIBANA_PASSWORD | head -n 1`
kubect1 -n ccp create secret generic ccp-kibana-basic-auth --from-literal=auth=$KIBANA_SECRET

# Creating ingress
kubect1 apply -f kibana_ingress.yaml

EOF
```

c. Run the script.

```
sudo apt-get install htpasswd
chmod 700 ccp_kibana_ingress_setup.sh
./ccp_kibana_ingress_setup.sh
```

d. Restart the Kibana pod.

```
kubectl -n ccp delete $(kubectl -n ccp get pods -o name | grep ccp-efk-kibana)
```

e. Access the Kibana UI and view the data using a web browser.

```
http://<INGRESS IP>/kibana
```

For more information on customizing the Kibana UI, see 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 Log in 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":
"ccp-efk.localLogForwarding.enabled=false,ccp-efk.localLogForwarding.elasticsearchHost='$EFK_SERVER',ccp-efk.localLogForwarding.elasticsearchPort='$EFK_PORT'"'}'
$MGMT_HOST/2/clusters/$CLUSTER_UUID/helmcharts
```

Renewing Kubernetes API Certificates

By default, the Kubernetes API certificates have a validity period of one year, after which you must manually renew these certificates.

To manually renew the Kubernetes API certificates for a cluster:

Step 1 SSH to a cluster master node.

Step 2 Check if the certificates are expired using following command:

```
sudo kubeadm alpha certs check-expiration
```

Step 3 On each master node of the cluster, create a backup of existing certificates and configuration files of the cluster. Perform following steps:

a) Create a backup directory.

```
mkdir /home/ccpuser/previous-certs
cd /etc/kubernetes/pki
mv *.* /home/ccpuser/previous-certs/
```

Step 4 Run the following commands:

a) Check for expired certificates.

```
sudo kubeadm alpha certs check-expiration
```

The command shows the expiration date and time of all the certificates.

b) Renew the certificates.

```
sudo kubeadm alpha certs renew <apiserver | all>
```

c) Verify that the certificates are renewed.

```
sudo kubeadm alpha certs check-expiration
```

d) Generate Kubernetes configuration files using following commands.

```
kubeadm alpha kubeconfig user --org system:masters --client-name kubernetes-admin >
/etc/kubernetes/admin.conf
kubeadm alpha kubeconfig user --client-name system:kube-controller-manager >
/etc/kubernetes/controller-manager.conf
kubeadm alpha kubeconfig user --client-name system:kube-scheduler >
/etc/kubernetes/scheduler.conf
kubeadm alpha kubeconfig user --org system:nodes --client-name system:node:$(hostname) >
/etc/kubernetes/kubelet.conf
```

e) Copy the newly generate admin.conf file to /etc/kubernetes/node.conf.

```
sudo cp /etc/kubernetes/admin.conf /etc/kubernetes/node.conf
```

f) Update the /etc/kubernetes/node.conf with local node IP address.

```
vi /etc/kubernetes/node.conf
```

g) Export your new admin.conf file to your host.

```
cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
chown $(id -u):$(id -g) $HOME/.kube/config
chmod 777 $HOME/.kube/config
export KUBECONFIG=.kube/config
```

h) Reboot the cluster master node.

```
sudo reboot
```

i) Verify that the master node has rejoined the Kubernetes cluster.

```
kubectl get nodes
```

Step 5 Repeat Step 3 and Step 4 for each cluster master node.
