



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 1](#)
- [Downloading Kubeconfig File, on page 1](#)
- [Sample Scenarios, on page 2](#)

Workflow of Deploying Applications

Task	Related Section
Create Kubernetes clusters using the Cisco Container Platform web interface.	Creating Kubernetes Clusters
Download the kubeconfig file that contains the cluster information and the certificates required to access clusters.	Downloading Kubeconfig File, on page 1
Use the kubectl utility to deploy the application and test the scenario.	Sample Scenarios, on page 2

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

This scenario describes deploying and configuring a pod with persistent volume.

Step 1 Go to the following URL:

<https://github.com/kubernetes/examples/tree/master/staging/volumes/vsphere>

Step 2 Download the following yaml files:

- vsphere-volume-sc-fast.yaml
- vsphere-volume-pvcsc.yaml
- vsphere-volume-pvcscpod.yaml

Step 3 Open the **kubect** utility.

Step 4 Configure the Kubernetes cluster.

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

Step 5 Create the storage class.

```
$ kubectl create -f vsphere-volume-sc-fast.yaml
```

Step 6 Verify if the storage cluster is created.

```
$ kubectl describe storageclass fast

Name: fast
IsDefaultClass: No
Annotations: <none>
Provisioner: kubernetes.io/vsphere-volume
Parameters: diskformat=zeroedthick,fstype=ext3
No events.```
```

Step 7 Create the persistent volume claim to request for storage.

```
$ kubectl create -f vsphere-volume-pvcsc.yaml
```

Step 8 Verify if the persistent volume claim (pvc) is created.

```
$ kubectl describe pvc pvcsc001

Name: pvcsc001
Namespace: default
StorageClass: fast
Status: Bound
Volume: pvc-83295256-f8e0-11e6-8263-005056b2349c
Labels: <none>
Capacity: 2Gi
Access Modes: RWO
Events:````
FirstSeen LastSeen Count From              SubObjectPath Type Reason      Message
1m         1m         1    persistentvolume Normal Provisioning Successfully provisioned
```

```
-controller                               Succeeded   volume pvc-83295256-f8e0
                                             -11e6-8263-005056b2349c
                                             using
```

```
kubernetes.io/vsphere-volume
```

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

Step 9 Verify if the persistent volume claim is created:

```
$ kubectl describe pv pvc-83295256-f8e0-11e6-8263-005056b2349c
```

```
Name: pvc-83295256-f8e0-11e6-8263-005056b2349c
Labels: <none>
StorageClass: fast
Status: Bound
Claim: default/pvcsc001
Reclaim Policy: Delete
Access Modes: RWO
Capacity: 2Gi
Message:
Source:
Type: vSphereVolume (a Persistent Disk resource in vSphere)
VolumePath: [datastore1] kubevols/kubernetes-dynamic-pvc-83295256-f8e0-11e6-8263-005056b2349c.vmdk
FSType: ext3
No events.
```

Note VMDK is created inside the *kubevols* folder in the datastore, which is specified in the `vsphere cloudprovider config` file that is created during the setup of Kubernetes cluster on vSphere.

Step 10 Create a pod that uses persistent volume claim with storage class.

```
$ kubectl create -f vsphere-volume-pvcscpod.yaml
```

Step 11 Verify if the pod is up and running.

```
$ kubectl get pod pvpod
```

NAME	READY	STATUS	RESTARTS	AGE
pvpod	1/1	Running	0	48m

Step 12 While the pod is starting, 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](#).

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=ingress-nginx -o wide
NAMESPACE      NAME                                READY  STATUS  RESTARTS  AGE  IP              NODE
ingressnginx   nginx-                              1/1    Running  0          3d   10.10.45.235   test-clusterwc5729f9ce2
                ingresscontroller
                -66974b775-jnmpl
```

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-6qmvm    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.

- a) Access the load balancer IP address 10.10.45.235, which is obtained in Step2.
- b) 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>
...
```
