[6] Deployment & Pod Scale

1. Pod Scale out

□Pod Scale Out

- Expand the number of running pods in “eshop-order” to 5.
  - □□ □□□□: k8s
  - namespace: devops
  - deployment: eshop-order

Reference

docs□□ reference □□ → scale

Kubectl Reference Docs

□□

```
# □□ □□□□ □□
[user@console ~]$ kubectl config use-context k8s

# devops□□ namespace□ □□□ □□
[user@console ~]$ kubectl get namespaces devops

# namespace□ devops□ □□□ deployment□ □□□ □□
# □□□□ □ 2□□ pod□ □□□□ □□□ □□
[user@console ~]$ kubectl get deployments.app -n devops
NAME          READY
eshop-order   2/2
```

```
# Deployment□ ReplicaSet□ 5□ □□□□
[user@console ~]$ kubectl scale deployment eshop-order -n devops --replicas=5

# □□□□ □ 5□□ pod□ □□□□ □□□ □□
[user@console ~]$ kubectl get deployments.app -n devops
NAME           READY
eshop-order    5/5
```

2. deployment □□□ scaling□□

□deployment □□□ scaling□□

- create a deployment as follows:
  - □□ □□□□: k8s
  - TASK:
    - name: webserver
    - 2 replicas
    - label: app_env_stage=dev
    - container name: webserver
    - container image: nginx:1.14
  - Scale Out deployment
    - Scale the deployment webserver to 3 pods

□□

```
# □□□ □□ deployment □□
# webserver.yaml □□□ YAML □□ □□
[user@console ~]$ kubectl create deployment webserver --image=nginx:1.14 --replicas=2 --dry-run=client -o
yaml > webserver.yaml

# webserver.yaml □□ □□
[user@console ~]$ vi webserver.yaml

matchLabels:
□app: webserver
□□□□ □□□ □□□□ □□□□
```

```
matchLabels:
  app_env_stage: dev
  □□□□

□□ metadata□ □□□□ labels □□□ □□□□

container□□□ name webserver□ □□□□

:wq

# □□□ yaml□□ □□
[user@console ~]$ kubectl apply -f webserver.yaml

# □□
[user@console ~]$ kubectl get deployments -o wide

# 3□□ Scale Out
[user@console ~]$ kubectl scale deployments webserver --replicas=3

# □□ (label □□□□ □□□□)
[user@console ~]$ kubectl get pods --show-labels
```

□□□□ (10□)

Task: Scale the deployment webserver to 3 pods.

- □□

  ```
  kubectl get deployments
  --> webserver□□ deployment□ □□□ □□
  ```

  ```
  kubectl scale deployment webserver --replicas=3
  ```

□□□□ (14□)

Task: Scale the deployment presentation to 6 pods.

- □□

```
# presentation□□□ deployment□ □□□ □□
kubectl get deployments presentation
```

```
kubectl scale deployment presentation --replicas=6
```

□□□□ (65□)

Create a deployment spec file that will: Launch 7 replicas of the nginx Image with the label app_runtime_stage=dev deployment name: kual00201.

Save a copy of this spec file to /opt/KUAL00201/spec_deployment.yaml (or /opt/KUAL00201/spec_deployment.json). When you are done, clean up (delete) any new Kubernetes API object that you produced during this task.

- □□

```
kubectl create deploy kual00201 --image=nginx --dry-run=client -o yaml >
/opt/KUAL00201/spec_deployment.yaml
```

```
vi /opt/KUAL00201/spec_deployment.yaml


apiVersion: apps/v1
kind: Deployment
metadata:
  creationTimestamp: null
  labels:
    app_runtime_stage: dev
  name: kual00201
spec:
  replicas: 7
  selector:
    matchLabels:
      app_runtime_stage: dev
  strategy: {}
  template:
    metadata:
      creationTimestamp: null
      labels:
```

```
        app_runtime_stage: dev
    spec:
      containers:
      - image: nginx
        name: nginx
        resources: {}
status: {}


:wq
```

```
kubectl apply -f /opt/KUAL00201/spec_deployment.yaml
```

```
# □□□□□□ □□
kubectl get deploy kual00201
```

```
# □□ □ □□□□
rm -rf /opt/KUAL00201/spec_deployment.yaml


kubectl delete deploy kual00201
```

Revision #1
Created 30 May 2023 05:26:47 by □□
Updated 20 June 2023 13:20:12 by □□