CKA Q&As

Certified Kubernetes Administrator (CKA) Program

Pass Linux Foundation CKA Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/cka.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Linux Foundation Official Exam Center

Instant Download After Purchase
100% Money Back Guarantee
365 Days Free Update
800,000+ Satisfied Customers

QUESTION 1

Create a namespace called \\'development\\' and a pod with image nginx called nginx on this namespace.

Correct Answer: Check the answer in explanation.

kubectl create namespace development kubectl run nginx --image=nginx --restart=Never -n development

QUESTION 2

SIMULATION

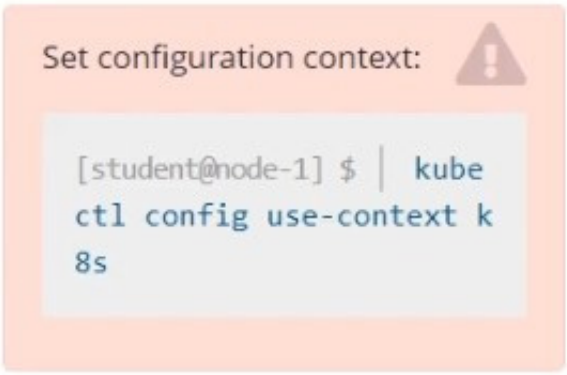

Context

You have been asked to create a new ClusterRole for a deployment pipeline and bind it to a specific ServiceAccount scoped to a specific namespace.

Task

Create a new ClusterRole named deployment-clusterrole, which only allows to create the following resource types:

1.

Deployment

2.

StatefulSet

3.

DaemonSet

Create a new ServiceAccount named cicd-token in the existing namespace app-team1.

Bind the new ClusterRole deployment-clusterrole lo the new ServiceAccount cicd-token , limited to the namespace app-team1.

Correct Answer: Check the answer in explanation.

```
The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ kubectl create clusterrole deployment-clusterrole --verb=create --resource=Deployment,StatefulSet,DaemonS
et
clusterrole.rbac.authorization.k8s.io/deployment-clusterrole created
student@node-1:~$ kubectl create sa cicd-token --namespace app-team1
serviceaccount/cicd-token created
student@node-1:~$ kubectl create clusterrolebinding deploy-b --clusterrole=deployment-clusterrole --serviceaccount=app-team
1:cicd-token
clusterrolebinding.rbac.authorization.k8s.io/deploy-b created
student@node-1:~$
```

Task should be complete on node k8s -1 master, 2 worker for this connect use command [student@node-1] > ssh k8s kubectl create clusterrole deployment-clusterrole --verb=create -- resource=deployments,statefulsets,daemonsets kubectl create serviceaccount cicd-token --namespace=app-team1 kubectl create rolebinding deployment-clusterrole --clusterrole=deployment-clusterrole -- serviceaccount=default:cicd-token --namespace=app-team1

QUESTION 3

SIMULATION

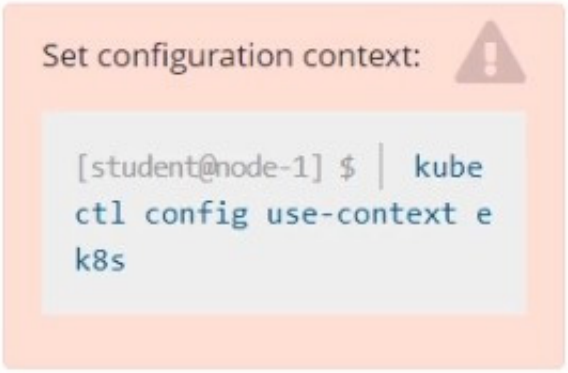

Set the node named ek8s-node-1 as unavailable and reschedule all the pods running on it.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context ek8s
Switched to context "ek8s".
student@node-1:~$ kubectl get nodes
NAME           STATUS   ROLES                  AGE   VERSION
ek8s-master-0  Ready    control-plane,master   67d   v1.23.1
ek8s-node-0    Ready    <none>                 67d   v1.23.1
ek8s-node-1    Ready    <none>                 67d   v1.23.1
student@node-1:~$ kubectl drain ek8s-node-1 --ignore-daemonsets
node/ek8s-node-1 cordoned
error: unable to drain node "ek8s-node-1" due to error:cannot delete Pods with local storage (use --delete-emptydir-data to
 override): kube-system/metrics-server-7cb5455c67-m6qvd, continuing command...
There are pending nodes to be drained:
 ek8s-node-1
cannot delete Pods with local storage (use --delete-emptydir-data to override): kube-system/metrics-server-7cb5455c67-m6qvd
student@node-1:~$ kubectl drain ek8s-node-1 --ignore-daemonsets --delete-emptydir-data
node/ek8s-node-1 already cordoned
WARNING: ignoring DaemonSet-managed Pods: kube-system/kube-flannel-ds-chvkf, kube-system/kube-proxy-7pf29
evicting pod kube-system/metrics-server-7cb5455c67-m6qvd
evicting pod default/nginx-5cb786cffd-vjbs8
pod/nginx-5cb786cffd-vjbs8 evicted
pod/metrics-server-7cb5455c67-m6qvd evicted
node/ek8s-node-1 drained
student@node-1:~$ kubectl get nodes
NAME           STATUS                     ROLES                  AGE   VERSION
ek8s-master-0  Ready                      control-plane,master   67d   v1.23.1
ek8s-node-0    Ready                      <none>                 67d   v1.23.1
ek8s-node-1    Ready,SchedulingDisabled   <none>                 67d   v1.23.1
student@node-1:~$
```

QUESTION 4

SIMULATION

Create a pod as follows: Name: non-persistent-redis container Image: redis Volume with name: cache-control Mount path: /data/redis The pod should launch in the staging namespace and the volume must not be persistent.

Correct Answer: Check the answer in explanation.

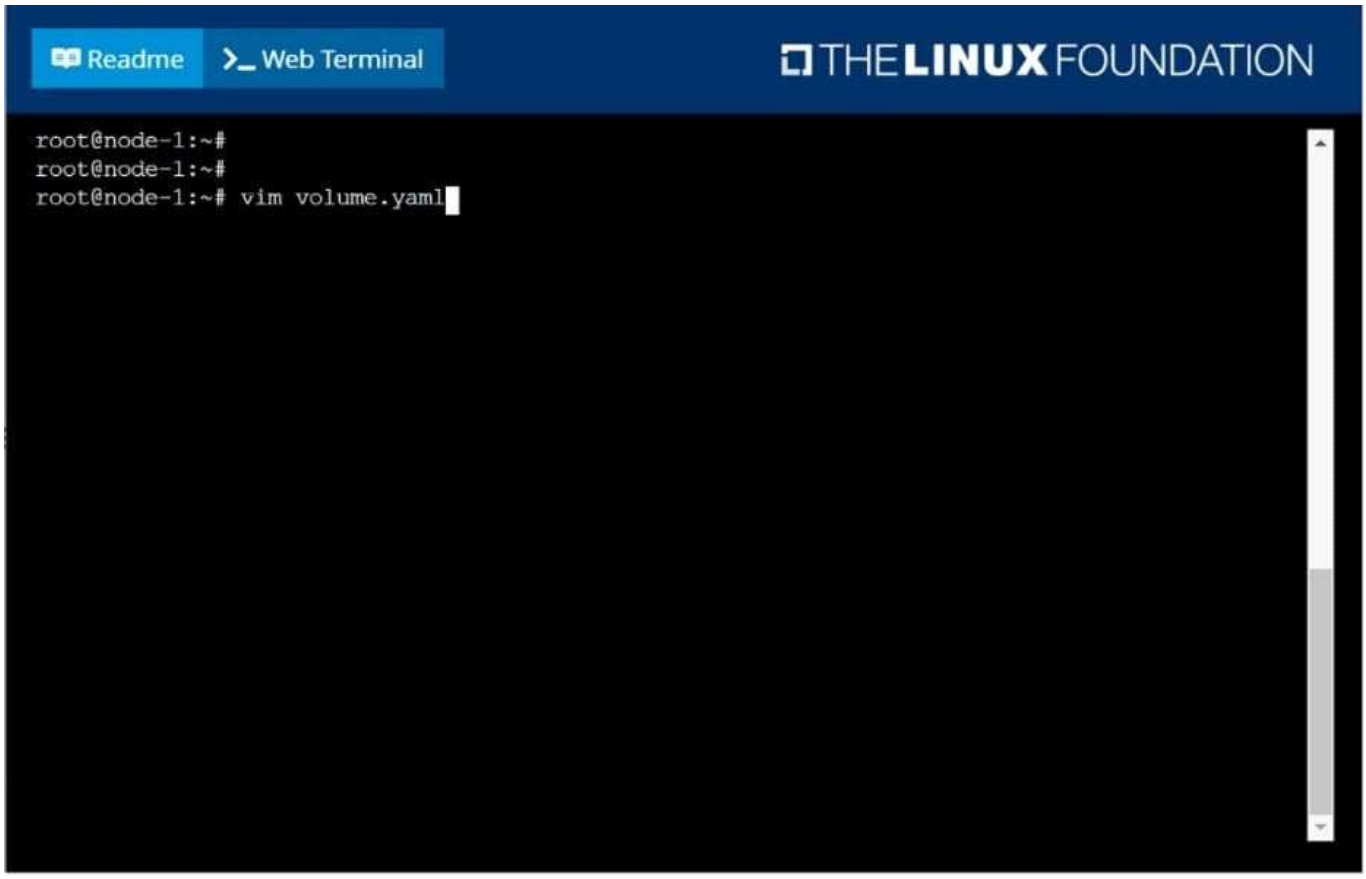

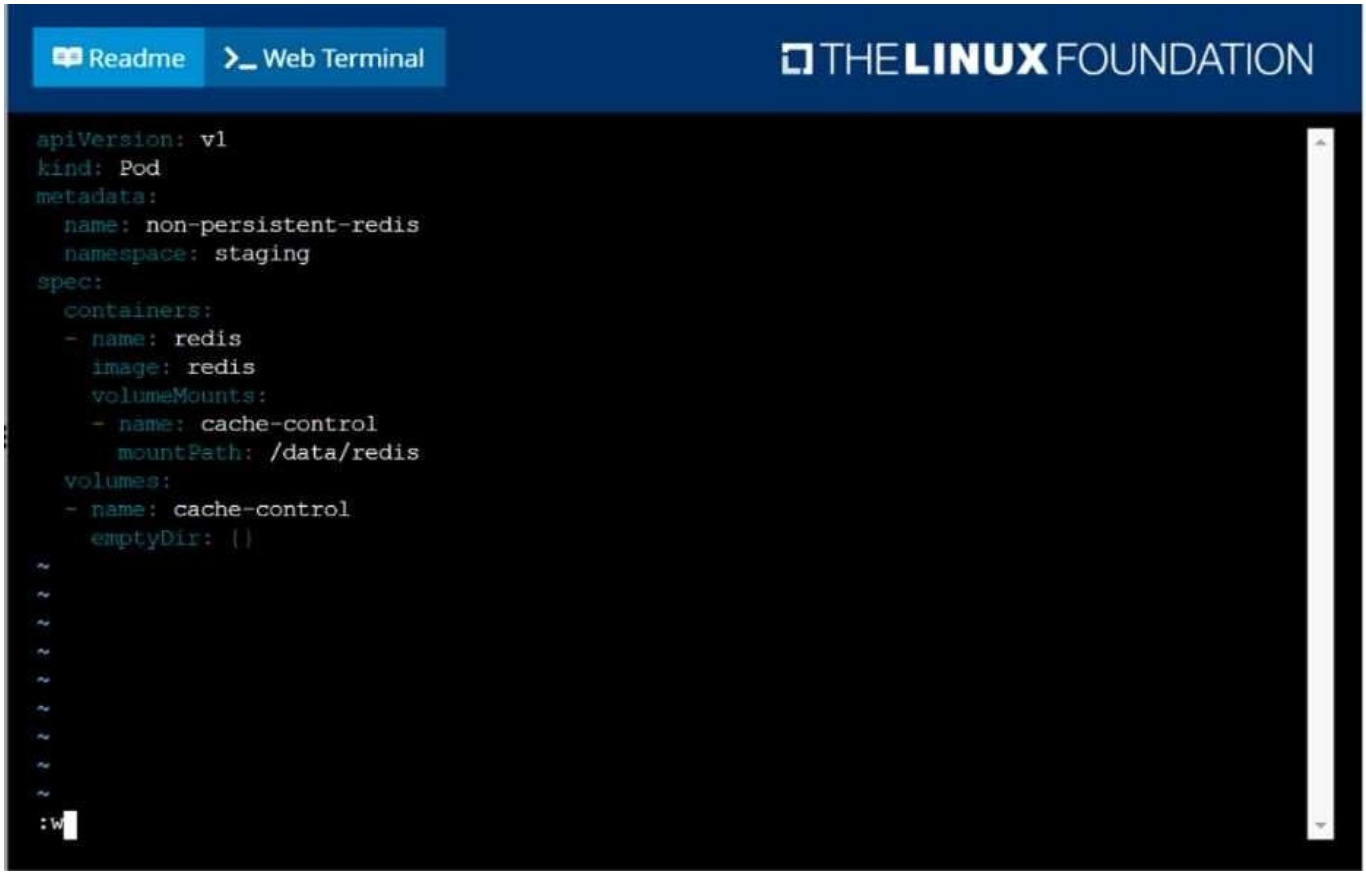
Readme
Web Terminal
THE LINUX FOUNDATION
apiVersion: v1
kind: Pod
metadata:
name: non-persistent-redis
namespace: staging
spec:
containers:
- name: redis
image: redis
volumeMounts:
- name: cache-control
mountPath: /data/redis
volumes:
- name: cache-control
emptyDir: {}
:w

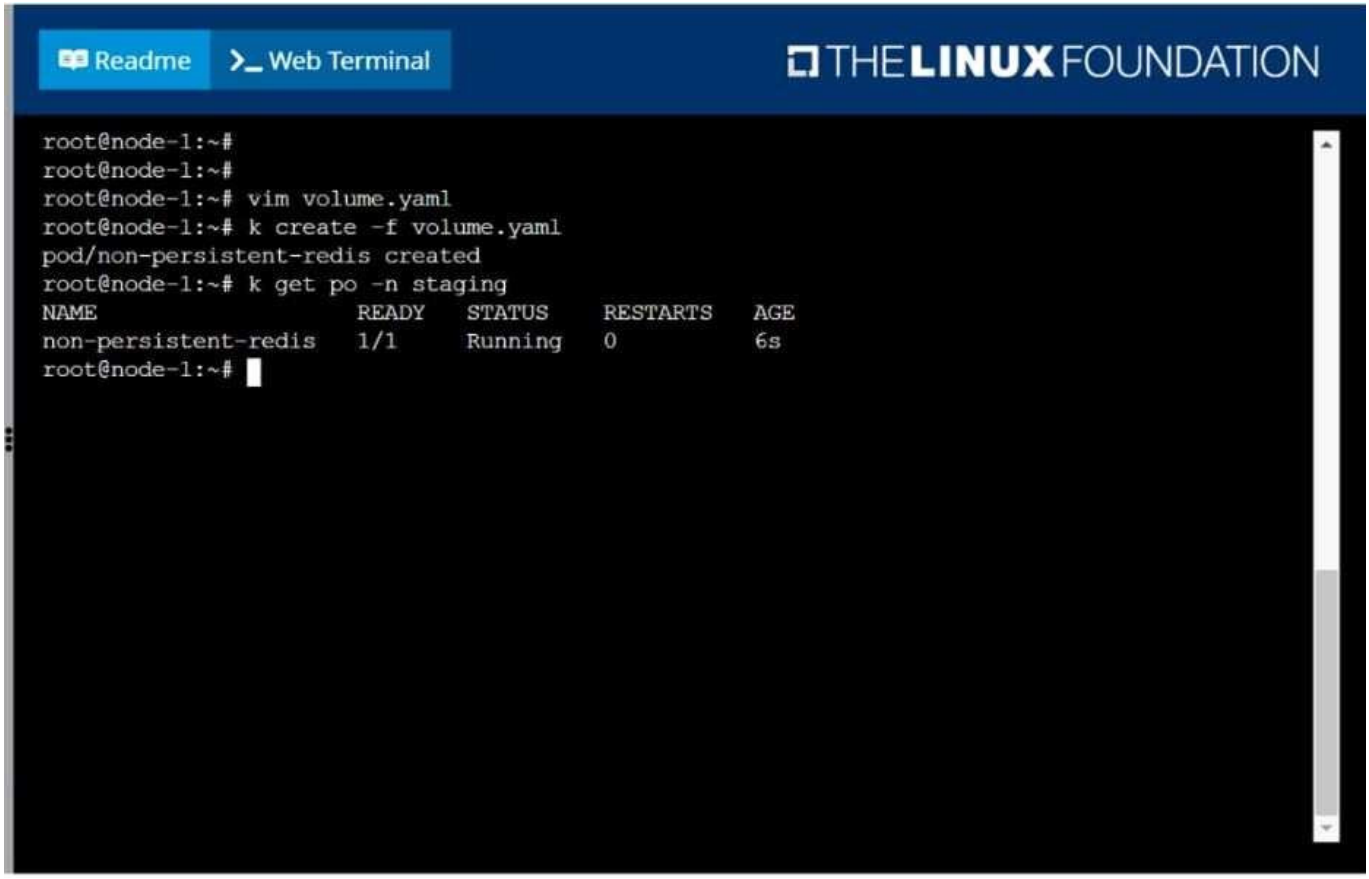
Readme
Web Terminal
THE LINUX FOUNDATION
root@node-1:~#
root@node-1:~#
root@node-1:~# vim volume.yaml
root@node-1:~# k create -f volume.yaml
pod/non-persistent-redis created
root@node-1:~# k get po -n staging
NAME READY STATUS RESTARTS AGE
non-persistent-redis 1/1 Running 0 6s
root@node-1:~#

QUESTION 5

Create a pod that having 3 containers in it? (Multi-Container)

Correct Answer: Check the answer in explanation.

image=nginx, image=redis, image=consul Name nginx container as "nginx-container" Name redis container as "redis-container" Name consul container as "consul-container" Create a pod manifest file for a container and append container section for rest of the images kubectl run multi-container --generator=run-pod/v1 --image=nginx -- dry-run -o yaml > multi-container.yaml # then vim multi-container.yaml apiVersion: v1 kind: Pod metadata: labels: run: multi-container name: multi-container spec: containers:

-

image: nginx name: nginx-container

-

image: redis name: redis-container

-

image: consul name: consul-container restartPolicy: Always

Latest CKA Dumps CKA PDF Dumps CKA Braindumps