

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.certbus.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 busybox pod and add "sleep 3600" command

Correct Answer: Check the answer in explanation.

Solution

```
kubectl run busybox --image=busybox --restart=Never -- /bin/sh -c "sleep 3600"
```

QUESTION 2

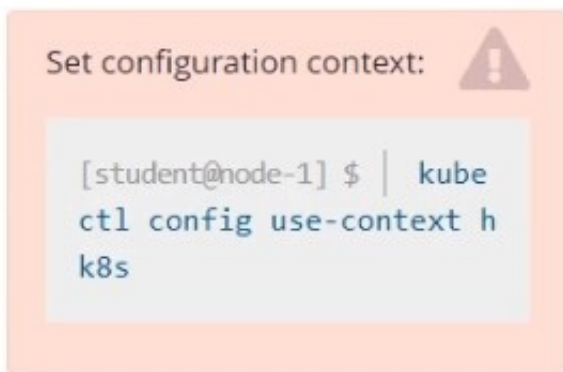
Create a pod that echo "hello world" and then exists. Have the pod deleted automatically when it's completed

Correct Answer: Check the answer in explanation.

```
kubectl run busybox --image=busybox -it --rm --restart=Never -- /bin/sh -c '\`echo hello world\`' kubectl get po # You shouldn't see pod with the name "busybox"
```

QUESTION 3

SIMULATION



Task

Create a persistent volume with name app-data , of capacity 1Gi and access mode ReadOnlyMany. The type of volume is hostPath and its location is /srv/app-data .

Correct Answer: Check the answer in explanation.

```
#vi pv.yaml apiVersion: v1 kind: PersistentVolume metadata: name: app-config spec: capacity: storage: 1Gi accessModes:
```

```
-ReadOnlyMany hostPath: path: /srv/app-config # kubectl create -f pv.yaml
```

QUESTION 4

Create 2 nginx image pods in which one of them is labelled with env=prod and another one labelled with env=dev and verify the same.

Correct Answer: Check the answer in explanation.

Solution

kubectl run --generator=run-pod/v1 --image=nginx -- labels=env=prod nginx-prod --dry-run -o yaml > nginx-prodpod.yaml Now, edit nginx-prod-pod.yaml file and remove entries like "creationTimestamp":

```
null" "dnsPolicy: ClusterFirst"
```

```
vim nginx-prod-pod.yaml
```

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
labels:
```

```
env: prod
```

```
name: nginx-prod
```

```
spec:
```

```
containers:
```

```
-
```

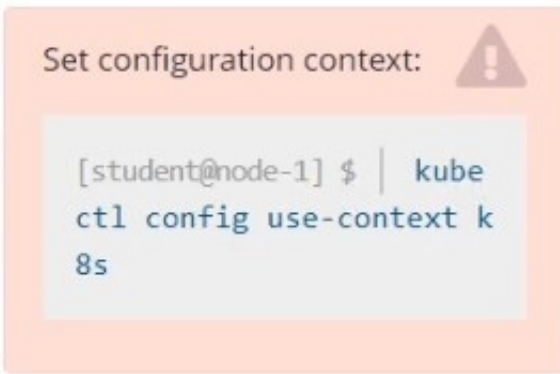
```
image: nginx name: nginx-prod restartPolicy: Always # kubectl create -f nginx-prod-pod.yaml kubectl run --generator=run-pod/v1 --image=nginx -labels=env=dev nginx-dev --dry-run -o yaml > nginx-dev-pod.yaml apiVersion: v1 kind: Pod metadata: labels: env: dev name: nginx-dev spec: containers:
```

```
-
```

```
image: nginx name: nginx-dev restartPolicy: Always # kubectl create -f nginx-prod-dev.yaml Verify : kubectl get po --show-labels kubectl get po -l env=prod kubectl get po -l env=dev
```

QUESTION 5

SIMULATION



Task

Check to see how many nodes are ready (not including nodes tainted NoSchedule) and write the number to /opt/KUSC00402/kusc00402.txt.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl get nodes  
NAME                STATUS    ROLES    AGE   VERSION  
k8s-master-0       Ready    control-plane,master   67d   v1.23.1  
k8s-node-0         Ready    <none>   67d   v1.23.1  
k8s-node-1         Ready    <none>   67d   v1.23.1  
student@node-1:~$  
student@node-1:~$  
student@node-1:~$  
student@node-1:~$ echo "2" > /opt/KUSC00402/kusc00402.txt  
student@node-1:~$ cat /opt/KUSC00402/kusc00402.txt  
2  
student@node-1:~$
```

QUESTION 6

List all the pods sorted by name

Correct Answer: Check the answer in explanation.

kubect1 get pods --sort-by=.metadata.name

QUESTION 7


List pod logs named "frontend" and search for the pattern "started" and write it to a file "/opt/error- logs"

Correct Answer: Check the answer in explanation.

Kubectl logs frontend | grep -i "started" > /opt/error-logs

QUESTION 8


SIMULATION

```
Set configuration context:   
  
[student@node-1] $ | kube  
ctl config use-context m  
k8s
```


Task

Given an existing Kubernetes cluster running version 1.20.0, upgrade all of the Kubernetes control plane and node components on the master node only to version 1.20.1.

Be sure to drain the master node before upgrading it and uncordon it after the upgrade.

```
You can ssh to the master   
node using:  
  
[student@node-1] $ | ssh  
mk8s-master-0  
  
You can assume elevated  
privileges on the master node  
with the following command:  
  
[student@mk8s-master-0] |  
$  
sudo -i
```

You are also expected to upgrade kubelet and kubectl on the master node.

Do not upgrade the worker nodes, etcd, the container manager, the CNI plugin, the DNS service or any other addons. 

Correct Answer: Check the answer in explanation.

```

student@node-1:~$ kubectl config use-context mk8s
Switched to context "mk8s".
student@node-1:~$ kubectl config use-context mk8s
Switched to context "mk8s".
student@node-1:~$ kubectl get nodes
NAME                STATUS    ROLES    AGE   VERSION
mk8s-master-0       Ready    control-plane,master   67d   v1.22.1
mk8s-node-0         Ready    <none>    67d   v1.22.1
student@node-1:~$ kubectl drain mk8s-master-0 --ignore-daemonsets
node/mk8s-master-0 cordoned
WARNING: ignoring DaemonSet-managed Pods: kube-system/kube-flannel-ds-jxzmk, kube-system/kube-proxy-9rzg9
evicting pod kube-system/coredns-78fcd69978-tt2b8
evicting pod default/nginx-74b46d4cfc-dfkvs
evicting pod kube-system/coredns-78fcd69978-nbkmz
pod/nginx-74b46d4cfc-dfkvs evicted
pod/coredns-78fcd69978-tt2b8 evicted
pod/coredns-78fcd69978-nbkmz evicted
node/mk8s-master-0 drained
student@node-1:~$ kubectl get nodes
NAME                STATUS    ROLES    AGE   VERSION
mk8s-master-0       Ready,SchedulingDisabled control-plane,master   67d   v1.22.1
mk8s-node-0         Ready    <none>    67d   v1.22.1
student@node-1:~$ ssh mk8s-master-0
Warning: Permanently added '10.250.5.55' (ECDSA) to the list of known hosts.
Welcome to Ubuntu 20.04.3 LTS (GNU/Linux 5.11.0-1028-aws x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

System information as of Mon Apr 25 09:30:48 UTC 2022

System load:  1.98           Users logged in:  0
Usage of /:   83.2% of 67.79GB IPv4 address for cni0: 10.244.0.1
Memory usage: 2%           IPv4 address for docker0: 172.17.0.1
Swap usage:   0%            IPv4 address for eth0: 10.250.5.55
Processes:   85

30 updates can be applied immediately.
15 of these updates are standard security updates.
To see these additional updates run: apt list --upgradable

The list of available updates is more than a week old.
To check for new updates run: sudo apt update

student@mk8s-master-0:~$ sudo -i
root@mk8s-master-0:~# apt install kubeadm=1.22.2-00 kubelet=1.22.2-00 kubectl=1.22.2-00
Reading package lists... Done
Building dependency tree
Reading state information... Done
The following packages will be upgraded:
  kubeadm kubectl kubelet
3 upgraded, 0 newly installed, 0 to remove and 27 not upgraded.
Need to get 39.6 MB of archives.
After this operation, 0 B of additional disk space will be used.
Get:1 https://packages.cloud.google.com/apt/kubernetes-xenial/main amd64 kubelet amd64 1.22.2-00 [21.9 MB]
Get:2 https://packages.cloud.google.com/apt/kubernetes-xenial/main amd64 kubectl amd64 1.22.2-00 [9038 kB]
Get:3 https://packages.cloud.google.com/apt/kubernetes-xenial/main amd64 kubeadm amd64 1.22.2-00 [8718 kB]
Fetched 39.6 MB in 13s (3156 kB/s)
(Reading database ... 33901 files and directories currently installed.)
Preparing to unpack .../kubelet_1.22.2-00_amd64.deb ...
Unpacking kubelet (1.22.2-00) over (1.22.1-00) ...
Preparing to unpack .../kubectl_1.22.2-00_amd64.deb ...
Unpacking kubectl (1.22.2-00) over (1.22.1-00) ...
Preparing to unpack .../kubeadm_1.22.2-00_amd64.deb ...
Unpacking kubeadm (1.22.2-00) over (1.22.1-00) ...
Setting up kubectl (1.22.2-00) ...
Setting up kubelet (1.22.2-00) ...
Setting up kubeadm (1.22.2-00) ...
root@mk8s-master-0:~# apt install kubeadm=1.22.2-00 kubelet=1.22.2-00 kubectl=1.22.2-00
Reading package lists... Done
Building dependency tree
Reading state information... Done
kubeadm is already the newest version (1.22.2-00).
kubectl is already the newest version (1.22.2-00).
kubelet is already the newest version (1.22.2-00).
0 upgraded, 0 newly installed, 0 to remove and 27 not upgraded.
root@mk8s-master-0:~# kubeadm upgrade plan
[upgrade/config] Making sure the configuration is correct:
[upgrade/config] Reading configuration from the cluster...
[upgrade/config] FYI: You can look at this config file with 'kubectl -n kube-system get cm kubeadm-config -o yaml'
[preflight] Running pre-flight checks.
[upgrade] Running cluster health checks
[upgrade] Fetching available versions to upgrade to
[upgrade/versions] Cluster version: v1.22.1
[upgrade/versions] kubeadm version: v1.22.2
    
```



```

COMPONENT    CURRENT    TARGET
kubeadm      1 x v1.22.1  v1.22.9
              1 x v1.22.2  v1.22.9

Upgrade to the latest version in the v1.22 series:

COMPONENT    CURRENT    TARGET
kube-apiserver  v1.22.1  v1.22.9
kube-controller-manager v1.22.1  v1.22.9
kube-scheduler  v1.22.1  v1.22.9
kube-proxy      v1.22.1  v1.22.9
CoreDNS         v1.8.4    v1.8.4
etcd            3.5.0-0   3.5.0-0

You can now apply the upgrade by executing the following command:

    kubeadm upgrade apply v1.22.9

Note: Before you can perform this upgrade, you have to update kubeadm to v1.22.9.

-----

The table below shows the current state of component configs as understood by this version of kubeadm.
Configs that have a "yes" mark in the "MANUAL UPGRADE REQUIRED" column require manual config upgrade or
resetting to kubeadm defaults before a successful upgrade can be performed. The version to manually
upgrade to is denoted in the "PREFERRED VERSION" column.

API GROUP          CURRENT VERSION    PREFERRED VERSION    MANUAL UPGRADE REQUIRED
-----
kubeproxy.config.k8s.io  v1alpha          v1alpha              no
kubeadm.config.k8s.io    v1beta1          v1beta1              no

-----

root@mk8s-master-0:~# kubeadm upgrade apply v1.22.2
[upgrade/config] Making sure the configuration is correct:
[upgrade/config] Reading configuration from the cluster...
[upgrade/config] FYI: You can look at this config file with 'kubectl -n kube-system get cm kubeadm-config -o yaml'
[preflight] Running pre-flight checks.
[upgrade] Running cluster health checks
[upgrade/version] You have chosen to change the cluster version to "v1.22.2"
[upgrade/versions] Cluster version: v1.22.1
[upgrade/versions] kubeadm version: v1.22.2
[upgrade/confirm] Are you sure you want to proceed with the upgrade? [y/N]: y
[upgrade/prepull] Pulling images required for setting up a Kubernetes cluster
[upgrade/prepull] This might take a minute or two, depending on the speed of your internet connection
[upgrade/prepull] You can also perform this action in beforehand using 'kubeadm config images pull'
[upgrade/apply] Upgrading your Static Pod-hosted control plane to version "v1.22.2"...
Static pod: kube-apiserver-mk8s-master-0 hash: b1d9f2b63ce85cb6310a6d8f6f728f03
Static pod: kube-controller-manager-mk8s-master-0 hash: 91af4173de8872b5f7aec58b2fc0f1fc
Static pod: kube-scheduler-mk8s-master-0 hash: d98fe109788b5b498301dd6c53afcfa9
[upgrade/etcd] Upgrading to TLS for etcd
Static pod: etcd-mk8s-master-0 hash: 6828726f91dba72616d11ac4a737e533
[upgrade/staticpods] Preparing for "etcd" upgrade
[upgrade/staticpods] Current and new manifests of etcd are equal, skipping upgrade
[upgrade/etcd] Waiting for etcd to become available
[upgrade/staticpods] Writing new Static Pod manifests to "/etc/kubernetes/tmp/kubeadm-upgraded-manifests804306747"
[upgrade/staticpods] Preparing for "kube-apiserver" upgrade
[upgrade/staticpods] Renewing apiserver certificate
[upgrade/staticpods] Renewing apiserver-kubelet-client certificate
[upgrade/staticpods] Renewing front-proxy-client certificate
[upgrade/staticpods] Renewing apiserver-etcd-client certificate
[upgrade/staticpods] Moved new manifest to "/etc/kubernetes/manifests/kube-apiserver.yaml" and backed up old manifest to
/etc/kubernetes/tmp/kubeadm-backup-manifests-2022-04-25-15-11-18/kube-apiserver.yaml"
[upgrade/staticpods] Waiting for the kubelet to restart the component
[upgrade/staticpods] This might take a minute or longer depending on the component/version gap (timeout 5m0s)
Static pod: kube-scheduler-mk8s-master-0 hash: d98fe109788b5b498301dd6c53afcfa9
Static pod: kube-scheduler-mk8s-master-0 hash: acb75f76060c8873ac4bf8b2fcl9466
[apiclient] Found 1 Pods for label selector component=kube-scheduler
[upgrade/staticpods] Component "kube-scheduler" upgraded successfully!
[upgrade/postupgrade] Applying label node-role.kubernetes.io/control-plane="" to Nodes with label node-role.kubernetes.io
aster="" (deprecated)
[upload-config] Storing the configuration used in ConfigMap "kubeadm-config" in the "kube-system" Namespace
[kubelet] Creating a ConfigMap "kubelet-config-1.22" in namespace kube-system with the configuration for the kubelets in
e cluster
[kubelet-start] Writing kubelet configuration to file "/var/lib/kubelet/config.yaml"
[bootstrap-token] configured RBAC rules to allow Node Bootstrap tokens to get nodes
[bootstrap-token] configured RBAC rules to allow Node Bootstrap tokens to post CSRs in order for nodes to get long term
tificate credentials
[bootstrap-token] configured RBAC rules to allow the csrapprover controller automatically approve CSRs from a Node Boot
p Token
[bootstrap-token] configured RBAC rules to allow certificate rotation for all node client certificates in the cluster
[addons] Applied essential addon: CoreDNS
[addons] Applied essential addon: kube-proxy

[upgrade/successful] SUCCESS! Your cluster was upgraded to "v1.22.2". Enjoy!

[upgrade/kubelet] Now that your control plane is upgraded, please proceed with upgrading your kubelets if you haven't al
dy done so.
root@mk8s-master-0:~# systemctl restart kubelet
root@mk8s-master-0:~# exit
logout
student@mk8s-master-0:~$ exit
logout
Connection to 10.250.5.55 closed.
student@node-1:~$ kubectl uncordon mk8s-master-0
node/mk8s-master-0 uncordoned
student@node-1:~$ kubectl get nodes
NAME                STATUS    ROLES    AGE   VERSION
mk8s-master-0      Ready    control-plane,master   67d   v1.22.2
mk8s-node-0        Ready    <none>    67d   v1.22.1
student@node-1:~$
    
```


QUESTION 9

Get IP address of the pod - "nginx-dev"

Correct Answer: Check the answer in explanation.

Kubect1 get po -o wide Using JsonPath kubect1 get pods -o=jsonpath='{range .items[*]}{.metadata.name}{"\t"}{.status.podIP}{"\n"}{end}'

QUESTION 10

SIMULATION



Task Create a new PersistentVolumeClaim

1.

Name: pv-volume

2.

Class: csi-hostpath-sc

3.

Capacity: 10Mi

Create a new Pod which mounts the PersistentVolumeClaim as a volume:

1.

Name: web-server

2.

Image: nginx

3.

Mount path: /usr/share/nginx/html

Configure the new Pod to have ReadWriteOnce access on the volume.

Finally, using kubectl edit or kubectl patch expand the PersistentVolumeClaim to a capacity of 70Mi and record that change.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context ok8s
Switched to context "ok8s".
student@node-1:~$
```

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: pv-volume
spec:
  storageClassName: csi-hostpath-sc
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Mi
```

```
student@node-1:~$ kubectl config use-context ok8s
Switched to context "ok8s".
student@node-1:~$ vim pvc.yml
student@node-1:~$ kubectl get pv,pvc
No resources found
student@node-1:~$ kubectl create -f pvc.yml
persistentvolumeclaim/pv-volume created
student@node-1:~$ kubectl get pv,pvc
```

NAME	STORAGECLASS	REASON	AGE	CAPACITY	ACCESS MODES	RECLAIM POLICY	STATUS	CLAIM
persistentvolume/pvc-6b6c71cb-558d-4b47-a0db-3951737097eb	csi-hostpath-sc		3s	10Mi	RWO	Delete	Bound	default/pv-volume

```
NAME                STATUS    VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLAS
persistentvolumeclaim/pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi       RWO             csi-hostpat
h-sc                 3s
student@node-1:~$
```

```
apiVersion: v1
kind: Pod
metadata:
  name: web-server
spec:
  volumes:
    - name: task-pv-storage
      persistentVolumeClaim:
        claimName: pv-volume
  containers:
    - name: web-server
      image: nginx
      volumeMounts:
        - mountPath: "/usr/share/nginx/html"
          name: task-pv-storage
```

```
student@node-1:~$ kubectl config use-context ok8s
Switched to context "ok8s".
student@node-1:~$ vim pvc.yml
student@node-1:~$ kubectl get pv,pvc
No resources found
student@node-1:~$ kubectl create -f pvc.yml
persistentvolumeclaim/pv-volume created
student@node-1:~$ kubectl get pv,pvc
NAME                                     CAPACITY  ACCESS MODES  RECLAIM POLICY  STATUS  CLAIM
persistentvolume/pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           Delete          Bound   default/pv-volume
csi-hostpath-sc
NAME                                     STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLAS
S
persistentvolumeclaim/pv-volume         Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           csi-hostpat
h-sc 3s
student@node-1:~$
student@node-1:~$
student@node-1:~$ vim pod.yml
student@node-1:~$ kubectl create -f pod.yml
pod/web-server created
student@node-1:~$ kubectl get pods
NAME                READY  STATUS             RESTARTS  AGE
csi-hostpath-socat-0  1/1    Running            0          6h12m
csi-hostpathplugin-0  9/9    Running            0          6h12m
web-server           0/1    ContainerCreating  0          5s
student@node-1:~$ kubectl get pods -w
NAME                READY  STATUS             RESTARTS  AGE
csi-hostpath-socat-0  1/1    Running            0          6h12m
csi-hostpathplugin-0  9/9    Running            0          6h12m
web-server           0/1    ContainerCreating  0          10s
web-server           1/1    Running            0          22s
^Cstudent@node-1:~$ kubectl get pods
NAME                READY  STATUS             RESTARTS  AGE
csi-hostpath-socat-0  1/1    Running            0          6h12m
csi-hostpathplugin-0  9/9    Running            0          6h12m
web-server           1/1    Running            0          27s
student@node-1:~$ kubectl edit pvc pv-volume
```

```
# Please edit the object below. Lines beginning with a '#' will be ignored,
# and an empty file will abort the edit. If an error occurs while saving this file will be
# reopened with the relevant failures.
#
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  annotations:
    pv.kubernetes.io/bind-completed: "yes"
    pv.kubernetes.io/bound-by-controller: "yes"
    volume.beta.kubernetes.io/storage-provisioner: hostpath.csi.k8s.io
    volume.kubernetes.io/storage-provisioner: hostpath.csi.k8s.io
  creationTimestamp: "2022-04-25T15:37:42Z"
  finalizers:
  - kubernetes.io/pvc-protection
  name: pv-volume
  namespace: default
  resourceVersion: "42413"
  uid: 6b6c71cb-558d-4b47-a0db-3951737097eb
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 70Mi
  storageClassName: csi-hostpath-sc
  volumeMode: Filesystem
  volumeName: pvc-6b6c71cb-558d-4b47-a0db-3951737097eb
status:
  accessModes:
  - ReadWriteOnce
  capacity:
    storage: 10Mi
:WQ
```

```
student@node-1:~$ kubectl get pvc
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pv-volume           Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           csi-hostpath-sc  2m46s
student@node-1:~$ kubectl get pvc
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pv-volume           Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           csi-hostpath-sc  2m56s
student@node-1:~$ kubectl get pvc
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pv-volume           Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           csi-hostpath-sc  3m
student@node-1:~$ kubectl get pvc
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pv-volume           Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           csi-hostpath-sc  3m5s
student@node-1:~$ kubectl get pvc
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pv-volume           Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  10Mi      RWO           csi-hostpath-sc  3m9s
student@node-1:~$ kubectl get pvc
NAME                STATUS  VOLUME                                     CAPACITY  ACCESS MODES  STORAGECLASS  AGE
pv-volume           Bound   pvc-6b6c71cb-558d-4b47-a0db-3951737097eb  70Mi      RWO           csi-hostpath-sc  3m42s
student@node-1:~$ kubectl config use-context k8s
```

QUESTION 11

Create a nginx pod with label env=test in engineering namespace .

Correct Answer: Check the answer in explanation.

```
kubectl run nginx --image=nginx --restart=Never --labels=env=test --namespace=engineering --dry-run -o yaml > nginx-pod.yaml  
kubectl run nginx --image=nginx --restart=Never --labels=env=test --namespace=engineering --dry-run -o yaml | kubectl create -n engineering -f YAML File:
```

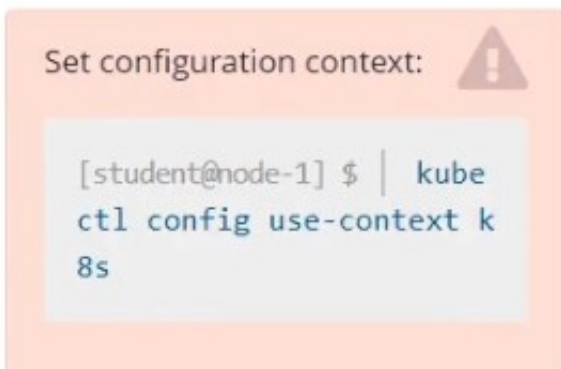
apiVersion: v1 kind: Pod metadata: name: nginx namespace: engineering labels: env: test spec: containers:

-name: nginx image: nginx imagePullPolicy: IfNotPresent restartPolicy: Never

```
kubectl create -f nginx-pod.yaml
```

QUESTION 12

SIMULATION



Task

Scale the deployment presentation to 6 pods.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ vim ping.yml
student@node-1:~$ kubectl create -f ping.yml
ingress.networking.k8s.io/ping created
student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ kubectl get deployments.apps
NAME                READY   UP-TO-DATE   AVAILABLE   AGE
front-end           2/2     2             2           6h2m
presentation        2/2     2             2           6h1m
student@node-1:~$ kubectl scale deployment presentation --replicas=3
deployment.apps/presentation scaled
student@node-1:~$ kubectl get deployments.apps
NAME                READY   UP-TO-DATE   AVAILABLE   AGE
front-end           2/2     2             2           6h2m
presentation        2/3     3             2           6h1m
student@node-1:~$ kubectl get pods
NAME                READY   STATUS    RESTARTS   AGE
big-corp-app        1/1     Running   0           5h58m
foo                 1/1     Running   0           5h58m
front-end-6bc87b9748-n7v8h  1/1     Running   0           3m47s
front-end-6bc87b9748-zmb8g  1/1     Running   0           3m45s
overloaded-cpu-98b9se  1/1     Running   0           5h57m
overloaded-cpu-ab2d3s  1/1     Running   0           5h57m
overloaded-cpu-kipb9a  1/1     Running   0           5h57m
presentation-684cd7ccb6-4gf56  1/1     Running   0           6h1m
presentation-684cd7ccb6-6zjls  1/1     Running   0           13s
presentation-684cd7ccb6-vshxj  1/1     Running   0           6h1m
student@node-1:~$ █
```

QUESTION 13

List "nginx-dev" and "nginx-prod" pod and delete those pods

Correct Answer: Check the answer in explanation.

Solution

kubectl get pods -o wide kubectl delete po "nginx-dev" kubectl delete po "nginx-prod"

QUESTION 14

List all the pods sorted by name

Correct Answer: Check the answer in explanation.

Solution

kubectl get pods --sort-by=.metadata.name

QUESTION 15

Print pod name and start time to "/opt/pod-status" file

Correct Answer: Check the answer in explanation.

Solution

```
kubect1 get pods -o=jsonpath=\{\{range .items[*]\}\{.metadata.name\}\{\t\}\{.status.podIP\}\{\n\}\{end\}\}
```

[CKA PDF Dumps](#)

[CKA VCE Dumps](#)

[CKA Braindumps](#)