

## Linux Foundation CKA Test Engine Dumps Training With 68 Questions [Q21-Q38]



Linux Foundation CKA Test Engine Dumps Training With 68 Questions  
CKA Questions Pass on Your First Attempt Dumps for Kubernetes Administrator Certified

The CKA certification is valid for three years, and candidates can renew their certification by retaking the exam or by earning a higher-level certification. Certified Kubernetes Administrator (CKA) Program Exam certification is recognized by industry leaders, including AWS, Google Cloud, and Red Hat. Certified Kubernetes Administrator (CKA) Program Exam certification also provides access to a network of certified professionals and resources, including training, events, and community support. Overall, the CKA certification is an excellent way to demonstrate your expertise in Kubernetes and advance your career in the cloud-native ecosystem.

**NO.21** List pod logs named `&#8220;frontend&#8221;` and search for the pattern `&#8220;started&#8221;` and write it to a file `&#8220;/opt/error-logs&#8221;`;

See the solution below.

Explanation

Kubectl logs frontend | grep -i &#8220;started&#8221; > /opt/error-logs

**NO.22** Score: 4%



Task

Schedule a pod as follows:

\* Name: nginx-kusc00401

\* Image: nginx

\* Node selector: disk=ssd

Solution:

#yaml

apiVersion: v1

kind: Pod

metadata:

name: nginx-kusc00401

spec:

containers:

&#8211; name: nginx

image: nginx

imagePullPolicy: IfNotPresent

nodeSelector:

disk: spinning

#

```
kubectl create -f node-select.yaml
```

**NO.23** Delete the pod without any delay (force delete)

```
Kubectl delete po &#8220;POD-NAME&#8221; &#8211;--grace-period=0 &#8211;--force
```

**NO.24** Create a snapshot of the etcd instance running at `127.0.0.1:2379`, saving the snapshot to the file path

```
/srv/data/etcd-snapshot.db.
```

The following TLS certificates/key are supplied for connecting to the server with etcdctl:

\* CA certificate: `/opt/KUCM00302/ca.crt`

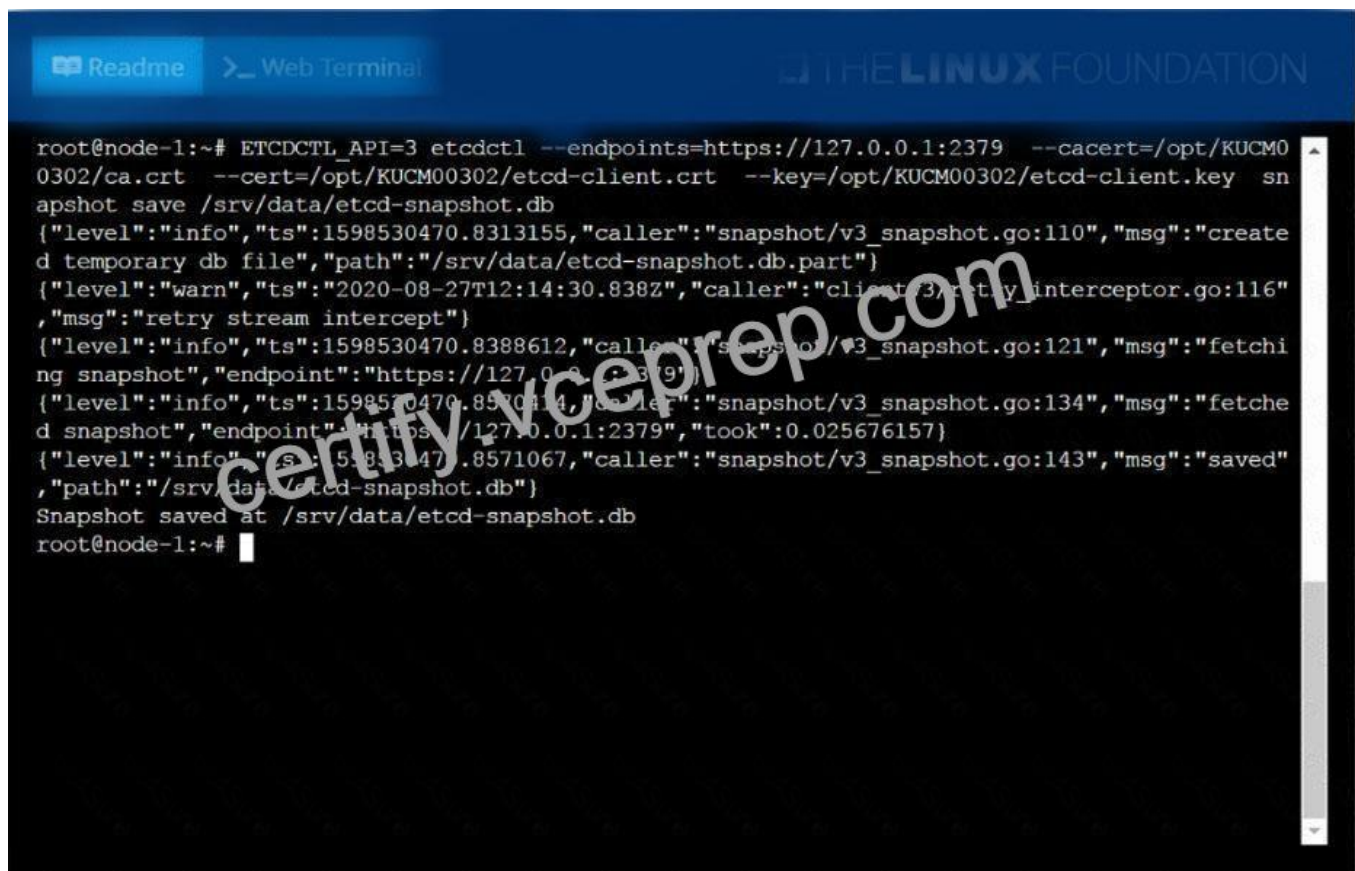
\* Client certificate: `/opt/KUCM00302/etcd-client.crt`

\* Client key: `/opt/KUCM00302/etcd-client.key`

See the solution below.

Explanation

solution



```
root@node-1:~# ETCDCTL_API=3 etcdctl --endpoints=https://127.0.0.1:2379 --cacert=/opt/KUCM00302/ca.crt --cert=/opt/KUCM00302/etcd-client.crt --key=/opt/KUCM00302/etcd-client.key snapshot save /srv/data/etcd-snapshot.db
{"level":"info","ts":1598530470.8313155,"caller":"snapshot/v3_snapshot.go:110","msg":"created temporary db file","path":"/srv/data/etcd-snapshot.db.part"}
{"level":"warn","ts":"2020-08-27T12:14:30.838Z","caller":"client/v3_retry_interceptor.go:116","msg":"retry stream intercept"}
{"level":"info","ts":1598530470.8388612,"caller":"snapshot/v3_snapshot.go:121","msg":"fetching snapshot","endpoint":"https://127.0.0.1:2379"}
{"level":"info","ts":1598530470.8570414,"caller":"snapshot/v3_snapshot.go:134","msg":"fetched snapshot","endpoint":"https://127.0.0.1:2379","took":0.025676157}
{"level":"info","ts":1598530470.8571067,"caller":"snapshot/v3_snapshot.go:143","msg":"saved","path":"/srv/data/etcd-snapshot.db"}
Snapshot saved at /srv/data/etcd-snapshot.db
root@node-1:~#
```

**NO.25** Change the Image version back to 1.17.1 for the pod you just updated and observe the changes  
kubectl set image pod/nginx nginx=nginx:1.17.1 kubectl describe po nginx kubectl get po nginx -w # watch it

**NO.26** Verify certificate expiry date for ca certificate in /etc/kubernetes/pki  
openssl x509 -in ca.crt -noout -text | grep -i validity -A 4

**NO.27** Create a namespace called 'development' and a pod with image nginx called nginx on this namespace.  
kubectl create namespace development kubectl run nginx --image=nginx --restart=Never -n development

**NO.28** Score:7%



## Context

An existing Pod needs to be integrated into the Kubernetes built-in logging architecture (e. g. kubectl logs).

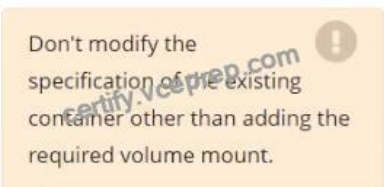
Adding a streaming sidecar container is a good and common way to accomplish this requirement.

## Task

Add a sidecar container named sidecar, using the busybox Image, to the existing Pod big-corp-app. The new sidecar container has to run the following command:

```
/bin/sh -c tail -n+1 -f /var/log/big-corp-app.log
```

Use a Volume, mounted at /var/log, to make the log file big-corp-app.log available to the sidecar container.



See the solution below.

## Explanation

Solution:

```
#
```

```
kubectl get pod big-corp-app -o yaml
```

```
#
```

```
apiVersion: v1
```

```
kind: Pod
```

```
metadata:
```

```
name: big-corp-app
```

```
spec:
```

```
containers:
```

```
&#8211; name: big-corp-app
```

```
image: busybox
```

```
args:
```

```
&#8211; /bin/sh
```

```
&#8211; -c
```

```
&#8211; >
```

```
i=0;
```

```
while true;
```

```
do
```

```
echo &#8220;$(date) INFO $i&#8221; >> /var/log/big-corp-app.log;
```

```
i=$((i+1));
```

```
sleep 1;
```

```
done
```

```
volumeMounts:
```

```
&#8211; name: logs
```

```
mountPath: /var/log
```

```
&#8211; name: count-log-1

image: busybox

args: [/bin/sh, -c, &#8216;tail -n+1 -f /var/log/big-corp-app.log&#8217;]

volumeMounts:

&#8211; name: logs

mountPath: /var/log

volumes:

&#8211; name: logs

emptyDir: {

}

#

kubectl logs big-corp-app -c count-log-1
```

**NO.29** Which one is the correct configuration?

- \* #Panorama
- \* &Panorama
- \* \$Panorama
- \* @Panorama

**NO.30** For this item, you will have to ssh to the nodes ik8s-master-0 and ik8s-node-0 and complete all tasks on these nodes. Ensure that you return to the base node (hostname: node-1) when you have completed this item.

Context

As an administrator of a small development team, you have been asked to set up a Kubernetes cluster to test the viability of a new application.

Task

You must use kubeadm to perform this task. Any kubeadm invocations will require the use of the

&#8211;ignore-preflight-errors=all option.

\* Configure the node ik8s-master-0 as a master node. .

\* Join the node ik8s-node-0 to the cluster.

See the solution below.

## Explanation

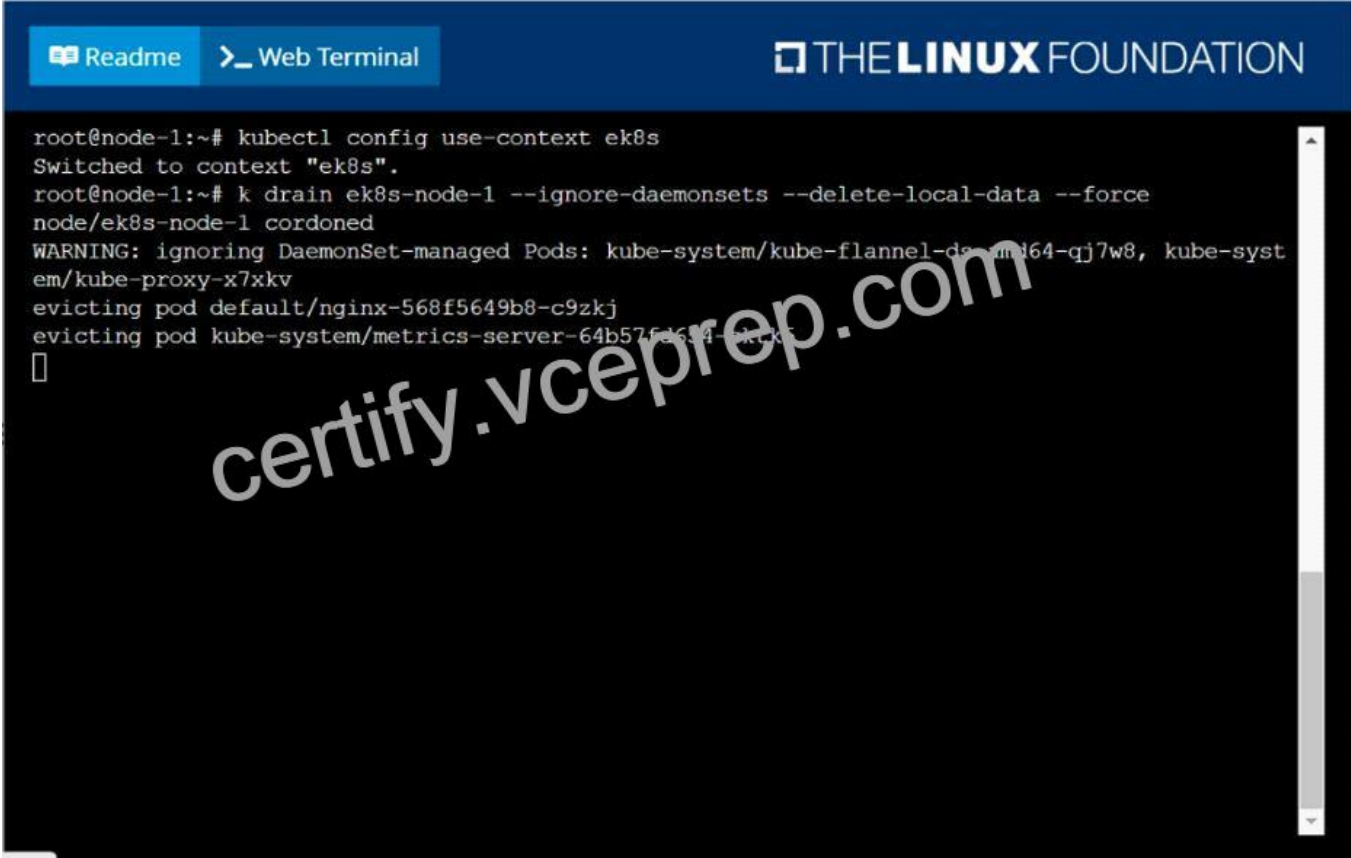
solution

You must use the kubeadm configuration file located at /etc/kubeadm.conf when initializing your cluster.

You may use any CNI plugin to complete this task, but if you don't have your favourite CNI plugin's manifest URL at hand, Calico is one popular option: <https://docs.projectcalico.org/v3.14/manifests/calico.yaml> Docker is already installed on both nodes and has been configured so that you can install the required tools.

**NO.31** Set the node named ek8s-node-1 as unavailable and reschedule all the pods running on it.

solution



The screenshot shows a terminal window with a dark background and white text. At the top, there are two buttons: 'Readme' and 'Web Terminal'. On the right side, the logo for 'THE LINUX FOUNDATION' is visible. The terminal output shows the following commands and their results:

```
root@node-1:~# kubectl config use-context ek8s
Switched to context "ek8s".
root@node-1:~# k drain ek8s-node-1 --ignore-daemonsets --delete-local-data --force
node/ek8s-node-1 cordoned
WARNING: ignoring DaemonSet-managed Pods: kube-system/kube-flannel-ds-amd64-qj7w8, kube-system/kube-proxy-x7xkv
evicting pod default/nginx-568f5649b8-c9zkj
evicting pod kube-system/metrics-server-64b57f2614-ekkk
[]
```

A large watermark 'certify.vceprep.com' is overlaid diagonally across the terminal output.

**NO.32** Check nodes which are ready and print it to a file /opt/nodestatus

```
* JSONPATH='{range .items[*]}{@.metadata.name}:{range
```

```
@.status.conditions[*]}{@.type}={@.status};{end}{end}&#8217;
```

```
&& kubectl get nodes -o jsonpath=&#8221;$JSONPATH&#8221; | grep
```

```
&#8220;Ready=True&#8221; > /opt/node-status
```

```
//Verify
```

```
cat /opt/node-status
* JSONPATH='{range .items[*]}{@.metadata.name}:{range
@.status.conditions[*]}{@.type}={@.status};{end}{end}&#8217;

//Verify
```

```
cat /opt/node-status
```

**NO.33** Create an nginx pod and list the pod with different levels of verbosity  
See the solution below.

Explanation

```
// create a pod
```

```
kubectl run nginx &#8211;image=nginx &#8211;restart=Never &#8211;port=80
```

```
// List the pod with different verbosity
```

```
kubectl get po nginx &#8211;v=7
```

```
kubectl get po nginx &#8211;v=8
```

```
kubectl get po nginx &#8211;v=9
```

**NO.34** Create a pod named kucc8 with a single app container for each of the

following images running inside (there may be between 1 and 4 images specified):

nginx + redis + memcached.

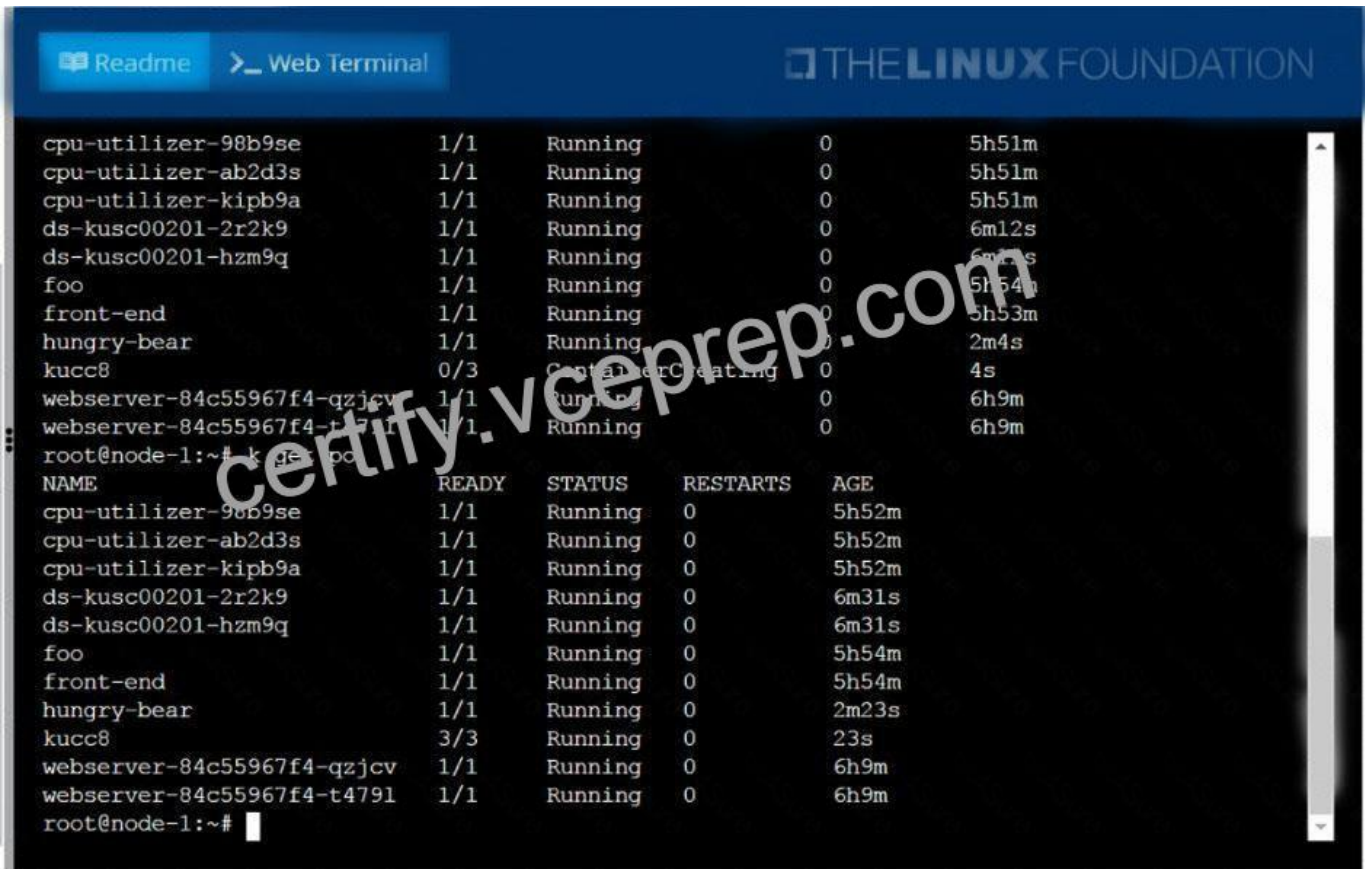
See the solution below.

Explanation

solution







```
Readme > Web Terminal THE LINUX FOUNDATION

cpu-utilizer-98b9se      1/1    Running    0        5h51m
cpu-utilizer-ab2d3s     1/1    Running    0        5h51m
cpu-utilizer-kipb9a     1/1    Running    0        5h51m
ds-kusc00201-2r2k9     1/1    Running    0        6m12s
ds-kusc00201-hzm9q     1/1    Running    0        6m12s
foo                     1/1    Running    0        5h54m
front-end               1/1    Running    0        5h53m
hungry-bear             1/1    Running    0        2m4s
kucc8                   0/3    ContainerCreating 0        4s
webserver-84c55967f4-qzjcv 1/1    Running    0        6h9m
webserver-84c55967f4-t479l 1/1    Running    0        6h9m
root@node-1:~# kubectl po
NAME                READY   STATUS    RESTARTS   AGE
cpu-utilizer-98b9se 1/1     Running   0          5h52m
cpu-utilizer-ab2d3s 1/1     Running   0          5h52m
cpu-utilizer-kipb9a 1/1     Running   0          5h52m
ds-kusc00201-2r2k9 1/1     Running   0          6m31s
ds-kusc00201-hzm9q 1/1     Running   0          6m31s
foo                 1/1     Running   0          5h54m
front-end           1/1     Running   0          5h54m
hungry-bear         1/1     Running   0          2m23s
kucc8               3/3     Running   0          23s
webserver-84c55967f4-qzjcv 1/1     Running   0          6h9m
webserver-84c55967f4-t479l 1/1     Running   0          6h9m
root@node-1:~#
```

**NO.35** Create a pod as follows:

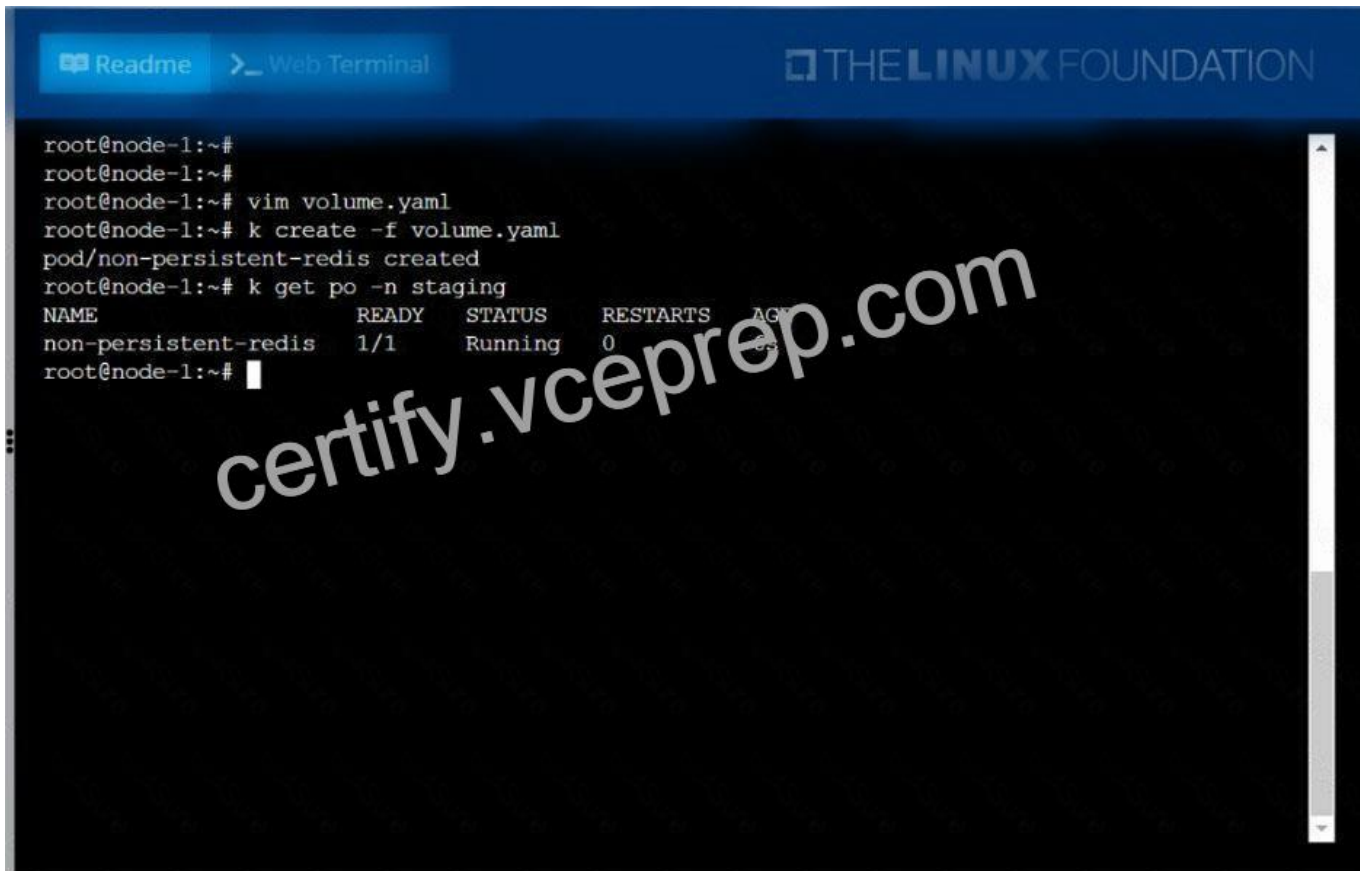
- \* Name:non-persistent-redis
- \* container Image:redis
- \* Volume with name:cache-control
- \* Mount path:/data/redis

The pod should launch in thestagingnamespace and the volumemust notbe persistent.  
See the solution below.

Explanation

solution





```
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           3s  
root@node-1:~#
```

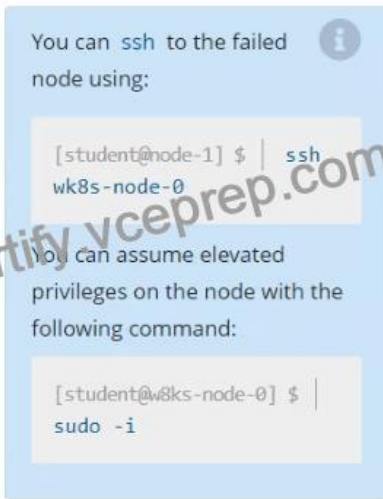
NO.36 Score: 13%



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

Task

A Kubernetes worker node, named wk8s-node-0 is in state NotReady. Investigate why this is the case, and perform any appropriate steps to bring the node to a Ready state, ensuring that any changes are made permanent.



See the solution below.

Explanation

Solution:

```
sudo -i
```

```
systemctl status kubelet
```

```
systemctl start kubelet
```

```
systemctl enable kubelet
```

**NO.37** Create a pod that having 3 containers in it? (Multi-Container)

See the solution below.

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:

&#8211; image: nginx

name: nginx-container

&#8211; image: redis

name: redis-container

&#8211; image: consul

name: consul-container

restartPolicy: Always

**NO.38** List the nginx pod with custom columns POD\_NAME and POD\_STATUS

See the solution below.

Explanation

kubectl get po -o=custom-columns=&#8221;POD\_NAME:.metadata.name,

POD\_STATUS:.status.containerStatuses[].state&#8221;

The CKA certification is recognized globally and is highly regarded by employers in the IT industry. Holding a CKA certification demonstrates an individual's ability to manage Kubernetes clusters effectively, which is a highly sought-after skill in today's job market. The CKA certification also provides individuals with access to a global community of certified professionals who can share knowledge and best practices in Kubernetes administration. The CKA exam is a valuable investment for individuals who want to enhance their career prospects and stay up-to-date with the latest technologies in the industry.

**CKA Practice Test Pdf Exam Material:** <https://www.vceprep.com/CKA-latest-vce-prep.html>