

# Linux-Foundation

## Exam Questions CKA

Certified Kubernetes Administrator (CKA) Program



### NEW QUESTION 1

Create a pod with environment variables as var1=value1. Check the environment variable in pod

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

```
kubectl run nginx --image=nginx --restart=Never --env=var1=value1
# then
kubectl exec -it nginx -- env
# or
kubectl exec -it nginx -- sh -c 'echo $var1'
# or
kubectl describe po nginx | grep value1
```

### NEW QUESTION 2

Create a deployment as follows:

- > Name:nginx-random
- > Exposed via a service nginx-random
- > Ensure that the service & pod are accessible via their respective DNS records
- > The container(s) within any pod(s) running as a part of this deployment should use the nginx image

Next, use the utility nslookup to lookup the DNS records of the service & pod and write the output to /opt/KUNW00601/service.dns and /opt/KUNW00601/pod.dns respectively.


- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:


Solution:

F:\Work\Data Entry Work\Data Entry\20200827\CKA\17 C.JPG



```
root@node-1:~#
root@node-1:~# k create deploy nginx-random --image=nginx
deployment.apps/nginx-random created
root@node-1:~# k expose deploy nginx-random --name=nginx-random --port=80 --target-port=80
service/nginx-random exposed
root@node-1:~# vim dns.yaml
```


F:\Work\Data Entry Work\Data Entry\20200827\CKA\17 D.JPG



```

apiVersion: v1
kind: Pod
metadata:
  name: busybox1
  labels:
    name: busybox
spec:
  containers:
  - image: busybox:1.28
    command:
    - sleep
    - "3600"
    name: busybox
  
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\17 E.JPG



```

root@node-1:~# k create deploy nginx-random --image=nginx
deployment.apps/nginx-random created
root@node-1:~# k expose deploy nginx-random --name=nginx-random --port=80 --target-port=80
service/nginx-random exposed
root@node-1:~# vim dns.yaml
root@node-1:~# k create -f dns.yaml
pod/busybox1 created
root@node-1:~# k get po -o wide | grep nginx-random
nginx-random-6d5766bbdc-ptzv2 1/1 Running 0 103s 10.244.2.16 k8s-node-1 <none> <none>
root@node-1:~# k exec -it busybox1 -- nslookup nginx-random
Server: 10.96.0.10
Address 1: 10.96.0.10 kube-dns.kube-system.svc.cluster.local

Name: nginx-random
Address 1: 10.111.37.132 nginx-random.default.svc.cluster.local
root@node-1:~# k exec -it busybox1 -- nslookup nginx-random > /opt/KUNW00601/service.dns
root@node-1:~# k exec -it busybox1 -- nslookup 10-244-2-16.default.pod
Server: 10.96.0.10
Address 1: 10.96.0.10 kube-dns.kube-system.svc.cluster.local

Name: 10-244-2-16.default.pod
Address 1: 10.244.2.16 10-244-2-16.nginx-random.default.svc.cluster.local
root@node-1:~# k exec -it busybox1 -- nslookup 10-244-2-16.default.pod > /opt/KUNW00601/pod.dns

```

**NEW QUESTION 3**

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

- A. Mastered
- B. Not Mastered

Answer: A

**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"

**NEW QUESTION 4**

Given a partially-functioning Kubernetes cluster, identify symptoms of failure on the cluster. Determine the node, the failing service, and take actions to bring up the failed service and restore the health of the cluster. Ensure that any changes are made permanently.

You can ssh to the relevant nodes (k8s-master-0 or k8s-node-0) using:

[student@node-1] \$ ssh <nodename>

You can assume elevated privileges on any node in the cluster with the following command:

[student@nodename] \$ | sudo ?Ci


- A. Mastered
- B. Not Mastered

Answer: A

**Explanation:**

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\23 C.JPG



The screenshot shows a terminal window with the following content:

```

root@node-1:~#
root@node-1:~# kubectl config use-context bk8s
Switched to context "bk8s".
root@node-1:~# ssh bk8s-master-0
Welcome to Ubuntu 16.04.6 LTS (GNU/Linux 4.4.0-1109-aws x86_64)

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

 * Are you ready for Kubernetes 1.19? It's nearly here! Try RC3 with
   sudo snap install microk8s --channel=1.19/candidate --classic
   https://microk8s.io/ has docs and details.


4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@bk8s-master-0:~$ sudo -i
root@bk8s-master-0:~# vim /var/lib/kubelet/config.yaml

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\23 D.JPG



The screenshot shows a terminal window displaying the configuration for a kubelet:

```

authorization:
  mode: Webhook
  webhook:
    cacheAuthorizedTTL: 0s
    cacheUnauthorizedTTL: 0s
clusterDNS:
- 10.96.0.10
clusterDomain: cluster.local
cpuManagerReconcilePeriod: 0s
evictionPressureTransitionPeriod: 0s
fileCheckFrequency: 0s
healthzBindAddress: 127.0.0.1
healthzPort: 10248
httpCheckFrequency: 0s
imageMinimumGCAge: 0s
kind: KubeletConfiguration
nodeStatusReportFrequency: 0s
nodeStatusUpdateFrequency: 0s
rotateCertificates: true
runtimeRequestTimeout: 0s
staticPodPath: /etc/kubernetes/manifests
streamingConnectionIdleTimeout: 0s
syncFrequency: 0s
volumeStatsAggPeriod: 0s
:WC

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\23 E.JPG

```

Readme > Web Terminal THE LINUX FOUNDATION

https://microk8s.io/ has docs and details.

4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@bk8s-master-0:~$ sudo -i
root@bk8s-master-0:~# vim /var/lib/kubelet/config.yaml
root@bk8s-master-0:~# systemctl restart kubelet
root@bk8s-master-0:~# systemctl enable kubelet
root@bk8s-master-0:~# kubectl get nodes

NAME          STATUS    ROLES    AGE   VERSION
bk8s-master-0 Ready    master   77d   v1.18.2
bk8s-node-0   Ready    <none>   77d   v1.18.2
root@bk8s-master-0:~#
root@bk8s-master-0:~# exit
logout
student@bk8s-master-0:~$ exit
logout
Connection to 10.250.4.77 closed.
root@node-1:~#

```

**NEW QUESTION 5**

Create a deployment as follows:

- > Name:nginx-app
- > Using containernginxwithversion 1.11.10-alpine
- > The deployment should contain3replicas

Next, deploy the application with newversion1.11.13-alpine, byperforming a rolling update.  
 Finally, rollback that update to theprevious version1.11.10-alpine.

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

solution  
 F:\Work\Data Entry Work\Data Entry\20200827\CKA\7 B.JPG


```

Readme > Web Terminal THE LINUX FOUNDATION

root@node-1:~# k create deploy nginx-app --image=nginx:1.11.10-alpine --dry-run=client -o y
aml > app.yaml
root@node-1:~# vim app.yaml

```


F:\Work\Data Entry Work\Data Entry\20200827\CKA\7 C.JPG



```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-app
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx-app
  template:
    metadata:
      labels:
        app: nginx-app
    spec:
      containers:
        - image: nginx:1.11.10-alpine
          name: nginx
'app.yaml
  
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\7 D.JPG



```

root@node-1:~# k create deploy nginx-app --image=nginx:1.11.10-alpine --dry-run=client -o y
aml > app.yaml
root@node-1:~# vim app.yaml
root@node-1:~# k create -f app.yaml
deployment.apps/nginx-app created
root@node-1:~#
root@node-1:~#
root@node-1:~# k set image deploy nginx-app nginx=nginx:1.11.13-alpine --record
deployment.apps/nginx-app image updated
root@node-1:~# k rollout undo deploy nginx-app
deployment.apps/nginx-app rolled back
root@node-1:~#
  
```

**NEW QUESTION 6**

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

kubect1 get pods -o wide  
 kubectl delete po ??nginx-dev??kubectl delete po ??nginx-prod??

**NEW QUESTION 7**

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**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 -nengineering -f ?C
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
```

### NEW QUESTION 8

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


- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\19 B.JPG



### NEW QUESTION 9

Ensure a single instance of podnginx is running on each node of the Kubernetes cluster where nginx also represents the Image name which has to be used. Do not override any taints currently in place.

Use DaemonSet to complete this task and use ds-kusc00201 as DaemonSet name.

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 B.JPG

```
Readme Web Terminal THE LINUX FOUNDATION  
root@node-1:~# vim ds.yaml  
:
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 C.JPG

```
Readme Web Terminal THE LINUX FOUNDATION  
apiVersion: apps/v1  
kind: DaemonSet  
metadata:  
  name: fluentd-elasticsearch  
  namespace: kube-system  
  labels:  
    k8s-app: fluentd-logging  
spec:  
  selector:  
    matchLabels:  
      name: fluentd-elasticsearch  
  template:  
    metadata:  
      labels:  
        name: fluentd-elasticsearch  
    spec:  
      tolerations:  
        # this toleration is to have the daemonset runnable on master nodes  
        # remove it if your masters can't run pods  
        - key: node-role.kubernetes.io/master  
          effect: NoSchedule  
      containers:  
        - name: nginx  
          image: nginx  
-- INSERT -- 17,19 All
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 D.JPG



```

Readme Web Terminal THE LINUX FOUNDATION

root@node-1:~#
root@node-1:~# kubectl config use-context wk8s
Switched to context "wk8s".
root@node-1:~# ssh wk8s-node-1
Welcome to Ubuntu 16.04.6 LTS (GNU/Linux 4.4.0-1109-aws x86_64)

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

 * Are you ready for Kubernetes 1.19? It's nearly here! Try RC3 with
   sudo snap install microk8s --channel=1.19/candidate --classic

   https://microk8s.io/ has docs and details.

4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@wk8s-node-1:~$ sudo -i
root@wk8s-node-1:~# vim /var/lib/kubelet/config.yaml

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\21 D.JPG

```

Readme Web Terminal THE LINUX FOUNDATION

clientCAfile: /etc/kubernetes/pki/ca.crt
authorization:
  mode: Webhook
  webhook:
    cacheAuthorizedTTL: 0s
    cacheUnauthorizedTTL: 0s
clusterDNS:
- 10.96.0.10
clusterDomain: cluster.local
cpuManagerReconcilePeriod: 0s
evictionPressureTransitionPeriod: 0s
fileCheckFrequency: 0s
healthzBindAddress: 127.0.0.1
healthzPort: 10248
httpCheckFrequency: 0s
imageMinimumGCAge: 0s
kind: KubeletConfiguration
nodeStatusReportFrequency: 0s
nodeStatusUpdateFrequency: 0s
rotateCertificates: true
runtimeRequestTimeout: 0s
staticPodPath: /etc/kubernetes/manifests
streamingConnectionIdleTimeout: 0s
syncFrequency: 0s
:WC

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\21 E.JPG



```

Readme Web Terminal THE LINUX FOUNDATION

https://microk8s.io/ has docs and details.

4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@wk8s-node-1:~$ sudo -i
root@wk8s-node-1:~# vim /var/lib/kubelet/config.yaml
root@wk8s-node-1:~# cd /etc/kubernetes/manifests
root@wk8s-node-1:/etc/kubernetes/manifests#
root@wk8s-node-1:/etc/kubernetes/manifests# vim pod.yaml
root@wk8s-node-1:/etc/kubernetes/manifests# systemctl restart kubelet
root@wk8s-node-1:/etc/kubernetes/manifests# systemctl enable kubelet
root@wk8s-node-1:/etc/kubernetes/manifests# exit
logout
student@wk8s-node-1:~$ exit
logout
Connection to 10.250.5.39 closed.
root@node-1:~# k get po
NAME                READY   STATUS    RESTARTS   AGE
webtool-wk8s-node-1 1/1     Running   0           11s
root@node-1:~#

```

**NEW QUESTION 10**

From the pod labelname=cpu-utilizer, find podsrunning high CPU workloads and write the name of the pod consumingmost CPU to thefile/opt/KUTR00102/KUTR00102.txt(which already exists).

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

solution  
 F:\Work\Data Entry Work\Data Entry\20200827\CKA\16 B.JPG


```

Readme Web Terminal THE LINUX FOUNDATION

root@node-1:~# k top po -l name=cpu-utilizer
NAME                CPU (cores)  MEMORY (bytes)
cpu-utilizer-98b9se 60m          7Mi
cpu-utilizer-ab2d3s 14m          7Mi
cpu-utilizer-kipb9a 45m          7Mi
root@node-1:~# vim /opt/KUTR00102/KUTR00102.txt

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\16 C.JPG



### NEW QUESTION 12

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 `--ignore-preflight-errors=alloption`.

- > Configure the node `ik8s-master-0` as a master node.
- > Join the node `ik8s-node-0` to the cluster.

- A. Mastered
- B. Not Mastered

Answer: A

#### 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.

### NEW QUESTION 15

Create a snapshot of the `etcd` instance running at `https://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`

- A. Mastered
- B. Not Mastered

Answer: A

#### Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\18 C.JPG

[Readme](#)
[Web Terminal](#)
THE LINUX FOUNDATION

```

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":"create temporary db file","path":"/srv/data/etcd-snapshot.db.part"}
{"level":"warn","ts":"2020-08-27T12:14:30.838Z","caller":"clientv3/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:~#
    
```

**NEW QUESTION 16**

List all the pods sorted by created timestamp

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 19**

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

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

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

**NEW QUESTION 22**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **CKA Practice Exam Features:**

- \* CKA Questions and Answers Updated Frequently
- \* CKA Practice Questions Verified by Expert Senior Certified Staff
- \* CKA Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* CKA Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The CKA Practice Test Here](#)**