



Linux-Foundation

Exam Questions CKA

Certified Kubernetes Administrator (CKA) Program

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate


Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1


Score: 7%




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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

SOLUTION:

```
[student@node-1] > ssh ek8s
kubectl cordon k8s-master
kubectl drain k8s-master --delete-local-data --ignore-daemonsets --force
apt-get install kubeadm=1.20.1-00 kubelet=1.20.1-00 kubectl=1.20.1-00 --disableexcludes=kubernetes kubeadm upgrade apply 1.20.1 --etcd-upgrade=false
systemctl daemon-reload systemctl restart kubelet kubectl uncordon k8s-master
```

NEW QUESTION 2


List all persistent volumes sorted by capacity, saving the full kubectl output to /opt/KUCC00102/volume_list. Use kubectl's own functionality for sorting the output, and do not manipulate it any further.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
 F:\Work\Data Entry Work\Data Entry\20200827\CKA\2 C.JPG



NEW QUESTION 3

Get list of all pods in all namespaces and write it to file "/opt/pods-list.yaml"

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

kubectl get po --all-namespaces > /opt/pods-list.yaml

NEW QUESTION 4

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.

Don't modify the specification of the existing container other than adding the required volume mount.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```
#
kubectl get pod big-corp-app -o yaml
#
apiVersion: v1 kind: Pod metadata:
name: big-corp-app spec:
containers:
- name: big-corp-app image: busybox
args:
- /bin/sh
- -c
- > i=0;
while true; do
echo "$(date) INFO $i" >> /var/log/big-corp-app.log; i=$((i+1));
sleep 1; done
volumeMounts:
- name: logs mountPath: /var/log
- name: count-log-1 image: busybox
args: [/bin/sh, -c, 'tail -n+1 -f /var/log/big-corp-app.log'] volumeMounts:
- name: logs mountPath: /var/log volumes:
- name: logs emptyDir: {
}
#
kubectl logs big-corp-app -c count-log-1
```

NEW QUESTION 5

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.

You can ssh to the failed node using:

```
[student@node-1] $ | ssh Wk8s-node-0
```

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

```
[student@w8ks-node-0] $ | sudo -i
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\20 C.JPG

```

Readme > Web Terminal THE LINUX FOUNDATION

root@node-1:~# kubectl config use-context wk8s
Switched to context "wk8s".
root@node-1:~# k get nodes
NAME             STATUS    ROLES    AGE   VERSION
wk8s-master-0   Ready     master   77d   v1.18.2
wk8s-node-0     NotReady  <none>   77d   v1.18.2
wk8s-node-1     Ready     <none>   77d   v1.18.2
root@node-1:~# ssh wk8s-node-0

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\20 D.JPG

```

Readme > Web Terminal THE LINUX FOUNDATION

wk8s-node-0     NotReady  <none>   77d   v1.18.2
wk8s-node-1     Ready     <none>   77d   v1.18.2
root@node-1:~# ssh wk8s-node-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@wk8s-node-0:~$ sudo -i
root@wk8s-node-0:~# systemctl restart kubelet
root@wk8s-node-0:~# systemctl enable kubelet

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\20 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-0:~$ sudo -i
root@wk8s-node-0:~# systemctl restart kubelet
root@wk8s-node-0:~# systemctl enable kubelet
Created symlink from /etc/systemd/system/multi-user.target.wants/kubelet.service to /lib/systemd/system/kubelet.service.
root@wk8s-node-0:~# exit
logout
student@wk8s-node-0:~$ exit
logout
Connection to 10.250.5.34 closed.
root@node-1:~# k get nodes
NAME                STATUS    ROLES    AGE   VERSION
wk8s-master-0      Ready    master   77d   v1.18.2
wk8s-node-0        Ready    <none>   77d   v1.18.2
wk8s-node-1        Ready    <none>   77d   v1.18.2
root@node-1:~# █

```

NEW QUESTION 6

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

- A. Mastered
- B. Not Mastered

Answer: A

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

```

NEW QUESTION 7

List all the pods sorted by name

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 8

Create a deployment spec file that will:

- > Launch 7 replicas of the nginx Image with the labelapp_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.


- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\10 B.JPG



The screenshot shows a web terminal interface with a dark blue header. On the left, there are two buttons: 'Readme' and 'Web Terminal'. On the right, the 'THE LINUX FOUNDATION' logo is displayed. The terminal window shows the following commands and output:

```
root@node-1:~# k create deploy kual00201 --image=nginx --dry-run=client -o yaml > /opt/KUAL
00201/spec_deployment.yaml
root@node-1:~# vim /opt/KUAL00201/spec_deployment.yaml
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\10 C.JPG

```
Readme Web Terminal THE LINUX FOUNDATION

apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app_runtime_stage: dev
  name: kual00201
spec:
  replicas: 7
  selector:
    matchLabels:
      app_runtime_stage: dev
  template:
    metadata:
      labels:
        app_runtime_stage: dev
    spec:
      containers:
      - image: nginx
        name: nginx
~
~
~
~
~
~
"/opt/KUAL00201/spec_deployment.yaml" 19L, 320C written
```

NEW QUESTION 9

Perform the following tasks:

- > Add an init container to hungry-bear (which has been defined in spec file /opt/KUCC00108/pod-spec-KUCC00108.yaml)
- > The init container should create an empty file named /workdir/calm.txt
- > If /workdir/calm.txt is not detected, the pod should exit
- > Once the spec file has been updated with the init container definition, the pod should be created

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 B.JPG

Readme Web Terminal THE **LINUX** FOUNDATION

```

root@node-1:~# vim ds.yaml
iroot@node-1:~# k create -f ds.yaml
daemonset.apps/ds-kusc00201 created
root@node-1:~# k get ds
NAME           DESIRED   CURRENT   READY   UP-TO-DATE   AVAILABLE   NODE SELECTOR   AGE
ds-kusc00201   2         2         2       2             2           <none>          4s
root@node-1:~# vim /opt/KUCC00108/pod-spec-KUCC00108.yaml

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 C.JPG

Readme Web Terminal THE **LINUX** FOUNDATION

```

apiVersion: v1
kind: Pod
metadata:
  name: hungry-bear
spec:
  volumes:
  - name: workdir
    emptyDir: {}
  containers:
  - name: checker
    image: alpine
    command: ["/bin/sh", "-c", "if [ -f /workdir/calm.txt ];
              then sleep 100000; else exit 1; fi"]
    volumeMounts:
    - name: workdir
      mountPath: /workdir
  initContainers:
  - name: create
    image: alpine
    command: ["/bin/sh", "-c", "touch /workdir/calm.txt"]
    volumeMounts:
    - name: workdir
      mountPath: /workdir
:W

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 D.JPG

```

Readme > Web Terminal THE LINUX FOUNDATION

root@node-1:~# vim ds.yaml
iroot@node-1:~# k create -f ds.yaml
daemonset.apps/ds-kusc00201 created
root@node-1:~# k get ds
NAME           DESIRED   CURRENT   READY   UP-TO-DATE   AVAILABLE   NODE SELECTOR   AGE
ds-kusc00201   2         2         2       2             2           <none>          4s
root@node-1:~# vim /opt/KUCC00108/pod-spec-KUCC00108.yaml
root@node-1:~# k create -f /opt/KUCC00108/pod-spec-KUCC00108.yaml
pod/hungry-bear created
root@node-1:~#

```

NEW QUESTION 10

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:


```

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

```

NEW QUESTION 10

Score: 13%



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.

You can ssh to the failed node using:

```
[student@node-1] $ | ssh
wk8s-node-0
```

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

```
[student@wk8s-node-0] $ |
sudo -i
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:
 sudo -i
 systemctl status kubelet systemctl start kubelet systemctl enable kubelet

NEW QUESTION 14

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=all option.

- > 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 apt has been configured so that you can install the required tools.

NEW QUESTION 15

List all the pods sorted by name

- A. Mastered
- B. Not Mastered


Answer: A

Explanation:

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

NEW QUESTION 19


Score: 7%



Task

Create a new nginx Ingress resource as follows:

- Name: ping
- Namespace: ing-internal
- Exposing service hi on path /hi using service port 5678



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```
vi ingress.yaml
#
apiVersion: networking.k8s.io/v1 kind: Ingress
metadata: name: ping
namespace: ing-internal spec:
rules:
- http:
paths:
- path: /hi pathType: Prefix backend: service:
name: hi port:
number: 5678
#
kubectl create -f ingress.yaml
```

NEW QUESTION 20

Create a busybox pod that runs the command “env” and save the output to “envpod” file

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
kubectl run busybox --image=busybox --restart=Never --rm -it -- env > envpod.yaml
```

NEW QUESTION 22

Create a busybox pod and add “sleep 3600” command

- A. Mastered
- B. Not Mastered


Answer: A

Explanation:

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

NEW QUESTION 27

Score: 4%



Task

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 .

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```
kubectl run kucc8 --image=nginx --dry-run -o yaml > kucc8.yaml
```

```
# vi kucc8.yaml apiVersion: v1 kind: Pod metadata:
```

```
creationTimestamp: null name: kucc8
```

```
spec: containers:
```

```
- image: nginx name: nginx
```

```
- image: redis name: redis
```

```
- image: memcached
```

```
name: memcached
```

```
- image: consul name: consul
```

```
#
```

```
kubectl create -f kucc8.yaml
```

```
#12.07
```

NEW QUESTION 28

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
```



```
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-c9z kj
evicting pod kube-system/metrics-server-64b57fd654-cktk5
█
```

NEW QUESTION 29

Ensure a single instance of pod nginx 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
i
```

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



```

apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: ds-kusc00201
spec:
  selector:
    matchLabels:
      name: fluentd-elasticsearch
  template:
    metadata:
      labels:
        name: fluentd-elasticsearch
    spec:
      containers:
      - name: nginx
        image: nginx
  
```

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



```

root@node-1:~# vim ds.yaml
root@node-1:~# k create -f ds.yaml
daemonset.apps/ds-kusc00201 created
root@node-1:~# k get ds

```

NAME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE	NODE SELECTOR	AGE
ds-kusc00201	2	2	2	2	2	<none>	4s

```

root@node-1:~#

```

NEW QUESTION 30

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 -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
```

NEW QUESTION 31

Score: 7%



Task

Create a new NetworkPolicy named allow-port-from-namespace in the existing namespace echo. Ensure that the new NetworkPolicy allows Pods in namespace my-app to connect to port 9000 of Pods in namespace echo.

Further ensure that the new NetworkPolicy:

- does not allow access to Pods, which don't listen on port 9000
- does not allow access from Pods, which are not in namespace my-app

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```
#network.yaml
apiVersion: networking.k8s.io/v1 kind: NetworkPolicy
metadata:
name: allow-port-from-namespace namespace: internal
spec: podSelector: matchLabels: {
}
policyTypes:
- Ingress ingress:
- from:
- podSelector: {
}
ports:
- protocol: TCP port: 8080
#spec.podSelector namespace pod kubectl create -f network.yaml
```

NEW QUESTION 35

Score: 4%



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:

- Deployment
- StatefulSet
- DaemonSet

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

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
```

NEW QUESTION 40

.....

Relate Links

100% Pass Your CKA Exam with Exambible Prep Materials

<https://www.exambible.com/CKA-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>