



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

Score: 5%



Task

Monitor the logs of pod bar and:

- Extract log lines corresponding to error file-not-found
- Write them to /opt/KUTR00101/bar

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

kubectl logs bar | grep 'unable-to-access-website' > /opt/KUTR00101/bar cat /opt/KUTR00101/bar

NEW QUESTION 3

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 4

Score: 7%



Task
 Reconfigure the existing deployment front-end and add a port specification named http exposing port 80/tcp of the existing container nginx.
 Create a new service named front-end-svc exposing the container port http.
 Configure the new service to also expose the individual Pods via a NodePort on the nodes on which they are scheduled.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:
 kubectl get deploy front-end
 kubectl edit deploy front-end -o yaml
 #port specification named http
 #service.yaml apiVersion: v1
 kind: Service metadata:
 name: front-end-svc labels:
 app: nginx spec: ports:
 - port: 80 protocol: tcp name: http selector: app: nginx
 type: NodePort
 # kubectl create -f service.yaml
 # kubectl get svc
 # port specification named http
 kubectl expose deployment front-end --name=front-end-svc --port=80 --target-port=80 --type=NodePort

NEW QUESTION 5

Create and configure the service front-end-service so it's accessible through NodePort and routes to the existing pod named front-end.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
 F:\Work\Data Entry Work\Data Entry\20200827\CKA\8 B.JPG

Readme
Web Terminal
THE LINUX FOUNDATION

```

root@node-1:~# k expose po
error: resource(s) were provided, but no name, label selector, or --all flag specified
See 'kubectl expose -h' for help and examples
root@node-1:~# k expose po fron-end --name=front-end-service --port=80 --target-port=80 --t
ype=NodePort
Error from server (NotFound): pods "fron-end" not found
root@node-1:~# k expose po front-end --name=front-end-service --port=80 --target-port=80 --
type=NodePort
service/front-end-service exposed
root@node-1:~# k get svc
NAME                TYPE                CLUSTER-IP      EXTERNAL-IP      PORT(S)          AGE
front-end-service   NodePort            10.103.221.227  <none>           80:31828/TCP     3s
kubernetes          ClusterIP           10.96.0.1       <none>           443/TCP          77d
root@node-1:~#
    
```

NEW QUESTION 6

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.



- 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 7

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 8

Get IP address of the pod – “nginx-dev”

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 9

Create a file:
/opt/KUCC00302/kucc00302.txt that lists all pods that implement service baz in namespace development.
The format of the file should be one pod name per line.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
F:\Work\Data Entry Work\Data Entry\20200827\CKA\11 B.JPG


```

Readme  Web Terminal  THE LINUX FOUNDATION

Name:          baz
Namespace:     development
Labels:        <none>
Annotations:   <none>
Selector:      name=foo
Type:          ClusterIP
IP:            10.104.252.175
Port:          <unset> 80/TCP
TargetPort:    9376/TCP
Endpoints:     10.244.1.5:9376,10.244.2.3:9376,10.244.2.6:9376
Session Affinity: None
Events:        <none>
root@node-1:~# k get po -l name=foo -n development
NAME                    READY   STATUS    RESTARTS   AGE
pod-kucc00302-847878    1/1     Running   0           6h35m
pod-kucc00302-983457    1/1     Running   0           6h35m
pod-kucc00302-985953    1/1     Running   0           6h35m
root@node-1:~# k get po -l name=foo -n development -o NAME
pod/pod-kucc00302-847878
pod/pod-kucc00302-983457
pod/pod-kucc00302-985953
root@node-1:~# k get po -l name=foo -n development -o NAME > /opt/KUCC00302/kucc00302.txt
root@node-1:~# vim /opt/KUCC00302/kucc00302.txt
root@node-1:~# vim /opt/KUCC00302/kucc00302.txt
root@node-1:~#

```

NEW QUESTION 10

List the nginx pod with custom columns POD_NAME and POD_STATUS

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

kubectl get po -o=custom-columns="POD_NAME:.metadata.name, POD_STATUS:.status.containerStatuses[.state]"

NEW QUESTION 10

Score: 4%



Task

Schedule a pod as follows:

- Name: nginx-kusc00401
- Image: nginx
- Node selector: disk=ssd

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```

#yaml apiVersion: v1 kind: Pod metadata:
name: nginx-kusc00401 spec:
containers:
- name: nginx image: nginx
imagePullPolicy: IfNotPresent nodeSelector:

```

```
disk: spinning
#
kubectl create -f node-select.yaml
```

NEW QUESTION 15

Get list of all the pods showing name and namespace with a jsonpath expression.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
kubectl get pods -o=jsonpath="{.items[*]['metadata.name'
, 'metadata.namespace']}"
```

NEW QUESTION 18

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 22

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.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:
sudo -i

systemctl status kubelet systemctl start kubelet systemctl enable kubelet

NEW QUESTION 25

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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```
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
```

NEW QUESTION 27


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
F:\Work\Data Entry Work\Data Entry\20200827\CKA\5 B.JPG




F:\Work\Data Entry Work\Data Entry\20200827\CKA\5 C.JPG

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
 F:\Work\Data Entry Work\Data Entry\20200827\CKA\14 B.JPG



NEW QUESTION 36

Check the Image version of nginx-dev pod using jsonpath

- A. Mastered
- B. Not Mastered


Answer: A

Explanation:

kubect1 get po nginx-dev -o jsonpath='{.spec.containers[].image}'

NEW QUESTION 37

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

The availability of service hi 
 can be checked using the
 following command, which
 should return hi :

```
[student@node-1] $ | curl
-kL <INTERNAL_IP>/hi
```

- 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 40

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 45

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 47

Create a Kubernetes secret as follows:

- > Name: super-secret
- > password: bob

Create a pod named pod-secrets-via-file, using the redis Image, which mounts a secret named super-secret at /secrets.
 Create a second pod named pod-secrets-via-env, using the redis Image, which exports password as
 CONFIDENTIAL

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
 F:\Work\Data Entry Work\Data Entry\20200827\CKA\12 B.JPG

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

```

root@node-1:~# k create -f secret.yaml
pod/pod-secrets-via-file created
root@node-1:~# vim secret1.yaml
root@node-1:~# k create -f secret1.yaml
pod/pod-secrets-via-env created
root@node-1:~# k get po
NAME                                READY   STATUS    RESTARTS   AGE
cpu-utilizer-98b9se                 1/1     Running   0           6h25m
cpu-utilizer-ab2d3s                 1/1     Running   0           6h25m
cpu-utilizer-kipb9a                 1/1     Running   0           6h25m
ds-kusc00201-2r2k9                  1/1     Running   0           40m
ds-kusc00201-hzm9q                  1/1     Running   0           40m
foo                                  1/1     Running   0           6h28m
front-end                           1/1     Running   0           6h27m
hungry-bear                         1/1     Running   0           36m
kucc8                                3/3     Running   0           34m
nginx-app-848cfcf495-9prjh          1/1     Running   0           19m
nginx-app-848cfcf495-g12kh          1/1     Running   0           19m
nginx-app-848cfcf495-pg2c8          1/1     Running   0           19m
nginx-kusc00101                     1/1     Running   0           26m
pod-secrets-via-env                 1/1     Running   0           4s
pod-secrets-via-file                1/1     Running   0           106s
webserver-84c55967f4-qzjcv          1/1     Running   0           6h43m
webserver-84c55967f4-t4791         1/1     Running   0           6h43m
root@node-1:~#

```

NEW QUESTION 49

Score:7%



Task

Create a new PersistentVolumeClaim

- Name: pv-volume
- Class: csi-hostpath-sc
- Capacity: 10Mi

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

- Name: web-server
- Image: nginx
- 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.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```


vi pvc.yaml storageclass pvc apiVersion: v1
kind: PersistentVolumeClaim metadata:
name: pv-volume spec: accessModes:
- ReadWriteOnce volumeMode: Filesystem resources:
requests: storage: 10Mi

```

```
storageClassName: csi-hostpath-sc
# vi pod-pvc.yaml apiVersion: v1 kind: Pod metadata:
name: web-server spec:
containers:
- name: web-server image: nginx volumeMounts:
- mountPath: "/usr/share/nginx/html"
name: my-volume volumes:
- name: my-volume persistentVolumeClaim: claimName: pv-volume
# craete
kubectl create -f pod-pvc.yaml
#edit
kubectl edit pvc pv-volume --record
```

NEW QUESTION 53

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 56

Configure the kubelet systemd- managed service, on the node labelled with name=wk8s-node-1, to launch a pod containing a single container of Image httpd named webtool automatically. Any spec files required should be placed in the /etc/kubernetes/manifests directory on the node.

You can ssh to the appropriate node using:

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

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

```
[student@wk8s-node-1] $ | sudo -i
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\21 C.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
:WG

```

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 60

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 62

.....

Relate Links

100% Pass Your CKA Exam with Examible 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/>