

CKA Dumps

Certified Kubernetes Administrator (CKA) Program

<https://www.certleader.com/CKA-dumps.html>



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.

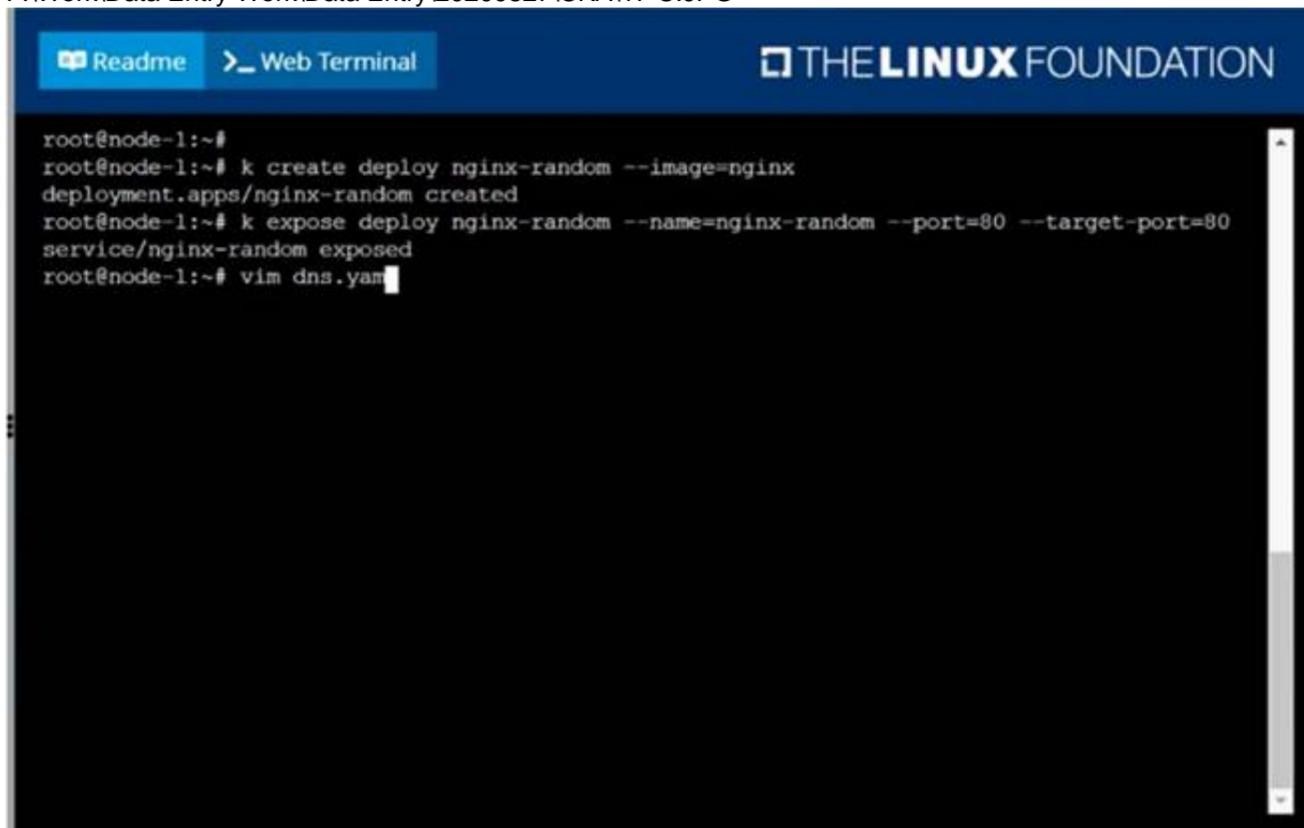
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

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



The screenshot shows a terminal window with a dark background and white text. At the top, there is a blue header bar with 'Readme' and 'Web Terminal' buttons on the left, and 'THE LINUX FOUNDATION' logo on the right. The terminal content shows the following commands and output:

```
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: PersistentVolume
metadata:
  name: app-data
spec:
  capacity:
    storage: 2Gi
  accessModes:
    - ReadWriteMany
  hostPath:
    path: /srv/app-data
  storageClassName: share
```

"app-data.yaml" 12L, 194C

* 2. Save the file and create the persistent volume. Image for post

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl create -f pv.yaml
persistentvolume/pv created
```

* 3. View the persistent volume.

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl get pv
```

| NAME | CAPACITY | ACCESS MODES | RECLAIM POLICY | STATUS | CLAIM | STORAGECLASS | REASON | AGE |
|----------|----------|--------------|----------------|-----------|-------|--------------|--------|-----|
| app-data | 2Gi | RWX | Retain | Available | | shared | | 31s |

> Our persistent volume status is available meaning it is available and it has not been mounted yet. This status will change when we mount the persistentVolume to a persistentVolumeClaim.

PersistentVolumeClaim

In a real ecosystem, a system admin will create the PersistentVolume then a developer will create a PersistentVolumeClaim which will be referenced in a pod. A PersistentVolumeClaim is created by specifying the minimum size and the access mode they require from the persistentVolume.

Challenge

> Create a Persistent Volume Claim that requests the Persistent Volume we had created above. The claim should request 2Gi. Ensure that the Persistent Volume Claim has the same storageClassName as the persistentVolume you had previously created.

kind: PersistentVolumeClaim

apiVersion: v1

metadata: name: app-data

spec: accessModes: - ReadWriteMany

resources: requests: storage: 2Gi storageClassName: shared

* 2. Save and create the pvc

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl create -f app-data.yaml persistentvolumeclaim/app-data created
```

* 3. View the pvc Image for post

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl get pvc
```

| NAME | STATUS | VOLUME | CAPACITY | ACCESS MODES | STORAGECLASS |
|------|--------|--------|----------|--------------|--------------|
| pv | Bound | pv | 512m | RWX | shared |

* 4. Let's see what has changed in the pv we had initially created.

Image for post

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl get pv
```

| NAME | CAPACITY | ACCESS MODES | RECLAIM POLICY | STATUS | CLAIM | STORAGECLASS | REASON | AGE |
|------|----------|--------------|----------------|--------|------------|--------------|--------|-----|
| pv | 512m | RWX | Retain | Bound | default/pv | shared | | 16m |

Our status has now changed from available to bound.

* 5. Create a new pod named myapp with image nginx that will be used to Mount the Persistent Volume Claim with the path /var/app/config.

Mounting a Claim

apiVersion: v1

kind: Pod

metadata: creationTimestamp: null name: app-data

spec: volumes: - name: config persistentVolumeClaim: claimName: app-data

containers: - image: nginx name: app volumeMounts: - mountPath: "/srv/app-data" name: config

NEW QUESTION 4

Check the image version in pod without the describe command

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 5

Perform the following tasks:

- > Add an init container tohungry-bear(which has beendefined in spec file /opt/KUCC00108/pod-spec-KUCC00108.yaml)
- > The init container should createan empty file named/workdir/calm.txt
- > If/workdir/calm.txtis notdetected, the pod should exit
- > Once the spec file has beenupdatedwith the init containerdefinition, the pod should becreated

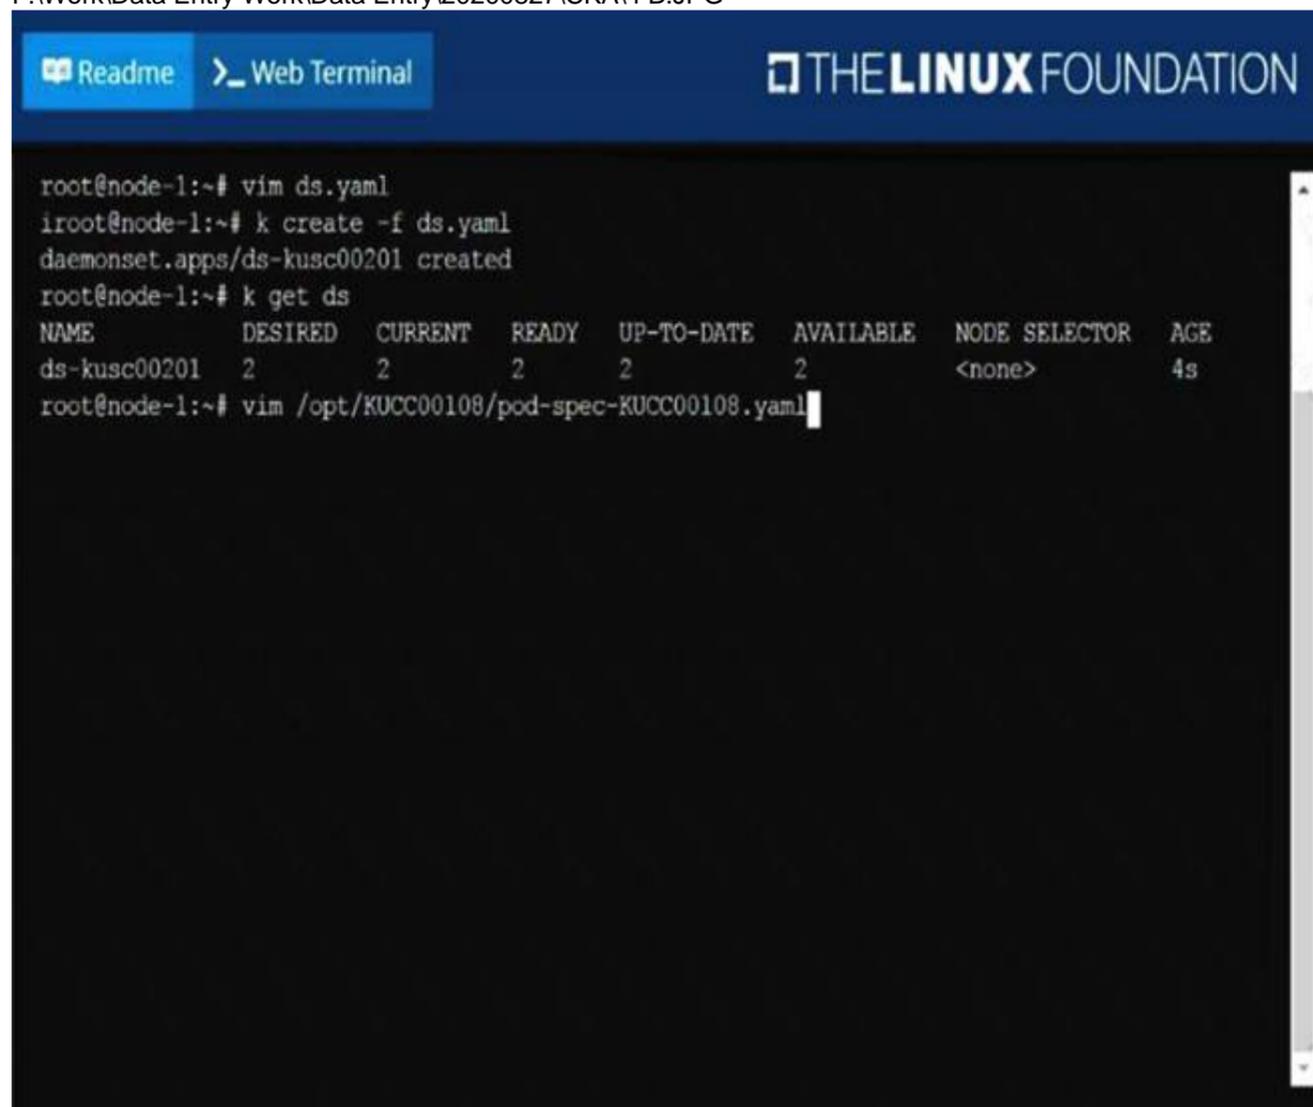
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

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



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

```

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
:wc

```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 D.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:~# 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 6

Check the Image version of nginx-dev pod using jsonpath

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 7

List all the pods sorted by name

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 8

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 9

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

`kubect 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 10

Check to see how many worker nodes are ready (not including nodes `taintedNoSchedule`) and write the number to `/opt/KUCC00104/kucc00104.txt`.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\15 B.JPG

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your CKA Exam with Our Prep Materials Via below:

<https://www.certleader.com/CKA-dumps.html>