

CKA Dumps

Certified Kubernetes Administrator (CKA) Program

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



NEW QUESTION 1

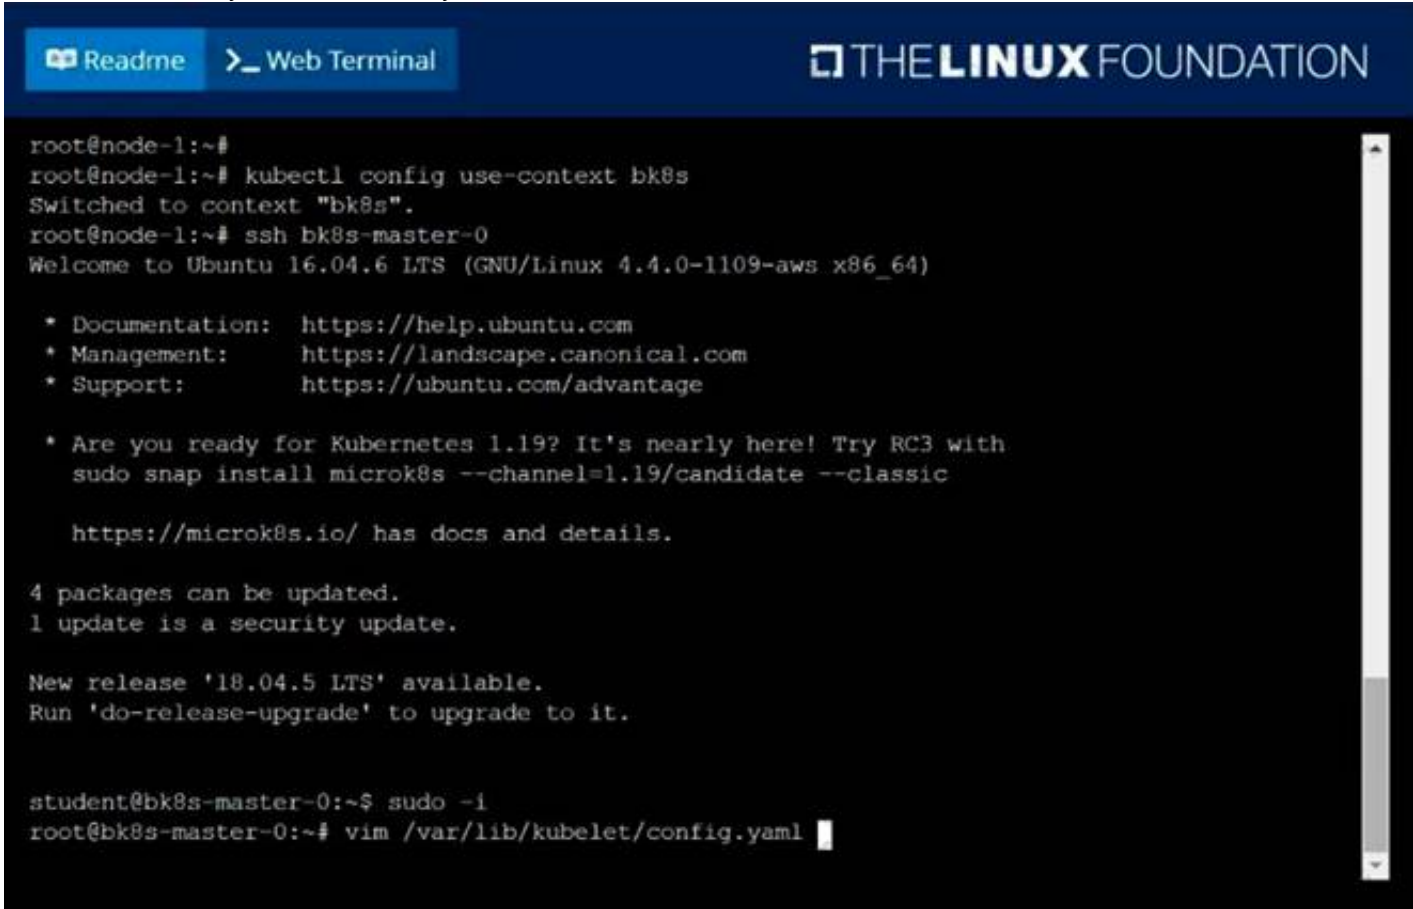
Given a partially-functioningKubernetes cluster, identifysymptoms of failure on the cluster.
Determine the node, the failingservice, and take actions to bring upthe failed service and restore thehealth of the cluster. Ensure that anychanges are made permanently.
You canssh to the relevant Inodes (bk8s-master-0orbk8s-node-0) using:
[student@node-1] \$ ssh<nodename>
You can assume elevatedprivileges on any node in thecluster with the followingcommand:
[student@nodename] \$ | sudo ?Ci

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

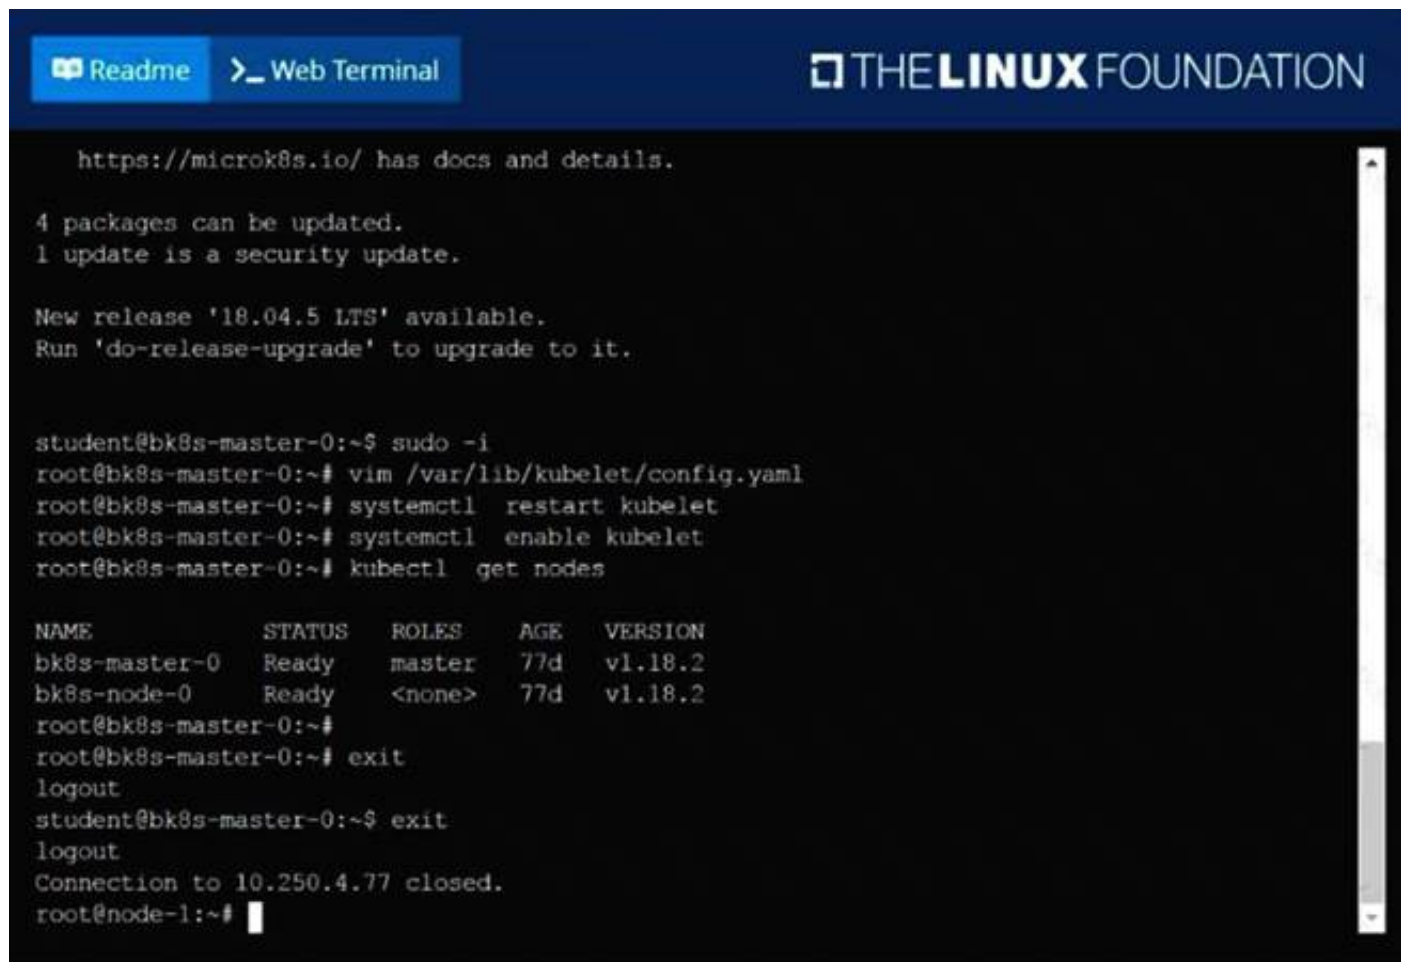
solution
F:\Work\Data Entry Work\Data Entry\20200827\CKA\23 C.JPG



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



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



The screenshot shows a terminal window with a dark background. At the top, there are two buttons: 'Readme' and 'Web Terminal', and the 'THE LINUX FOUNDATION' logo. The terminal output shows a user logging into a master node and performing several commands to set up Kubernetes. It includes checking for updates, upgrading to the latest LTS version, restarting and enabling kubelet, and finally checking the status of the nodes. A table is displayed showing the status of the master and node.

```
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:~# kubectll get nodes

NAME             STATUS    ROLES    AGE   VERSION
bk8s-master-0    Ready     master   77d   vl.18.2
bk8s-node-0      Ready     <none>   77d   vl.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 2

Create a pod as follows:

- > Name:mongo
- > Using Image:mongo
- > In anew Kubernetes namespacenamed:my-website

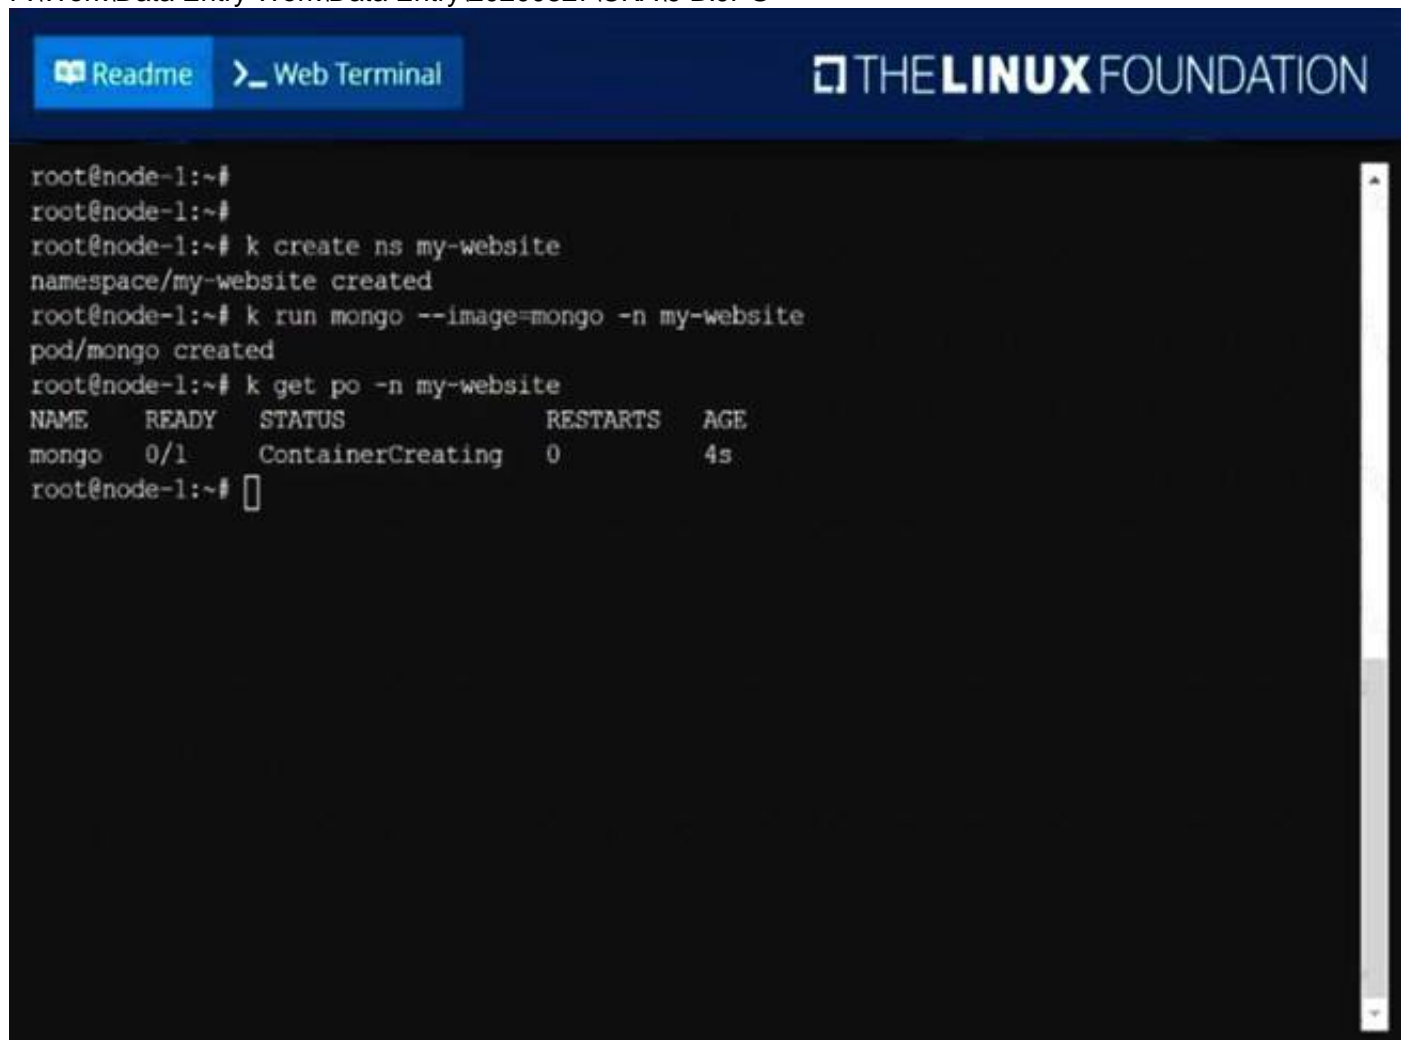
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\9 B.JPG



The screenshot shows a terminal window with a dark background. At the top, there are two buttons: 'Readme' and 'Web Terminal', and the 'THE LINUX FOUNDATION' logo. The terminal output shows a user logging into a node and performing several commands to create a namespace and a pod. It includes creating the 'my-website' namespace, creating a pod named 'mongo' using the 'mongo' image, and checking the status of the pod. A table is displayed showing the status of the pod.

```
root@node-1:~#
root@node-1:~#
root@node-1:~# k create ns my-website
namespace/my-website created
root@node-1:~# k run mongo --image=mongo -n my-website
pod/mongo created
root@node-1:~# k get po -n my-website

NAME    READY   STATUS             RESTARTS   AGE
mongo   0/1     ContainerCreating   0           4s
root@node-1:~#
```

NEW QUESTION 3

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 4

Create a persistent volume with nameapp-data, of capacity2Giandaccess modeReadWriteMany. Thetype of volume ishostPathand itslocation is/srv/app-data.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution
Persistent Volume
A persistent volume is a piece of storage in aKubernetes cluster. PersistentVolumes are a cluster-level resource like nodes, which don't belong to any namespace. It is provisioned by the administrator and has a particular file size. This way, a developer deploying their app on Kubernetes need not knowthe underlying infrastructure. When the developer needs a certain amount of persistent storage for their application, the system administrator configures the cluster so that they consume the PersistentVolume provisioned in an easy way.
Creating PersistentVolume
kind: PersistentVolumeapiVersion: v1metadata:name:app-dataspec:capacity: # defines the capacity of PV we are creatingstorage:2Gi#the amount of storage we are trying to claimaccessModes: # defines the rights of the volumewe are creating-ReadWriteManyhostPath:path: "/srv/app-data" # path to which we are creating the volume
Challenge
> Create a Persistent Volume namedapp-data, with access modeReadWriteMany, storage classname shared,2Giof storage capacity and the host path/srv/app-data.

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: app-data
spec:
  capacity:
    storage: 2Gi
  accessModes:
    - ReadWriteMany
  hostPath:
    path: /srv/app-data
  storageClassName: shared
```

"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 willchange 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. Ensurethat the Persistent Volume Claim has the same storageClassName as the persistentVolume you had previously created.

kind: PersistentVolumeapiVersion: v1metadata:name:app-data spec:
accessModes:-ReadWriteManyresources:
requests:storage:2Gi storageClassName:shared

* 2. Save and create the pvc

njerry191@cloudshell:~(extreme-clone-2654111)\$ kubectl create -f app-data.yaml persistentvolumeclaim/app-data created
* 3. View the pvc Image for post

```
njerry191@cloudshell:~ (extreme-clone-2654111)$ 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-2654111)$ 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-pvc
    persistentVolumeClaim:
      claimName: app-data
containers:
- image: nginx
  name: app
  volumeMounts:
  - mountPath: /srv/app-data
    name: config-pvc
```

NEW QUESTION 5

List all the pods sorted by name

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 6

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 7

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 8

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

ReadmeWeb Terminal

THELINUX 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

ReadmeWeb Terminal

THELINUX 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
:~
```

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

ReadmeWeb 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 9

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 10

Create and configure the servicefront-end-serviceso it's accessiblethroughNodePortand routes to theexisting pod namedfront-end.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

ReadmeWeb 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 10

Get IP address of the pod ?C ??nginx-dev??

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Kubect1 get po -o wide Using JsonPath

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

NEW QUESTION 11

Scale the deploymentwebserverto6pods.

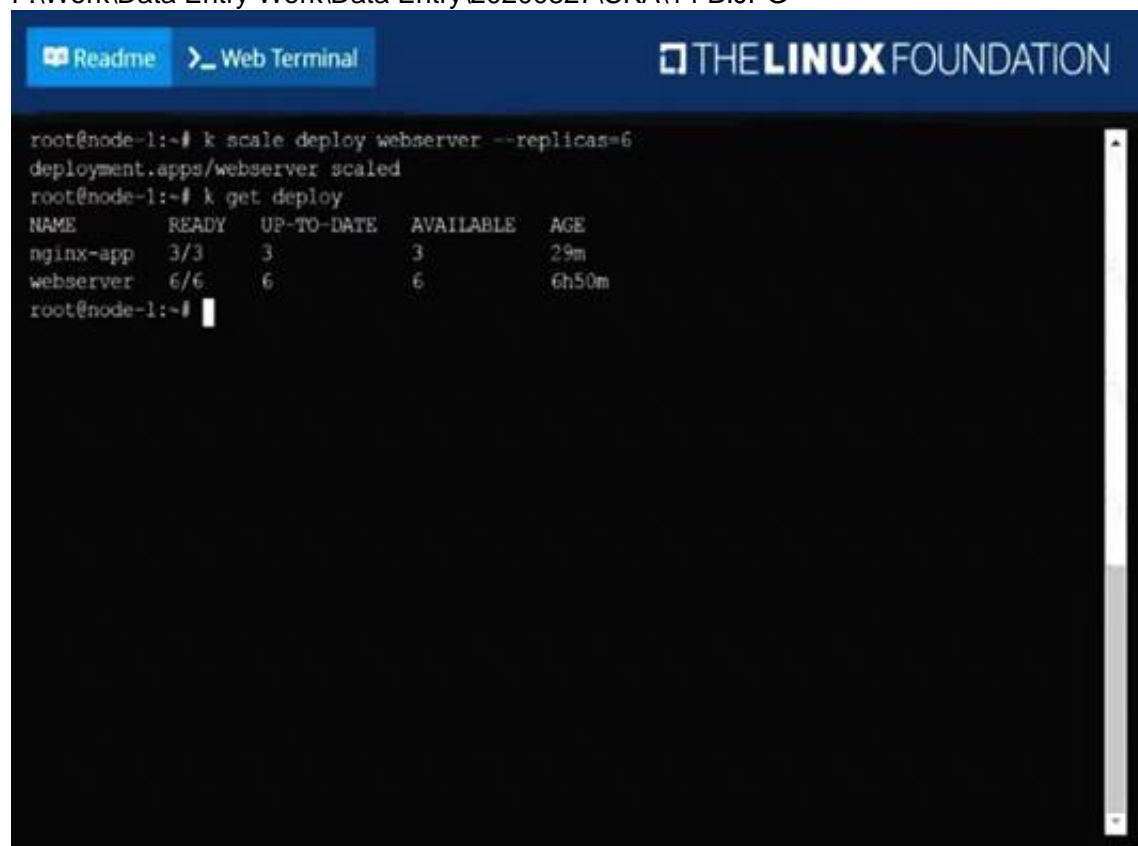
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

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



The screenshot shows a terminal window titled 'Web Terminal' with 'THE LINUX FOUNDATION' logo. The terminal output is as follows:

```
root@node-1:~# k scale deploy webserver --replicas=6
deployment.apps/webserver scaled
root@node-1:~# k get deploy
NAME       READY   UP-TO-DATE   AVAILABLE   AGE
nginx-app  3/3     3            3           29m
webserver  6/6     6            6           6h50m
root@node-1:~#
```

NEW QUESTION 15

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 17

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 19

Schedule a pod as follows:

- > Name: nginx-kusc00101
- > Image: nginx
- > Node selector: disk=ssd


- A. Mastered
B. Not Mastered

Answer: A

Explanation:

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\6 B.JPG



F:\Work\Data Entry Work\Data Entry\20200827\CKA\6 C.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>