

Linux-Foundation

Exam Questions CKA

Certified Kubernetes Administrator (CKA) Program



NEW QUESTION 1

CORRECT TEXT

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

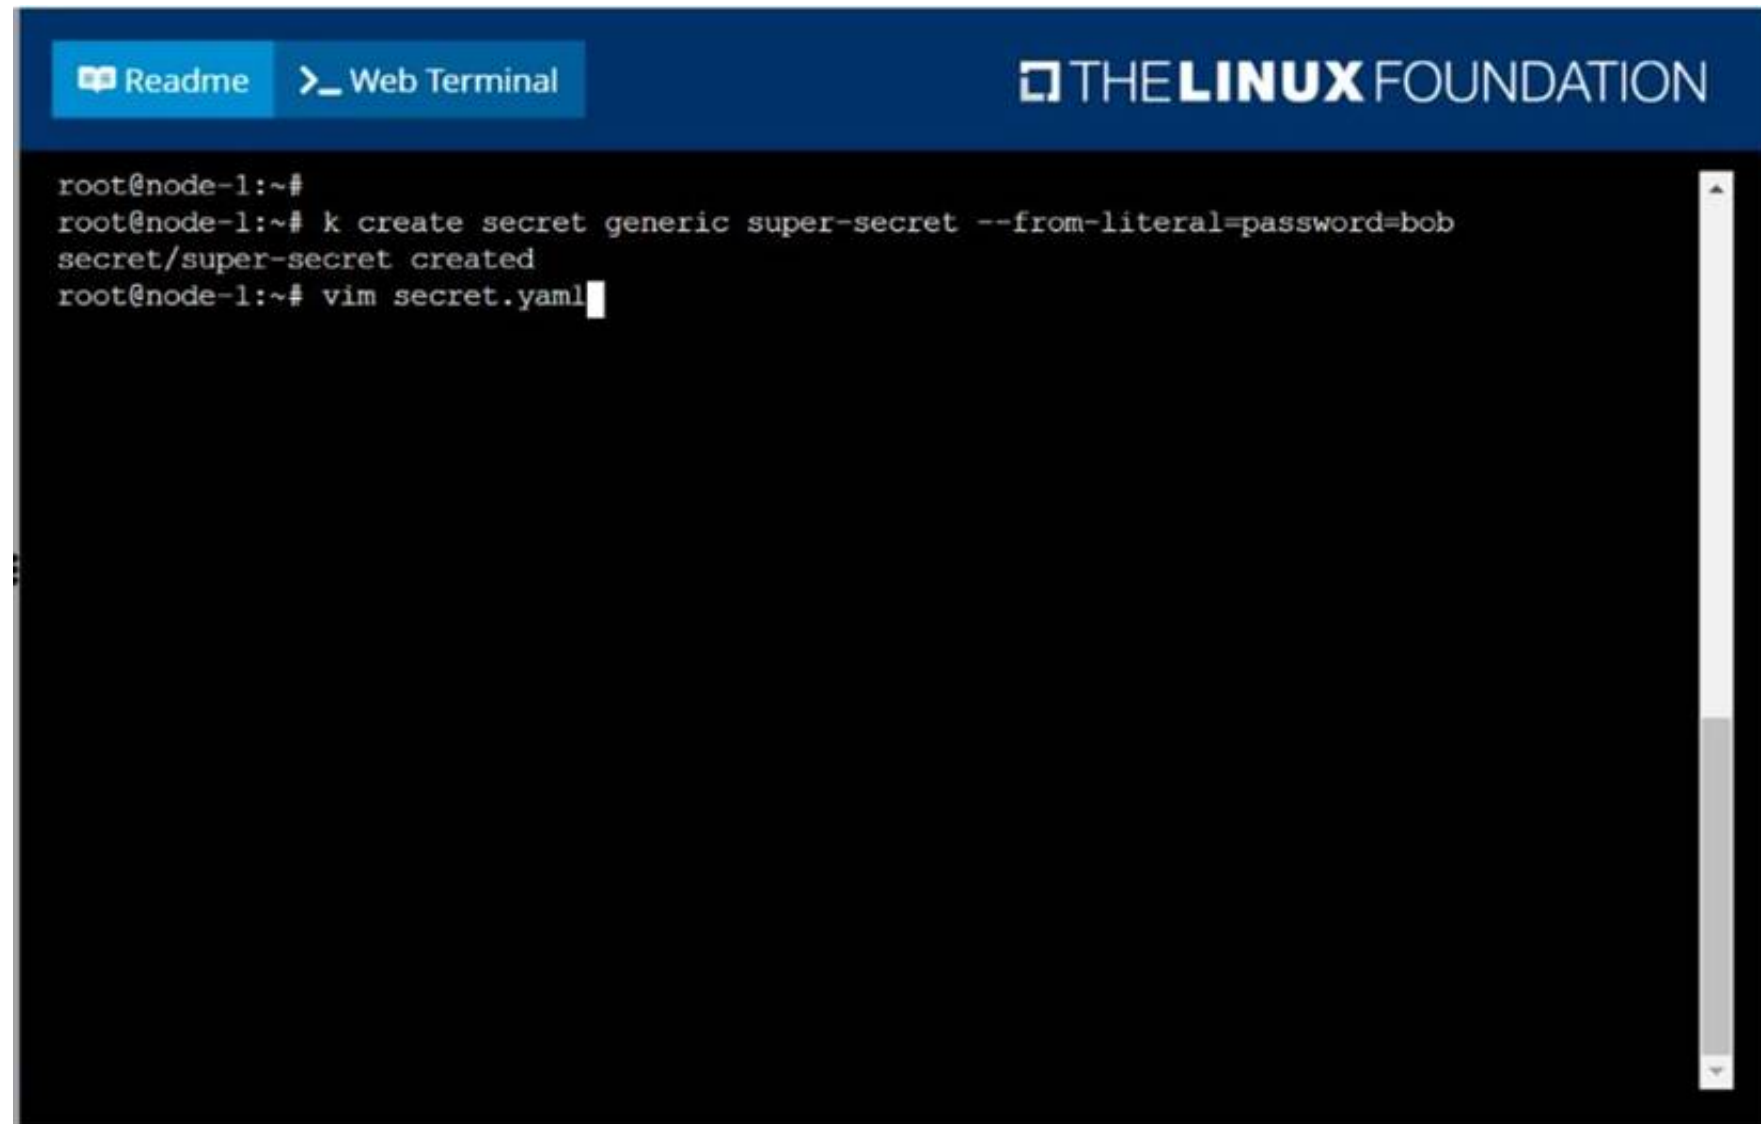
A. Mastered

B. Not Mastered

Answer: A


Explanation:

solution




```
root@node-1:~#  
root@node-1:~# k create secret generic super-secret --from-literal=password=bob  
secret/super-secret created  
root@node-1:~# vim secret.yaml
```

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



F:\Work\Data Entry Work\Data Entry\20200827\CKA\12 C.JPG



F:\Work\Data Entry Work\Data Entry\20200827\CKA\12 D.JPG

NEW QUESTION 2

CORRECT TEXT

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 3

CORRECT TEXT

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 4

CORRECT TEXT

Score: 5%



Task

From the pod label name=cpu-utilizer, find pods running high CPU workloads and write the name of the pod consuming most CPU to the file /opt/KUTR00401/KUTR00401.txt (which already exists).

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:
kubectl top -l name=cpu-user -A
echo 'pod name' >> /opt/KUT00401/KUT00401.txt

NEW QUESTION 5

CORRECT TEXT

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 6

CORRECT TEXT

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

Readme
Web Terminal



```

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-c9zkj
evicting pod kube-system/metrics-server-64b57fd654-cktk5

```

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

NEW QUESTION 7

CORRECT TEXT


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 8

CORRECT TEXT

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:

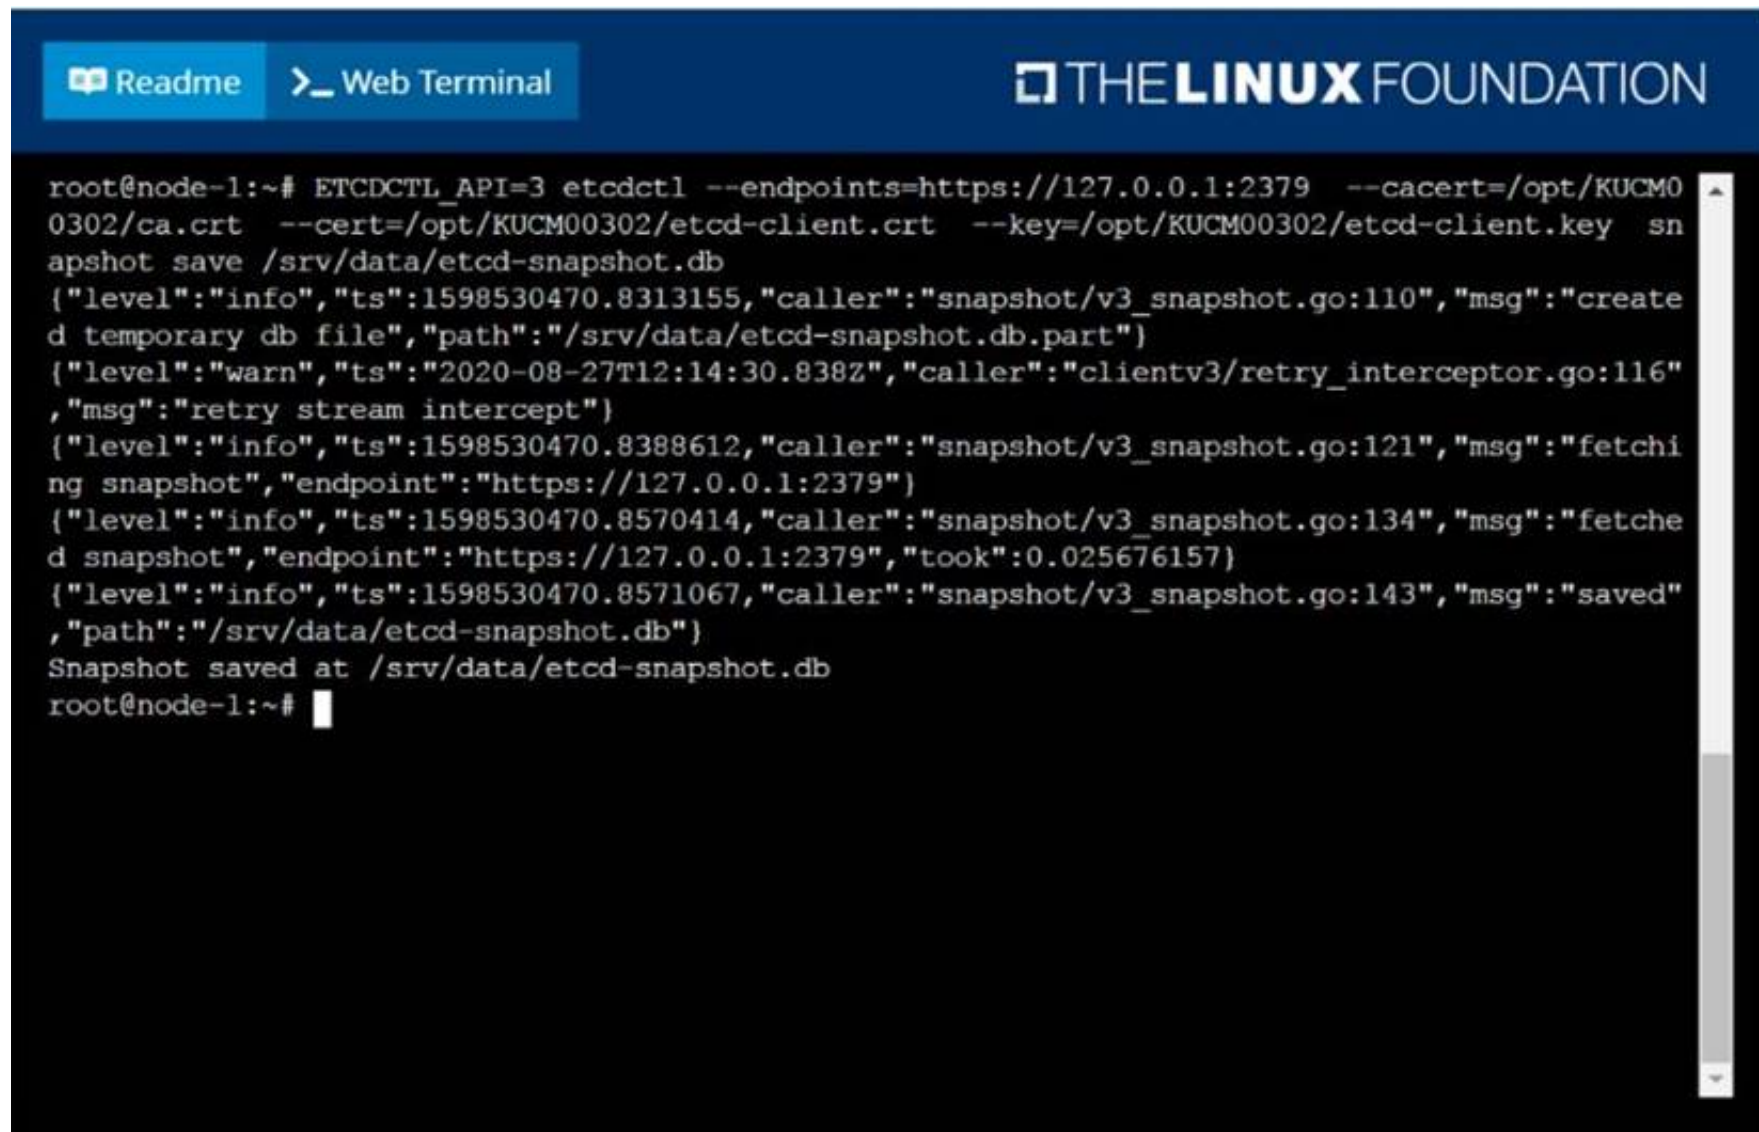
- ? CA certificate: /opt/KUCM00302/ca.crt
- ? Client certificate: /opt/KUCM00302/etcd-client.crt
- ? Client key: /opt/KUCM00302/etcd-client.key

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution



```
root@node-1:~# ETCDCCTL_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 d 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:~#
```

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

NEW QUESTION 9

CORRECT TEXT

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}'{"\n"}
```

NEW QUESTION 10

CORRECT TEXT

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 I nodes (bk8s-master-0 or bk8s-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 -i
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

solution

Readme
Web Terminal
THE **LINUX** FOUNDATION

```

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 C.JPG

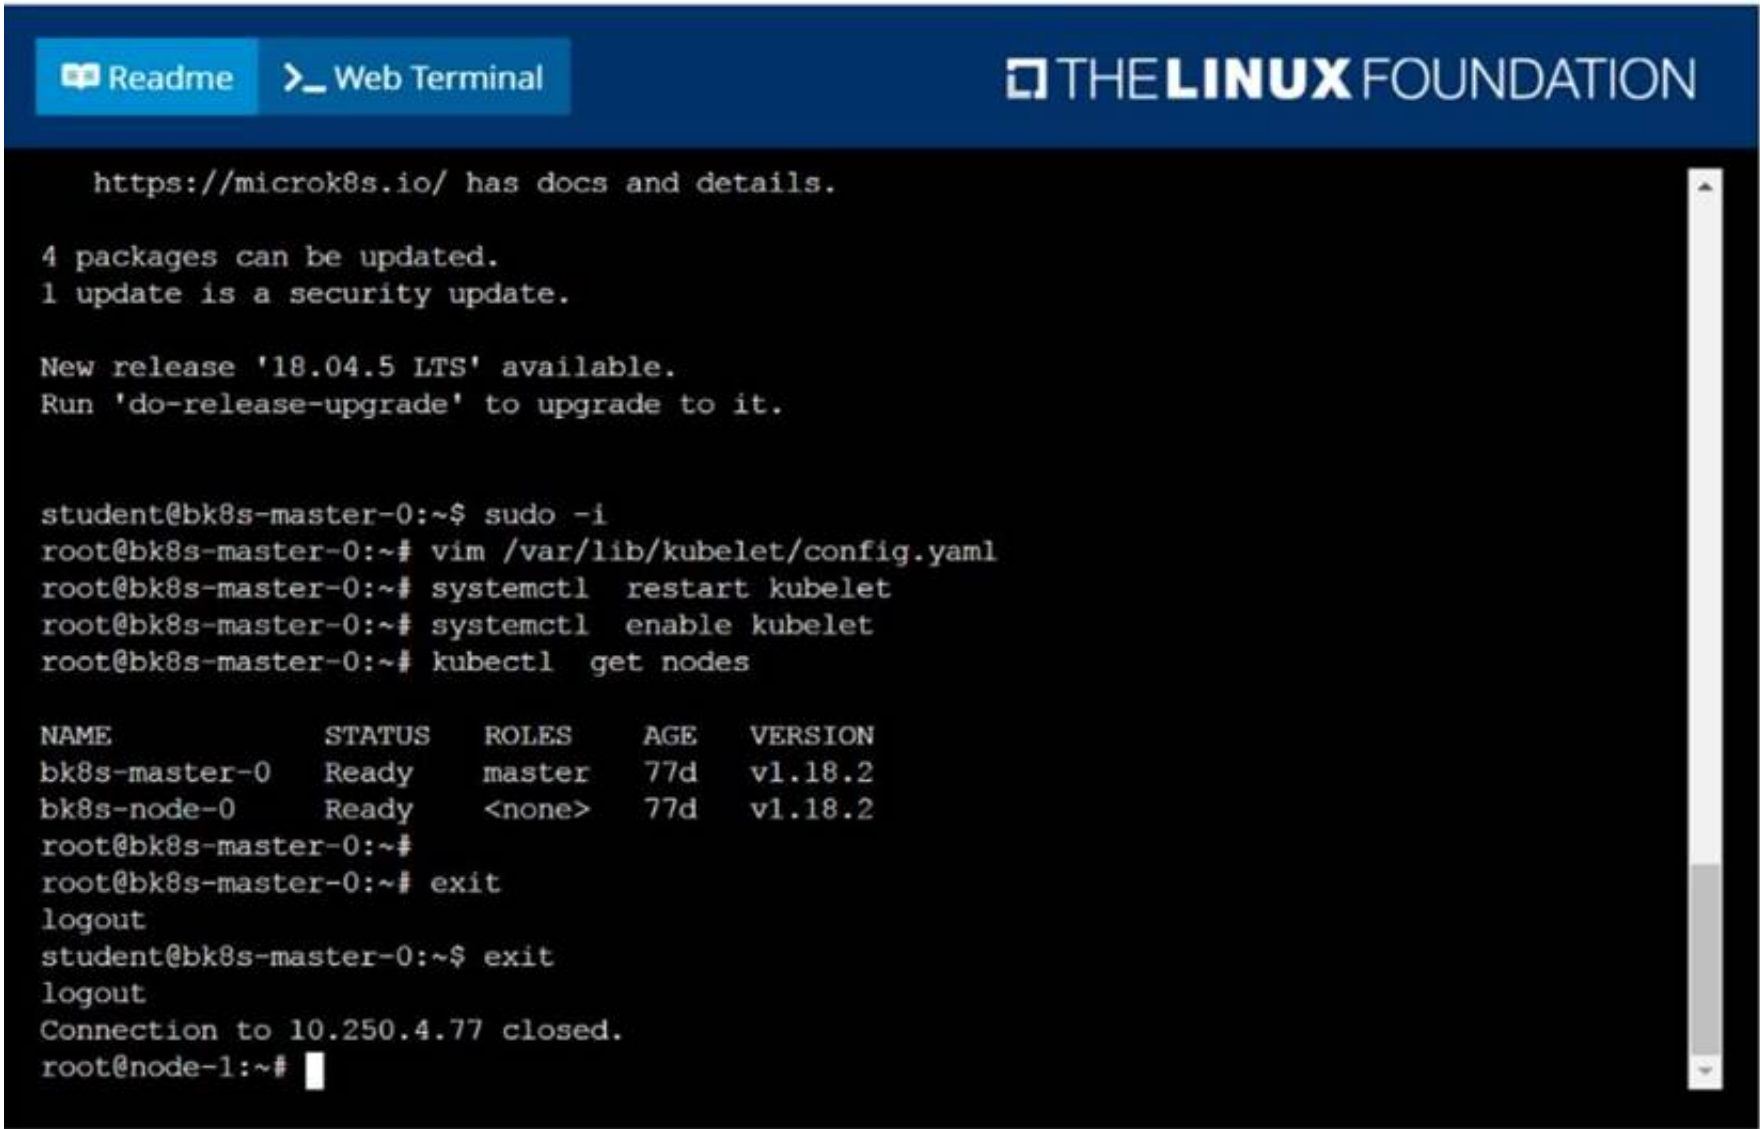
Readme
Web Terminal
THE **LINUX** FOUNDATION

```

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
:wg

```

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



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

NEW QUESTION 10
 CORRECT TEXT
 Score: 4%



Task
 Scale the deployment presentation to 6 pods.

- A. Mastered
- B. Not Mastered

Answer: A


Explanation:
 Solution:
 kubectl get deployment
 kubectl scale deployment.apps/presentation --replicas=6

NEW QUESTION 11
 CORRECT TEXT
 Create a deployment as follows:
 ? Name: nginx-app
 ? Using container nginx with version 1.11.10-alpine
 ? The deployment should contain 3 replicas
 Next, deploy the application with new version 1.11.13-alpine, by performing a rolling update.
 Finally, rollback that update to the previous version 1.11.10-alpine.

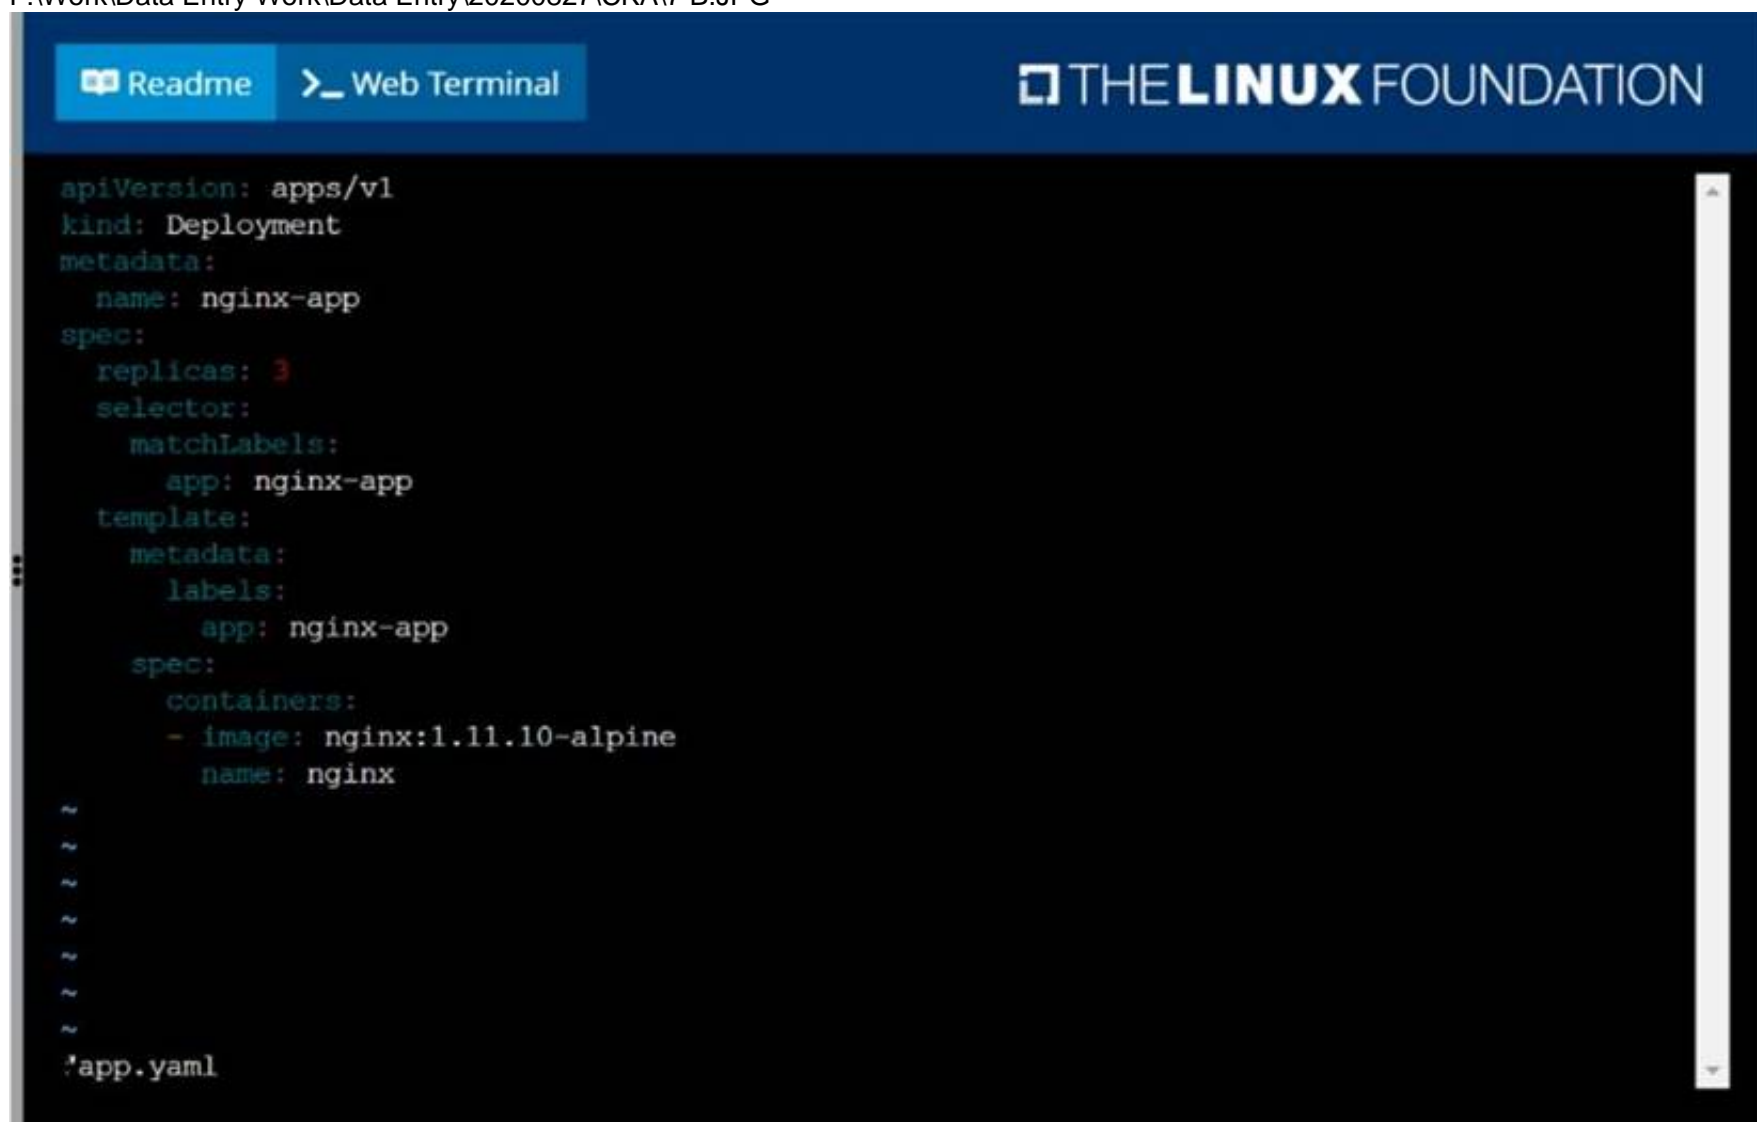
- A. Mastered
- B. Not Mastered

Answer: A

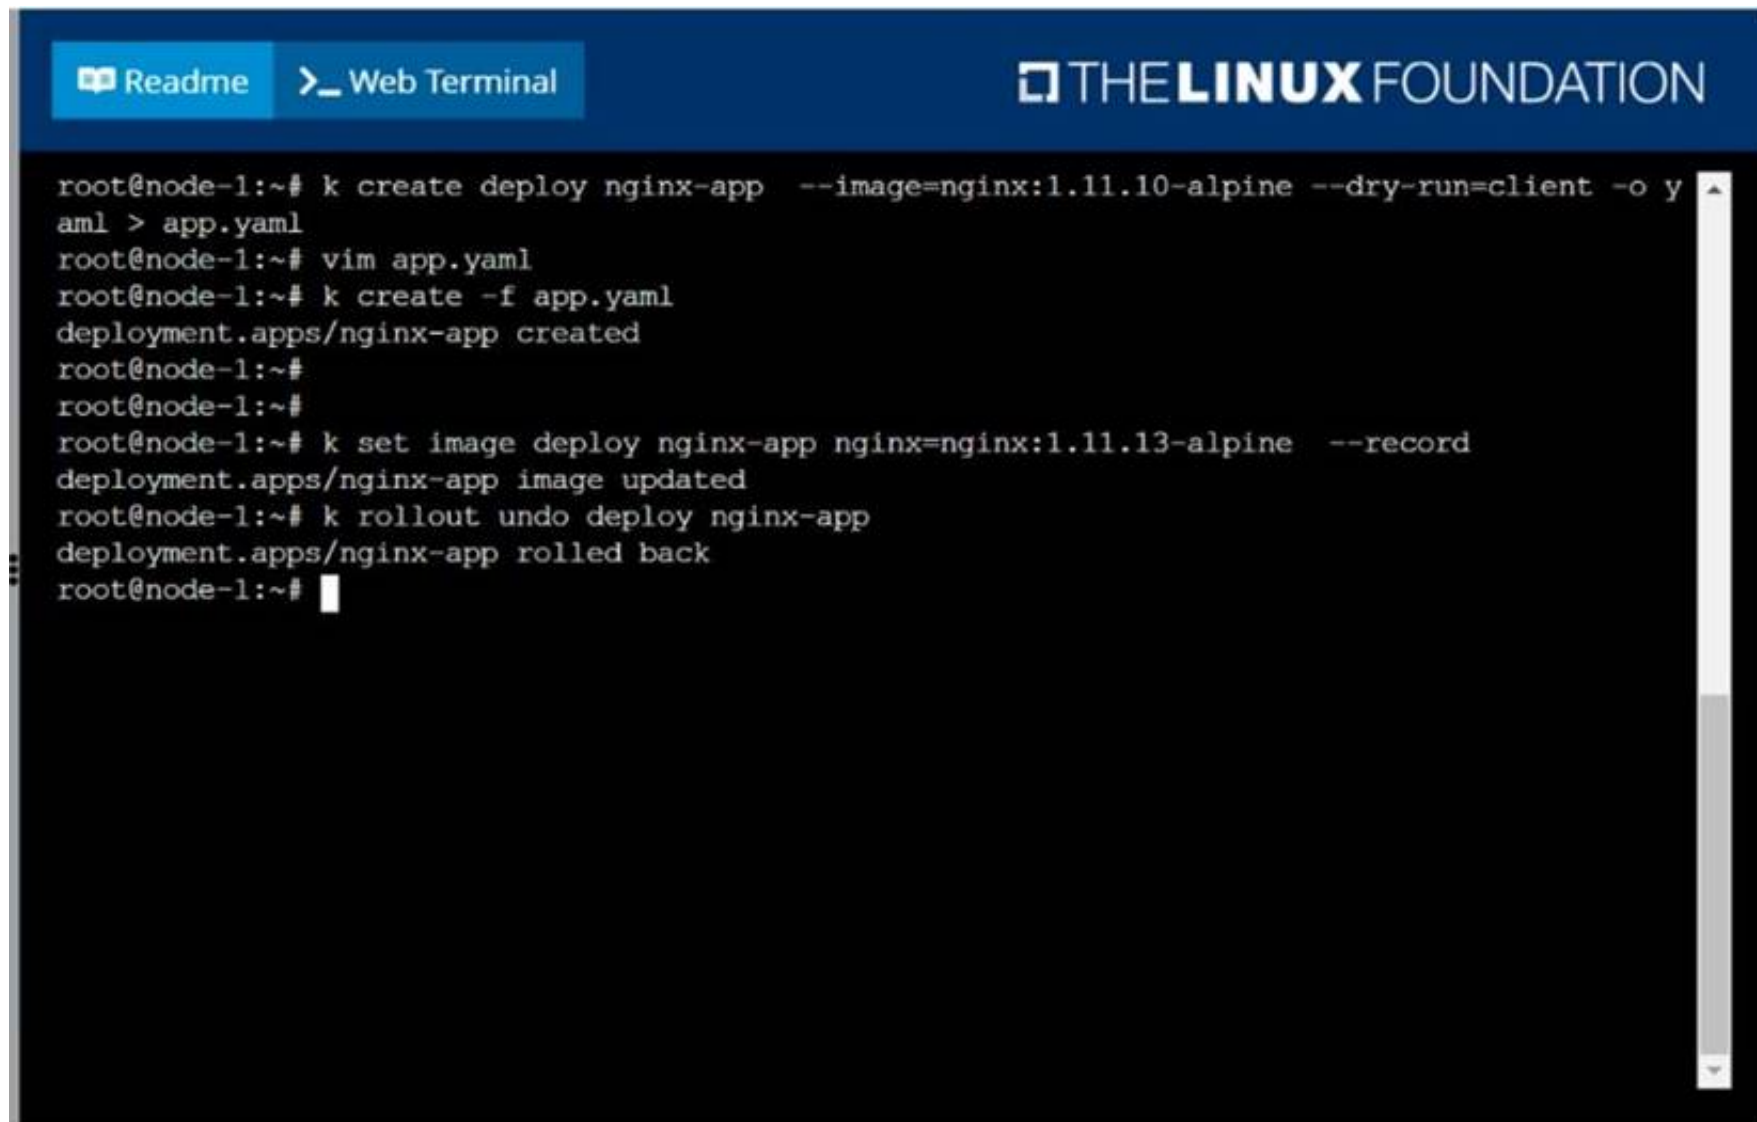
Explanation:
solution



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



F:\Work\Data Entry Work\Data Entry\20200827\CKA\7 C.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:~#
```

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

NEW QUESTION 16

CORRECT TEXT

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



CORRECT TEXT

You can ssh to the failed node using:

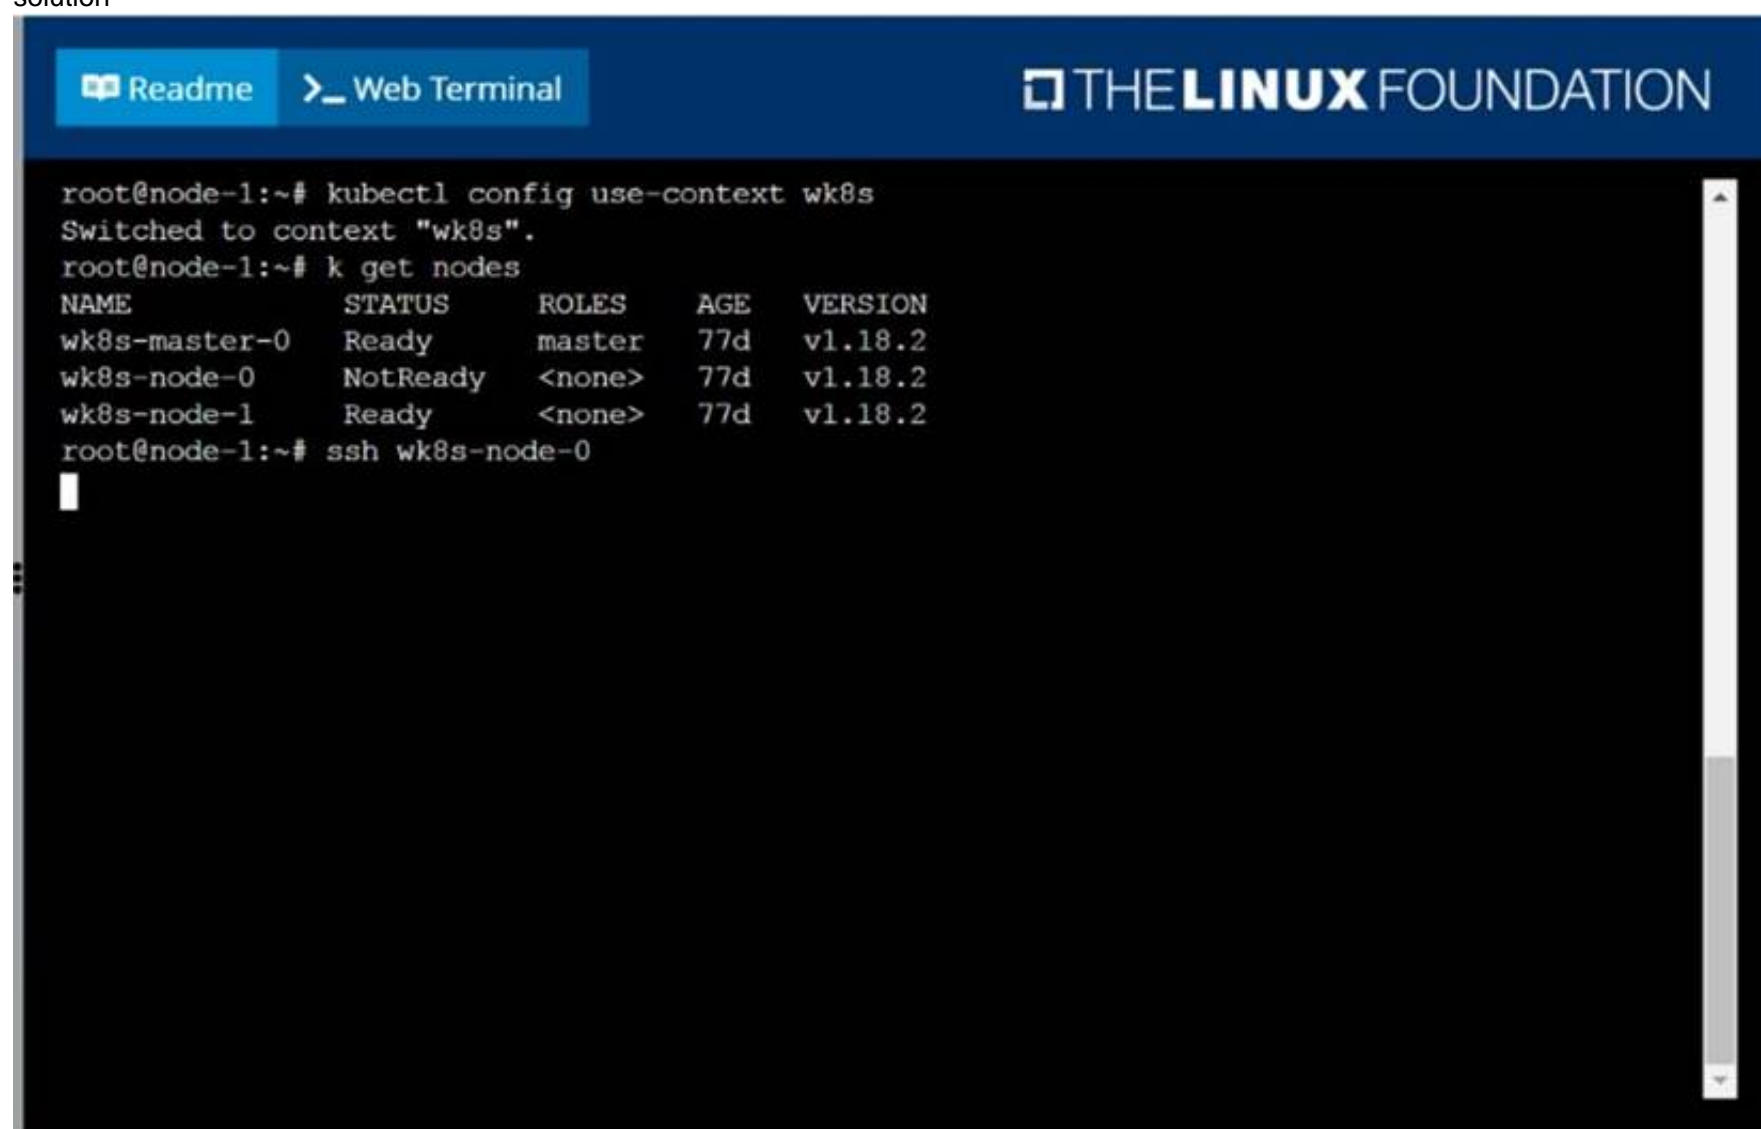
You can assume elevated privileges on the node with the following command:

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

- A. Mastered
B. Not Mastered

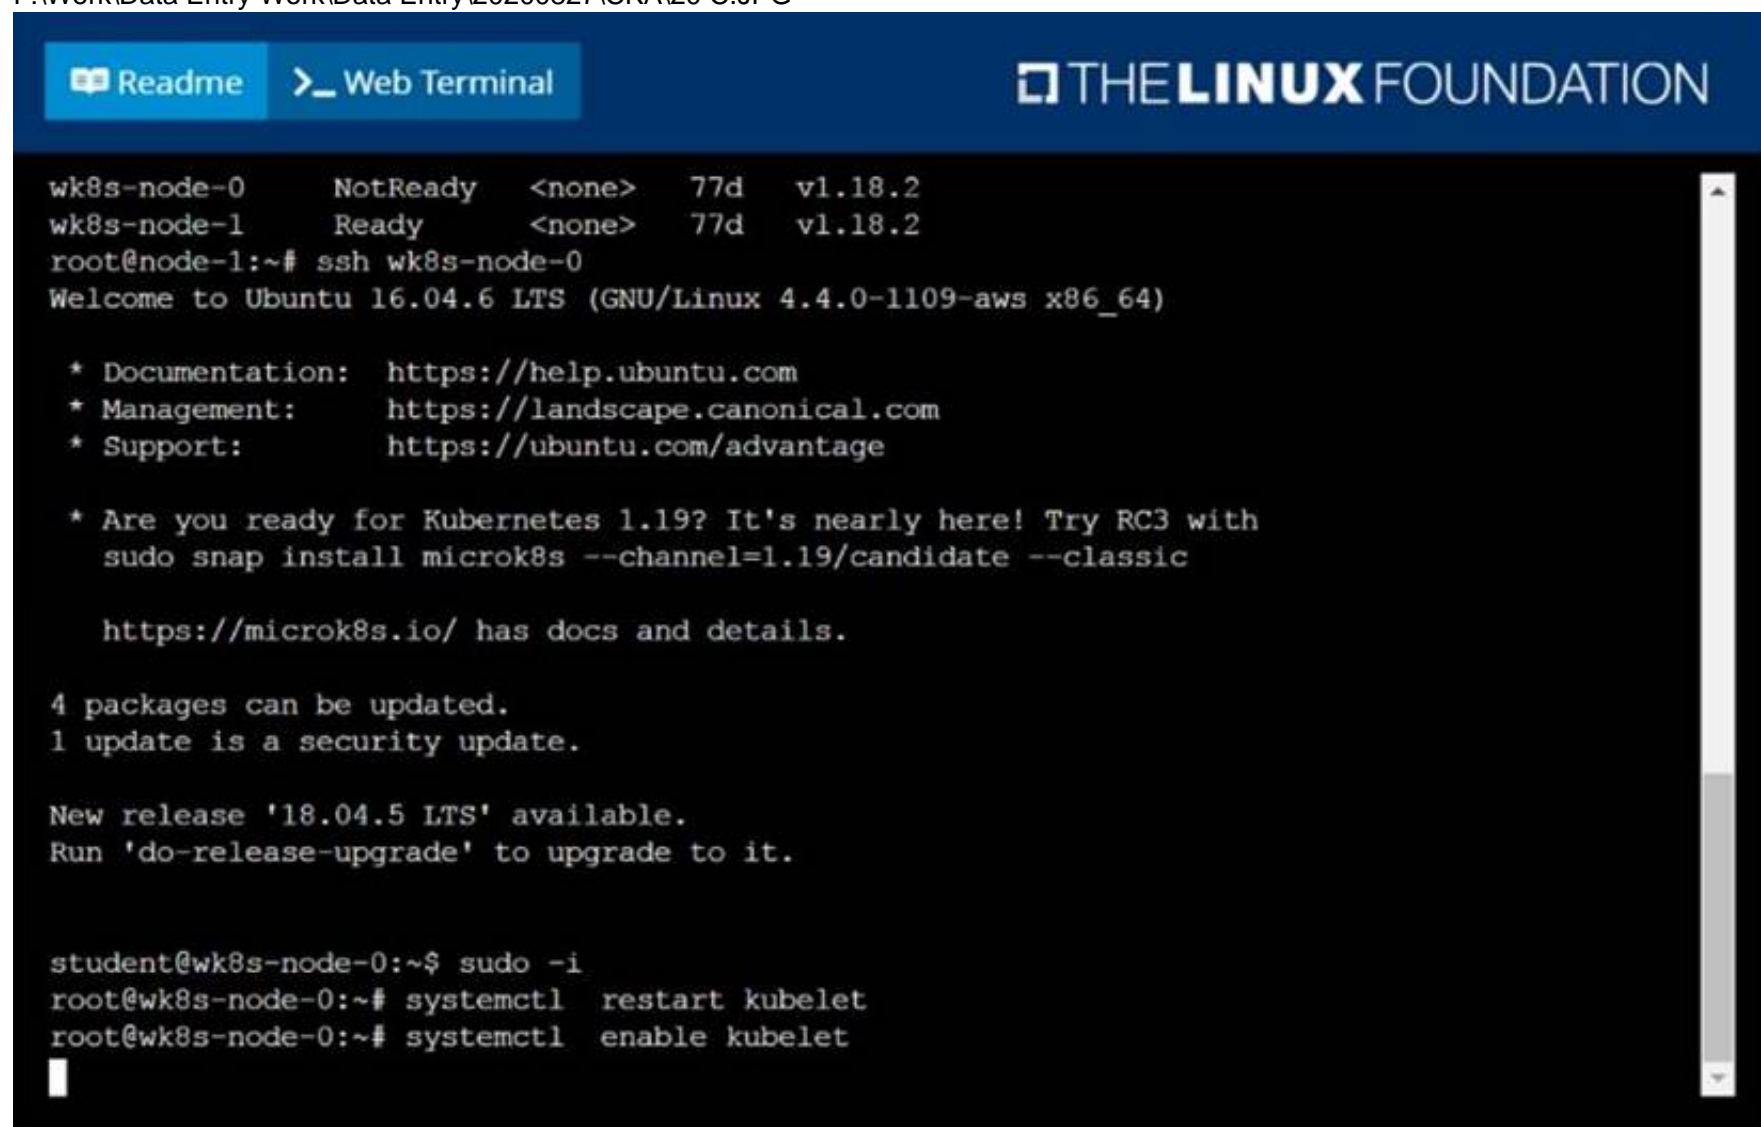
Answer: A

Explanation:
solution



```
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 C.JPG



```
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 D.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/sy
temd/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:~#
```

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

NEW QUESTION 23

CORRECT TEXT

List all the pods showing name and namespace with a json path expression

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 27

CORRECT TEXT

List all the pods sorted by name

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 32

CORRECT TEXT

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 35

CORRECT TEXT

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 39

CORRECT TEXT

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 43

.....

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](#)