Exam Questions CKA

Certified Kubernetes Administrator (CKA) Program

https://www.2passeasy.com/dumps/CKA/

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 cansshto 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

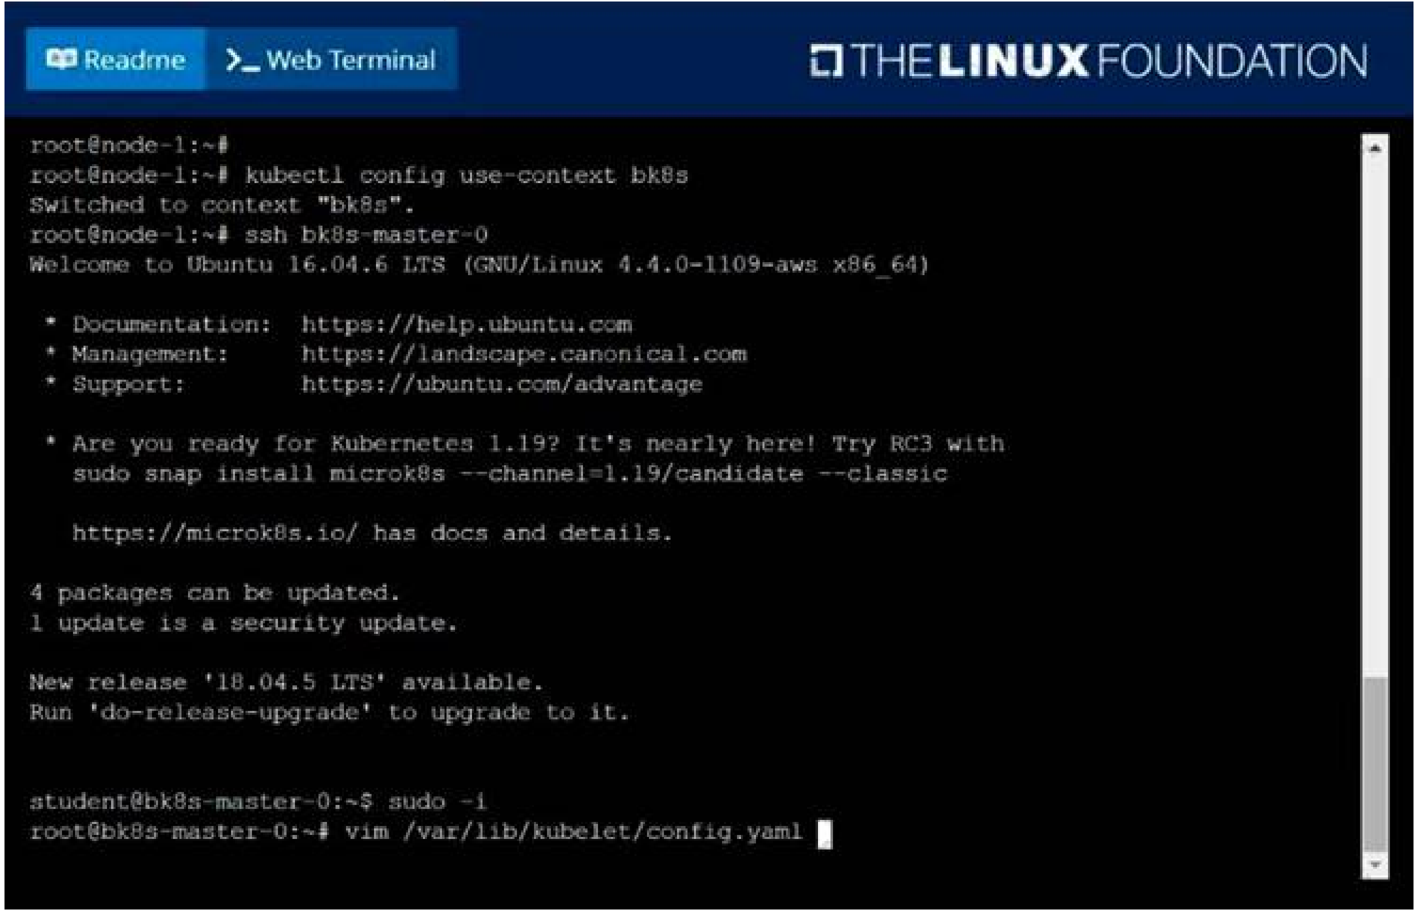

```
root@node-1:~#
root@node-1:~# kubectl config use-context bk8s
Switched to context "bk8s".
root@node-1:~# ssh bk8s-master-0
Welcome to Ubuntu 16.04.6 LTS (GNU/Linux 4.4.0-1109-aws x86_64)

 * 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 D.JPG

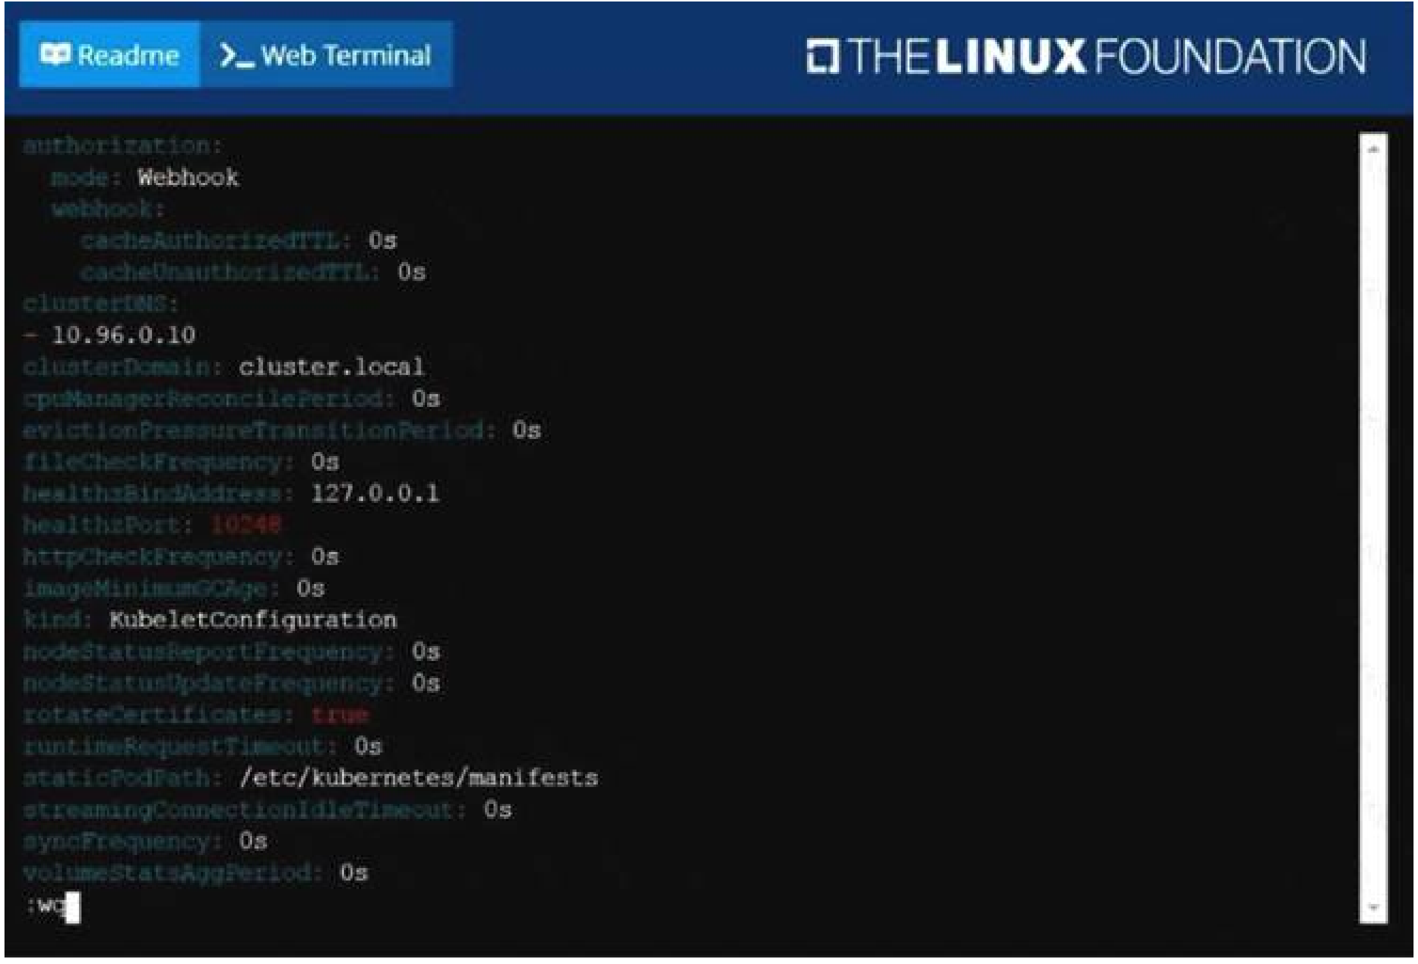

```
authorization:
  mode: Webhook
  webhook:
    cacheAuthorizedTTL: 0s
    cacheUnauthorizedTTL: 0s
clusterDNS:
- 10.96.0.10
clusterDomain: cluster.local
cpuManagerReconcilePeriod: 0s
evictionPressureTransitionPeriod: 0s
fileCheckFrequency: 0s
healthzBindAddress: 127.0.0.1
healthzPort: 10248
httpCheckFrequency: 0s
imageMinimumGCAge: 0s
kind: KubeletConfiguration
nodeStatusReportFrequency: 0s
nodeStatusUpdateFrequency: 0s
rotateCertificates: true
runtimeRequestTimeout: 0s
staticPodPath: /etc/kubernetes/manifests
streamingConnectionIdleTimeout: 0s
syncFrequency: 0s
volumeStatsAggPeriod: 0s
:wq
```

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

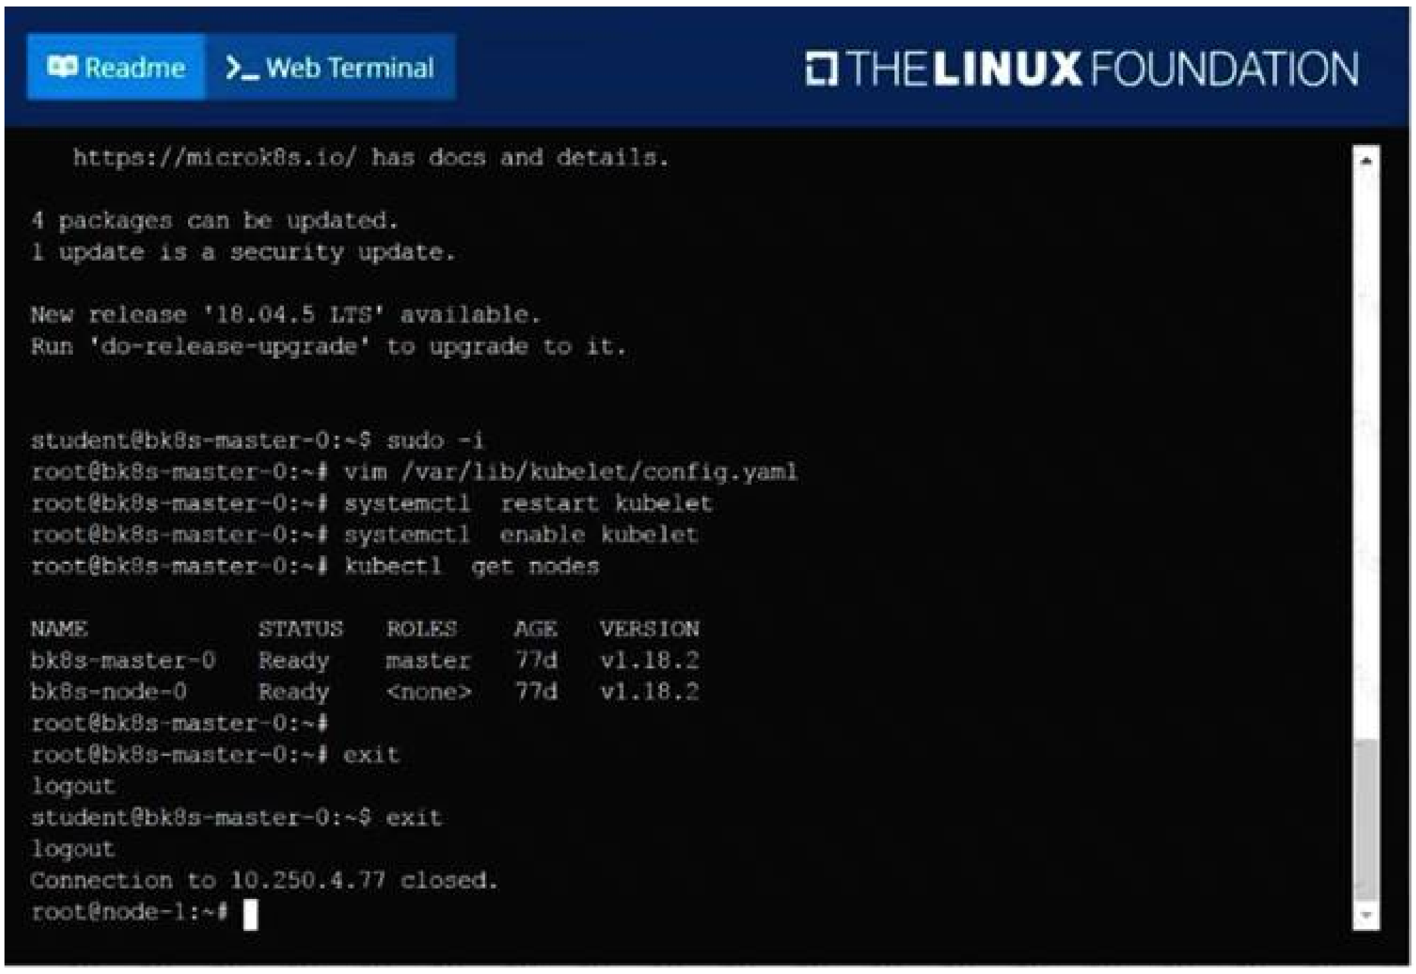

```
   https://microk8s.io/ has docs and details.

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:~# kubectl  get nodes

NAME            STATUS   ROLES    AGE   VERSION
bk8s-master-0   Ready    master   77d   v1.18.2
bk8s-node-0     Ready    <none>   77d   v1.18.2
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

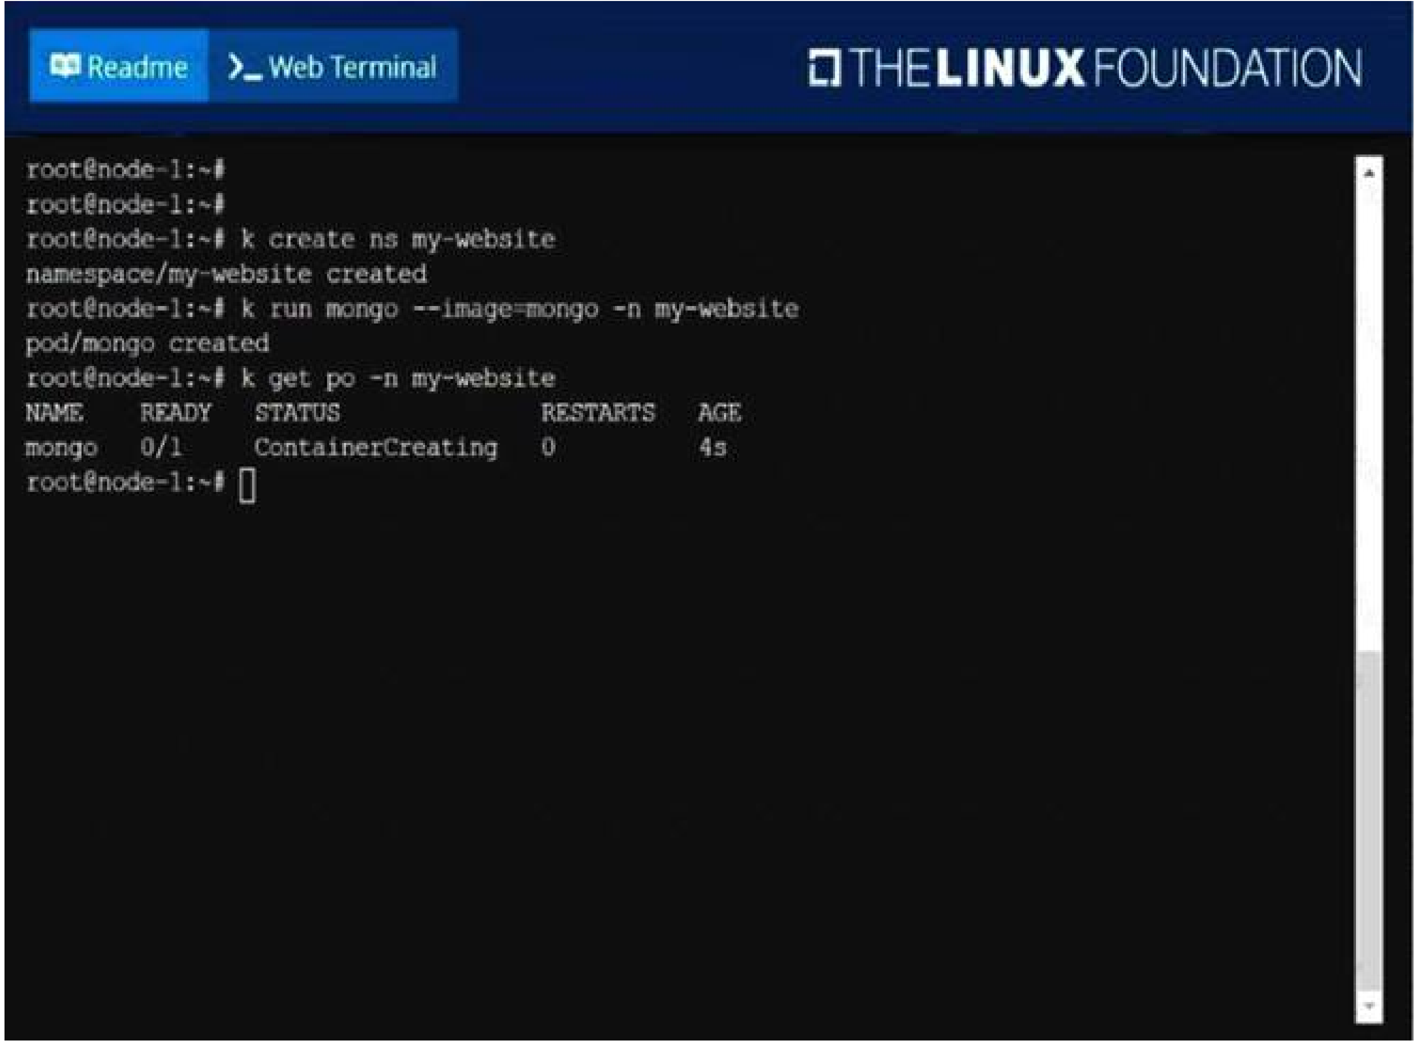

```
root@node-1:~#
root@node-1:~#
root@node-1:~# k create ns my-website
namespace/my-website created
root@node-1:~# k run mongo --image=mongo -n my-website
pod/mongo created
root@node-1:~# k get po -n my-website
NAME    READY   STATUS              RESTARTS   AGE
mongo   0/1     ContainerCreating   0          4s
root@node-1:~#
```

NEW QUESTION 3
Create a deployment as follows:

- Name:nginx-app
- Using containernginxwithversion 1.11.10-alpine
- The deployment should contain3replicas

Next, deploy the application with newversion1.11.13-alpine, byperforming a rolling update.
Finally, rollback that update to theprevious version1.11.10-alpine.

A. Mastered
B. Not Mastered

Answer: A

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

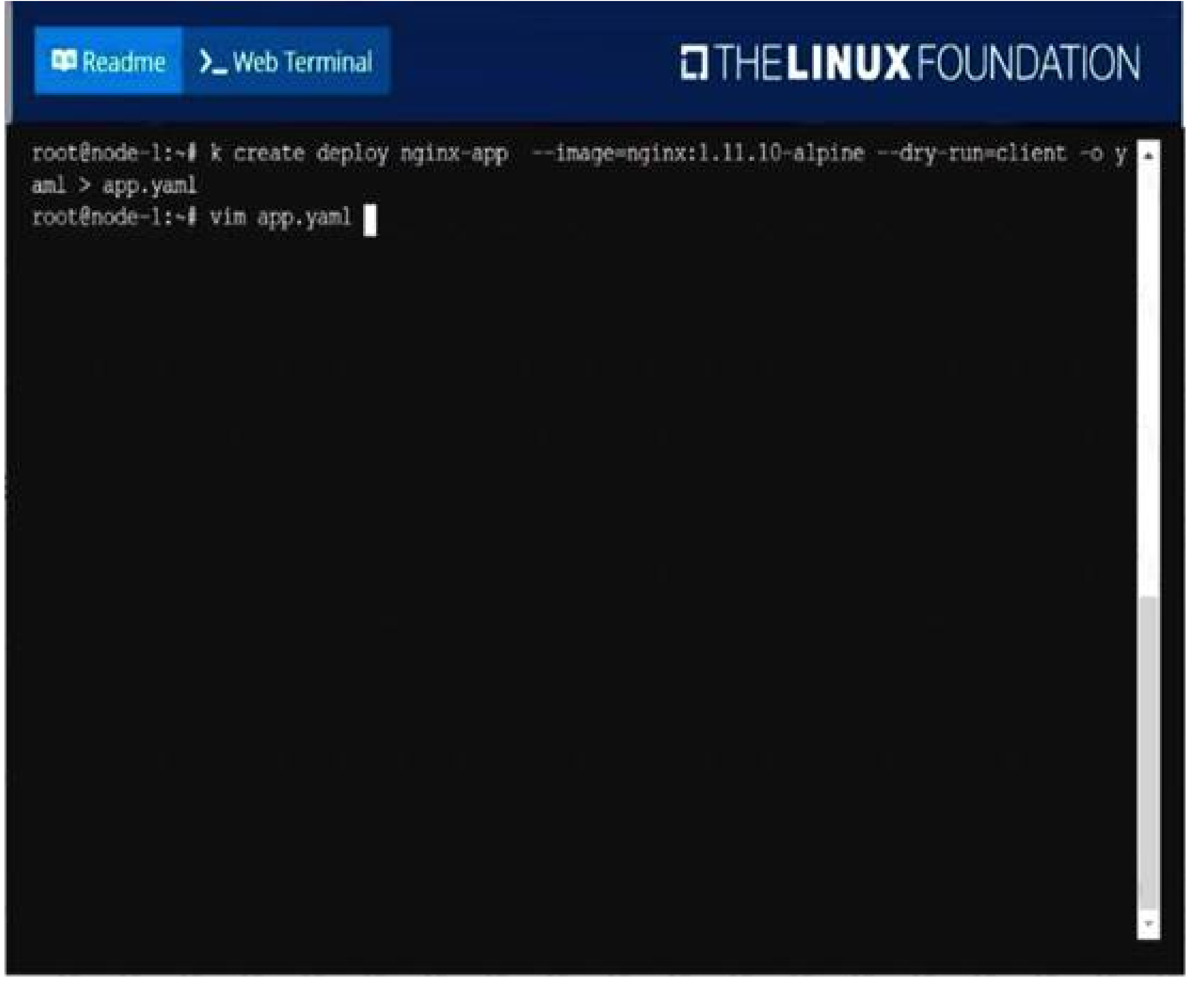

```
root@node-1:~# k create deploy nginx-app --image=nginx:1.11.10-alpine --dry-run=client -o y
aml > app.yaml
root@node-1:~# vim app.yaml
```

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

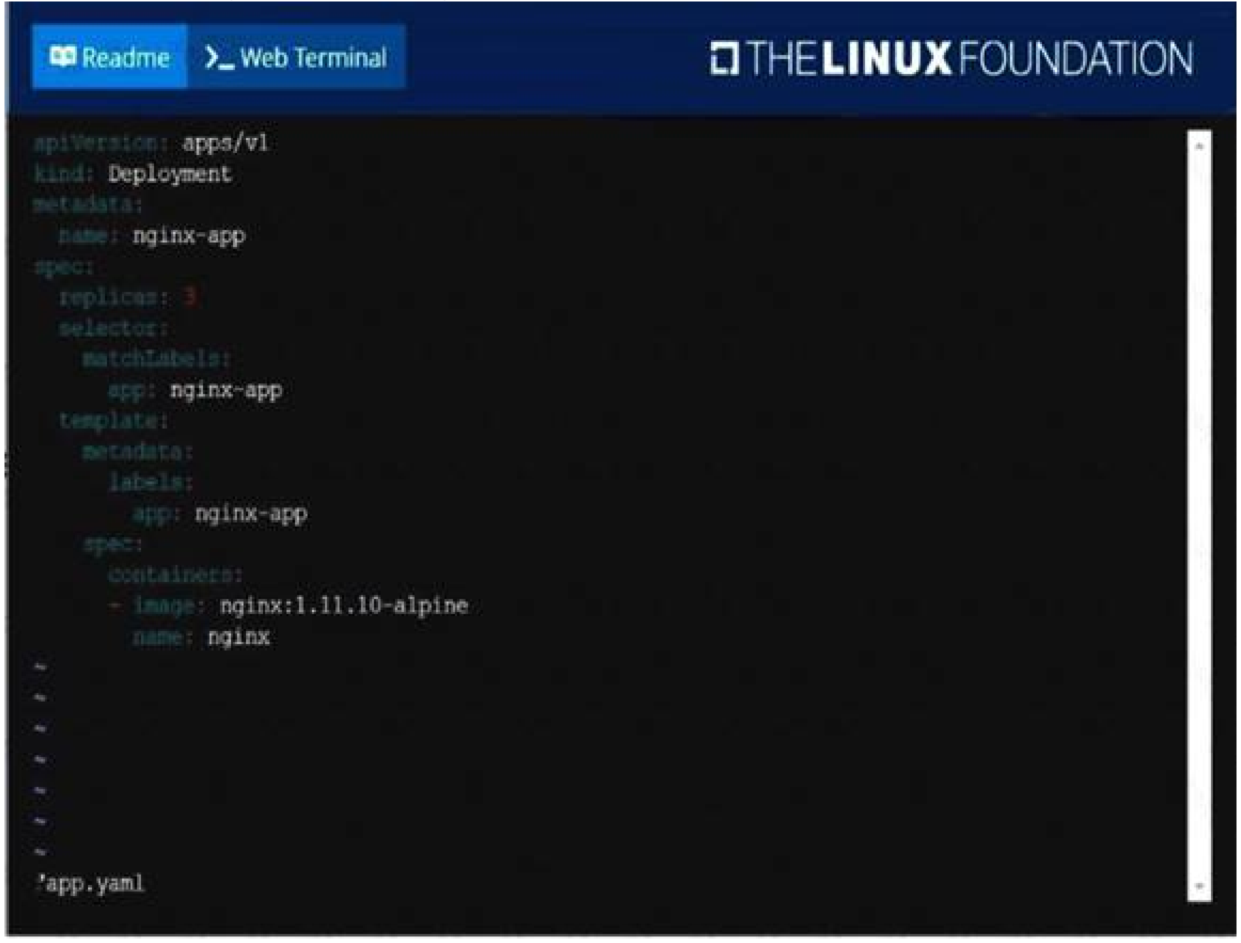

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-app
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx-app
  template:
    metadata:
      labels:
        app: nginx-app
    spec:
      containers:
      - image: nginx:1.11.10-alpine
        name: nginx
~
~
~
~
~
~
~
'app.yaml
```

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

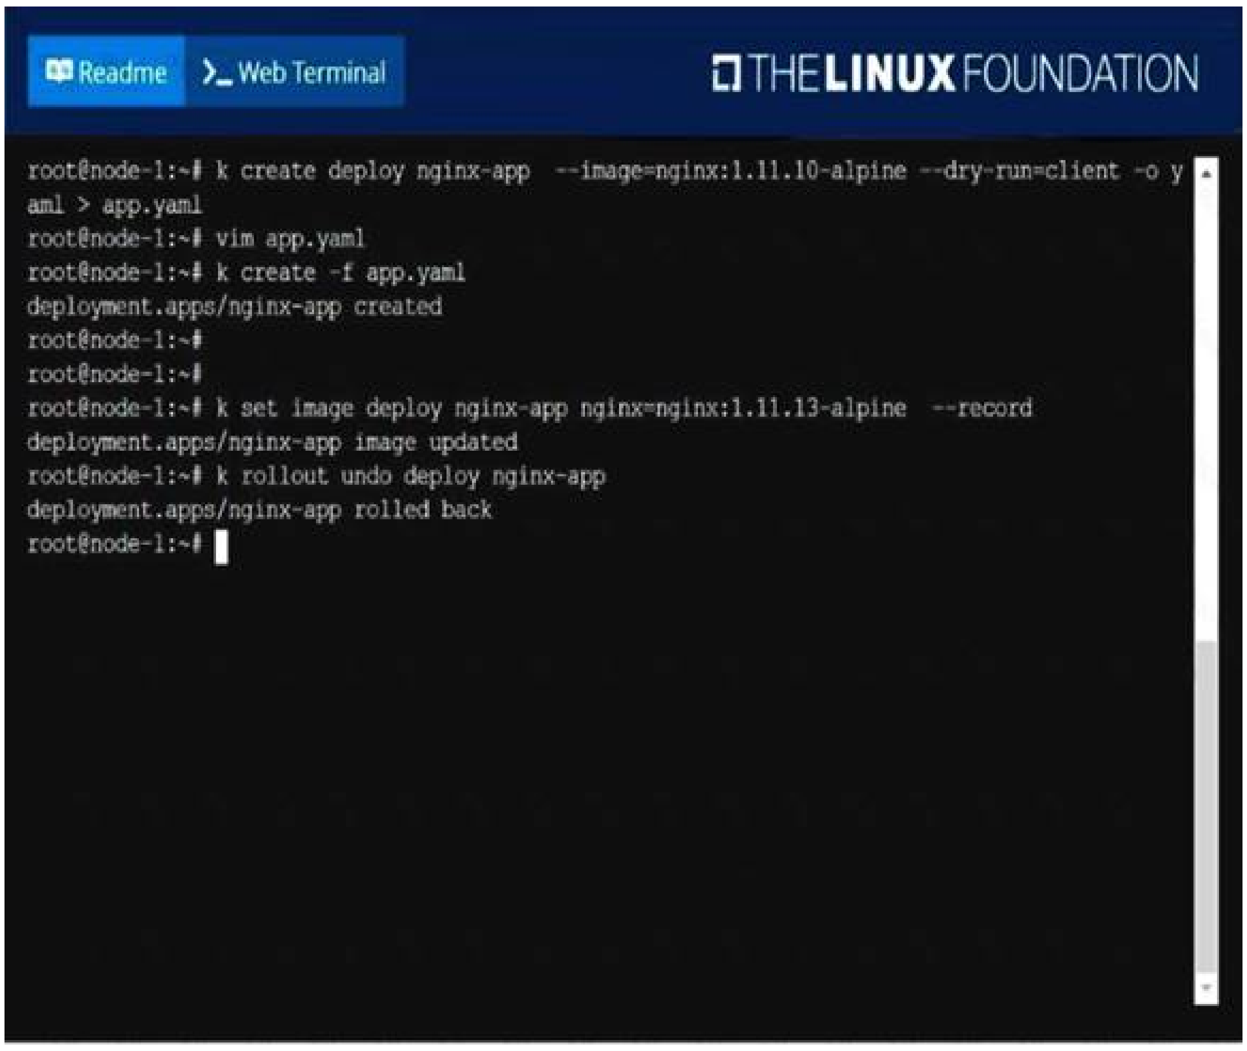

```
root@node-1:~# k create deploy nginx-app  --image=nginx:1.11.10-alpine --dry-run=client -o y
aml > app.yaml
root@node-1:~# vim app.yaml
root@node-1:~# k create -f app.yaml
deployment.apps/nginx-app created
root@node-1:~#
root@node-1:~#
root@node-1:~# k set image deploy nginx-app nginx=nginx:1.11.13-alpine  --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:~#
```

NEW QUESTION 4
Create a pod with image nginx called nginx and allow traffic on port 80

A. Mastered
B. Not Mastered

Answer: A

Explanation:
kubectlrun nginx --image=nginx --restart=Never --port=80

NEW QUESTION 5
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 -nengineering-f ?C
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 6
Set the node named ek8s-node-1as 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

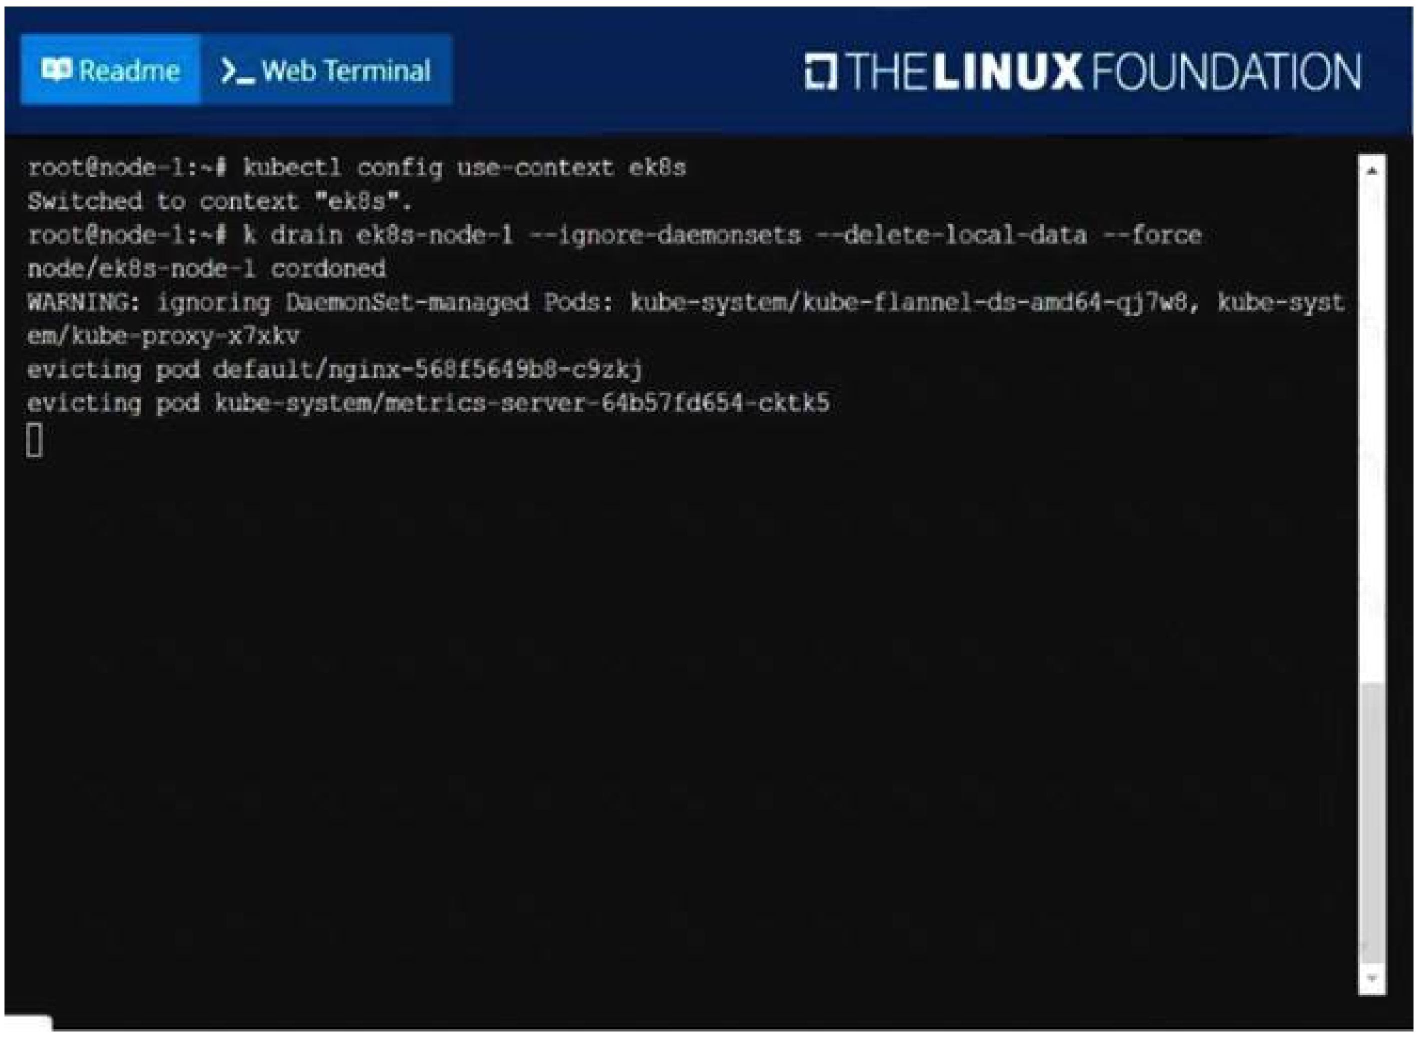

```
root@node-1:~# kubectl config use-context ek8s
Switched to context "ek8s".
root@node-1:~# k drain ek8s-node-1 --ignore-daemonsets --delete-local-data --force
node/ek8s-node-1 cordoned
WARNING: ignoring DaemonSet-managed Pods: kube-system/kube-flannel-ds-amd64-qj7w8, kube-syst
em/kube-proxy-x7xkv
evicting pod default/nginx-568f5649b8-c9zkj
evicting pod kube-system/metrics-server-64b57fd654-cktk5
```

NEW QUESTION 7

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 tying 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.

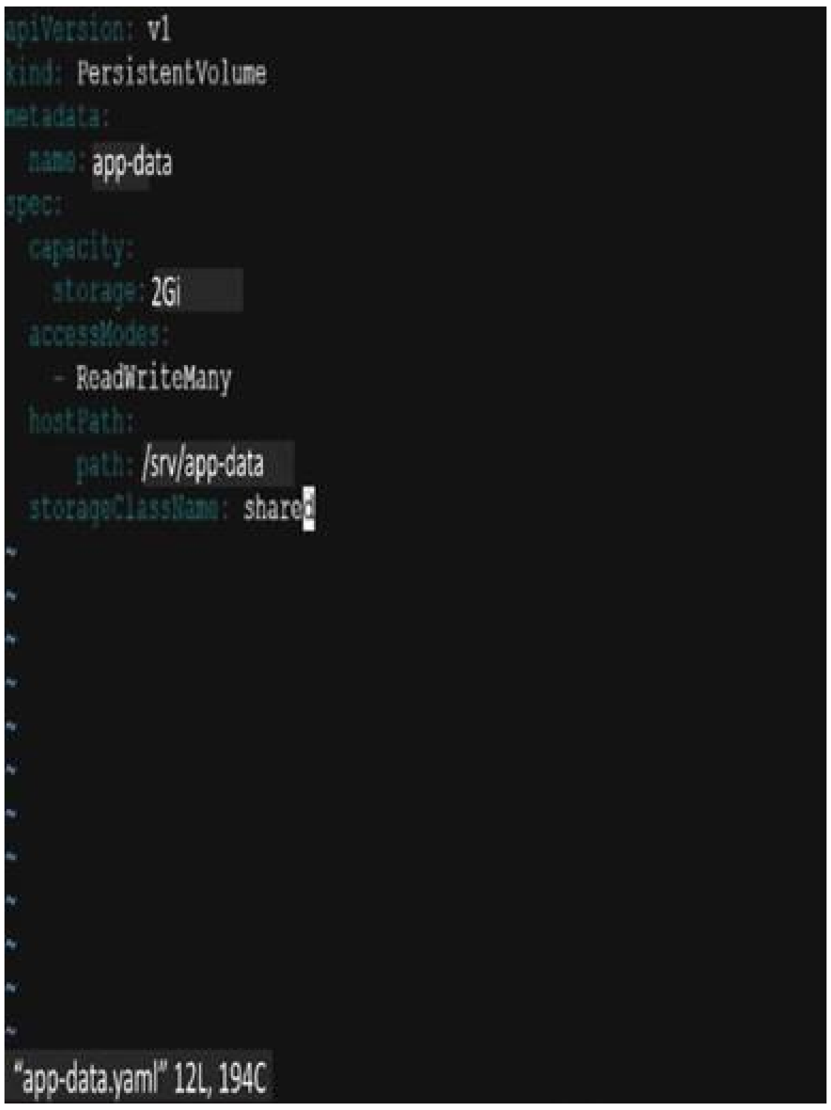

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: app-data
spec:
  capacity:
    storage: 2Gi
  accessModes:
    - ReadWriteMany
  hostPath:
      path: /srv/app-data
  storageClassName: shared
~
~
~
~
~
~
~
~
~
~
~
~
"app-data.yaml" 12L, 194C
```

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

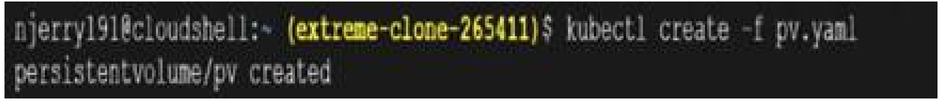

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl create -f pv.yaml
persistentvolume/pv created
```

* 3. View the persistent volume.

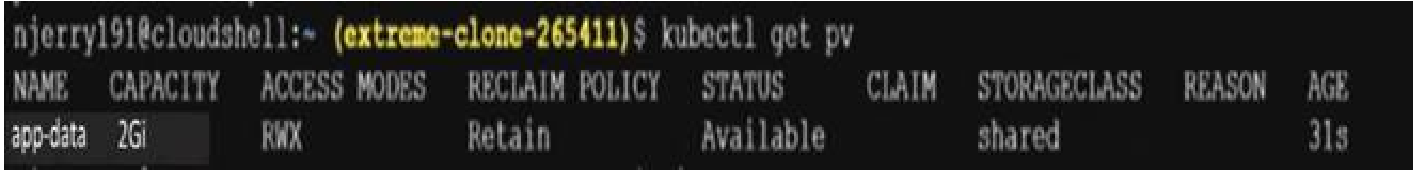

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl get pv
NAME      CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS      CLAIM   STORAGECLASS   REASON   AGE
app-data  2Gi        RWX            Retain           Available           shared                  31s
```

➢ 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)$ kubect1 create -f app-data.yaml persistentvolumeclaim/app-data created

* 3. View the pvc Image for post

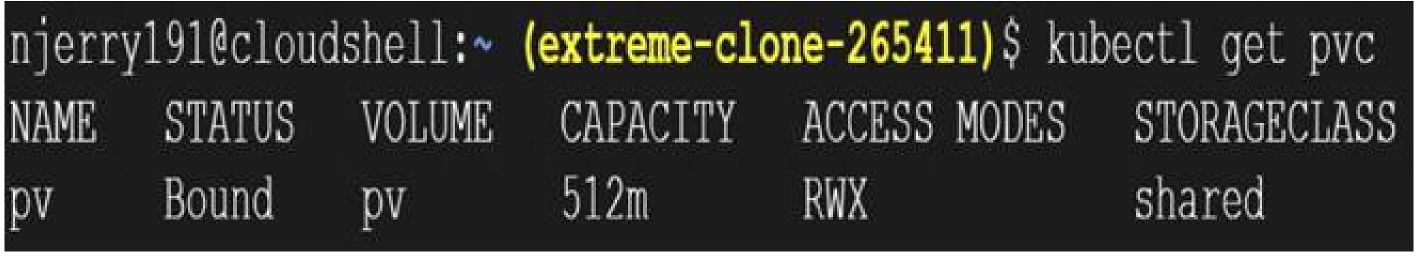

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl get pvc
NAME   STATUS   VOLUME   CAPACITY   ACCESS MODES   STORAGECLASS
pv     Bound    pv       512m       RWX            shared
```

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

Image for post

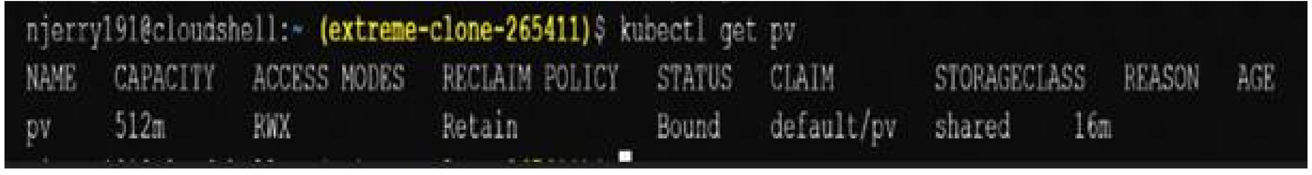

```
njerry191@cloudshell:~ (extreme-clone-265411)$ kubectl get pv
NAME   CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM        STORAGECLASS   REASON   AGE
pv     512m       RWX            Retain           Bound    default/pv   shared         16m
```

Our status has now changed fromavailabletobound.

* 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: v1kind: Podmetadata:creationTimestamp: nullname: app-dataspec:volumes:- name:congigpvcpersistenVolumeClaim:claimName: app-datacontainers:- image: nginxname: appvolumeMounts:- mountPath: "/srv/app-data"name: configpvc

NEW QUESTION 8

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 9
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
Create a pod as follows:

- Name:non-persistent-redis
- container Image:redis
- Volume with name:cache-control
- Mount path:/data/redis

The pod should launch in thestagingnamespace and the volumemust notbe persistent.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution
F:\Work\Data Entry Work\Data Entry\20200827\CKA\13 B.JPG

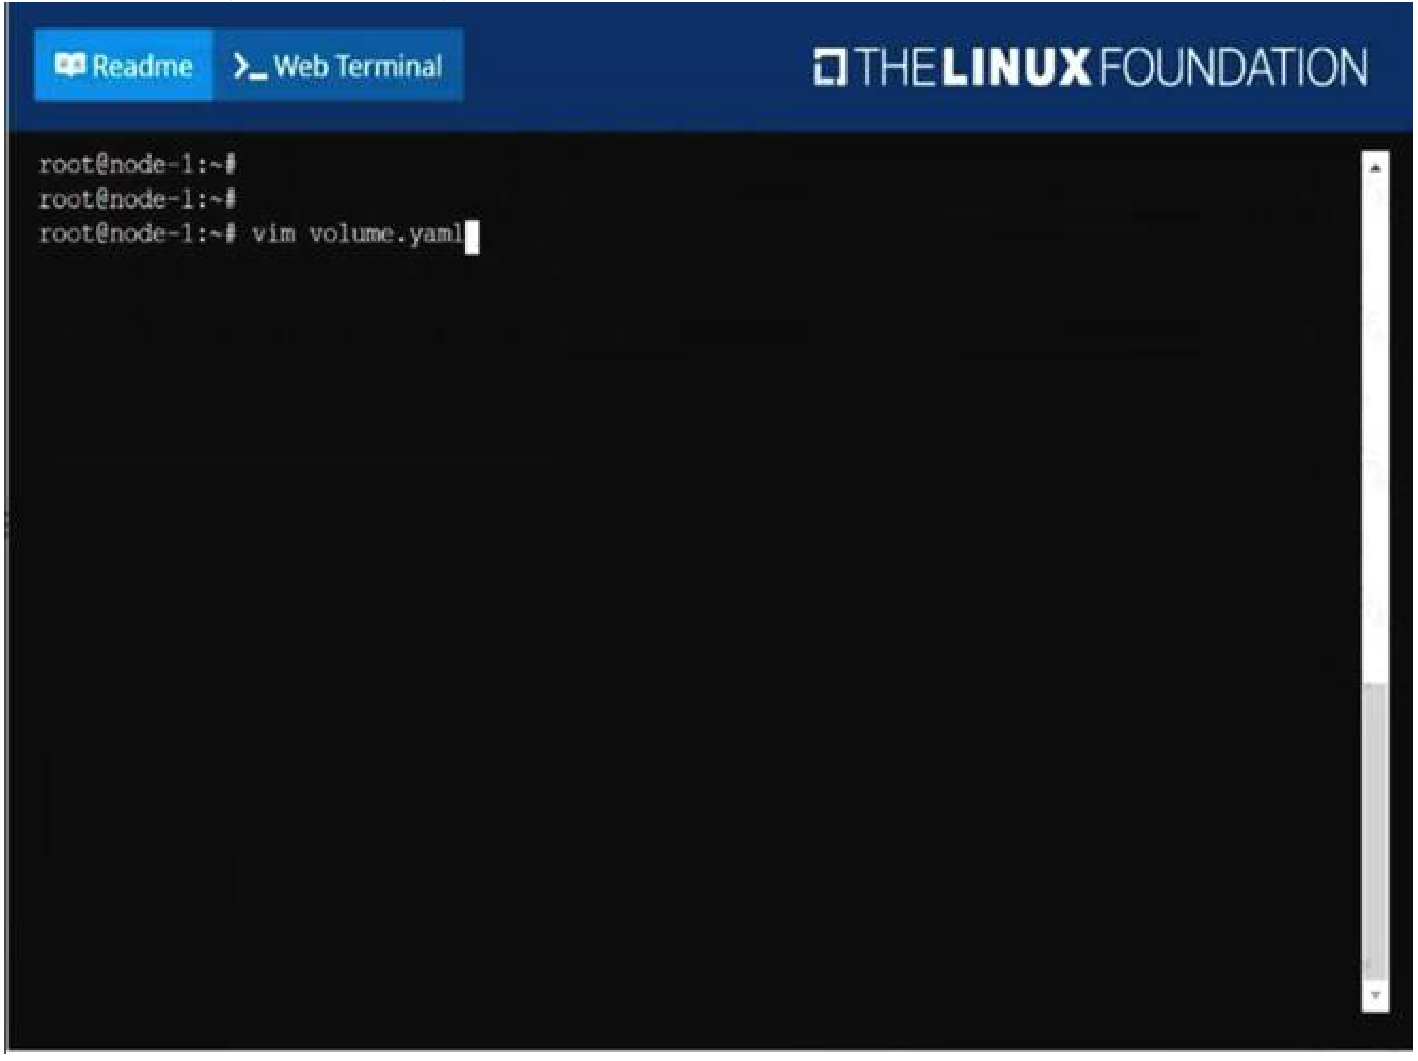

```
root@node-1:~#
root@node-1:~#
root@node-1:~# vim volume.yaml
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\13 C.JPG

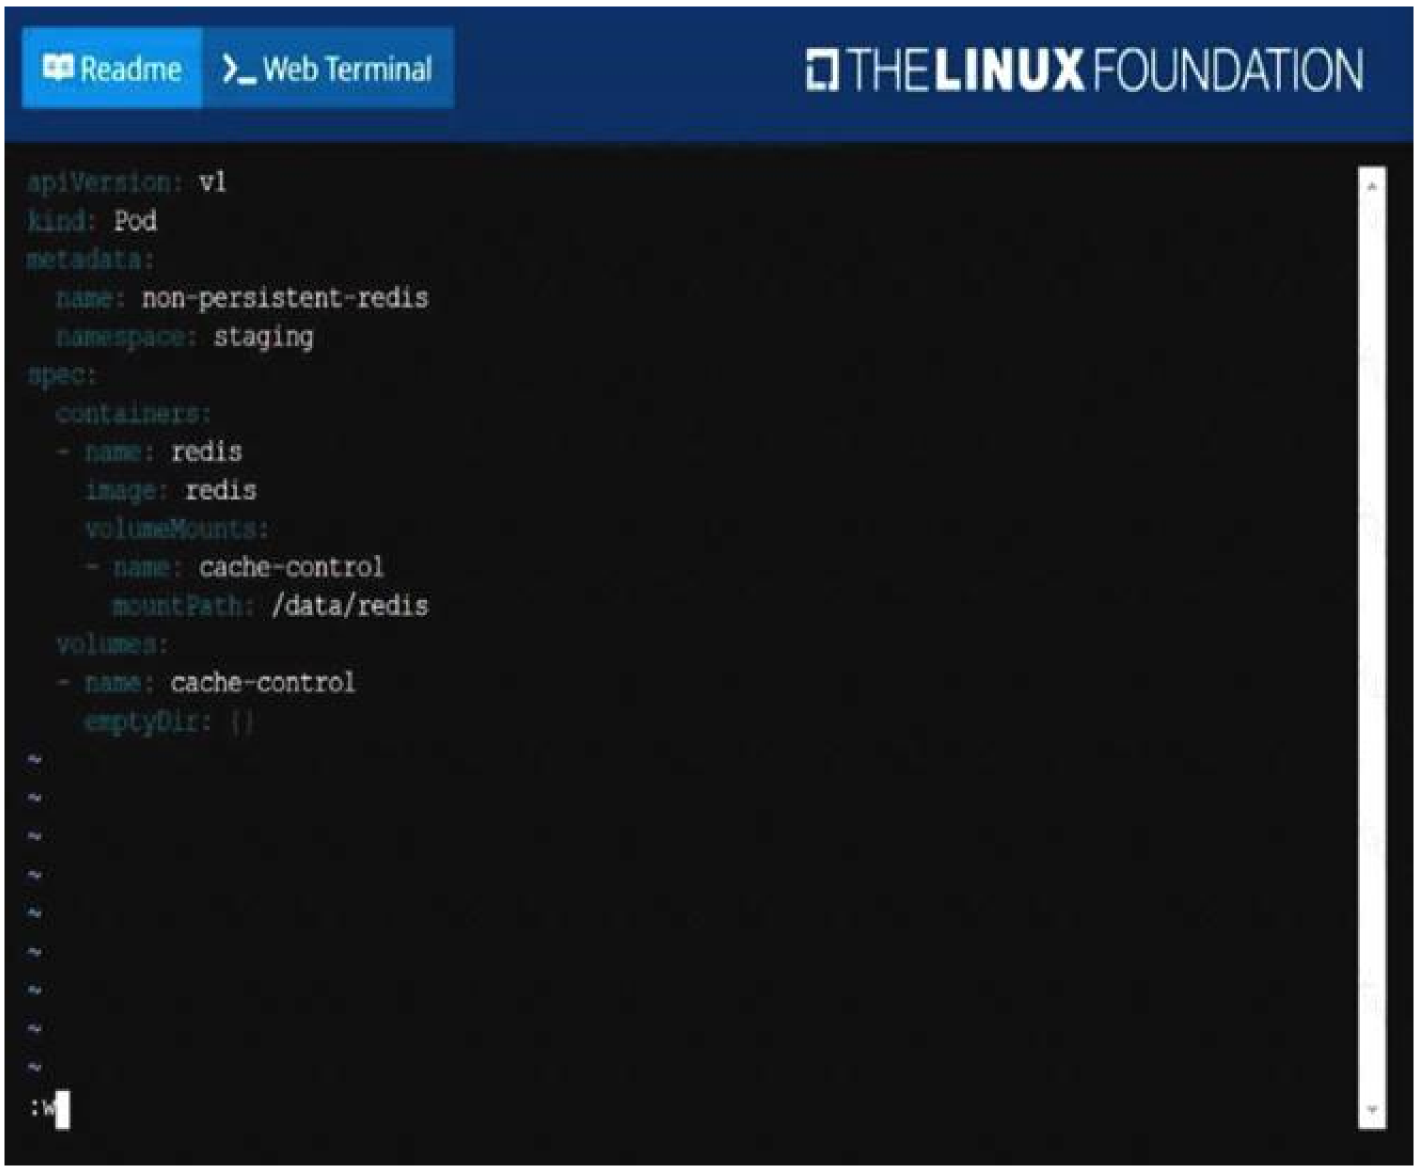

```
apiVersion: v1
kind: Pod
metadata:
  name: non-persistent-redis
  namespace: staging
spec:
  containers:
  - name: redis
    image: redis
    volumeMounts:
    - name: cache-control
      mountPath: /data/redis
  volumes:
  - name: cache-control
    emptyDir: {}
~
~
~
~
~
~
~
~
~
:w
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\13 D.JPG

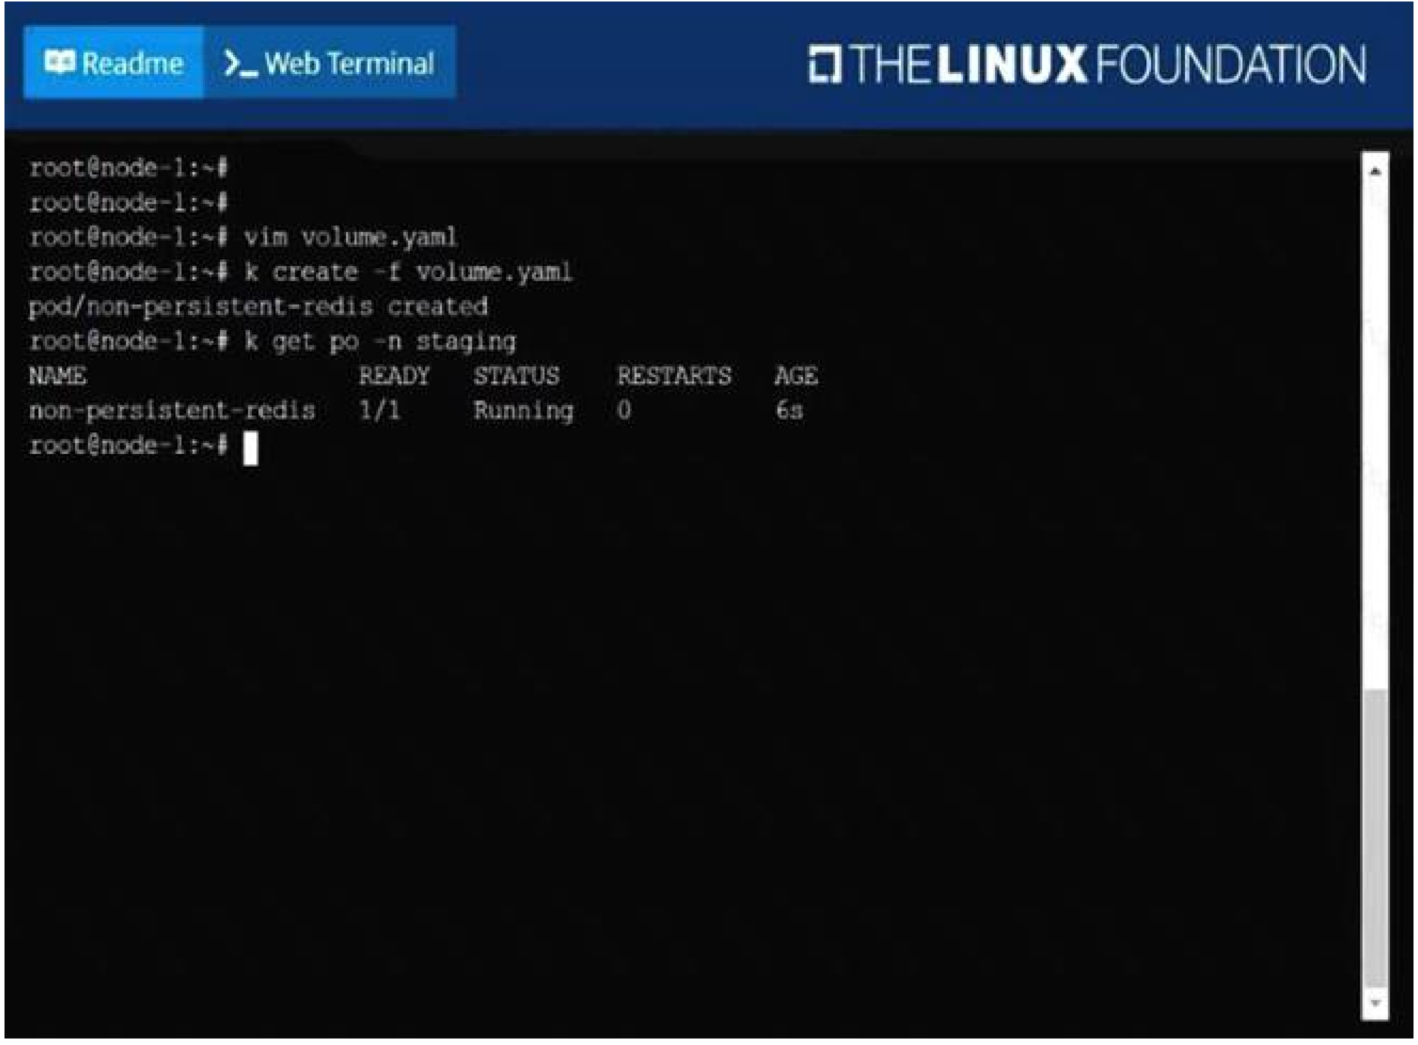

```
root@node-1:~#
root@node-1:~#
root@node-1:~# vim volume.yaml
root@node-1:~# k create -f volume.yaml
pod/non-persistent-redis created
root@node-1:~# k get po -n staging
NAME                   READY   STATUS    RESTARTS   AGE
non-persistent-redis   1/1     Running   0          6s
root@node-1:~#
```

NEW QUESTION 10
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 ?C-rm -it -- env > envpod.yaml

NEW QUESTION 15
Create a Kubernetes secret asfollows:

- Name: super-secret
-

password: bob
Create a pod namedpod-secrets-via-file, using theredisImage, which mounts a secret namedsuper-secretat /secrets.
Create a second pod namedpod-secrets-via-env, using theredisImage, which exportspasswordas CONFIDENTIAL

A. Mastered
B. Not Mastered

Answer: A

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

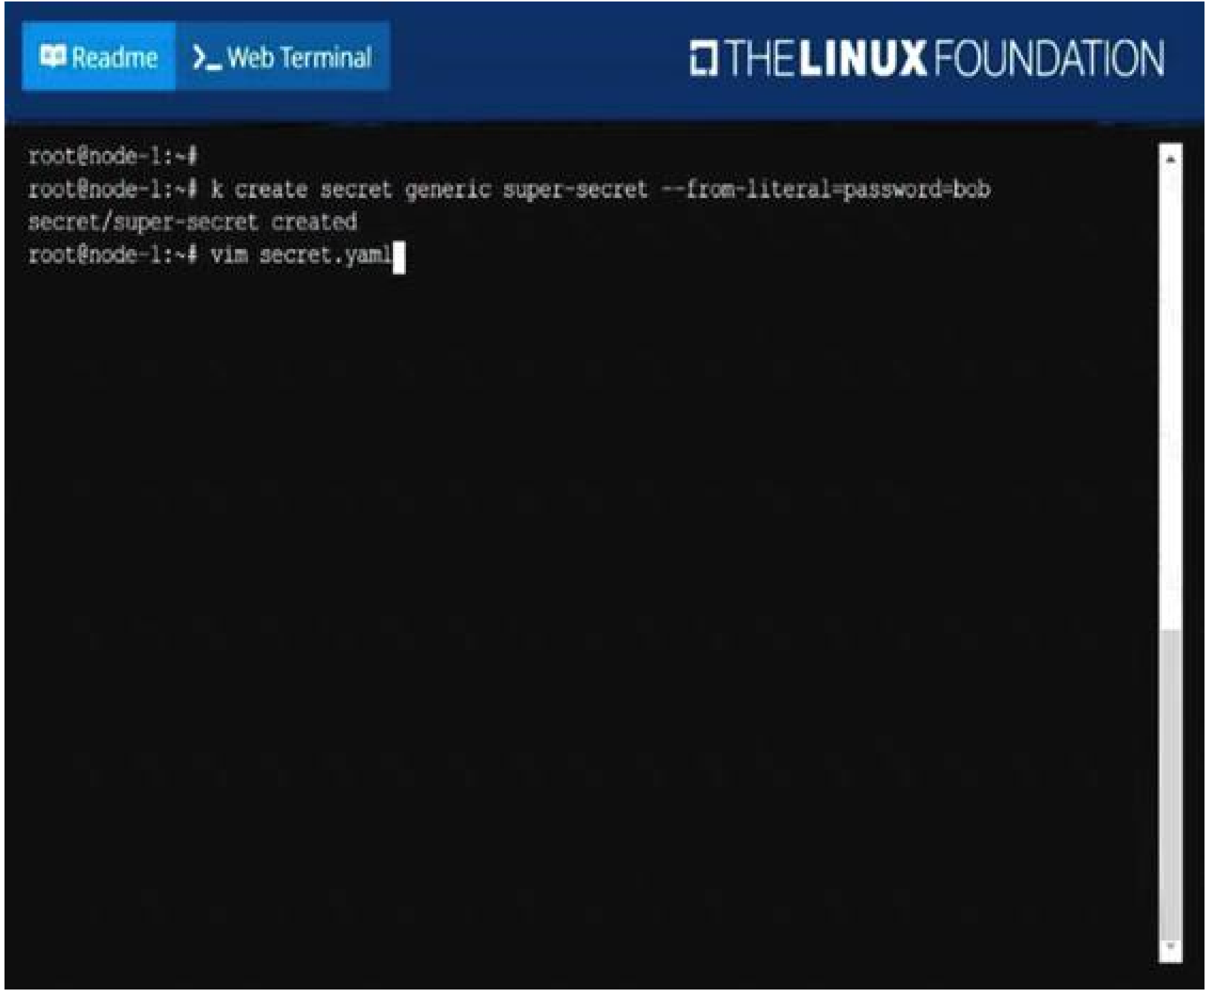

```
root@node-1:~#
root@node-1:~# k create secret generic super-secret --from-literal=password=bob
secret/super-secret created
root@node-1:~# vim secret.yaml
```

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

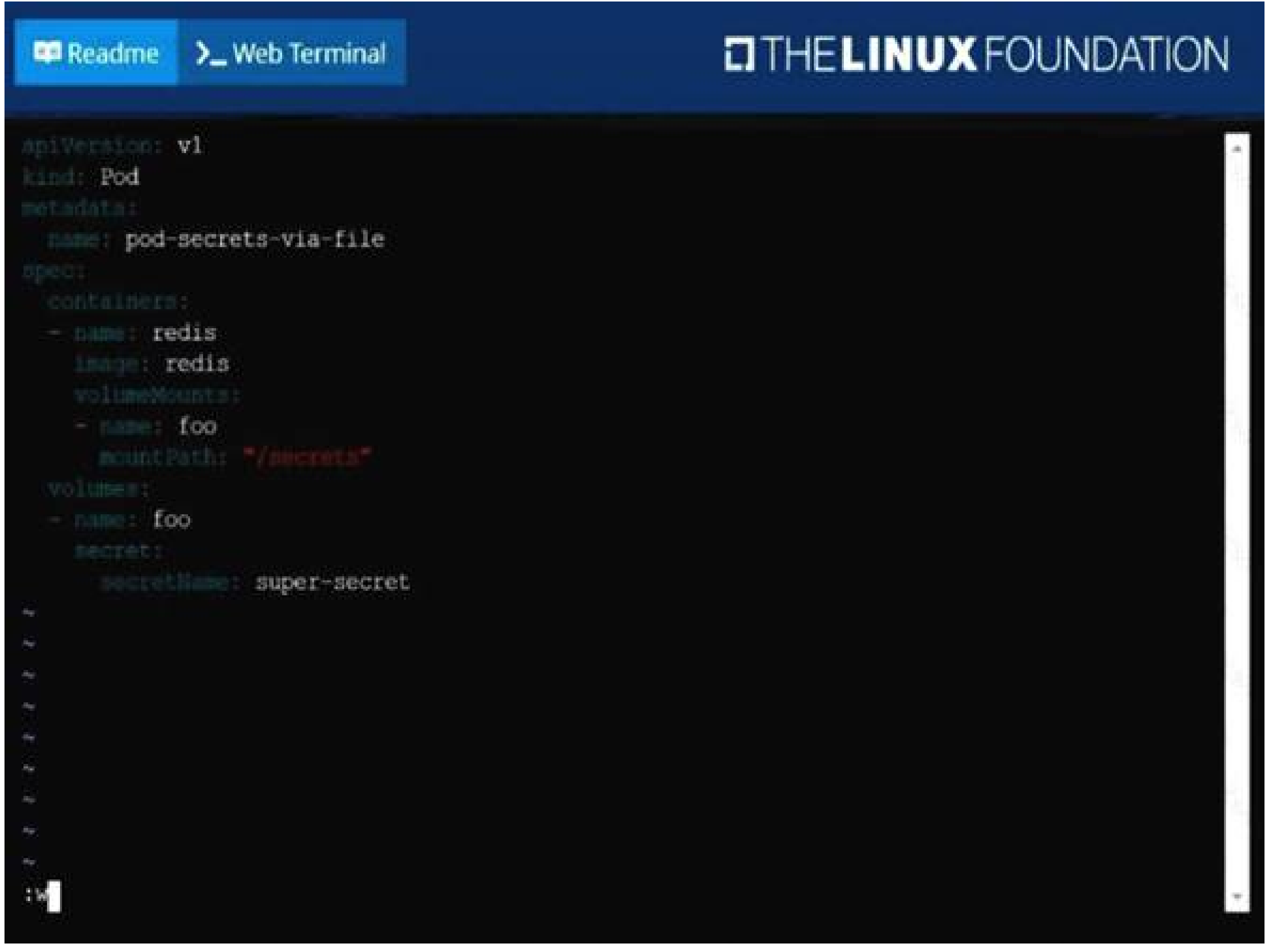

```
apiVersion: v1
kind: Pod
metadata:
  name: pod-secrets-via-file
spec:
  containers:
  - name: redis
    image: redis
    volumeMounts:
    - name: foo
      mountPath: "/secrets"
  volumes:
  - name: foo
    secret:
      secretName: super-secret
~
~
~
~
~
~
~
~
~
:w
```

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

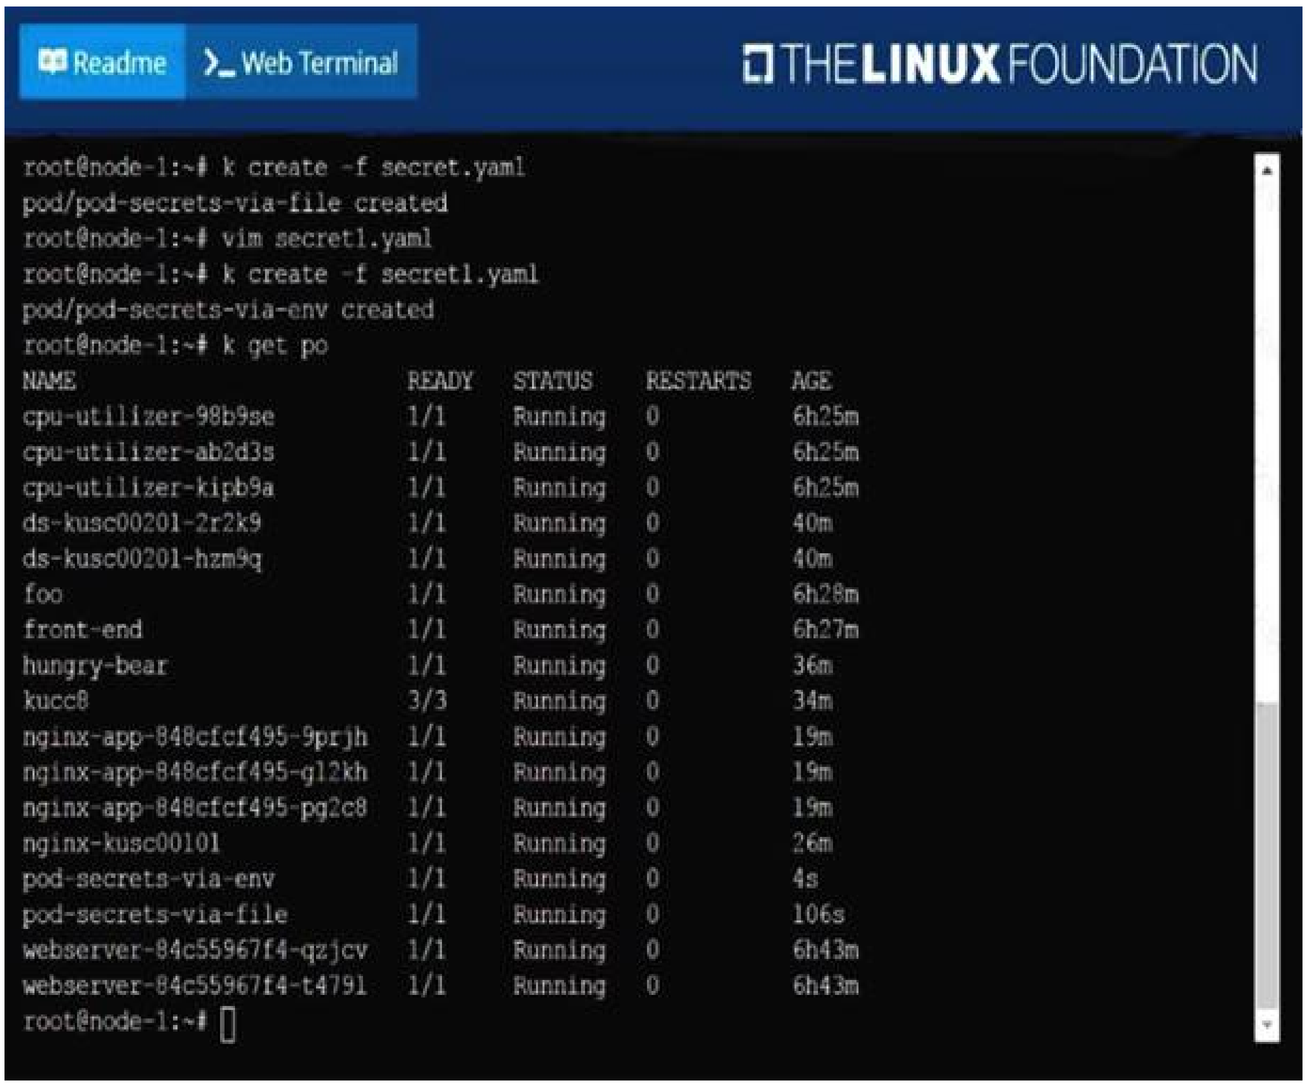

```
root@node-1:~# k create -f secret.yaml
pod/pod-secrets-via-file created
root@node-1:~# vim secret1.yaml
root@node-1:~# k create -f secret1.yaml
pod/pod-secrets-via-env created
root@node-1:~# k get po
NAME                          READY   STATUS    RESTARTS   AGE
cpu-utilizer-98b9se           1/1     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 20
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}{"\n"}'

NEW QUESTION 24
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

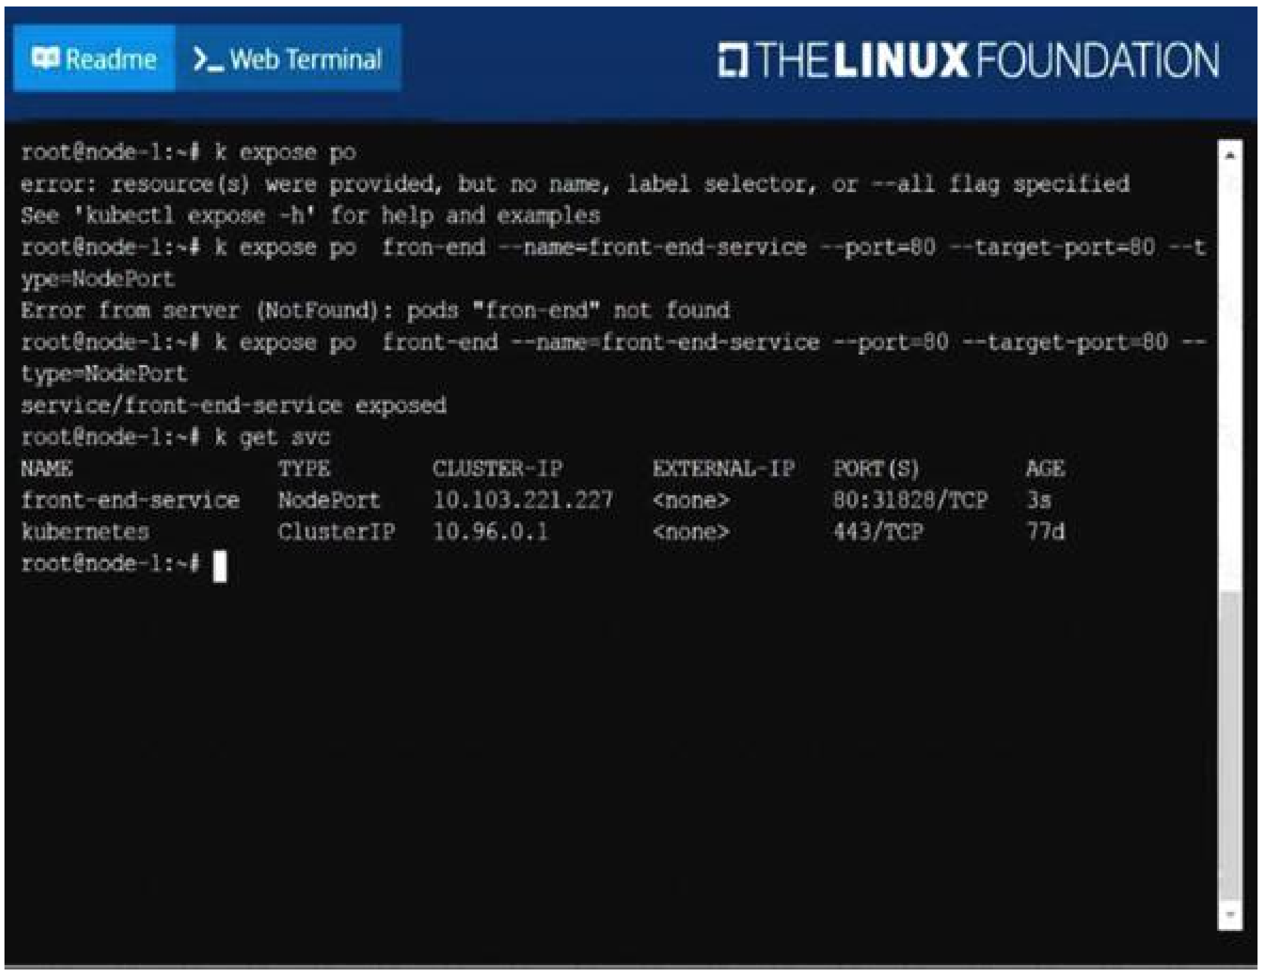

```
root@node-1:~# k expose po
error: resource(s) were provided, but no name, label selector, or --all flag specified
See 'kubectl expose -h' for help and examples
root@node-1:~# k expose po  fron-end --name=front-end-service --port=80 --target-port=80 --t
ype=NodePort
Error from server (NotFound): pods "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 28
For this item, you will havetosshto the nodesik8s-master-0andik8s-node-0and complete all tasks on thesenodes. Ensure that you return tothe base node (hostname:node-1) when you havecompleted this item.
Context
As an administrator of a smalldevelopment team, you have beenasked to set up a Kubernetes clusterto test the viability of a newapplication.
Task
You must usekubeadmto performthis task. Anykubeadminvocationswill require the use of the
--ignore-preflight-errors=alloption.

- Configure thenodeik8s-master-Oas a masternode. .
- Join the nodeik8s-node-otothe cluster.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution
You must use thekubeadmconfiguration file located at/etc/kubeadm.confwhen initializingyour cluster.
You may use any CNI pluginto complete this task, but ifyou don't have your favouriteCNI plugin's manifest URL athand, Calico is one popularoption:https://docs.projectcalico.org/v3.14/manifests/calico.yaml
Docker is already installedon both nodes andapthasbeen configured so that you caninstall the required tools.

NEW QUESTION 32
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 35
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 36
Create a snapshot of theetcdinstance running athttps://127.0.0.1:2379, saving thesnapshot to the file path /srv/data/etcd-snapshot.db.
The following TLScertificates/key are suppliedfor connecting to the server withetcdctl:

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

A. Mastered
B. Not Mastered

Answer: A

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

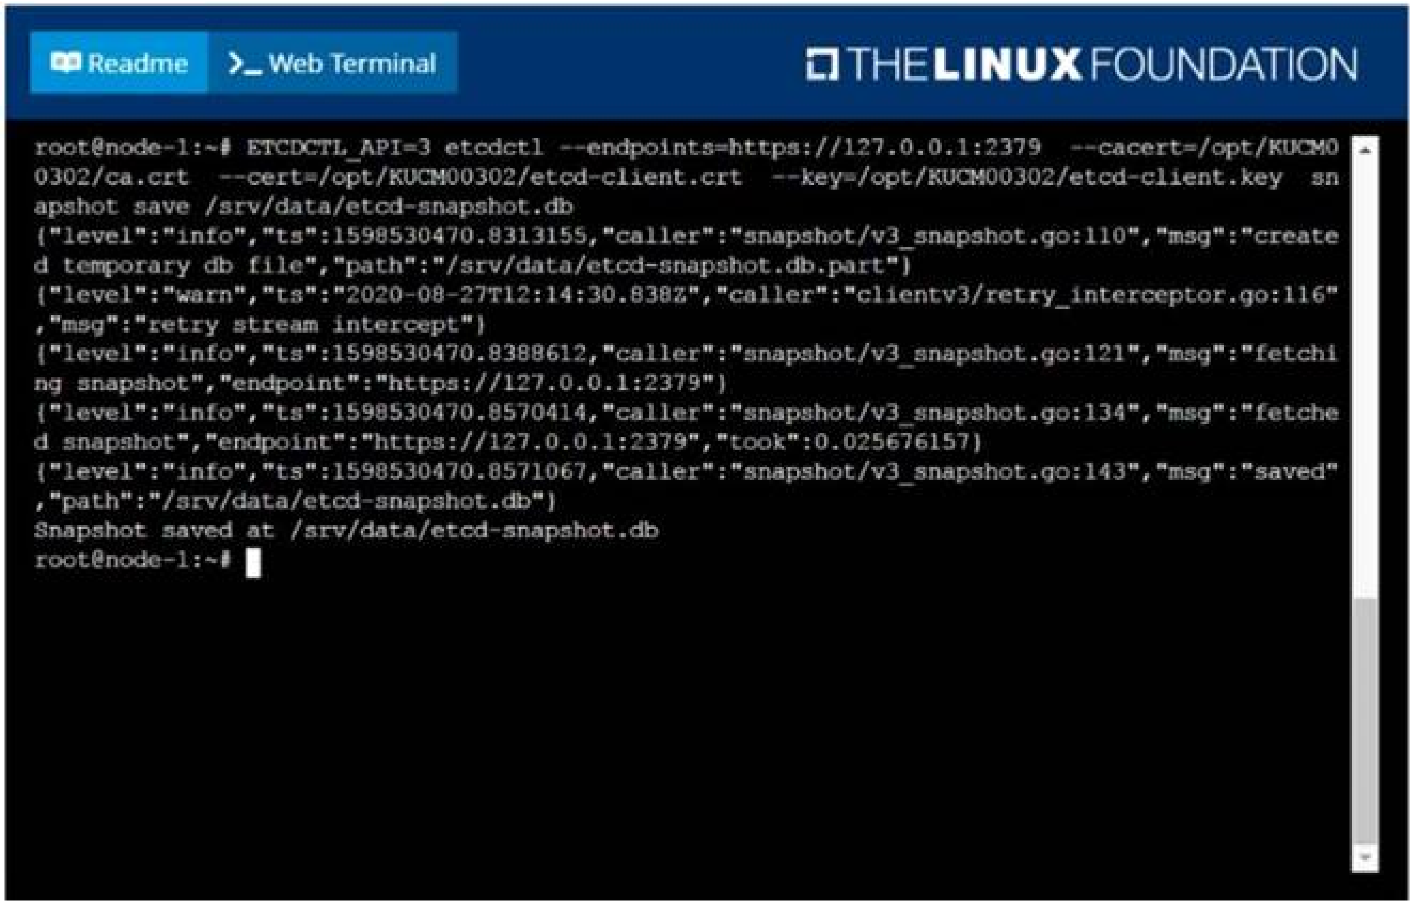

```
root@node-1:~# ETCDCTL_API=3 etcdctl --endpoints=https://127.0.0.1:2379  --cacert=/opt/KUCM0
0302/ca.crt  --cert=/opt/KUCM00302/etcd-client.crt  --key=/opt/KUCM00302/etcd-client.key  sn
apshot save /srv/data/etcd-snapshot.db
{"level":"info","ts":1598530470.8313155,"caller":"snapshot/v3_snapshot.go:110","msg":"create
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":"fetchi
ng snapshot","endpoint":"https://127.0.0.1:2379"}
{"level":"info","ts":1598530470.8570414,"caller":"snapshot/v3_snapshot.go:134","msg":"fetche
d 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:~#
```

NEW QUESTION 39
Scale the deploymentwebserverto6pods.

A. Mastered
B. Not Mastered

Answer: A

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

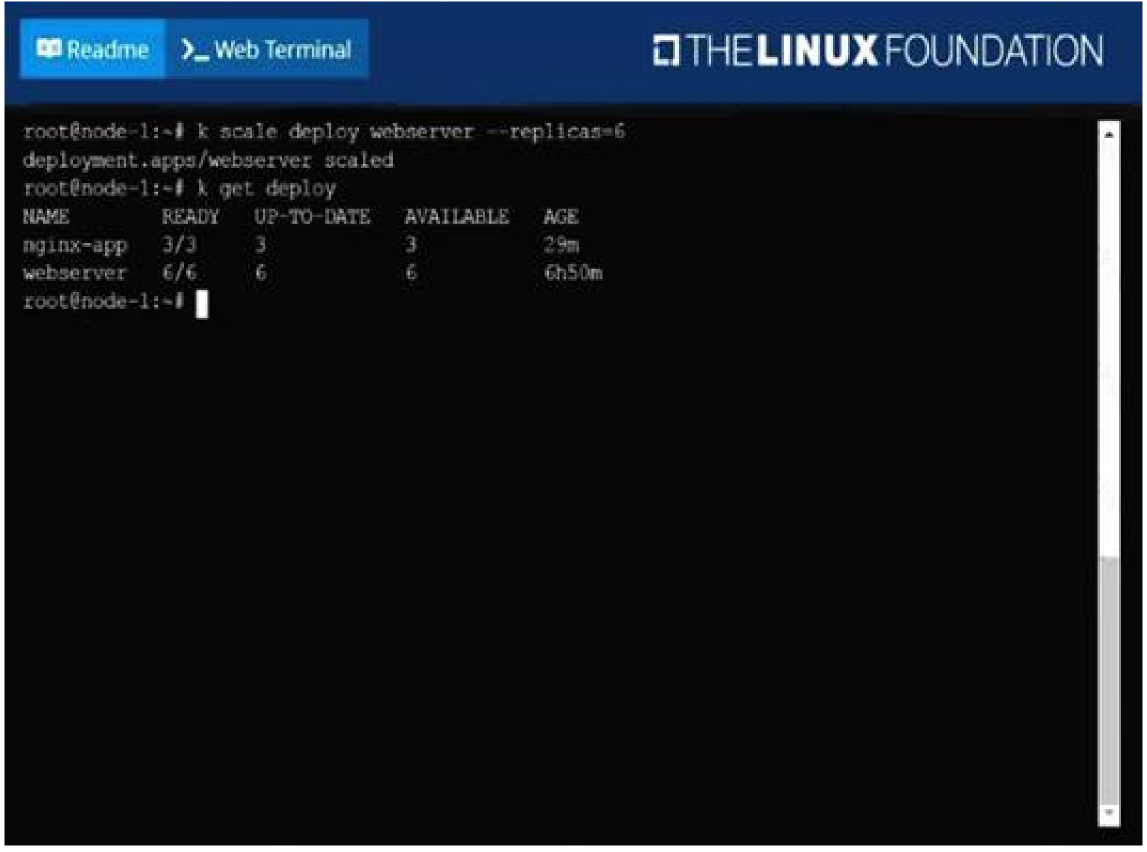

```
root@node-1:~# k scale deploy webserver --replicas=6
deployment.apps/webserver scaled
root@node-1:~# k get deploy
NAME        READY   UP-TO-DATE   AVAILABLE   AGE
nginx-app   3/3     3            3           29m
webserver   6/6     6            6           6h50m
root@node-1:~#
```

NEW QUESTION 43
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 46
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 50
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

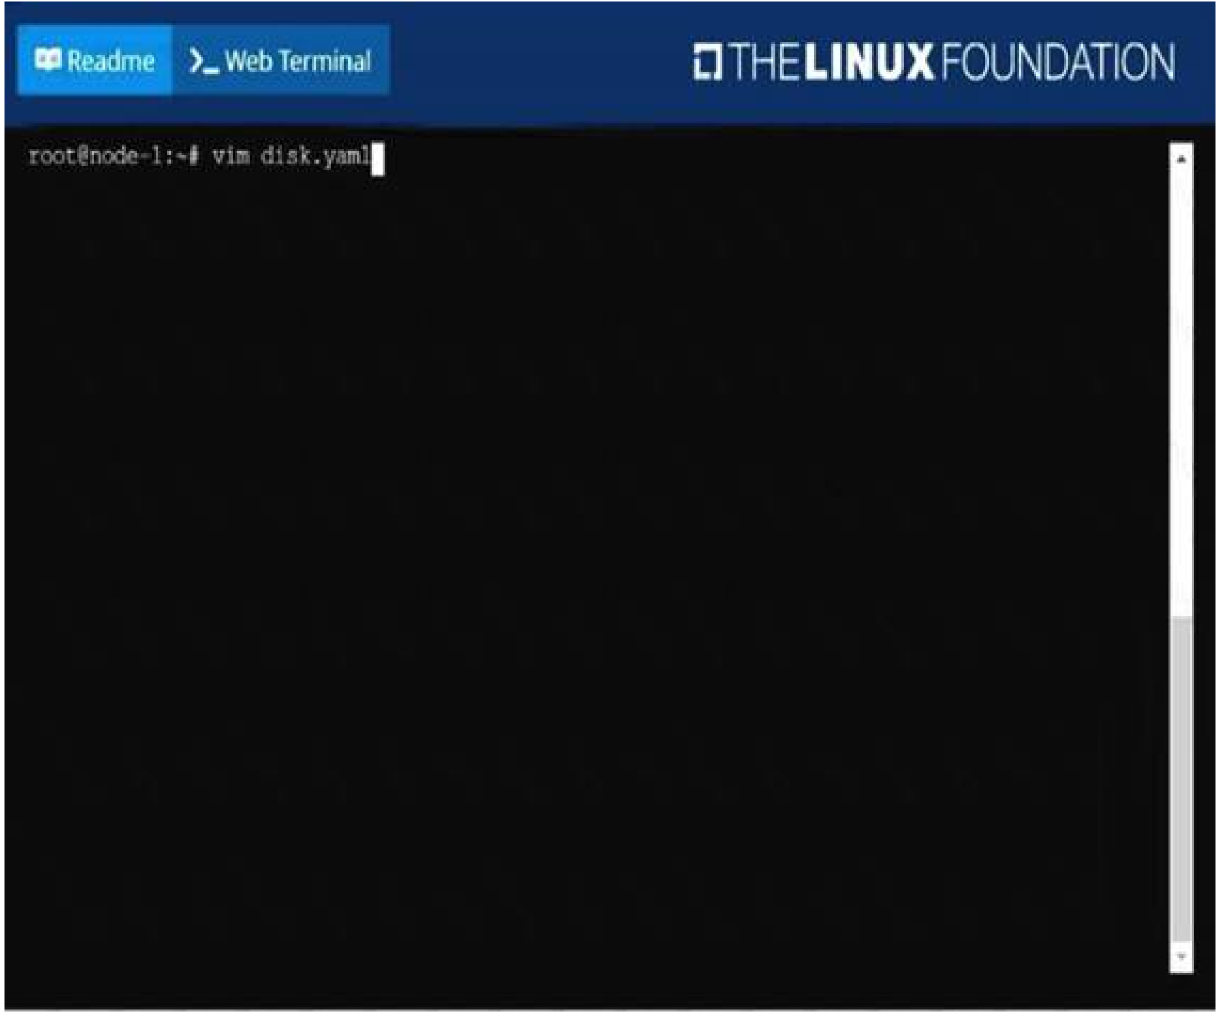

```
root@node-1:~# vim disk.yaml
```

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

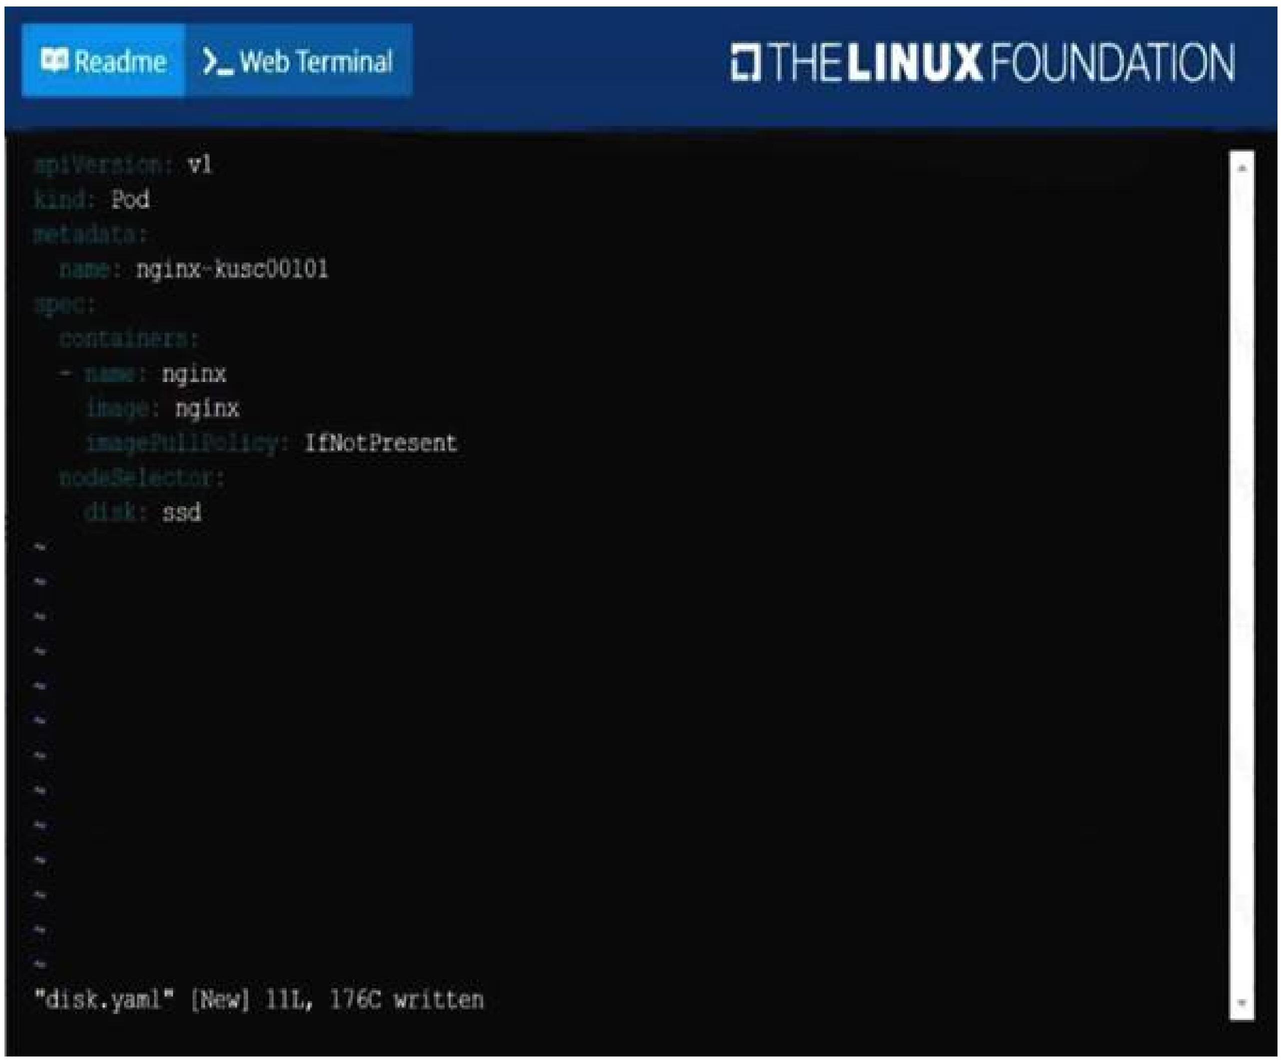

```
apiVersion: v1
kind: Pod
metadata:
  name: nginx-kusc00101
spec:
  containers:
  - name: nginx
    image: nginx
    imagePullPolicy: IfNotPresent
  nodeSelector:
    disk: ssd
~
~
~
~
~
~
~
~
~
~
~
~
~
"disk.yaml" [New] 11L, 176C written
```

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

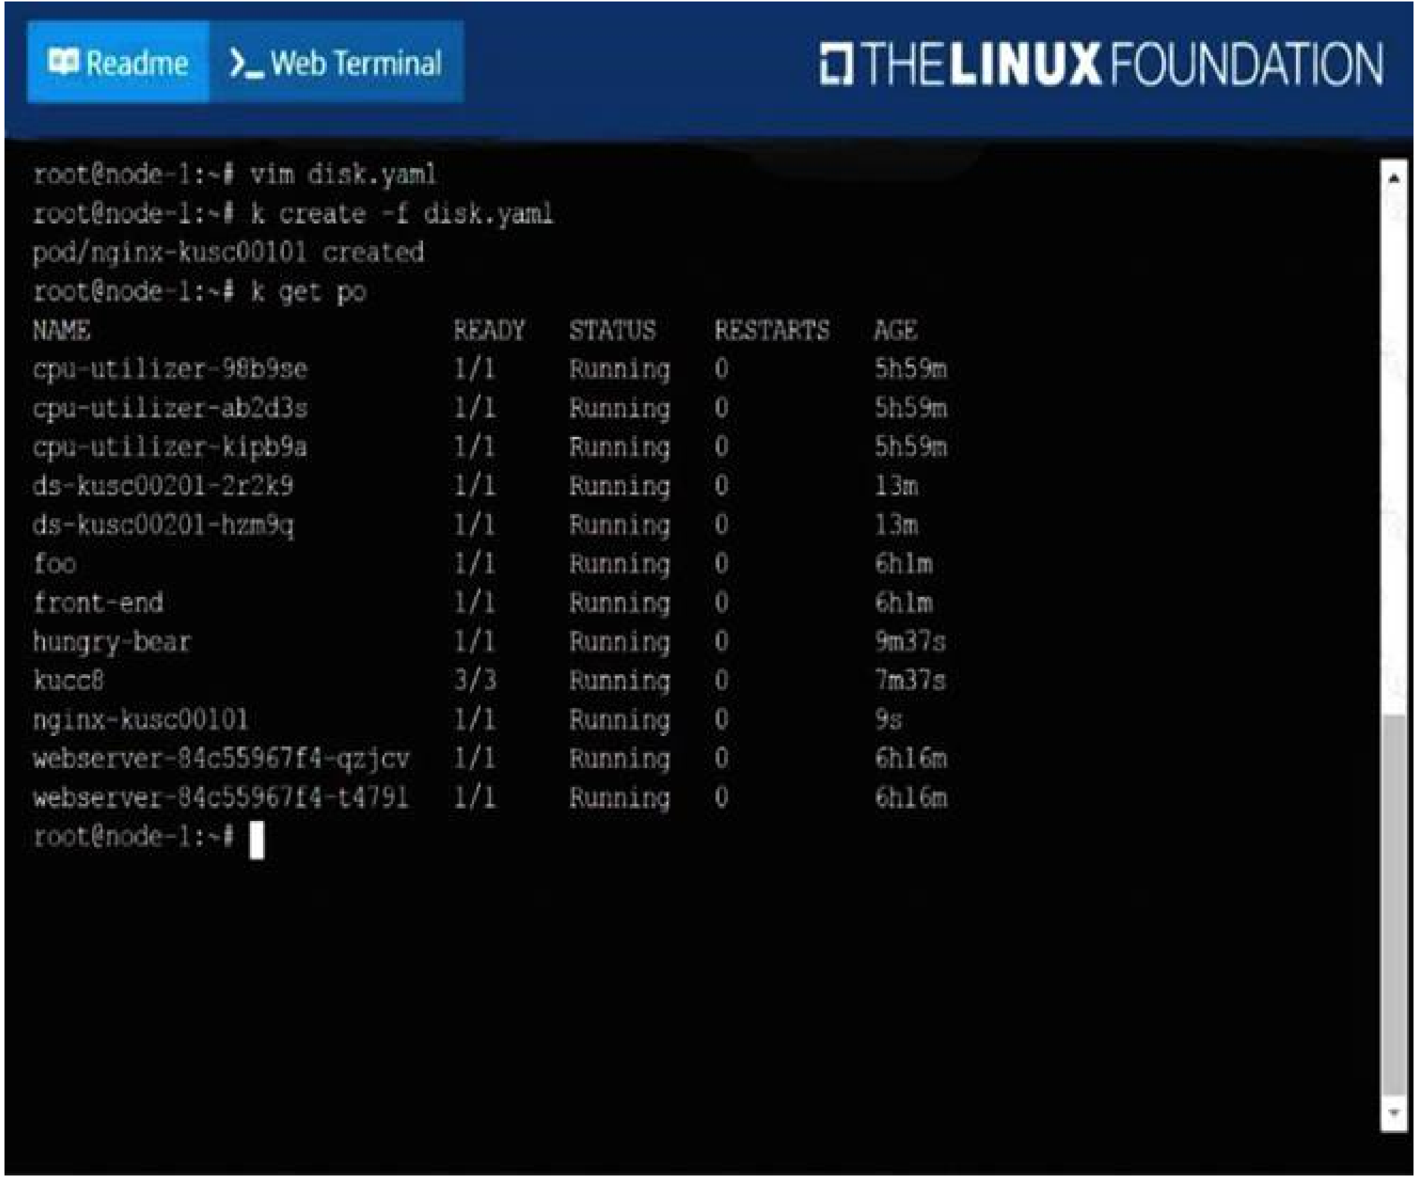

```
root@node-1:~# vim disk.yaml
root@node-1:~# k create -f disk.yaml
pod/nginx-kusc00101 created
root@node-1:~# k get po
NAME                           READY   STATUS    RESTARTS   AGE
cpu-utilizer-98b9se            1/1     Running   0          5h59m
cpu-utilizer-ab2d3s            1/1     Running   0          5h59m
cpu-utilizer-kipb9a            1/1     Running   0          5h59m
ds-kusc00201-2r2k9             1/1     Running   0          13m
ds-kusc00201-hzm9q             1/1     Running   0          13m
foo                            1/1     Running   0          6h1m
front-end                      1/1     Running   0          6h1m
hungry-bear                    1/1     Running   0          9m37s
kucc8                          3/3     Running   0          7m37s
nginx-kusc00101                1/1     Running   0          9s
webserver-84c55967f4-qzjcv     1/1     Running   0          6h16m
webserver-84c55967f4-t4791     1/1     Running   0          6h16m
root@node-1:~# 
```

NEW QUESTION 55
Create an nginx pod and list the pod with different levels of verbosity

A. Mastered
B. Not Mastered

Answer: A

Explanation:
// create a pod
kubectl run nginx --image=nginx --restart=Never --port=80
// List the pod with different verbosity kubectl get po nginx --v=7
kubectl get po nginx --v=8 kubectl get po nginx --v=9

NEW QUESTION 56
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

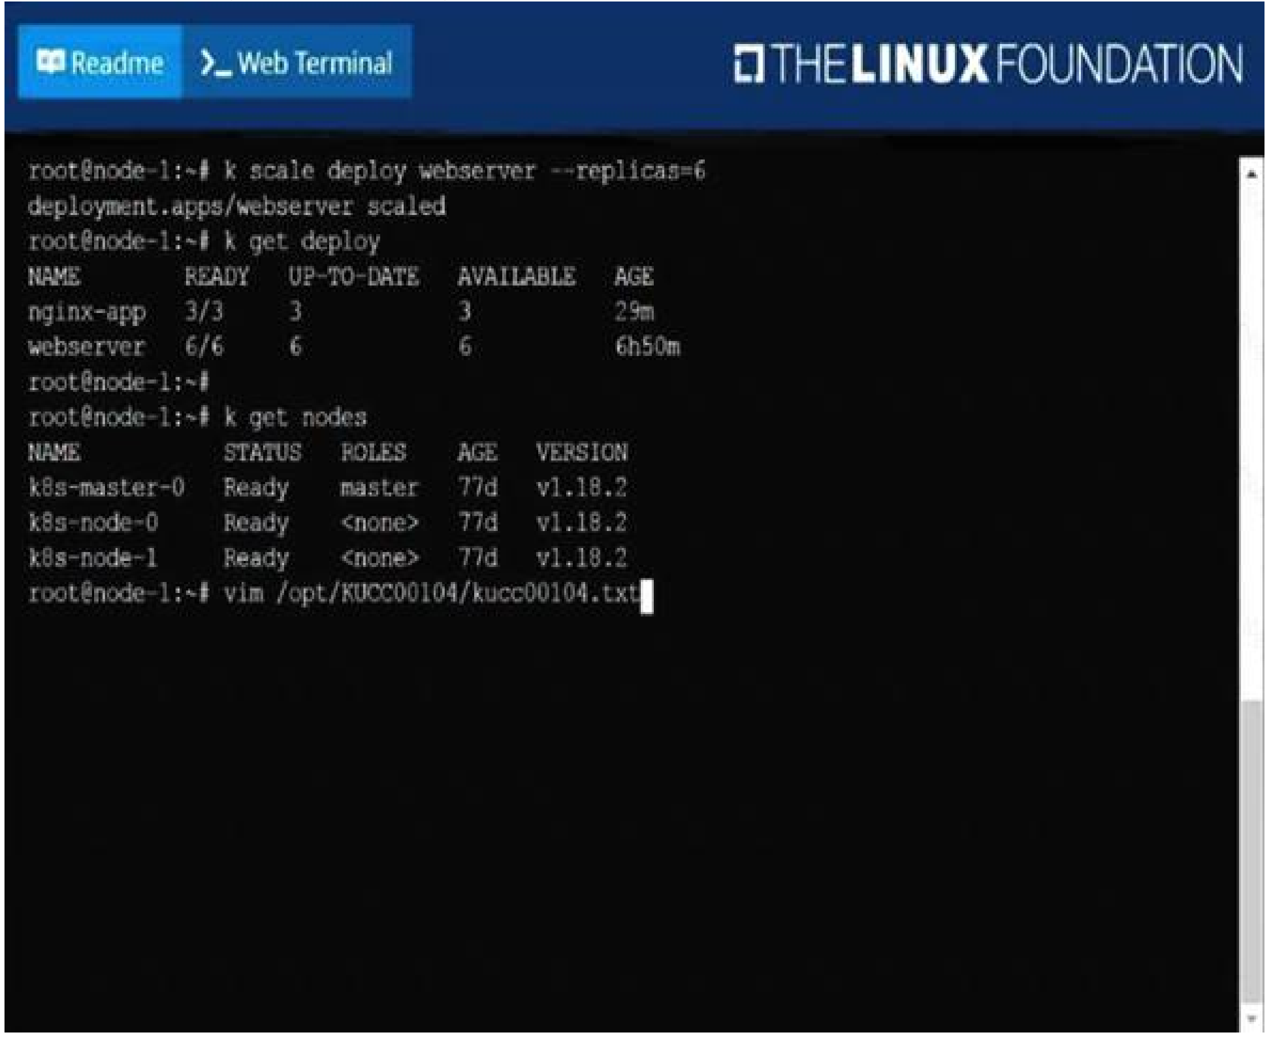

```
root@node-1:~# k scale deploy webserver --replicas=6
deployment.apps/webserver scaled
root@node-1:~# k get deploy
NAME         READY   UP-TO-DATE   AVAILABLE   AGE
nginx-app    3/3     3            3           29m
webserver    6/6     6            6           6h50m
root@node-1:~#
root@node-1:~# k get nodes
NAME           STATUS   ROLES    AGE   VERSION
k8s-master-0   Ready    master   77d   v1.18.2
k8s-node-0     Ready    <none>   77d   v1.18.2
k8s-node-1     Ready    <none>   77d   v1.18.2
root@node-1:~# vim /opt/KUCC00104/kucc00104.txt
```

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

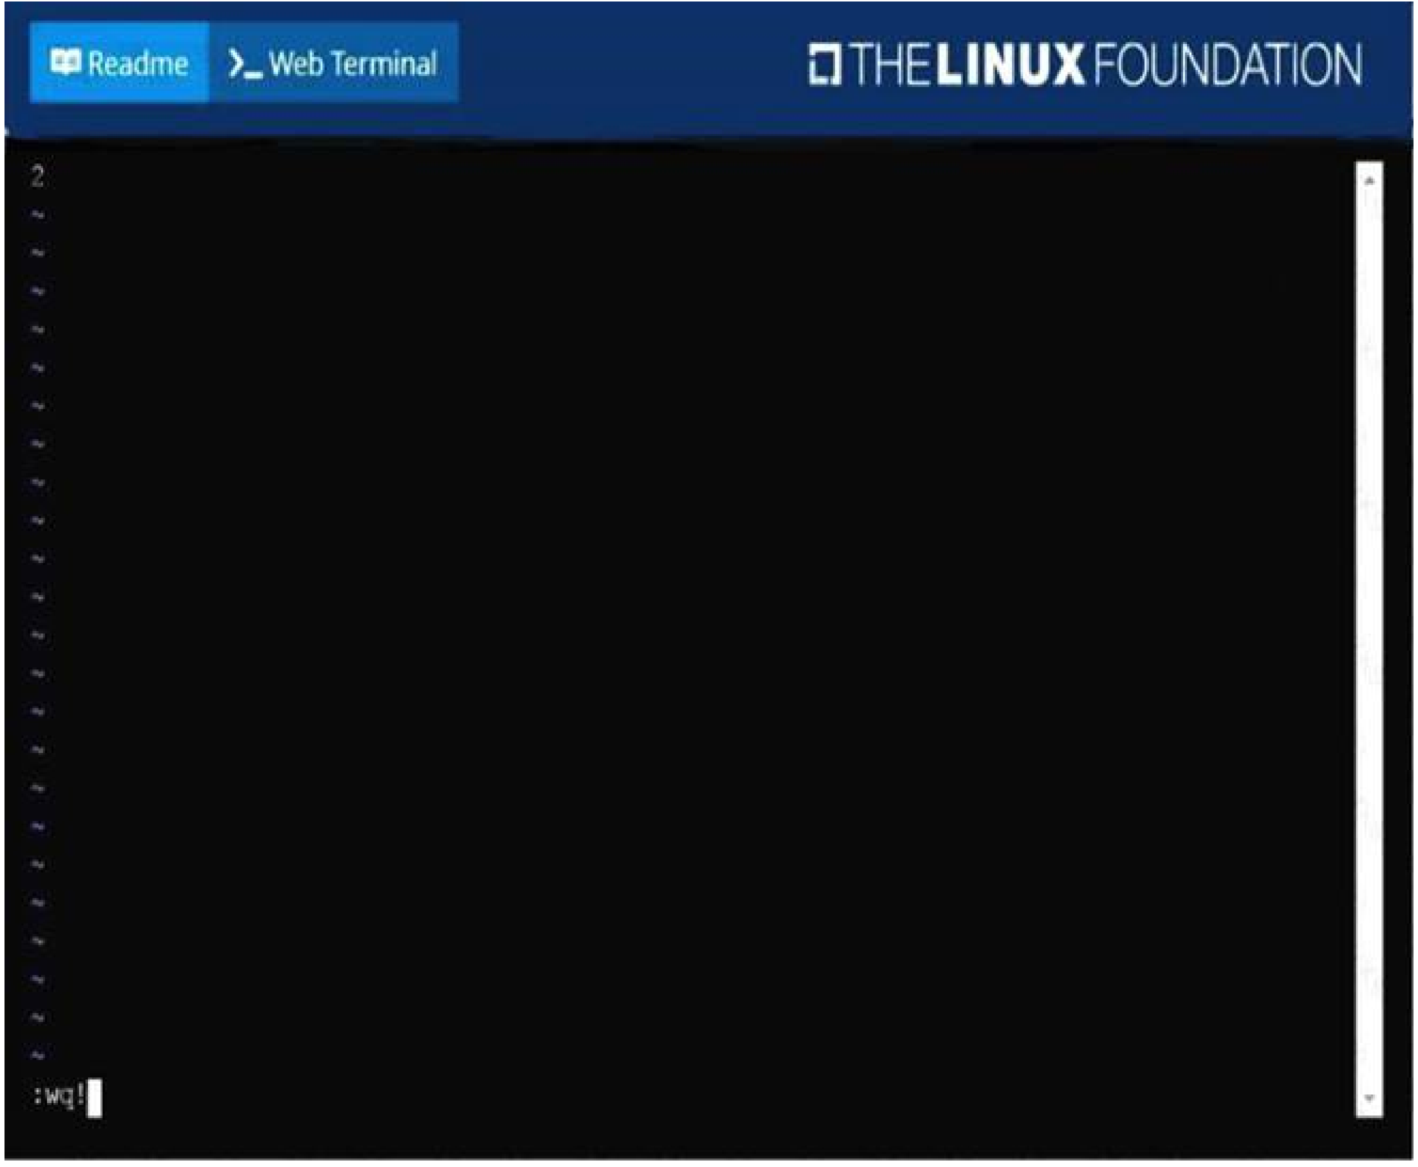

```
2
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
~
:wq!
```

NEW QUESTION 61
......