CKA Dumps

Certified Kubernetes Administrator (CKA) Program

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

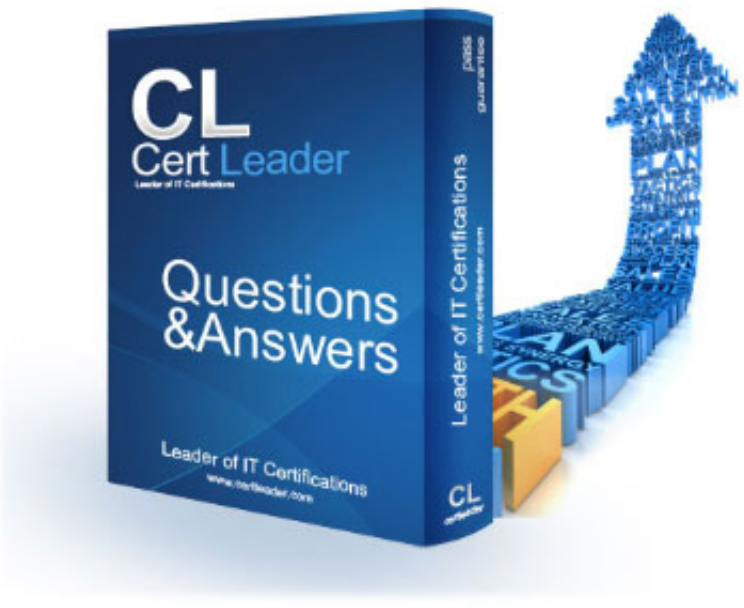

NEW QUESTION 1
CORRECT TEXT
Create a pod with image nginx called nginx and allow traffic on port 80

A. Mastered
B. Not Mastered

Answer: A

Explanation:
kubectl run nginx --image=nginx --restart=Never --port=80

NEW QUESTION 2
CORRECT TEXT
Score: 4%

Task
Check to see how many nodes are ready (not including nodes tainted NoSchedule) and write the number to /opt/KUSC00402/kusc00402.txt.

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Solution:
kubectl describe nodes | grep ready|wc -l
kubectl describe nodes | grep -i taint | grep -i noschedule |wc -l
echo 3 > /opt/KUSC00402/kusc00402.txt
#
kubectl get node | grep -i ready |wc -l
taintsnoSchedule
kubectl describe nodes | grep -i taints | grep -i noschedule |wc -l
#
echo 2 > /opt/KUSC00402/kusc00402.txt

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
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

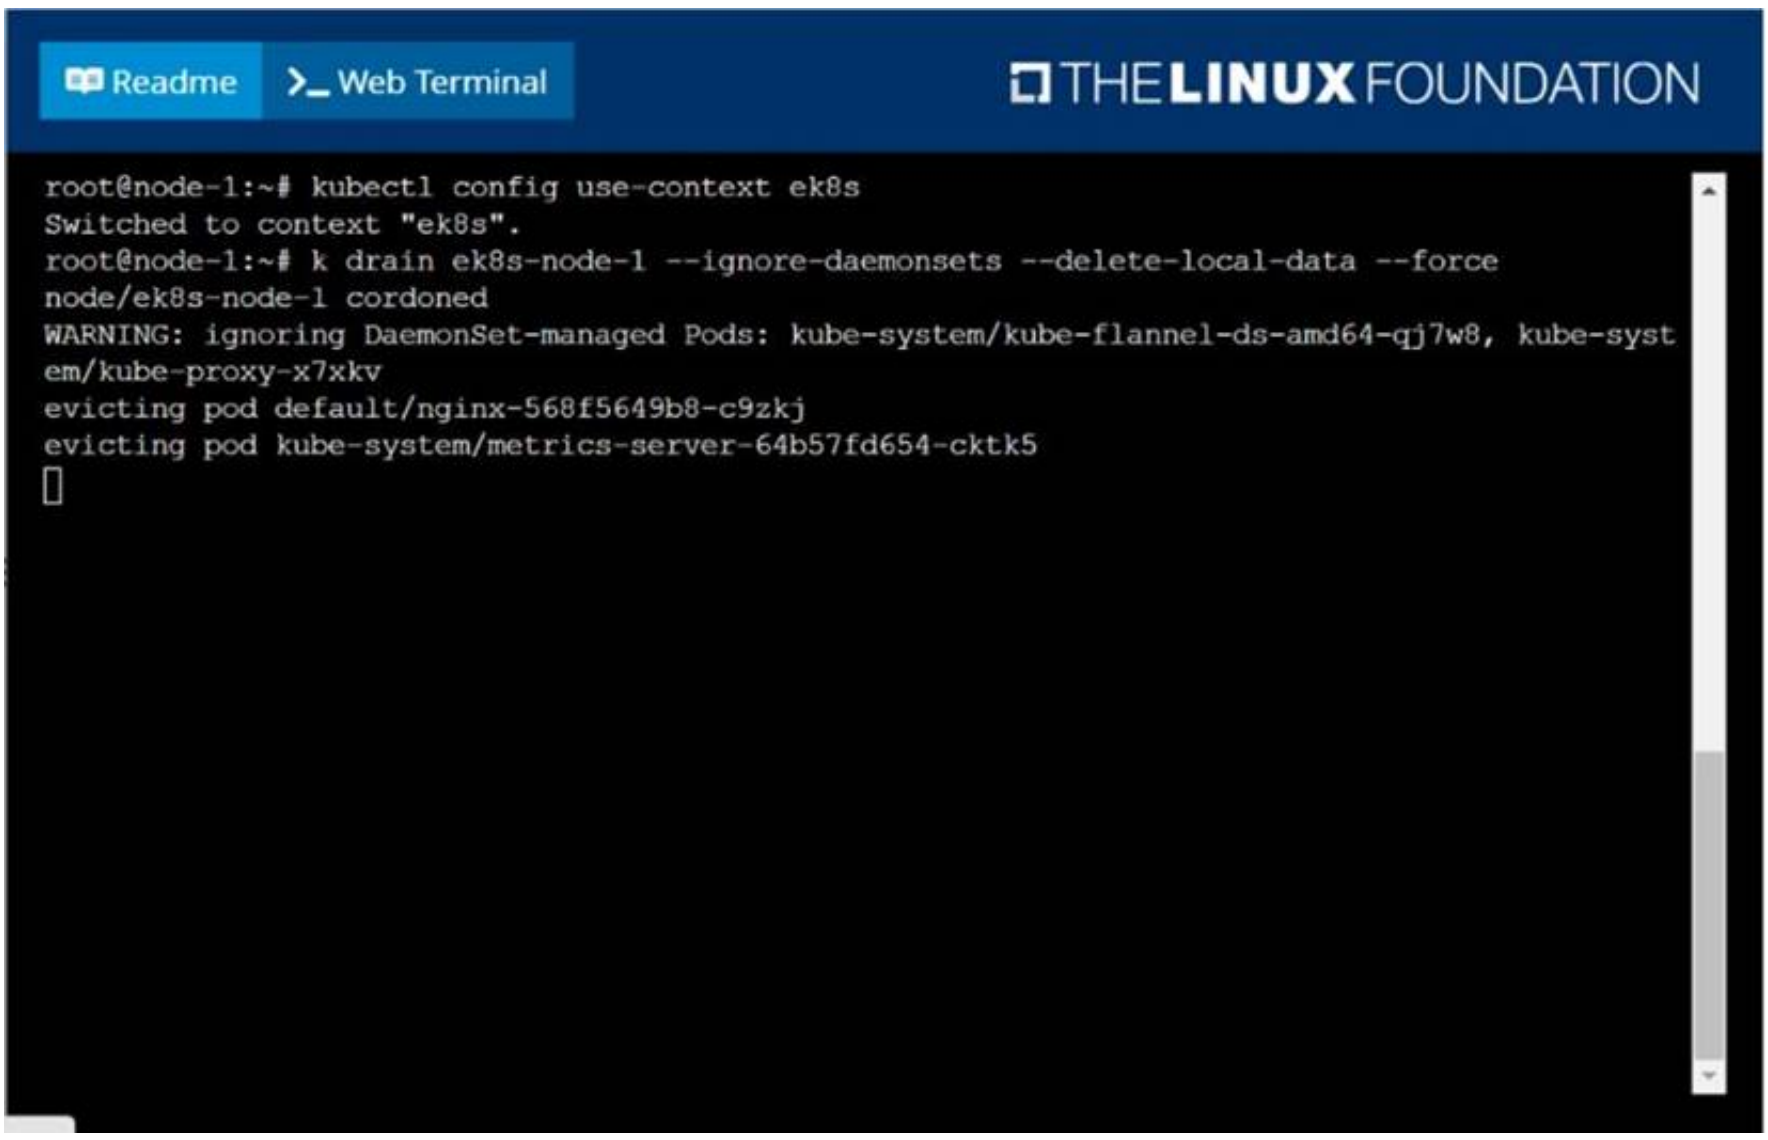

F:\Work\Data Entry Work\Data Entry\20200827\CKA\19 B.JPG

NEW QUESTION 6
CORRECT TEXT
Scale the deployment webserver to 6 pods.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution

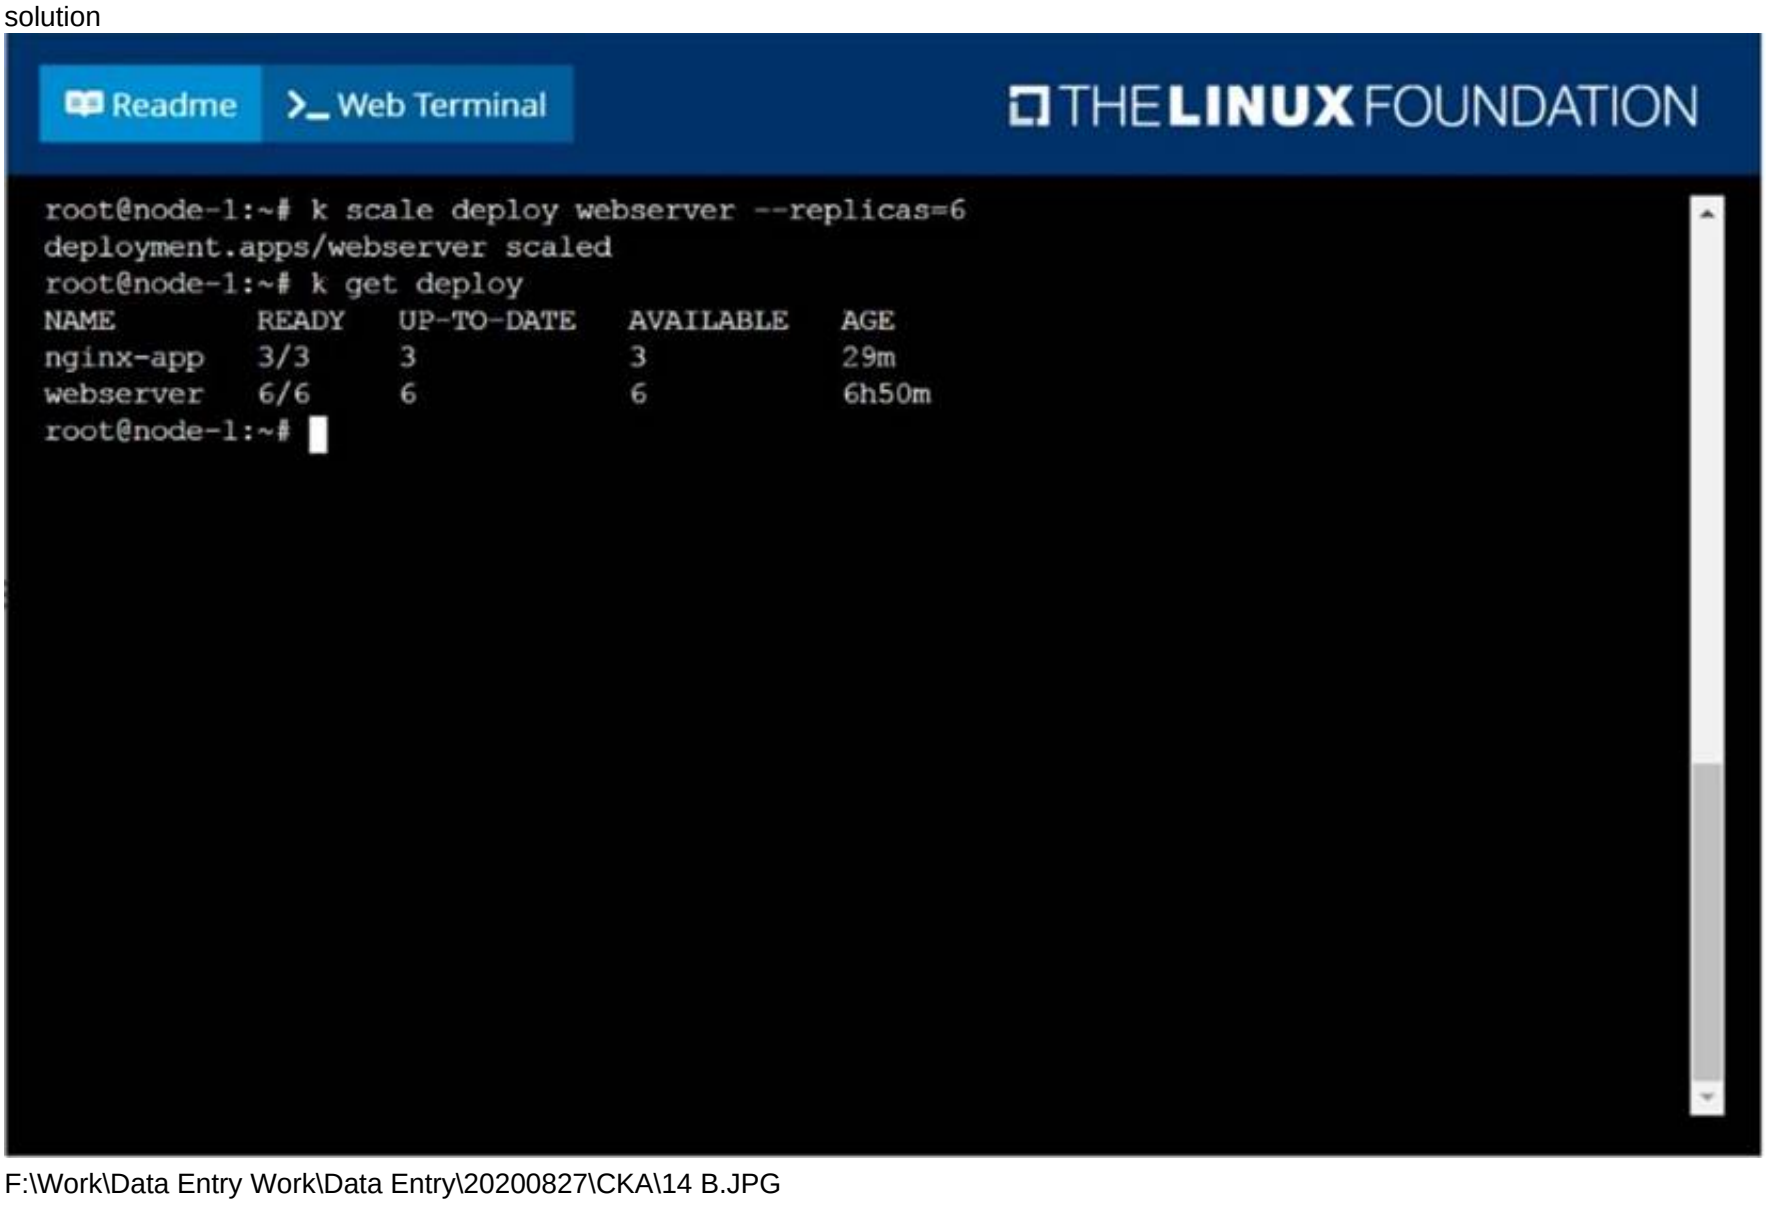

F:\Work\Data Entry Work\Data Entry\20200827\CKA\14 B.JPG

NEW QUESTION 7

CORRECT TEXT
Check the image version in pod without the describe command

A. Mastered
B. Not Mastered

Answer: A

Explanation:
kubectl get po nginx -o
jsonpath='{.spec.containers[].image}{"\n"}'

NEW QUESTION 8
CORRECT TEXT
Create a pod as follows:
? Name: mongo
? Using Image: mongo
? In a new Kubernetes namespace named: my-website

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution

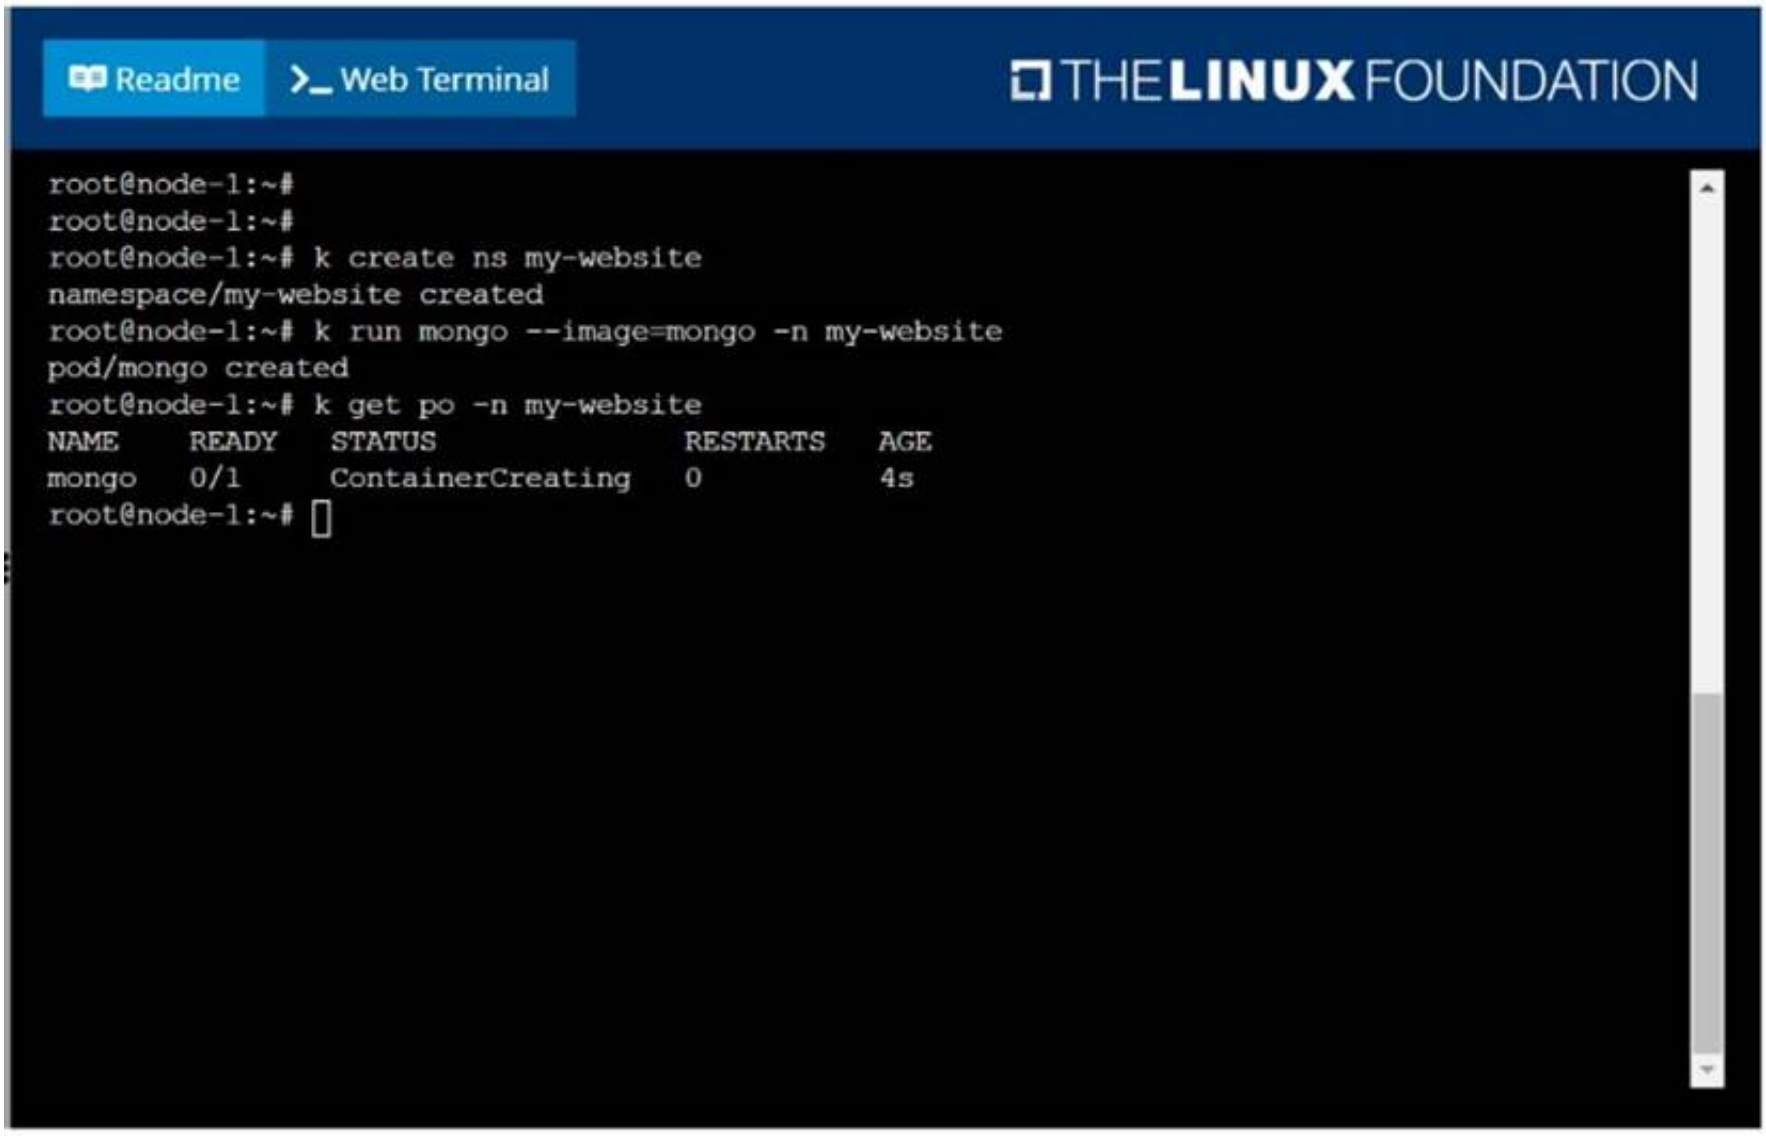

F:\Work\Data Entry Work\Data Entry\20200827\CKA\9 B.JPG

NEW QUESTION 9
CORRECT TEXT
Create a pod that echo “hello world” and then exists. Have the pod deleted automatically when it’s completed

A. Mastered
B. Not Mastered

Answer: A

Explanation:
kubectl run busybox --image=busybox -it --rm --restart=Never --
/bin/sh -c 'echo hello world'
kubectl get po # You shouldn't see pod with the name "busybox"

NEW QUESTION 10
CORRECT TEXT
Ensure a single instance of pod nginx is running on each node of the Kubernetes cluster where nginx also represents the Image name which has to be used. Do not override any taints currently in place.
Use DaemonSet to complete this task and use ds-kusc00201 as DaemonSet name.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution

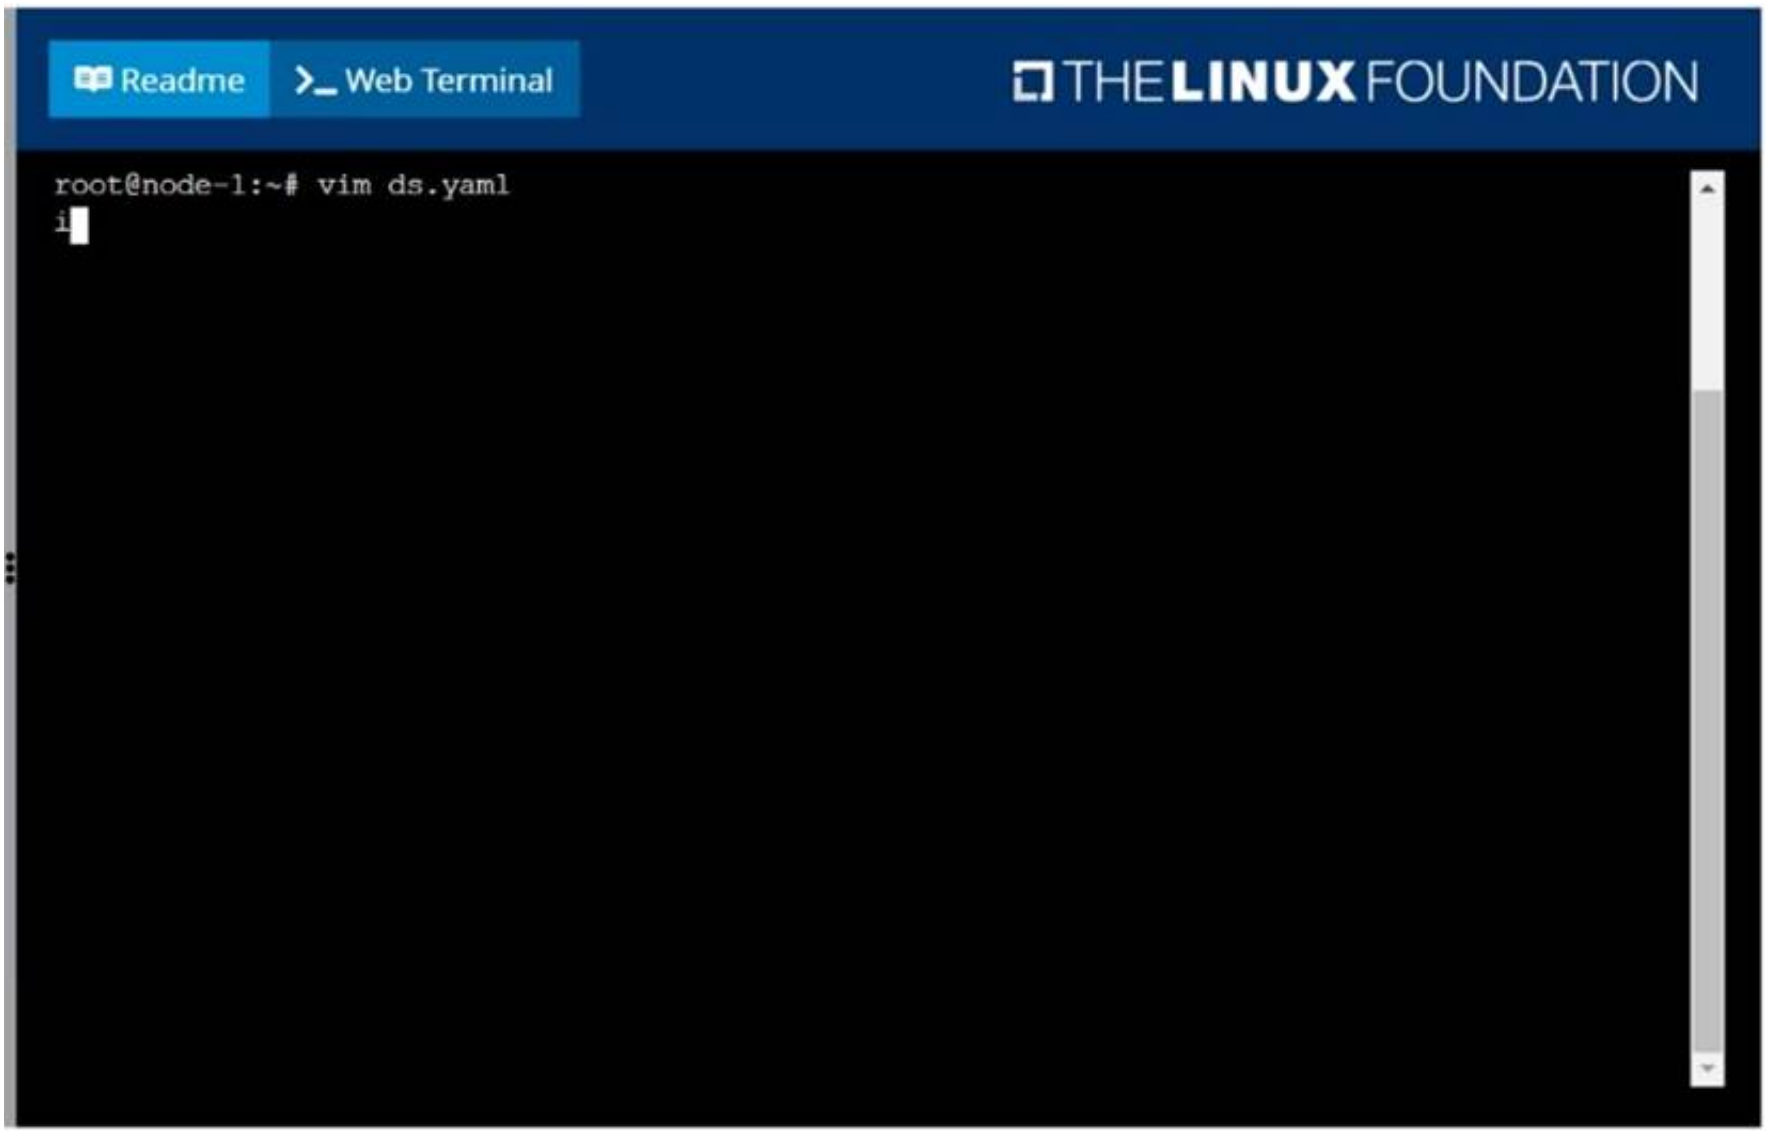

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 B.JPG

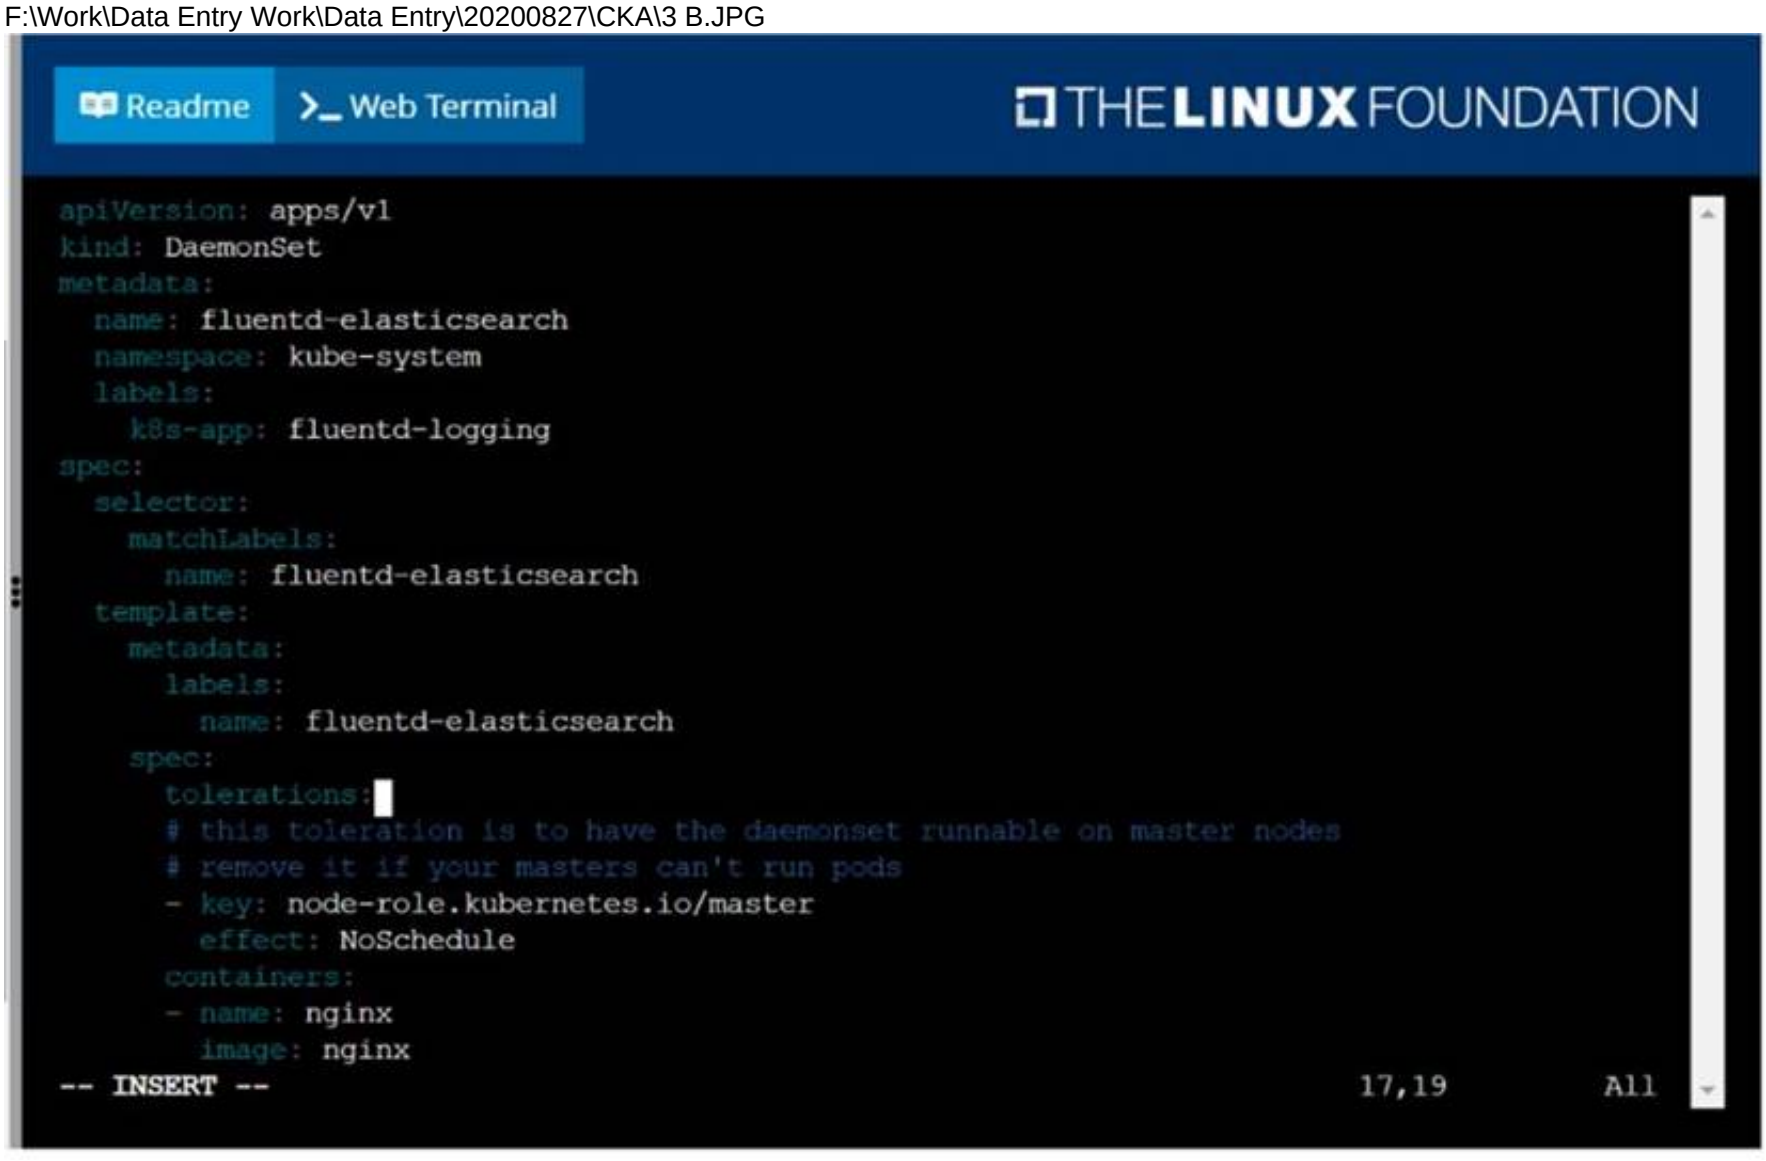

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 C.JPG

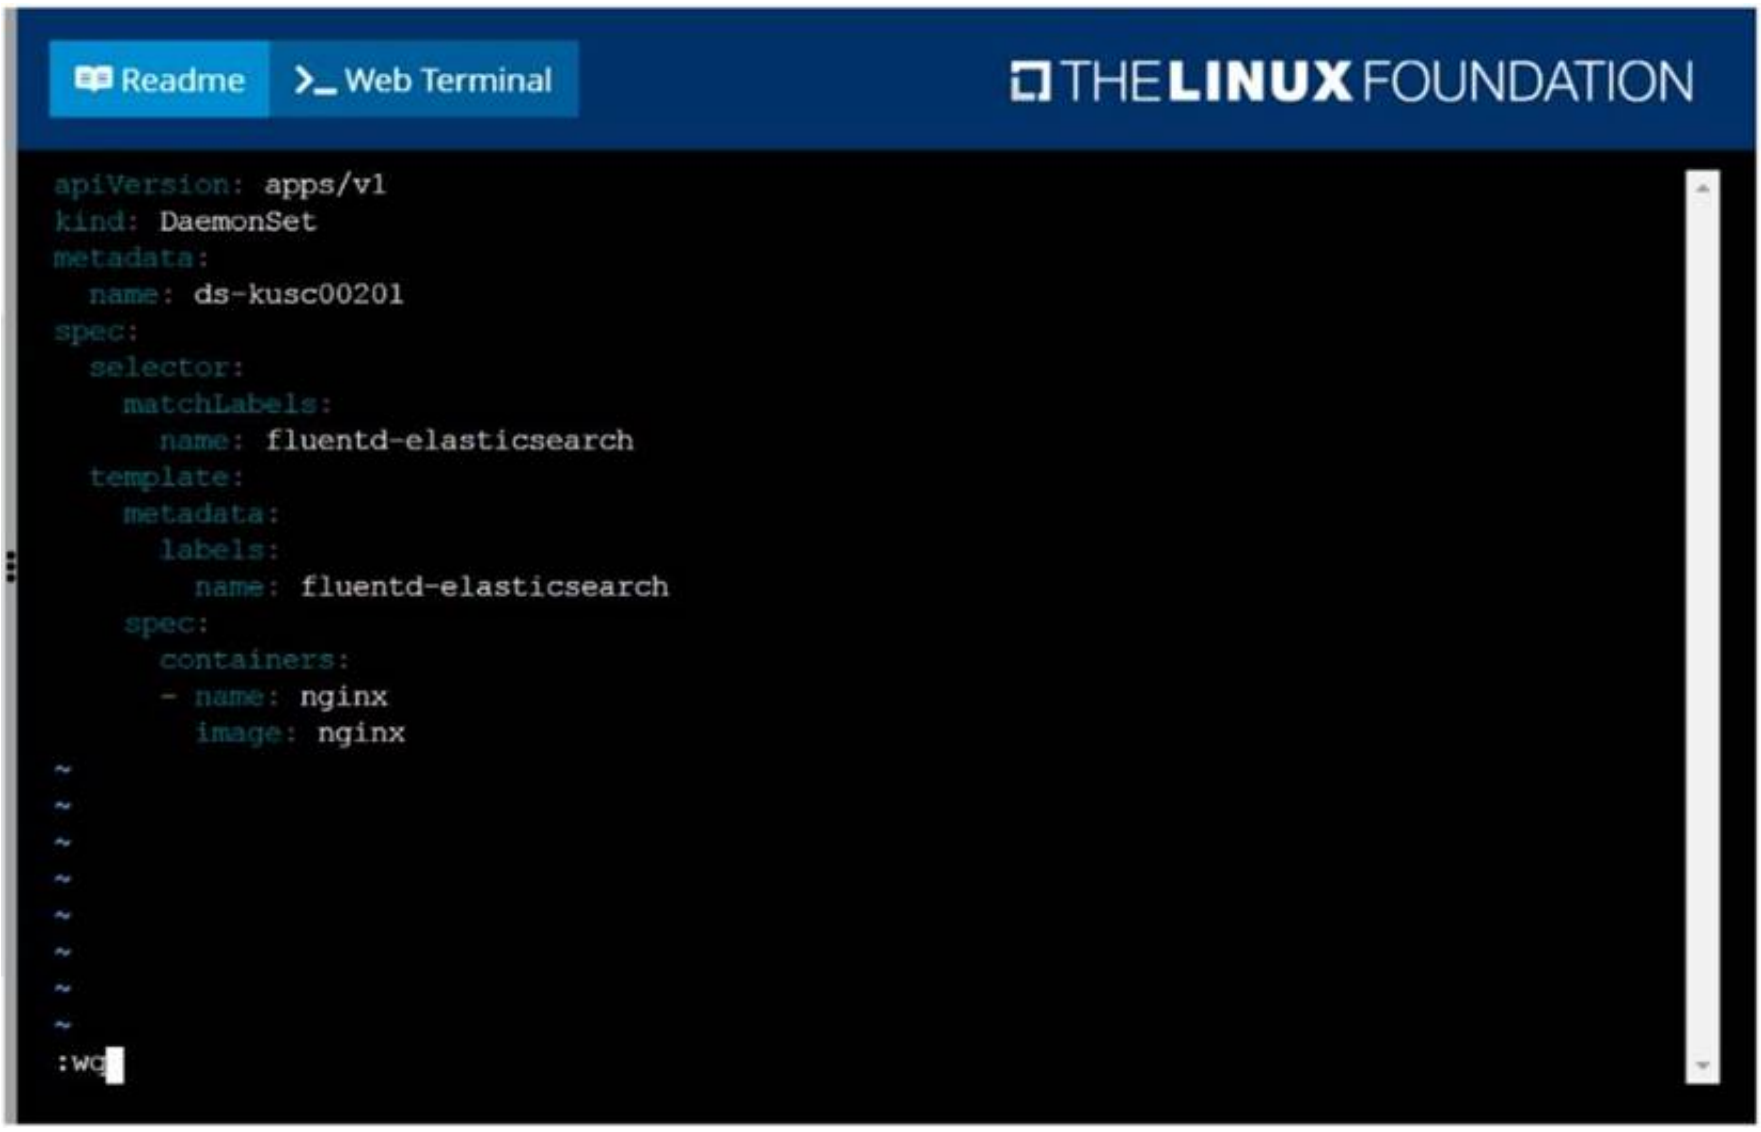

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 D.JPG

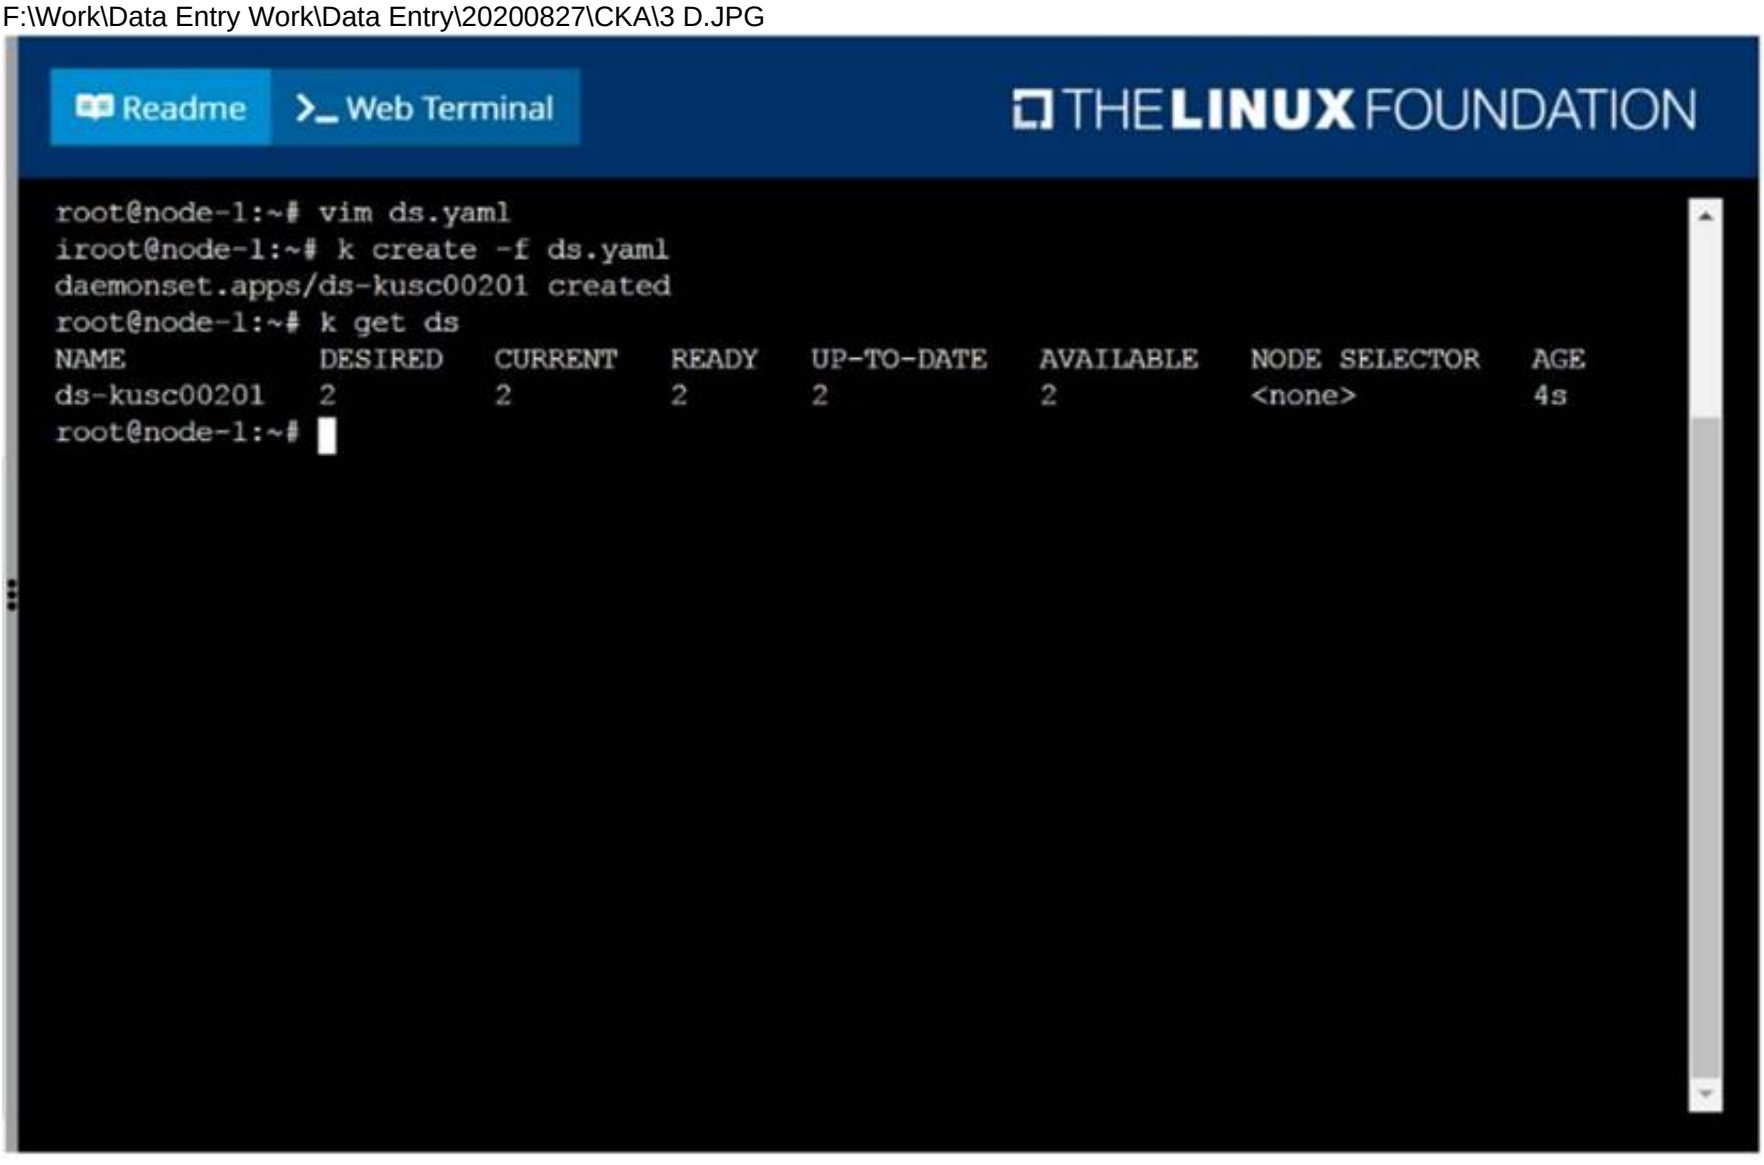

F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 E.JPG

NEW QUESTION 10
CORRECT TEXT
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 12
CORRECT TEXT
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.

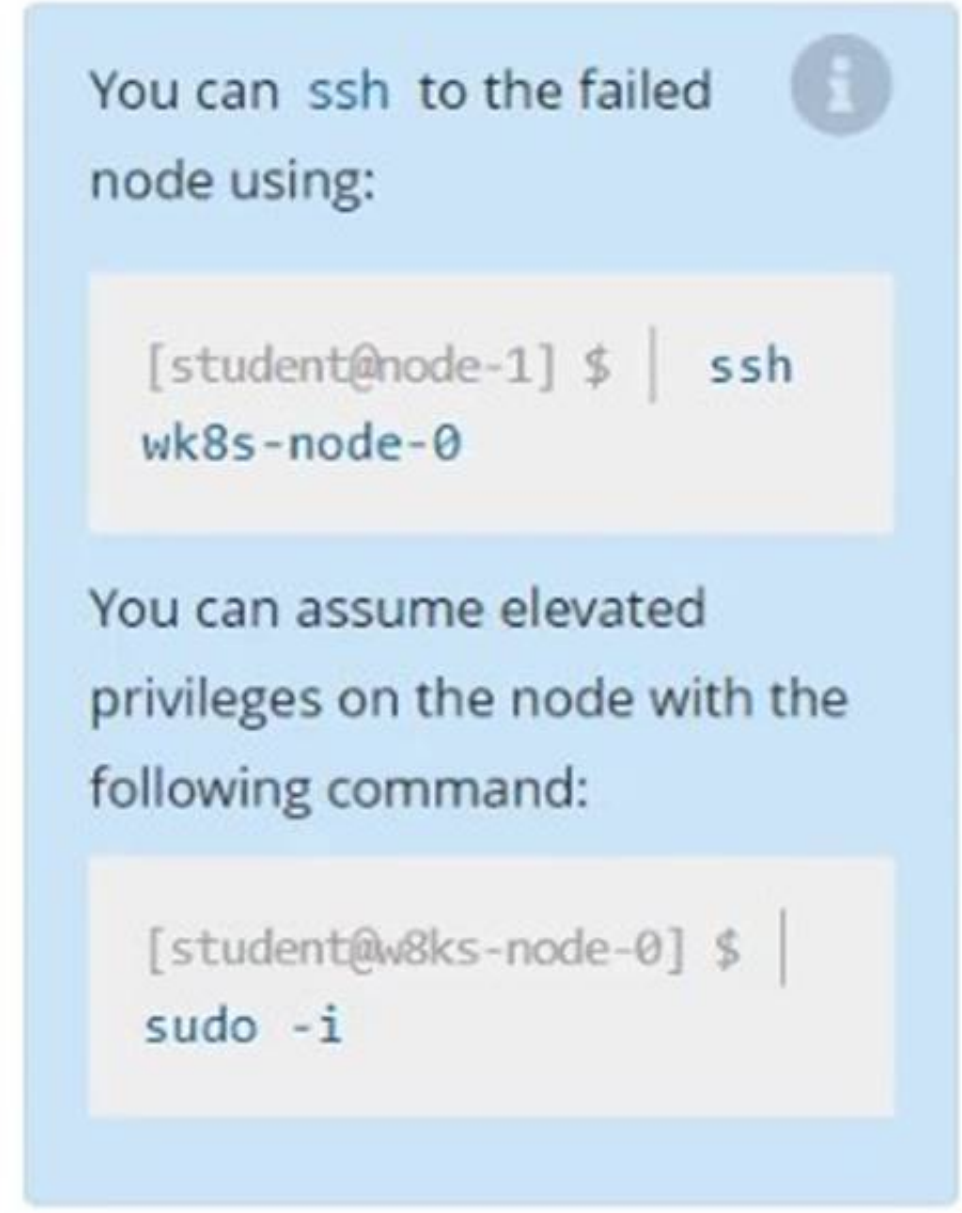

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Solution:
sudo -i
systemctl status kubelet
systemctl start kubelet
systemctl enable kubelet

NEW QUESTION 15
CORRECT TEXT
Score: 4%

Task
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:
kubectl run kucc8 --image=nginx --dry-run -o yaml > kucc8.yaml
vi kucc8.yaml
apiVersion: v1
kind: Pod
metadata:
creationTimestamp: null
name: kucc8
spec:
containers:
- image: nginx
name: nginx
- image: redis
name: redis
- image: memcached
name: memcached
- image: consul
name: consul
#
kubectl create -f kucc8.yaml
#12.07

NEW QUESTION 20
CORRECT TEXT
Task Weight: 4%

Task
Scale the deployment webserver to 3 pods.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Solution:

```
student@node-1:~$ kubectl  scale  deploy webserver  --replicas=3
deployment.apps/webserver scaled
student@node-1:~$ kubectl  scale  deploy webserver  --replicas=3
```

NEW QUESTION 25
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 26
CORRECT TEXT
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 30
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

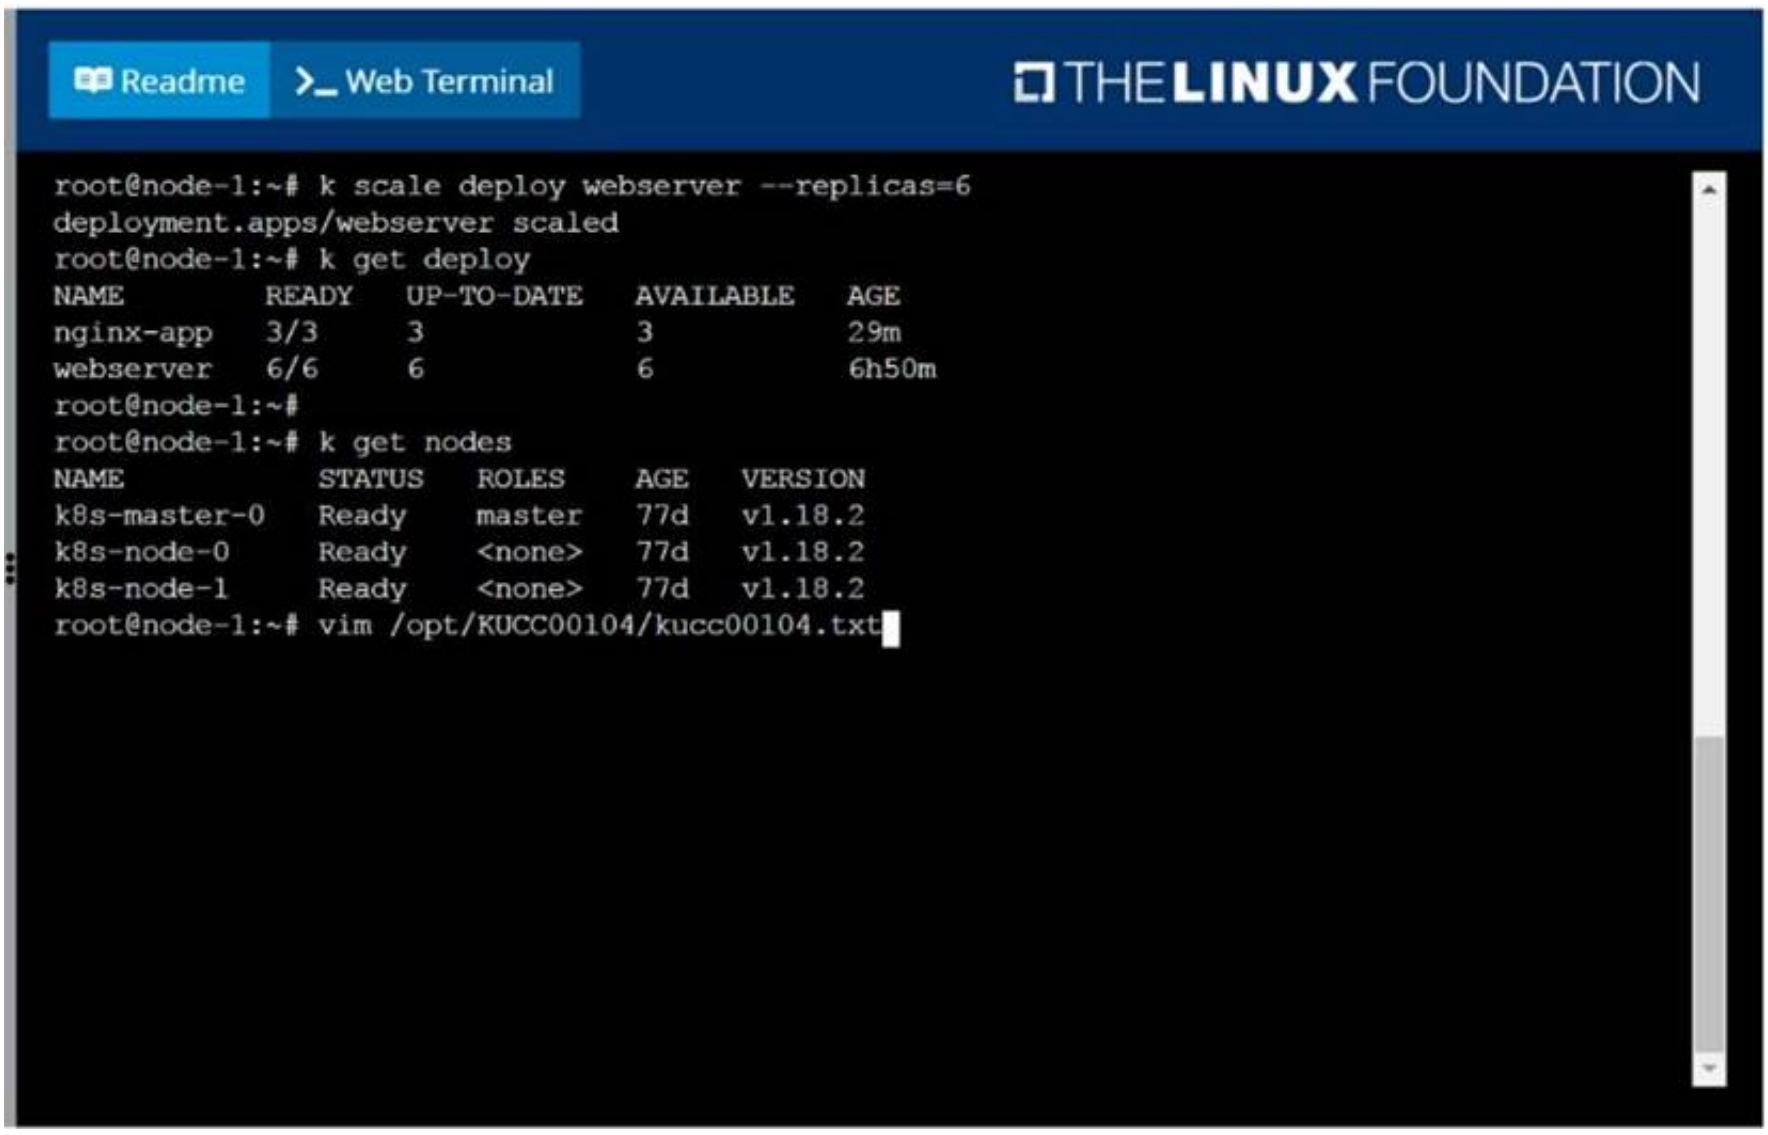

F:\Work\Data Entry Work\Data Entry\20200827\CKA\15 B.JPG

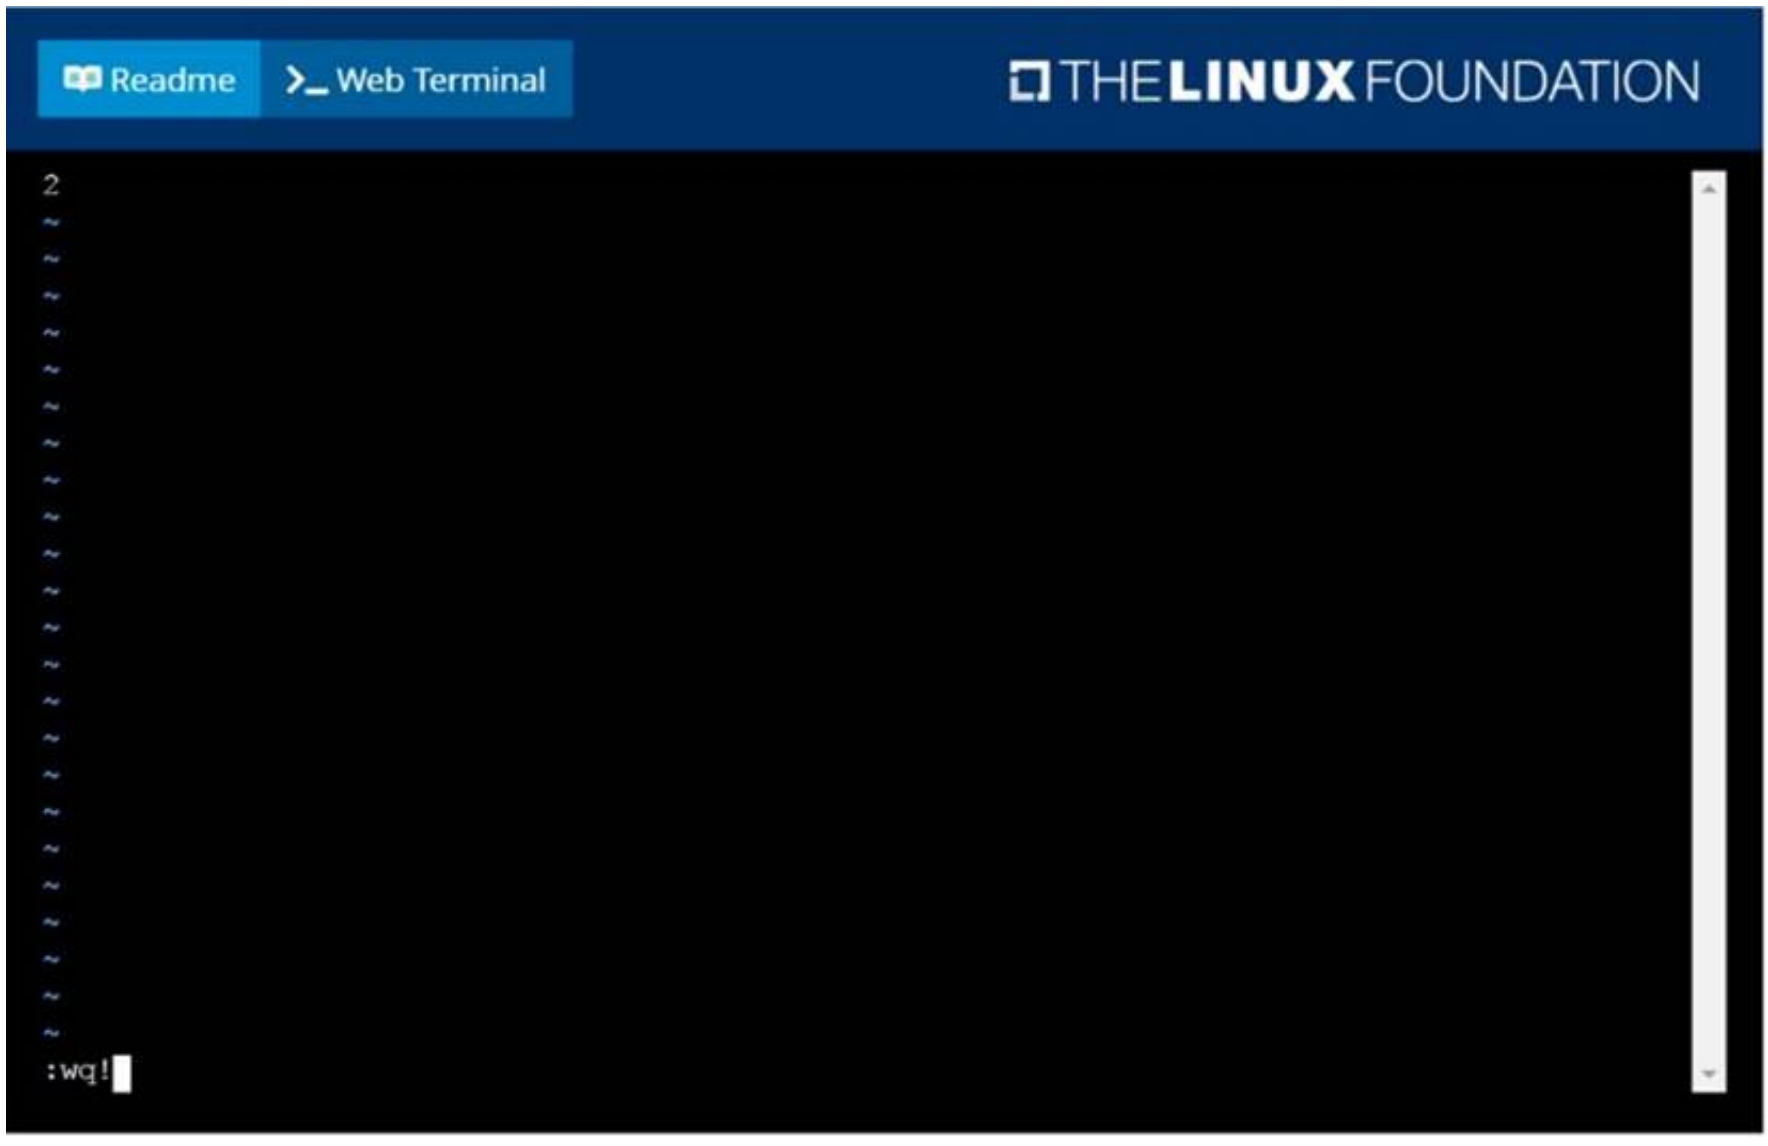

F:\Work\Data Entry Work\Data Entry\20200827\CKA\15 C.JPG

NEW QUESTION 35
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 40
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 41
CORRECT TEXT
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

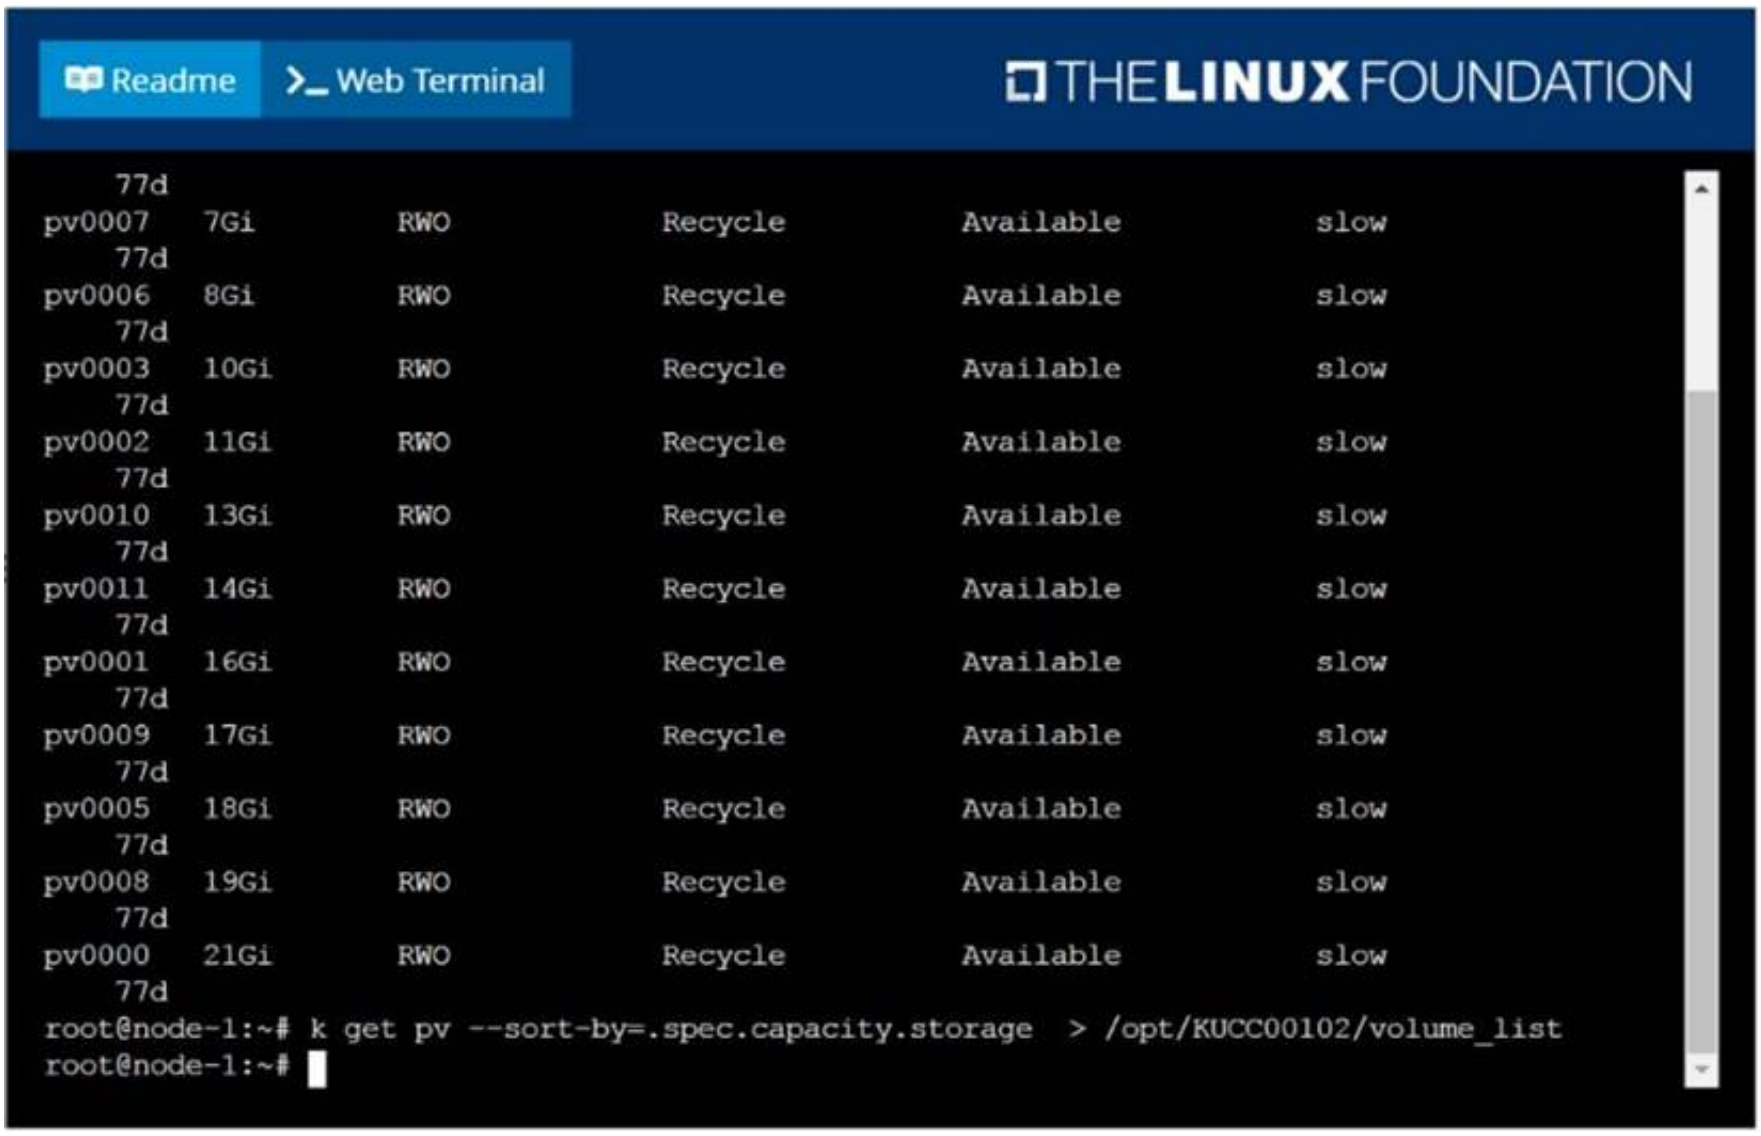

F:\Work\Data Entry Work\Data Entry\20200827\CKA\2 C.JPG

NEW QUESTION 46
CORRECT TEXT
Create a persistent volume with name app-data, of capacity 2Gi and access mode ReadWriteMany. The type of volume is hostPath and its location is /srv/app-data.

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution
Persistent Volume
A persistent volume is a piece of storage in a Kubernetes 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 know the 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 Persistent Volume
kind: PersistentVolumeapiVersion: v1metadata: name:app-dataspec: capacity: # defines the capacity of PV we are creating storage: 2Gi #the amount of storage we are tying to claim accessModes: # defines the rights of the volume we are creating - ReadWriteMany hostPath: path: "/srv/app-data" # path to which we are creating the volume
Challenge
? Create a Persistent Volume named app-data, with access mode ReadWriteMany, storage classname shared, 2Gi of 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.

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

Image for post

* 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: PersistentVolumeapiVersion: v1metadata: name:app-data

spec:

accessModes: - ReadWriteMany resources:

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

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

Image for post

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

```
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
```

Image for post

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: v1kind: Podmetadata: creationTimestamp: null name: app-dataspec: volumes: - name:congigpvc persistenVolumeClaim: claimName: app-data containers: - image: nginx name: app volumeMounts: - mountPath: "/srv/app-data " name: configpvc

NEW QUESTION 51

CORRECT TEXT

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 56
CORRECT TEXT
Schedule a pod as follows:
? Name: nginx-kusc00101
? Image: nginx
? Node selector: disk=ssd

A. Mastered
B. Not Mastered

Answer: A

Explanation:
solution

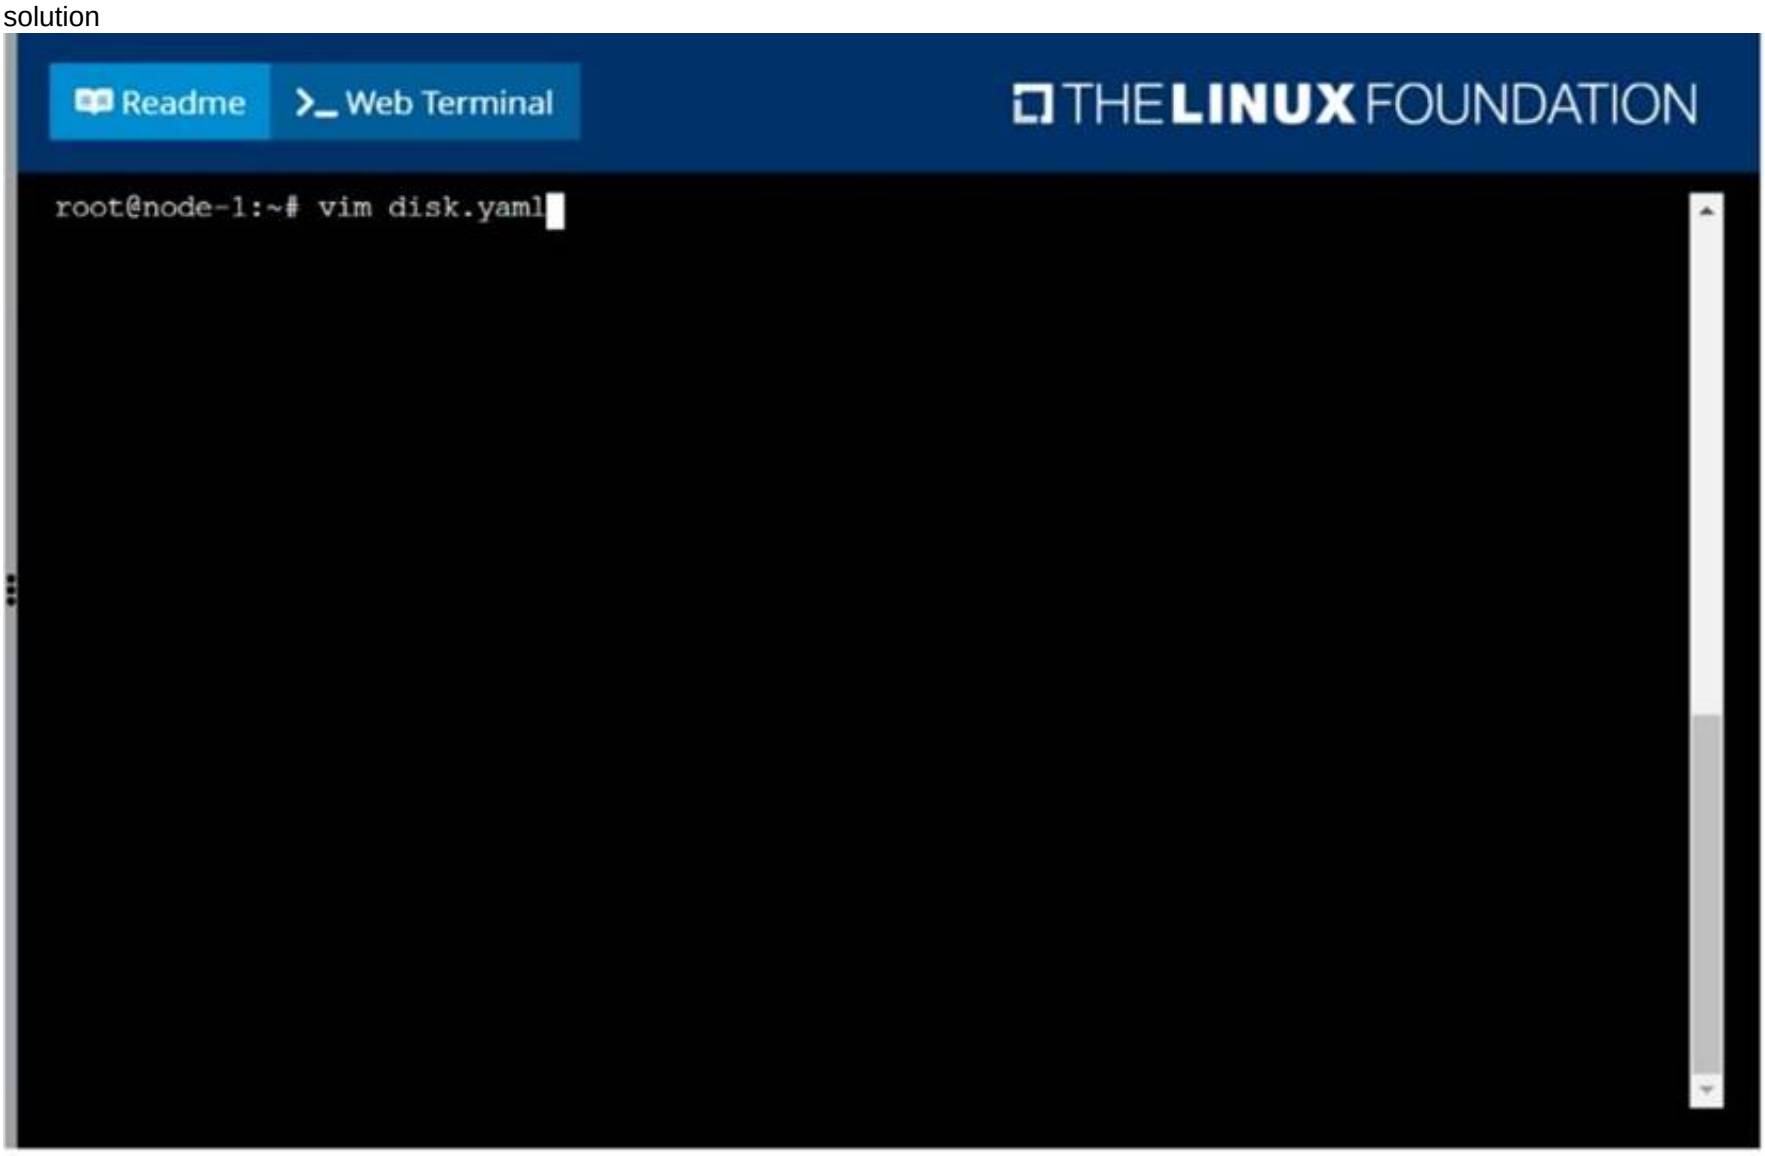

F:\Work\Data Entry Work\Data Entry\20200827\CKA\6 B.JPGF:\Work\Data

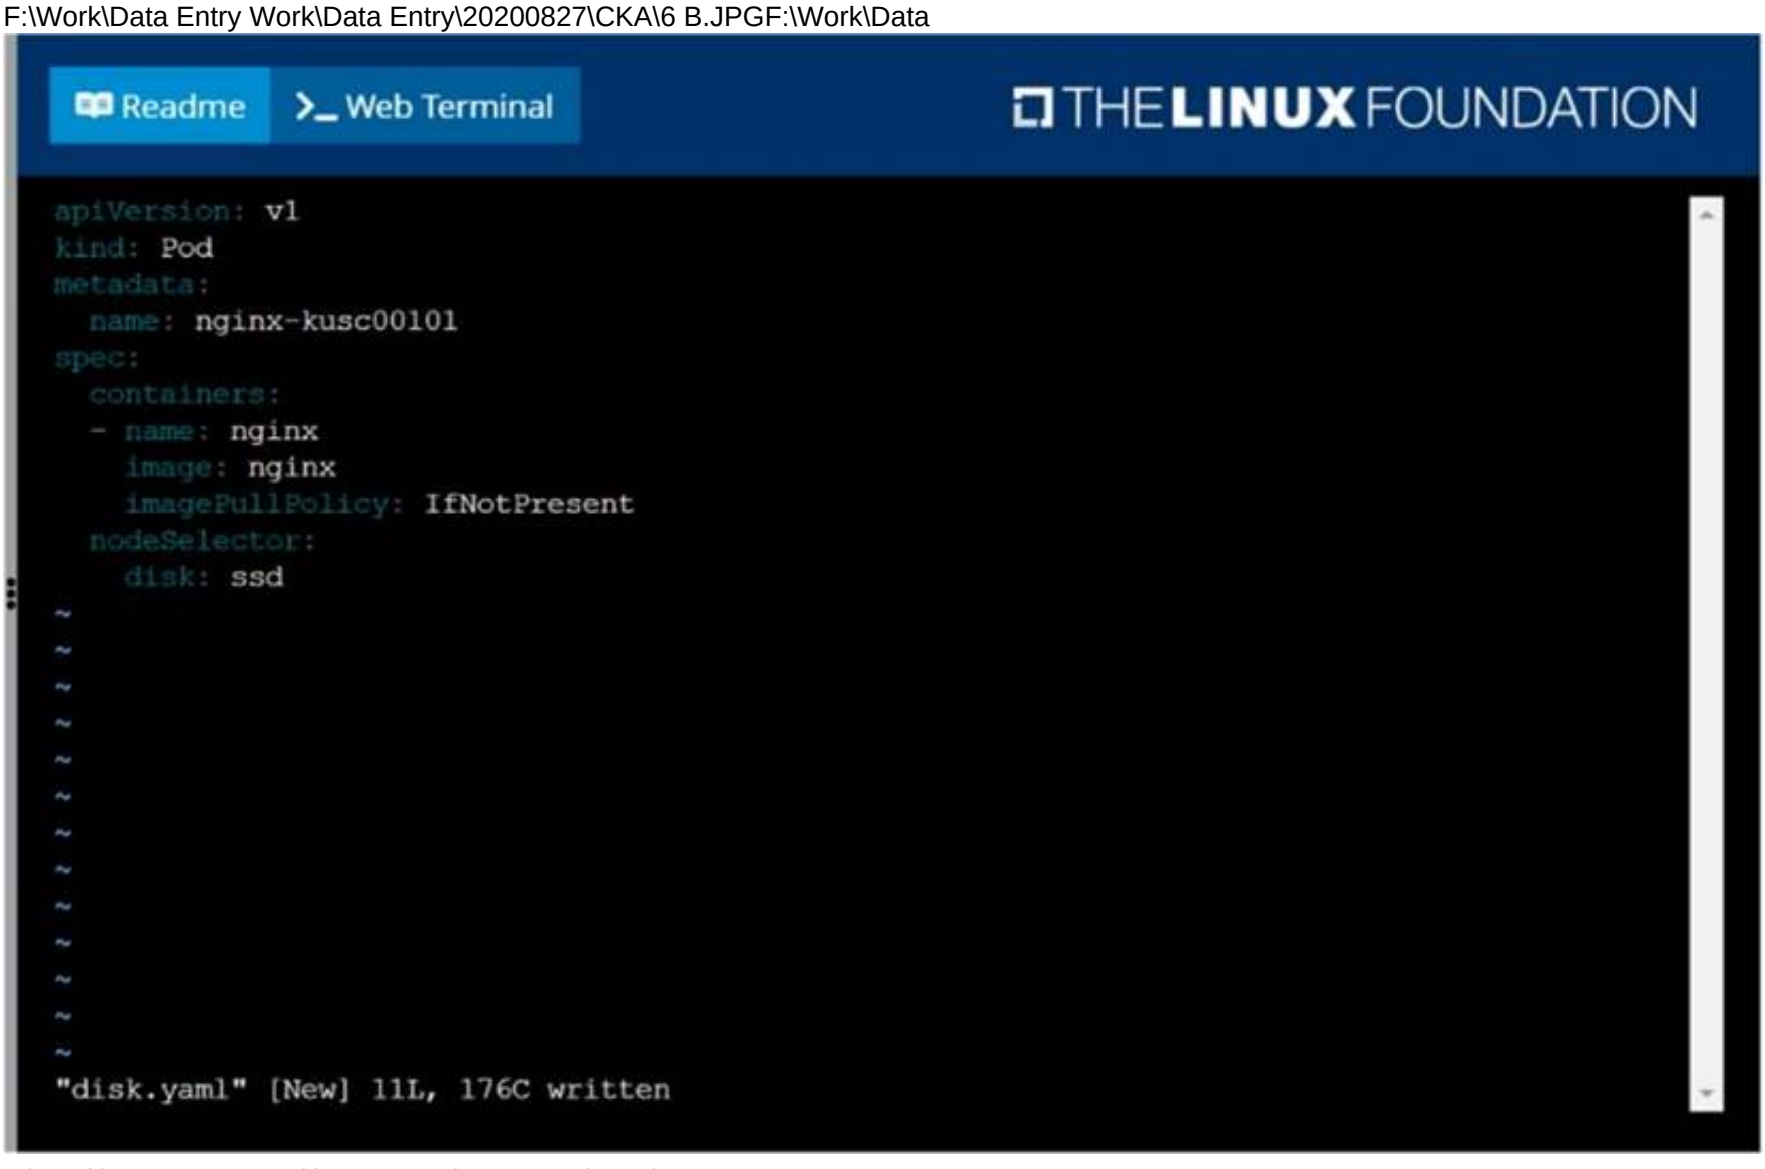

F:\Work\Data Entry Work\Data Entry\20200827\CKA\6 C.JPG

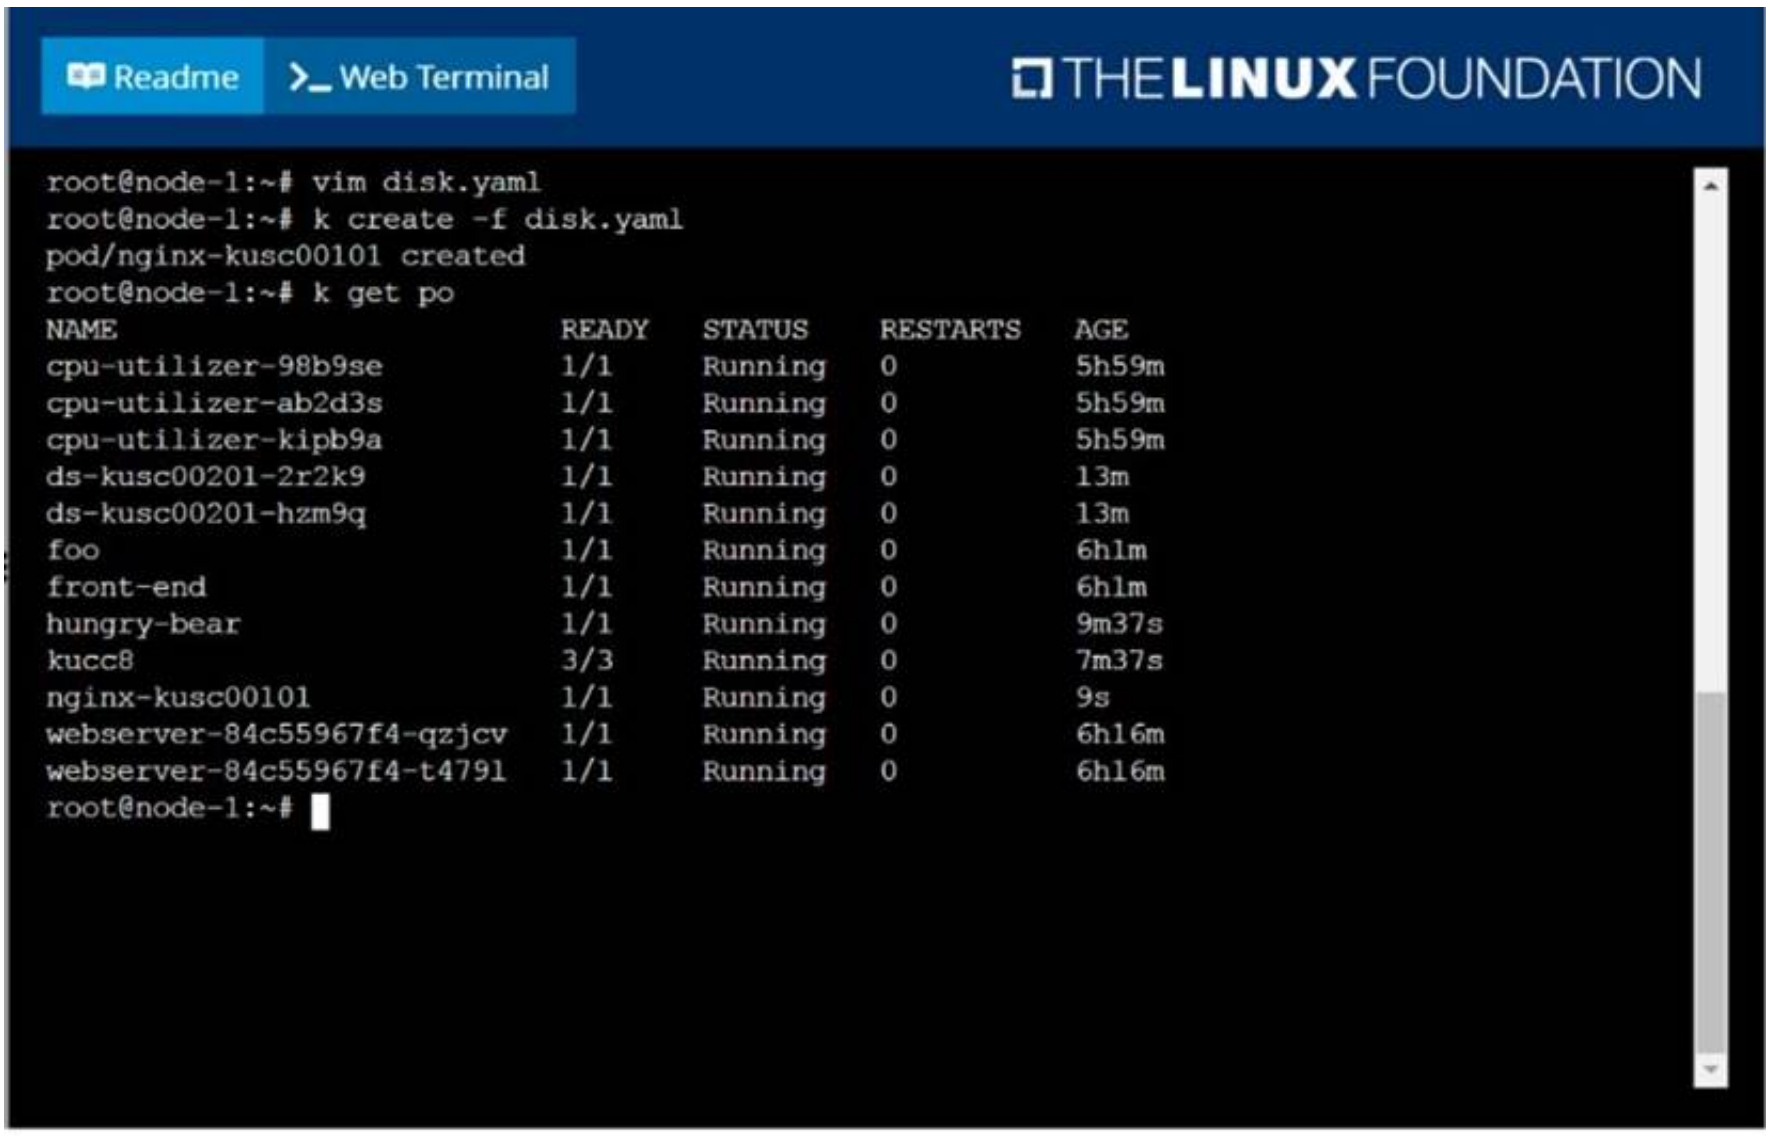

F:\Work\Data Entry Work\Data Entry\20200827\CKA\6 D.JPG

NEW QUESTION 61
CORRECT TEXT
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 65
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 66
CORRECT TEXT
Score: 4%

Task
Create a persistent volume with name app-data , of capacity 1Gi and access mode ReadOnlyMany. The type of volume is hostPath and its location is /srv/app-data .

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Solution:
#vi pv.yaml
apiVersion: v1
kind: PersistentVolume
metadata:
name: app-config
spec:
capacity:
storage: 1Gi
accessModes:
- ReadOnlyMany
hostPath:
path: /srv/app-config
#
kubectl create -f pv.yaml

NEW QUESTION 69
CORRECT TEXT
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 72
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 75
.....

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