

# Linux-Foundation

## Exam Questions CKA

Certified Kubernetes Administrator (CKA) Program



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

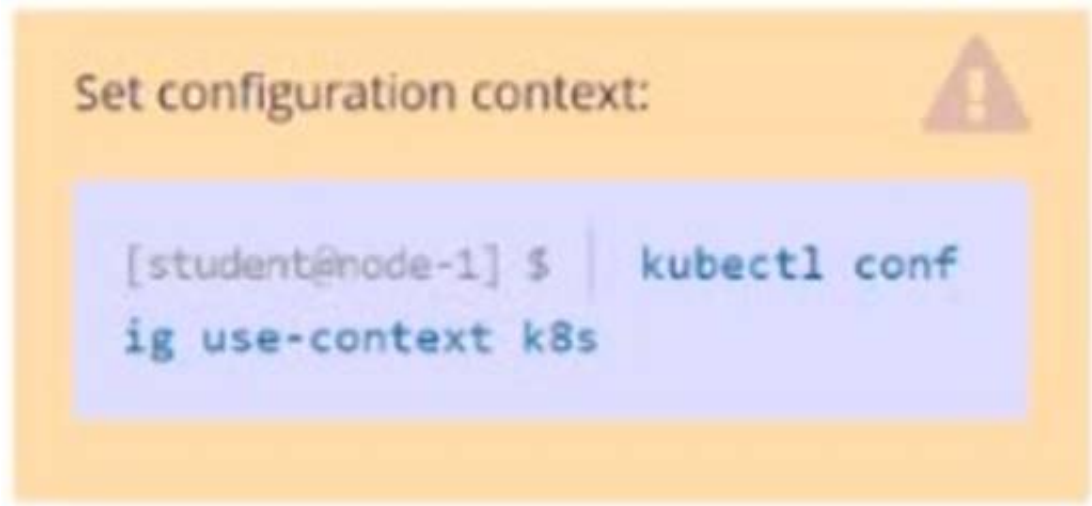
Explanation:

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

NEW QUESTION 2

CORRECT TEXT

Task Weight: 4%



Task

Schedule a Pod as follows:

- Name: kucc1
- App Containers: 2
- Container Name/Images: o nginx  
o consul

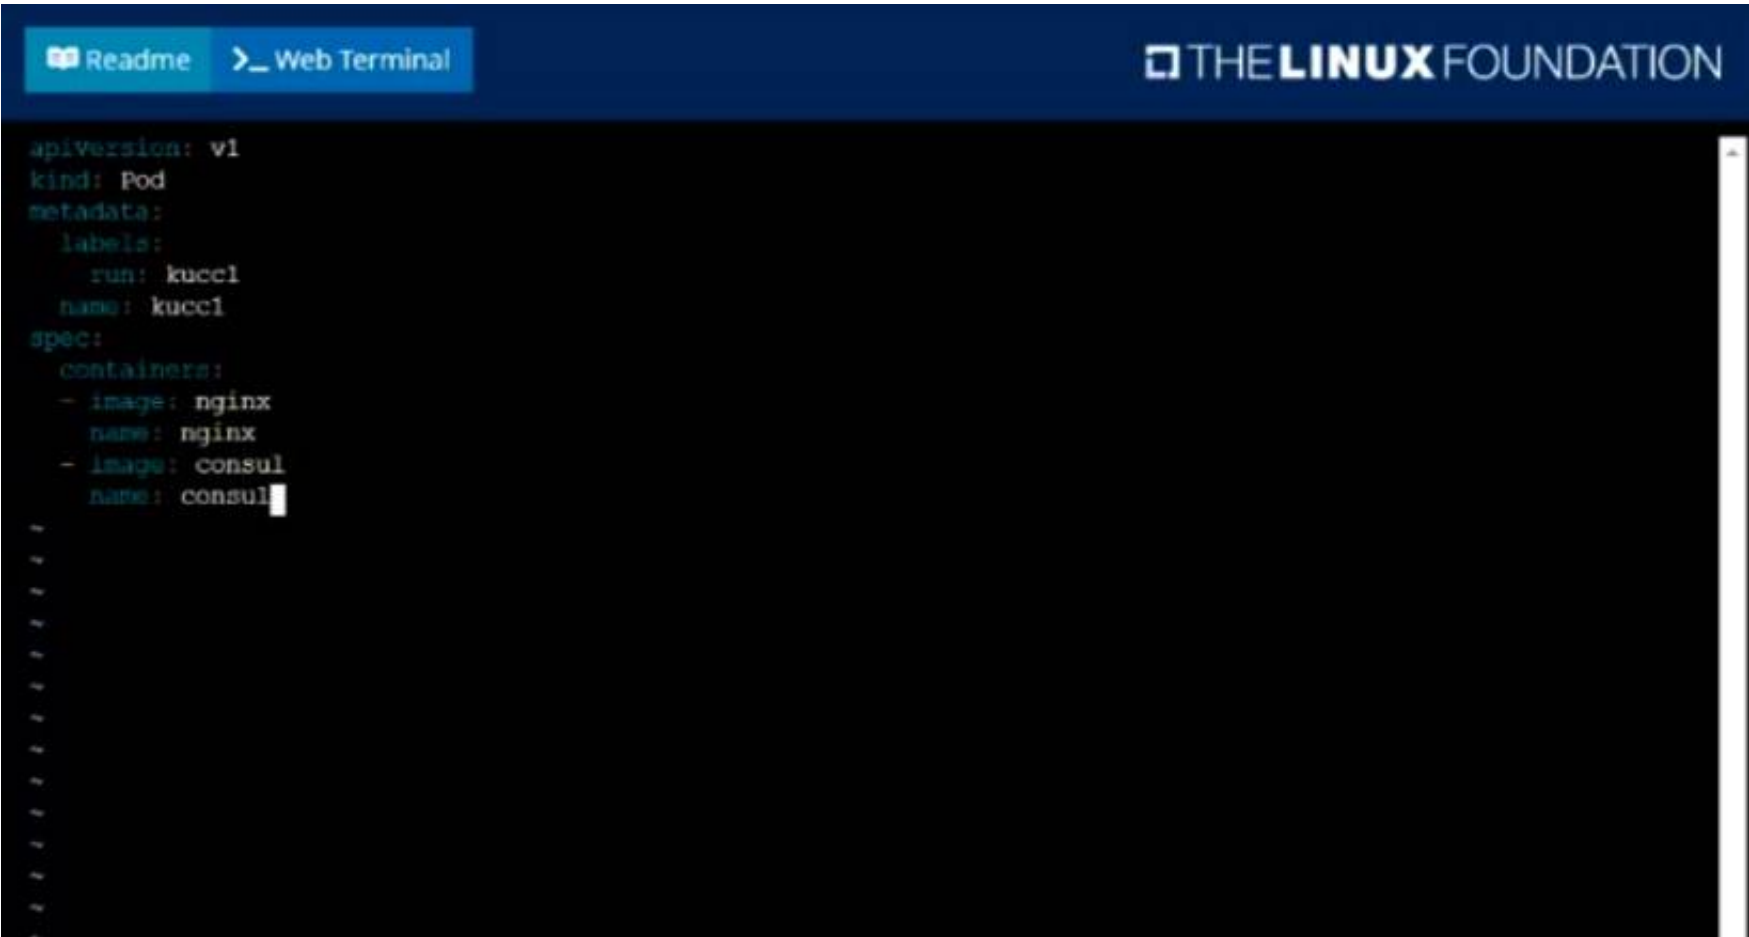
- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Solution:

```
student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ kubectl run kucc1 --image=nginx --dry-run=client -o yaml > aa.y
```



Graphical user interface, text, application  
Description automatically generated

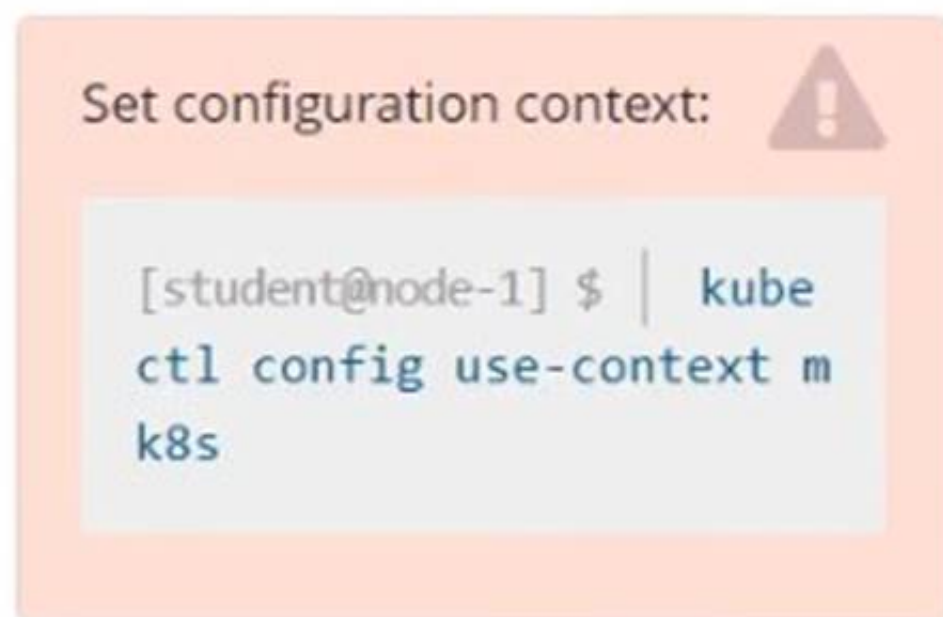
```
student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ kubectl run kucc1 --image=nginx --dry-run=client -o yaml > aa.yaml
student@node-1:~$ vim aa.yaml
student@node-1:~$ kubectl create -f aa.yaml
pod/kucc1 created
student@node-1:~$ kubectl get pods
NAME                                READY   STATUS              RESTARTS   AGE
ll-factor-app                       1/1     Running             0           6h34m
cpu-loader-98b9se                   1/1     Running             0           6h33m
cpu-loader-ab2d3s                   1/1     Running             0           6h33m
cpu-loader-kipb9a                   1/1     Running             0           6h33m
foobar                              1/1     Running             0           6h34m
front-end-6bc87b9748-24rcm          1/1     Running             0           5m4s
front-end-6bc87b9748-hd5wp          1/1     Running             0           5m2s
kucc1                               0/2     ContainerCreating   0           3s
nginx-kusc00401                     1/1     Running             0           2m28s
webserver-84c89dfd75-2d1jn          1/1     Running             0           6h38m
webserver-84c89dfd75-8d8x2          1/1     Running             0           6h38m
webserver-84c89dfd75-z5zz4          1/1     Running             0           3m51s
student@node-1:~$
```

Text Description automatically generated

### NEW QUESTION 3

CORRECT TEXT

Score: 7%



Task

Given an existing Kubernetes cluster running version 1.20.0, upgrade all of the Kubernetes control plane and node components on the master node only to version 1.20.1.

Be sure to drain the master node before upgrading it and uncordon it after the upgrade.

You can ssh to the master node using:

```
[student@node-1] $ | ssh
mk8s-master-0
```

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

```
[student@mk8s-master-0] |
$
sudo -i
```

You are also expected to upgrade kubelet and kubectl on the master node.

Do not upgrade the worker nodes, etcd, the container manager, the CNI plugin, the DNS service or any other addons.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

SOLUTION:

```
[student@node-1] > ssh ek8s
kubectl cordon k8s-master
kubectl drain k8s-master --delete-local-data --ignore-daemonsets --force
apt-get install kubeadm=1.20.1-00 kubelet=1.20.1-00 kubectl=1.20.1-00 --
disableexcludes=kubernetes
kubeadm upgrade apply 1.20.1 --etcd-upgrade=false
systemctl daemon-reload
systemctl restart kubelet kubectl
uncordon k8s-master
```

**NEW QUESTION 4**

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 5**

CORRECT TEXT

Score: 7%



Task

Create a new NetworkPolicy named allow-port-from-namespace in the existing namespace echo. Ensure that the new NetworkPolicy allows Pods in namespace my-app to connect to port 9000 of Pods in namespace echo.

Further ensure that the new NetworkPolicy:

- does not allow access to Pods, which don't listen on port 9000
- does not allow access from Pods, which are not in namespace my-app

- A. Mastered  
B. Not Mastered

**Answer:** A

#### Explanation:

Solution:  
#network.yaml  
apiVersion: networking.k8s.io/v1  
kind: NetworkPolicy  
metadata:  
name: allow-port-from-namespace  
namespace: internal  
spec:  
podSelector:  
matchLabels: {  
}  
policyTypes:  
- Ingress  
ingress:  
- from:  
- podSelector: {  
}  
ports:  
- protocol: TCP  
port: 8080  
#spec.podSelector namespace pod  
kubectl create -f network.yaml

#### NEW QUESTION 6

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



- 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 7**

CORRECT TEXT

Score: 7%



Task  
Reconfigure the existing deployment front-end and add a port specification named http exposing port 80/tcp of the existing container nginx.  
Create a new service named front-end-svc exposing the container port http.  
Configure the new service to also expose the individual Pods via a NodePort on the nodes on which they are scheduled.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Solution:  
kubectl get deploy front-end  
kubectl edit deploy front-end -o yaml  
#port specification named http  
#service.yaml  
apiVersion: v1  
kind: Service  
metadata:  
name: front-end-svc  
labels:  
app: nginx  
spec:  
ports:  
- port: 80  
protocol: tcp  
name: http  
selector:  
app: nginx  
type: NodePort  
# kubectl create -f service.yaml  
# kubectl get svc  
# port specification named http  
kubectl expose deployment front-end --name=front-end-svc --port=80 --target-port=80 --type=NodePort

### NEW QUESTION 8

CORRECT TEXT

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 9

CORRECT TEXT

Create a file:

/opt/KUCC00302/kucc00302.txt that lists all pods that implement service baz in namespace development.

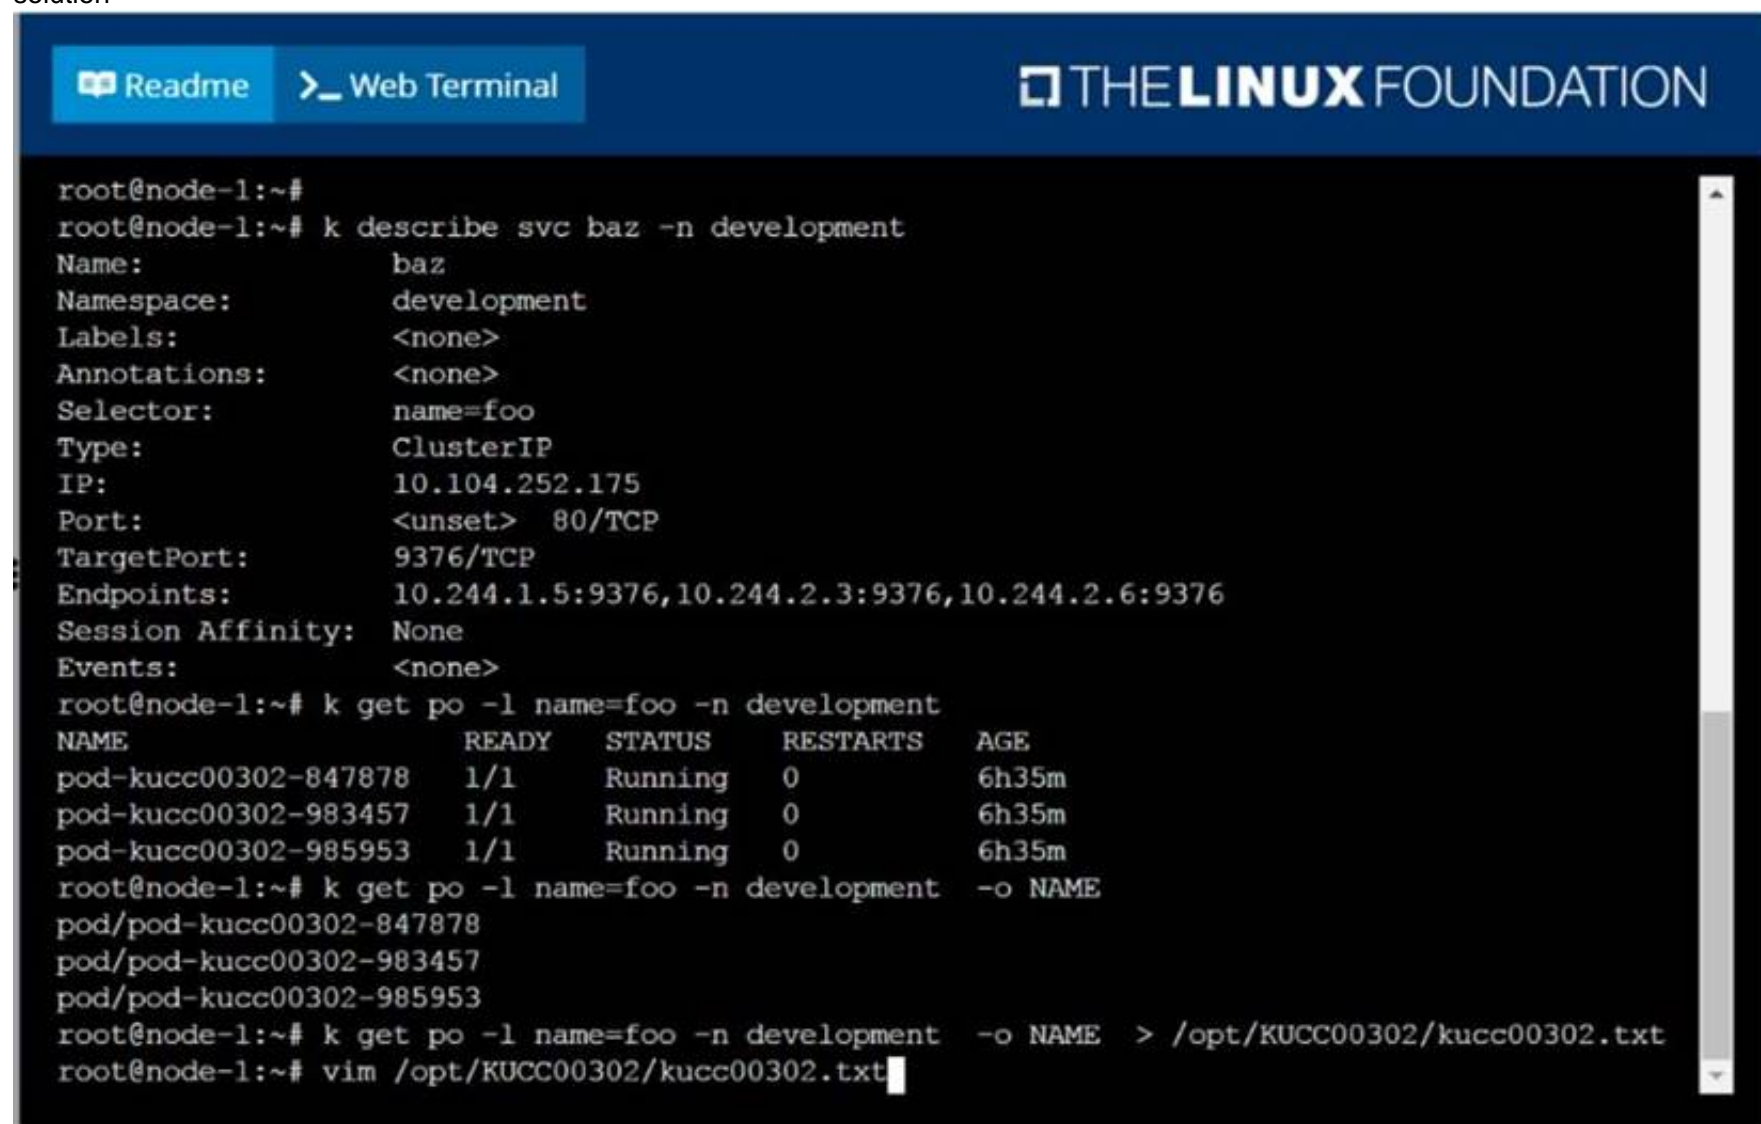
The format of the file should be one pod name per line.

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:


solution




```
root@node-1:~#
root@node-1:~# k describe svc baz -n development
Name:          baz
Namespace:     development
Labels:        <none>
Annotations:   <none>
Selector:      name=foo
Type:          ClusterIP
IP:            10.104.252.175
Port:          <unset> 80/TCP
TargetPort:    9376/TCP
Endpoints:     10.244.1.5:9376,10.244.2.3:9376,10.244.2.6:9376
Session Affinity: None
Events:        <none>
root@node-1:~# k get po -l name=foo -n development
NAME                                READY   STATUS    RESTARTS   AGE
pod-kucc00302-847878                1/1     Running   0           6h35m
pod-kucc00302-983457                1/1     Running   0           6h35m
pod-kucc00302-985953                1/1     Running   0           6h35m
root@node-1:~# k get po -l name=foo -n development -o NAME
pod/pod-kucc00302-847878
pod/pod-kucc00302-983457
pod/pod-kucc00302-985953
root@node-1:~# k get po -l name=foo -n development -o NAME > /opt/KUCC00302/kucc00302.txt
root@node-1:~# vim /opt/KUCC00302/kucc00302.txt
```

F:\Work\Data Entry Work\Data Entry\20200827\CKA\11 B.JPG





F:\Work\Data Entry Work\Data Entry\20200827\CKA\11 C.JPG



F:\Work\Data Entry Work\Data Entry\20200827\CKA\11 D.JPG

**NEW QUESTION 10**

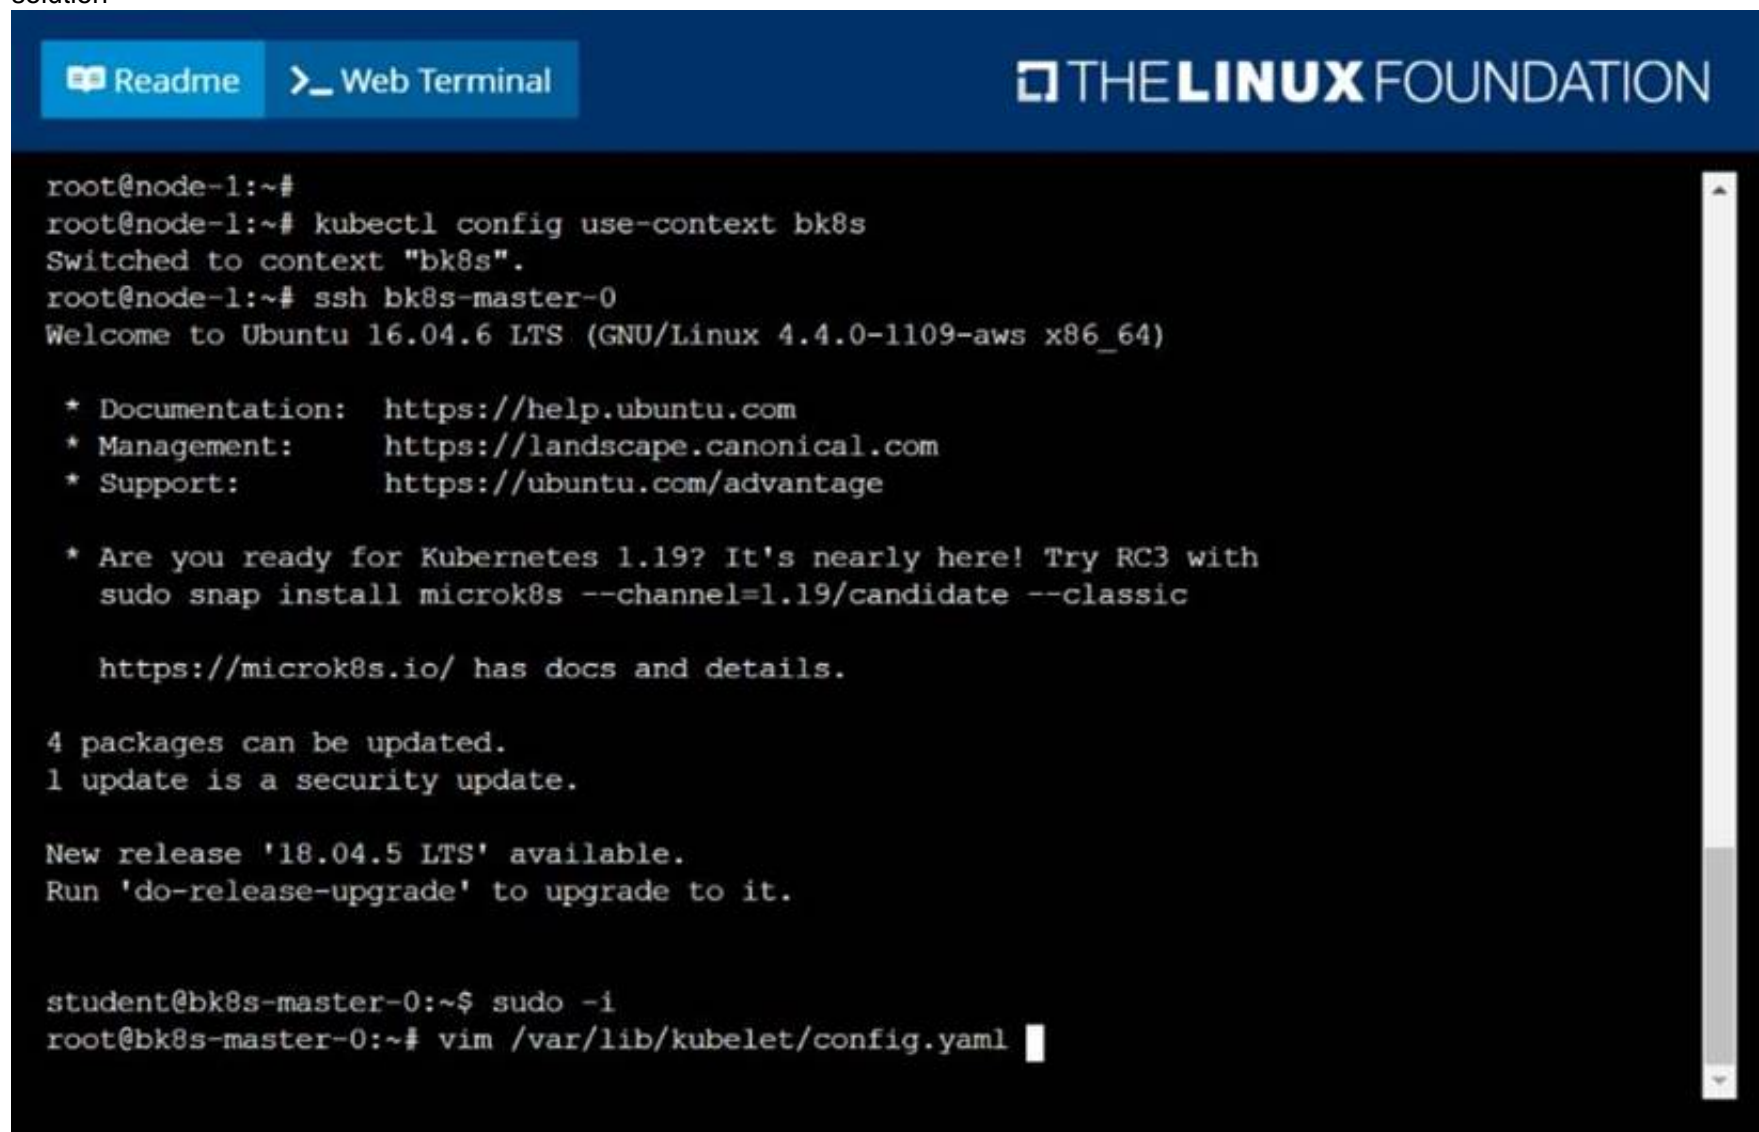
**CORRECT TEXT**

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



The screenshot shows a web terminal interface with a dark blue header. On the left, there are two buttons: 'Readme' and 'Web Terminal'. On the right, the 'THE LINUX FOUNDATION' logo is displayed. The terminal window shows a series of commands and their outputs. The user is initially at a root prompt on a node-1 machine. They run 'kubectl config use-context bk8s', which switches the context to 'bk8s'. Then, they run 'ssh bk8s-master-0', which connects them to the master node. The master node is running Ubuntu 16.04.6 LTS. The terminal output includes system messages about updates and a new release. Finally, the user runs 'sudo -i' to become root, and then 'vim /var/lib/kubelet/config.yaml' to edit the kubelet configuration file.

```
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

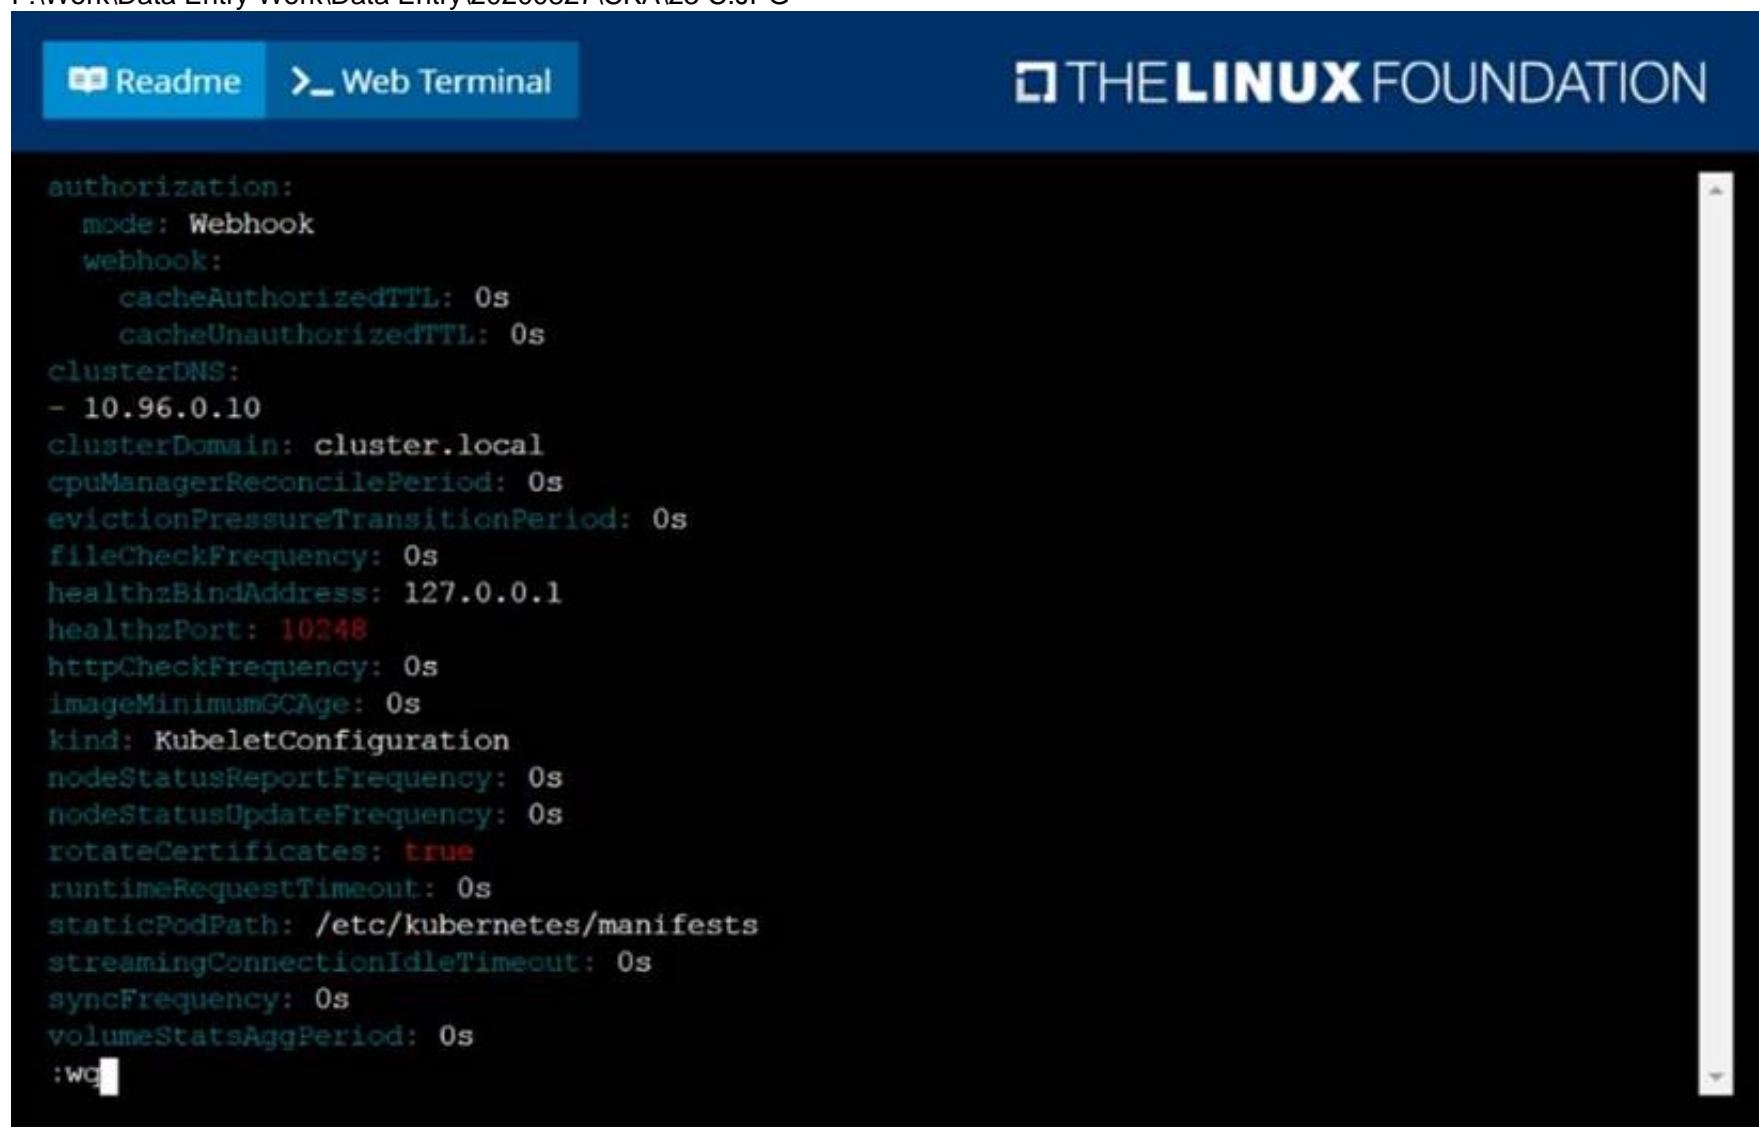
   https://microk8s.io/ has docs and details.

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



The screenshot shows a web terminal interface with a dark blue header. On the left, there are two buttons: 'Readme' and 'Web Terminal'. On the right, the 'THE LINUX FOUNDATION' logo is displayed. The terminal window shows the contents of the kubelet configuration file, which is a YAML document. The configuration includes settings for authorization, cluster DNS, cluster domain, CPU manager, eviction pressure, file check frequency, healthz bind address, healthz port, HTTP check frequency, image minimum GC age, kind, node status report frequency, node status update frequency, rotate certificates, runtime request timeout, static pod path, streaming connection idle timeout, sync frequency, and volume stats aggregation period.

```
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

ReadmeWeb 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@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:~# kubect1 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:~#
```

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

#### NEW QUESTION 10

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:

kubect1 run busybox --image=busybox -it --rm --restart=Never --  
/bin/sh -c 'echo hello world'  
kubect1 get po # You shouldn't see pod with the name "busybox"

#### NEW QUESTION 15

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

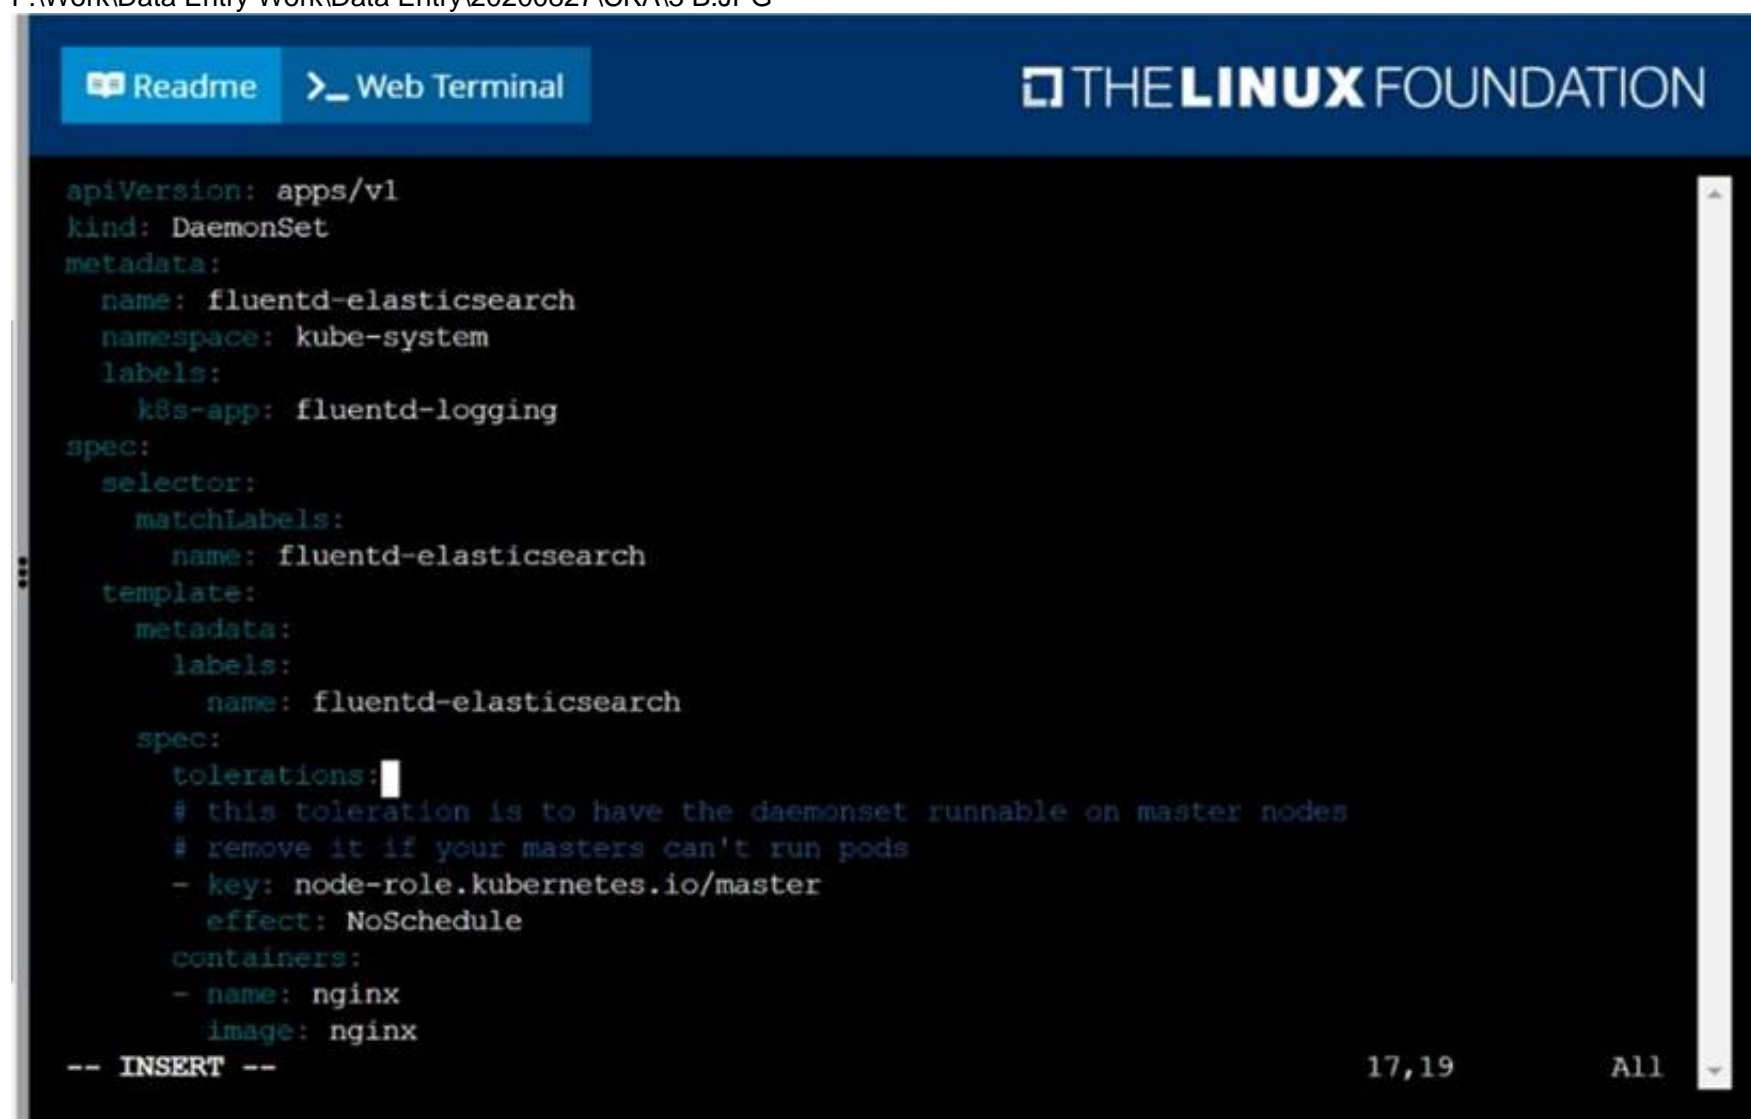
#### Explanation:

solution

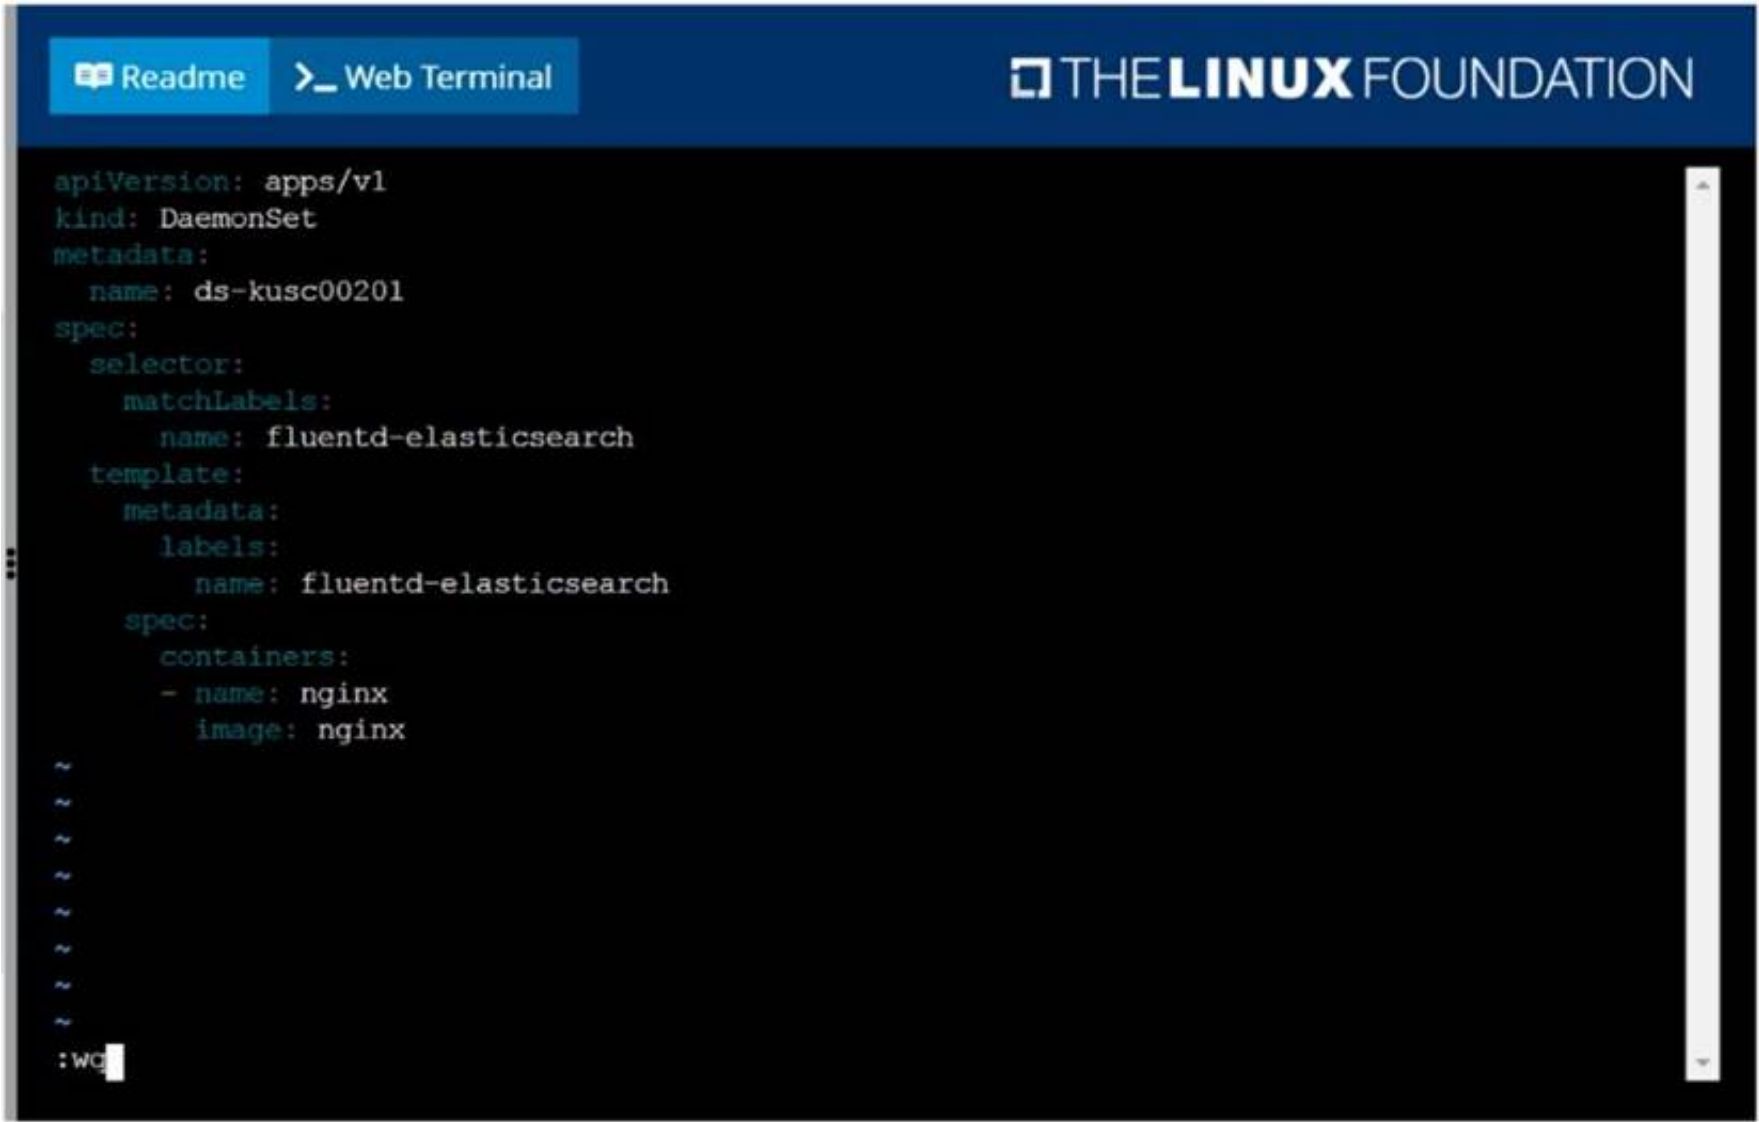




F:\Work\Data Entry Work\Data Entry\20200827\CKA\3 B.JPG



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 18**

CORRECT TEXT

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 the staging namespace and the volume must not be persistent.

- A. Mastered
- B. Not Mastered

Answer: A



solution



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

NEW QUESTION 20

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 22

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 25

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 26

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**  
List the nginx pod with custom columns POD\_NAME and POD\_STATUS

- ```
kubectl get po -o=custom-columns="POD_NAME:.metadata.name, POD_STATUS:.status.containerStatuses[].state"
```

visit - <https://www.surepassexam.com>

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 40

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-prod-pod.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 45

.....

## 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](#)**