CKA Q&As

Certified Kubernetes Administrator (CKA) Program

Pass Linux Foundation CKA Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.pass4itsure.com/cka.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Linux Foundation Official Exam Center

QUESTION 1

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

Correct Answer: Check the answer in explanation.

Solution

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

QUESTION 2

List the nginx pod with custom columns POD_NAME and POD_STATUS

Correct Answer: Check the answer in explanation.

Solution

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

QUESTION 3

SIMULATION

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/KUTR00102/KUTR00102.txt (which already exists).

Correct Answer: Check the answer in explanation.

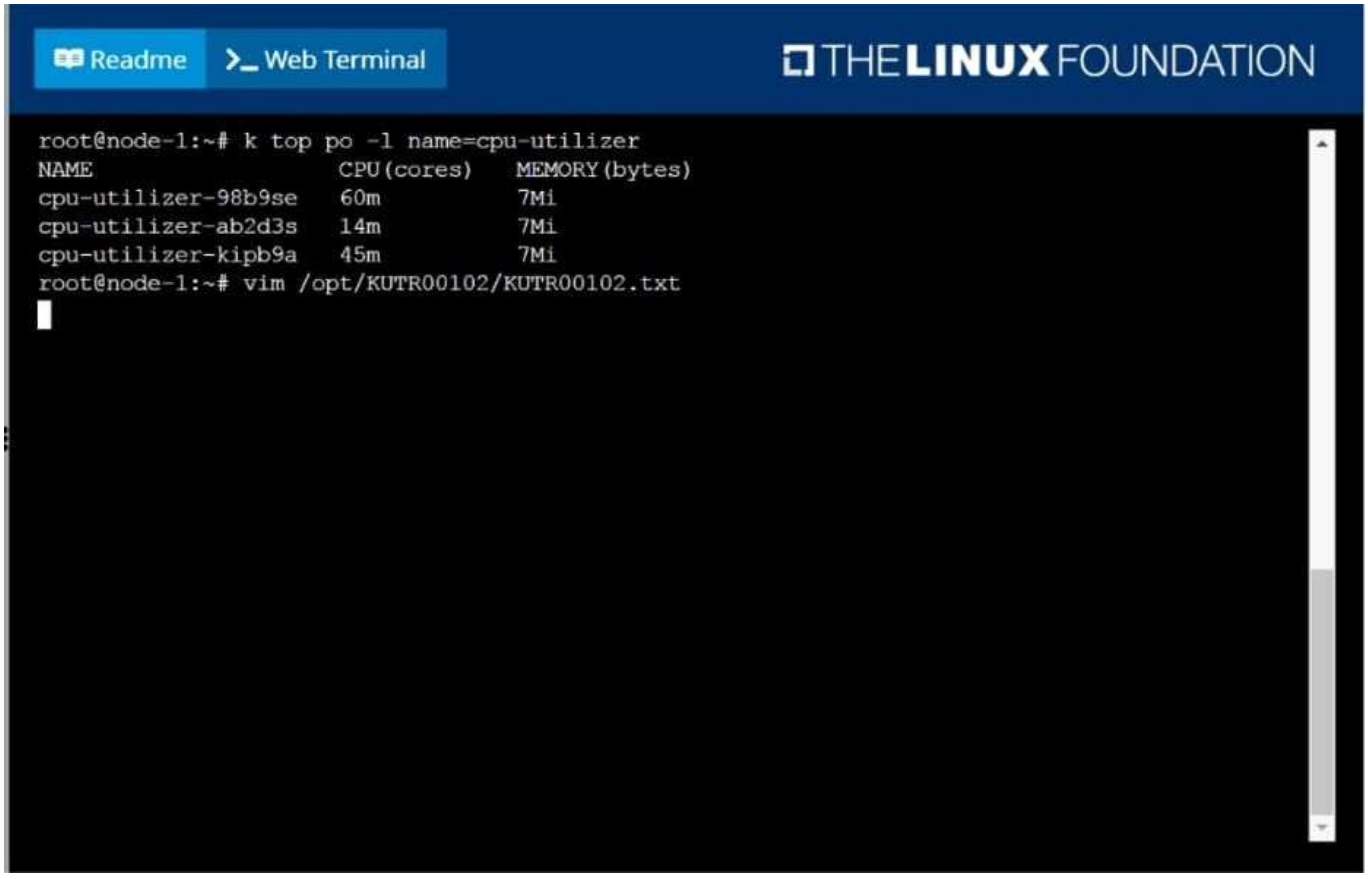
Readme
Web Terminal
THE LINUX FOUNDATION
root@node-1:~# k top po -l name=cpu-utilizer
NAME CPU(cores) MEMORY(bytes)
cpu-utilizer-98b9se 60m 7Mi
cpu-utilizer-ab2d3s 14m 7Mi
cpu-utilizer-kipb9a 45m 7Mi
root@node-1:~# vim /opt/KUTR00102/KUTR00102.txt

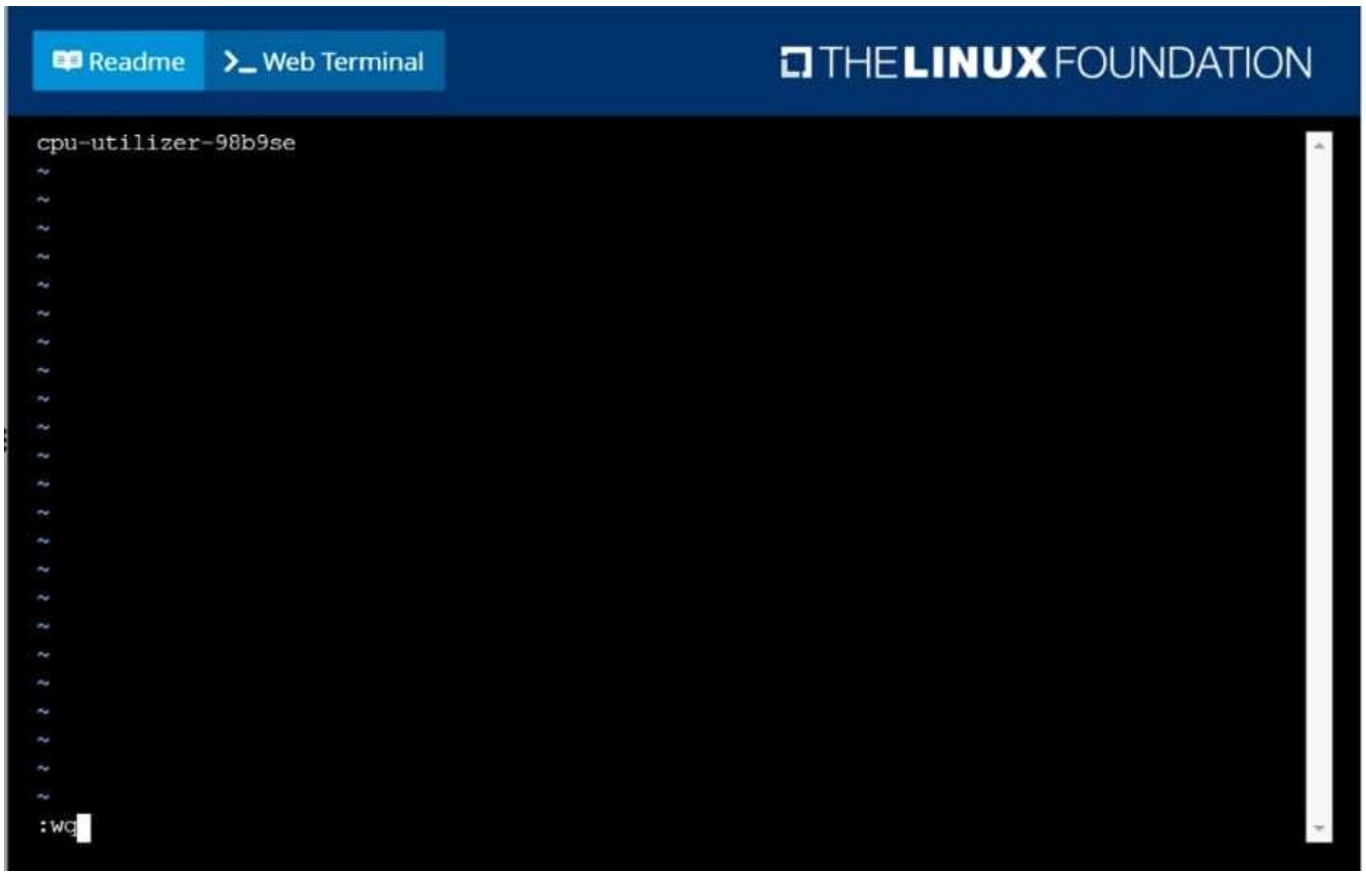
Readme
Web Terminal
THE LINUX FOUNDATION
cpu-utilizer-98b9se
:wq

QUESTION 4

Print pod name and start time to "/opt/pod-status" file

Correct Answer: Check the answer in explanation.

Solution

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

QUESTION 5

SIMULATION

Create a pod as follows: Name: mongo Using Image: mongo In a new Kubernetes namespace named: my-website

Correct Answer: Check the answer in explanation.

Solution

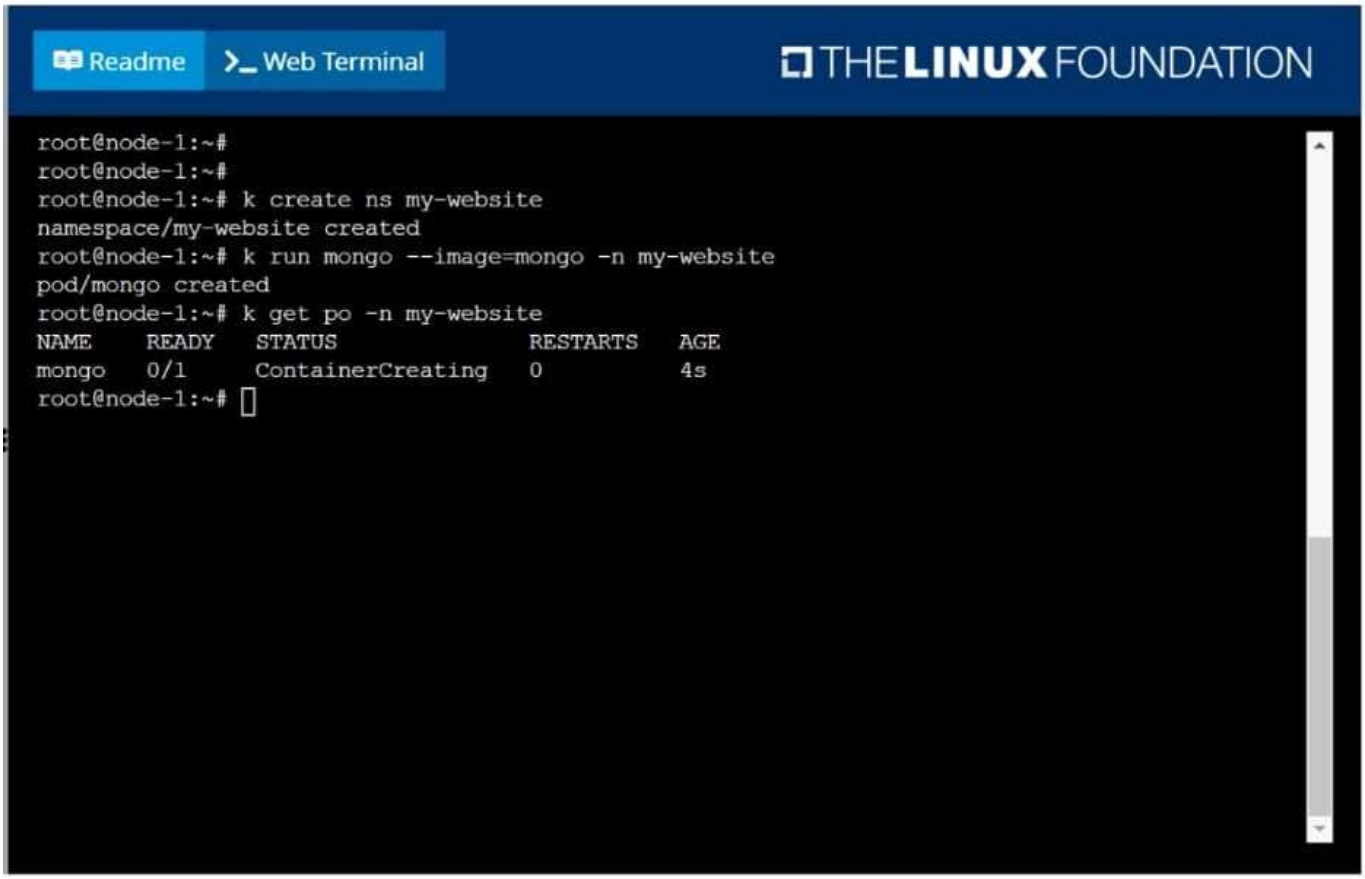

QUESTION 6

Check the image version in pod without the describe command

Correct Answer: Check the answer in explanation.

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

QUESTION 7

SIMULATION

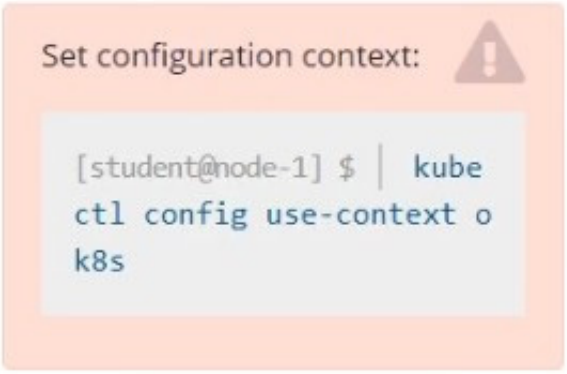

Task Create a new PersistentVolumeClaim

1.

Name: pv-volume

2.

Class: csi-hostpath-sc

3.

Capacity: 10Mi

Create a new Pod which mounts the PersistentVolumeClaim as a volume:

1.

Name: web-server

2.

Image: nginx

3.

Mount path: /usr/share/nginx/html

Configure the new Pod to have ReadWriteOnce access on the volume.

Finally, using kubectl edit or kubectl patch expand the PersistentVolumeClaim to a capacity of 70Mi and record that change.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context ok8s
Switched to context "ok8s".
student@node-1:~$
```

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: pv-volume
spec:
  storageClassName: csi-hostpath-sc
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Mi
```

```
student@node-1:~$ kubectl config use-context ok8s
Switched to context "ok8s".
student@node-1:~$ vim pvc.yml
student@node-1:~$ kubectl get pv,pvc
No resources found
student@node-1:~$ kubectl create -f pvc.yml
persistentvolumeclaim/pv-volume created
student@node-1:~$ kubectl get pv,pvc
NAME                                                        CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM
        STORAGECLASS      REASON   AGE
persistentvolume/pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            Delete           Bound    default/pv-
volume   csi-hostpath-sc            3s

NAME                                STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLAS
S      AGE
persistentvolumeclaim/pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpat
h-sc   3s
student@node-1:~$
```

```
apiVersion: v1
kind: Pod
metadata:
  name: web-server
spec:
  volumes:
    - name: task-pv-storage
      persistentVolumeClaim:
        claimName: pv-volume
  containers:
    - name: web-server
      image: nginx
      volumeMounts:
        - mountPath: "/usr/share/nginx/html"
          name: task-pv-storage
```

```
student@node-1:~$ kubectl config use-context ok8s
Switched to context "ok8s".
student@node-1:~$ vim pvc.yml
student@node-1:~$ kubectl get pv,pvc
No resources found
student@node-1:~$ kubectl create -f pvc.yml
persistentvolumeclaim/pv-volume created
student@node-1:~$ kubectl get pv,pvc
NAME                                                        CAPACITY   ACCESS MODES   RECLAIM POLICY   STATUS   CLAIM
        STORAGECLASS      REASON   AGE
persistentvolume/pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            Delete           Bound    default/pv-
volume   csi-hostpath-sc            3s

NAME                                STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLAS
S      AGE
persistentvolumeclaim/pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpat
h-sc   3s
student@node-1:~$
student@node-1:~$
student@node-1:~$ vim pod.yml
student@node-1:~$ kubectl create -f pod.yml
pod/web-server created
student@node-1:~$ kubectl get pods
NAME                  READY   STATUS              RESTARTS   AGE
csi-hostpath-socat-0  1/1     Running             0          6h12m
csi-hostpathplugin-0  9/9     Running             0          6h12m
web-server            0/1     ContainerCreating   0          5s
student@node-1:~$ kubectl get pods -w
NAME                  READY   STATUS              RESTARTS   AGE
csi-hostpath-socat-0  1/1     Running             0          6h12m
csi-hostpathplugin-0  9/9     Running             0          6h12m
web-server            0/1     ContainerCreating   0          10s
web-server            1/1     Running             0          22s
^Cstudent@node-1:~$ kubectl get pods
NAME                  READY   STATUS    RESTARTS   AGE
csi-hostpath-socat-0  1/1     Running   0          6h12m
csi-hostpathplugin-0  9/9     Running   0          6h12m
web-server            1/1     Running   0          27s
student@node-1:~$ kubectl edit pvc pv-volume
```

```
# Please edit the object below. Lines beginning with a '#' will be ignored,
# and an empty file will abort the edit. If an error occurs while saving this file will be
# reopened with the relevant failures.
#
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  annotations:
    pv.kubernetes.io/bind-completed: "yes"
    pv.kubernetes.io/bound-by-controller: "yes"
    volume.beta.kubernetes.io/storage-provisioner: hostpath.csi.k8s.io
    volume.kubernetes.io/storage-provisioner: hostpath.csi.k8s.io
  creationTimestamp: "2022-04-25T15:37:42Z"
  finalizers:
  - kubernetes.io/pvc-protection
  name: pv-volume
  namespace: default
  resourceVersion: "42412"
  uid: 6b6c71cb-558d-4b47-a0db-3951737097eb
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 70Mi
  storageClassName: csi-hostpath-sc
  volumeMode: Filesystem
  volumeName: pvc-6b6c71cb-558d-4b47-a0db-3951737097eb
status:
  accessModes:
  - ReadWriteOnce
  capacity:
    storage: 10Mi
:wq
```

```
student@node-1:~$ kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS      AGE
pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpath-sc   2m46s
student@node-1:~$ kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS      AGE
pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpath-sc   2m56s
student@node-1:~$ kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS      AGE
pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpath-sc   3m
student@node-1:~$ kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS      AGE
pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpath-sc   3m5s
student@node-1:~$ kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS      AGE
pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   10Mi       RWO            csi-hostpath-sc   3m9s
student@node-1:~$ kubectl get pvc
NAME        STATUS   VOLUME                                     CAPACITY   ACCESS MODES   STORAGECLASS      AGE
pv-volume   Bound    pvc-6b6c71cb-558d-4b47-a0db-3951737097eb   70Mi       RWO            csi-hostpath-sc   3m42s
student@node-1:~$ kubectl config use-context k8s
```

QUESTION 8

SIMULATION

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

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context wk8s
Switched to context "wk8s".
student@node-1:~$ kubectl get nodes
NAME           STATUS     ROLES                  AGE   VERSION
wk8s-master-0  Ready      control-plane,master   67d   v1.23.1
wk8s-node-0    NotReady   <none>                 67d   v1.23.1
wk8s-node-1    Ready      <none>                 67d   v1.23.1
student@node-1:~$ kubectl describe nodes wk8s-node-0
```

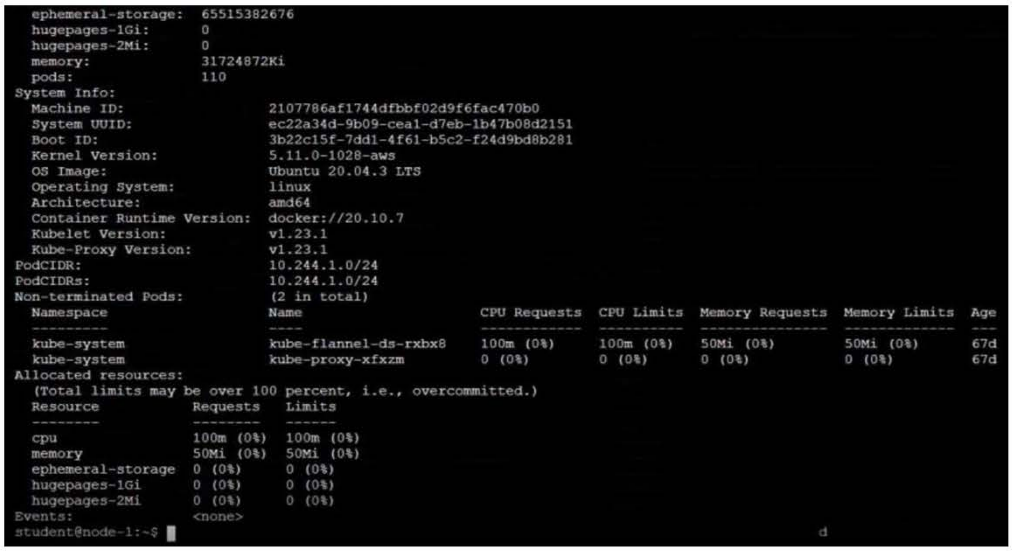

```
  ephemeral-storage:  65515382676
  hugepages-1Gi:      0
  hugepages-2Mi:      0
  memory:             31724872Ki
  pods:               110
System Info:
  Machine ID:                 2107786af1744dfbbf02d9f6fac470b0
  System UUID:                ec22a34d-9b09-cea1-d7eb-1b47b08d2151
  Boot ID:                    3b22c15f-7dd1-4f61-b5c2-f24d9bd8b281
  Kernel Version:             5.11.0-1028-aws
  OS Image:                   Ubuntu 20.04.3 LTS
  Operating System:           linux
  Architecture:               amd64
  Container Runtime Version:  docker://20.10.7
  Kubelet Version:            v1.23.1
  Kube-Proxy Version:         v1.23.1
PodCIDR:                      10.244.1.0/24
PodCIDRs:                     10.244.1.0/24
Non-terminated Pods:          (2 in total)
  Namespace                   Name                CPU Requests  CPU Limits  Memory Requests  Memory Limits  Age
  ---------                   ----                ------------  ----------  ---------------  -------------  ---
  kube-system                 kube-flannel-ds-rxbx8  100m (0%)     100m (0%)   50Mi (0%)        50Mi (0%)      67d
  kube-system                 kube-proxy-xfxzm    0 (0%)        0 (0%)      0 (0%)           0 (0%)         67d
Allocated resources:
  (Total limits may be over 100 percent, i.e., overcommitted.)
  Resource           Requests    Limits
  --------           --------    ------
  cpu                100m (0%)   100m (0%)
  memory             50Mi (0%)   50Mi (0%)
  ephemeral-storage  0 (0%)      0 (0%)
  hugepages-1Gi      0 (0%)      0 (0%)
  hugepages-2Mi      0 (0%)      0 (0%)
Events:              <none>
student@node-1:~$
```

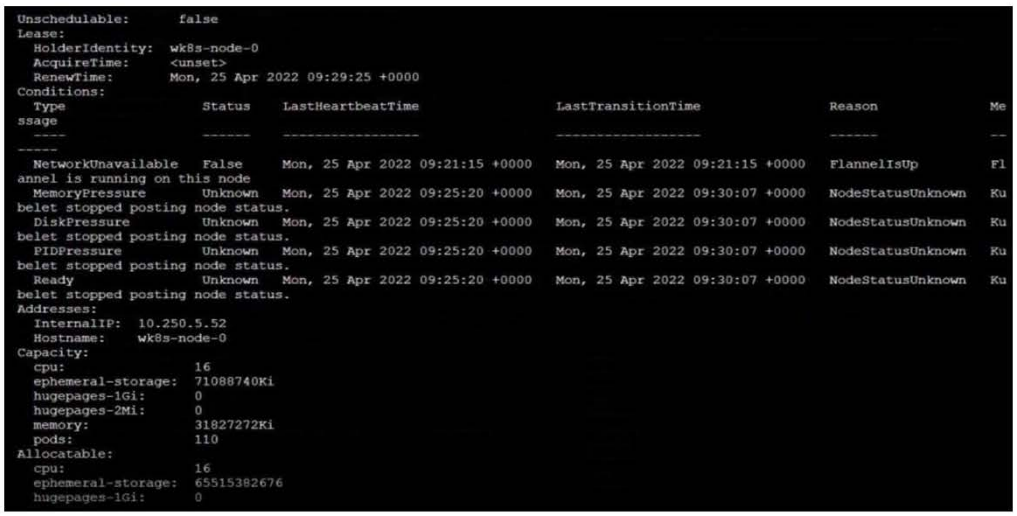

```
Unschedulable:      false
Lease:
  HolderIdentity:  wk8s-node-0
  AcquireTime:     <unset>
  RenewTime:       Mon, 25 Apr 2022 09:29:25 +0000
Conditions:
  Type                 Status    LastHeartbeatTime                 LastTransitionTime                Reason              Me
ssage
  ----                 ------    -----------------                 ------------------                ------              --
-----
  NetworkUnavailable   False     Mon, 25 Apr 2022 09:21:15 +0000   Mon, 25 Apr 2022 09:21:15 +0000   FlannelIsUp         Fl
annel is running on this node
  MemoryPressure       Unknown   Mon, 25 Apr 2022 09:25:20 +0000   Mon, 25 Apr 2022 09:30:07 +0000   NodeStatusUnknown   Ku
belet stopped posting node status.
  DiskPressure         Unknown   Mon, 25 Apr 2022 09:25:20 +0000   Mon, 25 Apr 2022 09:30:07 +0000   NodeStatusUnknown   Ku
belet stopped posting node status.
  PIDPressure          Unknown   Mon, 25 Apr 2022 09:25:20 +0000   Mon, 25 Apr 2022 09:30:07 +0000   NodeStatusUnknown   Ku
belet stopped posting node status.
  Ready                Unknown   Mon, 25 Apr 2022 09:25:20 +0000   Mon, 25 Apr 2022 09:30:07 +0000   NodeStatusUnknown   Ku
belet stopped posting node status.
Addresses:
  InternalIP:  10.250.5.52
  Hostname:    wk8s-node-0
Capacity:
  cpu:                16
  ephemeral-storage:  71088740Ki
  hugepages-1Gi:      0
  hugepages-2Mi:      0
  memory:             31827272Ki
  pods:               110
Allocatable:
  cpu:                16
  ephemeral-storage:  65515382676
  hugepages-1Gi:      0
```

```
  ephemeral-storage:  65515382676
  hugepages-1Gi:      0
  hugepages-2Mi:      0
  memory:             31724872Ki
  pods:               110
System Info:
  Machine ID:                 2107786af1744dfbbf02d9f6fac470b0
  System UUID:                ec22a34d-9b09-cea1-d7eb-1b47b08d2151
  Boot ID:                    3b22c15f-7dd1-4f61-b5c2-f24d9bd8b281
  Kernel Version:             5.11.0-1028-aws
  OS Image:                   Ubuntu 20.04.3 LTS
  Operating System:           linux
  Architecture:               amd64
  Container Runtime Version:  docker://20.10.7
  Kubelet Version:            v1.23.1
  Kube-Proxy Version:         v1.23.1
PodCIDR:                      10.244.1.0/24
PodCIDRs:                     10.244.1.0/24
Non-terminated Pods:          (2 in total)
  Namespace                   Name                    CPU Requests  CPU Limits  Memory Requests  Memory Lim
  ---------                   ----                    ------------  ----------  ---------------  ----------
  kube-system                 kube-flannel-ds-rxbx8   100m (0%)     100m (0%)   50Mi (0%)        50Mi (0%)
  kube-system                 kube-proxy-xfxzm        0 (0%)        0 (0%)      0 (0%)           0 (0%)
Allocated resources:
  (Total limits may be over 100 percent, i.e., overcommitted.)
  Resource           Requests   Limits
  --------           --------   ------
  cpu                100m (0%)  100m (0%)
  memory             50Mi (0%)  50Mi (0%)
  ephemeral-storage  0 (0%)     0 (0%)
  hugepages-1Gi      0 (0%)     0 (0%)
  hugepages-2Mi      0 (0%)     0 (0%)
Events:              <none>
student@node-1:~$ kubectl config use-context wk8s
Switched to context "wk8s".
student@node-1:~$ kubectl get nodes
NAME           STATUS     ROLES                  AGE   VERSION
wk8s-master-0  Ready      control-plane,master   67d   v1.23.1
wk8s-node-0    NotReady   <none>                 67d   v1.23.1
wk8s-node-1    Ready      <none>                 67d   v1.23.1
student@node-1:~$ ssh wk8s-node-0
Warning: Permanently added '10.250.5.52' (ECDSA) to the list of known hosts.
Welcome to Ubuntu 20.04.3 LTS (GNU/Linux 5.11.0-1028-aws x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/advantage

  System information as of Mon Apr 25 09:31:01 UTC 2022

  System load:  2.05               Processes:               40
  Usage of /:   83.2% of 67.79GB   Users logged in:         0
  Memory usage: 0%                 IPv4 address for docker0: 172.17.0.1
  Swap usage:   0%                 IPv4 address for eth0:    10.250.5.52


30 updates can be applied immediately.
15 of these updates are standard security updates.
To see these additional updates run: apt list --upgradable


The list of available updates is more than a week old.
To check for new updates run: sudo apt update

student@wk8s-node-0:~$ sudo -i
root@wk8s-node-0:~# systemctl enable --now kubelet
Created symlink /etc/systemd/system/multi-user.target.wants/kubelet.service → /lib/systemd/system/kubelet.se
root@wk8s-node-0:~# systemctl restart kubelet
root@wk8s-node-0:~# systemctl status kube
```

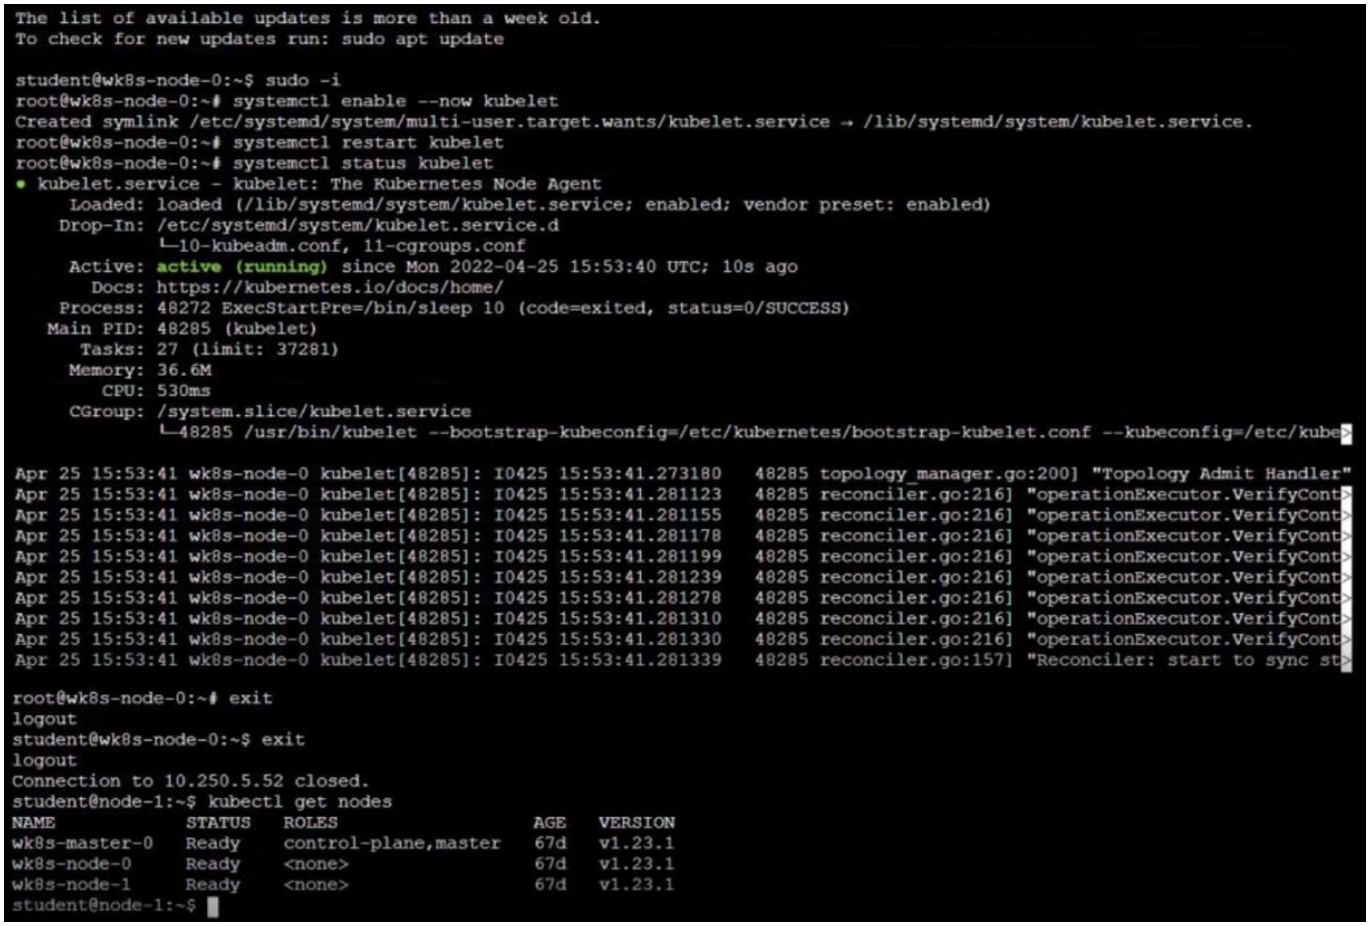

QUESTION 9

Create a pod that having 3 containers in it? (Multi-Container)

Correct Answer: Check the answer in 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

QUESTION 10

SIMULATION

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

Correct Answer: Check the answer in explanation.

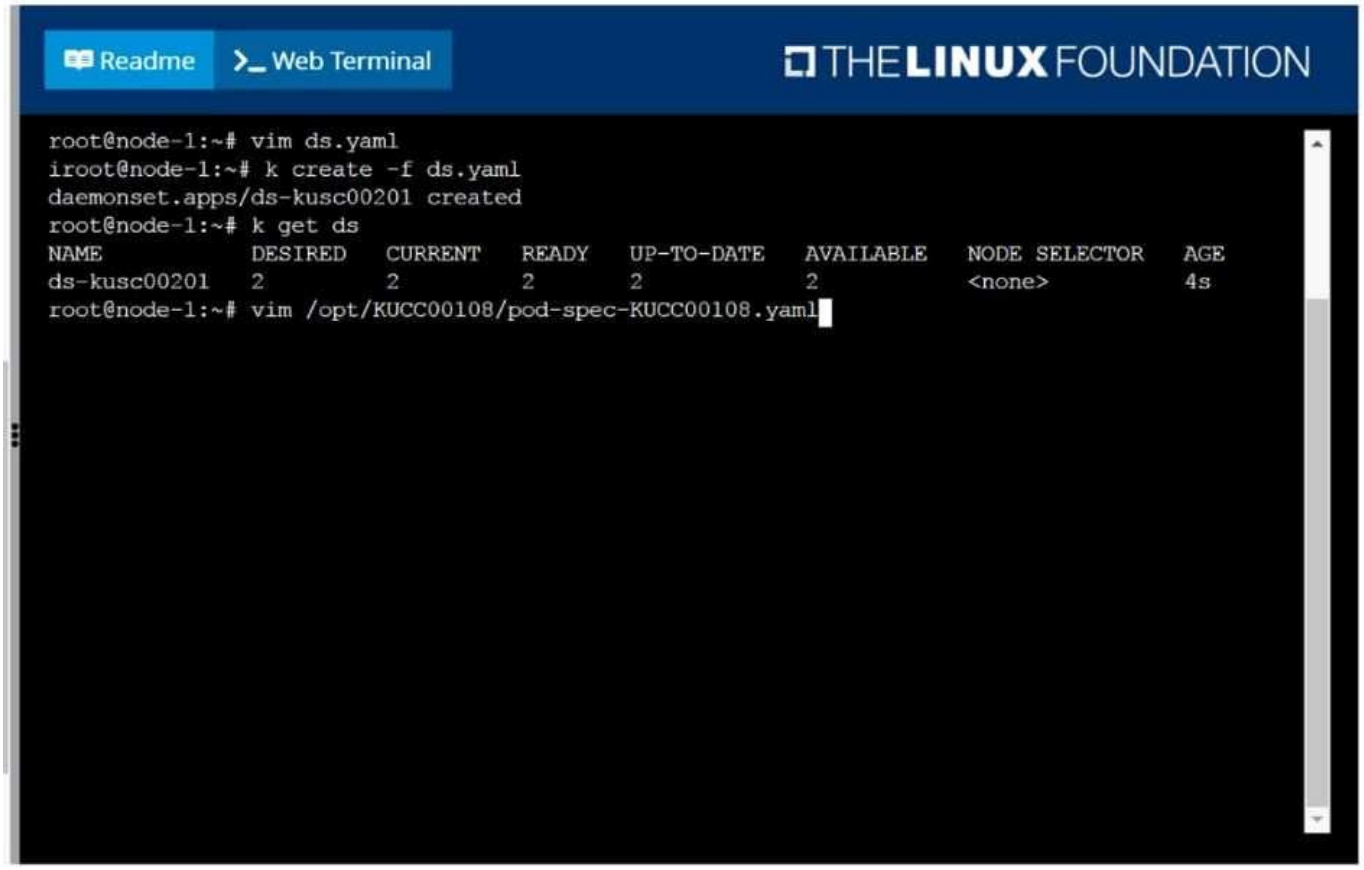

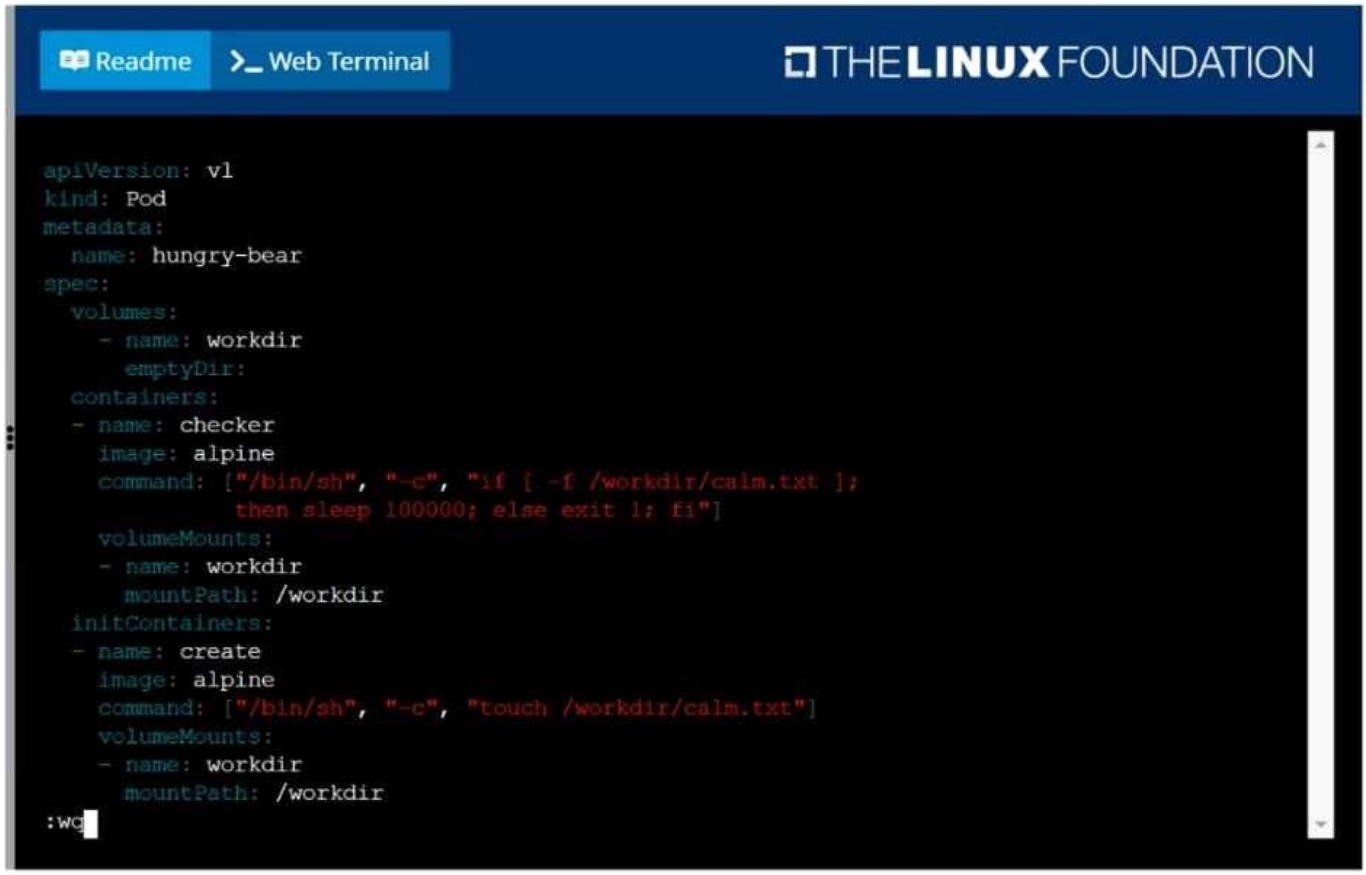
Readme
Web Terminal
THE LINUX FOUNDATION
apiVersion: v1
kind: Pod
metadata:
name: hungry-bear
spec:
volumes:
- name: workdir
emptyDir:
containers:
- name: checker
image: alpine
command: ["/bin/sh", "-c", "if [-f /workdir/calm.txt];
then sleep 100000; else exit 1; fi"]
volumeMounts:
- name: workdir
mountPath: /workdir
initContainers:
- name: create
image: alpine
command: ["/bin/sh", "-c", "touch /workdir/calm.txt"]
volumeMounts:
- name: workdir
mountPath: /workdir
:wq

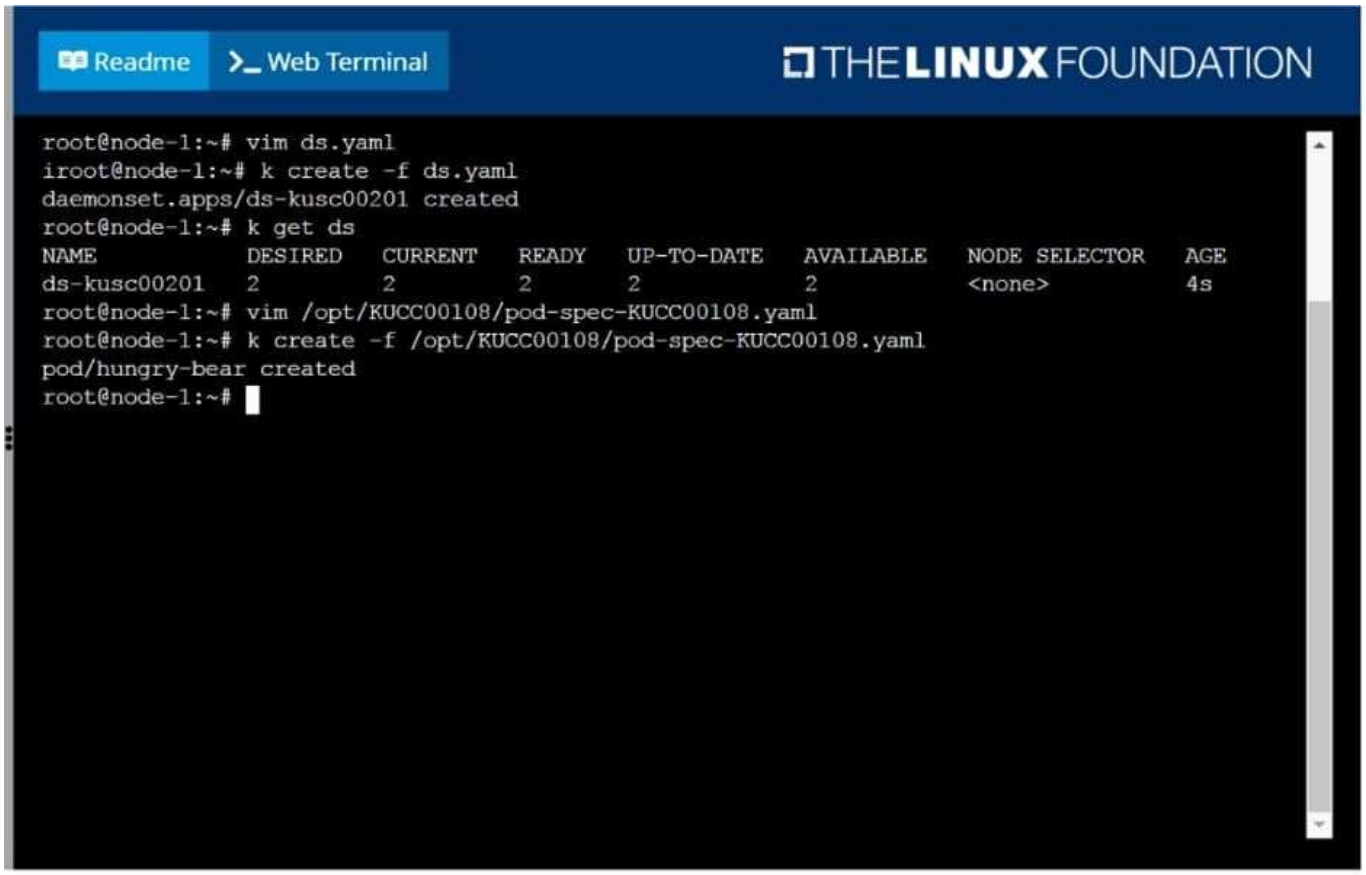
Readme
Web Terminal
THE LINUX FOUNDATION
root@node-1:~# vim ds.yaml
iroot@node-1:~# k create -f ds.yaml
daemonset.apps/ds-kusc00201 created
root@node-1:~# k get ds
NAME DESIRED CURRENT READY UP-TO-DATE AVAILABLE NODE SELECTOR AGE
ds-kusc00201 2 2 2 2 2 <none> 4s
root@node-1:~# vim /opt/KUCC00108/pod-spec-KUCC00108.yaml
root@node-1:~# k create -f /opt/KUCC00108/pod-spec-KUCC00108.yaml
pod/hungry-bear created
root@node-1:~#

QUESTION 11

SIMULATION 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

Correct Answer: Check the answer in explanation.

Solution

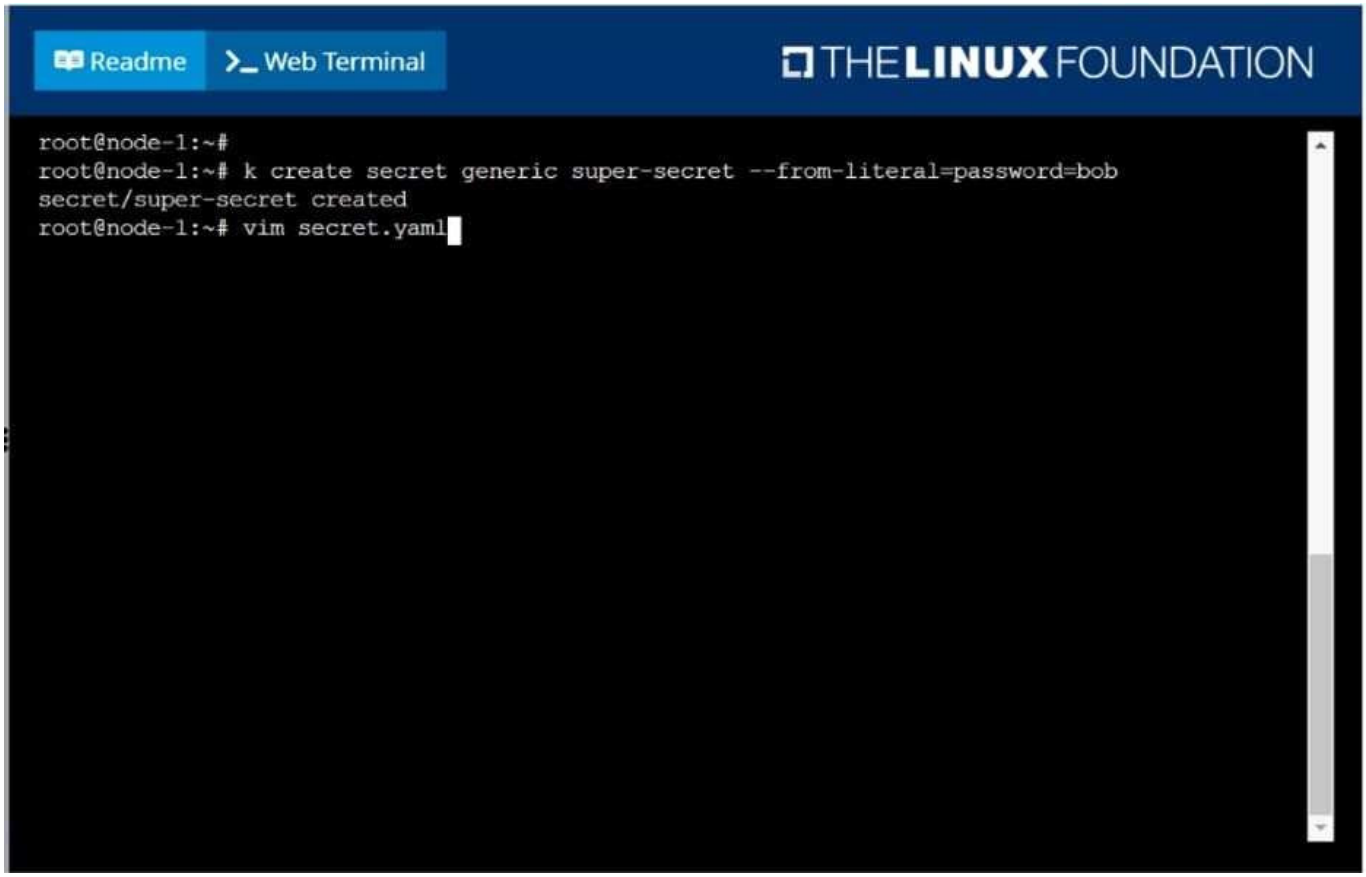

Readme | Web Terminal — THE LINUX FOUNDATION

```
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
```

Readme | Web Terminal — THE LINUX FOUNDATION

```
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-gl2kh    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:~#
```

QUESTION 12

SIMULATION

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.

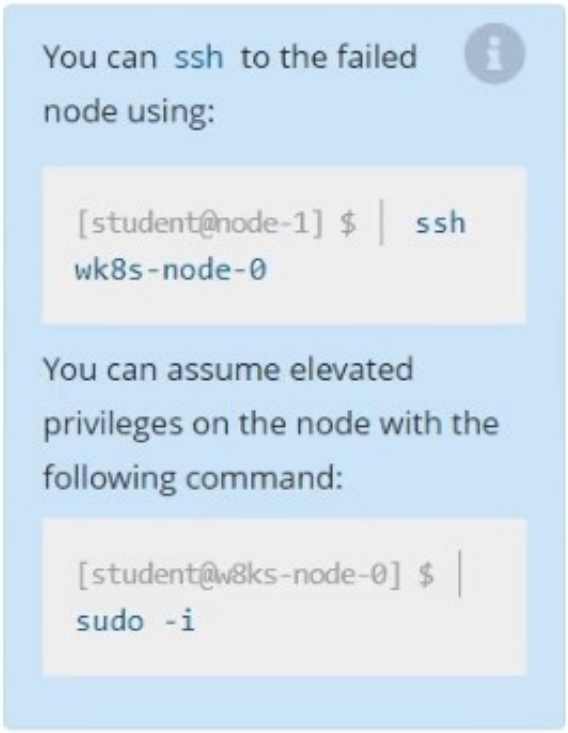

You can ssh to the failed node using:

[student@node-1] $ | ssh Wk8s-node-0

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

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

Correct Answer: Check the answer in explanation.

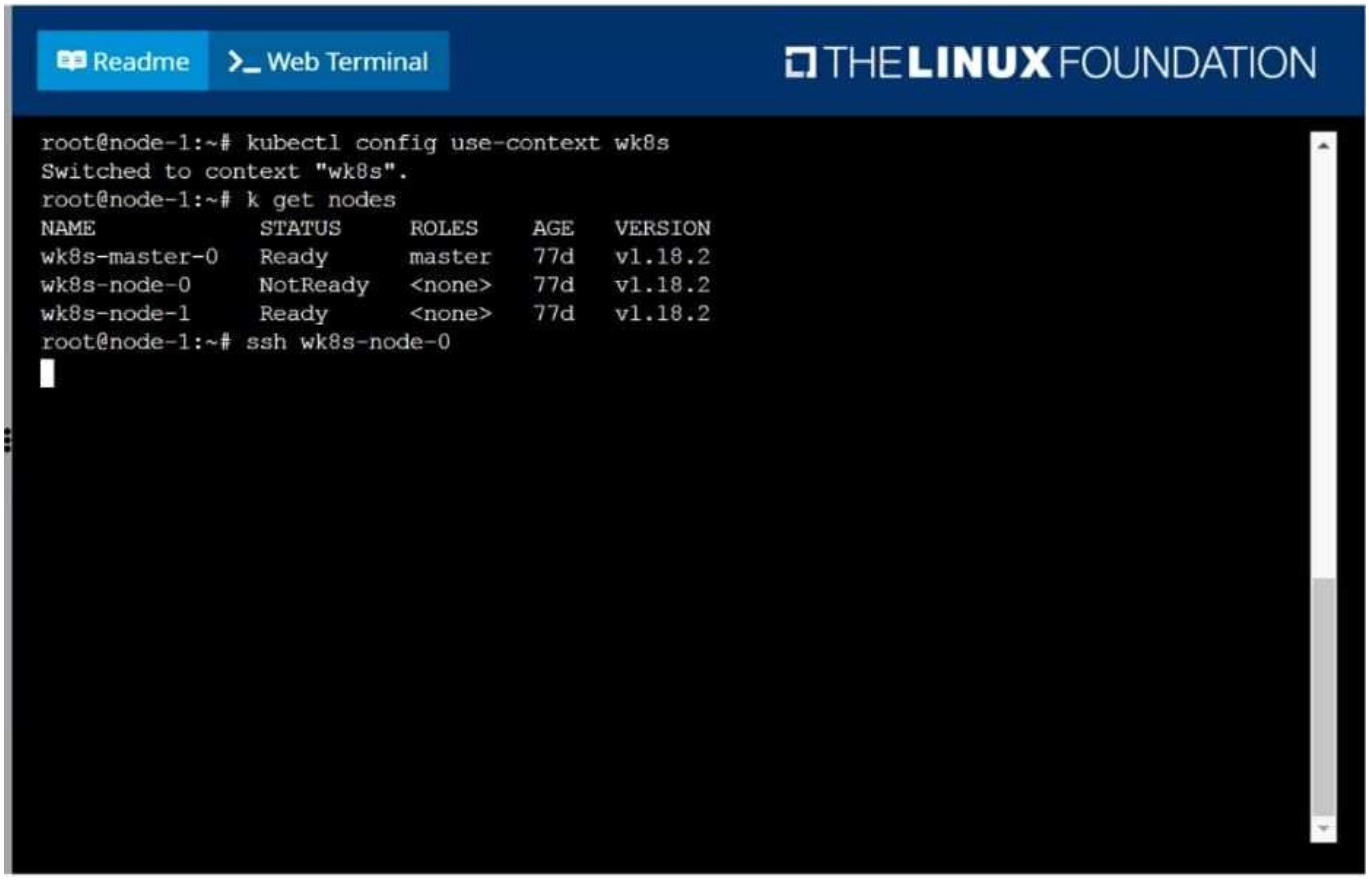
Readme
Web Terminal
THE LINUX FOUNDATION
root@node-1:~# kubectl config use-context wk8s
Switched to context "wk8s".
root@node-1:~# k get nodes
NAME STATUS ROLES AGE VERSION
wk8s-master-0 Ready master 77d v1.18.2
wk8s-node-0 NotReady <none> 77d v1.18.2
wk8s-node-1 Ready <none> 77d v1.18.2
root@node-1:~# ssh wk8s-node-0

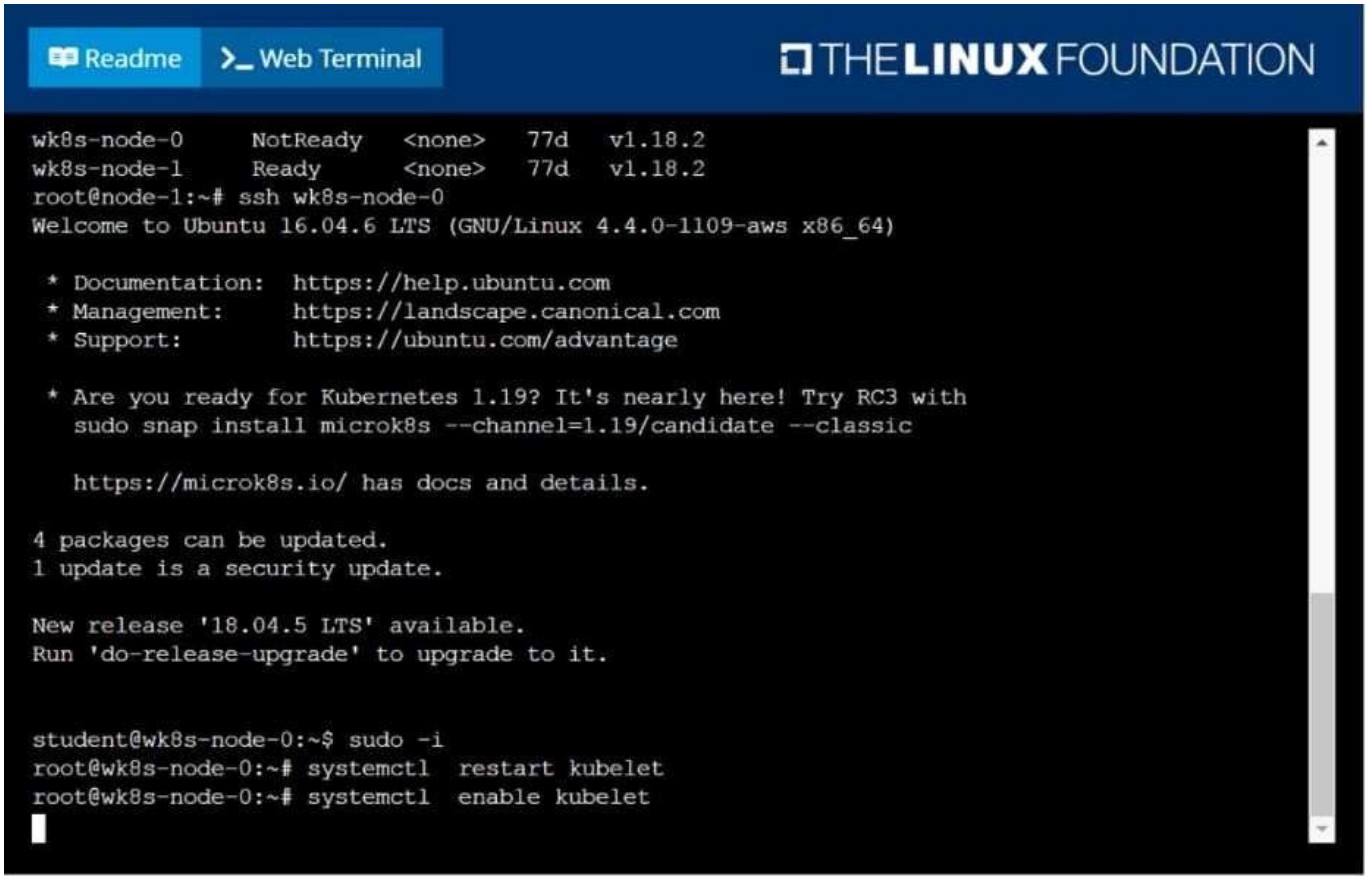
Readme
Web Terminal
THE LINUX FOUNDATION
wk8s-node-0 NotReady <none> 77d v1.18.2
wk8s-node-1 Ready <none> 77d v1.18.2
root@node-1:~# ssh wk8s-node-0
Welcome to Ubuntu 16.04.6 LTS (GNU/Linux 4.4.0-1109-aws x86_64)
* Documentation: https://help.ubuntu.com
* Management: https://landscape.canonical.com
* 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

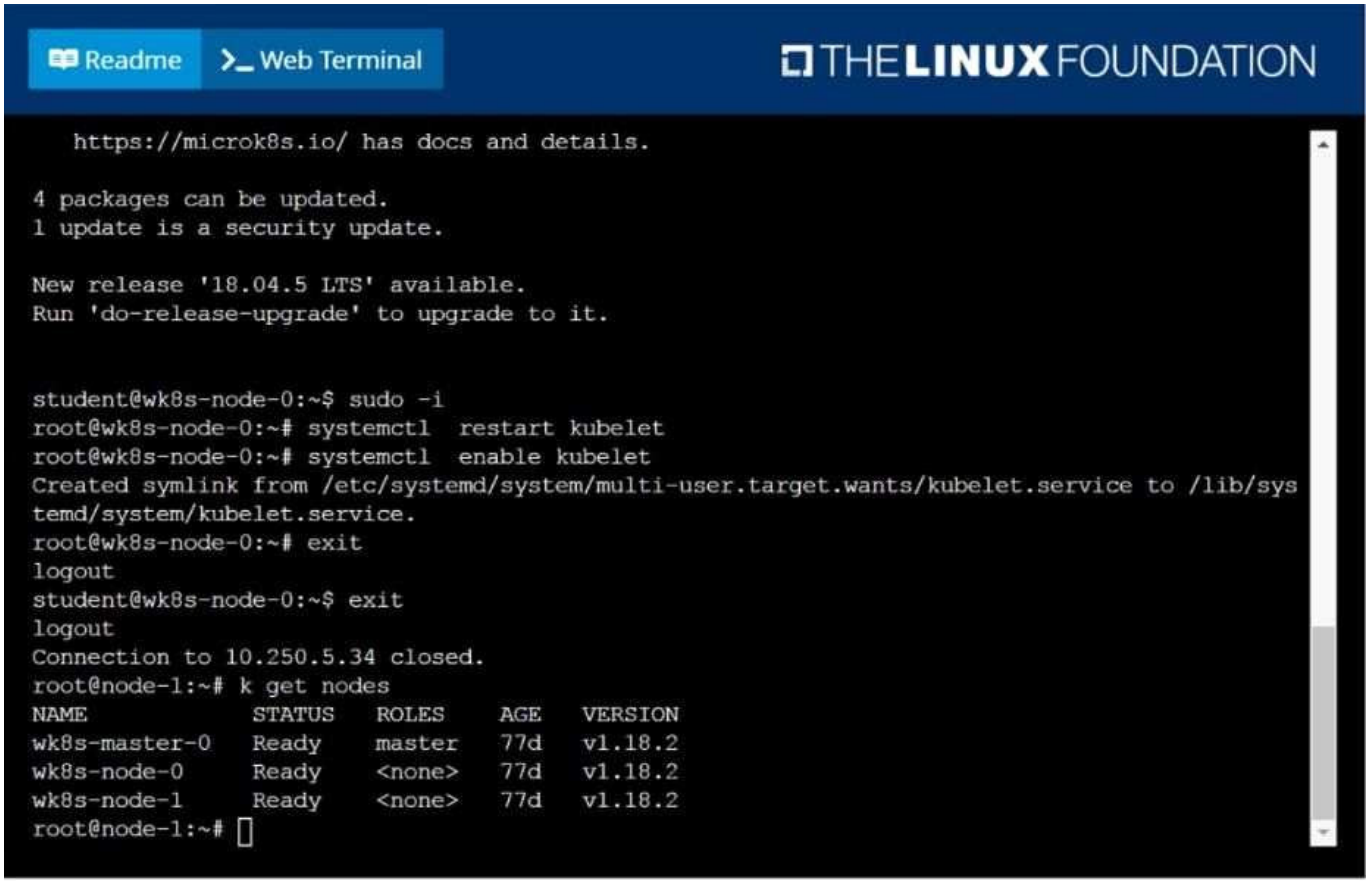

QUESTION 13

SIMULATION

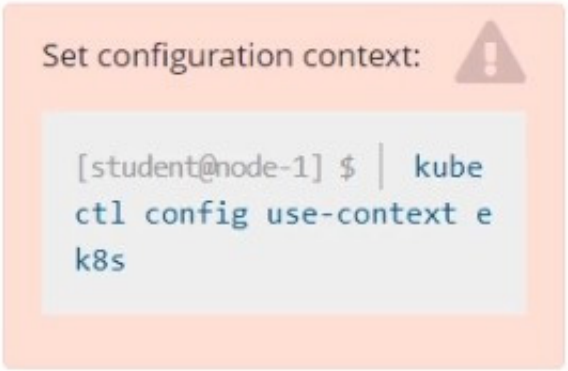

Set the node named ek8s-node-1 as unavailable and reschedule all the pods running on it.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context ek8s
Switched to context "ek8s".
student@node-1:~$ kubectl get nodes
NAME           STATUS   ROLES                  AGE   VERSION
ek8s-master-0  Ready    control-plane,master   67d   v1.23.1
ek8s-node-0    Ready    <none>                 67d   v1.23.1
ek8s-node-1    Ready    <none>                 67d   v1.23.1
student@node-1:~$ kubectl drain ek8s-node-1 --ignore-daemonsets
node/ek8s-node-1 cordoned
error: unable to drain node "ek8s-node-1" due to error:cannot delete Pods with local storage (use --delete-emptydir-data to
 override): kube-system/metrics-server-7cb5455c67-m6qvd, continuing command...
There are pending nodes to be drained:
 ek8s-node-1
cannot delete Pods with local storage (use --delete-emptydir-data to override): kube-system/metrics-server-7cb5455c67-m6qvd
student@node-1:~$ kubectl drain ek8s-node-1 --ignore-daemonsets --delete-emptydir-data
node/ek8s-node-1 already cordoned
WARNING: ignoring DaemonSet-managed Pods: kube-system/kube-flannel-ds-chvkf, kube-system/kube-proxy-7pf29
evicting pod kube-system/metrics-server-7cb5455c67-m6qvd
evicting pod default/nginx-5cb786cffd-vjbs8
pod/nginx-5cb786cffd-vjbs8 evicted
pod/metrics-server-7cb5455c67-m6qvd evicted
node/ek8s-node-1 drained
student@node-1:~$ kubectl get nodes
NAME           STATUS                     ROLES                  AGE   VERSION
ek8s-master-0  Ready                      control-plane,master   67d   v1.23.1
ek8s-node-0    Ready                      <none>                 67d   v1.23.1
ek8s-node-1    Ready,SchedulingDisabled   <none>                 67d   v1.23.1
student@node-1:~$
```

QUESTION 14

SIMULATION

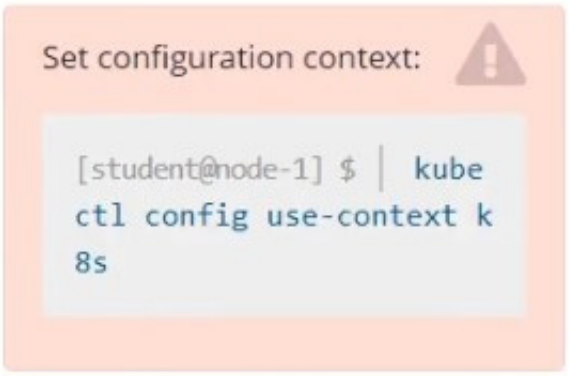

Task

Scale the deployment presentation to 6 pods.

Correct Answer: Check the answer in explanation.

```
student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ vim ping.yml
student@node-1:~$ kubectl create -f ping.yml
ingress.networking.k8s.io/ping created
student@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
student@node-1:~$ kubectl get deployments.apps
NAME           READY   UP-TO-DATE   AVAILABLE   AGE
front-end      2/2     2            2           6h2m
presentation   2/2     2            2           6h1m
student@node-1:~$ kubectl scale deployment presentation --replicas=3
deployment.apps/presentation scaled
student@node-1:~$ kubectl get deployments.apps
NAME           READY   UP-TO-DATE   AVAILABLE   AGE
front-end      2/2     2            2           6h2m
presentation   2/3     3            2           6h1m
student@node-1:~$ kubectl get pods
NAME                            READY   STATUS    RESTARTS   AGE
big-corp-app                    1/1     Running   0          5h58m
foo                             1/1     Running   0          5h58m
front-end-6bc87b9748-n7v8h      1/1     Running   0          3m47s
front-end-6bc87b9748-zmb8g      1/1     Running   0          3m45s
overloaded-cpu-98b9se           1/1     Running   0          5h57m
overloaded-cpu-ab2d3s           1/1     Running   0          5h57m
overloaded-cpu-kipb9a           1/1     Running   0          5h57m
presentation-684cd7ccb6-4gf56   1/1     Running   0          6h1m
presentation-684cd7ccb6-6zjls   1/1     Running   0          13s
presentation-684cd7ccb6-vshxj   1/1     Running   0          6h1m
student@node-1:~$
```

QUESTION 15

SIMULATION

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.

Correct Answer: Check the answer in explanation.

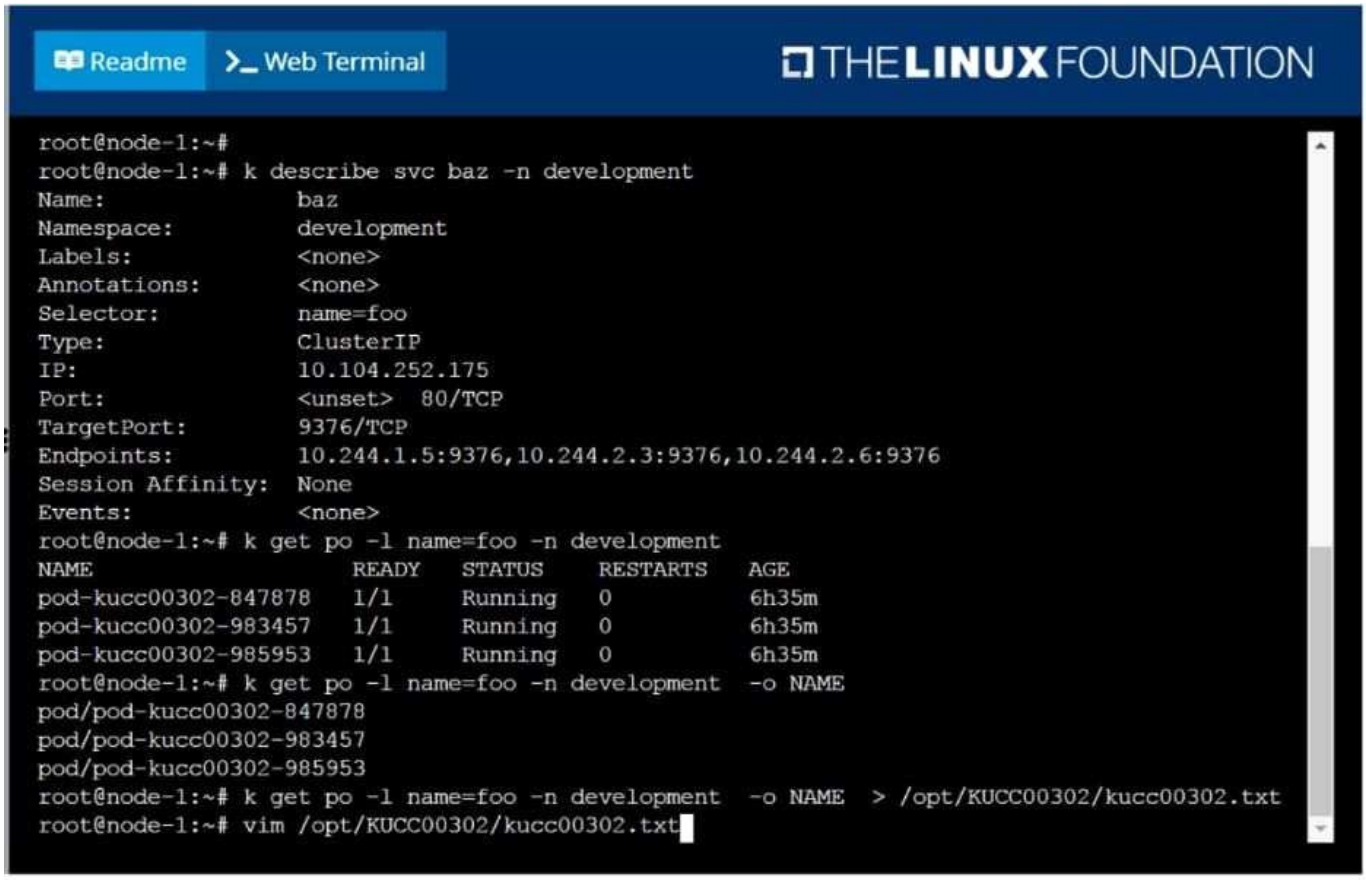
Readme
Web Terminal
THE LINUX FOUNDATION
root@node-1:~#
root@node-1:~# k describe svc baz -n development
Name: baz
Namespace: development
Labels: <none>
Annotations: <none>
Selector: name=foo
Type: ClusterIP
IP: 10.104.252.175
Port: <unset> 80/TCP
TargetPort: 9376/TCP
Endpoints: 10.244.1.5:9376,10.244.2.3:9376,10.244.2.6:9376
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

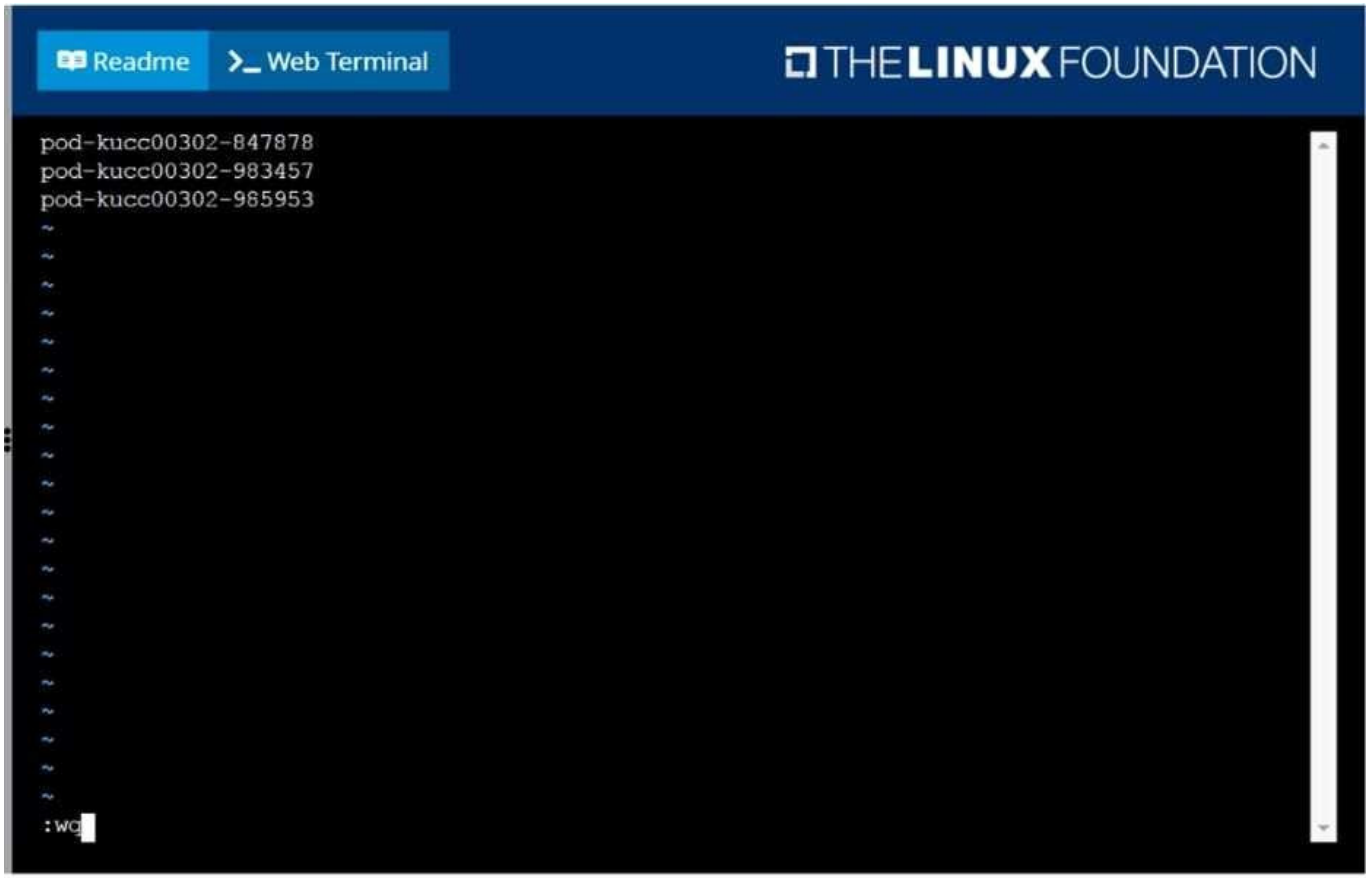
Readme
Web Terminal
THE LINUX FOUNDATION
pod-kucc00302-847878
pod-kucc00302-983457
pod-kucc00302-985953
:wq