

QUESTIONS & ANSWERS

Kill your exam at first Attempt



CNCF

CKA

Certified Kubernetes Administrator

<https://killexams.com/pass4sure/exam-detail/CKA>




Question: 11

CORRECT TEXT

Create and configure the service front-end-service so it's accessible through NodePort and routes to the existing pod named front-end.

Answer: solution



The screenshot shows a web terminal interface with a dark background. At the top, there are tabs for 'Readme' and 'Web Terminal', and a logo for 'THE LINUX FOUNDATION'. The terminal content shows the following sequence of commands and output:

```
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:~#
```

Question: 12

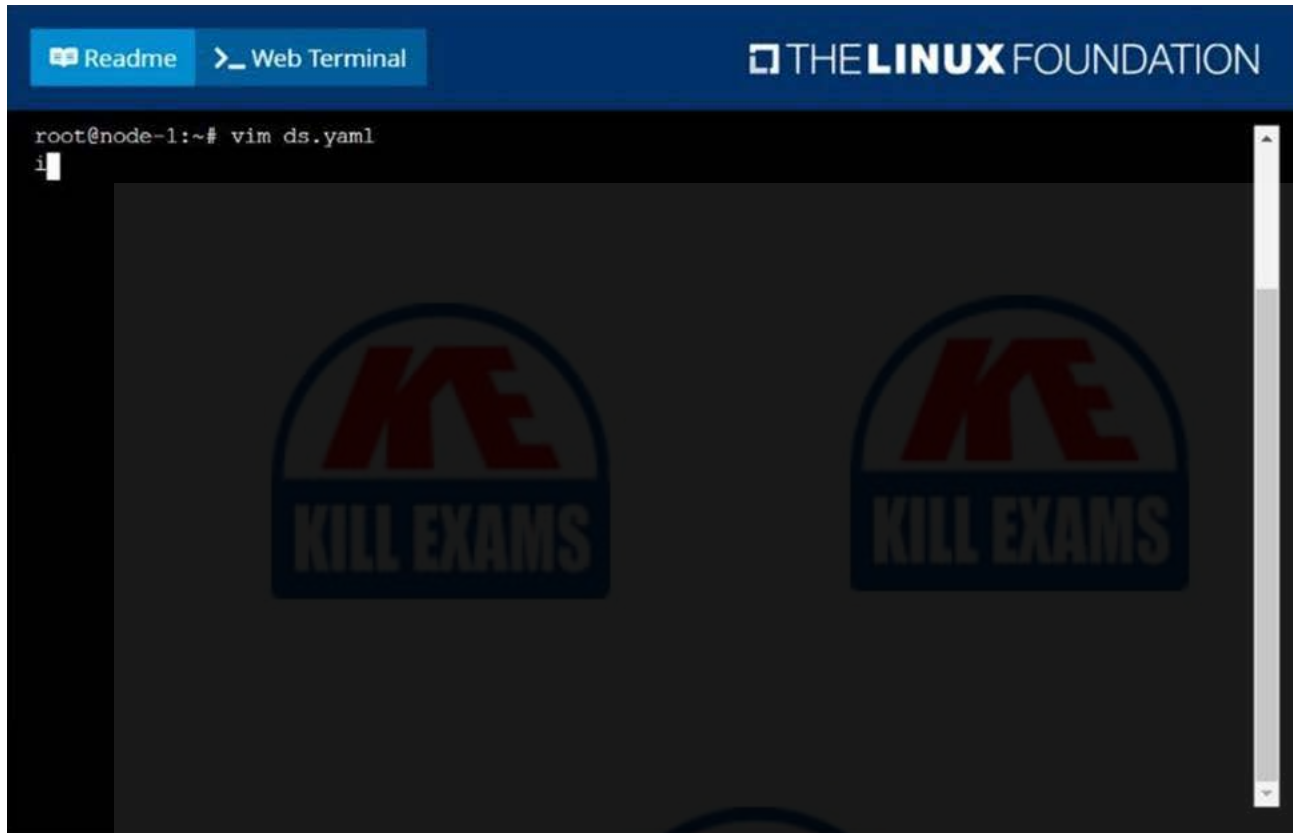
CORRECT TEXT

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.

Answer: solution



```
apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: fluentd-elasticsearch
  namespace: kube-system
  labels:
    k8s-app: fluentd-logging
spec:
  selector:
    matchLabels:
      name: fluentd-elasticsearch
  template:
    metadata:
      labels:
        name: fluentd-elasticsearch
    spec:
      tolerations:
        # this toleration is to have the daemonset runnable on master nodes
        # remove it if your masters can't run pods
        - key: node-role.kubernetes.io/master
          effect: NoSchedule
      containers:
        - name: nginx
          image: nginx
-- INSERT --
```

17,19 All

```
apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: ds-kusc00201
spec:
  selector:
    matchLabels:
      name: fluentd-elasticsearch
  template:
    metadata:
      labels:
        name: fluentd-elasticsearch
    spec:
      containers:
        - name: nginx
          image: nginx
```

```
~
~
~
~
~
~
~
~
~
:wc
```

```
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:~#
```

Question: 13

CORRECT TEXT

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).

Answer: solution


```
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
```

cpu-utilizer-98b9se

:WQ

Question: 14

CORRECT TEXT

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

Answer: solution



The screenshot shows a web terminal interface with a dark background. At the top, there is a blue header bar with a "Readme" button and a "Web Terminal" button. The terminal content shows a series of commands and their outputs:

```
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
```

The terminal also features a large, semi-transparent watermark in the center that reads "KILL EXAMS" with a stylized "KE" logo above it.

```
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
:WG
```

```
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: 15


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.

Answer: 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
```

pod-kucc00302-847878
pod-kucc00302-983457
pod-kucc00302-985953

```
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
```

```
root@node-1:~# vim /opt/KUCC00302/kucc00302.txt
```

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

Question: 16

CORRECT TEXT

Create a deployment spec file that will:

Launch 7 replicas of the nginx Image with the labelapp_runtime_stage=dev


deployment name: kual00201

Save a copy of this spec file to /opt/KUAL00201/spec_deployment.yaml

(or /opt/KUAL00201/spec_deployment.json).

When you are done, clean up (delete) any new Kubernetes API object that you produced during this task.

Answer: solution



The screenshot shows a web terminal interface with a dark background. At the top, there is a blue header bar with a "Readme" button and a "Web Terminal" button. To the right of the buttons is the "THE LINUX FOUNDATION" logo. The terminal window displays the following commands and output:

```
root@node-1:~# k create deploy kual00201 --image=nginx --dry-run=client -o yaml > /opt/KUAL00201/spec_deployment.yaml
root@node-1:~# vim /opt/KUAL00201/spec_deployment.yaml
```

A large, semi-transparent watermark is overlaid on the terminal window. It features a stylized "KE" logo inside a circle, with the text "KILL EXAMS" written below it. The watermark is repeated multiple times across the screen.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app_runtime_stage: dev
  name: kual00201
spec:
  replicas: 7
  selector:
    matchLabels:
      app_runtime_stage: dev
  template:
    metadata:
      labels:
        app_runtime_stage: dev
    spec:
      containers:
      - image: nginx
        name: nginx
~
~
~
~
~
"/opt/KUAL00201/spec_deployment.yaml" 19L, 320C written
```

Question: 17

CORRECT TEXT

Create a Kubernetes secret as follows:

Name: super-secret

password: bob

Create a pod named pod-secrets-via-file, using the redis Image, which mounts a secret named super-secret at /secrets.

Create a second pod named pod-secrets-via-env, using the redis Image, which exports password as CONFIDENTIAL

Answer: solution

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

```
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
```

```
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: 18

CORRECT TEXT

Score: 7%

Set configuration context: 

```
[student@node-1] $ | kube
ctl config use-context m
k8s
```



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.

Answer: SOLUTION:

```
[[email protected]] > 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 --disable-excludes=kubernetes
```

```
kubeadm upgrade apply 1.20.1 --etcd-upgrade=false
```

```
systemctl daemon-reload
```

```
systemctl restart kubelet
```

```
kubectl uncordon k8s-master
```

Question: 19

CORRECT TEXT

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:

```
[email protected] $ | ssh Wk8s-node-0
```

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

```
[email protected] $ | sudo Ci
```

Answer: solution



The screenshot shows a web terminal window with a dark blue header. On the left, there are tabs for 'Readme' and 'Web Terminal'. On the right, it says 'THE LINUX FOUNDATION'. The terminal content shows the following commands and output:

```
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
```

The terminal also features a large, semi-transparent watermark in the center that reads 'KILL EXAMS' with a stylized 'KE' logo above it.

```
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
```

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

4 packages can be updated.
1 update is a security update.

New release '18.04.5 LTS' available.
Run 'do-release-upgrade' to upgrade to it.

student@wk8s-node-0:~$ sudo -i
root@wk8s-node-0:~# systemctl restart kubelet
root@wk8s-node-0:~# systemctl enable kubelet
Created symlink from /etc/systemd/system/multi-user.target.wants/kubelet.service to /lib/sy
stemd/system/kubelet.service.
root@wk8s-node-0:~# exit
logout
student@wk8s-node-0:~$ exit
logout
Connection to 10.250.5.34 closed.
root@node-1:~# k get nodes
NAME          STATUS    ROLES    AGE   VERSION
wk8s-master-0 Ready     master   77d   v1.18.2
wk8s-node-0   Ready     <none>    77d   v1.18.2
wk8s-node-1   Ready     <none>    77d   v1.18.2
root@node-1:~#
```

Question: 20

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.

Answer: solution



```
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
```

```
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:~#
```

For More exams visit <https://killexams.com/vendors-exam-list>

