LinuxFoundation.CKA.v2022-12-07.q68

Exam Code:	CKA
Exam Name:	Certified Kubernetes Administrator (CKA) Program Exam
Certification Provider:	Linux Foundation
Free Question Number:	68
Version:	v2022-12-07
# of views:	1765
# of Questions views:	680
https://www.freeqas.com/qa/Linux-Foundation/CKA/LinuxFoundation.CKA.v2022-12-07.q68.html	

NEW QUESTION: 1
Pause the rollout of the deployment
Answer:
kubectl rollout pause deploy webapp

NEW QUESTION: 2
Create a Pod nginx and specify both CPU, memory requests and limits together and verify.
A. kubectl run nginx-request --image=nginx --restart=Always --dryrun -o yaml > nginx-request.yml
// add the resources section and create
apiVersion: v1
kind: Pod
metadata:
labels:
run: nginx
name: nginx
resources:
requests:
memory: "100Mi"
cpu: "0.4"
limits:
memory: "200Mi"
cpu: "7"
restartPolicy: Always
k kubectl apply -f nginx-request.yaml
// Verify
Kubectl top po
B. kubectl run nginx-request --image=nginx --restart=Always --dryrun -o yaml > nginx-request.yml

// add the resources section and create
apiVersion: v1
kind: Pod
metadata:
labels:
run: nginx
name: nginx-request
spec:
containers:
- image: nginx
name: nginx
resources:
requests:
memory: "100Mi"
cpu: "0.5"
limits:
memory: "200Mi"
cpu: "1"
restartPolicy: Always
k kubectl apply -f nginx-request.yaml
// Verify
Kubectl top po
Answer: B (LEAVE A REPLY)

NEW QUESTION: 3
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 .
Answer:
See the solution below.
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: 4
Change the Image version back to 1.17.1 for the pod you just updated and observe the changes
Answer:
kubectl set image pod/nginx nginx=nginx:1.17.1 kubectl describe po nginx kubectl get po nginx -w # watch it

NEW QUESTION: 5
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:
See the solution below.
Explanation
solution

Readme | >_ Web Terminal

THE LINUX FOUNDATION

```
root@node-1:~# k create deploy kual00201  --image=nginx --dry-run=client -o yaml > /opt/KUAL
00201/spec_deployment.yaml
root@node-1:~# vim /opt/KUAL00201/spec_deployment.yaml
```

freeqas.com

NEW QUESTION: 6

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

Answer:

See the solution below.

Explanation

solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\8 B.JPG

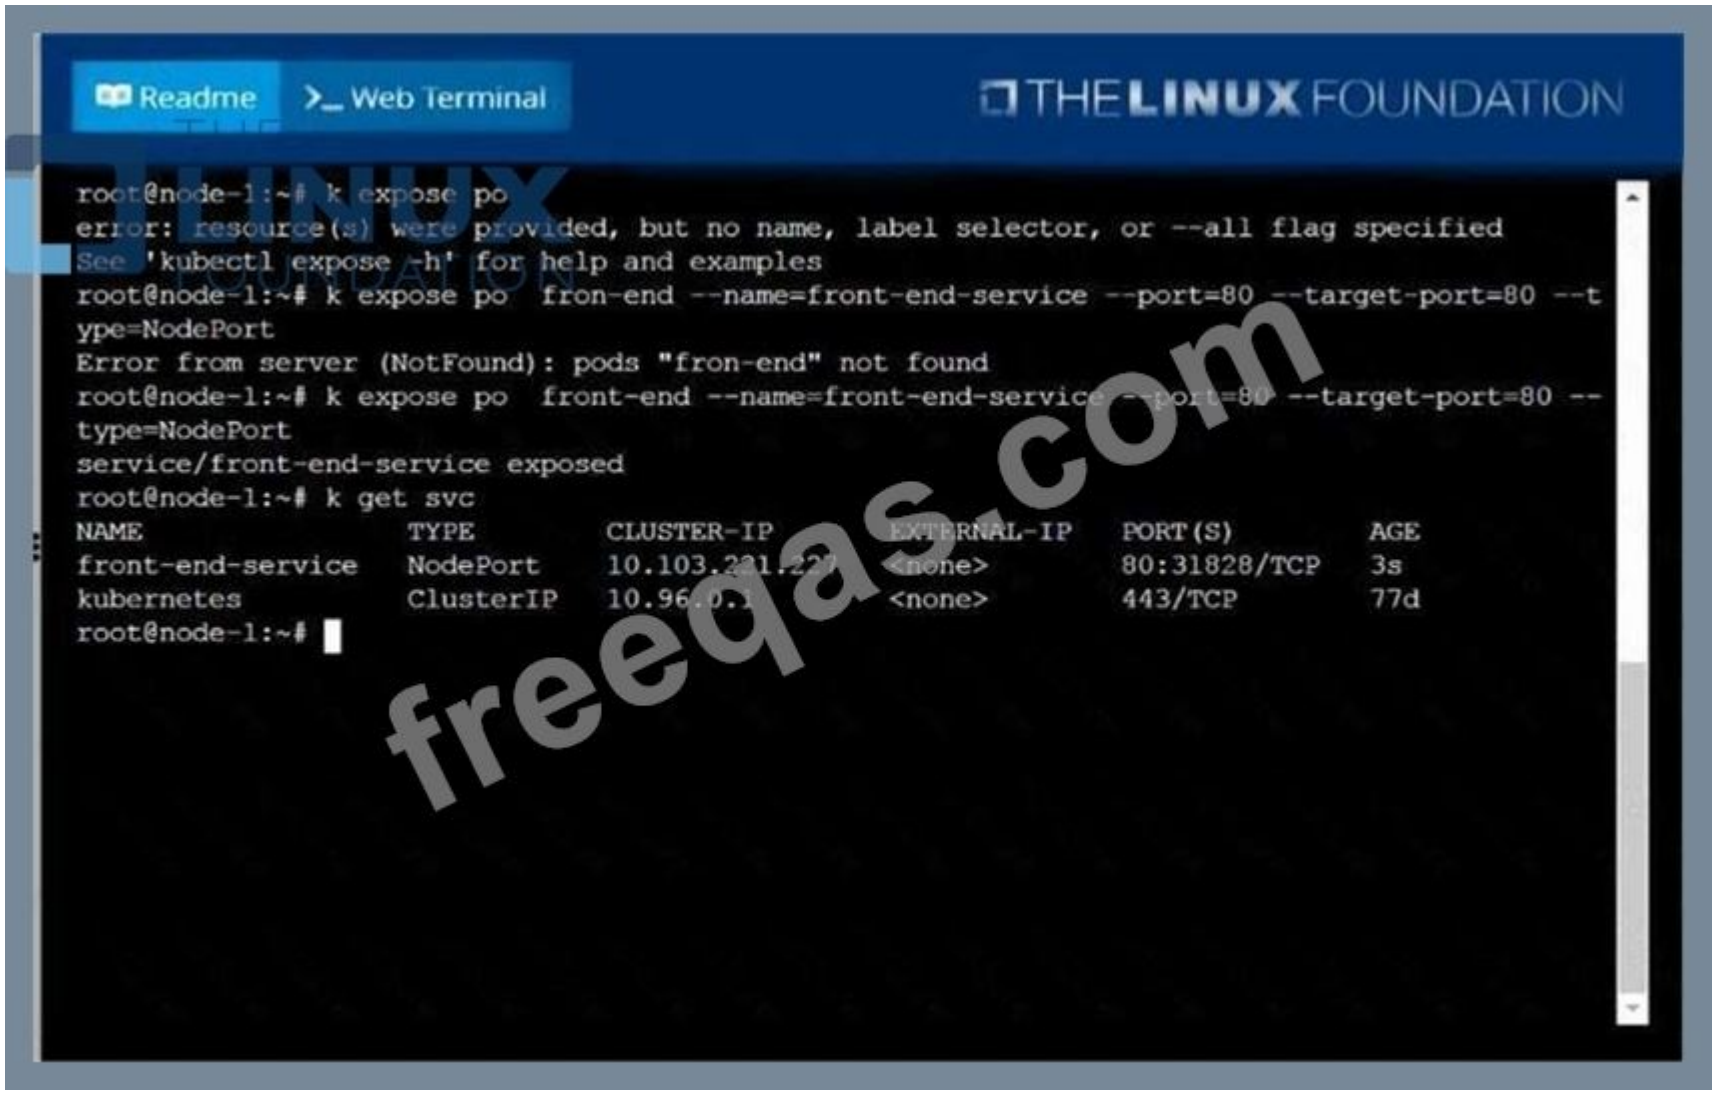

freeqas.com

NEW QUESTION: 7

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

Answer:

See the solution below.

Explanation

solution

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

freeqas.com

NEW QUESTION: 8
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
Answer:
solution

Readme | Web Terminal

THE LINUX FOUNDATION

```
root@node-1:~#
root@node-1:~# kubectl config use-context wk8s
Switched to context "wk8s".
root@node-1:~# ssh wk8s-node-1
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-1:~$ sudo -i
root@wk8s-node-1:~# vim /var/lib/kubelet/config.yaml
```

freeqas.com

THE LINUX FOUNDATION

Readme | >_ Web Terminal

THE LINUX FOUNDATION

```
    clientCAFile: /etc/kubernetes/pki/ca.crt
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
:wq
```

freeqas.com

Readme | >_ Web Terminal

THE LINUX FOUNDATION

```
root@node-1:~# ssh wk8s-node-1
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-1:~$ sudo -i
root@wk8s-node-1:~# vim /var/lib/kubelet/config.yaml
root@wk8s-node-1:~# cd /etc/kubernetes/manifests
root@wk8s-node-1:/etc/kubernetes/manifests#
root@wk8s-node-1:/etc/kubernetes/manifests# vim pod.yaml
```

freeqas.com


Readme
>_ Web Terminal
THE LINUX FOUNDATION
apiVersion: v1
kind: Pod
metadata:
name: webtool
spec:
containers:
- name: webtool
image: httpd
:w
freeqas.com

freedas.com

NEW QUESTION: 9

Create a pod as follows:

* Name:mongo
* Using Image:mongo
* In anew Kubernetes namespacenamed:my-website

Answer:

See the solution below.

Explanation

solution

NEW QUESTION: 10
List all the pods showing name and namespace with a json path expression
Answer:
kubectl get pods -o=jsonpath="{.items[*]['metadata.name',
'metadata.namespace']}"

NEW QUESTION: 11
Create a pod with environment variables as var1=value1.Check the environment variable in pod
Answer:
kubectl run nginx --image=nginx --restart=Never --env=var1=value1
then
kubectl exec -it nginx -- env
or
kubectl exec -it nginx -- sh -c 'echo $var1'
or
kubectl describe po nginx | grep value1

NEW QUESTION: 12

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

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
```

freeqas.com

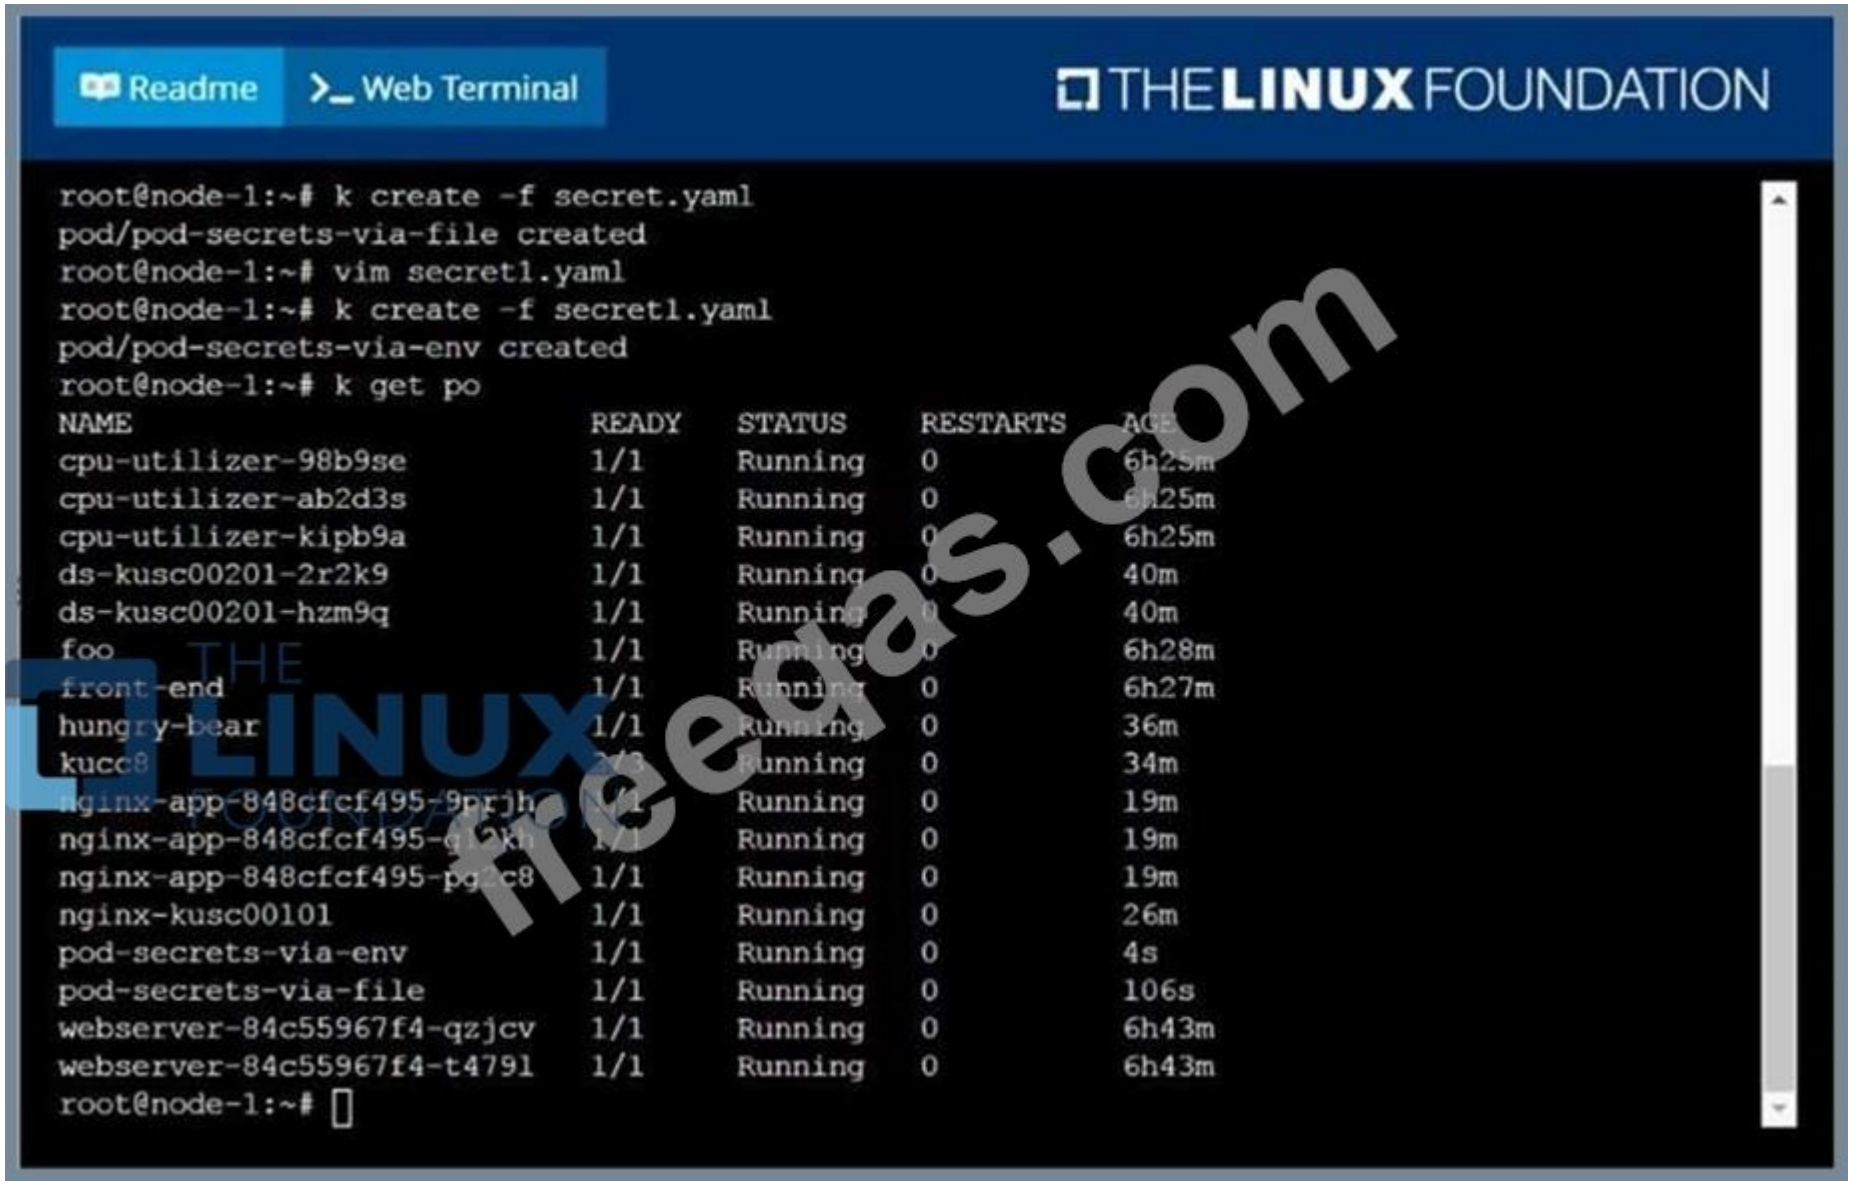

NEW QUESTION: 13
An Administrator is configuring Authentication Enforcement and they would like to create an exemption rule to exempt a specific group from authentication. Which authentication enforcement object should they select?
A. default-web-form
B. default-browser-challenge
C. default-no-captive-port
D. default-authentication-bypass
Answer: C (LEAVE A REPLY)

NEW QUESTION: 14
Print pod name and start time to "/opt/pod-status" file
Answer:
kubect1 get pods -o=jsonpath='{range .items[*]}{.metadata.name}{"\t"}{.status.podIP}{"\n"}{end}'

NEW QUESTION: 15

Get list of persistent volumes and persistent volume claim in the cluster
Answer:
kubectl get pv kubectl get pvc

NEW QUESTION: 16
Delete the pod without any delay (force delete)
Answer:
Kubect1 delete po "POD-NAME" --grace-period=0 --force

Valid CKA Dumps shared by PrepPdf.com for Helping Passing CKA Exam! PrepPdf.com now offer the **newest CKA exam dumps**, the PrepPdf.com CKA exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CKA dumps with Test Engine here: https://www.preppdf.com/Linux-Foundation/CKA-prepaway-exam-dumps.html (**85** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 17
What are the differences between using a service versus using an application for Security Policy match?
A. Use of a "service" enables the firewall to take immediate action with the first observed packet based on port numbers. Use of an application allows the firewall to take action after enough packets allow for App-ID identification regardless of the ports being used
B. Use of a "service" enables the firewall to take immediate action with the first observed packet based on port numbers. Use of an "application allows the firewall to take immediate action if the port being used is a member of the application standard port list
C. Use of a "service" enables the firewall to take action after enough packets allow for App-ID identification
D. There are no differences between "service" or "application." Use of an "application simplifies configuration by allowing use of a friendly application name instead of port numbers.
Answer: (SHOW ANSWER)

NEW QUESTION: 18
Create a snapshot of the etcd instance running at https://127.0.0.1:2379, saving the snapshot to the file path /srv/data/etcd-snapshot.db.
The following TLS certificates/key are supplied for connecting to the server with etcdctl:
CA certificate: /opt/KUCM00302/ca.crt
Client certificate: /opt/KUCM00302/etcd-client.crt
Client key: Topt/KUCM00302/etcd-client.key
Answer:
solution

NEW QUESTION: 19
Check logs of each container that "busyboxpod-{1,2,3}"
A. kubectl logs busybox -c busybox-container-1
kubectl logs busybox -c busybox-container-3
kubectl logs busybox -c busybox-container-3
B. kubectl logs busybox -c busybox-container-1
kubectl logs busybox -c busybox-container-2
kubectl logs busybox -c busybox-container-3
Answer: B (LEAVE A REPLY)

NEW QUESTION: 20
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 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
Answer:
See the solution below.
Explanation
solution

F:\Work\Data Entry Work\Data Entry\20200827\CKA\20 C.JPG

F:\Work\Data Entry Work\Data Entry\20200827\CKA\20 D.JPG

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
```

freeqas.com

F:\Work\Data Entry Work\Data Entry\20200827\CKA\20 E.JPG

freeqas.com

NEW QUESTION: 21

Check to see how many worker nodes are ready (not including nodes tainted NoSchedule) and write the number to /opt/KUCC00104/kucc00104.txt.

Answer:

See the solution below.

Explanation

solution


Readme
Web Terminal
THE LINUX FOUNDATION
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
freeqas.com

freeqas.com

NEW QUESTION: 22
Fix a node that shows as non-ready
A. Kubectl get nodes
// Check which node shows a not ready
kubectl describe nodes "node-name"
// Login to the node which shows as not ready and check the
process for kubelet, docker , kube-proxy.
// systemctl status kubelet (or) ps -aux | grep -i "processname"
// If the process is not started, then start using
systemctl start kubelet / docker
// Verify
ps -auxxww | grep -i "process-name"
kubectl get nodes
B. Kubectl get nodes

// Check which node shows a not ready
kubectl describe nodes "node-name"
// Login to the node which shows as not ready and check the
systemctl start kubelet / docker
// Verify
ps -auxxww | grep -i "process-name"
kubectl get nodes
Answer: A (LEAVE A REPLY)

NEW QUESTION: 23
Create a deployment as follows:
* Name:nginx-random
* Exposed via a servicenginx-random
* Ensure that the service & podare accessible via theirrespective DNS records
* The container(s) within anypod(s) running as a part of thisdeployment should use thenginxImage Next, use the utilitynslookupto lookup the DNS records of the service &pod and write the output to /opt/KUNW00601/service.dnsand/opt/KUNW00601/pod.dnsrespectively.
Answer:
See the solution below.
Explanation
Solution:

Readme >_ Web Terminal THE LINUX FOUNDATION

```
root@node-1:~#
root@node-1:~# k create deploy nginx-random --image=nginx
deployment.apps/nginx-random created
root@node-1:~# k expose deploy nginx-random --name=nginx-random --port=80 --target-port=80
service/nginx-random exposed
root@node-1:~# vim dns.yam
```

freeqas.com

NEW QUESTION: 24

Delete the above pod and create again from the same yaml file
and verifies there is no "test-file.txt" in the path /data/redis
(Since non-persistent storage "emptyDir" is used).
Answer:
kubectl delete pod test-redis kubectl create -f test-redis.yaml kubectl exec -it test-redis /bin/sh cat /data/redis/file.txt // file doesn't exist

NEW QUESTION: 25
Create an nginx pod which loads the secret as environment variables
A. // create a yml file
kubectl run nginx --image=nginx --restart=Never --dry-run -o
yaml > nginx.yml
// add env section below and create
vim nginx.yaml
apiVersion: v1
kind: Pod
metadata:
labels:
run: nginx
name: nginx
spec:
containers:
- image: nginx
name: nginx
envFrom:
- secretRef:
name: my-secret
restartPolicy: Never
kubectl apply -f nginx.yaml
//verify
kubectl exec -it nginx - env
B. // create a yml file
kubectl run nginx --image=nginx --restart=Never --dry-run -o
yaml > nginx.yml
// add env section below and create
vim nginx.yaml
run: nginx
name: nginx
spec:
containers:
- image: nginx
name: nginx

envFrom:
- secretRef:
name: my-secret
restartPolicy: Never
kubectl apply -f nginx.yaml
//verify
kubectl exec -it nginx - env
Answer: A (LEAVE A REPLY)

NEW QUESTION: 26
Create a busybox pod which executes this command sleep 3600 with the service account admin and verify
A. kubectl run busybox --image=busybox --restart=Always --dry-run
-o yaml -- /bin/sh -c "sleep 3600" > busybox.yml
// Edit busybox.yaml file
apiVersion: v1
kind: Pod
metadata:
creationTimestamp: null
labels:
run: busybox
name: busybox
spec:
serviceAccountName: admin
containers:
- args:
- /bin/sh
- -c
- sleep 3600
image: busybox
name: busybox
restartPolicy: Always
// verify
K kubectl describe po busybox
B. kubectl run busybox --image=busybox --restart=Always --dry-run
-o yaml -- /bin/sh -c "sleep 3600" > busybox.yml
// Edit busybox.yaml file
apiVersion: v1
kind: Pod
metadata:
creationTimestamp: null
labels:

run: busybox
name: busybox
spec:
serviceAccountName: admin
containers:
- args:
- /bin/sh
- -c
- sleep 3800
image: busybox
name: busybox
restartPolicy: Always
// verify
K kubectl describe po busybox
Answer: A (LEAVE A REPLY)

NEW QUESTION: 27
Monitor the logs of pod foo and:
* Extract log lines corresponding unable-to-access-website
* Write them to/opt/KULM00201/foo

Answer:
See the solution below.
Explanation
solution


Readme
Web Terminal
THE LINUX FOUNDATION
student@node-1:~$
student@node-1:~$ sudo -i
root@node-1:~# alias k=kubectl
root@node-1:~#
FOUNDATION
freeqas.com

freeqas.com

NEW QUESTION: 28
What file type upload is supported as part of the basic WildFire service?
A. ELF
B. PE
C. BAT
D. VBS
Answer: (SHOW ANSWER)

NEW QUESTION: 29
Check the image version in pod without the describe command
Answer:
See the solution below.
Explanation
kubectl get po nginx -o
jsonpath='{.spec.containers[].image}{"\n"}'

NEW QUESTION: 30
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. 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:
- 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=dev
B. 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
Answer: (SHOW ANSWER)

NEW QUESTION: 31
Scale the deployment to 5 replicas
Answer:
kubectl scale deployment webapp -replicas=5 //Verify kubectl get deploy kubectl get po,rs

Valid CKA Dumps shared by PrepPdf.com for Helping Passing CKA Exam! PrepPdf.com now offer the **newest CKA exam dumps**, the PrepPdf.com CKA exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CKA dumps with Test Engine here: https://www.preppdf.com/Linux-Foundation/CKA-prepaway-exam-dumps.html (**85** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 32
Create a deployment as follows:
Name: nginx-random
Exposed via a service nginx-random
Ensure that the service & pod are accessible via their respective DNS records The container(s) within any pod(s) running as a part of this deployment should use the nginx Image Next, use the utility nslookup to look up the DNS records of the service & pod and write the output to /opt/KUNW00601/service.dns and /opt/KUNW00601/pod.dns respectively.
Answer:
See the solution below.
Explanation
Solution:
F:\Work\Data Entry Work\Data Entry\20200827\CKA\17 C.JPG

freeqas.com

F:\Work\Data Entry Work\Data Entry\20200827\CKA\17 D.JPG

freeqas.com

F:\Work\Data Entry Work\Data Entry\20200827\CKA\17 E.JPG

NEW QUESTION: 33
Monitor the logs of pod foo and:
Extract log lines corresponding to error
unable-to-access-website
Write them to /opt/KULM00201/foo

Answer:
solution

NEW QUESTION: 34

To protect your firewall and network from single source denial of service (DoS) attacks that can overwhelm its packet buffer and cause legitimate traffic to drop, you can configure:
A. BGP (Border Gateway Protocol)
B. PGP (Packet Gateway Protocol)
C. PBP (Protocol Based Protection)
D. PBP (Packet Buffer Protection)
Answer: (SHOW ANSWER)

NEW QUESTION: 35
A bootstrap USB flash drive has been prepared using a Linux workstation to load the initial configuration of a Palo Alto Networks firewall. The USB flash drive was formatted using file system ntfs and the initial configuration is stored in a file named init-cfg.txt.
The contents of Init-cfg.txt in the USB flash drive are as follows:
type=static
ip-address=10.5.107.19
default-gateway=10.5.107.1
netmask=255.255.255.0
Ipv6-address=2001:400:100::1/64
ipv6-default-gateway=2001:400:100::2
hostname=Ca-FW-DC1
panorama-server=10.5.107.20
panorama-server-2=10.5.107.21
tplname=FINANCE TG4
dgname=finance_dg
dns-primary=10.5.6.6
op-command-modes multi-vsys.jumbo-frame
dhcp-send-hostname=no
dhcp-send-client-id=no
dhcp-accept-server-hostname=no
dhcp-accept-server-domain=no
The USB flash drive has been inserted in the firewalls' USB port, and the firewall has been powered on.Upon boot, the firewall fails to begin the bootstrapping process. The failure is caused because:
A. nit-cfg bit is an incorrect filename the correct filename should be init-ofg.xml
B. The USB drive has been formatted with an unsupported file system
C. The USB must be formatted using the ext4 file system
D. The bootstrap.xml file is a required file, but it is missing
E. There must be commas between the parameter names and their values instead of the equal symbols
Answer: C (LEAVE A REPLY)

NEW QUESTION: 36
Create an nginx pod with containerPort 80 and with a PersistentVolumeClaim "task-pv-claim" and has a mouth path "/usr/share/nginx/html"
A. vim nginx-pvc-pod.yaml

apiVersion: v1
kind: Pod
metadata:
name: task-pv-pod
spec:
volumes:
- name: task-pv-storage
persistentVolumeClaim:
claimName: task-pv-claim
containers:
- name: task-pv-container
image: nginx
ports:
- containerPort: 80
name: "http"
volumeMounts:
- mountPath: "/usr/share/nginx/html"
name: task-pv-storage
kubectl apply -f nginx-pvc-pod.yaml
// Verify
kubectl describe po "POD-Name" | grep -i volumes -A5
Volumes:
task-pv-storage:
Type: PersistentVolumeClaim (a reference to a
PersistentVolumeClaim in the same namespace)
ClaimName: task-pv-claim
ReadOnly: false
B. vim nginx-pvc-pod.yaml
apiVersion: v1
kind: Pod
metadata:
name: task-pv-pod
spec:
volumes:
- name: task-pv-storage
persistentVolumeClaim:
claimName: task-pv-claim
containers:
- name: task-pv-container
image: nginx
ports:

- containerPort: 60
name: "http"
volumeMounts:
- mountPath: "/usr/share/nginx/html"
name: task-pv-storage
kubectl apply -f nginx-pvc-pod.yaml
// Verify
kubectl describe po "POD-Name" | grep -i volumes -A4
Volumes:
task-pv-storage:
Type: PersistentVolumeClaim (a reference to a
PersistentVolumeClaim in the same namespace)
ClaimName: task-pv-claim
ReadOnly: false
Answer: (SHOW ANSWER)

NEW QUESTION: 37
Create an nginx pod and list the pod with different levels of verbosity
Answer:
See the solution below.
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: 38
A Kubernetes worker node, named .Investigate why this is the case,
andperform any appropriate steps tobring the node to a state,ensuring that any changes are madepermanent.
You cansshto the failednode using:
[student@node-1] $ | sshWk8s-node-0
You can assume elevatedprivileges on the node with thefollowing command:
[student@w8ks-node-0] $ |sudo -i
Answer:
See the solution below.
Explanation
solution

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
```

freeqas.com

THE LINUX FOUNDATION

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
```

freeqas.com

THE LINUX FOUNDATION

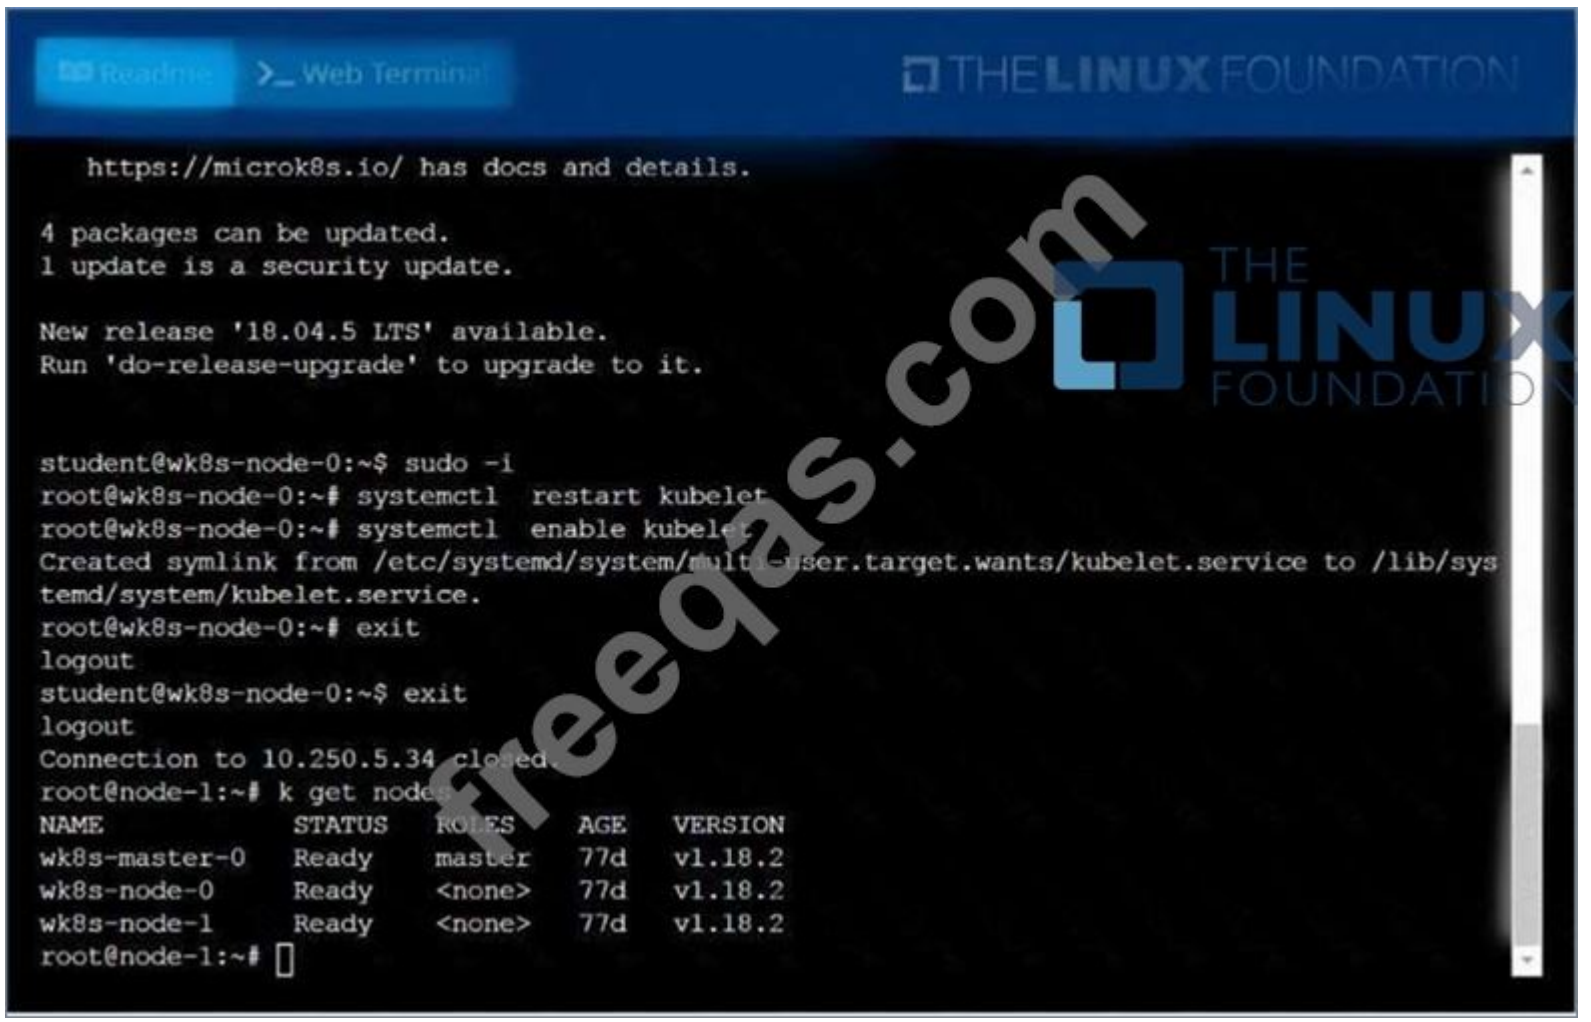

freeqas.com

NEW QUESTION: 39
Create a pod with image nginx called nginx and allow traffic on port 80
Answer:
kubectl run nginx --image=nginx --restart=Never --port=80

NEW QUESTION: 40
Create the nginx pod with version 1.17.4 and expose it on port 80
Answer:
kubectl run nginx --image=nginx:1.17.4 --restart=Never -- port=80

NEW QUESTION: 41
A bootstrap USB flash drive has been prepared using a Windows workstation to load the initial configuration of a Palo Alto Networks firewall that was previously being used in a lab. The USB flash drive was formatted using file system FAT32 and the initial configuration is stored in a file named init-cfg.txt. The firewall is currently running PAN-OS 10.0 and using a lab config. The contents of init-cgf.txt in the USB flash drive are as follows:
type=dhcp-client
Ip-address=
default-gateway=
netmask=

Ipv6-address=
Ipv6-default-gateway=
hostname=Ca-FW-DC1
panorama-server=10.5.107.20
panorama-server-2=10.5.107.21
tplname=FINANCE_TG4
dgname=finance_dg
dns-primary=10.5.6.6
dns-secondary=10.5.6.7
op-command-modes-multi-vsys.jumbo-frame
dhcp-send-hostname=yes
dhcp-send-client-id=yes
dhcp-accept-server-hostname=yes
dhcp-accept-server-domain=yes
The USB flash drive has been inserted in the firewalls' USB port, and the firewall has been restarted using command> request restart system
Upon restart, the firewall fails to begin the bootstrapping process. The failure is caused because:
A. The USB must be formatted using the exi3 file system, FAT32 is
B. The hostname is a required parameter, but it is missing in init-cfg.txt
C. PAN-OS version must be 9.1 x at a minimum, but the firewall is running 10.0x
D. Firewall must be in factory default state or have all private data deleted for bootstrapping
E. The bootstrap xml file is a required file, but it is missing
Answer: (SHOW ANSWER)

NEW QUESTION: 42
Check the history of deployment
Answer:
kubectl rollout history deployment webapp

NEW QUESTION: 43
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

Answer:

See the solution below.
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: 44

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

Answer:

solution

freeqas.com

NEW QUESTION: 45
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.
Answer:
solution

NEW QUESTION: 46

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).

Answer:

See the solution below.

Explanation

Solution:
kubectl top -l name=cpu-user -A
echo 'pod name' >> /opt/KUT00401/KUT00401.txt

Valid CKA Dumps shared by PrepPdf.com for Helping Passing CKA Exam! PrepPdf.com now offer the **newest CKA exam dumps**, the PrepPdf.com CKA exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CKA dumps with Test Engine here: https://www.preppdf.com/Linux-Foundation/CKA-prepaway-exam-dumps.html (**85** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 47
Create a file:
/opt/KUCC00302/kucc00302.txt that lists all pods that implement service in namespace development.
The format of the file should be one pod name per line.
Answer:
See the solution below.
Explanation
solution

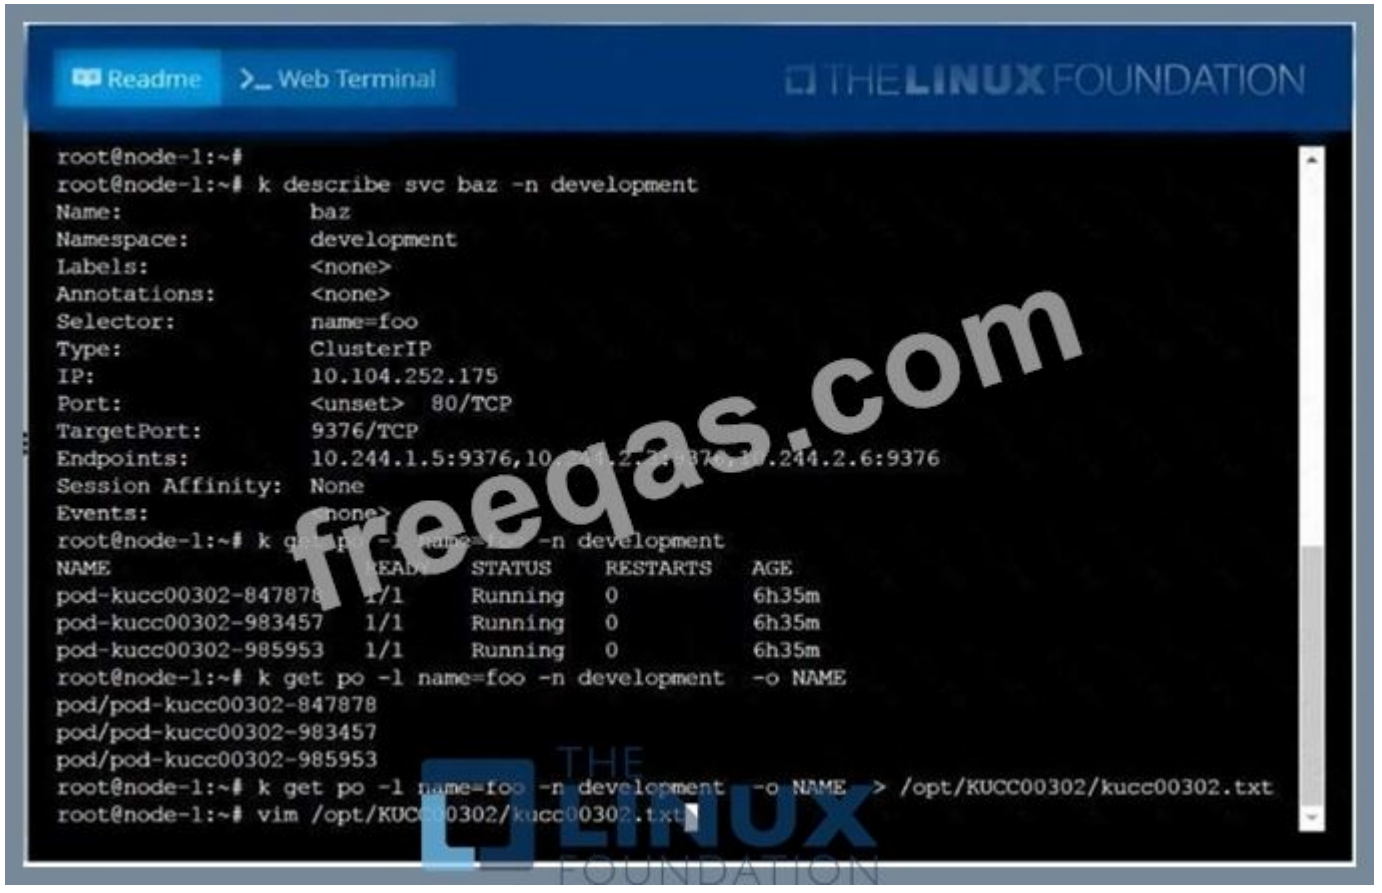


Readme
>_ Web Terminal
THE LINUX FOUNDATION
pod-kucc00302-847878
pod-kucc00302-983457
pod-kucc00302-985953
:wq
freeqas.com

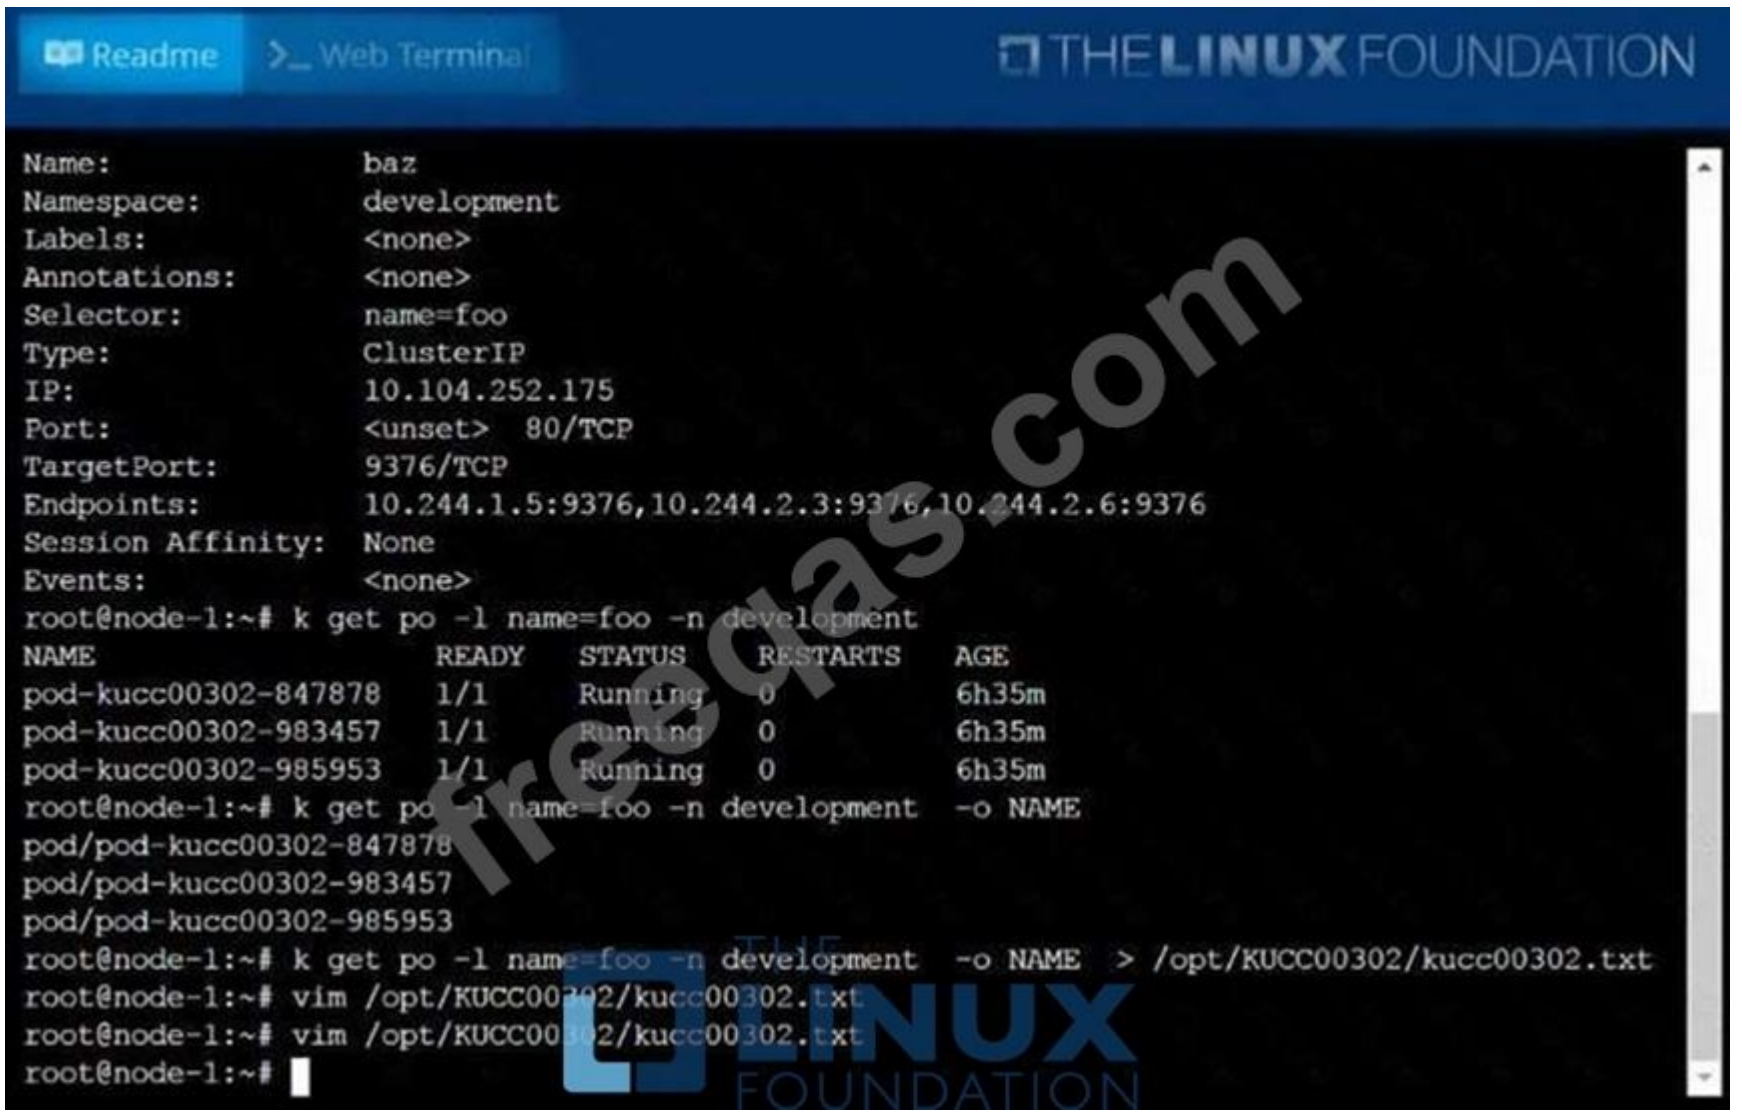

NEW QUESTION: 48
List all the pods sorted by name
Answer:
kubectl get pods --sort-by=.metadata.name

NEW QUESTION: 49
Create a nginx pod with label env=test in engineering namespace
Answer:
See the solution below.
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 - 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: 50
List all the pods showing name and namespace with a json path expression
Answer:
kubectl get pods -o=jsonpath="{.items[*]['metadata.name', 'metadata.namespace']}"

NEW QUESTION: 51
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
Answer:
solution

Readme | >_ Web Terminal

THE LINUX FOUNDATION

```
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
```

freeqas.com


Readme
Web Terminal
THE LINUX FOUNDATION
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
staticPodPath: /etc/kubernetes/manifests
syncFrequency: 0s
volumeStatsAggPeriod: 0s
:wq
freeqas.com

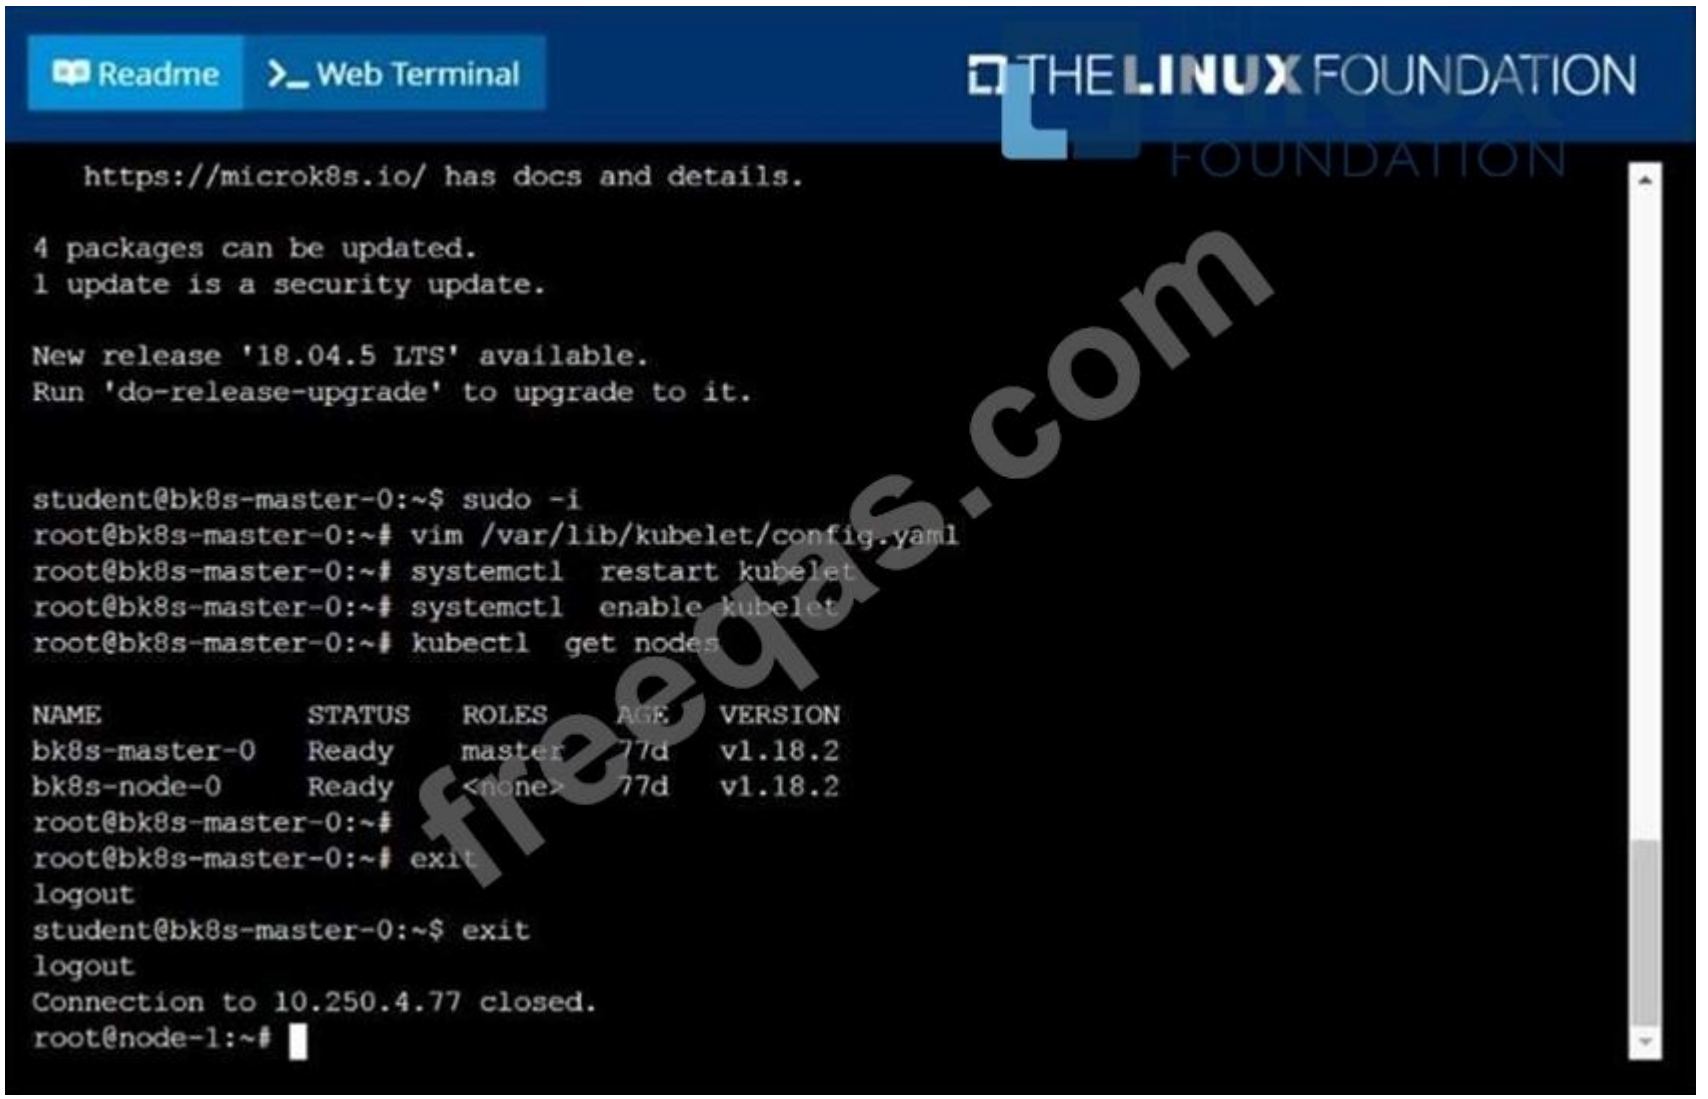

freeqas.com

NEW QUESTION: 52

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.

Answer:

See the solution below.

Explanation

solution

freeqas.com

NEW QUESTION: 53
Deployment
a. Create a deployment of webapp with image nginx:1.17.1 with
container port 80 and verify the image version
A. // Create initial YAML file with -dry-run option
kubectl create deploy webapp --image=nginx:1.17.1 --dryrun=client -o yaml > webapp.yaml vim webapp.yaml apiVersion: apps/v1 kind: Deployment metadata:
labels:
app: webapp
name: webapp
spec: replicas: 1 containers: - image: nginx:1.17.1 name: nginx kubectl create -f webapp.yaml -record=true //Verify Image Version kubectl describe deploy webapp | grep -i "Image" Using JsonPath kubectl get deploy -o=jsonpath='{range.items [*]}{.[*]} {.metadata.name}{"\t"} {.spec.template.spec.containers[*].i mage}{"\n"}'
B. // Create initial YAML file with -dry-run option

kubectl create deploy webapp --image=nginx:1.17.1 --dryrun=client -o yaml > webapp.yaml vim webapp.yaml apiVersion: apps/v1 kind: Deployment metadata:
labels:
app: webapp
name: webapp
spec: replicas: 1 selector: matchLabels: app: webapp template: metadata: labels: app: webapp spec: containers: - image: nginx:1.17.1 name: nginx kubectl create -f webapp.yaml -record=true //Verify Image Version kubectl describe deploy webapp | grep -i "Image" Using JsonPath kubectl get deploy -o=jsonpath='{range.items [*]}{.[*]} {.metadata.name}{"\t"}{.spec.template.spec.containers[*].i mage}{"\n"}'
Answer: (SHOW ANSWER)

NEW QUESTION: 54
Perform the following tasks:
Add an init container to hungry-bear defined in spec file
/opt/KUCC00108/pod-spec-KUC
The init container should create /workdir/calm.txt
If /workdir/calm.txt is not
Once the spec file has been definition, the pod should be created
Answer:
See the solution below.
Explanation
solution
F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 B.JPG

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
```

freeqas.com

F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 C.JPG

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
```

freeqas.com

F:\Work\Data Entry Work\Data Entry\20200827\CKA\4 D.JPG

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:~#
```

freeqas.com

NEW QUESTION: 55

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

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
```

freeqas.com


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

freeqas.com

NEW QUESTION: 56
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

Readme >_ Web Terminal THE LINUX FOUNDATION

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

freeqas.com

Readme >_ Web Terminal THE LINUX FOUNDATION

```
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
```

freeqas.com

Readme | Web Terminal

THE LINUX FOUNDATION

```
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
~
~
~
~
~
~
~
:wq
```

freeqas.com

freeqas.com

NEW QUESTION: 57

Create a deployment as follows:

* Name: nginx-app
* Using container nginx with version 1.11.10-alpine
* The deployment should contain 3 replicas

Next, deploy the application with new version 1.11.13-alpine, by performing a rolling update.

Finally, rollback that update to the previous version 1.11.10-alpine.

Answer:

See the solution below.

Explanation

solution

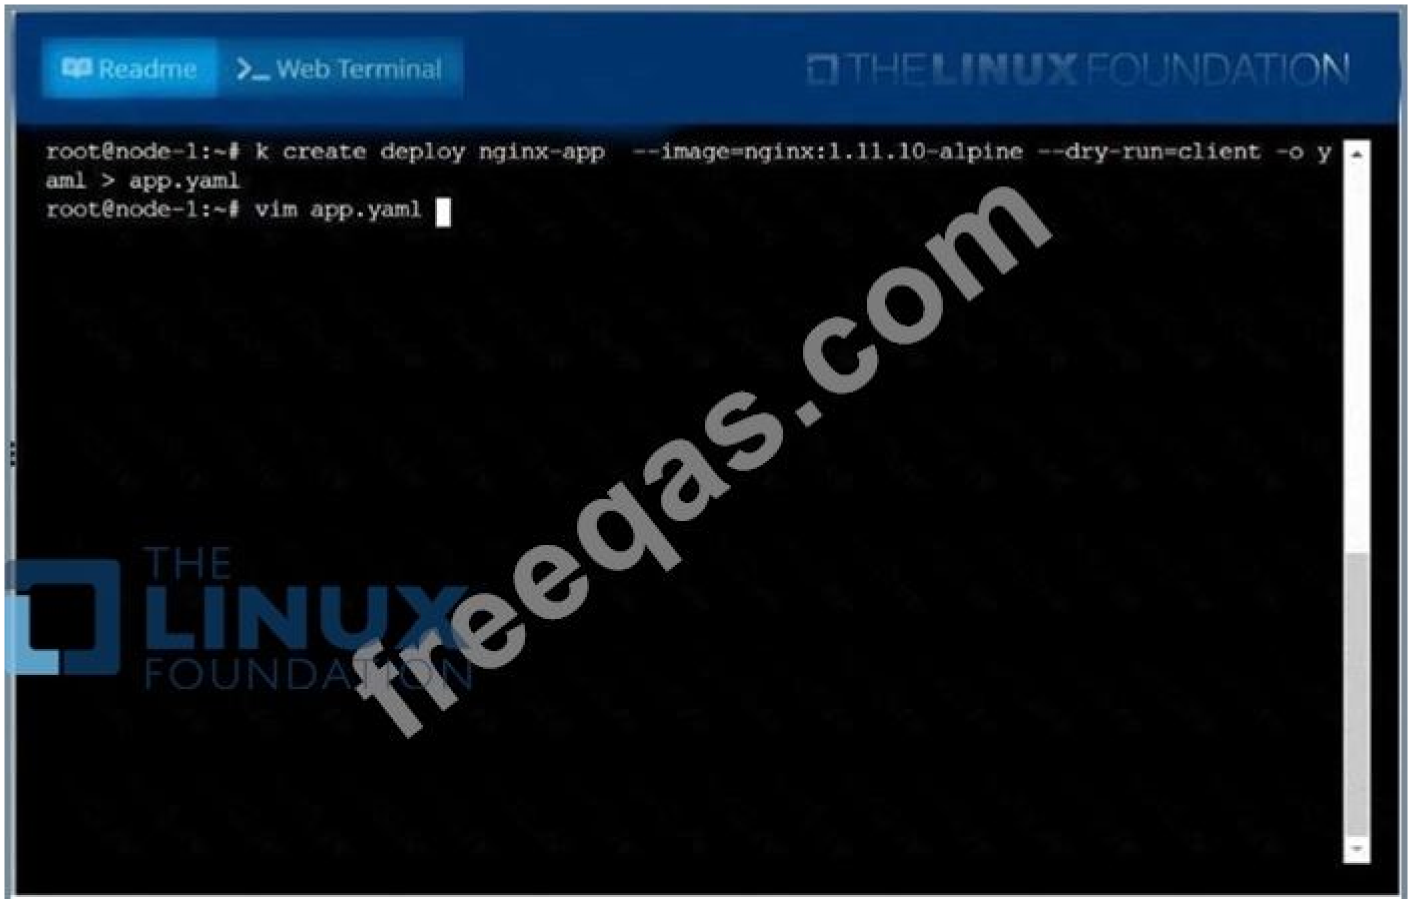
Readme
Web Terminal
THE LINUX FOUNDATION
root@node-1:~# k create deploy nginx-app --image=nginx:1.11.10-alpine --dry-run=client -o y
aml > app.yaml
root@node-1:~# vim app.yaml


Readme
Web Terminal
THE LINUX FOUNDATION
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

freeqas.com

Readme >_ Web Terminal THE LINUX FOUNDATION

```
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:~#
```

freeqas.com

NEW QUESTION: 58

Ensure a single instance of podnginxis running on each node of theKubernetes cluster wherenginxalso represents the Image name whichhas to be used. Do not override anytaints currently in place.

UseDaemonSetto complete thistask and useds-kusc00201asDaemonSet name.

Answer:

See the solution below.

Explanation

solution


Readme
>_ Web Terminal
THE LINUX FOUNDATION
root@node-1:~# vim ds.yaml
i
LINUX
FOUNDATION
freeqas.com

Readme | >_ Web Terminal

THE LINUX FOUNDATION

```
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
```

freeqas.com

NEW QUESTION: 59
List all service account and create a service account called "admin"

A. kubectl get sa
kubectl get sa --all-namespaces
kubectl create sa admin
//Verify
kubectl get sa admin -o yaml
B. kubectl get sa
kubectl get sa --all-namespaces
//Verify
kubectl get sa admin -o yaml
Answer: A (LEAVE A REPLY)

NEW QUESTION: 60
Create a persistent volume with nameapp-data, of capacity2Giandaccess modeReadWriteMany. Thetype of volume ishostPathand itslocation is/srv/app-data.
Answer:
See the solution below.
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 named ReadWriteMany, storage classname
shared,2Giof storage capacity and the host path

```
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:

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: "

NEW QUESTION: 61
Create a configmap called cfgvolume with values var1=val1,
var2=val2 and create an nginx pod with volume nginx-volume which
reads data from this configmap cfgvolume and put it on the path
/etc/cfg
A. // first create a configmap cfgvolume
kubectl create cm cfgvolume --from-literal=var1=val1 --fromliteral=var2=val2
// verify the configmap
kubectl describe cm cfgvolume
// create the config map
kubectl create -f nginx-volume.yml
vim nginx-configmap-pod.yaml
apiVersion: v1
kind: Pod
- name: nginx-volume
configMap:
name: cfgvolume
containers:
- image: nginx
name: nginx
volumeMounts:
- name: nginx-volume

mountPath: /etc/cfg
restartPolicy: Always
k kubectl apply -f nginx-configmap-pod.yaml
/ // Verify
// exec into the pod
kubectl exec -it nginx -- /bin/sh
// check the path
cd /etc/cfg
B. // first create a configmap cfgvolume
kubectl create cm cfgvolume --from-literal=var1=val1 --fromliteral=var2=val2
// verify the configmap
kubectl describe cm cfgvolume
// create the config map
kubectl create -f nginx-volume.yml
vim nginx-configmap-pod.yaml
apiVersion: v1
kind: Pod
metadata:
labels:
run: nginx
name: nginx
spec:
volumes:
- name: nginx-volume
configMap:
name: cfgvolume
containers:
- image: nginx
name: nginx
volumeMounts:
- name: nginx-volume
mountPath: /etc/cfg
restartPolicy: Always
k kubectl apply -f nginx-configmap-pod.yaml
/ // Verify
// exec into the pod
kubectl exec -it nginx -- /bin/sh
// check the path
cd /etc/cfg
Answer: B (LEAVE A REPLY)

Valid CKA Dumps shared by PrepPdf.com for Helping Passing CKA Exam! PrepPdf.com now offer the **newest CKA exam dumps**, the PrepPdf.com CKA exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CKA dumps with Test Engine here: https://www.preppdf.com/Linux-Foundation/CKA-prepaway-exam-dumps.html (**85** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)

NEW QUESTION: 62
Score: 4%

Task
Scale the deployment presentation
Answer:
See the solution below.
Explanation
Solution:
kubectl get deployment
kubectl scale deployment.apps/presentation --replicas=6

NEW QUESTION: 63
Watch the job that runs 10 times one by one and verify 10 pods are created and delete those after it's completed
Answer:
kubectl get job -w kubectl get po kubectl delete job hello-job

NEW QUESTION: 64
Create a namespace called 'development' and a pod with image nginx called nginx on this namespace.
Answer:
kubectl create namespace development
kubectl run nginx --image=nginx --restart=Never -n development

NEW QUESTION: 65
Get list of all pods in all namespaces and write it to file "/opt/pods-list.yaml"
Answer:

See the solution below.
Explanation
kubectl get po -all-namespaces > /opt/pods-list.yaml

NEW QUESTION: 66
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
Answer:
solution

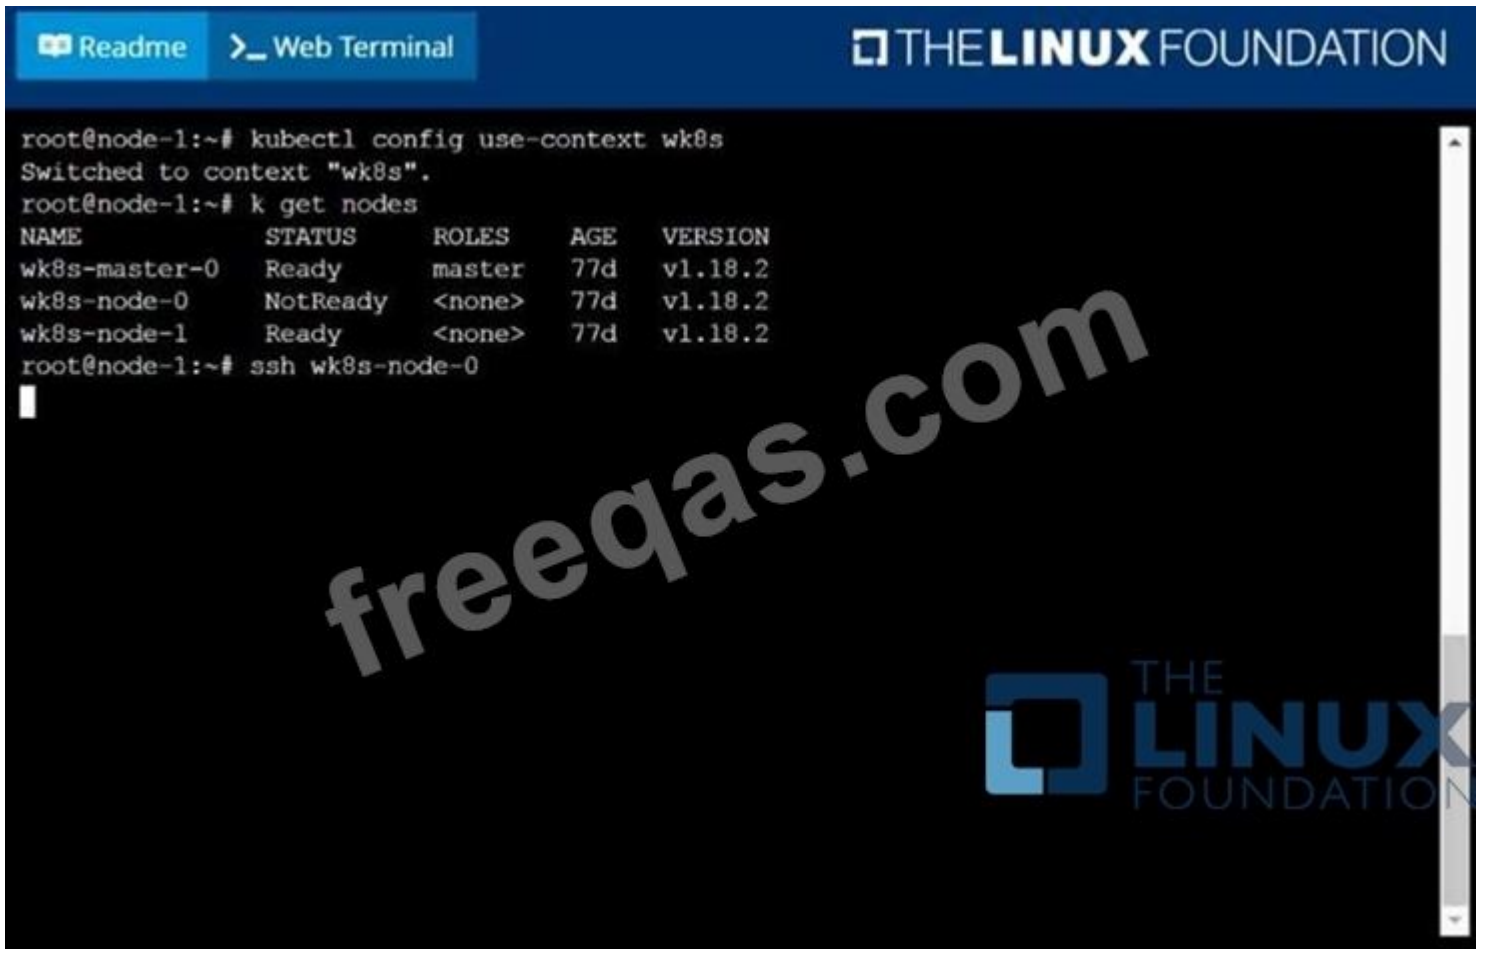

Readme | >_ Web Terminal | THE LINUX FOUNDATION

```
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
```

freeqas.com

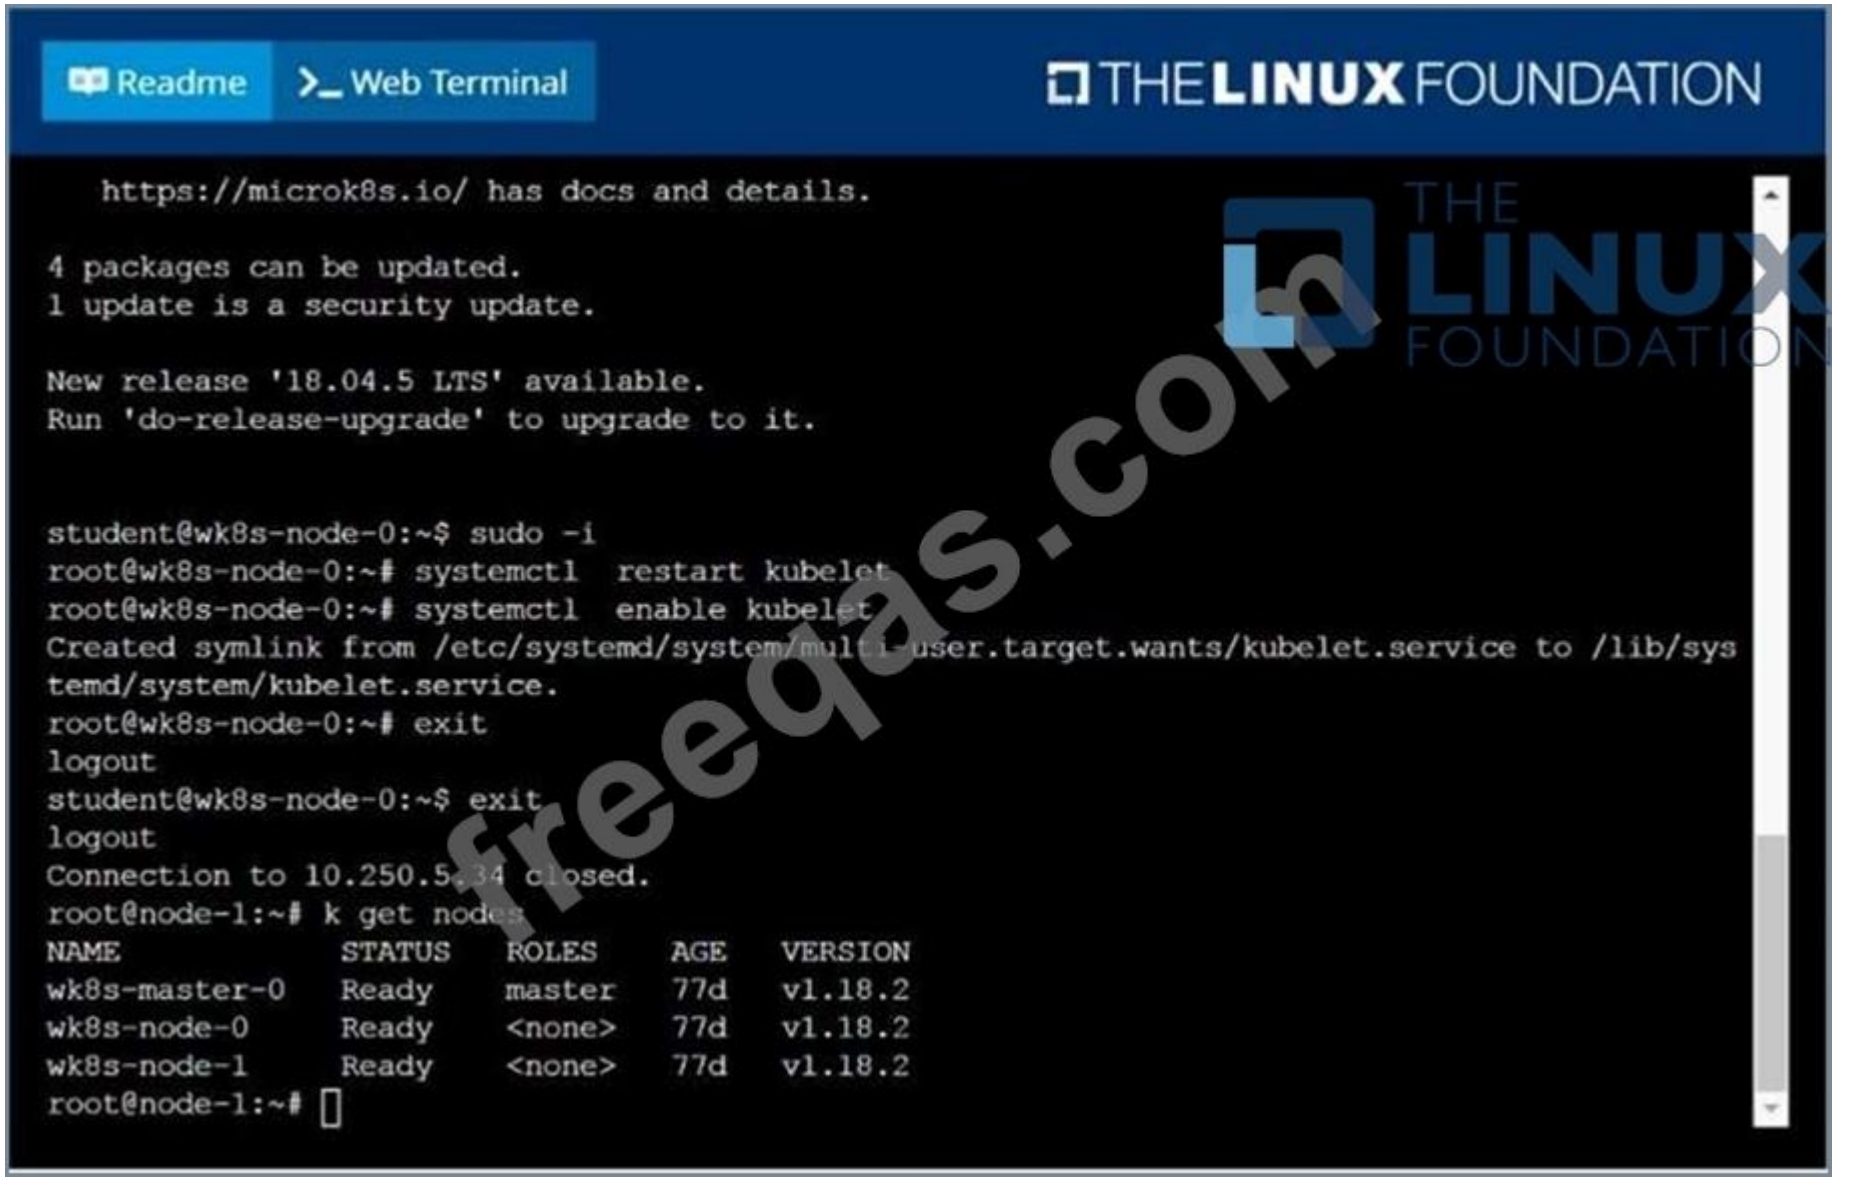

freeqas.com

NEW QUESTION: 67
Create a pod with environment variables as var1=value1.Check the environment variable in pod
A. kubectl run nginx --image=nginx --restart=Never --env=var1=value1
then
kubectl exec -it nginx -- env
or
kubectl exec -it nginx -- sh -c 'echo $var1'
or
kubectl describe po nginx | grep value1
B. kubectl run nginx --image=nginx --restart=Never --env=var1=value1
then
kubectl exec -it nginx -- env
or
kubectl describe po nginx | grep value1
Answer: (SHOW ANSWER)

NEW QUESTION: 68
Modify "hello-job" and make it run 10 times one after one and 5 times parallelism: 5
A. kubectl create job hello-job --image=busybox --dry-run -o yaml
-- echo "Hello I am from job" > hello-job.yaml
// edit the yaml file to add completions: 16 and
kubectl create -f hello-job.yaml
YAML File:
apiVersion: batch/v1
kind: Job
metadata:
name: hello-job
spec:
completions: 16
parallelism: 5
template:
metadata:
spec:
containers:
- command:
- echo
- Hello I am from job
image: busybox
name: hello-job
restartPolicy: Never
B. kubectl create job hello-job --image=busybox --dry-run -o yaml
-- echo "Hello I am from job" > hello-job.yaml
// edit the yaml file to add completions: 10 and
kubectl create -f hello-job.yaml
YAML File:
apiVersion: batch/v1
kind: Job
metadata:
name: hello-job
spec:
completions: 10
parallelism: 5
template:
metadata:
spec:
containers:

- command:
- echo
- Hello I am from job
image: busybox
name: hello-job
restartPolicy: Never
Answer: B (LEAVE A REPLY)

Valid CKA Dumps shared by PrepPdf.com for Helping Passing CKA Exam! PrepPdf.com now offer the **newest CKA exam dumps**, the PrepPdf.com CKA exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com CKA dumps with Test Engine here: https://www.preppdf.com/Linux-Foundation/CKA-prepaway-exam-dumps.html (**85** Q&As Dumps, **40%OFF** **Special Discount:** **Exam-Tests**)