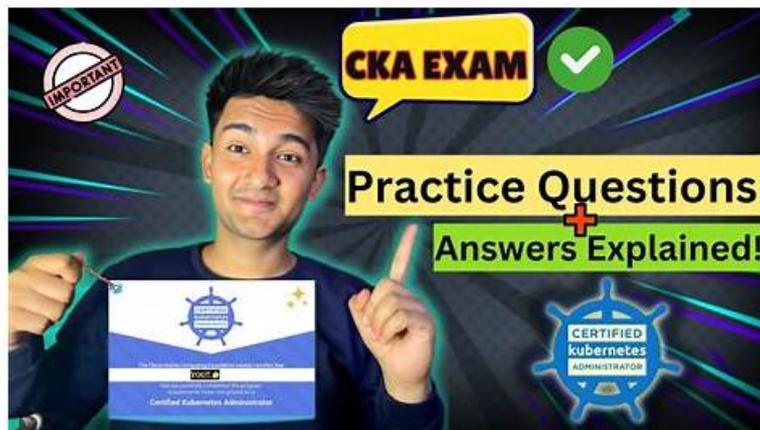


CKA Latest Exam Pattern - CKA Real Torrent



BTW, DOWNLOAD part of Free4Torrent CKA dumps from Cloud Storage: https://drive.google.com/open?id=1OA2fyHjC_exrkheXmylwUBEPBjR0vFUo

According to the needs of all people, the experts and professors in our company designed three different versions of the CKA certification training materials for all customers. The three versions are very flexible for all customers to operate. According to your actual need, you can choose the version for yourself which is most suitable for you to preparing for the coming exam. All the CKA Training Materials of our company can be found in the three versions. It is very flexible for you to use the three versions of the CKA latest questions to preparing for your coming exam.

Linux Foundation CKA (Certified Kubernetes Administrator) Program is a certification exam designed to test the skills of individuals who work with Kubernetes. Kubernetes is an open-source platform that is used to automate the deployment, scaling, and management of containerized applications. The CKA Exam is designed to validate an individual's understanding of Kubernetes and their ability to deploy, configure, and troubleshoot Kubernetes clusters.

>> CKA Latest Exam Pattern <<

CKA Real Torrent | CKA Pass4sure Pass Guide

Free4Torrent's CKA exam training materials are proved to be effective by some professionals and examinees that have passed CKA exam, Free4Torrent's CKA exam dumps are almost the same with real exam paper. It can help you pass CKA certification exam. After you purchase our CKA VCE Dumps, if you fail CKA certification exam or there are any problems of CKA test training materials, we will give a full refund to you. We believe that our Free4Torrent's CKA vce dumps will help you.

Linux Foundation Certified Kubernetes Administrator (CKA) Program Exam Sample Questions (Q25-Q30):

NEW QUESTION # 25

Score: 5%



Task

Monitor the logs of pod bar and:

- * Extract log lines corresponding to error file-not-found
- * Write them to /opt/KUTR00101/bar

Answer:

Explanation:

Solution:

```
kubectl logs bar | grep 'unable-to-access-website' > /opt/KUTR00101/bar cat /opt/KUTR00101/bar
```

NEW QUESTION # 26

Create a nginx pod with label env=test in engineering namespace

Answer:

Explanation:

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 # 27

You have a Kubernetes cluster with three nodes. Node1 and Node2 are in the 'default' availability zone, while Node3 is in the 'us-east-1a' availability zone. You want to ensure that pods are spread across all three nodes, considering the availability zones.

Describe how to configure the cluster to achieve this goal, specifically addressing how to leverage 'nodeSelector' and/or 'affinity' to enforce desired node placement. Explain the rationale behind your chosen approach.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

Step 1: Configure Node Labels

Label each node with its corresponding availability zone:

For Node1 and Node2 (in 'default' availability zone):

```
kubectl label node availability-zone=default
```

For Node3 (in availability zone):

```
kubectl label node availability-zone=us-east-1a
```

Step 2: Use Node Selector

Use 'nodeSelector' in your Deployment or Pod definition to specify the desired availability zone:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:1.14.2
        nodeSelector:
          availability-zone: us-east-1a
```

This ensures pods with the 'nginx-deployment' label will be scheduled only on Node3. Step 3: Use Affinity (Optional) You can also use 'affinity' for more fine-grained control. For example, to ensure that pods are spread across different availability zones:

```

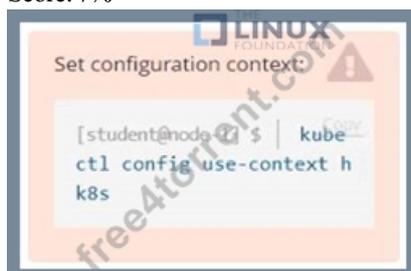
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
        - name: nginx
          image: nginx:1.14.2
      affinity:
        podAntiAffinity:
          preferredDuringSchedulingIgnoredDuringExecution:
            - weight: 100
              podAffinityTerm:
                labelSelector:
                  matchExpressions:
                    - key: availability-zone
                      operator: In
                      values:
                        - default
                        - us-east-1a
                topologyKey: topology.kubernetes.io/zone

```

This configuration will prefer scheduling pods in different availability zones. Rationale: Node Selector: Provides a simple mechanism to direct pods to specific nodes based on labels. Affinity: Offers more advanced options, including 'podAntiAffinity' to spread pods across nodes (or availability zones) and 'podAffinity' to ensure pods are scheduled on the same node. Availability Zone: Distributes pods across different zones for high availability, as failures in one zone won't impact pods scheduled in other zones. ,

NEW QUESTION # 28

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:

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 # 29

You have a 'NodePort' service named 'my-service' exposed on port 30000. You need to configure a NetworkPolicy to only allow access to the service from specific nodes in the cluster based on their node labels.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a NetworkPolicy:

- Define a NetworkPolicy that allows traffic from nodes with specific labels to the service.

```
apiVersion: networking.k8s.io/v1
```

```
kind: NetworkPolicy
```

```
metadata:
```

```
  name: allow-specific-nodes
```

```
  namespace:
```

```
spec:
```

```
  podSelector:
```

```
    matchLabels:
```

```
      app: my-service # Match the service label
```

```
  ingress:
```

```
  - from:
```

```
    - ipBlock:
```

```
      cidr: 10.244.0.0/16 # Example: CIDR for the nodes you want to allow access
```

```
      except:
```

```
        - ipBlock:
```

```
          cidr: 10.244.0.1/32 # Exclude specific node IPs if necessary
```

```
    - ipBlock:
```

```
      cidr: 172.17.0.0/16 # Another CIDR for nodes
```

