

Use Linux Foundation CKAD Practice Exam Software (Desktop and Web-Based) For Self Evaluation



BTW, DOWNLOAD part of Lead1Pass CKAD dumps from Cloud Storage: <https://drive.google.com/open?id=1Ju1bszpk1tIJVxXkxE9yY67AzPdpX6WS>

Since our company's establishment, we have devoted mass manpower, materials and financial resources into CKAD exam materials and until now, we have a bold idea that we will definitely introduce our study materials to the whole world and make all people that seek fortune and better opportunities have access to realize their life value. Our CKAD Practice Questions, therefore, is bound to help you pass though the exam and win a better future. We will also continuously keep a pioneering spirit and are willing to tackle any project that comes your way.

Linux Foundation Certified Kubernetes Application Developer (CKAD) certification exam is a performance-based exam designed to test the knowledge and skills of developers who want to build and deploy cloud-native applications using Kubernetes. CKAD exam is intended for developers who have a solid understanding of Kubernetes fundamentals and have hands-on experience developing, deploying, and testing Kubernetes applications.

Linux Foundation Certified Kubernetes Application Developer (CKAD) Certification Exam is a rigorous test designed to evaluate the expertise of developers working with Kubernetes. Kubernetes is an open-source container orchestration system that automates the deployment, scaling, and management of containerized applications. It is widely used by organizations to manage their cloud-native applications, and the demand for certified Kubernetes developers is increasing day by day. The CKAD Certification Exam validates the skills and knowledge of developers in designing, building, configuring, and deploying cloud-native applications on Kubernetes.

>> CKAD Valid Test Registration <<

100% Pass Quiz Linux Foundation - Authoritative CKAD Valid Test Registration

As you all know that the Linux Foundation Certified Kubernetes Application Developer Exam (CKAD) exam is the most challenging exam, since it's difficult to find preparation material for passing the Linux Foundation CKAD exam. Lead1Pass provides you with the most complete and comprehensive preparation material for the Linux Foundation CKAD Exam that will thoroughly prepare you to attempt the CKAD exam and pass it with 100% success guaranteed.

Linux Foundation Certified Kubernetes Application Developer Exam Sample Questions (Q153-Q158):

NEW QUESTION # 153

Context

Set configuration context:

```
[student@node-1 ~]$ kubectl  
config use-context k8s
```

Given a container that writes a log file in format A and a container that converts log files from format A to format B, create a deployment that runs both containers such that the log files from the first container are converted by the second container, emitting logs in format B.

Task:

- * Create a deployment named deployment-xyz in the default namespace, that:
- * Includes a primary lfcncf/busybox:1 container, named logger-dev
- * includes a sidecar lfcncf/fluentd:v0.12 container, named adapter-zen
- * Mounts a shared volume /tmp/log on both containers, which does not persist when the pod is deleted
- * Instructs the logger-dev container to run the command

```
while true; do  
  echo "i luv cncf" >  
  tmp/log/input.log;  
  sleep 10;  
done
```

which should output logs to /tmp/log/input.log in plain text format, with example values:

```
i luv cncf  
i luv cncf  
i luv cncf
```

* The adapter-zen sidecar container should read /tmp/log/input.log and output the data to /tmp/log/output.* in Fluentd JSON format. Note that no knowledge of Fluentd is required to complete this task: all you will need to achieve this is to create the ConfigMap from the spec file provided at /opt/KDMC00102/fluentd-configmap.yaml, and mount that ConfigMap to /fluentd/etc in the adapter-zen sidecar container

Answer:

Explanation:

Solution:

Readme

Web Terminal



THE LINUX FOUNDATION

```
student@node-1:~$ kubectl create deployment deployment-xyz --image=lfcncf/busybox:1 --dry-run=client -o yaml > deployment_xyz.yaml  
student@node-1:~$ vim deployment_xyz.yaml
```

ReadmeWeb Terminal

THE LINUX FOUNDATION

```
apiVersion: apps/v1
kind: Deployment
metadata:
  creationTimestamp: null
  labels:
    app: deployment-xyz
  name: deployment-xyz
spec:
  replicas: 1
  selector:
    matchLabels:
      app: deployment-xyz
  strategy: {}
  template:
    metadata:
      creationTimestamp: null
      labels:
        app: deployment-xyz
    spec:
      containers:
      - image: lfccncf/busybox:1
        name: busybox
        resources: {}
status: {}
~
~
"deployment xyz.yml" 24L, 434C3,1All
```

ReadmeWeb Terminal

THE LINUX FOUNDATION

```
kind: Deployment
metadata:
  labels:
    app: deployment-xyz
  name: deployment-xyz
spec:
  replicas: 1
  selector:
    matchLabels:
      app: deployment-xyz
  template:
    metadata:
      labels:
        app: deployment-xyz
    spec:
      volumes:
      - name: myvol1
        emptyDir: {}
      containers:
      - image: lfccncf/busybox:1
        name: logger-dev
        volumeMounts:
        - name: myvol1
          mountPath: /tmp/log
      - image: lfccncf/fluentd:v0.12
        name: adapter-zen
3 lines yanked27,22Bot
```

ReadmeWeb Terminal

THE LINUX FOUNDATION

```
replicas: 1
selector:
  matchLabels:
    app: deployment-xyz
template:
  metadata:
    labels:
      app: deployment-xyz
  spec:
    volumes:
    - name: myvol1
      emptyDir: {}
    containers:
    - image: lfccncf/busybox:1
      name: logger-dev
      command: ["/bin/sh","-c","tail -f /dev/null"]
      volumeMounts:
      - name: myvol1
        mountPath: /tmp/log
    - image: lfccncf/fluentd:v0.12
      name: adapter-zen
      command: ["/bin/sh","-c","tail -f /tmp/log/input.log >> /tmp/log/output.log"]
      volumeMounts:
      - name: myvol1
        mountPath: /tmp/log
29,83Bot
```

Readme
Web Terminal

```

metadata:
  labels:
    app: deployment-xyz
spec:
  volumes:
    - name: myvol1
      emptyDir: {}
    - name: myvol2
      configMap:
        name: logconf
  containers:
    - image: lfccncf/busybox:1
      name: logger-dev
      command: ["sh", "-c", "while true; do echo 'i luv cncf' >> /tmp/log/input.log; sleep 10; done"]
      volumeMounts:
        - name: myvol1
          mountPath: /tmp/log
    - image: lfccncf/fluentd:v0.12
      name: adapter-zen
      command: ["/bin/sh", "-c", "tail -f /tmp/log/input.log >> /tmp/log/output.log"]
      volumeMounts:
        - name: myvol1
          mountPath: /tmp/log
        - name: myvol2
          mountPath: /fluentd/etc
  
```

37,33
Bot

```

student@node-1:~$ kubectl create -f deployment-xyz.yml
deployment.apps/deployment-xyz created
student@node-1:~$ kubectl get deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 0/1     1            0           5s
student@node-1:~$ kubectl get deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 0/1     1            0           9s
student@node-1:~$ kubectl rollout deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 1/1     1            1           12s
student@node-1:~$
  
```

NEW QUESTION # 154

You have a Kubernetes cluster with a namespace called 'dev' and a deployment named 'app-deployment' in that namespace. You need to create a new Role that allows users in the 'developers' group to only scale the 'app-deployment' deployment. They should not be able to access any other resources in the 'dev' namespace. Implement the RBAC configuration for this scenario.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a Role:

- Create a YAML file named 'scale-app-role.yaml' with the following content:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: scale-app-role
  namespace: dev
rules:
- apiGroups: ["apps"]
  resources: ["deployments"]
  verbs: ["get", "list", "watch", "update", "patch", "scale"]
  resourceNames: ["app-deployment"]

```

2. Create a RoleBinding: - Create a YAML file named 'scale-app-rolebinding.yaml' with the following content:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: scale-app-rolebinding
  namespace: dev
subjects:
- kind: Group
  name: developers
roleRef:
  kind: Role
  name: scale-app-role
  apiGroup: rbac.authorization.k8s.io

```

3. Apply the configuration: - Apply the Role and RoleBinding using the following commands: `bash kubectl apply -f scale-app-role.yaml` `bash kubectl apply -f scale-app-rolebinding.yaml` 4. Verify the configuration: - You can verify the configuration by using the following command: `bash kubectl auth can-i --list --as=user:testuser--group=developers--namespace=dev` - Replace 'testuser' with the name of a user in the 'developers' group. The output should show only the following permissions: - 'apps/deployments': 'get', 'list', 'watch', 'update', 'patch', 'scale' 5. Test the permissions: - Try to scale the Sapp-deployment deployment using the 'kubectl' command as a user in the 'developers' group. - Try to perform other actions on the deployment or other resources in the 'devs' namespace. You should only be able to scale the Sapp-deployment deployment.

NEW QUESTION # 155

You are running a critical application in Kubernetes that requires high availability and IOW latency. The application uses a statefulset With 3 replicas, each consuming a large amount of memory. You need to define resource requests and limits for the pods to ensure that the application operates smoothly and doesn't get evicted due to resource constraints.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Determine Resource Requirements:

- Analyze the application's memory usage. Determine the average memory consumption per pod and the peak memory usage.
- Consider the resources available on your Kubernetes nodes.
- Define realistic requests and limits based on the application's needs and available node resources.

2. Define Resource Requests and Limits in the StatefulSet:

- Update the StatefulSet YAML configuration with resource requests and limits for the container.
- requests: Specifies the minimum amount of resources the pod will request
- limits: Specifies the maximum amount of resources the pod can use.


```

apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: my-critical-app
spec:
  serviceName: "my-critical-app"
  replicas: 3
  selector:
    matchLabels:
      app: my-critical-app
  template:
    metadata:
      labels:
        app: my-critical-app
    spec:
      containers:
      - name: my-critical-app
        image: my-app-image:latest
        resources:
          requests:
            memory: "2Gi"
            cpu: "1" # 1 CPU core
          limits:
            memory: "4Gi"
            cpu: "2" # 2 CPU cores

```

3. Apply the StatefulSet Configuration: - Apply the updated StatefulSet configuration to your Kubernetes cluster: `bash kubectl apply -f my-critical-app-statefulset.yaml` 4. Monitor Resource Usage: - Use 'kubectl describe pod' to monitor the resource usage of the pods. - Ensure that the pods are utilizing the requested resources and not exceeding the limits.

NEW QUESTION # 156



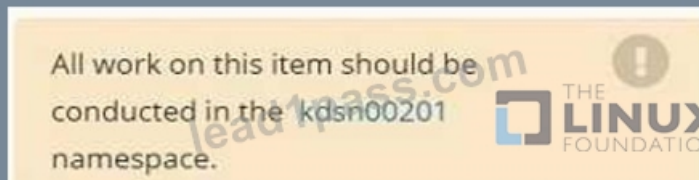
Set configuration context:

```
[student@node-1] $ kubectl config use-context nk8s
```

THE LINUX FOUNDATION

Task

You have rolled out a new pod to your infrastructure and now you need to allow it to communicate with the web and storage pods but nothing else. Given the running pod `kdsn00201` -newpod edit it to use a network policy that will allow it to send and receive traffic only to and from the web and storage pods.



All work on this item should be conducted in the `kdsn00201` namespace.

THE LINUX FOUNDATION

All required NetworkPolicy resources are already created and ready for use as appropriate. You should not create, modify or delete any network policies whilst completing this item.



Answer:

Explanation:

See the solution below.

Explanation

apiVersion: networking.k8s.io/v1

kind: NetworkPolicy

metadata:

name: internal-policy

namespace: default

spec:

podSelector:

matchLabels:

name: internal

policyTypes:

- Egress

- Ingress

ingress:

- {}

egress:

- to:

- podSelector:

matchLabels:

name: mysql

ports:

- protocol: TCP

port: 3306

- to:

- podSelector:

matchLabels:

name: payroll

ports:

- protocol: TCP

port: 8080

- ports:

- port: 53

protocol: UDP

- port: 53

protocol: TCP

NEW QUESTION # 157

You need to implement a mechanism for automatically rolling out new versions of your application pods. This process should be

triggered by a change in the application's container image tag in a Docker Hub repository.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Configure the Deployment for Rolling Updates:

- Update your application deployment to specify a 'rollingUpdate' strategy
- Set 'maxUnavailable' and 'maxSurge' to control the rolling update process-
- Include a 'strategy.type' to 'Rollingupdates
- Set 'imagePullPolicy' to 'Always' to ensure that new images are always pulled from the Docker Hub repository.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: your-application-deployment
  namespace: your-application-namespace
spec:
  replicas: 3
  selector:
    matchLabels:
      app: your-application
  template:
    metadata:
      labels:
        app: your-application
    spec:
      containers:
        - name: your-application
          image: your-docker-hub-account/your-image:latest
          imagePullPolicy: Always
  strategy:
    type: RollingUpdate
    rollingUpdate:
      maxUnavailable: 1
      maxSurge: 0
```

2. Apply the Deployment: - Apply the updated deployment using 'kubectl apply -f your-application-deployment.yaml' 3. Push a New Image to Docker Hub: - Update your application's container image in the Docker Hub repository and push the new image With a different tag. For example, update the tag from "latest to 'v2'. 4. Monitor the Deployment: - Observe the rolling update process using 'kubectl get pods -l app=your-application'. You should see new pods with the updated image being created and old pods being terminated. 5. Verify the Update: - Once the rolling update is complete, use 'kubectl describe deployment your-application-deployment' to verify that the 'updatedReplicas' field matches the 'replicas' field. This confirms that the update was successful ,

NEW QUESTION # 158

.....

Preparing for the CKAD exam can be a daunting task, but with real CKAD exam questions, it can be a lot easier. The importance of actual Linux Foundation Certified Kubernetes Application Developer Exam (CKAD) questions cannot be overemphasized. CKAD Real Questions are crucial for passing the CKAD exam. When candidates have access to the updated Linux Foundation CKAD practice test questions, they are better prepared to succeed.

CKAD Valid Exam Vce: <https://www.lead4pass.com/Linux-Foundation/CKAD-practice-exam-dumps.html>

- New CKAD Valid Test Registration | Pass-Sure CKAD: Linux Foundation Certified Kubernetes Application Developer

[illegible]

DOWNLOAD the newest Lead1Pass CKAD PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1Julbszpk1tIJVxXkxE9yY67AzPdpx6WS>