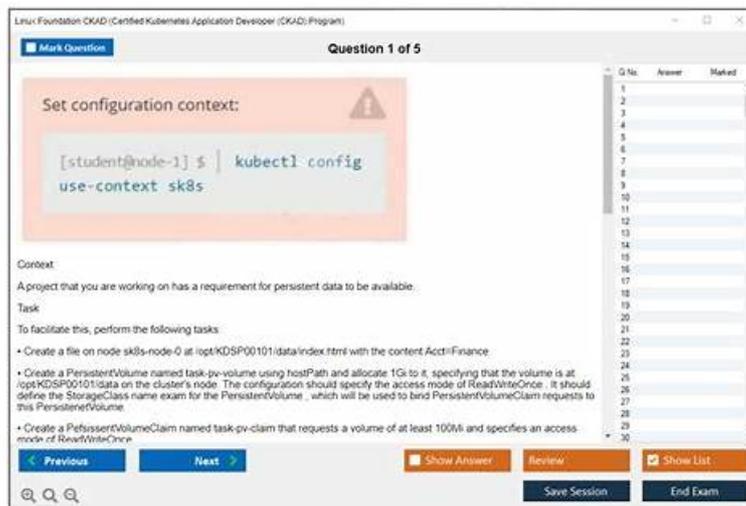


# CKAD exam braindumps & CKAD guide torrent



P.S. Free 2025 Linux Foundation CKAD dumps are available on Google Drive shared by BraindumpsPrep: <https://drive.google.com/open?id=11UVbRWHJTn12QvlgF3sgS3kArd7HGW0C>

Are you often regretful that you have purchased an inappropriate product? Unlike other platforms for selling test materials, in order to make you more aware of your needs, CKAD test preps provide sample questions for you to download for free. You can use the sample questions to learn some of the topics about CKAD learn torrent and familiarize yourself with the CKAD quiz torrent in advance. If you feel that the CKAD quiz torrent is satisfying to you, you can choose to purchase our complete question bank. After the payment, you will receive the email sent by the system within 5-10 minutes.

CKAD test questions have a mock examination system with a timing function, which provides you with the same examination environment as the real exam. Although some of the hard copy materials contain mock examination papers, they do not have the automatic timekeeping system. Therefore, it is difficult for them to bring the students into a real test state. With CKAD Exam Guide, you can perform the same computer operations as the real exam, completely taking you into the state of the actual exam, which will help you to predict the problems that may occur during the exam, and let you familiarize yourself with the exam operation in advance and avoid rushing during exams.

>> CKAD New Soft Simulations <<

## Free PDF 2026 Linux Foundation CKAD: Authoritative Linux Foundation Certified Kubernetes Application Developer Exam New Soft Simulations

Whether you are a newcomer or an old man with more experience, CKAD study materials will be your best choice for our professional experts compiled them based on changes in the examination outlines over the years and industry trends. CKAD test torrent not only help you to improve the efficiency of learning, but also help you to shorten the review time of up to several months to one month or even two or three weeks, so that you use the least time and effort to get the maximum improvement. And with our CKAD Exam Questions, your success is guaranteed.

The CKAD Exam is a hands-on, performance-based exam that tests the candidate's ability to deploy, configure, and manage Kubernetes applications. CKAD exam is designed to be challenging and comprehensive, covering all aspects of Kubernetes application development, including Kubernetes basics, application design and deployment, troubleshooting, and automation.

## Linux Foundation Certified Kubernetes Application Developer Exam Sample Questions (Q133-Q138):

### NEW QUESTION # 133

You are building a new web application that utilizes a microservice architecture- One of the microservices, 'recommendation-service', is responsible for providing personalized product recommendations to users.

This service uses a machine learning model for generating recommendations based on user purchase history and browsing behavior.

The model is trained offline and its weights are stored in a 'model-store' service.

Design a multi-container Pod for the 'recommendation-service' that incorporates the following considerations:

- The Pod should include a primary container for the 'recommendation-service' application.
- The Pod should include a secondary container that runs the 'model-store' service to provide access to the trained model weights.
- Both containers should share a common volume to ensure that the model weights are available to the 'recommendation-service' container-
- The 'recommendation-service' should be able to access the model weights from the 'model-store' container without relying on a network call to another service-
- The 'recommendation-service' container should be configured to periodically update the model weights from the 'model-store' container when a new version of the model is available.

**Answer:**

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create the Deployment YAML:

- Define a Deployment with the name 'recommendation-service'
- Set the replicas to for redundancy and scalability.
- Specify the labels 'app: recommendation-service' for selecting the Pods in the Deployment.
- Create a 'template' section to define the Pod specification

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: recommendation-service
spec:
  replicas: 3
  selector:
    matchLabels:
      app: recommendation-service
  template:
    metadata:
      labels:
        app: recommendation-service
    spec:
      containers:
        - name: recommendation-service
          image: example/recommendation-service:latest
          ports:
            - containerPort: 8080
          volumeMounts:
            - name: model-store
              mountPath: /model
          env:
            - name: MODEL_PATH
              value: /model/latest-model.weights
          command: ["python", "manage.py", "runserver", "0.0.0.0:8080"]
        - name: model-store
          image: example/model-store:latest
          volumeMounts:
            - name: model-store
              mountPath: /model
      volumes:
        - name: model-store
          emptyDir: {}
```

2. Deploy the Resources: - Apply the Deployment using 'kubectl apply -f deployment-yaml' 3. Verify the Deployment: - Check the status of the Deployment using 'kubectl get deployments recommendation-service' and ensure that three Pods are running. 4. Configure the 'recommendation-service' - Modify the 'recommendation-service' application to load the model weights from the specified path 'CLmodeVlatest-modelI\_weightS). - Implement a mechanism within the 'recommendation-service' to periodically check

for updated model weights in the shared volume. 5. Configure the 'model-store service': - Ensure that the model-store service is properly configured to store and retrieve the model weights. - Implement a mechanism in the 'model-store' service to notify the 'recommendation-service' when a new model version is available. This notification can be achieved using a shared volume or a separate messaging system. 6. Test the Application: - Send requests to the 'recommendation-service' to generate recommendations. - Monitor the 'model-store' service and the shared volume to verify that the model weights are being updated correctly and the recommendation-service is using the latest model version. Important Considerations: - Ensure that the 'recommendation-service' application is properly configured to access and load the model weights from the shared volume. - Implement a robust model management strategy, including versioning and rollback mechanisms, to ensure that the recommendation-service always uses the appropriate model. - Consider using a dedicated model store service that provides a dedicated API for retrieving and updating model weights. This can simplify the communication between the 'recommendation-service' and the model store. - Monitor the performance and resource usage of both services to ensure optimal performance.,

### NEW QUESTION # 134

You are running a critical application that requires high availability and minimal downtime during updates. Your current deployment strategy uses a single Deployment with 3 replicas. You need to ensure that during updates, only one pod is unavailable at any given time, minimizing service disruption. Design a deployment strategy that meets this requirement and allows for seamless updates.

#### Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define Rolling Update Strategy:

- In your Deployment configuration, specify the rolling update strategy with 'maxUnavailable: 1 s' and 'maxSurge: 0'. This ensures that during updates, only one pod is taken down at a time, while the remaining two continue serving traffic.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: critical-app-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: critical-app
  template:
    metadata:
      labels:
        app: critical-app
    spec:
      containers:
        - name: critical-app
          image: your-image-repo:latest
      strategy:
        type: RollingUpdate
        rollingUpdate:
          maxUnavailable: 1
          maxSurge: 0
```

2. Use Liveness and Readiness Probes: - Configure liveness and readiness probes for your application containers. Liveness probes check the health of running containers and restart them if unhealthy. Readiness probes check if a container is ready to receive traffic. - This ensures that only healthy pods are marked as ready, and traffic is routed only to ready pods.

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: critical-app-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: critical-app
  template:
    metadata:
      labels:
        app: critical-app
    spec:
      containers:
      - name: critical-app
        image: your-image-repo:latest
        livenessProbe:
          tcpSocket:
            port: 8080
          initialDelaySeconds: 15
          periodSeconds: 20
          failureThreshold: 3
        readinessProbe:
          tcpSocket:
            port: 8080
          initialDelaySeconds: 5
          periodSeconds: 10
          failureThreshold: 2
      strategy:
        type: RollingUpdate
        rollingUpdate:
          maxUnavailable: 1
          maxSurge: 0

```

3. Implement Horizontal Pod Autoscaling (HPA): - Set up HPA to automatically scale the number of pods based on CPU or memory utilization- This ensures that the application can handle increased traffic during updates without compromising performance. - You can configure the desired minimum and maximum replicas for the HPA based on your application's requirements. 4. Use Service with Session Affinity: - Configure your Service to use 'ClientIP' or 'Cookie' session affinity. This ensures that client connections are consistently routed to the same pod during the rolling update, minimizing disruption for users.

```

apiVersion: v1
kind: Service
metadata:
  name: critical-app-service
spec:
  selector:
    app: critical-app
  ports:
  - protocol: TCP
    port: 80
    targetPort: 8080
  sessionAffinity: ClientIP

```

5. Use Daemonsets for System Components: - If you have any system components (like monitoring agents or log collectors) that need to run on every node in the cluster, use DaemonSets instead of Deployments. - DaemonSets ensure that these components are always running on all nodes, even during node restarts or updates, ensuring continuous monitoring and logging.

### NEW QUESTION # 135

You are running a web application on a Kubernetes cluster, and you want to ensure that the container running your application is protected from potential security vulnerabilities. You are specifically concerned about unauthorized access to the container's filesystem. Explain how you would implement AppArmor profiles to restrict access to the container's filesystem.

**Answer:**

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define the AppArmor Profile:

- Create a new AppArmor profile file, for example, 'nginx-apparmor.conf', within your Kubernetes configuration directory.
- Within this file, define the restrictions for the container.
- For instance, to allow access to specific directories and files:

```
# include common AppArmor profile
include /etc/apparmor.d/abstractions/base/nginx.apparmor
# Allow access to specific directories
/var/www/html r,
/etc/nginx r,
# Allow access to specific files
/etc/nginx/nginx.conf r,
/usr/sbin/nginx r,
# Deny access to all other files and directories
Deny
```

2. Load the AppArmor Profile:

- Use the create configmap' command to create a ConfigMap containing your AppArmor profile:

Bash

```
kubectl create configmap nginx-apparmor-profile --from-file=nginx-apparmor.conf
```

3. Apply the Profile to Your Deployment:

- Update your Deployment YAML file to include the AppArmor profile:

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-deployment
spec:
  replicas: 2
  selector:
    matchLabels:
      app: nginx
  template:
    metadata:
      labels:
        app: nginx
    spec:
      containers:
      - name: nginx
        image: nginx:latest
        securityContext:
          # Enable AppArmor and specify the profile name
          appArmor: nginx-apparmor-profile
          # ... (rest of your Deployment YAML)
```

4. Restart the Pods: - Apply the updated Deployment YAML using 'kubectl apply -f nginx-deployment.yaml' - The updated deployment will restart the pods with the new AppArmor profile. 5. Verify the Profile: - Check the status of the pods with 'kubectl describe pod' - Look for the "Security Context" section and verify that the AppArmor profile is correctly applied. 6. Test the Restrictions: - Try to access files or directories that are not allowed by your AppArmor profile. - This will help you confirm that the profile is effectively restricting access.

#### NEW QUESTION # 136

You have a microservice application that relies on a Redis cache for data retrieval. Design a multi-container Pod that incorporates a Redis sidecar container to provide local caching within the Pod. Ensure that the main application container can access the Redis sidecar container within the same Pod Namespace Without needing to communicate with an external Redis cluster.

**Answer:**

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define the Pod YAML: Create a Pod YAML file that includes both the main application container and the Redis sidecar container.

```
apiVersion: v1
kind: Pod
metadata:
  name: my-app-pod
spec:
  containers:
  - name: my-app
    image: my-app-image:latest
    ports:
    - containerPort: 8080
    env:
    - name: REDIS_HOST
      value: redis
  - name: redis
    image: redis:latest
    ports:
    - containerPort: 6379
```

2. Configure Environment Variables: Set an environment variable 'REDIS HOST' within the main application container to point to the Redis sidecar containers hostname- In Kubernetes, containers within the same Pod can communicate with each other using their container names. 3. Connect Application to Redis: Modify the application code to connect to the Redis instance using the 'REDIS HOST' environment variable. For example, using a Python application with the 'redis-py' library: `python import redis r = redis-Redis(host=os.environ.get('REDIS_HOST'), port=6379) # Perform Redis operations (e.g., r.set('key', 'value'))` 4. Deploy the Pod: Apply the Pod YAML using `kubectl apply -f my-app-pod.yaml` 5. Verify Connectivity: Check the logs of the main application container to ensure it's successfully connecting to the Redis sidecar container Note: This approach provides local caching within the Pod, reducing external network calls and improving performance. It's important to consider potential data consistency issues if multiple Pods share the same Redis instance.

**NEW QUESTION # 137**

Refer to Exhibit.



Set Configuration Context:

```
[student@node-1] $ | kubectl
```

```
Config use-context k8s
```

Context

You sometimes need to observe a pod's logs, and write those logs to a file for further analysis.

Task

Please complete the following:

\* Deploy the counter pod to the cluster using the provided YAMLSpec file at `/opt/KDOB00201/counter.yaml`

\* Retrieve all currently available application logs from the running pod and store them in the file `/opt/KDOB00201/log_Output.txt`,

which has already been created

**Answer:**

Explanation:

Solution:

To deploy the counter pod to the cluster using the provided YAML spec file, you can use the `kubectl apply` command. The `apply` command creates and updates resources in a cluster.

```
kubectl apply -f /opt/KDOB00201/counter.yaml
```

This command will create the pod in the cluster. You can use the `kubectl get pods` command to check the status of the pod and ensure that it is running.

```
kubectl get pods
```

To retrieve all currently available application logs from the running pod and store them in the file `/opt/KDOB00201/log_Output.txt`, you can use the `kubectl logs` command. The `logs` command retrieves logs from a container in a pod.

```
kubectl logs -f <pod-name> > /opt/KDOB00201/log_Output.txt
```

Replace `<pod-name>` with the name of the pod.

You can also use `-f` option to stream the logs.

```
kubectl logs -f <pod-name> > /opt/KDOB00201/log_Output.txt &
```

This command will retrieve the logs from the pod and write them to the `/opt/KDOB00201/log_Output.txt` file.

Please note that the above command will retrieve all logs from the pod, including previous logs. If you want to retrieve only the new logs that are generated after running the command, you can add the `--since` flag to the `kubectl logs` command and specify a duration, for example `--since=24h` for logs generated in the last 24 hours.

Also, please note that, if the pod has multiple containers, you need to specify the container name using `-c` option.

```
kubectl logs -f <pod-name> -c <container-name> > /opt/KDOB00201/log_Output.txt
```

The above command will redirect the logs of the specified container to the file.

```
student@node-1:~$ kubectl create -f /opt/KDOB00201/counter.yaml
pod/counter created
student@node-1:~$ kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
counter       1/1     Running   0           10s
liveness-http 1/1     Running   0           6h45m
nginx-101     1/1     Running   0           6h46m
nginx-configmap 1/1     Running   0           107s
nginx-secret  1/1     Running   0           7m21s
poller        1/1     Running   0           6h46m
student@node-1:~$ kubectl logs counter
1: 2b305101817ae25ca60ae46510fb6d11
2: 3648cf2eae95ab680dba8f195f891af4
3: 65c8bbd4dbf70bf81f2a0984a3a44ede
4: 40d3a9c8e46f5533bb4828fbc5c8d038
5: 390442d2530a90c3602901e3fe999ac8
6: b71d95187417e139effb33af77681040
7: 66ae55a6491e756d2d0549ad6ab90a7
8: ff2b3d583b64125d2f9129c443bb37ff
9: b6c6a12b6e77944ed8baaf6c242dae4
10: bfcc9a894a0604fc4b814b37d0a200a4
student@node-1:~$ kubectl logs counter > /opt/KDOB00201/log_output.txt
student@node-1:~$
student@node-1:~$ kubectl logs counter > /opt/KDOB00201/log_output.txt
student@node-1:~$ kubectl logs counter > /opt/KDOB00201/log_output.txt
student@node-1:~$ cat /opt/KDOB00201/log_output.txt
```

```
student@node-1:~$ kubectl logs counter > /opt/KDOB00201/log_output.txt
student@node-1:~$ cat /opt/KDOB00201/log_output.txt
1: 2b305101817ae25ca60ae46510fb6d11
2: 3648cf2eae95ab680dba8f195f891af4
3: 65c8bbd4dbf70bf81f2a0984a3a44ede
4: 40d3a9c8e46f5533bb4828fbc5c8d038
5: 390442d2530a90c3602901e3fe999ac8
6: b71d95187417e139effb33af77681040
7: 66ae55a6491e756d2d0549ad6ab90a7
8: ff2b3d583b64125d2f9129c443bb37ff
9: b6c6a12b6e77944ed8baaf6c242dae4
10: bfcc9a894a0604fc4b814b37d0a200a4
11: 5493cd16a1790ab5fb9512b0e9d4c5dd1
12: 03f169e93e6143438e6dfefeb3cc9e9d
13: 764b37fe611373c42d0b47154041f6eb
14: 1a56fbc1896b0e6394136166281839e
15: ecc492eb17715de090c47345a98d98d7
16: 7971a6bec9eb44b0b8bbfc71aa3fbc74
17: 9e001be10174812ce9197ae19f77cd6
18: 231b22e3434272e4e9e0051174313f
19: ec7e1a5d314da9a0ad45d53be9a7acae
20: 0becdd8ee02cd442029e8162ed1c1197c
21: d6851ea4354621cb95bcb1ced997102
22: 7ed9a30ea8bf0d86206569481442af44
23: 29b8416ddc63dfcb987ab3c8198e9fe
24: 1f2062001df51a108ab25010f506716f
student@node-1:~$
```

