

# CKAD復習攻略問題、CKAD模擬トレーニング



2026年JPTestKingの最新CKAD PDFダンプおよびCKAD試験エンジンの無料共有: [https://drive.google.com/open?id=1dfZ\\_FVxnd24\\_BzXuybOUo-UjjKTflcUn](https://drive.google.com/open?id=1dfZ_FVxnd24_BzXuybOUo-UjjKTflcUn)

我々社のLinux Foundation CKAD問題集を使用して試験に合格しないで全額での返金を承諾するのは弊社の商品に不自信ではなく、行為でもって我々の誠意を示します。Linux Foundation CKAD問題集の专业化であれば、アフタサービスの細心であれば、我々JPTestKingはお客様を安心して購買して利用させます。お客様の満足は我々の進む力です。

CKAD試験は、Kubernetesアプリケーション開発における候補者の熟練度を示す業界で認められた認定試験です。クラウドネイティブアプリケーション開発と展開のスキルを向上させたいITプロフェッショナルに最適な認定試験です。この認定試験に合格することにより、Kubernetesリソース、アプリケーション設計・開発、デバッグ、トラブルシューティング、セキュリティに関する知識が証明されます。

CKAD試験の準備をするために、候補者は、Kubernetesの基礎を強く理解し、生産環境でKubernetesとの仕事の経験を持つことをお勧めします。Linux Foundationは、候補者が試験の準備を支援するために、幅広いトレーニングリソースとコースを提供しています。さらに、ブログ、フォーラム、オンラインミートアップなどの多くのコミュニティリソースを利用でき、候補者が試験を受けた他の人の経験から学習するのに役立ちます。

>> CKAD復習攻略問題 <<

## 検証するCKAD復習攻略問題 & 合格スムーズCKAD模擬トレーニング | 効果的なCKAD受験料過去問

Linux FoundationのCKAD試験に合格するのは難しいですが、合格できるのはあなたの能力を証明できるだけでなく、国際的な認可を得られます。Linux FoundationのCKAD試験の準備は重要です。我々JPTestKingの研究したLinux FoundationのCKADの復習資料は科学的方法でああなたの圧力を減少します。

Linux Foundation CKAD試験は、Kubernetesスキルを証明し、専門知識を認められる開発者にとって優れた機会です。経験豊富なKubernetesプロフェッショナルであろうと、初心者であろうと、この認定はキャリアを次のレベルに引き上げ、急速に進化するクラウドネイティブ開発の世界で新しい機会を開くことができます。

## Linux Foundation Certified Kubernetes Application Developer Exam 認定CKAD 試験問題 (Q195-Q200):

### 質問 # 195

You're managing a Kubernetes cluster with various applications. You want to implement a mechanism that automatically scales deployments based on CPU utilization. The scaling should be triggered when CPU utilization exceeds 70% and should scale down to 50% utilization.

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define the Horizontal Pod Autoscaler (HPA) YAMLI

- Create an HPA YAML file named 'auto-scaler.yaml' with the following contents:

```
apiVersion: autoscaling/v2beta2
kind: HorizontalPodAutoscaler
metadata:
  name: auto-scaler
  namespace: your-application-namespace
spec:
  scaleTargetRef:
    apiVersion: apps/v1
    kind: Deployment
    name: your-deployment
  minReplicas: 1
  maxReplicas: 5
  metrics:
  - type: Resource
    resource:
      name: cpu
      target:
        type: Utilization
        averageUtilization: 70
    targetCPUUtilizationPercentage: 50
```

2. Apply the HPA: - Apply the HPA YAML file using 'kubectl apply -f auto-scaler.yaml'. 3. Test the Auto-scaler - Monitor the CPU utilization of your deployment. When it exceeds 70%, the HPA will automatically scale up the deployment. - Observe the deployment scaling down when CPU utilization drops below 50%.

#### 質問 # 196

You are building a web application that uses a set of environment variables for configuration. These variables are stored in a ConfigMap named 'app-config'. How would you ensure that the web application pods always use the latest version of the ConfigMap even when the ConfigMap is updated?

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create the ConfigMap: Define the ConfigMap with your desired environment variables.

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: app-config
data:
  API_KEY: "your_api_key"
  DB_HOST: "db.example.com"
  DB_PORT: "5432"
```

2. Update the Deployment: Modify your Deployment YAML file to: - Use a 'volumeMount' to mount the ConfigMap into the container. - Specify a 'volume' using a 'configMap' source, referencing the 'app-config' ConfigMap. - Set 'imagePullPolicy: Always' to ensure the pod always pulls the latest container image.

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: web-app
spec:
  replicas: 3
  selector:
    matchLabels:
      app: web-app
  template:
    metadata:
      labels:
        app: web-app
    spec:
      containers:
      - name: web-app
        image: your-image-name:latest
        imagePullPolicy: Always
        envFrom:
        - configMapRef:
            name: app-config
        volumeMounts:
        - name: app-config-volume
          mountPath: "/etc/config"
      volumes:
      - name: app-config-volume
        configMap:
          name: app-config

```

3. Apply the changes: Use 'kubectl apply -f deployment-yaml' to update the Deployment. 4. Update the ConfigMap: Whenever you need to update the configuration, modify the 'app-config' ConfigMap using 'kubectl apply -f configmap-yaml'. 5. Verify changes: Observe the pods for the 'web-app' Deployment. They should automatically restart and pick up the new environment variables from the updated ConfigMap. By setting 'imagePullPolicy: Always', your pods will always pull the latest container image. This ensures that the pod's container always uses the latest code. Additionally, the 'volumeMount' and 'volume' definitions mount the 'app-config' ConfigMap into the container's '/etc/config' directory, making the environment variables accessible within the container. When you update the ConfigMap, the pod will detect the change and automatically restart, loading the new configuration from the updated ConfigMap.

#### 質問 # 197



#### Context

A user has reported an application is unresponsive due to a failing livenessProbe.

#### Task

Perform the following tasks:

\* Find the broken pod and store its name and namespace to /opt/KDOB00401/broken.txt in the format:



The output file has already been created

\* Store the associated error events to a file `/opt/KDOB00401/error.txt`, The output file has already been created. You will need to use the `-o` wide output specifier with your command

\* Fix the issue.



正解:

解説:

See the solution below.

Explanation

Solution:

Create the Pod:

```
kubectl create
```

```
-fhttp://k8s.io/docs/tasks/configure-pod-container/
exec-liveness.yaml
```

Within 30 seconds, view the Pod events:

```
kubectl describe pod liveness-exec
```

The output indicates that no liveness probes have failed yet:

```
FirstSeen LastSeen Count From SubobjectPath Type Reason Message
```

```
-----
24s 24s 1 {default-scheduler } Normal Scheduled Successfully assigned liveness-exec to worker0
23s 23s 1 {kubelet worker0} spec.containers{liveness} Normal Pulling pulling image
"gr.io/google_containers/busybox"
23s 23s 1 {kubelet worker0} spec.containers{liveness} Normal Pulled Successfully pulled image
"gr.io/google_containers/busybox"
23s 23s 1 {kubelet worker0} spec.containers{liveness} Normal Created Created container with docker id
86849c15382e; Security:[seccomp=unconfined]
23s 23s 1 {kubelet worker0} spec.containers{liveness} Normal Started Started container with docker id
86849c15382e
```

After 35 seconds, view the Pod events again:

```
kubectl describe pod liveness-exec
```

At the bottom of the output, there are messages indicating that the liveness probes have failed, and the containers have been killed

and recreated.

FirstSeen LastSeen Count From SubobjectPath Type Reason Message

```
-----  
37s 37s 1 {default-scheduler } Normal Scheduled Successfully assigned liveness-exec to worker0  
36s 36s 1 {kubelet worker0} spec.containers{liveness} Normal Pulling pulling image  
"gcr.io/google_containers/busybox"  
36s 36s 1 {kubelet worker0} spec.containers{liveness} Normal Pulled Successfully pulled image  
"gcr.io/google_containers/busybox"  
36s 36s 1 {kubelet worker0} spec.containers{liveness} Normal Created Created container with docker id  
86849c15382e; Security:[seccomp=unconfined]  
36s 36s 1 {kubelet worker0} spec.containers{liveness} Normal Started Started container with docker id  
86849c15382e  
2s 2s 1 {kubelet worker0} spec.containers{liveness} Warning Unhealthy Liveness probe failed: cat: can't open  
'/tmp/healthy': No such file or directory  
Wait another 30 seconds, and verify that the Container has been restarted:  
kubectl get pod liveness-exec  
The output shows that RESTARTS has been incremented:  
NAME READY STATUS RESTARTS AGE  
liveness-exec 1/1 Running 1 m
```

## 質問 # 198

You have a Deployment named 'database-deployment' that runs a PostgreSQL database container. You want to enforce the following security restrictions:

- The container should only be allowed to run with the I-IID 1000.
- The container should be able to access a specific hostPath volume mounted at '/db-data' for storing database data.
- The container should not be allowed to escalate privileges.
- The container should only have the 'NET\_BIND\_SERVICE' capability, allowing it to listen on specific ports.

You need to define a SecurityContext in the Deployment configuration to enforce these restrictions.

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define the SecurityContext

- Create a 'securityContext' section within the 'spec-template-spec-containers' block for your 'database-deployment container-
- Set 'runAsUsers' to '1000' to enforce running as UID 1000.
- Set 'allowPrivilegeEscalation' to 'false' to disable privilege escalation-
- In the 'capabilities' section
- Set 'drop' to an array containing all capabilities except 'NET\_BIND\_SERVICE'
- Set 'add' to an array containing 'NET\_BIND\_SERVICE'
- Define a 'volumeMount' to mount the '/db-data' hostPath volume.

Solution (Step by Step) :

1. Define the SecurityContext:

- Create a 'securityContext' section within the block for your 'database-deployment container.
- Set 'runAsUser' to '1000' to enforce running as UID 1000.
- Set 'allowPrivilegeEscalation' to 'false' to disable privilege escalation.
- In the 'capabilities' section:
- Set 'drop' to an array containing all capabilities except 'NET\_BIND\_SERVICE'
- Set 'add' to an array containing
- Define a 'volumeMount' to mount the '/db-data' hostPath volume.

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: database-deployment
spec:
  replicas: 1
  selector:
    matchLabels:
      app: database
  template:
    metadata:
      labels:
        app: database
    spec:
      containers:
        - name: postgres
          image: postgres:latest
          securityContext:
            runAsUser: 1000
            allowPrivilegeEscalation: false
          capabilities:
            drop: ["ALL"]
            add: ["NET_BIND_SERVICE"]
          volumeMounts:
            - name: db-data
              mountPath: /var/lib/postgresql/data
              readOnly: false
      volumes:
        - name: db-data
          hostPath:
            path: /db-data

```



2. Create the Deployment: - Apply the Deployment YAML file using 'kubectl apply -f database-deployment.yaml'. - The 'securityContext' restricts the container's behavior and capabilities. - Setting 'runAsUser' to '1000' forces the container to run as the specified UID. - 'allowPrivilegeEscalation' set to 'false' prevents the container from gaining higher privileges. - The 'capabilities' section controls specific capabilities. 'drop' removes unwanted capabilities, while 'add' grants specific capabilities. In this case, the container is allowed to use the capability, enabling it to bind to specific ports. - The 'volumeMount' defines the mount point for the hostPath volume, providing access to the specified directory for database data. This configuration ensures that the 'database-deployment' container runs with the specific UID, cannot escalate privileges, and only has the 'NET BIND SERVICE' capability, while accessing the hostPath volume for database data. This provides a secure environment for your database container.

### 質問 # 199

You're working on a Kubernetes application that involves retrieving data from a database. You have a Deployment With multiple pods, each accessing the database directly. To improve the application's performance and reliability, you want to implement an adapter pattern that introduces a service layer to handle database interactions. This layer should be responsible for connection pooling, caching, and error handling, making the application more resilient to database outages.

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a Service Account:

- Create a service account for the application. This will be used by the service layer to access the database.

```

apiVersion: v1
kind: ServiceAccount
metadata:
  name: db-service-account

```

- Apply the service account to the cluster: 'kubectl apply -f db-service-account.yaml' 2. Create a Role and RoleBinding: - Create a role that grants the necessary permissions to access the database.

```

apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  name: db-access-role
  namespace:
rules:
- apiGroups: ["database.example.com"] # Replace with actual API group
  resources: ["databases", "databases/credentials"]
  verbs: ["get", "list", "create", "update", "delete", "watch", "patch"]

```

- Create a role binding that associates the role with the service account

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: db-access-binding
  namespace:
subjects:
- kind: ServiceAccount
  name: db-service-account
  namespace:
roleRef:
  kind: Role
  name: db-access-role
apiGroup: rbac.authorization.k8s.io

```

- Apply the role and role binding to the cluster: - 'kubectl apply -f db-access-role.yaml' - 'kubectl apply -f db-access-binding.yaml'
3. Create the Service Layer Deployment: - Deploy the service layer component. This can be a containerized application that handles database interactions.

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: db-service
spec:
  replicas: 1 # Adjust as needed
  selector:
    matchLabels:
      app: db-service
  template:
    metadata:
      labels:
        app: db-service
    spec:
      serviceAccountName: db-service-account
      containers:
      - name: db-service
        image: # Replace with your service image
        ports:
        - containerPort: 8080 # Adjust based on your service's port
        env:
        - name: DATABASE_HOST
          value: # Replace with your database host
        - name: DATABASE_PORT
          value: # Replace with your database port
        - name: DATABASE_USER
          value: # Replace with your database user
        - name: DATABASE_PASSWORD
          valueFrom:
            secretKeyRef:
              name: db-credentials
              key: password # Replace with the key for your database password

```

- Apply the deployment: 'kubectl apply -f db-service.yaml'
4. Create a Secret for Database Credentials: - Create a secret to store sensitive database credentials.

```

apiVersion: v1
kind: Secret
metadata:
  name: db-credentials
stringData:
  password: # Replace with your database password

```

- Apply the secret 'kubectl apply -f db-credentials.yaml'
5. Create a Service for the Service Layer: - Create a service to expose the service layer to the application pods.

```

apiVersion: v1
kind: Service
metadata:
  name: db-service
spec:
  selector:
    app: db-service
  ports:
  - protocol: TCP
    port: 8080 # Adjust based on your service's port
    targetPort: 8080 # Match the containerPort in the Deployment

```

- Apply the service: 'kubectl apply -f db-service.yaml'
6. Update the Application Deployment: - Update the Deployment for your main application to use the service layer.

