

CKAD復習過去問 & CKAD勉強の資料

はじめに

この度は『Certified Kubernetes App Developer (CKAD) 開発者 解く問クイック』をお手に取っていただき、ありがとうございます。

本書は、隙間時間を使って効率よく学習ができるように構成された問題集です。
まずは自分の苦手な分野（レベル）から解き始めるもよし、最初から順に解くもよし。
合格を目指して、一緒に頑張りましょう。今日も一日、元気に行こうよ！

他の媒体・プラットフォームもご紹介します。
メダル実績機能やテーマ変更機能などございますので
ぜひ、ご活用ください！

- ・Android 版アプリリンクはこちら
- ・iOS 版アプリリンクはこちら
- ・Web 版リンクはこちら

さらに、Fast2test CKADダンプの一部が現在無料で提供されています：<https://drive.google.com/open?id=1P2CBrt0BDZyTg0hb9eEbSaB6yBR-l7>

CKAD試験トレントの3つのバージョンを提供しており、PDFバージョン、PCバージョン、APPオンラインバージョンが含まれています。各バージョンの機能と使用方法は異なり、実際の状況に適した最も便利なバージョンを選択できます。たとえば、PDFバージョンは、CKADテストトレントをダウンロードして印刷するのに便利です、学習の閲覧に適しています。PDFバージョンを使用している場合は、ペーパーで急流CKADガイドを印刷できます。CKAD試験問題のPCバージョンは、Linux Foundation Certified Kubernetes Application Developer Exam実際の試験環境を刺激します。

Linux Foundation CKAD Exam Syllabus Topics:

Section	Weight	Objectives
Topic 1: Application Observability and Maintenance	15%	- Understand API deprecations - Implement probes and health checks - Monitor applications using CLI tools - Work with container logs - Debug applications in Kubernetes
Topic 2: Application Design and Build	20%	- Define, build and modify container images - Choose and use appropriate workload resources - Understand multi-container Pod design patterns - Utilize persistent and ephemeral volumes

Topic 3: Application Deployment	20%	- Use Helm package manager - Implement deployment strategies - Work with Kustomize - Manage Deployments, rolling updates and rollbacks
Topic 4: Services and Networking	20%	- Expose applications via Services - Troubleshoot network access - Use Ingress rules - Understand and apply NetworkPolicies
Topic 5: Application Environment, Configuration and Security	25%	- Use ServiceAccounts - Create and consume Secrets - Work with ConfigMaps - Apply application security settings - Understand authentication, authorization and admission control - Configure resource requirements, limits and quotas - Use custom resources and extensions

>> CKAD復習過去問 <<

CKAD更新される学習資料、有効なCKADpdf問題集、Linux Foundation Certified Kubernetes Application Developer Exam勉強資料

テストの気分が悪い場合は、毎回CKADのソフトテストエンジンまたはアプリテストエンジンを選択する必要があります。これらの2つのバージョンには、実際のテストシーンをシミュレートする機能が1つあります。時間指定試験を設定し、何度も練習することができます。Linux Foundation CKADダンプトレントで試験のベースを感じ、テストする時間を確保できます。あなたがしたいことをしなければならない時間と機会を利用すべきです。CKADダンプトレントファイルを使用すると、テストの雰囲気を保つことができます。

Linux Foundation Certified Kubernetes Application Developer Exam 認定 CKAD 試験問題 (Q185-Q190):

質問 # 185

You are designing a microservice architecture where a frontend application needs to communicate with multiple backend services. The services are deployed as Pods in a Kubernetes cluster- To streamline communication and security, you decide to implement a sidecar proxy pattern Explain the benefits of using a sidecar proxy in this scenario and illustrate how you would implement it using a container image like Envoy Proxy.

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Understand Sidecar Proxy Pattern:

- In the sidecar proxy pattern, a proxy container runs alongside your main application container within the same Pod.
- The proxy acts as an intermediary, handling network traffic between your application and other services.

2. Benefits of Using a Sidecar Proxy:

- Traffic Management:
 - Routing requests to different backend services.
 - Load balancing across multiple instances of a service.
- Security:
 - Enforcing access control and authentication.
 - Handling SSL termination.
- Observability:
 - Monitoring and logging network traffic.
- Simplified Development:
 - Separating networking concerns from application logic.

3. Implementing with Envoy Proxy:

- Choose Envoy Proxy:
- Envoy is a popular open-source proxy designed for high-performance network communication.
- Create a Deployment YAML:
- Define a Deployment for your application, including a main container for your application code and a sidecar container for Envoy Proxy.

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-app-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: my-app
  template:
    metadata:
      labels:
        app: my-app
    spec:
      containers:
        - name: my-app
          image: my-app-image:latest
          ports:
            - containerPort: 8080
        - name: envoy
          image: envoyproxy/envoy:v1.23.0
          ports:
            - containerPort: 9901
          command: ["/usr/local/bin/envoy", "-c", "/etc/envoy/envoy.yaml"]
          volumeMounts:
            - name: config-volume
              mountPath: /etc/envoy
      volumes:
        - name: config-volume
          configMap:
            name: envoy-config
```

4. Configure Envoy: - Create a ConfigMap: - Create a ConfigMap to hold the Envoy configuration (envoy.yaml). - Define the routes, listeners, and clusters for your services.

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: envoy-config
data:
  envoy.yaml: |-
    static_resources:
      listeners:
        - name: listener_0
          address:
            socket_address:
              port_value: 9901
          filter_chains:
            - filters:
                - name: envoy.filters.network.http_connection_manager
                  typed_config:
                    "@type": type.googleapis.com/envoy.config.filter.network.http_connection_manager.v3.HttpConnectionManager
                    stat_prefix: ingress_http
                    route_config:
                      name: local_route
                      virtual_hosts:
                        - name: my_app_host
                          domains: [""]
                          routes:
                            - match:
                                prefix: "/service1"
                              route:
                                cluster: service1_cluster
                            - match:
                                prefix: "/service2"
                              route:
                                cluster: service2_cluster
                    http_filters:
                      - name: envoy.filters.http.router
      clusters:
        - name: service1_cluster
          connect_timeout: 0.25s
          type: strict_dns
          lb_policy: ROUND_ROBIN
          hosts:
            - socket_address:
                address: service1
                port_value: 8080
        - name: service2_cluster
          connect_timeout: 0.25s
          type: strict_dns
          lb_policy: ROUND_ROBIN
          hosts:
            - socket_address:
                address: service2
                port_value: 8080
```

5. Deploy the Configuration: - Apply the Deployment and ConfigMap using 'kubectl apply -f deployment.yaml' and 'kubectl apply -f configmap.yaml' 6. Test the Setup: - Access your frontend application from outside the cluster. - Verify that traffic is routed correctly to the backend services through the Envoy proxy. 7. Optional: Service Discovery: - For more dynamic environments, you can integrate Envoy with service discovery mechanisms like Kubernetes Service or Consul to automatically update its configuration based on service changes. This configuration assumes that your backend services are named "service1" and "service2" with ports 8080. Adjust the configuration based on your specific services and their port numbers.,

質問 # 186

You are building a Kubernetes application that requires persistent storage for its data. The application needs to be able to access the data even if the pod is restarted or deleted. You have a PersistentVolumeClaim (PVC) defined for this purpose.

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a PersistentVolume (PV):

- Define a PV with a suitable storage class, access modes (ReadWriteOnce), and a capacity that meets your application's storage requirements.

- Example:

```
apiVersion: v1
kind: PersistentVolume
metadata:
  name: my-pv
spec:
  capacity:
    storage: 1Gi
  accessModes:
    - ReadWriteOnce
  hostPath:
    path: "/data/my-pv"
  persistentVolumeReclaimPolicy: Retain
```

2. Create a PersistentVolumeClaim (PVC): - Define a PVC with the desired storage class and access modes. - Specify the desired storage capacity. - Example:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: my-pvc
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
  storageClassName: my-storage-class
```

3. Create a Deployment With the PVC: - In the Deployment YAML, define a volume mount that uses the PVC you created - Specify the volume mount path within the container. - Example:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-app
spec:
  replicas: 1
  selector:
    matchLabels:
      app: my-app
  template:
    metadata:
      labels:
        app: my-app
    spec:
      containers:
        - name: my-app-container
          image: my-app-image:latest
          volumeMounts:
            - name: my-data
              mountPath: /data
      volumes:
        - name: my-data
          persistentVolumeClaim:
            claimName: my-pvc

```

4. Create the Deployment: - Apply the Deployment YAML using 'kubectl apply -f my-app-deployment.yaml' 5. Verify the Deployment - Check the status of the Deployment using 'kubectl get deployments my-app' - Verify that the Pod is running and using the PersistentVolumeClaim - You can also check the pod's logs for confirmation that the data is stored in the mounted volume.

質問 # 187

You must connect to the correct host . Failure to do so may result in a zero score.

[candidate@base] \$ ssh ckad00021

Task

Create a Cronjob named grep that executes a Pod running the following single container:

```

name: busybox
image: busybox:stable
command: ["grep", "-i", "nameserv
er", "/etc/resolv.conf"]

```

Configure the CronJob to:

- * execute Once every 30 minutes
- * keep 96 completed Job
- * keep 192 failed Job
- * never restart podsterminate pods after 8 seconds

Manually create and execute once job

named grep-test from the grep Cronjob

正解:

解説:

See the Explanation below for complete solution.

Explanation:

ssh ckad00021

Below is the clean, CKAD-friendly way (YAML + apply + verify + manual job).

1) Create the CronJob grep

Create a file (anywhere, e.g. in your home):

```
cat <<'EOF' > grep-cronjob.yaml
```

```
apiVersion: batch/v1
```

```
kind: CronJob
```

```
metadata:
```

```
name: grep
```

```
spec:
```

```
schedule: "*/30 * * * *
```

```
successfulJobsHistoryLimit: 96
```

```
failedJobsHistoryLimit: 192
```

```
jobTemplate:
```

```
spec:
```

```
activeDeadlineSeconds: 8
```

```
template:
```

```
spec:
restartPolicy: Never
containers:
- name: busybox
image: busybox:stable
command: ["grep", "-i", "nameserver", "/etc/resolv.conf"]
EOF
Apply it:
kubectl apply -f grep-cronjob.yaml
Verify:
kubectl get cronjob grep
kubectl describe cronjob grep
Confirm the key fields quickly:
kubectl get cronjob grep -o jsonpath='{.spec.schedule} {"\n"} {.spec.successfulJobsHistoryLimit} {"\n"} {.spec.failedJobsHistoryLimit} {"\n"}'
kubectl get cronjob grep -o jsonpath='{.spec.jobTemplate.spec.activeDeadlineSeconds} {"\n"} {.spec.jobTemplate.spec.template.spec.restartPolicy} {"\n"}'
2) Manually create and execute the one-off Job grep-test from the CronJob Create the Job from the CronJob:
kubectl create job --from=cronjob/grep grep-test
Watch it:
kubectl get jobs grep-test
kubectl get pods -l job-name=grep-test
Get logs (most important proof):
POD=$(kubectl get pods -l job-name=grep-test -o jsonpath='{.items[0].metadata.name}') kubectl logs "$POD" You should see
one or more nameserver ... lines from /etc/resolv.conf.
```

質問 # 188

You have a web application that requires a dedicated sidecar container to manage logging and monitoring. The sidecar container should be deployed alongside every pod of the application. You need to ensure that the sidecar container is always available alongside the application pods, even if the main application container experiences failures. Which Kubernetes resource is most suitable for this scenario and why?

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Choose DaemonSet The most suitable Kubernetes resource for this scenario is a DaemonSet.
2. Daemonset Functionality: Daemonsets ensure that a pod is running on every node in your cluster. This is ideal for sidecar containers because they need to be present alongside the main application pod on each node.
3. Daemonset Benefits:
 - Guaranteed Availability: Daemonsets guarantee that the sidecar container is always available on the same node as the main application pod, even if the application pod is restarted or fails.
 - Pod Management: DaemonSets manage the lifecycle of the sidecar container, ensuring its availability and resource allocation.
 - Node-Level Deployment: Daemonsets deploy pods on all nodes, ensuring consistent functionality across the cluster
4. Implementation Example:

```

apiVersion: apps/v1
kind: DaemonSet
metadata:
  name: logging-sidecar
spec:
  selector:
    matchLabels:
      app: logging-sidecar
  template:
    metadata:
      labels:
        app: logging-sidecar
    spec:
      containers:
        - name: logging-sidecar
          image: your-logging-image:latest
          # ... other container configuration
        - name: your-application
          image: your-app-image:latest
          # ... other container configuration

```

This DaemonSet definition specifies a pod with two containers: the 'logging-sidecar' and 'your-application'. The 'logging-sidecar' is your sidecar container, and 'your-application' represents your main application. - Important: The Daemonset will ensure that a pod with these containers is deployed on every node of your Kubernetes cluster. 5. Deployment and Monitoring: - Deployment: Use 'kubectl apply -f logging-sidecar.yaml' to deploy the DaemonSet. - Monitoring: Observe the pods created by the Daemonset using 'kubectl get pods'. You should see a pod with the 'logging-sidecar' and 'your-application' containers running on each node. 6. Conclusion: - Using a DaemonSet to manage your sidecar container ensures its consistent availability alongside the main application pods, guaranteeing logging and monitoring capabilities even in case of pod failures.

質問 # 189

You have a Deployment running a web application built With a Node.js container. The application currently uses an older version of the Node.js runtime, and you need to upgrade to a newer version. Describe the steps involved in modifying the container image to include the new Node.js runtime without rebuilding the entire application.

正解:

解説:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a Dockerfile:

- Create a new 'Dockerfile' With the following content

```

FROM node:16-alpine # Use the desired Node.js version
COPY --from=existing-image:latest /app /app
WORKDIR /app
CMD ["npm", "start"]

```

- Replace With the name of the existing Docker image used by your Deployment. - This Dockerfile uses a multi-stage build approach. It starts with a new Node.js base image and copies the application code from the existing image. This allows you to update the runtime without rebuilding the entire application. 2. Build the New Image: - Build the image using the Dockerfile: docker build -t updated-image:latest 3. Update the Deployment - Modify your Deployment YAML file to use the newly built image:

```

apiVersion: apps/v1
kind: Deployment
metadata:
  name: my-node-app
spec:
  replicas: 3
  selector:
    matchLabels:
      app: my-node-app
  template:
    metadata:
      labels:
        app: my-node-app
    spec:
      containers:
        - name: my-node-app
          image: updated-image:latest # Use the new image name
          ports:
            - containerPort: 8080
          restartPolicy: Always

```

ちなみに、Fast2test CKADの一部をクラウドストレージからダウンロードできます：<https://drive.google.com/open?id=1P2CBrt0BDzyyTg0lvb9eEbSaB6yBR-17>