

CKAD시험응시료 & CKAD시험패스덤프공부자료

안녕하세요
조훈(Hoon Jo) 입니다.

[클라우드△솔루션_아키텍처]



더읽기 ▾

커리큘럼 총 13 개 2시간 16분의 수업

이 강의는 영상, 수업 노트가 제공됩니다. 미리보기를 통해 콘텐츠를 확인해보세요. 모두 보기

^	섹션 0. 쿠버네티스 시험(CKA, CKAD, CKS) 업데이트 - 2022.1.31 기준	2 강 · 44분
⊙	X.13.000-1.쿠버네티스 시험(CKA, CKAD, CKS)을 위한 공부법	27:49
⊙	X.13.000-2.Killer.sh 난이도 설명	16:35

2026 KoreaDumps 최신 CKAD PDF 버전 시험 문제집과 CKAD 시험 문제 및 답변 무료 공유:
https://drive.google.com/open?id=1yXFzqCZlc0yljydnE_6xW67GzeP3NHL

우리사이트가 다른 덤프사이트보다 우수한 점은 바로 자료들이 모두 전면적이고 적중률과 정확입니다. 때문에 우리KoreaDumps를 선택함으로Linux Foundation인증CKAD시험준비에는 최고의 자료입니다. 여러분이 성공을 위한 최고의 자료입니다.

CKAD 시험은 Kubernetes 플랫폼에서 클라우드 네이티브 애플리케이션을 생성하고 배포 할 때 개발자의 실제 기술을 테스트하도록 설계되었습니다. 이 시험은 컨테이너 오케스트레이션 기술, Kubernetes API 프리미티브 및 Kubernetes 아키텍처의 핵심 개념을 포함하여 Kubernetes 애플리케이션을 설계, 구축 및 문제 해결하는 개발자의 기능을 평가합니다. 인증은 Kubernetes 애플리케이션 개발 기술을 향상시키고 잠재적 고용주에 대한 숙련도를 보여 주려는 개발자를 대상으로합니다. CKAD 인증은 또한 인증 된 Kubernetes 관리자 (CKA) 인증과 같은 Kubernetes에서 고급 인증을 추구하려는 개발자에게 필수 전제 조건입니다.

>> CKAD시험응시료 <<

적중률 좋은 CKAD시험응시료 덤프문제자료

IT인증자격증은 국제적으로 승인받는 자격증이기 때문에 많이 취득해두시면 취업이나 승진이나 이직이나 모두 편해집니다. 다른 사람이 없는 자격증을 내가 가지고 있다는것은 실력을 증명해주는 수단입니다. Linux Foundation인증 CKAD시험은 널리 승인받는 자격증의 시험과목입니다. Linux Foundation인증 CKAD덤프로Linux Foundation인증 CKAD시험공부를 하시면 시험패스 난이도가 낮아지고 자격증 취득율이 높고 올라갑니다.자격증을 많이 취득하여 취업이나 승진의 문을 두드려 보시면 빈틈없이 달힌 문도 활짝 열릴것입니다.

CKAD 인증 시험을 준비하기 위해 개발자는 교육 과정, 학습 가이드 및 실습 시험을 포함하여 Linux Foundation에서 제공하는 다양한 리소스를 활용할 수 있습니다. Linux Foundation은 또한 개발자가 다른 CKAD 후보와 연결하고 시험 준비 방법에 대한 팁과 조언을 공유 할 수있는 커뮤니티 포럼을 제공합니다. 올바른 준비와 헌신으로 개발자는 CKAD 인증을 받고 Kubernetes 응용 프로그램 개발 분야에서 다음 단계로 경력을 쌓을 수 있습니다.

최신 Kubernetes Application Developer CKAD 무료샘플문제 (Q152-Q157):

질문 # 152

You have a Kustomization file that uses a resource patch to modify the deployment of an Nginx service. The patch uses the field to set the CPLJ request for the container to 500m. However, you've noticed that this patch is no longer working as expected. You've

been informed that the field has been deprecated and replaced with a new field structure in newer Kubernetes API versions. Explain how to update the Kustomization file to accommodate this change, ensuring compatibility with both older and newer Kubernetes versions.

정답:

설명:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Identify the New Field Structure: Research the updated field structure for container resource definitions in the newer Kubernetes API version. The new structure likely utilizes nested resource fields for each container, like instead of a flat structure.
2. Update the Kustomization Patch: Modify the resource patch in your Kustomization file to use the updated field structure. If the newer field structure is 'spec-template-spec-containers[L]resources.requests.cpu', update your patch accordingly. This could involve changing the patch's path or using a different patch strategy, such as a strategic merge patch.

```
# kustomization.yaml
resources:
- deployment.yaml
patchesStrategicMerge:
- patch.yaml
# patch.yaml
---
spec:
  template:
    spec:
      containers:
      - name: nginx
        resources:
          requests:
            cpu: 500m
```

3. Consider Conditional Patches: If you need to support both older and newer Kubernetes versions, utilize conditional patches in your Kustomization file. This allows you to apply different patches based on the Kubernetes API version detected. You can use Kustomize's 'patchJson6902' strategy With a conditional statement to apply the correct patch depending on the API version.

```
# kustomization.yaml
resources:
- deployment.yaml
patchesStrategicMerge:
- patch.yaml
patchesJson6902:
- patch: |-
  {
    "op": "replace",
    "path": "/spec/template/spec/containers/0/resources/requests/cpu",
    "value": "500m"
  }
  target:
  kind: Deployment
  apiVersion: apps/v1
- patch: |-
  {
    "op": "replace",
    "path": "/spec/template/spec/containers/0/resources/requests/cpu",
    "value": "500m"
  }
  target:
  kind: Deployment
  apiVersion: apps/v1beta1
```

4. Test the Updated Kustomization: Deploy your Kustomization to a cluster running both older and newer Kubernetes versions. Validate that the CPU requests are correctly applied to the Nginx deployment containers in each version. Verify that the patches are being applied appropriately based on the detected Kubernetes API version.
5. Document Changes: Ensure that the updated Kustomization file and any conditional logic are well-documented to prevent future confusion or errors when deploying to different Kubernetes environments. By following these steps, you can successfully update your Kustomization file to accommodate the deprecated field structure and ensure compatibility with different Kubernetes API versions. This will allow you to manage and configure your deployments effectively, even as Kubernetes evolves.

You are designing a container image for a Python application that uses a specific version of a Python library (requests). You want to ensure that this specific library version is always used, regardless of the host system's installed version. Explain how you would achieve this within your Dockerfile.

정답 :

설명 :

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Install Library in Dockerfile:

- Utilize the 'COPY' instruction in your Dockerfile to copy a requirements file containing the exact library version you need.
- Use the 'RUN' instruction to install the library from the requirements file.

- Example:

Dockerfile

```
FROM python:3.5
```

```
COPY requirements.txt
```

```
RUN pip install -r requirements.txt
```

```
COPY
```

```
CMD ["python", "app.py"]
```

2. Create Requirements File ('requirements.txt'):

- Create a 'requirements.txt' file within your project directory.
- Add the specific version of the 'requests' library to this file.

- Example:

```
requests==2.28.1
```

3. Build the Docker Image:

- Construct your Docker image using the Dockerfile.
- Run the following command: 'docker build -t your-image-name .'

4. Run the Container:

- Launch the container in Kubernetes.
- Verify that the 'requests' library with the specified version is successfully used within the container.

질문 # 154

Refer to Exhibit.



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.p.yaml, and mount that ConfigMap to /fluentd/etc in the adapter-zen sidecar container

정답:

설명:

Solution:

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

```
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: lfcncf/busybox:1
        name: busybox
        resources: {}
status: {}
```

```
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: myvoll
        emptyDir: {}
      containers:
      - image: lfccncf/busybox:1
        name: logger-dev
        volumeMounts:
        - name: myvoll
          mountPath: /tmp/log
      - image: lfccncf/fluentd:v0.12
        name: adapter-zen
```

3 lines yanked

27,22

Bot

```
replicas: 1
selector:
  matchLabels:
    app: deployment-xyz
template:
  metadata:
    labels:
      app: deployment-xyz
  spec:
    volumes:
    - name: myvoll
      emptyDir: {}
    containers:
    - image: lfccncf/busybox:1
      name: logger-dev
      command: ["/bin/sh", "-c", "while true; do echo 'i luv cncf' >> /tmp/log/input.log; sleep 10; done"]
      volumeMounts:
      - name: myvoll
        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: myvoll
        mountPath: /tmp/log
```

29,83

Bot

```

metadata:
  labels:
    app: deployment-xyz
spec:
  volumes:
  - name: myvol1
    emptyDir: {}
  - name: myvol2
    configMap:
      name: logconf
  containers:
  - image: lfcncf/busybox:1
    name: logger-dev
    command: ["/bin/sh", "-c", "cat /dev/urandom | true ]; do echo 'i luv -cf' >> /tmp/log/input.log; sleep 10; done"]
    volumeMounts:
    - name: myvol1
      mountPath: /tmp/log
  - image: lfcncf/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

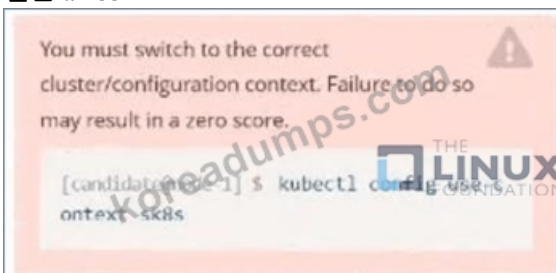
Bot

```

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 scale deployment
NAME          READY   UP-TO-DATE   AVAILABLE   AGE
deployment-xyz 1/1     1             1           12s
student@node-1:~$

```

질문 # 155



Task:

Update the Deployment app-1 in the frontend namespace to use the existing ServiceAccount app.

정답:

설명:

See the solution below.

Explanation:

Solution:

```

terminal - candidate@node-1 -
File Edit View Terminal Tabs Help
The programs included with the Ubuntu system are free software;
the exact distribution terms for each program are described in the
individual files in /usr/share/doc/*/copyright.

Ubuntu comes with ABSOLUTELY NO WARRANTY, to the extent permitted by
applicable law.

candidate@node-1:~$ vi ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context sk8s
Switched to context "sk8s".
candidate@node-1:~$ vim .vimrc
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl apply -f ~/spicy-pikachu/backend-deployment.yaml
deployment.apps/backend-deployment configured
candidate@node-1:~$ kubectl get pods -n staging
NAME                                READY   STATUS    RESTARTS   AGE
backend-deployment-59d449b99d-cxct6 1/1     Running   0           20s
backend-deployment-59d449b99d-h2zjq 0/1     Running   0           9s
backend-deployment-78976f74f5-b8c85 1/1     Running   0           6h40m
backend-deployment-78976f74f5-flfsj 1/1     Running   0           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME    READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment 3/3      3             3           6h40m
candidate@node-1:~$ kubectl get deploy -n staging
NAME    READY   UP-TO-DATE   AVAILABLE   AGE
backend-deployment 3/3      3             3           6h41m
candidate@node-1:~$ vim ~/spicy-pikachu/backend-deployment.yaml
candidate@node-1:~$ kubectl config use-context k8s
Switched to context "k8s".
candidate@node-1:~$ kubectl set serviceaccount deploy app-1 app -n frontend
deployment.apps/app-1 serviceaccount updated
candidate@node-1:~$

```

질문 # 156

You're running a MYSQL database pod in a Kubernetes cluster. You need to ensure that the pod is always running on a specific node, regardless of node failures or maintenance events. This node has specific hardware or software requirements that the MySQL database requires. How do you achieve this?

정답:

설명:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a Node Affinity: Define a node affinity rule for your MySQL pod that specifically targets the desired node. You'll use 'nodeselector' or 'nodeAffinity' in your pod definition.

```

apiVersion: v1
kind: Pod
metadata:
  name: mysql-pod
spec:
  affinity:
    nodeAffinity:
      requiredDuringSchedulingIgnoredDuringExecution:
        nodeSelectorTerms:
          - matchExpressions:
              - key: "kubernetes.io/hostname"
                operator: In
                values:
                  - "your-specific-node-name"
  containers:
    - name: mysql
      image: mysql:latest
      ports:
        - containerPort: 3306

```

2. Apply the Pod Definition: Apply the YAML configuration to your Kubernetes cluster using 'kubectl apply -f mysql-pod.yaml'
3. Verify Pod Placement: Use 'kubectl get pods -l app=mysql' to verify that the pod is running on the intended node (i.e., "your-specific-node-name").
4. Handle Node Failure: While this ensures the pod starts on the desired node, if that node fails, the pod will not be automatically rescheduled. To address this, consider using:
 - Node Selectors: You can combine 'nodeselector' with 'nodeAffinity' to prioritize your specific node. This ensures that the pod tries to schedule on your preferred node first-
 - Taint and Tolerations: You can taint the specific node with a unique key and then add a toleration to your MySQL pod to tolerate that taint. This allows the pod to be scheduled on that node and only that node.
5. Deployment for Scalability: If you need to run multiple

