

Latest CKAD Test Answers | Reliable CKAD Exam Topics



What's more, part of that Prep4away CKAD dumps now are free: <https://drive.google.com/open?id=14ReHhztT7W13qQyL3WmpmxDpJBPsdGyF>

In the complicated and changeable information age, have you ever been tried hard to find the right training materials of CKAD exam certification? We feel delighted for you to find Prep4away, and more delighted to find the reliable CKAD Exam Certification training materials. It will help you get your coveted CKAD exam certification.

The time for CKAD test certification is approaching. If you do not prepare well for the Linux Foundation certification, please choose our CKAD exam test engine. You just need to spend 20-30 hours for study and preparation, then confident to attend the actual test. If you have any question about CKAD study pdf, please contact us at any time. The online chat button is at the right bottom of the Prep4away page. Besides, we guarantee money refund policy in case of failure.

>> Latest CKAD Test Answers <<

Actual CKAD Exam Dumps Will Be the Best Choice to Prepare for Your Exam

As the most important element that almost all the candidates will take into consider, the pass rate of our CKAD exam questions is high as 98% to 100%, which is unique in the market and no one has made it. And also the exam passing guarantee that makes our CKAD Study Guide superior in the market. As the best seller, our CKAD learning braindumps are very popular among the candidates. Many of the loyal customers are introduced by their friends or classmates.

The CKAD Exam is a performance-based exam that requires candidates to complete a set of tasks within a given time frame. CKAD exam is conducted online and candidates are required to use a terminal and a web browser to complete the tasks. CKAD exam is designed to test the candidate's ability to work with Kubernetes in a hands-on environment and to complete tasks that are similar to those encountered in real-world Kubernetes application development scenarios.

Linux Foundation CKAD Exam is an ideal certification for IT professionals who want to enhance their skills in cloud-native application development and deployment. Linux Foundation Certified Kubernetes Application Developer Exam certification validates a candidate's knowledge of Kubernetes resources, application design and development, debugging, troubleshooting, and security. To prepare for the exam, candidates can take advantage of various resources provided by the Linux Foundation and join the Kubernetes community to gain insights into best practices.

Linux Foundation Certified Kubernetes Application Developer Exam Sample Questions (Q94-Q99):

NEW QUESTION # 94

Refer to Exhibit.

Set Configuration Context:

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

Config use-context k8s

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.

Answer:

Explanation:

To allow a pod to send and receive traffic only to and from specific pods, you can use network policies in Kubernetes.

First, you will need to create a network policy that defines the allowed traffic. You can create a network policy yaml file with the following rules:

```
apiVersion: networking.k8s.io/v1
```

```
kind: NetworkPolicy
```

```
metadata:
```

```
  name: newpod-network-policy
```

```
  namespace: default
```

```
spec:
```

```
  podSelector:
```

```
    matchLabels:
```

```
      app: kdsn00201-newpod
```

```
  ingress:
```

```
    - from:
```

```
      - podSelector:
```

```
        matchLabels:
```

```
          app: web
```

```
      - podSelector:
```

```
        matchLabels:
```

```
          app: storage
```

This policy will only allow incoming traffic to the pod with the label app=kdsn00201-newpod from pods with the label app=web or app=storage. If you have different labels on your web and storage pods please update the matchLabels accordingly.

Once you have created the network policy, you can apply it to the cluster by running the following command:

```
kubectl apply -f <network-policy-file>.yaml
```

This will apply the network policy to the cluster, and the newpod will only be able to send and receive traffic to and from the web and storage pods.

Please note that, NetworkPolicy resource is not available by default, you need to enable the NetworkPolicy feature on your

Kubernetes cluster. This feature is enabled by default on some clusters and must be explicitly enabled on others. You can check if NetworkPolicy is available by running the command `kubectl api-versions | grep networking`. Also, you need to ensure that the pods that you want to allow traffic to and from are running on the same namespace.

NEW QUESTION # 95

You have a Deployment named 'my-app-deployment' running a Flask application. You want to add a liveness probe that checks if the Flask application is responding on port '5000' and a readiness probe that checks if the application is ready to receive requests. Implement these probes using Kustomize.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Create a base Deployment configuration:

2. Create a 'kustomization.yaml' file:

3. Create 'patches/liveness-probe.yaml':

4. Create 'patches/readiness-probe.yaml':

5. Apply the Kustomize configuration: `bash kustomize . | kubectl apply -t -`

- Liveness probe: This probe checks if the application is still alive and running. It uses a TCP socket to connect to port '5000' and waits for 15 seconds before making the first check. It checks every 20 seconds, and if it fails 3 times in a row, the pod is restarted.

- Readiness probe: This probe checks if the application is ready to receive requests. It also uses a TCP socket to connect to port '5000'. It checks every 10 seconds and waits for 5

seconds before the first check. If it fails 2 times in a row, the pod is marked as unhealthy and excluded from receiving traffic. Note: Make sure your Flask application is actually listening on port '5000' and responding to requests. ,

NEW QUESTION # 96

You have a Deployment named 'wordpress-deployment' that runs 3 replicas of a WordPress container. You need to implement a persistent volume claim (PVC) for each pod that stores the website data, and you want to ensure that the data persists even if the pod is deleted or restarted. The PVC should be created using a storage class named 'standard' with a capacity of 10Gi.

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1). Create a Storage Class:

- Create a 'standard' storage class:

- Apply the YAML file: `bash kubectl apply -f standard-storage-class.yaml` 2. Create a Persistent Volume Claim - Create a PVC named 'wordpress-pvc' with a request for 10Gi storage and using the 'standard' storage class:

- Apply the YAML file: `bash kubectl apply -f wordpress-pvc.yaml` 3. Update the Deployment - Update the 'wordpress-deployment' YAML file to mount the PVC to each pod:

- Apply the updated YAML file: `bash kubectl apply -f wordpress-deployment.yaml` 4. Verify the Deployment - Check the status of the deployment using `'kubectl get deployments wordpress-deployment'` to confirm the rollout and updated replica count. - Use `'kubectl describe pods -l app=wordpress'` to confirm that each pod is using the 'wordpress-pvc' and the website data is stored in the persistent volume. - You can now access the WordPress website through the service that is associated with the Deployment. 5. Test Data Persistence: - Delete or restart one of the pods in the deployment. - Observe that the website data remains intact because the PVC is persistent and the data is stored in the underlying volume.,

NEW QUESTION # 97

You are tasked with designing a multi-container Pod that hosts both a web server and a database. The web server should be able to connect to the database within the pod- How would you implement this design, including networking considerations?

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1. Define the Pod YAML:- Create a Pod definition that includes two containers: one for the web server and one for the database.

2. Configure Networking: The key to allowing the web server to connect to the database is to use the pod's internal network. Since containers within a pod share the same network namespace, you can configure the webserver to connect to the database using the name 'db'. 3. Environment Variables: Set an environment variable (CDB_HOST) within the webserver container to point to the database container by its name. This ensures the web server can correctly connect to the database within the pod. 4. Pod

Deployment: Apply the YAML to create the pod using `'kubectl apply -f web-db-pod.yaml'` 5. Verification: To check the pod's

status: - Run `'kubectl get pods'` - Check the logs of the web server container to confirm it can connect to the database. 6. Important

Note: In this example, we're using the default pod networking within Kubernetes. For more complex applications, consider using a service to expose the database container This will allow access to the database from outside the pod.,

NEW QUESTION # 98

Context

A pod is running on the cluster but it is not responding.

Task

The desired behavior is to have Kubernetes restart the pod when an endpoint returns an HTTP 500 on the /healthz endpoint. The service, probe-pod, should never send traffic to the pod while it is failing. Please complete the following:

* The application has an endpoint, /started, that will indicate if it can accept traffic by returning an HTTP 200.

If the endpoint returns an HTTP 500, the application has not yet finished initialization.

* The application has another endpoint /healthz that will indicate if the application is still working as expected by returning an HTTP

200. If the endpoint returns an HTTP 500 the application is no longer responsive.

* Configure the probe-pod provided to use these endpoints

* The probes should use port 8080

Answer:

Explanation:

See the solution below.

Explanation

Solution:

apiVersion: v1

kind: Pod

metadata:

labels:

test: liveness

name: liveness-exec

spec:

containers:

- name: liveness

image: k8s.gcr.io/busybox

args:

- /bin/sh

- -c

- touch /tmp/healthy; sleep 30; rm -rf /tmp/healthy; sleep 600

livenessProbe:

exec:

command:

- cat

- /tmp/healthy

initialDelaySeconds: 5

periodSeconds: 5

In the configuration file, you can see that the Pod has a single Container. The periodSeconds field specifies that the kubelet should perform a liveness probe every 5 seconds. The initialDelaySeconds field tells the kubelet that it should wait 5 seconds before performing the first probe. To perform a probe, the kubelet executes the command `cat /tmp/healthy` in the target container. If the command succeeds, it returns 0, and the kubelet considers the container to be alive and healthy. If the command returns a non-zero value, the kubelet kills the container and restarts it.

When the container starts, it executes this command:

```
/bin/sh -c "touch /tmp/healthy; sleep 30; rm -rf /tmp/healthy; sleep 600"
```

For the first 30 seconds of the container's life, there is a `/tmp/healthy` file. So during the first 30 seconds, the command `cat /tmp/healthy` returns a success code. After 30 seconds, `cat /tmp/healthy` returns a failure code.

Create the Pod:

```
kubectl apply -f https://k8s.io/examples/pods/probe/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 "k8s.gcr.io/busybox"  
23s 23s 1 {kubelet worker0} spec.containers{liveness} Normal Pulled Successfully pulled image  
"k8s.gcr.io/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
```

