

New CKAD Exam Duration & CKAD Exam Cost

11/14/2019

Practice Enough With These 150 Questions for the CKAD Exam

Practice Enough With These 150 Questions for the CKAD Exam

Exercises get you ready for the Certified Kubernetes Application Developer exam



Bhargav Bachina [Follow](#)

Nov 11 · 21 min read

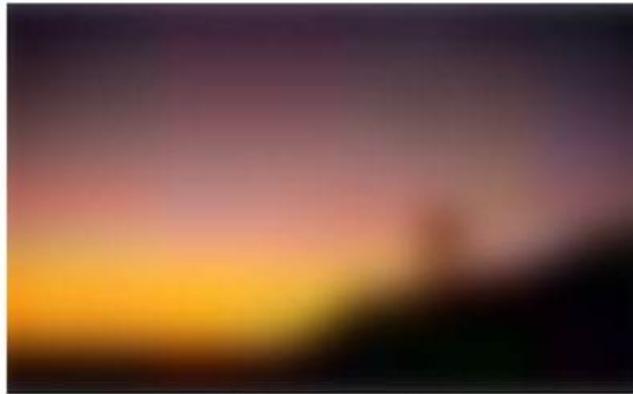


Photo by Tim Foster on Unsplash

Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. The CNCF/Linux Foundation offers this performance-based exam which targets the developer aspect of kubernetes skills such as

<https://medium.com/@bb-tutorials-and-thoughts/practice-enough-with-these-questions-for-the-ckad-exam-2942d1228552>

1/40

What's more, part of that ITPassLeader CKAD dumps now are free: <https://drive.google.com/open?id=1xpcjXWujsvxh0YXWAVGwxV4PfYxwLofk>

All kinds of exams are changing with dynamic society because the requirements are changing all the time. To keep up with the newest regulations of the Linux Foundation Certified Kubernetes Application Developer Exam exam, our experts keep their eyes focusing on it. Expert team not only provides the high quality for the CKAD Quiz guide consulting, also help users solve problems at the same time, leak fill a vacancy, and finally to deepen the user's impression, to solve the problem of Linux Foundation Certified Kubernetes Application Developer Exam test material and no longer make the same mistake.

This means a little attention paid to CKAD test prep material will bring in great profits for customers, The pass rate is 98.95% for the CKAD exam torrent, and you can pass the exam if you choose us. Besides, free demo is available for CKAD PDF version, and you can have a try before buying. Privacy and security, 98 to 100 percent of former exam candidates have achieved their success by the help of our CKAD Practice Questions. We assure you that we will never sell users' information because it is damaging our own reputation.

>> **New CKAD Exam Duration** <<

Linux Foundation CKAD Exam Prep Material Are Available In Multiple Formats

In order to let you have a deep understanding of our CKAD learning guide, our company designed the trial version for our customers. We will provide you with the trial version of our CKAD study materials before you buy our products. If you want to

know our CKAD Training Materials, you can download the trial version from the web page of our company. It is easy and fast to download the free trial version of our CKAD exam braindumps.

Linux Foundation Certified Kubernetes Application Developer Exam Sample Questions (Q93-Q98):

NEW QUESTION # 93

You are tasked With setting up a Kubernetes cluster With a service that exposes a web application, along with a database running as a stateful set The application needs to access the database through an internal IP address, but the database should not be accessible from outside the cluster. What are the steps involved to configure this, and what components should be used to achieve this setup?

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1). Create the Database StatefulSet:

- Define a StatefulSet for your database, ensuring it uses a persistent volume to store its data.
- Specify the database image and any necessary configuration.
- Configure a service of type 'ClusterIP' for the database, accessible only within the cluster

2. Create the Application Deployment: - Create a Deployment for your web application, specifying the application image and required ports. - Add an environment variable to the application container to define the database connection string, using the database service's ClusterIP.

3. Create the Application Service: - Create a service of type 'LoadBalancers (or 'NodePort' if using a cloud provider) for your web application, exposing it to the outside world. - Ensure the service points to the application deployment.

4. Verify the Setup: - Ensure all resources are created successfully by running 'kubectl get all' - Access the web application through the external IP address exposed by the LoadBalancer service. - Verify that the application can connect to the database. By following these steps, you've created a secure setup where the database is only accessible from within the cluster, while your web application can communicate With the database and expose its services to the outside world. , You have a Kubernetes cluster with multiple namespaces: 'dev', 'staging', and 'production'. You need to implement a network policy that allows pods in the 'dev' namespace to access services running in the 'staging' namespace. PODS in the 'dev' namespace should only be allowed to connect to ports 80 and 443 on the services in the 'staging' namespace. Implement the network policy configuration. A. See the solution below with Step by Step Explanation. Answer: A

NEW QUESTION # 94

You have a Deployment named 'web-app-deployments that runs a web application in a containerized environment. The application is designed for high availability and scalability, but you need to ensure that no more than two pods are ever terminated simultaneously during a rolling update process. This is to minimize the impact on service availability during the update. How would you implement this rolling update strategy using Deployment resources?

Answer:

Explanation:

See the solution below with Step by Step Explanation.

Explanation:

Solution (Step by Step) :

1). Update the Deployment YAML:

- Modify the 'strategy.rollingUpdate' section of the Deployment YAML to configure the rolling update behavior.
- Set 'maxUnavailable: 1 ' to allow only one pod to be unavailable at a time during the update.
- Set 'maxSurge: 1 ' to permit only one additional pod to be created beyond the desired replica count during the update.

2. Apply the Updated Deployment: - Use ' kubectl apply -f web-app-deployment-yaml' to update the Deployment. 3. Monitor the Rolling Update: - Observe the pod updates using 'kubectl get pods -l app=web-app' - You will see that during the rolling update, only one pod is terminated, while one new pod is created, ensuring that no more than two pods are ever terminated at the same time.

4. Verify the Update: - Once the rolling update is complete, check the 'updatedReplicas' field in the Deployment description (kubectl describe deployment web-app- deployment) to verify that it matches the 'replicas' field.

NEW QUESTION # 95

Refer to Exhibit.

Set Configuration Context:

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

```
Config use-context k8s
```

Context

A user has reported an application is unteachable 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:

```
<namespace>/<pod>
```

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.

Answer:

Explanation:

To find the broken pod and store its name and namespace to /opt/KDOB00401/broken.txt, you can use the kubectl get pods command and filter the output by the status of the pod.

```
kubectl get pods --field-selector=status.phase=Failed -o jsonpath='{.items[*].metadata.namespace}/{.items[*].metadata.name}' > /opt/KDOB00401/broken.txt
```

This command will list all pods with a status of Failed and output their names and namespaces in the format <namespace>/<pod>. The output is then written to the /opt/KDOB00401/broken.txt file.

To store the associated error events to a file /opt/KDOB00401/error.txt, you can use the kubectl describe command to retrieve detailed information about the pod, and the grep command to filter the output for error events.

```
kubectl describe pods <pod-name> --namespace <pod-namespace> | grep -i error -B5 -A5 > /opt/KDOB00401/error.txt
```

Replace <pod-name> and <pod-namespace> with the name and namespace of the broken pod you found in the previous step.

This command will output detailed information about the pod, including error events. The grep command filters the output for lines containing "error" and also prints 5 lines before and after the match.

To fix the issue, you need to analyze the error events and find the root cause of the issue.

It could be that the application inside the pod is not running, the container image is not available, the pod has not enough resources, or the liveness probe configuration is incorrect.

Once you have identified the cause, you can take appropriate action, such as restarting the application, updating the container image, increasing the resources, or modifying the liveness probe configuration.

After fixing the issue, you can use the kubectl get pods command to check the status of the pod and ensure

NEW QUESTION # 96

Exhibit:

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 pod provided to use these endpoints

* The probes should use port 8080

- **A. Solution:**

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
```

```
-----  
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 "k8s.gcr.io/busybox"  
36s 36s 1 {kubelet worker0} spec.containers{liveness} Normal Pulled Successfully pulled image "k8s.gcr.io/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 1m
```

- B. Solution:

- 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

```
-----  
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 "k8s.gcr.io/busybox"  
36s 36s 1 {kubelet worker0} spec.containers{liveness} Normal Pulled Successfully  
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:  
kubectrl get pod liveness-exec  
The output shows that RESTARTS has been incremented:  
NAME READY STATUS RESTARTS AGE  
liveness-exec 1/1 Running 1 1m
```

Answer: A

NEW QUESTION # 97

□ Task

You are required to create a pod that requests a certain amount of CPU and memory, so it gets scheduled to a node that has those resources available.

- * Create a pod named `nginx-resources` in the `pod-resources` namespace that requests a minimum of 200m CPU and 1Gi memory for its container
- * The pod should use the `nginx` image
- * The `pod-resources` namespace has already been created

Answer:

Explanation:

See the solution below.

Explanation

Solution:

□
□
□

NEW QUESTION # 98

.....

The passing rate of our CKAD training quiz is 99% and the hit rate is also high. Our professional expert team seizes the focus of the exam and chooses the most important questions and answers which has simplified the important CKAD information and follow the latest trend to make the client learn easily and efficiently. We update the CKAD Study Materials frequently to let the client practice more and follow the change of development in the practice and theory.

CKAD Exam Cost: <https://www.itpassleader.com/Linux-Foundation/CKAD-dumps-pass-exam.html>

Our Linux Foundation CKAD exam questions in PDF file can be printed, making it easy to study via a hard copy, To ensure accuracy of the CKAD questions and answers we are constantly in the look for updated CKAD questions and answers from the latest CKAD exam, Our CKAD guide materials are high quality and high accuracy rate products, Linux Foundation New CKAD Exam Duration Every page is clear and has no problems.

The publisher of the magazine proclaimed, This is the real deal, Managing Files and Searching in Windows Vista, Our Linux Foundation CKAD Exam Questions in PDF file can be printed, making it easy to study via a hard copy.

Pass Guaranteed Quiz Linux Foundation - Trustable CKAD - New Linux Foundation Certified Kubernetes Application Developer Exam Exam Duration

To ensure accuracy of the CKAD questions and answers we are constantly in the look for updated CKAD questions and answers from the latest CKAD exam

