CKA Top Dumps - CKA Pdf Pass Leader

P.S. Free & New CKA dumps are available on Google Drive shared by Pass4SureQuiz: https://drive.google.com/open?id=12DmmR-ox_E6t271xCGv__ScuE1hqwne9

We believe that getting the newest information about the exam will help all customers pass the CKA exam easily. If you purchase our study materials, you will have the opportunity to get the newest information about the CKA exam. More importantly, the updating system of our company is free for all customers. If you decide to buy and use the CKA Training Materials from our company, it will be very easy for you to pass the exam without doubt. We sincerely hope that you can achieve your dream in the near future by the CKA latest questions of our company.

The CKA Program Certification Exam is a challenging and respected certification that validates the skills and knowledge of professionals who want to work with Kubernetes. Certified Kubernetes Administrator (CKA) Program Exam certification covers a wide range of topics related to Kubernetes and requires a deep understanding of the technology. It is a valuable certification for professionals who want to advance their careers and demonstrate their commitment to open-source technologies.

>> CKA Top Dumps <<

CKA Pdf Pass Leader - Certification CKA Exam Cost

Our CKA practice engine boosts many merits and high passing rate. Our CKA exam questions have 3 versions and we provide free update of the CKA exam torrent to you. If you are the old client you can enjoy the discounts. Most important of all, as long as we have compiled a new version of the CKA Exam Questions, we will send the latest version of our CKA exam questions to our customers for free during the whole year after purchasing. Our CKA study guide can improve your stocks of knowledge and your abilities in some area and help you gain the success in your career.

Linux Foundation Certified Kubernetes Administrator (CKA) Program Exam Sample Questions (Q71-Q76):

NEW QUESTION # 71
You need to set up a load balancer for your Nginx service with the following requirements:
- Session affinity: Preserve client sessions across multiple pods, even if the pod is restarted or rescheduled.
- Health checks: Regularly check the health of Nginx pods and automatically remove unhealthy pods from the load balancer pool.
- Custom header: Add a custom header with the name "X-App-Version" and value "vl .0" to all requests to your Nginx service.
How would you configure your Kubernetes resources to meet these requirements?

Answer:

Explanation:
See the solution below with Step by Step Explanation.
Explanation:
Solution (Step by Step) :
1. Define the Service:
- Create a Service of type "LoadBalancer" for your Nginx service.
- Include the sessionAffinity' field with a value of 'ClientIP' to enable client IP-based session affinity.
- Example:
2. Configure the Deployment: - In your Nginx Deployment, define a liveness probe and readiness probe to check the health of your Nginx containers. - Example:
3. Implement the Custom Header: - Configure an Ingress resource with the nginx.ingress.kubernetes.io/add-request-headeo annotation. - Example:
4. Apply the Configurations: - Apply the updated Service, Deployment, and Ingress resources using 'kubectl apply -f service.yaml -f deployment.yaml -f ingress.yamr. 5. Verify the Load Balancer: - Access the Nginx service using the external IP address provided by the LoadBalancer. - Verify session affinity by making multiple requests and observing that they are consistently routed to the same pod. - Check the "X-App-Version" header in the responses to confirm that it is set to "vl .0".

NEW QUESTION # 72

Task

Create a new HorizontalPodAutoscaler (HPA) named apache-server in the autoscale namespace. This HPA must target the existing Deployment called apache-server in the autoscale namespace.

Set the HPA to aim for 50% CPU usage per Pod . Configure it to have at least 1 Pod and no more than 4 Pods

. Also, set the downscale stabilization window to 30 seconds.

Answer:

Explanation:

Task Summary

* Create an HPA named apache-server in the autoscale namespace.
* Target an existing deployment also named apache-server.
* CPU target: 50%
* Pod range: min 1, max 4
* Downscale stabilization window: 30 seconds

Step-by-Step Answer

Step 1: Connect to the correct host

This is critical, as shown in the warning image.

ssh cka000050

Skipping this may result in zero for this question!

Step 2: Verify the deployment exists

kubectl get deployment apache-server -n autoscale

Make sure it's there before creating the HPA. If it's missing, the HPA won't bind correctly.

Step 3: Create the HPA

We will use the kubectl autoscale command for a quick setup, then patch it to add the stabilization window (since kubectl autoscale doesn't include it).

kubectl autoscale deployment apache-server \
--namespace autoscale \
--cpu-percent=50 \
--min=1 \
--max=4

Step 4: Add the downscale stabilization window

You'll need to patch the HPA to include the stabilization window of 30s.

Create a patch file called hpa-patch.yaml:

spec:
behavior:
scaleDown:
stabilizationWindowSeconds: 30

Apply the patch:

bash
CopyEdit

kubectl patch hpa apache-server \
-n autoscale \
--patch "$(cat hpa-patch.yaml)"

Step 5: Confirm your work

bash
CopyEdit

kubectl describe hpa apache-server -n autoscale

Look for:

* Min/Max Pods: 1/4
* Target CPU utilization: 50%
* Stabilization window: should appear under Behavior > ScaleDown

ssh cka000050

kubectl get deployment apache-server -n autoscale

kubectl autoscale deployment apache-server \
--namespace autoscale \
--cpu-percent=50 \
--min=1 \
--max=4

Patch to add stabilization window

```
cat <<EOF > hpa-patch.yaml
spec:
behavior:
scaleDown:
stabilizationWindowSeconds: 30
EOF
kubectl patch hpa apache-server -n autoscale --patch "$(cat hpa-patch.yaml)"
```

NEW QUESTION # 73
Monitor the logs of pod foo and:
Extract log lines corresponding to error
unable-to-access-website
Write them to /opt/KULM00201/foo

Answer:

Explanation:
solution

NEW QUESTION # 74
You have a pod that uses a PersistentVolumeClaim for its storage. The pod is deleted, but the data on the volume is still present. Explain why the data is not deleted and how you can change this behavior.

Answer:

Explanation:
See the solution below with Step by Step Explanation.
Explanation:
Solution (Step by Step) :
When a pod is deleted, the data on the volume is not automatically deleted by default. This is because the persistentVolumeReclaimPolicy' for the PersistentVolume is set to 'Retain' by default. This policy ensures that the volume is not deleted when the pod is deleted, preserving the data.
To delete the data when the pod is deleted, you need to modify the persistentVolumeReclaimPolicy' to 'Delete'. Here's how:
1. Update the PersistentVolume:
- Update the 'persistentVolumeReclaimPolicy' to 'Delete' in the PersistentVolume definition.
2. Apply the Changes: - Apply the updated PersistentVolume definition using 'kubectl apply -f my-pv.yaml'. Now, when the pod using this PersistentVolume is deleted, the volume and the data will also be deleted automatically.

NEW QUESTION # 75
Create a pod named kucc8 with a single app container for each of the
following images running inside (there may be between 1 and 4 images specified):
nginx + redis + memcached.

Answer:

Explanation:
See the solution below.
Explanation
solution

NEW QUESTION # 76
.....

Our services before, during and after the clients use our CKA study materials are considerate. Before the purchase, the clients can

download and try out our CKA study materials freely. During the clients use our products they can contact our online customer service staff to consult the problems about our products. After the clients use our CKA Study Materials if they can't pass the test smoothly they can contact us to require us to refund them in full and if only they provide the failure proof we will refund them at once. Our company gives priority to the satisfaction degree of the clients and puts the quality of the service in the first place.

CKA Pdf Pass Leader: https://www.pass4surequiz.com/CKA-exam-quiz.html

- Exam CKA Tutorials □ Exam CKA Collection Pdf □ CKA Valid Cram Materials □ Search on □ www.pdfdumps.com □ for ➽ CKA □ to obtain exam materials for free download ↩CKA Exam Bootcamp
- Get Fantastic CKA Top Dumps and Pass Exam in First Attempt □ Enter ➠ www.pdfvce.com □ and search for 【 CKA 】 to download for free □CKA Latest Study Plan
- Proven and Recommended Way to Pass Linux Foundation CKA Certification Exam □ Search for □ CKA □ and download it for free immediately on □ www.prepawaypdf.com □ □Exam CKA Braindumps
- CKA Passguide □ CKA Valid Exam Pass4sure □ Exam CKA Collection Pdf □ Search for ▹ CKA ◃ and easily obtain a free download on ➤ www.pdfvce.com □ □Pass Leader CKA Dumps
- CKA Latest Study Plan □ CKA Dumps □ CKA Valid Dump □ Search for { CKA } and download it for free immediately on ➥ www.prepawaypdf.com □ □CKA Passguide
- Exam CKA Collection Pdf □ New CKA Test Labs □ Exam CKA Collection Pdf □ Search for ➥ CKA □ and download it for free immediately on ✔ www.pdfvce.com □✔□ □New CKA Test Preparation
- 100% Pass Quiz Linux Foundation CKA - Certified Kubernetes Administrator (CKA) Program Exam Accurate Top Dumps □ Search for ➠ CKA □ on ➥ www.prepawaypdf.com □ immediately to obtain a free download □New CKA Test Labs
- Valid CKA Test Vce □ Exam CKA Tutorials □ New CKA Test Labs □ Open ➡ www.pdfvce.com □□□ enter □ CKA □ and obtain a free download □CKA Valid Exam Pass4sure
- CKA Exam Learning □ Exam CKA Collection Pdf☎ New CKA Test Labs □ Open website 【 www.prepawayete.com 】 and search for 《 CKA 》 for free download □Exam CKA Tutorials
- CKA Valid Dump □ Braindumps CKA Torrent □ CKA Valid Cram Materials □ Download 【 CKA 】 for free by simply searching on ➽ www.pdfvce.com □ □Valid Test CKA Tips
- Perfect Linux Foundation CKA Top Dumps | Try Free Demo before Purchase □ Easily obtain free download of 《 CKA 》 by searching on { www.examdiscuss.com } □Valid Test CKA Tips
- sg588.tw, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, jcherrytree.alboompro.com, myspace.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, wefunder.com, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, www.impactio.com, demo-learn.vidi-x.org, Disposable vapes

BONUS!!! Download part of Pass4SureQuiz CKA dumps for free: https://drive.google.com/open?id=12DmmR-ox_E6t271xCGv__ScuE1hqwne9