KCNA Q&As

Kubernetes and Cloud Native Associate (KCNA)

Pass Linux Foundation KCNA Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/kcna.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Linux Foundation Official Exam Center

Instant Download After Purchase
100% Money Back Guarantee
365 Days Free Update
800,000+ Satisfied Customers

QUESTION 1

Which of the following best describes a cloud-native app?

A. An application where all logic is coded into a single large binary.

B. An application that publishes an HTTPS web front-end.

C. An application that takes advantages of cloud computing fromworks and their loosely coupled cloud services.

D. An application that leverages services that are native to public cloud platforms such as Azure, GCP, and/or AWS.

Correct Answer: C

Explanation: Cloud-native apps leverage cloud computing frameworks and tend to be microservices based, where individual components of the app are coded as individual.

QUESTION 2

What is etcd used for in Kubernetes?

A. Integration with cloud platforms

B. Network routing for the cluster

C. Kubernetes API security

D. Backend object storage for the Kubernetes API

Correct Answer: D

Explanation: etcd serves as a distributed object store that backs the Kubernetes API.

QUESTION 3

The 4C\\'s of Cloud Native security

A. Chroot, Compute, Cluster and Container

B. Cluster, Cloud, Compute, and Containers

C. Code, Containers, Compute, and Cloud

D. Cloud, Clusters, Containers, and Code

Correct Answer: D

Explanation: https://kubernetes.io/docs/concepts/security/overview/

QUESTION 4

How to get the logs of the previously terminated nginx container from the web pod?

A. kubectl logs -p -c nginx web

B. kubectl logs nginx

C. kubectl logs -p -c web nginx

D. kubectl logs -f -c nginx web

Correct Answer: A

Explanation: https://kubernetes.io/docs/reference/generated/kubectl/kubectl- commands#logs

QUESTION 5

What is horizontal scaling?

A. Creating a Deployment

B. Adding resources to existing apps and servers

C. Moving workloads from one server to another

D. Adding additional replicas of apps and servers

Correct Answer: D

Explanation: https://kubernetes.io/docs/tasks/run-application/horizontal-pod-autoscale/

In Kubernetes, a *HorizontalPodAutoscaler* automatically updates a workload resource (such as a Deployment or StatefulSet), with the aim of automatically scaling the workload to match demand.

Horizontal scaling means that the response to increased load is to deploy more Pods. This is different from *vertical* scaling, which for Kubernetes would mean assigning more resources (for example: memory or CPU) to the Pods that are already running for the workload.

If the load decreases, and the number of Pods is above the configured minimum, the HorizontalPodAutoscaler instructs the workload resource (the Deployment, StatefulSet, or other similar resource) to scale back down.

Horizontal pod autoscaling does not apply to objects that can't be scaled (for example: a DaemonSet.)

The HorizontalPodAutoscaler is implemented as a Kubernetes API resource and a controller. The resource determines the behavior of the controller. The horizontal pod autoscaling controller, running within the Kubernetes control plane, periodically adjusts the desired scale of its target (for example, a Deployment) to match observed metrics such as average CPU utilization, average memory utilization, or any other custom metric you specify.

KCNA Practice Test KCNA Exam Questions KCNA Braindumps