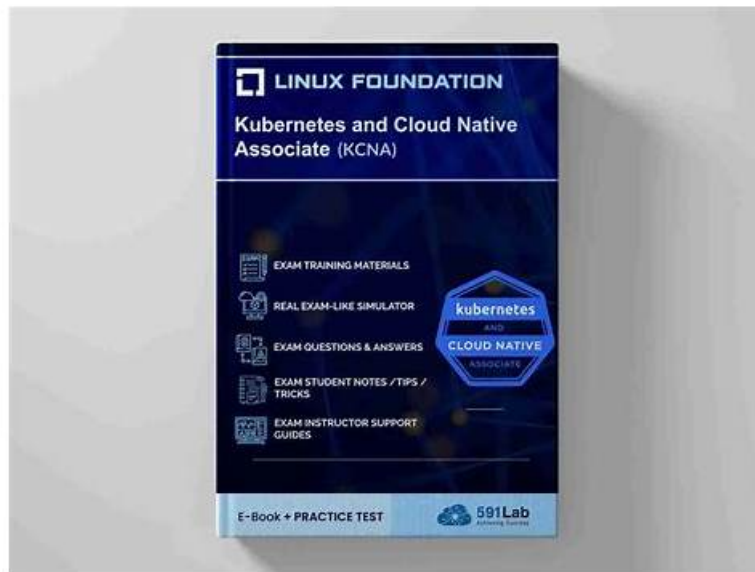


Quiz 2026 Linux Foundation KCNA: Kubernetes and Cloud Native Associate Accurate Latest Dumps Free



P.S. Free & New KCNA dumps are available on Google Drive shared by ITExamSimulator: https://drive.google.com/open?id=1_pqjDCVXrtlzmY6AQXILVVyUEIuDFwt

The passing rate of our KCNA exam materials are very high and about 99% and so usually the client will pass the exam successfully. But in case the client fails in the exam unfortunately we will refund the client immediately in full at one time. The refund procedures are very simple if you provide the KCNA exam proof of the failure marks we will refund you immediately. Clients always wish that they can get immediate use after they buy our KCNA Test Questions because their time to get prepared for the exam is limited. Our KCNA test torrent won't let the client wait for too much time and the client will receive the mails in 5-10 minutes sent by our system. Then the client can log in and use our software to learn immediately. It saves the client's time.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Exam is a certification program designed to test one's knowledge and skills in the field of Kubernetes and Cloud Native technologies. Kubernetes and Cloud Native Associate certification is globally recognized and is a valuable asset for professionals looking to advance their careers in the cloud computing industry. KCNA exam covers a wide range of topics including Kubernetes architecture, deployment, networking, security, and troubleshooting. It also covers cloud-native technologies such as containerization, microservices, and serverless computing.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Certification Exam is a globally recognized certification program that validates one's knowledge and skills in Kubernetes and cloud-native technologies. Kubernetes is an open-source container orchestration system that automates the deployment, scaling, and management of containerized applications. Cloud-native technologies, on the other hand, refer to a set of tools and practices that enable the development and deployment of applications in a cloud-native environment.

The KCNA Exam is a vendor-neutral certification, meaning that it is not tied to any specific cloud provider or technology. This makes it an attractive option for professionals who work with different cloud platforms and want to showcase their skills in a way that is recognized across the industry. KCNA exam covers a broad range of topics, including Kubernetes architecture, deployment, networking, and security.

>> **KCNA Latest Dumps Free** <<

KCNA Valid Exam Questions & KCNA Online Test

By seeing your goofs you can work on your show continually for the Linux Foundation KCNA approach. You can give vast phony tests to make them ideal for Linux Foundation KCNA and can check their past given exams. Linux Foundation KCNA Dumps will give reliable free updates to our clients generally all the Kubernetes and Cloud Native Associate.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions

(Q88-Q93):

NEW QUESTION # 88

Which mechanism can be used to automatically adjust the amount of resources for an application?

- A. Kubernetes Event-driven Autoscaling (KEDA)
- B. Vertical Pod Autoscaler (VPA)
- C. Horizontal Pod Autoscaler (HPA)
- D. Cluster Autoscaler

Answer: C

Explanation:

The verified answer in the PDF is A (HPA), and that aligns with the common Kubernetes meaning of "adjust resources for an application" by scaling replicas. The Horizontal Pod Autoscaler automatically changes the number of Pod replicas for a workload (typically a Deployment) based on observed metrics such as CPU utilization, memory (in some configurations), or custom/external metrics. By increasing replicas under load, the application gains more total CPU/memory capacity available across Pods; by decreasing replicas when load drops, it reduces resource consumption and cost.

It's important to distinguish what each mechanism adjusts:

* HPA adjusts replica count (horizontal scaling).

* VPA adjusts Pod resource requests/limits (vertical scaling), which is literally "amount of CPU /memory per pod," but it often requires restarts to apply changes depending on mode.

* Cluster Autoscaler adjusts the number of nodes in the cluster, not application replicas.

* KEDA is event-driven autoscaling that often drives HPA behavior using external event sources (queues, streams), but it's not the primary built-in mechanism referenced in many foundational Kubernetes questions.

Given the wording and the provided answer key, the intended interpretation is: "automatically adjust the resources available to the application" by scaling out/in the number of replicas. That's exactly HPA's role.

For example, if CPU utilization exceeds a target (say 60%), HPA computes a higher desired replica count and updates the workload. The Deployment then creates more Pods, distributing load and increasing available compute.

So, within this question set, the verified correct choice is A (Horizontal Pod Autoscaler).

NEW QUESTION # 89

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

Answer: D

Explanation:

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

NEW QUESTION # 90

You want to configure a CI/CD pipeline to deploy a microservice to Kubernetes. Which of the following steps are essential in the pipeline?

- A. Build the container image
- B. Push the container image to a registry
- C. Monitor the deployed application
- D. Apply Kubernetes configuration files (e.g., Deployment, Service)
- E. Run automated tests

Answer: A,B,D,E

Explanation:

All the listed steps are essential for a robust CI/CD pipeline. Building the container image encapsulates your application, pushing it to

a registry ensures easy access, applying Kubernetes configurations defines the deployment, and automated tests validate the application's functionality before deployment.

NEW QUESTION # 91

How is application data maintained in containers?

- A. Store data in separate folders.
- **B. Store data into volumes.**
- C. Store data into sidecar containers.
- D. Store data into data folders.

Answer: B

Explanation:

Container filesystems are ephemeral: the writable layer is tied to the container lifecycle and can be lost when containers are recreated. Therefore, maintaining application data correctly means storing it in volumes, making D the correct answer. In Kubernetes, volumes provide durable or shareable storage that is mounted into containers at specific paths. Depending on the volume type, the data can persist across container restarts and even Pod rescheduling.

Kubernetes supports many volume patterns. For transient scratch data you might use `emptyDir` (ephemeral for the Pod's lifetime). For durable state, you typically use `PersistentVolumes` consumed by `PersistentVolumeClaims` (PVCs), backed by storage systems via CSI drivers (cloud disks, SAN/NAS, distributed storage). This decouples the application container image from its state and enables rolling updates, rescheduling, and scaling without losing data.

Options A and B ("folders") are incomplete because folders inside the container filesystem do not guarantee persistence. A folder is only as durable as the underlying storage; without a mounted volume, it lives in the container's writable layer and will disappear when the container is replaced. Option C is incorrect because "sidecar containers" are not a data durability mechanism; sidecars can help ship logs or sync data, but persistent data should still be stored on volumes (or external services like managed databases).

From an application delivery standpoint, the principle is: containers should be immutable and disposable, and state should be externalized. Volumes (and external managed services) make this possible. In Kubernetes, this is a foundational pattern enabling safe rollouts, self-healing, and portability: the platform can kill and recreate Pods freely because data is maintained independently via volumes.

Therefore, the verified correct choice is D: Store data into volumes.

NEW QUESTION # 92

In CNCF, who develops specifications for industry standards around container formats and runtimes?

- A. Container Network Interface (CNI)
- B. Linux Foundation Certification Group (LFCG)
- C. Container Runtime Interface (CRI)
- **D. Open Container Initiative (OCI)**

Answer: D

Explanation:

The organization responsible for defining widely adopted standards around container formats and runtime specifications is the Open Container Initiative (OCI), so A is correct. OCI defines the image specification (how container images are structured and stored) and the runtime specification (how to run a container), enabling interoperability across tooling and vendors. This is foundational to the cloud-native ecosystem because it allows different build tools, registries, runtimes, and orchestration platforms to work together reliably.

Within Kubernetes and CNCF-adjacent ecosystems, OCI standards are the reason an image built by one tool can be pushed to a registry and pulled/run by many different runtimes. For example, a Kubernetes node running `containerd` or `CRI-O` can run OCI-compliant images consistently. OCI standardization reduces fragmentation and vendor lock-in, which is a core motivation in open source cloud-native architecture.

The other options are not correct for this question. CNI (Container Network Interface) is a standard for configuring container networking, not container image formats and runtimes. CRI (Container Runtime Interface) is a Kubernetes-specific interface between kubelet and container runtimes—it enables pluggable runtimes for Kubernetes, but it is not the industry standard body for container format/runtime specifications.

"LFCG" is not the recognized standards body here.

In short: OCI defines the "language" for container images and runtime behavior, which is why the same image can be executed across environments. Kubernetes relies on those standards indirectly through runtimes and tooling, but the specification work is

