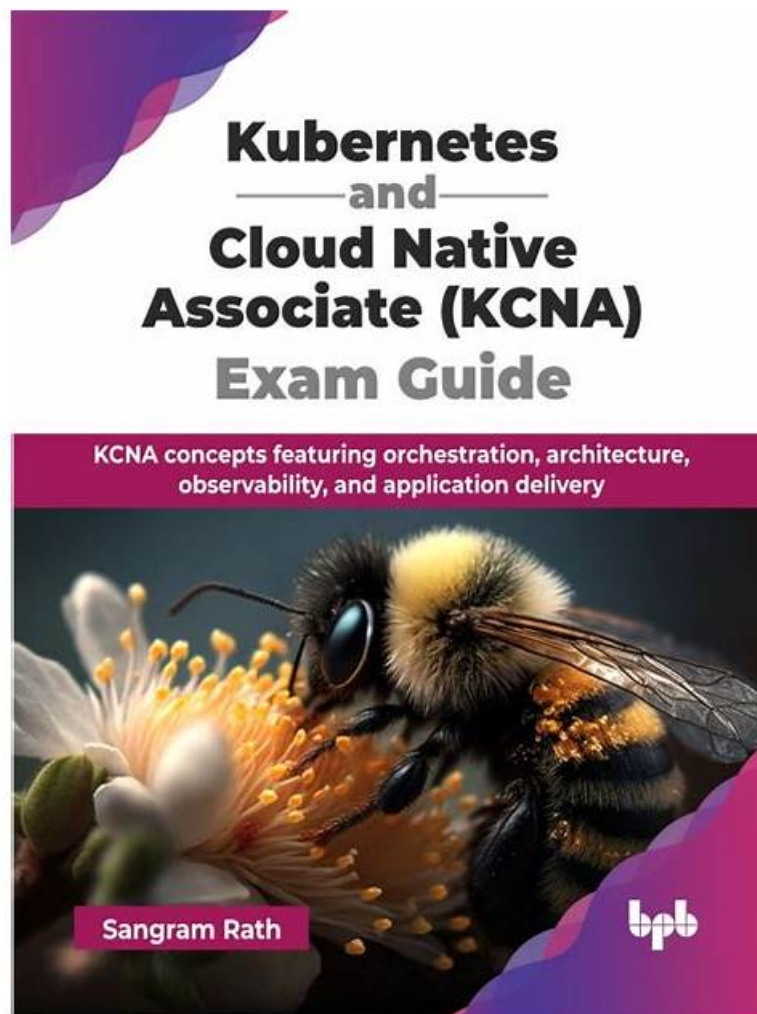


KCNA Exam Learning | KCNA Exam Book



P.S. Free 2026 Linux Foundation KCNA dumps are available on Google Drive shared by Braindumpsqa:
<https://drive.google.com/open?id=13uazR6BfOXbyfg3qylRLL1wnKXh3Dn8>

The Linux Foundation desktop practice test software and web-based Understanding Kubernetes and Cloud Native Associate KCNA practice test both simulate the actual exam environment and identify your mistakes. With these two Linux Foundation KCNA practice exams, you will get the actual KCNA Exam environment. Whereas the Braindumpsqa PDF file is ideal for restriction-free test preparation. You can open this PDF file and revise KCNA real exam questions at any time.

The KCNA Exam Tests candidates' knowledge of Kubernetes, the leading container orchestration system, and related cloud-native tools such as Prometheus, Fluentd, and Istio. KCNA exam also covers topics such as containerization, microservices architecture, and cloud-native infrastructure. Candidates who pass the exam will be certified as Kubernetes and Cloud Native Associates, demonstrating their proficiency in cloud-native technologies and their ability to work on cloud-native projects.

>> **KCNA Exam Learning** <<

100% Pass Quiz High Hit-Rate Linux Foundation - KCNA - Kubernetes and Cloud Native Associate Exam Learning

Braindumpsqa has been going through all ups and downs tested by the market, and now our KCNA exam questions have become perfectly professional. We never circumvent the difficulties of our KCNA study materials happened on the road as long as there is bright at the end, and it is the satisfactory results you want. And we have helped so many of our customers achieve their certifications according to our KCNA learning guide.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q197-Q202):

NEW QUESTION # 197

Which persona is normally responsible for defining, testing, and running an incident management process?

- A. Project Managers
- B. Quality Engineers
- C. Site Reliability Engineers
- D. Application Developers

Answer: C

Explanation:

The role most commonly responsible for defining, testing, and running an incident management process is Site Reliability Engineers (SREs), so A is correct. SRE is an operational engineering discipline focused on ensuring reliability, availability, and performance of services in production. Incident management is a core part of that mission: when outages or severe degradations occur, someone must coordinate response, restore service quickly, and then drive follow-up improvements to prevent recurrence.

In cloud native environments (including Kubernetes), incident response involves both technical and process elements. On the technical side, SREs ensure observability is in place-metrics, logs, traces, dashboards, and actionable alerts-so incidents can be detected and diagnosed quickly. They also validate operational readiness: runbooks, escalation paths, on-call rotations, and post-incident review practices. On the process side, SREs often establish severity classifications, response roles (incident commander, communications lead, subject matter experts), and "game day" exercises or simulated incidents to test preparedness.

Project managers may help coordinate schedules and communication for projects, but they are not typically the owners of operational incident response mechanics. Application developers are crucial participants during incidents, especially for debugging application-level failures, but they are not usually the primary maintainers of the incident management framework. Quality engineers focus on testing and quality assurance, and while they contribute to preventing defects, they are not usually the owners of real-time incident operations.

In Kubernetes specifically, incidents often span multiple layers: workload behavior, cluster resources, networking, storage, and platform dependencies. SREs are positioned to manage the cross-cutting operational view and to continuously improve reliability through error budgets, SLOs/SLIs, and iterative hardening. That's why the correct persona is Site Reliability Engineers.

NEW QUESTION # 198

You're managing a Kubernetes cluster with several deployments using resource requests and limits. Which of the following strategies can help you effectively manage resource utilization and potentially reduce costs?

- A. Use Horizontal Pod Autoscaler (HPA) to adjust the number of pods based on resource utilization and load.
- B. Increase resource requests for all deployments to ensure enough resources are always available.
- C. Configure resource reservation to ensure a minimum amount of resources is always available.
- D. Set resource limits lower than requests to encourage efficient resource use.
- E. Utilize resource quotas to restrict the total resources that can be consumed by namespaces.

Answer: A,C,E

Explanation:

The most effective strategies for managing resource utilization and cost in Kubernetes involve dynamic scaling and resource allocation. Horizontal Pod Autoscaler (HPA) allows you to adjust the number of pods based on resource utilization, reducing overprovisioning. Resource quotas restrict the total resources that can be consumed by namespaces, preventing resource exhaustion and potential cost spikes. Resource reservation guarantees a minimum amount of resources for critical applications, ensuring they perform well even under high load. Increasing requests without adjusting limits can lead to overprovisioning, while setting limits lower than requests can limit pod performance.

NEW QUESTION # 199

Your organization is migrating a monolithic application to a serverless architecture on Google Cloud Platform. Which Google Cloud service would be the most suitable for running your application's backend logic?

- A. Google Cloud Storage

- **B. Google Cloud Functions**
- C. Google Cloud SQL
- D. Google Kubernetes Engine (GKE)
- E. Google App Engine

Answer: B

Explanation:

Google Cloud Functions is the ideal choice for running backend logic in a serverless environment on Google Cloud. It offers a fully managed serverless execution environment for event-driven code. App Engine (A) is more suited for web applications. GKE (B) is a container orchestration service. Cloud Storage (D) is for data storage, and Cloud SQL (E) is for relational databases.

NEW QUESTION # 200

You are working on a cloud-native application that needs to interact with a database service. The application requires high availability and fault tolerance. Which open standard would you use to ensure reliable database connections and handle failures gracefully?

- A. OpenTelemetry
- B. CloudEvents
- C. OpenTracing
- D. Kubernetes Ingress
- **E. Service Mesh**

Answer: E

Explanation:

A Service Mesh is a crucial element in achieving high availability and fault tolerance for database connections. It provides features like load balancing, circuit breakers, and retries, which are essential for handling failures and ensuring the application remains operational even if a database instance becomes unavailable.

NEW QUESTION # 201

What is CRD?

- **A. Custom Resource Definition**
- B. Customized RUST Definition
- C. Custom Restricted Definition
- D. Custom RUST Definition

Answer: A

Explanation:

A CRD is a CustomResourceDefinition, making A correct. Kubernetes is built around an API-driven model: resources like Pods, Services, and Deployments are all objects served by the Kubernetes API. CRDs allow you to extend the Kubernetes API by defining your own resource types. Once a CRD is installed, the API server can store and serve custom objects (Custom Resources) of that new type, and Kubernetes tooling (kubectl, RBAC, admission, watch mechanisms) can interact with them just like built-in resources.

CRDs are a core building block of the Kubernetes ecosystem because they enable operators and platform extensions. A typical pattern is: define a CRD that represents the desired state of some higher-level concept (for example, a database cluster, a certificate request, an application release), and then run a controller (often called an "operator") that watches those custom resources and reconciles the cluster to match. That controller may create Deployments, StatefulSets, Services, Secrets, or cloud resources to implement the desired state encoded in the custom resource.

The incorrect answers are made-up expansions. CRDs are not related to Rust in Kubernetes terminology, and "custom restricted definition" is not the standard meaning.

So the verified meaning is: CRD = CustomResourceDefinition, used to extend Kubernetes APIs and enable Kubernetes-native automation via controllers/operators.

NEW QUESTION # 202

.....

