

KCNA Learning Materials: Kubernetes and Cloud Native Associate & KCNA Exam braindumps



BONUS!!! Download part of TestInsides KCNA dumps for free: <https://drive.google.com/open?id=1LYK3YCl6xs2YnzYB66G4lXFtQWs7bO5n>

All of our KCNA exam questions have high pass rate as 99% to 100% and they are valid. We revise our KCNA study guide aperiodicity. You may rest assured that what you purchase are the latest and high-quality KCNA preparation materials. We guarantee our KCNA practice prep will be good value for money, every user will benefit from our KCNA Exam Guide. If you fail exams we will refund the full test dumps cost to you soon. Every extra penny deserves its value. Our KCNA test questions will be your best choice.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Certification Exam is a popular certification program that validates the skills and knowledge of professionals in the field of Kubernetes and cloud native technologies. Kubernetes and Cloud Native Associate certification exam is designed to test the proficiency of candidates in using Kubernetes and cloud native technologies to develop, deploy, and manage scalable and resilient applications.

>> KCNA Exam Practice <<

Free PDF Linux Foundation - Perfect KCNA Exam Practice

Finding 60 exam preparation material that suits your learning preferences, timetable, and objectives is essential to prepare successfully for the test. You can prepare for the Linux Foundation KCNA test in a short time and attain the Kubernetes and Cloud Native Associate certification exam with the aid of our updated and valid exam questions. We emphasize quality over quantity, so we provide you with Linux Foundation KCNA Actual Exam questions to help you succeed without overwhelming you.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q62-Q67):

NEW QUESTION # 62

You are deploying a pod that requires access to a specific storage volume attached to a particular node. Which Kubernetes feature can you utilize to guarantee the pod is scheduled only on that specific node?

- A. Taints and tolerations
- B. Pod anti-affinity
- C. Pod affinity
- D. Node affinity
- E. Node anti-affinity

Answer: A,D

Explanation:

You can achieve this by using either •nodeAffinity• or •taints and tolerationS: 'nodeAffinity•• Similar to the previous question, define

'requiredDuringSchedulingIgnoredDuringExecution' to enforce scheduling on the specific node. 'Taints and Tolerations': Apply a taint on the specific node that reflects the volume availability. Then, configure the pod to tolerate that specific taint, ensuring it can only be scheduled on the node with the matching taint. While 'podAffinity' and 'podAntiAffinity' are useful for grouping or distributing pods, they do not directly guarantee scheduling on a specific node based on volume availability.

NEW QUESTION # 63


How does service logical group set of pods?

- A. Using IP address
- **B. Using label and selectors**
- C. Using hostname

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/services-networking/service/>



NEW QUESTION # 64

Which of the following is NOT a key characteristic of a serverless architecture?

- **A. Pre-provisioned and managed servers-**
- B. Event-driven execution model.
- C. Focus on business logic rather than infrastructure management
- D. Pay-per-use pricing model.
- E. Automatic scaling based on workload demand

Answer: A

Explanation:

Serverless architectures are designed to abstract away the need for managing servers. They rely on cloud providers to handle provisioning, scaling, and infrastructure maintenance. Option C, Pre-provisioned and managed servers, contradicts the core principle of serverless computing, which is to avoid managing server infrastructure.

NEW QUESTION # 65

You are implementing a GitOps workflow for a complex application with multiple microservices. To manage dependencies between these microservices and ensure their correct deployment order, which approach would be most suitable?

- A. Use a single Git repository for all microservices and define deployment dependencies using Helm charts.
- B. Use a separate Git repository for each microservice and manually coordinate their deployments.
- C. Use a separate Git repository for each microservice and configure a CI/CD pipeline to manage their dependencies.
- D. Use a Kubernetes Operator to automate the deployment and management of the microservices, ensuring their dependencies are met.

- E. Use a single Git repository for all microservices and leverage a GitOps tool like ArgoCD or Flux to manage dependencies with features like resource dependencies and deployment order.

Answer: E

Explanation:

Option C provides the most suitable approach for managing dependencies between multiple microservices in a GitOps workflow. Using a single repository for all microservices and leveraging a GitOps tool like ArgoCD or Flux allows you to define resource dependencies and specify the deployment order for microservices, ensuring a consistent and predictable deployment process.

NEW QUESTION # 66

You have a Kubernetes cluster with multiple applications deployed. Each application is instrumented to emit logs, metrics, and traces. You want to use a single dashboard to visualize the performance of all applications in a unified view. What are the possible approaches to achieve this?

- A. Use Prometheus to aggregate metrics from all applications and create a single dashboard.
- B. Use a custom application to collect and aggregate data from Prometheus, Jaeger, and Loki, and then visualize the data on a custom dashboard.
- C. Use Loki for log aggregation and create a single dashboard in Grafana to visualize logs, metrics, and traces.
- D. Use Grafana to create a single dashboard that queries data from Prometheus and Jaeger.
- E. Configure Kubernetes to forward logs, metrics, and traces to a centralized observability platform like CloudWatch or Stackdriver.

Answer: A,B,C,D,E

Explanation:

All of the provided options can contribute to achieving a unified dashboard for visualizing the performance of multiple applications. Each option has its strengths and weaknesses: A: Use Prometheus to aggregate metrics from all applications and create a single dashboard. Prometheus is a powerful tool for collecting and aggregating metrics. You can use Prometheus's query language to fetch data from multiple applications and create a centralized dashboard in Grafana or a custom application. B: Use Grafana to create a single dashboard that queries data from Prometheus and Jaeger. Grafana is a popular dashboarding tool that can visualize data from multiple sources. It can query metrics from Prometheus and tracing data from Jaeger to create a unified view. C: Use Loki for log aggregation and create a single dashboard in Grafana to visualize logs, metrics, and traces. Loki is a log aggregation system that can collect logs from various sources, including Kubernetes. By integrating Loki with Grafana, you can visualize logs, metrics, and traces on a single dashboard. D: Use a custom application to collect and aggregate data from Prometheus, Jaeger, and Loki, and then visualize the data on a custom dashboard. You can build a custom application to collect data from Prometheus, Jaeger, and Loki and then create a custom dashboard using a framework like React or Vue.js. This allows you to have full control over the data aggregation and visualization process. E: Configure Kubernetes to forward logs, metrics, and traces to a centralized observability platform like CloudWatch or Stackdriver. Cloud-based observability platforms like Amazon CloudWatch or Google Stackdriver provide a centralized platform for collecting, aggregating, and visualizing data from multiple applications. These platforms often have pre-built dashboards and alerting capabilities, making it easy to monitor and analyze data from different applications in a unified view. The best approach depends on your specific needs, resources, and preferred tools. You can choose a combination of these options to meet your requirements.

NEW QUESTION # 67

.....

The meaning of qualifying examinations is, in some ways, to prove the candidate's ability to obtain KCNA qualifications that show your ability in various fields of expertise. If you choose our KCNA learning guide materials, you can create more unlimited value in the limited study time, learn more knowledge, and take the exam that you can take. Through qualifying KCNA examinations, this is our KCNA real questions and the common goal of every user, we are trustworthy helpers, so please don't miss such a good opportunity.

Latest KCNA Exam Labs: <https://www.testinsides.top/KCNA-dumps-review.html>

- Free PDF Quiz Linux Foundation - KCNA - Unparalleled Kubernetes and Cloud Native Associate Exam Practice ☐ Easily obtain free download of ⇒ KCNA ⇐ by searching on ☐ www.pass4test.com ☐ KCNA Test Pattern
- Pass Linux Foundation KCNA Exam Easily With Questions And Answers ☐ Easily obtain free download of 「 KCNA 」 by searching on ☐ www.pdfvce.com ☐ KCNA Test Pattern

- BTW, DOWNLOAD part of TestInsides KCNA dumps from Cloud Storage: <https://drive.google.com/open?id=1LYK3YCl6xs2YnzYB66G4lXFtOWs7bO5n>

BTW, DOWNLOAD part of TestInsides KCNA dumps from Cloud Storage: <https://drive.google.com/open?id=1LYK3YCl6xs2YnzYB66G4lXFtOWs7bO5n>