



Free Questions for KCNA by ebraindumps

Shared by Hatfield on 12-12-2023

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

What framework allows developers to write code without worrying about the servers and operating systems they will run on?

Options:

- A- Virtualization
- B- Docker
- C- Serverless
- D- Kubernetes

Answer:

C

Question 2

Question Type: MultipleChoice

Which role is responsible of creating service level indicator 'SLI', service level objective 'SLO', & Service Level Agreements 'SLA'

Options:

- A- Site reliability engineer 'SRE'
- B- DevOps
- C- GitOps
- D- Security and compliance engineer
- E- Developer

Answer:

A

Explanation:

<https://www.atlassian.com/incident-management/kpis/sla-vs-slo-vs-sli>

How does this impact SREs?

For those of you following Google's model and using [Site Reliability Engineering \(SRE\) teams](#) to bridge the gap between development and operations, SLAs, SLOs, and SLIs are foundational to success. SLAs help teams set boundaries and error budgets. SLOs help prioritize work. And SLIs tell SREs when they need to freeze all launches to save an endangered error budget—and when they can loosen up the reins.

Question 3

Question Type: MultipleChoice

Which tool is built on the GitOps toolkit?

Options:

- A- Jenkins-X
- B- GitHub Workflow & Actions
- C- Flux
- D- Jenkins
- E- ArgoCD
- F- Travis CI

Answer:

C

Explanation:

<https://fluxcd.io/#gitops-toolkit>

GitOps Toolkit

The set of APIs and controllers that make up the runtime for Flux. You can use the GitOps Toolkit to extend Flux, and to build your own systems for continuous delivery.

Note: Argo CD is a GitOps tool and not using GitOps toolkit

Question 4

Question Type: MultipleChoice

In distributed system tracing, is the term used to refer to a request as it passes through a single component of the distributed system?

Options:

A- Log

B- Span

C- Trace

D- Bucket

Answer:

B

Explanation:

https://www.splunk.com/en_us/data-insider/what-is-distributed-tracing.html

How does distributed tracing work?

To quickly grasp how distributed tracing works, it's best to look at how it handles a request. Tracing starts the moment an end user interacts with an application. When the user sends an initial request — an HTTP request, to use a common example — it is assigned a unique trace ID. As the request moves through the host system, every operation performed on it (called a “span” or a “child span”) is tagged with that first request ID, as well as its own unique ID, plus the ID of the operation that originally generated the current request (called the “parent span”).

Each span is a single step on the request's journey and is encoded with important information relating to the microservice process that is performing that operation. These include:

- The service name and address of the process handling the request.
- Logs and events that provide context about the process's activity.
- Tags to query and filter requests by session ID, database host, HTTP method, and other identifiers.

Question 5

Question Type: MultipleChoice

Which of the following is an example of vertical scaling?

Options:

- A- Using cluster autoscaler
- B- Adding more resources (memory and/or cpu) to a kubernetes node
- C- Adding more nodes to kubernetes cluster
- D- Adding more replica pods to a deployment

Answer:

B

Explanation:

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

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

Question 6

Question Type: MultipleChoice

What tool allows you to create self-managing, self-scaling, self-healing storage?

Options:

A- Persistent Volume

B- Persistent Volume Claim

C- Storage Class

D- Rook

E- Volume

Answer:

D

Explanation:

<https://rook.io/>

Storage Operators for Kubernetes

Rook turns distributed storage systems into self-managing, self-scaling, self-healing storage services. It automates the tasks of a storage administrator: deployment, bootstrapping, configuration, provisioning, scaling, upgrading, migration, disaster recovery, monitoring, and resource management.

- Rook uses the power of the Kubernetes platform to deliver its services via a Kubernetes Operator for each storage provider.

Question 7

Question Type: MultipleChoice

What tool allows us to build useful visual representations of prometheus data?

Options:

- A- Grafana
- B- kubectl
- C- Distributed system tracing
- D- Rook
- E- Kibana

Answer:

A

Explanation:

<https://prometheus.io/>

Great visualization

Prometheus has multiple modes for visualizing data: a built-in expression browser, Grafana integration, and a console template language.

Question 8

Question Type: MultipleChoice

What is container runtime?

Options:

- A- The amount of time it takes a container to execute
- B- A container image format
- C- Another term of kubelet or kubectl
- D- Software that runs containers

Answer:

D

Explanation:

<https://www.aquasec.com/cloud-native-academy/container-security/container-runtime/>

What Is a Container Runtime?

A container runtime, also known as container engine, is a software component that manages containers on a host operating system. In a [containerized architecture](#), the container runtime is responsible for loading container images from a repository, monitoring container health, isolating system resources for use of a container, and managing container lifecycle.

Common container runtimes commonly work together with container orchestration tools. A container orchestrator is responsible for managing clusters of containers, taking care of container scalability, networking, and security. The container engine is responsible for managing the individual containers running on every compute node.

Common examples of container runtimes are runC, containerd, Docker, and Podman. There are three main types of container runtimes—low-level runtimes, and sandboxed or virtualized runtimes.

Question 9

Question Type: MultipleChoice

How would you return all the pod data in the json format using kubectl command?

Options:

- A- kubectl get pods -o json
- B- kubectl get pods --all-namespaces
- C- kubectl get pods -o wide
- D- kubectl get pods -o jsonpath

Answer:

A

Explanation:

Question 10

Question Type: MultipleChoice

What does the 'kops' acronym means?

Options:

- A- Kubernetes Open Platform Specification
- B- Kubernetes Operations
- C- Kubernetes Operators
- D- Kubernetes Operation Policy Specification

Answer:

B

Explanation:

<https://github.com/kubernetes/kops>

kOps - Kubernetes Operations

go report

A+

GO

reference

The easiest way to get a production grade Kubernetes cluster up and running.

What is kOps?

We like to think of it as `kubectl` for clusters.

`kops` will not only help you create, destroy, upgrade and maintain production-grade Kubernetes cluster, but it will also provision the necessary cloud infrastructure.

AWS (Amazon Web Services) is currently officially supported, with DigitalOcean support, and Azure and AliCloud in alpha.

To Get Premium Files for KCNA Visit

<https://www.p2pexams.com/products/kcna>

For More Free Questions Visit

<https://www.p2pexams.com/linux-foundation/pdf/kcna>

