


Valid KCNA Exam Syllabus & Valid Exam KCNA Practice



P.S. Free 2026 Linux Foundation KCNA dumps are available on Google Drive shared by DumpsTorrent:
https://drive.google.com/open?id=1eHqgJ_ovUgs5vliqXwOfOhW8t7LeMegM

The DumpsTorrent offers desktop Linux Foundation KCNA Practice Exam software for students to practice for the KCNA exam. This software mimics the actual Kubernetes and Cloud Native Associate (KCNA) exam and tracks the student's progress, records grades, and compares results. Available for Windows computers, it requires an internet connection only for license validation.

Linux Foundation KCNA Exam Syllabus Topics:

Section	Objectives
Topic 1: Cloud Security and Governance	- Security fundamentals in cloud native systems • 1. Basic Kubernetes security principles • 2. Identity and access concepts
Topic 2: Cloud Native Architecture	- Microservices and distributed systems • 1. Scalability and resilience patterns • 2. Service discovery and API gateways
Topic 3: Cloud Native Observability	- Monitoring and logging • 1. Observability tools overview • 2. Metrics, logs, and tracing concepts
Topic 4: Container Orchestration	- Container lifecycle and scheduling • 1. Container runtimes • 2. Workload management concepts
Topic 5: Kubernetes Fundamentals	- Core Kubernetes concepts • 1. Cluster architecture overview • 2. Pods, nodes, and control plane basics
Topic 6: Cloud Native Application Delivery	- CI/CD and deployment strategies • 1. Release strategies (blue/green, canary) • 2. GitOps concepts

>> Valid KCNA Exam Syllabus <<

Three Easy-to-Use and Compatible Formats of KCNA Exam Questions

There is a high demand for Linux Foundation Development certification, therefore there is an increase in the number of Linux Foundation KCNA exam candidates. Many resources are available on the internet to prepare for the Kubernetes and Cloud Native Associate exam. DumpsTorrent is one of the best certification exam preparation material providers where you can find newly released Linux Foundation KCNA Dumps for your exam preparation. With years of experience in compiling top-notch relevant Linux Foundation KCNA dumps questions, we also offer the Linux Foundation KCNA practice test (online and offline) to help you get familiar with the actual exam environment.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q264-Q269):

NEW QUESTION # 264

Which of the following is a good habit for cloud native cost efficiency?

- **A. Follow an automated approach to cost optimization, including visibility and forecasting.**
- B. Follow manual processes for cost analysis, including visibility and forecasting.
- C. Use only one cloud provider to simplify the cost analysis.
- D. Keep your legacy workloads unchanged, to avoid cloud costs.

Answer: A

Explanation:

The correct answer is A. In cloud-native environments, costs are highly dynamic: autoscaling changes compute footprint, ephemeral environments come and go, and usage-based billing applies to storage, network egress, load balancers, and observability tooling. Because of this variability, automation is the most sustainable way to achieve cost efficiency. Automated visibility (dashboards, chargeback/showback), anomaly detection, and forecasting help teams understand where spend is coming from and how it changes over time. Automated optimization actions can include right-sizing requests/limits, enforcing TTLs on preview environments, scaling down idle clusters, and cleaning unused resources.

Manual processes (B) don't scale as complexity grows. By the time someone reviews a spreadsheet or dashboard weekly, cost spikes may have already occurred. Automation enables fast feedback loops and guardrails, which is essential for preventing runaway spend caused by misconfiguration (e.g., excessive log ingestion, unbounded autoscaling, oversized node pools).

Option C is not a cost-efficiency "habit." Single-provider strategies may simplify some billing views, but they can also reduce leverage and may not be feasible for resilience/compliance; it's a business choice, not a best practice for cloud-native cost management. Option D is counterproductive: keeping legacy workloads unchanged often wastes money because cloud efficiency typically requires adapting workloads-right-sizing, adopting autoscaling, and using managed services appropriately.

In Kubernetes specifically, cost efficiency is tightly linked to resource management: accurate CPU/memory requests, limits where appropriate, cluster autoscaler tuning, and avoiding overprovisioning. Observability also matters because you can't optimize what you can't measure. Therefore, the best habit is an automated cost optimization approach with strong visibility and forecasting-A.

NEW QUESTION # 265

Which item is a Kubernetes node component?

- A. etcd
- B. kube-scheduler
- **C. kube-proxy**
- D. kubectrl

Answer: C

Explanation:

A Kubernetes node component is a component that runs on worker nodes to support Pods and node-level networking/operations. Among the options, kube-proxy is a node component, so C is correct.

kube-proxy runs on each node and implements parts of the Kubernetes Service networking model. It watches the API server for Service and endpoint updates and then programs node networking rules (iptables/IPVS, or equivalent) so traffic sent to a Service IP/port is forwarded to one of the backend Pod endpoints. This is essential for stable virtual IPs and load distribution across Pods. Why the other options are not node components:

* kube-scheduler is a control plane component; it assigns Pods to nodes but does not run on every node as part of node functionality.

* kubectrl is a client CLI tool used by humans/automation; it is not a cluster component.

* etcd is the control plane datastore; it stores cluster state and is not a per-node workload component.

Operationally, kube-proxy can be replaced by some modern CNI/eBPF dataplanes, but in classic Kubernetes architecture it

remains the canonical node-level component for Service rule programming. Understanding which components are node vs control plane is key for troubleshooting: node issues involve kubelet/runtime /kube-proxy/CNI; control plane issues involve API server/scheduler/controller-manager/etcd. So, the verified node component in this list is kube-proxy (C).

NEW QUESTION # 266

Which of the following sentences is true about namespaces in Kubernetes?

- A. You can create a namespace within another namespace in Kubernetes.
- B. You can create two resources of the same kind and name in a namespace.
- C. The default namespace exists when a new cluster is created.
- D. All the objects in the cluster are namespaced by default.

Answer: C

NEW QUESTION # 267

You're designing a serverless application that processes large amounts of data using AWS Lambda

a. You need to ensure that your Lambda functions can handle potential spikes in workload without compromising performance.

Which of the following approaches would be most effective in achieving this goal?

- A. Implement a message queue in front of Lambda functions to buffer incoming requests.
- B. Enable concurrency settings for Lambda functions to allow multiple invocations simultaneously.
- C. Configure Lambda functions to use a larger memory allocation.
- D. Use serverless frameworks like AWS Serverless Application Model (SAM) to manage Lambda deployments.
- E. Utilize AWS Elastic Beanstalk to auto-scale Lambda functions based on workload.

Answer: A,B

Explanation:

The most effective approaches to handle workload spikes are implementing a message queue (like SQS or SNS) to buffer requests and enabling concurrency settings on Lambda functions. Option A increases memory but doesn't address scalability. Option C is not relevant to Lambda auto-scaling. Option D is a deployment tool, not directly related to handling workload spikes. Option E directly enables parallel processing, making it a suitable solution.

NEW QUESTION # 268

Which of the following sentences is true about namespaces in Kubernetes?

- A. You can create a namespace within another namespace in Kubernetes.
- B. You can create two resources of the same kind and name in a namespace.
- C. All the objects in the cluster are namespaced by default.
- D. The default namespace exists when a new cluster is created.

Answer: D

Explanation:

The true statement is C: the default namespace exists when a new cluster is created. Namespaces are a Kubernetes mechanism for partitioning cluster resources into logical groups. When you set up a cluster, Kubernetes creates some initial namespaces (including default, and commonly kube-system, kube-public, and kube-node-lease). The default namespace is where resources go if you don't specify a namespace explicitly.

Option A is false because namespaces are not hierarchical; Kubernetes does not support "namespaces inside namespaces." Option B is false because within a given namespace, resource names must be unique per resource kind. You can't have two Deployments with the same name in the same namespace. You can have a Deployment named web in one namespace and another Deployment named web in a different namespace- namespaces provide that scope boundary. Option D is false because not all objects are namespaced. Many resources are cluster-scoped (for example, Nodes, PersistentVolumes, ClusterRoles, ClusterRoleBindings, and StorageClasses). Namespaces apply only to namespaced resources.

Operationally, namespaces support multi-tenancy and environment separation (dev/test/prod), RBAC scoping, resource quotas, and policy boundaries. For example, you can grant a team access only to their namespace and enforce quotas that prevent them from

