

# Mock KCNA Exams & Valid Braindumps KCNA Free



P.S. Free 2026 Linux Foundation KCNA dumps are available on Google Drive shared by Exam4Labs:  
<https://drive.google.com/open?id=1v-dNcKgyVLJJZVTnmcdMEoFRspGGU5SF>

They put all their efforts to maintain the top standard of Linux Foundation KCNA exam questions all the time. So you rest assured that with Linux Foundation KCNA exam dumps you will get everything thing that is mandatory to learn, prepare and pass the difficult Linux Foundation KCNA Exam with good scores. Take the best decision of your career and just enroll in the Linux Foundation KCNA certification exam and start preparation with Linux Foundation KCNA practice questions without wasting further time.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Certification Exam is a highly sought-after certification for professionals working in the field of cloud computing. Kubernetes and Cloud Native Associate certification exam is designed to test the candidate's knowledge and skills in Kubernetes and cloud-native technologies. KCNA exam is conducted by the Linux Foundation, which is a non-profit organization that supports the development of open-source software.

The Linux Foundation KCNA Exam is designed to test the candidate's knowledge and understanding of Kubernetes and cloud-native technologies. It covers a range of topics such as deploying, managing, and scaling applications in Kubernetes, understanding Kubernetes architecture and components, and configuring Kubernetes networking and security. KCNA exam also tests the candidate's knowledge of containerization technologies such as Docker and container orchestration tools like Helm.

>> [Mock KCNA Exams](#) <<

**Pass Guaranteed 2026 Linux Foundation Newest Mock KCNA Exams**

At the information age, knowledge is wealth as well as productivity. All excellent people will become outstanding one day as long as one masters skill. In order to train qualified personnel, our company has launched the KCNA Study Materials for job seekers. We are professional to help tens of thousands of the candidates get their KCNA certification with our high quality of KCNA exam questions and live a better life.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) certification exam is designed to test the candidate's knowledge and understanding of Kubernetes and cloud-native technologies. Kubernetes and Cloud Native Associate certification is ideal for individuals who are interested in pursuing a career in cloud computing and want to validate their skills and knowledge in this area.

## **Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q143-Q148):**

### **NEW QUESTION # 143**

You are running a large-scale, microservice-based application in Kubernetes, and you are using HorizontalPodAutoscalers (HPAs) for automatic scaling. Which of the following statements about resource utilization monitoring in this scenario is true?

- A. It is essential to monitor resource utilization at both the Pod and namespace levels to get a comprehensive view of resource usage and optimize scaling
- B. Monitoring resource utilization is only necessary during peak load periods, not during normal operations.
- C. Resource utilization monitoring is not critical in Kubernetes as the system automatically manages resource allocation and scaling.
- D. Monitoring resource utilization at the node level is sufficient for scaling individual microservices.
- E. Monitoring resource utilization at the Pod level provides sufficient insights for overall application health and scaling decisions-

### **Answer: A**

Explanation:

Monitoring resource utilization at both the Pod and namespace levels provides a comprehensive understanding of resource usage. Monitoring Pods allows you to identify resource-intensive services and fine-tune their scaling behavior. Monitoring namespaces helps you understand the overall resource consumption and optimize scaling across the entire application. This is crucial in large-scale microservice deployments where resources can be easily exhausted if not carefully monitored.

### **NEW QUESTION # 144**

What edge and service proxy tool is designed to be integrated with cloud native applications?

- A. CNI
- B. gRPC
- C. Envoy
- D. CoreDNS

### **Answer: C**

Explanation:

The correct answer is D: Envoy. Envoy is a high-performance edge and service proxy designed for cloud-native environments. It is commonly used as the data plane in service meshes and modern API gateways because it provides consistent traffic management, observability, and security features across microservices without requiring every application to implement those capabilities directly. Envoy operates at Layer 7 (application-aware) and supports protocols like HTTP/1.1, HTTP/2, gRPC, and more. It can handle routing, load balancing, retries, timeouts, circuit breaking, rate limiting, TLS termination, and mutual TLS (mTLS). Envoy also emits rich telemetry (metrics, access logs, tracing) that integrates well with cloud-native observability stacks.

Why the other options are incorrect:

CoreDNS (A) provides DNS-based service discovery within Kubernetes; it is not an edge/service proxy.

CNI (B) is a specification and plugin ecosystem for container networking (Pod networking), not a proxy.

gRPC (C) is an RPC protocol/framework used by applications; it's not a proxy tool. (Envoy can proxy gRPC traffic, but gRPC itself isn't the proxy.) In Kubernetes architectures, Envoy often appears in two places: (1) at the edge as part of an ingress/gateway layer, and (2) sidecar proxies alongside Pods in a service mesh (like Istio) to standardize service-to-service communication controls and telemetry. This is why it is described as "designed to be integrated with cloud native applications": it's purpose-built for dynamic service discovery, resilient routing, and operational visibility in distributed systems.

So the verified correct choice is D (Envoy).

## NEW QUESTION # 145

How to create a headless Service?

- A. By specifying .spec.clusterIP: None
- B. By specifying .spec.clusterIP: 0.0.0.0
- C. By specifying .spec.clusterIP: headless
- D. By specifying .spec.clusterIP: localhost

**Answer: A**

Explanation:

A headless Service is created by setting spec.clusterIP: None, so B is correct. Normally, a Service gets a ClusterIP, and kube-proxy (or an alternative dataplane) implements virtual-IP-based load balancing to route traffic from that ClusterIP to the backend Pods. A headless Service intentionally disables that virtual IP allocation. Instead of giving you a single stable VIP, Kubernetes publishes DNS records that resolve directly to the endpoints (the Pod IPs) behind the Service.

This is especially important for workloads that need direct endpoint discovery or stable per-Pod identities, such as StatefulSets. With a headless Service, clients can discover all Pod IPs (or individual Pod DNS names in StatefulSet patterns) and implement their own selection, quorum, or leader/follower logic. Kubernetes DNS (CoreDNS) responds differently for headless Services: rather than returning a single ClusterIP, it returns multiple A/AAAA records (one per endpoint) or SRV records for named ports, enabling richer service discovery behavior.

The other options are invalid. "headless" is not a magic value for clusterIP; the API expects either an actual IP address assigned by the cluster or the special literal None. 0.0.0.0 and localhost are not valid ways to request headless semantics. Kubernetes uses None specifically to signal "do not allocate a ClusterIP." Operationally, headless Services are used to: (1) expose each backend instance individually, (2) support stateful clustering and stable DNS names, and (3) avoid load balancing when the application or client library must choose endpoints itself. The key is that the Service still provides a stable DNS name, but the resolution yields endpoints, not a VIP.

## NEW QUESTION # 146

Fluentd is the only way to export logs from Kubernetes cluster or applications running in cluster

- A. False
- B. True

**Answer: A**

Explanation:

<https://github.com/cncf/landscape#trail-map>

## NEW QUESTION # 147

What are the characteristics for building every cloud-native application?

- A. Resiliency, Agility, Operability, Observability
- B. Resiliency, Operability, Observability, Availability
- C. Kubernetes, Operability, Observability, Availability
- D. Resiliency, Containerd, Observability, Agility

**Answer: A**

Explanation:

Cloud-native applications are typically designed to thrive in dynamic, distributed environments where infrastructure is elastic and failures are expected. The best set of characteristics listed is Resiliency, Agility, Operability, Observability, making D correct.

Resiliency means the application and its supporting platform can tolerate failures and continue providing service. In Kubernetes terms, resiliency is supported through self-healing controllers, replica management, health probes, and safe rollout mechanisms, but the application must also be designed to handle transient failures, retries, and graceful degradation.

Agility reflects the ability to deliver changes quickly and safely. Cloud-native systems emphasize automation, CI/CD, declarative configuration, and small, frequent releases—often enabled by Kubernetes primitives like Deployments and rollout strategies. Agility is about reducing the friction to ship improvements while maintaining reliability.

Operability is how manageable the system is in production: clear configuration, predictable deployments, safe scaling, and automation-friendly operations. Kubernetes encourages operability through consistent APIs, controllers, and standardized patterns for configuration and lifecycle.

Observability means you can understand what's happening inside the system using telemetry-metrics, logs, and traces-so you can troubleshoot issues, measure SLOs, and improve performance. Kubernetes provides many integration points for observability, but cloud-native apps must also emit meaningful signals.

Options B and C include items that are not "characteristics" (Containerd is a runtime; Kubernetes is a platform). Option A includes "availability," which is important, but the canonical cloud-native framing in this question emphasizes the four qualities in D as the foundational build characteristics.

## NEW QUESTION # 148

• • • • •

Valid Braindumps KCNA Free: <https://www.exam4labs.com/KCNA-practice-torrent.html>

P.S. Free 2026 Linux Foundation KCNA dumps are available on Google Drive shared by Exam4Labs: <https://drive.google.com/open?id=1v-dNcKgyVLJJZVTnmcdeMEOFrSpGGU5SF>