

# Linux Foundation KCNA Actual Test | Test KCNA Dump



BTW, DOWNLOAD part of PassSureExam KCNA dumps from Cloud Storage: <https://drive.google.com/open?id=1yruOeq32OtDxgtfW5JnBJpqt-WcGNYUx>

Our KCNA study prep has inspired millions of exam candidates to pursue their dreams and motivated them to learn more efficiently. Many customers get manifest improvement. KCNA simulating exam will inspire your potential. And you will be more successful with the help of our KCNA training guide. Just imagine that when you have the certification, you will have a lot of opportunities to come to the bigger companies and get a higher salary.

If you are interested in pursuing a career in cloud-native technologies, the Linux Foundation KCNA Certification Exam is a great way to get started. With this certification, you will be able to demonstrate your knowledge and expertise in Kubernetes and cloud-native technologies, which will help you stand out in a competitive job market.

The KCNA Exam is a performance-based exam that involves completing a set of tasks in a real-world environment. KCNA Exam is conducted online and requires candidates to have access to a Linux environment with Kubernetes and other cloud-native tools installed. KCNA exam is designed to test the candidate's ability to deploy and manage containerized applications using Kubernetes and other cloud-native technologies.

>> **Linux Foundation KCNA Actual Test** <<

## Earn The Badge Of Linux Foundation KCNA Certification Exam On The First Attempt

When preparing to take the Linux Foundation KCNA exam dumps, knowing where to start can be a little frustrating, but with PassSureExam Linux Foundation KCNA practice questions, you will feel fully prepared. Using our Linux Foundation KCNA practice test software, you can prepare for the increased difficulty on KCNA Exam day. Plus, we have various question types and difficulty levels so that you can tailor your Linux Foundation KCNA exam dumps preparation to your requirements.

## Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q112-Q117):

### NEW QUESTION # 112

You have a Kubernetes cluster with a pod that uses a resource limit of 100m CPU. What happens if the pod requests more CPU resources than the limit?

- A. The pod will be scheduled on a different node with more available resources.
- B. The pod will be automatically scaled down to 100m CPU.
- C. The pod will continue to run using the requested amount of CPU.
- D. The pod will be killed and restarted.
- E. The pod will be throttled and its performance will be impacted.

**Answer: E**

Explanation:

The pod will be throttled and its performance will be impacted. The CPU limit acts as a hard ceiling, preventing the pod from consuming more resources than specified. Even if the pod requests more CPU, it will be constrained to the limit, leading to performance degradation. The pod will not be killed, restarted, or automatically scaled down. It will continue to run, but its performance will be affected due to the resource constraint.

### NEW QUESTION # 113

In CNCF, who develops specifications for industry standards around container formats and runtimes?

- A. Container Network Interface (CNI)
- B. Linux Foundation Certification Group (LFCG)
- C. Container Runtime Interface (CRI)
- D. Open Container Initiative (OCI)

**Answer: D**

Explanation:

The organization responsible for defining widely adopted standards around container formats and runtime specifications is the Open Container Initiative (OCI), so A is correct. OCI defines the image specification (how container images are structured and stored) and the runtime specification (how to run a container), enabling interoperability across tooling and vendors. This is foundational to the cloud-native ecosystem because it allows different build tools, registries, runtimes, and orchestration platforms to work together reliably.

Within Kubernetes and CNCF-adjacent ecosystems, OCI standards are the reason an image built by one tool can be pushed to a registry and pulled/run by many different runtimes. For example, a Kubernetes node running containerd or CRI-O can run OCI-compliant images consistently. OCI standardization reduces fragmentation and vendor lock-in, which is a core motivation in open source cloud-native architecture.

The other options are not correct for this question. CNI (Container Network Interface) is a standard for configuring container networking, not container image formats and runtimes. CRI (Container Runtime Interface) is a Kubernetes-specific interface between kubelet and container runtimes-it enables pluggable runtimes for Kubernetes, but it is not the industry standard body for container format/runtime specifications. "LFCG" is not the recognized standards body here.

In short: OCI defines the "language" for container images and runtime behavior, which is why the same image can be executed across environments. Kubernetes relies on those standards indirectly through runtimes and tooling, but the specification work is owned by OCI. Therefore, the verified correct answer is A.

### NEW QUESTION # 114

You are building a distributed system with microservices deployed in a Kubernetes cluster. Each microservice has its own database. How would you manage the consistency and reliability of data across these microservices?

- A. Implement a distributed transaction manager
- B. Implement a distributed database system
- C. Use a shared database for all microservices
- D. Use a message broker to ensure data synchronization
- E. Use a combination of event sourcing and CQRS (Command Query Responsibility Segregation)

**Answer: A,B,D,E**

Explanation:

Managing data consistency and reliability in a distributed system with microservices is a complex challenge- There are various approaches to achieve this, including: Distributed Database System: Using a distributed database system like Cassandra or MongoDB allows for data replication and fault tolerance across different nodes. This ensures data availability and consistency. Message Broker A message broker can facilitate data synchronization between microservices- Messages are used to communicate data changes and ensure that all services are aware of the latest state. Distributed Transaction Manager: A distributed transaction manager ensures that multiple transactions across different microservices are completed as a single atomic operation. This helps maintain data consistency and integrity. Event Sourcing and CQRS: Event sourcing involves storing a sequence of events that represent changes to the system. CQRS (Command Query Responsibility Segregation) separates commands (which modify data) from queries (which retrieve data). This approach can help manage data consistency and provide a more robust architecture. The best approach depends on the specific requirements of your application, such as data consistency levels, performance needs, and scalability requirements.

### NEW QUESTION # 115

Consider the following Kubernetes resource configuration for a deployment:

```
kind: Deployment
metadata:
  name: my-app
spec:
  replicas: 3
  selector:
    matchLabels:
      app: my-app
  template:
    metadata:
      labels:
        app: my-app
    spec:
      containers:
      - name: my-app-container
        image: my-app:latest
        resources:
          requests:
            cpu: 100m
            memory: 256Mi
          limits:
            cpu: 200m
            memory: 512Mi
```

If each pod is consistently utilizing only 50% of the requested CPU and memory, what are the potential cost-saving measures you could take without impacting performance?

- A. Reduce the 'requests' values for CPU and memory to 50m and 128Mi, respectively.
- B. Increase the 'limits' values for CPU and memory to 400m and 1024Mi, respectively.
- C. Reduce the 'replicas' value to 2 to decrease the number of pods.
- D. Utilize a vertical pod autoscaler (VPA) to dynamically adjust the resource requests and limits based on actual usage.
- E. Leave the current configuration as it is; the application is already optimized.

**Answer: A,D**

Explanation:

The application is overprovisioned as pods are only utilizing 50% of requested resources. Reducing the 'requests' values to reflect actual usage (50m CPU and 128Mi memory) would optimize resource allocation without affecting performance. A VPA can further refine the resource allocation dynamically based on actual usage, ensuring optimal resource utilization. Increasing 'limits' would not reduce costs and might even increase them. Reducing replicas might impact performance, while leaving the configuration unchanged would continue overprovisioning and waste resources.

### NEW QUESTION # 116

You are building a CI/CD pipeline to deploy a Java application to Kubernetes. The application requires a specific Java version and some libraries to run correctly. How would you ensure that the correct runtime environment is available in the Kubernetes cluster?

- A. Use Kubernetes Secrets to store the Java version and library information securely.
- B. Use Kubernetes Init Containers to install the required dependencies before starting the main container.
- C. Use Kubernetes ConfigMaps to define the environment variables for Java and libraries.
- D. Pre-install the Java version and libraries on the Kubernetes nodes.
- E. Include the necessary Java version and libraries in the Dockerfile for the application image.

**Answer: B,E**

Explanation:

The most common approaches are to include the required Java and libraries within the Dockerfile to create a self-contained image, or to use Kubernetes Init Containers. Init containers run before the main application container, installing the dependencies, ensuring

