

# Top KCNA Dumps Fantastic Questions Pool Only at ValidVCE



DOWNLOAD the newest ValidVCE KCNA PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1x9NQeqdgya7gPWldYcTk4HFz8dGpL>

ValidVCE try hard to makes KCNA exam preparation easy with its several quality features. Our KCNA exam dumps come with 100% refund assurance. We are dedicated to your accomplishment, hence pledges you victory in KCNA exam in a single attempt. If for any reason, a user fails in KCNA exam then he will be refunded the money after the process. Also, we offer 1 year free updates to our KCNA Exam esteemed users; and these updates will be entitled to your account right from the date of purchase. Also the 24/7 Customer support is given to users, who can email us if they find any haziness in the KCNA exam dumps, our team will merely answer to your all KCNA exam product related queries.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Exam is a certification program designed to test the knowledge and skills of individuals who are working with Kubernetes and other cloud-native technologies. Kubernetes has become the de facto standard for container orchestration, and with the increasing adoption of cloud-native technologies, the demand for certified professionals is on the rise.

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.

## Buy ValidVCE Linux Foundation KCNA Exam Dumps Today and Get Free Updates for 1 year

If I tell you, you can get international certification by using KCNA preparation materials for twenty to thirty hours. You must be very surprised. However, you must believe that this is true! You can ask anyone who has used KCNA Actual Exam. We can receive numerous warm feedbacks every day. Our reputation is really good. After you have learned about the achievements of KCNA study questions, you will definitely choose us!

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Certification Exam is a professional certification exam that validates the skills and knowledge of individuals in the field of Kubernetes and cloud native technologies. KCNA Exam is designed for individuals who are interested in learning and mastering the fundamentals of Kubernetes and cloud native technologies, and who wish to demonstrate their expertise in these areas to potential employers.

## Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q113-Q118):

### NEW QUESTION # 113

What is ephemeral storage?

- A. Storage that may grow dynamically.
- B. Storage that is always provisioned locally.
- C. Storage used by multiple consumers (e.g., multiple Pods).
- **D. Storage space that need not persist across restarts.**

**Answer: D**

Explanation:

The correct answer is A: ephemeral storage is non-persistent storage whose data does not need to survive Pod restarts or rescheduling. In Kubernetes, ephemeral storage typically refers to storage tied to the Pod's lifetime—such as the container writable layer, emptyDir volumes, and other temporary storage types. When a Pod is deleted or moved to a different node, that data is generally lost.

This is different from persistent storage, which is backed by PersistentVolumes and PersistentVolumeClaims and is designed to outlive individual Pod instances. Ephemeral storage is commonly used for caches, scratch space, temporary files, and intermediate build artifacts—data that can be recreated and is not the authoritative system of record.

Option B is incorrect because "may grow dynamically" describes an allocation behavior, not the defining characteristic of ephemeral storage. Option C is incorrect because multiple consumers is about access semantics (ReadWriteMany etc.) and shared volumes, not ephemerality. Option D is incorrect because ephemeral storage is not "always provisioned locally" in a strict sense; while many ephemeral forms are local to the node, the definition is about lifecycle and persistence guarantees, not necessarily physical locality. Operationally, ephemeral storage is an important scheduling and reliability consideration. Pods can request/limit ephemeral storage similarly to CPU/memory, and nodes can evict Pods under disk pressure. Mismanaged ephemeral storage (logs written to the container filesystem, runaway temp files) can cause node disk exhaustion and cascading failures. Best practices include shipping logs off-node, using emptyDir intentionally with size limits where supported, and using persistent volumes for state that must survive restarts.

So, ephemeral storage is best defined as storage that does not need to persist across restarts/rescheduling, matching option A.

### NEW QUESTION # 114

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

- A. Container Runtime Interface (CRI)
- B. Linux Foundation Certification Group (LFCG)
- C. Container Network Interface (CNI)
- **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 # 115

Your organization is adopting a cloud-native approach and plans to migrate several legacy applications to Kubernetes. Which role would be primarily responsible for designing and implementing the overall Kubernetes infrastructure, including resource allocation, networking, and security policies?

- A. DevOps Engineer
- B. Site Reliability Engineer (SRE)
- C. Data Scientist
- D. Security Engineer
- E. Cloud Architect

**Answer: E**

Explanation:

A Cloud Architect is responsible for the overall design, implementation, and management of cloud-based infrastructure, including Kubernetes- They ensure the infrastructure meets the organization's needs for scalability, security, and performance.

### NEW QUESTION # 116

Imagine there is a requirement to run a database backup every day. Which Kubernetes resource could be used to achieve that?

- A. kube-scheduler
- B. CronJob
- C. Job
- D. Task

**Answer: B**

Explanation:

To run a workload on a repeating schedule (like "every day"), Kubernetes provides CronJob, making B correct. A CronJob creates Jobs according to a cron-formatted schedule, and then each Job creates one or more Pods that run to completion. This is the Kubernetes-native replacement for traditional cron scheduling, but implemented as a declarative resource managed by controllers in the cluster.

For a daily database backup, you'd define a CronJob with a schedule (e.g., "0 2 \* \* \*" for 2:00 AM daily), and specify the Pod template that performs the backup (invokes backup scripts/tools, writes output to durable storage, uploads to object storage, etc.). Kubernetes will then create a Job at each scheduled time. CronJobs also support operational controls like concurrencyPolicy (Allow/Forbid/Replace) to decide what happens if a previous backup is still running, startingDeadlineSeconds to handle missed schedules, and history limits to retain recent successful/failed Job records for debugging.

Option D (Job) is close but not sufficient for "every day." A Job runs a workload until completion once; you would need an external scheduler to create a Job every day. Option A (kube-scheduler) is a control plane component responsible for placing Pods onto nodes and does not schedule recurring tasks. Option C ("Task") is not a standard Kubernetes workload resource.

This question is fundamentally about mapping a recurring operational requirement (backup cadence) to Kubernetes primitives. The correct design is: CronJob triggers Job creation on a schedule; Job runs Pods to completion. Therefore, the correct answer is B.

