

# KCNA Reliable Test Bootcamp 100% Pass | Pass-Sure KCNA Vce Download: Kubernetes and Cloud Native Associate



What's more, part of that Lead2PassExam KCNA dumps now are free: [https://drive.google.com/open?id=1KUZLr5PQtgteTxM6uAna\\_pna\\_Pdi83jf](https://drive.google.com/open?id=1KUZLr5PQtgteTxM6uAna_pna_Pdi83jf)

Thousands of Kubernetes and Cloud Native Associate exam aspirants have already passed their Linux Foundation KCNA certification exam and they all got help from top-notch and easy-to-use Linux Foundation KCNA Exam Questions. You can also use the Lead2PassExam KCNA exam questions and earn the badge of Linux Foundation KCNA certification easily.

Nowadays, using electronic materials to prepare for the exam has become more and more popular, so now, you really should not be restricted to paper materials any more, our electronic KCNA exam torrent will surprise you with their effectiveness and usefulness. I can assure you that you will pass the KCNA Exam as well as getting the related certification under the guidance of our KCNA training materials as easy as pie. Just have a try on our KCNA exam questions, you will love them for sure!

>> **KCNA Reliable Test Bootcamp** <<

## KCNA Vce Download, Practical KCNA Information

You will obtain these updates entirely free if the Linux Foundation KCNA certification authorities issue fresh updates. Lead2PassExam ensures that you will hold the prestigious Linux Foundation KCNA certificate on the first endeavor if you work consistently, taking help from our remarkable, up-to-date, and competitive Linux Foundation KCNA dumps.

The Kubernetes and Cloud Native Associate certification is recognized globally and is highly respected in the industry. It is an excellent way for professionals to demonstrate their expertise in Kubernetes and cloud-native computing to potential employers. Kubernetes and Cloud Native Associate certification also provides access to a community of like-minded professionals who are passionate about cloud-native computing and Kubernetes.

## Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q44-Q49):

### NEW QUESTION # 44

You are implementing a new security policy for your Kubernetes cluster. The policy requires that all pods running in the cluster must authenticate with a specific identity provider before they are allowed to access any resources. Which Kubernetes component is responsible for enforcing this authentication policy?

- A. kube-proxy
- **B. kube-apiserver**
- C. kubect1
- D. etcd

- E. kubelet

**Answer: B**

Explanation:

The kube-apiserver component acts as the central control plane for Kubernetes. It handles all communication and requests from other components, including authentication and authorization. It enforces security policies by verifying credentials and granting access based on configured rules.

#### NEW QUESTION # 45

How many hosts are required to set up a highly available Kubernetes cluster when using an external etcd topology?

- A. Four hosts. One for a control plane node and three for etcd nodes.
- B. Four hosts. Two for control plane nodes and two for etcd nodes.
- C. Three hosts. The control plane nodes and etcd nodes share the same host.
- **D. Six hosts. Three for control plane nodes and three for etcd nodes.**

**Answer: D**

Explanation:

In a highly available (HA) Kubernetes control plane using an external etcd topology, you typically run three control plane nodes and three separate etcd nodes, totaling six hosts, making D correct. HA design relies on quorum-based consensus: etcd uses Raft and requires a majority of members available to make progress. Running three etcd members is the common minimum for HA because it tolerates one member failure while maintaining quorum (2/3).

In the external etcd topology, etcd is decoupled from the control plane nodes. This separation improves fault isolation: if a control plane node fails or is replaced, etcd remains stable and independent; likewise, etcd maintenance can be handled separately. Kubernetes API servers (often multiple instances behind a load balancer) talk to the external etcd cluster for storage of cluster state. Options A and B propose four hosts, but they break common HA/quorum best practices. Two etcd nodes do not form a robust quorum configuration (a two-member etcd cluster cannot tolerate a single failure without losing quorum). One control plane node is not HA for the API server/scheduler/controller-manager components. Option C describes a stacked etcd topology (control plane + etcd on same hosts), which can be HA with three hosts, but the question explicitly says external etcd, not stacked. In stacked topology, you often use three control plane nodes each running an etcd member. In external topology, you use three control plane + three etcd.

Operationally, external etcd topology is often used when you want dedicated resources, separate lifecycle management, or stronger isolation for the datastore. It can reduce blast radius but increases infrastructure footprint and operational complexity (TLS, backup/restore, networking). Still, for the canonical HA external-etcd pattern, the expected answer is six hosts: 3 control plane + 3 etcd.

#### NEW QUESTION # 46

Which of the following workload requires a headless Service while deploying into the namespace?

- A. Deployment
- **B. StatefulSet**
- C. DaemonSet
- D. CronJob

**Answer: B**

Explanation:

A StatefulSet commonly requires a headless Service, so A is the correct answer. In Kubernetes, StatefulSets are designed for workloads that need stable identities, stable network names, and often stable storage per replica. To support that stable identity model, Kubernetes typically uses a headless Service (spec.clusterIP:

None) to provide DNS records that map directly to each Pod, rather than load-balancing behind a single virtual ClusterIP.

With a headless Service, DNS queries return individual endpoint records (the Pod IPs) so that each StatefulSet Pod can be addressed predictably, such as pod-0.service-name.namespace.svc.cluster.local. This is critical for clustered databases, quorum systems, and leader/follower setups where members must discover and address specific peers. The StatefulSet controller then ensures ordered creation/deletion and preserves identity (pod-0, pod-1, etc.), while the headless Service provides discovery for those stable hostnames.

CronJobs run periodic Jobs and don't require stable DNS identity for multiple replicas. Deployments manage stateless replicas and

normally use a standard Service that load-balances across Pods. DaemonSets run one Pod per node, and while they can be exposed by Services, they do not intrinsically require headless discovery.

So while you can use a headless Service for other designs, StatefulSet is the workload type most associated with "requires a headless Service" due to how stable identities and per-Pod addressing work in Kubernetes.

#### NEW QUESTION # 47

What happens with a regular Pod running in Kubernetes when a node fails?

- A. A new Pod with the same UID is scheduled to another node after a while.
- B. A new Pod is scheduled on a different node only if it is configured explicitly.
- C. A new, near-identical Pod but with different UID is scheduled to another node.
- D. By default, a Pod can only be scheduled to the same node when the node fails.

**Answer: C**

Explanation:

B is correct: when a node fails, Kubernetes does not "move" the same Pod instance; instead, a new Pod object (new UID) is created to replace it—assuming the Pod is managed by a controller (Deployment/ReplicaSet, StatefulSet, etc.). A Pod is an API object with a unique identifier (UID) and is tightly associated with the node it's scheduled to via `spec.nodeName`. If the node becomes unreachable, that original Pod cannot be restarted elsewhere because it was bound to that node.

Kubernetes' high availability comes from controllers maintaining desired state. For example, a Deployment desires N replicas. If a node fails and the replicas on that node are lost, the controller will create replacement Pods, and the scheduler will place them onto healthy nodes. These replacement Pods will be "near-identical" in spec (same template), but they are still new instances with new UIDs and typically new IPs.

Why the other options are wrong:

A is incorrect because the UID does not remain the same—Kubernetes creates a new Pod object rather than reusing the old identity.

C is incorrect; pods are not restricted to the same node after failure. The whole point of orchestration is to reschedule elsewhere.

D is incorrect; rescheduling does not require special explicit configuration for typical controller-managed workloads. The controller behavior is standard. (If it's a bare Pod without a controller, it will not be recreated automatically.) This also ties to the difference between "regular Pod" vs controller-managed workloads: a standalone Pod is not self-healing by itself, while a

Deployment/ReplicaSet provides that resilience. In typical production design, you run workloads under controllers specifically so node failure triggers replacement and restores replica count.

Therefore, the correct outcome is B.

#### NEW QUESTION # 48

In Kubernetes, if the API version of feature is v2beta3, it means that:

- A. The software may contain bugs. Enabling a feature may expose bugs.
- B. The version will remain available for all future releases within a Kubernetes major version.
- C. The API may change in incompatible ways in a later software release without notice.
- D. The software is well tested. Enabling a feature is considered safe.

**Answer: C**

Explanation:

The correct answer is B. In Kubernetes API versioning, the stability level is encoded in the version string:

alpha, beta, and stable (v1). A version like v2beta3 indicates the API is in a beta stage. Beta APIs are more mature than alpha, but they are not fully guaranteed stable in perpetuity the way v1 stable APIs are intended to be. The key implication is that while beta APIs are generally usable, they can still undergo incompatible changes in future releases as the API design evolves.

Option B captures that meaning: a beta API may change in ways that break compatibility. This is why teams should treat beta APIs with some caution in production: verify upgrade plans, monitor deprecation notices, and be prepared to adjust manifests or client code when moving between Kubernetes versions.

Why the other options are incorrect:

\* A implies permanence across all future releases in a major version, which is not a beta guarantee.

Kubernetes has deprecation and graduation processes, but beta does not equal "forever."

\* C overstates safety; beta is typically "tested and enabled by default" for some features, but it's not the same as stable API guarantees.

\* D is too vague and misaligned. While any software may contain bugs, the defining point of "beta API" is about stability/compatibility guarantees, not merely "bugs." In practice, Kubernetes communicates API lifecycle clearly: alpha is

