

KCNA Prüfungsguide: Kubernetes and Cloud Native Associate & KCNA echter Test & KCNA sicherlich-zu-bestehen



2026 Die neuesten DeutschPrüfung KCNA PDF-Versionen Prüfungsfragen und KCNA Fragen und Antworten sind kostenlos verfügbar: https://drive.google.com/open?id=1uujpaVsB9YnaWtl-_7SujByPzJrcbml

Die Fragenkataloge zur Linux Foundation KCNA Zertifizierungsprüfung von DeutschPrüfung werden Ihnen zum Erfolg führen. Unsere Fragenkataloge werden von den Experten neuerlich erforscht. Und Sie können deshalb immer die neuesten Forschungsmaterialien bekommen. Wir garantieren Ihnen den Erfolg. Wir helfen Ihnen sehr gerne. Sie werden sicher die genauesten Fragen und Antworten zur Linux Foundation KCNA Zertifizierungsprüfung von uns bekommen. Wir aktualisieren ständig unsere Schulungsinstrumente, um den geänderten Prüfungsthemen anzupassen. Eigentlich liegt der Erfolg nicht weit entfernt. Wenn Sie DeutschPrüfung benutzen, können Sie sicher den Erfolg erlangen.

Die Zertifizierungsprüfung der Linux Foundation KCNA (Kubernetes und Cloud Native Associate) ist ein beliebtes Zertifizierungsprogramm, das die Fähigkeiten und Kenntnisse von Fachleuten im Bereich Kubernetes und Cloud Native Technologies validiert. Diese Zertifizierungsprüfung soll die Kandidaten in der Verwendung von Kubernetes- und Cloud -nativen Technologien zur Entwicklung, Bereitstellung und Verwaltung skalierbarer und belastbarer Anwendungen testen.

>> KCNA Pruefungssimulationen <<

Echte und neueste KCNA Fragen und Antworten der Linux Foundation KCNA Zertifizierungsprüfung

Die simulierten Prüfungen zu machen können Ihre Selbstbewusstsein erstarken. Mit der Simulations-Software Testing Engine von unserer Linux Foundation KCNA können Sie die realistische Atmosphäre dieser Prüfung erfahren. Diese Erfahrungen sind sehr wichtig für Sie bei der späteren echten Linux Foundation KCNA Prüfung. Neben Linux Foundation KCNA haben wir auch viele andere IT-Prüfungsunterlagen geforscht. Diese Prüfungshilfe können Sie auf unserer Webseite finden. Wenn Sie irgend bezügliche Fragen haben, können Sie einfach mit unserem 24/7 online Kundendienst Personal kommunizieren.

Linux Foundation Kubernetes and Cloud Native Associate KCNA Prüfungsfragen mit Lösungen (Q23-Q28):

23. Frage

Imagine you have a Kubernetes cluster running a large number of applications. You want to streamline the configuration of these applications and manage their dependencies consistently. Which open standard would you use to address this?

- A. Kubernetes Ingress
- B. open Policy Agent (OPA)
- C. Helm
- D. Kubernetes Resource Quotas
- E. Service Mesh

Antwort: C

Begründung:

Helm is a package manager for Kubernetes. It allows you to package your application's configuration, dependencies, and deployment logic into a Helm chart. This simplifies application deployment and management, making it easier to manage the configuration and dependencies of multiple applications consistently.

24. Frage

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

Antwort: B

Begründung:

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.

25. Frage

What is the main purpose of etcd in Kubernetes?

- **A. etcd stores all cluster data in a key value store.**
- B. etcd stores the containers running in the cluster for disaster recovery.
- C. etcd stores the YAML definitions for all the cluster components.
- D. etcd stores copies of the Kubernetes config files that live `/etc/`.

Antwort: A

Begründung:

The main purpose of etcd in Kubernetes is to store the cluster's state as a distributed key-value store, so A is correct. Kubernetes is API-driven: objects like Pods, Deployments, Services, ConfigMaps, Secrets, Nodes, and RBAC rules are persisted by the API server into etcd. Controllers, schedulers, and other components then watch the API for changes and reconcile the cluster accordingly. This makes etcd the "source of truth" for desired and observed cluster state.

Options B, C, and D are misconceptions. etcd does not store the running containers; that's the job of the kubelet/container runtime on each node, and container state is ephemeral. etcd does not store `/etc` configuration file copies. And while you may author objects as YAML manifests, Kubernetes stores them internally as API objects (serialized) in etcd—not as "YAML definitions for all components." The data is structured key/value entries representing Kubernetes resources and metadata.

Because etcd is so critical, its performance and reliability directly affect the cluster. Slow disk I/O or poor network latency increases API request latency and can delay controller reconciliation, leading to cascading operational problems (slow rollouts, delayed scheduling, timeouts). That's why etcd is typically run on fast, reliable storage and in an HA configuration (often 3 or 5 members) to maintain quorum and tolerate failures. Backups (snapshots) and restore procedures are also central to disaster recovery: if etcd is

lost, the cluster loses its state.

Security is also important: etcd can contain sensitive information (especially Secrets unless encrypted at rest). Proper TLS, restricted access, and encryption-at-rest configuration are standard best practices.

So, the verified correct answer is A: etcd stores all cluster data/state in a key-value store.

26. Frage

What sentence is true about CronJobs in Kubernetes?

- A. A CronJob creates one container on a repeating schedule.
- **B. A CronJob creates one or multiple Jobs on a repeating schedule.**
- C. The CronJob schedule format is different in Kubernetes and Linux.
- D. CronJobs are useful on Linux but are obsolete in Kubernetes.

Antwort: B

Begründung:

The true statement is A: a Kubernetes CronJob creates Jobs on a repeating schedule. CronJob is a controller designed for time-based execution. You define a schedule using standard cron syntax (minute, hour, day-of-month, month, day-of-week), and when the schedule triggers, the CronJob controller creates a Job object. Then the Job controller creates one or more Pods to run the task to completion.

Option B is incorrect because CronJobs do not "create one container"; they create Jobs, and Jobs create Pods (which may contain one or multiple containers). Option C is wrong because CronJobs are a core Kubernetes workload primitive for recurring tasks and remain widely used for periodic work like backups, batch processing, and cleanup. Option D is wrong because Kubernetes CronJobs intentionally use cron-like scheduling expressions; the format aligns with the cron concept (with Kubernetes-specific controller behavior around missed runs, concurrency, and history).

CronJobs also provide operational controls you don't get from plain Linux cron on a node:

concurrencyPolicy (Allow/Forbid/Replace) to manage overlapping runs

startingDeadlineSeconds to control how missed schedules are handled

history limits for successful/failed Jobs to avoid clutter

integration with Kubernetes RBAC, Secrets, ConfigMaps, and volumes for consistent runtime configuration consistent execution environment via container images, not ad-hoc node scripts Because the CronJob creates Jobs as first-class API objects, you get observability (events/status), predictable retries, and lifecycle management. That's why the accurate statement is A.

27. Frage

You are using ArgoCD to manage your Kubernetes cluster with GitOps. How can you configure ArgoCD to automatically update the cluster when a new image tag is pushed to a container registry?

- **A. Configure ArgoCD to use a webhook that triggers an update whenever a new image tag is pushed.**
- B. Use the *imagePullPolicy field in the container definition to automatically pull the latest image.
- C. Use the *imagePullSecrets field in the deployment configuration to provide the registry credentials to ArgoCD.
- D. Configure ArgoCD to use a Cron job that regularly checks for new image tags.
- E. Use the *imageUpdateStrategy field in the deployment configuration to define the update policy.

Antwort: A

Begründung:

ArgoCD can be configured to use webhooks, which are HTTP callbacks triggered by events like image pushes to container registries. This allows ArgoCD to automatically update the cluster whenever a new image tag is pushed, ensuring your deployments are always using the latest image versions.

28. Frage

.....

DeutschPrüfung verspricht den Kunden, dass Sie die Linux Foundation KCNA IT-Zertifizierungsprüfung 100% bestehen können. Die Qualität von DeutschPrüfung wird nach den IT-Experten überprüft. Das wichtigste Merkmal unserer Produkte ist ihre Relevanz. Der Schulungskurs dauert nur 20 Stunden. Und Sie werden die Linux Foundation KCNA Zertifizierungsprüfung dann mühelos bestehen. Wenn Sie DeutschPrüfung wählen, werden Sie dann sicher nicht bereuen. Denn es wird Ihnen Erfolg bringen.

KCNA Unterlage: <https://www.deutschpruefung.com/KCNA-deutsch-pruefungsfragen.html>

- [illegible]

Außerdem sind jetzt einige Teile dieser DeutschPrüfung KCNA Prüfungsfragen kostenlos erhältlich: <https://drive.google.com/open?id=1uujpaVsB9YnaWtl-7SujByPzJrcbml>