

KCNA Lernhilfe & KCNA Online Test



P.S. Kostenlose und neue KCNA Prüfungsfragen sind auf Google Drive freigegeben von ZertFragen verfügbar:
<https://drive.google.com/open?id=1IJYnLNgxZie7zYnrz0thrH-gwVkhtND5>

Sie können im Internet kostenlos die Lern Tipps und einen Teil der Prüfungsfragen und Antworten zur Linux Foundation KCNA Zertifizierungsprüfung von ZertFragen als Probe herunterladen.

Linux Foundation KCNA Exam Syllabus Topics:

Section	Weight	Objectives
Topic 1: Kubernetes Fundamentals	46%	- Kubernetes Resources • 1. Networking Resources (Services, Ingress) • 2. Workload Resources (Pods, Deployments, StatefulSets, DaemonSets, ReplicaSets, Jobs, CronJobs) • 3. Configuration Resources (ConfigMaps, Secrets) - Kubernetes Architecture • 1. Control Plane Components (API Server, etcd, Scheduler, Controller Manager) • 2. Worker Node Components (Kubelet, Kube-proxy, Container Runtime) - Containers • 1. Container Images and Registries • 2. Container Runtime Interface (CRI) • 3. Basic kubectl Commands - Scheduling • 1. Taints and Tolerations • 2. Node Selection and Affinity • 3. Resource Requests and Limits - Kubernetes API • 1. Declarative Management (Manifests/YAML) • 2. API Resource Structure and Versioning

Topic 2: Cloud Native Architecture	16%	- Cloud Native Landscape • 1. CNCF Role and Governance • 2. CNCF Project Categories (Sandbox, Incubating, Graduated) - Architecture Concepts • 1. Serverless and FaaS • 2. Autoscaling (HPA, VPA, Cluster Autoscaler) • 3. Microservices Architecture • 4. Elasticity and Resilience - Infrastructure and Practices • 1. Infrastructure as Code (IaC) • 2. Immutable Infrastructure • 3. DevOps Practices and Culture
Topic 3: Container Orchestration	22%	- Security • 1. Network Policies • 2. Pod Security Standards (Admission Control) • 3. RBAC (Role-Based Access Control) - Orchestration Fundamentals • 1. Self-healing and Rolling Updates • 2. Scheduling and Resource Management • 3. Service Discovery and Load Balancing - Networking • 1. CoreDNS and Service Networking • 2. Kubernetes Networking Model - Storage • 1. Storage Classes and Dynamic Provisioning • 2. Volumes, PersistentVolumes (PV), PersistentVolumeClaims (PVC) - Service Mesh • 1. Service Mesh Concepts (Istio, Linkerd) • 2. Sidecar Pattern and Traffic Management - Container Runtimes • 1. Docker, containerd, CRI-O
Topic 4: Cloud Native Application Delivery	8%	- CI/CD • 1. Artifact Management and Image Registries • 2. Continuous Integration and Continuous Delivery Pipelines - Deployment Strategies • 1. Blue/Green, Canary, Rolling Updates - GitOps • 1. GitOps Principles and Workflow • 2. Tools (Argo CD, Flux)

Topic 5: Cloud Native Observability	8%	- Tracing • 1. Distributed Tracing Concepts (OpenTelemetry, Jaeger) - Monitoring and Metrics • 1. Prometheus and Metrics Collection • 2. Dashboards and Visualization (Grafana) - Logging • 1. Kubernetes Logging Architecture • 2. Centralized Logging (Fluentd, Elasticsearch, Kibana)
-------------------------------------	----	---

>> KCNA Lernhilfe <<

KCNA Bestehen Sie Kubernetes and Cloud Native Associate! - mit höhere Effizienz und weniger Mühen

Wenn Sie die Produkte von ZertFragen kaufen, werden wir mit äußerster Kraft Ihnen helfen, die Linux Foundation KCNA Zertifizierungsprüfung zu bestehen. Außerdem bieten wir Ihnen einen einjährigen kostenlosen Update-Service. Wenn der Prüfungsplan von staatlicher Seite geändert werden, benachrichtigen wir die Kunden sofort. Wenn unsere Software neue Version hat, liefern wir den Kunden sofort. ZertFragen verspricht, dass Sie nur einmal die Linux Foundation KCNA Zertifizierungsprüfung bestehen können.

Linux Foundation Kubernetes and Cloud Native Associate KCNA Prüfungsfragen mit Lösungen (Q148-Q153):

148. Frage

Which of the following factors does scheduling take into account when selecting a Node?

- A. The number of existing Pods on a Node
- B. Services
- C. How many replicas there are in a Deployment
- **D. Resource requirements**

Antwort: D

Begründung:

Scheduling takes resource requirements into account in the form of resource requests.

149. Frage

In Kubernetes, what is the main purpose of applying the principle of least privilege when configuring access controls?

- **A. To ensure that users and applications have only the minimum level of access necessary to perform their tasks.**
- B. To simplify and standardize the management of user permissions across all namespaces in the cluster.
- C. To allow every user unrestricted access to all cluster resources for improved collaboration.
- D. To enforce strict separation of access boundaries between different workloads in the Kubernetes cluster.

Antwort: A

Begründung:

The principle of least privilege limits each user, service account, and application to only the permissions required for its assigned tasks. This reduces the potential impact of compromised credentials, configuration mistakes, and unauthorized actions.

150. Frage

What factors influence the Kubernetes scheduler when it places Pods on nodes?

- A. Node taints, node level, and Pod priority.
- B. Pod priority, container command, and node labels.
- C. Pod labels, node labels, and request labels.
- **D. Pod memory requests, node taints, and Pod affinity.**

Antwort: D

Begründung:

The Kubernetes scheduler chooses a node for a Pod by evaluating scheduling constraints and cluster state. Key inputs include resource requests (CPU/memory), taints/tolerations, and affinity/anti-affinity rules. Option A directly names three real, high-impact scheduling factors-Pod memory requests, node taints, and Pod affinity-so A is correct.

Resource requests are fundamental: the scheduler must ensure the target node has enough allocatable CPU/memory to satisfy the Pod's requests. Requests (not limits) drive placement decisions. Taints on nodes repel Pods unless the Pod has a matching toleration, which is commonly used to reserve nodes for special workloads (GPU nodes, system nodes, restricted nodes) or to protect nodes under certain conditions. Affinity and anti-affinity allow expressing "place me near" or "place me away" rules-e.g., keep replicas spread across failure domains or co-locate components for latency.

Option B includes labels, which do matter, but "request labels" is not a standard scheduler concept; labels influence scheduling mainly through selectors and affinity, not as a direct category called "request labels." Option C mixes a real concept (taints, priority) with "node level," which isn't a standard scheduling factor term. Option D includes "container command," which does not influence scheduling; the scheduler does not care what command the container runs, only placement constraints and resources.

Under the hood, kube-scheduler uses a two-phase process (filtering then scoring) to select a node, but the inputs it filters/scores include exactly the kinds of constraints in A. Therefore, the verified best answer is A.

151. Frage

How can you monitor the progress for an updated Deployment/DaemonSets/StatefulSets?

- A. kubectl rollout state
- B. kubectl rollout progress
- C. kubectl rollout watch
- **D. kubectl rollout status**

Antwort: D

Begründung:

To monitor rollout progress for Kubernetes workload updates (most commonly Deployments, and also StatefulSets and DaemonSets where applicable), the standard kubectl command is kubectl rollout status, which makes D correct.

Kubernetes manages updates declaratively through controllers. For a Deployment, an update typically creates a new ReplicaSet and gradually shifts replicas from the old to the new according to the strategy (e.g., RollingUpdate with maxUnavailable and maxSurge). For StatefulSets, updates may be ordered and respect stable identities, and for DaemonSets, an update replaces node-level Pods according to update strategy. In all cases, you often want a single command that tells you whether the controller has completed the update and whether the new replicas are available. kubectl rollout status queries the resource status and prints a progress view until completion or timeout.

The other commands listed are not the canonical kubectl subcommands. kubectl rollout watch, kubectl rollout progress, and kubectl rollout state are not standard rollout verbs in kubectl. The supported rollout verbs typically include status, history, undo, pause, and resume (depending on kubectl version and resource type).

Operationally, kubectl rollout status deployment/<name> is used during releases and troubleshooting to confirm that new Pods become ready, that old replicas scale down, and that the rollout does not stall due to readiness probe failures, insufficient resources, or policy constraints (e.g., PodDisruptionBudgets interacting with drains). It ties directly into Kubernetes' self-healing and declarative update model: you declare a new desired Pod template, and controllers reconcile toward it while rollout status reports that reconciliation's progress.

152. Frage

Which of the following is a lightweight tool that manages traffic flows between services, enforces access policies, and aggregates telemetry data, all without requiring changes to application code?

- **A. Linkerd**
- B. NetworkPolicy

P.S. Kostenlose und neue KCNA Prüfungsfragen sind auf Google Drive freigegeben von ZertFragen verfügbar:
<https://drive.google.com/open?id=1IJYnLNgxZie7zYnrz0thrH-gwVkhND5>