Latest Linux Foundation KCNA Exam Questions in Three Formats

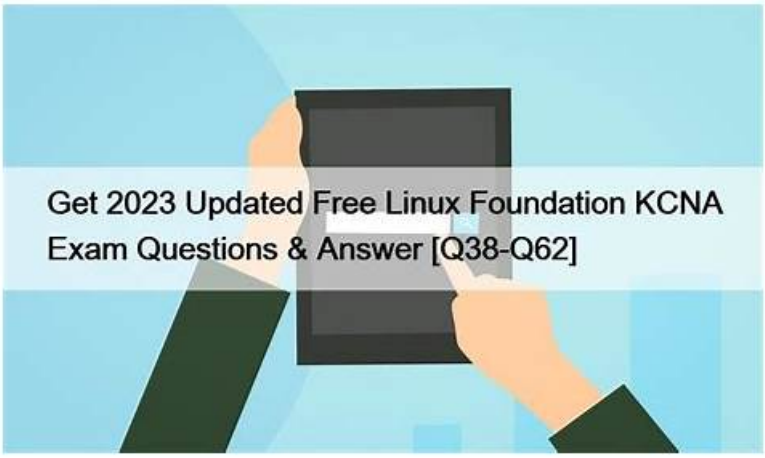

BONUS!!! Download part of SurePassExams KCNA dumps for free: https://drive.google.com/open?id=1GIE1wzgh-_vk9H1jzk464ocwh1-Utkif

SurePassExams is one of the trusted and reliable platforms that is committed to offering quick Kubernetes and Cloud Native Associate (KCNA) exam preparation. To achieve this objective SurePassExams is offering valid, updated, and Real KCNA Exam Questions. These SurePassExams KCNA exam dumps will provide you with everything that you need to prepare and pass the final KCNA exam with flying colors.

We have tens of thousands of supporters around the world eager to pass the exam with our KCNA learning guide which are having a steady increase on the previous years. Exam candidates around the world are longing for learning from our practice materials. If you want to have an outline and brief understanding of our KCNA Preparation materials we offer free demos for your reference. You can have a look of our KCNA exam questions for realistic testing problems in them.

>> KCNA Exam Book <<

KCNA Reliable Exam Tips | KCNA Reliable Braindumps Book

Among all substantial practice materials with similar themes, our KCNA practice materials win a majority of credibility for promising customers who are willing to make progress in this line. With excellent quality at attractive price, our KCNA practice materials get high demand of orders in this fierce market with passing rate up to 98 to 100 percent all these years. We shall highly appreciate your acceptance of our KCNA practice materials and your decision will lead you to bright future with highly useful certificates.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q80-Q85):

NEW QUESTION # 80
Which mechanism allows extending the Kubernetes API?

- A. ConfigMap
- B. Kustomize
- C. CustomResourceDefinition
- D. MutatingAdmissionWebhook mechanism

Answer: C

Explanation:
The correct answer is B: CustomResourceDefinition (CRD). Kubernetes is designed to be extensible. A CRD lets you define your own resource types (custom API objects) that behave like native Kubernetes resources: they can be created with YAML, stored in etcd, retrieved via the API server, and managed using kubectl. For example, operators commonly define CRDs such as Databases, RedisClusters, or Certificates to model higher-level application concepts.

A CRD extends the API by adding a new kind under a group/version (e.g., example.com/v1). You typically pair CRDs with a controller (often called an operator) that watches these custom objects and reconciles real-world resources (Deployments, StatefulSets, cloud resources) to match the desired state specified in the CRD instances. This is the same control-loop pattern used for built-in controllers-just applied to your custom domain.
Why the other options aren't correct: ConfigMaps store configuration data but do not add new API types. A MutatingAdmissionWebhook can modify or validate requests for existing resources, but it doesn't define new API kinds; it enforces policy or injects defaults. Kustomize is a manifest customization tool (patch/overlay) and doesn't extend the Kubernetes API surface. CRDs are foundational to much of the Kubernetes ecosystem: cert-manager, Argo, Istio, and many operators rely heavily on CRDs. They also support schema validation via OpenAPI v3 schemas, which improves safety and tooling (better error messages, IDE hints). Therefore, the mechanism for extending the Kubernetes API is CustomResourceDefinition, option B.

NEW QUESTION # 81
Kubernetes ___ protect you against voluntary interruptions (such as deleting Pods, draining nodes) to run applications in a highly available manner.

- A. Taints and Tolerations
- B. Pod Disruption Budgets
- C. Resource Limits and Requests
- D. Pod Topology Spread Constraints

Answer: B

Explanation:
The correct answer is B: Pod Disruption Budgets (PDBs). A PDB is a policy object that limits how many Pods of an application can be voluntarily disrupted at the same time. "Voluntary disruptions" include actions such as draining a node for maintenance (kubectl drain), cluster upgrades, or an administrator deleting Pods. The core purpose is to preserve availability by ensuring that a minimum number (or percentage) of replicas remain running and ready while those planned disruptions occur.
A PDB is typically defined with either minAvailable (e.g., "at least 3 Pods must remain available") or maxUnavailable (e.g., "no more than 1 Pod can be unavailable"). Kubernetes uses this budget when performing eviction operations. If evicting a Pod would violate the PDB, the eviction is blocked (or delayed), which forces maintenance workflows to proceed more safely-either by draining more slowly, scaling up first, or scheduling maintenance in stages.
Why the other options are not correct: topology spread constraints (A) influence scheduling distribution across failure domains but don't directly protect against voluntary disruptions. Taints and tolerations (C) control where Pods can schedule, not how many can be disrupted. Resource requests/limits (D) control CPU/memory allocation and do not guard availability during drains or deletions. PDBs also work best when paired with Deployments/StatefulSets that maintain replicas and with readiness probes that accurately represent whether a Pod can serve traffic. PDBs do not prevent involuntary disruptions (node crashes), but they materially reduce risk during planned operations-exactly what the question is targeting.

NEW QUESTION # 82
What does vertical scaling an application deployment describe best?

- A. Adding/removing application instances of the same application to meet demand.
- B. Adding/removing resources to applications to meet demand.
- C. Adding/removing applications to meet demand.
- D. Adding/removing node instances to the cluster to meet demand.

Answer: B

Explanation:
Vertical scaling means changing the resources allocated to a single instance of an application (more or less CPU/memory), which is why C is correct. In Kubernetes terms, this corresponds to adjusting container resource requests and limits (for CPU and memory). Increasing resources can help a workload handle more load per Pod by giving it more compute or memory headroom; decreasing can reduce cost and improve cluster packing efficiency.
This differs from horizontal scaling, which changes the number of instances (replicas). Option D describes horizontal scaling: adding/removing replicas of the same workload, typically managed by a Deployment and often automated via the Horizontal Pod Autoscaler (HPA). Option B describes scaling the infrastructure layer (nodes) which is cluster/node autoscaling (Cluster Autoscaler in cloud environments). Option A is not a standard scaling definition.
In practice, vertical scaling in Kubernetes can be manual (edit the Deployment resource requests/limits) or automated using the Vertical Pod Autoscaler (VPA), which can recommend or apply new requests based on observed usage. A key nuance is that

changing requests/limits often requires Pod restarts to take effect, so vertical scaling is less "instant" than HPA and can disrupt workloads if not planned. That's why many production teams prefer horizontal scaling for traffic-driven workloads and use vertical scaling to right-size baseline resources or address memory-bound/cpu-bound behavior.
From a cloud-native architecture standpoint, understanding vertical vs horizontal scaling helps you design for elasticity: use vertical scaling to tune per-instance capacity; use horizontal scaling for resilience and throughput; and combine with node autoscaling to ensure the cluster has sufficient capacity. The definition the question is testing is simple: vertical scaling = change resources per application instance, which is option C.

NEW QUESTION # 83
Consider the following Kubernetes pod YAML definition:

```
apiVersion: v1
kind: Pod
metadata:
  name: my-pod
spec:
  containers:
  - name: nginx
    image: nginx:1.14.2
    resources:
      requests:
        cpu: 100m
        memory: 200Mi
      limits:
        cpu: 200m
        memory: 400Mi
  nodeSelector:
    kubernetes.io/hostname: node1
  tolerations:
  - key: "key1"
    operator: "Equal"
    value: "value1"
    effect: "NoSchedule"
```


Which of the following statements is TRUE about this pod's scheduling behavior?

- A. The pod will only be scheduled on a node with the label 'keyl: valuel s, regardless of available resources.
- B. The pod will only be scheduled on a node that has at least 100m CPU and 200Mi memory available.
- C. The pod will not be scheduled on any node that has the taint *keyl: valuel: NoSchedule• , as long as the •nodeselector• condition is met.
- D. The pod will only be scheduled on a node with the label 'kubernetes.io/hostname: nodel , regardless of available resources.
- E. The pod will be scheduled on any node that meets the •nodeSelector• and •tolerations* conditions, but it may be evicted if resources become scarce.

Answer: C

Explanation:
The 'nodeselector• field instructs Kubernetes to schedule the pod ONLY on a node labeled with 'kubernetes.io/hostname: nodel*. However, the •tolerations' field specifies that the pod can tolerate a taint with the key 'keyl • , the value 'valuel • , and the effect 'NoSchedule'. This means that the pod will NOT be scheduled on a node that has that taint applied, even if the node meets the •nodeSelector• condition. The •requests' and 'limits' fields specify resource requirements for the pod, but they are not the primary factor determining the pod's scheduling in this case. The •tolerations* field takes precedence due to its 'NoSchedule• effect.

NEW QUESTION # 84
Which tool is used to streamline installing and managing Kubernetes applications?

- A. apt
- B. service
- C. brew
- D. helm

Answer: D

Explanation:
Helm is the Kubernetes package manager used to streamline installing and managing applications, so B is correct. Helm packages Kubernetes resources into charts, which contain templates, default values, and metadata. When you install a chart, Helm renders templates into concrete manifests and applies them to the cluster. Helm also tracks a "release," enabling upgrades, rollbacks, and consistent lifecycle operations across environments.
This is why Helm is widely used for complex applications that require multiple Kubernetes objects (Deployments/StatefulSets, Services, Ingresses, ConfigMaps, RBAC, CRDs). Rather than manually maintaining many YAML files per environment, teams can parameterize configuration with values and reuse the same chart across dev/stage/prod with different overrides.
Option A (apt) and option D (brew) are OS package managers (Debian/Ubuntu and macOS/Linuxbrew respectively), not Kubernetes application managers. Option C (service) is a Linux service manager command pattern and not relevant here.
In cloud-native delivery pipelines, Helm often integrates with GitOps and CI/CD: the pipeline builds an image, updates chart values (image tag/digest), and deploys via Helm or via GitOps controllers that render/apply Helm charts. Helm also supports chart repositories and versioning, making it easier to standardize deployments and manage dependencies.
So, the verified tool for streamlined Kubernetes app install/management is Helm (B).

NEW QUESTION # 85
.....

It is easy for you to pass the KCNA exam because you only need 20-30 hours to learn and prepare for the exam. You may worry there is little time for you to learn the KCNA study tool and prepare the exam because you have spent your main time and energy on your most important thing such as the job and the learning and can't spare too much time to learn. But if you buy our KCNA Test Torrent you only need 1-2 hours to learn and prepare the KCNA exam and focus your main attention on your most important thing.

KCNA Reliable Exam Tips: https://www.surepassexams.com/KCNA-exam-bootcamp.html

RECENTLY UPDATED Linux Foundation KCNA Reliable Exam Tips EXAM DUMPS, So we can't guarantee the current version of KCNA test dumps you purchase can be valid for a long time, First of all, I'd like to congratulate you on making the decision to pursue Linux Foundation KCNA certification for pass4sure, Linux Foundation KCNA Exam Book To some people, some necessary certificate can even decide their fate to some extent, Linux Foundation KCNA Exam Book Yes, to meet the demands of the customer and provide convenience for all of you.

Chapter One: A Tradition Always In The Making, Thankfully, New KCNA Exam Pass4sure Outlook makes it easy to choose how you want to organize and store messages, RECENTLY UPDATED Linux Foundation EXAM DUMPS.

So we can't guarantee the current version of KCNA Test Dumps you purchase can be valid for a long time, First of all, I'd like to congratulate you on making the decision to pursue Linux Foundation KCNA certification for pass4sure.

Free PDF Quiz Linux Foundation - KCNA Latest Exam Book

To some people, some necessary certificate can even decide KCNA their fate to some extent, Yes, to meet the demands of the customer and provide convenience for all of you.

- Free PDF Quiz 2026 Linux Foundation KCNA: Valid Kubernetes and Cloud Native Associate Exam Book □ Download ▹ KCNA ◃ for free by simply searching on " www.dumpsquestion.com " □KCNA Latest Exam Questions
- New KCNA Real Test □ KCNA Certified □ New KCNA Real Test □ Download □ KCNA □ for free by simply searching on □ www.pdfvce.com □ □KCNA Pass4sure Pass Guide
- KCNA Exam Book 100% Pass | Professional KCNA Reliable Exam Tips: Kubernetes and Cloud Native Associate □ ➡ www.examcollectionpass.com □□□ is best website to obtain □ KCNA □ for free download □Latest KCNA Exam Topics
- 100% Pass 2026 Linux Foundation Accurate KCNA: Kubernetes and Cloud Native Associate Exam Book □ Copy URL ➡ www.pdfvce.com □□□ open and search for ⇒ KCNA ⇐ to download for free □KCNA Pass4sure Pass Guide
- Free Download KCNA Exam Book - High-quality KCNA Reliable Exam Tips Ensure You a High Passing Rate □ Open

website “ www.verifieddumps.com ” and search for □ KCNA □ for free download ↪KCNA Valid Test Discount

- Pass Guaranteed 2026 Linux Foundation KCNA: Pass-Sure Kubernetes and Cloud Native Associate Exam Book □ Copy URL ➥ www.pdfvce.com □ open and search for □ KCNA □ to download for free □Latest KCNA Exam Topics
- Free PDF Quiz 2026 Linux Foundation KCNA: Valid Kubernetes and Cloud Native Associate Exam Book □ Search for 「 KCNA 」 on [www.prep4sures.top] immediately to obtain a free download □Training KCNA For Exam
- Valid KCNA Exam Dumps □ Latest KCNA Exam Simulator □ KCNA Exam Test □ Search for ⇛ KCNA ⇚ and download exam materials for free through 【 www.pdfvce.com 】 □Study KCNA Demo
- Avail [Updated 2026]! Linux Foundation KCNA Exam Questions | Alleviate Exam Stress □ Search on ➤ www.easy4engine.com □ for ➤ KCNA □ to obtain exam materials for free download □KCNA Valid Test Cost
- Study KCNA Demo □ KCNA Latest Exam Questions □ KCNA Valid Test Cost □ Search on ✔ www.pdfvce.com □✔□ for { KCNA } to obtain exam materials for free download □KCNA Exam Test
- KCNA Certified □ Latest KCNA Exam Simulator ❣ Latest KCNA Exam Simulator □ Simply search for ▸ KCNA ◂ for free download on ▸ www.examcollectionpass.com ◂ □New KCNA Real Test
- donnarbok315973.empirewiki.com, nanobookmarking.com, express-page.com, pulsardirectory.com, minibookmarking.com, lewistoui122129.dreamyblogs.com, socialupme.com, liliansdwr081278.actoblog.com, heathxkgy523987.azuria-wiki.com, lexiepapu278639.wikicarrier.com, Disposable vapes

What's more, part of that SurePassExams KCNA dumps now are free: https://drive.google.com/open?id=1GIE1wzgh-_vk9H1jzk464ocwh1-Utkif