

Latest Online Linux Foundation KCNA Practice Tests



P.S. Free & New KCNA dumps are available on Google Drive shared by PassLeaderVCE: <https://drive.google.com/open?id=1krfD8dgDVUKbdx20GvENnXavat5D56j>

With the Linux Foundation KCNA practice test, users can reduce stress, and improve their confidence to succeed. The desktop-based Kubernetes and Cloud Native Associate (KCNA) practice test software is compatible with Windows only. But the web-based KCNA Practice Test is compatible with all operating systems.

Passing KCNA exam is not very simple. KCNA exam requires a high degree of professional knowledge of IT, and if you lack this knowledge, PassLeaderVCE can provide you with a source of IT knowledge. PassLeaderVCE's expert team will use their wealth of expertise and experience to help you increase your knowledge, and can provide you practice questions and answers KCNA certification exam. PassLeaderVCE will not only do our best to help you pass the KCNA Certification Exam for only one time, but also help you consolidate your IT expertise. If you select PassLeaderVCE, we can not only guarantee you 100% pass KCNA certification exam, but also provide you with a free year of exam practice questions and answers update service. And if you fail to pass the examination carelessly, we can guarantee that we will immediately 100% refund your cost to you.

>> **KCNA Exam Details** <<

Easy To Use and Compatible PassLeaderVCE Linux Foundation KCNA Questions Formats

Infinite striving to be the best is man's duty. We have the responsibility to realize our values in the society. Of course, you must have enough ability to assume the tasks. Then our KCNA study materials can give you some help. First of all, you can easily pass the exam and win out from many candidates. The KCNA certificate is hard to get. If you really crave for it, our KCNA study materials are your best choice. We know it is hard for you to make decisions. You will feel sorry if you give up trying.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q140-Q145):

NEW QUESTION # 140

Which Kubernetes resource creates Kubernetes Jobs?

- A. JobFactory
- **B. CronJob**
- C. JobDeployment
- D. Task

Answer: B

Explanation:

<https://kubernetes.io/docs/concepts/workloads/controllers/cron-jobs/>

CronJob

FEATURE STATE: Kubernetes v1.21 [stable]

A CronJob creates Jobs on a repeating schedule.

One CronJob object is like one line of a *crontab* (cron table) file. It runs a job periodically on a given schedule, written in **Cron** format.

NEW QUESTION # 141

What is the purpose of the "securityContext" field in a pod specification? Choose all that apply.

- A. To define the user ID and group ID for the containers within the pod.
- B. To configure the resource limits for the containers within the pod.
- C. To control access to the Kubernetes API server.
- D. To specify the network namespace that the pod should be created in.
- E. To define the security context of the pods, such as SELinux and AppArmor.

Answer: A,E

Explanation:

The "securityContext" field in a pod specification is used to define the security context of the containers within the pod. This includes settings related to user ID and group ID, SELinux and AppArmor profiles, and other security-related options. It does not control access to the Kubernetes API server, configure resource limits, or specify the network namespace.

NEW QUESTION # 142

Which of the following command is used to get detailed information about the pod?

- A. kubectl info
- B. kubectl explain
- C. kubectl describe
- D. kubectl get

Answer: C

Explanation:

<https://kubernetes.io/docs/reference/generated/kubectl/kubectl-commands#describe>



NEW QUESTION # 143

In which framework do the developers no longer have to deal with capacity, deployments, scaling and fault tolerance, and OS?

- **A. Serverless**
- B. Mesos
- C. Docker Swarm
- D. Kubernetes

Answer: A

Explanation:

Serverless is the model where developers most directly avoid managing server capacity, OS operations, and much of the deployment/scaling/fault-tolerance mechanics, which is why D is correct. In serverless computing (commonly Function-as-a-Service, FaaS, and managed serverless container platforms), the provider abstracts away the underlying servers. You typically deploy code (functions) or a container image, define triggers (HTTP events, queues, schedules), and the platform automatically provisions the required compute, scales it based on demand, and handles much of the availability and fault tolerance behind the scenes.

It's important to compare this to Kubernetes: Kubernetes does automate scheduling, self-healing, rolling updates, and scaling, but it still requires you (or your platform team) to design and operate cluster capacity, node pools, upgrades, runtime configuration, networking, and baseline reliability controls. Even in managed Kubernetes services, you still choose node sizes, scale policies, and operational configuration. Kubernetes reduces toil, but it does not eliminate infrastructure concerns in the same way serverless does. Docker Swarm and Mesos are orchestration platforms that schedule workloads, but they also require managing the underlying capacity and OS-level aspects. They are not "no longer have to deal with capacity and OS" frameworks.

From a cloud native viewpoint, serverless is about consuming compute as an on-demand utility. Kubernetes can be a foundation for a serverless experience (for example, with event-driven autoscaling or serverless frameworks), but the pure framework that removes the most operational burden from developers is serverless.

NEW QUESTION # 144

There is an application running in a logical chain: Gateway API # Service # EndpointSlice # Container.
What Kubernetes API object is missing from this sequence?

- A. Docker
- B. Firewall
- C. Proxy
- **D. Pod**

Answer: D

Explanation:

In Kubernetes, application traffic flows through a well-defined set of API objects and runtime components before reaching a running container. Understanding this logical chain is essential for grasping how Kubernetes networking works internally.

The given sequence is: Gateway API # Service # EndpointSlice # Container. While this looks close to correct, it is missing a critical Kubernetes abstraction: the Pod. Containers in Kubernetes do not run independently; they always run inside Pods. A Pod is the smallest deployable and schedulable unit in Kubernetes and serves as the execution environment for one or more containers that share networking and storage resources.

The correct logical chain should be:

Gateway API # Service # EndpointSlice # Pod # Container

The Gateway API defines how external or internal traffic enters the cluster. The Service provides a stable virtual IP and DNS name, abstracting a set of backend workloads. EndpointSlices then represent the actual network endpoints backing the Service, typically mapping to the IP addresses of Pods. Finally, traffic is delivered to containers running inside those Pods.

Option A (Proxy) is incorrect because while proxies such as kube-proxy or data plane proxies play a role in traffic forwarding, they are not Kubernetes API objects that represent application workloads in this logical chain. Option B (Docker) is incorrect because Docker is a container runtime, not a Kubernetes API object, and Kubernetes is runtime-agnostic. Option D (Firewall) is incorrect because firewalls are not core Kubernetes workload or networking API objects involved in service-to-container routing. Option C (Pod) is the correct answer because Pods are the missing link between EndpointSlices and containers. EndpointSlices point to Pod IPs, and containers cannot exist outside of Pods. Kubernetes documentation clearly states that Pods are the fundamental unit of execution and networking, making them essential in any accurate representation of application traffic flow within a cluster.

NEW QUESTION # 145

.....

