Pass Guaranteed 2026 Linux Foundation KCNA–Valid New Test Answers

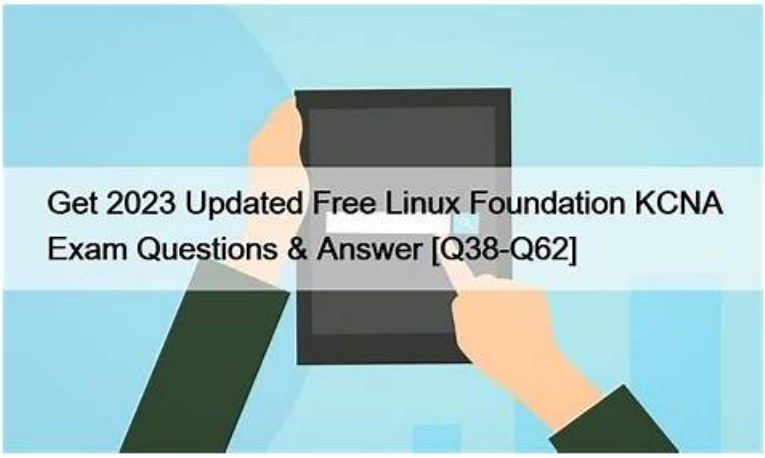

DOWNLOAD the newest BraindumpsPrep KCNA PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1IakcSoNV4dkhQYSycYIBEQii4jX2jMGL

No doubt the Kubernetes and Cloud Native Associate (KCNA) certification is one of the most challenging certification exams in the market. This Linux Foundation KCNA certification exam gives always a tough time to Kubernetes and Cloud Native Associate (KCNA) exam candidates. The BraindumpsPrep understands this hurdle and offers recommended and real Linux Foundation KCNA exam practice questions in three different formats.

Linux Foundation KCNA Exam validates one's knowledge of cloud-native technologies and adds credibility to their resume. Passing the exam demonstrates one's proficiency in Kubernetes-based application development and management and opens up new career opportunities in the cloud-native space. Kubernetes and Cloud Native Associate certification provides a reliable way to showcase one's skills and expertise in cloud-native technologies to employers and potential clients.

>> New KCNA Test Answers <<

100% Pass Quiz 2026 Linux Foundation Pass-Sure New KCNA Test Answers

Dear,do you tired of the study and preparation for the KCNA actual test? Here, we advise you to try the Linux Foundation KCNA online test which can simulate the real test environment and give an excellent study experience. You see, you can set the test time and get the score immediately after each test by using KCNA Online Test engine. With the interactive and intelligent functions of BraindumpsPrep KCNA online test, you will be interested in the study. Besides, the valid questions & verified answers can also ensure the 100% pass rate.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q121-Q126):

NEW QUESTION # 121
You're running a microservices application in Kubernetes, and you're concerned about the cost associated with resource usage. Which of the following approaches is most effective for reducing cost without impacting performance?

- A. Design your microservices to be stateless and idempotent, allowing them to be easily scaled and restarted without affecting application state.
- B. Utilize serverless functions to replace all microservices, as they scale automatically and only consume resources when active.
- C. Configure resource quotas for each namespace containing microservices to restrict their resource usage.
- D. Implement a distributed tracing system to monitor the resource consumption of individual microservices and identify areas for optimization.
- E. Use a horizontal pod autoscaler (HPA) to scale each microservice independently based on its specific resource needs.

Answer: A,D,E

Explanation:
Cost optimization for microservices involves understanding resource consumption and ensuring efficient scaling. Distributed tracing helps pinpoint resource-intensive microservices for optimization. HPA allows individual microservices to scale independently, ensuring optimal resource utilization. Designing microservices to be stateless and idempotent simplifies scaling and restarts, reducing potential resource waste. While serverless functions offer scaling benefits, they might not be suitable for all microservices. Resource quotas can restrict usage but don't address dynamic scaling needs.

NEW QUESTION # 122
Which key-value store is used to persist Kubernetes cluster data?

- A. ControlPlaneStore
- B. Redis
- C. etcd
- D. ZooKeeper

Answer: C

Explanation:
Kubernetes stores its cluster state (API objects) in etcd, making A correct. etcd is a distributed, strongly consistent key-value store that serves as the source of truth for the Kubernetes control plane. When you create or update objects such as Pods, Deployments, ConfigMaps, Secrets, or Nodes, the kube-apiserver validates the request and then persists the desired state into etcd. Controllers and the scheduler watch the API for changes (which ultimately reflect etcd state) and reconcile the cluster to match that desired state.
etcd's consistency guarantees are crucial. Kubernetes relies on accurate, up-to-date state to make scheduling decisions, enforce RBAC/admission policies, coordinate leader elections, and ensure controllers behave correctly. etcd uses the Raft consensus algorithm to replicate data among members and requires quorum for writes, enabling fault tolerance when deployed in HA configurations (commonly three or five members).
The other options are incorrect in Kubernetes' standard architecture. ZooKeeper is a distributed coordination system used by some other platforms, but Kubernetes does not use it as its primary datastore. Redis is an in-memory data store used for caching or messaging, not as Kubernetes' authoritative state store. "ControlPlaneStore" is not a standard Kubernetes component.
Operationally, etcd health is one of the most important determinants of cluster reliability. Slow disk I/O or unstable networking can degrade etcd performance and cause API latency spikes. Backup and restore procedures for etcd are critical disaster-recovery practices, and securing etcd (TLS, access restrictions) is essential because it may contain sensitive data (e.g., Secrets-often base64-encoded, and optionally encrypted at rest depending on configuration).
Therefore, the verified Kubernetes datastore is etcd, option A.

NEW QUESTION # 123
Which of the following is a definition of Hybrid Cloud?

- A. A cloud native architecture that uses services running in different public and private clouds, including on-premises data centers.
- B. A combination of services running in public and private data centers, only including data centers from the same cloud provider.
- C. A cloud native architecture that uses services running in public clouds, excluding data centers in different availability zones.
- D. A combination of services running in public and private data centers, excluding serverless functions.

Answer: A

Explanation:
A hybrid cloud architecture combines public cloud and private/on-premises environments, often spanning multiple infrastructure domains while maintaining some level of portability, connectivity, and unified operations. Option C captures the commonly accepted definition: services run across public and private clouds, including on-premises data centers, so C is correct.
Hybrid cloud is not limited to a single cloud provider (which is why A is too restrictive). Many organizations adopt hybrid cloud to meet regulatory requirements, data residency constraints, latency needs, or to preserve existing investments while still using public cloud elasticity. In Kubernetes terms, hybrid strategies often include running clusters both on-prem and in one or more public clouds, then standardizing deployment through Kubernetes APIs, GitOps, and consistent security/observability practices.
Option B is incorrect because excluding data centers in different availability zones is not a defining property; in fact, hybrid

deployments commonly use multiple zones/regions for resilience. Option D is a distraction: serverless inclusion or exclusion does not define hybrid cloud. Hybrid is about the combination of infrastructure environments, not a specific compute model.
A practical cloud-native view is that hybrid architectures introduce challenges around identity, networking, policy enforcement, and consistent observability across environments. Kubernetes helps because it provides a consistent control plane API and workload model regardless of where it runs. Tools like service meshes, federated identity, and unified monitoring can further reduce fragmentation.
So, the most accurate definition in the given choices is C: hybrid cloud combines public and private clouds, including on-premises infrastructure, to run services in a coordinated architecture.

NEW QUESTION # 124
If a Pod was waiting for container images to download on the scheduled node, what state would it be in?

- A. Unknown
- B. Pending
- C. Failed
- D. Succeeded

Answer: B

Explanation:
If a Pod is waiting for its container images to be pulled to the node, it remains in the Pending phase, so D is correct. Kubernetes Pod "phase" is a high-level summary of where the Pod is in its lifecycle. Pending means the Pod has been accepted by the cluster but one or more of its containers has not started yet. That can occur because the Pod is waiting to be scheduled, waiting on volume attachment/mount, or-very commonly-waiting for the container runtime to pull the image.
When image pulling is the blocker, kubectl describe pod <name> usually shows events like "Pulling image ..." and "Successfully pulled image ..." or failures like ImagePullBackOff/ErrImagePull. Even if the node has been assigned (scheduler has set spec.nodeName), the Pod can still be Pending while kubelet and the runtime perform preparation steps.
Why the other phases don't apply:
Succeeded is for run-to-completion Pods that have finished successfully (typical for Jobs).
Failed means the Pod terminated and at least one container terminated in failure (and won't be restarted, depending on restartPolicy).
Unknown is used when the node can't be contacted and the Pod's state can't be reliably determined (rare in healthy clusters).
A subtle but important Kubernetes detail: status "Waiting" reasons like ImagePullBackOff are container states inside .status.containerStatuses, while the Pod phase can still be Pending. So, "waiting for images to download" maps to Pod Pending, with container waiting reasons providing the deeper diagnosis.
Therefore, the verified correct answer is D: Pending.

NEW QUESTION # 125
Which organizational persona creates Service Level Agreements 'SLA', Service Level Objectives 'SLO', and Service Level Indicator 'SLI'?

- A. DevSecOps
- B. Site Reliability Engineer (SRE)
- C. Security and Compliance Engineer
- D. DevOps
- E. Developer

Answer: B

Explanation:
SREs create SLAs, SLOs, and SLIs to define and implement standards for application and infra-structure reliability.

NEW QUESTION # 126
.....

It is universally acknowledged that KCNA certification can help present you as a good master of some knowledge in certain areas, and it also serves as an embodiment in showcasing one’s personal skills. However, it is easier to say so than to actually get the

KCNA certification. We have to understand that not everyone is good at self-learning and self-discipline, and thus many people need outside help to cultivate good study habits, especially those who have trouble in following a timetable. Buy our KCNA Exam Questions, we will help you pass the KCNA exam without difficulty.

KCNA Exam Simulations: https://www.briandumpsprep.com/KCNA-prep-exam-braindumps.html

DOWNLOAD the newest BraindumpsPrep KCNA PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1IakcSoNV4dkhQYSycYIBEQii4jX2jMGL