

Providing You Updated KCNA New Braindumps Pdf with 100% Passing Guarantee



What's more, part of that ActualTestsQuiz KCNA dumps now are free: <https://drive.google.com/open?id=1IIdM68fqCKFI5B3cF5scqNwEmkR5PkD>

As is known to us, people who want to take the KCNA exam include different ages, different fields and so on. It is very important for company to design the KCNA exam prep suitable for all people. However, our company has achieved the goal. We can promise that the KCNA test questions from our company will be suitable all people. There are many functions about our study materials beyond your imagination. You can purchase our KCNA reference guide according to your own tastes. We believe that the understanding of our study materials will be very easy for you. We hope that you can choose the KCNA test questions from our company, because our products know you better.

Before purchasing KCNA prep torrent, you can log in to our website for free download. During your installation, KCNA exam torrent hired dedicated experts to provide you with free online guidance. During your studies, KCNA exam torrent also provides you with free online services for 24 hours, regardless of where and when you are, as long as an email, we will solve all the problems for you. At the same time, if you fail to pass the exam after you have purchased KCNA prep torrent, you just need to submit your transcript to our customer service staff and you will receive a full refund.

>> **KCNA New Braindumps Pdf** <<

2026 Linux Foundation KCNA –High Pass-Rate New Braindumps Pdf

Our company is a multinational company with sales and after-sale service of KCNA exam torrent compiling departments throughout the world. In addition, our company has become the top-notch one in the fields, therefore, if you are preparing for the exam in order to get the related certification, then the Kubernetes and Cloud Native Associate exam question compiled by our company is your solid choice. All employees worldwide in our company operate under a common mission: to be the best global supplier of electronic KCNA Exam Torrent for our customers through product innovation and enhancement of customers' satisfaction. Wherever you are in the world we will provide you with the most useful and effectively KCNA guide torrent in this website, which will help you to pass the exam as well as getting the related certification with a great ease.

The Linux Foundation KCNA exam covers a wide range of topics related to cloud native computing, such as containerization, microservices, and orchestration. It also tests candidates on their understanding of Kubernetes architecture, deployment, and management. KCNA Exam is comprised of 50 multiple choice questions and candidates have 90 minutes to complete it.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q125-Q130):

NEW QUESTION # 125

You are running a multi-cluster Kubernetes environment with Istio deployed. You want to enable mutual TLS authentication between

services across different clusters. You have already configured Istio's 'mutualTLS' setting in the control plane. What additional step is required to enforce this security measure for inter-cluster communication?

- A. Configure the 'gateway' setting in Istio to allow only encrypted traffic between clusters
- B. Create a separate Kubernetes namespace for inter-cluster communication and configure Istio within that namespace
- C. Deploy a dedicated Istio sidecar proxy in each cluster to handle inter-cluster communication
- **D. Manually configure the 'istio.io/auth' annotation on all services in both clusters**
- E. Deploy a dedicated Istio control plane for each cluster and configure cross-cluster communication

Answer: D

Explanation:

While configuring 'mutualTLS' in the Istio control plane enables the security feature, you need to explicitly enable it for individual services by applying the 'istio.io/auth' annotation with 'peerAuthentication: mutualTLS: { mode: STRICT Y' to each service. This ensures that services in both clusters will only communicate with each other over encrypted channels- Options B, C, D, and E are not directly related to enforcing mutual TLS authentication for inter-cluster communication.

NEW QUESTION # 126

You have a Kubernetes cluster running on AWS. You need to ensure that only approved container images are used in your cluster. Which Kubernetes feature can you use to enforce this policy?

- A. Network Policies
- **B. Admission Controllers**
- C. Service Accounts
- D. Resource Quotas
- E. Pod Security Policies

Answer: B

Explanation:

Admission Controllers in Kubernetes can be used to enforce policies for container images. You can configure an Admission Controller to check if the image is present in an approved image registry or if it meets certain security criteria. This helps prevent unauthorized or insecure images from being deployed to your cluster.

NEW QUESTION # 127

What does "Continuous Integration" mean?

- A. The continuous integration of changes from one environment to another.
- B. The continuous integration of new tools to support developers in a project.
- **C. The continuous integration and testing of code changes from multiple sources via automation.**
- D. The continuous integration and testing of code changes from multiple sources manually.

Answer: C

Explanation:

The correct answer is B: Continuous Integration (CI) is the practice of frequently integrating code changes from multiple contributors and validating them through automated builds and tests. The "continuous" part is about doing this often (ideally many times per day) and consistently, so integration problems are detected early instead of piling up until a painful merge or release window.

Automation is essential. CI typically includes steps like compiling/building artifacts, running unit and integration tests, executing linters, checking formatting, scanning dependencies for vulnerabilities, and producing build reports. This automation creates fast feedback loops that help developers catch regressions quickly and maintain a releasable main branch.

Option A is incorrect because manual integration/testing does not scale and undermines the reliability and speed that CI is meant to provide. Option C confuses CI with deployment promotion across environments (which is more aligned with Continuous Delivery/Deployment). Option D is unrelated: adding tools can support CI, but it isn't the definition.

In cloud-native application delivery, CI is tightly coupled with containerization and Kubernetes: CI pipelines often build container images from source, run tests, scan images, sign artifacts, and push to registries. Those validated artifacts then flow into CD processes that deploy to Kubernetes using manifests, Helm, or GitOps controllers. Without CI, Kubernetes rollouts become riskier because you lack consistent validation of what you're deploying.

So, CI is best defined as automated integration and testing of code changes from multiple sources, which matches option B.

NEW QUESTION # 128

What is a Kubernetes Service Endpoint?

- A. It is an IP address that we can access from the Internet.
- B. It is a name of special Pod in kube-system namespace.
- **C. It is an object that gets IP addresses of individual Pods assigned to it.**
- D. It is the API endpoint of our Kubernetes cluster.

Answer: C

Explanation:

A Kubernetes Service routes traffic to a dynamic set of backends (usually Pods). The set of backend IPs and ports is represented by endpoint-tracking resources. Historically this was the Endpoints object; today Kubernetes commonly uses EndpointSlice for scalability, but the concept remains the same: endpoints represent the concrete network destinations behind a Service. That's why D is correct: a Service endpoint is an object that contains the IP addresses (and ports) of the individual Pods (or other backends) associated with that Service.

When a Service has a selector, Kubernetes automatically maintains endpoints by watching which Pods match the selector and are Ready, then publishing those Pod IPs into Endpoints/EndpointSlices. Consumers don't usually use endpoints directly; instead they call the Service DNS name, and kube-proxy (or an alternate dataplane) forwards traffic to one of the endpoints. Still, endpoints are critical because they are what make Service routing accurate and up to date during scaling events, rolling updates, and failures.

Option A confuses this with the Kubernetes API server endpoint (the cluster API URL). Option B is incorrect; there's no special "Service Endpoint Pod." Option C describes an external/public IP concept, which may exist for LoadBalancer Services, but "Service endpoint" in Kubernetes vocabulary is about the backend destinations, not the public endpoint.

Operationally, endpoints are useful for debugging; if a Service isn't routing traffic, checking Endpoints/EndpointSlices shows whether the Service actually has backends and whether readiness is excluding Pods.

This ties directly into Kubernetes service discovery and load balancing: the Service is the stable front door; endpoints are the actual backends.

NEW QUESTION # 129

How many hosts are required to set up a highly available Kubernetes cluster when using an external etcd topology?

- A. Four hosts. Two for control plane nodes and two for etcd nodes.
- B. Four hosts. One for a control plane node and three for etcd nodes.
- C. Three hosts. The control plane nodes and etcd nodes share the same host.
- **D. Six hosts. Three for control plane nodes and three for etcd nodes.**

Answer: D

Explanation:

In a highly available (HA) Kubernetes control plane using an external etcd topology, you typically run three control plane nodes and three separate etcd nodes, totaling six hosts, making D correct. HA design relies on quorum-based consensus: etcd uses Raft and requires a majority of members available to make progress. Running three etcd members is the common minimum for HA because it tolerates one member failure while maintaining quorum (2/3).

In the external etcd topology, etcd is decoupled from the control plane nodes. This separation improves fault isolation: if a control plane node fails or is replaced, etcd remains stable and independent; likewise, etcd maintenance can be handled separately.

Kubernetes API servers (often multiple instances behind a load balancer) talk to the external etcd cluster for storage of cluster state. Options A and B propose four hosts, but they break common HA/quorum best practices. Two etcd nodes do not form a robust quorum configuration (a two-member etcd cluster cannot tolerate a single failure without losing quorum). One control plane node is not HA for the API server/scheduler/controller-manager components. Option C describes a stacked etcd topology (control plane + etcd on same hosts), which can be HA with three hosts, but the question explicitly says external etcd, not stacked. In stacked topology, you often use three control plane nodes each running an etcd member. In external topology, you use three control plane + three etcd.

Operationally, external etcd topology is often used when you want dedicated resources, separate lifecycle management, or stronger isolation for the datastore. It can reduce blast radius but increases infrastructure footprint and operational complexity (TLS, backup/restore, networking). Still, for the canonical HA external-etcd pattern, the expected answer is six hosts: 3 control plane + 3 etcd.

