KCNA Exam Material & Test KCNA Sample Online

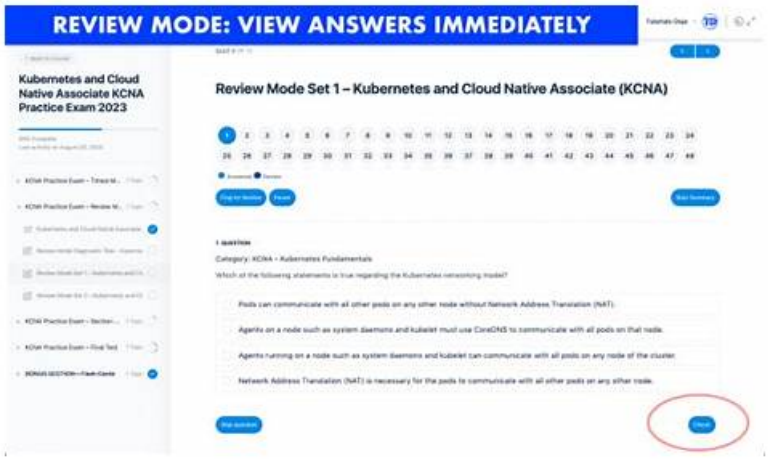

BTW, DOWNLOAD part of Pass4Leader KCNA dumps from Cloud Storage: https://drive.google.com/open?id=15bG5YNhEb1RmyiLyoKiv9yPsg7-543RB

For certificates who will attend the exam, some practice is evitable. But sometimes, time for preparation is quite urgent. KCNA exam braindumps of us will help you to use the least time to pass the exam. If you choose the KCNA exam dumps of us, you just need to spend about 48 to 72 hours to practice and you can pass the exam successfully. In addition, KCNA Exam Dumps are verified by experienced experts, and the accuracy and correctness can be guaranteed. And we pass guarantee and money back guarantee if can't pass the exam.

Linux Foundation KCNA (Kubernetes and Cloud Native Associate) Certification Exam is a globally recognized certification program designed for IT professionals who want to demonstrate their skills and knowledge in container orchestration and cloud-native technologies. KCNA Exam provides a comprehensive assessment of an individual's understanding of Kubernetes and other cloud-native technologies, including containerization, microservices, and networking.

>> KCNA Exam Material <<

KCNA Kubernetes and Cloud Native Associate Exam Material - Free PDF Realistic Linux Foundation KCNA

We are a team of certified professionals with lots of experience in editing KCNA exam questions. Every candidate should have more than 11 years' education experience in this filed of KCNA study guide. We have rather a large influence over quite a quantity of candidates. We are more than more popular by our high passing rate and high quality of our KCNA Study Guide. Our education team of professionals will give you the best of what you deserve. If you are headache about your KCNA certification exams, our KCNA training materials will be your best select.

Linux Foundation Kubernetes and Cloud Native Associate Sample Questions (Q161-Q166):

NEW QUESTION # 161
You are monitoring a Kubernetes cluster using Prometheus. You notice a sudden spike in the number of requests to a specific pod, followed by a corresponding increase in pod CPU utilization. Which Prometheus query can you use to identify the specific endpoint being heavily accessed?

- A. kube_pod_status_phase{pod=""}
- B. kube_pod_container_status_restart_count{pod=""}
- C. kube_deployment_status_replicas{deployment=""}
- D. kube_pod_container_resource_requests_cpu_cores{pod=""}
- E. kube_http_server_requests_seconds_bucket{job="", method="GET", uri="/specific_endpoint"}

Answer: E

Explanation:
The correct answer is E . The query 'kube_http_server_requests_seconds_bucket{job="", method="GET", uri="/specific_endpoint"}" targets HTTP requests made to a specific endpoint within a pod (identified by the •job', •method', and 'uri' labels). This query allows you to see the number of requests made to the '/specific_endpoint• and identify spikes in traffic to that endpoint. The other options are not relevant to identifying specific endpoints being heavily accessed. Option A shows CPU resource requests for a pod, option B shows the pod's phase, option C shows the number of replicas in a deployment, and option D shows the restart count for a pod's containers. These queries do not provide information about specific endpoints and traffic patterns.

NEW QUESTION # 162
Consider the following Istio configuration snippet:

```
apiVersion: networking.istio.io/v1alpha3
kind: VirtualService
metadata:
  name: productcatalog
  namespace: default
spec:
  hosts:
  - productcatalog.example.com
  gateways:
  - istio-system/gateway
  http:
  - match:
    - uri:
      prefix: /products
    route:
    - destination:
        host: productcatalog.example.com
        port:
          number: 80
    - destination:
        host: productcatalog-v2.example.com
        port:
          number: 80
      weight: 50
```


What does this configuration achieve?

- A. It configures a virtual service that only routes traffic to the 'productcatalog-v2.example.com' service for requests that match the '/products' path
- B. It creates a virtual service that routes all requests to the 'productcatalog.example-com' service only if they match the URI prefix '/products'
- C. It enables Istio's 'fault injection' feature to simulate failures and test the service's resilience.
- D. It defines a virtual service named 'productcatalog' in the 'default' namespace and forwards all traffic to the 'productcatalog.example-com' service-
- E. It exposes the 'productcatalog' service on port 80 using the 'istio-system/gateway' gateway and routes traffic to both 'productcatalog.example-com' and 'productcatalog-v2.example.com' with a 50/50 weight distributiom

Answer: E

Explanation:
This Istio VirtualService configuration exposes the 'productcatalog' service on port 80 using the 'istio- system/gateway' gateway It routes traffic to both 'productcatalog.example.com' and 'productcatalog-v2.example.com' with a 50/50 weight distribution for requests matching the 'Iproducts' prefix. This configuration demonstrates Istio's ability to implement dynamic traffic routing and load balancing between services.

NEW QUESTION # 163
You are using Istio to manage traffic flow in a microservices architecture. Your application consists of three services: 'service-A', 'service-B', and 'service-C', where 'service-A' depends on 'service-B' and 'service-B' depends on 'service-C'. You want to implement a canary rollout for 'service-B' to test a new version before rolling it out to all users. How would you configure Istio to achieve this?

- A. Deploy the new version of 'service-B' with a different Kubernetes namespace and use Istio to route traffic to the new namespace
- B. Use the 'istio.io/canary' annotation on 'service-B' and configure the rollout percentage
- C. Use the 'istio.io/route' annotation on 'service-B' to specify a weighted distribution of traffic to different versions
- D. Create a new Istio VirtualService for 'service-B' and route a percentage of traffic to the new version
- E. Configure 'service-A' to automatically retry requests if they fail to 'service-B' and route to the new version

Answer: D

Explanation:
Istio's VirtualService allows you to define routing rules that can direct a percentage of traffic to a specific version of a service- This enables you to perform canary rollouts by gradually introducing the new version of 'service-B' to a subset of users- Option A is not a valid Istio configuration option- Option C creates unnecessary complexity Option D is related to retries, not canary rollouts. Option E is closer but doesn't specify the percentage split.

NEW QUESTION # 164
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. Admission Controllers
- B. Pod Security Policies
- C. Service Accounts
- D. Resource Quotas
- E. Network Policies

Answer: A

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 # 165
You are building a serverless application on Google Cloud Functions that needs to interact with a Google Cloud SQL database. How can you securely connect your Cloud Functions to the Cloud SQL instance?

- A. Use Google Cloud Functions' built-in database connector to access Cloud SQL.
- B. Configure a Cloud SQL proxy service to forward connections to the Cloud SQL instance.
- C. Use a service account with appropriate permissions to connect to Cloud SQL from Cloud Functions.
- D. Utilize Google Cloud IAM to grant Cloud Functions access to the Cloud SQL instance.
- E. Use a dedicated virtual machine (VM) to host a database connector that communicates with both Cloud Functions and Cloud SQL.

Answer: C

Explanation:

The most secure and recommended approach is to use a service account with appropriate permissions to connect to Cloud SQL from Cloud Functions. This method provides granular control over access and ensures that Cloud Functions only have the necessary permissions to interact with the database. Option A is not a standard feature. Option B is an alternative but adds complexity. Option C introduces extra infrastructure and potential security risks. Option D, while important for overall security, doesnt specifically address database connections.

NEW QUESTION # 166

.....

Only with high quality and valid information of our KCNA exam braindumps, our candidates can successfully pass their exams. At the same time, own to our professional experts constantly improvement on the design of the KCNA study materials, we have developed three versions of layouts: PDF, Software and APP online. Though the content of them are the same, the different layouts provide lots of conveniences out of your imagination. Just have a try and you will love our KCNA Practice Engine.