Free Questions for 1Z0-1084-23 by certscare

Shared by Gardner on 29-01-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

How are cloud native application versions deployed to an OKE cluster when using a blue/green deployment strategy?

Options:

A- Current applications are slowly replaced with new application versions.

B- New application versions are deployed in minor increments to a select group of people.

C- Both old and new application versions are deployed to production at the same time.

Answer:

C

Explanation:

Blue/Green deployment strategy allows releasing a new version of an application using two identical environments where one of them is active at a given time.The current version of the application is provisioned on the active environment, whereas the new version gets deployed to the standby environment1.The traffic is shifted from the active to the standby environment by updating the ingress resource2. Therefore, both old and new application versions are deployed to production at the same time, but only one of them receives

the traffic. Verified Reference:Announcing new deployment strategies for OCI DevOps Service,Blue-Green OKE Deployment

Question 2

Question Type: MultipleChoice

You have been asked to update an OKE cluster to a network configuration that has the least attack surface while the deployed applications are still directly available for access from the Internet. Which is a valid OKE cluster network configuration that meets this requirement? (Choose the best answer.)

Options:

A- Private subnets for nodes, the Kubemetes API endpoint, and load balancers

B- Private subnets for nodes; public subnets for the Kubemetes API endpoint and load balancers

C- Private subnets for nodes and the Kubemetes API endpoint; public subnets for load balancers

D- Private subnet for the Kubemetes API endpoint; public subnets for nodes and load balancers

Answer:

C

Explanation:

The valid OKE cluster network configuration that meets the requirement of having the least attack surface while still allowing direct access to the deployed applications from the Internet is: Private subnets for nodes and the Kubernetes API endpoint; public subnets for load balancers. By placing the nodes and the Kubernetes API endpoint in private subnets, they are not directly accessible from the Internet, reducing the attack surface. The load balancers, on the other hand, are placed in public subnets, allowing them to be accessed from the Internet and serve as the entry point for accessing the deployed applications. This configuration ensures that the critical components of the cluster, such as the nodes and the API endpoint, are protected within the private network, while still providing accessibility to the applications through the load balancers. It helps to enhance security by limiting direct access to the internal components of the cluster while maintaining the availability of the deployed applications.

Question 3

Question Type: MultipleChoice

You deployed a Python application to an Oracle Container Engine for Kubernetes (OKE) cluster. However, while testing you found a bug, which you rectified and then created a new Docker image. You now need to ensure that if this new image does not work once deployed, you should be able to roll back to the previous version. Using kubect1, which strategy should you use?

Options:

A- Blue/Green Deployment

B- Canary Deployment

C- Rolling Update

D- A/B Testing

Answer:

C

Explanation:

A rolling update is a deployment strategy that gradually replaces the old version of an application with the new version without any downtime4.OKE supports rolling updates by using the kubectl rollout command4.A rolling update allows you to roll back to the previous version if something goes wrong with the new version4. Therefore, using a rolling update strategy with kubectl ensures that you can roll back to the previous version of your Python application if the new image does not work once deployed. Verified Reference:Deploy Oracle Container Engine for Kubernetes

Question 4

Question Type: MultipleChoice

A company is developing a new application that needs to process transactions in real time. The company wants to ensure that all transactions are processed in order and that no transaction is lost. Which of these is a correct strategy for leveraging OCI Queue in this scenario?

Options:

A- Use a separate queue for each type of transaction.

B- Use a single queue to process all transactions.

C- Use a separate queue for each application instance.

D- Use a priority queue to prioritize requests.

Answer:

B

Explanation:

OCI Queue is a service for enabling asynchronous (decoupled) communication in a serverless manner3.Queue handles high-volume transactional data that requires independent processing without loss or duplication3.Queue supports ordering of messages within a queue by using the FIFO (first-in-first-out) delivery option3. Therefore, using a single queue to process all transactions ensures that all transactions are processed in order and that no transaction is lost. Verified Reference:Overview of Queue

Question 5

Question Type: MultipleChoice

You are developing a real-time monitoring application for a fleet of vehicles, which will be deployed on Oracle Cloud Infrastructure (OCI). You need to choose between using OCI Queue or OCI Streaming to handle the real-time data feeds from the vehicles. Based on the scenario described, which is the most appropriate choice for handling real-time data feeds?

Options:

A- OCI Streaming, because it is designed for high-volume, continuous ingestion and processing of data, making it the best choice for a fleet of vehicles

B- OCI Streaming, because it offers exactly-once message delivery, which is necessary for real-time applications

C- OCI Queue, because it is optimized for low-latency messaging and ideal for real-time applications

D- OCI Queue, because it provides at-least-once message delivery, which is critical for real-time monitoring applications

Answer:

A

Explanation:

OCI Streaming is a fully managed, scalable, and durable messaging solution for ingesting continuous, high-volume streams of data that you can consume and process in real-time1.Streaming is suitable for any use case in which data is produced and processed continually and sequentially in a publish-subscribe messaging model1.Streaming can handle millions of messages per second with low latency2. Therefore, OCI Streaming is the most appropriate choice for handling real-time data feeds from a fleet of vehicles. Verified Reference:Overview of Streaming,Container Engine for Kubernetes

Question 6

Question Type: MultipleChoice

You are tasked with developing an application that requires the use of Oracle Cloud Infrastructure (OCI) APIs to POST messages to a stream in the OCI Streaming service. Which statement is incorrect? (Choose the best answer.)

Options:

A- The Content-Type header must be set to application/json

B- The request must include an authorization signing string including (but not limited to) x-content-sha256, content-type, and content-length headers.

C- The request does not require an Authorization header.

D- An HTTP 401 will be returned if the client's clock is skewed more than 5 minutes from the server's.

Answer:

C

Explanation:

The statement that is incorrect is: 'The request does not require an Authorization header.' In order to POST messages to a stream in the OCI Streaming service using OCI APIs, the request does require an Authorization header. The Authorization header is used to provide authentication and ensure the request is authorized to access the stream. The correct approach is to include the Authorization header in the request, along with other required headers such as x-content-sha256, content-type, and content-length. Therefore, the incorrect statement is that the request does not require an Authorization header.\

Question 7

Question Type: MultipleChoice

As a developer, you have been tasked with implementing a microservices-based application. Which THREE technologies are best suited to accomplish the task? (Choose three.)

Options:

A- Terraform

B- Big Data

C- Anomaly Detection

D- Service Mesh

E- Docker

F- Kubemetes

Answer:

D, E, F

Explanation:

The three technologies best suited for implementing a microservices-based application are: Service Mesh: A service mesh is a dedicated infrastructure layer that provides features like service discovery, load balancing, encryption, authentication, and observability for microservices. It helps in managing the communication and interactions between microservices in a scalable and secure manner. Kubernetes: Kubernetes is an open-source container orchestration platform that enables the deployment, scaling, and management of containerized applications. It provides features like automated scaling, service discovery, load balancing, and self-healing capabilities, which are essential for managing microservices in a distributed environment. Docker: Docker is a popular containerization platform that allows packaging applications and their dependencies into lightweight containers. It provides a consistent and portable environment for

running microservices, enabling easy deployment and scalability. Docker also facilitates isolation and resource efficiency, making it an ideal choice for deploying microservices. While Big Data, Anomaly Detection, and Terraform are valuable technologies, they are not specifically focused on enabling the implementation of microservices-based applications.

Question 8

Question Type: MultipleChoice

What is the difference between continuous delivery and continuous deployment in the DevOps methodology? (Choose the best answer.)

Options:

A- Continuous delivery involves automation of developer tasks, whereas continuous deployment involves manual operational tasks.

B- Continuous delivery requires automatic linting, whereas continuous deployment testing must be run manually.

C- Continuous delivery utilizes automatic deployment to a development environment, whereas continuous deployment involves automatic deployment to a production environment.

D- Continuous delivery is a process that Initiates deployment manually, whereas continuous deployment is based on automating the deployment process.

Answer:

C

Explanation:

The two correct differences between continuous delivery and continuous deployment in the DevOps lifecycle are: Continuous delivery is a process that initiates deployment manually, while continuous deployment is based on automating the deployment process. In continuous delivery, the software is ready for deployment, but the decision to deploy is made manually by a human. On the other hand, continuous deployment automates the deployment process, and once the software passes all the necessary tests and quality checks, it is automatically deployed without human intervention. Continuous delivery utilizes automatic deployment to a development environment, while continuous deployment involves automatic deployment to a production environment. In continuous delivery, the software is automatically deployed to a development or staging environment for further testing and validation. However, the actual deployment to the production environment is performed manually. In continuous deployment, the software is automatically deployed to the production environment, eliminating the need for manual intervention in the deployment process. These differences highlight the level of automation and human involvement in the deployment process between continuous delivery and continuous deployment approaches in the DevOps lifecycle.

Question 9

Question Type: MultipleChoice

Your organization is developing serverless applications with Oracle Functions. Many functions will need to store state data in a database, which will require using appropriate credentials. However, your corporate security standards mandate encryption of secret information, such as database passwords. How would you address this security requirement?

Options:

A- Use OCI Console to enter the password in the function configuration section in the provided input field.

B- Leverage application-level configuration variables to store passwords because they are automatically encrypted by Oracle Functions.

C- Use the OCI Vault service to auto-encrypt the password and then set an application-level configuration variable to reference the auto-decrypted password inside your function container.

D- Encrypt the password using the OCI Vault service and then decrypt this password in your function code with the generated key.

Answer:

D

Explanation:

The best way to store and use secret information, such as database passwords, in Oracle Functions is to use the OCI Vault service. The OCI Vault service provides encryption and decryption capabilities for sensitive data. You can use the OCI Vault service to encrypt the password and store it as an application-level configuration variable. Then, you can use the generated key to decrypt the password in your function code when you need to access the database. Verified Reference:Oracle Functions: Using Key Management To Encrypt And Decrypt Configuration Variables

Question 10

Question Type: MultipleChoice

As a Cloud Native developer, you have written a web service for your company. However, your security team has suggested that your web service should address Distributed Denial-of-Service (DDoS) attack. You are time-constrained and you need to ensure that this is implemented as soon as possible. What should you do in this scenario? (Choose the best answer.)

Options:

A- Use a third party service integration to Implement DDoS attack mitigation.

B- Re-write your web service and implement rate limiting.

C- Use the OCI Virtual Cloud Network (VCN) segregation to control DDoS.

D- Use the OCI API Gateway service and configure rate limiting.

Answer:

D

Explanation:

The correct answer in this scenario is to use the OCI API Gateway service and configure rate limiting. Using the OCI API Gateway service and configuring rate limiting is an effective approach to address Distributed Denial-of-Service (DDoS) attacks. By implementing rate limiting, you can control the number of requests that can be made to your web service within a specific time frame. This helps to prevent overload and ensures that your service can handle legitimate traffic while mitigating the impact of DDoS attacks. By leveraging the OCI API Gateway service, you can easily configure rate limiting rules to restrict the number of requests per second or per minute. This allows you to set appropriate thresholds and safeguard your web service from being overwhelmed by excessive requests. The API Gateway acts as a protective layer, filtering out malicious traffic and ensuring the smooth operation of your service. While options like OCI Virtual Cloud Network (VCN) segregation and third-party service integrations may contribute to overall security, they do not specifically address DDoS attacks as efficiently as rate limiting. VCN segregation focuses more on network segmentation and isolation, while third-party service integration may introduce additional dependencies and complexities. Re-writing your web service and implementing rate limiting is a viable option, but it may not be feasible considering the time constraints mentioned. Leveraging the OCI API Gateway service provides a quicker and easier solution to implement DDoS attack mitigation through rate limiting.

Question 11

Question Type: MultipleChoice

Which is NOT a valid use case for leveraging the Oracle Cloud Infrastructure (OCI) Events service?

Options:

A- Capturing the OCI Monitoring service alarms and invoking autoscaling of compute instances.

B- Publishing a notification when long-lived tasks complete, such as an OCI Autonomous Database backup completion.

C- Triggering a notification action when a function completes its execution.

D- Triggering a function deployed in Oracle Functions when new files are uploaded to an OCI Object Storage bucket.

E- Publishing all the OCI resource events in a specific compartment to the OCI Streaming service for later analysis.

Answer:

A

Explanation:

The use case that is NOT a valid use case for leveraging the Oracle Cloud Infrastructure (OCI) Events service is 'Capturing the OCI Monitoring service alarms and invoking autoscaling of compute instances.' The OCI Events service is designed to provide event-driven architecture and enable automated responses to events occurring within the Oracle Cloud Infrastructure. It allows you to react to changes and activities happening within your OCI resources. The Events service can be used to trigger actions based on events like file uploads, resource changes, or task completions. However, capturing the OCI Monitoring service alarms and invoking autoscaling of compute instances is not a direct functionality provided by the OCI Events service. Autoscaling based on monitoring metrics is typically handled by the OCI Autoscaling service, which is specifically designed for that purpose. The OCI Monitoring service provides monitoring and alerting capabilities, while the Autoscaling service handles the dynamic scaling of compute instances based on predefined policies and thresholds.