Free Questions for 1Z0-1084-24

Shared by Powers on 26-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Your company has recently deployed a new web application that uses Oracle Functions. Your manager instructs you to implement monitoring metrics to manage your systems more effectively. You know that Oracle Functions automatically monitors functions on your behalf and reports metrics via Oracle Cloud Infrastructure (OCI) Monitoring. Which TWO metrics are collected and made available by this feature? (Choose two.)

Options:

A- Amount of CPU used by a function
B- Length of time a function runs
C- Number of times a function Is removed
D- Amount of RAM used by a function
E- Number of times a function is invoked

Answer:

A, D

Explanation:

The correct answers are: Amount of RAM used by a function: Oracle Functions collects and reports the amount of memory (RAM) used by a function during its execution. This metric helps in monitoring and optimizing the resource consumption of functions. Length of time a function runs: Oracle Functions captures and provides the duration of function executions. This metric allows you to track the performance and responsiveness of your functions and identify any potential bottlenecks or delays. These metrics provide valuable insights into the resource utilization and performance of your functions, enabling you to monitor and optimize their behavior in the Oracle Cloud Infrastructure (OCI) environment.

Question 2

Question Type: MultipleChoice

You developed a microservices-based application that runs in an Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) cluster. It has multiple endpoints that need to be exposed to the public internet. What is the most cost-effective way to expose multiple application

endpoints without adding unnecessary complexity to the application?

Options:

A- Use a NodePort service type in Kubernetes for each of your service endpoints using the node's public IP address to access the applications.
B- Create a separate load balancer instance for each service using the lowest 100 Mbps option.
C- Use a ClusterIP service type in Kubernetes for each of your service endpoints using a load balancer to expose the endpoints.
D- Deploy an Ingress Controller and use it to expose each endpoint with its own routing endpoint.

Answer:

D

Explanation:

An Ingress Controller is a Kubernetes resource that provides advanced routing and load balancing for your applications running on a Kubernetes cluster1.An Ingress Controller allows you to define rules that specify how to route traffic to different services in your cluster based on the host name or path of the incoming request1.By deploying an Ingress Controller and using it to expose multiple application endpoints, you can achieve the following benefits1:

Cost-effectiveness: You only need to create one load balancer instance per cluster, instead of one per service, which reduces the cost of exposing your applications.

Simplicity: You only need to manage one set of routing rules for all your services, instead of configuring each service separately, which simplifies the application deployment and maintenance.

Flexibility: You can use different types of Ingress Controllers, such as NGINX or Traefik, that offer various features and customization options for your routing needs.

Question 3

Question Type: MultipleChoice

Which of these is a valid use case for OCI Queue?

Options:

A- Managing network traffic between services
B- Storing and retrieving large files
C- Sending real-time streaming data
D- Building decoupled and scalable systems

Answer:

D

Explanation:

OCI Queue is a fully managed serverless service that helps decouple systems and enable asynchronous operations2.Queue handles high-volume transactional data that requires independently processed messages without loss or duplication2.A valid use case for OCI Queue is building decoupled and scalable systems, such as event-driven architectures or microservices-based applications2. For example, you can use Queue to decouple your application and build an event-driven architecture.Decoupling ensures that individual application components can scale independently and that you can future-proof your design so that as new application components are built, they can publish or subscribe to the queue2.

Question 4

Question Type: MultipleChoice

Which two "Action Type" options are NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition? (Choose two.)

Options:

A- Email
B- Streaming
C- Slack
D- Functions
E- Notifications

Answer:

A, C

Explanation:

The two 'Action Type' options that are NOT available in an Oracle Cloud Infrastructure (OCI) Events rule definition are: Email (Correct) Slack (Correct) The available 'Action Type' options in OCI Events rule definition include Functions, Notifications, and Streaming. However, email and Slack are not directly supported as action types in OCI Events. Instead, you can use Notifications to send notifications to various notification channels, including email and Slack, through the OCI Notifications service.

Question 5

Question Type: MultipleChoice

What can you use to dynamically make Kubernetes resources discoverable to public DNS servers? (Choose the best answer.)

Options:

A- kubeDNS
B- DynDNS
C- CoreDNS
D- ExternalDNS

Answer:

D

Explanation:

To dynamically make Kubernetes resources discoverable to public DNS servers, you can use ExternalDNS. ExternalDNS is a Kubernetes add-on that automates the management of DNS records for your Kubernetes services and ingresses. It can be configured to monitor the changes in your Kubernetes resources and automatically update DNS records in a supported DNS provider. By integrating ExternalDNS with your Kubernetes cluster, you can ensure that the DNS records for your services and ingresses are automatically created, updated, or deleted based on changes in your Kubernetes resources. This allows your Kubernetes resources to be discoverable by external systems through public DNS servers.

Question 6

Question Type: MultipleChoice

A DevOps engineer is troubleshooting the Meshifyd application, which is running in an Oracle Cloud Infrastructure (OCI) environment. The engineer has set up the OCI Logging service to store access logs for the application but notices that the logs from the Meshifyd application are not showing up in the logging service. The engineer suspects that there might be an issue with the logging configuration. Which two statements are potential reasons for logs from the Meshifyd application not showing up in the OCI Logging service?

Options:

A- The logconfig.json file has incorrect or missing OCID for the custom log in the logobjectId field.
B- The OCI Logging service is set up to pre access logs by creating a log group and custom log within the same compartment.
C- The logconfig.json file has incorrect or missing information in the application namespace in the paths field.
D- The logconfig.json file has incorrect or missing information in the application namespace in the src field.
E- The logconfig.json file has incorrect or missing OCID for the custom log group in the logGroupObjectId field.

Answer:

A, E

Explanation:

The logconfig.json file is a configuration file that specifies how the Unified Monitoring Agent collects and uploads custom logs to the OCI Logging service2.The logconfig.json file contains an array of objects, each representing a custom log configuration2.Each custom log configuration object has the following fields2:

logGroupObjectId: The OCID of the log group where the custom log is stored.

logObjectId: The OCID of the custom log.

paths: An array of paths to files or directories containing the custom logs.

src: A regular expression that matches the files containing the custom logs.

parser: A parser definition that specifies how to parse the custom logs. If the logconfig.json file has incorrect or missing OCID for the custom log in the logobjectId field, or incorrect or missing

OCID for the custom log group in the logGroupObjectId field, then the Unified Monitoring Agent will not be able to upload the custom logs to the OCI Logging service2. Therefore, these are potential reasons for logs from the Meshifyd application not showing up in the OCI Logging service. Verified Reference:Unified Monitoring Agent Configuration File

Question 7

Question Type: MultipleChoice

Kubernetes includes various elements such as compute, network, and storage. Compute is essentially CPU (units) and memory (bytes). Within an OKE cluster, what is considered to be the smallest unit of deployment with respect to compute?

Options:

A- Container
B- Service
C- Pod
D- Namespace
E- Deployment resource

Answer:

C

Explanation:

A pod is the smallest and simplest unit in the Kubernetes object model that you create or deploy2. A pod represents a single instance of a running process in your cluster. Pods contain one or more containers, such as Docker containers. When you create a pod, you define how much CPU and memory (RAM) each container needs.A pod can also include storage volumes, IP addresses, options that govern how the container(s) should run, and more2.Pods are the basic building blocks of larger Kubernetes constructs such as deployments, replica sets, and services2.

Question 8

Question Type: MultipleChoice

Which open source engine is used by Oracle Cloud Infrastructure (OCI) to power Oracle Functions?

Options:

A- Knative
B- Kubeless
C- Apache OpenWhisk
D- Fn Project

Answer:

D

Explanation:

Fn Project is the open source engine that is used by OCI to power Oracle Functions1.Fn Project is an open source, container native, serverless platform that can be run anywhere - any cloud or on-premises1. Fn Project is easy to use, extensible, and performant.You can download and install the open source distribution of Fn Project, develop and test a function locally, and then use the same tooling to deploy that function to Oracle Functions1. Verified Reference:Overview of Functions

Question 9

Question Type: MultipleChoice

You have two microservices, A and B, running in production. Service A relies on APIs from service B. You want to test changes to service A without deploying all of its dependencies, which include service B. Which approach should you take to test service A?

Options:

A- Test using API mocks.
B- Test the APIs in private environments.
C- Test against production APIs.
D- There is no need to explicitly test APIs.

Answer:

A

Explanation:

API mocking is a technique that simulates the behavior of real APIs without requiring the actual implementation or deployment of the dependent services1.API mocking allows you to test changes to service A without deploying all of its dependencies, such as service B, by creating mock responses for the APIs that service A relies on1.API mocking has several benefits, such as1:

Faster testing: You can test your service A without waiting for service B to be ready or available, which reduces the testing time and feedback loop.

Isolated testing: You can test your service A in isolation from service B, which eliminates the possibility of external factors affecting the test results or causing errors.

Controlled testing: You can test your service A with different scenarios and edge cases by creating mock responses that mimic various situations, such as success, failure, timeout, etc.

Question 10

Question Type: MultipleChoice

Which technique is used for testing the entire user flow as well as the moving parts of a cloud native app, ensuring that there are no high-level discrepancies?

Options:

A- Contract Testing
B- Integration Testing
C- Unit Testing
D- Component Testing
E- End-to-end Testing

Answer:

E

Explanation:

End-to-end testing is a technique that involves checking the entire user flow as well as the

moving parts of a cloud native app, ensuring that there are no high-level discrepancies3.End-to-end testing simulates real user scenarios and validates the functionality, performance, reliability, and security of the app from start to finish3.End-to-end testing has several benefits, such as3:

Comprehensive testing: You can test your app as a whole and verify that all the components work together as expected.

User-centric testing: You can test your app from the user's perspective and ensure that it meets the user's needs and expectations.

Quality assurance: You can test your app in a realistic environment and identify any issues or defects before releasing it to the users.

Question 11

Question Type: MultipleChoice

Having created a Container Engine for Kubernetes (OKE) cluster, you can use Oracle Cloud Infrastructure (OCI) Logging to view and search the logs of applications running on the worker node compute instances in the cluster. Which task is NOT required to collect and parse application logs? (Choose the best answer.)

Options:

A- Create a dynamic group with a rule that includes all worker nodes In the cluster.
B- Set the OCI Logging option to Enabled for the cluster.
C- Enable monitoring for all worker nodes in the cluster.
D- Configure a custom log in OCI Logging with the appropriate agent configuration.

Answer:

C

Explanation:

The correct answer is: Enable monitoring for all worker nodes in the cluster. Enabling monitoring for all worker nodes in the cluster is not required to collect and parse application logs using Oracle Cloud Infrastructure (OCI) Logging. Monitoring is a separate feature that allows you to collect metrics and monitor the health and performance of the worker nodes. To collect and parse application logs, you need to perform the following tasks: Set the OCI Logging option to Enabled for the cluster: This enables the OCI Logging service for the cluster. Create a dynamic group with a rule that includes all worker nodes in the cluster: This helps in targeting the logs generated by

the worker nodes. Configure a custom log in OCI Logging with the appropriate agent configuration: This involves specifying the log source, log path, and log format to parse and collect the application logs. By completing these tasks, you can collect and parse the application logs generated by the applications running on the worker node compute instances in the OKE cluster.

P2P exams