1Z0-1084-22$^{\text{Q\&As}}$

Oracle Cloud Infrastructure 2022 Developer Professional

Pass Oracle 1Z0-1084-22 Exam with 100% Guarantee

Following Questions and Answers are all new published by Oracle Official Exam Center

QUESTION 1

A programmer Is developing a Node is application which will run in a Linux server on their on-premises data center. This application will access various Oracle Cloud Infrastructure (OC1) services using OCI SDKs. What is the secure way to access OCI services with OCI Identity and Access Management (JAM)?

A. Create a new OCI IAM user associated with a dynamic group and a policy that grants the desired permissions to OCI services. Add the on-premises Linux server in the dynamic group.

B. Create an OCI IAM policy with the appropriate permissions to access the required OCI services and assign the policy to the on-premises Linux server.

C. Create a new OCI IAM user, add the user to a group associated with a policy that grants the desired permissions to OCI services. In the on-premises Linux server, generate the keypair used for signing API requests and upload the public key to the IAM user.

D. Create a new OCI IAM user, add the user to a group associated with a policy that grants the desired permissions to OCI services. In the on-premises Linux server, add the user name and password to a file used by Node.js authentication.

Correct Answer: C

Before using Oracle Functions, you have to set up an Oracle Cloud Infrastructure API signing key. The instructions in this topic assume:

-

you are using Linux

-

you are following Oracle\\'s recommendation to provide a passphrase to encrypt the private key For more Detials Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionssetupapikey.htm

QUESTION 2

In order to effectively test your cloud-native applications, you might utilize separate environments (development, testing, staging, production, etc.). Which Oracle Cloud Infrastructure (OC1) service can you use to create and manage your infrastructure?

A. OCI Compute

B. OCI Container Engine for Kubernetes

C. OCI Resource Manager

D. OCI API Gateway

Correct Answer: C

Resource Manager is an Oracle Cloud Infrastructure service that allows you to automate the process of provisioning your Oracle Cloud Infrastructure resources. Using Terraform, Resource Manager helps you install, configure, and manage resources through the "infrastructure-as-code" model.

QUESTION 3

Which two statements are true for service choreography?

A. Service choreographer is responsible for invoking other services.

B. Services involved in choreography communicate through messages/messaging systems.

C. Service choreography relies on a central coordinator.

D. Service choreography should not use events for communication.

E. Decision logic in service choreography is distributed.

Correct Answer: BE

Service Choreography Service choreography is a global description of the participating services, which is defined by exchange of messages, rules of interaction and agreements between two or more endpoints. Choreography employs a decentralized approach for service composition. the decision logic is distributed, with no centralized point.

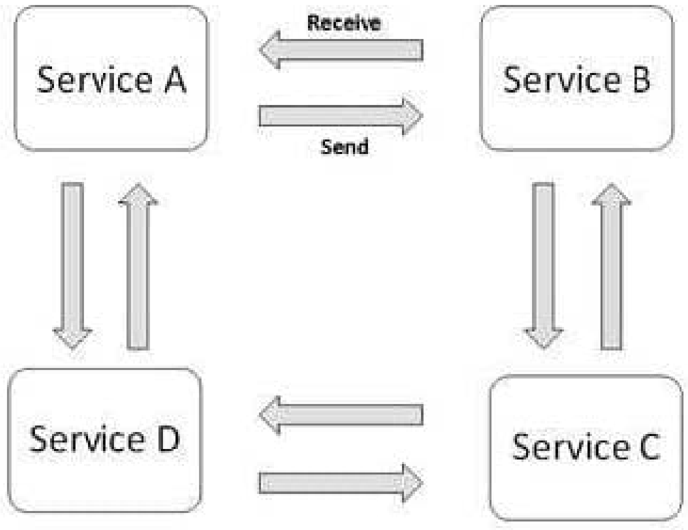

Choreography, in contrast, does not rely on a central coordinator. and all participants in the choreography need to be aware of the business process, operations to execute, messages to exchange, and the timing of message exchanges.

QUESTION 4

What is the difference between blue/green and canary deployment strategies?

A. In blue/green, application Is deployed In minor increments to a select group of people. In canary, both old and new applications are simultaneously in production.

B. In blue/green, both old and new applications are in production at the same time. In canary, application is deployed Incrementally to a select group of people.

C. In blue/green, current applications are slowly replaced with new ones. In

D. In blue/green, current applications are slowly replaced with new ones. In canary, both old and new applications are In production at the same time.

Correct Answer: B

Blue-green deployment is a technique that reduces downtime and risk by running two identical production environments called Blue and Green. At any time, only one of the environments is live, with the live environment serving all production traffic. For this example, Blue is currently live and Green is idle. https://docs.cloudfoundry.org/devguide/deploy-apps/blue-green.html Canary deployments are a pattern for rolling out releases to a subset of users or servers. The idea is to first deploy the change to a small subset of servers, test it, and then roll the change out to the rest of the servers. ... Canaries were once regularly used in coal mining as an early warning system. https://octopus.com/docs/deployment-patterns/canary-deployments

QUESTION 5

Which two handle Oracle Functions authentication automatically?

A. Oracle Cloud Infrastructure SDK

B. cURL

C. Oracle Cloud Infrastructure CLl

D. Signed HTTP Request

E. Fn Project CLI

Correct Answer: CE

Fn Project CLI you can create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure and specify --provider oracle This option enables Oracle Functions to perform authentication and authorization using Oracle Cloud Infrastructure request signing, private keys, user groups, and policies that grant permissions to those user groups.

QUESTION 6

You are developing a distributed application and you need a call to a path to always return a specific JSON content deploy an Oracle Cloud Infrastructure API Gateway with the below API deployment specification.

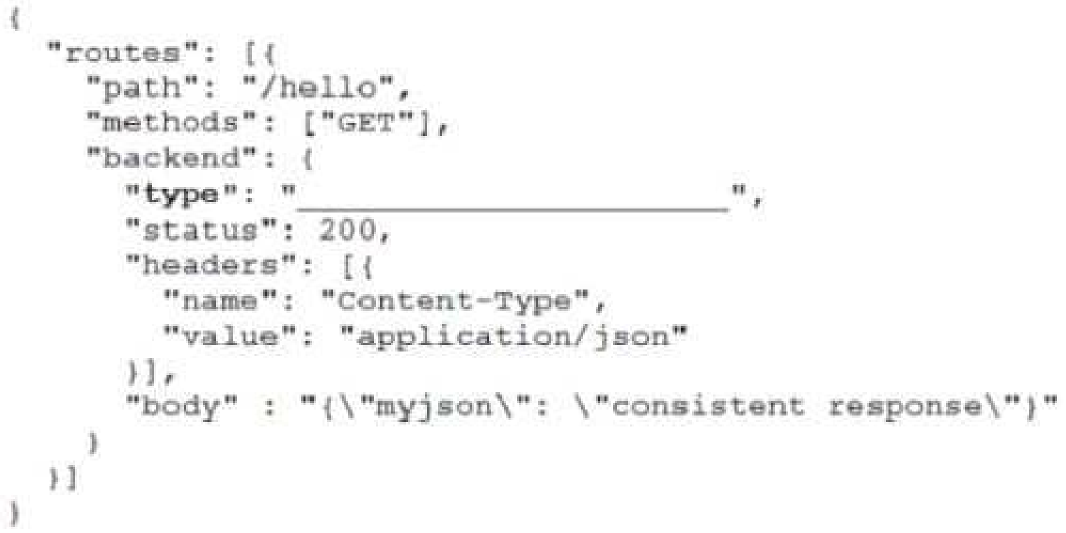

```
{
  "routes": [{
    "path": "/hello",
    "methods": ["GET"],
    "backend": {
      "type": "______________________",
      "status": 200,
      "headers": [{
        "name": "Content-Type",
        "value": "application/json"
      }],
      "body" : "{\"myjson\": \"consistent response\"}"
    }
  }]
}
```

What is the correct value for type?

A. STOCK_RESPONSE_BACKEND

B. CONSTANT_BACKEND

C. JSON_BACKEND

D. HTTP_BACKEND

Correct Answer: A

"type": "STOCK_RESPONSE_BACKEND" indicates that the API gateway itself will act as the back end and return the stock response you define (the status code, the header fields and the body content). https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingstockresponses.htm

QUESTION 7

Which one of the statements describes a service aggregator pattern?

A. It is implemented in each service separately and uses a streaming service

B. It involves implementing a separate service that makes multiple calls to other backend services

C. It uses a queue on both sides of the service communication

D. It involves sending events through a message broker

Correct Answer: B

this pattern isolates an operation that makes calls to multiple back-end microservices, centralizing its logic into a specialized microservice.

QUESTION 8

Which two statements are true for serverless computing and serverless architectures?

A. Long running tasks are perfectly suited for serverless

B. Serverless function state should never be stored externally

C. Application DevOps team is responsible for scaling

D. Serverless function execution is fully managed by a third party

E. Applications running on a FaaS (Functions as a Service) platform

Correct Answer: BE

Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a- Service platform. It is built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs. The serverless and elastic architecture of Oracle Functions means there\\'s no infrastructure administration or software administration for you to perform. You don\\'t provision or maintain compute instances, and operating system software patches and upgrades are applied automatically. Oracle Functions simply ensures your app is highly-available, scalable, secure, and monitored Applications built with a serverless infrastructure will scale automatically as the user base grows or usage increases. If a function needs to be run in multiple instances, the vendor\\'s servers will start up, run, and end them as they are needed. Oracle Functions is based on Fn Project. Fn Project is an open source, container native, serverless platform that can be run anywhere - any cloud or on-premises. Serverless architectures are not built for long-running processes. This limits the kinds of applications that can cost-effectively run in a serverless architecture. Because serverless providers charge for the amount of time code is running, it may cost more to run an application with long-running processes in a serverless infrastructure compared to a traditional one.

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Concepts/functionsconcepts.htm
https://www.cloudflare.com/learning/serverless/why-use-serverless/

QUESTION 9

Who is responsible for patching, upgrading and maintaining the worker nodes in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE)?

A. It Is automated

B. Independent Software Vendors

C. Oracle Support

D. The user

Correct Answer: D

After a new version of Kubernetes has been released and when Container Engine for Kubernetes supports the new version, you can use Container Engine for Kubernetes to upgrade master nodes running older versions of Kubernetes. Because Container Engine for Kubernetes distributes the Kubernetes Control Plane on multiple Oracle-managed master nodes (distributed across different availability domains in a region where supported) to ensure high availability, you\\'re able to upgrade the Kubernetes version running on master nodes with zero downtime. Having upgraded master nodes to a new version of Kubernetes, you can subsequently create new node pools running the newer version. Alternatively, you can continue to create new node pools that will run older versions of Kubernetes (providing those older versions are compatible with the Kubernetes version running on the master nodes). Note that you upgrade master nodes by performing an `in-place\\' upgrade, but you upgrade worker nodes by performing an `out-of-place\\' upgrade. To upgrade the version of Kubernetes running on worker nodes in a node pool, you replace the original node pool with a new node pool that has new worker nodes running the appropriate Kubernetes version. Having \\'drained\\' existing worker nodes in the original node pool to prevent new pods starting and to delete existing pods, you can then delete the original node pool.

QUESTION 10

You are a consumer of Oracle Cloud Infrastructure (OCI) Streaming service. Which API should you use to read and process the stream?

A. ListMessages

B. GetMessages

C. GetObject

D. ReadMessages

Correct Answer: B

https://docs.cloud.oracle.com/en-us/iaas/Content/Streaming/Concepts/streamingoverview.htm Building consumers to read and process messages from a stream using the GetMessages API.

QUESTION 11

You are developing a serverless application with Oracle Functions. Your function needs to store state in a database. Your corporate security Standards mandate encryption of secret information like database passwords. As a function developer, which approach should you follow to satisfy this security requirement?

A. Use the Oracle Cloud Infrastructure Console and enter the password in the function configuration section in the provided input field.

B. Use Oracle Cloud Infrastructure Key Management to auto-encrypt the password. It will inject the auto-decrypted password inside your function container.

C. Encrypt the password using Oracle Cloud Infrastructure Key Management. Decrypt this password in your function code with the generated key.

D. All function configuration variables are automatically encrypted by Oracle Functions.

Correct Answer: A

Passing Custom Configuration Parameters to Functions

he code in functions you deploy to Oracle Functions will typically require values for different parameters. Some pre-defined parameters are available to your functions as environment variables. But you\\'ll often want your functions to use

parameters that you\\'ve defined yourself. For example, you might create a function that reads from and writes to a database. The function will require a database connect string, comprising a username, password, and hostname. You\\'ll

probably want to define username, password, and hostname as parameters that are passed to the function when it\\'s invoked.

Using the Console

To specify custom configuration parameters to pass to functions using the Console:

Log in to the Console as a functions developer.

In the Console, open the navigation menu. Under Solutions and Platform, go to Developer Services and click Functions.

Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that\\'s specified in the Fn Project CLI context (see 6. Create an Fn Project CLI Context to Connect to Oracle

Cloud Infrastructure). Select the compartment specified in the Fn Project CLI context (see 6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure). The Applications page shows the applications defined in the

compartment. Click the name of the application containing functions to which you want to pass custom configuration parameters:

To pass one or more custom configuration parameters to every function in the application, click Configuration to see the Configuration section for the application. To pass one or more custom configuration parameters to a particular function,

click the function\\'s name to see the Configuration section for the function. In the Configuration section, specify details for the first custom configuration parameter:

Key: The name of the custom configuration parameter. The name must only contain alphanumeric characters and underscores, and must not start with a number. For example, username Value: A value for the custom configuration parameter.

The value must only contain printable unicode characters. For example, jdoe

Click the plus button to save the new custom configuration parameter. Oracle Functions combines the key-value pairs for all the custom configuration parameters (both application-wide and function-specific) in the application into a single,

serially-encoded configuration object with a maximum allowable size of 4Kb. You cannot save the new custom configuration parameter if the size of the serially-encoded configuration object would be greater than 4Kb. (Optional) Enter

additional custom configuration parameters as required.

QUESTION 12

You have deployed a Python application on Oracle Cloud Infrastructure Container Engine for Kubernetes. However, during testing you found a bug that you rectified and created a new Docker image. You need to make sure that if this new

Image doesn\\'t work then you can roll back to the previous version.

Using kubectl, which deployment strategies should you choose?

A. Rolling Update

B. Canary Deployment

C. Blue/Green Deployment

D. A/B Testing

Correct Answer: C

Canary deployments are a pattern for rolling out releases to a subset of users or servers. The idea is to first deploy the change to a small subset of servers, test it, and then roll the change out to the rest of the servers. The canary deployment serves as an early warning indicator with less impact on downtime: if the canary deployment fails, the rest of the servers aren\\'t impacted. Blue-green deployment is a technique that reduces downtime and risk by running two identical production environments called Blue and Green. At any time, only one of the environments is live, with the live environment serving all production traffic. For this example, Blue is currently live and Green is idle. A/B testing is a way to compare two versions of a single variable, typically by testing a subject\\'s response to variant A against variant B, and determining which of the two variants is more effective A rolling update offers a way to deploy the new version of your application gradually across your cluster.

QUESTION 13

Which pattern can help you minimize the probability of cascading failures in your system during partial loss of connectivity or a complete service failure?

A. Retry pattern

B. Anti-corruption layer pattern

C. Circuit breaker pattern

D. Compensating transaction pattern

Correct Answer: C

A cascading failure is a failure that grows over time as a result of positive feedback. It can occur when a portion of an overall system fails, increasing the probability that other portions of the system fail. the circuit breaker pattern prevents the service from performing an operation that is likely to fail. For example, a client service can use a circuit breaker to prevent further remote calls over the network when a downstream service is not functioning properly. This can also prevent the network from becoming congested by a sudden spike in failed retries by one service to another, and it can also prevent cascading failures. Self-healing circuit breakers check the downstream service at regular intervals and reset the circuit breaker when the downstream service starts functioning properly.
https://blogs.oracle.com/developers/getting-started-with-microservices-part-three

QUESTION 14

Your Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) administrator has created an OKE cluster with one node pool in a public subnet. You have been asked to provide a log file from one of the nodes for troubleshooting

purpose.

Which step should you take to obtain the log file?

A. ssh into the node using public key.

B. ssh into the nodes using private key.

C. It is impossible since OKE is a managed Kubernetes service.

D. Use the username open and password to login.

Correct Answer: B

Kubernetes cluster is a group of nodes. The nodes are the machines running applications. Each node can be a physical machine or a virtual machine. The node\\'s capacity (its number of CPUs and amount of memory) is defined when the

node is created. A cluster comprises:

- one or more master nodes (for high availability, typically there will be a number of master nodes)

- one or more worker nodes (sometimes known as minions) Connecting to Worker Nodes Using SSH

If you provided a public SSH key when creating the node pool in a cluster, the public key is installed on all worker nodes in the cluster. On UNIX and UNIX-like platforms (including Solaris and Linux), you can then connect through SSH to the

worker nodes using the ssh utility (an SSH client) to perform administrative tasks.

Note the following instructions assume the UNIX machine you use to connect to the worker node:

Has the ssh utility installed.

Has access to the SSH private key file paired with the SSH public key that was specified when the cluster was created.

How to connect to worker nodes using SSH depends on whether you specified public or private subnets for the worker nodes when defining the node pools in the cluster. Connecting to Worker Nodes in Public Subnets Using SSH Before you

can connect to a worker node in a public subnet using SSH, you must define an ingress rule in the subnet\\'s security list to allow SSH access. The ingress rule must allow access to port 22 on worker nodes from source 0.0.0.0/0 and any

source port To connect to a worker node in a public subnet through SSH from a UNIX machine using the ssh utility:

1- Find out the IP address of the worker node to which you want to connect. You can do this in a number of ways: Using kubectl. If you haven\\'t already done so, follow the steps to set up the cluster\\'s kubeconfig configuration file and (if necessary) set the KUBECONFIG environment variable to point to the file. Note that you must set up your own kubeconfig file. You cannot access a cluster using a kubeconfig file that a different user set up. See Setting Up Cluster Access. Then in a terminal window, enter kubectl get nodes to see the public IP addresses of worker nodes in node pools in the cluster. Using the Console. In the Console, display the Cluster List page and then select the cluster to which the worker node belongs. On the Node Pools tab, click the name of the node pool to which the worker node belongs. On the Nodes tab, you see the public IP address of every worker node in the node pool. Using the REST API. Use the ListNodePools operation to see the public IP addresses of worker nodes in a node pool. 2- In the terminal window, enter ssh opc@ to connect to the worker node, where is the IP address of the worker node that you made a note of earlier. For example, you might enter ssh opc@192.0.2.254. Note that if the SSH private key is not stored in the file or in the path that the ssh utility expects (for example, the ssh utility might expect the private key to be stored in ~/.ssh/id_rsa), you must explicitly specify the private key filename and location in one of two ways: Use the -i option to specify the filename and location of the private key. For example, ssh -i ~/.ssh/my_keys/my_host_key_filename opc@192.0.2.254 Add the private key filename and location to an SSH configuration file, either the client configuration file (~/.ssh/config) if it exists, or the system-wide client configuration file (/etc/ssh/ssh_config). For example, you might add the following: Host 192.0.2.254 IdentityFile ~/.ssh/my_keys/my_host_key_filename For more about the ssh utility\\'s configuration file, enter man ssh_config Note also that permissions on the private key file must allow you read/write/execute access, but prevent other users from accessing the file. For example, to set appropriate permissions, you might enter chmod 600 ~/.ssh/my_keys/my_host_key_filename. If permissions are not set correctly and the private key file is accessible to other users, the ssh utility will simply ignore the private key file.

QUESTION 15

You have a containerized app that requires an Autonomous Transaction Processing (ATP) Database. Which option is not valid for o from a container in Kubernetes?

A. Enable Oracle REST Data Services for the required schemas and connect via HTTPS.

B. Create a Kubernetes secret with contents from the instance Wallet files. Use this secret to create a volume mounted to the appropriate path in the application deployment manifest.

C. Use Kubernetes secrets to configure environment variables on the container with ATP instance OCID, and OCI API credentials. Then use the CreateConnection API endpoint from the service runtime.

D. Install the Oracle Cloud Infrastructure Service Broker on the Kubernetes cluster and deploy serviceinstance and serviceBinding resources for ATP. Then use the specified binding name as a volume in the application deployment manifest.

Correct Answer: A

https://blogs.oracle.com/developers/creating-an-atp-instance-with-the-oci-service-broker https://blogs.oracle.com/cloud-infrastructure/integrating-oci-service-broker-with-autonomous- transaction-processing-in-the-real-world