

1Z0-1084-22^{Q&As}

Oracle Cloud Infrastructure 2022 Developer Professional

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

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/1z0-1084-22.html>

100% Passing Guarantee
100% Money Back Assurance

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

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

In a Linux environment, what is the default locations of the configuration file that Oracle Cloud Infrastructure CLI uses for profile information?

- A. /etc/.oci/config
- B. /usr/local/bin/config
- C. \$HOME/.oci/config
- D. /usr/bin/oci/config

Correct Answer: C

By default, the Oracle Cloud Infrastructure CLI configuration file is located at `~/.oci/config`. You might already have a configuration file as a result of installing the Oracle Cloud Infrastructure CLI.

QUESTION 2

You are developing a serverless application with Oracle Functions and Oracle Cloud Infrastructure Object Storage- Your function needs to read a JSON file object from an Object Storage bucket named "input-bucket" in compartment "qa-compartment". Your corporate security standards mandate the use of Resource Principals for this use case.

Which two statements are needed to implement this use case?

- A. Set up a policy with the following statement to grant read access to the bucket:allow dynamic-group read-file-dg to read objects in compartment qa-compartment where target .bucket .name=\\ input-bucket *
- B. Set up the following dynamic group for your function\\'s OCID: Name: read-file-dg Rule: resource.id = `ocid1.fnfunc.oc1.phx.aaaaaaakeaobctakezjz5i4ujj7g25q7sx5mvr55pms6f4da\\`
- C. Set up a policy to grant all functions read access to the bucket:allow all functions in compartment qa-compartment to read objects in target.bucket.name=\\input-bucket\\
- D. Set up a policy to grant your user account read access to the bucket:allow user XYZ to read objects in compartment qa-compartment where target .bucket, name-\\input-bucket\\
- E. No policies are needed. By default, every function has read access to Object Storage buckets in the tenancy

Correct Answer: AB

When a function you\\'ve deployed to Oracle Functions is running, it can access other Oracle Cloud Infrastructure resources. For example:

-You might want a function to get a list of VCNs from the Networking service.

- You might want a function to read data from an Object Storage bucket, perform some operation on the data, and then write the modified data back to the Object Storage bucket. To enable a function to access another Oracle Cloud Infrastructure resource, you have to include the function in a dynamic group, and then create a policy to grant the dynamic group access to that resource. <https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionsaccessingociresources.htm>

QUESTION 3

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

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

Correct Answer: DE

<https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Concepts/eventsoverview.htm> ACTIONS Event Rules must also specify an action to trigger when the filter finds a matching event. Actions are responses you define for event matches. You set up select Oracle Cloud Infrastructure services that the Events service has established as actions. The resources for these services act as destinations for matching events. When the filter in the rule finds a match, the Events service delivers the matching event to one or more of the destinations you identified in the rule. The destination service that receives the event then processes the event in whatever manner you defined. This delivery provides the automation in your environment. You can only deliver events to certain Oracle Cloud Infrastructure services with a rule. Use the following services to create actions: Notifications Streaming Functions

QUESTION 4

Which is NOT a supported SDK Oracle Cloud Infrastructure (OCI)?

- A. Go SDK
- B. Java SDK
- C. NET SDK
- D. Ruby SDK
- E. Python SDK

Correct Answer: C

<https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/sdks.htm>

Software Development Kits (SDKs) Build and deploy apps that integrate with Oracle Cloud Infrastructure services. Each SDK provides the tools you need to develop an app, including code samples and documentation to create, test, and troubleshoot. In addition, if you want to contribute to the development of the SDKs, they are all open source and available on GitHub. SDK for Java Python SDK Ruby SDK Go SDK

QUESTION 5

Which Oracle Cloud Infrastructure (OCI) load balancer shape is used by default in OCI container Engineer for

Kubernetes?

- A. 400 Mbps
- B. 8000 Mbps
- C. There is no default. The shape has to be specified.
- D. 100 Mbps

Correct Answer: D

Specifying Alternative Load Balancer Shapes The shape of an Oracle Cloud Infrastructure load balancer specifies its maximum total bandwidth (that is, ingress plus egress). By default, load balancers are created with a shape of 100Mbps. Other shapes are available, including 400Mbps and 8000Mbps. <https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingloadbalancer.htm>

QUESTION 6

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 7

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 includes service B. Which approach should you take to test service A?

- A. Test against production APIs.

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

Correct Answer: B

Testing using API mocks Developers are frequently tasked with writing code that integrates with other system components via APIs. Unfortunately, it might not always be desirable or even possible to actually access those systems during development. There could be security, performance or maintenance issues that make them unavailable or they might simply not have been developed yet. This is where mocking comes in: instead of developing code with actual external dependencies in place, a mock of those dependencies is created and used instead. Depending on your development needs this mock is made "intelligent" enough to allow you to make the calls you need and get similar results back as you would from the actual component, thus enabling development to move forward without being hindered by eventual unavailability of external systems you depend on

QUESTION 8

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 9

A developer using Oracle Cloud Infrastructure (OCI) API Gateway must authenticate the API requests to their web application. The authentication process must be implemented using a custom scheme which accepts string parameters from the API caller. Which method can the developer use in this scenario?

- A. Create an authorizer function using request header authorization.
- B. Create an authorizer function using token-based authorization.
- C. Create a cross account functions authorizer.

D. Create an authorizer function using OCI Identity and Access Management based authentication

Correct Answer: B

Having deployed the authorizer function, you enable authentication and authorization for an API deployment by including two different kinds of request policy in the API deployment specification:

An authentication request policy for the entire API deployment that specifies: The OCID of the authorizer function that you deployed to Oracle Functions that will perform authentication and authorization. The request attributes to pass to the

authorizer function. Whether unauthenticated callers can access routes in the API deployment.

An authorization request policy for each route that specifies the operations a caller is allowed to perform, based on the caller's access scopes as returned by the authorizer function. Using the Console to Add Authentication and Authorization

Request Policies To add authentication and authorization request policies to an API deployment specification using the Console:

Create or update an API deployment using the Console, select the From Scratch option, and enter details on the Basic Information page. For more information, see [Deploying an API on an API Gateway by Creating an API Deployment](#) and

[Updating API Gateways and API Deployments](#). In the API Request Policies section of the Basic Information page, click the Add button beside Authentication and specify:

Application in : The name of the application in Oracle Functions that contains the authorizer function. You can select an application from a different compartment. Function Name: The name of the authorizer function in

Oracle Functions. Authentication Token: Whether the access token is contained in a request header or a query parameter.

Authentication Token Value: Depending on whether the access token is contained in a request header or a query parameter, specify:

Header Name: If the access token is contained in a request header, enter the name of the header. Parameter Name: If the access token is contained in a query parameter, enter the name of the query parameter.

<https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingauthzauthn.htm>

QUESTION 10

You have been asked to create a stateful application deployed in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE) that requires all of your worker nodes to mount and write data to persistent volumes. Which two OCI storage services should you use?

- A. Use OCI File Services as persistent volume.
- B. Use GlusterFS as persistent volume.
- C. Use OCI Block Volume backed persistent volume.
- D. Use open source storage solutions on top of OCI.
- E. Use OCI Object Storage as persistent volume.

Correct Answer: AC

A PersistentVolume (PV) is a piece of storage in the cluster that has been provisioned by an administrator. PVs are volume plugins like Volumes, but have a lifecycle independent of any individual Pod that uses the PV. A PersistentVolumeClaim (PVC) is a request for storage by a user. It is similar to a Pod. Pods consume node resources and PVCs consume PV resources. If you intend to create Kubernetes persistent volumes, sufficient block volume quota must be available in each availability domain to meet the persistent volume claim. Persistent volume claims must request a minimum of 50 gigabytes. You can define and apply a persistent volume claim to your cluster, which in turn creates a persistent volume that's bound to the claim. A claim is a block storage volume in the underlying IaaS provider that's durable and offers persistent storage, enabling your data to remain intact, regardless of whether the containers that the storage is connected to are terminated. With Oracle Cloud Infrastructure as the underlying IaaS provider, you can provision persistent volume claims by attaching volumes from the Block Storage service.

QUESTION 11

You have written a Node.js function and deployed it to Oracle Functions. Next, you need to call this function from a microservice written in Java deployed on Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE). Which can help you to achieve this?

- A. Use the OCI CLI with kubectl to invoke the function from the microservice.
- B. Oracle Functions does not allow a microservice deployed on OKE to invoke a function.
- C. OKE does not allow a microservice to invoke a function from Oracle Functions.
- D. Use the OCI Java SDK to invoke the function from the microservice.

Correct Answer: D

You can invoke a function that you've deployed to Oracle Functions in different ways:

Using the Fn Project CLI.

Using the Oracle Cloud Infrastructure CLI.

Using the Oracle Cloud Infrastructure SDKs.

Making a signed HTTP request to the function's invoke endpoint. Every function has an invoke endpoint.

QUESTION 12

Which two are characteristics of microservices?

- A. Microservices are hard to test in isolation.
- B. Microservices can be independently deployed.
- C. All microservices share a data store.
- D. Microservices can be implemented in limited number of programming languages.
- E. Microservices communicate over lightweight APIs.

Correct Answer: BE

<https://www.techjini.com/blog/microservices/>

QUESTION 13

A pod security policy (PSP) is implemented in your Oracle Cloud Infrastructure Container Engine for Kubernetes cluster. Which rule can you use to prevent a container from running as root using PSP?

- A. NoPrivilege
- B. RunOnlyAsUser
- C. MustRunAsNonRoot
- D. forbiddenRoot

Correct Answer: C

Require the container to run without root privileges.

rule: \\MustRunAsNonRoot\\

Reference: <https://kubernetes.io/docs/concepts/policy/pod-security-policy/>

QUESTION 14

Which statement is incorrect with regards to the Oracle Cloud Infrastructure (OCI) Notifications service?

- A. Notification topics may be assigned as the action performed by an OCI Events configuration.
- B. OCI Alarms can be configured to publish to a notification topic when triggered.
- C. An OCI function may subscribe to a notification topic.
- D. A subscription can forward notifications to an HTTPS endpoint.
- E. A subscription can integrate with PagerDuty events.
- F. It may be used to receive an email each time an OCI Autonomous Database backup is completed.

Correct Answer: F

QUESTION 15

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/>

[1Z0-1084-22 VCE Dumps](#)

[1Z0-1084-22 Practice Test](#)

[1Z0-1084-22 Study Guide](#)