

1Z0-1084-21^{Q&As}

Oracle Cloud Infrastructure Developer 2021 Associate

Pass Oracle 1Z0-1084-21 Exam with 100% Guarantee

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

<https://www.pass2lead.com/1z0-1084-21.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

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 2

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 3

What is one of the differences between a microservice and a serverless function?

- A. Microservices are used for long running operations and serverless functions for short running operations.
- B. Microservices always use a data store and serverless functions never use a data store.
- C. Microservices are stateless and serverless functions are stateful.
- D. Microservices are triggered by events and serverless functions are not.

Correct Answer: A

microservice is larger and can do more than a function. A function is a relatively small bit of code that performs only one action in response to an event. In many cases, microservices can be decomposed into a number of smaller stateless functions. The difference between microservices and functions is not simply the size. Functions are stateless, and they require no knowledge about or configuration of the underlying server--hence, the term serverless.

<https://developer.oracle.com/java/fn-project-introduction.html>

QUESTION 4

How do you perform a rolling update in Kubernetes?

- A. `kubect1 rolling-update`
- B. `kubect1 upgrade --image=*image:v2`
- C. `kubect1 update -c`
- D. `kubect1 rolling-update --image=image`

Correct Answer: A

<https://docs.oracle.com/en/cloud/iaas/wercker-cloud/wercm/quickstarts/platforms/kubernetes/>

QUESTION 5

As a cloud-native developer, you have written a web service for your company. You have used Oracle Cloud Infrastructure (OCI) API Gateway service to expose the HTTP backend. However, your security team has suggested that your web service should handle Distributed Denial-of-Service (DDoS) attack. You are time-constrained and you need to make sure that this is implemented as soon as possible. What should you do in this scenario?

- A. Use OCI virtual cloud network (VCN) segregation to control DDoS.

- B. Use a third party service integration to implement a DDoS attack mitigation,
- C. Use OCI API Gateway service and configure rate limiting.
- D. Re-write your web service and implement rate limiting.

Correct Answer: C

Having created an API gateway and deployed one or more APIs on it, you'll typically want to limit the rate at which front-end clients can make requests to back-end services. For example, to:

- maintain high availability and fair use of resources by protecting back ends from being overwhelmed by too many requests
- prevent denial-of-service attacks
- constrain costs of resource consumption
- restrict usage of APIs by your customers' users in order to monetize APIs You apply a rate limit globally to all routes in an API deployment specification. If a request is denied because the rate limit has been exceeded, the response header specifies when the request can be retried. You can add a rate-limiting request policy to an API deployment specification by: using the Console editing a JSON file

[1Z0-1084-21 Practice Test](#)

[1Z0-1084-21 Study Guide](#)

[1Z0-1084-21 Braindumps](#)