



1Z0-1084-20^{Q&As}

Oracle Cloud Infrastructure Developer 2020 Associate

Pass Oracle 1Z0-1084-20 Exam with 100% Guarantee

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

<https://www.passapply.com/1z0-1084-20.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

You are developing a polyglot serverless application using Oracle Functions. Which language cannot be used to write your function code?

- A. PL/SQL
- B. Python
- C. Node.js
- D. Java

Correct Answer: A

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. With Oracle Functions, you can write code in Java, Python, Node, Go, and Ruby (and for advanced use cases, bring your own Dockerfile, and Graal VM). You can then deploy your code, call it directly or trigger it in response to events, and get billed only for the resources consumed during the execution.

QUESTION 2

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 3

In the sample Kubernetes manifest file below, what annotations should you add to create a private load balancer In oracle Cloud infrastructure Container Engine for Kubernetes?



```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    <Fill in>
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: nginx
```

```
apiVersion: v1
kind: Service
metadata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    <Fill in>
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: nginx
```

A. service.beta.kubernetes.io/oci-load-balancer-private: "true"

B. service.beta.kubernetes.io/oci-load-balancer-private: "true" service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocid1.subnet.oc1..aaaaa....vdfw"

C. service.beta.kubernetes.io/oci-load-balancer-internal: "true"

D. service.beta.kubernetes.io/oci-load-balancer-internal: "true" service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocid1.subnet.oc1..aaaaa....vdfw"

Correct Answer: D

https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingloadbalancer.htm?TocPath=Services%7CExample%20Network%20Resource%20Configuration%7CUpgrading%20the%20Version%20of%20Kubernetes%20Running%20on%20a%20Master%20Node%7C_____2

2 Creating Internal Load

Balancers in Public and Private Subnets You can create Oracle Cloud Infrastructure load balancers to control access to



services running on a cluster: When you create a `custom` cluster, you select an existing VCN that contains the network resources to be used by the new cluster. If you want to use load balancers to control traffic into the VCN, you select existing public or private subnets in that VCN to host the load balancers. When you create a `quick cluster`, the VCN that's automatically created contains a public regional subnet to host a load balancer. If you want to host load balancers in private subnets, you can add private subnets to the VCN later.

Alternatively, you can create an internal load balancer service in a cluster to enable other programs running in the same VCN as the cluster to access services in the cluster. You can host internal load balancers in public subnets and private subnets. To create an internal load balancer hosted on a public subnet, add the following annotation in the metadata section of the manifest file: `service.beta.kubernetes.io/oci-load-balancer-internal: "true"` To create an internal load balancer hosted on a private subnet, add both following annotations in the metadata section of the manifest file: `service.beta.kubernetes.io/oci-load-balancer-internal: "true"` `service.beta.kubernetes.io/oci-load-balancersubnet1: "ocid1.subnet.oc1..aaaaa....vdfw"` where `ocid1.subnet.oc1..aaaaa....vdfw` is the OCID of the private subnet.

QUESTION 4

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
-

QUESTION 5

With the volume of communication that can happen between different components in cloud-native applications, it is vital to not only test functionality, but also service resiliency. Which statement is true with regards to service resiliency?

- A. Resiliency is about recovering from failures without downtime or data loss.



- B. A goal of resiliency is not to bring a service to a functioning state after a failure.
- C. Resiliency testing can be only done in a test environment.
- D. Resiliency is about avoiding failures.

Correct Answer: D

Resiliency and Availability Resiliency and availability refers to the ability of a system to continue operating, despite the failure or suboptimal performance of some of its components. In the case of Oracle Functions: The control plane is a set of components that manages function definitions. The data plane is a set of components that executes functions in response to invocation requests. For resiliency and high availability, both the control plane and data plane components are distributed across different availability domains and fault domains in a region. If one of the domains ceases to be available, the components in the remaining domains take over to ensure that function definition management and execution are not disrupted. When functions are invoked, they run in the subnets specified for the application to which the functions belong. For resiliency and high availability, best practice is to specify a regional subnet for an application (or alternatively, multiple AD- specific subnets in different availability domains). If an availability domain specified for an application ceases to be available, Oracle Functions runs functions in an alternative availability domain.

[1Z0-1084-20 PDF Dumps](#)

[1Z0-1084-20 Practice Test](#)

[1Z0-1084-20 Study Guide](#)