



# 1Z0-1084-22<sup>Q&As</sup>

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.geekcert.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**

You are working on a cloud native e-commerce application on Oracle Cloud Infrastructure (OCI). Your application architecture has multiple OCI services, including Oracle Functions. You need to trigger these functions directly from other OCI services, without having to run custom code. Which OCI service cannot trigger your functions directly?

- A. OCI Events Service
- B. OCI Registry
- C. OCI API Gateway
- D. Oracle Integration

Correct Answer: B

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. 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 invoke a function that you've deployed to Oracle Functions from:

-

The Fn Project CLI.

-

The Oracle Cloud Infrastructure SDKs.

-Signed HTTP requests to the function's invoke endpoint. Every function has an invoke endpoint.

-Other Oracle Cloud services (for example, triggered by an event in the Events service) or from external services. so 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. Below are the oracle services that can trigger Oracle functions -Events Service

-Notification Service -API Gateway Service -Oracle Integration service(using OCI Signature Version 1 security policy) so OCI Registry services cannot trigger your functions directly

---

**QUESTION 2**

You created a pod called "nginx" and its state is set to Pending. Which command can you run to see the reason why the "nginx" pod is in the pending state?

- A. kubectl logs pod nginx



- B. kubectl describe pod nginx
- C. kubectl get pod nginx
- D. Through the Oracle Cloud Infrastructure Console

Correct Answer: B

#### Debugging Pods

The first step in debugging a pod is taking a look at it. Check the current state of the pod and recent events with the following command:

```
kubectl describe pods ${POD_NAME}
```

Look at the state of the containers in the pod. Are they all Running? Have there been recent restarts? Continue debugging depending on the state of the pods.

My pod stays pending

If a pod is stuck in Pending it means that it can not be scheduled onto a node. Generally this is because there are insufficient resources of one type or another that prevent scheduling. Look at the output of the kubectl describe ... command

above. There should be messages from the scheduler about why it can not schedule your pod.

<https://kubernetes.io/docs/tasks/debug-application-cluster/debug-pod-replication-controller/>

### QUESTION 3

A service you are deploying to Oracle infrastructure (OCI) Container Engine for Kubernetes (OKE) uses a docker image from a private repository Which configuration is necessary to provide access to this repository from OKE?

- A. Add a generic secret on the cluster containing your identity credentials. Then specify a registrycredentials property in the deployment manifest.
- B. Create a docker-registry secret for OCIR with API key credentials on the cluster, and specify the imagepullsecret property in the application deployment manifest.
- C. Create a docker-registry secret for OCIR with identity Auth Token on the cluster, and specify the image pull secret property in the application deployment manifest.
- D. Create a dynamic group for nodes in the cluster, and a policy that allows the dynamic group to read repositories in the same compartment.

Correct Answer: C

**Pulling Images from Registry during Deployment** During the deployment of an application to a Kubernetes cluster, you'll typically want one or more images to be pulled from a Docker registry. In the application's manifest file you specify the images to pull, the registry to pull them from, and the credentials to use when pulling the images. The manifest file is commonly also referred to as a pod spec, or as a deployment.yaml file (although other filenames are allowed). If you want the application to pull images that reside in Oracle Cloud Infrastructure Registry, you have to perform two steps:

-



You have to use kubectl to create a Docker registry secret. The secret contains the Oracle Cloud Infrastructure credentials to use when pulling the image. When creating secrets, Oracle strongly recommends you use the latest version of kubectl To create a Docker registry secret: 1- 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. 2- In a terminal window, enter: \$ kubectl create secret docker-registry --docker-server=.ocir.io --docker-username=\\'\\' --docker-password=\\'\\' --docker-email=\\'\\' where: is a name of your choice, that you will use in the manifest file to refer to the secret . For example, ocirsecret is the key for the Oracle Cloud Infrastructure Registry region you're using. For example, iad. See Availability by Region. ocir.io is the Oracle Cloud Infrastructure Registry name. is the auto-generated Object Storage namespace string of the tenancy containing the repository from which the application is to pull the image (as shown on the Tenancy Information page). For example, the namespace of the acme-dev tenancy might be ansh81vru1zp. Note that for some older tenancies, the namespace string might be the same as the tenancy name in all lower-case letters (for example, acme-dev). is the username to use when pulling the image. The username must have access to the tenancy specified by . For example, jdoe@acme.com . If your tenancy is federated with Oracle Identity Cloud Service, use the format oracleidentitycloudservice/ is the auth token of the user specified by . For example, klj64r{1sJSSF-;)K8 is an email address. An email address is required, but it doesn't matter what you specify. For example, jdoe@acme.com

-

You have to specify the image to pull from Oracle Cloud Infrastructure Registry, including the repository location and the Docker registry secret to use, in the application's manifest file.

---

#### QUESTION 4

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 5

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.

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

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

[1Z0-1084-22 Braindumps](#)