



# 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.pass4itsure.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**

Which one of the following is NOT a valid backend-type supported by Oracle Cloud Infrastructure (OCI) API Gateway?

- A. STOCK\_RESPONSE\_BACKEND
- B. ORACLE\_FUNCTIONS\_BACKEND
- C. ORACLE\_STREAMS\_BACKEND
- D. HTTP\_BACKEND

Correct Answer: C

In the API Gateway service, a back end is the means by which a gateway routes requests to the back-end services that implement APIs. If you add a private endpoint back end to an API gateway, you give the API gateway access to the VCN associated with that private endpoint. You can also grant an API gateway access to other Oracle Cloud Infrastructure services as back ends. For example, you could grant an API gateway access to Oracle Functions, so you can create and deploy an API that is backed by a serverless function. API Gateway service to create an API gateway, you can create an API deployment to access HTTP and HTTPS URLs. <https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayusinghttpbackend.htm> API Gateway service to create an API gateway, you can create an API deployment that invokes serverless functions defined in Oracle Functions. <https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayusingfunctionsbackend.htm> API Gateway service, you can define a path to a stock response back end <https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingstockresponses.htm>

---

**QUESTION 2**

Which two statements accurately describe an Oracle Functions application?

- A. A small block of code invoked in response to an Oracle Cloud Infrastructure (OCI) Events service
- B. A Docker image containing all the functions that share the same configuration
- C. An application based on Oracle Functions, Oracle Cloud Infrastructure (OCI) Events and OCI API Gateway services
- D. A common context to store configuration variables that are available to all functions in the application
- E. A logical group of functions

Correct Answer: DE

Applications in the Function services In Oracle Functions, an application is:

1.  
a logical grouping of functions
2.  
a common context to store configuration variables that are available to all functions in the application

When you define an application in Oracle Functions, you specify the subnets in which to run the functions in the



application.

---

### QUESTION 3

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 4

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 5

Given a service deployed on Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE), which annotation



should you add in the sample manifest file below to specify a 400 Mbps load balancer?

```
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-value: 400Mbps
- B. service.beta.kubernetes.io/oci-load-balancer-size: 400Mbps
- C. service.beta.kubernetes.io/oci-load-balancer-shape: 400Mbps
- D. service, beta, kubernetes . io/oci-load-balancer-kind: 400Mbps

Correct Answer: C

---

## QUESTION 6

Per CAP theorem, in which scenario do you NOT need to make any trade-off between the guarantees?

- A. when there are no network partitions
- B. when the system is running in the cloud
- C. when the system is running on-premise
- D. when you are using load balancers

Correct Answer: A

## CAP THEOREM

"CONSISTENCY, AVAILABILITY and PARTITION TOLERANCE are the features that we want in our distributed system together"

Of three properties of shared-data systems (Consistency, Availability and tolerance to network Partitions) only two can be achieved at any given moment in time.

---

## QUESTION 7



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>

---

### QUESTION 8

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 9

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 10**

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 11**

You are developing a serverless application with Oracle Functions. You have created a function in compartment named prod. When you try to invoke your function you get the following error. Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions.oc1.phx.aaaaaaac... does not exist or Oracle Functions is not authorized to use it How can you resolve this error?

- A. Create a policy:Allow function-family to use virtual-network-family in compartment prod
- B. Create a policy:Allow any-user to manage function-family and virtual-network-family in compartment prod
- C. Create a policy:Allow service FaaS to use virtual-network-family in compartment prod
- D. Deleting the function and redeploying it will fix the problem

Correct Answer: C

Invoking a function returns a FunctionInvokeSubnetNotAvailable message and a 502 error (due to a DHCP Options issue)

When you invoke a function that you've deployed to Oracle Functions, you might see the following error message:

```
{"code":"FunctionInvokeSubnetNotAvailable","message":"dhcp options ocid1.dhcpoptions..... does not exist or Oracle Functions is not authorized to use it"}
```

Fn: Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions..... does not exist or Oracle Functions is not authorized to use it If you see this error:

Double-check that a policy has been created to give Oracle Functions access to network resources.





### Service Access to Network Resources

When Oracle Functions users create a function or application, they have to specify a VCN and a subnet in which to create them. To enable the Oracle Functions service to create the function or application in the specified VCN and subnet,

you must create an identity policy to grant the Oracle Functions service access to the compartment to which the network resources belong. To create a policy to give the Oracle Functions service access to network resources:

Log in to the Console as a tenancy administrator.

Create a new policy in the root compartment:

Open the navigation menu. Under Governance and Administration, go to Identity and click Policies. Follow the instructions in To create a policy, and give the policy a name (for example, functions- service-network-access).

Specify a policy statement to give the Oracle Functions service access to the network resources in the compartment:

Allow service FaaS to use virtual-network-family in compartment For example:

Allow service FaaS to use virtual-network-family in compartment acme-network Click Create.

Double-check that the set of DHCP Options in the VCN specified for the application still exists.

---

### QUESTION 12

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

The 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 13

Which header is NOT required when signing GET requests to Oracle Cloud Infrastructure APIs?

- A. date or x-date
- B. (request-target)
- C. content-type
- D. host

Correct Answer: C

For GET and DELETE requests (when there's no content in the request body), the signing string must include at least these headers: (request-target) (as described in draft-cavage-http-signatures-08) host date or x-date (if both are



included, Oracle uses x-date)

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

---

#### QUESTION 14

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 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 Practice Test](#)

[1Z0-1084-22 Exam  
Questions](#)

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