



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.pass4itsure.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**

Which two statements accurately describe Oracle SQL Developer Web on Oracle Cloud Infrastructure (OCI) Autonomous Database?

- A. It is available for databases with dedicated Exadata infrastructure only.
- B. After provisioning into an OCI compute Instance, it can automatically connect to the OCI Autonomous Databases instances.
- C. It is available for databases with both dedicated and shared Exadata infrastructure.
- D. It provides a development environment and a data modeler interface for OCI Autonomous Databases.
- E. It must be enabled via OCI Identity and Access Management policy to get access to the Autonomous Databases instances.

Correct Answer: AD

Oracle SQL Developer Web in Autonomous Data Warehouse provides a development environment and a data modeler interface for Autonomous Databases. SQL Developer Web is available for databases with both dedicated Exadata infrastructure and shared Exadata infrastructure.

<https://docs.cloud.oracle.com/enus/iaas/Content/Database/Tasks/adbtools.htm>

QUESTION 2

Which concept is NOT related to Oracle Cloud Infrastructure Resource Manager?

- A. Job
- B. Stack
- C. Queue
- D. Plan

Correct Answer: C

<https://docs.cloud.oracle.com/en-us/iaas/Content/ResourceManager/Concepts/resourcemanager.htm> Following are brief descriptions of key concepts and the main components of Resource Manager. CONFIGURATION Information to codify your infrastructure. A Terraform configuration can be either a solution or a file that you write and upload. JOB Instructions to perform the actions defined in your configuration. Only one job at a time can run on a given stack; further, you can have only one set of Oracle Cloud Infrastructure resources on a given stack. To provision a different set of resources, you must create a separate stack and use a different configuration. Resource Manager provides the following job types: Plan: Parses your Terraform configuration and creates an execution plan for the associated stack. The execution plan lists the sequence of specific actions planned to provision your Oracle Cloud Infrastructure resources. The execution plan is handed off to the apply job, which then executes the instructions. Apply. Applies the execution plan to the associated stack to create (or modify) your Oracle Cloud Infrastructure resources. Depending on the number and type of resources specified, a given apply job can take some time. You can check status while the job runs. Destroy. Releases resources associated with a stack. Released resources are not deleted. For example, terminates a Compute instance controlled by a stack. The stack's job history and state remain after running a destroy job. You can monitor the status and review the results of a destroy job by inspecting the stack's log files. Import State. Sets the provided Terraform state file as the current state of the stack. Use this job to migrate local Terraform



environments to Resource Manager. STACK The collection of Oracle Cloud Infrastructure resources corresponding to a given Terraform configuration. Each stack resides in the compartment you specify, in a single region; however, resources on a given stack can be deployed across multiple regions. An OCID is assigned to each stack.

QUESTION 3

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 4

A leading insurance firm is hosting its customer portal in Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes with an OCI Autonomous Database. Their support team discovered a lot of SQL injection attempts and cross-site scripting attacks to the portal, which is starting to affect the production environment. What should they implement to mitigate this attack?

- A. Network Security Lists
- B. Network Security Groups
- C. Network Security Firewall
- D. Web Application Firewall

Correct Answer: D

Oracle Cloud Infrastructure Web Application Firewall (WAF) is a cloud-based, Payment Card Industry (PCI) compliant, global security service that protects applications from malicious and unwanted internet traffic. WAF can protect any internet facing endpoint, providing consistent rule enforcement across a customer's applications. WAF provides you with the ability to create and manage rules for internet threats including Cross- Site Scripting (XSS), SQL Injection and other OWASP-defined vulnerabilities. Unwanted bots can be mitigated while tactically allowed desirable bots to enter. Access rules can limit based on geography or the signature of the request.

QUESTION 5

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-networkaccess).

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 6

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/>

QUESTION 7

As a cloud-native developer, you are designing an application that depends on Oracle Cloud Infrastructure (OCI) Object Storage wherever the application is running. Therefore, provisioning of storage buckets should be part of your Kubernetes deployment process for the application. Which should you leverage to meet this requirement?

- A. OCI Service Broker for Kubernetes
- B. OCI Container Engine for Kubernetes
- C. Open Service Broker API
- D. Oracle Functions

Correct Answer: A

<https://blogs.oracle.com/cloud-infrastructure/introducing-service-broker-for-kubernetes> OCI Service Broker for Kubernetes is an implementation of the Open Service Broker API. OCI Service Broker for Kubernetes is specifically for



interacting with Oracle Cloud Infrastructure services from Kubernetes clusters. It includes three service broker adapters to bind to the following Oracle Cloud Infrastructure services: Object Storage Autonomous Transaction Processing Autonomous Data Warehouse

QUESTION 8

Which two handle Oracle Functions authentication automatically?

- A. Oracle Cloud Infrastructure SDK
- B. cURL
- C. Oracle Cloud Infrastructure CLI
- D. Signed HTTP Request
- E. Fn Project CLI

Correct Answer: CE

Fn Project CLI you can create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure and specify --provider oracle This option enables Oracle Functions to perform authentication and authorization using Oracle Cloud Infrastructure request signing, private keys, user groups, and policies that grant permissions to those user groups.

QUESTION 9

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 10

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 --dockerusername=/' --docker-password=/' --dockeremail=/'` 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, `acmedev`). `.` 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, `kjj64r{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 11

You are developing a distributed application and you need a call to a path to always return a specific JSON content deploy an Oracle Cloud Infrastructure API Gateway with the below API deployment specification.

```
{
  "routes": [{
    "path": "/hello",
    "methods": ["GET"],
    "backend": {
      "type": "_____",
      "status": 200,
      "headers": [{
        "name": "Content-Type",
        "value": "application/json"
      }],
      "body" : "{\"myjson\": \"consistent response\"}"
    }
  ]
}
```

What is the correct value for type?

- A. STOCK_RESPONSE_BACKEND
- B. CONSTANT_BACKEND
- C. JSON_BACKEND
- D. HTTP_BACKEND

Correct Answer: A

"type": "STOCK_RESPONSE_BACKEND" indicates that the API gateway itself will act as the back end and return the stock response you define (the status code, the header fields and the body content). <https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingstockresponses.htm>

QUESTION 12

What are two of the main reasons you would choose to implement a serverless architecture?

- A. No need for integration testing
- B. Reduced operational cost
- C. Improved In-function state management



- D. Automatic horizontal scaling
- E. Easier to run long-running operations

Correct Answer: BD

QUESTION 13

You are using Oracle Cloud Infrastructure (OCI) Resource Manager to manage your infrastructure lifecycle and wish to receive an email each time a Terraform action begins. How should you use the OCI Events service to do this without writing any code?

- A. Create an OCI Notifications topic and email subscription with the destination email address. Then create an OCI Events rule matching "Resource Manager Stack - Update" condition, and select the notification topic for the corresponding action.
- B. Create an OCI Notification topic and email subscription with the destination email address. Then create an OCI Events rule matching "Resource Manager job - Create" condition, and select the notification topic for the corresponding action.
- C. Create a rule in OCI Events service matching the "Resource Manager Stack - Update" condition. Then select "Action Type: Email" and provide the destination email address.
- D. Create an OCI Email Delivery configuration with the destination email address. Then create an OCI Events rule matching "Resource Manager Job - Create" condition, and select the email configuration for the corresponding action.

Correct Answer: B

1.

Create Notifications Topic and Subscription If a suitable Notifications topic doesn't already exist, then you must log in to the Console as a tenancy administrator and create it. Whether you use an existing topic or create a new one, add an email address as a subscription so that you can monitor that email account for notifications

2.

Using the Console to Create a Rule Use the Console to create a rule with a pattern that matches bucket creation events emitted by Object Storage. Specify the Notifications topic you created as an action to deliver matching events. To test your rule, create a bucket. Object Storage emits an event which triggers the action. Check the email specified in the subscription to receive your notification

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

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

QUESTION 14

You have created a repository in Oracle Cloud Infrastructure Registry in the us-ashburn-1 (iad) region in your tenancy with a namespace called "heyoci.

Which three are valid tags for an image named "myapp"?

- A. iad.ocir.io/heyoci/myproject/myapp:0.0.1



- B. us-ashburn-l.ocirJo/heyoci/myapp:0.0.2-beta
- C. us-ashburn-l.ocir.io/heyoci/myproject/myapp:0.0.2-beta
- D. us-ashburn-l.ocir.io/myproject/heyoci/myapp:latest
- E. iad.ocir.io/myproject/heyoci/myapprlatest
- F. iad.ocir.io/heyoci/myapp:0.0.2-beta
- G. iad.ocir.io/heyoci/myapp:latest

Correct Answer: AFG

Give a tag to the image that you're going to push to Oracle Cloud Infrastructure Registry by entering:

docker tag

where:

uniquely identifies the image, either using the image's id (for example, 8e0506e14874), or the image's name and tag separated by a colon (for example, acme-web-app:latest). is in the format .ocir.io///: where: 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 that owns the repository to which you want to push 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). Note also that your user must have access to the tenancy. (if specified) is the name of a repository to which you want to push the image (for example, project01). Note that specifying a repository is optional (see About Repositories). is the name you want to give the image in Oracle Cloud Infrastructure Registry (for example, acme-web-app). is an image tag you want to give the image in Oracle Cloud Infrastructure Registry (for example, version2.0.test). For example, for convenience you might want to group together multiple versions of the acme-web-app image in the acme-dev tenancy in the Ashburn region into a repository called project01. You do this by including the name of the repository in the image name when you push the image, in the format .ocir.io///:. For example, iad.ocir.io/ ansh81vru1zp/project01/acme-web-app:4.6.3. Subsequently, when you use the docker push command, the presence of the repository in the image's name ensures the image is pushed to the intended repository. If you push an image and include the name of a repository that doesn't already exist, a new private repository is created automatically. For example, if you enter a command like docker push iad.ocir.io/ansh81vru1zp/project02/acme-web-app:7.5.2 and the project02 repository doesn't exist, a private repository called project02 is created automatically. If you push an image and don't include a repository name, the image's name is used as the name of the repository. For example, if you enter a command like docker push iad.ocir.io/ansh81vru1zp/acme-web-app:7.5.2 that doesn't contain a repository name, the image's name (acme-web-app) is used as the name of a private repository. [https:// docs.cloud.oracle.com/en-us/iaas/Content/Registry/Concepts/registrywhatisarepository.htm](https://docs.cloud.oracle.com/en-us/iaas/Content/Registry/Concepts/registrywhatisarepository.htm)

QUESTION 15

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>

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

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

[1Z0-1084-20 Exam Questions](#)