

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.certbus.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 tasked with developing an application that requires the use of Oracle Cloud Infrastructure (OCI) APIs to POST messages to a stream in the OCI Streaming service.

Which statement is incorrect?

- A. The request must include an authorization signing string including (but not limited to) x-contentsha256, content-type, and content-length headers.
- B. The Content-Type header must be Set to application/json
- C. An HTTP 401 will be returned if the client's clock is skewed more than 5 minutes from the server's.
- D. The request does not require an Authorization header.

Correct Answer: A

Emits messages to a stream. There's no limit to the number of messages in a request, but the total size of a message or request must be 1 MiB or less. The service calculates the partition ID from the message key and stores messages that share a key on the same partition. If a message does not contain a key or if the key is null, the service generates a message key for you. The partition ID cannot be passed as a parameter. POST /20180418/streams//messages Host: streaming-api.us-phoenix-1.oraclecloud.com { "messages": { { "key": null, "value": "VGhllHF1aWNrIGJyb3dulGZveCBqdW1wZWQgb3ZlciB0aGUgbGF6eSBkb2cu" }, { "key": null, "value": "UGFjayBteSBib3ggd2l0aCBmaXZlIGRvemVulGxpcXVvciBqdWdzLg==" } } } <https://docs.cloud.oracle.com/en-us/iaas/api/#/en/streaming/20180418/Message/PutMessages>

QUESTION 2

In order to effectively test your cloud-native applications, you might utilize separate environments (development, testing, staging, production, etc.). Which Oracle Cloud Infrastructure (OCI) service can you use to create and manage your infrastructure?

- A. OCI Compute
- B. OCI Container Engine for Kubernetes
- C. OCI Resource Manager
- D. OCI API Gateway

Correct Answer: C

Resource Manager is an Oracle Cloud Infrastructure service that allows you to automate the process of provisioning your Oracle Cloud Infrastructure resources. Using Terraform, Resource Manager helps you install, configure, and manage resources through the "infrastructure-as-code" model.

QUESTION 3

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/myapp:latest
- 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 4

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 5

What is the minimum of storage that a persistent volume claim can obtain in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE)?

- A. 50 GB
- B. 10 GB
- C. 1 GB D. 1 TB

Correct Answer: A

The minimum amount of persistent storage that a PVC can request is 50 gigabytes. If the request is for less than 50 gigabytes, the request is rounded up to 50 gigabytes. <https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingpersistentvolumeclaim.htm>

QUESTION 6

You are implementing logging in your services that will be running in Oracle Cloud Infrastructure Container Engine for Kubernetes. Which statement describes the appropriate logging approach?

- A. Each service logs to its own log file.

- B. All services log to an external logging system.
- C. All services log to standard output only.
- D. All services log to a shared log file.

Correct Answer: C

Application and systems logs can help you understand what is happening inside your cluster. The logs are particularly useful for debugging problems and monitoring cluster activity. Most modern applications have some kind of logging mechanism; as such, most container engines are likewise designed to support some kind of logging. The easiest and most embraced logging method for containerized applications is to write to the standard output and standard error streams.

<https://kubernetes.io/docs/concepts/cluster-administration/logging/> <https://blogs.oracle.com/developers/5-best-practices-for-kubernetes-security>

QUESTION 7

You have two microservices, A and B running in production. Service A relies on APIs from service B. You want to test changes to service A without deploying all of its dependencies, which includes service B.

Which approach should you take to test service A?

- A. Test against production APIs.
- B. Test using API mocks.
- C. There is no need to explicitly test APIs.
- D. Test the APIs in private environments.

Correct Answer: B

Testing using API mocks Developers are frequently tasked with writing code that integrates with other system components via APIs. Unfortunately, it might not always be desirable or even possible to actually access those systems during development. There could be security, performance or maintenance issues that make them unavailable ? or they might simply not have been developed yet. This is where mocking comes in: instead of developing code with actual external dependencies in place, a mock of those dependencies is created and used instead. Depending on your development needs this mock is made "intelligent" enough to allow you to make the calls you need and get similar results back as you would from the actual component, thus enabling development to move forward without being hindered by eventual unavailability of external systems you depend on

QUESTION 8

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 9

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 10

What is the minimum amount of storage that a persistent volume claim can obtain in Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE)?

- A. 1 TB
- B. 10 GB
- C. 1 GB

D. 50 GB

Correct Answer: D

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

QUESTION 11

A developer using Oracle Cloud Infrastructure (OCI) API Gateway must authenticate the API requests to their web application. The authentication process must be implemented using a custom scheme which accepts string parameters from the API caller. Which method can the developer use In this scenario?

- A. Create an authorizer function using request header authorization.
- B. Create an authorizer function using token-based authorization.
- C. Create a cross account functions authorizer.
- D. Create an authorizer function using OCI Identity and Access Management based authentication

Correct Answer: B

Having deployed the authorizer function, you enable authentication and authorization for an API deployment by including two different kinds of request policy in the API deployment specification: An authentication request policy for the entire API deployment that specifies: The OCID of the authorizer function that you deployed to Oracle Functions that will perform authentication and authorization. The request attributes to pass to the authorizer function. Whether unauthenticated callers can access routes in the API deployment. An authorization request policy for each route that specifies the operations a caller is allowed to perform, based on the caller's access scopes as returned by the authorizer function. Using the Console to Add Authentication and Authorization Request Policies To add authentication and authorization request policies to an API deployment specification using the Console: Create or update an API deployment using the Console, select the From Scratch option, and enter details on the Basic Information page. For more information, see Deploying an API on an API Gateway by Creating an API Deployment and Updating API Gateways and API Deployments. In the API Request Policies section of the Basic Information page, click the Add button beside Authentication and specify: Application in : The name of the application in Oracle Functions that contains the authorizer function. You can select an application from a different compartment. Function Name: The name of the authorizer function in Oracle Functions. Authentication Token: Whether the access token is contained in a request header or a query parameter. Authentication Token Value: Depending on whether the access token is contained in a request header or a query parameter, specify:

Header Name: If the access token is contained in a request header, enter the name of the header. Parameter Name: If the access token is contained in a query parameter, enter the name of the query parameter.

<https://docs.cloud.oracle.com/en-us/iaas/Content/APIGateway/Tasks/apigatewayaddingauthzauthn.htm>

QUESTION 12

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 13

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 14

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 15

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

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

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

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