

**100%** Money Back  
**Guarantee**

**Vendor:**Oracle

**Exam Code:**1Z0-1084-21

**Exam Name:**Oracle Cloud Infrastructure Developer  
2021 Associate

**Version:**Demo

## QUESTION 1

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 2

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 3

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 4

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 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.

---

## QUESTION 6

Which statement accurately describes Oracle Cloud Infrastructure (OCI) Load Balancer integration with OCI Container Engine for Kubernetes (OKE)?

- A. OKE service provisions an OCI Load Balancer instance for each Kubernetes service with LoadBalancer type in the YAML configuration.
- B. OCI Load Balancer instance provisioning is triggered by OCI Events service for each Kubernetes service with LoadBalancer type in the YAML configuration.
- C. OCI Load Balancer instance must be manually provisioned for each Kubernetes service that requires traffic balancing.
- D. OKE service provisions a single OCI Load Balancer instance shared with all the Kubernetes services with LoadBalancer type in the YAML configuration.

Correct Answer: D

If you are running your Kubernetes cluster on Oracle Container Engine for Kubernetes (commonly known as OKE), you can have OCI automatically provision load balancers for you by creating a Service of type LoadBalancer instead of (or in addition to) installing an ingress controller like Traefik or Voyage YAML file

```

apiVersion: v1
kind: Service
metadata:
  name: bobs-bookstore-oci-lb-service
  namespace: bob
  annotations:
    service.beta.kubernetes.io/oci-load-balancer-shape: 400Mbps
spec:
  ports:
  - name: http
    port: 31111
    protocol: TCP
    targetPort: 31111
  selector:
    weblogic.clusterName: cluster-1
    weblogic.domainUID: bobs-bookstore
  sessionAffinity: None
  type: LoadBalancer

```

When you apply this YAML file to your cluster, you will see the new service is created. After a short time (typically less than a minute) the OCI Load Balancer will be provisioned.

```

$ kubectl -n bob get svc

```

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
bobs-bookstore-admin-server	ClusterIP	None	<none>	
bobs-bookstore-admin-server-external	NodePort	10.96.224.13	<none>	7001:32401/TCP
bobs-bookstore-cluster-cluster-1	ClusterIP	10.96.86.113	<none>	8888/TCP,8001/TCP,31111/TCP
bobs-bookstore-managed-server1	ClusterIP	None	<none>	8888/TCP,8001/TCP,31111/TCP
bobs-bookstore-managed-server2	ClusterIP	None	<none>	8888/TCP,8001/TCP,31111/TCP
bobs-bookstore-oci-lb-service	LoadBalancer	10.96.121.216	132.145.235.215	31111:31671/TCP

<https://oracle.github.io/weblogic-kubernetes-operator/faq/oci-lb/>

## QUESTION 7

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 8

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 9

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 10

Your organization uses a federated identity provider to login to your Oracle Cloud Infrastructure (OCI) environment. As a developer, you are writing a script to automate some operation and want to use OCI CLI to do that. Your security team doesn't allow storing private keys on local machines.

How can you authenticate with OCI CLI?

- A. Run `oci setup keys` and provide your credentials
- B. Run `oci session refresh --profile`
- C. Run `oci session authenticate` and provide your credentials
- D. Run `oci setup oci-cli-rc --file path/to/target/file`

Correct Answer: C

Token-based authentication for the CLI allows customers to authenticate their session interactively, then use the CLI for a single session without an API signing key. This enables customers using an identity provider that is not SCIM-supported to use a federated user account with the CLI and SDKs.

Starting a Token-based CLI Session

To use token-based authentication for the CLI on a computer with a web browser:

In the CLI, run the following command. This will launch a web browser.

```
oci session authenticate
```

In the browser, enter your user credentials. This authentication information is saved to the .config file.

---

### QUESTION 11

You encounter an unexpected error when invoking the Oracle Function named "myfunction" in application "myapp". Which can you use to get more information on the error?

- A. `fn --debug invoke myapp myfunction`
- B. `DEBUG=1 fn invoke myapp myfunction`
- C. `fn --verbose invoke myapp myfunction`
- D. Call Oracle support with your error message

Correct Answer: B

Troubleshooting Oracle Functions

If you encounter an unexpected error when using an Fn Project CLI command, you can find out more about the problem by starting the command with the string `DEBUG=1` and running the command again.

For example:

```
$ DEBUG=1 fn invoke helloworld-app helloworld-func
```

Note that `DEBUG=1` must appear before the command, and that `DEBUG` must be in upper case.

---

### QUESTION 12

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": "VGhIIHF1aWNrIGJyb3dulGZveCBqdW1wZWQgb3ZlciB0aGUgbGF6eSBkb2cu" }, { "key": null, "value": "UGFjayBteSBib3ggd2l0aCBmaXZlIGRvemVulGxpcXVvciBqdWdzLg==" } } } <https://docs.cloud.oracle.com/en-us/iaas/api/#/en/streaming/20180418/Message/PutMessages>