



1Z0-1084-21^{Q&As}

Oracle Cloud Infrastructure Developer 2021 Associate

Pass Oracle 1Z0-1084-21 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/1z0-1084-21.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

In the sample Kubernetes manifest file below, what annotations should you add to create a private load balancer In oracle Cloud infrastructure Container Engine for Kubernetes?

```
apiVersion: v1
kind: Service
metacata:
  name: my-nginx-svc
  labels:
    app: nginx
  annotations:
    <Fill in>
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: nginx
```

```
apiVersion: v1
kind: Service
metacata:
  name: my-nginx-svc
  labels:
    app: nginx
  annctations:
    <Fill in>
spec:
  type: LoadBalancer
  ports:
    - port: 80
  selector:
    app: nginx
```

- A. service.beta.kubernetes.io/oci-load-balancer-private:"true"
- B. service.beta.kubernetes.io/oci-load-balancer-private: "true" service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocid1.subnet.oc1..aaaaa....vdfw"
- C. service.beta.kubernetes.io/oci-load-balancer-internal: "true"
- D. service.beta.kubernetes.io/oci-load-balancer-internal: "true" service.beta.kubernetes.io/oci-load-balancer-subnet1: "ocid1.subnet.oc1..aaaaa....vdfw"



Correct Answer: D

https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingloadbalancer.htm?TocPath=Services%7CExample%20Network%20Resource%20Configuration%7CUpgrading%20the%20Version%20of%20Kubernetes%20Running%20on%20a%20Master%20Node%7C_____2 Creating Internal Load Balancers in Public and Private Subnets You can create Oracle Cloud Infrastructure load balancers to control access to services running on a cluster: When you create a `custom` cluster, you select an existing VCN that contains the network resources to be used by the new cluster. If you want to use load balancers to control traffic into the VCN, you select existing public or private subnets in that VCN to host the load balancers. When you create a `quick cluster`, the VCN that's automatically created contains a public regional subnet to host a load balancer. If you want to host load balancers in private subnets, you can add private subnets to the VCN later.

Alternatively, you can create an internal load balancer service in a cluster to enable other programs running in the same VCN as the cluster to access services in the cluster. You can host internal load balancers in public subnets and private subnets. To create an internal load balancer hosted on a public subnet, add the following annotation in the metadata section of the manifest file: `service.beta.kubernetes.io/oci-load-balancer-internal: "true"` To create an internal load balancer hosted on a private subnet, add both following annotations in the metadata section of the manifest file: `service.beta.kubernetes.io/oci-load-balancer-internal: "true"` `service.beta.kubernetes.io/oci-load-balancersubnet1: "ocid1.subnet.oc1..aaaaa....vdfw"` where `ocid1.subnet.oc1..aaaaa....vdfw` is the OCID of the private subnet.

QUESTION 2

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 3

You have written a Node.js function and deployed it to Oracle Functions. Next, you need to call this function from a microservice written in Java deployed on Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE).

Which can help you to achieve this?

- A. Use the OCI CLI with kubectl to invoke the function from the microservice.
- B. Oracle Functions does not allow a microservice deployed on OKE to invoke a function.
- C. OKE does not allow a microservice to invoke a function from Oracle Functions.
- D. Use the OCI Java SDK to invoke the function from the microservice.

Correct Answer: D

You can invoke a function that you've deployed to Oracle Functions in different ways:

Using the Fn Project CLI.

Using the Oracle Cloud Infrastructure CLI.

Using the Oracle Cloud Infrastructure SDKs.

Making a signed HTTP request to the function's invoke endpoint. Every function has an invoke endpoint.

QUESTION 4

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 5

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 6

What is the difference between blue/green and canary deployment strategies?

- A. In blue/green, application is deployed in minor increments to a select group of people. In canary, both old and new applications are simultaneously in production.
- B. In blue/green, both old and new applications are in production at the same time. In canary, application is deployed incrementally to a select group of people.
- C. In blue/green, current applications are slowly replaced with new ones. In
- D. In blue/green, current applications are slowly replaced with new ones. In canary, both old and new applications are in production at the same time.

Correct Answer: B

Blue-green deployment is a technique that reduces downtime and risk by running two identical production environments called Blue and Green. At any time, only one of the environments is live, with the live environment serving all production traffic. For this example, Blue is currently live and Green is idle. <https://docs.cloudfoundry.org/devguide/deploy-apps/blue-green.html> Canary deployments are a pattern for rolling out releases to a subset of users or servers. The idea is to first deploy the change to a small subset of servers, test it, and then roll the change out to the rest of the servers. ... Canaries were once regularly used in coal mining as an early warning system. <https://octopus.com/docs/deployment-patterns/canary-deployments>



QUESTION 7

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 8

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 9

You are building a cloud native, serverless travel application with multiple Oracle Functions in Java, Python and Node.js. You need to build and deploy these functions to a single applications named travel-app. Which command will help you complete this task successfully?

- A. `oci fn function deploy --ap travel-ap --all`
- B. `fn deploy --ap travel-ap -- all`
- C. `oci fn application --application-name-ap deploy --all`
- D. `fn function deploy --all --application-name travel-ap`

Correct Answer: B

check the steps for Creating, Deploying, and Invoking a Helloworld Function <https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionscreatingfirst.htm> in step 7 that will deploy the funcation 7Enter the following single Fn Project command to build the function and its dependencies as a Docker image called helloworld-func, push the image to the specified Docker registry, and deploy the function to Oracle Functions in the helloworld-app: `$ fn -v deploy --app helloworld-app` The `-v` option simply shows more detail about what Fn Project commands are doing (see Using the Fn Project CLI with Oracle Functions).

QUESTION 10

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 11

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 12

A pod security policy (PSP) is implemented in your Oracle Cloud Infrastructure Container Engine for Kubernetes cluster Which rule can you use to prevent a container from running as root using PSP?

- A. NoPrivilege
- B. RunOnlyAsUser
- C. MustRunAsNonRoot
- D. forbiddenRoot

Correct Answer: C

Require the container to run without root privileges.

```
rule: \\MustRunAsNonRoot\\
```

Reference: <https://kubernetes.io/docs/concepts/policy/pod-security-policy/>

QUESTION 13

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 14

Which testing approaches is a must for achieving high velocity of deployments and release of cloud- native applications?

- A. Integration testing
- B. A/B testing
- C. Automated testing
- D. Penetration testing

Correct Answer: C

Oracle Cloud Infrastructure provides a number of DevOps tools and plug-ins for working with Oracle Cloud Infrastructure services. These can simplify provisioning and managing infrastructure or enable automated testing and continuous delivery. A/B Testing While A/B testing can be combined with either canary or blue-green deployments, it is a very different thing. A/B testing really targets testing the usage behavior of a service or feature and is typically used to validate a hypothesis or to measure two versions of a service or feature and how they stack up against each other in terms of performance, discoverability and usability. A/B testing often leverages feature flags (feature toggles), which allow you to dynamically turn features on and off. Integration Testing Integration tests are also known as end-to-end (e2e) tests. These are long-running tests that exercise the system in the way it is intended to be used in production. These are the most valuable tests in demonstrating reliability and thus increasing confidence. Penetration Testing Oracle regularly performs penetration and vulnerability testing and security assessments against the Oracle cloud infrastructure, platforms, and applications. These tests are intended to validate and improve the overall security of Oracle Cloud Services.

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/>

[1Z0-1084-21 PDF Dumps](#)

[1Z0-1084-21 Practice Test](#)

[1Z0-1084-21 Study Guide](#)