

1Z0-1084-22^{Q&As}

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.pass2lead.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

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 2

In a Linux environment, what is the default location of the configuration file that Oracle Cloud Infrastructure CLI uses for profile information?

- A. /etc/oci/config
- B. /usr/local/bin/config
- C. \$HOME/.oci/config
- D. /usr/bin/oci/config

Correct Answer: C

By default, the Oracle Cloud Infrastructure CLI configuration file is located at `~/.oci/config`. You might already have a configuration file as a result of installing the Oracle Cloud Infrastructure CLI.

QUESTION 3

With the volume of communication that can happen between different components in cloud-native applications, it is vital to not only test functionality, but also service resiliency. Which statement is true with regards to service resiliency?

- A. Resiliency is about recovering from failures without downtime or data loss.
- B. A goal of resiliency is not to bring a service to a functioning state after a failure.
- C. Resiliency testing can be only done in a test environment.
- D. Resiliency is about avoiding failures.

Correct Answer: D

Resiliency and Availability

Resiliency and availability refers to the ability of a system to continue operating, despite the failure or sub-optimal performance of some of its components.

In the case of Oracle Functions:

The control plane is a set of components that manages function definitions. The data plane is a set of components that executes functions in response to invocation requests. For resiliency and high availability, both the control plane and data plane components are distributed across different availability domains and fault domains in a region. If one of the domains ceases to be available, the components in the remaining domains take over to ensure that function definition management and execution are not disrupted. When functions are invoked, they run in the subnets specified for the application to which the functions belong. For resiliency and high availability, best practice is to specify a regional subnet for an application (or alternatively, multiple AD-specific subnets in different availability domains). If an availability domain specified for an application ceases to be available, Oracle Functions runs functions in an alternative availability domain.

QUESTION 4

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/en-us/iaas/Content/Database/Tasks/adbtools.htm>

QUESTION 5

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:

- 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.
- In a terminal window, enter: `$ kubectl create secret docker-registry --docker-server=.ocir.io --docker-username=\\' --docker-password=\\' --docker-email=\\'` 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, `acme-dev`). `.` 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, `kj64r{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.