

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.leads4pass.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

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 2

How do you perform a rolling update in Kubernetes?

- A. `kubect1 rolling-update`
- B. `kubect1 upgrade --image=*image:v2`
- C. `kubect1 update -c`
- D. `kubect1 rolling-update --image=image`

Correct Answer: A

<https://docs.oracle.com/en/cloud/iaas/wercker-cloud/wercm/quickstarts/platforms/kubernetes/>

QUESTION 3

You are building a container image and pushing it to the Oracle Cloud Infrastructure Registry (OCIR). You need to make sure that these get deleted from the repository.

Which action should you take?

- A. Create a group and assign a policy to perform lifecycle operations on images.
- B. Set global policy of image retention to "Retain All Images".
- C. In your compartment, write a policy to limit access to the specific repository.
- D. Edit the tenancy global retention policy.

Correct Answer: D

Deleting an Image When you no longer need an old image or you simply want to clean up the list of image tags in a repository, you can delete images from Oracle Cloud Infrastructure Registry. Your permissions control the images in Oracle Cloud Infrastructure Registry that you can delete. You can delete images from repositories you've created, and from repositories that the groups to which you belong have been granted access by identity policies. If you belong to the Administrators group, you can delete images from any repository in the tenancy. Note that as well deleting individual images, you can set up image retention policies to delete images automatically based on selection criteria you specify (see [Retaining and Deleting Images Using Retention Policies](#)). Note: In each region in a tenancy, there's a global image retention policy. The global image retention policy's default selection criteria retain all images so that no images are automatically deleted.

However, you can change the global image retention policy so that images are deleted if they meet the criteria you specify. A region's global image retention policy applies to all repositories in the region, unless it is explicitly overridden by one or more custom image retention policies. You can set up custom image retention policies to override the global image retention policy with different criteria for specific repositories in a region. Having created a custom image retention policy, you apply the custom retention policy to a repository by adding the repository to the policy. The global image retention policy no longer applies to repositories that you add to a custom retention policy.

QUESTION 4

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 5

Your Oracle Cloud Infrastructure Container Engine for Kubernetes (OKE) administrator has created an OKE cluster with one node pool in a public subnet. You have been asked to provide a log file from one of the nodes for troubleshooting purpose.

Which step should you take to obtain the log file?

- A. ssh into the node using public key.
- B. ssh into the nodes using private key.
- C. It is impossible since OKE is a managed Kubernetes service.
- D. Use the username open and password to login.

Correct Answer: B

Kubernetes cluster is a group of nodes. The nodes are the machines running applications. Each node can be a physical machine or a virtual machine. The node's capacity (its number of CPUs and amount of memory) is defined when the node is created. A cluster comprises: - one or more master nodes (for high availability, typically there will be a number of master nodes) - one or more worker nodes (sometimes known as minions) Connecting to Worker Nodes Using SSH If you provided a public SSH key when creating the node pool in a cluster, the public key is installed on all worker nodes in the cluster. On UNIX and UNIX-like platforms (including Solaris and Linux), you can then connect through SSH to the worker nodes using the ssh utility (an SSH client) to perform administrative tasks. Note the following instructions assume the UNIX machine you use to connect to the worker node: Has the ssh utility installed. Has access to the SSH private key file paired with the SSH public key that was specified when the cluster was created. How to connect to worker nodes using SSH depends on whether you specified public or private subnets for the worker nodes when defining the node pools in the cluster. Connecting to Worker Nodes in Public Subnets Using SSH Before you can connect to a worker node in a public subnet using SSH, you must define an ingress rule in the subnet's security list to allow SSH access. The ingress rule must allow access to port 22 on worker nodes from source 0.0.0.0/0 and any source port To connect to a worker node in a public subnet through SSH from a UNIX machine using the ssh utility: 1- Find out the IP address of the worker node to which you want to connect. You can do this in a number of ways: Using kubectl. 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. See Setting Up Cluster Access. Then in a terminal window, enter kubectl get nodes to see the public IP addresses of worker nodes in node pools in the cluster. Using the Console. In the Console, display the Cluster List page and then select the cluster to which the worker node belongs. On the Node Pools tab, click the name of the node pool to which the worker node belongs. On the Nodes tab, you see the public IP address of every worker node in the node pool. Using the REST API. Use the ListNodePools operation to see the public IP addresses of worker nodes in a node pool. 2- In the terminal window, enter ssh opc@ to connect to the worker node, where is the IP address of the worker node that you made a note of earlier. For example, you might enter ssh opc@192.0.2.254. Note that if the SSH private key is not stored in the file or in the path that the ssh utility expects (for example, the ssh utility might expect the private key to be stored in ~/.ssh/id_rsa), you must explicitly specify the private key filename and location in one of two ways: Use the -i option to specify the filename and location of the private key. For example, ssh -i ~/.ssh/my_keys/my_host_key_filename opc@192.0.2.254 Add the private key filename and location to an SSH

configuration file, either the client configuration file (~/.ssh/config) if it exists, or the system-wide client

configuration file (/etc/ssh/ssh_config). For example, you might add the following:

```
Host 192.0.2.254 IdentityFile ~/.ssh/my_keys/my_host_key_filename
```

For more about the ssh utility's configuration file, enter man ssh_config Note also that permissions on the private key file must allow you read/write/execute access, but prevent other users from accessing the file.

For example, to set appropriate permissions, you might enter `chmod 600 ~/.ssh/my_keys/`

`my_host_key_filename`. If permissions are not set correctly and the private key file is accessible to other users, the ssh utility will simply ignore the private key file.

QUESTION 6

Which two are benefits of distributed systems?

- A. Privacy
- B. Security
- C. Ease of testing
- D. Scalability
- E. Resiliency

Correct Answer: DE

distributed systems of native-cloud like functions that have a lot of benefit like Resiliency and availability Resiliency and availability refers to the ability of a system to continue operating, despite the failure or suboptimal 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. Concurrency and Scalability Concurrency refers to the ability of a system to run multiple operations in parallel using shared resources. Scalability refers to the ability of the system to scale capacity (both up and down) to meet demand. In the case of Functions, when a function is invoked for the first time, the function's image is run as a container on an instance in a subnet associated with the application to which the function belongs. When the function is executing inside the container, the function can read from and write to other shared resources and services running in the same subnet (for example, Database as a Service). The function can also read from and write to other shared resources (for example, Object Storage), and other Oracle Cloud Services. If Oracle Functions receives multiple calls to a function that is currently executing inside a running container, Oracle Functions automatically and seamlessly scales horizontally to serve all the incoming requests. Oracle Functions starts multiple Docker containers, up to the limit specified for your tenancy. The default limit is 30 GB of RAM reserved for function execution per availability domain, although you can request an increase to this limit. Provided the limit is not exceeded, there is no difference in response time (latency) between functions executing on the different containers.

QUESTION 7

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>

QUESTION 8

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 9

You have a containerized app that requires an Autonomous Transaction Processing (ATP) Database. Which option is not valid for o from a container in Kubernetes?

A. Enable Oracle REST Data Services for the required schemas and connect via HTTPS.

B. Create a Kubernetes secret with contents from the instance Wallet files. Use this secret to create a volume mounted to the appropriate path in the application deployment manifest.

C. Use Kubernetes secrets to configure environment variables on the container with ATP instance OCID, and OCI API credentials. Then use the CreateConnection API endpoint from the service runtime.

D. Install the Oracle Cloud Infrastructure Service Broker on the Kubernetes cluster and deploy serviceinstance and serviceBinding resources for ATP. Then use the specified binding name as a volume in the application deployment manifest.

Correct Answer: A

<https://blogs.oracle.com/developers/creating-an-atp-instance-with-the-oci-service-broker> <https://blogs.oracle.com/cloud-infrastructure/integrating-oci-service-broker-with-autonomous-transaction-processing-in-the-real-world>

QUESTION 10

Which pattern can help you minimize the probability of cascading failures in your system during partial loss of connectivity or a complete service failure?

- A. Retry pattern
- B. Anti-corruption layer pattern
- C. Circuit breaker pattern
- D. Compensating transaction pattern

Correct Answer: C

A cascading failure is a failure that grows over time as a result of positive feedback. It can occur when a portion of an overall system fails, increasing the probability that other portions of the system fail. The circuit breaker pattern prevents the service from performing an operation that is likely to fail. For example, a client service can use a circuit breaker to prevent further remote calls over the network when a downstream service is not functioning properly. This can also prevent the network from becoming congested by a sudden spike in failed retries by one service to another, and it can also prevent cascading failures. Self-healing circuit breakers check the downstream service at regular intervals and reset the circuit breaker when the downstream service starts functioning properly.

<https://blogs.oracle.com/developers/gettingstarted-with-microservices-part-three>

QUESTION 11

What can you use to dynamically make Kubernetes resources discoverable to public DNS servers?

- A. ExternalDNS
- B. CoreDNS
- C. DynDNS
- D. kubeDNS

Correct Answer: A

ExternalDNS allows you to control DNS records dynamically via Kubernetes resources in a DNS provider-agnostic way

<https://github.com/kubernetes-sigs/external-dns/blob/master/README.md> <https://github.com/kubernetes-sigs/external-dns/blob/master/docs/tutorials/oracle.md>

QUESTION 12

Which Oracle Cloud Infrastructure (OCI) load balancer shape is used by default in OCI container Engine for Kubernetes?

- A. 400 Mbps
- B. 8000 Mbps
- C. There is no default. The shape has to be specified.
- D. 100 Mbps

Correct Answer: D

Specifying Alternative Load Balancer Shapes The shape of an Oracle Cloud Infrastructure load balancer specifies its maximum total bandwidth (that is, ingress plus egress). By default, load balancers are created with a shape of 100Mbps. Other shapes are available, including 400Mbps and 8000Mbps. <https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Tasks/contengcreatingloadbalancer.htm>

QUESTION 13

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 14

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 15

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>