

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.leads4pass.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 working on a cloud native e-commerce application on Oracle Cloud Infrastructure (OCI). Your application architecture has multiple OCI services, including Oracle Functions. You need to trigger these functions directly from other OCI services, without having to run custom code. Which OCI service cannot trigger your functions directly?

- A. OCI Events Service
- B. OCI Registry
- C. OCI API Gateway
- D. Oracle Integration

Correct Answer: B

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. With Oracle Functions, you can write code in Java, Python, Node, Go, and Ruby (and for advanced use cases, bring your own Dockerfile, and Graal VM).

You can invoke a function that you've deployed to Oracle Functions from:

-

The Fn Project CLI.

-

The Oracle Cloud Infrastructure SDKs.

-

Signed HTTP requests to the function's invoke endpoint. Every function has an invoke endpoint.

-

Other Oracle Cloud services (for example, triggered by an event in the Events service) or from external services. so You can then deploy your code, call it directly or trigger it in response to events, and get billed only for the resources consumed during the execution. Below are the oracle services that can trigger Oracle functions -Events Service -Notification Service -API Gateway Service -Oracle Integration service(using OCI Signature Version 1 security policy) so OCI Registry services cannot trigger your functions directly

QUESTION 2

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 3

You need to execute a script on a remote instance through Oracle Cloud Infrastructure Resource Manager. Which option can you use?

- A. Use /bin/sh with the full path to the location of the script to execute the script.
- B. It cannot be done.
- C. Download the script to a local desktop and execute the script.
- D. Use remote-exec

Correct Answer: D

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. With Resource Manager, you can use Terraform's remote exec functionality to execute scripts or commands on a remote computer. You can also use this technique for other provisioners that require access to the remote resource.

QUESTION 4

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
  arnotations:
    <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 5

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 6

You are developing a serverless application with Oracle Functions. Your function needs to store state in a database. Your corporate security Standards mandate encryption of secret information like database passwords.

As a function developer, which approach should you follow to satisfy this security requirement?

- A. Use the Oracle Cloud Infrastructure Console and enter the password in the function configuration section in the provided input field.
- B. Use Oracle Cloud Infrastructure Key Management to auto-encrypt the password. It will inject the auto-decrypted password inside your function container.

C. Encrypt the password using Oracle Cloud Infrastructure Key Management. Decrypt this password in your function code with the generated key.

D. All function configuration variables are automatically encrypted by Oracle Functions.

Correct Answer: A

Passing Custom Configuration Parameters to Functions The code in functions you deploy to Oracle Functions will typically require values for different parameters. Some pre-defined parameters are available to your functions as environment variables. But you'll often want your functions to use parameters that you've defined yourself. For example, you might create a function that reads from and writes to a database. The function will require a database connect string, comprising a username, password, and hostname. You'll probably want to define username, password, and hostname as parameters that are passed to the function when it's invoked. Using the Console To specify custom configuration parameters to pass to functions using the Console: Log in to the Console as a functions developer. In the Console, open the navigation menu. Under Solutions and Platform, go to Developer Services and click Functions. Select the region you are using with Oracle Functions. Oracle recommends that you use the same region as the Docker registry that's specified in the Fn Project CLI context (see 6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure). Select the compartment specified in the Fn Project CLI context (see 6. Create an Fn Project CLI Context to Connect to Oracle Cloud Infrastructure). The Applications page shows the applications defined in the compartment. Click the name of the application containing functions to which you want to pass custom configuration parameters: To pass one or more custom configuration parameters to every function in the application, click Configuration to see the Configuration section for the application. To pass one or more custom configuration parameters to a particular function, click the function's name to see the Configuration section for the function. In the Configuration section, specify details for the first custom configuration parameter: Key: The name of the custom configuration parameter. The name must only contain alphanumeric characters and underscores, and must not start with a number. For example, username Value: A value for the custom configuration parameter. The value must only contain printable unicode characters. For example, jdoe Click the plus button to save the new custom configuration parameter. Oracle Functions combines the key-value pairs for all the custom configuration parameters (both application-wide and function-specific) in the application into a single, serially-encoded configuration object with a maximum allowable size of 4Kb. You cannot save the new custom configuration parameter if the size of the serially-encoded configuration object would be greater than 4Kb. (Optional) Enter additional custom configuration parameters as required.

QUESTION 7

You are developing a serverless application with Oracle Functions. You have created a function in compartment named prod. When you try to invoke your function you get the following error. Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions.oc1.phx.aaaaaaac... does not exist or Oracle Functions is not authorized to use it How can you resolve this error?

- A. Create a policy:Allow function-family to use virtual-network-family in compartment prod
- B. Create a policy:Allow any-user to manage function-family and virtual-network-family in compartment prod
- C. Create a policy:Allow service FaaS to use virtual-network-family in compartment prod
- D. Deleting the function and redeploying it will fix the problem

Correct Answer: C

Invoking a function returns a FunctionInvokeSubnetNotAvailable message and a 502 error (due to a DHCP

Options issue)

When you invoke a function that you've deployed to Oracle Functions, you might see the following error

message:

```
{"code":"FunctionInvokeSubnetNotAvailable","message":"dhcp options ocid1.dhcpoptions..... does not exist or Oracle Functions is not authorized to use it"}
```

Fn: Error invoking function. status: 502 message: dhcp options ocid1.dhcpoptions..... does not exist or Oracle Functions is not authorized to use it If you see this error:

Double-check that a policy has been created to give Oracle Functions access to network resources.

Service Access to Network Resources

When Oracle Functions users create a function or application, they have to specify a VCN and a subnet in which to create them. To enable the Oracle Functions service to create the function or application in the specified VCN and subnet, you must create an identity policy to grant the Oracle Functions service access to the compartment to which the network resources belong. To create a policy to give the Oracle Functions service access to network resources:

Log in to the Console as a tenancy administrator.

Create a new policy in the root compartment:

Open the navigation menu. Under Governance and Administration, go to Identity and click Policies. Follow the instructions in To create a policy, and give the policy a name (for example, functions- service-networkaccess).

Specify a policy statement to give the Oracle Functions service access to the network resources in the compartment:

Allow service FaaS to use virtual-network-family in compartment For example:

Allow service FaaS to use virtual-network-family in compartment acme-network Click Create.

Double-check that the set of DHCP Options in the VCN specified for the application still exists.

QUESTION 8

How do you perform a rolling update in Kubernetes?

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

Correct Answer: A

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

QUESTION 9

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/>

QUESTION 10

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 11

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 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": "VGhllHF1aWNrIGJyb3dulGZveCBqdW1wZWQgb3ZlciB0aGUgbGF6eSBkb2cu" }, { "key": null, "value": "UGFjayBteSBib3ggd2l0aCBmaXZlIGRvemVulGxpcXVvciBqdWdzLg==" } } } <https://docs.cloud.oracle.com/en-us/iaas/api/#/en/streaming/20180418/Message/PutMessages>

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

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 15

As a cloud-native developer, you are designing an application that depends on Oracle Cloud Infrastructure (OCI) Object Storage wherever the application is running. Therefore, provisioning of storage buckets should be part of your Kubernetes deployment process for the application. Which should you leverage to meet this requirement?

- A. OCI Service Broker for Kubernetes
- B. OCI Container Engine for Kubernetes
- C. Open Service Broker API

D. Oracle Functions

Correct Answer: A

<https://blogs.oracle.com/cloud-infrastructure/introducing-service-broker-for-kubernetes> OCI Service Broker for Kubernetes is an implementation of the Open Service Broker API. OCI Service Broker for Kubernetes is specifically for interacting with Oracle Cloud Infrastructure services from Kubernetes clusters. It includes three service broker adapters to bind to the following Oracle Cloud Infrastructure services: Object Storage Autonomous Transaction Processing Autonomous Data Warehouse

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

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

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