# 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.passapply.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 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/j son

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": "VGhlIHF1aWNrIGJyb3duIGZveCBqdW1wZWQgb3ZlciB0aGUgbGF6eSBkb2cu" }, { "key": null, "value": "UGFjayBteSBib3ggd2l0aCBmaXZlIGRvemVuIGxpcXVvciBqdWdzLg==" } } } https://docs.cloud.oracle.com/en-us/iaas/api/#/en/streaming/20180418/Message/PutMessages

**QUESTION 2**

You are using Oracle Cloud Infrastructure (0CI) Resource Manager to manage your infrastructure lifecycle and wish to receive an email each time a Terraform action begins. How should you use the OCI Events service to do this without writing any code?

A. Create an OCI Notifications topic and email subscription with the destination email address. Then create an OCI Events rule matching "Resource Manager Stack - Update" condition, and select the notification topic for the corresponding action.

B. Create an OCI Notification topic and email subscription with the destination email address. Then create an OCI Events rule matching "Resource Manager job - Create" condition, and select the notification topic for the corresponding action.

C. Create a rule in OCI Events service matching the "Resource Manager Stack - Update" condition.Then select "Action Type: Email" and provide the destination email address.

D. Create an OCI Email Delivery configuration with the destination email address. Then create an OCI Events rule matching "Resource Manager Job - Create" condition, and select the email configuration for the corresponding action.

Correct Answer: B

1.

Create Notifications Topic and Subscription If a suitable Notifications topic doesn\\'t already exist, then you must log in

to the Console as a tenancy administrator and create it. Whether you use an existing topic or create a new one, add an email address as a subscription so that you can monitor that email account for notifications

2.

Using the Console to Create a Rule Use the Console to create a rule with a pattern that matches bucket creation events emitted by Object Storage. Specify the Notifications topic you created as an action to deliver matching events. To test your rule, create a bucket. Object Storage emits an event which triggers the action. Check the email specified in the subscription to receive your notification

https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Concepts/eventsgetstarted.htm
https://docs.cloud.oracle.com/en-us/iaas/Content/Events/Concepts/filterevents.htm

**QUESTION 3**

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 4**

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 5**

Which is NOT a supported SDK on Oracle Cloud Infrastructure (OCI)?

A. Ruby SDK

B. Java SDK

C. Python SDK

D. Go SDK

E. .NET SDK

Correct Answer: E

https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/sdks.htm

**QUESTION 6**

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 7**

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 8**

Given a service deployed on Oracle Cloud infrastructure Container Engine for Kubernetes (OKE), which annotation should you add in the sample manifest file 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. lo/oci-load-balancer-kind: 400Mbps

B. service, beta, kubernetes. lo/oci-load-balancer-value: 4 00Mbps

C. service . beta. kubernetes . lo/oci-load-balancer-shape: 400Mbps

D. service . beta . kubernetes . lo/oci-load-balancer-size: 400Mbps

Correct Answer: C

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.

To specify an alternative shape for a load balancer, add the following annotation in the metadata section of the manifest file:

service.beta.kubernetes.io/oci-load-balancer-shape: where value is the bandwidth of the shape (for example, 100Mbps, 400Mbps, 8000Mbps).

For example:

apiVersion: v1

kind: Service

metadata:

name: my-nginx-svc

labels:

app: nginx

annotations:

service.beta.kubernetes.io/oci-load-balancer-shape: 400Mbps spec:

type: LoadBalancer

ports:

-port: 80 selector: app: nginx https://github.com/oracle/oci-cloud-controller-manager/blob/master/docs/load-balancer-annotations.md

**QUESTION 9**

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 10**

A programmer Is developing a Node is application which will run in a Linux server on their on-premises

data center. This application will access various Oracle Cloud Infrastructure (OC1) services using OCI

SDKs.

What is the secure way to access OCI services with OCI Identity and Access Management (JAM)?

A. Create a new OCI IAM user associated with a dynamic group and a policy that grants the desired permissions to OCI services. Add the on-premises Linux server in the dynamic group.

B. Create an OCI IAM policy with the appropriate permissions to access the required OCI services and assign the policy to the on-premises Linux server.

C. Create a new OCI IAM user, add the user to a group associated with a policy that grants the desired permissions to OCI services. In the on-premises Linux server, generate the keypair used for signing API requests and upload the public key to the IAM user.

D. Create a new OCI IAM user, add the user to a group associated with a policy that grants the desired permissions to OCI services. In the on-premises Linux server, add the user name and password to a file used by Node.js authentication.

Correct Answer: C

Before using Oracle Functions, you have to set up an Oracle Cloud Infrastructure API signing key. The instructions in this topic assume:

-you are using Linux

- you are following Oracle\\'s recommendation to provide a passphrase to encrypt the private key For more Detials Set up an Oracle Cloud Infrastructure API Signing Key for Use with Oracle Functions

https://docs.cloud.oracle.com/en-us/iaas/Content/Functions/Tasks/functionssetupapikey.htm

**QUESTION 11**

What is the minimum amount of storage that a persistent volume claim can obtain In Oracle Cloud Infrastructure Container Engine for Kubemetes (OKE)?

A. 1 TB

B. 10 GB

C. 1 GB

D. 50 GB

Correct Answer: D

https://docs.cloud.oracle.com/en-us/iaas/Content/ContEng/Concepts/contengprerequisites.htm

**QUESTION 12**

A service you are deploying to Oracle infrastructure (OCI) Container En9ine 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: 1- 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. 2- In a terminal window, enter: $ kubectl create secret docker-registry --docker-server=.ocir.io --dockerusername=\\'/\\' --docker-password=\\'\\' --dockeremail=\\'\\' 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, acmedev). 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, k]j64r{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.

**QUESTION 13**

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 Kubermetes?

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

apiVersion: v1
kind: Service
metadata:
 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%2 0Network%20Resource%20Configuration%7CUpgrading%20the%20Version%20of%20Kubernetes%2 0Running%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..aaaaaa....vdfw" where ocid1.subnet.oc1..aaaaaa....vdfw is the OCID of the private subnet.

**QUESTION 14**

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 15**

You are developing a polyglot serverless application using Oracle Functions. Which language cannot be used to write your function code?

A. PL/SQL

B. Python

C. Node.js

D. Java

Correct Answer: A

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 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.