

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.pass2lead.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 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 2

A programmer is developing a Node.js application which will run in a Linux server on their on-premises data center. This application will access various Oracle Cloud Infrastructure (OCI) services using OCI SDKs.

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

- 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 Details 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 3

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 4

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 5

You are a consumer of Oracle Cloud Infrastructure (OCI) Streaming service. Which API should you use to read and process the stream?

- A. ListMessages
- B. GetMessages
- C. GetObject
- D. ReadMessages

Correct Answer: B

<https://docs.cloud.oracle.com/en-us/iaas/Content/Streaming/Concepts/streamingoverview.htm> Building consumers to read and process messages from a stream using the GetMessage API.

QUESTION 6

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 7

What is the minimum amount of storage that a persistent volume claim can obtain In Oracle Cloud Infrastructure Container Engine for Kubernetes (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 8

You are deploying an API via Oracle Cloud Infrastructure (OCI) API Gateway and you want to implement request policies to control access Which is NOT available in OCI API Gateway?

- A. Limiting the number of requests sent to backend services

- B. Enabling CORS (Cross-Origin Resource Sharing) support
- C. Providing authentication and authorization
- D. Controlling access to OCI resources

Correct Answer: D

In the API Gateway service, there are two types of policy:

-

a request policy describes actions to be performed on an incoming request from a caller before it is sent to a back end

-

a response policy describes actions to be performed on a response returned from a back end before it is sent to a caller

You can use request policies to:

- limit the number of requests sent to back-end services
- enable CORS (Cross-Origin Resource Sharing) support
- provide authentication and authorization

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

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 11

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 12

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. io/oci-load-balancer-kind: 400Mbps
- B. service, beta, kubernetes. io/oci-load-balancer-value: 4 00Mbps
- C. service . beta. kubernetes . io/oci-load-balancer-shape: 400Mbps
- D. service . beta . kubernetes . io/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 13

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.

QUESTION 14

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 15

You have created a repository in Oracle Cloud Infrastructure Registry in the us-ashburn-1 (iad) region in your tenancy with a namespace called "heyoci.

Which three are valid tags for an image named "myapp"?

- A. iad.ocir.io/heyoci/myproject/myapp:0.0.1
- B. us-ashburn-1.ocir.io/heyoci/myapp:0.0.2-beta
- C. us-ashburn-1.ocir.io/heyoci/myproject/myapp:0.0.2-beta
- D. us-ashburn-1.ocir.io/myproject/heyoci/myapp:latest
- E. iad.ocir.io/myproject/heyoci/myapp:latest
- F. iad.ocir.io/heyoci/myapp:0.0.2-beta
- G. iad.ocir.io/heyoci/myapp:latest

Correct Answer: AFG

Give a tag to the image that you're going to push to Oracle Cloud Infrastructure Registry by entering:

docker tag

where:

uniquely identifies the image, either using the image's id (for example, 8e0506e14874), or the image's name and tag separated by a colon (for example, acme-web-app:latest). is in the format .ocir.io///: where: 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 that owns the repository to which you want to push 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). Note also that your user must have access to the tenancy. (if specified) is the name of a repository to which you want to push the image (for example, project01). Note that specifying a repository is optional (see About Repositories). is the name you want to give the image in Oracle Cloud Infrastructure Registry (for example, acme-web-app). is an image tag you want to give the image in Oracle Cloud Infrastructure Registry (for example, version2.0.test). For example, for convenience you might want to group together multiple versions of the acme-web-app image in the acme-dev tenancy in the Ashburn region into a repository called project01. You do this by including the name of the repository in the image name when you push the image, in the format .ocir.io///. For example, iad.ocir.io/ansh81vru1zp/project01/acme-web-app:4.6.3. Subsequently, when you use the docker push command, the presence of the repository in the image's name ensures the image is pushed to the intended repository. If you push an image and include the name of a repository that doesn't already exist, a new private repository is created automatically. For example, if you enter a command like docker push iad.ocir.io/ansh81vru1zp/project02/acme-web-app:7.5.2 and the project02 repository doesn't exist, a private repository called project02 is created automatically. If you push an image and don't include a repository name, the image's

name is used as the name of the repository. For example, if you enter a command like docker push iad.ocir.io/ansh81vru1zp/acme-web-app:7.5.2 that doesn't contain a repository name, the image's name (acme-web-app) is used as the name of a private repository. [https:// docs.cloud.oracle.com/en-us/iaas/Content/Registry/Concepts/registrywhatisarepository.htm](https://docs.cloud.oracle.com/en-us/iaas/Content/Registry/Concepts/registrywhatisarepository.htm)

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

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

[1Z0-1084-21 Exam Questions](#)