



# 1Z0-1084-21<sup>Q&As</sup>

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.pass4itsure.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

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 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

You are processing millions of files in an Oracle Cloud Infrastructure (OCI) Object Storage bucket. Each time a new file is created, you want to send an email to the customer and create an order in a database. The solution should perform and minimize cost, Which action should you use to trigger this email?

A. Schedule a cron job that monitors the OCI Object Storage bucket and emails the customer when a new file is created.

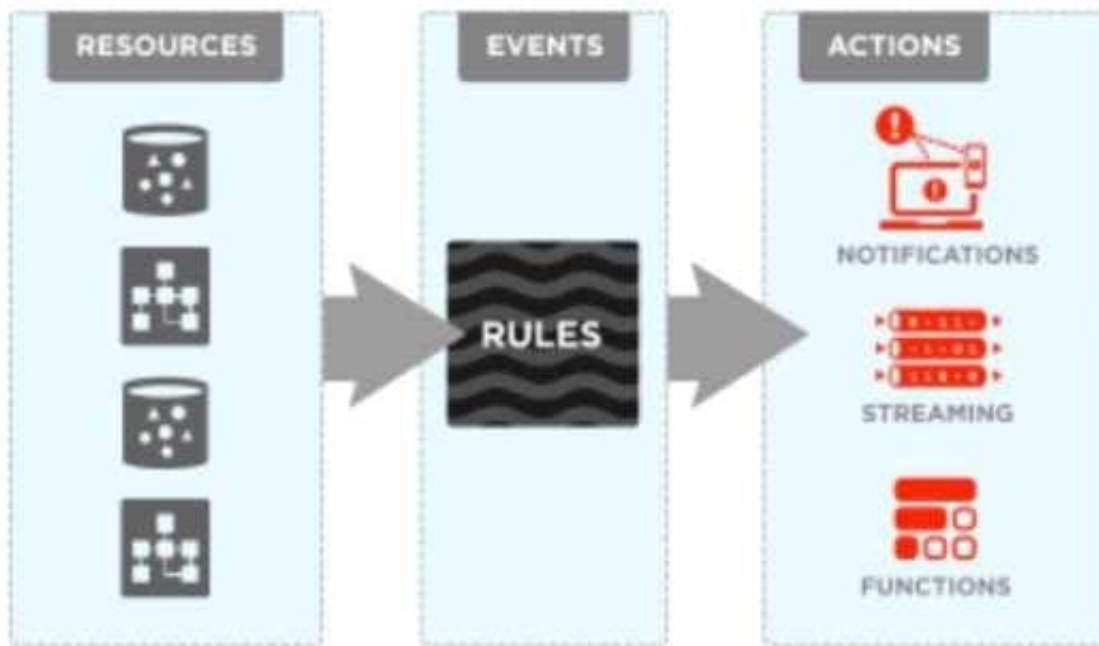
B. Use OCI Events service and OCI Notification service to send an email each time a file is created.

C. Schedule an Oracle Function that checks the OCI Object Storage bucket every minute and emails the customer when a file is found.

D. Schedule an Oracle Function that checks the OCI Object Storage bucket every second and emails the customer when a file is found.

Correct Answer: B

Oracle Cloud Infrastructure Events enables you to create automation based on the state changes of resources throughout your tenancy. Use Events to allow your development teams to automatically respond when a resource changes its state. Here are some examples of how you might use Events: Send a notification to a DevOps team when a database backup completes. Convert files of one format to another when files are uploaded to an Object Storage bucket. You can only deliver events to certain Oracle Cloud Infrastructure services with a rule. Use the following services to create actions: Notifications Streaming Functions



#### QUESTION 4

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 5

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>

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

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

[1Z0-1084-21 Study Guide](#)