



# 1Z0-997<sup>Q&As</sup>

Oracle Cloud Infrastructure 2019 Architect Professional

## Pass Oracle 1Z0-997 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/1z0-997.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 have multiple IAM users who launch different types of compute Instances and block volumes every day. As a result, your Oracle cloud Infrastructure (OCI) tenancy quickly hit the service limit and you can no longer create any new instances. As you are cleaning up environment, you notice that the majority of the Instances and block volumes are untagged. Therefore, It is difficult to pinpoint the owner of these resources verify if they are safe to terminate. Because of this, your company has issued a new mandate, which requires adding compute instances. Which option is the simplest way to implement this new requirement?

- A. Create a policy to automatically tag a resource with the user name.
- B. Create a policy using IAM requiring users to tag specific resources. This will allow a user to launch compute instances only if certain tags were defined.
- C. Create tag variables to automatically tag a resource with the user name.
- D. Create a default tag for each compartment, which ensure that appropriate tags are applied at resource creation
- E. Create tag variables for each compartment to automatically tag a resource with the user name.

Correct Answer: C

**Tag Variables** You can use a variable to set the value of a defined tag. When you add the tag to a resource, the variable resolves to the data it represents. You can use tag variables in defined tags and default tags.

**Supported Tag Variables** The following tag variables are supported. `${iam.principal.name}` The name of the principal that tagged the resource `${iam.principal.type}` The type of principal that tagged the resource. `${oci.datetime}` The date and time that the tag was created. Consider the following example: `Operations.CostCenter=" ${iam.principal.name} at ${oci.datetime} "` Operations is the namespace, CostCenter is the tag key, and the tag value contains two tag variables `${iam.principal.name}` and `${oci.datetime}` . When you add this tag to a resource, the variable resolves to your user name (the name of the principal that applied the tag) and a time date stamp for when you added the tag. `user_name at 2019-06-18T18:00:57.604Z` The variable is replaced with data at the time you apply the tag. If you later edit the tag, the variable is gone and only the data remains. You can edit the tag value in all the ways you would edit any other tag value. To create a tag variable, you must use a specific format. `${}` Type a dollar sign followed by open and close curly brackets. The tag variable goes between the curly brackets. You can use tag variables with other tag variables and with string values. Tag defaults let you specify tags to be applied automatically to all resources, at the time of creation, in a specific compartment. This feature allows you to ensure that appropriate tags are applied at resource creation without requiring the user who is creating the resource to have access to the tag namespaces. <https://docs.cloud.oracle.com/en-us/iaas/Content/Tagging/Tasks/managingtagdefaults.htm>

---

## QUESTION 2

You are working with a social media company as a solution architect. The media company wants to collect and analyze large amounts of data being generated from their websites and social media feeds to gain insights and continuously improve the user experience. In order to meet this requirement, you have developed a microservices application hosted on Oracle Container Engine for Kubernetes. The application will process the data and store the result to an Autonomous Data Warehouse (ADW) instance. Which Oracle Cloud Infrastructure (OCI) service can you use to collect and process a large volume of unstructured data in real time?

- A. OCI Events
- B. OCI Streaming



C. OCI Resource Manager

D. OCI Notifications

Correct Answer: B

---

### QUESTION 3

An OCI Architect is working on a solution consisting of analysis of data from clinical trials of a pharmaceutical company. The data is being stored in OCI Autonomous Data Warehouse (ADW) having 8 CPU Cores and 70 TB of storage. The architect is planning to setup autoscaling to respond to dynamic changes in the workload. Which of the following needs to be considered while configuring auto scaling? Choose two

A. Enabling auto scaling does not change the concurrency and parallelism settings

B. Auto scaling also scales IO throughput linearly along with CPU

C. The database memory SGA and PGA will not get affected by the changes in the number of CPUs during auto scaling

D. The maximum CPU cores that will be automatically allocated for this database is 16 CPUs

Correct Answer: AB

Auto scaling is enabled by default when you create an Autonomous Database instance or you can use Scale Up/Down on the Oracle Cloud Infrastructure console to enable or disable auto scaling. With auto scaling enabled the database can use up to three times more CPU and IO resources than specified by the number of OCPUs currently shown in the Scale Up/Down dialog. When auto scaling is enabled, if your workload requires additional CPU and IO resources the database automatically uses the resources without any manual intervention required. Enabling auto scaling does not change the concurrency and parallelism settings for the predefined services IO throughput depends on the number of CPUs you provision and scales linearly with the number of CPUs.

---

### QUESTION 4

You are working as a solution architect for an online retail store to create a portal to allow the users to pay for their groceries using credit cards. Since the application is not fully compliant with the Payment Card Industry Data Security Standard (PCI DSS), your company is looking to use a third party payment service to process credit card payments. The third party service allows a maximum of 5 public IP addresses at a time. However, your website is using Oracle Cloud Infrastructure (OCI) Instance Pool Auto Scaling policy to create up to 15 Instances during peak traffic demand, which are launched in VCN private subnets and attached to an OCI public Load Balancer. Upon user payment, the portal connects to the payment service over the Internet to complete the transaction. What solution can you implement to make sure that all compute Instances can connect to the third party system to process the payments at peak traffic demand?

A. Route credit card payment request from the compute instances through the NAT Gateway. On the third-party services, whitelist the public IP associated with the NAT Gateway.

B. Whitelist the Internet Gateway Public IP on the third party service and route all payment requests through the Internet Gateway.

C. Create an OCI Command Line Interface (CLI) script to automatically reserve public IP address for the compute instances. On the third services, whitelist the Reserved public IP.



D. Route payment request from the compute instances through the OCI Load Balancer, which will then be routed to the third party service.

Correct Answer: D

You can OCI Load Balancer for this solution which can you the Public IPs of Load balancer to Traffic to third party services which allows a maximum of Spelunk IP addresses 5 public IP addresses at a time However, your website is using Oracle Cloud Infrastructure (OCI) Instance Pool Auto Scaling policy to create up to 15 Instances during peak traffic demand

---

## QUESTION 5

As a part of migration exercise for an existing on premises application to Oracle Cloud Infrastructure (OCI), you are required to transfer a 7 TB file to OCI Object Storage. You have decided to upload functionality of Object Storage.

Which two statements are true?

- A. Active multipart upload can be checked by listing all parts that have been uploaded, however It is not possible to list information for individual object part in an active multipart upload
- B. It is possible to split this file into multiple parts using the APIs provided by Object Storage.
- C. It is possible to split this file into multiple parts using rclone tool provided by Object Storage.
- D. After initiating a multipart upload by making a CreateMultiPartUpload REST API Call, the upload remains active until you explicitly commit it or abort.
- E. Contiguous numbers need to be assigned for each part so that Object Storage constructs the object by ordering, part numbers in ascending order

Correct Answer: AD

You can check on an active multipart upload by listing all parts that have been uploaded. (You cannot list information for an individual object part in an active multipart upload.) After you finish creating object parts, initiate a multipart upload by making a CreateMultipartUpload REST API call. Provide the object name and any object metadata. Object Storage responds with a unique upload ID that you must include in any requests related to this multipart upload. Object Storage also marks the upload as active. The upload remains active until you explicitly commit it or abort it.

[1Z0-997 PDF Dumps](#)

[1Z0-997 VCE Dumps](#)

[1Z0-997 Practice Test](#)