



# 1Z0-997-22<sup>Q&As</sup>

Oracle Cloud Infrastructure 2022 Architect Professional

## Pass Oracle 1Z0-997-22 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/1z0-997-22.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 notice that a majority of your Oracle Cloud Infrastructure (OCI) resources like compute instances, block volumes, and load balancers are not tagged. You have received a mandate from your CIO to add a predefined set of tags to identify owners for respective OCI resources. E.g. if Chris and Larry each create compute instances in a compartment, the instances that Chris creates include tags that contain his name as the value, while the instances that Larry creates have his name.

Which option is the simplest way to implement this new tagging requirement?

- A. Create a default tag for each compartment, which ensure that appropriate tags are applied at the time of resource creation.
- B. Create an OCI Identity and Access Management policy requiring users to tag resources with their user name.
- C. Create an OCI Identity and Access Management policy to automatically tag a resource with the user name.
- D. Create tag variables to automatically tag a resource with the user name.

Correct Answer: D

---

**QUESTION 2**

Which of the two options are true for an autonomous database in dedicated infrastructure deployment? (Choose two.)

- A. You can modify maintenance schedule of the AVM after provisioning, to match your organization maintenance schedules.
- B. The new resource model consists of autonomous exadata infrastructure, autonomous container database and autonomous database.
- C. Unlike autonomous database in shared infrastructure, you can customize the maintenance schedule of the autonomous databases in dedicated infrastructure in OCI public cloud.
- D. The new resource model consists of exadata infrastructure, autonomous Exadata VM cluster, autonomous container database.
- E. Network selection, License model and certificate management are resources configured at AVM level.

Correct Answer: DE

---

**QUESTION 3**

You work for a large bank where your main application is a payment processing gateway API. You deployed the application on Oracle Container Engine for Kubernetes (OKE) and used API Gateway with several policies to control the access of the API endpoint.

However, your customers are complaining about the unavailability of the API endpoint. Upon checking, you noticed that the Gateway URL is throwing Service Unavailable error. You need to check the backend latency and backend responses when this error started last night.



What should you do to get this data? (Choose the best answer.)

- A. Check with the application owner and search the log file for the container to get the metrics from the log file.
- B. Go to Governance Menu and click on Audit to see the Audit log for the API Gateway. Filter it using Start and End date with a 503 response status.
- C. Go to Developer Services and click on API Gateway. Go to the detail page of the gateway and select Metrics. Change the Start and End time to filter the metrics.
- D. Go to Monitoring and click on Service Metrics. Choose the Metric Namespace as oci\_apigateway. Change the Start and End time accordingly. Add a Dimension and select httpStatusCode: 503. Check the backend latency and backend responses metric.

Correct Answer: D

Explanation: <https://medium.com/oracledevs/using-oci-monitoring-healthchecks-to-schedule-execution-of-serverless-functions-on-oracle-cloud-ef233f887a5>

---

#### QUESTION 4

You developed a microservices based application that runs on Oracle Cloud Infrastructure (OCI) Container Engine for Kubernetes (OKE). It has multiple endpoints that needs to be exposed to the public internet.

What Is the most cost-effective way to expose multiple application endpoints without adding complexity to the application?

- A. Use clusterIP service type in Kubernetes for each of your service endpoint and use a load balancer to expose the endpoints.
- B. Use separate load balancer instance for each service but use the 100 Mbps load balancer option.
- C. Deploy an Ingress controller and use it to expose each endpoint with its own routing endpoint.
- D. Use NodePort service type in Kubernetes for each of your service endpoint and use node's public IP address to access the applications.

Correct Answer: C

---

#### QUESTION 5

You are working as a solutions architect for an online retail store in Frankfurt which uses multiple compute instance VMs spread among three availability domains in the eu-frankfurt-1 region.

You noticed the website is having very high traffic, so you enabled autoscaling to ensure none of your application but, you observed that one of the availability domains is not receiving any traffic.

What could be wrong in this situation?

- A. Autoscaling only works with single availability domains.
- B. You have to manually add all three availability domains to your load balancer configuration.



- C. Autoscaling can be enabled for multiple availability domains only in uk-london t region.
- D. Autoscaling is using an Instance Pool configured to create instances in two availability Domains.
- E. You forgot to attach a load balancer to your instance pool configuration.

Correct Answer: D

Autoscaling lets you automatically adjust the number of Compute instances in an instance pool based on performance metrics such as CPU utilization. This helps you provide consistent performance for your end users during periods of high demand, and helps you reduce your costs during periods of low demand. you can associate a load balancer with an instance pool. If you do this, when you add an instance to the instance pool, the instance is automatically added to the load balancer's backend set . After the instance reaches a healthy state (the instance is listening on the configured port number), incoming traffic is automatically routed to the new instance. Instance pools let you provision and create multiple Compute instances based off the same configuration, within the same region. By default, the instances in a pool are distributed across all fault Domains in a best-effort manner based on capacity. If capacity isn't available in one fault domain, the instances are placed in other fault domains to allow the instance pool to launch successfully. In a high availability scenario, you can require that the instances in a pool are evenly distributed across each of the fault domains that you specify. When sufficient capacity isn't available in one of the fault domains, the instance pool will not launch or scale successfully, and a work request for the instance pool will return an "out of capacity" error. To fix the capacity error, either wait for capacity to become available, or use the UpdateInstancePool operation to update the placement configuration (the availability domain and fault domain) for the instance pool. during create the instance pool you can select the location where you want to place the instances" In the Availability Domain list, select the availability domain to launch the instances in. If you want the instances in the pool to be placed evenly in one or more fault domains, select the Distribute instances evenly across selected fault domains check box. Then, select the fault domains to place the instances in.

[Latest 1Z0-997-22 Dumps](#)

[1Z0-997-22 PDF Dumps](#)

[1Z0-997-22 Exam Questions](#)