

# 1Z0-997-21<sup>Q&As</sup>

Oracle Cloud Infrastructure 2021 Architect Professional

## Pass Oracle 1Z0-997-21 Exam with 100% Guarantee

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

<https://www.pass2lead.com/1z0-997-21.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

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 2

You have deployed a multi-tier application with multiple compute instances in Oracle Cloud Infrastructure.

You want to back up these volumes and have decided to use '\Volume Groups\' feature. The Block volume and Compute instances exist in different compartments within your tenancy. Periodically, a few child compartments are moved under different parent compartments, and you notice that sometimes volume group backup fails.

What could be the cause?

- A. The Identity and Access Management policy allowing backup failed to move when the compartment was moved.
- B. You are exceeding your volume group backup quota configured.
- C. You have the same block volume attached to multiple compute instances; if these compute instances are in different compartments then all concerned compartments must be moved at the same time.
- D. A compute instance with multiple block volumes attached cannot move when a compartment is moved.

Correct Answer: A

---

### QUESTION 3

Your Oracle database is deployed on-premises and has produced 100 TB database backup locally. You have a disaster recovery plan that requires you to create redundant database backups in Oracle Cloud Infrastructure (OCI).

Once the initial backup is completed, the backup must be available for retrieval in less than 30 minutes to support the Recovery Time Objective (RTO) of your solution. Which is the most cost effective option to meet these requirements?

- A. Setup an IPsec VPNConnect between on-premises data center and OCI. Then to use OCI CLI command to upload database backups to OCI Object Storage Archive tier as the final destination.
- B. Use OCI Storage Gateway to transfer the backup files to OCI Object Storage Archive tier as the final destination.
- C. Setup a FastConnect connection between on-premises data center and OCI. Then to use OCI CLI command to upload database backups to OCI Object Storage Standard tier as the final destination.
- D. Use OCI Storage Gateway to transfer the backup files to OCI Object Storage Standard tier as the final destination.

Correct Answer: D

---

### QUESTION 4

A global retailer has decided to re-design its e-commerce platform to have a micro-services architecture.

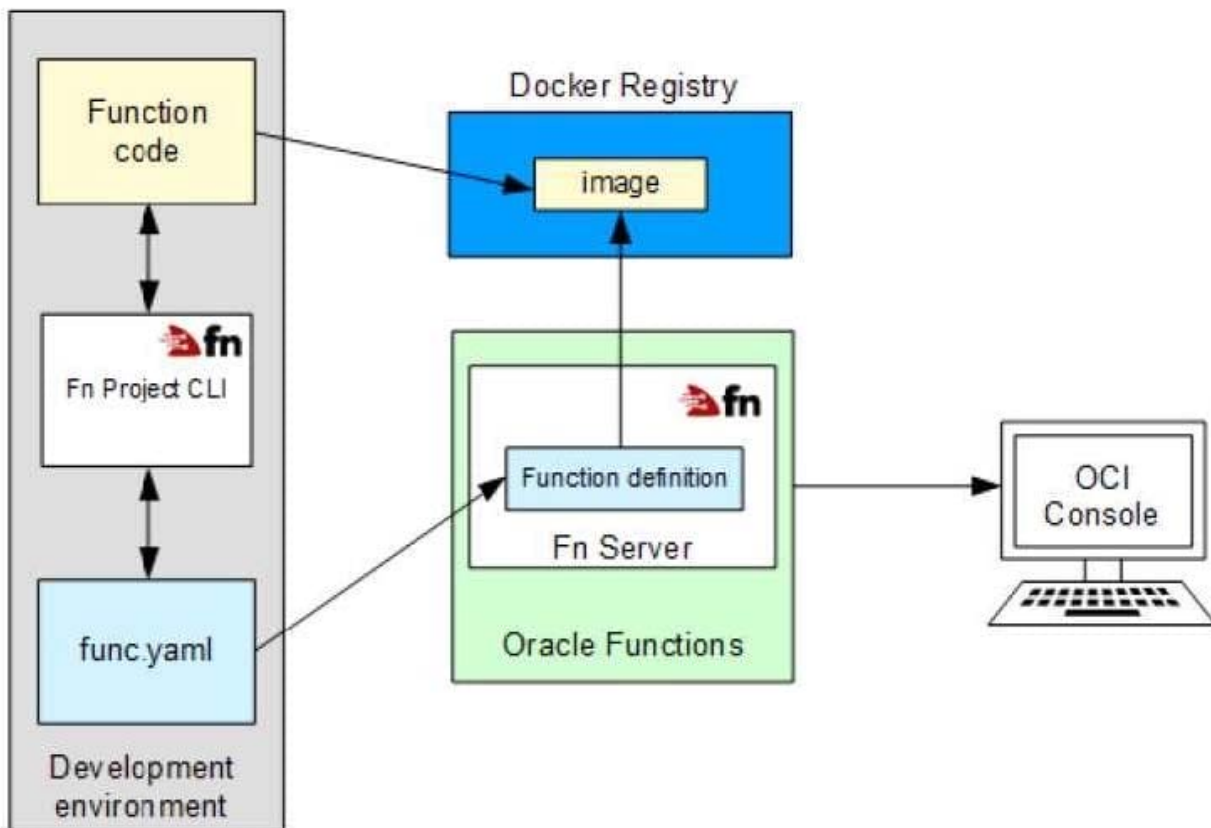
They would like to decouple application architecture into smaller, independent services using Oracle Cloud Infrastructure (OCI). They have decided to use both containers and servers technologies to run these application instances.

Which option should you recommend to build this new platform?

- A. Install a kubernetes cluster on OCI and use OCI event service.
- B. Use Oracle Container Engine for kubernetes, OCI Registry and OCI Functions.
- C. Use OCI Resource Manager to automate compute Instances provisioning and use OCI Streaming service.
- D. Use OCI functions, OCI object storage and OCI event service.

Correct Answer: B

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.



Oracle Cloud Infrastructure Container Engine for Kubernetes is a fully-managed, scalable, and highly available service that you can use to deploy your containerized applications to the cloud. Use Container Engine for Kubernetes (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud-native applications. You specify the compute resources that your applications require, and Container Engine for Kubernetes provisions them on Oracle Cloud Infrastructure in an existing OCI tenancy.

#### QUESTION 5

You are advising the database administrator responsible for managing non-production environment for Oracle Autonomous Database running on Oracle Cloud Infrastructure. You need to help the database administrator ensure that the non-production environments have a copy of the current data from the production environment in a manner that is most time-efficient. Which method should you recommend? (Choose the best answer.)

- A. Take a full database backup of the production Autonomous database and create the non- production database from it.
- B. Create a metadata clone of the production Autonomous Database and create the non-production database from it.
- C. Create a full clone of the production Autonomous Database and create the non-production database from it.
- D. Take a Data Pump export of the production Autonomous database and import into the non- production database.

Correct Answer: C

<https://www.oracle.com/database/technologies/datawarehouse-bigdata/adb-faqs.html>

[Latest 1Z0-997-21 Dumps](#)

[1Z0-997-21 PDF Dumps](#)

[1Z0-997-21 Braindumps](#)