

1Z0-997^{Q&As}

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.leads4pass.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

An upcoming e-commerce company has deployed their online shopping application on OCI. The application was deployed on compute instances with autoscaling configuration for application servers fronted by a load balancer and OCI Autonomous Transaction Processing (ATP) in the backend. In order to promote their e-commerce platform 50% discount was announced on all the products for a limited period. During the day 1 of promotional period it was observed that the application is running slow and company's hotline is flooded with complaints. What could be two possible reasons for this situation?

- A. The health check on some of the backend servers has failed and the load balancer has taken those servers temporarily out of rotation
- B. As part of autoscaling, the load balancer shape has dynamically changed to a larger shape to handle more incoming traffic and the system was slow for a short time during this change
- C. The health check on some of the backend servers has failed and the load balancer was rebooting these servers.
- D. The autoscaling has already scaled to the maximum number of instances specified in the configuration and there is no room of scaling

Correct Answer: AD


QUESTION 2

All three Data Guard Configuration are fully supported on Oracle Cloud infrastructure (OCI). You want to deploy a maximum availability architecture (MAA) for database workload. Which option should you consider while designing your Data Guard configuration to ensure best RTO and PRO without causing any data loss?

- A. Configure "Maximum Protection" mode which provides zero data loss If the primary database fails.
- B. Configure "Maximum Performance" mode In SYNC mode between two availability domains (same region) which provides, the highest level of data protection that is possible without affecting the performance of the primary database.
- C. Configure "Maximum Scalability" mode which provides the highest level of scalability without compromising the availability of the primary database.
- D. Configure "Maximum Availability" mode in SYNC mode between two availability domains (same

Correct Answer: D

<https://docs.cloud.oracle.com/en-us/iaas/Content/Resources/Assets/whitepapers/best-practices-for-dr-onoci.pdf> All three Data Guard configurations are fully supported on Oracle Cloud Infrastructure. However, because of a high risk of production outage, we don't recommend using the maximum protection mode for your Data Guard configuration. We recommend using the maximum availability mode in SYNC mode between two availability domains (same region), and using the maximum availability mode in ASYNC mode between two regions. This architecture provides you the best RTO and RPO without causing any data loss. We recommend building this architecture in daisy-chain mode: the primary database ships redo logs to the first standby database in another availability domain in SYNC mode, and then the first standby database ships the redo logs to another region in ASYNC mode. This method ensures that your primary database is not doing the double work of shipping redo logs, which can cause performance impact on a production workload.



This configuration offers the following benefits: No data loss within a region. No overhead on the production database to maintain standbys in another region. Option to configure lagging on the DR site if needed for business reasons. Option to configure multiple standbys in different regions without any additional overhead on the production database. A typical use case is a CDN application Bottom of Form

QUESTION 3

A large London based eCommerce company is running Oracle DB System Virtual RAC database on Oracle Cloud Infrastructure (OCI) for their eCommerce application activity. They are launching a new product soon, which is expected to sell in large quantities all over the world. The application architecture should have minimal cost, no data loss, no performance impacts during the database backup windows and should have minimal downtime.

- A. Launch a new VM RAC database in another availability domain, launch a compute instance, deploy Oracle GoldenGate on it and then configure it to replicate the data from the eCommerce Database over to the new RAC database using GoldenGate. Take backups from the new VM RAC database.
- B. Turn off automated backups from the eCommerce database, implement Oracle Data Guard with the Standby database deployed on another availability domain, take backups from the standby database.
- C. Launch a new VM RAC database in another availability domain, launch a compute instance, deploy Oracle GoldenGate on it and then configure bi-directional replication from the eCommerce Database over to the new VM RAC database using GoldenGate. Take backups from the new VM RAC database.

D. Turn off automatic backups from the eCommerce database, implement Oracle Active Data Guard with the standby database deployed on another availability domain, and take backups from the standby database.

Correct Answer: C

Active Data Guard or GoldenGate are used for disaster recovery when fast recovery times or additional levels of data protection are required. And offload queries and backup to standby system. Oracle GoldenGate to support a disaster recovery site is to have a working bi-directional data flow, from the primary system to the live-standby system and vice versa.

DataGuard and Automatic Backup You can enable the Automatic Backup feature on a database with the standby role in a Data Guard association. However, automatic backups for that database will not be created until it assumes the primary role.

QUESTION 4

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 solve one 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 us-east-1 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.

QUESTION 5

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.

[1Z0-997 PDF Dumps](#)

[1Z0-997 Practice Test](#)

[1Z0-997 Exam Questions](#)