



1Z0-1110-22^{Q&As}

Oracle Cloud Infrastructure Data Science 2022 Professional

Pass Oracle 1Z0-1110-22 Exam with 100% Guarantee

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

<https://www.geekcert.com/1z0-1110-22.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 are asked to prepare data for a custom-built model that requires transcribing Spanish video recordings into a readable text format with profane words identified. Which Oracle Cloud service would you use?

- A. OCI Translation
- B. OCI Language
- C. OCI Anomaly Detection
- D. OCI Speech

Correct Answer: D

QUESTION 2

You are a data scientist designing an air traffic control model, and you choose to leverage Oracle AutoML. You understand that the Oracle AutoML pipeline consists of multiple stages and automatically operates in a certain sequence. What is the correct sequence for the Oracle AutoML pipeline?

- A. Adaptive sampling, Feature selection, Algorithm selection, Hyperparameter tuning.
- B. Adaptive sampling, Algorithm selection, Feature selection, Hyperparameter tuning.
- C. Algorithm selection, Feature selection, Adaptive sampling, Hyperparameter tuning.
- D. Algorithm selection, Adaptive sampling, Feature selection, Hyperparameter tuning.

Correct Answer: D

QUESTION 3

You are building a model and need input that represents data as morning, afternoon, or evening. However, the data contains a time stamp. What part of the Data Science life cycle would you be in when creating the new variable?

- A. Model type selection
- B. Model validation
- C. Data access
- D. Feature engineering

Correct Answer: B

QUESTION 4

You realize that your model deployment is about to reach its utilization limit. What would you do to avoid the issue



before requests start to fail?

- A. Update the deployment to use fewer instances.
- B. Delete the deployment.
- C. Reduce the load balancer bandwidth limit so that fewer requests come in.
- D. Update the deployment to use a larger virtual machine (more CPUs/memory).
- E. Update the deployment to add more instances.

Correct Answer: E

QUESTION 5

You want to ensure that all stdout and stderr from your code are automatically collected and logged, without implementing additional logging in your code. How would you achieve this with Data Science Jobs?

- A. Data Science Jobs does not support automatic log collection and storing.
- B. On job creation, enable logging and select a log group. Then, select either log or the option to enable automatic log creation.
- C. You can implement custom logging in your code by using the Data Science Jobs logging.
- D. Make sure that your code is using the standard logging library and then store all the logs to Check Storage at the end of the job.

Correct Answer: C

[1Z0-1110-22 PDF Dumps](#)

[1Z0-1110-22 Practice Test](#)

[1Z0-1110-22 Brindumps](#)