



**Free Questions for 1Z0-1110-22 by ebraindumps**

**Shared by Key on 12-12-2023**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

## Question 1

---

**Question Type:** MultipleChoice

---

You are given the task of writing a program that sorts document images by language. Which Oracle AI service would you use?

### Options:

---

- A- Oracle Digital Assistant
- B- OCI Speech
- C- OCI Vision
- D- OCI Language

### Answer:

---

C

## Question 2

---

**Question Type:** MultipleChoice

---

You have a data set with fewer than 1000 observations, and you are using Oracle AutoML to build a classifier. While visualizing the results of each stage of the Oracle AutoML pipeline, you notice that no visualization has been generated for one of the stages. Which stage is not visualized?

**Options:**

---

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

**Answer:**

---

C

## Question 3

---

**Question Type: MultipleChoice**

---

Select two reasons why it is important to rotate encryption keys when using Oracle Cloud In-frastructure (OCI) Vault to store credentials or other secrets.?

### Options:

---

- A- Key rotation allows you to encrypt no more than five keys at a time.
- B- Key rotation reduces risk if a key is ever compromised.
- C- Key rotation improves encryption efficiency.
- D- Periodically rotating keys make it easier to reuse key.
- E- Periodically rotating keys limits the amount of data encrypted by one key version.

### Answer:

---

B, D

## Question 4

---

**Question Type:** MultipleChoice

---

You have received machine learning model training code, without clear information about the optimal shape to run the training on. How would you proceed to identify the optimal compute shape for your model training that provides a balanced cost and processing time?

## Options:

---

- A-** Start with the strangest compute shape Jobs support and monitor the Job Run metrics and time required to complete the model training. Tune the model so that it utilizes as much compute resources as possible, even at an increased cost.
- B-** Start with a random compute shape and monitor the utilization metrics and time required to finish the model training Perform model training optimizations and performance tests in advance to identify the right compute shape before running the model training as a job.
- C-** Start with a smaller shape and monitor the Job Run metrics and time required to complete the model training: If the compute shape is not fully utilized, tune the model parameters, and rerun the job. Repeat the process until the shape resources are fully utilized.
- D-** Start with a smaller shape and monitor the utilization metrics and time required to complete the model training. If the compute shape is fully utilized, change to compute that has more resources and re-run the job. Repeat the process until the processing time does not improve.

## Answer:

---

D

## Question 5

---

### Question Type: MultipleChoice

---

Using Oracle AutoML, you are tuning hyperparameters on a supported model class and have specified a time budget. AutoML terminates computation once the time budget is exhausted. What would you expect AutoML to return in case the time budget is

exhausted before hyperparameter tuning is completed?

**Options:**

---

- A- A hyperparameter configuration with a minimum learning rate is returned.
- B- The current best-known hyperparameter configuration is returned.
- C- A random hyperparameter configuration is returned.
- D- The last generated hyperparameter configuration is returned.

**Answer:**

---

B

## Question 6

---

**Question Type: MultipleChoice**

---

You are a data scientist working for a utilities company. You have developed an algorithm that detects anomalies from a utility reader in the grid. The size of the model artifact is about 2 GB, and you are trying to store it in the model catalog. Which THREE interfaces would you use to save the model artifact into the model catalog?

### Options:

---

- A- Console
- B- Accelerated Data Science (ADS) Software Development Kit (SDK)
- C- Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI)
- D- OCI Python SDK
- E- Git CLI
- F- ODSC CLI

### Answer:

---

A, B, D

## Question 7

---

### Question Type: MultipleChoice

---

Which Oracle Accelerated Data Science (ADS) classes can be used for easy access to data sets from reference libraries and index websites, such as scikit-learn?

**Options:**

---

- A- ADSTurner
- B- DatasetFactory
- C- SecretKeeper
- D- Dataset Browser

**Answer:**

---

D

## Question 8

---

**Question Type: MultipleChoice**

---

As a data scientist, you are tasked with creating a model training job that is expected to take different hyperparameter values on every run. What is the most efficient way to set those pa-rameters with Oracle Data Science Jobs?

**Options:**

---

- A- Create a new job every time you need to run your code and pass the parameters as en-vironment variables.



- B-** Create your code to expect different parameters as command line arguments, and create it new job every time you run the code.
- C-** Create your code to expect different parameters either as environment variables or as command line arguments, which are set on every job run with different values.
- D-** Create a new no by setting the required parameters in your code, and create a new job for mery code change.

**Answer:**

---

C

## Question 9

---

**Question Type:** MultipleChoice

---

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?

**Options:**

---

- A-** Data Science Jots does not support automatic fog collection and storing.
- B-** On job creation, enable logging and select a log group. Then, select either log or the op-tion 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.

**Answer:**

---

C

## Question 10

---

**Question Type: MultipleChoice**

---

You have created a model, and you want to use the Accelerated Data Science (ADS) SDK to deploy this model. Where can you save the artifacts to deploy this model with ADS?

**Options:**

---

**A-** Model Depository

**B-** Model Catalog

**C-** OCI Vault

**D-** Data Science Artifactory

**Answer:**

---

B

## Question 11

---

**Question Type:** MultipleChoice

---

You have an embarrassingly parallel or distributed batch job on a large amount of data running using Data Science Jobs. What would be the best approach to run the workload?

**Options:**

---

- A-** Create the job in Data Science Jobs and then start the number of simultaneous job runs required for your workload.
- B-** Create the job in Data Science Jobs and start a job run. When it is done, start a new job run until you achieve the number of runs required.
- C-** Reconfigure the job run because Data science jobs does not support embarrassingly parallel.
- D-** Create a new job for every job run that you have to run in parallel, because the Date Science Jobs service can have only one job run per job.

**Answer:**

---

A

**To Get Premium Files for 1Z0-1110-22 Visit**

**<https://www.p2pexams.com/products/1z0-1110-22>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/oracle/pdf/1z0-1110-22>**

