


Oracle 1z0-1110-22 Test Engine Dumps Training With 58 Questions [Q19-Q37]



Oracle 1z0-1110-22 Test Engine Dumps Training With 58 Questions [Q19-Q37]

Oracle 1z0-1110-22 Test Engine Dumps Training With 58 Questions

1z0-1110-22 Questions Pass on Your First Attempt Dumps for Oracle Cloud Infrastructure Certified

The Oracle 1z0-1110-22 exam is a certification exam designed for professionals who want to demonstrate their expertise in Oracle Cloud Infrastructure data science. This exam is intended to validate the knowledge and skills required to design, implement, and manage data science solutions using Oracle Cloud Infrastructure Data Science. Passing this exam will prove that you have the necessary skills to enable machine learning models and analyze big data.

The Oracle 1z0-1110-22 exam covers a wide range of topics related to data science, including machine learning, data visualization, data preparation, and data analysis. Candidates are expected to have a solid understanding of statistical analysis, programming languages such as Python and R, and data modeling techniques. They should also be familiar with cloud-based data science tools and technologies, including the Oracle Cloud Infrastructure, Oracle Autonomous Database, Oracle Analytics Cloud, and Oracle Machine Learning.

QUESTION 19

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?

- * Console
- * Accelerated Data Science (ADS) Software Development Kit (SDK)
- * Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI)
- * OCI Python SDK
- * Git CLI
- * ODSC CLI

QUESTION 20

What preparation steps are required to access an Oracle AI service SDK from a Data Science notebook session?

- * Call the Accented Data Science (ADS) command to enable AI integration
- * Create and upload the API signing key and config file
- * Import the REST API
- * Create and upload execute.py and runtime.yaml

QUESTION 21

You are creating an Oracle Cloud Infrastructure (OCI) Data Science job that will run on a recurring basis in a production environment. This job will pick up sensitive data from an Object Storage bucket, train a model, and save it to the model catalog. How would you design the authentication mechanism for the job?

- * Package your personal OC file and keys in the job artifact.
- * Use the resource principal of the job run as the signer in the job code, ensuring there is a dynamic group for this job run with appropriate access to Object Storage and the model catalog.
- * Store your personal OCI config file and kays in the Vault, and access the Vault through the job nun resource principal
- * Create a pre-authenticated request (PAA) for the Object Storage bucket, and use that in the job code.

QUESTION 22

You are a data scientist designing an air traffic control model, and you choose to leverage Or-acle 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?

- * Adaptive sampling, Feature selection, Algorithm selection, Hyperparameter tuning.
- * Adaptive sampling, Algorithm selection, Feature selection, Hyperparameter tuning.
- * Algorithm selection, Feature selection, Adaptive sampling, Hyperparameter tuning.
- * Algorithm selection, Adaptive sampling. Feature selection, Hyperparameter tuning.

QUESTION 23

You have trained three different models on your data set using Oracle AutoML. You want to visualize the behavior of each of the models, including the baseline model, on the test set. Which class should be used from the Accelerated Data Science (ADS) SDK to visually compare the models?

- * ADS Explainer
- * ADS Evaluator
- * ADSTuner
- * EvaluationMetrics

QUESTION 24

When preparing your model artifact to save it to the Oracle Cloud Infrastructure (OCI) Data Science model catalog, you create a score.py file. What is the purpose of the score.py file?

- * Define the compute scaling strategy.
- * Configure the deployment infrastructure.
- * Define the inference server dependencies.
- * Execute the inference logic code

QUESTION 25

You are a data scientist working inside a notebook session and you attempt to pip install a package from a public repository that is not included in your conda environment. After running this command, you get a network timeout error. What might be missing from your networking configuration?

- * Service Gateway with private subnet access.
- * NAT Gateway with public internet access.
- * FastConnect to an on-premises network.
- * Primary Virtual Network Interface Card (VNIC).

QUESTION 26

You are a data scientist working for a manufacturing company, you have developed a forecasting model to predict the sales demand in the upcoming months. You created a model artifact that contained custom logic requiring third party libraries. When you deployed the model, it failed to run because you did not include all the third-party dependencies in the model artifact.?

- * Score.py
- * Requirement.txt
- * Model_artifact_validate.py
- * Runtime.yaml

QUESTION 27

Which TWO statements are true about published conda environments?

- * The odsc conda init command is used to configure the location of published conda environments.
- * They can be used in Data Science Jobs and model deployments.
- * Your notebook session acts as the source to share published conda environment with team members.
- * You can only create published conda environment by modifying a Data Science conda
- * They are curated by Oracle Cloud Infrastructure (OCI) Data Science.

QUESTION 28

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 parameters with Oracle Data Science Jobs?

- * Create a new job every time you need to run your code and pass the parameters as environment variables.
- * Create your code to expect different parameters as command line arguments, and create it new job every time you run the code.
- * 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.
- * Create a new job by setting the required parameters in your code, and create a new job for every code change.

QUESTION 29

Six months ago, you created and deployed a model that predicts customer churn for a call center. Initially, it was yielding quality predictions. However, over the last two months, users have been questioning the credibility of the predictions. Which TWO methods customer churn would you employ to verify the accuracy of the model?

- * Redeploy the model
- * Retrain the model
- * Operational monitoring
- * Validate the model using recent data
- * Drift monitoring

QUESTION 30

The Oracle AutoML pipeline automates hyperparameter tuning by training the model with different parameters in parallel. You have created an instance of Oracle AutoML as ora-cle_automl and now you want an output with all the different trials performed by Oracle Au-toML. Which of the following command gives you the results of all the trials?

- * Oracle.automl.visualize_algorithm_selection_trails()
- * Oracle.automl.visualize_adaptive_sampling_trails()
- * Oracle.automl.print_trails()
- * Oracle.automl.visualize_tuning_trails()

QUESTION 31

You have created a conda environment in your notebook session. This is the first time you are working with published conda environments. You have also created an Object Storage bucket with permission to manage the bucket Which TWO commands are required to publish the conda environment?

- * odsc conda publish –slug <SLUG>
- * odsc conda create –file manifest.yaml
- * odsc conda init -b <your-bucket-name> -a <api_key or resource_principal>
- * odsc conda list –override

QUESTION 32

You want to make your model more parsimonious to reduce the cost of collecting and processing data. You plan to do this by removing features that are highly correlated. You would like to create a heat map that displays the correlation so that you can identify candidate features to remove. Which Accelerated Data Science (ADS) SDK method would be appropriate to display the correlation between Continuous and Categorical features?

- * Corr{ }
- * Correlation_ratio_plot{ }
- * Pearson_plot{ }
- * Cramersv_plot{ }

QUESTION 33

You are using a third-party Continuous Integration/Continuous Delivery (CI/CD) tool to create a pipeline for preparing and training models. How would you integrate a third-party tool outside Oracle Cloud Infrastructure (OCI) to access Data Science Jobs?

- * Third-party software can access Data Science Jobs by using any of the OCI Software Development Kits (SDKs).
- * Data Science Jobs does not accept code from third-party tools, therefore you need to run the pipeline externally.
- * Third-party tools use authentication keys to create and run.
- * Data Science Jobs Data Science Jobs is not accessible from outside OCI.

QUESTION 34

You have just received a new data set from a colleague. You want to quickly find out summary information about the data set, such as the types of features, total number of observations, and data distributions, Which Accelerated Data Science (ADS) SDK method from the AD&Dataset class would you use?

- * `Show_in_notebook{ }`
- * `To_xgb{ }`
- * `Compute{ }`
- * `Show_corr{ }`

QUESTION 35

While reviewing your data, you discover that your data set has a class imbalance. You are aware that the Accelerated Data Science (ADS) SDK provides multiple built-in automatic transformation tools for data set transformation. Which would be the right tool to correct any imbalance between the classes?

- * `sample()`
- * `suggeste_recoomendations()`
- * `auto_transform()`
- * `visualize_transforms()`

QUESTION 36

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?

- * Data Science Jots does not support automatic fog collection and storing.
- * On job creation, enable logging and select a log group. Then, select either log or the op-tion to enable automatic log creation.
- * You can implement custom logging in your code by using the Data Science Jobs logging.
- * 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.

QUESTION 37

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.?

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

1z0-1110-22 Practice Test Pdf Exam Material: <https://www.actualtests4sure.com/1z0-1110-22-test-questions.html>]