


[Q20-Q41] 1z0-1110-22 100% Guarantee Download 1z0-1110-22 Exam PDF Q&A [Dec 19, 2023]



1z0-1110-22 100% Guarantee Download 1z0-1110-22 Exam PDF Q&A [Dec 19, 2023]
Get 1z0-1110-22 Actual Free Exam Q&As to Prepare for Your Oracle Certification

Oracle 1z0-1110-22 (Oracle Cloud Infrastructure Data Science 2022 Professional) Certification Exam is a valuable credential for data scientists who work on the Oracle Cloud Infrastructure platform. It tests the knowledge and skills required to use Oracle Cloud Infrastructure tools and services effectively for data science tasks, and demonstrates an individual's expertise in this area. Obtaining this certification can help data scientists to advance their careers and gain recognition for their expertise in the field.

If you're looking to become a professional data scientist in the Oracle Cloud Infrastructure environment, the Oracle Cloud Infrastructure Data Science 2022 Professional certification is the certification for you. Oracle Cloud Infrastructure Data Science 2022 Professional certification validates your knowledge and expertise in machine learning and data science, making you one of the elite professionals in the field of data science. It's a comprehensive exam that covers all aspects of data science in the Oracle Cloud Infrastructure environment, ensuring you have an in-depth understanding of the subject matter.

NO.20 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?

- * Create a new job every time you need to run your code and pass the parameters as en-vironment 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 no by setting the required parameters in your code, and create a new job for mery code change.

NO.21 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

NO.22 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()

NO.23 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

NO.24 The Accelerated Data Science (ADS) model evaluation classes support different types of ma-chine learning modeling techniques Which THREE types of modeling techniques are supported by ADS Evaluators?

- * Recurrent Neural Network
- * Regression Analysis
- * Principal Component Analysis
- * Binary Classification
- * Multiclass Classification
- * K-Means Clustering

NO.25 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.

NO.26 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 condo 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).

NO.27 Where do calls to stdout and stderr from score.py go in a model deployment?

- * The predict log in the Oracle Cloud Infrastructure (OCI) Logging service as defined in the deployment.
- * The OCI Cloud Shell, which can be accessed from the console.
- * The file that was defined for them on the Virtual machine (VM).
- * The OCI console.

NO.28 After you have created and opened a notebook session, you want to use the Accelerated Data Science (ADS) SDK to access your data and get started with exploratory data analysis. From which TWO places can you access the ADS SDK?

- * Oracle Autonomous Data Warehouse
- * Oracle Machine Learning
- * Conda environments in Oracle Cloud infrastructure (OCI) Data Science.
- * Python Package Index (PyPi)
- * Oracle Big Data Service

NO.29 You have built a machine model to predict whether a bank customer is going to default on a loan. You want to use Local Interpretable Model-Agnostic Explanations (LIME) to understand a specific prediction. What is the key idea behind LIME?

- * Model-agnostic techniques are more interpretable than techniques that are dependent on the types of models.
- * Local explanation techniques are model agnostic, while global explanation techniques are not.
- * Global behavior of a machine learning model may be complex, while the local behavior may be approximated with a simpler surrogate model.
- * Global and local behaviors of machine learning models are similar.

NO.30 You are attempting to save a model from a notebook session to the model catalog by using the Accelerated Data Science (ADS) SDK, with resource principal as the authentication signer, and you get a 404 authentication error. Which TWO should you look for to ensure permissions are set up correctly?

- * The networking configuration allows access to the Oracle Cloud Infrastructure (OCI) services through a Service Gateway.
- * The model artifact is saved to the block volume of the notebook session.
- * The dynamic group's matching rule exists for notebook sessions in this compartment.
- * The policy for the dynamic group grants permissions for the model catalog in this compartment.
- * The policy for your user group grants permissions for the model catalog in this compartment.

NO.31 For your next data science project, you need access to public geospatial images. Which Oracle Cloud service provides free access to those images?

- * Oracle Big Data Service
- * Oracle Analytics Cloud
- * Oracle Cloud Infrastructure (OCI) Data Science
- * Oracle Open Data

NO.32 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()

NO.33 As you are working in your notebook session, you find that your notebook session does not have enough compute CPU and memory for your workload. How would you scale up your notebook session without losing your work?

- * Ensure your files and environments are written to the block volume storage under the /home/datascience directory, deactivate the notebook session, and activate the notebook larger compute shape selected.
- * Down your files and data to your local machine, delete your notebook session, provision tebook session on a larger compute shape, and upload your files from your local the new notebook session.
- * Deactivate your notebook session, provision a new notebook session on larger compute shape, and re-create all your file changes.
- * Create a temporary bucket in Object Storage, write all your files and data to Object Storage, delete tur ctebook session, provision a new notebook session on a larger com-pute shape, and copy your flies and data from your temporary bucket onto your new notebook session.

NO.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{ }

NO.35 You are using Oracle Cloud Infrastructure Anomaly Detection to train a model to detect anomalies in pump sensor dat a. How does the required False Alarm Probability settings affect an anomaly detection model?

- * It changes the sensitivity of the model to detect anomalies.
- * It is used to disable the reporting of false alarm.
- * It Adds a score to each signal indicating the probability that it is false alarm.
- * It determines how many false alarms occur before an error message is generated.

NO.36 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?

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

NO.37 You are a data scientist using Oracle AutoML to produce a model and you are evaluating the score metric for the model. Which of the following TWO prevailing metrics would you use for evaluating multiclass classification model?

- * Recall
- * Mean squared error
- * F1 Score
- * R-Squared
- * Explained variance score

NO.38 Which of the following TWO non-open source JupyterLab extensions has Oracle Cloud In-frastructure (OCI) Data Science developed and added to the notebook session experience?

- * Environment Explorer
- * Table of Contents
- * Command Palette
- * Notebook Examples

* Terminal

NO.39 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?

- * ADSTurner
- * DatasetFactory
- * SecretKeeper
- * Dataset Browser

NO.40 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

NO.41 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 Jobs does not support automatic log collection and storing.
- * On job creation, enable logging and select a log group. Then, select either log or the option 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.

1z0-1110-22 Questions Truly Valid For Your Oracle Exam: <https://www.actualtests4sure.com/1z0-1110-22-test-questions.html>