

[Q29-Q47] 1z0-1110-22 Certification - The Ultimate Guide [Updated 2023]



[Q29-Q47] 1z0-1110-22 Certification - The Ultimate Guide [Updated 2023]

1z0-1110-22 Certification - The Ultimate Guide [Updated 2023]

1z0-1110-22 Practice Exam and Study Guides - Verified By VCEPrep

Oracle 1z0-1110-22 Exam Syllabus Topics:

Topic 1- Create and train models using OCI and Open source Libraries- Discuss Accelerated Data Science (ADS) SDK Capabilities
Topic 2- Configure and manage source code in Code Repositories (Git)- Configure your tenancy for OCI Data Science
Topic 3- Create and Manage Spark Applications using Data Flow and OCI Data Science- Create and manage Conda environments
Topic 4- Obtain Global & Local Model Explanations- Access data from different sources
Topic 5- Monitor & Log using MLOps Practices- Use OCI AI Services for ML Solutions
Topic 6- Implement end-to-end Machine Learning Lifecycle- Design and Set up OCI Data Science Workspace
Topic 7- Create and Use automated ML capability from Oracle AutoML- Configure your tenancy for Data Science
Topic 8- Create and Export a Dataset using OCI Data Labeling- Discuss general MLOps Architecture in OCITopic 9- Explain core OCI Open Data Service concepts- Create & Manage Jobs for custom tasks
Topic 10- Create and manage Projects and Notebook sessions- Discuss OCI Data Science Overview & Concepts
Topic 11- Manage Models using Model Catalog- Deploy & Invoke a Cataloged Model

NO.29 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_trials()
- * Oracle.automl.visualize_tuning_trails()

NO.30 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

NO.31 The feature type TechJob has the following registered validators:

```
Tech-Job.validator.register(name='is_tech_job', handler=is_tech_job_default_handler)
Tech-Job.validator.register(name='is_tech_job', handler= is_tech_job_open_handler,
condition=('job_family',)) TechJob.validator.register(name='is_tech_job', handler=
is_tech_job_closed_handler, condition=('job_family': 'IT'))
```

When you run `is_tech_job(job_family='Engineering')`, what does the feature type validator system do?

- * Execute the `is_tech_job_default_handler` sales handler.
- * Throw an error because the system cannot determine which handler to run.
- * Execute the `is_tech_job_closed_handler` handler.
- * Execute the `is_tech_job_open_handler` handler.

NO.32 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.33 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 manages permissions for the model catalog in this compartment.
- * The policy for your user group grants manage permissions for the model catalog in this compartment.

NO.34 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.

NO.35 You want to evaluate the relationship between feature values and model predictions. You suspect that some of the features are correlated. Which model explanation technique would you recommend?

- * Accumulated Local Effects.
- * Local Interpretable Model-Agnostic Explanations.
- * Feature Dependence Explanations.
- * Feature Permutation Importance Explanations.

NO.36 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.37 Which of the following TWO non-open source JupyterLab extensions has Oracle Cloud Infrastructure (OCI) Data Science developed and added to the notebook session experience?

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

NO.38 As a data scientist, you are working on a global health data set that has data from more than 50 countries. You want to encode three features, such as `‘countries’`, `‘race’`, and `‘body organ’` as categories. Which option would you use to encode the categorical feature?

- * `DataFrameLabelEncode()`
- * `auto_transform()`
- * `OneHotEncoder()`
- * `show_in_notebook()`

NO.39 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.

NO.40 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?

- * 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.

NO.41 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 `ADSDataset` class would you use?

- * `Show_in_notebook{ }`

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

NO.42 You are a data scientist with a set of text and image files that need annotation, and you want to use Oracle Cloud Infrastructure (OCI) Data Labeling. Which of the following THREE an-notation classes are supported by the tool.?

- * Object Detection
- * Named Entity Extraction
- * Classification (single/multi label)
- * Key-Point and Landmark
- * Polygonal Segmentation
- * Semantic Segmentation

NO.43 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.44 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.45 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.46 You are a data scientist leveraging the Oracle Cloud Infrastructure (OCI) Language AI service for various types of text analyses. Which TWO capabilities can you utilize with this tool?

- * Table extraction
- * Punctuation correction
- * Sentence diagramming
- * Topic classification
- * Sentiment analysis

NO.47 You have developed a model training code that regularly checks for new data in Object Storage and retrains the model. Which statement best describes the Oracle Cloud Infrastructure (OCI) services that can be accessed from Data Science Jobs?

- * Data Science Jobs can access OCI resources only via the resource principal.
- * Some OCI services require authorizations not supported by Data Science Jobs.
- * Data Science Jobs cannot access all OCI services.
- * Data Science Jobs can access all OCI services.

Ultimate Guide to the 1z0-1110-22 - Latest Edition Available Now: <https://www.vceprep.com/1z0-1110-22-latest-vce-prep.html>