Oracle 1z0-1110-25 Exam Review - 1z0-1110-25 Valid Test Tips

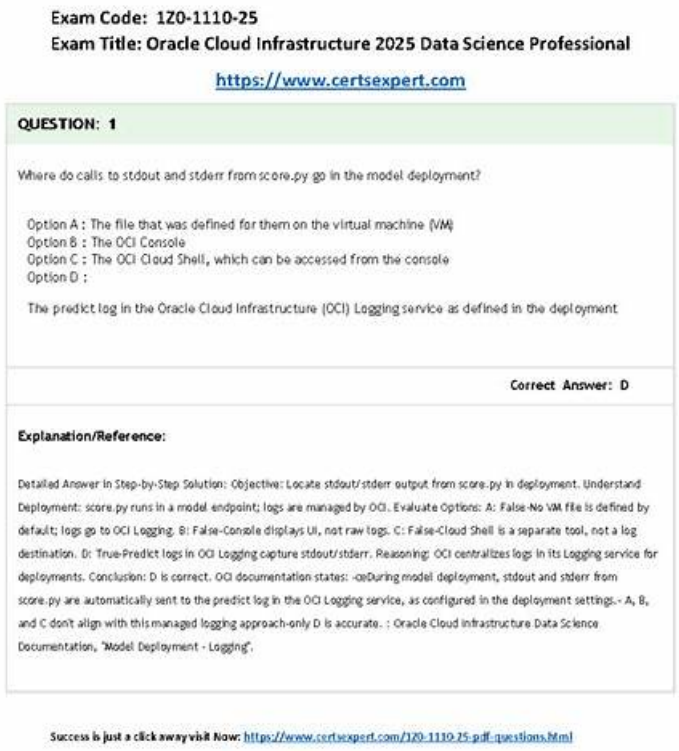

Exam Code: 1Z0-1110-25
Exam Title: Oracle Cloud Infrastructure 2025 Data Science Professional

https://www.certsexpert.com

QUESTION: 1

Where do calls to stdout and stderr from score.py go in the model deployment?

Option A : The file that was defined for them on the virtual machine (VM)
Option B : The OCI Console
Option C : The OCI Cloud Shell, which can be accessed from the console
Option D :

The predict log in the Oracle Cloud Infrastructure (OCI) Logging service as defined in the deployment

Correct Answer: D

Explanation/Reference:

Detailed Answer in Step-by-Step Solution: Objective: Locate stdout/stderr output from score.py in deployment. Understand Deployment: score.py runs in a model endpoint; logs are managed by OCI. Evaluate Options: A: False-No VM file is defined by default; logs go to OCI Logging. B: False-Console displays UI, not raw logs. C: False-Cloud Shell is a separate tool, not a log destination. D: True-Predict logs in OCI Logging capture stdout/stderr. Reasoning: OCI centralizes logs in its Logging service for deployments. Conclusion: D is correct. OCI documentation states: -oeDuring model deployment, stdout and stderr from score.py are automatically sent to the predict log in the OCI Logging service, as configured in the deployment settings.- A, B, and C don't align with this managed logging approach-only D is accurate. : Oracle Cloud Infrastructure Data Science Documentation, "Model Deployment - Logging".

Success is just a click away visit Now: https://www.certsexpert.com/1Z0-1110-25-pdf-questions.html

BONUS!!! Download part of Lead2Passed 1z0-1110-25 dumps for free: https://drive.google.com/open?id=15Rpqw_HmZUZnUUOj9-B38z14AYnLQpvO

Our 1z0-1110-25 latest preparation materials provide users with three different versions, including a PDF version, a software version, and an online version. Although involved three versions of the 1z0-1110-25 teaching content is the same, but for all types of users can realize their own needs, whether it is which version of 1z0-1110-25 Learning Materials, believe that can give the user a better 1z0-1110-25 learning experience. Below, I would like to introduce you to the main advantages of our research materials, and I'm sure you won't want to miss it.

Oracle 1z0-1110-25 Exam Syllabus Topics:

Topic	Details
Topic 1	• Use Related OCI Services: This final section measures the competence of Machine Learning Engineers in utilizing OCI-integrated services to enhance data science capabilities. It includes creating Spark applications through OCI Data Flow, utilizing the OCI Open Data Service, and integrating other tools to optimize data handling and model execution workflows.

Topic 2	• Implement End-to-End Machine Learning Lifecycle: This section evaluates the abilities of Machine Learning Engineers and includes an end-to-end walkthrough of the ML lifecycle within OCI. It involves data acquisition from various sources, data preparation, visualization, profiling, model building with open-source libraries, Oracle AutoML, model evaluation, interpretability with global and local explanations, and deployment using the model catalog.
Topic 3	• Create and Manage Projects and Notebook Sessions: This part assesses the skills of Cloud Data Scientists and focuses on setting up and managing projects and notebook sessions within OCI Data Science. It also covers managing Conda environments, integrating OCI Vault for credentials, using Git-based repositories for source code control, and organizing your development environment to support streamlined collaboration and reproducibility.
Topic 4	• OCI Data Science - Introduction & Configuration: This section of the exam measures the skills of Machine Learning Engineers and covers foundational concepts of Oracle Cloud Infrastructure (OCI) Data Science. It includes an overview of the platform, its architecture, and the capabilities offered by the Accelerated Data Science (ADS) SDK. It also addresses the initial configuration of tenancy and workspace setup to begin data science operations in OCI.
Topic 5	• Apply MLOps Practices: This domain targets the skills of Cloud Data Scientists and focuses on applying MLOps within the OCI ecosystem. It covers the architecture of OCI MLOps, managing custom jobs, leveraging autoscaling for deployed models, monitoring, logging, and automating ML workflows using pipelines to ensure scalable and production-ready deployments.

1z0-1110-25 Valid Test Tips | 1z0-1110-25 Exam Preparation

With the rapid development of IT technology, the questions in the IT certification exam are also changing. Therefore, Lead2Passed also keeps updating test questions and answers. And if you purchase Lead2Passed Oracle 1z0-1110-25 Practice Test materials, we will provide you with free updates for a year. As long as the questions updates, Lead2Passed will immediately send the latest questions and answers to you which guarantees that you can get the latest materials at any time. Lead2Passed can not only help you pass the test, but also help you learn the latest knowledge. Never pass up a good chance to have the substantial materials.

Oracle Cloud Infrastructure 2025 Data Science Professional Sample Questions (Q100-Q105):

NEW QUESTION # 100
You loaded data into Oracle Cloud Infrastructure (OCI) Data Science. To transform the data, you want to use the Accelerated Data Science (ADS) SDK. When you applied the get_recommendations() tool to the ADSDataset object, it showed you user-detected issues with all the recommended changes to apply to the dataset. Which option should you use to apply all the recommended transformations at once?

- A. visualize_transforms()
- B. get_transformed_dataset()
- C. auto_transform()
- D. fit_transform()

Answer: C

Explanation:
Detailed Answer in Step-by-Step Solution:
* Objective: Apply all recommended transformations from get_recommendations() in ADS.
* Understand ADS Tools: get_recommendations() suggests fixes (e.g., missing values).
* Evaluate Options:
* A: Returns transformed data-Not for applying-incorrect.
* B: Sklearn-style, not ADS-specific-incorrect.
* C: auto_transform()-Applies all recommendations-correct.

* D: Visualizes, doesn't apply-incorrect.
* Reasoning: auto_transform() executes the fixes suggested by get_recommendations().
* Conclusion: C is correct.
OCI documentation states: "After get_recommendations() identifies issues, use auto_transform() (C) on the ADSDataset to apply all recommended transformations at once." A retrieves, B is external, D visualizes- only C aligns with OCI's ADS transformation workflow.
Oracle Cloud Infrastructure ADS SDK Documentation, "Data Transformation Methods".

NEW QUESTION # 101
You are a data scientist leveraging Oracle Cloud Infrastructure (OCI) to create a model and need some additional Python libraries for processing genome sequencing data. Which of the following THREE statements are correct with respect to installing additional Python libraries to process the data?

- A. You can install private or custom libraries from your own internal repositories
- B. You cannot install a library that's not preinstalled in the provided image
- C. OCI Data Science allows root privileges in notebook sessions
- D. You can install any open-source package available in a publicly accessible Python Package Index (PyPI) repository
- E. You can only install libraries using yum and pip as a normal user

Answer: A,B,D

Explanation:
Detailed Answer in Step-by-Step Solution:
* Objective: Identify correct statements about installing Python libraries in OCI Data Science.
* Understand Notebook Sessions: Run in a managed environment with specific permissions.
* Evaluate Options:
* A: False-No root privileges; users operate as datascience with limited sudo.
* B: True-pip install from PyPI works with internet access (e.g., NAT Gateway).
* C: False-Yum isn't available; pip is the primary tool as a normal user.
* D: False-Misstated; youcaninstall non-preinstalled libraries-likely a typo (intended opposite).
* E: True-Custom repos are supported with proper network config.
* Correct Interpretation: Assuming D's intent was "Youcaninstall..." (common exam error), B, D (corrected), E are true.
* Conclusion: B, D (corrected), E are correct.
OCI documentation states: "In notebook sessions, you can install Python libraries from PyPI (B) or private repositories (E) using pip, but root privileges (A) are not granted-users operate asdatascience." Yum (C) isn' t supported, and D's phrasing contradicts capability-corrected, it's true you can install beyond preinstalled.
B, D (adjusted), E align with OCI's flexibility.
Oracle Cloud Infrastructure Data Science Documentation, "Installing Libraries in Notebook Sessions".

NEW QUESTION # 102
You have an image classification model in the model catalog which is deployed as an HTTP endpoint using model deployments. Your tenancy administrator is seeing increased demands and has asked you to increase the load balancing bandwidth from the default of 10Mbps. You are provided with the following information:
* Payload size in KB = 1024
* Estimated requests per second = 120 requests/second (Monday through Friday, in every month, in every year)
* Buffer percentage = 20%What is the optimal load balancing bandwidth to redeploy your model?

- A. 452 Mbps
- B. 7052 Mbps
- C. 52 Mbps
- D. 1152 Mbps

Answer: D

Explanation:
Detailed Answer in Step-by-Step Solution:
* Objective: Calculate optimal bandwidth for model deployment.
* Given Data:
* Payload size = 1024 KB = 1024 * 8 = 8192 Kb (kilobits).

* Requests/sec = 120.
* Buffer = 20% = 0.2.
* Calculate Base Bandwidth:
* Bits/sec = Payload * Requests = 8192 Kb * 120 = 983,040 Kb/s = 983.04 Mbps.
* Add Buffer:
* Total = Base * (1 + Buffer) = 983.04 * 1.2 = 1179.648 Mbps.
* Evaluate Options: Closest to 1179.648 Mbps is 1152 Mbps (D)-realistic rounding.
* Conclusion: D is correct.
OCI documentation advises: "Calculate bandwidth as payload size (in bits) * requests/sec, then add a buffer (e.
g., 20%) for peak loads." Here, 1024 KB = 8192 Kb, * 120 = 983.04 Mbps, * 1.2 = 1179.648 Mbps. D (1152 Mbps) is the closest practical option-452 (A) and 52 (B) are too low, 7052 (C) excessive.
Oracle Cloud Infrastructure Data Science Documentation, "Model Deployment - Load Balancing".

NEW QUESTION # 103
Which statement best describes Oracle Cloud Infrastructure Data Science Jobs?

- A. Jobs let you define and run all Oracle Cloud DevOps workloads.
- B. Jobs let you define and run repeatable tasks on fully managed third-party cloud infrastructures.
- C. Jobs let you define and run repeatable tasks on customer-managed infrastructure.
- D. Jobs let you define and run repeatable tasks on fully managed infrastructure.

Answer: D

Explanation:
Detailed Answer in Step-by-Step Solution:
* Understand OCI Data Science Jobs: This service automates ML tasks (e.g., training, evaluation) with configurable, repeatable executions.
* Key Characteristics: Jobs run on OCI's infrastructure, managed by Oracle, not the customer or third parties, and are specific to Data Science, not general DevOps.
* Evaluate Options:
* A: Correct-Jobs are defined by users (e.g., via scripts) and executed on OCI's fully managed compute resources.
* B: Incorrect-Infrastructure is managed by OCI, not the customer.
* C: Incorrect-No third-party cloud integration; it's OCI-specific.
* D: Incorrect-Jobs are for Data Science tasks (e.g., ML training), not all DevOps workloads (e.
g., CI/CD pipelines).
* Reasoning: "Fully managed" means OCI handles provisioning and scaling, aligning with A.
* Conclusion: A accurately reflects the service's purpose and operation.
OCI Data Science Jobs "allow users to define and execute repeatable machine learning tasks, such as model training or batch processing, on fully managed OCI infrastructure." This eliminates customer management (B), third-party clouds (C), or broad DevOps scope (D). The documentation emphasizes automation and management by OCI, making A the precise description.
Oracle Cloud Infrastructure Data Science Documentation, "Overview of Jobs" section.

NEW QUESTION # 104
Which OCI Data Science interaction method can function without the need of scripting?

- A. CLI
- B. Language SDKs
- C. OCI Console
- D. REST APIs

Answer: C

Explanation:
Detailed Answer in Step-by-Step Solution:
* Objective: Identify the OCI Data Science interaction method that doesn't require scripting.
* Understand Interaction Methods: OCI provides multiple ways to interact with Data Science services- some are GUI-based, others script-based.
* Evaluate Options:
* A. OCI Console: A web-based graphical interface allowing users to manage resources (e.g., create notebook sessions, deploy

models) via point-and-click-no scripting needed.
* B. CLI: Command Line Interface requires writing commands (scripts) to execute tasks (e.g., oci data-science notebook-session create).
* C. Language SDKs: Software Development Kits (e.g., Python SDK) require coding to interact programmatically (e.g., oci.data_science.DataScienceClient).
* D. REST APIs: Application Programming Interfaces require scripted HTTP requests (e.g., using curl or a programming language).
* Reasoning: Only the OCI Console (A) offers a no-code, user-friendly interface, while B, C, and D rely on scripting or programming.
* Conclusion: A is the correct answer as it eliminates the need for scripting.
The OCI Console is described in the documentation as "a browser-based interface that allows users to manage OCI Data Science resources, such as creating notebook sessions or jobs, without writing code or scripts." In contrast, the CLI (B) requires command-line scripts, SDKs (C) need programming (e.g., Python), and REST APIs (D) involve scripted API calls. The Console's GUI distinguishes it as the only option functioning without scripting, aligning with Oracle's design for accessibility to non-programmers.
Oracle Cloud Infrastructure Data Science Documentation, "Getting Started with OCI Console" section.

NEW QUESTION # 105
.....

If you are applying for the 1z0-1110-25 certification exam, it is great to show your dedication to it. You cannot take it for granted because the Oracle Cloud Infrastructure 2025 Data Science Professional (1z0-1110-25) certification test is tough and you have to pay a good sum for appearing in it. You will lose money and time by studying with 1z0-1110-25 Exam Preparation material that is not updated. So, to avoid your loss and failure in the 1z0-1110-25 exam, you must prepare with actual Oracle 1z0-1110-25 questions from Lead2Passed.

P.S. Free & New 1z0-1110-25 dumps are available on Google Drive shared by Lead2Passed: https://drive.google.com/open?id=15Rpqw_HmZUZnUUOj9-B38z14AYnLQpvO