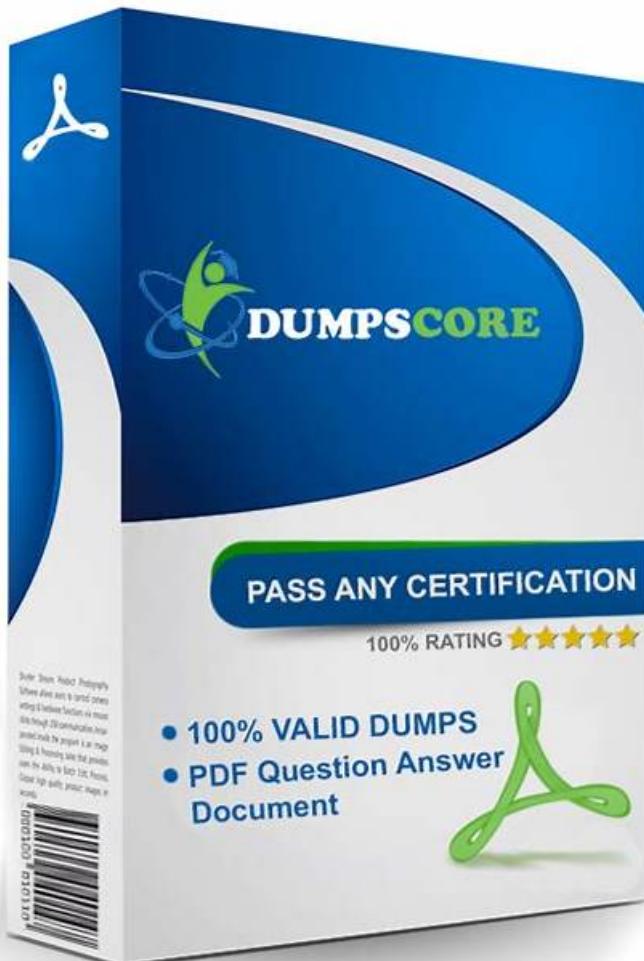


Reliable 1z0-1110-25 Test Guide, New 1z0-1110-25 Dumps Ebook



P.S. Free & New 1z0-1110-25 dumps are available on Google Drive shared by iPassleader: https://drive.google.com/open?id=1k-6LaPvmWJ3U3jd3dLzy_etD7ldEwfu

Dreaming to be a certified professional in this line? Our 1z0-1110-25 study materials are befitting choices. We made real test materials in three accessible formats for your inclinations. (PDF, APP, software). Our website is an excellent platform, which provides the questions of these versions of our 1z0-1110-25 Exam Questions compiled by experts. By browsing this website, all there versions of our 1z0-1110-25 practice engine can be chosen according to your taste or preference.

Oracle 1z0-1110-25 Exam Syllabus Topics:

| Topic | Details |
|---------|--|
| Topic 1 | <ul style="list-style-type: none">• 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 2 | <ul style="list-style-type: none"> • 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 3 | <ul style="list-style-type: none"> • 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. |
| Topic 4 | <ul style="list-style-type: none"> • 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 5 | <ul style="list-style-type: none"> • 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. |

>> Reliable 1z0-1110-25 Test Guide <<

Unparalleled Oracle Reliable 1z0-1110-25 Test Guide Are Leading Materials & Trustworthy 1z0-1110-25: Oracle Cloud Infrastructure 2025 Data Science Professional

We have been studying for many years since kindergarten. I believe that you must have your own opinions and requirements in terms of learning. Our 1z0-1110-25 learning guide has been enriching the content and form of the product in order to meet the needs of users. No matter what kind of learning method you like, you can find the best one for you at 1z0-1110-25 Exam Materials. And our 1z0-1110-25 study braindumps contain three different versions: the PDF, Software and APP online.

Oracle Cloud Infrastructure 2025 Data Science Professional Sample Questions (Q69-Q74):

NEW QUESTION # 69

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? Pick THREE.

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

Answer: B,C,E

Explanation:

Detailed Answer in Step-by-Step Solution:

* Objective: Prevent deployment failure due to high utilization.

* Evaluate Options:

- * A: More instances-Scales capacity-correct.
- * B: Delete-Stops service, not a solution.
- * C: Fewer instances-Worsens utilization.
- * D: Larger VM-Increases resource capacity-correct.

- * E: Reduce bandwidth-Limits load-correct.
- * Reasoning: A and D boost capacity, E controls demand-proactive fixes.
- * Conclusion: A, D, E are correct.

OCI documentation advises: "To handle high utilization, increase instances (A), use a larger compute shape (D), or adjust load balancer bandwidth (E) to manage request volume." B stops service, C reduces capacity- only A, D, E prevent failure per OCI's scaling options.

Oracle Cloud Infrastructure Data Science Documentation, "Model Deployment Scaling".

NEW QUESTION # 70

Which of these protects customer data at rest and in transit in a way that allows customers to meet their security and compliance requirements for cryptographic algorithms and key management?

- A. Data encryption
- B. Identity Federation
- C. Customer isolation
- D. Security controls

Answer: A

Explanation:

Detailed Answer in Step-by-Step Solution:

- * Objective: Identify protection for data at rest/transit with customer control.

* Evaluate Options:

- * A: Controls-Broad, not specific to encryption.
- * B: Isolation-Separates tenants, not crypto-focused.
- * C: Encryption-Secures data, allows key management-correct.
- * D: Federation-Auth sharing, not data protection.

* Reasoning: C provides crypto control (e.g., Vault keys).

* Conclusion: C is correct.

OCI documentation states: "Data encryption (C) protects data at rest and in transit, with customer-managed keys in OCI Vault meeting compliance needs." A and B are broader, D is unrelated-only C fits per OCI's security model.

Oracle Cloud Infrastructure Security Documentation, "Data Encryption".

NEW QUESTION # 71

You are a data scientist trying to load data into your notebook session. You understand that Accelerated Data Science (ADS) SDK supports loading various data formats. Which of the following THREE are ADS- supported data formats?

- A. JSON
- B. DOCX
- C. Pandas DataFrame
- D. XML
- E. Raw Images

Answer: A,C,D

Explanation:

Detailed Answer in Step-by-Step Solution:

- * Objective: Identify three data formats supported by ADS SDK for loading data.

* Understand ADS SDK: Facilitates data loading into notebook sessions via DatasetFactory.

* Evaluate Options:

- * A. DOCX: Not natively supported-requires conversion (e.g., to text).
- * B. Pandas DataFrame: Supported-core format for data manipulation in ADS.
- * C. JSON: Supported-common structured data format.
- * D. Raw Images: Not directly supported-image data needs preprocessing (e.g., via Vision).
- * E. XML: Supported-parseable structured format.

* Reasoning: ADS focuses on tabular/structured data-B, C, E align; A and D require external handling.

* Conclusion: B, C, E are correct.

OCI documentation states: "ADS SDK's DatasetFactory supports loading data from formats like Pandas DataFrames (B), JSON (C), and XML (E), enabling easy integration into notebook sessions." DOCX (A) isn't natively handled, and raw images (D) require

preprocessing outside ADS-B, C, E match the supported list.
Oracle Cloud Infrastructure ADS SDK Documentation, "Supported Data Formats".

NEW QUESTION # 72

Which encryption is used for Oracle Data Science?

- A. 256-bit Advanced Encryption Standard (AES-256)
- B. Rivest Shamir Adleman (RSA)
- C. Triple DES (TDES)
- D. Data Encryption Standard (DES)
- E. Twofish

Answer: A

Explanation:

Detailed Answer in Step-by-Step Solution:

- * Objective: Identify encryption standard for OCI Data Science.
- * Understand OCI Encryption: Applies to data at rest and in transit.
- * Evaluate Options:
 - * A: AES-256-Industry-standard, OCI default-correct.
 - * B: DES-Outdated, weak-incorrect.
 - * C: TDES-Older, less secure-incorrect.
 - * D: Twofish-Not OCI standard-incorrect.
 - * E: RSA-Asymmetric, not primary for data at rest-incorrect.
- * Reasoning: AES-256 is OCI's go-to for Data Science resources.
- * Conclusion: A is correct.

OCI documentation states: "Data Science services encrypt data at rest using AES-256 (A), ensuring high security for notebooks, jobs, and models." B, C, D, and E are either outdated or not used-only A matches OCI's encryption policy.

Oracle Cloud Infrastructure Data Science Documentation, "Data Encryption".

NEW QUESTION # 73

You are working as a Data Scientist for a healthcare company. You have a series of neurophysiological data on OCI Data Science and have developed a convolutional neural network (CNN) classification model. It predicts the source of seizures in drug-resistant epileptic patients. You created a model artifact with all the necessary files. When you deployed the model, it failed to run because you did not point to the correct conda environment in the model artifact. Where would you provide instructions to use the correct conda environment?

- A. requirements.txt
- B. model_artifact_validate.py
- C. score.py
- D. runtime.yaml

Answer: D

Explanation:

Detailed Answer in Step-by-Step Solution:

- * Objective: Determine where to specify the conda environment for an OCI model deployment.
- * Understand Model Deployment: Requires artifacts like score.py and runtime.yaml to define runtime settings.
- * Evaluate Options:
 - * A. score.py: Contains inference logic (e.g., load_model(), predict())-not for environment specs.
 - * B. runtime.yaml: Defines deployment runtime, including conda environment path-correct.
 - * C. requirements.txt: Lists pip dependencies-not used in OCI for conda environments.
 - * D. model_artifact_validate.py: Not a standard artifact; doesn't exist in OCI deployment.
- * Reasoning: runtime.yaml specifies the conda env (e.g., slug: pyspark30_p37_cpu_v2)-failure to set this causes deployment errors.
- * Conclusion: B is correct.

OCI documentation states: "The runtime.yaml file in a model artifact specifies the runtime environment, including the conda environment path (e.g., ENVIRONMENT_SLUG: pyspark30_p37_cpu_v2), ensuring the deployed model uses the correct dependencies." score.py (A) handles inference, requirements.txt (C) is for pip (not conda in OCI), and D isn't valid-only B addresses the conda issue per OCI's deployment process.

NEW QUESTION # 74

During nearly ten years, our company has kept on improving ourselves, and now we have become the leader in this field. And now our 1z0-1110-25 training materials have become the most popular 1z0-1110-25 practice materials in the international market. There are so many advantages of our 1z0-1110-25 Study Materials, and as long as you free download the demos on our website, then you will know that how good quality our 1z0-1110-25 exam questions are in! You won't regret for your wise choice if you buy our 1z0-1110-25 learning guide!

New 1z0-1110-25 Dumps Ebook: <https://www.ipassleader.com/Oracle/1z0-1110-25-practice-exam-dumps.html>

What's more, part of that iPassleader 1z0-1110-25 dumps now are free: https://drive.google.com/open?id=1k-6LaPvmlWJ3U3jd3dLzy_etD7ldEfwu