

New 1z0-1110-25 Exam Papers & 1z0-1110-25 Training Materials

Sample Questions for Oracle 1Z0-1110-25 Exam By Roberson - Page 1

P2P
exams

Free Questions for 1Z0-1110-25

Shared by Roberson on 10-04-2025

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)

P2P
exams

DOWNLOAD the newest RealVCE 1z0-1110-25 PDF dumps from Cloud Storage for free: <https://drive.google.com/open?id=1uY9N80iD4HOk9lcRYol-7K4DdiVneamM>

The modern Oracle world is changing its dynamics at a fast pace and has become so competitive. To stay updated and competitive in the market you have to learn new in-demand skills. With one Oracle 1z0-1110-25 exam certificate you can do this task nicely. With the Oracle 1z0-1110-25 Certification Exam successful candidates can validate their knowledge, increase marketability, enhance academic performance, improve reputation and increase earning power and other personal and professional benefits, etc.

Taking these mock exams is important because it tells you where you stand. People who are confident about their knowledge and expertise can take these 1z0-1110-25 practice tests and check their scores to know where they lack. This is good practice to be a pro and clear your Oracle Cloud Infrastructure 2025 Data Science Professional (1z0-1110-25) exam with amazing scores. RealVCE practice tests simulate the real 1z0-1110-25 exam questions environment.

>> [New 1z0-1110-25 Exam Papers](#) <<

New 1z0-1110-25 Exam Papers | Reliable 1z0-1110-25: Oracle Cloud Infrastructure 2025 Data Science Professional 100% Pass

The Oracle 1z0-1110-25 practice exam software of RealVCE has questions that have a striking resemblance to the queries of the Oracle Cloud Infrastructure 2025 Data Science Professional (1z0-1110-25) real questions. It has a user-friendly interface. You don't require an active internet connection to run it once the 1z0-1110-25 Practice Test software is installed on Windows computers

and laptops.

Oracle Cloud Infrastructure 2025 Data Science Professional Sample Questions (Q55-Q60):

NEW QUESTION # 55

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 would you employ to verify accuracy and lower customer churn?

- A. Operational monitoring
- B. Redeploy the model
- C. Drift monitoring
- D. Retrain the model
- E. Validate the model using recent data

Answer: C,D

Explanation:

Detailed Answer in Step-by-Step Solution:

- * Objective: Address declining model performance and improve churn prediction.
- * Analyze Issue: Poor predictions suggest data drift or model staleness-common ML challenges.
- * Evaluate Options:
 - * A. Drift monitoring: Tracks changes in data distribution-identifies root cause of accuracy drop- correct.
 - * B. Redeploy the model: Repeats deployment without fixing the issue-ineffective alone.
 - * C. Operational monitoring: Tracks system health (e.g., latency), not prediction quality.
 - * D. Retrain the model: Updates model with new data-directly improves accuracy-correct.
 - * E. Validate with recent data: Checks performance but doesn't fix-diagnostic, not corrective.
- * Reasoning: A diagnoses drift (cause), D retrains to adapt (solution)-best combo to verify and lower churn.
- * Conclusion: A and D are correct.

OCI documentation advises: "Drift monitoring (A) detects shifts in data distribution that degrade model performance, while retraining (D) with fresh data restores accuracy." Redeployment (B) doesn't address drift, operational monitoring (C) focuses on infra, and validation (E) is a check-not a fix. A and D align with OCI's model maintenance strategy.

Oracle Cloud Infrastructure Data Science Documentation, "Model Monitoring and Retraining".

NEW QUESTION # 56

As a data scientist for a hardware company, you have been asked to predict the revenue demand for the upcoming quarter. You develop a time series forecasting model to analyze the data. Select the correct sequence of steps to predict the revenue demand values for the upcoming quarter.

- A. Verify, prepare model, deploy, save, predict
- B. Prepare model, verify, save, deploy, predict
- C. Prepare model, deploy, verify, save, predict
- D. Predict, deploy, save, verify, prepare model

Answer: B

Explanation:

Detailed Answer in Step-by-Step Solution:

- * Prepare Model: Build and train the time series model using historical data.
- * Verify: Validate the model's accuracy (e.g., using metrics like MAE or RMSE).
- * Save: Store the trained model (e.g., in the OCI Model Catalog).
- * Deploy: Make the model available for predictions (e.g., via OCI Model Deployment).
- * Predict: Generate revenue forecasts for the upcoming quarter.
- * Evaluate Options: D follows this logical flow; others (e.g., A starts with "verify" before preparation) don't.

In OCI Data Science, the workflow for time series forecasting involves preparing the model (training), verifying its performance, saving it to the catalog, deploying it, and then predicting. This sequence is standard for ML deployment in OCI, as per the documentation. (Reference: Oracle Cloud Infrastructure Data Science Documentation, "Time Series Forecasting Workflow").

NEW QUESTION # 57

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?

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

Answer: C

Explanation:

Detailed Answer in Step-by-Step Solution:

* Objective: Fix network timeout for pip install in a notebook.

* Evaluate Options:

* A: FastConnect-On-premises link, not public internet.

* B: VNIC-Default, not the issue.

* C: NAT Gateway-Grants internet access-correct.

* D: Service Gateway-OCI services, not PyPI.

* Reasoning: C enables outbound traffic to public repos.

* Conclusion: C is correct.

OCI documentation states: "A NAT Gateway (C) is required for notebook sessions in private subnets to access public internet repositories like PyPI." A, B, and D don't provide this-only C resolves the timeout.

Oracle Cloud Infrastructure Data Science Documentation, "Notebook Networking".

NEW QUESTION # 58

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 could you use to save the model artifact into the model catalog?

- A. Console
- B. Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI)
- C. Git CLI
- D. ODSC CLI
- E. OCI Python SDK
- F. Accelerated Data Science (ADS) Software Development Kit (SDK)

Answer: A,E,F

Explanation:

Detailed Answer in Step-by-Step Solution:

* Objective: Identify interfaces to save a 2 GB model to the Model Catalog.

* Evaluate Options:

* A: OCI CLI-Supports Data Science tasks-possible but not primary.

* B: ADS SDK-Designed for model catalog ops-correct.

* C: ODSC CLI-Not standard; likely typo for OCI CLI.

* D: Console-GUI for catalog uploads-correct.

* E: OCI Python SDK-Programmatic catalog access-correct.

* F: Git CLI-Version control, not catalog-related.

* Reasoning: B, D, E are OCI's primary interfaces; A is valid but less emphasized.

* Conclusion: B, D, E are correct (A plausible but not top-tier).

OCI documentation lists "ADS SDK (B), OCI Console (D), and OCI Python SDK (E) as primary methods to save models to the Model Catalog." OCI CLI (A) works but isn't highlighted, C isn't real, and F is unrelated- B, D, E are the standard trio.

Oracle Cloud Infrastructure Data Science Documentation, "Model Catalog Interfaces".

NEW QUESTION # 59

You're going to create an Oracle Cloud Infrastructure Anomaly Detection model for multivariate data. Where do you need to store the training data?

- A. Autonomous Data Warehouse
- **B. Object Storage Bucket**
- C. MySQL database
- D. Your local machine

Answer: B

Explanation:

Detailed Answer in Step-by-Step Solution:

* Understand OCI Anomaly Detection: This service trains models to detect anomalies in multivariate data (e.g., sensor readings), requiring data to be accessible within OCI's ecosystem.

* Assess Storage Requirements: The training data must be in a scalable, OCI-compatible location that the Anomaly Detection service can access programmatically.

* Evaluate Options:

* A. Your local machine: Data on a local machine isn't directly accessible to OCI services without upload, making it impractical for cloud-based training.

* B. MySQL database: While OCI supports MySQL, Anomaly Detection doesn't natively integrate with it for training data; it prefers file-based input.

* C. Autonomous Data Warehouse: This is a database for analytics, not the default storage for Anomaly Detection training data, which expects CSV/JSON files.

* D. Object Storage Bucket: OCI Object Storage is a scalable, durable storage service that Anomaly Detection uses to ingest training data (e.g., CSV files).

* Reasoning: Object Storage is the standard for large-scale data in OCI services, offering seamless integration with Anomaly Detection via APIs or SDKs.

* Conclusion: D is the correct choice as it aligns with the service's architecture.

The OCI Anomaly Detection service requires training data to be uploaded to an Object Storage Bucket in formats like CSV or JSON. This is explicitly outlined in the official documentation, which states that users must "upload the training dataset to an OCI Object Storage bucket" before creating a data asset for model training. Local storage (A) isn't viable for cloud processing, and databases like MySQL (B) or Autonomous Data Warehouse (C) aren't supported as primary inputs. Object Storage (D) provides the scalability and accessibility needed for multivariate anomaly detection workflows.

Oracle Cloud Infrastructure Anomaly Detection Documentation, "Preparing Training Data" section.

NEW QUESTION # 60

.....

Our team regularly modified it to provide you with the real and updated 1z0-1110-25 pdf exam questions every time. The applicants are informed of these new changes till three months after purchase from the RealVCE. The RealVCE gives its applicants a Oracle 1z0-1110-25 web-based practice test software that doesn't require installation. Oracle 1z0-1110-25 Practice Test is compatible with all operating systems, including iOS, Mac, and Windows. You can use this Oracle 1z0-1110-25 practice test on any browser on any device anywhere. You need to sign in to a verified account on our website to use the entire premium Oracle 1z0-1110-25 practice test questions.

1z0-1110-25 Training Materials: https://www.realvce.com/1z0-1110-25_free-dumps.html

The advantages of 1z0-1110-25 study materials are numerous and they are all you need, First and foremost, we have high class operation system so we can assure you that you can start to prepare for the 1z0-1110-25 exam with our 1z0-1110-25 study materials only 5 to 10 minutes after payment, Oracle New 1z0-1110-25 Exam Papers You can use the practice test software to check your learning outcomes, Oracle New 1z0-1110-25 Exam Papers The free demo is a part of our complete on-sale exam dump.

Anticipate how Intent-Based Networking will change your enterprise architecture, IT operations, and business, The Custom Installation Procedure, The advantages of 1z0-1110-25 Study Materials are numerous and they are all you need!

Free Oracle 1z0-1110-25 Exam Questions Updates and Demos

First and foremost, we have high class operation system so we can assure you that you can start to prepare for the 1z0-1110-25 exam with our 1z0-1110-25 study materials only 5 to 10 minutes after payment.

You can use the practice test software to check your learning 1z0-1110-25 outcomes, The free demo is a part of our complete on-sale exam dump, Here are a few more examples we can add to the list.

