Download Oracle 1Z0-1110-25 Exam Dumps Free

Shared by Davenport on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which Oracle Cloud Infrastructure (OCI) Data Science policy is invalid?

Options:

A- Allow group DataScienceGroup to use virtual-network-family in compartment DataScience
B- Allow group DataScienceGroup to use data-science-model-sessions in compartment DataScience
C- Allow dynamic-group DataScienceDynamicGroup to manage data-science-projects in compartment DataScience
D- Allow dynamic-group DataScienceDynamicGroup to manage data-science-family in compartment DataScience

Answer:

B

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Spot the invalid OCI IAM policy.

Understand Policies: Use verbs (e.g., use, manage) and resource types.

Evaluate Options:

A: Valid---virtual-network-family is a resource type.

B: Invalid---No data-science-model-sessions resource type; likely meant data-science-models.

C: Valid---data-science-projects is correct.

D: Valid---data-science-family aggregates Data Science resources.

Reasoning: B uses a nonexistent resource type---error in syntax.

Conclusion: B is incorrect.

OCI IAM documentation lists valid resource types: ''data-science-models, data-science-projects, data-science-family---but not data-science-model-sessions (B).'' A, C, D are correct---B is invalid due to this typo/misunderstanding.

: Oracle Cloud Infrastructure IAM Documentation, 'Data Science Resource Types'.

Question 2

Question Type: MultipleChoice

Which OCI service enables you to build, train, and deploy machine learning models in the cloud?

Options:

A- Oracle Cloud Infrastructure Data Catalog
B- Oracle Cloud Infrastructure Data Integration
C- Oracle Cloud Infrastructure Data Science
D- Oracle Cloud Infrastructure Data Flow

Answer:

C

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Identify the OCI service for ML model lifecycle.

Evaluate Options:

A: Data Catalog---Metadata management, not ML.

B: Data Integration---ETL, not ML.

C: Data Science---Full ML lifecycle---correct.

D: Data Flow---Spark processing, not full ML.

Reasoning: C supports building, training, deploying models.

Conclusion: C is correct.

OCI documentation states: ''OCI Data Science (C) provides tools to build, train, and deploy machine learning models in the cloud, including notebooks and model catalog.'' A, B, and D serve other purposes---only C fits the ML lifecycle per OCI's offerings.

: Oracle Cloud Infrastructure Data Science Documentation, 'Service Overview'.

Question 3

Question Type: MultipleChoice

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?

Options:

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

Answer:

A, 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'.

Question 4

Question Type: MultipleChoice

As a data scientist, you use the Oracle Cloud Infrastructure (OCI) Language service to train custommodels. Which types of custom models can be trained?

Options:

A- Image classification, Named Entity Recognition (NER)
B- Text classification, Named Entity Recognition (NER)
C- Sentiment Analysis, Named Entity Recognition (NER)
D- Object detection, Text classification

Answer:

B

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Identify custom model types for OCI Language service.

Understand OCI Language: Focuses on text analysis, not images.

Evaluate Options:

A: Image classification---Incorrect; Language is text-based.

B: Text classification, NER---Both text tasks---correct.

C: Sentiment---Pretrained, not custom; NER ok---incorrect.

D: Object detection---Image-based, incorrect.

Reasoning: B aligns with OCI Language's custom text capabilities.

Conclusion: B is correct.

OCI documentation states: ''OCI Language supports training custom models for text classification and Named Entity Recognition (NER) (B) using user data.'' A and D involve images, C includes pretrained sentiment---only B matches OCI Language's custom model scope.

: Oracle Cloud Infrastructure Language Documentation, 'Custom Model Types'.

Question 5

Question Type: MultipleChoice

You are running a pipeline in the OCI Data Science service and want to override some of the pipeline's default settings. Which of the following statements about overriding pipeline defaults is true?

Options:

A- Pipeline defaults can be overridden only during pipeline creation.
B- Pipeline defaults can be overridden only by the Administrator.
C- Pipeline defaults can be overridden before starting the pipeline execution.
D- Pipeline defaults cannot be overridden once the pipeline has been created.

Answer:

C

Explanation:

Detailed Answer in Step-by-Step Solution:

Understand OCI Data Science Pipelines: Pipelines automate ML workflows with configurable steps.

Check Override Mechanism: Defaults (e.g., compute shape, storage) can be modified before execution via the OCI Console, SDK, or CLI.

Evaluate Options:

A: False---Overrides can occur post-creation, before running.

B: False---Any authorized user, not just admins, can override defaults.

C: True---Settings can be adjusted before execution starts.

D: False---Defaults can be changed post-creation, pre-execution.

Conclusion: C is correct as it reflects the flexibility of pipeline configuration.

OCI Data Science Pipelines allow users to override default settings (e.g., compute resources, environment variables) before execution, as noted in the official documentation. This can be done via the UI or programmatically, offering flexibility beyond creation time (A) and without admin-only restrictions (B). (Reference: Oracle Cloud Infrastructure Data Science Pipelines Documentation, 'Configuring Pipelines').

Question 6

Question Type: MultipleChoice

Which statement about resource principals is true?

Options:

A- When you authenticate using a resource principal, you need to create and manage credentials to access OCI resources.
B- A resource principal is not a secure way to authenticate to resources, compared to the OCI configuration and API key approach.
C- The Data Science service does not provide authentication via a notebook session's or job run's resource principal to access other OCI resources.
D- A resource principal is a feature of IAM that enables resources to be authorized principal actors.

Answer:

D

Explanation:

Detailed Answer in Step-by-Step Solution:

Define Resource Principals: They allow OCI resources (e.g., notebook sessions) to authenticate to other OCI services without user credentials.

Evaluate Options:

A: False---Resource principals eliminate manual credential management.

B: False---They're secure, leveraging IAM policies, not less secure than API keys.

C: False---Data Science supports resource principals for accessing resources (e.g., Object Storage).

D: True---Resource principals are an IAM feature authorizing resources as actors.

Reasoning: D captures the essence of resource principals as an IAM mechanism.

Conclusion: D is correct.

OCI documentation states: ''A resource principal is an IAM feature that enables OCI resources, such as compute instances or notebook sessions, to act as principal actors and authenticate to other OCI services using policies.'' This refutes A (no credentials needed), B (secure method), and C (supported in Data Science), making D the accurate statement.

: Oracle Cloud Infrastructure IAM Documentation, 'Resource Principals'.

Question 7

Question Type: MultipleChoice

Which step is a part of the AutoML pipeline?

Options:

A- Feature Extraction
B- Model saved to Model Catalog
C- Model Deployment
D- Feature Selection

Answer:

D

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Identify a step in OCI's AutoML pipeline.

Understand AutoML: Automates model building---includes preprocessing, selection, and tuning.

Evaluate Options:

A: Feature Extraction (e.g., PCA) isn't explicitly part of OCI AutoML---too specific.

B: Saving to Model Catalog is post-AutoML, not a pipeline step.

C: Deployment is a separate action after AutoML---incorrect.

D: Feature Selection (e.g., choosing relevant features) is a core AutoML step---correct.

Reasoning: OCI AutoML automates feature selection, algorithm choice, and tuning---D fits.

Conclusion: D is correct.

OCI AutoML's pipeline includes ''feature selection, algorithm selection, adaptive sampling, and hyperparameter tuning,'' per the documentation. Extraction (A) isn't highlighted, while saving (B) and deployment (C) are post-process actions---only Feature Selection (D) is an integral automated step.

: Oracle Cloud Infrastructure Data Science Documentation, 'AutoML Pipeline'.

Question 8

Question Type: MultipleChoice

Which Web Application Firewall (WAF) service component must be configured to allow, block, or log network requests when they meet specified criteria?

Options:

A- Protection rules
B- Bot Management
C- Origin
D- Web Application Firewall policy

Answer:

A

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Identify the WAF component that controls request actions based on criteria.

Understand WAF Components:

Protection Rules: Define conditions and actions (e.g., allow, block, log).

Bot Management: Handles bot traffic, not general request rules.

Origin: Backend server endpoint, not rule-based.

WAF Policy: Umbrella config, but rules specify actions.

Evaluate Options:

A: Protection rules---Set specific criteria and actions---correct.

B: Bot Management---Bot-specific, not general requests.

C: Origin---Defines source, not actions.

D: WAF policy---Broad config, not the granular rules.

Reasoning: Protection rules directly manage request behavior---fit the requirement.

Conclusion: A is correct.

OCI documentation states: ''Protection rules (A) in WAF define conditions (e.g., IP, URL) and actions (allow, block, log) for incoming requests.'' Bot Management (B) targets bots, Origin (C) is a target server, and WAF Policy (D) encompasses rules but isn't the action specifier---only A aligns with OCI's WAF configuration.

: Oracle Cloud Infrastructure WAF Documentation, 'Protection Rules'.

Question 9

Question Type: MultipleChoice

You have an embarrassingly parallel or distributed batch job on a large amount of data that you consider running using Data Science Jobs. What would be the best approach to run the workload?

Options:

A- Create the job in Data Science Jobs and start a job run. When it is done, start a new job run until you achieve the number of runs required

B- Create the job in Data Science Jobs and then start the number of simultaneous job runs required for your workload

C- Reconfigure the job run because Data Science Jobs does not support embarrassingly parallel workloads

D- Create a new job for every job run that you have to run in parallel, because the Data Science

Jobs service can have only one job run per job

Answer:

B

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Optimize embarrassingly parallel workload in OCI Jobs.

Evaluate Options:

A: Sequential runs---Inefficient for parallel tasks.

B: Simultaneous runs---Maximizes parallelism---correct.

C: False---Jobs support parallelism.

D: One job per run---Misstates capability, wasteful.

Reasoning: B leverages OCI's parallel run support.

Conclusion: B is correct.

OCI documentation states: ''For embarrassingly parallel tasks, create one Job and launch multiple simultaneous Job Runs (B) to process data efficiently.'' A is slow, C is incorrect, and Dovercomplicates---B is the best approach.

: Oracle Cloud Infrastructure Data Science Documentation, 'Parallel Job Execution'.

Question 10

Question Type: MultipleChoice

Which of the following programming languages are most widely used by data scientists?

Options:

A- C and C++
B- Python, R, and SQL
C- Java and JavaScript

Answer:

B

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Identify top languages for data science.

Evaluate Options:

A: C/C++---Low-level, less common for data tasks.

B: Python (ML, libraries), R (stats), SQL (data)---Industry standards.

C: Java (enterprise), JavaScript (web)---Not data-focused.

Reasoning: B aligns with data science tools (e.g., pandas, ggplot).

Conclusion: B is correct.

OCI documentation highlights ''Python, R, and SQL as the most widely used languages in Data Science for modeling, analysis, and data querying.'' C/C++ (A) and Java/JS (C) are less prevalent---B matches OCI's notebook support and industry trends.

: Oracle Cloud Infrastructure Data Science Documentation, 'Supported Languages'.

Question 11

Question Type: MultipleChoice

You are attempting to save a model from a notebook session to the model catalog by using 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?

Options:

A- The dynamic groups matching rule exists for notebook sessions in the compartment
B- The model artifact is saved to the block volume of the notebook session
C- The policy for the dynamic group grants manage permissions for the model catalog in this compartment
D- The networking configuration allows access to the Oracle Cloud Infrastructure services through a service gateway

E- The policy for your user group grants manage permissions for the model catalog in this compartment

Answer:

A, C

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Troubleshoot a 404 authentication error when saving a model using ADS SDK with resource principal.

Understand Resource Principal: Allows notebook sessions to act as principals via dynamic groups and policies---no user credentials needed.

Analyze 404 Error: Indicates an authorization failure---likely missing permissions or misconfigured resource principal.

Evaluate Options:

A: True---Dynamic group must include notebook sessions (e.g., resource.type = 'datasciencenotebooksession') to authenticate.

B: False---Block volume stores artifacts locally, but saving to the catalog is a permission issue, not storage.

C: True---Policy must grant manage data-science-models to the dynamic group for catalog access.

D: False---Service gateway ensures network access, but 404 is auth-related, not connectivity.

E: False---Resource principal uses dynamic group policies, not user group policies.

Reasoning: A (group inclusion) and C (policy permission) are critical for resource principal auth---others are tangential.

Conclusion: A and C are correct.

OCI documentation states: ''To use resource principal with ADS SDK for model catalog operations, ensure (1) a dynamic group includes the notebook session with a matching rule (e.g., all {resource.type = 'datasciencenotebooksession'}) and (2) a policy grants the dynamic group manage data-science-models permissions in the compartment.'' B is unrelated (storage location), D is network-focused, and E applies to user auth---not resource principal. A 404 error flags missing auth, fixed by A and C.

: Oracle Cloud Infrastructure Data Science Documentation, 'Using Resource Principals with ADS SDK'.

Question 12

Question Type: MultipleChoice

As a data scientist, you require a pipeline to train ML models. When can a pipeline run be initiated?

Options:

A- Pipeline can be initiated once it is created.
B- Pipeline can be initiated during the pipeline run state.
C- Pipeline can be initiated after the active state.
D- Pipeline can be initiated before the active state.

Answer:

A

Explanation:

Detailed Answer in Step-by-Step Solution:

Objective: Determine when an OCI Data Science pipeline can start.

Understand Pipelines: They're workflows with defined steps, executed on demand or scheduled.

Evaluate Options:

A: Once created, a pipeline can be run immediately---correct.

B: ''During run state'' implies it's already running---illogical.

C: ''After active state'' is unclear; pipelines run when triggered, not post-state.

D: ''Before active state'' is vague---creation precedes running.

Reasoning: Pipelines are executable post-creation via UI/CLI---simplest interpretation is A.

Conclusion: A is correct.

OCI Data Science documentation states: ''After a pipeline is created, you can initiate a pipeline run immediately or schedule it using the OCI Console, CLI, or SDK.'' B, C, and D misalign with this---running starts post-creation (A), not during/after ambiguous states.

: Oracle Cloud Infrastructure Data Science Documentation, 'Pipelines - Running a Pipeline'.