

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : **AI-300**

Title : Operationalizing Machine
Learning and Generative AI
Solutions

Version : DEMO

1.A GenAIOps Engineer is planning the production rollout strategy for a new foundation model. The primary requirement is to ensure that the deployment process allows for rapid recovery and auditability of the model's history throughout its lifecycle.

Which capability is most essential for tracking model lineage and ensuring that the specific artifacts used in production are reproducible?

- A. Model registration with unique versioning and metadata tracking.
- B. Overwriting the previous production model file directly.
- C. Manual renaming of model files on the main storage account.
- D. Storing model files in an unversioned public cloud storage bucket.

Answer: A

2.A Data Scientist is using a notebook for experimentation and exploration. They are training a baseline model on a Compute Instance and need to log the training metrics back to the Azure ML workspace using MLflow.

```
from azure.ai.ml import MLClient
from azure.identity import DefaultAzureCredential
import mlflow
ml_client = MLClient(DefaultAzureCredential(), sub_id, rg_name, ws_name)
# Missing MLflow tracking setup
with mlflow.start_run():
    mlflow.log_metric("accuracy", 0.92)
```

Which two programmatic steps must be included in the notebook before calling `mlflow.start_run()` to route the MLflow tracking data correctly to the workspace? (Choose two.)

- A. Pass the retrieved URI to `mlflow.set_tracking_uri()`.
- B. Instantiate a `MachineLearningServices` local SQLite database.
- C. Set the `MLFLOW_TRACKING_URI` environment variable to a local directory path.
- D. Use `ml_client.workspaces.get()` to retrieve the `mlflow_tracking_uri` property.
- E. Execute `mlflow.server()` directly within the notebook cell.

Answer: A, D

3.A Solution Architect is designing a new Azure Machine Learning workspace to support multiple data science squads. The squads frequently build complex machine learning pipelines containing distinct steps for data preprocessing, model training, and evaluation. To ensure reusability and enforce organizational standards, the architect mandates that all pipeline steps must be modularized and shared centrally.

Which workspace asset type must be explicitly created to fulfill this modular pipeline design requirement?

- A. A component asset containing the step logic, environment, and interface.
- B. A model asset tagged with the specific pipeline step execution metadata.
- C. An environment asset referencing a standard Conda configuration file.
- D. A datastore asset configured to share intermediate pipeline outputs globally.

Answer: A

4.A Quality Engineer wants to share a pipeline component across workspaces to standardize data

preprocessing. They are using the Azure CLI to promote a local component definition to the central organizational registry.

```
az ml component create --file prep_data.yml --registry-name shared-ai-registry
```

When executing this CLI command to share the component, which two operational constraints and behaviors apply to the infrastructure? (Choose two)

- A. The command automatically iterates through all workspaces associated with the user's subscription and creates a local duplicate of the component in each.
- B. The registry must be located in the exact same Azure region as the user's local workspace to allow component creation.
- C. The component's execution environment is forced to compile locally on the CLI host before registration completes.
- D. The physical code snapshot defined in prep_data.yml is uploaded directly to the registry's managed storage, not the local workspace's storage.
- E. The identity executing this command must have appropriate RBAC permissions (e.g., AzureML Registry Contributor) on the target registry resource.

Answer: D, E

5. An AI Quality Engineer is examining continuous monitoring in Foundry to establish a complete post-production observability strategy. The quality goal is to ensure the generative AI system is resilient to new adversarial attacks, tracks operational health, and maintains textual response quality.

The team is reviewing the available post-production monitoring components in Microsoft Foundry to define their telemetry architecture.

Which of the following components should be implemented to achieve these comprehensive continuous monitoring goals? (Select all that apply.)

- A. Continuous evaluation.
- B. Scheduled red teaming.
- C. Provisioned throughput capacity pooling.
- D. Operational metrics.
- E. Azure Monitor alerts

Answer: A, B, D, E