

KTest

高品質 より良いサービス



練習問題

<https://www.ktest.jp>

1年間無料アップデート

Exam : **AI-300**

Title : Operationalizing Machine
Learning and Generative AI
Solutions

Version : DEMO

1. Case Study

This is a case study. Case studies are not timed separately from other exam sections. You can use as much exam time as you would like to complete each case study. However, there might be additional case studies or other exam sections. Manage your time to ensure that you can complete all the exam sections in the time provided. Pay attention to the Exam Progress at the top of the screen so you have sufficient time to complete any exam sections that follow this case study.

To answer the case study questions, you will need to reference information that is provided in the case. Case studies and associated questions might contain exhibits or other resources that provide more information about the scenario described in the case. Information provided in an individual question does not apply to the other questions in the case study.

A Review Screen will appear at the end of this case study. From the Review Screen, you can review and change your answers before you move to the next exam section. After you leave this case study, you will NOT be able to return to it.

To start the case study

To display the first question in this case study, select the "Next" button. To the left of the question, a menu provides links to information such as business requirements, the existing environment, and problem statements. Please read through all this information before answering any questions. When you are ready to answer a question, select the "Question" button to return to the question.

Background

Fabrikam Inc. is a mid-sized healthcare analytics company that provides population health dashboards and predictive insights to regional hospital systems across the United States. Fabrikam Inc. customers rely on near real time analytics to monitor patient flow, staffing needs, and readmission risks. They use multiple traditional forecasting machine learning models for predictions. Fabrikam Inc. has an established Microsoft Azure footprint. The company uses Jupyter Notebooks that run on a local server as the primary development environment. The data science team is experiencing scalability, asset management and code management issues with the current development platform. Fabrikam Inc. plans to migrate to a cloud-based development environment to mitigate the issues.

Additionally, the company plans to implement a Retrieval-Augmented Generation (RAG)-based chat application for client support.

Leadership requires the application to be developed and deployed with a low operational risk.

Current Environment

Fabrikam Inc. operates a single Azure subscription that has the following components:

Azure Data Lake Storage Gen2 that contains de-identified clinical and operational datasets

Azure AI Search indexing curated analytical documents and reference materials

A small set of Python-based training scripts maintained by data scientists Azure OpenAI Service with deployed foundational models

A Microsoft Foundry resource for building a RAG-based solution

Evaluation data has manually defined expected responses.

The current challenges faced by the data science team include the following:

Model training jobs are run manually from notebooks.

Experiment tracking is inconsistent

Model versions are registered without standardized metadata.

Deployment is performed manually by data scientists, with limited rollback capability.

The team has no standardized evaluation process for generative AI outputs.

The environment currently allows public network access. Authentication relies on user accounts rather than managed identities.

Compute targets are manually created and shared across experiments. This has led to resource contention during peak usage.

Business Requirements

Fabrikam Inc. has the following business requirements for the modernization initiative:

Provide a conversational interface that answers analytics questions by using internal documents and datasets.

Ensure that sensitive healthcare-related data is not exposed outside the Fabrikam Inc. Azure tenant.

Enable repeatable and auditable model training and deployment processes.

Support experimentation to compare prompt strategies and fine-tuned models.

Align the model with the ranked preferences and optimize behavior for the long term.

Minimize disruption to existing analytics workloads during rollout.

Technical Requirements

To support the business goals, Fabrikam Inc. identifies these technical requirements:

Use Azure Machine Learning workspaces to centrally manage data assets, models, and environments.

Implement experiment tracking and model versioning for all training jobs.

Orchestrate training and evaluation by using pipelines rather than manually running notebooks.

Deploy traditional machine learning models with support for staged rollout and rollback.

Improve RAG-based solution output quality.

Use the existing evaluation datasets that are based on real data with input-output pairs.

Apply advanced fine-tuning techniques only when prompt engineering is insufficient

Issues and Constraints

Fabrikam Inc. must comply with internal security policies that require the company to restrict network access and avoid long-lived secrets. The data science team has limited Azure DevOps experience, so solutions must favor managed services and automation over custom infrastructure.

Cost predictability is important. Leadership prefers serverless or managed compute options where possible but is willing to approve dedicated compute for stable production workloads.

Problem Statement

Fabrikam Inc. must design and implement an Azure-based AI operations solution that enables reliable training, evaluation, deployment, and iteration of generative AI models. The solution must support experimentation and gradual rollout while ensuring governance, security, and operational stability. The data science and platform teams must collaborate to deliver this solution by using Azure Machine

Learning and Microsoft Foundry capabilities.

You need to standardize how Fabrikam Inc. manages machine learning assets.

Which action should you perform first?

- A. Register assets in the Azure Machine Learning registry.
- B. Create a shared Azure Machine Learning workspace.
- C. Deploy a managed online endpoint.
- D. Create a new Microsoft Foundry project.

Answer: B

2. Case Study

This is a case study. Case studies are not timed separately from other exam sections. You can use as much exam time as you would like to complete each case study. However, there might be additional case studies or other exam sections. Manage your time to ensure that you can complete all the exam sections in the time provided. Pay attention to the Exam Progress at the top of the screen so you have sufficient time to complete any exam sections that follow this case study.

To answer the case study questions, you will need to reference information that is provided in the case. Case studies and associated questions might contain exhibits or other resources that provide more information about the scenario described in the case. Information provided in an individual question does not apply to the other questions in the case study.

A Review Screen will appear at the end of this case study. From the Review Screen, you can review and change your answers before you move to the next exam section. After you leave this case study, you will NOT be able to return to it.

To start the case study

To display the first question in this case study, select the "Next" button. To the left of the question, a menu provides links to information such as business requirements, the existing environment, and problem statements. Please read through all this information before answering any questions. When you are ready to answer a question, select the "Question" button to return to the question.

Background

Fabrikam Inc. is a mid-sized healthcare analytics company that provides population health dashboards and predictive insights to regional hospital systems across the United States. Fabrikam Inc. customers rely on near real time analytics to monitor patient flow, staffing needs, and readmission risks. They use multiple traditional forecasting machine learning models for predictions. Fabrikam Inc. has an established Microsoft Azure footprint. The company uses Jupyter Notebooks that run on a local server as the primary development environment. The data science team is experiencing scalability, asset management and code management issues with the current development platform. Fabrikam Inc. plans to migrate to a cloud-based development environment to mitigate the issues.

Additionally, the company plans to implement a Retrieval-Augmented Generation (RAG)-based chat application for client support.

Leadership requires the application to be developed and deployed with a low operational risk.

Current Environment

Fabrikam Inc. operates a single Azure subscription that has the following components:

Azure Data Lake Storage Gen2 that contains de-identified clinical and operational datasets

Azure AI Search indexing curated analytical documents and reference materials

A small set of Python-based training scripts maintained by data scientists Azure OpenAI Service with deployed foundational models

A Microsoft Foundry resource for building a RAG-based solution

Evaluation data has manually defined expected responses.

The current challenges faced by the data science team include the following:

Model training jobs are run manually from notebooks.

Experiment tracking is inconsistent

Model versions are registered without standardized metadata.

Deployment is performed manually by data scientists, with limited rollback capability.

The team has no standardized evaluation process for generative AI outputs.

The environment currently allows public network access. Authentication relies on user accounts rather than managed identities.

Compute targets are manually created and shared across experiments. This has led to resource contention during peak usage.

Business Requirements

Fabrikam Inc. has the following business requirements for the modernization initiative:

Provide a conversational interface that answers analytics questions by using internal documents and datasets.

Ensure that sensitive healthcare-related data is not exposed outside the Fabrikam Inc. Azure tenant.

Enable repeatable and auditable model training and deployment processes.

Support experimentation to compare prompt strategies and fine-tuned models.

Align the model with the ranked preferences and optimize behavior for the long term.

Minimize disruption to existing analytics workloads during rollout.

Technical Requirements

To support the business goals, Fabrikam Inc. identifies these technical requirements:

Use Azure Machine Learning workspaces to centrally manage data assets, models, and environments.

Implement experiment tracking and model versioning for all training jobs.

Orchestrate training and evaluation by using pipelines rather than manually running notebooks.

Deploy traditional machine learning models with support for staged rollout and rollback.

Improve RAG-based solution output quality.

Use the existing evaluation datasets that are based on real data with input-output pairs.

Apply advanced fine-tuning techniques only when prompt engineering is insufficient

Issues and Constraints

Fabrikam Inc. must comply with internal security policies that require the company to restrict network

access and avoid long-lived secrets. The data science team has limited Azure DevOps experience, so solutions must favor managed services and automation over custom infrastructure. Cost predictability is important. Leadership prefers serverless or managed compute options where possible but is willing to approve dedicated compute for stable production workloads.

Problem Statement

Fabrikam Inc. must design and implement an Azure-based AI operations solution that enables reliable training, evaluation, deployment, and iteration of generative AI models. The solution must support experimentation and gradual rollout while ensuring governance, security, and operational stability. The data science and platform teams must collaborate to deliver this solution by using Azure Machine Learning and Microsoft Foundry capabilities.

You need to isolate training workloads while remaining cost-aware to address Fabrikam Inc.'s issues, constraints, and technical requirements.

What should you implement?

- A. Training jobs that run on a single shared compute cluster
- B. Fixed-size compute cluster
- C. Dedicated compute clusters per experiment
- D. Managed compute targets with autoscaling

Answer: D

3.HOTSPOT -

A team trains an MLflow model that scores customer churn risk. The model will be consumed by different downstream systems.

One system requests predictions synchronously during customer interactions.

Another system submits files containing millions of records for scheduled scoring.

You need to deploy the model by using managed inference options that match each usage pattern.

Which option should you use for each usage pattern? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Managed inference options

Requirement

Low-latency synchronous predictions

Deployment option

Batch endpoint
Job-based training pipeline
Real-time endpoint
Registered model artifact

High-volume scheduled scoring

Batch endpoint
Online endpoint with autoscaling
Managed compute cluster
Model registry version

Answer:

Managed inference options

Requirement

Low-latency synchronous predictions

Deployment option

Batch endpoint
Job-based training pipeline
Real-time endpoint
Registered model artifact

High-volume scheduled scoring

Batch endpoint
Online endpoint with autoscaling
Managed compute cluster
Model registry version

4. You manage an Azure Machine learning workspace. You develop a machine learning model. You must deploy the model to use a low-priority VM with a pricing discount. You need to deploy the model. Which compute target should you use?

-
- A. Azure Container Instances (ACI)
 - B. Azure Machine Learning compute clusters
 - C. Local deployment
 - D. Azure Kubernetes Service (AKS)

Answer: B

5. A team manages an Azure Machine Learning workspace where they deploy models to online endpoints.

The team needs to introduce a new version of a model to production without disrupting existing users.

The team must validate the new version before full rollout.

You need to reduce risk during deployment.

What should you do?

- A. Deploy the model to a batch endpoint.
- B. Split traffic between deployments.
- C. Replace the existing endpoint.
- D. Route all traffic to the new deployment.

Answer: B