

Microsoft AI-300

Operationalizing Machine Learning and Generative AI Solutions

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/ai-300>

Latest Version: 6.0

Question: 1

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

Question: 2

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

Question: 3

You need to recommend an experiment-tracking strategy that ensures consistent experiment results.

What should you recommend?

- A. Azure Machine Learning job output logs
- B. MLflow experiment tracking

- C. Application Insights logs
- D. Azure Monitor alerts

Answer: B

Question: 4

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.

a. 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

Question: 5

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

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>