Question #1 *Topic 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 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.

Suggested Answer: *B*

Community vote distribution

03f65a8 1 month ago

Selected Answer: B

Request is: ML asset management standardization

upvoted 1 times

uncledana 2 months, 1 week ago

Selected Answer: B

https://learn.microsoft.com/en-us/azure/machine-learning/concept-workspace?view=azureml-api-2&utm_source=chatgpt.com

upvoted 2 times

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

Suggested Answer: *D*

Community vote distribution

03f65a8 1 month ago

Selected Answer: D

- workload isolation
- autoscaling
- managed service

upvoted 1 times

Question #3 *Topic 2*

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.

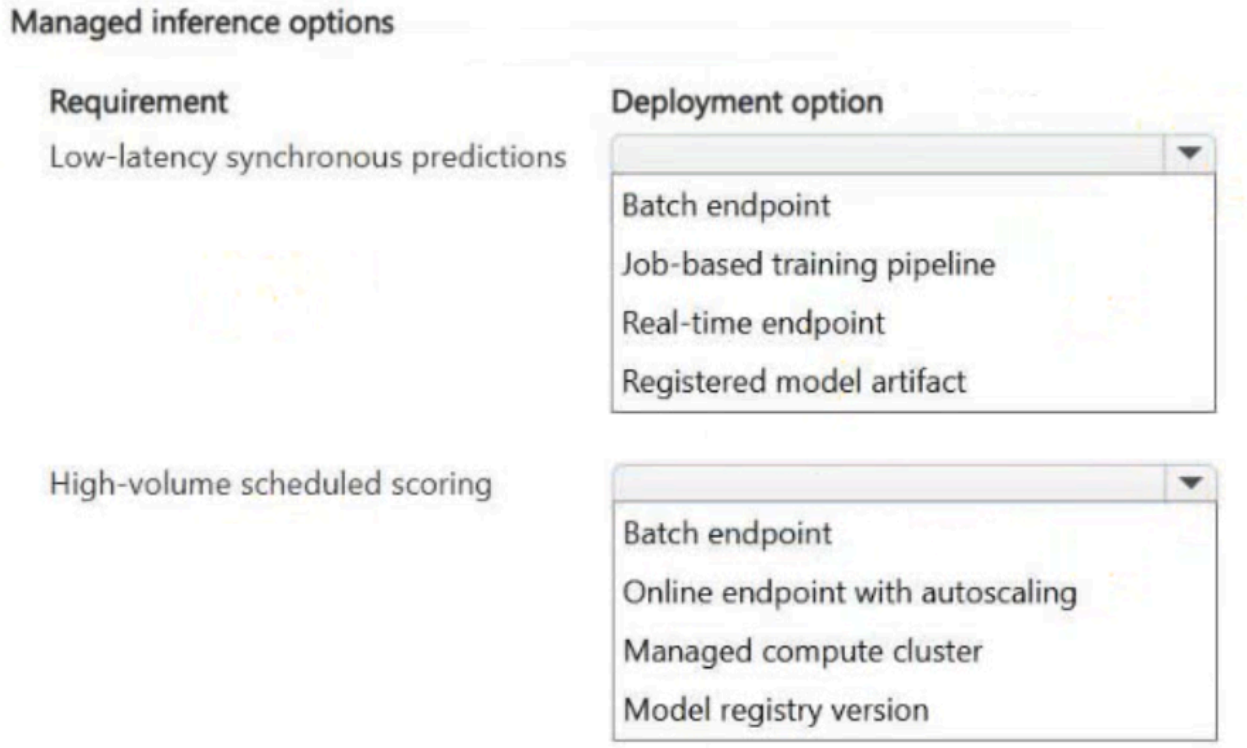

Suggested Answer:

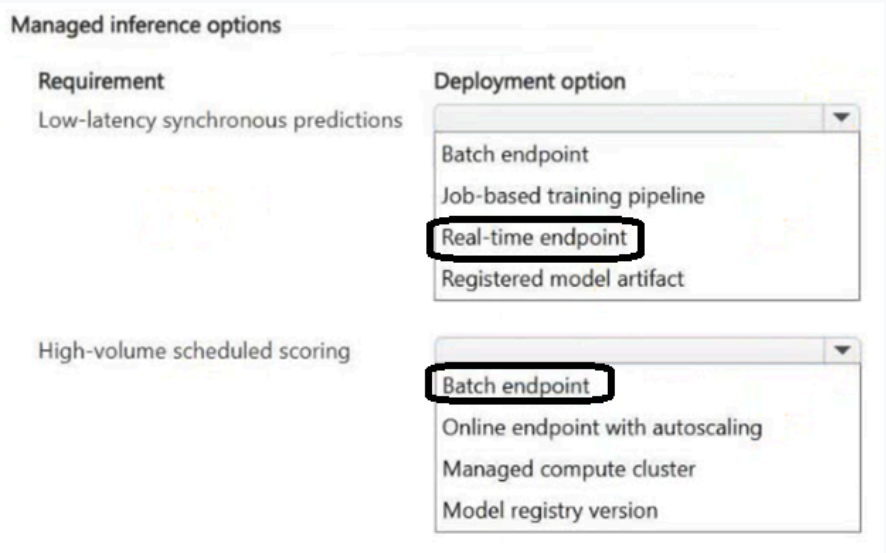

uncledana 2 months, 1 week ago

Answer is correct - low latency requires a real time endpoint, for schedules always use a batch endpoint

upvoted 3 times

Question #4 *Topic 2*

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)

Suggested Answer: *B*

Community vote distribution

B (100%)

Peeking 1 week, 5 days ago

Selected Answer: B

Azure batch deployments support low priority virtual machines (VMs) to reduce the cost of batch inference workloads. Low priority VMs enable a large amount of compute power to be used for a low cost. Low priority virtual machines take advantage of surplus capacity in Azure. When you specify low priority VMs in your pools, Azure can use this surplus, when available.

```
from azure.ai.ml.entities import AmlCompute

compute_name = "low-pri-cluster"
compute_cluster = AmlCompute(
name=compute_name,
description="Low priority compute cluster",
min_instances=0,
max_instances=2,
tier='LowPriority'
)

ml_client.begin_create_or_update(compute_cluster)
```

upvoted 1 times

Question #5 *Topic 2*

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.

Suggested Answer: *B*

Community vote distribution

ugwuozorcollinsemezie 4 weeks, 1 day ago

Selected Answer: B

Blue/Green deployment.

upvoted 2 times

Peeking 1 month ago

Selected Answer: B

Some percent and gradually increase traffic to the new model.

upvoted 1 times

uncledana 2 months, 1 week ago

Selected Answer: B

Feature flag - classic

upvoted 1 times

Question #6 *Topic 2*

You have a deployment of an Azure OpenAI Service base model.
You plan to fine-tune the model.
You need to prepare a file that contains training data.
Which file format should you use?

A. CSV

B. TSV

C. JSONL

D. JSON

Suggested Answer: *C*

Community vote distribution

C (100%)

snashe 1 month, 3 weeks ago

Selected Answer: C

JSONL is best option.

upvoted 3 times

Question #7 *Topic 2*

You have a deployment of an Azure OpenAI Service base model.
You plan to fine-tune the model.
You need to prepare a file that contains training data for multi-turn chat.
Which file encoding method should you use?

A. ISO-8859-1

B. UTF-16

C. UTF-8

D. ASCII

Suggested Answer: *C*

Community vote distribution

C (100%)

uncledana 2 months, 1 week ago

Selected Answer: C

standard for JSONL

upvoted 2 times

Question #8 *Topic 2*

You are fine-tuning a base language model to analyze customer feedback.
You label examples of support tickets. You must improve classification accuracy by configuring and fine-tuning the base model in Microsoft Foundry.
You need to configure and run fine-tuning.
What should you do first?

A. Use prompt flow to generate multiple prompt templates for evaluation.

B. Deploy the base model to an online endpoint before starting fine-tuning.

C. Enable tracing for all inference calls in the evaluation pipeline.

D. Format the dataset as a JSONL file with prompt-completion pairs and upload the file.

Suggested Answer: *D*

Community vote distribution

D (100%)

Peeking 1 week, 5 days ago

Selected Answer: D

You should prepare the labelled training and validation dataset first.

For Microsoft Foundry fine-tuning, the first step is to format your labelled support-ticket examples as JSONL in the chat/completions conversational format, where each example maps the customer feedback or ticket text to the expected classification label.

upvoted 1 times

03f65a8 1 month ago

Selected Answer: D

1st step is dataset formating

upvoted 1 times

Peeking 1 month ago

Selected Answer: D

Data preparation is first in every model fine tuning

upvoted 2 times

cayenne06 1 month, 3 weeks ago

Selected Answer: D

This is the first step for fine-tuning

upvoted 2 times

uncledana 1 month, 3 weeks ago

Selected Answer: D

D - https://learn.microsoft.com/en-us/azure/foundry/openai/how-to/fine-tuning?tabs=oai-sdk&pivots=programming-language-python

upvoted 2 times

Question #9 *Topic 2*

A team is working in Microsoft Foundry to test and compare large language model (LLM) prompt variants in a development environment.
The team requires consistent inputs to evaluate prompt variants without relying on live user traffic.
You need to create a controlled evaluation of input data.
Which action should you perform first?

A. Generate synthetic interaction data.

B. Configure content filters.

C. Apply a blocklist.

D. Enable observability metrics.

Suggested Answer: *A*

Community vote distribution

A (100%)

naim149 1 day, 21 hours ago

Selected Answer: A

The team needs consistent, controlled input data to compare prompt variants without live traffic. Microsoft Foundry evaluation supports synthetic/simulated data to test applications in controlled scenarios before deployment.

upvoted 1 times

Question #10 *Topic 2*

DRAG DROP -

A team maintains Infrastructure as Code (IaC) templates to provision Azure Machine Learning resources.

Provisioning must be triggered by changes in the templates and executed without manual intervention.

You need to automate resource provisioning.

Which action should you take for each requirement? To answer, move the appropriate actions to the correct requirements. You may use each action once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

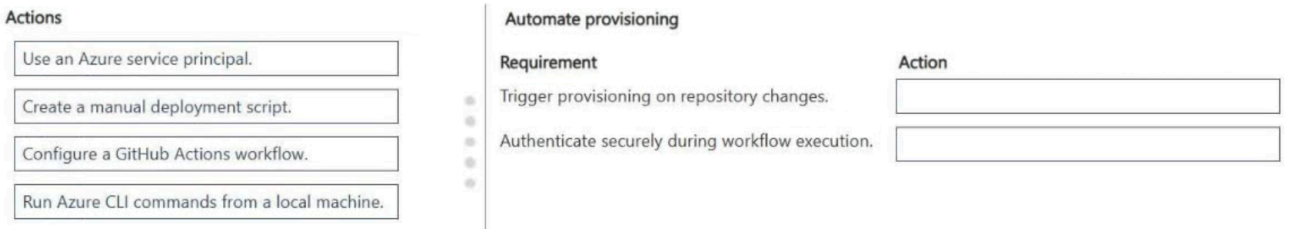

Suggested Answer:

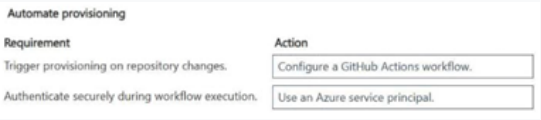

naim149 1 day, 21 hours ago

Trigger provisioning on repository changes. Configure a GitHub Actions workflow.

Authenticate securely during workflow execution. Use an Azure service principal.

upvoted 1 times

Question #11 *Topic 2*

An organization maintains separate Azure Machine Learning workspaces for development and production.
Both environments must use the same validated assets without duplicating them.
Assets must be shared across workspaces while maintaining centralized governance and version control.
You need to enable reuse of assets across workspaces without copying them.
What should you do?

A. Enable workspace-level Git integration and sync assets between repositories.

B. Publish the asset as a pipeline component.

C. Create a shared Azure Machine Learning environment that includes the asset.

D. Publish the asset to an Azure Machine Learning registry.

Suggested Answer: *D*

Community vote distribution

naim149 1 day, 21 hours ago

Selected Answer: D

Azure Machine Learning registries are designed to share validated assets across multiple workspaces, including dev/test/prod, while keeping centralized governance and version control. Microsoft describes registries as decoupling ML assets from individual workspaces and making assets such as models, environments, components, and datasets available across the organization.

upvoted 1 times

Question #12 *Topic 2*

An Azure Machine Learning workspace processes sensitive training data.
The workspace must NOT be accessible from the public internet.
You need to restrict network access.
Which configuration should you implement?

A. Azure Firewall rules

B. Private endpoints

C. Network security groups

D. Service endpoints

Suggested Answer: *B*

Community vote distribution

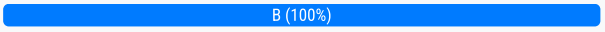

uncledana 2 months ago

Selected Answer: B

Private is the way

upvoted 3 times

Question #13 *Topic 2*

A team is experimenting with traditional models for a classification workflow in Azure Machine Learning.
The team requires a consistent way to manage assets that are created during experimentation.
You need to ensure that artifacts can be reused and governed across projects.
Which asset should you register?

A. Model

B. Component

C. Environment

D. Pipeline

Suggested Answer: *A*

Community vote distribution

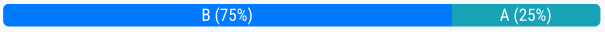

03f65a8 1 month ago

Selected Answer: A

Model - the most important reuseable artifact in AML. (and the only artifact in the list :))

upvoted 1 times

rveney 1 month, 2 weeks ago

Selected Answer: B

Correct answer: B. Component

If the team wants a consistent, reusable, and governed way to manage assets created during experimentation, the right asset to register is a component.

upvoted 3 times

Question #14 *Topic 2*

HOTSPOT -

You manage an Azure Machine Learning workspace named workspace1 by using the Python SDK v2.

The default datastore of workspace1 contains a folder named sample_data. The folder structure contains the following content:

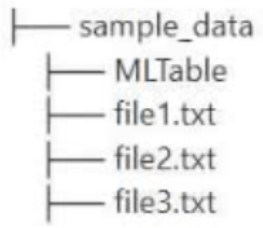

You write Python SDK v2 code to materialize the data from the files in the sample_data folder into a Pandas data frame.

You need to complete the Python SDK v2 code to use the MLTable folder as the materialization blueprint.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point

Answer Area

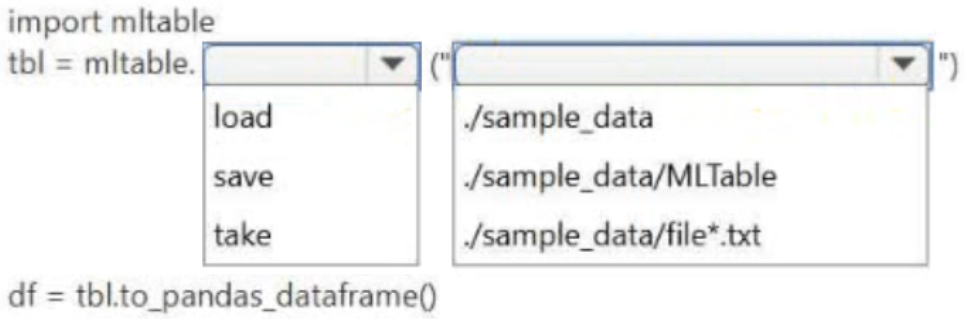

Suggested Answer:

Answer Area

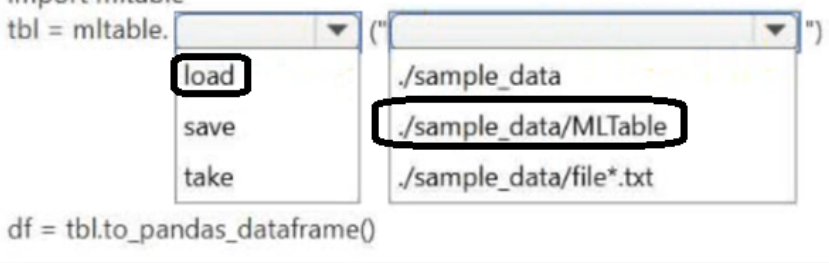

03f65a8 1 month ago

I disagree.

mltable.load() expects the folder that contains the relevant files, not the MLTable file itself.

upvoted 1 times

uncledana 2 months ago

correct

upvoted 2 times

Question #15 *Topic 2*

HOTSPOT -

You manage an Azure Machine Learning workspace named workspace1 by using the Python SDK v2. You create a General Purpose v2 Azure storage account named mlstorage1. The storage account includes a publicly accessible container named mlcontainer1. The container stores 10 blobs with files in the CSV format.

You must develop Python SDK v2 code to create a data asset referencing all blobs in the container named mlcontainer1.

You need to complete the Python SDK v2 code.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

```
sample_dataset = Data(
    path=[ wasbs / abfss / azureml ]://mlstorage1.blob.core.windows.net/mlcontainer1,
    type=AssetTypes.[ URI_FILE / URI_FOLDER / URI_MLTABLE ],
    description="sample_dataset",
    name="sample_dataset",
    version='1.0'
)
```

Suggested Answer:

Answer Area

```
sample_dataset = Data(
    path=[ wasbs / abfss / azureml ]://mlstorage1.blob.core.windows.net/mlcontainer1,
    type=AssetTypes.[ URI_FILE / URI_FOLDER / URI_MLTABLE ],
    description="sample_dataset",
    name="sample_dataset",
    version='1.0'
)
```

rveney 1 month, 2 weeks ago

Path: wasbs

Type: URI_FOLDER

upvoted 2 times

cayenne06 1 month, 3 weeks ago

I think it's wasbs for blob storage access. Abfss is for Data Lake Gen 2 and azureml for internal storage of azure ml.

upvoted 1 times

Question #16 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear on the review screen.
You manage an Azure Machine Learning workspace. The Python script named script.py reads an argument named training_data. The training_data argument specifies the path to the training data in a file named dataset1.csv.
You plan to run the script.py Python script as a command job that trains a machine learning model.
You need to provide the command to pass the path for the dataset as a parameter value when you submit the script as a training job.
Solution: python train.py --training_data training_data
Does the solution meet the goal?

A. Yes

B. No

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #17 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear on the review screen.
You manage an Azure Machine Learning workspace. The Python script named script.py reads an argument named training_data. The training_data argument specifies the path to the training data in a file named dataset1.csv.
You plan to run the script.py Python script as a command job that trains a machine learning model.
You need to provide the command to pass the path for the dataset as a parameter value when you submit the script as a training job.
Solution: python script.py --trainingdata ${{inputs.training_data}}
Does the solution meet the goal?

A. Yes

B. No

Suggested Answer: *A*

Community vote distribution

B (100%)

uncledana 2 months ago

Selected Answer: B

See the typo - should match exactly as the parameter: python script.py --training_data ${{inputs.training_data}}

upvoted 2 times

Question #18 *Topic 2*

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear on the review screen.
You manage an Azure Machine Learning workspace. The Python script named script.py reads an argument named training_data. The training_data argument specifies the path to the training data in a file named dataset1.csv.
You plan to run the script.py Python script as a command job that trains a machine learning model.
You need to provide the command to pass the path for the dataset as a parameter value when you submit the script as a training job.
Solution: python script.py dataset1.csv
Does the solution meet the goal?

A. Yes

B. No

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #1 *Topic 3*

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 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

Suggested Answer: *B*

Currently there are no comments in this discussion, be the first to comment!

Question #2 *Topic 3*

HOTSPOT -

A machine learning model is deployed to production in Azure Machine Learning and is actively serving predictions for a business application. The model was trained by using a historical dataset that represented expected input patterns at the time of deployment.

The team working on the model must ensure the following:

Changes in input data distribution are detected.

Appropriate actions are triggered when predefined thresholds are exceeded.

You need to configure monitoring to meet the requirements.

Which configuration should you use for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

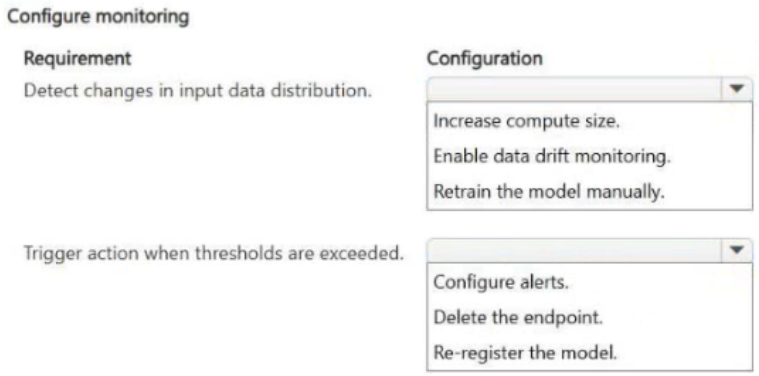

Suggested Answer:

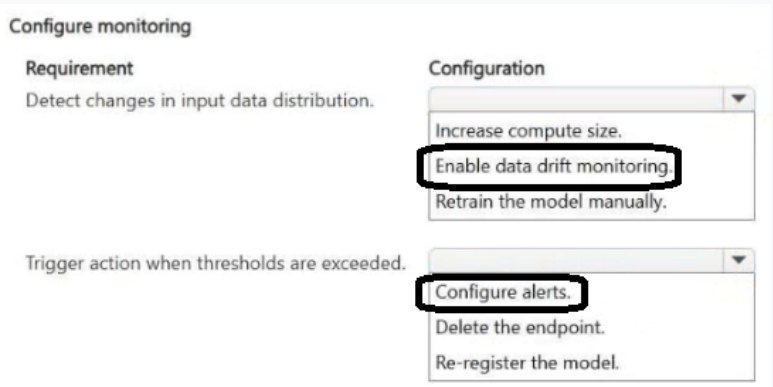

Currently there are no comments in this discussion, be the first to comment!

Question #3 *Topic 3*

You create a binary classification model. You use the Fairlearn package to assess model fairness.
You must eliminate the need to retrain the model.
You need to implement the Fairlearn package.
Which algorithm should you use?

A. fairiearn.reductions.ExponentiatedGradient

B. fairlearn.preprocessing.CorrelationRemover

C. fairlearn.reductions.GridSearch

D. fairlearn.postprocessing.ThresholdOptimizer

Suggested Answer: *D*

Currently there are no comments in this discussion, be the first to comment!

Question #4 *Topic 3*

DRAG DROP -

A real-time endpoint is deployed in Azure Machine Learning to serve predictions to a web application.

Users report intermittent failures and unexpected responses when calling the endpoint.

You need to identify the appropriate troubleshooting action for each reported issue.

Which troubleshooting action should you perform for each issue? To answer, move the appropriate troubleshooting actions to the correct issues. You may use each troubleshooting action once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

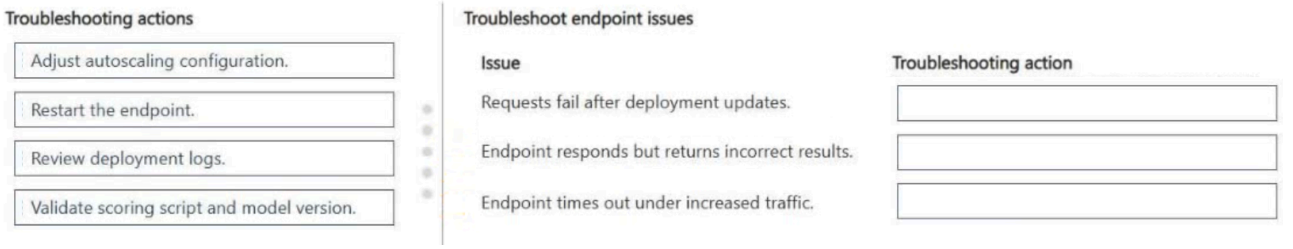

Suggested Answer:

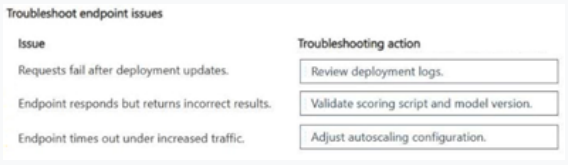

Currently there are no comments in this discussion, be the first to comment!

Question #5 *Topic 3*

A team manages an Azure Machine Learning workspace and deploys a model to an endpoint.
A deployed online endpoint shows inconsistent response times during periods of high traffic.
You need to identify potential performance degradation.
Which three metrics should you monitor? Each correct answer presents part of the solution. Choose three.
NOTE: Each correct selection is worth one point.

A. Feature count

B. Requests per minute

C. Connections active

D. Dataset size

E. Request latency

Suggested Answer: *BCE*

Currently there are no comments in this discussion, be the first to comment!

Question #6 *Topic 3*

DRAG DROP -

A team deploys a machine learning model to a managed online endpoint. The team monitors model performance and data quality metrics in production.

When monitoring thresholds are exceeded, the team requires an automated operational response that notifies downstream systems.

You need to configure the monitoring solution to meet the requirements.

Which configuration should you associate with each requirement as a first step? To answer, move the appropriate configurations to the correct requirements. You may use each configuration once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

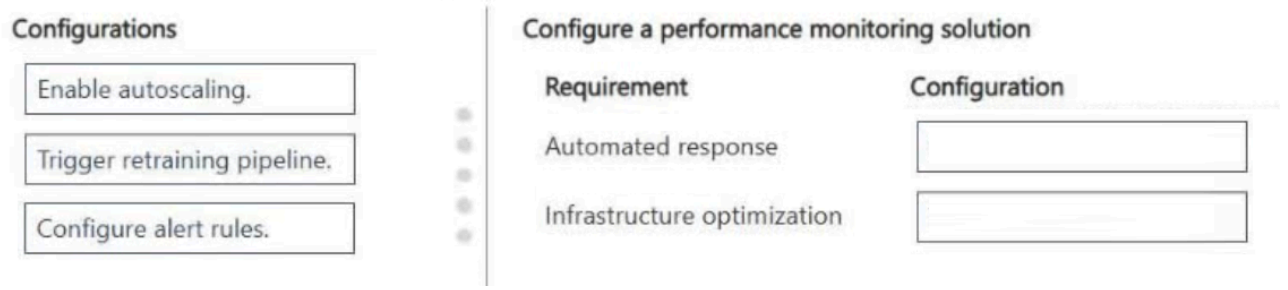

Suggested Answer:

Configure a performance monitoring solution

Requirement	Configuration
Automated response	Configure alert rules.
Infrastructure optimization	Trigger retraining pipeline.

03f65a8 1 month ago

Configure alert rules

Enable autoscaling

upvoted 1 times

cayenne06 1 month, 2 weeks ago

the right choice is Enable autoscaling, because it automatically adjusts capacity based on traffic/load. Trigger retraining pipeline is for improving model quality after drift or performance degradation, not for optimizing endpoint infrastructure.

upvoted 1 times

uncledana 2 months ago

Configure alert rules

Enable autoscaling

upvoted 2 times

Question #7 *Topic 3*

HOTSPOT -

You use Azure Machine Learning to train models across multiple experiments by using the same workspace.

You must record training runs in a centralized location to compare results from different jobs.

During training, performance values must be captured so they appear in the experiment run history.

You need to configure experiment tracking.

What should you configure for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

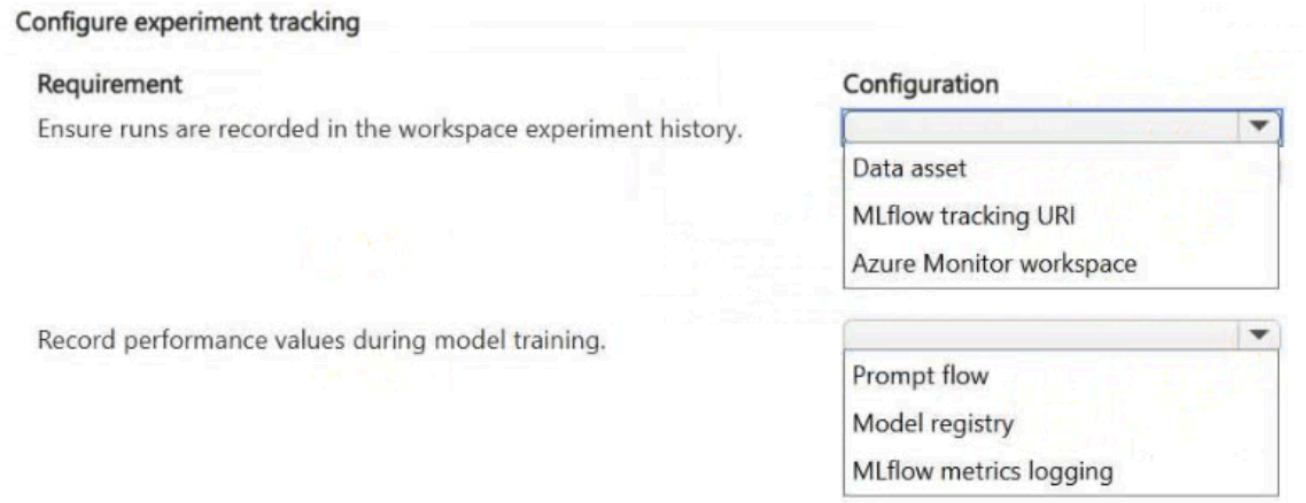

Suggested Answer:

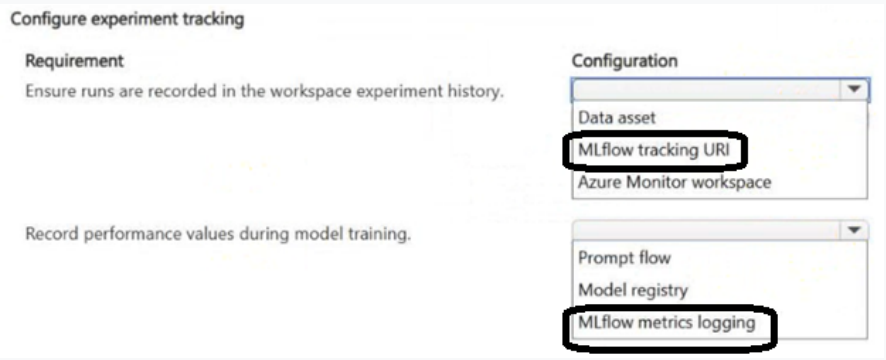

Currently there are no comments in this discussion, be the first to comment!

Question #8 *Topic 3*

A data science team trains a classification model that predicts loan approval outcomes.
Before registering the model, the team must ensure the following:
Predictions must not disproportionately impact protected groups.
Prediction errors can be evaluated across different data segments.
You need to assess whether the model meets Responsible AI expectations.
Which two approaches should you use? Each correct answer presents part of the solution. Choose two.
NOTE: Each correct selection is worth one point.

A. Analyze error rates across the global cohort.

B. Measure endpoint latency under load.

C. Validate inference schema compatibility.

D. Evaluate feature importance for prediction transparency.

E. Analyze error rates across defined demographic cohorts.

Suggested Answer: *DE*

Currently there are no comments in this discussion, be the first to comment!

Question #9 *Topic 3*

A team deploys a model to a real-time endpoint in Azure Machine Learning. You deploy some updates to the endpoint.
The endpoint returns errors after the new deployment is released.
You need to restore the service as quickly as possible.
What should you do first?

A. Roll back traffic to the previous deployment.

B. Delete the endpoint and immediately redeploy it.

C. Change the authentication type to Azure Machine Learning token-based authentication.

D. Increase the compute size.

Suggested Answer: *A*

Currently there are no comments in this discussion, be the first to comment!

Question #10 *Topic 3*

HOTSPOT -

You train a model in Azure Machine Learning.

You plan to capture experiment details for later comparison. The training code must log parameters and metrics for each run.

You review the following training script.

```
import mlflow
import mlflow.sklearn
from sklearn.linear_model import LogisticRegression
from sklearn.datasets import load_iris

X, y = load_iris(return_X_y=True)

mlflow.set_experiment("classification-experiment")

with mlflow.start_run():
    model = LogisticRegression(max_iter=200)
    model.fit(X, y)

    mlflow.log_param("max_iter", 200)
    y_pred = model.predict(X)
    accuracy = accuracy_score(y, y_pred)
    mlflow.log_metric("accuracy", accuracy)

    mlflow.sklearn.log_model(model, "model")
```

You need to verify whether the training script meets the experiment tracking requirement. For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

MLflow experiment tracking

Statement	Yes	No
The code ensures that training parameters and metrics are tracked for each run.	○	○
The code uses the experiment name to define the artifact path of the logged model.	○	○
The code associates the run with a named experiment.	○	○

Suggested Answer:

MLflow experiment tracking

Statement	Yes	No
The code ensures that training parameters and metrics are tracked for each run.	●	○
The code uses the experiment name to define the artifact path of the logged model.	●	○
The code associates the run with a named experiment.	●	○

03f65a8 4 weeks, 1 day ago

Yes No Yes:
mlflow.set_experiment("classification-experiment")
--> experiment association, not artefact path definition

mlflow.sklearn.log_model(model, "model")
--> "model" is the artefact path.
upvoted 1 times

uncledana 1 month, 3 weeks ago
YNY - makes sense, no experiment name is classification-experiment but artefact name is 'model'
upvoted 2 times

Question #11 *Topic 3*

A data science team completes multiple training runs within an experiment by using MLflow.
The team wants to store a selected model in Azure Machine Learning so that it can be versioned and deployed later.
The model must be versioned centrally for reuse across environments.
You need to version the trained model.
Which two actions should you perform? Each correct answer presents part of the solution. Choose two.
NOTE: Each correct selection is worth one point.

A. Locate and capture the model artifacts from the outputs of the training run.

B. Register the model in the Azure Machine Learning workspace.

C. Tag the training experiment with a name.

D. Export the model files to local storage.

Suggested Answer: *AB*

Currently there are no comments in this discussion, be the first to comment!