

AI-300 Latest Braindumps Sheet, Reliable AI-300 Exam Pattern



Research has found that stimulating interest in learning may be the best solution. Therefore, the AI-300 prepare guide' focus is to reform the rigid and useless memory mode by changing the way in which the AI-300 exams are prepared. Our Soft version of AI-300 practice materials combine knowledge with the latest technology to greatly stimulate your learning power. By simulating enjoyable learning scenes and vivid explanations, users will have greater confidence in passing the qualifying AI-300 exams.

If you want to make one thing perfect and professional, then the first step is that you have to find the people who are good at them. In this AI-300 exam braindumps field, our experts are the core value and truly helpful with the greatest skills. So our AI-300 practice materials are perfect paragon in this industry full of elucidating content for exam candidates of various degrees to use for reference. Just come to buy our AI-300 study guide!

>> AI-300 Latest Braindumps Sheet <<

Pass Guaranteed Quiz Microsoft - AI-300 - Operationalizing Machine Learning and Generative AI Solutions –Reliable Latest Braindumps Sheet

We are a team of IT professionals that provide our customers with the up-to-date AI-300 study guide and the current certification exam information. Our exam collection contains the latest questions, accurate AI-300 Exam Answers and some detailed explanations. You will find everything you want to overcome the difficulties of AI-300 practice exam and questions. You will get high

mark followed by our materials.

Microsoft Operationalizing Machine Learning and Generative AI Solutions Sample Questions (Q75-Q80):

NEW QUESTION # 75

A financial services company is deploying Microsoft Foundry to host generative AI workloads that process regulated customer data. The Microsoft Foundry environment must prevent any public network exposure while still allowing services managed by Microsoft Foundry to communicate with dependent Azure resources.

Security auditors require that all traffic to and from the Microsoft Foundry resource remain on private networks, with no public endpoints available.

You need to configure the Microsoft Foundry environment so that network access is restricted while maintaining full platform functionality.

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. Disable all inbound network access.
- B. Use API key authentication for all model endpoints.
- C. Disable public network access to the Microsoft Foundry resource.
- D. Deploy the Microsoft Foundry resource in a separate Azure subscription.
- E. Configure a managed virtual network for the Microsoft Foundry resource.

Answer: A,E

Explanation:

To host generative AI workloads in a Microsoft Foundry environment with strictly private communication and no public network exposure, you must configure a Managed Virtual Network (Managed VNet) with specific isolation settings and disable all public inbound access.

[A]

Enable Managed Virtual Network Isolation

During the creation of your Azure AI Foundry hub, navigate to the Networking tab.

Select the Private with Approved Outbound isolation mode. This mode ensures that all outbound traffic from the managed compute resources is restricted to only the destinations you explicitly approve, such as dependent Azure resources.

Once enabled, this isolation mode cannot be disabled.

[E]

Disable Public Inbound Access

In the Networking tab of your Foundry resource, set Public network access to Disabled.

This action blocks all traffic from the public internet, ensuring the resource is only accessible through private connections.

Reference:

<https://learn.microsoft.com/en-us/azure/foundry/how-to/managed-virtual-network>

NEW QUESTION # 76

Hotspot Question

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	Deployment option
Low-latency synchronous predictions	<ul style="list-style-type: none"> Batch endpoint Job-based training pipeline Real-time endpoint Registered model artifact
High-volume scheduled scoring	<ul style="list-style-type: none"> Batch endpoint Online endpoint with autoscaling Managed compute cluster Model registry version

Answer:

Explanation:

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



NEW QUESTION # 77

Drag and Drop Question

A company is standardizing generative AI development across multiple teams.

Each team requires an isolated workspace. Governance and shared connections must be centrally managed.

You need to implement a Microsoft Foundry environment structure that supports centralized governance and team isolation.

Which type of configuration should you use for each requirement? 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.

Configurations

- A shared Microsoft Foundry project
- A shared Microsoft Foundry resource
- A Microsoft Foundry project per team
- A Microsoft Foundry resource per team

Workload configuration options

Requirement

- Isolated workspace for each team
- Security and governance across teams

Configuration

-
-

Answer:

Explanation:

Configurations

- A shared Microsoft Foundry project
- A shared Microsoft Foundry resource
- A Microsoft Foundry project per team
- A Microsoft Foundry resource per team

Workload configuration options

Requirement

- Isolated workspace for each team
- Security and governance across teams

Configuration

- A Microsoft Foundry project per team
- A shared Microsoft Foundry resource

NEW QUESTION # 78

Your model requires access to external APIs using sensitive credentials during inference. You must ensure credentials are not exposed in code, logs, or environment variables. What should you implement?

- A. Hardcode credentials
- B. Encrypt credentials locally
- C. Store in config files
- D. Use Azure Key Vault with managed identity

Answer: D

Explanation:

Azure Key Vault with managed identity ensures secure access to sensitive credentials without exposing them in code or configuration. Managed identities eliminate the need for hardcoded secrets. Other approaches, such as config files or environment variables, increase the risk of accidental exposure.

NEW QUESTION # 79

You must ensure full reproducibility of experiments including dataset, code, and environment across multiple runs and workspaces. Which combination of practices is MOST appropriate?

- A. Logging metrics only
- B. Git only
- C. Dataset versioning only
- D. Environment + dataset + code versioning

Answer: D

Explanation:

Reproducibility requires versioning all components: code, datasets, and environments. Missing any of these elements prevents exact replication of experiments. For example, the same code with different data or dependencies can produce different results, making debugging and auditing difficult.

NEW QUESTION # 80

