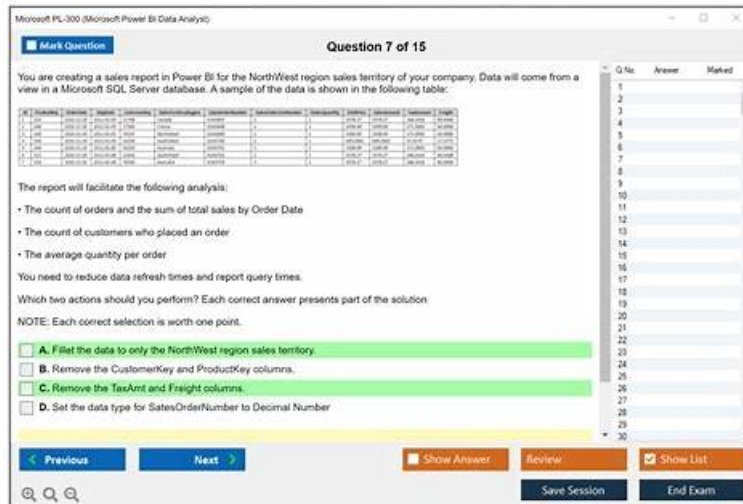


Free PDF Perfect Microsoft - Practice AI-300 Exam Online



The latest AI-300 exam prep is created by our IT experts and certified trainers who are dedicated to Microsoft braindumps pdf for a long time. All questions of our AI-300 PDF VCE are written based on the real questions. Besides, we always check the updating of AI-300 exam questions to make sure exam preparation smoothly.

Success in the Operationalizing Machine Learning and Generative AI Solutions (AI-300) certification exam helps people update their skills. Many aspirants don't find updated Microsoft AI-300 practice test questions and fail the final test. This failure in the Microsoft AI-300 Exam leads to a loss of money and time. If you are also planning to attempt the Operationalizing Machine Learning and Generative AI Solutions (AI-300) exam and are confused about where to prepare yourself for it then you are at the right place.

>> Practice AI-300 Exam Online <<

Marvelous Practice AI-300 Exam Online - Unparalleled Source of AI-300 Exam

The crucial thing when it comes to appearing a competitive exam like AI-300 knowing your problem-solving skills. And to do that you are going to need help from a AI-300 practice questions or braindumps. This is exactly what is delivered by our AI-300 test materials. The AI-300 Exam Dumps cover every topic of the actual Microsoft certification exam. The AI-300 exam questions are divided into various groups and the candidate can solve these questions to test his skills and knowledge.

Microsoft Operationalizing Machine Learning and Generative AI Solutions Sample Questions (Q16-Q21):

NEW QUESTION # 16

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. Configure a managed virtual network for the Microsoft Foundry resource.

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

Answer: A,C

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 # 17

Drag and Drop Question

An organization operates a generative AI application in production by using Microsoft Foundry.

The application serves live user traffic and is updated by a data scientist team regularly as prompts and models evolve.

The application intermittently times out during production use, which requires ongoing visibility into runtime behavior.

The team must also validate model quality and safety before releasing new updates to avoid introducing regressions.

You need to apply the correct mechanisms for continuous runtime monitoring and for release time validation.

Which mechanisms should you use for each requirement? To answer, move the appropriate mechanisms to the correct requirements.

You may use each mechanism 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.

Answer:

Explanation:

NEW QUESTION # 18

You need to run large-scale inference jobs on millions of records periodically. Jobs are not latency-sensitive but must be cost-efficient and scalable. Which deployment option is MOST appropriate?

- **A. Batch endpoint**
- B. Notebook execution
- C. Local endpoint
- D. Managed online endpoint

Answer: A

Explanation:

Batch endpoints are optimized for large-scale, asynchronous inference workloads. They efficiently process large datasets and scale based on demand, making them cost-effective for non-real-time scenarios. Online endpoints are designed for low-latency use cases and are more expensive for batch processing.

NEW QUESTION # 19

Drag and Drop Question

A team deploys a generative AI application that uses a model deployed in Microsoft Foundry. The application must support latency monitoring under production load.

You need to enable performance observability.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Answer:

Explanation:

NEW QUESTION # 20

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 work in Microsoft Foundry with a prompt flow.

You must manually evaluate prompts and compare results across prompt variants.

You need to capture the inputs, outputs, token usage, and latencies for each flow run for the evaluation.

Solution: Create prompt variants and compare their outputs in the Evaluation experience.

Does the solution meet the goal?

- A. No
- B. Yes

Answer: A

Explanation:

Correct:

* In Microsoft Foundry, turn on Tracing for the prompt flow of the project and execute test runs to produce trace data.

Incorrect:

* Create prompt variants and compare their outputs in the Evaluation experience.

* Use the prompt flow SDK to enable tracing for the flow before executing runs. Then run the flow to generate traceable results.

Note:

In Azure AI Foundry, you can capture and compare these metrics by enabling Tracing and using the Bulk Test feature. This allows you to systematically evaluate different prompt variants against a common dataset.

Steps to Evaluate and Compare Prompt Variants

*-> 1. Enable Tracing

Navigate to your Prompt Flow project.

Locate the Tracing toggle at the top of the flow authoring page.

Switch it to On.

This ensures every execution captures latency, token counts, and node-level inputs/outputs.

2. Create Prompt Variants

Within your flow, identify the LLM node you want to test.

Click Variants to create multiple versions of your prompt (e.g., Variant_0, Variant_1).

This allows you to test different instructions or few-shot examples side-by-side.

3. Run a Bulk Test (Evaluation)

4. Analyze the Results

Reference:

<https://www.linkedin.com/pulse/streamlining-generative-ai-development-azure-foundry-tracing- taneja-mbwze>

NEW QUESTION # 21

.....

It is because of our high quality AI-300 preparation software, PDF files and other relevant products, we have gathered more than 50,000 customers who have successfully passed the Microsoft AI-300 in one go. You can also attain the same success rate by using our high standard AI-300 Preparation products. Thousands of satisfied customers can't be wrong. You must try our products to

