

AI-300 aktueller Test, Test VCE-Dumps für Operationalizing Machine Learning and Generative AI Solutions



Alle IT-Fachleute sind mit der Microsoft AI-300 Zertifizierungsprüfung vertraut. Sie alle träumen davon, ein Zertifikat zu bekommen. Sie können Ihren Traum verwirklichen und eine gute Berufskarriere machen. Durch die Schulungsunterlagen zur Microsoft AI-300 Zertifizierungsprüfung von PrüfungFrage können Sie bekommen, was Sie wollen.

IT-Fachleute sind sehr beliebt. Aber die Konkurrenz ist zugleich auch sehr heftig. So beteiligen sich viele IT-Fachleute an der autoritären Microsoft AI-300 IT-Zertifizierungsprüfung, um Ihre Position zu konsolidieren. Und unser PrüfungFrage bietet speziell Bequemlichkeiten für den Microsoft AI-300 Kandidaten.

>> AI-300 Deutsche Prüfungsfragen <<

AI-300 Prüfungs & AI-300 PDF

Wenn Sie die Unterlagen von PrüfungFrage kaufen, bekommen Sie einjährigen kostenlosen Aktualisierungsservice. Wenn die Dumps aktualisiert sind, werden wir PrüfungFrage Ihnen die neuesten Versionen per E-Mail senden. Sie können auch an uns E-Mails schreiben, die neuesten Prüfungsunterlagen zur Microsoft AI-300 Zertifizierung zu fordern. Und PrüfungFrage kann Ihnen die Aktualisierungsservice innerhalb einem Jahr kostenlos bieten, obwohl Sie diese Microsoft AI-300 Prüfung erfolgreich machen.

Microsoft Operationalizing Machine Learning and Generative AI Solutions AI-300 Prüfungsfragen mit Lösungen (Q54-Q59):

54. Frage

A company requires that only models meeting predefined performance thresholds are registered and deployed. The solution must be fully automated within the ML workflow. What should you implement?

- A. Azure Monitor alerts
- **B. Conditional logic in pipeline**
- C. Manual approval gate
- D. GitHub pull request checks

Antwort: B

Begründung:

Conditional logic in pipelines allows automated decisions based on evaluation metrics, ensuring only models that meet performance thresholds are registered and deployed. This maintains full automation. Manual approval steps interrupt automation and are not scalable in continuous integration and deployment workflows.

55. Frage

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. UTF-8
- B. ISO-8859-1
- C. ASCII
- D. UTF-16

Antwort: A

Begründung:

For preparing a multi-turn training data file for the Azure OpenAI Service, you should use UTF-8 with a Byte Order Mark (BOM) encoding.

File Format Requirements

Format: The file must be in JSON Lines (JSONL) format, where each individual line is a valid JSON object representing one training example.

Encoding: Specifically, Azure OpenAI requires the JSONL file to be encoded in UTF-8 with BOM.

Structure: For multi-turn conversations, each line must contain a messages array with multiple role ("system", "user", "assistant") and content pairs to represent the dialogue history.

Reference:

<https://dev.to/icebeam7/fine-tuning-a-model-with-azure-open-ai-studio-39p7>

56. Frage

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. Use API key authentication for all model endpoints.
- B. Disable all inbound network access.
- 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.

Antwort: B,E

Begründung:

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>

57. Frage

An organization validates generative AI applications during CI/CD Microsoft Foundry.

Evaluation must run automatically and block releases when quality thresholds are NOT met.

Manual evaluation is no longer acceptable.

Evaluation must use both predefined quality metrics and custom safety checks.

You need to implement an automated evaluation workflow that supports both built-in and custom metrics.

What should you do?

- A. Enable application tracing to collect runtime telemetry.
- **B. Implement an evaluation step by using GitHub Actions.**
- C. Monitor latency metrics during model inference.
- D. Review evaluation results manually after deployment.

Antwort: B

Begründung:

To implement an automated evaluation step in GitHub Actions for Microsoft Foundry AI, you can use the Microsoft Foundry Evaluation GitHub Action (or the Azure AI Evaluation SDK).

This setup allows you to run both built-in metrics (like groundedness or coherence) and custom safety checks, then fail the build if scores fall below your defined thresholds.

Implementation Steps

1. Define Your Evaluators

You need to configure which metrics to use. Microsoft Foundry supports two main types:

Built-in Metrics: Pre-trained models that score quality (coherence, fluency) and safety (hate, violence, self-harm).

Custom Metrics: Python-based evaluators you define to check domain-specific requirements.

2. Configure the GitHub Actions Workflow

Create a .yaml file in your .github/workflows directory. This workflow will:

Trigger on a pull request or commit.

Authenticate with Azure/Foundry.

Run Evaluation using the microsoft/ai-agent-evals action.

Enforce Thresholds to block the release if quality is insufficient.

Key Components for "Block Release" Logic

To ensure the release is blocked, your workflow must include a gating step that interprets the evaluation results:

Reference:

<https://learn.microsoft.com/en-us/training/modules/automated-evaluation-genaiops>

58. Frage

A team develops and manages a conversational assistant by using Microsoft Foundry.

The team must be able to validate that the assistant does not produce hateful responses before the application is exposed to any users.

You need to evaluate the model output for hateful responses as part of a repeatable validation process.

Which evaluator should you configure first?

- A. Protected material
- B. Groundedness
- C. Indirect attacks
- **D. Content safety**

Antwort: D

Begründung:

You should use the Hate and Unfairness Evaluator (a specific type of Content Safety Evaluator) in Microsoft Azure AI Foundry.

This tool is designed to identify and measure the severity of hateful content toward protected groups, ensuring your assistant aligns with responsible AI standards.

Why Use the Content Safety Evaluator?

Targeted Detection: It identifies language attacking or discriminating against people based on race, religion, gender, and other identity factors.

Severity Scoring: It provides a numerical severity score (often 0-7) to help you understand the level of risk in a response.

Reasoning: It often includes a "reason" or "explanation" column that explains why a specific response was flagged.

myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, myportal.utt.edu.tt, mattierwmz554124.blogspot.com,
Disposable vapes