
Exam : AB-731

**Title : AI Transformation Leader
(beta)**

Version : DEMO

1.Which business requirement most closely relates to grounding a generative AI model?

- A. supporting multiple languages
- B. measuring the number of user interactions per day
- C. enabling users to interact by using natural language queries
- D. ensuring that verified company data sources are used for response generation

Answer: D

Explanation:

Grounding in generative AI means ensuring model outputs are based on trusted, relevant information sources rather than only on the model's general training data. In a business context, grounding is about aligning responses with verified enterprise knowledge (policies, product documentation, internal procedures, approved FAQs, etc.) so the system is more accurate, consistent, and defensible. That is exactly what option D describes: "ensuring that verified company data sources are used for response generation."

In Microsoft AI solution patterns, grounding is commonly achieved using retrieval-augmented generation (RAG). With RAG, the system retrieves relevant passages from approved company repositories (for example, indexed documents or knowledge bases) and supplies them as context to the model during response generation. This reduces hallucinations, improves factual correctness, and makes answers more relevant to the organization's reality—critical when AI is used for customer support, employee helpdesks, compliance guidance, or executive reporting. The other options do not directly address grounding.

A relates to localization/multilingual capability, B is a usage/telemetry metric, and C is an interaction method (natural language interface). They can all be important requirements, but none of them ensure outputs are anchored to verified company data—the core purpose of grounding.

2.HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point

Answer Area

Statements	Yes	No
For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license.	☐	☐
To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license.	☐	☐
To use the Analyst agent, you need a Microsoft 365 Copilot license.	☐	☐

Answer:

Answer Area

Statements	Yes	No
For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license.	☐	☒
To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license.	☒	☐
To use the Analyst agent, you need a Microsoft 365 Copilot license.	☒	☐

Explanation:

Top of Form

Bottom of Form

The key distinction here is between Copilot Chat capabilities available with a standard Microsoft 365 subscription and the full Microsoft 365 Copilot add-on that enables richer, in-context experiences grounded in organizational data.

Mobile access to organizational data does not inherently require a Microsoft 365 Copilot license.

Microsoft's Microsoft 365 Copilot app (and related mobile experiences) can provide Copilot Chat for work/school accounts with a Microsoft 365 license, so simply accessing organizational content on a mobile device is not the same as having the paid Copilot add-on. The statement claims a Copilot license is required just to access org data from mobile, which is too broad—there are mobile Microsoft 365 apps that access org data without the Copilot add-on.

Reasoning over organizational data via Microsoft Graph is a core value proposition of Microsoft 365 Copilot. Microsoft documents explain that Microsoft 365 Copilot connects LLMs to your organization's content and context through Microsoft Graph and generates responses "anchored" in organizational data. That deeper integration is tied to the Microsoft 365 Copilot experience (an add-on license).

Analyst is a "reasoning agent" within Microsoft 365 Copilot and Microsoft states that users with a Microsoft 365 Copilot license can use Analyst (with defined usage limits).

Therefore, the Analyst agent requires the Microsoft 365 Copilot license.

3.HOTSPOT

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Microsoft Foundry helps organizations securely build and manage generative AI solutions in a governed environment.	☐	☐
Microsoft Foundry provides built-in scalability to enable organizations to expand AI workloads as usage increases.	☐	☐
Microsoft Foundry can be used for image recognition and computer vision tasks.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Microsoft Foundry helps organizations securely build and manage generative AI solutions in a governed environment.	☒	☐
Microsoft Foundry provides built-in scalability to enable organizations to expand AI workloads as usage increases.	☒	☐
Microsoft Foundry can be used for image recognition and computer vision tasks.	☒	☐

Explanation:

Microsoft Foundry is positioned as a unified platform experience for building, optimizing, and governing AI applications and agents. Microsoft explicitly emphasizes "fleetwide security and governance" and the ability to build and manage AI in a unified environment, which directly supports statement 1 being Yes: it is designed to help organizations build and operate generative AI solutions with centralized governance controls (for example, environment setup, data isolation, access control, and operational management). For statement 2, Foundry supports scaling as demand increases. Microsoft documentation for Foundry-

related model usage notes that as usage grows, Foundry can automatically increase quotas by moving users to higher tiers (and allows requesting additional quota). This is a concrete scalability mechanism tied to increased workload demand, so the statement is Yes.

For statement 3, Foundry is not limited to text-only generative AI. Microsoft provides “Azure Vision in Foundry Tools,” which delivers computer vision capabilities such as analyzing images, reading text (OCR), and other image-processing features. That means Foundry can be used for image recognition/computer vision workloads, so the statement is Yes.

4. Your company uses a non-reasoning generative AI model to create textual content. You discover that the model’s responses are inconsistent and do NOT meet expectations. You need to improve the prompts.

What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Provide the prompts with extensive examples of the expected output.
- B. Add the context, sources, and expectations to the prompts.
- C. Use technical terms in the prompts to enhance AI comprehension.
- D. Add only a single concise requirement to the prompts.

Answer: A, B

Explanation:

When a non-reasoning generative AI model produces inconsistent outputs, the most reliable improvement is to make the prompt more specific, constrained, and demonstrative of what “good” looks like.

A is correct because adding high-quality examples is a form of few-shot prompting. Examples act like “training wheels” at inference time: they show the model the desired structure, tone, level of detail, formatting rules, and boundaries. This reduces ambiguity and variance, especially for tasks like marketing copy, summaries, policy text, or customer replies. The more your examples resemble real target outputs (including edge cases), the more consistent the model’s completions become.

B is correct because adding context, relevant source material, and explicit expectations narrows the model’s degrees of freedom. Including the intended audience, purpose, constraints (length, voice, banned claims), and trusted reference content (approved facts, product specs, policy excerpts) helps the model stay aligned and reduces hallucinations and off-brand language. This is also where you specify acceptance criteria such as “must include 3 bullet points,” “use UK English,” or “cite only provided text.”

C is not best: technical jargon can confuse or bias output if it’s not aligned to the task; clarity beats jargon.

D is not best: a single concise requirement is usually under-specified and often increases variability.

5. You have a historical dataset that contains 1,000 records. You need an AI solution that can analyze the data to identify patterns and predict future outcomes.

What should you include in the solution?

- A. Microsoft Foundry
- B. Azure Document Intelligence in Foundry Tools
- C. Azure Machine Learning
- D. Azure Content Understanding in Foundry Tools

Answer: C

Explanation:

The requirement describes a predictive analytics /machine learning scenario: using historical data to learn patterns and then predict future outcomes. The Microsoft service that directly supports the end-to-end machine learning lifecycle—data preparation, model training, evaluation, deployment, and MLOps—is Azure Machine Learning, which is why C is the best choice. Azure Machine Learning is explicitly designed to help data scientists and engineers train and deploy models and manage the ML project lifecycle, making it the right fit for building a predictive model from your dataset.

The other options focus on different problem classes: Azure Document Intelligence is for extracting structured data from documents (OCR, key-value pairs, tables), not for general predictive modeling. Azure Content Understanding is for deriving structured insights from multimodal content (documents, images, audio, video) into a user-defined schema; it's not the primary service for training predictive models from a tabular historical dataset. Microsoft Foundry is a broader platform for building AI apps/agents and orchestrating models/tools, but the specific need here is classical ML training and prediction—handled most directly by Azure Machine Learning.