Pass2Certify.com
Prepare, Practice, & Pass.

Microsoft

AB-731

ExamName: AI Transformation Leader

Exam Version: 6.0

Questions & Answers Sample PDF

(Preview content before you buy)

Check the full version using the link below.

https://pass2certify.com/exam/ab-731

Unlock Full Features:

Stay Updated: 90 days of free exam updates
Zero Risk: 30-day money-back policy
Instant Access: Download right after purchase
Always Here: 24/7 customer support team

Question 1. (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:

Question 2. (Multi Select)

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.

Question 3. (Single Select)

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.

Question 4. (DRAGDROP)

Match the business scenario to the appropriate AI solution design approach. Each solution may be used once, more than once, or not at all.

Answer:

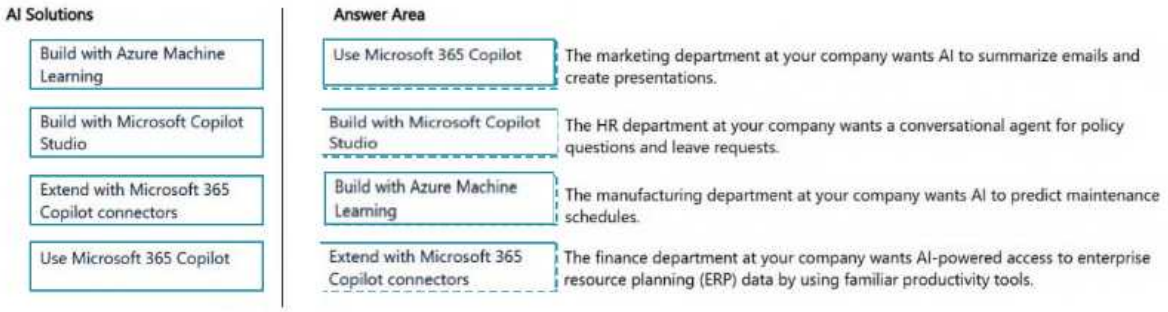

Question 5. (Single Select)

You need to create a custom Azure Machine Learning model. The data used to train the model is consistent and uniform. What should you do first?

A: Prepare the training data.
B: Evaluate the model.
C: Train the model.
D: Tune hyperparameters.
E: Deploy the model.

Answer: A

Explanation:
Even when training data is already consistent and uniform, the first step in building a custom Azure Machine Learning model is still to prepare the training data. “Consistent” data reduces the amount of cleaning you may need, but preparation is broader than cleaning: you still must confirm the schema, validate data types, handle missing values (if any), ensure label quality (for supervised learning), select/engineer features, and split data into training/validation/test sets. Those actions determine whether training will be stable and whether evaluation metrics will be meaningful.
If you skip preparation and go directly to training (C), the model might learn from the wrong columns, inconsistent labels, or poorly partitioned data, producing misleading results. Evaluation (B) comes after training because you need a trained model to score and measure. Hyperparameter tuning (D) is an optimization activity that presupposes you already have a working training pipeline and a baseline model to improve. Deployment (E) is last, after you have validated performance and selected the model candidate. Azure Machine Learning commonly operationalizes these steps through pipelines, where data preparation is a foundational stage that precedes training and evaluation (and can also be iterated as you refine features and quality).

Need more info? Check the link below:

https://pass2certify.com/exam/ab-731

Thanks for Being a Valued Pass2Certify User!

Guaranteed Success Pass Every Exam with Pass2Certify.

Save $15 instantly with promo code

SAVEFAST

Sales: sales@pass2certify.com

Support: support@pass2certify.com