



Microsoft AB-100

Agentic AI Business Solutions Architect

Thank You for Downloading AB-100 Updated Exam Questions

<https://certs4sale.com/microsoft/ab-100-pdf-exam-dumps>

Question 1

MULTIPLE CHOICE

Your organization is planning to implement an AI-powered solution for their sales team using Dynamics 365 Sales and Microsoft 365 Copilot for Sales. The sales director has expressed concerns about how the AI will make recommendations and wants to ensure the solution aligns with responsible AI principles. Before moving forward, you need to assess what aspects of the solution design should be prioritized to address these concerns.

Which two of the following should you include in your responsible AI design approach? (Choose two.)

- A. Review the solution for adherence to responsible AI principles and design access controls on grounding data to ensure model recommendations are based on accurate, authorized information
- B. Implement Model Context Protocol to automatically filter out any AI-generated recommendations that don't align with company policy
- C. Design audit trails for changes to models and data to provide transparency and traceability of how the AI system evolves and makes decisions
- D. Replace all grounding data with synthetic data to eliminate bias concerns entirely
- E. Use Computer Use agents to automate all sales recommendations without human oversight

Answer: A, C

Explanation:

The correct answers are reviewing adherence to responsible AI principles and designing audit trails for changes to models and data.

Why these are correct: Responsible AI design requires assessing alignment with ethical principles like fairness, transparency, and accountability. Audit trails provide transparency by documenting how models and data change over time, enabling organizations to trace decisions and identify potential issues. Access controls on grounding data ensure only authorized, accurate information influences AI recommendations—this is a governance and security control essential to responsible AI.

Why the others are incorrect:

- **Model Context Protocol automatic filtering:** MCP enables extensibility and interoperability for agents, but it doesn't automatically enforce policy compliance; that requires intentional design and governance.
- **Synthetic data replacement:** While synthetic data can help with privacy, it doesn't eliminate bias and creates new risks around data quality and relevance for grounding. The focus should be on data quality assessment and governance, not wholesale replacement.
- **Computer Use without oversight:** Computer Use in Copilot Studio enables agents to automate tasks across apps, but responsible AI requires human oversight and explainability, not full autonomy without controls.

Question 2**SINGLE CHOICE**

A financial services organization is planning to implement AI agents across their Dynamics 365 Finance and Supply Chain Management systems. They want to automate invoice processing, demand forecasting, and compliance reporting. The organization has identified that data quality issues exist in their historical supplier and customer records, and they need to prioritize which AI components to build versus extend from existing Microsoft solutions.

As the AI Business Solutions Architect, what should be your primary focus during the planning phase?

- A. Build custom foundational AI models trained exclusively on internal data to maintain full control over the algorithms and training data
- B. Assess the accuracy, relevance, timeliness, and cleanliness of the data available for grounding. Then evaluate whether to build, buy, or extend prebuilt agents for each use case based on data readiness and business impact
- C. Deploy Microsoft 365 Copilot across all Dynamics 365 modules first and defer any data quality improvements until after the agents are in production
- D. Create a prompt library immediately and begin training the organization's staff on prompt engineering before completing any technical assessment

Answer: B

Explanation:

The correct answer focuses on the planning phase's core requirement: assessing data quality for grounding and using a build/buy/extend decision framework. The scenario explicitly mentions data quality issues, making data assessment critical. The AB-100 exam emphasizes that grounding data must be evaluated for accuracy, relevance, timeliness, cleanliness, and availability before designing agent solutions. Additionally, the 'build, buy, or extend' analysis is a key ROI and design decision point when planning multi-system implementations.

Building all custom models (option A) skips necessary assessment and wastes resources. Deploying Copilot before addressing data quality (option C) violates best practices and will result in poor agent performance. Creating prompts before planning (option D) puts the cart before the horse and ignores the foundational planning requirements.

<https://www.validexamdumps.com>

Question 3

MULTIPLE CHOICE

A retail organization wants to create a multi-agent solution that includes a customer service agent in Copilot Studio, a sales enablement agent in Microsoft 365 Copilot for Sales, and an inventory management agent built with Microsoft Foundry. Each agent will need to access different data sources and operate under different business rules.

Which of the following should be included in your agent design strategy to ensure proper orchestration and governance across this multi-agent solution? (Select two)

- A. Define clear solution rules and constraints for each agent within Copilot Studio and Microsoft Foundry, including data access boundaries and response limitations
- B. Implement a model router to intelligently route requests to the most suitable agent based on the nature of the user query and the agent's capabilities
- C. Configure all agents to use the same large language model and knowledge base to simplify maintenance and reduce licensing costs
- D. Require all agents to operate independently without any integration points to avoid complexity and potential failures
- E. Use only prebuilt agents from Microsoft and avoid any custom extensibility to maintain supportability

Answer: A, B

Explanation:

The two correct answers address critical multi-agent orchestration requirements. Defining solution rules and constraints for each platform (Copilot Studio, Microsoft 365 Copilot, Microsoft Foundry) ensures agents operate safely within organizational policies and data governance boundaries—a key requirement in the design phase. Implementing a model router is a specific technique mentioned in the exam objectives for intelligently routing requests to the most suitable model or agent based on context, improving efficiency and user experience.

Option C is incorrect because different agents serve different purposes and may benefit from different models or knowledge bases optimized for their domain. Option D is incorrect because agents often need to integrate to provide comprehensive solutions. Option E is incorrect because custom extensibility (through Model Context Protocol, Computer Use, and agent customization) is often necessary to meet business requirements beyond prebuilt capabilities.

Question 4**SINGLE CHOICE**

A manufacturing company is implementing a custom AI agent in Copilot Studio that will assist production planners with demand forecasting and supply chain optimization. The organization has invested significantly in preparing training data, but they're uncertain whether to use standard natural language processing (NLP) Azure Conversational Language Understanding (CLU), or generative AI orchestration for the core logic.

What should guide your decision on which natural language understanding approach to recommend?

- A. Always default to generative AI orchestration because it is the most advanced and will provide the best results for all scenarios
- B. Evaluate the use case complexity, the need for structured entity extraction and intent classification, the availability of training data, and the required response latency to determine the best approach for the scenario
- C. Use standard NLP for all supply chain optimization tasks because it is the most reliable and cost-effective solution for enterprise workloads
- D. Implement all three approaches simultaneously so the agent can choose the best one at runtime without additional guidance

Answer: B

Explanation:

The correct answer reflects the nuanced decision-making required in the design phase of AI solutions. Standard NLP, Azure CLU, and generative AI orchestration each have different strengths: NLP excels at pattern matching with low latency, CLU is optimized for structured entity and intent classification with labeled training data, and generative AI orchestration handles complex reasoning and open-ended scenarios. The AB-100 exam emphasizes that architects must evaluate specific requirements like complexity, data availability, performance needs, and cost to make appropriate technology choices.

Option A is incorrect because generative AI orchestration adds latency and cost that may be unnecessary for simpler, well-structured tasks. Option C is incorrect for the opposite reason—not all scenarios require the lightweight approach of NLP. Option D is incorrect because running all three simultaneously wastes resources and creates unnecessary complexity without clear business justification.

Question 5**SINGLE CHOICE**

Your organization has deployed a conversational agent in Copilot Studio that handles customer support inquiries. After two months in production, users report that the agent frequently provides outdated product information, sometimes contradicts recent policy changes, and occasionally hallucinates details about return procedures. The telemetry shows high user abandonment after the agent's first response.

As the solution architect, what steps should you take to address these issues?

- A. Immediately retrain the underlying large language model from scratch using proprietary data to eliminate hallucinations
- B. Monitor agent performance metrics and telemetry data, analyze the specific failure patterns, review and update the grounding data sources for accuracy and timeliness, validate the prompt best practices, and iterate on the solution design based on user feedback
- C. Disable the agent and return to human-only support until the organization can upgrade to a newer LLM version
- D. Remove the agent's access to external knowledge sources and limit it to only pre-scripted responses to prevent hallucinations

Answer: B

Explanation:

The correct answer represents the proper post-deployment monitoring and tuning process outlined in the AB-100 exam. The scenario presents classic symptoms of poor grounding data (outdated product info), stale knowledge sources (policy changes), and prompt issues (hallucinations). The deployment phase requires analyzing telemetry, monitoring metrics, assessing user feedback, and iteratively tuning the agent. This includes validating that grounding data is accurate, relevant, and timely, and ensuring prompts follow practices.

Option A is incorrect because retraining a large LLM is expensive, time-consuming, and usually unnecessary; the problem is typically in grounding and prompting, not the model itself. Option C is incorrect because it abandons the solution rather than improving it through proper tuning. Option D creates a brittle solution that cannot leverage the agent's strengths and fails to address the root causes of hallucination (poor grounding and prompts).

THANK YOU FOR DOWNLOADING AB-100 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our AB-100 Exam Product

Visit Our Site to Purchase the Full Set of Actual AB-100 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/microsoft/ab-100-pdf-exam-dumps>