



Free Microsoft AB-620 Practice Questions

Shared by Macias on 13-08-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

You run the same fixed test set three times in Copilot Studio.

During evaluation, you observe the following:

The same interaction fails in all three runs.

Score values range from 0.58 to 0.61.

The reasoning states that the response partially matches the expected answer.

The knowledge source that is used is internal documentation.

No tools are invoked.

Conclusions	Evaluation result conclusions	
☐ A pattern of failure across repeated runs.	Evaluation results	
☐ The specific underlying cause of the failure	☐ Evidence of a recurring issue.	
☐ The response partially matches the expected answer	☐ Information provided about response quality.	
☐ The response partially matches the expected answer	☐ Evaluation limitations.	

Answer:

See the Answer in the Premium Version!

Question 2

Question Type: MultipleChoice

An agent uses a flow that calls an external service which can occasionally fail or time out. When a failure occurs, the agent must meet the following requirements.

- * Must not terminate silently.
- * Must send a notification containing the error details.

You need to configure the agent flow so that failures are handled in a controlled and predictable way. Solution: Terminate the flow immediately on error.

Does the solution meet the goal?

Options:

A- Yes

B- No

Answer:

B

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: Terminating immediately on error makes the failure explicit to the runtime but does not send the required notification. The solution therefore still allows an operator-facing silent failure even though the run status is marked failed. A proper pattern uses Configure run after to enter a notification or catch scope on failure and timeout, collects safe diagnostic details, sends the administrative alert, and only then terminates with the intended status. If notification delivery fails, a secondary monitoring or logging path should still preserve the incident. The error message should include a run or correlation identifier, the failed action, and a concise remediation hint without exposing credentials or sensitive business data. Retry behavior should be configured separately according to idempotency and service characteristics. The proposed immediate termination satisfies neither the communication requirement nor controlled recovery, so the answer is No. Tests should force each expected failure state and confirm both notification delivery and final run status. Study Guide alignment: Plan and configure agent solutions > Create and monitor agent flows in Copilot Studio > Implement error handling in agent flows.

=====

Question 3

Question Type: MultipleChoice

You create a test set in Copilot Studio to evaluate an agent that answers policy questions and retrieves data from connected knowledge sources. The evaluation threshold is set to 0.80.

You run the evaluation and review the following result for one test case:

Expected response: Digital products are non-refundable after download.

Actual response: Digital products cannot be refunded once downloaded.

Result: Pass

Score: 0.92

Reasoning: The actual response semantically matches the expected response.

Knowledge sources used: RefundPolicy_KB

Tools invoked: None

You need to determine whether additional corrective action is required based on the evaluation results.

Which two actions should you perform? Each correct answer presents a complete solution. Choose two.

NOTE: Each correct selection is worth one point.

Options:

- A- Rerun the evaluation because the score is below 1.00.
- B- Confirm that the connected knowledge source is functioning as expected.
- C- Investigate external tool configuration for this test case.
- D- Modify the test case because the response is not an exact match.
- E- Accept the result because semantic matching is sufficient for a passing evaluation.

Answer:

B, E

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: The test passed because the score of 0.92 exceeds the configured threshold of 0.80 and the evaluation method recognizes semantic equivalence. The actual answer conveys the same policy meaning as the expected response even though the wording is different, so an exact textual match is unnecessary. The recorded knowledge source is RefundPolicy_KB, which is the expected grounding path; confirming that source behavior is therefore appropriate. No tool was expected or invoked, so investigating tool configuration would add noise rather than address a defect. A score below 1.00 is not itself a failure, and rewriting the expected answer to mirror every valid paraphrase would undermine a meaning-based evaluation. The result can be accepted while still being retained as regression evidence. Reviewers should confirm that the cited or retrieved policy content supports the non-refund statement and that the chosen threshold reflects the business risk. For fixed codes or legally prescribed phrases, exact match might be preferable, but this natural-language policy answer is correctly evaluated by semantic similarity. Study Guide alignment: Test and manage agents > Evaluate agent performance > Review test results.

=====

Question 4

Question Type: MultipleChoice

A company uses an agent in Copilot Studio that produces structured responses using a reusable instruction template. The responses follow the required tone and format. During a compliance review, the team discovers that the responses were generated using a model that is not on the approved list.

The company requires that the reusable instruction template continues to be used and the model used for prompt-driven responses meets governance requirements.

The solution must ensure that prompt-driven responses are generated using an approved model.

You need to configure the solution that meets the requirements.

Which two actions should you perform? Each correct answer presents part of the solution. Choose two.

NOTE: Each correct selection is worth one point.

Options:

- A- Ensure the reusable instruction template is the one applied to the agent that generates the structured response.
- B- Select an approved model from the Microsoft Foundry model catalog in the configuration that is used for prompt-driven response generation.
- C- Enable citations so generated responses can be traced and comply with governance requirements.
- D- Configure Azure AI Search as the grounding source so the approved model requirement is enforced during response generation.
- E- Publish the agent so that the model used by prompt-driven responses is updated automatically.
- F- Add a line to the instruction template that requires the agent to use only approved models when generating responses.

Answer:

A, B

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: The approved solution preserves the reusable prompt template and changes the model configuration that actually executes it. Applying the established template retains the required tone, structure, and repeatable instruction set. Selecting an approved model in the prompt configuration ensures that prompt-driven responses are produced by a model permitted under the organization's governance policy. An instruction that tells the model to use an approved model cannot control the underlying runtime; model selection is a platform configuration, not a behavior the model can enforce on itself. Citations address provenance, not model approval. Azure AI Search supplies grounding data but also does not choose the foundation model. Publishing is necessary to expose completed changes to users, yet it does not automatically substitute a compliant model. The team should also test the prompt against representative inputs after changing models, because output quality, formatting adherence, latency, and token behavior can vary across models even when the prompt remains identical. Governance records should capture the selected model, prompt version, owner, and validation evidence. Study Guide alignment: Integrate and extend agents in Copilot Studio > Integrate agents with Azure > Configure custom prompts to use the Foundry model catalog.

=====

Question 5

Question Type: MultipleChoice

You need to enable Operations Concierge to delegate shipment tracking inquiries according to Fabrikam Inc.'s defined architecture and technical requirements.

Which integration approach should you use?

Options:

- A- Configure the primary agent to handle shipment tracking through generative responses after publishing.
- B- Configure the partner capability as an enterprise knowledge source with indexed retrieval enabled.
- C- Configure the partner capability as an MCP tool and invoke it as a transactional operation.
- D- Configure the partner capability as an A2A integration and supply its endpoint information.

Answer:

D

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: The partner capability is described as an agent available only through a standardized agent-to-agent endpoint. That is the defining use case for an Agent2Agent connection. In Copilot Studio, the primary agent is configured with the external agent's A2A message endpoint and the applicable authentication details, after which orchestration can delegate shipment-tracking requests to it. An MCP server exposes tools or resources rather than an independently orchestrated partner agent, so option C changes the integration contract. An indexed knowledge source is also unsuitable because shipment status is dynamic and the requirement is delegation, not retrieval from a copied corpus. Generative responses in the primary agent would have no authoritative live shipment source and would violate the defined architecture. The A2A description should identify shipment tracking precisely so that it does not compete with unrelated agents or tools. Because external-agent support can be preview functionality, Fabrikam should also validate regional availability, authentication, error behavior, latency, and governance before production use. Study Guide alignment: Integrate and extend agents in Copilot Studio > Configure multi-agent collaboration from Copilot Studio > Create a multi-agent solution by using A2A protocol.

=====

Question 6

Question Type: MultipleChoice

You need to ensure that every AI-generated response from the agent in Copilot Studio includes a disclaimer that complies with the company's security and governance policies.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

Options:

- A- Add a greeting message that includes the disclaimer.
- B- Add a prompt modification in the generative answers node settings.
- C- Create a disclaimer topic that always runs first.
- D- Disable generative answers and use only pre-authored responses.
- E- Edit every topic's first message to include the disclaimer.

Answer:

B, C

Question 7

Question Type: MultipleChoice

A company requires an agent in Copilot Studio to answer questions using content stored in a third-party knowledge base system and reuse that content for knowledge grounding.

The agent must ground answers using indexed enterprise knowledge from a non-Microsoft system.

You need to choose the correct integration approach.

What should you use?



Options:

- A- Call the third-party API from a topic for each question.
- B- Add a Copilot Studio tool as a real-time knowledge source.
- C- Add a Copilot Studio tool and call it from a topic.
- D- Add a Copilot connector knowledge source for authorized content.

Answer:

D

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: A Copilot connector knowledge source is the appropriate integration for indexed enterprise content held in a non-Microsoft system. The **connector** brings **authorized content** into the Microsoft Graph search index and makes it available for **knowledge grounding**, while access controls determine which results a user may receive. Calling the third-party API from a topic on every question is a tool pattern and would require the agent to implement retrieval, ranking, and response handling itself. A Copilot Studio tool is also an executable capability, not automatically a knowledge source, so options B and C do not meet the indexed-grounding requirement. After the tenant administrator configures and synchronizes the connector, the builder adds the corresponding enterprise knowledge source to the agent, supplies clear metadata, and tests retrieval and citations. Content lifecycle, crawl status, schema mapping, and permission trimming must be monitored because a technically successful connection can still return stale, poorly ranked, or overexposed content. The design should avoid uploading a second copy that would drift from the connector-managed index. Study Guide alignment: Integrate and extend agents in Copilot Studio > Connect to enterprise knowledge sources > Connect to Copilot connectors.

=====

Question 8

Question Type: MultipleChoice

A company is building an agent topic in Copilot Studio. The topic must meet the following requirements:

- * Collect structured user input in a single step.
- * Store the submitted values for use later in the conversation.

You need to configure the topic based on the requirements.

Which two actions should you perform? Each correct answer presents part of the solution. Choose two.

NOTE: Each correct selection is worth one point.

Options:

- A- Use the latest Adaptive Card schema version in all channels to ensure full feature support
- B- Paste the card JSON into a Message node to collect the user's response.
- C- Add input fields to the card and include an Action.Submit button.
- D- Add an Ask with Adaptive Card node to the topic.
- E- Use Action.Execute for the button so the card can run commands across channels.

Answer:

C, D

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: The Ask with Adaptive Card node is the interactive topic element that waits for a card submission and stores the user's responses as variables. The card must contain input controls with unique identifiers and an Action.Submit button so the host returns the collected values to Copilot Studio. A Message node can display a card, but it is not the correct choice when the topic must collect and bind structured input in one step. Action.Execute is not consistently supported across Copilot Studio channels; Microsoft documents host-specific schema limits, and the default web chat does not support that action. Using the latest schema indiscriminately can therefore break portability. The card should target the lowest schema version required by every intended channel, validate required fields, and provide clear labels and error messages. After submission, the topic variables should be typed and checked before they are passed to tools or flows. Testing is required in each production channel because rendering and supported Adaptive Card features differ between Teams, web chat, and customer-service hosts. Study Guide alignment: Plan and configure agent solutions > Configure topics > Configure adaptive cards; Manage variables.

=====

Question 9

Question Type: MultipleChoice

A team is standardizing ALM for agents in Copilot Studio.

The team requires a custom solution that includes agents. The solution must support ALM scenarios, such as managing and transporting agent assets across environments and pipeline deployments.

You need to create an ALM-ready custom solution in Copilot Studio.

What should you do?

Options:

- A- Export the agent directly and then import it into another environment.
- B- Publish the agent to a channel and use that channel configuration as the ALM container.
- C- Create a new custom solution from the Solutions area.
- D- Export the custom solution as a managed solution and import it into production.

Answer:

C

Explanation:

Comprehensive and Detailed Explanation From Microsoft AB-620 Study Guide: The required first action is to create a custom solution from the Solutions area. A solution is the ALM carrier that groups the agent with its related components and enables controlled export, import, pipeline deployment, dependency management, and versioning across environments. Exporting an agent directly is not the supported replacement for solution-based ALM. A channel configuration determines where users can interact with a published agent; it is not a container for development assets. Exporting a managed solution and importing it into production is a later release step and cannot occur until the custom solution exists and contains the agent and its dependencies. After creation, the team should add the existing agent, associated flows, connection references, environment-variable definitions, prompts, and any other required components. Development normally occurs in an unmanaged solution, while a managed build is deployed to Test and Prod. This separation supports repeatable promotion, protects downstream environments from unmanaged edits, and provides a clearer rollback and support boundary. Study Guide alignment: Test and manage agents > Implement application lifecycle management

for agents in Copilot Studio > Create a solution.

=====



To Get Premium Files for AB-620 Visit

<https://www.p2pexams.com/products/ab-620>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/ab-620>

20%
DISCOUNT

P2P
exams