



Accelerate Your Certification with Microsoft AB-620 Practice Questions

Last chance to prepare smart! Get your hands on free Microsoft Designing and Building Integrated AI Agent Solutions in Copilot Studio exam PDF questions. Study real AB-620 dumps with verified answers and fast-track your certification success with [PrepBolt](https://prepbolt.com/AB-620.html) AB-620 exam pdf questions and answers.

Thank you for Downloading AB-620 exam PDF Demo

<https://prepbolt.com/AB-620.html>

QUESTIONS & ANSWERS
DEMO VERSION
(LIMITED CONTENT)

Question 1

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

An agent in Copilot Studio already exists in a development environment.

You need to prepare the agent so it can be moved to a production environment by using the supported solution-based approach.

You need to create an exportable package that contains the existing agent.

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

NOTE: Each correct selection is worth one point.

Options:

- A- Remove agent topics directly from the solution to prevent export errors.
- B- Convert the solution to a managed solution so it can be exported.
- C- Export the unmanaged solution that contains the agent.
- D- Publish the agent to a channel and export the channel configuration.
- E- Create a custom solution and add the existing agent to the solution.

Answer:

C, E

Explanation:

Comprehensive and Detailed Explanation From [Microsoft AB-620 Study Guide]

First, create a custom solution in the development environment and add the existing agent by selecting Add existing Agent Agent. The solution acts as the application lifecycle management package that contains the agent and its dependent components. Required objects such as agent flows, environment variables, connection references, and custom connectors should also be included when applicable.

After the agent and its dependencies are included, export the solution from the development environment. Microsoft's documented procedure starts from the list of unmanaged solutions and exports the solution containing the agent. During the export process, the maker can select the appropriate package type. Production deployments commonly use a managed export, but the source

development solution remains unmanaged; it does not need to be converted before export.

Removing topics directly from the solution is specifically discouraged because changing or removing agent components outside the standard Copilot Studio authoring interface can cause export or import failures. Publishing the agent to a channel is also unrelated to creating a solution package. After the solution is imported into production, the imported agent must be configured, tested, and published in that environment.

Relevant study-guide area: Test and manage agents Implement application lifecycle management for agents in Copilot Studio Create a solution and add existing agents to a solution. See Export and import agents using solutions.

Question 2

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

Users can add content to a paragraph field in a team-managed project, but they cannot @mention other users.

Identify the cause.

Options:

- A- Global permissions
- B- Field settings
- C- Issue-level security settings
- D- Notification settings
- E- Screen settings
- F- Project permissions

Answer:

A

Explanation:

Comprehensive and Detailed Explanation From [Atlassian Jira Cloud Documentation]

The cause is the Jira Browse users and groups global permission. Atlassian states that users with this permission can view user names when working with user-picker fields or @mentioning users. Without it, Jira cannot provide the affected users with the complete user directory required to search for and select another person through an @mention.

The fact that users can enter and save ordinary text in the paragraph field confirms that the field is present, editable, and available on the applicable work-item layout. Therefore, field settings and screen settings are not responsible. Notification settings determine how users receive updates after an event and do not control whether names appear in the mention picker.

Project permissions and issue-level security can affect whether a particular person can access the project or work item. However, a general inability to search for or mention other users indicates that the initiating users lack the site-wide user-browsing capability.

A Jira administrator should review Settings System Global permissions and grant Browse users and groups to an appropriate group containing the affected users, subject to the organization's privacy policies.

Relevant documentation topic: Jira Cloud administration Global permissions Browse users and groups. See What are global permissions, and what do they control?.

Question 3

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

A company uses an agent in Copilot Studio to answer questions based on enterprise content.

The content is indexed in Azure AI Search. Responses must be generated by a Microsoft Foundry model.

To meet the business needs, the solution must meet the following requirements:

Users must be able to trace each answer back to its source.

The agent must retrieve information from a centralized data index.

Responses must be generated by an enterprise-approved foundation model.

You need to configure generative answers for the agent.

What should you configure for each requirement? To answer, move the appropriate configurations to the correct requirements. You may use each configuration once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Configurations	Evaluations	
Upload documents directly into the agent.	Requirement	Configuration
Configure Azure AI Search as the grounding data source.	Users trace each answer back to its source.	
Instruct the agent to cite sources in knowledge settings.	The agent retrieves information from a centralized data index.	
Deploy a model in Foundry to an agent and connect to Copilot Studio.	Responses are generated by an enterprise-approved foundational model.	

Answer:

See the Answer in the Premium Version!

Question 4

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

A company plans to use Copilot connectors to ground answers from a third-party system for an agent in Copilot Studio.

To meet the business needs, responsibilities for configuring and using Copilot connectors must be separated based on the following requirements:

For tenant administration, the connector must be available at the tenant level.

For agent runtime behavior, answers must be grounded from authorized sources.

For builder configuration, you must add the connector to the agent.

You need to determine how Copilot connector setup responsibilities are handled across the tenant, runtime, and builder stages.

What should you configure for each requirement? To answer, move the appropriate configuration to the correct requirement. You may use each configuration once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Configurations	Copilot connector setup responsibilities.	
	Requirement	Configuration
Call the connector knowledge source.	Tenant administration	
Add the Copilot connector as a knowledge source.	Agent runtime behavior	
Create a new Copilot connector directly in Copilot Studio.	Builder configuration	
Configure the Copilot connector in Microsoft 365 Admin Center.		
Configure the Copilot connector in Power Platform Admin Center.		

Answer:

See the Answer in the Premium Version!

Question 5

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

A company integrates a Model Context Protocol (MCP) service that exposes ticket-management operations.

The MCP service requires authenticated requests and exposes multiple callable operations that must be selectively governed for agent use.

To meet the business needs, the MCP configuration must meet the following requirements:

Only the allowed capabilities must be callable during conversations.

Requests must succeed without embedding credentials in topic steps.

The agent must be configured to send requests to the correct backend.

You need to configure MCP tools for the agent.

Which setting should you configure for each requirement? To answer, move the appropriate settings to the correct requirements. You may use each setting once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Settings

Server URL

API key header X-API-Key

CreateTicket, EscalateIncident, GetTicket, UpdateSLA

Agent-to-agent integration prerequisites in Copilot Studio

Requirement	Setting
Only the allowed capabilities must be callable during conversations.	
Requests must succeed without embedding credentials in topic steps.	
The agent must be configured to send requests to the correct backend.	

Answer:

See the Answer in the Premium Version!

Question 6

Question Type: Hotspot

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

You deploy an agent in Microsoft Copilot Studio that triggers a cloud flow to update customer records in Microsoft Dataverse.

You observe the following after deploying the agent:

Several runs fail because of a temporary authentication issue.

The flow continues to trigger and generate additional failed runs.

You must:

Stop additional failed executions while troubleshooting.

Process a failed execution after resolving the authentication issue.

Validate whether the most recent executions succeeded immediately after resolution.

You need to use the appropriate monitoring and management actions to restore normal flow operation and validate successful execution.

Which action should you perform for each requirement?

Flow monitoring actions

Requirement

Stop additional failed executions while troubleshooting.

Process a failed execution after resolving the authentication issue.

Validate if the most recent executions succeeded immediately after resolution.

Action

Delete the flow.
Disable the flow.
Cancel the most recent run.
Modify the trigger condition.

Restart the flow.
Refresh the flow.
Resubmit the run.
Recreate the trigger.

Republish the flow.
Review the Activity history.
Export the flow to a solution.
Inspect the Analytics summary.

Answer:

See the Answer in the Premium Version!

Question 7

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

A company is preparing an agent flow so that it can be invoked by an agent during conversations.

The agent flow must meet the following requirements:

The agent must be able to trigger the flow.

The agent flow must be verified.

You need to prepare an agent flow so that it can be used by the agent.

In which order should you perform the actions to prepare the agent flow? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

Agent flow readiness steps

Configure the invocation trigger that allows the agent to call the agent flow.

Make the agent flow available for agent usage.

Validate the agent flow behavior with representative test runs.

Configure the actions that implement the required behavior.

Create a new agent flow artifact in the environment.

Answer Area

1

2

3

4

5

Answer:

See the Answer in the Premium Version!

Question 8

Question Type: MultipleChoice

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

A company is configuring generative answers for an agent in Copilot Studio.

The agent must use Azure AI Search to retrieve enterprise content, and all responses must be generated by a Microsoft Foundry model.

The agent sometimes generates answers that are not grounded in Azure AI Search content. The company requires that the agent only generate responses when relevant indexed content is retrieved.

You need to configure the agent so that its generative answer behavior meets the organization's requirements.

What should you do?

Options:

A- Use Azure AI Search as a fallback source and enable answer expansion.

B- Upload documents into the agent and disable fallback responses.

- C- Require grounding for generative answers and use Azure AI Search as the retrieval source.
- D- Enable citations and configure a public model endpoint.
- E- Enable conversational boosting and select a Foundry-approved model.

Answer:

C

Explanation:

Comprehensive and Detailed 150 to 250 words of Explanation From [Microsoft AB-620 study guide]: Azure AI Search must be configured as the authoritative retrieval source, with ungrounded or general-knowledge responses disabled. In the retrieval-augmented generation process, Copilot Studio first reformulates the user's question, searches the connected vector index, and supplies relevant results to the Microsoft Foundry-hosted model for response generation. Requiring grounding prevents the model from answering solely from its pretrained knowledge when Azure AI Search returns no sufficiently relevant content. In that situation, the agent should return no generated answer and follow its configured fallback or escalation behavior. Selecting a Foundry model alone does not enforce grounding because model selection controls the generation layer, not the retrieval requirement. Citations provide traceability but likewise do not prevent an unsupported answer. Uploading duplicate documents would introduce another knowledge store and undermine Azure AI Search as the controlled enterprise index. The Azure AI Search connection should identify the intended vector index and use semantic ranking where appropriate. Testing should confirm that valid queries return supported answers and citations, while out-of-scope queries produce no ungrounded response. Study Guide alignment: Integrate and extend agents in Copilot Studio Connect to Azure AI Search; Integrate agents with Azure Configure generative answers by using Azure AI Search with Foundry.

Question 9

Question Type: DragDrop

Case Study: Mix Questions

Mix Questions

AB-620 Mix Questions IN THIS CASE STUDY

A company configures an agent flow that interacts with multiple systems to complete user requests.

The company must integrate external services, retrieve data, and notify users as part of the same automated flow run. To support this scenario, the agent flow must be able to do the following:

Call an external API.

Retrieve stored data from a supported internal data source.

Send a Teams message.

You need to select a connector or action type for each requirement.

What should you use? To answer, move the appropriate connectors or actions to the correct requirements. You may use each connector or action once, more than once, or not at all. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Connector or action	Connectors and actions	
Trigger a notification.	Requirement	Connector or action
Make an HTTPS request.	Call an external API.	
Add a standard connector.	Retrieve stored data from a supported internal data source.	
	Send a Teams message.	

Answer:

See the Answer in the Premium Version!

Thank You for trying AB-620 PDF Demo

To try our AB-620 practice exam software visit
link below

<https://prepbolt.com/AB-620.html>

Start Your AB-620 Preparation

Use Coupon “**SAVE50**” for extra 50% discount on the purchase of Practice Test Software. Test your AB-620 preparation with actual exam questions.