DUMPSARENA

Designing and Building Integrated AI Agent Solutions in Copilot Studio

Microsoft AB-620

Version Demo

Total Demo Questions: 11

Total Premium Questions: 117

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

Topic Break Down

Topic	No. of Questions
Topic 1, Design agent solutions	30
Topic 2, Build agent solutions	36
Topic 3, Integrate agent solutions	50
Topic 4, Mix Questions	1
Total	117

QUESTION NO: 1 - (DRAG DROP)

You need to connect Operations Concierge to Fabrikam Inc. ' s Azure AI Search knowledge index while complying with security requirements.

Which configuration should you use 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. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point

ANSWER:

Explanation:

An Azure AI Search knowledge connection requires separate settings for authorization, the search service, and the index to query. For authenticated governed access, select the service principal as the authentication method. A service principal represents an application identity in Microsoft Entra ID, allowing Operations Concierge to access Azure AI Search through a managed, governed identity rather than relying on an individual maker's personal credentials. This supports centrally controlled permissions and is appropriate when the organization needs to manage access to its enterprise knowledge source.

Enter the name of the Azure AI Search index to identify the curated vector index that contains the policy and procedure content. An Azure AI Search service can contain multiple indexes, each with different documents, fields, vector configurations, and search behavior. Specifying the index name ensures that retrieval for policy and procedure questions is directed to the intended curated corpus. The index is therefore the setting that selects the knowledge collection used to ground responses.

Provide the Azure AI Search Endpoint URL to identify the Azure AI Search service instance itself. The endpoint is the service-level address used to reach the Azure AI Search resource that hosts the indexes. Together, the endpoint and index name identify both the correct service and the exact searchable index within that service, while service-principal authentication establishes the governed identity used for access. Microsoft documents Azure AI Search authentication and authorization through Microsoft Entra ID at Role-based access control in Azure AI Search. For service and index concepts, see Indexes in Azure AI Search.

QUESTION NO: 2

An employee-facing agent answers questions by using a SharePoint knowledge source that contains departmental documents.

The company requires that employees authenticate before using the agent and that employees cannot receive content from SharePoint locations for which they do not have permission.

You need to configure the solution to meet the security 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.

A. Select Authenticate with Microsoft for the agent.

B. Maintain appropriate SharePoint permissions for the source content.

C. Copy the departmental documents to a public website knowledge source.

D. Add instructions that tell the agent not to reveal confidential information.

E. Configure the SharePoint connection by using an unrestricted shared account.

ANSWER: A B

Explanation:

Select Authenticate with Microsoft so that the agent identifies the employee who is making the request. Anonymous access cannot support authorization decisions based on an organizational user identity.

Maintain the appropriate SharePoint permissions on the source sites and documents. Knowledge retrieval must remain aligned with the access controls of the underlying content so that users are not granted access merely because the agent can use the source.

Copying all documents into a public website would remove the intended permission boundary. Instructions that tell the agent not to disclose restricted content are not a replacement for source authorization.

QUESTION NO: 3 - (SIMULATION)

A team is preparing to assess an agent in Copilot Studio before expanding access to additional users.

The team requires an evaluation that provides controlled, repeatable, and consistently measured test runs.

The evaluation must:

Be triggered directly by the team to control when testing occurs.

Determine success based on clearly established acceptance criteria.

Rely on a predefined set of inputs that ensures consistency across repeated runs.

ANSWER: See the explanation for the answer

Explanation:

Configure the evaluation as follows:

An on-demand evaluation gives the team control over execution timing. Defined expectations provide measurable acceptance criteria, and a curated prompt list supplies the same repeatable inputs for every evaluation run.

QUESTION NO: 4 - (HOTSPOT)

A company uses an agent flow that occasionally requires human input before continuing execution.

Some automated actions must pause until a human provides a decision or additional information. The flow must be configured to:

â€¢ Capture a human response for use in later steps.

â€¢ Continue processing within the same flow run after the response is submitted.

â€¢ Wait for a manual decision before proceeding.

You need to configure a human-in-the-loop agent flow.

Which setting should you configure for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

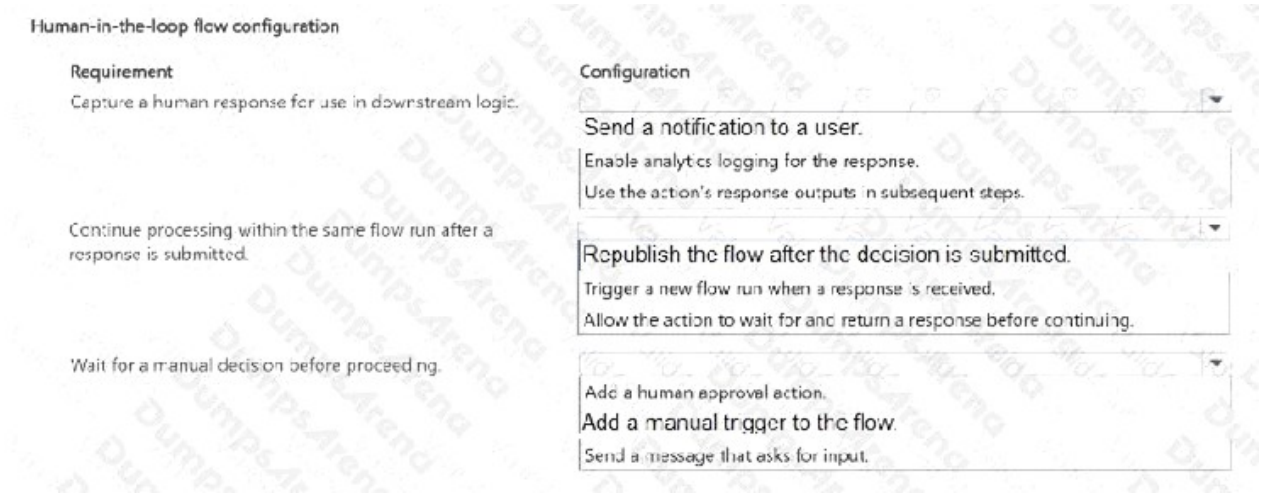

ANSWER:

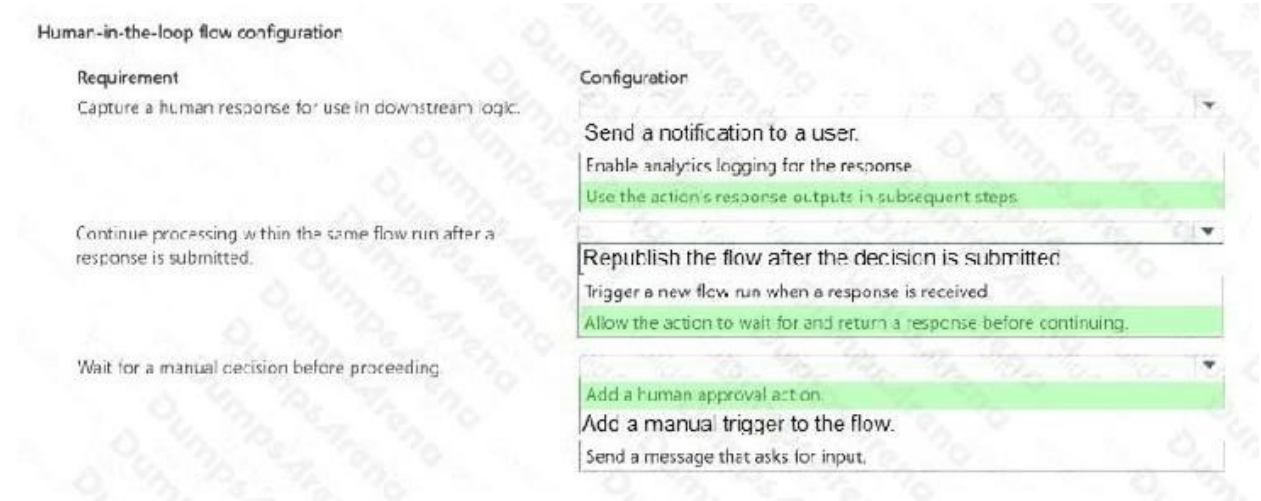

Explanation:

A human-in-the-loop agent flow is designed to suspend automation at a controlled point, collect input from a person, and then resume the same running flow using that input. For the requirement to capture information for downstream logic, configure the flow to **use the action's response outputs in subsequent steps**. The submitted response becomes output data from the human-interaction action, allowing later conditions, variable assignments, messages, and business operations to use the person's decision, comments, or supplied information.

To continue the existing flow execution after the human submits a response, configure the interaction so it will **allow the action to wait for and return a response before continuing**. This behavior maintains the flow run's context while the flow

is paused. Once a response arrives, the action completes and the following steps in that same run can execute. This is the essential wait-and-resume pattern for human-in-the-loop orchestration rather than a disconnected notification process. Microsoft describes approvals as actions that wait for an approver's response before proceeding with the workflow; see Modern approvals in Power Automate.

For a requirement that specifically calls for a manual decision before processing proceeds, configure the flow to **add a human approval action**. An approval provides a formal decision checkpoint, enabling the human participant to approve, reject, or otherwise provide the configured decision outcome. The flow can then use that returned decision to direct the remaining execution. This fits the governance purpose of human oversight: automation does not pass the decision boundary until an authorized person has acted. For further background on building approval processes and using their returned outcomes in flows, review Approvals overview in Power Automate.

QUESTION NO: 5 - (DRAG DROP)

You need to connect Operations Concierge to Fabrikam Inc. ' s Azure AI Search knowledge index while complying with security requirements.

Which configuration should you use 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. You may need to move the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point

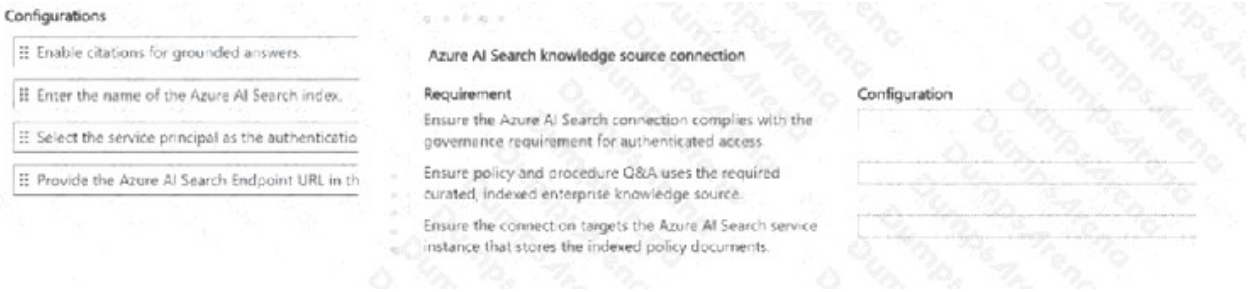

ANSWER:

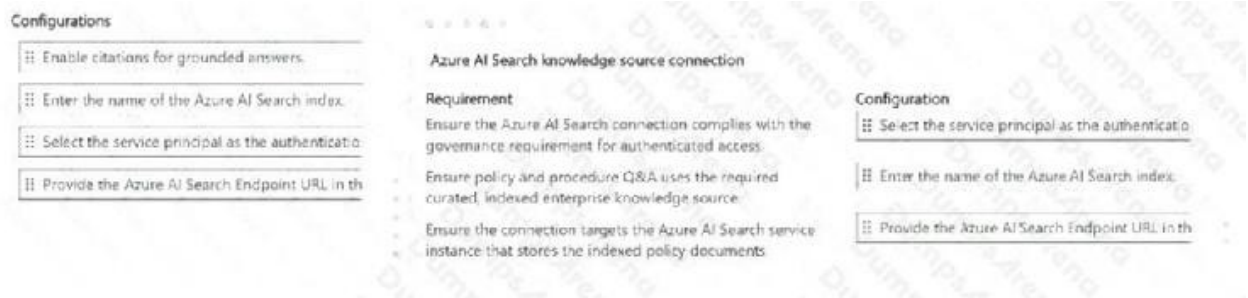

Explanation:

An Azure AI Search knowledge source connection requires separate configuration for identity, the target search service, and the particular index to query. For the governance requirement covering authenticated access, select the service principal as the authentication method. A service principal is an application identity, which supports governed, centrally managed access to Azure resources rather than tying access to an individual user. Its permissions can be assigned through Azure role-based access control and constrained to the least privilege needed by the agent connection.

For the requirement that policy and procedure questions use the curated, indexed enterprise knowledge source, enter the Azure AI Search index name. An Azure AI Search service can contain multiple indexes, so the index name is what identifies the intended curated vector or hybrid-search corpus that contains the approved policy documents. This ensures retrieval is directed to the specific enterprise knowledge collection prepared for Operations Concierge.

For the requirement to connect to the Azure AI Search service instance that stores the indexed policy documents, provide the Azure AI Search endpoint URL. The endpoint identifies the actual Azure AI Search resource, including the service host to which Copilot Studio sends search queries. Together, the endpoint URL identifies the service, the index name identifies the knowledge collection within that service, and service-principal authentication supplies the governed application identity used

to access it. Microsoft documents Azure AI Search connectivity and the use of supported connections in Copilot Studio at Add Azure AI Search as a knowledge source. Guidance on service principals and managed application identities in Microsoft Entra ID is available at Application objects and service principals.

QUESTION NO: 6

A company is building an asset-support agent in Copilot Studio.

Users enter asset identifiers in the following format: two uppercase letters, a hyphen, and six digits. For example, AB-123456.

The agent must reliably recognize the identifier from a user message and use the value in a subsequent action.

You need to configure the agent to meet the requirements.

What should you do?

A. Create a custom entity that uses a regular expression for the asset identifier.

B. Create a closed-list entity that contains every possible asset identifier.

C. Add the asset identifier format to the agent instructions.

D. Add an Adaptive Card as an agent knowledge source.

ANSWER: A

Explanation:

Create a custom entity that uses a regular expression matching the required asset identifier format. A regular-expression entity is appropriate when a value follows a predictable pattern and the agent must extract that value from natural-language input. The extracted entity value can then be stored in a variable and supplied to a subsequent tool or flow action.

A closed-list entity is suitable for a limited set of known values, such as product categories, but is not suitable for potentially many asset identifiers. A knowledge source can provide information about assets but does not reliably extract an identifier from a user message. An Adaptive Card is useful when the agent must collect structured input through a form, but it is not required when the identifier can be recognized directly from conversational text.

QUESTION NO: 7

A company is configuring an agent in Copilot Studio that uses a generative answers node inside multiple topics.

The company requires responses to meet the following requirements:

Must use an executive summary format.

Must apply only within a specific topic.

Must continue using configured knowledge sources.

Must comply with connector security controls.

You need to configure a custom prompt to meet the formatting requirements.

Solution: Configure custom instructions on the existing generative answers node at the topic level so that other topics are not affected.

Does the solution meet the goal?

A. Yes

B. No

ANSWER: A

Explanation:

Yes. Custom instructions configured on a generative answers node control how that particular node composes its grounded response. Therefore, an instruction such as “Provide an executive summary with key findings, risks, and recommended next actions” can establish the required presentation format while remaining limited to the topic that contains that node. This is the appropriate scope when the company does not want the formatting behavior applied to other topics in the agent.

Because the existing generative answers node remains in place, it continues to retrieve and ground responses using the knowledge sources configured for that node. The custom instruction supplements response generation; it does not replace the node’s knowledge retrieval configuration. Connector access and security controls also continue to be enforced by the relevant connector authentication, user permissions, and the environment’s data-loss-prevention policies. The instruction should focus on the desired output structure and not instruct the agent to bypass source selection, authentication, or policy controls. Microsoft documents that generative answers nodes can be configured with knowledge sources and instructions to tailor generated responses, including within a topic’s authoring canvas. See Generative answers in Copilot Studio and Create and edit topics in Copilot Studio.

QUESTION NO: 8

A company is building an agent in Copilot Studio that collects a customer reference number in one topic.

The agent transfers the conversation to a separate topic that retrieves account information. The customer reference number must remain available after the topic transition and throughout the current conversation.

You need to configure storage for the customer reference number.

What should you use?

A. A topic variable

B. A global variable

C. A system variable

D. An environment variable

ANSWER: B

Explanation:

Use a global variable. Global variables are available across topics for the duration of the current conversation. This allows the first topic to store the customer reference number and the account-information topic to use the same value without asking the user to enter it again.

A topic variable is limited to the topic in which it is created and is therefore unsuitable after the conversation transitions to another topic. System variables contain values supplied by Copilot Studio, such as channel or user context, and are not intended for storing a customer reference collected during the conversation. An environment variable is an application lifecycle management configuration value rather than a per-conversation value.

QUESTION NO: 9 - (SIMULATION)

A team deploys Copilot Studio solutions by using Power Platform Pipelines.

You must configure a deployment strategy that meets the following requirements:

Ensure that the solutions deploy in dev, test, and production environments in order.

Identify missing connection references or environment variables before the deployment process begins.

Manage all pipeline stages and security from a single, unified location for all linked environments.

You need to configure the pipeline to meet the requirements.

ANSWER: See the explanation for the answer

Explanation:

Configure the pipeline deployment controls as follows:

Sequential progression enforces promotion through the defined Dev, Test, and Production stages. Deployment input validation checks required environment-specific inputs, including connection references and environment variables, before a deployment starts. Centralized configuration provides a single host-pipeline location to govern stages, linked environments, and pipeline security.

QUESTION NO: 10

Which two actions should you perform before making the agent available on both channels? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

A. Configure prompt modifications to enforce tone, disclaimers, and refusal behavior at the system level.

B. Manually add disclaimers to each topic before publishing.

C. Embed the web channel and then rely on channel-level settings to enforce content moderation.

D. Configure Power Platform DLP policies to restrict unauthorized data connectors.

E. Publish the agent and then enable Responsible AI filters individually for each channel.

ANSWER: A D

Explanation:

Configure prompt modifications to enforce tone, disclaimers, and refusal behavior at the system level. is appropriate because agent-level instructions provide centralized guidance for generative orchestration and responses. These instructions can establish the agent's intended tone, response style, scope, formatting expectations, and handling of requests that should not be answered. Applying this guidance before channel deployment helps ensure that users receive consistent behavior regardless of where they access the agent, rather than relying on channel-specific implementation details. Microsoft documents how agent instructions influence the agent's behavior and response generation in Configure agent instructions.

Configure Power Platform DLP policies to restrict unauthorized data connectors. is also a necessary governance action before making an agent broadly available. Power Platform data loss prevention policies classify connectors and control which connectors can be used together, helping prevent unintended movement of organizational data to unapproved services. Because agents can use connectors, knowledge sources, and actions to retrieve or send data, validating applicable DLP policies before deployment ensures that the agent operates within the organization's data-boundary

requirements. This is particularly important when the same agent is exposed through multiple channels. See Data loss prevention policies for Microsoft's connector-governance guidance.

QUESTION NO: 11

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.

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:

Create a custom solution and add the existing agent to the solution. is required because solutions are the Power Platform application lifecycle management container used to move Copilot Studio agents and their components between environments. In the development environment, create or open an unmanaged solution, then use *Add existing* to include the agent. This establishes the solution as the source package and allows relevant dependencies, such as flows, connection references, environment variables, and custom connectors, to be identified and included as appropriate. Microsoft documents this as the supported approach for moving existing agents through environments.

Export the unmanaged solution that contains the agent. is the next required action for creating the deployable package. The maker exports the unmanaged source solution from the development environment and chooses the intended export package type. For a production deployment, a managed package is commonly selected during export, while the development solution remains unmanaged as the editable source. The exported solution can then be imported into the production environment, where configuration values and connections can be finalized before the agent is tested and published. See Use solutions to export and import agents and Solution concepts for application lifecycle management.