

Agentic AI Business Solutions Architect

Microsoft AB-100

Version Demo

Total Demo Questions: 14

Total Premium Questions: 140

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Design AI solutions	22
Topic 2, Design agentic AI solutions	66
Topic 3, Design business solutions	19
Topic 4, Design governance solutions	31
Topic 5, Mix Questions	2
Total	140

QUESTION NO: 1

A company plans to make a Microsoft Copilot Studio agent available in Microsoft Teams.

You need to prepare the agent for organizational deployment. The solution must ensure that the agent is available through Teams and that only intended users can access it.

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

NOTE: Each correct selection is worth one point.

- A. Publish the agent to the Microsoft Teams channel.
- B. Configure access to the agent for the intended Microsoft Entra ID users or groups.
- C. Export all Teams chat messages to a public website.
- D. Assign all employees the Power Platform administrator role.
- E. Create a separate Dataverse environment for every Teams user.

ANSWER: A B

Explanation:

Publish the agent to the Microsoft Teams channel makes the agent available for use in Teams. Publishing to a channel is required in addition to building and testing the agent in Copilot Studio.

Configure access to the agent for the intended Microsoft Entra ID users or groups limits the audience to authorized users. Access control should be designed together with permissions to underlying data sources and actions so that users cannot retrieve or update data beyond their existing authorization.

Publishing and access management are separate deployment activities: the Teams channel provides the user experience, while identity-based access controls govern who can use the agent. See [Publish an agent to Microsoft Teams](#) and [Share agents in Microsoft Copilot Studio](#).

QUESTION NO: 2 - (DRAG DROP)

A company has an AI solution that uses a Microsoft Copilot Studio agent.

You need to monitor the agent's performance. The solution must meet the following requirements: Monitor the agent's telemetry in near-real-time (NRT).

Download transcripts of full conversations. Monitor the agent's usage and performance.

What should you use for each requirement? To answer, drag the appropriate options to the correct requirements. Each option may be used once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Options	Answer Area
Application Insights	Monitor the agent's telemetry in NRT:
Copilot Studio	Download transcripts of full conversations:
Log Analytics	Monitor the agent's usage and performance:
Microsoft Power Apps	

ANSWER:

Options	Answer Area	Option
Application Insights	Monitor the agent's telemetry in NRT:	Application Insights
Copilot Studio	Download transcripts of full conversations:	Copilot Studio
Log Analytics	Monitor the agent's usage and performance:	Copilot Studio
Microsoft Power Apps		

Explanation:

Application Insights is the correct service for monitoring a Copilot Studio agent's telemetry in near-real-time. When an agent is configured to send telemetry to Azure Application Insights, operational events can be reviewed shortly after they occur. This supports active monitoring and investigation of the agent's behavior, including conversation-related telemetry, errors, latency, dependencies, traces, and custom events. Because this telemetry is available for live operational analysis and can be queried or visualized through Azure Monitor tooling, it is the appropriate choice when the requirement specifically calls for near-real-time telemetry. Microsoft documents the integration and its telemetry capabilities in the [Copilot Studio Application Insights documentation](#).

Copilot Studio is the correct place to download transcripts of complete conversations. The agent's transcript capability is designed around individual sessions and preserves the end-to-end exchange between a user and the agent. This gives administrators and makers an agent-aware way to review and export conversation records for auditing, quality review, troubleshooting, and analysis. The transcript experience includes the conversational context needed to understand how an interaction proceeded, rather than requiring a reviewer to reconstruct it from lower-level diagnostic events. Details on reviewing and downloading conversation transcripts are provided in [Copilot Studio conversation transcript analytics](#).

Copilot Studio is also the correct choice for monitoring the agent's usage and performance. Its built-in analytics are tailored to agent outcomes and adoption, such as session and user activity, engagement, resolution-related measures, escalation patterns, and trends over time. These metrics help stakeholders assess whether the agent is being used, whether it is handling conversations effectively, and where further optimization may be needed. This is the agent-level business and performance monitoring experience described in the [Copilot Studio analytics summary documentation](#).

QUESTION NO: 3

A company has an AI business solution that uses Microsoft Copilot Studio agents. You need to recommend prompt best practices to improve the effectiveness of agent interactions. Which two actions should you include in the recommendation? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Measure system resource usage during prompt processing.
- B. Analyze the prompt length distribution.
- C. Use clear and specific instructions in the prompts.
- D. Regularly test and refine the prompts based on user input.
- E. Track the duration of the average user session.

ANSWER: C D

Explanation:

Use clear and specific instructions in the prompts because an agent needs explicit direction about its role, task, expected output, constraints, tone, and any required source-grounding behavior. Clear instructions reduce ambiguity and help produce responses that are more relevant, consistent, and aligned with the organization's intended user experience. In Copilot Studio, instructions are a key mechanism for shaping how an agent responds, including guidance on how it should use available knowledge and conduct conversations.

Regularly test and refine the prompts based on user input because effective prompt design is an iterative process. Testing with representative user questions and reviewing actual conversation outcomes identifies gaps such as unclear instructions,

missing context, or response patterns that do not meet business requirements. The team can then refine the agent's instructions and evaluate the changes using realistic scenarios. This feedback-driven approach improves response quality over time and helps ensure the agent remains effective as user needs and business content evolve. Microsoft guidance for prompt engineering emphasizes being clear and specific, while Copilot Studio guidance supports testing and iterating agent behavior before and after deployment.

References: [Prompt engineering techniques - Azure OpenAI](#); [Generative answers in Microsoft Copilot Studio](#).

QUESTION NO: 4

You need to recommend a solution to integrate a Microsoft Copilot agent with a Microsoft Dynamics 365 Contact Center chat channel.

The agent must respond to customer questions and hand off the conversation to a live customer service representative when the customer requests an escalation.

What should you recommend?

- A. Configure the Escalate topic.
- B. Configure the Conversation Start topic.
- C. Build an agent flow.
- D. Call a Microsoft Power Automate connector.
- E. Configure a skill.

ANSWER: A

Explanation:

Configure the Escalate topic. Copilot Studio includes a system Escalate topic specifically for recognizing when a customer asks to speak to a person, such as by requesting an agent, representative, or human support. This topic can be customized to match the organization's escalation language and to invoke the appropriate handoff behavior for the connected engagement channel.

For a Dynamics 365 Contact Center chat deployment, the handoff process preserves the customer's conversation context and routes the engagement to a live customer service representative. The representative can use the history of the bot interaction to continue assisting the customer without requiring them to restate the issue. This directly fulfills both requirements: the Copilot agent handles routine customer questions, and it transfers the conversation when escalation is requested.

Microsoft documents the Escalate system topic as the mechanism for configuring agent-to-human handoff. Its behavior and available routing configuration depend on the channel to which the agent is connected; Dynamics 365 customer service and contact-center channels support this live-agent engagement model. See [Configure escalation in Microsoft Copilot Studio](#) and [Configure Copilot Studio bots in Dynamics 365 Customer Service](#).

QUESTION NO: 5 - (HOTSPOT)

A company has a Microsoft Copilot Studio prompt-and-response agent.

You need to ensure that the agent meets the following requirements:

Provides effective and relevant responses

Provides conversational outcomes

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

NOTE: Each correct selection is worth one point.

Provides effective and relevant responses:

Generated answer rate and quality
Reactions
Tool use

Provides conversational outcomes:

Satisfaction
Tool use
Topics by outcome

ANSWER:

Provides effective and relevant responses:

Generated answer rate and quality
Reactions
Tool use

Provides conversational outcomes:

Satisfaction
Tool use
Topics by outcome

Explanation:

Use **Generated answer rate and quality** to measure whether the prompt-and-response agent provides effective and relevant responses. Prompt-and-response agents rely on generative AI to formulate answers from available knowledge and instructions. Therefore, evaluating generated-answer performance is the appropriate way to understand whether the agent is responding consistently and whether those responses meet an acceptable quality threshold. This metric gives the solution team an operational view of the agent's generative-answer behavior and helps them identify where knowledge, instructions, grounding, or orchestration may need refinement. It directly supports ongoing monitoring of response usefulness in an AI-first conversational experience. Microsoft documents generative-answer analytics and quality monitoring in its [Copilot Studio analytics overview](#).

Use **Topics by outcome** to measure conversational outcomes. An outcome is the meaningful end state reached through a conversation or topic interaction, such as successful resolution, escalation, abandonment, or another configured completion result. Viewing topics grouped by outcome enables architects and makers to determine whether the agent's conversations are achieving intended business goals, rather than merely producing a response. This supports analysis of the customer journey through the agent and helps identify where topic design, escalation handling, or follow-up actions need improvement. Outcome analysis is especially valuable when the agent is expected to resolve support requests, guide users through a process, or complete a business task. Microsoft describes topic-level analysis and outcome reporting in the [Copilot Studio topic analytics documentation](#).

QUESTION NO: 6

A company uses Microsoft Foundry agents.

You need to ensure that an agent can dynamically use external tools at runtime without updating the agent. What should you include in the solution?

- A. a Model Context Protocol (MCP) server
- B. a Microsoft Foundry hub
- C. Microsoft Copilot Studio

ANSWER: A

Explanation:

An MCP server is the appropriate solution because the Model Context Protocol provides a standardized client-server mechanism through which a Microsoft Foundry agent can connect to external capabilities as tools. The server publishes the tools it exposes, including their names, descriptions, input schemas, and execution interfaces. At runtime, the agent can discover those available tools and invoke the suitable one based on the user's request and the model's tool-selection process.

This decouples the agent from individual external tool implementations. Tool capabilities can be managed on the MCP server, allowing the connected agent to use the server's currently available tools without requiring the agent itself to be redesigned for each tool integration. MCP is therefore well suited to scenarios where external systems, APIs, and actions need to be made available dynamically through a consistent protocol.

Microsoft Foundry supports MCP as an agent tool integration, enabling agents to access remote MCP servers and use their exposed functionality during execution. For implementation details, see [Model Context Protocol \(MCP\) tools in Microsoft Foundry](#) and the [Microsoft Foundry Agent Service overview](#).

QUESTION NO: 7

A company plans to deploy a Microsoft Foundry agent

You need to recommend an application lifecycle management (ALM) process to ensure that the agent evaluates against baseline accuracy metrics before being deployed. What should you recommend?

- A. Use Observability in Foundry Control Plane with evaluation and drift monitoring.
- B. Deploy each new agent version directly to production.
- C. Enable Application Insights and use Azure Monitor.
- D. Configure GitHub Actions for new agent versions.

ANSWER: A

Explanation:

Use Observability in Foundry Control Plane with evaluation and drift monitoring. This approach supports an ALM quality gate in which the agent is evaluated on representative test data before promotion to production. Azure AI Foundry evaluations allow teams to define and run evaluators that measure application quality against task-specific criteria, including accuracy-related metrics where appropriate. Evaluation results can be compared with an established baseline and used as release evidence, so a version that does not meet the required threshold is not promoted.

Foundry observability also provides the operational context needed to maintain that quality standard after release. It captures traces and application behavior that can be analyzed through evaluations and monitoring. This enables the team to detect changes in agent behavior, data, or quality over time and investigate regressions using production evidence. Combining predeployment evaluation with ongoing monitoring creates a repeatable lifecycle process: establish a baseline, evaluate each candidate version, approve deployment only when the quality bar is met, and continue observing the deployed agent for degradation or drift signals.

Microsoft documents the evaluation workflow in [Azure AI Foundry evaluation concepts](#) and describes tracing, monitoring, and evaluation capabilities in [Azure AI Foundry observability concepts](#).

QUESTION NO: 8 - (HOTSPOT)

You are designing a testing solution for Microsoft Copilot Studio agents.

You need to validate prompt engineering best practices to ensure that the agents generate accurate and contextually relevant responses. Which prompt validation techniques and metrics should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Prompt validation techniques:

- Exclude domain-specific terminology from the prompts.
- Use prompts that have varied phrasing
- Use only simple, one-word prompts.

Metrics:

- The number of words generated per response
- Response relevance and accuracy
- The response generation time

ANSWER:

Explanation:

Prompt validation for a Microsoft Copilot Studio agent should reflect how people actually communicate with an agent in production. Users can express the same request through different wording, sentence structures, synonyms, levels of detail, and terminology. Selecting **Use prompts that have varied phrasing** provides a practical validation technique because it tests whether the prompt design consistently recognizes the same underlying intent and produces an appropriate result despite ordinary language variation. This makes validation more representative of real conversations and helps establish that the agent is robust rather than dependent on one narrowly worded test input.

The metric should directly assess the stated objective: responses must be both accurate and contextually relevant. Selecting **Response relevance and accuracy** measures whether the generated response addresses what the user asked, uses the supplied conversation and business context appropriately, and provides factually or operationally correct information. This can be evaluated with representative test cases, expected-answer criteria, and reviewer assessment of whether the response is useful for the intended scenario. Together, varied-phrasing tests and relevance-and-accuracy evaluation validate the quality of the agent's prompt behavior across realistic user inputs. Microsoft's guidance on testing copilots and assessing generative AI output supports using representative test inputs and evaluating answer quality: [Test your agent in Microsoft Copilot Studio](#) and [Evaluate generative AI applications](#).

QUESTION NO: 9

You are designing an AI business solution that contains the following components:

- A Microsoft Power Automate workflow
- A Microsoft Copilot Studio agent
- A Microsoft Dataverse database
- A Microsoft Power Apps app

As part of the application lifecycle management (ALM) process, you plan to package the components, so that they can be deployed to other environments as a group. You need to recommend a solution that supports versioning, dependencies, and deployments. What should you include in the recommendation?

- A. Azure DevOps
- B. GitHub Actions
- C. Microsoft Power Platform solutions

ANSWER: C

Explanation:

Microsoft Power Platform solutions are the correct recommendation because a solution is the native ALM packaging mechanism for Power Platform components. A solution can contain Power Apps apps, Power Automate cloud flows, Dataverse tables and other customizations, and Copilot Studio agents. Packaging these assets together establishes a single deployable unit, allowing the application to move consistently between development, test, and production environments.

Solutions also provide dependency tracking. When a component requires other components, such as a flow relying on Dataverse tables or connection references, those relationships are maintained as part of the solution architecture. Solution version numbers can be set and updated as the application evolves, supporting controlled release management. Teams typically build in unmanaged solutions during development, then export managed solutions for controlled deployment to downstream environments. This approach helps preserve application integrity and supports repeatable ALM practices.

For implementation details, Microsoft documents solution concepts, component management, dependencies, and solution layering in [Solution concepts for application lifecycle management](#). Microsoft also describes how solutions are used to manage and deploy Copilot Studio agents in [Use solutions with Microsoft Copilot Studio](#).

QUESTION NO: 10

You need to recommend a Microsoft Power Platform business solution that consolidates data from multiple internal and external data sources. The solution must meet the following requirements: Provide the data as a centralized source for multiple AI systems, including Microsoft Copilot Studio agents, Dynamics 365 applications, and external AI models.

Support built-in data classification and protection policies. Provide data for grounding and analytics.

What should you include in the recommendation?

- A. Microsoft Dataverse
- B. Azure Data Lake Storage
- C. a Microsoft Power BI semantic model
- D. Azure Cosmos DB

ANSWER: A

Explanation:

Microsoft Dataverse is the appropriate recommendation because it is the Power Platform's managed business-data platform and the native data foundation for Dynamics 365 applications. It centralizes operational business information in relational tables with standardized metadata, relationships, business rules, and security, enabling multiple applications and AI experiences to use consistent, governed data. Copilot Studio agents can use Dataverse knowledge and data sources to retrieve current business information for grounding, while Dynamics 365 uses Dataverse as its underlying business-data store. Dataverse also exposes data through supported APIs and integration capabilities, allowing appropriately authorized external systems and AI models to access data.

Its governance model includes role-based security, column-level security, auditing, and environment-level administration. These controls support protected handling of business data and work with wider Microsoft Purview and Power Platform governance capabilities. For analytical requirements, Dataverse integrates with Power BI and supports Azure Synapse Link, allowing operational data to be made available for analytics without abandoning the governed business-data layer. Microsoft describes Dataverse as a secure cloud-based data platform for storing and managing data used by business applications. See the [Microsoft Dataverse overview](#) and [Power Platform security documentation](#).

QUESTION NO: 11 - (DRAG DROP)

A company has a cloud-based AI solution that uses Azure OpenAI models.

You need to design a monitoring solution that meets the following requirements: Monitors performance metrics and operational health for the models

Monitors AI apps and agents for compliance Uses Azure-native capabilities

Minimizes development effort

What should you use for each requirement? To answer, drag the appropriate options to the correct requirements. Each option may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Options	Answer Area
☐ Azure API Management ☐ Azure Monitor ☐ Azure Policy ☐ Azure Stream Analytics ☐ Microsoft Defender ☐ Microsoft Purview	Monitors AI app and agents for compliance: Monitors performance metrics and operational health:

ANSWER:

Options	Answer Area
☐ Azure API Management ☐ Azure Monitor ☐ Azure Policy ☐ Azure Stream Analytics ☐ Microsoft Defender ☐ Microsoft Purview	Monitors AI app and agents for compliance: Microsoft Purview Monitors performance metrics and operational health: Azure Monitor

Explanation:

Use **Microsoft Purview** to monitor AI apps and agents for compliance, and use **Azure Monitor** to monitor performance metrics and operational health. This design uses managed Microsoft services, which satisfies the Azure-native requirement while avoiding the effort of building a custom telemetry, governance, and reporting solution.

Azure Monitor is the central Azure observability service. It collects and analyzes platform metrics, resource logs, application telemetry, and diagnostic data. For a solution based on Azure OpenAI models, it supports tracking operational signals such as availability, request volume, latency, errors, throttling-related indicators, and the health of dependent Azure resources. Alert rules, Log Analytics queries, workbooks, and dashboards allow teams to identify and respond to operational conditions without implementing a separate monitoring platform. Microsoft describes these capabilities in the [Azure Monitor overview](#).

Microsoft Purview is appropriate for the compliance-monitoring requirement because it is Microsoft's governance and compliance solution. It helps organizations discover and govern data, apply classifications and policies, assess risk, and maintain visibility into how sensitive or regulated information is managed. Those capabilities support compliance oversight for AI applications and agents that use organizational data and Azure OpenAI services. Purview's managed governance features reduce the amount of custom development needed to establish compliance processes and reporting. More information is available in the [Microsoft Purview documentation](#) and the [Microsoft Purview AI security posture management overview](#).

QUESTION NO: 12 - (SIMULATION)

Which Copilot Studio analytics metrics should you recommend to assist the executives with their specific responsibilities? To answer, drag the appropriate metrics to the correct executives. Each metric may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer Area

Metrics

- Effectiveness
- Satisfaction
- Tool use
- Use

Answer Area

CFO concerns about Copilot Studio credit usage

CTO concerns about poor feedback on AI agent responses

ANSWER: See the explanation for the answer

Explanation:

Configure the metrics as follows:

Reasoning: Effectiveness analytics help assess conversation outcomes, such as resolved versus abandoned interactions, which supports evaluating whether credit consumption is producing useful results. Satisfaction analytics measure user feedback and sentiment about agent responses, directly addressing poor response feedback.

QUESTION NO: 13

A company has several Microsoft Foundry applications in production.

The finance team needs to receive an alert when monthly Azure spending approaches an approved limit. The architects must also be able to allocate the costs of each application to the appropriate business unit.

Which two actions should you include in the design? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Configure Azure Cost Management budgets and alerts.
- B. Apply Azure resource tags for the application and business unit.
- C. Configure an Azure AI Search synonym map.
- D. Create a Microsoft Purview retention policy.
- E. Use an unmanaged Power Platform solution in production.

ANSWER: A B

Explanation:

Configure **Azure Cost Management budgets and alerts** to define spending thresholds and notify stakeholders when actual or forecasted costs approach the approved limit. Budgets support proactive cost governance and allow the organization to respond before monthly consumption exceeds the planned amount.

Apply **Azure resource tags** to resources that support each application. Tags such as CostCenter, BusinessUnit, Application, and Environment provide metadata that can be used to group, filter, and analyze costs in Azure Cost Management. Consistent tagging enables chargeback or showback reporting across the application portfolio.

For more information, see [Create and manage Azure budgets](#) and [Use tags to organize Azure resources](#).

QUESTION NO: 14

A company has Microsoft 365 Copilot agents.

You need to design a security solution for the agents. The solution must meet the following requirements: Identify and mitigate potential risks that relate to AI use.

Protect AI apps and the sensitive data processed or generated by the agents.

Support responsible AI governance by retaining and logging interactions, detecting policy violations, and investigating incidents.

Which two components should you include in the design? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Microsoft Purview
- B. Azure AI Content Safety
- C. role-based access control (RBAC) in Microsoft Foundry
- D. Microsoft Defender

ANSWER: A D

Explanation:

Microsoft Purview and **Microsoft Defender** provide the complementary compliance and security capabilities required for a Microsoft 365 Copilot agent solution. Microsoft Purview supports data security and governance for AI interactions through capabilities such as sensitivity labels, data loss prevention, audit, retention, eDiscovery, and communication compliance. These controls help protect sensitive information that agents access, process, or generate, while preserving records needed for monitoring, policy enforcement, and investigations. Purview can also apply controls to Copilot-related data flows according to the organization's data classification and compliance requirements.

Microsoft Defender provides security operations capabilities for identifying, investigating, and responding to threats affecting identities, devices, cloud applications, and collaboration services used with AI. Defender signals and incident investigation features help security teams detect suspicious activity and mitigate risks associated with AI use, including risks involving compromised accounts, unsafe application behavior, or data exposure. Together, the two services establish layered protection: Purview governs information and compliance evidence, while Defender supplies threat detection and incident response. See [Microsoft Purview documentation](#) and [Microsoft Defender XDR documentation](#).