References:

https://learn.microsoft.com/en-us/microsoft-copilot-studio/publication-add-bot-to-microsoft-teams

https://www.microsoft.com/insidetrack/blog/how-our-employees-are-extending-enterprise-ai-with-custom-retrieval-agents/

https://learn.microsoft.com/en-us/microsoft-365-copilot/extensibility/copilot-studio-experience

Question 8

Domain: Design AI-powered business solutions

Scenario: A development team is building a complex, multi-step agent using the Microsoft Agent Framework. This agent needs to receive a high-level user request, break it down into the required sequence of internal API calls and knowledge base, and manage the execution flow of these steps to achieve the final outcome.

Within the Agentic Core, which specific sub-component is primarily responsible for analyzing the user's intent, determining the optimal sequence of required tools/APIs and data sources (RAG), and managing the logical execution flow of these steps?

A. The Safety System

B. The Model Context Protocol (MCP) Interface

C. The Large Language Model (LLM)

D. The Planner and Orchestrator right

Explanation:

Correct Answer: D

Option D: The Planner and Orchestrator is correct because this component is the "brain" of the Agentic Core. Its function is to take the user prompt, use the LLM to assist in decomposition, build a

logical plan (Planner) of steps to solve the request, and then manage the execution of those steps using the available Tools and data sources (Orchestrator).

Option A: The Safety System is incorrect because the Safety System's primary role is to enforce security, compliance, and responsible AI policies by monitoring the inputs and outputs, not to decompose the task or manage the operational sequence.

Option B: The Model Context Protocol (MCP) Interface is incorrect because the MCP is the standardized protocol used for communication between the Agentic Core and the Tools/APIs; it is not the component responsible for creating or managing the execution plan.

Option C: The Large Language Model (LLM) is incorrect because while the LLM provides the reasoning and natural language capabilities, the LLM itself is a resource utilized by the Planner to generate the plan; the Planner/Orchestrator is the component that handles the management and execution of the plan.

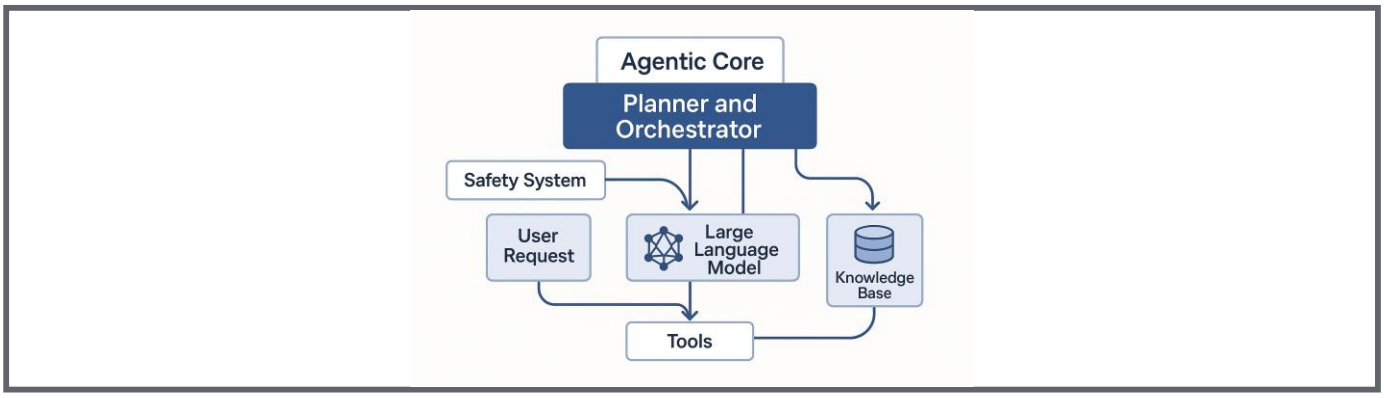

References:

https://learn.microsoft.com/en-us/agent-framework/media/agent.svg

https://techcommunity.microsoft.com/blog/educatordeveloperblog/ai-agents-planning-and-orchestration-with-the-planning-design-pattern---part-7/4399204

https://www.microsoft.com/en-us/microsoft-365/planner/microsoft-planner

Question 9 Unattempted

Domain: Design AI-powered business solutions

An AI Agent needs to securely access Azure resources and external services (like an MCP Server) within a Microsoft-centric agentic solution. The diagram illustrates a key mechanism for managing the agent's identity and permissions.

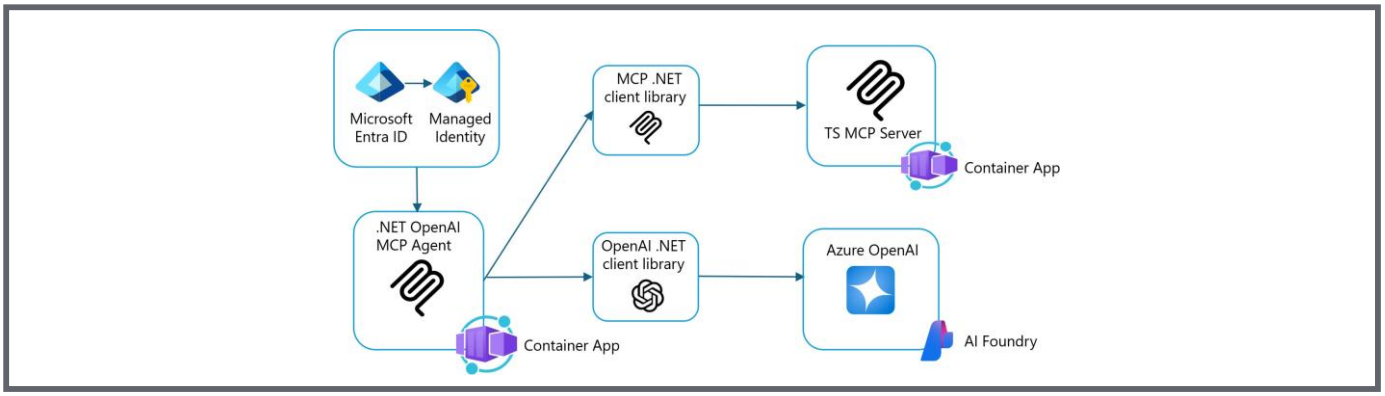

Note: Drag and drop the following components into the correct sequential order to represent the authentication flow for the .NET OpenAI MCP Agent accessing secured resources, as depicted in the diagram

Correct Answer

1. C. Microsoft Entra ID
2. A. Managed Identity
3. B. .NET OpenAI MCP Agent

Explanation:

Correct Answers: C, A and B

Step 1 → C. Microsoft Entra ID

Step 2 → A. Managed Identity

Step 3 → B .NET OpenAI MCP Agent

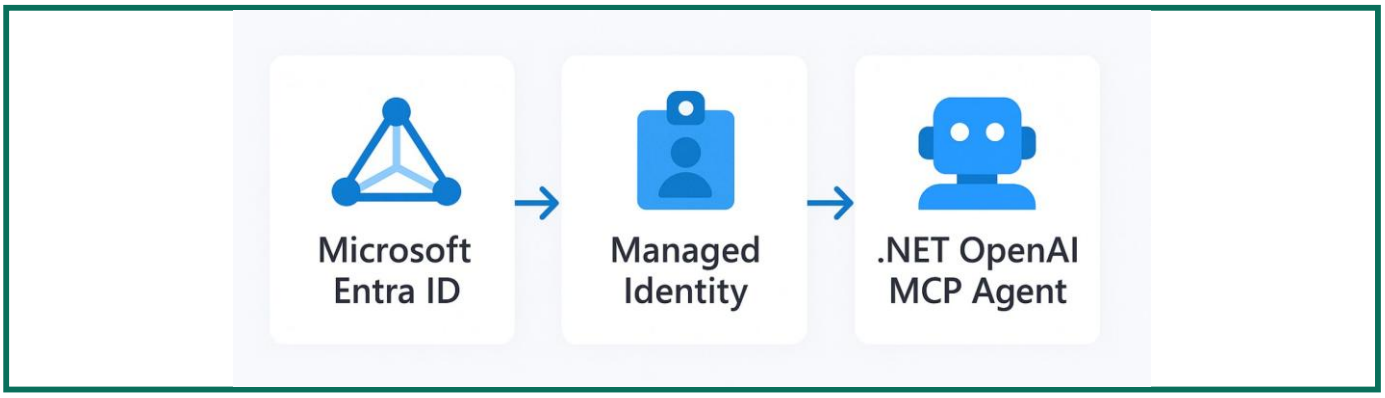

Step 1: Microsoft Entra ID: This is the foundational cloud-based identity and access management service that centrally manages user identities, groups, and permissions. For Azure resources, it serves as the ultimate authority for issuing and validating identities. In the context of Managed Identities, Microsoft Entra ID is where the identity itself is registered and managed.

Step 2: Managed Identity: A Managed Identity is an identity registered with Microsoft Entra ID that Azure resources (like the Container App hosting the .NET OpenAI MCP Agent) can use. It eliminates the need for developers to manage credentials directly. Microsoft Entra ID provisions and manages this identity. The Agent's hosting environment then uses this Managed Identity to authenticate to other Azure services (like Azure OpenAI) or external services that trust Entra ID.

Step 3: .NET OpenAI MCP Agent: The .NET OpenAI MCP Agent, deployed within an Azure Container App, is configured to use the Managed Identity assigned to its hosting environment. This allows the Agent to automatically obtain access tokens from Microsoft Entra ID (via the Managed Identity) and present them when making calls to secured services, ensuring that its requests are authenticated and authorized without hardcoding secrets.

References:

https://learn.microsoft.com/en-us/azure/developer/ai/intro-agents-mcp

https://learn.microsoft.com/en-us/azure/developer/ai/build-mcp-server-ts?tabs=github-codespaces

https://learn.microsoft.com/en-us/azure/developer/ai/build-openai-mcp-server-dotnet?tabs=github-codespaces

Question 10

Domain: Design AI-powered business solutions

The .NET OpenAI MCP Agent needs to interact with a custom business service exposed via a "TS MCP Server" (Tool Service Model Context Protocol Server). The diagram shows the communication path.

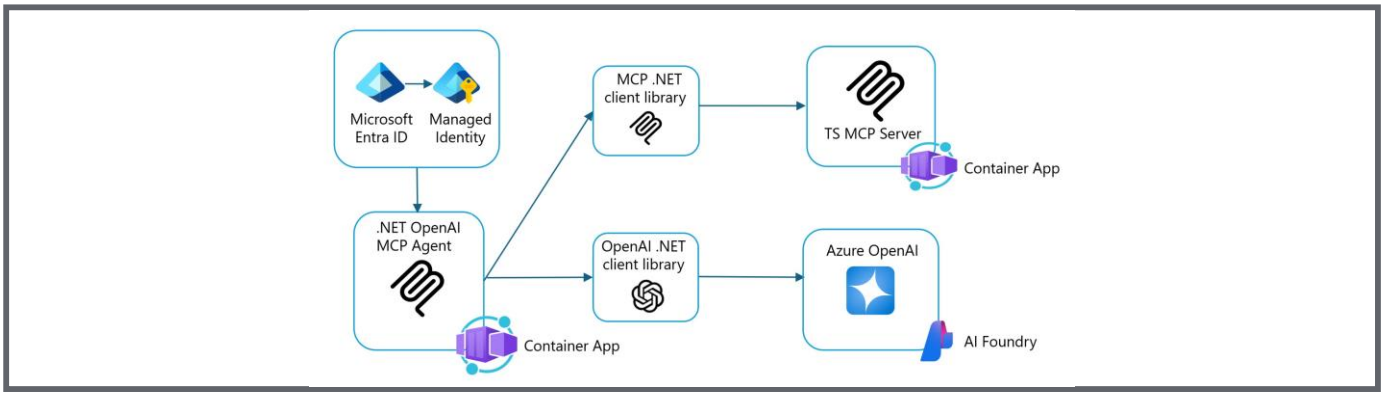

Drag and drop the following components into the correct sequential order, starting from the .NET OpenAI MCP Agent and ending with the TS MCP Server, to illustrate the communication path for invoking a tool/service.

Correct Answer

1. B. .NET OpenAI MCP Agent
2. A. MCP .NET client library
3. C. TS MCP Server

Explanation:

Correct Answers : B, A and C

Step 1 → B .NET OpenAI MCP Agent

Step 2 → A. MCP .NET client library

Step 3 → C. TS MCP Server

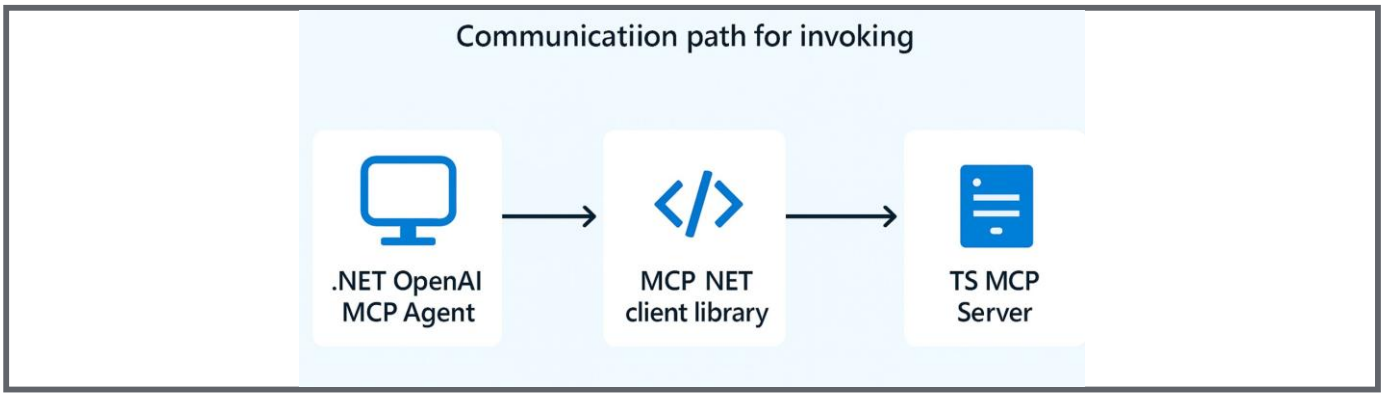

Step 1: .NET OpenAI MCP Agent: This is the initiating component. When the Agent's Planner (its "brain") determines that a specific task requires calling an external tool or service, it prepares the necessary arguments and initiates the call.

Step 2: MCP .NET client library: The .NET OpenAI MCP Agent does not directly call the raw API endpoint. Instead, it utilizes an MCP .NET client library. This library serves as an abstraction layer, handling the complexities of the Model Context Protocol (MCP) by translating the agent's high-level request into the standardized MCP format for communication with the tool service. This ensures consistent interaction regardless of the tool's underlying implementation.

Step 3: TS MCP Server: The MCP .NET client library then sends the MCP-formatted request to the TS MCP Server (Tool Service Model Context Protocol Server). This server, which exposes the custom business service capabilities, receives the request, processes it, performs the desired action (e.g., retrieving data, updating records), and then sends a response back through the MCP client library to the agent.

References:

https://learn.microsoft.com/en-us/azure/developer/ai/intro-agents-mcp

https://learn.microsoft.com/en-us/azure/developer/ai/build-mcp-server-ts?tabs=github-codespaces

https://learn.microsoft.com/en-us/azure/developer/ai/build-openai-mcp-server-dotnet?tabs=github-codespaces

Question 11 Unattempted

Domain: Deploy AI-powered business solutions

An organization wants to deploy a new tool that their AI agent can use. This tool, an "MCP Tool Service," is implemented using Azure Functions to leverage serverless capabilities and serve as a remote MCP server, as described by Microsoft's guidance.

Note: Drag and drop the following components into the correct sequential order to represent the architectural flow, from a user's prompt to the execution of a tool service hosted on Azure Functions as a remote MCP server, as per Microsoft's recommendations.

Correct Answer

1. D. User Prompt
2. B. Agent (e.g., .NET OpenAI MCP Agent)
3. C. Model Context Protocol (MCP) Definition
4. A. Azure Function App (hosting MCP Server)

Explanation:

Correct Answers : D, B, C and A

Step 1 - D. User Prompt

Step 2 - B. Agent (e.g., .NET OpenAI MCP Agent)

Step 3 - C. Model Context Protocol (MCP) Definition

Step 4 - A. Azure Function App (hosting MCP Server)

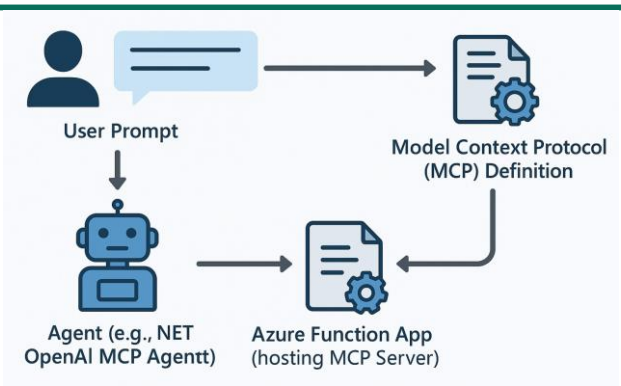

Step 1: User Prompt: The entire agentic workflow begins when an end-user provides a natural language prompt or request to the AI agent, specifying a task they want to accomplish.

Step 2: Agent (e.g., .NET OpenAI MCP Agent): The Agent receives the user's prompt. Its internal components, particularly the Planner and Orchestrator (often working with an LLM), analyze the intent of the prompt and determine if any external tools are required to fulfill the request.

Step 3: Model Context Protocol (MCP) Definition: Once the Agent identifies that a tool is needed, it consults the MCP Definition for that specific tool. This definition (often registered in an Agent Manifest or a tool registry) provides the agent with a standardized, machine-readable description of the tool's capabilities, its input parameters, and how to invoke it via the Model Context Protocol. This is crucial for the agent to know what the tool can do and how to call it.

Step 4: Azure Function App (hosting MCP Server): Based on the MCP Definition, the Agent constructs an MCP-compliant request and sends it. The Azure Function App is where the actual MCP Server is hosted. It receives this request, processes it by executing the underlying business logic of the tool (e.g., checking inventory, updating a database), and then returns an MCP-compliant response back to the Agent, allowing the Agent to complete the user's request.

Reference:

https://techcommunity.microsoft.com/blog/appsonazureblog/build-ai-agent-tools-using-remote-mcp-with-azure-functions/4401059

Question 12 Unattempted

Domain: Deploy AI-powered business solutions

Your organization is planning to deploy a suite of AI-powered agents and applications across multiple departments. You want to ensure the applications follow a strong ALM (Application Lifecycle Management) strategy to maintain quality, traceability, and governance throughout development, testing, deployment, and updates.

According to Microsoft's best practices for ALM in Dynamics 365 and Power Platform applications, which FOUR of the following practices should be adopted as part of a robust ALM strategy? (Select 4)

Note: Drag the correct answer and drop into the corresponding answer area.

Correct Answer

A. Use version control and branching strategies in a system like Azure DevOps to manage code and solution changes throughout development and release cycles

C. Use managed solutions for non-development environments (e.g., Test, Production) and unmanaged solutions only in development environments

D. Configure release pipelines and build automation so that deployable packages are tested in non-production environments before being marked as release candidates for production

E. Use multiple development environments in parallel to support the current production version (for patches) while developing the next major version

Explanation:

Correct Answers: A, C, D and E

Option A: Use version control and branching strategies in a system like Azure DevOps to manage code and solution changes throughout development and release cycles is correct because – Version control (like Git in Azure DevOps) is fundamental to ALM. It provides traceability, enables collaborative development, supports rollback capabilities, and is essential for managing changes across different development and release cycles.

Option C: Use managed solutions for non-development environments (e.g., Test, Production) and unmanaged solutions only in development environments is correct because this is a critical best practice in Power Platform ALM. Unmanaged solutions provide flexibility for active development, allowing direct component editing. Managed solutions are deployed to non-development environments (like Test, Staging, and Production) to ensure governance, controlled deployment, and proper cleanup/uninstallation of solution components.

Option D: Configure release pipelines and build automation so that deployable packages are tested in non-production environments before being marked as release candidates for production is correct because automation through CI/CD (Continuous Integration/Continuous Delivery) pipelines in tools like Azure DevOps is key to a robust ALM strategy. It ensures consistent builds, automates testing, reduces manual errors, and provides a reliable process for promoting validated solution packages through environments.

Option E: Use multiple development environments in parallel to support the current production version (for patches) while developing the next major version is correct because this parallel

environment strategy (e.g., one environment/branch for hotfixes for the current production version and another for developing the next major feature release) is crucial for agility and stability. It allows critical patches to be developed and deployed quickly without being impacted by potentially unstable new features under development, thus minimizing disruption to the live system.

Option B: Deploy all customizations directly to the production environment to speed up innovation and avoid re-testing in sandbox environments is incorrect because this practice violates core ALM principles. Deploying directly to production without proper testing in lower environments (e.g., dev, test, UAT) significantly increases the risk of introducing errors, causing system instability, and impacting business operations. All changes should follow a controlled, validated path.

References:

https://learn.microsoft.com/en-us/dynamics365/guidance/implementation-guide/application-lifecycle-management-product

https://learn.microsoft.com/en-us/power-platform/alm/overview-alm

https://learn.microsoft.com/en-us/power-platform/alm/devops-build-tools

Question 13 Unattempted

Domain: Deploy AI-powered business solutions

Your company is onboarding employees to use Microsoft Copilot effectively for everyday productivity tasks. Many users are unsure which type of prompt to use for different activities such as creating new content, catching up on work, asking questions, or editing existing files.

You decide to test their understanding by providing a matching exercise that aligns each prompt type with a real-world example. You need to match each Copilot prompt type with the correct example or use cases.

Note: This is a Dra-Drop Matching question and for this you need to match prompt types with the correct examples by drag and drop the appropriate service into the corresponding answer area.

Correct Answers

A. Catch up

"What questions were asked during the meeting?"

B. Create

"Create a short presentation about time management"

C. Ask

"Give me ideas for a team building activity"

D. Edit

"Add an image of a target with arrows to this slide"

Explanation:

Correct Answers : 1-B, 2-C, 3-D and 4-A

Step 1 → Option B (Catch up → "What questions were asked during the meeting?")

"Catch up" prompts are designed to quickly bring users up to speed on information they might have missed from meetings, emails, chats, or documents. For example, Copilot in applications like Teams or Outlook can efficiently extract key information, such as questions asked or decisions made, enabling users to get a concise overview without needing to review all content manually. This example directly asks for specific information from a past event, fitting the "Catch up" category perfectly.

Step 2 → Option C (Create → "Create a short presentation about time management.")

"Create" prompts instruct Copilot to generate entirely new content from scratch. This can include various formats such as documents, emails, presentations, project plans, or outlines, based on the user's natural-language instructions. The prompt "Create a short presentation about time management" clearly requests the generation of a new PowerPoint presentation, directly aligning with the "Create" prompt type.

Step 3 → Option D (Ask → "Give me ideas for a team building activity.")

"Ask" prompts are open-ended questions aimed at seeking knowledge, suggestions, creative ideas, or general guidance. Unlike "Catch up" (which summarizes existing content) or "Edit" (which modifies existing content), "Ask" prompts are for exploratory purposes and decision support. The request "Give me ideas for a team building activity" is a classic example of asking Copilot for creative suggestions and brainstorming, making it an ideal match for the "Ask" prompt type.

Step 4 → Option A (Edit → "Add an image of a target with arrows to this slide.")

"Edit" prompts are used to modify existing content. These modifications can range from rewriting text, improving clarity, adjusting tone, restructuring paragraphs, to inserting new elements like images or design features into an existing file. The instruction "Add an image of a target with arrows to this slide" explicitly directs Copilot to make a change to an already existing PowerPoint slide, which is a clear editing action.

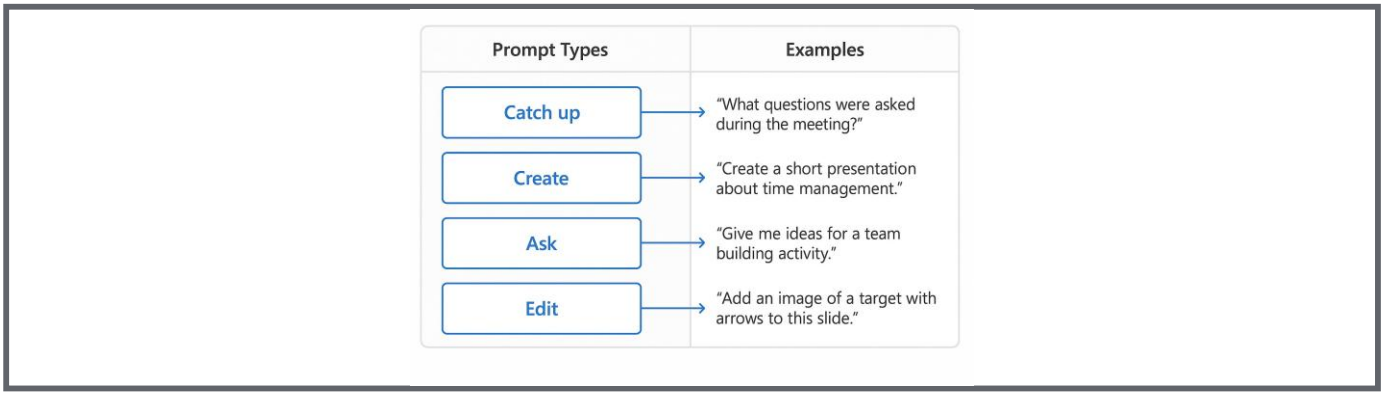

References:

https://support.microsoft.com/en-au/topic/learn-about-copilot-prompts-f6c3b467-f07c-4db1-ae54-ffac96184dd5

https://learn.microsoft.com/en-us/copilot/security/prompting-tips

Question 14

Domain: Plan AI-powered business solutions

Scenario:

Your organization creates a new AI Center of Excellence (CoE) to guide enterprise-wide adoption of generative AI. A project team submits a proposal requesting immediate development of a generative AI model. They argue that identifying use cases and validating data quality can wait until after the prototype is built, since the CoE can "fix the data later."

You are asked whether this approach aligns with Microsoft's recommended AI adoption lifecycle, which starts with identifying use cases, selecting domain-specific data, preparing and validating that data, designing and training solutions, and then monitoring and adapting them over time.

According to Microsoft's AI adoption guidance, is it appropriate to skip identifying use cases and validating domain-specific data before beginning AI model development? [Select Yes or No]

A. Yes

B. No right

Explanation:

Correct Answer: B

Microsoft's generative AI adoption framework – as shown in the diagram – emphasizes a sequenced lifecycle:

1. **Identify use cases**
2. **Prepare, validate, and aggregate the required data**
3. **Design, train, and validate AI solutions**

4. **Monitor and adapt**

The Microsoft Learn module clearly states that a Center of Excellence ensures organizations start with aligned business use cases and validated domain-specific data before any model development begins.

Skipping these early steps introduces high risk, creates misaligned solutions, and prevents effective contextualization of AI models.

Therefore, beginning model development without first identifying use cases and validating data does not follow Microsoft's recommended AI planning and adoption process.

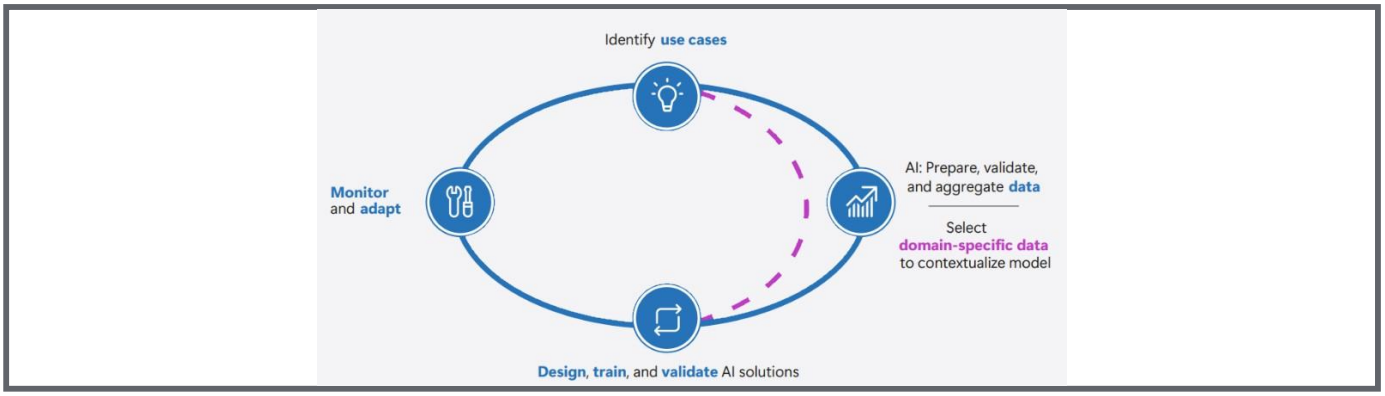

References:

https://learn.microsoft.com/en-us/training/modules/intro-ai-center-excellence/2-how-center-excellence-assists-planning-adoption-generative-ai

https://learn.microsoft.com/en-us/training/modules/intro-ai-center-excellence/1-introduction-generative-ai-center-excellence

https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/scenarios/ai/center-of-excellence

Question 15

Domain: Plan AI-powered business solutions

Scenario:

A key stakeholder in your organization is championing the immediate deployment of Microsoft Copilot across all business units. Their primary justification is the belief that AI will automatically generate superior, data-driven decisions starting on day one, irrespective of the current state of organizational data quality or the alignment of existing business processes. You are tasked with providing an accurate assessment of this claim based on industry guidance, particularly Microsoft's "AI for Business" principles.

Based on Microsoft's "AI for Business" guidance, can you confidently state that deploying AI tools like Copilot will lead to better decisions with having a foundation of quality business data and well-defined, aligned workflows first? [Select Yes or No]

A. Yes right

B. No

Explanation:

Correct Answer: A

While Microsoft's "AI for Business" guidance indeed emphasizes that AI can enable better decision-making, it critically states that this capability is not automatic or instantaneous. Instead, it explicitly requires several foundational elements:

Processing large volumes of relevant, quality business data: AI tools like Copilot depend heavily on access to accurate, comprehensive, and well-structured organizational data to derive meaningful insights. Without quality data, the AI lacks the necessary inputs to provide intelligent, contextual, or reliable outputs.

Uncovering patterns: AI's strength lies in identifying patterns and correlations within data that human analysis might miss. This process is futile if the underlying data is incomplete, inaccurate, or unstructured.

Aligning the technology to business workflows: For AI-generated insights to translate into genuinely "better decisions," they must be integrated into and relevant to existing business processes. Deploying AI in a vacuum, without considering how it fits into and enhances current workflows, will diminish its practical value and impact on decision-making.

Therefore, deploying AI tools like Copilot without an existing foundation of quality data and aligned processes will not automatically guarantee improved decisions. The AI will lack the essential context and reliable inputs needed to generate meaningful insights and drive effective outcomes.

References:

https://www.microsoft.com/en-in/microsoft-copilot/copilot-101/ai-for-business#Customerexperience

https://www.microsoft.com/en-in/microsoft-365/business-insights-ideas/resources/grow-your-small-business-with-artificial-intelligence

Question 16

Domain: Design AI-powered business solutions

View Case Study

A bank is developing a customer service copilot using Copilot Studio. The copilot has defined topics for "Checking Account Balance" and "Reporting a Lost Card." A user types, "I need to speak to someone about refinancing my mortgage," which falls outside the scope of any defined topic. The bank wants to ensure that these out-of-scope inquiries are handled gracefully and logged for future topic development without abruptly ending the conversation.

Which feature in Copilot Studio must be correctly configured to meet this requirement for managing unexpected user input?

A. System-Level Conversation Starters

B. Topic-Level Input Variables

C. The Fallback Topic right

D. Custom Action Endpoints

Explanation:

Correct Answer: C

Option C: The Fallback Topic is correct because the Fallback Topic in Copilot Studio is specifically designed to manage instances where the user's input does not match the trigger phrases of any defined system or custom topics. By configuring the Fallback Topic, the copilot can execute a graceful response (e.g., "I'm sorry, I can only help with accounts and lost cards right now. I'll log your mortgage

question.") and keep the conversation flowing or offer to transfer to a human, preventing an abrupt failure.

Option A: System-Level Conversation Starters is incorrect because this feature helps initiate a conversation (e.g., with a welcome message), but it does not manage or intercept input that occurs during the conversation and falls outside the topic scope.

Option B: Topic-Level Input Variables is incorrect because variables are used to capture and store specific data within a matching topic flow (e.g., capturing an account number), not to manage input that doesn't trigger any topic flow at all.

Option D: Custom Action Endpoints is incorrect because custom actions are used to perform external tasks (e.g., look up a balance) within a topic flow, not to handle topic routing or out-of-scope inputs.

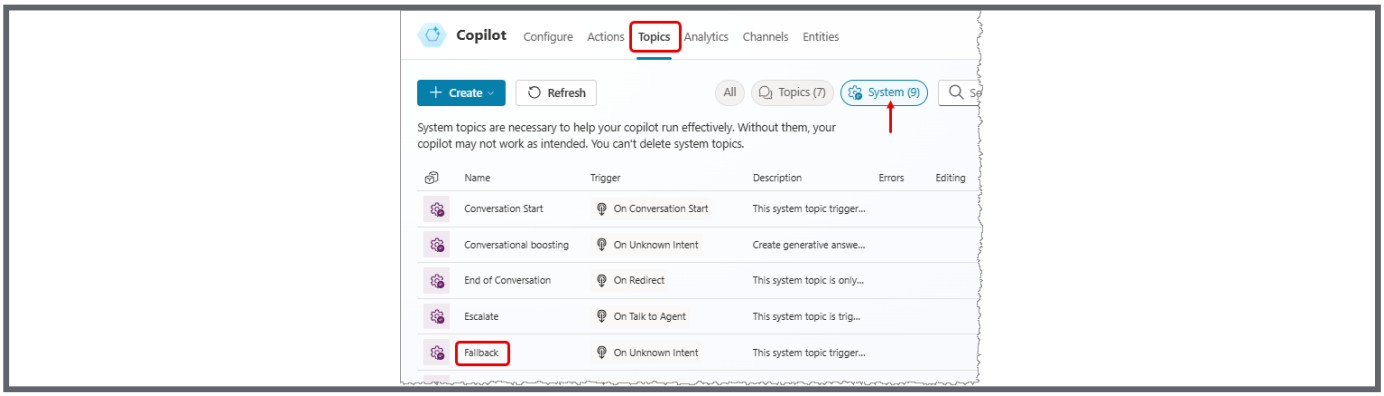

References:

Configure the system fallback topic - Microsoft Copilot Studio

Create and edit topics - Microsoft Copilot Studio

Question 17

Domain: Design AI-powered business solutions

View Case Study

A car rental company needs an internal agent to perform a very specific, limited task: interpreting a damage report from a mechanic and extracting only the Car VIN (Vehicle Identification Number) and

the Repair Cost. The list of possible damage types is large but finite and well-documented. The company requires 100% reliable extraction and needs to minimize both development cost and the risk of hallucination.

Which AI approach is the most appropriate and cost-effective choice for designing this specific, high-precision extraction task?

A. Deploy a large, custom Azure OpenAI model with a complex few-shot prompt for extraction

B. Implement a Standard Natural Language Processing (NLP) solution, specifically using Azure Conversational Language Understanding (CLU) with a custom entity model for VIN and Cost right

C. Use a multi-agent system to coordinate two different LLMs for validation

D. Design of a Generative AI orchestration layer with Retrieval-Augmented Generation (RAG)

Explanation:

Correct Answer: B

Option B: Implement a Standard Natural Language Processing (NLP) solution, specifically using Azure Conversational Language Understanding (CLU) with a custom entity model for VIN and Cost is correct because for highly specific, high-precision data extraction tasks where the required outputs (VIN, Cost) are known entities, Standard NLP (like Azure CLU) is superior to Generative AI. CLU's custom entity models are deterministic, providing far greater accuracy and reliability for structured extraction while being significantly more cost-effective and having zero risk of hallucination compared to a large LLM.

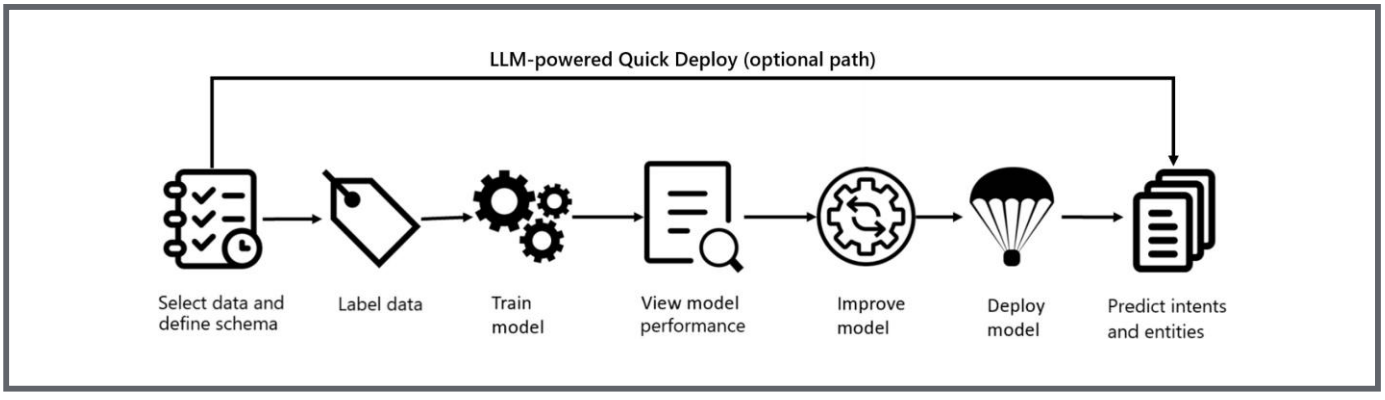

Option A: Deploy a large, custom Azure OpenAI model with a complex few-shot prompt for extraction is incorrect because while an LLM can perform extraction, it is prone to hallucination, less deterministic, and more expensive for this specific, structured task than a dedicated NLP entity model.

Option C: Use a multi-agent system to coordinate two different LLMs for validation is incorrect because this introduces unnecessary architectural complexity and cost for a simple, single-task extraction requirement.

Option D: Design of a Generative AI orchestration layer with Retrieval-Augmented Generation (RAG) is incorrect because RAG is used to ground an LLM in large knowledge bases for answering questions, not for reliable, high-precision extraction of known entities from a single document.

References:

Custom named entity recognition - Foundry Tools | Microsoft Learn

https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/scenarios/ai/

https://learn.microsoft.com/en-us/azure/ai-services/language-service/named-entity-recognition/overview

Question 18

Domain: Design AI-powered business solutions

View Case Study

An engineering firm is building a generative AI agent to answer highly complex, cross-document technical queries about project blueprints. The agent must synthesize information from tens of thousands of PDF and CAD files, often requiring references from multiple files to formulate a single answer.

The primary challenge is ensuring the LLM has the necessary and relevant context during the query. Which design strategy is most effective for providing this context strategy to the Generative AI agent?

A. Indexing the full text of all files in a single, large vector store without chunking or metadata

B. Implementing a Retrieval-Augmented Generation (RAG) system using Azure AI Search to vectorize document chunks, index relevant metadata, and retrieve only the top relevant document passages at query time right

C. Pre-training a custom Large Language Model (LLM) on all the PDF and CAD file contents

D. Using a simple text classification model to tag documents as "relevant" or "irrelevant" before querying

---

Explanation:

Correct Answer: B

Option B: Implementing a Retrieval-Augmented Generation (RAG) system using Azure AI Search to vectorize document chunks, index relevant metadata, and retrieve only the top relevant document passages at query time is correct because RAG is the industry standard for grounding LLMs in vast amounts of external knowledge. Using Azure AI Search allows the system to efficiently index the complex technical data, retrieve the most relevant context (document chunks) based on semantic similarity to the user's query, and inject only that context into the LLM's prompt. This prevents the LLM from exceeding its token limit and drastically improves the accuracy of synthesized answers.

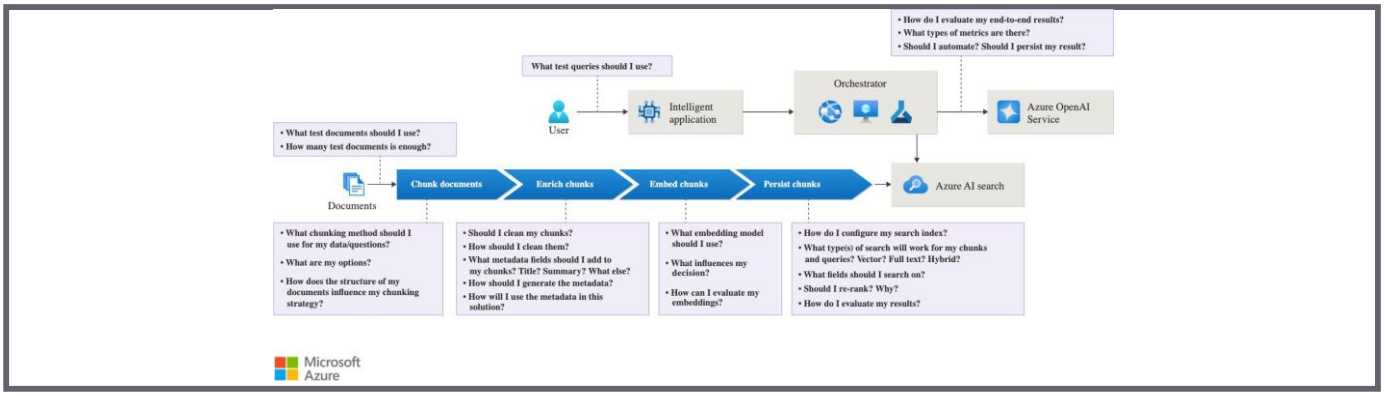

Option A: Indexing the full text of all files in a single, large vector store without chunking or metadata is incorrect because searching the full text is inefficient, and without chunking and metadata, the system cannot precisely retrieve relevant passages—it can only retrieve entire large documents, which quickly exceeds the LLM's context window.

Option C: Pre-training a custom Large Language Model (LLM) on all the PDF and CAD file contents is incorrect because this is extremely expensive, time-consuming, and difficult to update. RAG provides a dynamic, updatable, and much more cost-effective way to ground the LLM in the current knowledge base.

Option D: Using a simple text classification model to tag documents as "relevant" or "irrelevant" before querying is incorrect because classification provides a coarse filter. It does not provide the fine-grained, semantic retrieval of specific text passages required to answer a complex, cross-document query.

References:

Retrieval Augmented Generation (RAG) in Azure AI Search

Design and Develop a RAG Solution - Azure Architecture Center | Microsoft Learn

Question 19

Domain: Design AI-powered business solutions

View Case Study

A software company is designing three distinct AI agents:

Agent X: Answers single, simple queries from a knowledge base (e.g., "What is the Wi-Fi password?").

Agent Y: Automatically detects code vulnerabilities, suggests fixes, and submits a pull request without human intervention.

Agent Z: Translates a natural language command into a series of structured API calls (e.g., "Set my OOO for tomorrow") and executes the steps.

Which choice correctly maps the required agent function to its most appropriate design type?

A. Agent X: Autonomous Agent; Agent Y: Prompt/Response Agent; Agent Z: Task Agent

B. Agent X: Task Agent; Agent Y: Autonomous Agent; Agent Z: Prompt/Response Agent

C. Agent X: Prompt/Response Agent; Agent Y: Autonomous Agent; Agent Z: Task Agent right

D. Agent X: Prompt/Response Agent; Agent Y: Task Agent; Agent Z: Autonomous Agent

Explanation:

Correct Answer: C

Option C: Agent X: Prompt/Response Agent; Agent Y: Autonomous Agent; Agent Z: Task Agent is correct because

Agent X (Prompt/Response Agent): This agent simply takes an input (prompt) and generates an immediate, single output (response) based on static knowledge. It requires no external actions or complex planning.

Agent Y (Autonomous Agent): This agent operates on its own, perceives the environment (detects vulnerabilities), plans a series of actions (suggests, submits pull request), and executes them to achieve a high-level goal, all without requiring specific, step-by-step human guidance.

Agent Z (Task Agent): This agent is designed to execute a defined, multi-step transaction (translating a command into API calls, executing OOO status). It requires the use of external tools/actions to complete the structured process, making it a classic task agent.

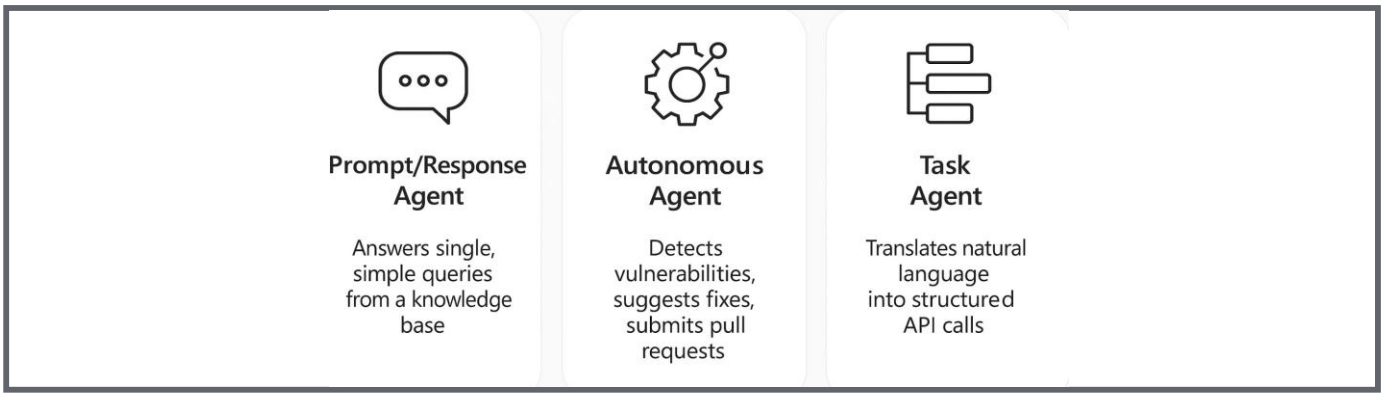

Option A: Agent X: Autonomous Agent; Agent Y: Prompt/Response Agent; Agent Z: Task Agent is incorrect because Agent X is not autonomous; it merely responds to a prompt.

Option B: Agent X: Task Agent; Agent Y: Autonomous Agent; Agent Z: Prompt/Response Agent is incorrect because Agent X does not perform a multi-step task, and Agent Z executes a multi-step task, making it a task agent, not a simple prompt/response agent.

Option D: Agent X: Prompt/Response Agent; Agent Y: Task Agent; Agent Z: Autonomous Agent is incorrect because Agent Y's ability to self-initiate detection, plan, and execute a fix makes it autonomous, not a simple task agent. Agent Z requires specific human initiation of a task, making it a task agent.

References:

What is Microsoft Foundry? - Microsoft Foundry | Microsoft Learn

AI Agent Orchestration Patterns - Azure Architecture Center | Microsoft Learn

Foundry Models sold directly by Azure

Azure OpenAI in Microsoft Foundry Models content filtering

Question 20

Domain: Design AI-powered business solutions

View Case Study

A firm is deploying a mission-critical financial reporting copilot built in Copilot Studio that connects to sensitive Azure SQL data via a custom connector. As the architect, you must review the design using the Power Platform Well-Architected Framework. The business requires that the copilot and its related data flows must prevent unauthorized data exfiltration and ensure data access conforms to company policies.

Which pillar of the Power Platform Well-Architected Framework is primarily addressed by implementing Data Loss Prevention (DLP) policies that restrict the custom connector's usage to specific environments?

A. Cost Optimization

B. Reliability

C. Performance Efficiency

D. Security right

Explanation:

Correct Answer: D

Option D: Security is correct because the primary goal of implementing Data Loss Prevention (DLP) policies in the Power Platform is to safeguard organizational data by preventing unauthorized sharing (exfiltration) of sensitive information to external or unsecured services. This directly falls under the Security pillar of the Power Platform Well-Architected Framework, which emphasizes protecting systems and data, and controlling access.

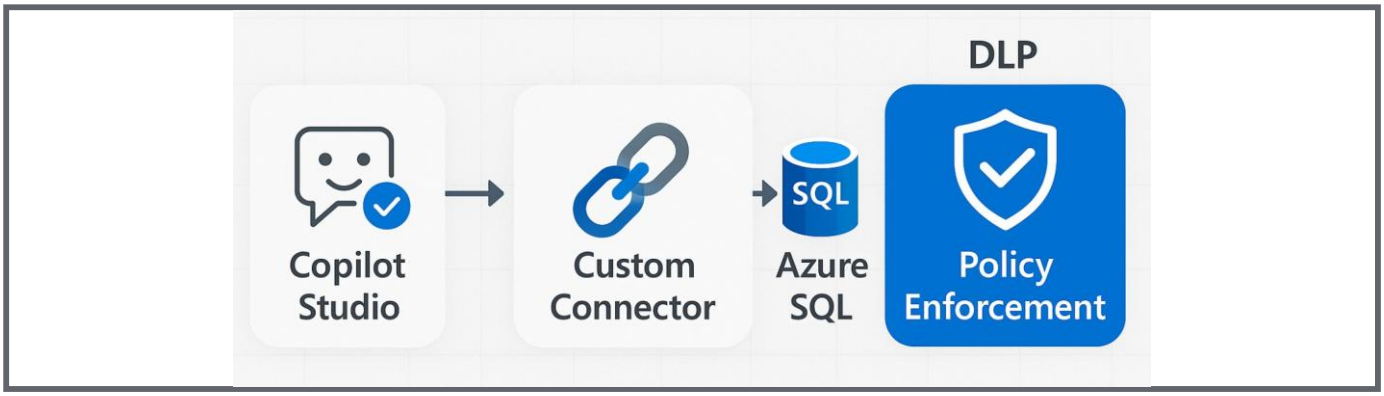

Option A: Cost Optimization is incorrect because DLP policies are security controls and do not directly optimize the financial expenditure associated with running the AI or Power Platform services.

Option B: Reliability is incorrect because Reliability focuses on the system's ability to recover from failures and consistently deliver its intended function. DLP policies are about data protection, not operational consistency.

Option C: Performance Efficiency is incorrect because Performance Efficiency relates to the system's ability to handle load and scale efficiently. DLP policies are security governance measures and do not inherently speed up or slow down the execution of the copilot.

References:

https://learn.microsoft.com/en-us/power-platform/well-architected/security/principles

https://www.microsoft.com/en-in/security/business/security-101/what-is-data-loss-prevention-dlp

https://learn.microsoft.com/en-us/power-platform/guidance/adoption/dlp-strategy

Question 21

Domain: Plan AI-powered business solutions

View Case Study

A diversified holding company is launching its AI initiative. Following the strategic guidance in the Cloud Adoption Framework (CAF) for Azure, the Chief Strategy Officer (CSO) needs to form a cross-functional team responsible for setting the initial AI vision, defining the value streams, and approving