

Microsoft.AB-100.v2026-06-11.q39

Exam Code:	AB-100
Exam Name:	Agentic AI Business Solutions Architect
Certification Provider:	Microsoft
Free Question Number:	39
Version:	v2026-06-11
# of views:	110
# of Questions views:	406
https://www.freecram.net/torrent/Microsoft.AB-100.v2026-06-11.q39.html	

NEW QUESTION: 1

A company has an AI business solution.
You need to extend the solution so that Microsoft 365 Copilot can invoke external logic hosted in Azure services.
What should you include in the solution?

- A. Microsoft Copilot Studio skills
- B. Microsoft Power Platform connectors
- C. custom engine agents

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:
The correct answer is B. Microsoft Power Platform connectors .
In this scenario, the requirement is to enable Microsoft 365 Copilot to invoke external logic hosted in Azure services . That means Copilot needs a standardized way to call external systems, APIs, or service endpoints.
This is exactly what Microsoft Power Platform connectors are designed for.
Why this is correct:
Power Platform connectors provide a structured integration layer that allows Microsoft copilots and business solutions to interact with external services. When Azure-hosted logic is exposed through APIs, connectors can be used to make that logic callable from Copilot experiences.
This is important in AI-powered business solution architecture because copilots often need to do more than answer questions. They need to perform actions such as:

- * calling Azure Functions
- * interacting with Azure-hosted APIs
- * retrieving data from external business systems
- * triggering workflows
- * passing parameters into enterprise services

Connectors act as the operational bridge between the conversational interface and the external execution logic.
From an agentic AI business solutions standpoint, this is the preferred enterprise pattern because it provides:

- * reusability
- * governance
- * security

- * low-code extensibility
- * integration consistency across Microsoft business applications

Why the other options are incorrect:

A). Microsoft Copilot Studio skills

Skills are associated with extending agent capabilities, but they are not the primary answer when the requirement is specifically for Microsoft 365 Copilot to invoke Azure-hosted external logic. The practical integration mechanism in this case is the connector layer.

C). custom engine agents

Custom engine agents are used for more advanced orchestration scenarios where you bring your own backend agent runtime or external agent system. That is more complex than what the question asks.

The question is not asking for a custom agent orchestration framework; it is asking how Microsoft 365 Copilot can invoke external Azure logic. Connectors are the more direct and appropriate solution.

Expert reasoning:

When a Microsoft exam question says:

- * Microsoft 365 Copilot
- * invoke external logic
- * hosted in Azure services

the safest and strongest answer is usually the integration mechanism that exposes actions cleanly to the Copilot environment. That mechanism is Microsoft Power Platform connectors .

NEW QUESTION: 2

A company has an ecommerce support portal that uses Microsoft Dataverse.

You are designing a Microsoft Copilot Studio agent for the portal. The agent must meet the following requirements:

- * Respond with a default help message when the user input is unclear.
- * Initiate external processes, such as retrieving the order status, when users make specific requests.

Generative orchestration will be enabled for the solution. You need to recommend a feature for each requirement.

What should you recommend? To answer, drag the appropriate features to the correct requirements. Each feature 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.

Features

A tool (connector)

A skill

A trigger phrase

The Escalate topic

The Fallback topic

Answer Area

Respond with a default help message when the user input is unclear:

Initiate external processes when requested:

Answer:

Features

A tool (connector)

A skill

A trigger phrase

The Escalate topic

The Fallback topic

Answer Area

Respond with a default help message when the user input is unclear:

The Fallback topic

Initiate external processes when requested:

A tool (connector)

Explanation:

Features

A tool (connector)

A skill

A trigger phrase

The Escalate topic

The Fallback topic

Answer Area

Respond with a default help message when the user input is unclear:

The Fallback topic

Initiate external processes when requested:

A tool (connector)

In Microsoft Copilot Studio, especially when generative orchestration is enabled, the agent selects the most appropriate capability based on user intent.

For the first requirement, when a user's input is vague, ambiguous, or does not clearly map to a known action, the correct feature is the Fallback topic . This topic is specifically intended to handle situations where the copilot cannot confidently determine the user's intent. It provides a safe default response such as a help message, clarification request, or guidance on what the user can ask next.

For the second requirement, when the user asks for something that requires interacting with an outside system-such as checking order status-the correct feature is a tool (connector) . In Copilot Studio, tools /connectors are used to call external services, APIs, workflows, or business systems. Since the portal uses Microsoft Dataverse , and the scenario involves starting an external process based on the user's request, a connector-backed tool is the right design choice.

Why the other options are not correct here:

- * A trigger phrase is used to help match utterances to a topic, but it does not itself provide the default unclear-input handling or execute external business processes.
- * A skill is not the best answer in this context; Copilot Studio commonly uses tools/connectors/actions for external operations.
- * The Escalate topic is used when handing off to a human or another escalation path, not for unclear input by default.

NEW QUESTION: 3

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:

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

Explanation:

Answer Area

Prompt validation techniques: Use prompts that have varied phrasing.

Metrics: Response relevance and accuracy

To validate prompt engineering for Microsoft Copilot Studio agents, the goal is to confirm that the agent responds correctly even when users ask the same thing in different ways, and to measure whether the responses are actually useful and correct.

For the prompt validation technique, the correct choice is Use prompts that have varied phrasing. This is a core best practice because real users do not ask questions in one fixed form. They may use different wording, sentence structure, synonyms, or levels of detail. Testing with varied phrasing checks whether the prompt design is robust and whether the agent can still produce the right response across natural language variation.

For the metric, the correct choice is Response relevance and accuracy. Since the requirement is to ensure responses are accurate and contextually relevant, this is the most appropriate measure. It directly evaluates whether the output answers the user's need correctly and in the right context.

Why the other options are not correct:

* Exclude domain-specific terminology from the prompts is not a best practice in business AI solutions.

In many enterprise scenarios, domain-specific terms are essential for accuracy.

- * Use only simple, one-word prompts does not reflect real-world usage and would weaken testing coverage.
- * The number of words generated per response does not tell you whether the response is correct or contextually appropriate.
- * The response generation time is a performance metric, not the best metric for validating prompt quality.

NEW QUESTION: 4

A company plans to deploy a Microsoft Dynamics 365 Contact Center agent.

You need to ensure that the agent can transfer the conversation to a live customer service representative.

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

NOTE: Each correct selection is worth one point.

- A.** Microsoft Foundry
- B.** Microsoft Copilot Studio
- C.** Microsoft 365 Agents Toolkit
- D.** an Azure AI Bot Service skill
- E.** Customer engagement hub

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answers are B. Microsoft Copilot Studio and E. Customer engagement hub .

This question focuses on enabling a Dynamics 365 Contact Center agent to hand off a conversation to a live customer service representative . That requires both:

- * the tool used to build and configure the conversational agent
 - * the service environment where live customer engagement and routing occur
- Why B. Microsoft Copilot Studio is correct Microsoft Copilot Studio is the platform used to build, configure, and manage the contact center agent experience. It enables you to define conversation flows, escalation logic, triggers, and handoff behavior.

In this case, the requirement is specifically that the agent must be able to transfer the conversation to a live representative. Copilot Studio is where that escalation or transfer behavior is designed as part of the agent experience.

Why E. Customer engagement hub is correct

The Customer engagement hub provides the operational environment for customer service interactions and live-agent engagement within Dynamics 365. Once the AI agent determines that escalation is required, the live representative needs an environment to receive and continue that engagement.

From a business solutions architecture perspective, this makes sense:

- * Copilot Studio defines the agent and transfer logic
 - * Customer engagement hub supports the human service experience after transfer
- Together, they satisfy the end-to-end requirement for AI-to-human handoff.

Why the other options are incorrect

A). Microsoft Foundry

Foundry supports AI model and agent development scenarios, but it is not the specific component needed for live-agent transfer in Dynamics 365 Contact Center.

C). Microsoft 365 Agents Toolkit

This is not the core component for enabling Dynamics 365 Contact Center handoff to a live service representative.

D). an Azure AI Bot Service skill

Bot skills can extend capabilities, but they are not the primary required components for enabling the standard transfer from a Dynamics 365 Contact Center agent to a live customer service representative.

Expert reasoning:

For Contact Center escalation questions, think in two layers:

- * agent authoring/orchestration # Microsoft Copilot Studio
- * human service environment / live representative experience # Customer engagement hub

NEW QUESTION: 5

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: ([SHOW ANSWER](#))

When deploying a Microsoft Foundry agent , the platform provides built #in:

- * Evaluation pipelines
- * Baseline accuracy checks
- * Drift monitoring
- * Observability dashboards

These features allow you to validate the agent against baseline metrics BEFORE deployment , which is exactly what the question requires.

This is the only option that directly addresses:

- * ALM
- * Pre#deployment evaluation
- * Accuracy validation
- * Automated quality gates

NEW QUESTION: 6

A company has a Microsoft Copilot Studio agent that uses custom connectors to interact with enterprise APIs.

You need to recommend an application lifecycle management (ALM) process to ensure that the connectors are deployed consistently across development, test, and production environments and meet governance and traceability requirements.

What should you recommend?

- A.** Export and import the connectors between the environments as unmanaged solutions.
- B.** Manage the connectors as solution components and deploy the components by using ALM pipelines.
- C.** Deploy the APIs as Azure Functions.
- D.** Maintain connector definitions in environment variables.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 7

A company has an AI solution built by using Microsoft Copilot Studio and Power Platform. The solution is used by the company ' s sales, marketing, and customer service teams.

You are performing a return on AI investment (ROAI) analysis to evaluate the impact of the solution.

You need to identify which measurable business drivers to include in the analysis.

Which two business drivers should you identify? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A.** the reduced average case resolution time
- B.** market capitalization
- C.** economic market predictability

D. increased employee productivity

E. brand awareness

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answers are A. the reduced average case resolution time and D. increased employee productivity .

This question is asking for measurable business drivers for a ROAI analysis of a Copilot Studio and Power Platform solution used by operational teams.

For ROAI, the strongest business drivers are those that are:

- * directly attributable to the AI solution

- * operationally measurable

- * tied to business outcomes

- * relevant across teams

Why A is correct

Reduced average case resolution time is a strong measurable driver because it reflects a direct operational improvement in customer service and support workflows.

This metric can be quantified clearly by comparing:

- * baseline resolution time before AI

- * resolution time after deployment

That makes it ideal for ROAI because faster case resolution often leads to:

- * lower service cost

- * higher throughput

- * better customer experience

- * more efficient staffing

Why D is correct

Increased employee productivity is another core ROAI driver because AI solutions in sales, marketing, and customer service are often deployed specifically to reduce manual work and improve output per employee.

This can be measured through indicators such as:

- * more tasks completed per agent or employee

- * reduced manual effort

- * increased throughput

- * faster response cycles

- * more time spent on higher-value work

From an AI business solutions perspective, productivity improvement is one of the most common and valid drivers in ROAI analysis.

Why the other options are incorrect

B). market capitalization

This is too broad and influenced by many external factors. It is not a practical direct business driver for evaluating the specific impact of one AI business solution.

C). economic market predictability

This is not a direct business driver created by the solution and is too external to the organization's operational AI ROI calculation.

E). brand awareness

Brand awareness can matter strategically, but it is less directly attributable and less operationally measurable than resolution time and productivity for this kind of internal business solution.

Expert reasoning

For ROAI questions, prefer metrics that are:

- * operational

- * attributable
- * measurable before and after deployment

That leads to:

- * reduced average case resolution time
- * increased employee productivity

NEW QUESTION: 8

You need to design a Microsoft 365 Copilot solution to optimize employee productivity. The solution must meet the following requirements:

Ensure that the employees can query content stored in a subset of Microsoft SharePoint Online sites and in Teams by using natural language-based prompt actions.

Ensure that employees receive contextually relevant responses in Microsoft 365 Copilot.

What should you include in the design?

- A.** Build a Microsoft Power Automate desktop flow to read the SharePoint content and post the responses to Teams.
- B.** Modify SharePoint settings.
- C.** Create a custom REST API that crawls the SharePoint content.
- D.** Configure Microsoft Graph access.

Answer: (SHOW ANSWER)

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answer is D. Configure Microsoft Graph access .

Microsoft 365 Copilot grounds its responses in Microsoft 365 data through the Microsoft Graph . If employees need to query content from a subset of SharePoint Online sites and Teams using natural-language prompts, the solution must ensure Copilot can access and use the right Microsoft 365 content context through Graph-connected permissions and data access patterns.

Why D is correct

Microsoft Graph is the core data and context layer for Microsoft 365 Copilot. It connects Copilot to organizational content such as:

- * SharePoint sites
- * Teams messages and files
- * OneDrive content
- * Outlook data
- * calendar and collaboration context

Because the requirement is to provide contextually relevant responses in Microsoft 365 Copilot, the design must rely on the platform's native grounding mechanism. That mechanism is Graph-based access to Microsoft 365 content.

From an AI business solutions perspective, this is the right design because it ensures:

- * natural-language prompts can retrieve relevant organizational knowledge
- * responses are grounded in authorized enterprise content
- * access remains aligned to Microsoft 365 permissions
- * employees only see content they are allowed to access

This is especially important when only a subset of SharePoint sites should be included. The relevance and security model depend on the Microsoft 365 content graph and its permission-aware access behavior.

Why the other options are incorrect

A). Build a Microsoft Power Automate desktop flow to read the SharePoint content and post the responses to Teams This is not how Microsoft 365 Copilot should be designed for grounding enterprise content. It is overly manual, indirect, and does not provide native contextual grounding for Copilot responses.

B). Modify SharePoint settings

SharePoint settings may affect site permissions or content availability, but they do not by themselves enable Microsoft 365 Copilot's natural-language grounding across SharePoint and Teams.

C). Create a custom REST API that crawls the SharePoint content

This adds unnecessary custom complexity and bypasses the native Microsoft 365 Copilot architecture. The requirement is best met through Microsoft Graph-based access, not by building a parallel crawler.

Expert reasoning

For Microsoft 365 Copilot questions:

- * if the requirement is to query Microsoft 365 content with natural language
- * and return contextually relevant responses from SharePoint and Teams
- * the key design element is usually Microsoft Graph

NEW QUESTION: 9

A company has a Microsoft Foundry generative AI model.

You need to evaluate the model 's output to measure the overall quality and coherence of generated responses. The evaluation must use GPT-4o as a judge and return a numeric score for each output.

Which type of metric should you use?

- A.** Groundedness
- B.** AI quality (NLP)
- C.** risk and safety
- D.** AI quality (AI assisted)

Answer: ([SHOW ANSWER](#))

The requirement says the evaluation must:

- * measure overall quality and coherence of generated responses
- * use GPT-4o as a judge
- * return a numeric score for each output

That maps directly to AI quality (AI assisted) .

Why D is correct:

AI-assisted quality metrics use a strong model such as GPT-4o to evaluate generated outputs against qualitative dimensions like coherence, relevance, completeness, and overall response quality, and then assign scores.

Why the other options are not correct:

- * A. Groundedness checks whether the answer is supported by the source/context, not overall quality and coherence.
- * B. AI quality (NLP) usually refers to traditional NLP-style metrics rather than judge-model scoring with GPT-4o.
- * C. risk and safety evaluates harmfulness, toxicity, jailbreak risk, or policy concerns, not overall response quality.

NEW QUESTION: 10

You need to design a shared prompt library that will be used across multiple business units. The solution must meet the following requirements:

Ensure consistent AI responses with reusable formats.

Support governance and version control.

Minimize administrative effort.

Minimize ongoing costs.

What should you recommend for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Ensure consistent AI responses:

Delegate department-specific prompt templates.
Define standardized prompt templates.
Maintain a prompt history.

Support governance and version control:

Define standardized prompt templates.
Store prompts in a Git repository.
Categorize prompts by business function.

Answer:

Ensure consistent AI responses:

Delegate department-specific prompt templates.
Define standardized prompt templates.
Maintain a prompt history.

Support governance and version control:

Define standardized prompt templates.
Store prompts in a Git repository.
Categorize prompts by business function.

Explanation:

Ensure consistent AI responses # Define standardized prompt templates; Support governance and version control # Store prompts in a Git repository The requirement says the shared prompt library must ensure consistent AI responses with reusable formats .

The best way to achieve that is to create standardized prompt templates .

Standardized templates help multiple business units use the same structure for:

- * system instructions
- * response style
- * output formatting
- * safety guidance
- * business terminology
- * escalation behavior

This is important in AI-powered business solutions because prompt variability often leads to inconsistent outputs. A shared template approach creates repeatability and reduces drift across departments.

From an enterprise AI design perspective, standardized prompt templates improve:

- * response consistency
- * quality control
- * reusability

- * onboarding speed
- * compliance alignment

They also minimize administrative effort because teams do not need to invent new prompt structures from scratch.

Why "Store prompts in a Git repository" is correct

The requirement says the solution must support governance and version control while minimizing effort and cost.

A Git repository is the most appropriate answer because it provides:

- * version history
- * change tracking
- * rollback capability
- * collaboration support
- * approval workflows
- * auditability

This directly supports governance. It allows the organization to know:

- * who changed a prompt
- * what changed
- * when it changed
- * why it changed

That is extremely valuable for enterprise AI operations, especially when prompts influence business decisions, customer interactions, compliance behavior, or internal process automation.

Git also minimizes ongoing cost because it uses standard development and content management practices rather than requiring a separate expensive governance platform for prompt assets.

Why the other options are incorrect

Delegate department-specific prompt templates

This reduces consistency instead of improving it. Different departments creating their own templates introduces fragmentation, style drift, and governance challenges.

Maintain a prompt history

Keeping a history is useful, but it does not by itself ensure consistent reusable formats. It is more of a recordkeeping activity than a standardization mechanism.

Categorize prompts by business function

Categorization is helpful for organization and discoverability, but it does not provide version control or governance by itself.

Define standardized prompt templates for governance/version control

Standardization improves consistency, but version control specifically requires a system for tracking and managing changes. That is why Git is the better answer for the governance requirement.

Expert reasoning

A shared enterprise prompt library should usually follow this pattern:

- * Consistency and reuse # standard templates
- * Governance and versioning # Git-based repository
- * Low cost and low admin overhead # reuse existing tooling and centralized controls

NEW QUESTION: 11

A financial services company uses Microsoft Dynamics 365 Finance.

Currently, the company 's support staff manually reviews customer transaction histories to detect potential fraud cases before escalating the cases.

You need to recommend an automation solution for the review process. The solution must ensure that escalations reach a human analyst for final decision making. What should you recommend?

- A.** Use Microsoft 365 Copilot in Word to automatically finalize fraud detection policies.
- B.** Configure a task agent to generate fraud risk scores for the human analyst to review.
- C.** Deploy an autonomous agent that closes non-fraud cases automatically.

D. Export the data to a data lake for analysis in Microsoft Power BI.

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 12

You are designing an AI strategy for Microsoft Dynamics 365 finance and operations apps. You are evaluating the use of Microsoft Copilot Studio to provide in-app help and guidance based on generative AI general knowledge.

You need to recommend which knowledge sources to include in the generative help and guidance agent. The solution must minimize the risk of generating inaccurate responses.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Custom knowledge sources:

- Must be uploaded to the agent
- Must be excluded from the agent
- Are not supported

AI general knowledge:

- Must be enabled for the agent
- Must be disabled for the agent
- Is not supported

Answer:

Custom knowledge sources:

- Must be uploaded to the agent
- Must be excluded from the agent
- Are not supported

AI general knowledge:

- Must be enabled for the agent
- Must be disabled for the agent
- Is not supported

Explanation:

Custom knowledge sources # Must be uploaded to the agent; AI general knowledge # Must be disabled for the agent The requirement says the solution must minimize the risk of generating inaccurate responses . In generative AI business solutions, the best way to reduce hallucinations or loosely grounded answers is to rely on approved, scoped, domain-specific knowledge rather than broad general knowledge.

That means the agent should use:

- * custom knowledge sources that are intentionally provided to the agent
- * AI general knowledge disabled , so the agent does not answer from broad open-ended model knowledge when enterprise-specific guidance is needed In finance and operations scenarios, this is especially important because users may ask about:
 - * internal business processes
 - * financial controls
 - * operational procedures
 - * policy-specific guidance

* company-configured ERP workflows

These are areas where inaccurate answers can create operational or compliance risk. Grounding the agent on uploaded custom sources provides a more controlled knowledge base.

Why custom knowledge sources must be uploaded

If the goal is in-app help and guidance for Dynamics 365 finance and operations apps, the safest and most reliable approach is to provide the agent with curated documentation such as:

* internal process guides

* approved SOPs

* finance workflow documentation

* policy documents

* ERP-specific instructions

Uploading these custom knowledge sources ensures the agent answers from enterprise-approved material rather than from generic model knowledge.

From an AI governance perspective, this improves:

* factual grounding

* relevance

* auditability

* trustworthiness

* alignment to company-specific processes

Why AI general knowledge must be disabled

General AI knowledge can be useful in broad assistant scenarios, but in a finance and operations help context it increases the risk that the agent may generate:

* overly generic responses

* answers not aligned to company policy

* incorrect interpretations of internal process steps

* content that sounds plausible but is not operationally correct

Because the question explicitly says to minimize the risk of inaccurate responses , the safer design is to disable general knowledge and constrain responses to approved custom sources.

NEW QUESTION: 13

A company has a Microsoft Copilot Studio agent for customer support. You are reviewing and validating the following prompts:

* A prompt that has instructions to " help the customer as best you can "

* A prompt that helps retrieve product information from a knowledge base You need to ensure that the agent delivers consistent and accurate responses.

What should you do for each prompt? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

A prompt that has instructions to "help the customer as best you can":



Add filler words to make the prompt sound more natural and conversational.
Keep the prompt vague to enable model flexibility.
Rewrite the prompt with clear and task- specific instructions.

Add several open-ended questions to give the model broader context.
Use responses with only reference sources and limit the response scope.
Remove the knowledge source so that the model responds freely with general product information.

Answer:

Answer Area

A prompt that has instructions to "help the customer as best you can":

Add filler words to make the prompt sound more natural and conversational.

Keep the prompt vague to enable model flexibility.

Rewrite the prompt with clear and task-specific instructions.

Add several open-ended questions to give the model broader context.

Use responses with only reference sources and limit the response scope.

Remove the knowledge source so that the model responds freely with general product information.

Explanation:

Answer Area

A prompt that has instructions to "help the customer as best you can": Rewrite the prompt with clear and task-specific instructions.

A prompt that helps retrieve product information from a knowledge base: Use responses with only reference sources and limit the response scope.



This question is about improving prompt quality so a Microsoft Copilot Studio agent gives consistent and accurate answers.

For the first prompt, "help the customer as best you can" is too vague. It does not tell the model exactly what task to perform, what boundaries to follow, or what kind of response is expected. The correct improvement is to rewrite the prompt with clear and task-specific instructions. Clear prompts reduce ambiguity and make agent behavior more predictable and repeatable.

For the second prompt, the agent is retrieving product information from a knowledge base. To keep answers accurate and grounded, the best practice is to use responses with only reference sources and limit the response scope. That ensures the model stays tied to approved knowledge and does not invent unsupported product details.

Why the other options are not correct:

- * Add filler words to make the prompt sound more natural and conversational does not improve accuracy or consistency.
- * Keep the prompt vague to enable model flexibility increases inconsistency.
- * Add several open-ended questions to give the model broader context can make responses less focused.
- * Remove the knowledge source so that the model responds freely with general product information would reduce reliability and increase hallucination risk.

NEW QUESTION: 14

A company has a canvas app named App1 in a Microsoft Power Platform environment named Env1.

Env1 uses a customer-managed key for data encryption. App1 connects to multiple data sources to retrieve and update customer and order information.

You need to recommend a solution to add Microsoft Copilot components to App1. The solution must NOT modify the current security or encryption configurations of Env1.

What should you include in the recommendation?

A. Duplicate App1 and republish the app in Env1.

- B.** Enable Copilot features for Env1.
- C.** Modify the data sources of App1 to make them compatible with Copilot.
- D.** Move App1 to a new environment that uses Microsoft-managed keys.

Answer: ([SHOW ANSWER](#))

The deciding factor here is that Env1 uses a customer-managed key (CMK) for encryption, while the requirement is to add Microsoft Copilot components without changing the existing security or encryption setup of that environment.

In Power Platform, Copilot and generative AI capabilities have environment-level support requirements. A CMK-enabled environment can limit or prevent the use of certain Copilot capabilities . Because of that, the clean recommendation is to place the app in a different environment that supports Copilot features , specifically one using Microsoft-managed keys , instead of changing Env1.

Why D is correct:

- * It allows Copilot components to be used in a compatible environment.
- * It avoids changing the existing encryption/security configuration of Env1 .
- * It satisfies the requirement that Env1's current configuration must remain untouched.

Why the other options are not correct:

- * A. Duplicate App1 and republish the app in Env1 This keeps the app inside the same CMK-based environment, so it does not solve the compatibility issue.
- * B. Enable Copilot features for Env1 If the environment's encryption model is the blocker, merely enabling features does not resolve that limitation.
- * C. Modify the data sources of App1 to make them compatible with Copilot The issue is not mainly about the app's data sources. It is about the environment's support and encryption configuration.

NEW QUESTION: 15

A company has Microsoft Copilot Studio agents.

The company plans to deploy custom connectors across development, test, and production environments.

You need to design an application lifecycle management (ALM) process to ensure consistency and prevent direct editing in production. 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.** Include agents and connectors in a solution.
- B.** Move the agents between the environments by using data export and import.
- C.** Manually rebuild the agents in each environment.
- D.** Deploy managed solutions to production.
- E.** Deploy unmanaged solutions to production.

Answer: ([SHOW ANSWER](#))

For Microsoft Copilot Studio ALM across development, test, and production , the right approach is to package the app components properly and use controlled deployment in production.

A). Include agents and connectors in a solution is correct because solutions are the standard Power Platform ALM container for moving components consistently across environments. Putting both the agents and the custom connectors into the solution ensures they travel together and remain aligned.

D). Deploy managed solutions to production is also correct because managed solutions are intended for production use. They help prevent direct editing in production and support stronger governance and release control.

Why the other options are not correct:

- * B. Move the agents between environments by using data export and import is not the recommended ALM approach for Copilot Studio components.
- * C. Manually rebuild the agents in each environment creates inconsistency and increases effort and risk.
- * E. Deploy unmanaged solutions to production would allow easier direct modification, which conflicts with the requirement to prevent direct editing in production.

NEW QUESTION: 16

A company uses Microsoft 365 and Dynamics 365

You need to recommend a solution to automatically summarize email threads, generate suggested replies in Microsoft Outlook, and provide meeting preparation summaries that include relevant customer relationship management (CRM) data.

Solution: You recommend Microsoft 365 Copilot for Sales.

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

Valid AB-100 Dumps shared by EduDump.com for Helping Passing AB-100 Exam! EduDump.com now offer the **newest AB-100 exam dumps**, the EduDump.com AB-100 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AB-100 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/AB-100/premium/> (**99** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 17

You are evaluating a Microsoft Copilot Studio agent that supports Microsoft Dynamics 365 Customer Service representatives.

You need to recommend a testing solution that meets the following requirements:

Evaluates agent effectiveness during active sessions

Validates whether the agent delivers accurate and helpful responses

Provides measurable, actionable insights for continuous improvement

What should you recommend?

A. Track resolution, deflection, and accuracy by using dashboards and use scripts to ensure consistent responses.

B. Perform load testing to validate how the agent scales under a high chat volume.

C. Review historical tickets to find agents that have the shortest resolution times.

D. Measure uptime and page load times.

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answer is A. Track resolution, deflection, and accuracy by using dashboards and use scripts to ensure consistent responses .

This question is about evaluating a Copilot Studio agent in live support operations , not just testing technical uptime or infrastructure performance. The requirements emphasize three things:

* effectiveness during active sessions

* response accuracy and helpfulness

* measurable insights for continuous improvement

That combination points to operational quality metrics and analytics dashboards.

Why A is correct

Tracking resolution , deflection , and accuracy directly measures how well the agent performs in real support conversations:

* Resolution shows whether the issue is successfully handled

* Deflection shows whether the agent reduces human workload appropriately

* Accuracy shows whether responses are correct and helpful

Using dashboards gives leaders and support teams measurable, ongoing visibility into agent behavior. Adding scripts for consistent testing further supports repeatable evaluation and improvement.

From an AI business solutions perspective, this is the right recommendation because it combines:

* business outcome measurement

- * quality validation
- * operational analytics
- * continuous improvement feedback loops

This is exactly how enterprise copilots should be managed after deployment.

Why the other options are incorrect

B). Perform load testing to validate how the agent scales under a high chat volume Load testing is useful for scalability and capacity planning, but it does not directly validate whether responses are accurate, helpful, or effective during active sessions from a business-outcome perspective.

C). Review historical tickets to find agents that have the shortest resolution times This may give some retrospective insight, but it does not directly evaluate the Copilot Studio agent during active sessions, and shortest resolution time alone does not prove response quality or helpfulness.

D). Measure uptime and page load times

These are infrastructure and availability metrics. They are important for system health, but they do not evaluate conversational effectiveness or answer quality.

Expert reasoning

For Copilot evaluation questions:

- * if the goal is business effectiveness in active sessions , use resolution/deflection/accuracy
- * if the goal is system scale , use load testing
- * if the goal is infrastructure reliability , use uptime and latency

NEW QUESTION: 18

A company has a Microsoft Foundry generative AI model.

You need to evaluate the model ' s output to measure the overall quality and coherence of generated responses. The evaluation must use GPT-4o as a judge and return a numeric score for each output.

Which type of metric should you use?

- A.** AI quality (NLP)
- B.** AI quality (AI assisted)
- C.** Groundedness
- D.** risk and safety

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 19

You are designing a Microsoft Copilot Studio agent that uses a custom Microsoft Foundry model to generate responses. You need to ensure that the agent can securely connect to and invoke the custom model during user interactions. What should you include in the design?

- A.** Create a connection to Microsoft Foundry in the agent
- B.** Create a custom engine agent
- C.** Add the Microsoft Foundry model as a Copilot Studio skill.
- D.** Configure the agent to use classic orchestration.

Answer: ([SHOW ANSWER](#)**)**

NEW QUESTION: 20

A company uses Microsoft 365 and Dynamics 365
You need to recommend a solution to automatically summarize email threads, generate suggested replies in Microsoft Outlook and provide meeting preparation summaries that include relevant customer relationship management (CRM) data.
Solution: You recommend a classic Microsoft Dataverse workflow.
Does this meet the goal?

- A. Yes
 - B. No
- Answer: ([SHOW ANSWER](#))

NEW QUESTION: 21

A company has a Microsoft Power Platform environment.
You need to build two agents named Agent1 and Agent2. The solution must meet the following requirements:

- * Agent1 must be extendable by using the Semantic Kernel and must connect to multiple business apps and APIs.
- * Agent2 must connect directly to data stored in Microsoft Dataverse and must be embeddable in a Microsoft Power Apps canvas app.

What should you use to build each agent? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

ANSWER AREA



Agent1:

Microsoft foundation
Azure Logic Apps
Copilot in Power Apps
Microsoft Copilot Studio

Agent2:

Microsoft Foundry
Azure Logic Apps
Copilot in Power Apps
Microsoft Copilot Studio

Answer:

Answer Area

Agent1:

Agent2:

Microsoft

Explanation:

Verified Answer : =

* Agent1 # Microsoft Foundry

* Agent2 # Copilot in Power Apps

Comprehensive and Detailed Explanation from Agentic AI Topics:

For Agent1 , the requirement is that it must be extendable by using Semantic Kernel and connect to multiple business apps and APIs . The best fit is Microsoft Foundry because Foundry-based agents are designed for extensibility and developer-oriented orchestration, including integration patterns that work well with Semantic Kernel and external tools/APIs.

For Agent2 , the requirement is that it must connect directly to Microsoft Dataverse and be embeddable in a Power Apps canvas app . The best fit is Copilot in Power Apps , because it is designed for Power Platform-native experiences, works naturally with Dataverse-backed app data , and is intended for embedding AI experiences inside canvas apps .

Why the other options are not the best match:

* Azure Logic Apps is for workflow orchestration, not the primary platform for building these agents.

* Microsoft Copilot Studio is strong for conversational agents, but the wording here points more directly to Power Apps-native embedding for Agent2 and Semantic Kernel extensibility for Agent1.

NEW QUESTION: 22

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: ([SHOW ANSWER](#))

For Microsoft Copilot Studio agents, effective prompting depends on giving the model strong guidance and continuously improving that guidance from real interactions.

C). Use clear and specific instructions in the prompts is correct because well-written prompts reduce ambiguity, improve intent interpretation, and make the agent's behavior more reliable and consistent.

D). Regularly test and refine the prompts based on user input is also correct because prompt quality improves through iteration. Reviewing how users actually interact with the agent helps identify unclear wording, missing instructions, and failure cases.

Why the others are not correct:

- * A. Measure system resource usage during prompt processing is an operations/performance activity, not a prompt best practice.
- * B. Analyze the prompt length distribution may be interesting analytically, but it is not a primary best practice for improving agent interactions.
- * E. Track the duration of the average user session is a usage metric, not a prompt-engineering best practice.

NEW QUESTION: 23

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 (LEAVE A REPLY)

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answers are A. Microsoft Purview and D. Microsoft Defender .

This question is asking for an enterprise security and governance design for Microsoft 365 Copilot agents

. The requirements span three major control areas:

- * identify and mitigate AI-related risks
 - * protect AI apps and sensitive data
 - * retain/log interactions, detect policy violations, and investigate incidents
- No single tool in the list fully covers all of those needs. The best solution is the combination of Microsoft Purview and Microsoft Defender .

Why A. Microsoft Purview is correct

Microsoft Purview is the strongest match for the requirements around:

- * protecting sensitive data
- * governance of AI usage
- * retaining and logging interactions
- * detecting policy violations
- * supporting investigation and compliance processes

Purview is central to Microsoft's information protection, compliance, insider risk, auditing, and data governance capabilities. In the context of Microsoft 365 Copilot agents, Purview helps organizations:

- * classify and label sensitive data
- * apply data loss prevention controls
- * retain records and interactions
- * audit activity
- * investigate policy issues

* support responsible AI governance practices

From an AI business solutions perspective, this is essential because copilots often process sensitive enterprise information, and organizations need visibility into how that information is used, exposed, and governed.

Why D. Microsoft Defender is correct

Microsoft Defender addresses the requirement to identify and mitigate potential risks that relate to AI use and to protect AI apps .

Defender is the broader security layer that helps monitor and protect applications, detect threats, identify vulnerabilities, and support incident response. In AI-enabled enterprise solutions, Defender helps secure the application environment and detect risk patterns that could affect AI systems or the data they use.

This is important because AI security is not only about content and compliance. It is also about:

* threat detection

* app protection

* attack surface awareness

* suspicious activity monitoring

* incident investigation

Defender complements Purview by focusing more on the security posture and threat protection side of the solution.

Why the other options are incorrect

B). Azure AI Content Safety

Azure AI Content Safety is valuable for filtering harmful or unsafe AI-generated or user-supplied content.

However, it does not fully address the broader requirements here around enterprise data protection, interaction retention, policy logging, governance, and incident investigation. It is useful, but not the best two-part answer.

C). role-based access control (RBAC) in Microsoft Foundry

RBAC is important for access management, but this option is too narrow and also not the best fit for Microsoft 365 Copilot agents in this question. It does not cover the required governance, retention, policy violation detection, or investigation capabilities.

Expert reasoning

A good way to solve this kind of question is to separate the requirements into two control domains:

* data governance, retention, policy, compliance # Microsoft Purview

* threat protection, risk mitigation, app security, investigation support # Microsoft Defender That pairing gives the most complete answer across the listed options.

NEW QUESTION: 24

A company plans to deploy a Microsoft Copilot Studio agent to enhance customer support.

The company stores customer data across ServiceNow, Microsoft Dynamics 365 Finance, Dynamics 365 Supply Chain Management, and Excel files in SharePoint Online.

You need to recommend a solution to ensure that the agent can deliver accurate and timely responses.

What should you recommend?

A. Enable incremental indexing in Azure AI Search.

B. Implement a model router for query handling.

C. Create custom prompts.

D. Implement Microsoft Power Platform connectors.

Answer: ([SHOW ANSWER](#))

The agent must deliver accurate and timely responses while customer data is spread across several systems:

* ServiceNow

* Dynamics 365 Finance

* Dynamics 365 Supply Chain Management

* Excel files in SharePoint Online

The most appropriate recommendation is Microsoft Power Platform connectors because they are the standard low-code integration mechanism for bringing together data and actions from multiple enterprise systems inside Microsoft Copilot Studio .

Why D is correct:

* Connectors let the agent access data across different business systems

* They reduce development effort compared with custom integration patterns

* They help the agent ground responses on the latest data from connected sources Why the other options are not the best fit:

* A. Enable incremental indexing in Azure AI Search Useful in search-based architectures, but the core issue here is connecting multiple business systems into the agent solution.

* B. Implement a model router for query handling A router helps distribute requests, but it does not solve enterprise data access and grounding across these systems.

* C. Create custom prompts Prompting alone does not integrate the source systems or ensure current enterprise data access.

NEW QUESTION: 25

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: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answer is A. Microsoft Dataverse .

This question is asking for a Microsoft Power Platform business solution that can act as a centralized data foundation across multiple AI and business application workloads. The requirements are very specific:

* consolidate data from multiple internal and external sources

* serve as a centralized source for Copilot Studio agents , Dynamics 365 , and external AI models

* support built-in data classification and protection policies

* provide data for grounding and analytics

Among the options, Microsoft Dataverse is the best fit.

Why A is correct

Microsoft Dataverse is the native business data platform for Microsoft Power Platform and Dynamics 365. It is designed to act as a structured, centralized, governed source of business data. That makes it the strongest answer when the scenario explicitly involves:

* Copilot Studio

* Dynamics 365

* broader Power Platform

* governed enterprise business data

* AI grounding and analytics

Dataverse supports these needs because it provides:

* a common business data model

- * secure centralized storage
- * integration across Power Platform and Dynamics 365
- * metadata-rich tables and relationships
- * role-based security
- * support for business rules and governance
- * compatibility with analytics and AI-based experiences

From an AI business solutions perspective, Dataverse is especially strong because it can act as the single source of truth for enterprise business data that powers both transactional applications and AI systems.

Why Dataverse fits the AI requirements

For AI systems, especially Copilot and agent scenarios, centralized structured business data is essential for:

- * grounding responses in current operational data
- * supporting retrieval across customer, sales, finance, or service records
- * enabling governed access to sensitive information
- * providing high-quality data for downstream reporting and analytics

Dataverse also aligns well with the requirement for built-in data classification and protection policies , because it works within Microsoft's enterprise governance ecosystem and supports security, auditing, and compliance-oriented controls better than the other listed options in a Power Platform business context.

Why the other options are incorrect

B). Azure Data Lake Storage

Azure Data Lake Storage is excellent for large-scale analytics and raw data storage, but it is not the best Power Platform business solution answer here. It lacks the same native business application integration and governed operational data model that Dataverse provides for Copilot Studio and Dynamics 365 scenarios.

C). a Microsoft Power BI semantic model

A semantic model is useful for reporting and analytics, but it is not the central operational data platform for multiple AI systems. It sits more at the reporting layer than the transactional and grounding layer.

D). Azure Cosmos DB

Cosmos DB is a scalable NoSQL database, but it is not the native Microsoft Power Platform business data platform for Dynamics 365 and Copilot Studio integration. It also does not provide the same built-in business data modeling and governance experience expected here.

Expert reasoning

When the question combines:

- * Power Platform
- * Dynamics 365
- * Copilot Studio
- * centralized business data
- * governance
- * AI grounding

the best answer is almost always Microsoft Dataverse .

NEW QUESTION: 26

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

Answer: (SHOW ANSWER)

NEW QUESTION: 27

You are designing end-to-end test scenarios for a business solution that uses Microsoft Dynamics 365 Sales and Dynamics 365 Finance. You need to ensure that the business solution meets the following test requirements:

- * Properly exchanges data between the Dynamics 365 apps
- * Aligns with defined user workflows and business processes

Which type of testing should you use for each requirement? To answer, drag the appropriate testing types to the correct requirements. Each testing type 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.

Testing types

⋮ Drift

⋮ Exploratory

⋮ Integration

⋮ Performance

⋮ User acceptance

Answer Area

Properly exchanges data between the Dynamics 365 apps:

Aligns with defined user workflows and business processes:

Answer:

Testing types

⋮ Drift

⋮ Exploratory

⋮ Integration

⋮ Performance

⋮ User acceptance

Answer Area

Properly exchanges data between the Dynamics 365 app: ⋮ Integration

Aligns with defined user workflows and business processes: ⋮ User acceptance

Explanation:

Testing types

Drift

Exploratory

Integration

Performance

User acceptance

Answer Area

Properly exchanges data between the Dynamics 365 apps: Integration

Aligns with defined user workflows and business processes: User acceptance

For end-to-end validation of a solution that uses Dynamics 365 Sales and Dynamics 365 Finance , the testing type should match the goal of each requirement.

For properly exchanging data between the Dynamics 365 apps , the correct testing type is Integrat ion testing.

This verifies that the systems connect correctly, pass data accurately, and maintain consistency across app boundaries.

For aligning with defined user workflows and business processes , the correct testing type is User acceptance testing. This focuses on whether the solution supports real business tasks the way users expect and whether it fits the intended operational process.

Why the other options are not correct here:

- * Drift is about changes over time, often in model or behavior consistency.
- * Exploratory is useful for uncovering unexpected issues, but it is not the primary match for the stated requirement.
- * Performance focuses on speed, scale, and responsiveness, not workflow fit.

NEW QUESTION: 28

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:



Explanation:

Verified Answer : =

* CFO concerns about Copilot Studio credit usage # Effectiveness

* CTO concerns about poor feedback on AI agent responses # Satisfaction Comprehensive and Detailed Explanation from Agentic AI Topics:

The correct mapping is based on what each executive is trying to measure in Copilot Studio analytics .

For the CFO , the concern is about abandoned interactions compared to resolved conversations and whether Copilot Studio credits are being used efficiently. That aligns with Effectiveness , because effectiveness metrics focus on c onversation outcomes such as whether interactions are successfully resolved or dropped before completion.

For the CTO , the concern is specifically about user feedback on the quality of AI agent responses . That maps to Satisfaction , because satisfaction met rics capture user sentiment and feedback about the interaction quality.

Why the other metrics are not the best fit:

* Use is more about adoption and interaction volume.

* Tool use focuses on how often tools/actions are invoked, not on abandonment versus resolution or feedback quality.

NEW QUESTION: 29

A company uses a fine-tuned Microsoft Foundry model that requires frequent updates as new customer feedback becomes available.

You need to design an application lifecycle management (ALM) process that meets the following requirements:

* Data changes must be tracked and versioned.

* The model must be retrained consistently by using approved training data.

Which two actions should you include in the design?

NOTE: Each correct selection is worth one point.

A. Associate the storage location to the fine-tuning job.

B. Create a content filter.

C. Store the training data in Azure Files.

D. Upload the training data to Microsoft Foundry data files.

E. Store the training data in Azure Blob Storage that has version control enabled.

Answer: D,E (LEAVE A REPLY)

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics Designing an ALM process for fine #tuned Microsoft Foundry models requires two critical capabilities:

* Version-controlled training data

* A consistent, governed pipeline for retraining

Let's break down the reasoning using modern Agentic AI lifecycle , data governance , and model retraining best practices .

E). Store the training data in Azure Blob Storage that has version control enabled - # Correct This directly satisfies the requirement:

"Data changes must be tracked and versioned."

Azure Blob Storage with versioning provides:

- * Automatic version history for every training dataset
- * Immutable snapshots for audit and rollback
- * Governance controls for approved data
- * Integration with CI/CD pipelines for model retraining

In an agentic AI lifecycle, data versioning is mandatory because:

- * Training data evolves frequently
- * Retraining must be reproducible
- * Regulatory audits require traceability
- * Model drift must be monitored

Blob Storage with versioning is the Microsoft#recommended approach for enterprise AI ALM.

D). Upload the training data to Microsoft Foundry data files - # Correct Foundry fine #tuning jobs require training data to be stored in Foundry data files .

This ensures:

- * The fine #tuning job always uses the approved dataset
- * The model retraining pipeline is consistent
- * The data is validated and formatted correctly
- * The training job references a stable, governed data source

This aligns with the requirement:

"The model must be retrained consistently by using approved training data." In agentic AI systems, the training pipeline must be deterministic.

Uploading the data to Foundry data files ensures that the fine#tuning job always uses the correct dataset version.

Why the other options are NOT correct

A). Associate the storage location to the fine-tuning job - Not sufficient This does not provide:

- * Data versioning
- * Governance
- * Tracking of changes

It simply points the job to a location, not a controlled ALM process.

B). Create a content filter - Not related to ALM or training data

Content filters are for safety , not:

- * Versioning
- * Data governance
- * Retraining consistency

They do not help with the ALM requirements.

C). Store the training data in Azure Files - Not appropriate

Azure Files does not provide:

- * Built#in versioning
- * Immutable snapshots
- * ALM integration for ML pipelines

Blob Storage is the correct choice for AI training data.

Final Answer: D, E

- * D. Upload the training data to Microsoft Foundry data files

* E. Store the training data in Azure Blob Storage that has version control enabled These two actions together create a governed, versioned, repeatable ALM pipeline for fine #tuned Foundry models

NEW QUESTION: 30

A company has a Microsoft Copilot Studio agent that has been in production for three months.

The agent has received positive feedback from users.

You need to identify the number of questions unanswered by the agent and the number of abandoned sessions between the users and the agent.

Which Copilot Studio insights should you use? To answer, drag the appropriate insights to the correct requirements. Each insight may be used once, more than once, or not at all.

NOTE: Each correct selection is worth one point.

Answer:

Explanation:

The number of unanswered questions # Generated answer rate and quality; The number of abandoned sessions # Conversation outcomes Why "Generated answer rate and quality" is correct The requirement is to identify the number of questions unanswered by the agent . In Copilot Studio, unanswered-question behavior is tied to how often the agent successfully generates answers and the quality of those answers.

The Generated answer rate and quality insight is the right place to evaluate whether the agent is:

- * answering user questions
- * failing to generate answers
- * producing low-quality responses
- * missing knowledge coverage

From an AI business solutions standpoint, unanswered questions are a direct signal of knowledge gaps, grounding gaps, or prompt-response weaknesses. This is exactly what generated answer analytics are meant to surface.

Why "Conversation outcomes" is correct

The requirement also asks for the number of abandoned sessions between users and the agent.

Abandonment is a conversation-level outcome , not a reaction or survey result. The Conversation outcomes insight tracks what happened to the interaction, such as whether the conversation was:

- * resolved
- * escalated
- * abandoned
- * otherwise completed unsuccessfully

That makes it the correct metric for identifying abandoned sessions.

Why the other insights are not correct

Reactions

Reactions capture signals like positive or negative user feedback, but they do not directly measure unanswered questions or abandoned sessions.

Survey results

Survey results reflect user satisfaction feedback after interactions, but they do not directly quantify unanswered questions or abandonment counts.

NEW QUESTION: 31

You need to design a Microsoft Copilot Studio agent for customer support.

The agent must securely retrieve product warranty data from a REST API. The solution must minimize development effort. What should you include in the design?

- A.** Use a Microsoft Power Automate desktop flow to screen scrape the warranty data.
- B.** Export the agent as a managed solution and customize the agent in Power Apps.
- C.** Add the warranty data to the Fallback topic.
- D.** Create a custom connector in Copilot Studio and use the connector to call the API.

Answer: ([SHOW ANSWER](#))

The requirement is to build a Microsoft Copilot Studio agent that can securely retrieve product warranty data from a REST API while minimizing development effort .

The best choice is D. Create a custom connector in Copilot Studio and use the connector to call the API.

Why this is correct:

A custom connector is the standard low-code way to integrate Copilot Studio with a REST API. It lets the agent call the external warranty service securely, using supported authentication methods, without requiring heavy custom development. This directly satisfies both key requirements:

- * secure retrieval of warranty data
- * minimal development effort

Why the other options are incorrect:

- * A. Use a Microsoft Power Automate desktop flow to screen scrape the warranty data. Screen scraping is brittle, less secure, and unnecessary when a REST API already exists.
- * B. Export the agent as a managed solution and customize the agent in Power Apps. This does not directly solve secure API access and adds unnecessary complexity.
- * C. Add the warranty data to the Fallback topic. The Fallback topic is for unclear user input, not for live secure retrieval of external data.

Valid AB-100 Dumps shared by EduDump.com for Helping Passing AB-100 Exam! EduDump.com now offer the **newest AB-100 exam dumps**, the EduDump.com AB-100 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AB-100 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/AB-100/premium/> (**99** Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)

NEW QUESTION: 32

A company uses Microsoft Dynamics 365 finance and operations apps.

The company plans to use Microsoft Copilot in-app help and guidance to generate responses for internal business processes.

You need to add an additional knowledge source for the business processes. The solution must NOT add new topics to the Copilot agent for the finance and operations apps.

Which knowledge source should you add?

- A.** Microsoft Dataverse
- B.** a public website
- C.** Azure AI Search

D. a file upload

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answer is D. a file upload .

In this scenario, the company is using Microsoft Copilot in-app help and guidance in Dynamics 365 finance and operations apps and wants to extend the business-process knowledge available to Copilot without creating new topics .

That requirement matters because the solution is not asking for topic authoring or dialog design. It is asking for an additional knowledge source that Copilot can use directly for response generation.

A file upload is the correct choice because this approach allows you to provide business-process documentation directly as knowledge content for Copilot to reference, without requiring the creation of new conversational topics. This is especially useful for internal process guides, SOPs, policy documents, and operational instructions that already exist in document form.

Why D is correct

A file upload is appropriate when:

- * the knowledge already exists as documentation
- * you want Copilot to use that content directly
- * you do not want to build separate topics
- * the source is intended for internal business-process guidance

From an AI business solutions perspective, this is a strong pattern because it speeds up enablement. Instead of modeling every internal process as a structured dialog tree, the organization can ground Copilot on uploaded documentation and let it generate answers from that knowledge source.

This is a common enterprise AI design principle:

- * use topics when you need controlled conversation flow or task orchestration
- * use knowledge sources when you need grounded answers from existing documentation Since the requirement explicitly says must NOT add new topics , the best fit is a knowledge source such as file upload .

Why the other options are incorrect

A). Microsoft Dataverse

Dataverse is ideal for structured business data and transactional records, but this question is about adding a knowledge source for business processes in in-app help and guidance. That points more toward document- based knowledge than table-based operational data.

B). a public website

A public website can be a knowledge source in some Copilot scenarios, but here the requirement is centered on internal business-process guidance in finance and operations apps. Internal process content is more appropriately supplied through uploaded files than through a public website source.

C). Azure AI Search

Azure AI Search is useful in more advanced custom retrieval architectures, but it is not the direct built-in answer for this scenario. The question is asking for the supported knowledge-source addition without topic creation, and file upload is the simpler and more direct fit.

Expert reasoning

The deciding phrase in the question is:

- * "must NOT add new topics"

That means the solution should rely on knowledge grounding , not conversational authoring. Among the available options, file upload is the most direct and supported way to provide internal business-process knowledge for Copilot in this scenario.

NEW QUESTION: 33

A company extends Copilot in Microsoft Dynamics 365 Customer Service.

You need to recommend an automated application lifecycle management (ALM) process so that the Copilot components can be safely developed, tested, and promoted to production.

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

NOTE: Each correct selection is worth one point.

- A.** Use an unmanaged solution in production.
- B.** Rebuild the agents in each environment.
- C.** Use Microsoft Power Platform pipelines.
- D.** Include the components in a solution.
- E.** Store the agent transcripts in source control.

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answers are C. Use Microsoft Power Platform pipelines and D. Include the components in a solution .

This question is about implementing a proper ALM process for Copilot components in Microsoft Dynamics

365 Customer Service so they can be:

- * developed safely
- * tested consistently
- * promoted to production in a controlled way

That directly aligns with standard Power Platform ALM practices.

Why D. Include the components in a solution is correct

In Power Platform and Copilot-related environments, components should be packaged into a solution so they can be managed and transported across environments in a structured way.

Including the components in a solution enables:

- * dependency tracking
- * packaging of related assets together
- * environment-to-environment movement
- * better governance and change control
- * cleaner release management

From a business solutions architecture perspective, this is foundational. Without solutions, Copilot components are much harder to move consistently and govern properly across dev, test, and production.

Why C. Use Microsoft Power Platform pipelines is correct

Once the components are organized into a solution, Microsoft Power Platform pipelines provide the automated mechanism to promote them across environments.

Pipelines help with:

- * standardized deployments
- * safe promotion from development to test to production
- * reduced manual deployment errors
- * traceability of releases
- * repeatable and governed ALM execution

This is exactly what the question is asking for: an automated ALM process .

From an agentic AI business solutions perspective, automation in deployment is especially important because Copilot components can influence business workflows, customer interactions, and service operations. That means changes must be promoted in a disciplined and auditable way.

Why the other options are incorrect

A). Use an unmanaged solution in production

This is not recommended as a best practice for controlled enterprise production ALM. Production deployments should be governed and managed carefully, and unmanaged solutions are not the preferred pattern for that.

B). Rebuild the agents in each environment

This is inefficient, error-prone, and not an ALM best practice. It destroys consistency and traceability because each environment may end up with slight differences.

E). Store the agent transcripts in source control

Transcripts may be useful for analysis or audit in some contexts, but they are not a core ALM action for safely developing, testing, and promoting Copilot components.

Expert reasoning

For Microsoft business application ALM questions, the best-practice pattern is usually:

- * package artifacts in a solution

- * move them with Power Platform pipelines

That gives the cleanest answer for automated, governed promotion across environments.

NEW QUESTION: 34

A company has an Azure environment that supports multiple business units.

The company plans to implement an AI solution that will perform sentiment analysis on customer product reviews. You need to evaluate the potential cost of the solution to support return on AI investment (ROAI) analysis. What should you use?

A. Total Cost of Ownership (TCO) Calculator

B. Azure Reservations

C. Azure pricing calculator

D. Azure Monitor

Answer: ([SHOW ANSWER](#))

The requirement is to evaluate the potential cost of an AI solution before deployment so the company can support a return on AI investment (ROAI) analysis .

The correct choice is C. Azure pricing calculator .

Why C is correct:

The Azure pricing calculator is used to estimate the expected cost of Azure services before implementation.

For an AI solution such as sentiment analysis on customer product reviews , it helps estimate likely spending across the relevant Azure components and usage patterns. That makes it the right tool for cost forecasting as part of ROAI planning.

Why the other options are not correct:

- * A. Total Cost of Ownership (TCO) Calculator TCO is mainly used to compare broader infrastructure costs, often for migration or environment-level cost comparisons, not detailed service pricing estimation for a planned Azure AI solution.


- * B. Azure Reservations Reservations are a purchasing/discount mechanism, not the primary tool for estimating potential solution cost.

- * D. Azure Monitor Azure Monitor tracks operational metrics and runtime health after deployment; it is not used for pre-deployment cost estimation.


NEW QUESTION: 35

Which tools should you recommend to assist the CISO and the CIO with their specific responsibilities? To answer, drag the appropriate tools to the correct executives. Each tool 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:



Explanation:

Verified Answer : =

* CISO # Azure Resource Graph Explorer

* CIO # Microsoft Purview

Comprehensive and Detailed Explanation from Agentic AI Topics:

The CISO is responsible for discovering and inventorying AI resources for auditing across multiple Azure subscriptions. The best tool for that is Azure Resource Graph Explorer , because it is designed to query Azure resources at scale across subscriptions and provide visibility into what AI-related resources exist for governance and audit purposes.

The CIO is responsible for ensuring that appropriate security labels are assigned to the data used by the AI agents . The best tool for that is Microsoft Purview , because Purview supports data governance, classification, sensitivity labeling, and compliance management across enterprise data assets.

Why the other options are not correct:

* Azure Blob Storage is storage, not the primary tool for discovery/inventory or labeling governance.

* Copilot Studio is for building and managing agents, not for Azure-wide AI resource inventory or enterprise data labeling.

NEW QUESTION: 36

A company deploys agents that generate responses by using Azure OpenAI resources. The agents are deployed to both the United States and Europe. You need to recommend a governance solution that meets the following requirements:

Enforces the deployment of the resources to only approved Azure regions Provides continuous compliance verification of the resources

Enforces the deployment of the resources to only approved regions:

Azure Monitor

Azure Policy

Microsoft Defender for Cloud

Microsoft Purview

Microsoft Sentinel

Provides continuous compliance verification of the resources:

Azure Monitor

Azure Policy

Microsoft Defender for Cloud

Microsoft Purview

Microsoft Sentinel

Answer:

Enforces the deployment of the resources to only approved regions:

Azure Monitor

Azure Policy

Microsoft Defender for Cloud

Microsoft Purview

Microsoft Sentinel

Provides continuous compliance verification of the resources:

Azure Monitor

Azure Policy

Microsoft Defender for Cloud

Microsoft Purview

Microsoft Sentinel

Explanation:

Enforces deployment to only approved Azure regions # Azure Policy; Provides continuous compliance verification # Microsoft Defender for Cloud Why Azure Policy is correct The requirement is to enforce that Azure OpenAI resources can be deployed only in approved Azure regions .

That is exactly what Azure Policy is designed to do. Azure Policy allows organizations to create and assign rules that govern resource deployment and configuration. For regional restrictions, you can define a policy that permits deployments only in allowed locations and denies deployments elsewhere.

From an AI business solutions and cloud governance perspective, Azure Policy is the right preventive control because it acts at deployment time. It helps enforce organizational standards before noncompliant resources are created.

Typical policy use cases include:

- * restricting allowed Azure regions
- * enforcing approved SKUs
- * requiring tags
- * limiting resource types
- * ensuring security configuration standards

This is especially important for AI deployments where geography may affect:

- * regulatory compliance
- * data residency
- * internal governance
- * customer contract obligations

Why Microsoft Defender for Cloud is correct

The second requirement is to provide continuous compliance verification of the resources.

That points to Microsoft Defender for Cloud .

Defender for Cloud continuously assesses Azure resources against security and compliance standards. It provides visibility into resource posture, identifies misconfigurations, and tracks compliance status over time.

This makes it well suited for ongoing verification because it supports:

- * continuous assessment
- * compliance dashboards
- * security posture monitoring
- * recommendations for remediation
- * regulatory standard mapping

In enterprise AI deployments, this is critical because governance is not only about blocking bad deployments.

It is also about continuously validating that deployed resources remain compliant as environments evolve.

Why the other options are incorrect

Azure Monitor

Azure Monitor is used for telemetry, logging, metrics, and observability. It is not the primary service for enforcing allowed regions or for formal continuous compliance governance.

Microsoft Purview

Microsoft Purview focuses on data governance, data cataloging, classification, and compliance across data estates. It is not the main control for Azure resource deployment region enforcement.

Microsoft Sentinel

Microsoft Sentinel is a SIEM/SOAR platform for security analytics and threat detection. It is not the service used to enforce deployment locations, and it is not the primary tool for continuous Azure resource compliance verification.

Azure Policy for continuous verification

Azure Policy does provide compliance views, but in this question, the stronger mapping for continuous compliance verification is Microsoft Defender for Cloud , which is specifically designed for continuous security posture and compliance assessment across resources.

Expert reasoning

Use this exam pattern:

- * Prevent or restrict how Azure resources are deployed # Azure Policy
- * Continuously assess and verify cloud compliance posture # Microsoft Defender for Cloud

NEW QUESTION: 37

A company has an AI solution that uses Azure OpenAI models.

You need to recommend a governance solution that monitors and audits changes to model configurations and data usage. The solution must minimize administrative effort.

What should you include in the recommendation?

- A.** Azure Monitor
- B.** Azure Stream Analytics
- C.** Azure API Management

D. Azure Policy

E. Microsoft Purview

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answer is E. Microsoft Purview .

This question is centered on governance , specifically the need to:

- * monitor changes to model configurations
- * audit data usage
- * minimize administrative effort

That combination points most strongly to Microsoft Purview .

Why E is correct

Microsoft Purview is Microsoft's core platform for data governance, compliance, auditing, information protection, and lifecycle oversight . When an organization is using Azure OpenAI models and needs a governance-oriented solution for monitoring and auditing how data is used, Purview is the best fit among the listed options.

From an AI business solutions perspective, governance is broader than infrastructure monitoring. It includes:

- * understanding how sensitive data is handled
 - * tracking access and usage patterns
 - * supporting audit and compliance needs
 - * helping investigate data exposure concerns
 - * enforcing information governance practices across AI-enabled workloads
- Purview is especially strong when the requirement includes auditing data usage because that is a governance and compliance concern, not just a performance or telemetry concern.

It also minimizes administrative effort because it provides centralized governance capabilities rather than requiring the company to stitch together multiple lower-level services for oversight.

Why the other options are incorrect

A). Azure Monitor

Azure Monitor is useful for telemetry, logs, metrics, and operational monitoring. It helps observe system performance and activity, but it is not the best primary governance solution for auditing data usage and broader compliance oversight.

B). Azure Stream Analytics

This service is used for real-time stream processing and analytics. It does not address governance and audit requirements for Azure OpenAI model configurations and data usage.

C). Azure API Management

API Management helps publish, secure, and manage APIs. It is valuable for access mediation and control, but it is not the main governance and auditing platform for data usage and model-configuration oversight.

D). Azure Policy

Azure Policy is very strong for enforcing resource configuration standards and compliance rules at deployment and configuration time. However, the question also emphasizes auditing data usage , which is better aligned to Purview's governance capabilities. Policy is more about enforcement of resource state; Purview is stronger for governance, auditing, and data oversight.

Expert reasoning

Use this exam shortcut:

- * Need operational logs and metrics # Azure Monitor
 - * Need deployment/configuration enforcement # Azure Policy
 - * Need data governance, auditing, compliance, and information oversight # Microsoft Purview
- Because the question emphasizes both changes and data usage auditing with a governance lens, Microsoft Purview is the strongest answer.

NEW QUESTION: 38

A company has a Microsoft Copilot Studio agent that provides answers based on a knowledge base for customer support.

Users report that, occasionally, the agent provides inaccurate answers.

You need to use metrics from the Analytics tab in Copilot Studio to identify the cause of the inaccuracies.

Which two options should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. survey results

B. session information and session outcomes

C. topic usage and topics with low resolution

D. engagement, resolution, and escalation rates

E. quality of generated answers

Answer: ([SHOW ANSWER](#))

Comprehensive and Detailed Explanation From Agentic AI Business Solutions Topics:

The correct answers are B. session information and session outcomes and E. quality of generated answers .

This scenario is focused on a knowledge base-driven Copilot Studio agent where users report that the agent sometimes gives inaccurate answers . The question asks which Analytics tab metrics should be used to identify the cause of those inaccuracies.

That means you need metrics that help you examine:

- * how the answer was generated

- * what happened in the conversation when the bad answer occurred

Why E. quality of generated answers is correct

This is the most direct metric for this scenario.

Because the agent is answering from a knowledge base , the problem is tied to the quality of the generated response itself. The quality of generated answers metric helps assess whether the generated responses are relevant, useful, and accurate enough for the user's request.

From an AI business solutions perspective, this metric is essential because it helps diagnose problems such as:

- * weak grounding from the knowledge source

- * irrelevant retrieval

- * poor answer formulation

- * hallucination-like behavior

- * mismatch between user question and available source content

If the issue is inaccurate answers, the first place to investigate is the quality signal tied to generated answers.

Why B. session information and session outcomes is correct

To find the cause of inaccuracies, you also need to inspect the broader conversational context. Session information and session outcomes help you see:

- * what the user asked

- * how the agent responded

- * whether the conversation was resolved

- * whether the user abandoned, escalated, or retried

- * where the conversation broke down

This is important because an inaccurate answer may not come only from poor generation quality. It may also come from:

- * the way the user phrased the request

- * lack of sufficient grounding context

- * repeated failed attempts in a session

- * escalation after an unhelpful answer

* patterns in unsuccessful conversations

In other words, quality of generated answers tells you about answer quality, while session information and outcomes help you understand the operational context in which those inaccuracies appear.

Together, these two give the strongest diagnostic view.

Why the other options are incorrect

A). survey results

Survey results can tell you whether users were happy or unhappy, but they do not directly help identify the cause of inaccurate knowledge-based responses. They are more of a feedback signal than a root-cause metric.

C). topic usage and topics with low resolution

This is more relevant for agents built around explicit topics and topic flows. The scenario specifically describes an agent that provides answers based on a knowledge base , so generated-answer analytics are more appropriate than topic-resolution analysis.

D). engagement, resolution, and escalation rates

These are useful high-level operational KPIs, but they are not the best metrics for diagnosing why answers are inaccurate. They show outcome trends, not the direct cause of answer-quality issues.

NEW QUESTION: 39

A company has Microsoft Foundry agents that generate responses by using Azure OpenAI resources. The agents are deployed to both the United States and Europe.

A company mandate states that the agents and their grounding data must adhere to data residency and movement regulations.

You need to recommend a governance solution for the agents.

What should you include in the recommendation?

A. Microsoft Defender for Cloud

B. Azure Policy

C. Azure Monitor

D. Microsoft Purview

Answer: ([SHOW ANSWER](#))

<https://learn.microsoft.com/en-us/azure/governance/policy/overview>

Valid AB-100 Dumps shared by EduDump.com for Helping Passing AB-100 Exam! EduDump.com now offer the **newest AB-100 exam dumps**, the EduDump.com AB-100 exam **questions have been updated** and **answers have been corrected** get the **newest** EduDump.com AB-100 dumps with Test Engine here: <https://www.edudump.com/exams/Microsoft/AB-100/premium/> (99 Q&As Dumps, **35%OFF** Special Discount Code: **freecram**)