QUESTION 1
In which scenario is Azure Machine Learning most likely to deliver strategic value for an organization?

A. Using historical sales data to forecast demand across product categories.
B. Digitizing a paper-based process to reduce errors.
C. Entering customer feedback into a spreadsheet to understand sentiment.
D. Sending personalized emails to customers based on the customer location.

Answer: A
Explanation:
Azure Machine Learning (Azure ML) delivers strategic value by transforming historical sales data into a competitive advantage through advanced demand forecasting. By identifying complex patterns in past consumer behavior, it helps businesses optimize high-stakes operational areas such as inventory management, production planning, and resource allocation…..
Benefits
Inventory Optimization: Businesses can maintain leaner inventory levels, drastically reducing storage costs and minimizing the risk of both stockouts and overstocking.
Financial Performance: Improved forecast accuracy directly protects margins by reducing the need for emergency shipping, overtime labor, and waste from unsold goods.
Strategic Growth: Accurate long-term forecasts provide a reliable roadmap for planning product launches, marketing promotions, and market expansion.
Operational Agility: Azure ML's Automated Machine Learning (AutoML) allows companies to quickly adapt to market shifts—like seasonal trends or unexpected disruptions—by continuously learning from new data.
Reference:
https://learn.microsoft.com/en-us/azure/architecture/ai-ml/idea/next-order-forecasting

QUESTION 2
Hotspot Question
Select the answer that correctly completes the sentence.

Answer Area

The primary goal of generative AI is [ ]

- to analyze trends and classify data sources.
- to create new content, such as text, images, or code.
- to make predictions based on historical data.

Answer:

Answer Area

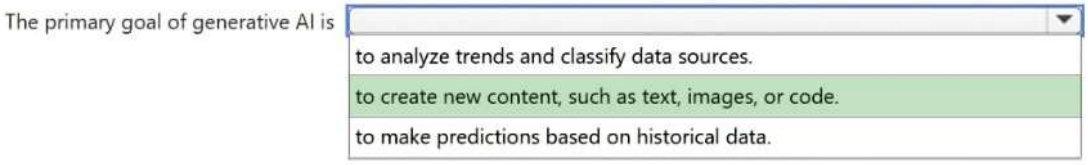

Explanation:
Box: to create new content, such as text, images, or code.
The primary goal of generative AI is _____.
Generative AI (GenAI) is a type of artificial intelligence designed to create new, original content—including text, images, videos, audio, and code—by learning patterns from large, existing datasets. Unlike traditional AI that analyzes or classifies data, GenAI produces unique, human-like outputs, such as written stories, realistic images, or computer code.
Reference:
https://www.ai21.com/glossary/foundational-llm/generative-ai/

QUESTION 3
Your company discovers that several employees use personal ChatGPT accounts to assist with work tasks. You are concerned about proprietary data being shared externally.
You need to evaluate the business value of rolling out Microsoft 365 Copilot.
Which capability is a key benefit of using Copilot instead of a personal ChatGPT account?

A. analyzing and producing reports based on complex data
B. generating ideas and solving issues
C. drafting documents, emails, presentations, and marketing materials
D. accessing internal data in accordance with existing Microsoft 365 policies

Answer: D
Explanation:
A major, defining advantage of Microsoft 365 Copilot over a personal ChatGPT account is its deep, native integration with an organization's internal data—including emails, documents, chats, and meetings—while strictly adhering to existing Microsoft 365 security, compliance, and privacy policies.

Here is a breakdown of why this is a critical differentiator:
1. Access to Internal Data ("Grounding")
Microsoft 365 Copilot: Accesses your organization's data via Microsoft Graph. It can summarize, analyze, and create content based on your Word documents, emails in Outlook, spreadsheets in Excel, and meetings in Teams.
Personal ChatGPT: Does not have access to your private company files, emails, or internal systems unless you manually copy and paste that information into the chat.

2. Adherence to Security and Compliance Policies
Microsoft 365 Copilot: Inherits your organization's existing security configurations, such as sensitivity labels, Data Loss Prevention (DLP) policies, and identity-based access controls. If you do not have permission to view a file, Copilot will not use that file to answer your prompt.
Personal ChatGPT: Operates outside your corporate security boundary. Using a personal account to analyze company data can risk leaking confidential information to a third-party, which is typically against corporate security policies.
Reference:
https://www.microsoft.com/en-us/microsoft-365-copilot/copilot-vs-chatgpt-enterprise

QUESTION 4
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
A manufacturer can use Azure Vision in Foundry Tools to identify product defects on an assembly line.	○	○
A logistics company can use Azure Vision in Foundry Tools to recognize package shipping labels.	○	○
The HR department at your company can only use Azure Vision in Foundry Tools to extract written content from Microsoft Word files.	○	○

Answer:

Answer Area

Statements	Yes	No
A manufacturer can use Azure Vision in Foundry Tools to identify product defects on an assembly line.	○	○
A logistics company can use Azure Vision in Foundry Tools to recognize package shipping labels.	○	○
The HR department at your company can only use Azure Vision in Foundry Tools to extract written content from Microsoft Word files.	○	○

Explanation:
Box 1: Yes
Yes - A manufacturer can use Azure Vision in Foundry Tools to identify product defects on an assembly line.
A manufacturer can use Azure Vision in Foundry Tools (formerly part of Azure AI Services, now integrated within the AI Foundry toolkit) to identify product defects on an assembly line. This solution automates visual inspections to detect anomalies such as surface scratches, cracks, misalignments, or missing components in real time.
Box 2: Yes
Yes - A logistics company can use Azure Vision in Foundry Tools to recognize package shipping labels.
A logistics company can use Azure Vision in Foundry Tools (part of the broader Azure AI services suite) to recognize, interpret, and digitize package shipping labels. By integrating Azure's advanced AI with Palantir Foundry, firms can automate manual data entry, track shipments, and improve operational efficiency.
Box 3: No
No - The HR department at your company can only use Azure Vision in Foundry Tools to extract written content from Microsoft Word files.
Azure Vision in Foundry Tools is primarily designed for images, while its sibling tool, Document Intelligence, handles Microsoft Word files.
Reference:
https://datalabs.io/azure-ai-for-smart-manufacturing-defect-detection-with-computer-vision
https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/overview
https://learn.microsoft.com/en-us/azure/ai-services/computer-vision/concept-ocr

QUESTION 5
Hotspot Question
Select the answer that correctly completes the sentence.

Answer Area

When you use Microsoft 365 Copilot connectors to connect external content to [ ], your users can find, summarize, and learn from line-of-business (LOB) data by using natural language prompts.

- Azure AI Search
- Microsoft Graph
- Microsoft Purview
- Microsoft SharePoint

Answer:

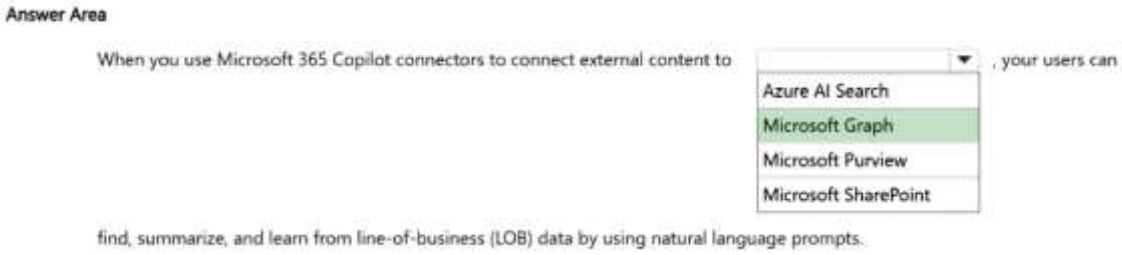

Explanation:
Box: Microsoft Graph
When you use Microsoft 365 Copilot connectors to connect to external content to _____, your users can find, summarize, and learn from line-of-business (LOB) data by using natural language prompts.
By using Microsoft Graph connectors, you can index your third-party data sources—such as file shares, ServiceNow, or SQL databases—directly into the Microsoft Graph.
Once indexed, this LOB data becomes part of the "knowledge base" that Microsoft 365 Copilot can access. This allows users to ask questions like, "What is the status of the tickets in ServiceNow?" or "Summarize the project specs from our internal wiki," and receive grounded, context-aware responses.
Reference:
https://learn.microsoft.com/en-us/graph/connecting-external-content-experiences

QUESTION 6
Your company manages an online catalog of office supplies.
You plan to use a generative AI solution to create product descriptions for your company's website. The solution must meet the following requirements:
- Ensure that the descriptions can be posted immediately after they are created.
- Enable the selection and inclusion of product details in each description.
- Be fast and simple for non-technical staff to use.

What is the best type of solution to use? More than one answer choice may achieve the goal. Select the BEST answer.

A. custom Azure Machine Learning model
B. the Researcher agent in Microsoft 365 Copilot
C. an interactive AI agent
D. a fine-tuned large language model (LLM)

Answer: B
Explanation:
Using the Researcher agent within Microsoft 365 Copilot provides a highly effective solution for creating and immediately posting product descriptions. It allows non-technical staff to generate tailored, brand-aligned content by leveraging both internal product data and web research, allowing for immediate publication.
Reference:
https://learn.microsoft.com/en-us/dynamics365/business-central/ai-overview

QUESTION 7
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

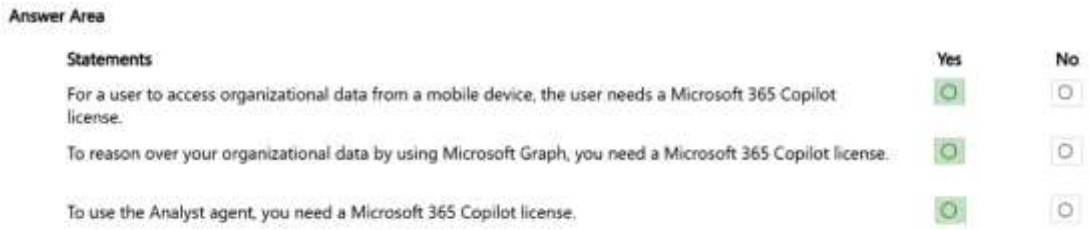
Answer Area

Statements	Yes	No
For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license.	○	○
To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license.	○	○
To use the Analyst agent, you need a Microsoft 365 Copilot license.	○	○

Explanation:
Box 1: Yes
Yes - For a user to access organizational data from a mobile device, the user needs a Microsoft 365 Copilot license.
To access, summarize, and query organizational data (such as emails, chats, documents in SharePoint/OneDrive, and calendar items) via Microsoft 365 Copilot on a mobile device, a user must have a Microsoft 365 Copilot license assigned to them.

This license is an add-on to a qualifying base subscription (such as Microsoft 365 E3, E5, Business Standard, or Business Premium).
Box 2: Yes
Yes - To reason over your organizational data by using Microsoft Graph, you need a Microsoft 365 Copilot license.
To use the advanced AI reasoning capabilities of Microsoft 365 Copilot—specifically to analyze, summarize, and query your organizational data (emails, chats, documents, meetings) via Microsoft Graph—you need a Microsoft 365 Copilot license.
Box 3: Yes
Yes - To use the Analyst agent, you need a Microsoft 365 Copilot license
To use the Analyst agent, you generally need a Microsoft 365 Copilot add-on license.
While a basic version of Copilot Chat is available for many Microsoft 365 and Office 365 subscribers at no extra cost, advanced "Frontier" agents like Analyst and Researcher are specifically built for deep reasoning and multi-step tasks, which are reserved for licensed Copilot users.
Reference:
https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-minimum-requirements
https://learn.microsoft.com/en-us/copilot/faq
https://it.osu.edu/news/2025/07/22/new-microsoft-365-copilot-agents-available-research-and-analysis

QUESTION 8
Your company plans to adopt AI across multiple business units.
You need to ensure that all AI projects align with the company's business strategy and are implemented responsibly.
What is the best approach to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

A. Allow each department to deploy its own AI tools and workflows.
B. Delegate AI decision-making to the company's IT department.
C. Outsource AI development to an external vendor.
D. Establish an AI council to provide guidance, oversight, and coordination.

Answer: D
Explanation:
An AI council is a cross-functional, board-level advisory body designed to align AI initiatives with corporate strategy, ensuring projects are ethically, legally, and fiscally responsible. It provides oversight, manages risks, and fosters, cross-departmental coordination, crucial for driving adoption and avoiding siloed, unaligned AI projects.
Benefits
Enhanced Decision-Making: Coordinated, expert-driven input leads to faster, better-aligned decisions.
Trusted AI: Builds, trust through transparent, non-biased, and, accountable, systems.
Value Realization: Ensures AI investments deliver measurable value to the organization.
Reference:
https://cognitivepath.com/ai-councils

QUESTION 9
Your company has a Microsoft 365 subscription and uses Microsoft 365 Copilot Chat.
Some users need to build and use declarative agents that can access work data.

Which type of license should you recommend for the users?

A. a Microsoft 365 Copilot add-on license
B. Microsoft Copilot Studio user license
C. a Copilot Chat pay-as-you-go plan

Answer: A
Explanation:
To use declarative agents that access work data (such as SharePoint or Graph connectors), users generally need a Microsoft 365 Copilot add-on license.

While Microsoft 365 Copilot Chat is included in most business subscriptions at no extra cost, its native capabilities are primarily limited to web-grounding and basic instructions. Accessing organizational data via agents typically requires one of the following licensing paths:

1. Microsoft 365 Copilot Add-on License
This is the most direct method. It grants full access to:
Declarative agents grounded in tenant data without additional usage fees.
Copilot Studio for authoring and managing these agents.
Embedded Copilot features in Word, Excel, Outlook, and Teams.

2. Pay-As-You-Go (Consumption) Model
If users do not have a full Copilot add-on license, organizations can enable metered usage (consumption-based billing).

Usage-based billing: Interactions with agents that access tenant data (SharePoint, connectors) consume "Copilot credits".
Requirement: This requires an Azure subscription and a billing policy set up in the Microsoft 365 admin center.
Reference:
https://support.microsoft.com/en-gb/topic/how-copilot-chat-works-with-and-without-a-microsoft-365-copilot-license-5810b659-fbe0-48ee-9fe6-d731fe86cdeb

QUESTION 10
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Allowing AI models to make autonomous decisions supports the Microsoft responsible AI principle of accountability.	○	○
Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles.	○	○
Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security.	○	○

Answer:

Answer Area

Statements	Yes	No
Allowing AI models to make autonomous decisions supports the Microsoft responsible AI principle of accountability.	○	○
Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles.	○	○
Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security.	○	○

Explanation:
Box 1: No
No - Allowing AI models to make autonomous decisions support Microsoft AI principle of accountability.
Microsoft's principle of accountability actually mandates that humans, not AI models, remain the final authority for how a system operates. While AI can perform automated tasks, the accountability principle requires that the people who design and deploy these systems take responsibility for their impact and maintain meaningful control.

Box 2: Yes
Yes - Regularly testing AI models for fairness and inclusiveness helps ensure they align with Microsoft's Responsible AI principles.
Regularly testing AI models for fairness and inclusiveness is a foundational practice within Microsoft's Responsible AI Standard, which acts as a guide for developing and deploying AI systems. This continuous testing ensures that AI applications do not reinforce historical biases and perform equitably across different demographic groups, including race, gender, age, and background.
Box 3: Yes
Yes - Protecting user data and limiting access to personal information supports the Microsoft responsible AI principles of privacy and security.
Protecting user data and limiting access to personal information are, in fact, foundational to Microsoft's Responsible AI principles of Privacy and Security. Microsoft's AI framework mandates that AI systems are developed and deployed in a manner that respects user privacy and maintains strict data security, aiming for AI systems that are "secure by design".
Reference:
https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai
https://techcommunity.microsoft.com/blog/nonprofittechies/the-importance-of-responsible-ai-a-comprehensive-guide/4404347

QUESTION 11
Your company plans to implement a proof of concept (PoC) agent that uses Azure OpenAI.
The solution must start small and provide flexibility to scale usage as demand grows.
Which pricing model should you use?

A. Microsoft 365 Copilot
B. Batch API
C. Provisioned (PTUs)
D. Standard (On-Demand)

Answer: D
Explanation:
The Standard (On-Demand) tier is the best choice for this scenario because it follows a pay-as-you-go consumption model. This allows a company to start a Proof of Concept (PoC) with virtually zero upfront cost or commitment, paying only for the tokens processed. As demand grows, the service provides the flexibility to scale without needing to manage complex capacity planning early on.
Reference:
https://azure.microsoft.com/en-us/products/ai-foundry/models/openai

QUESTION 12
Hotspot Question
Select the answer that correctly completes the sentence.

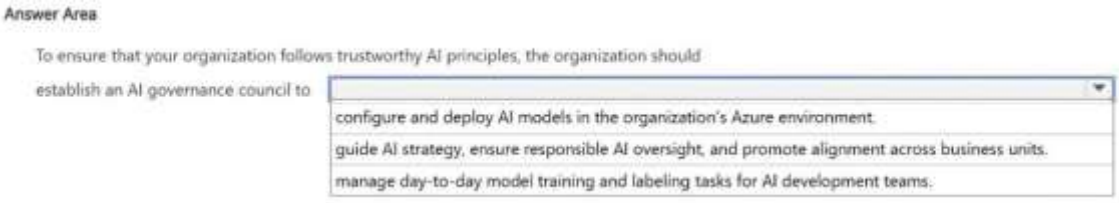

Answer:

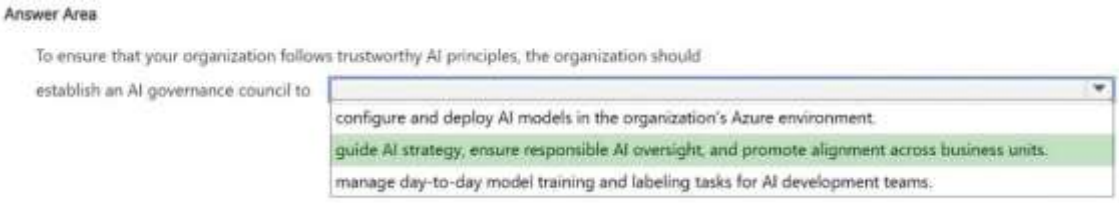

Explanation:
Box: guide AI strategy, ensure responsible AI oversight, and promote alignment across business units.
To ensure that your organization follows trustworthy AI principles, the organization should establish an AI governance council to _____.

An AI governance council should be a cross-functional, multidisciplinary body that oversees the development, deployment, and evaluation of AI solutions to ensure they are ethical, secure, and aligned with organizational values. Key functions include establishing accountability, managing risks, ensuring regulatory compliance (e.g., privacy), and promoting transparency and fairness.
Core Responsibilities
The council's mandate typically covers the entire AI lifecycle:
*-> Strategy & Policy: Developing internal AI usage policies and identifying high-risk use cases.
*-> Risk Oversight: Conducting impact assessments for bias, privacy, and security before deployment.
Continuous Monitoring: Tracking live models for "drift" (performance degradation over time) or emerging ethical issues.
Culture & Literacy: Driving organization-wide training to build AI fluency and a responsible AI culture.
Audit & Reporting: Maintaining transparent records and "model cards" to demonstrate accountability to stakeholders and regulators.
Reference:
https://athena-solutions.com/ai-governance-framework-2025/

QUESTION 13
Your company deploys an AI-powered loan approval solution that enables applicants to request an explanation as to why their loan application was denied.
Which Microsoft responsible AI principle is this an example of?

A. fairness
B. privacy and security
C. transparency
D. inclusiveness

Answer: C
Explanation:
According to Microsoft's guidelines, transparency means that AI systems should be understandable, and users should be able to understand the system's decisions or recommendations. Providing an explanation for a loan denial allows applicants to understand how the AI arrived at its decision.
Reference:
https://www.linkedin.com/pulse/deep-dive-responsible-ai-digitalbricksai-typie

QUESTION 14
Your company wants to ensure that AI solutions are used responsibly and align with company values and compliance requirements.
You need to establish governance principles for AI use.
Which two actions should you perform? Select the two BEST answers. Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Allow each department to tailor governance processes for its own AI initiatives.
B. Define accountability norms for AI decisions across business and technical teams.
C. Create a process to review AI initiatives for responsible AI alignment.
D. Assign governance ownership primarily to the AI engineering and data science teams.
E. Focus governance efforts on AI systems that handle regulated or sensitive data.

Answer: BC
Explanation:
To establish robust governance for AI, a company must define accountability norms that bridge the gap between business objectives and technical execution. Effective frameworks ensure that specific individuals—not just committees—are responsible for AI outcomes across the entire lifecycle.
Core Accountability Strategies
Establish Clear Ownership: Assign a specific person or team as the "accountable owner" for each AI application. This includes defining decision rights for who can approve, modify, or retire a system.
To establish an effective AI governance framework, you should implement a structured review process that translates high-level ethical principles into operational practice. This process serves as a "brake pedal" for the organization, ensuring innovation remains safe and aligned with corporate values.
Reference:

https://www.databricks.com/blog/ai-governance-best-practices-how-build-responsible-and-effective-ai-programs

QUESTION 15
Your company is deploying Microsoft 365 Copilot. The deployment must provide users with access to the Researcher agent to search across data in Microsoft SharePoint.
You need to recommend a licensing plan for the solution.
What should you recommend?

A. a Microsoft 365 Copilot per-user add-on license
B. pay-as-you-go
C. a Microsoft 365 subscription entitlement
D. a usage-based consumption license in Azure

Answer: A
Explanation:
To use the Researcher agent for searching across Microsoft SharePoint data, you need the Microsoft 365 Copilot add-on license.

The Researcher agent is a first-party reasoning agent included in the core experience for users with a paid Copilot license. It is designed for multi-step, in-depth research grounded in your Microsoft Graph data, including SharePoint, OneDrive, emails, and Teams.

Required Licensing Components
To deploy this capability, your organization must have:
1. A Qualifying Base License:
Enterprise: Microsoft 365 E3 or E5, or Office 365 E3 or E5.
Business: Microsoft 365 Business Standard or Business Premium.
*-> 2. Microsoft 365 Copilot Add-on:
This paid license (typically ~$30/user/month) unlocks the Researcher agent and the ability to search internal tenant data like SharePoint.
Reference:
https://learn.microsoft.com/en-us/copilot/microsoft-365/faq-researcher

QUESTION 16
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
A barrier to AI adoption can include data limitations and a lack of AI readiness and skills.	○	○
Organizations often struggle with AI adoption because they prioritize technology selection before defining clear business use cases.	○	○
A lack of cross-functional collaboration is a common barrier to AI adoption.	○	○

Answer:

Answer Area

Statements	Yes	No
A barrier to AI adoption can include data limitations and a lack of AI readiness and skills.	○	○
Organizations often struggle with AI adoption because they prioritize technology selection before defining clear business use cases.	○	○
A lack of cross-functional collaboration is a common barrier to AI adoption.	○	○

Explanation:
Box 1: Yes
Yes - A barrier to AI adoption can include data limitations and lock of AI readiness and skills.
Those are two of the most significant hurdles organizations face today. In fact, a 2025 PEX Report found that 52% of professionals cite data quality and availability as their primary challenge, closely followed by a lack of internal expertise

at 49%.
Box 2: Yes
Yes - Organizations often struggle with AI adoptions because they prioritize technology selection before defining clear business use cases.
It is essentially the "hammer looking for a nail" problem; organizations often get dazzled by the "hammer" (AI) and start swinging before they even know what they are trying to build.
Starting with technology rather than business value is a primary reason why 88% to 95% of AI pilots fail to deliver measurable results. This "technology-first" trap often leads to "random acts of AI"—expensive experiments that succeed in a lab but fail to solve any real operational bottlenecks.
Box 3: Yes
Yes - A lack of cross-functional collaboration is a common barrier to AI adoption.
That is spot on—it's often the "human" silos, not the hardware, that stall AI progress. While technical hurdles like data quality are significant, research consistently identifies a lack of cross-functional collaboration as a primary "hidden" barrier to successful AI adoption.
Reference:
https://www.ibm.com/think/insights/ai-adoption-challenges
https://www.rapidops.com/blog/why-ai-fails/
https://www.gartner.com/peer-community/post/preferred-tactics-building-effective-collaboration-cross-functional-teams-involved-ai-governance-risk-management-e-g-joint

QUESTION 17
What is considered a best practice when forming an AI adoption team in an enterprise environment?

A. Include primarily IT and project management staff initially to streamline deployment, adding governance and compliance roles later.
B. Include procurement and vendor management specialists early to evaluate AI tools, involving business teams once a platform is selected.
C. Include representatives from legal, leadership, and business units to align AI initiatives with organizational prioritie
D. Include only data scientists and engineers at first to validate technical feasibility, then add other stakeholders later.

Answer: C
Explanation:
Forming a cross-functional AI adoption team is a foundational best practice for enterprise environments.
A diverse "AI Center of Excellence" (CoE) or steering committee ensures that technical capabilities do not develop in isolation from regulatory requirements or business goals.
Key Representatives & Their Roles
*-> Executive Leadership: Champions the vision, secures budget, and ensures the AI strategy aligns with high-level corporate priorities.
*-> Legal & Compliance: Manages risk related to data privacy (e.g., GDPR), intellectual property, and evolving AI regulations to maintain stakeholder trust.
*- Business Units: Identify high-value use cases, define success metrics (KPIs), and ensure the AI tools actually solve operational pain points.
IT & Data Science: Provides the technical architecture, manages data pipelines, and handles the actual deployment and monitoring of models.
Change Management: Focuses on the "human" side of adoption, including upskilling employees and addressing fears about job displacement.
Reference:
https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/scenarios/ai/center-of-excellence

QUESTION 18
Your company is developing an AI-powered customer support agent.
You need to ensure that the solution follows Microsoft responsible AI principles.
Which two actions should you perform? Select the two BEST answers. Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Retain all customer conversations.

B. Test the agent to ensure that responses are inclusive and culturally sensitive.
C. Enable the agent to operate independently.
D. Ensure that the agent can be used for multiple purposes.
E. Provide a clear disclaimer that users are interacting with an AI solution.

Answer: BE
Explanation:
[B]
Under Microsoft's Responsible AI framework, this testing specifically addresses the principles of Fairness and Inclusiveness. To operationalize these, you should focus on the following testing and development areas:
Key Testing Focus Areas
*-> Cultural Sensitivity: Ensure the agent respects different values, perspectives, and international contexts to avoid causing offense or misunderstanding.
*-> Fairness and Bias Mitigation: Test the agent to ensure it treats all users equitably and does not reinforce societal stereotypes or discriminate based on protected characteristics like race, gender, or religion.
Accessibility: Validate that the agent is usable by people of all abilities, providing equal power and engagement regardless of their background.
Global Community Engagement: Involve diverse users and underserved communities in the pre-deployment validation and feedback process to identify representation gaps.
[E]
To align with Microsoft's Responsible AI principles, providing a clear disclaimer that users are interacting with an AI solution is a core requirement of the Transparency principle.
Under this principle and the Microsoft Responsible AI Standard, developers must ensure that AI systems are understandable and that users are not deceived into believing they are interacting with a human.
Reference:
https://learn.microsoft.com/en-us/legal/ai-code-of-conduct
https://www.microsoft.com/en-us/ai/principles-and-approach

QUESTION 19
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

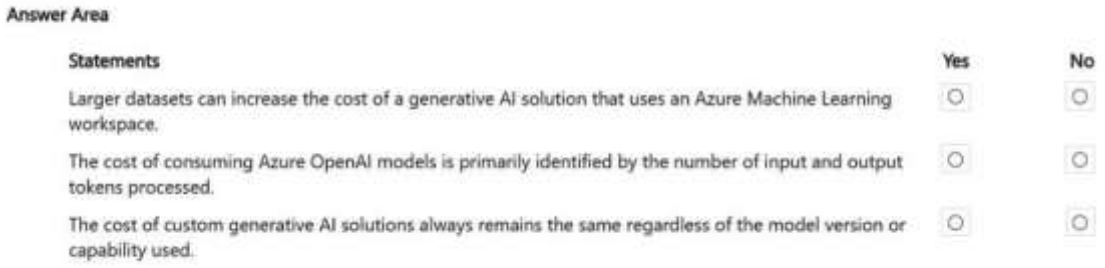

Answer:

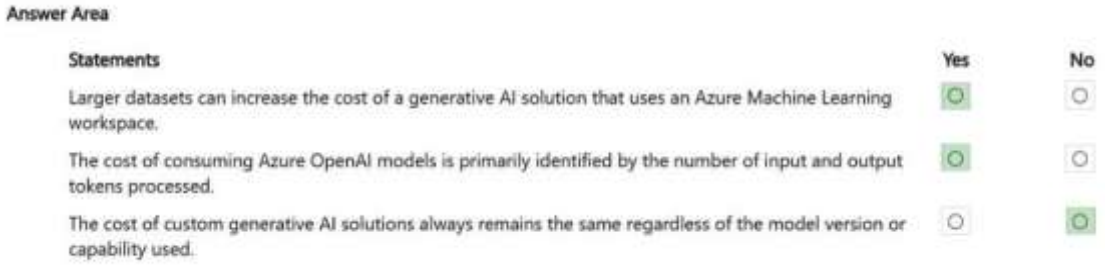

Explanation:
Box 1: Yes
Yes - Larger datasets can increase the cost of generative AI solution that uses an Azure Machine Learning workspace. larger datasets can increase the cost of a generative AI solution using an Azure Machine Learning (AML) workspace. While the AML workspace itself has no additional charge, larger datasets increase expenses across several underlying Azure services, particularly through higher storage, compute, and data transfer requirements.
Box 2: Yes

Yes - The cost of consuming Azure OpenAI models is primarily identified by the number of input and output tokens processed.
The cost of consuming Azure OpenAI models is primarily determined by the number of input and output tokens processed.

This consumption-based, pay-as-you-go model calculates costs based on the total volume of text (or image/audio data) sent to the model (input) and the text generated by the model (output), usually billed per 1,000 or 1 million tokens, depending on the model and pricing page updates.
Box 3: No
No - The cost of custom generative AI solutions always remains the same regardless of the model version or capability used
he cost of custom generative AI solutions varies significantly based on the model's version, complexity, and capability. While a basic chatbot might range from $40,000 to $150,000, advanced enterprise-grade platforms with high-risk reasoning can exceed $400,000 to $500,000+.
Reference:
https://www.doit.com/blog/optimizing-ml-costs-with-azure-machine-learning
https://www.finout.io/blog/azure-openai-pricing
https://medium.com/@dejanmarkovic_53716/custom-ai-solutions-cost-guide-2025-pricing-insights-revealed-cf19442261ec

QUESTION 20
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Using incomplete or poor-quality data during generative AI model training can increase costs.	○	○
AI models rely on training data to learn patterns and identify relationships to produce outputs.	○	○
Generative AI models trained on non-representative datasets can produce inaccurate or unbalanced results.	○	○

Answer:

Answer Area

Statements	Yes	No
Using incomplete or poor-quality data during generative AI model training can increase costs.	●	○
AI models rely on training data to learn patterns and identify relationships to produce outputs.	●	○
Generative AI models trained on non-representative datasets can produce inaccurate or unbalanced results.	●	○

Explanation:
Box 1: Yes
Yes - Using incomplete or poor-quality data during generative AI model training can increase costs.
Using incomplete or poor-quality data during generative AI (GenAI) model training significantly increases costs, acting as a major cause of project failure and inefficiencies. This phenomenon is driven by the "garbage in, garbage out" principle, where flawed inputs lead to, at best, unreliable outputs and, at worst, extensive, costly, and time-consuming remediation.
Box 2: Yes
Yes - AI models rely on training data to learn patterns and identify relationships to produce outputs.
At their core, AI models function like pattern-recognition engines; they don't "know" things in the human sense, but rather calculate the statistical likelihood of what should come next based on the data they've processed.
The quality and variety of that training data directly dictate how nuanced and accurate those relationships become. This is why we see such a massive leap between models trained on small, specific datasets versus Large Language Models (LLMs) trained on the vast diversity of the internet.
Box 3: Yes
Yes - Generative AI models trained on non-representative datasets can produce inaccurate or unbalanced results.
When generative AI models are trained on non-representative datasets, they often inherit and amplify existing societal prejudices, leading to systematic distortions known as representation bias. These models fail to generalize fairly across broader populations, resulting in outputs that marginalize or inaccurately depict minority groups.

QUESTION 21
Your company uses a fine-tuned generative AI solution trained on data that is representative of the general population. You discover that some of the generated responses include inappropriate or exclusionary language based on ableist assumptions.
You need to prevent the inappropriate responses. Your solution must minimize costs.
What should you do?

A. Apply a newer version of the generative AI model.
B. Apply a content-moderation filter.
C. Create a new version of the solution that is trained on only inclusive and representative content.
D. Create a new version of the solution that is trained on only exclusionary content.

Answer: B
Explanation:
Implementing a content moderation filter is a critical strategy for fine-tuned generative AI to prevent the output of inappropriate or exclusionary language, including content rooted in ableist assumptions. These filters serve as a "digital safety layer" that intercepts biased or harmful material before it reaches the user.

How Filters Address Ableism and Exclusion
Contextual Detection: Modern filters using Large Language Models (LLMs) and Natural Language Processing (NLP) can detect subtle discriminatory phrases and slurs that traditional keyword filters might miss.
Policy-Driven Guardrails: Developers can use tools like the Lakera Guard Content Safety template to apply specific policies that flag and block ableist speech in real-time.
Customizable Classifiers: Services such as Azure AI Content Safety allow for the detection of "abusive, derogatory, or discriminatory language" through trained statistical models.
Personalized Moderation: Emerging tools are being designed specifically to help disabled users filter out ableist hate and harassment tailored to their unique experiences.
Reference:
https://www.lakera.ai/blog/content-moderation

QUESTION 22
An organization is exploring artificial intelligence solutions to automate content creation tasks such as drafting emails, generating marketing visuals, and producing software code suggestions. Leadership wants to understand the core technology capability that enables these use cases.
Which of the following best describes generative AI?

A. A system that only classifies existing data into predefined categories
B. A technology that creates new content such as text, images, or code based on learned patterns
C. A database system designed to store unstructured data
D. A rules-based automation tool that follows fixed instructions

Answer: B
Explanation:
A technology that creates new content such as text, images, or code based on learned patterns is correct because generative AI systems learn from large datasets and produce original outputs such as written content, visuals, audio, video, or code using models like large language models and diffusion models.
Reference:
https://learn.microsoft.com/en-us/training/modules/understand-foundations-generative-ai-business-leaders/1-introduction

QUESTION 23
A marketing team wants to automatically create product descriptions and campaign email drafts. Which generative AI capability best meets this business need?

A. Image classification
B. Natural language content generation
C. Predictive demand forecasting

D. Anomaly detection

Answer: B
Explanation:
Natural language content generation is correct because Natural language content generation enables the creation of product descriptions, emails, blogs, and other written materials using prompts. This directly supports marketing automation and content scaling.
References:
https://learn.microsoft.com/en-us/training/modules/understand-foundations-generative-ai-business-leaders/3-explore-business-value-generative-ai-solutions
https://www.microsoft.com/en-us/power-platform/blog/power-automate/generative-ai-prompts-to-automate-content-processing/

QUESTION 24
Which benefit of generative AI enables organizations to accelerate content creation across departments such as marketing, HR, and communications?

A. Generating drafts of business content such as emails, reports, and job descriptions
B. Monitoring real-time network intrusion attempts
C. Managing enterprise device firmware updates
D. Automating warehouse robotics operations

Answer: A
Explanation:
Generative AI can produce first drafts of emails, reports, job descriptions, and other business documents. This accelerates content creation and reduces manual effort across organizational functions.
Reference:
https://learn.microsoft.com/en-us/training/modules/build-effective-generative-ai-solutions-organization/1-introduction

QUESTION 25
An organization wants to enhance employee productivity by using generative AI within tools such as Word, Excel, PowerPoint, Outlook, and Teams.
The solution must assist users by generating content, summarizing meetings, analyzing data, and drafting communications within their daily workflow.
Which solution should the organization implement?

A. Azure Machine Learning
B. Microsoft 365 Copilot
C. Azure AI Vision
D. Microsoft Defender for Cloud

Answer: B
Explanation:
Microsoft 365 Copilot embeds generative AI capabilities directly into Microsoft 365 applications, enabling users to generate content, summarize meetings, analyze spreadsheets, and draft communications within their daily productivity tools such as Word, Excel, PowerPoint, Outlook, and Teams.
References:
https://learn.microsoft.com/en-us/training/modules/business-value-microsoft-copilot-solutions/1-introduction?ns-enrollment-type=learningpath&ns-enrollment-id=learn.wwl.drive-value-generative-ai-solutions
https://www.microsoft.com/en-in/microsoft-365-copilot/in-apps-for-work#tabs-pill-bar-oc58d2_tab1

QUESTION 26
Your company plans to build a generative AI solution based on internal data.
You recommend using Microsoft Foundry as a starting point to develop and manage the solution.
What is a key benefit of using Microsoft Foundry for this project?

A. Provides a scalable platform for developing and deploying generative AI solutions.
B. Removes the need to select or configure the underlying AI model.

C. Enables business users to build generative AI solutions.
D. Offers a low-code platform for developing generative AI solutions.

Answer: C
Explanation:
Microsoft Foundry is an enterprise-grade platform specifically designed to help teams build, deploy, and manage generative AI solutions grounded in their own internal data.

While it is a powerful tool for this purpose, its target audience and complexity are important to distinguish:
*-> Building on Internal Data: The platform excels at this through Foundry IQ and Retrieval-Augmented Generation (RAG). It allows you to securely connect AI models to internal "knowledge bases"—such as SharePoint, OneLake, or custom databases—so the AI provides responses based specifically on your company's context and data.
Target User: Contrary to being a tool solely for general business users, it is primarily an interoperable platform for developers, data scientists, and IT professionals. It provides deep technical tools like SDKs, CLI, and MLOps pipelines for scaling AI from a prototype to a full production application.
*-> Accessibility for Business Users: While its primary focus is developers, it does include low-code/no-code interfaces and visual "playgrounds". These allow non-technical contributors to experiment with models, test prompts, and participate in the development process without deep coding knowledge.
Reference:
https://www.softwebsolutions.com/resources/what-is-azure-ai-foundry

QUESTION 27
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer:

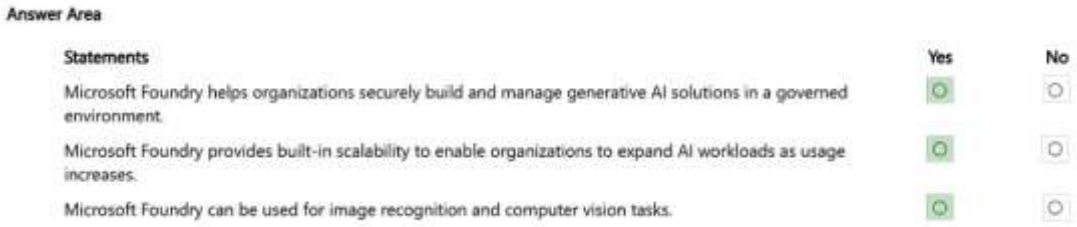

Explanation:
Box 1: Yes
Yes - Microsoft Foundry helps organizations securely build and manage generative AI solutions governed environment.
Microsoft Foundry is a unified, interoperable platform designed to help organizations build, optimize, and manage generative AI applications and autonomous agents within a secure, governed environment. It acts as a central "AI app and agent factory" that brings together models, data, and tools, allowing businesses to move from prototyping to production while maintaining safety and compliance.
Box 2: Yes
Yes - Microsoft Foundry provided built-in scalability to enable organizations to expand AI workloads as usage increases.
Microsoft Foundry acts as an enterprise-grade, unified platform for AI app and agent development, designed to enable organizations to build, deploy, and scale AI workloads efficiently. It provides built-in, automated scalability through several key mechanisms that allow organizations to expand their AI usage without manual infrastructure management.
Box 3: Yes
Yes - Microsoft Foundry can be used for image recognition and computer vision tasks.
Microsoft Foundry (part of Azure AI Services/Tools) offers Azure Vision, a comprehensive suite for image recognition and computer vision tasks. It provides prebuilt APIs and tools for analyzing images, detecting objects, OCR, and facial

recognition, allowing developers to build intelligent, agentic applications without deep machine learning expertise.
Reference:
https://learn.microsoft.com/en-us/azure/ai-foundry/what-is-foundry
https://azure.microsoft.com/en-us/products/ai-foundry

QUESTION 28
Hotspot Question
Select the answer that correctly completes the sentence.

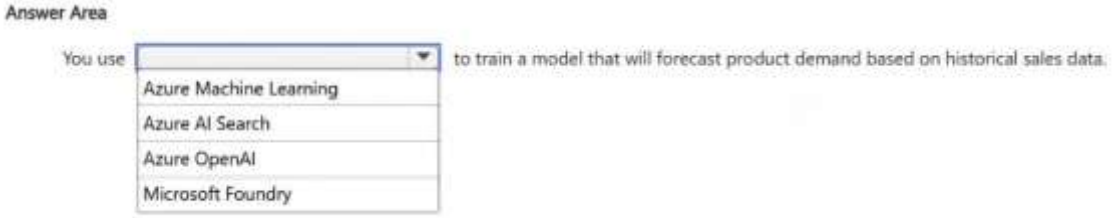

Answer:

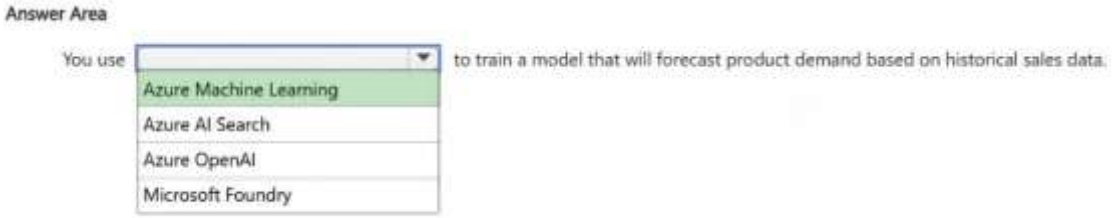

Explanation:
Box: Azure Machine Learning
You use _____ to train a model that will forecast product demand based on historical sales data.
Using Azure Machine Learning to forecast product demand based on historical sales data is best accomplished using Automated Machine Learning (AutoML) for Time-Series Forecasting. This approach allows you to train, evaluate, and deploy a high-quality model, often without writing extensive code, by automatically testing various algorithms and preprocessing data.
Reference:
https://learn.microsoft.com/en-us/azure/machine-learning/concept-automl-forecasting-methods

QUESTION 29
Which business requirement most closely relates to grounding a generative AI model?

A. supporting multiple languages
B. measuring the number of user interactions per day
C. enabling users to interact by using natural language queries
D. ensuring that verified company data sources are used for response generation

Answer: D
Explanation:
Ensuring that verified company data sources are used for response generation relates to grounding a generative AI model by anchoring its outputs in trusted, domain-specific, or enterprise-specific information. This process bridges the gap between the general knowledge a model has from its training data and the specific, up-to-date facts required for accurate, trustworthy business applications.
Reference:
https://decagon.ai/glossary/what-is-ai-grounding

QUESTION 30
You need to create a custom Azure Machine Learning model. The data used to train the model is consistent and uniform.
What should you do first?

A. Prepare the training data.
B. Evaluate the model.

C. Train the model.
D. Tune hyperparameters.
E. Deploy the model.

Answer: A
Explanation:
The first step in creating a custom Azure Machine Learning model trained on your data is to acquire and prepare the data. This involves activities such as:

Data Collection: Gathering the relevant data from its sources, such as databases, streaming sources, or Azure Blob storage.
Data Cleaning and Preprocessing: Even with consistent and uniform data, you will need to perform steps like handling missing values, removing duplicates, and ensuring standardization.
Data Transformation and Feature Engineering: Converting the raw data into a format suitable for the chosen machine learning algorithm and creating new features that can improve model performance.
Data Splitting: Dividing the dataset into separate training, validation, and testing sets so the model can be trained on one portion and evaluated on data it hasn't seen before.

Note:
Once the data is prepared and ready, the subsequent steps in Azure Machine Learning typically involve:
1. Setting up an Azure Machine Learning workspace if you don't already have one.
2. Creating a data asset within the workspace that points to your data in Azure storage.
3. Configuring compute resources for training the model.
4, Selecting an appropriate model algorithm and writing a training script (or using automated ML features).
5. Training and tuning the model using the prepared data and compute resources
Reference:
https://medium.com/@offpagework1.datatrained/building-custom-r-models-in-azure-machine-learning-is-easy-e548598c6325

QUESTION 31
Your company uses a non-reasoning generative AI model to create textual content.
You discover that the model's responses are inconsistent and do NOT meet expectations.
You need to improve the prompts.
What should you do? More than one answer choice may achieve the goal. Select the BEST answer.

A. Provide the prompts with extensive examples of the expected output.
B. Add the context, sources, and expectations to the prompts.
C. Use technical terms in the prompts to enhance AI comprehension.
D. Add only a single concise requirement to the prompts.

Answer: B
Explanation:
When a non-reasoning model produces inconsistent results, you can ground its output by transforming a vague request into a highly structured framework. Since these models rely on pattern prediction rather than true logical deduction, providing "missing" data directly in the prompt acts as a roadmap for the desired completion.

To move from inconsistent to reliable content, focus on these specific additions:
*-> Contextual Guardrails: Provide situational details, such as the intended audience (e.g., "tech-savvy software developers" vs. "elementary students") and domain-specific constraints. This narrows the model's focus to relevant training data patterns.
*-> External Sources & Grounding: Include specific facts, background documents, or source material within the prompt to prevent the model from guessing or "hallucinating" facts.
*-> Explicit Expectations: Clearly define the format (e.g., JSON, Markdown, bulleted list) and tone (e.g., professional, witty). Stating what "success" looks like—such as word count limits or mandatory sections—reduces ambiguity.
Few-Shot Prompting: Add 1–3 examples of the exact style and structure you want the model to mimic. This is often the most effective way to align a non-reasoning model's output with your expectations.
Persona Assignment: Instruct the model to "act as" a specific professional (e.g., "Senior Copy Editor" or "Skeptical Venture Capitalist") to influence the vocabulary and perspective of the generated text.
Reference:

https://cloud.google.com/discover/what-is-prompt-engineering

QUESTION 32
Hotspot Question
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

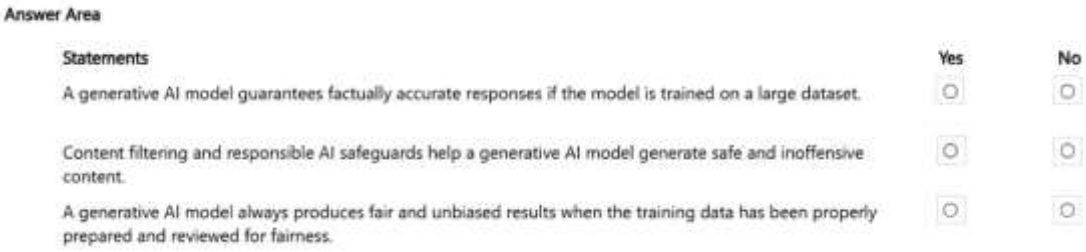

Answer:

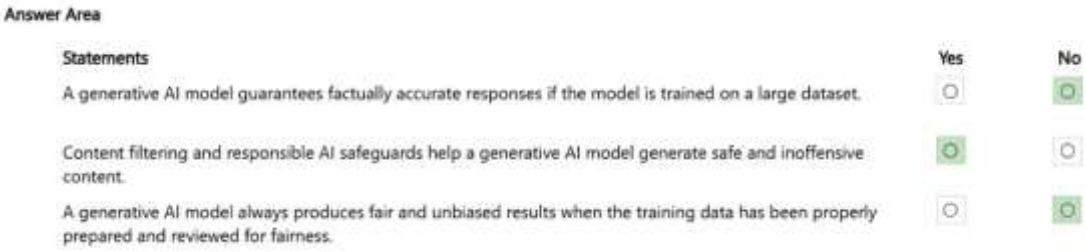

Explanation:
Box 1: No
No - A generative AI model guarantees factually accurate responses if the model is trained on a large dataset.
A large training dataset does not guarantee that a generative AI model will provide factually accurate responses. While larger, diverse datasets generally improve performance and reduce certain types of errors, they do not eliminate the fundamental tendency of these models to generate incorrect information, known as "hallucinations".
Box 2: Yes
Yes - Content filtering and responsible AI safeguards help a generative AI model generate safe an inoffensive content.
Content filtering and responsible AI safeguards (e.g., in Azure AI Foundry or Amazon Bedrock) act as essential, multi-layered, reactive mechanisms—covering both input and output—to detect and block harmful, illegal, or biased content. These systems use automated classifiers to, for example, filter for hate speech, sexual content, violence, and self-harm. They ensure safety by analyzing prompts and generating responses, often allowing for custom thresholds, to prevent models from generating unsafe or inappropriate output.
Box 3: No
No - A generative AI model always produce fair and unbiased results when the training data has been properly prepared and reviewed for fairness.
Even with perfectly prepared and reviewed training data, generative AI models can still produce biased results. While high-quality data is foundational, bias is a persistent challenge that can emerge from multiple sources throughout the AI lifecycle.
Reference:
https://mehmetozkaya.medium.com/limitations-of-large-language-models-llms-1790a14010db
https://monowar-mukul.medium.com/keeping-your-ai-safe-content-filters-in-azure-ai-foundry-9a87c8447e11
https://www.sap.com/resources/what-is-ai-bias