

DUMPS ARENA

AI Transformation Leader

Microsoft AB-731

Version Demo

Total Demo Questions: 10

Total Premium Questions: 100

Buy Premium PDF

<https://dumpsarena.co>

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

What is a key feature of Microsoft 365 Copilot that aligns with the Microsoft responsible AI principles of transparency, reliability, and safety?

- A. Automatically approves AI-generated content for company-wide publishing.
- B. Removes the need for a human to review AI-generated outputs.
- C. Provides grounded, verifiable responses based on organizational data.
- D. Enables users to select from an authorized catalog of AI models.

ANSWER: C**Explanation:**

The best answer is **C** because Microsoft 365 Copilot is designed to be *grounded* in your organization's Microsoft 365 data (for example, files, emails, chats, and calendars) and can provide responses that are tied back to enterprise content. This grounding approach supports **transparency** by helping users understand what information Copilot used, and it supports **reliability and safety** by reducing "hallucinations" and keeping outputs aligned to trusted organizational sources and permissions. In practice, this is a key mechanism Microsoft uses to operationalize responsible AI principles in Copilot experiences.

Option **A** is incorrect because automatically approving AI-generated content for broad publishing removes essential human oversight and governance controls, which conflicts with safe deployment practices. Option **B** is incorrect because Microsoft's guidance emphasizes that AI outputs should be reviewed as appropriate; removing human review would undermine accountability and risk management. Option **D** is incorrect because an "authorized model catalog" is not the core Copilot feature that directly maps to transparency/reliability/safety in Microsoft 365; Copilot's differentiator is grounding in Microsoft Graph and tenant data with security trimming, not end-user model selection.

References: [Microsoft 365 Copilot overview](#), [Responsible AI for Azure AI services](#).

QUESTION NO: 2

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: B C

Explanation:

The best governance principles for responsible AI start with clear accountability and a repeatable oversight process. **B** is correct because defining accountability norms across both business and technical teams ensures there are named owners for risk decisions, approvals, monitoring, and remediation. This is essential for auditability and for aligning model behavior with organizational values and regulatory obligations. **C** is also correct because establishing a formal process to review AI initiatives (for example, an intake/risk assessment, governance board review, and periodic re-evaluation) operationalizes responsible AI and helps catch issues like fairness, transparency, privacy, and security before deployment and during lifecycle changes.

A is not best because letting each department tailor governance typically creates inconsistent standards and gaps in enterprise compliance; governance should be centrally defined even if execution is federated. **D** is incorrect because governance cannot sit primarily with engineering/data science—legal, compliance, risk, and business stakeholders must be involved to ensure policy alignment and accountability. **E** is insufficient because responsible AI governance should apply broadly; even non-regulated use cases can create significant reputational or customer harm and still require oversight.

References: [Azure Machine Learning responsible AI](#), [Azure Machine Learning governance](#).

QUESTION NO: 3

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

Microsoft Copilot Studio user license – Microsoft does not offer a standalone “Copilot Studio user” license; Copilot Studio is a development environment for building agents, but execution of those agents still requires a Copilot add-on license tied to a Microsoft365 user account. Without the add-on, agents cannot access protected work data.

C. a Copilot Chat pay-as-you-go plan

Copilot Chat pay-as-you-go plan – This plan provides only conversational access to Copilot Chat and is intended for ad-hoc queries. It does not include the rights to develop or host declarative agents that interact with internal datasets, nor does it grant the required permissions to read/write work data.

Conclusion

To enable users to create and operate declarative agents that leverage work-related data, the appropriate entitlement is the Microsoft365 Copilot add-on license (OptionA).

References

Microsoft 365 Copilot licensing guide: <https://learn.microsoft.com/en-us/microsoft-365/copilot/licensing> [Build declarative agents with Copilot Studio: https://learn.microsoft.com/en-us/copilot-studio/overview-agents](https://learn.microsoft.com/en-us/copilot-studio/overview-agents)

ANSWER: A

Explanation:

The right recommendation is the **Microsoft 365 Copilot add-on license**. Declarative agents that can ground responses in (and act on) Microsoft 365 work data rely on Microsoft 365 Copilot capabilities such as Microsoft Graph grounding and the Microsoft 365 security/compliance boundary. In practice, the entitlement that enables a user to use Copilot with organizational data across apps like Teams, Outlook, SharePoint, and OneDrive is the Microsoft 365 Copilot add-on license.

Option B is incorrect because Copilot Studio licensing is not simply a generic “user license” that, by itself, grants Microsoft 365 Copilot work-data access in Copilot Chat. Copilot Studio is the tool to build agents, but access to Microsoft 365 work data in Copilot experiences is governed by Microsoft 365 Copilot licensing/entitlements and the user’s permissions.

Option C is incorrect because pay-as-you-go for Copilot Chat is aimed at metered usage and doesn’t inherently grant the same Microsoft 365 Copilot work-data grounding rights needed for users to build/use declarative agents over organizational content.

References: [Microsoft 365 Copilot licensing](#), [Microsoft Copilot Studio documentation](#)

QUESTION NO: 4

Your company is reviewing a new AI solution before deploying it. The company wants to ensure that the solution follows Microsoft responsible AI principles.

What is the best approach to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Design the AI solution to automatically approve or reject customer loan applications.
- B. Test the AI solution to identify and mitigate potential unfair or inconsistent outcomes in its outputs.
- C. Prioritize model performance when tuning the AI solution.
- D. Enable the AI solution to collect and store personal data.

ANSWER: B

Explanation:

The best approach is to **test the AI solution to identify and mitigate potential unfair or inconsistent outcomes**. Microsoft’s Responsible AI principles emphasize fairness, reliability & safety, transparency, accountability, and privacy & security. Pre-deployment evaluation is where you validate these principles in practice—by running structured assessments (for example, fairness/error analysis across groups, robustness testing, and reviewing model behavior for harmful or inconsistent outputs) and then applying mitigations (data improvements, model adjustments, human-in-the-loop controls, or policy/UX changes). This directly aligns with Microsoft guidance to assess and manage AI risks throughout the lifecycle, especially before release.

Option A is not the best approach because fully automating high-impact decisions (like loan approvals) without explicit governance and human oversight can undermine accountability and increase the risk of unfair outcomes. Option C is insufficient because optimizing performance metrics alone (accuracy, AUC, etc.) does not ensure fairness, transparency, or safety. Option D is incorrect because collecting/storing personal data is not a responsible-AI strategy by itself and can increase privacy and compliance risk unless there is a clear purpose, minimization, and governance.

References: [Responsible AI overview \(Microsoft Learn\)](#), [Responsible AI in Azure AI \(Microsoft Learn\)](#)

QUESTION NO: 5

Your company purchases Microsoft 365 Copilot for its sales department.

The sales department needs to find and summarize information across internal documents quickly.

From which two data sources can the sales department obtain results by default? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Microsoft OneDrive
- B. an on-premises file share
- C. Microsoft SharePoint
- D. Microsoft Sway
- E. a custom customer relationship management (CRM) system

ANSWER: A C**Explanation:**

By default, Microsoft 365 Copilot grounds its responses in Microsoft Graph content that's already part of the Microsoft 365 ecosystem and that the user has permission to access. For "internal documents," the two primary out-of-the-box repositories are OneDrive for Business and SharePoint Online. OneDrive (A) stores a user's work files and shared files, and Copilot can summarize and reference those documents without any extra setup beyond normal permissions. SharePoint (C) is the organization's document management platform (including files stored in Teams-backed SharePoint sites), and Copilot can search and summarize content across SharePoint sites the user can access.

The other options aren't available "by default." An on-premises file share (B) isn't natively indexed by Microsoft 365 Copilot; bringing that content in typically requires migration to SharePoint/OneDrive or use of additional integration approaches. Microsoft Sway (D) isn't listed as a standard supported grounding source for Copilot in Microsoft 365. A custom CRM system (E) also isn't automatically included; it generally requires custom integration (for example, via Microsoft Graph connectors or a purpose-built plugin/connector) and appropriate governance.

References: [Microsoft 365 Copilot data, privacy, and security](#), [Microsoft 365 Copilot overview](#).

QUESTION NO: 6

Your company is evaluating the use of Microsoft Copilot Studio to support business process automation and employee self-service.

Which two capabilities are directly supported in Copilot Studio? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. drafting and summarizing files in Microsoft Word and PowerPoint
- B. customizing agent behavior and responses
- C. using agents to identify and respond to security incidents
- D. building agents that connect to business data and automate user interactions
- E. configuring document security

ANSWER: B D

Explanation:

Copilot Studio is designed to let you build and customize copilots (agents) that can converse with users, call actions, and connect to enterprise data through connectors, Power Platform components, and APIs. Therefore, building agents that connect to business data and automate user interactions is a core, directly supported capability (D). Copilot Studio also explicitly supports customizing how an agent behaves and responds via topics, generative answers, instructions/system prompts, orchestration, and configuration of actions—this is a first-class part of authoring a copilot, not merely an implied sub-skill—so (B) is also directly supported.

Drafting/summarizing in Word and PowerPoint (A) is primarily a Microsoft 365 Copilot app experience rather than a Copilot Studio capability. Configuring document security (E) is handled by Microsoft Purview/AIP and related compliance tooling, not Copilot Studio. “Using agents to identify and respond to security incidents” (C) could be built as a custom solution by integrating with Microsoft Sentinel/Defender APIs and Power Automate, but it isn’t a direct, out-of-the-box Copilot Studio capability category in the way agent customization and business-data automation are.

References: [Microsoft Copilot Studio overview](#), [Configure instructions/behavior in Copilot Studio](#)

QUESTION NO: 7

You plan to meet with a group of stakeholders to discuss how generative AI can benefit your company. You need to provide the stakeholders with a relevant description of generative AI during the meeting. Which description should you use?

- A. Generative AI is designed to generate responses based on a user’s natural language prompts.
Generative AI is designed to generate responses based on a user’s natural language prompts.

Generative AI is a type of artificial intelligence that creates new content—such as text, images, or audio—by identifying patterns in existing data. The core mechanism through which users interact with these models is via natural language prompts, which provide the instructions or context necessary for the AI to generate a relevant response. This makes it a versatile tool for general business applications like content creation, summarization, and interactive assistance. Why other options are incorrect:

- B. Generative AI is designed to translate documents into other languages.
Generative AI is designed to translate documents into other languages. While generative AI can perform translation, this is a specific application (often handled by specialized Machine Translation models) rather than its defining purpose.

C. Generative AI is designed to predict future trends based on historical data.

Generative AI is designed to predict future trends based on historical data. This description typically refers to Predictive AI or predictive analytics, which focuses on forecasting outcomes rather than generating new, original content.

D. Generative AI is designed to recommend products based on user behavior.

Generative AI is designed to recommend products based on user behavior. This describes Recommendation Systems, a branch of traditional machine learning that uses historical user data to filter and suggest existing items, rather than creating new ones.

ANSWER: A

Explanation:

The best description is **A**. Generative AI refers to AI models that can *create new content* (for example, text, images, code, or audio) in response to an input. In business settings, the most common interaction pattern is that users provide a prompt (often natural language) and the model generates a response or content based on what it learned from training data. That makes option A the most accurate and broadly applicable description for stakeholders.

Option **B** is too narrow: translation is one task generative models can do, but it's not the defining purpose of generative AI. Option **C** describes predictive analytics/forecasting, which focuses on estimating future outcomes from historical data rather than generating novel content. Option **D** describes recommendation systems, which typically rank or suggest existing items based on behavior patterns; that's a different ML workload than generative content creation.

For Microsoft-aligned definitions, see the overview of generative AI and how prompts drive outputs in Azure AI and Microsoft Copilot concepts: <https://learn.microsoft.com/en-us/azure/ai-services/openai/overview> and <https://learn.microsoft.com/en-us/training/modules/fundamentals-generative-ai/>.

QUESTION NO: 8 - (HOTSPOT)

HOTSPOT -

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

A barrier to AI adoption can include data limitations and a lack of AI readiness and skills.

Yes

No

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

A barrier to AI adoption can include data limitations and a lack of AI readiness and skills.

Yes

No

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:

All three statements describe well-known, real-world blockers that show up repeatedly in AI transformation programs, so each one should be answered **Yes**. First, AI initiatives are heavily dependent on data. If an organization has limited data, poor data quality, siloed systems, or weak data governance, it becomes difficult to train models, evaluate results, and operationalize solutions. On top of that, many companies lack AI readiness in the form of skills (data engineering, ML/AI engineering, product ownership, and responsible AI knowledge), which is another common adoption barrier.

Second, organizations often struggle when they take a “technology-first” approach—picking an AI tool, model, or platform before they’ve clearly defined the business problem, success metrics, and the specific use case. This usually leads to pilots that don’t scale or solutions that don’t deliver measurable value. Microsoft’s guidance for AI strategy emphasizes starting from business outcomes and use cases, then selecting the right technology to support them.

Third, cross-functional collaboration is essential because AI solutions touch many groups: business stakeholders define value and requirements; data teams provide pipelines and quality controls; IT and platform teams handle deployment and operations; and security, privacy, and compliance teams ensure governance and responsible use. When these groups don’t work together, projects stall due to misaligned priorities, unclear ownership, and delays in data access or approvals. For more background on common AI adoption challenges and the importance of strategy, data, and organizational alignment, see Microsoft’s AI strategy and adoption resources such as [Cloud Adoption Framework: AI](#) and the overview of [Azure AI/ML architecture guidance](#).

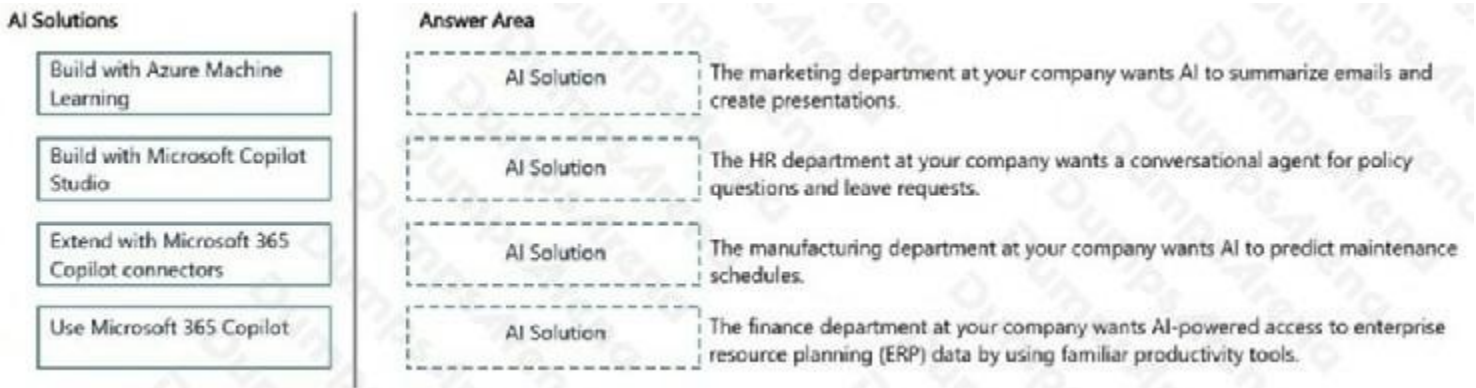
QUESTION NO: 9 - (DRAG DROP)

DRAG DROP -

Match the business scenario to the appropriate AI solution design approach.

To answer, drag the AI solution from the column on the left to its business scenario on the right. Each solution may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.



ANSWER:



Explanation:

This is a matching question about choosing the right Microsoft AI approach based on what the business is trying to accomplish. When the requirement is “do AI inside the Microsoft 365 apps people already use” (like Outlook and PowerPoint), the most direct solution is **Microsoft 365 Copilot**, because it’s designed to generate and summarize content across those apps. That’s why the marketing scenario (summarize emails + create presentations) maps to **Use Microsoft 365 Copilot**.

When the requirement is a **custom conversational agent** for a specific department (HR) that answers policy questions and can be extended with actions/workflows (like leave requests), the right design approach is **Build with Microsoft Copilot Studio**. Copilot Studio is intended for creating and managing custom copilots/chatbots connected to enterprise data and processes. See: [What is Microsoft Copilot Studio](#).

Predictive maintenance in manufacturing is a classic **machine learning prediction/forecasting** use case that typically uses historical sensor/telemetry data to predict failures or maintenance windows. That aligns to **Build with Azure Machine Learning**, which provides tooling for training, deploying, and managing ML models. Reference: [What is Azure Machine Learning](#).

Finally, finance wants AI-powered access to **ERP data** using familiar productivity tools. That implies bringing external line-of-business data into the Microsoft 365 Copilot experience. The appropriate approach is **Extend with Microsoft 365 Copilot connectors**, which are designed to connect and surface third-party or external data sources (like ERP systems) for Microsoft 365 search and Copilot grounding. Reference: [Microsoft Graph connectors overview](#).

QUESTION NO: 10 - (HOTSPOT)

HOTSPOT -

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

A generative AI model guarantees factually accurate responses if the model is trained on a large dataset.

Yes

No

Content filtering and responsible AI safeguards help a generative AI model generate safe and inoffensive content.

A generative AI model always produces fair and unbiased results when the training data has been properly prepared and reviewed for fairness.

ANSWER:

Answer Area

Statements

A generative AI model guarantees factually accurate responses if the model is trained on a large dataset.

Yes

No

Content filtering and responsible AI safeguards help a generative AI model generate safe and inoffensive content.

A generative AI model always produces fair and unbiased results when the training data has been properly prepared and reviewed for fairness.

Explanation:

You're being asked to judge three broad "always/guarantee" statements about generative AI. In Microsoft's Responsible AI guidance, those absolute claims are almost always false because generative models are probabilistic and can produce unexpected outputs.

Statement 1 (factual accuracy guaranteed with a large dataset): No. Even with huge training corpora, a generative model can still produce incorrect or fabricated information ("hallucinations"). The model is predicting likely text, not verifying truth against a source of record. Large datasets can improve general performance, but they don't create a factual guarantee. This aligns with Microsoft's guidance that generative AI outputs must be validated and that accuracy is a key risk area to manage rather than assume. See: [Azure OpenAI Service - Responsible AI overview](#).

Statement 2 (content filtering and safeguards help generate safe/inoffensive content): Yes. Content filters, moderation systems, prompt shields, and other responsible AI mitigations are specifically designed to reduce harmful, unsafe, or offensive outputs. They don't make a system perfect, but they do help significantly by detecting and blocking certain categories of content and misuse patterns. Microsoft documents these kinds of safety mechanisms for Azure OpenAI. See: [Azure OpenAI content filtering](#).

Statement 3 (always fair and unbiased after fairness review): No. Reviewing and preparing training data for fairness is important, but it cannot ensure the model will *always* be fair and unbiased. Bias can still appear due to subtle correlations in data, model generalization, prompt/context effects, and downstream usage. Microsoft frames fairness as something you assess and mitigate continuously, not something you can guarantee once. See: [Microsoft Responsible AI - Fairness](#).

So the correct pattern is **No, Yes, No** for statements 1–3.