

AI Transformation Leader

Microsoft AB-731

Version Demo

Total Demo Questions: 20

Total Premium Questions: 202

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Plan and manage an AI transformation	125
Topic 2, Implement responsible AI	55
Topic 3, Drive AI adoption and organizational change	22
Total	202

QUESTION NO: 1

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:

Define accountability norms for AI decisions across business and technical teams. Responsible AI governance requires clear ownership throughout the AI lifecycle. Business leaders, technical teams, risk personnel, legal stakeholders, and compliance functions must understand who is accountable for decisions such as approving an AI use case, accepting risk, validating controls, monitoring outcomes, and responding to incidents. Clear accountability makes governance operational rather than aspirational, supports traceability, and helps ensure that organizational values and compliance obligations are consistently applied.

Create a process to review AI initiatives for responsible AI alignment. A repeatable review process embeds responsible AI requirements into planning, development, deployment, and ongoing operation. The process can assess potential impacts and controls for fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability. It should be proportionate to the use case's potential impact and include periodic reassessment when a system, its data, or its operating context changes. Microsoft's responsible AI guidance emphasizes governance structures, clear accountability, and impact assessment as key mechanisms for managing AI risks. Microsoft also describes governing machine learning through policies and controls across the lifecycle. See [Microsoft responsible AI guidance](#) and [Responsible AI in Azure Machine Learning](#).

QUESTION NO: 2

Your company uses generative AI to assist with content creation and customer interactions.

You need to evaluate whether Azure Machine Learning can add value to the current customer management. For which use case should you use Machine Learning?

- A. predicting customer retention
- B. creating product descriptions from images
- C. generating marketing campaigns
- D. summarizing customer service transcripts

ANSWER: A

Explanation:

Azure Machine Learning is well suited to **predicting customer retention** because this is a predictive analytics problem based on historical business data. An organization can train a supervised machine learning model using customer attributes and behavior, such as tenure, purchasing history, product usage, service interactions, subscription status, and previous

retention outcomes. The trained model can then estimate the likelihood that an existing customer will remain with the organization or, equivalently, identify customers at risk of churn.

This prediction gives customer-management teams an actionable score that can be incorporated into CRM processes, dashboards, and retention workflows. For example, high-risk customers can be prioritized for proactive outreach, tailored offers, or additional support. Azure Machine Learning provides managed capabilities for preparing data, training and comparing models, tracking experiments, registering the selected model, and deploying it to an endpoint or batch scoring process. Automated ML can also accelerate development by testing suitable classification algorithms and tuning settings for a labeled retention dataset. Microsoft documents automated machine learning support for classification, regression, and forecasting scenarios in [Automated ML in Azure Machine Learning](#). The broader role of Azure Machine Learning in building, training, deploying, and managing machine learning models is described in the [Azure Machine Learning overview](#).

QUESTION NO: 3 - (SIMULATION)

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 text-to-image generator can be used to translate content into other languages.	☐	☐
A predictive analytics model can generate new marketing content for a company's online ads.	☐	☐
A generative AI chatbot can engage customers in personalized conversations and recommend products.	☐	☐

ANSWER: See the explanation for the answer

Explanation:

Selections:

Text-to-image tools generate images from prompts rather than perform language translation. Predictive analytics forecasts or classifies outcomes from data; it does not create new ad copy. Generative AI chatbots can conduct personalized natural-language interactions and provide product recommendations when connected to relevant product data and business rules.

QUESTION NO: 4

Your company is developing a custom agent that can submit purchase requests to an internal procurement system.

You need to apply appropriate controls before the agent can perform transactions.

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

NOTE: Each correct selection is worth one point.

- A. Require user confirmation before submitting a purchase request.
- B. Use least-privilege access for the procurement connection.
- C. Grant the agent administrator permissions in the procurement system.
- D. Allow the agent to submit requests without recording the transaction.
- E. Use a shared employee account for all agent actions.

ANSWER: A B

Explanation:

Require user confirmation before submitting a purchase request. Confirmation provides meaningful human oversight for an action that can create a business commitment. The user can review the request details before the agent invokes the procurement action.

Use least-privilege access for the procurement connection. The agent and its connection should have only the permissions needed for the defined task. Limiting permissions reduces the potential impact of an incorrect request, misuse, or compromised connection.

Action-taking agents require controls that address authorization, user oversight, and the potential consequences of automated actions. See [Actions in Microsoft Copilot Studio](#) and [Azure security design principles](#).

QUESTION NO: 5 - (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.

AI Solutions

Build with Azure Machine Learning

Build with Microsoft Copilot Studio

Extend with Microsoft 365 Copilot connectors

Use Microsoft 365 Copilot

Answer Area

AI Solution

The marketing department at your company wants AI to summarize emails and create presentations.

AI Solution

The HR department at your company wants a conversational agent for policy questions and leave requests.

AI Solution

The manufacturing department at your company wants AI to predict maintenance schedules.

AI Solution

The finance department at your company wants AI-powered access to enterprise resource planning (ERP) data by using familiar productivity tools.

ANSWER:

AI Solutions

Build with Azure Machine Learning

Build with Microsoft Copilot Studio

Extend with Microsoft 365 Copilot connectors

Use Microsoft 365 Copilot

Answer Area

Use Microsoft 365 Copilot

The marketing department at your company wants AI to summarize emails and create presentations.

Build with Microsoft Copilot Studio

The HR department at your company wants a conversational agent for policy questions and leave requests.

Build with Azure Machine Learning

The manufacturing department at your company wants AI to predict maintenance schedules.

Extend with Microsoft 365 Copilot connectors

The finance department at your company wants AI-powered access to enterprise resource planning (ERP) data by using familiar productivity tools.

Explanation:

The mappings are correct because each scenario is matched to the Microsoft AI capability designed for that type of business outcome. Microsoft 365 Copilot is integrated into familiar Microsoft 365 productivity experiences, including Outlook and PowerPoint. It can help users summarize email content and turn information into drafted presentation content, making it the appropriate approach for the marketing department’s productivity-focused request. Microsoft describes these in-app Copilot experiences in its [Copilot in Outlook documentation](#) and [Copilot in PowerPoint documentation](#).

Microsoft Copilot Studio is the appropriate design approach for the HR request because it supports creating custom copilots that can converse with employees, use organizational knowledge, and be connected to actions and business processes. This provides a foundation for answering policy questions and supporting leave-request workflows. The service is intended for building and managing organization-specific conversational experiences, as described in [What is Microsoft Copilot Studio](#).

Azure Machine Learning fits the manufacturing requirement because predictive maintenance depends on developing, training, operationalizing, and monitoring predictive models from historical operational, equipment, and sensor data. Azure Machine Learning supplies the managed machine learning platform capabilities needed for this lifecycle, including model training and deployment. See [What is Azure Machine Learning](#).

Finally, Microsoft 365 Copilot connectors are designed to make enterprise information from external systems available in Microsoft 365 search and Copilot experiences. Connecting ERP content allows finance users to access relevant business data through the productivity tools they already use. Microsoft's [Microsoft Graph connectors overview](#) explains how connectors bring external data into these Microsoft 365 experiences.

QUESTION NO: 6

Your company uses a retrieval-augmented solution to answer questions from a frequently updated set of technical procedures.

You need to ensure that revised procedures are available to users without retraining the underlying language model.

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

NOTE: Each correct selection is worth one point.

- A. Refresh the search index when source procedures change.
- B. Maintain metadata that identifies document versions and effective dates.
- C. Fine-tune the language model after every procedure update.
- D. Disable retrieval for revised procedures.
- E. Retain superseded procedures as the preferred source.

ANSWER: A B

Explanation:

Refresh the search index when source procedures change. A retrieval-augmented solution obtains current context from its retrieval system at the time of a user request. Updating or reindexing changed procedure content makes the revised information available for retrieval without modifying the underlying language model.

Maintain metadata that identifies document versions and effective dates. Metadata enables the retrieval layer to prioritize current procedures and can help exclude superseded documents from results. This reduces the risk that users receive answers based on obsolete guidance.

Azure AI Search supports indexing enterprise content and using metadata as searchable, filterable, and sortable fields. See [Azure AI Search overview](#) and [Retrieval-augmented generation in Azure AI Search](#).

QUESTION NO: 7

Your company plans to deploy Microsoft 365 Copilot to users who work with confidential financial documents.

You need to ensure that documents are protected according to their sensitivity classification when users access them through Microsoft 365 applications.

What should you recommend?

- A. Microsoft Purview sensitivity labels
- B. Microsoft Purview retention labels
- C. Microsoft Defender for Office 365 Safe Links
- D. Microsoft Teams meeting policies

ANSWER: A

Explanation:

Microsoft Purview sensitivity labels are the appropriate recommendation because they enable an organization to classify and protect sensitive content in Microsoft 365. Labels can apply protection settings such as encryption, access restrictions, content markings, and other controls to files and emails.

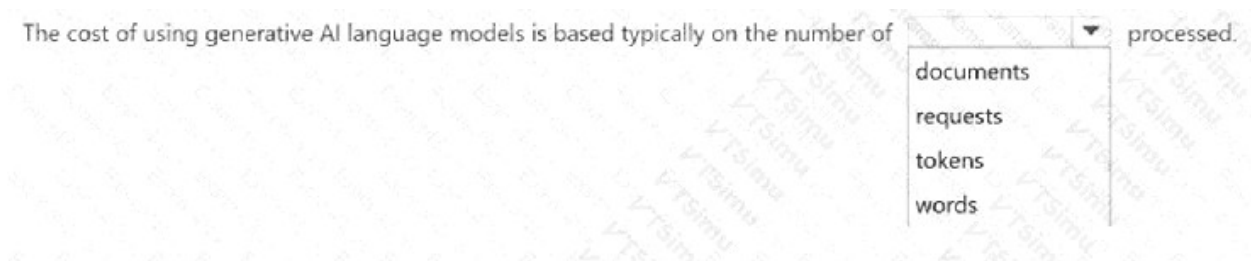
Microsoft 365 Copilot respects the existing permissions and protection applied to organizational data. Applying and managing sensitivity labels helps ensure that confidential financial documents retain the intended protection when they are stored, shared, and accessed through Microsoft 365 experiences. See [Microsoft Purview sensitivity labels](#) and [Microsoft 365 Copilot data, privacy, and security](#).

QUESTION NO: 8 - (SIMULATION)

- Select the answer that correctly completes the sentence.

The cost of using generative AI language models is based typically on the number of _____ processed.

The cost of using generative AI language models is based typically on the number of _____ processed.



ANSWER: See the explanation for the answer

Explanation:

Select **tokens**.

Generative AI language-model usage is typically priced according to the number of input and output **tokens** processed. Tokens are pieces of text (such as words, parts of words, and punctuation), rather than a fixed count of documents, requests, or words.

QUESTION NO: 9

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 priorities.
- D.** Include only data scientists and engineers at first to validate technical feasibility, then add other stakeholders later.

ANSWER: C

Explanation:

Include representatives from legal, leadership, and business units to align AI initiatives with organizational priorities. Enterprise AI adoption is an organizational transformation effort, rather than solely a technology implementation. Leadership sponsorship connects AI investments to strategic objectives, supplies decision-making authority, and helps remove organizational barriers. Business representatives identify meaningful scenarios, define expected outcomes, and ensure that

workflows and users can adopt the resulting solution. Legal and related governance stakeholders bring essential early input on privacy, regulatory requirements, intellectual property, security expectations, and responsible AI practices.

Bringing these perspectives together at the outset creates clear accountability and supports informed tradeoffs between value, feasibility, risk, and operational readiness. It also helps the organization establish appropriate policies, controls, and review processes before an AI capability is scaled into business operations. Microsoft's responsible AI guidance emphasizes governance across the AI lifecycle and identifies organizational accountability as an essential component of managing AI systems responsibly. Microsoft's AI adoption guidance likewise promotes alignment among business, technology, and governance stakeholders to enable sustainable adoption and measurable business value.

For further guidance, see [Microsoft Azure Architecture Center: Responsible AI](#) and [Microsoft Cloud Adoption Framework: AI strategy](#).

QUESTION NO: 10

Your company's HR department spends a large amount time answering repetitive questions and updating policy documents.

You need a solution that can automate employee interactions, respond to HR-related questions, and connect to internal business data sources.

What solution should you recommend?

- A. Microsoft Security Copilot
- B. Microsoft Copilot Studio
- C. Microsoft 365 Copilot
- D. Microsoft Copilot in Azure

ANSWER: B

Explanation:

Microsoft Copilot Studio is designed for creating, managing, and publishing custom copilots and agents for business-specific scenarios such as HR self-service. An HR team can use it to provide conversational answers to common employee questions, reducing the manual effort required to respond repeatedly to topics such as leave, benefits, onboarding, and company policies. Its low-code authoring experience enables teams to define agent behavior and conversation flows while using generative orchestration to handle users' natural-language requests.

Crucially, Copilot Studio can ground responses in organizational knowledge sources, including SharePoint content, Dataverse, websites, and uploaded files, helping ensure that policy-related answers are based on approved internal information. It also supports actions and connectors that let an agent retrieve or update information in business systems and APIs. The agent can then be published through channels employees already use, including Microsoft Teams and websites. This combination of custom conversational automation, knowledge grounding, enterprise data connectivity, and controlled publishing directly meets the HR department's stated requirements. See the [Microsoft Copilot Studio overview](#) and Microsoft guidance on [knowledge sources in Copilot Studio](#).

QUESTION NO: 11

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:

Microsoft 365 Copilot can ground responses in organizational data through Microsoft Graph, while honoring each user's existing permissions. Microsoft OneDrive is a default Microsoft 365 content location for work files. When a salesperson has permission to a file in OneDrive, Copilot can use that file's content to help find information, produce summaries, and create context-aware responses. Sharing and access controls remain in effect, so Copilot does not grant access to content that the user could not otherwise open.

Microsoft SharePoint is also a default source of organizational knowledge for Microsoft 365 Copilot. SharePoint document libraries commonly hold departmental and team documents, and files shared through Teams are typically stored in the associated SharePoint site. Consequently, Copilot can retrieve and summarize relevant SharePoint content that is available to the requesting salesperson. This Microsoft Graph grounding model lets users obtain useful answers from the files and information already present in their Microsoft 365 tenant without separately configuring these standard repositories. Microsoft documents that Copilot uses Microsoft Graph to access organizational content and only surfaces data a user is authorized to access. See [Microsoft 365 Copilot overview](#) and [Microsoft 365 Copilot data, privacy, and security](#).

QUESTION NO: 12

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

- A. Generative AI is designed to generate responses based on a user's natural language prompts.
- B. Generative AI is designed to translate documents into other languages.
- C. Generative AI is designed to predict future trends based on historical data.
- D. Generative AI is designed to recommend products based on user behavior.

ANSWER: A

Explanation:

Generative AI is designed to generate responses based on a user's natural language prompts. This is the most relevant high-level description because generative AI uses models trained to recognize patterns in large volumes of data and then create new content in response to an instruction. That content can include natural-language responses, summaries, drafts, images, code, or other outputs, depending on the model and the prompt. For a stakeholder discussion, this description clearly conveys the core interaction: a person provides an instruction or question, and the AI generates a useful response rather than merely retrieving a fixed answer. Microsoft describes generative AI as a type of AI that can create new content, including text, images, and code. Microsoft's guidance also explains that a prompt is the input used to instruct a generative AI model what to generate. These concepts make natural-language prompting and generated responses an accurate, accessible description of how generative AI can support business work such as drafting, summarizing, ideation, and analysis. See [Microsoft Learn: Fundamentals of generative AI](#) and [Microsoft Learn: Prompt engineering](#).

QUESTION NO: 13

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 ensure descriptions can be posted immediately after creation, enable selection/inclusion of product details, and be fast and simple for non-technical staff. What is the best type of solution to use? 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: C

Explanation:

An interactive AI agent is the best fit because it provides a straightforward, conversational experience for non-technical catalog staff while guiding them through the information needed to create an accurate product description. The agent can prompt staff to select or provide relevant attributes, such as product category, brand, dimensions, materials, key features, intended use, and required wording. It can then generate a structured, publication-ready description using those selected details.

An agent can also be connected to approved business data sources and business processes, allowing it to retrieve catalog information and invoke an action that sends the generated description to the appropriate publishing workflow. This supports the requirement for immediate posting without requiring users to manually move text between systems. Microsoft Copilot Studio is designed to let organizations create agents using low-code tools, connect them to knowledge and actions, and make them available to business users. Its conversational interface keeps the process quick and accessible while preserving staff control over the product information included in each description.

For more information, see [What is Microsoft Copilot Studio?](#) and [Add actions to custom copilots.](#)

QUESTION NO: 14 - (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	Yes	No
Prompt engineering involves specifying guidelines that help a generative AI model generate useful and accurate responses.	☐	☐
The instruction is the part of the prompt that includes examples to show the AI what its response should look like.	☐	☐
The context is the part of the prompt that helps a generative AI model understand which information to use when generating a response.	☐	☐

ANSWER:

Answer Area

Statements	Yes	No
Prompt engineering involves specifying guidelines that help a generative AI model generate useful and accurate responses.	☒	☐
The instruction is the part of the prompt that includes examples to show the AI what its response should look like.	☐	☒
The context is the part of the prompt that helps a generative AI model understand which information to use when generating a response.	☒	☐

Explanation:

Prompt engineering is the practice of deliberately designing and refining the information given to a generative AI model so that it can produce a useful, accurate, relevant, and appropriately formatted result. Guidelines such as the required tone, intended audience, output format, scope, constraints, and evaluation criteria all help make the expected response clear. Therefore, the statement about prompt engineering specifying guidelines that help a model generate useful and accurate responses is true.

A well-structured prompt commonly separates the task instruction from other supporting elements. The instruction states the action the model should take, such as drafting, summarizing, classifying, extracting, or rewriting information. Examples can also be included in a prompt to demonstrate a desired pattern, style, or input-and-output format. This distinction supports selecting No for the statement that defines the instruction itself as the part containing examples.

Context is the information supplied to establish the situation, relevant facts, source material, business background, or constraints that the model should consider while answering. Providing appropriate context helps ground the generated response in the intended information and reduces ambiguity about what the model should use. Consequently, the statement that context helps a generative AI model understand which information to use when generating a response is true.

Microsoft's prompt engineering guidance describes using clear instructions, contextual information, and examples to steer model behavior and improve output quality. See [Prompt engineering techniques in Azure OpenAI Service](#) and [Advanced prompt engineering](#) for further details.

QUESTION NO: 15

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 ensure descriptions can be posted immediately after creation, enable selection/inclusion of product details, and be fast and simple for non-technical staff. What is the best type of solution to use? 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: C

Explanation:

An interactive AI agent is the best fit because it provides a simple, guided experience for non-technical catalog staff while using generative AI to produce descriptions from the details they select. The agent can present structured inputs for product attributes such as brand, dimensions, materials, features, intended use, and required marketing language. It can then generate a description based on those selected values, allowing the user to request adjustments conversationally before using the final text. This interaction helps ensure that the generated description includes the relevant product information rather than relying on a generic request alone.

An agent can also be connected to business knowledge and actions. For example, it can retrieve approved product data from a catalog system and invoke a publishing workflow after the description is approved, supporting a streamlined path from creation to immediate posting. Microsoft Copilot Studio is designed to let organizations create agents that use knowledge, orchestration, and actions without requiring users to build or train models themselves. This makes an interactive agent appropriate where usability, guided data selection, and integration with business processes are central requirements. See [What is Microsoft Copilot Studio?](#) and [Add actions to a copilot.](#)

QUESTION NO: 16

You have a business unit that uses an AI solution to process loan applications.

You discover that the solution rejects the application of all applicants that are older than 60 years of age. Which Microsoft responsible AI principle is this violating?

- A. accountability

- B. transparency
- C. reliability and safety
- D. fairness

ANSWER: D

Explanation:

The correct principle is **fairness**. In Microsoft's Responsible AI framework, fairness requires AI systems to allocate opportunities, resources, and benefits equitably, while avoiding unjustified bias and discriminatory outcomes for people or groups. A loan-processing solution that automatically rejects every applicant older than 60 creates a systematic age-based disparity in access to credit. Age is a personal characteristic that must be considered carefully in high-impact decisions such as lending, where automated outcomes can materially affect an individual's financial opportunities.

Fairness is not limited to whether a model behaves consistently. A system can apply the same rule to every affected person and still produce an inequitable result. Responsible AI practice requires organizations to assess how a system performs across relevant demographic groups, identify harmful disparities, determine whether differences are justified by the business context and applicable requirements, and mitigate unjustified outcomes. In this scenario, rejecting all applicants above a defined age threshold is precisely the kind of outcome that requires fairness evaluation and remediation.

Microsoft describes fairness as ensuring AI systems do not allocate opportunities or resources in ways that create unfair outcomes. Learn more in [Microsoft Responsible AI](#) and [Microsoft Learn: Fairness in Responsible AI](#).

QUESTION NO: 17

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:

customizing agent behavior and responses is directly supported because Copilot Studio provides authoring and configuration features that determine how an agent interacts with users. Makers can define instructions, create topics and trigger phrases, configure generative orchestration and answers, specify knowledge sources, and add actions. These controls enable an organization to align an agent's conversational behavior, tone, grounding, and task execution with a particular employee self-service or business-process scenario. Microsoft describes Copilot Studio as a platform for creating and managing agents, including configuring their behavior through instructions and other authoring capabilities.

building agents that connect to business data and automate user interactions is also a core Copilot Studio capability. Agents can use knowledge sources and connectors to access organizational information, and can invoke actions such as Power Automate flows, connectors, and custom APIs to perform work. This enables practical automation scenarios, such as retrieving business information, initiating approval processes, updating systems of record, and guiding employees through self-service tasks within a conversational experience. These capabilities are central to Copilot Studio's purpose of extending agents beyond simple question answering into business workflows. For details, see the [Microsoft Copilot Studio overview](#) and [Configure agent instructions in Copilot Studio](#).

QUESTION NO: 18 - (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

Azure Vision in Foundry Tools can extract and analyze key phrases from PDF files.

Yes

☐

No

☐

Azure Vision in Foundry Tools can generate images based on natural language descriptions.

☐☐

Azure Document Intelligence in Foundry Tools can be used to automate the processing of invoices and credit notes.

☐☐

ANSWER:

Explanation:

The correct selections are No, No, and Yes, in statement order. Azure Vision is intended for computer-vision workloads, such as analyzing visual content, detecting objects, generating image captions, and reading text through optical character recognition. Although text can be read from a document or image in an appropriate vision workflow, identifying the important concepts in that text is a natural-language-analysis task. Key phrase extraction is provided by Azure AI Language, as documented in Microsoft's [key phrase extraction overview](#). Therefore, the statement that attributes extracting and analyzing key phrases from PDFs to Azure Vision is not true.

Azure Vision also does not create a new image solely from a textual description. Its role is to derive information from visual input that already exists. Text-to-image creation is performed by generative image models, such as the image-generation capabilities available through Azure OpenAI. Microsoft describes this separately in its [Azure OpenAI image-generation models documentation](#). Thus, the statement about Azure Vision generating images from natural-language descriptions is correctly marked No.

Azure Document Intelligence is designed for intelligent document processing: it extracts structured data from forms and business documents so organizations can validate, route, and automate downstream work. Its prebuilt invoice model recognizes common invoice fields, including supplier, customer, invoice number, dates, totals, and line items. The service can also be extended with custom extraction models for organization-specific document layouts, which supports workflows involving invoices and credit notes. Microsoft's [prebuilt invoice model documentation](#) describes this extraction capability. Accordingly, the statement about automating processing of invoices and credit notes is correctly marked Yes.

QUESTION NO: 19

Your company plans to deploy a generative AI solution that helps employees answer questions about benefits and leave policies.

You need to reduce the risk that the solution provides harmful or misleading responses.

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

NOTE: Each correct selection is worth one point.

- A. Evaluate responses by using representative employee questions and edge cases.
- B. Provide a process for users to escalate sensitive or unresolved requests to HR.
- C. Configure the solution to answer every question without exception.
- D. Remove all response limitations to improve user satisfaction.
- E. Allow the solution to modify HR policies automatically.

ANSWER: A B

Explanation:

Evaluate responses by using representative employee questions and edge cases. Testing with realistic and challenging scenarios helps identify unsupported answers, unsafe behavior, inconsistent results, and gaps in the knowledge source before the solution is deployed.

Provide a process for users to escalate sensitive or unresolved requests to HR. Human oversight is important when the solution cannot provide a reliable answer or when a request involves an exception, sensitive personal circumstances, or an employment decision. An escalation path limits inappropriate reliance on automated output.

These controls support a responsible deployment approach that combines evaluation, mitigations, and meaningful human oversight. See [Responsible use of AI on Azure](#) and [Evaluate generative AI applications](#).

QUESTION NO: 20

Your company wants to establish an enterprise process for selecting AI use cases.

You need to ensure that proposed use cases are evaluated consistently before development resources are assigned.

Which two criteria should the company include in the evaluation process? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Expected business value and measurable outcomes
- B. Data readiness and risk profile
- C. The number of available experimental features
- D. The preference of a single technical team
- E. The ability to avoid business-owner involvement

ANSWER: A B

Explanation:

Expected business value and measurable outcomes. A proposed use case should have a clear business objective and defined measures, such as reduced processing time, improved service quality, lower operating cost, or increased employee productivity. Measurable outcomes enable the organization to determine whether the investment delivers the intended value.

Data readiness and risk profile. The organization should determine whether required data is available, accurate, authorized for the intended use, and appropriately protected. It should also assess risks such as privacy impacts, harmful outputs, regulatory obligations, and the need for human oversight. This helps identify use cases that are both feasible and suitable for responsible deployment.

Microsoft guidance recommends aligning AI initiatives to business strategy while considering data, governance, risk, and operational readiness. See [AI strategy in the Cloud Adoption Framework](#) and [AI governance in the Cloud Adoption Framework](#).