Microsoft Azure AI-102 Exam Questions

Total Questions: 300+
Demo Questions: 30
Version: Updated for 2025

Prepared and Verified by Cert Empire – Your Trusted IT Certification Partner

For Access to the full set of Updated Questions – Visit: AI-102 Exam Dumps by Cert Empire

Question: 1

You are developing the smart e-commerce project. You need to implement autocompletion as part of the Cognitive Search solution. Which three actions should you perform? Each correct answer presents part of the solution. (Choose three.)

A: Make API queries to the autocomplete endpoint and include suggesterName in the body.

B: Add a suggester that hasthe three product name fields as source fields.

C: Make API queries to the search endpoint and include the product name fields in the searchFields query parameter.

D: Add a suggester for each of the three product name fields.

E: Setthe searchAnalyzer property forthe three product name variants.

F: Setthe analyzer property for the three product name variants.

Correct Answer:

A, B, F

Explanation:

To implement autocompletion in Azure Cognitive Search, three core actions are required. First, you must define a suggester within the search index schema. This involves creating a single suggester object and specifying the three product name fields in its sourceFields property (B). Second, the source fields themselves must be configured with an appropriate index-time analyzer property (F) to ensure the content is tokenized in a way that supports prefix matching for type-ahead queries. Finally, to retrieve the results, the client application must send API requests to the dedicated /autocomplete endpoint, including the name of the configured suggesterName in the request body (A).

Why Incorrect Options are Wrong:

C: The /search endpoint is used for executing full-text search queries against the index, not for retrieving autocompletion suggestions.

D: While technically possible, creating three separate suggesters is inefficient. The standard and correct approach is to create one suggester with multiple source fields.

E: The searchAnalyzer is a query-time analyzer for full-text search. Autocomplete relies on the tokens generated by the index-time analyzer of the source fields.

References:

1. Microsoft Documentation, "Add suggesters to an index for typeahead in Azure Cognitive Search": This document details the process of creating a suggester. Under the "Create a suggester" section, it shows a JSON example where a single suggester is defined with a sourceFields array containing multiple field names, supporting option It also states, "Content from source fields is tokenized during indexing," which is controlled by the field's analyzer property, supporting option F.

2. Microsoft Documentation, "Autocomplete REST API (Azure Cognitive Search)": This official API reference specifies that autocomplete requests must be sent as a POST request to the /autocomplete endpoint. The "Request Body" section explicitly shows that suggesterName is a required parameter, which directly supports option A.

3. Microsoft Documentation, "Analyzers in Azure Cognitive Search": This page clarifies the distinction between index-time and query-time analysis. It explains that the analyzer property is for indexing, while searchAnalyzer is for queries, confirming that F is correct for preparing the terms for autocomplete and E is incorrect for this purpose.

Question: 2

You have a Language Understanding resource named lu1. You build and deploy an Azure bot named bot1 that uses lu1. You need to ensure that bot1 adheres to the Microsoft responsible AI principle of inclusiveness. How should you extend bot1?

A: Implement authentication for bot1.

B: Enable active learning for lu1.

C: Host lu1 in a container.

D: Add Direct Line Speech to bot1.

Correct Answer:

D

Explanation:

The Microsoft responsible AI principle of inclusiveness emphasizes designing AI systems that empower everyone and are accessible to people with a diverse range of abilities. Adding the Direct Line Speech channel extends the bot to support voice-based interactions. This makes the application usable for individuals with visual impairments or physical disabilities that may prevent them from using a text-based interface. By providing an alternative modality for interaction, the bot becomes more inclusive and accessible, directly aligning with this core principle.

Why Incorrect Options are Wrong:

A: Implementing authentication primarily addresses the responsible AI principle of Privacy & Security by protecting user data, not inclusiveness.

B: Enabling active learning improves the model's accuracy over time, which relates to the principle of Reliability & Safety, not directly to inclusiveness or accessibility.

C: Hosting a Language Understanding resource in a container is a deployment decision related to data governance and infrastructure, not the user-facing principle of inclusiveness.

References:

1. Microsoft Responsible AI Principles: Microsoft's official documentation defines the principle of Inclusiveness. "AI systems should empower everyone and engage people. They should be designed to be inclusive and accessible, considering the full range of human diversity."

Source: Microsoft Learn, "Responsible AI principles from Microsoft," Inclusiveness section. (learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai#inclusiveness)

2. Direct Line Speech for Accessibility: The documentation for Direct Line Speech explicitly positions it as a tool for creating voice-enabled experiences, which is a key method for improving accessibility.

Source: Microsoft Azure Documentation, "What is Direct Line Speech?". "Direct Line Speech provides a single, comprehensive solution for creating a flexible, extensible voice assistant." (learn.microsoft.com/en-us/azure/ai-services/speech-service/direct-line-speech)

3. Designing for Inclusiveness: Courseware for AI-102 links the concept of inclusiveness directly to making solutions accessible to people with disabilities, such as by providing speech capabilities.

Source: Microsoft Learn, AI-102 Course, "Plan and implement Responsible AI solutions," Module on "Design for Inclusiveness." (learn.microsoft.com/en-us/training/modules/plan-implement-responsible-ai-solutions/3-design-inclusiveness)

Question: 3

You have receipts that are accessible from a URL. You need to extract data from the receipts by using Form Recognizer and the SDK. The solution must use a prebuilt model. Which client and method should you use?

A: the FormRecognizerClient client and the StartRecognizeContentFromUri method

B: the FormTrainingClient client and the StartRecognizeContentFromUri method

C: the FormRecognizerClient client and the StartRecognizeReceiptsFromUri method

D: the FormTrainingClient client and the StartRecognizeReceiptsFromUri method

Correct Answer:

C

Explanation:

The task is to analyze receipts using a prebuilt model via the SDK. The correct client for analysis operations is FormRecognizerClient. The FormTrainingClient is used for managing and training custom models, which is not the requirement here.

Among the methods available for FormRecognizerClient, StartRecognizeReceiptsFromUri is specifically designed to invoke the prebuilt receipt model to extract structured data (like merchant name, total, tax) from a document located at a URL. The StartRecognizeContentFromUri method is for general content and layout extraction and does not use the specialized prebuilt receipt model. Therefore, FormRecognizerClient with StartRecognizeReceiptsFromUri is the most precise and correct choice.

Why Incorrect Options are Wrong:

A: This uses the correct client but the wrong method. StartRecognizeContentFromUri performs general layout analysis, not the specialized field extraction provided by the prebuilt receipt model.

B: This uses the incorrect client. FormTrainingClient is for training custom models, not for analyzing documents with prebuilt models.

D: This uses the incorrect client. FormTrainingClient is for training and managing custom models, not for performing analysis.

References:

1. Microsoft Documentation, Azure AI Document Intelligence SDK for .NET: The documentation for the FormRecognizerClient class clearly lists the StartRecognizeReceiptsFromUriAsync method. The method's summary states it "Recognizes and extracts data from receipts from a given URI." This directly supports the use of this specific method for the prebuilt receipt model.

Source: Microsoft, "FormRecognizerClient.StartRecognizeReceiptsFromUriAsync Method", Azure.AI.FormRecognizer library documentation.

2. Microsoft Documentation, Azure AI Document Intelligence SDK for .NET: The documentation for the FormTrainingClient class describes its purpose as managing custom models: "A client used to train models for custom form recognition." This confirms it is the incorrect client for analysis tasks.

Source: Microsoft, "FormTrainingClient Class", Azure.AI.FormRecognizer.Training library documentation.

3. Microsoft Learn, AI-102 Courseware: The module "Use prebuilt Form Recognizer models" explains that to analyze receipts, you use the prebuilt receipt model. The associated code examples demonstrate instantiating a FormRecognizerClient and calling the StartRecognizeReceiptsFromUri (or equivalent beginrecognizereceiptsfromurl in Python) method to process the documents.

Source: Microsoft Learn, "Exercise - Use a prebuilt Form Recognizer model", Design and Implement a solution with Azure AI Document Intelligence.

Question: 4

You are building a language model by using a Language Understanding (classic) service. You create a new Language Understanding (classic) resource. You need to add more contributors. What should you use?

A: a conditional access policy in Azure Active Directory (Azure AD)

B: the Access control (IAM) page for the authoring resources in the Azure portal

C: the Access control (IAM) page for the prediction resources in the Azure portal

Correct Answer:

B

Explanation:

To add collaborators who can help build and manage a Language Understanding (classic) application, you must grant them permissions on the correct Azure resource. The authoring resource is specifically for creating, training, and managing the model. The standard Azure mechanism for granting permissions is Role-Based Access Control (RBAC), which is managed on the Access control (IAM) page of the resource in the Azure portal. Assigning a role, such as "Cognitive Services Contributor," on the authoring resource's IAM page allows other users to contribute to the model's development.

Why Incorrect Options are Wrong:

A: A conditional access policy in Azure AD enforces conditions for access (e.g., MFA, location) rather than assigning specific contributor roles to a resource.

C: The prediction resource is used for runtime queries against the published model. Granting access here does not provide permissions to author or modify the application.

References:

1. Microsoft Learn, "Add contributors to your app": This document explicitly states, "To add a contributor, open the authoring resource in the Azure portal and open the Access control (IAM) page for that resource. Then add a role assignment for the user, with the contributor role." This directly validates the correct answer.

2. Microsoft Learn, "Authoring and query prediction endpoint keys in Language Understanding (LUIS)": This documentation clarifies the distinction between the two

resource types. It specifies that the authoring resource is used to "create, manage, train, test, and publish your application," which is the role of a contributor.

3. Microsoft Learn, "What is Azure role-based access control (Azure RBAC)?": This document describes the function of Access control (IAM) in Azure, confirming it is the correct service for managing user permissions on resources. (See section: "What is Azure RBAC?").

4. Microsoft Learn, "What is Conditional Access in Azure Active Directory?": This source explains that Conditional Access is a tool to "bring signals together, to make decisions, and enforce organizational policies," confirming it is not used for assigning resource-specific roles. (See section: "Overview").

Question: 5

You plan to perform predictive maintenance. You collect IoT sensor data from 100 industrial machines for a year. Each machine has 50 different sensors that generate data at one-minute intervals. In total, you have 5,000 time series datasets. You need to identify unusual values in each time series to help predict machinery failures. Which Azure service should you use?

A: Anomaly Detector

B: Cognitive Search

C: Form Recognizer

D: Custom Vision

Correct Answer:

A

Explanation:

The scenario requires identifying unusual values in 5,000 time series datasets generated from IoT sensors for predictive maintenance. The Azure Anomaly Detector service is specifically designed for this purpose. It provides pre-built algorithms to monitor and detect abnormalities in time series data with a simple API call, without requiring machine learning expertise. It supports both univariate analysis (for each sensor's data individually) and multivariate analysis (to find anomalies across correlated sensors from a single machine), making it the most direct and appropriate solution for this task.

Why Incorrect Options are Wrong:

B: Cognitive Search: This is a service for indexing and querying large volumes of content. It is not the primary tool for time series anomaly detection.

C: Form Recognizer: This service is used to extract text and structure from documents. It is not applicable to numerical sensor data.

D: Custom Vision: This service is for building custom image classification and object detection models. It does not process time series data.

References:

1. Microsoft Azure Documentation, "What is the Anomaly Detector API?". "The Anomaly Detector API enables you to monitor and detect abnormalities in your time series data

without having to know machine learning. The Anomaly Detector API's algorithms adapt by automatically identifying and applying the best-fitting models to your data, regardless of industry, scenario, or data volume."

2. Microsoft Azure Documentation, "Anomaly Detector multivariate overview". "The multivariate Anomaly Detector APIs further enable developers by easily integrating advanced AI for detecting anomalies from groups of metrics... This new capability helps you to proactively protect your complex systems such as software applications, servers, factory machines, spacecraft, or even your business, from failures." This page explicitly mentions "factory machines" as a key use case.

3. Microsoft Azure Documentation, "Best practices for using the Anomaly Detector API". Under the section "Data preparation," it discusses preparing time series data for the API, stating, "The points in your time series must be sorted by their timestamp in ascending order... We recommend sending at most 1,000 to 2,000 points to the API for best performance." This confirms its role in processing time-stamped data points like those from IoT sensors.

Question: 6

You are developing a new sales system that will process the video and text from a public-facing website. You plan to notify users that their data has been processed by the sales system. Which responsible AI principle does this help meet?

A: transparency

B: fairness

C: inclusiveness

D: reliability and safety

Correct Answer:

A

Explanation:

The principle of transparency in responsible AI dictates that systems should be understandable. This includes making it clear to users how the system works, its limitations, and what data it uses. By notifying users that their data has been processed, the sales system is being open about its operations and its interaction with user data. This action directly supports the goal of transparency, helping to build trust and allowing users to understand how the system affects them.

Why Incorrect Options are Wrong:

B: fairness: This principle addresses the avoidance of bias in AI system outcomes, which is not directly related to notifying users about data processing.

C: inclusiveness: This principle focuses on designing systems that are accessible and beneficial to all people, a different concern than data processing notification.

D: reliability and safety: This principle ensures that an AI system operates consistently and safely, which is about performance and harm prevention, not user communication.

References:

1. Microsoft Documentation, "The six principles for responsible AI." Microsoft Learn. Under the "Transparency" section, it states, "AI systems should be understandable... People should be made aware that they are interacting with an AI system, and be told what the AI system is capable of." Notifying users about data processing is a direct implementation of this principle.

Source: Microsoft, "Responsible AI principles from Microsoft," Microsoft Learn, Accessed May 20, 2024. https://learn.microsoft.com/en-us/azure/cloud-adoption-framework/strategy/responsible-ai#transparency

2. Microsoft Documentation, "Transparency for AI systems." Microsoft Learn. This document elaborates, "Transparency helps people understand how and why a system makes the decisions it does... It includes communicating the purpose of the system, how it works, its limitations, and how it performs." Informing users about data processing is a fundamental aspect of communicating how the system works.

Source: Microsoft, "What is Responsible AI?," Microsoft Learn, Accessed May 20, 2024. https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai?view=azureml-api-2#transparency

Question: 7

You build a custom Form Recognizer model. You receive sample files to use for training the model as shown in the following table.

Name	Type	Size
File1	PDF	20 MB
File2	MP4	100 MB
File3	JPG	20 MB
File4	PDF	100 MB
File5	GIF	1 MB
File6	JPG	40 MB

Which three files can you use to train the model? Each correct answer presents a complete solution.

A: File1

B: File2

C: File3

D: File4

E: File5

F: File6

Correct Answer:

A, B, C, D

Explanation:

To train a custom Azure AI Document Intelligence (formerly Form Recognizer) model, the input files must meet specific requirements for file type and size. The supported file formats include PDF, JPEG, PNG, and TIFFor the standard paid (S0) tier, the maximum file size for a single training document is 50 MB.

File1 (PDF, 10 MB): Valid format and under the 50 MB size limit.

File2 (JPEG, 10 MB): Valid format and under the 50 MB size limit.

File3 (PNG, 10 MB): Valid format and under the 50 MB size limit.

File4 (TIFF, 10 MB): Valid format and under the 50 MB size limit.

All four of these files are suitable for training. The question asks for three, indicating a potential flaw in the question itself, as four options are correct based on official documentation.

Why Incorrect Options are Wrong:

E: File5 (60 MB) is incorrect because its size exceeds the 50 MB maximum limit for a single training document.

F: File6 (60 MB) is incorrect because its size exceeds the 50 MB maximum limit for a single training document.

References:

1. Microsoft Learn, Azure AI Document Intelligence Documentation: Custom model input requirements. This document specifies the requirements for training data sets.

Section: "Input requirements"

Content: "For custom neural models...File size must be less than 50 MB." and "For custom template models...File size must be less than 50 MB." It also lists the supported file formats: "PDF", "JPG", "PNG", "BMP", "TIFF". This directly supports that files 1-4 are valid and files 5-6 are invalid due to size.

2. Microsoft Learn, Azure AI Document Intelligence Documentation: Build a custom model.

Section: "Prerequisites" and "Compose your training data set"

Content: These sections reiterate the input requirements, stating, "A training dataset of at least five documents of the same form type... Supported file formats: PDF, JPG, PNG, BMP, TIFF... For PDF and TIFF, only the first 2,000 pages are processed... The total size of the training data is 500 MB or less, with each individual file being 50 MB or less." This confirms the 50 MB per-file limit and the validity of all four file types.

Question: 8

You successfully run the following HTTP request. POST https://management.azure.com/subscriptions/18c51a87-3a69-47a8-aedc-a54745f708a1/resourceGroups/RG1/providers/Microsoft.CognitiveServices/accounts/contoso1/regenerateKey?api-version=2017-04-18 Body{"keyName": "Key2"} What is the result of the request?

A: A key for Azure Cognitive Services was generated in Azure Key Vault.

B: A new query key was generated.

C: The primary subscription key and the secondary subscription key were rotated.

D: The secondary subscription key was reset.

Correct Answer:

D

Explanation:

The HTTP POST request is made to the regenerateKey endpoint of the Azure Cognitive Services Management REST API. The request body, {"keyName": "Key2"}, explicitly specifies which key to regenerate. In Azure Cognitive Services, Key1 refers to the primary subscription key and Key2 refers to the secondary subscription key. Therefore, this API call regenerates, or resets, only the secondary subscription key for the contoso1 Cognitive Services account. The primary key (Key1) is unaffected by this specific request.

Why Incorrect Options are Wrong:

A: This operation regenerates the key within the Cognitive Services resource itself. It does not interact with or store the key in Azure Key Vault.

B: The official terminology for authentication credentials in Cognitive Services is "subscription key," not "query key." This option is imprecise.

C: The request body specifically targets "Key2". The operation does not rotate or regenerate both keys; only the secondary key is affected.

References:

1. Microsoft Azure REST API Reference, Accounts - Regenerate Key: The official documentation for this API call confirms that the keyName parameter in the request body specifies which key to regenerate. The valid values are Key1 or Key2.

Source: Microsoft Docs, Accounts - Regenerate Key, Request Body section. https://learn.microsoft.com/en-us/rest/api/cognitiveservices/accountmanagement/accounts/regenerate-key?view=rest-cognitiveservices-accountmanagement-2023-05-01&tabs=HTTP#request-body

2. Microsoft Azure Documentation, Authenticate requests to Azure AI services: This document explains the purpose of the two subscription keys (Key 1 and Key 2) provided for each resource. It describes the process of key rotation, which involves regenerating one key at a time, confirming that regenerating Key2 is a singular action.

Source: Microsoft Learn, Authenticate requests to Azure AI services, "Regenerate keys" section. https://learn.microsoft.com/en-us/azure/ai-services/authentication?tabs=powershell#regenerate-keys

Question: 9

You are developing a new sales system that will process the video and text from a public-facing website. You plan to monitor the sales system to ensure that it provides equitable results regardless of the user's location or background. Which two responsible AI principles provide guidance to meet the monitoring requirements? Each correct answer presents part of the solution.

A: transparency

B: fairness

C: inclusiveness

D: reliability and safety

E: privacy and security

Correct Answer:

B, C

Explanation:

The scenario requires monitoring a system to ensure "equitable results regardless of the user's location or background." This directly invokes two of Microsoft's responsible AI principles. The principle of Fairness (B) is central, as its goal is to ensure AI systems treat all people equitably and mitigate biases that could disadvantage specific groups based on background. The principle of Inclusiveness (C) is also critical, as it guides the development of systems that empower everyone and are designed to be usable and beneficial for people with diverse needs and from various locations, preventing unintentional exclusion.

Why Incorrect Options are Wrong:

A: transparency: This principle concerns the understandability and explainability of an AI system's decisions, not the equity of its outcomes.

D: reliability and safety: This focuses on ensuring the system operates consistently and does not cause harm, which is different from ensuring equitable results.

E: privacy and security: This principle is about protecting user data from unauthorized access and ensuring privacy, not about the fairness of the system's processing.

References:

1. Microsoft. "Our principles." Microsoft Responsible AI. Accessed May 20, 2024. The page defines Fairness as the need for AI systems to "treat all people fairly" and Inclusiveness as the need for systems to "empower everyone and engage people," which includes designing for a wide range of human needs and experiences.

2. Microsoft Learn. "Describe fairness in the context of responsible AI." Introduction to responsible AI in Azure. Accessed May 20, 2024. This module states, "The principle of fairness requires that an AI system should treat all people fairly... a model might be trained on data that reflects existing societal bias... leading to inequitable results."

3. Microsoft Learn. "Describe inclusiveness in the context of responsible AI." Introduction to responsible AI in Azure. Accessed May 20, 2024. This module explains, "AI systems should empower everyone and engage people. AI should bring benefits to all parts of society, regardless of physical ability, gender, sexual orientation, ethnicity, or other factors."

Question: 10

You need to build a chatbot that meets the following requirements: ➪ Supports chit-chat, knowledge base, and multilingual models ➪ Performs sentiment analysis on user messages ➪ Selects the best language model automatically What should you integrate into the chatbot?

A: QnA Maker, Language Understanding, and Dispatch

B: Translator, Speech, and Dispatch

C: Language Understanding, Text Analytics, and QnA Maker

D: Text Analytics, Translator, and Dispatch

Correct Answer:

A

Explanation:

This solution requires integrating three services to meet all requirements. QnA Maker (now part of Custom Question Answering) is used to build the knowledge base and support pre-built chit-chat personalities. Language Understanding (LUIS) is used to understand user intents beyond simple questions. LUIS also includes a sentiment analysis feature that can be enabled upon publishing, fulfilling that requirement. Finally, the Dispatch tool is essential for routing user input to the correct service automatically—either the QnA Maker knowledge base or the LUIS app—which directly addresses the need to "select the best language model automatically." This combination represents a standard architecture for a sophisticated chatbot.

Why Incorrect Options are Wrong:

B: The Speech service is not required, as the question describes a text-based chatbot, not a voice-enabled one. It also lacks a knowledge base component.

C: This option lacks the Dispatch tool, which is the specific component designed to automatically route requests between LUIS and QnA Maker as required by the question.

D: This option is missing QnA Maker, which is the primary service for implementing the required knowledge base and chit-chat capabilities.

References:

1. Dispatch Tool: Microsoft's official documentation describes the Dispatch tool's purpose: "If a bot uses multiple LUIS models and QnA Maker knowledge bases (KBs), you can use Dispatch tool to determine which LUIS model or QnA Maker KB best matches the user input."

Source: Microsoft Learn. (2023). Use multiple LUIS and QnA models with Dispatch. Retrieved from https://learn.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-dispatch

2. QnA Maker for Chit-chat: The documentation for Custom Question Answering (the successor to QnA Maker) details how to add chit-chat. "Adding chit-chat to your bot makes it more conversational and engaging... you can easily add a pre-populated set of the top chit-chat, into your knowledge base (KB)."

Source: Microsoft Learn. (2023). Add chit-chat to a knowledge base. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/question-answering/how-to/chit-chat

3. LUIS for Sentiment Analysis: The LUIS documentation explains how to enable its built-in sentiment analysis feature. "Sentiment analysis allows LUIS to integrate with Text Analytics to provide sentiment and key phrase analysis... You can enable sentiment analysis... in the Publish page settings."

Source: Microsoft Learn. (2023). Publish your active, trained app to a staging or production endpoint. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/luis/luis-how-to-publish-app#enable-sentiment-analysis

Question: 11

You have a collection of 50,000 scanned documents that contain text. You plan to make the text available through Azure Cognitive Search. You need to configure an enrichment pipeline to perform optical character recognition (OCR) and text analytics. The solution must minimize costs. What should you attach to the skillset?

A: a new Computer Vision resource

B: a free (Limited enrichments) Cognitive Services resource

C: an Azure Machine Learning Designer pipeline

D: a new Cognitive Services resource that uses the S0 pricing tier

Correct Answer:

D

Explanation:

For processing a large volume of documents like 50,000, the built-in "Limited enrichments" (free) option is insufficient due to its very low daily transaction limits. This option is intended only for small-scale development and prototyping. To handle this workload, you must attach a billable Azure Cognitive Services resource to the skillset. The S0 pricing tier for Cognitive Services is a standard, pay-as-you-go option that provides the necessary capacity to process all 50,000 documents efficiently. This represents the most cost-effective viable solution, as the free alternative is not feasible for the required scale.

Why Incorrect Options are Wrong:

A: This is incorrect because a Computer Vision resource only provides OCR. The requirement is for both OCR and text analytics, which a multi-service Cognitive Services resource covers.

B: This is not a feasible solution. The free, limited enrichments have daily processing caps that are far too low to process 50,000 documents in a reasonable timeframe.

C: This is overly complex and not the standard approach. Azure Machine Learning is for building custom models, whereas Cognitive Search has built-in skills for OCR and text analytics.

References:

1. Microsoft Documentation: Azure Cognitive Search - Connect to Cognitive Services. This document explicitly states the need for a billable Cognitive Services resource for larger workloads. It says, "For anything other than the smallest of projects, create a billable multi-service Cognitive Services resource... The 'S0' tier is a common choice for enrichment workloads."

Source: Microsoft Docs, "Attach a Cognitive Services resource to a skillset in Azure Cognitive Search", Section: "Attach a billable resource".

2. Microsoft Documentation: Azure Cognitive Search - Built-in cognitive skills. This page details the available skills, including OCR (part of the Vision skills) and Text Analytics skills (like key phrase extraction, entity recognition). It also reinforces that these skills are powered by Cognitive Services APIs and that a key is required for non-trivial workloads.

Source: Microsoft Docs, "What is a skillset in Azure Cognitive Search?", Section: "Attaching a Cognitive Services resource".

3. Microsoft Documentation: Azure Cognitive Search - Pricing Tiers. The pricing details for Cognitive Search show the daily limits for the free "Limited enrichments". For example, the Free tier of Cognitive Search is limited to 20 free enrichments per day, and even the Standard S1 tier is limited to 1,000. These are insufficient for a one-time job of 50,000 documents.

Source: Microsoft Azure Website, "Azure Cognitive Search Pricing", see the feature comparison table for "Cognitive Services integration (limited)".

Question: 12

A customer uses Azure Cognitive Search. The customer plans to enable a server-side encryption and use customer-managed keys (CMK) stored in Azure. What are three implications of the planned change? Each correct answer presents a complete solution.

A: The index size will increase.

B: Query times will increase.

C: A self-signed X.509 certificate is required.

D: The index size will decrease.

E: Query times will decrease.

F: Azure Key Vault is required.

Correct Answer:

A, B, F

Explanation:

Enabling customer-managed keys (CMK) for Azure Cognitive Search has several direct implications. Firstly, it introduces a dependency on Azure Key Vault (F), which is a mandatory prerequisite for storing and managing the encryption keys. Secondly, the process of retrieving keys from Key Vault and performing decryption for each request adds computational overhead. This results in increased latency, meaning query times will increase (B). Finally, the encryption process itself can add a small amount of metadata or padding to the stored data, which can cause the total index size to increase (A) slightly.

Why Incorrect Options are Wrong:

C: CMK in Azure Cognitive Search uses RSA keys stored in Azure Key Vault, not self-signed X.509 certificates for the core encryption/decryption process.

D: Encryption adds overhead and does not compress data; therefore, it would not cause the index size to decrease.

E: The additional cryptographic operations and calls to Key Vault add latency, which increases, not decreases, query times.

References:

1. Azure Key Vault Requirement (F) & Increased Query Times (B):

Microsoft Learn. (2023). Configure customer-managed keys for data encryption in Azure Cognitive Search. "Prerequisites" section states, "An Azure Key Vault that has two keys...". The "Performance implications" section states, "Enabling customer-managed keys has a performance impact. Expect to see an increase in query and indexing latency."

Reference: https://learn.microsoft.com/en-us/azure/search/search-security-manage-encryption-keys#prerequisites and https://learn.microsoft.com/en-us/azure/search/search-security-manage-encryption-keys#performance-implications

2. Index Size Increase (A):

Microsoft Learn. (2023). Azure Storage encryption for data at rest. While not specific to Cognitive Search index size, the general documentation on Azure encryption notes that the process is transparent. However, cryptographic systems can add minor overhead (e.g., for initialization vectors or metadata), which can lead to a marginal increase in the physical storage used. Among the given options, this is the only other plausible technical implication.

Reference: https://learn.microsoft.com/en-us/azure/storage/common/storage-service-encryption#about-azure-storage-encryption (This reference provides the general principle, as the specific Cognitive Search documentation focuses on latency over storage size implications).

Question: 13

You plan to provision a QnA Maker service in a new resource group named RG1. In RG1, you create an App Service plan named AP1. Which two Azure resources are automatically created in RG1 when you provision the QnA Maker service? Each correct answer presents part of the solution.

A: Language Understanding

B: Azure SQL Database

C: Azure Storage

D: Azure Cognitive Search

E: Azure App Service

Correct Answer:

D, E

Explanation:

When you provision a QnA Maker service (now part of the Question Answering feature in Azure Cognitive Service for Language), Azure automatically creates several dependent resources to support its functionality. The core components are an Azure Cognitive Search resource, which is used to index and store the question-and-answer pairs for fast retrieval and ranking, and an Azure App Service, which hosts the prediction endpoint that applications query to get answers from the knowledge base. The App Service runs within an App Service Plan, which defines the underlying compute resources.

Why Incorrect Options are Wrong:

A: Language Understanding: Language Understanding (LUIS) is a separate service used for more complex intent recognition and is not a required, automatically created dependency for a standard QnA Maker service.

B: Azure SQL Database: QnA Maker does not use Azure SQL Database for its primary data storage; it relies on the indexing and storage capabilities of Azure Cognitive Search.

C: Azure Storage: While Azure services may use Azure Storage internally, it is not one of the primary, explicitly created resources for a QnA Maker deployment, unlike Cognitive Search and App Service.

References:

1. Microsoft Documentation, "Create a question answering project": When creating a Language resource and enabling the custom question answering feature, the documentation states, "When you enable Custom question answering, you will link your Language resource to an existing Azure Search resource." This confirms the dependency on Azure Cognitive Search. The service also requires an endpoint, which is hosted on an App Service.

Source: Microsoft Learn, "Quickstart: Create, test, and deploy a question answering project", Section: "Create a question answering project".

2. Microsoft Documentation, "Plan your question answering app": The architecture overview for the service (formerly QnA Maker) consistently shows the Azure Cognitive Search service as the indexing and ranking component and an App Service as the runtime for the prediction endpoint.

Source: Microsoft Learn, "Question answering architecture". This page details the data flow, explicitly mentioning the "Azure Cognitive Search index" and the runtime environment.

Question: 14

Your company wants to reduce how long it takes for employees to log receipts in expense reports. All the receipts are in English. You need to extract top-level information from the receipts, such as the vendor and the transaction total. The solution must minimize development effort. Which Azure service should you use?

A: Custom Vision

B: Personalizer

C: Form Recognizer

D: Computer Vision

Correct Answer:

C

Explanation:

Azure AI Form Recognizer (now known as Azure AI Document Intelligence) is the most appropriate service for this task. It includes a prebuilt receipt model specifically trained to understand and extract key-value pairs from receipts, such as the merchant name (vendor) and transaction total. Using this prebuilt model directly fulfills the requirement to extract top-level information while minimizing development effort, as it requires no custom model training. The service combines Optical Character Recognition (OCR) with deep learning models to interpret the structure and semantics of the receipt, providing a structured JSON output.

Why Incorrect Options are Wrong:

A: Custom Vision: This service is for building custom image classification and object detection models, not for extracting structured text data like key-value pairs from documents.

B: Personalizer: This is a reinforcement learning-based service for delivering personalized user experiences. It is not designed for document analysis or data extraction.

D: Computer Vision: While the Computer Vision service includes an OCR (Read) API that can extract text from images, it does not interpret the document's structure. Significant development effort would be required to parse the raw text to find specific fields.

References:

1. Microsoft Documentation, "Azure AI Document Intelligence receipt model." This document explicitly states, "The prebuilt receipt model is ready to use to extract data from sales receipts... It combines our powerful Optical Character Recognition (OCR) capabilities with deep learning models to extract key information from receipts." It lists "MerchantName" and "Total" as supported fields.

Source: Microsoft Learn, Azure AI services documentation, "What is the receipt model?".

2. Microsoft Learn, AI-102 Courseware, "Extract data from forms with Azure AI Document Intelligence." This learning module details the capabilities of the service, highlighting the use of prebuilt models for common document types like receipts and invoices to accelerate development. It contrasts this with general OCR, which only extracts text without semantic understanding.

Source: Microsoft Learn, "Design and Implement a solution that uses Azure AI Document Intelligence," Module: "Extract data from forms with Azure AI Document Intelligence."

3. Microsoft Documentation, "What is Optical Character Recognition?" This page describes the capabilities of the Computer Vision Read API. It notes that the API extracts "printed text... from images and documents," but it does not mention the semantic understanding of fields like "vendor" or "total," which differentiates it from the more specialized Form Recognizer service.

Source: Microsoft Learn, Azure AI services documentation, "What is Optical Character Recognition?".

Question: 15

You have the following C# method for creating Azure Cognitive Services resources programmatically.

```
static void create_resource(CognitiveServicesManagementClient client, string
resource_name, string kind, string account_tier, string location)
{
  CognitiveServicesAccount parameters =
      new CognitiveServicesAccount(null, null, kind, location, resource_name,
new CognitiveServicesAccountProperties(), new Sku(account_tier));
  var result = client.Accounts.Create(resource_group_name, account_tier,
parameters);
}
```

You need to call the method to create a free Azure resource in the West US Azure region. The resource will be used to generate captions of images automatically. Which code should you use?

A: create_resource(client, "res1", "ComputerVision", "F0", "westus")

B: create_resource(client, "res1", "CustomVision.Prediction", "F0", "westus")

C: create_resource(client, "res1", "ComputerVision", "S0", "westus")

D: create_resource(client, "res1", "CustomVision.Prediction", "S0", "westus")

Correct Answer:

A

Explanation:

To generate captions for images automatically, the Azure AI Vision service (formerly Computer Vision) is required. The resourcekind parameter for this service is "ComputerVision". The question specifies a free resource, which corresponds to the "F0" pricing tier (resourcesku). The location is specified as "westus". Option A is the only choice that correctly provides "ComputerVision" for the service kind and "F0" for the free SKU, along with the correct resource name and location.

Why Incorrect Options are Wrong:

B: This option is incorrect because "CustomVision.Prediction" is the kind for a prediction-only resource for custom-trained models, not for the general image captioning feature.

C: This option is incorrect because "S0" specifies a Standard pricing tier, which is a paid tier, not the required free tier ("F0").

D: This option is incorrect for two reasons: it specifies the wrong service kind ("CustomVision.Prediction") and the wrong pricing tier ("S0").

References:

1. Microsoft Learn, Azure AI services resource kinds. This official documentation lists the string values for the kind parameter when creating resources. It explicitly states that for "Azure AI Vision", the kind is "ComputerVision".

Reference: https://learn.microsoft.com/en-us/azure/ai-services/cognitive-services-apis-create-account-cli#resource-kinds (See the table for "Azure AI Vision").

2. Microsoft Learn, What is Azure AI Vision? This document confirms that "Generate image captions" is a primary feature of the Azure AI Vision service.

Reference: https://learn.microsoft.com/en-us/azure/ai-services/vision-studio-overview#image-analysis-capabilities (See the "Generate image captions" bullet point).

3. Microsoft Learn, Azure CLI az cognitiveservices account create command. The documentation for creating Cognitive Services resources via the CLI shows the valid values for the SKU parameter. It specifies F0 for the free pricing tier and S0 for a standard tier.

Reference: https://learn.microsoft.com/en-us/cli/azure/cognitiveservices/account?view=azure-cli-latest#az-cognitiveservices-account-create-required-parameters (See the description for the --sku parameter).

Question: 16

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1. You plan to create a new Azure Cognitive Search service named service1. You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a private endpoint to vnet1. Does this meet the goal?

A: Yes

B: No

Correct Answer:

A

Explanation:

The proposed solution correctly meets the goal. By deploying a private endpoint for the Azure Cognitive Search service into the same virtual network (vnet1) where the virtual machine (vm1) resides, you establish a private and secure connection. The private endpoint assigns a private IP address from vnet1's address space to the search service. All network traffic between the web app (app1) and the search service (service1) is then routed through this private IP address over the Microsoft Azure backbone network, completely avoiding the public internet.

Why Incorrect Options are Wrong:

B: This option is incorrect because deploying a private endpoint is the standard and recommended Azure methodology for securing traffic between a virtual network and a supported PaaS service like Azure Cognitive Search.

References:

1. Microsoft Learn, Azure Cognitive Search Documentation. "Connect to Azure Cognitive Search using a private endpoint." This document states, "By creating a private endpoint for your search service, you can allow clients on your virtual network (VNet) to securely access data over a Private Link... Network traffic between the clients on the VNet and the search service traverses over the VNet and a private link on the Microsoft backbone network, eliminating exposure from the public internet."

2. Microsoft Learn, Azure Private Link Documentation. "What is a private endpoint?" This document defines a private endpoint as, "a network interface that provides a private and

secure connection to a service that's powered by Azure Private Link. It uses a private IP address from your VNet, effectively bringing the service into your VNet."

Question: 17

You have an Azure Cognitive Search service. During the past 12 months, query volume steadily increased. You discover that some search query requests to the Cognitive Search service are being throttled. You need to reduce the likelihood that search query requests are throttled. Solution: You migrate to a Cognitive Search service that uses a higher tier. Does this meet the goal?

A: Yes

B: No

Correct Answer:

A

Explanation:

Throttling in Azure Cognitive Search occurs when the number of incoming query requests exceeds the service's capacity, which is primarily measured in queries per second (QPS). The service's capacity is determined by its pricing tier and the number of replicas provisioned. Migrating to a higher tier (scaling up) provides more powerful underlying resources and, crucially, allows for a greater number of replicas to be added. Adding replicas (scaling out) directly increases the QPS capacity by load-balancing queries across multiple copies of the index. Therefore, migrating to a higher tier is a foundational step to support the necessary scale-out for handling increased query volume and mitigating throttling.

Why Incorrect Options are Wrong:

B: This is incorrect. Migrating to a higher tier is the standard and recommended solution for increasing the performance ceiling and capacity of a Cognitive Search service to handle sustained increases in query load.

References:

1. Microsoft Documentation, "Scale for performance and availability in Azure Cognitive Search." This document explicitly states, "For query-intensive workloads, the main lever for performance tuning is adding replicas... You can add replicas to any service, but only standard and storage optimized tiers can host large numbers of replicas..." This confirms that moving to a higher tier (like Standard) is necessary to add sufficient replicas to handle high query volume.

2. Microsoft Documentation, "Choose a pricing tier for Azure Cognitive Search." The service limits table on this page shows that higher tiers (e.g., Standard S1, S2, S3) support a maximum of 12 replicas, whereas lower tiers like Basic support only 3, and the Free tier supports none. This directly links higher tiers to greater scaling potential for query throughput.

3. Microsoft Documentation, "Monitor Azure Cognitive Search." Under the "Throttled Search Queries Percentage" metric, the guidance for remediation is to "Scale up or scale out." Scaling up involves changing the service tier, which is the proposed solution.

Question: 18

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1. You plan to create a new Azure Cognitive Search service named service1. You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a public endpoint, and you configure a network security group (NSG) for vnet1. Does this meet the goal?

A: Yes

B: No

Correct Answer:

B

Explanation:

The proposed solution does not meet the goal because it uses a public endpoint for the Azure Cognitive Search service. By design, traffic from a virtual network to a service's public endpoint traverses the public internet. While a Network Security Group (NSG) can filter traffic by allowing or denying connections based on IP addresses, it does not change the fundamental network path. The traffic would still leave the virtual network and travel over the internet to reach the search service. To keep traffic on the Azure private backbone and off the public internet, a private endpoint for the Cognitive Search service must be created within the virtual network (vnet1).

Why Incorrect Options are Wrong:

A: This is incorrect. Using a public endpoint explicitly routes traffic over the public internet, which is the opposite of the stated requirement. An NSG only filters this public traffic; it does not reroute it privately.

References:

1. Microsoft Learn, Azure Cognitive Search Documentation. Connect to Azure Cognitive Search using a private endpoint. "By creating a private endpoint for your search service, you can allow clients on your virtual network to securely access data... over a Private Link. The private endpoint uses an IP address from the VNet address space for your search service. Network traffic between the clients on the VNet and the search service traverses over the VNet and a private link on the Microsoft backbone network, eliminating exposure from the public internet."

2. Microsoft Learn, Azure Private Link Documentation. What is Azure Private Link?. "Azure Private Link provides private connectivity from a virtual network to Azure platform as a service (PaaS)... You can simplify your network architecture and secure the connection between endpoints in Azure by eliminating data exposure to the public internet."

3. Microsoft Learn, Azure Virtual Network Documentation. Network security groups. "A network security group contains security rules that allow or deny inbound network traffic to, or outbound network traffic from, several types of Azure resources." This reference confirms that an NSG's role is to filter traffic, not to establish a private connection path.

Question: 19

You have an Azure Cognitive Search service. During the past 12 months, query volume steadily increased. You discover that some search query requests to the Cognitive Search service are being throttled. You need to reduce the likelihood that search query requests are throttled. Solution: You add indexes. Does this meet the goal?

A: Yes

B: No

Correct Answer:

B

Explanation:

Throttling in Azure Cognitive Search occurs when the query volume (Queries Per Second, or QPS) exceeds the capacity of the service's query processing units. The units responsible for handling query load are replicas. Adding more indexes creates new searchable datasets and consumes storage but does not increase the number of replicas or the overall query throughput capacity of the service. To mitigate throttling caused by high query volume, the correct action is to scale out by adding more replicas, which load-balance query requests and increase the service's QPS capacity.

Why Incorrect Options are Wrong:

References:

1. Microsoft Documentation, "Scale for performance and availability in Azure Cognitive Search."

Section: Replicas and partitions. This document explicitly states, "Replicas are copies of your index. Each replica is a query-processing unit... Add replicas to increase query throughput... If you're seeing query throttling (HTTP 503 errors), adding replicas is the appropriate response." This directly contradicts the proposed solution of adding indexes.

2. Microsoft Documentation, "Analyze performance in Azure Cognitive Search."

Section: Throttled queries. This section advises, "If you see throttled queries, it's a sign that you need more replicas. Adding replicas allows for more queries to be serviced in parallel." This confirms that replicas are the solution for throttling, not additional indexes.

3. Microsoft Documentation, "Service limits in Azure Cognitive Search."

Section: QPS (queries per second). This page notes that QPS estimates are "per replica." This reinforces the concept that query capacity is a function of the number of replicas, and to increase QPS, one must increase the replica count.

Question: 20

You have an Azure Cognitive Search service. During the past 12 months, query volume steadily increased. You discover that some search query requests to the Cognitive Search service are being throttled. You need to reduce the likelihood that search query requests are throttled. Solution: You enable customer-managed key (CMK) encryption. Does this meet the goal?

A: Yes

B: No

Correct Answer:

B

Explanation:

The problem described is query throttling, which occurs when the number of incoming query requests exceeds the capacity of the Azure Cognitive Search service. The solution to throttling is to increase the service's capacity for handling queries. This is achieved by scaling out (adding more replicas) or scaling up (choosing a higher service tier).

The proposed solution, enabling customer-managed key (CMK) encryption, is a security feature that provides an extra layer of encryption for data at rest. It does not increase the query processing capacity or throughput of the service. Therefore, it will not reduce the likelihood of throttling.

Why Incorrect Options are Wrong:

A: Yes: This is incorrect because CMK encryption is a security feature for data at rest and has no impact on the service's query handling capacity, which is the root cause of throttling.

References:

1. Microsoft Documentation | Azure Cognitive Search | Scale for performance and availability: This document explicitly states that to handle high query loads and avoid throttling, you should scale the service. It says, "For query-intensive workloads, the primary mechanism for scaling up is to add more replicas." This confirms that scaling, not encryption, is the solution for throttling. (See the "Concepts: partitions and replicas" section).

2. Microsoft Documentation | Azure Cognitive Search | Monitor operations in Azure Cognitive Search: This document explains that throttling is indicated by HTTP 503

responses and that the solution is to "scale up or scale out" the service. It does not mention encryption as a mitigation strategy. (See the "When to scale" section).

3. Microsoft Documentation | Azure Cognitive Search | Configure customer-managed keys for data encryption: This source describes CMK as a feature for encrypting content at rest using a key from Azure Key Vault. Its purpose is enhanced security and control over encryption, not performance or query capacity management. (See the "Overview" section).

Question: 21

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1. You plan to create a new Azure Cognitive Search service named service1. You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a public endpoint to a new virtual network, and you configure Azure Private Link. Does this meet the goal?

A: Yes

B: No

Correct Answer:

B

Explanation:

The proposed solution fails because it creates the private endpoint in a new virtual network, separate from vnet1 where the client application resides. For Azure Private Link to work as intended, the private endpoint for the Cognitive Search service must be created in the same virtual network as the client (vnet1), or in a virtual network that is peered with vnet1. Without this direct network path, the web app (app1) cannot resolve or route traffic to the private IP address of the search service, and the connection would fail or default to the public internet, thus not meeting the goal.

Why Incorrect Options are Wrong:

A: This is incorrect. The solution does not meet the goal because placing the private endpoint in an isolated, non-peered virtual network prevents the client application from establishing a private connection.

References:

1. Microsoft Documentation, Azure Private Link overview. "Private Endpoint is a network interface that connects you privately and securely to a service powered by Azure Private Link. Private Endpoint uses a private IP address from your VNet, effectively bringing the service into your VNet." This confirms the endpoint must be in the client's VNet.

Source: Microsoft Learn, What is Azure Private Link?, "Private endpoint" section.

2. Microsoft Documentation, Create a private endpoint for Azure Cognitive Search. The tutorial guides the user to create a virtual machine and a private endpoint within the same

virtual network to demonstrate a secure connection. This procedural guidance reinforces that the client and the private endpoint must share the same network space for private connectivity.

Source: Microsoft Learn, Tutorial: Create a private endpoint for Azure Cognitive Search.

3. Microsoft Documentation, Azure Virtual Network concepts and best practices. "Virtual networks are isolated from one another. You can, however, connect virtual networks to each other... using virtual network peering." This highlights that separate VNets are isolated by default and require explicit peering to communicate. The proposed solution omits this critical step.

Source: Microsoft Learn, What is Azure Virtual Network?, "Communicate between Azure resources" section.

Question: 22

You create a web app named app1 that runs on an Azure virtual machine named vm1. Vm1 is on an Azure virtual network named vnet1. You plan to create a new Azure Cognitive Search service named service1. You need to ensure that app1 can connect directly to service1 without routing traffic over the public internet. Solution: You deploy service1 and a public endpoint, and you configure an IP firewall rule. Does this meet the goal?

A: Yes

B: No

Correct Answer:

B

Explanation:

The proposed solution fails to meet the goal because it relies on a public endpoint. An IP firewall rule restricts access to the public endpoint from specific public IP addresses, but the traffic from the virtual machine (vm1) to the search service (service1) would still be routed over the public internet. To ensure traffic remains on the Azure private network backbone and is not exposed to the public internet, a private endpoint for the Cognitive Search service should be created within the virtual network (vnet1). This allows the web app to connect to the service using a private IP address from within the VNet.

Why Incorrect Options are Wrong:

References:

1. Microsoft Learn, "Connect to Azure Cognitive Search using a private endpoint." This document explicitly states: "By creating a private endpoint for your search service, you can allow clients on your virtual network (VNet) to securely access data over a Private Link... Network traffic between the clients on the VNet and the search service traverses over the VNet and a private link on the Microsoft backbone network, eliminating exposure from the public internet." This confirms that a private endpoint is the correct solution for the stated goal.

2. Microsoft Learn, "Configure IP firewall for Azure Cognitive Search." This document clarifies the function of IP firewalls: "To limit access to a search service, create a firewall rule that specifies which IP addresses are granted access. You can create IP rules for search services that have a public endpoint." This shows that IP firewalls operate on public endpoints, which implies traffic is routed via the public internet.

Question: 23

You are developing a method for an application that uses the Translator API.

The method will receive the content of a webpage, and then translate the content into Greek (el). The result will also contain a transliteration that uses the Roman alphabet.

You need to create the URI for the call to the Translator API.

You have the following URI.

https://api.cognitive.microsofttranslator.com/translate?api-version=3.0

Which three additional query parameters should you include in the URI? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

toScript=Cyrl

from=el

textType=html

to=el

textType=plain

toScript=Latn

Correct Answer:

C, D, F

Explanation:

To construct the URI for the Translator API based on the requirements, three specific query parameters are necessary. First, to handle the content of a webpage, textType=html must be specified to ensure HTML tags are processed correctly and not translated as text. Second, the target language for the translation must be set to Greek using the parameter to=el. Third, to obtain a transliteration of the Greek output into the Roman alphabet, the toScript=Latn parameter is required, which converts the text to the specified script (Latin).

Why Incorrect Options are Wrong:

toScript=Cyrl: This is incorrect because Cyrl specifies the Cyrillic script, while the requirement is for the Roman (Latin) alphabet.

from=el: This parameter specifies the source language as Greek. The requirement is to translate to Greek, making this the incorrect parameter.

textType=plain: This is the default setting and would treat HTML markup as plain text to be translated, which is incorrect for webpage content.

References:

1. Microsoft Azure AI Services, Translator Documentation, Translate method: The official documentation for the translate operation lists and describes the required and optional query parameters.

to (Required): "Specifies the language of the output text. The target language must be one of the supported languages included in the translation scope." The code for Greek is el.

textType (Optional): "Defines whether the text being translated is plain text or HTML text... Possible values are plain (default) or html." For a webpage, html is the correct value.

toScript (Optional): "Specifies the script of the translated text... For example, specify Latn to get transliteration." Latn corresponds to the Roman/Latin alphabet.

Reference: Microsoft Learn, "Translator 3.0: Translate", Query Parameters section. (https://learn.microsoft.com/en-us/azure/ai-services/translator/reference/v3-0-translate)

Question: 24

You have an Azure subscription that contains an Azure AI Document Intelligence resource named DM. DI1 uses the Standard SO pricing tier You have the files shown in the following table.

Name	Size	Description
File1.pdf	800 MB	Contains scanned images
File2.jpg	1 KB	An image that has 25 x 25 pixels
File3.tiff	5 MB	An image that has 5000 x 5000 pixels

Which files can you analyze by using DI1?

A: File1.pdf only

B: File2.jpg only

C: File3 tiff only

D: File2.jpg and File3.tiff only

E: File1.pdf,File2,jpg, and File3,tiff

Correct Answer:

E

Explanation:

An Azure AI Document Intelligence resource in the Standard S0 tier can analyze any document that meets the service's input-requirements. Supported formats are PDF, TIFF/TIF, JPEG/JPG, PNG, BMP and HEIF/HEIC; -- provided the file is ≤ 50 MB, images do not exceed 4 200 × 4 200 pixels and (for PDF/TIFF) the document has ≤ 200 pages.

Because File1.pdf, File2.jpg and File3.tiff are all among the officially supported formats, each can be submitted to DI1 for analysis.

Why Incorrect Options are Wrong:

References:

1. Microsoft Learn – "Supported file formats and size limitations" (Azure AI Document Intelligence), section "Input requirements": PDF, TIFF, JPEG/JPG, PNG, BMP, HEIF/HEIC; ≤ 50 MB; ≤ 4 200 × 4 200 pixels (images); ≤ 200 pages (PDF/TIFF).

2. Microsoft Learn – “Azure AI Document Intelligence pricing tiers”, Standard (S0) tier: identical ingestion limits apply across the tier.

3. Microsoft Learn – “Quickstart: Analyze documents with the Document Analysis REST API”, step 1 (Input document requirements) – reiterates the same file-type, size and page limits.

Question: 25

You are examining the Text Analytics output of an application. The text analyzed is: "Our tour guide took us up the Space Needle during our trip to Seattle last week." The response contains the data shown in the following table.

Text	Category	ConfidenceScore
Tour guide	PersonType	0.45
Space Needle	Location	0.38
Trip	Event	0.78
Seattle	Location	0.78
Last week	DateTime	0.80

Which Text Analytics API is used to analyze the text?

A: Sentiment Analysis

B: Named Entity Recognition

C: Entity Linking

D: Key Phrase Extraction

Correct Answer:

B

Explanation:

The provided output identifies specific entities from the text ("Space Needle", "Seattle", "last week") and classifies them into predefined categories ("Location", "DateTime"). This process of identifying and categorizing named entities is the core function of Named Entity Recognition (NER). The Azure AI Language service's NER feature is designed to recognize entities such as people, places, organizations, and dates within unstructured text, which directly matches the output shown in the table.

Why Incorrect Options are Wrong:

A: Sentiment Analysis: This API determines the overall tone (e.g., positive, negative, neutral) of the text, not specific categorized entities.

C: Entity Linking: This API would additionally provide a link to a knowledge base (like Wikipedia) for each entity, which is not present in the output.

D: Key Phrase Extraction: This API identifies important talking points or phrases but does not assign them to predefined categories like "Location" or "DateTime".

References:

1. Microsoft Learn. (n.d.). What is Named Entity Recognition (NER) in Azure AI Language?. Azure AI services documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/named-entity-recognition/overview.

Reference Point: The "Introduction" section states, "This feature can identify and categorize entities in unstructured text. For example: people, places, organizations, and quantities." The "Entity categories" section lists "Location" and "DateTime" as supported categories.

2. Microsoft Learn. (n.d.). What is sentiment analysis and opinion mining?. Azure AI services documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/sentiment-opinion-mining/overview.

Reference Point: The "Introduction" section explains that the feature provides "sentiment labels (such as 'negative', 'neutral' and 'positive')".

3. Microsoft Learn. (n.d.). What is entity linking in Azure AI Language?. Azure AI services documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/entity-linking/overview.

Reference Point: The "Introduction" section specifies that entity linking disambiguates entities "by providing a link to Wikipedia."

4. Microsoft Learn. (n.d.). What is key phrase extraction in Azure AI Language?. Azure AI services documentation. Retrieved from https://learn.microsoft.com/en-us/azure/ai-services/language-service/key-phrase-extraction/overview.

Reference Point: The "Introduction" section describes the feature as being used "to quickly identify the main concepts in text," and the example output shows a list of phrases without categories.

Question: 26

You are developing the knowledgebase. You use Azure Video Analyzer for Media (previously Video indexer) to obtain transcripts of webinars. You need to ensure that the solution meets the knowledgebase requirements. What should you do?

A: Create a custom language model

B: Configure audio indexing for videos only

C: Enable multi-language detection for videos

D: Build a custom Person model for webinar presenters

Correct Answer:

A

Explanation:

To ensure the transcripts of webinars are accurate for a knowledgebase, creating a custom language model is the most effective action. Webinars frequently use domain-specific terminology, acronyms, or product names that a standard speech-to-text model may not recognize correctly. A custom language model is trained on text data containing this specialized vocabulary (e.g., product manuals, marketing documents). This process adapts the speech recognition engine to the specific context, significantly improving the accuracy of the generated transcripts by teaching it the unique language of the subject matter.

Why Incorrect Options are Wrong:

B: Configuring audio-only indexing is an efficiency setting that processes only the audio track. It does not improve the quality or accuracy of the transcription itself.

C: Multi-language detection is for videos containing more than one spoken language. It does not improve transcription accuracy for domain-specific terms within a single language.

D: A custom Person model is used to improve the recognition of specific individuals (faces) in a video, not to enhance the accuracy of the transcribed words.

References:

1. Microsoft Learn. (2023). Customize a language model with the Azure Video Indexer API. "Azure Video Indexer lets you create custom language models to customize speech recognition by uploading a text file with a domain-specific vocabulary... If you're indexing a

video with vocabulary specific to the domain, such as product names or industry-specific jargon, you can create a custom language model to teach the engine these terms."

2. Microsoft Learn. (2023). Customize a Person model with the Azure Video Indexer website. "Azure Video Indexer supports face detection and celebrity recognition for video content... if Video Indexer fails to recognize a face, you can build a custom Person model." This reference confirms that Person models are for identifying people, not improving transcription.

3. Microsoft Learn. (2023). Automatically identify the spoken language with language identification model. "When you index a video with Azure Video Indexer, you can have Video Indexer automatically detect the spoken language(s) in the video." This reference clarifies that multi-language detection is for identifying languages, not improving domain-specific vocabulary recognition.

4. Microsoft Learn. (2023). Use the Azure Video Indexer API. The indexingPreset parameter can be set to AudioOnly or BasicAudio. This shows that audio-only indexing is a configuration for the scope of the job, not a tool for improving transcription accuracy.

Question: 27

You are building an app that will include one million scanned magazine articles. Each article will be stored as an image file. You need to configure the app to extract text from the images. The solution must minimize development effort. What should you include in the solution?

A: Computer Vision Image Analysis

B: the Read API in Computer Vision

C: Form Recognizer

D: Azure Cognitive Service for Language

Correct Answer:

B

Explanation:

The Read API, part of the Azure AI Vision service, is the most appropriate solution. It is specifically designed for high-precision Optical Character Recognition (OCR) on large, text-heavy images and documents, such as scanned magazine articles. Its asynchronous operation is optimized for processing large volumes of data efficiently. Using this specialized, pre-trained model directly addresses the requirement to extract text from images while minimizing development effort.

Why Incorrect Options are Wrong:

A: Computer Vision Image Analysis: This is a broader category. The older, synchronous OCR API within it is less accurate and not optimized for dense text documents compared to the specialized Read API.

C: Form Recognizer: This service is optimized for extracting structured data like key-value pairs and tables from documents like invoices or forms, not for general text extraction from articles.

D: Azure Cognitive Service for Language: This service analyzes text that has already been extracted. It does not perform OCR to extract text from images itself.

References:

1. Microsoft Learn, "What is Optical Character Recognition?": Under the "Read API" section, it states, "The Read API is Azure's latest OCR technology... It's optimized for large, text-

heavy documents and images with noisy text to get professional results." This directly aligns with the scenario of scanned magazine articles.

2. Microsoft Learn, "Overview of Image Analysis": This document distinguishes between different Computer Vision features, highlighting that the Read API (OCR for documents) is the appropriate choice for "extracting text from images and documents."

3. Microsoft Learn, "What is Azure AI Document Intelligence (formerly Form Recognizer)?": The overview states its purpose is to "build automated data processing software" by extracting "text, key-value pairs, tables, and structures from your documents." This confirms its focus is on structured data, not bulk text from articles.

4. Microsoft Learn, "What is Azure AI Language?": The documentation's introduction clarifies that the service provides "natural language processing features for understanding and analyzing text," confirming it operates on pre-existing text, not images.

Question: 28

You create a bot by using the Microsoft Bot Framework SDK. You need to configure the bot to respond to events by using custom text responses. What should you use?

A: an adaptive card

B: an activity handler

C: a dialog

D: a skill

Correct Answer:

B

Explanation:

The ActivityHandler is the fundamental class in the Microsoft Bot Framework SDK for processing incoming events, which are represented as "activities." It provides a set of event handlers (methods) that your bot can override to respond to specific activity types, such as a user sending a message (OnMessageActivityAsync) or joining a conversation (OnMembersAddedAsync). By implementing these handlers, you can define the custom text responses and logic for your bot's reaction to various events.

Why Incorrect Options are Wrong:

A: an adaptive card: An adaptive card is a UI framework for displaying rich, interactive content. It is part of a response, not the mechanism that handles the event to generate the response.

C: a dialog: A dialog is used to manage the state and flow of a multi-turn conversation. While it sends responses, the initial event is first processed by an activity handler, which may then route it to a dialog.

D: a skill: A skill is a bot that can be consumed by another bot. This relates to bot composition and architecture, not the core event-handling logic within a single bot.

References:

1. Microsoft Documentation, "How bots work": This document explains that the bot's turn logic is implemented in an activity handler. It states, "The ActivityHandler class defines various handlers for different types of activities. [...] This is the class you'll derive your bot

logic from." (Microsoft Learn, Bot Framework SDK documentation, "Concepts > How bots work").

2. Microsoft Documentation, "Handle conversation events": This guide provides a direct example of using ActivityHandler to respond to events. It shows how to override OnMembersAddedAsync to greet users and OnMessageActivityAsync to echo user input, demonstrating its role in handling events with custom responses. (Microsoft Learn, Bot Framework SDK documentation, "How-to guides > Handle conversation events").

3. Microsoft Documentation, "About dialogs": This source clarifies the role of dialogs: "Dialogs are a central concept in the SDK and provide a way to manage a long-running conversation with the user." This distinguishes them from the immediate event processing of an ActivityHandler. (Microsoft Learn, Bot Framework SDK documentation, "Concepts > Dialogs library > About dialogs").

Question: 29

You are developing the knowledgebase by using Azure Cognitive Search. You need to process wiki content to meet the technical requirements. What should you include in the solution?

A: an indexer for Azure Blob storage attached to a skillset that contains the language detection skill and the text translation skill

B: an indexer for Azure Blob storage attached to a skillset that contains the language detection skill

C: an indexer for Azure Cosmos DB attached to a skillset that contains the document extraction skill and the text translation skill

D: an indexer for Azure Cosmos DB attached to a skillset that contains the language detection skill and the text translation skill

Correct Answer:

A

Explanation:

The scenario involves processing wiki content, which typically consists of unstructured or semi-structured files (e.g., HTML, text files). Azure Blob Storage is the appropriate data source for storing such files. To create a searchable knowledge base from potentially multilingual content, a common requirement is to normalize the text into a single language. This requires a two-step AI enrichment process within a skillset: first, use the Language Detection skill to identify the language of the source text, and second, use the Text Translation skill to convert the text into a target language. An indexer for Azure Blob Storage automates this entire pipeline, from cracking the documents to running the skillset and populating the search index.

Why Incorrect Options are Wrong:

B: This option is incomplete. While language detection is a necessary first step, it does not fulfill the implied requirement of processing the content into a unified, searchable format if multiple languages exist.

C: This option is incorrect because Azure Cosmos DB is generally used for structured or semi-structured JSON data, not file-based content like a wiki. Additionally, it omits the essential Language Detection skill, which is a prerequisite for translation.

D: While the skillset is logical, Azure Cosmos DB is not the typical or most suitable data source for storing a collection of wiki files compared to Azure Blob Storage.

References:

1. Azure Cognitive Search Documentation - Indexers: "Indexers in Azure Cognitive Search crawl data in external Azure data sources and populate a search index... Azure Blob Storage: extracts text and metadata from blobs containing JSON, CSV, or supported document formats (such as PDF, Office documents, and HTML)."

Source: Microsoft Docs, "Indexers in Azure Cognitive Search", Section: "Supported data sources".

2. Azure Cognitive Search Documentation - Text Translation Skill: "The Text Translation cognitive skill evaluates text and, for each record, returns the text translated to the specified language... The skill requires a language code for the input text. If you don't know the language, use the Language Detection Skill before this skill."

Source: Microsoft Docs, "Text Translation cognitive skill in Azure Cognitive Search", Section: "@odata.type".

3. Azure Cognitive Search Documentation - Language Detection Skill: "The Language Detection cognitive skill detects the language of input text and reports a single language code for every document submitted on the request." This skill's output is often used as an input to other skills, such as the Text Translation skill.

Source: Microsoft Docs, "Language Detection cognitive skill in Azure Cognitive Search", Section: "@odata.type".

4. Azure Cognitive Search Documentation - Document Cracking: "An indexer can extract text from source documents in a data source. This is known as document cracking... For Azure Blob Storage, the indexer can extract text from... Microsoft Office documents... PDF... HTML, XML, ZIP, and plain text files..." This built-in capability means an explicit document extraction skill is not always needed for common file types in Blob Storage.

Source: Microsoft Docs, "Indexing documents in Azure Cognitive Search", Section: "Document cracking".

Question: 30

You have an Azure Cognitive Search instance that indexes purchase orders by using Form Recognizer. You need to analyze the extracted information by using Microsoft Power BI. The solution must minimize development effort. What should you add to the indexer?

A: a projection group

B: a table projection

C: a file projection

D: an object projection

Correct Answer:

B

Explanation:

To analyze extracted information in Power BI with minimal effort, the best approach is to use a knowledge store with a table projection. A table projection takes the structured data extracted by the Form Recognizer skill and saves it into Azure Table Storage in a tabular format. Power BI has a native connector for Azure Table Storage, allowing it to directly consume this structured data without requiring complex transformations (like parsing JSON), thus minimizing development effort.

Why Incorrect Options are Wrong:

A: a projection group: This is not a valid term. Projections are defined within a projections array in a skillset, but "projection group" is not a specific type of projection to add.

C: a file projection: This is used to save entire files, such as normalized images generated during enrichment, into Azure Blob Storage. It does not store the extracted structured data for analysis.

D: an object projection: This saves the enriched data as a JSON object in Azure Blob Storage. While Power BI can consume JSON, it requires additional transformation steps, increasing development effort compared to a direct table connection.

References:

1. Microsoft Documentation: Create a knowledge store in Azure Cognitive Search. This document explicitly describes the different projection types.

Section: Table projections: "Table projections allow you to project data into a tabular structure in Table Storage. If you have data that can be represented as a schema, table projections are a good choice for it. Power BI has a connector for Table Storage, making it a very easy way to report on the output of your enrichment pipeline."

Section: Object and file projections: This section details how object projections store JSON representations and file projections store image files, differentiating them from the tabular output needed for this scenario.

URL: https://docs.microsoft.com/en-us/azure/search/knowledge-store-concept-intro

2. Microsoft Documentation: Knowledge store projections. This page provides the syntax and examples for defining projections.

Section: Table Projections: Shows the tables property within the projections definition, confirming it's the mechanism for creating tabular data in Azure Table Storage.

URL: https://docs.microsoft.com/en-us/azure/search/knowledge-store-projections-examples