- **Vendor:** Microsoft
- **Exam Code:** AI-100
- **Exam Name:** Designing and Implementing an Azure AI Solution
- **New Updated Questions from Braindump2go (Updated in Dec./2020)**

QUESTION 98
Hotspot Question
You are designing an application to parse images of business forms and upload the data to a database. The upload process will occur once a week.
You need to recommend which services to use for the application. The solution must minimize infrastructure costs.
Which services should you recommend? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

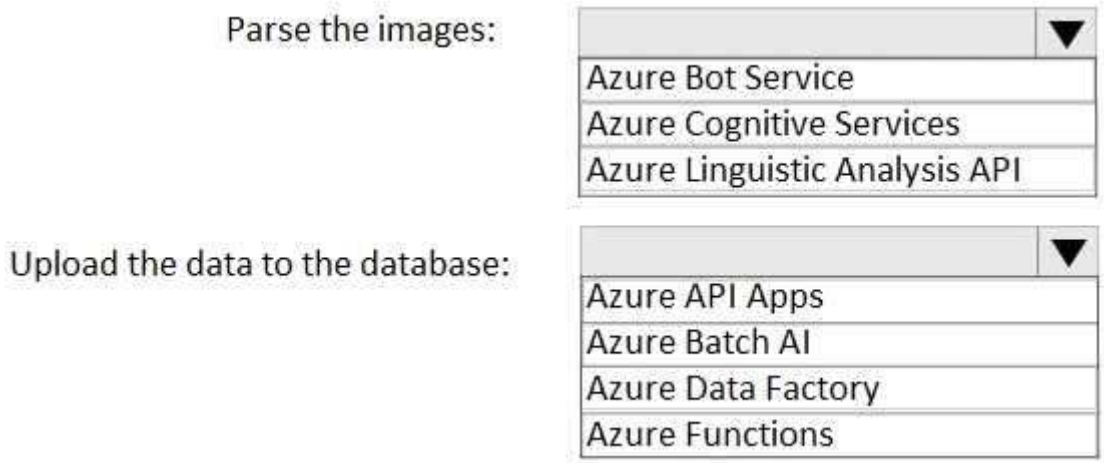

Answer:

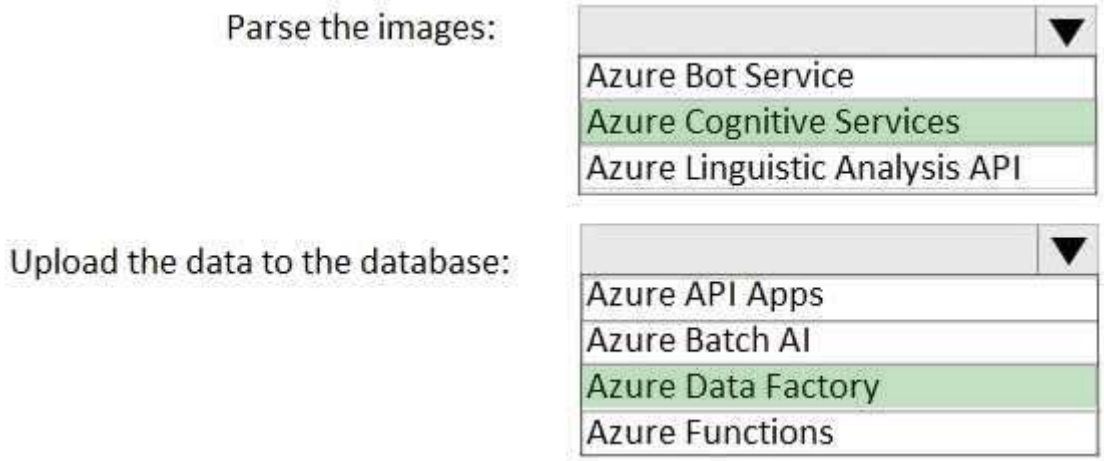

Explanation:
Box 1: Azure Cognitive Services

Azure Cognitive Services include image-processing algorithms to smartly identify, caption, index, and moderate your pictures and videos.
Not: Azure Linguistic Analytics API, which provides advanced natural language processing over raw text.
Box 2: Azure Data Factory
The Azure Data Factory (ADF) is a service designed to allow developers to integrate disparate data sources. It is a platform somewhat like SSIS in the cloud to manage the data you have both on-prem and in the cloud.
It provides access to on-premises data in SQL Server and cloud data in Azure Storage (Blob and Tables) and Azure SQL Database.
References:
https://azure.microsoft.com/en-us/services/cognitive-services/ https://www.jamesserra.com/archive/2014/11/what-is-azure-data-factory/

QUESTION 99
Hotspot Question
You plan to deploy an Azure Data Factory pipeline that will perform the following:
Move data from on-premises to the cloud.
Consume Azure Cognitive Services APIs.
You need to recommend which technologies the pipeline should use. The solution must minimize custom code.
What should you include in the recommendation? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

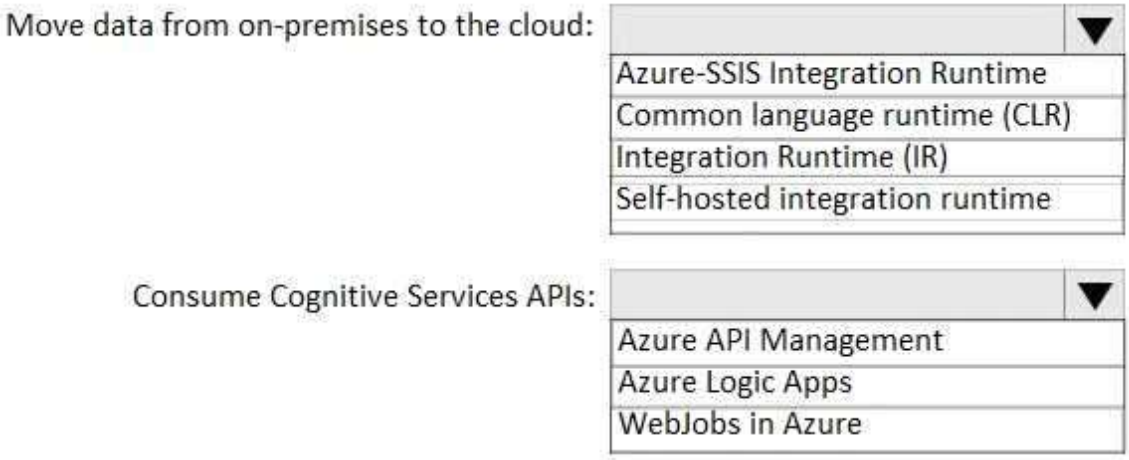

Answer:

Explanation:
Box 1: Self-hosted Integration Runtime
A self-hosted IR is capable of running copy activity between a cloud data stores and a data store in private network.
Not Azure-SSIS Integration Runtime, as you would need to write custom code.
Box 2: Azure Logic Apps

Azure Logic Apps helps you orchestrate and integrate different services by providing 100+ ready-to- use connectors, ranging from on-premises SQL Server or SAP to Microsoft Cognitive Services.
Incorrect:
Not Azure API Management: Use Azure API Management as a turnkey solution for publishing APIs to external and internal customers.
References:
https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-examples-and-scenarios

QUESTION 100
Hotspot Question
You need to build an interactive website that will accept uploaded images, and then ask a series of predefined questions based on each image.
Which services should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

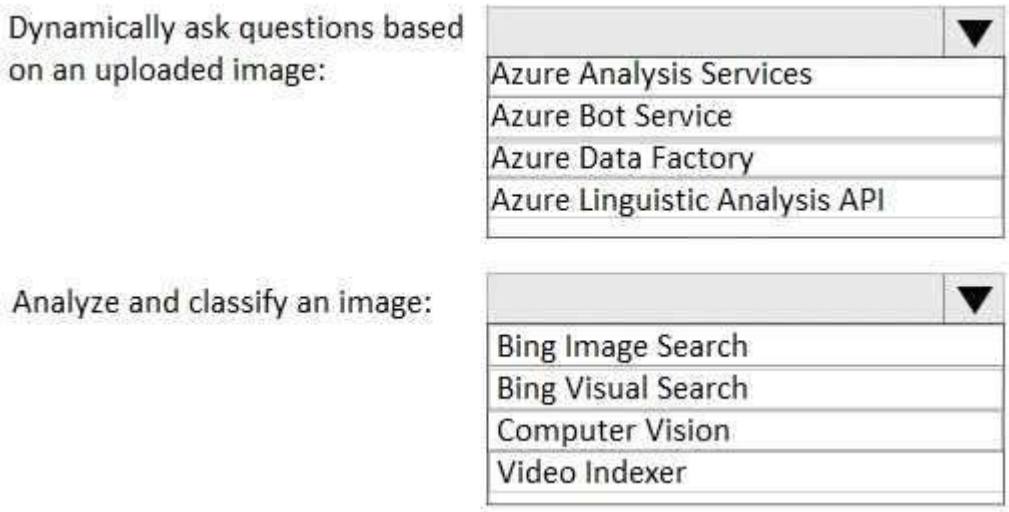

Answer:

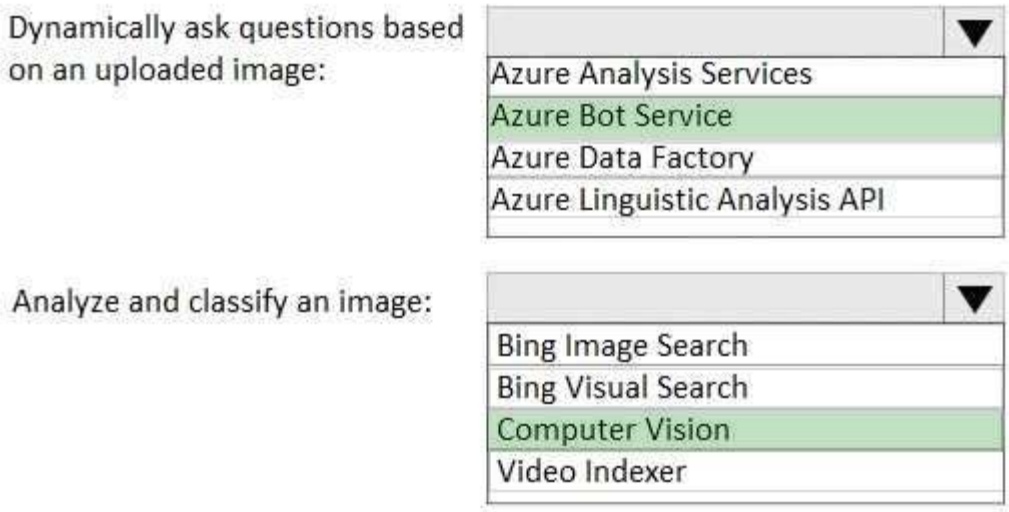

Explanation:
Box 1: Azure Bot Service
Box 2: Computer Vision
The Computer Vision Analyze an image feature, returns information about visual content found in an image. Use tagging, domain-specific models, and descriptions in four languages to identify content and label it with confidence. Use Object Detection to get location of thousands of objects within an image. Apply the adult/racy settings to help you detect potential adult content. Identify image types and color schemes in pictures.

References:
https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/

QUESTION 101
Hotspot Question
You are designing an AI solution that will be used to find buildings in aerial pictures. Users will upload the pictures to an Azure Storage account. A separate JSON document will contain for the pictures.
The solution must meet the following requirements:
- Store metadata for the pictures in a data store.
- Run a custom vision Azure Machine Learning module to identify the buildings in a picture and the position of the buildings' edges.
- Run a custom mathematical module to calculate the dimensions of the buildings in a picture based on the metadata and data from the vision module.
You need to identify which Azure infrastructure services are used for each component of the AI workflow. The solution must execute as quickly as possible. What should you identify? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

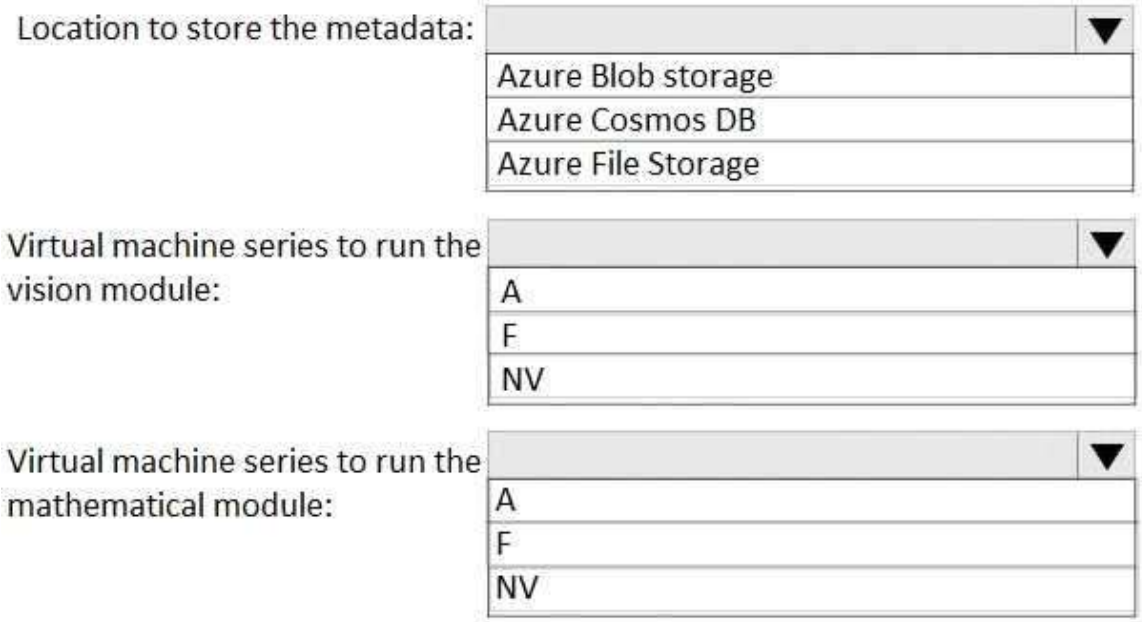

Answer:

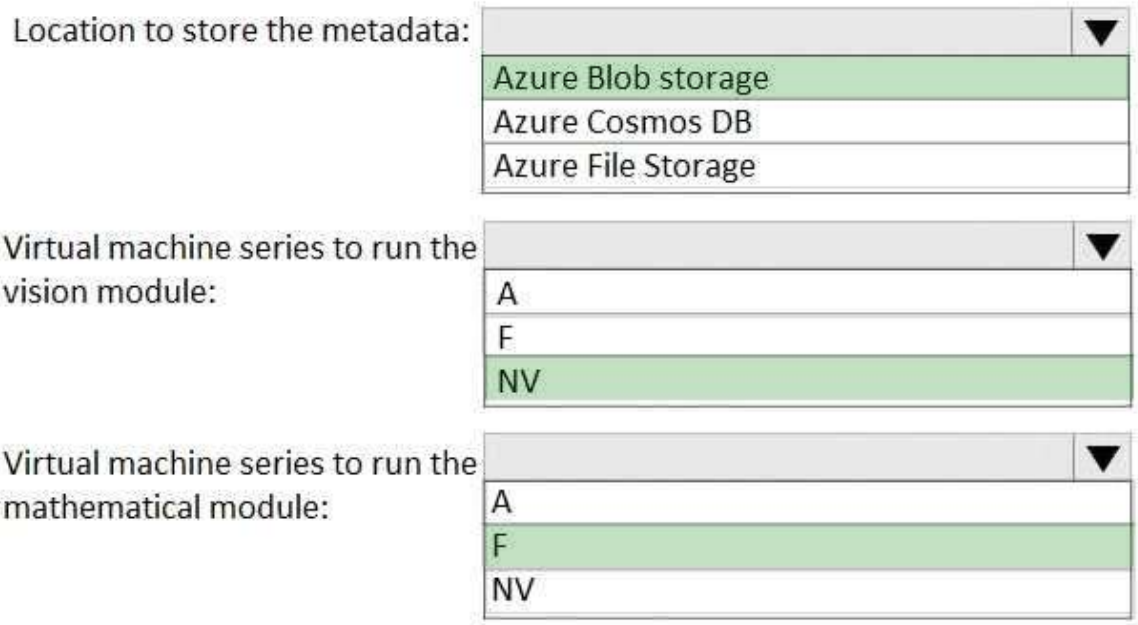

Explanation:
Box 1: Azure Blob Storage
Containers and blobs support custom metadata, represented as HTTP headers.
Box 2: NV
The NV-series enables powerful remote visualisation workloads and other graphics-intensive applications backed by the NVIDIA Tesla M60 GPU.
Note: The N-series is a family of Azure Virtual Machines with GPU capabilities. GPUs are ideal for compute and graphics-intensive workloads, helping customers to fuel innovation through scenarios like high-end remote visualisation, deep learning and predictive analytics.
Box 3: F
F-series VMs feature a higher CPU-to-memory ratio. Example use cases include batch processing, web servers, analytics and gaming.
Incorrect:
A-series VMs have CPU performance and memory configurations best suited for entry level workloads like development and test.
References:
https://azure.microsoft.com/en-in/pricing/details/virtual-machines/series/

QUESTION 102
Drag and Drop Question
You are designing an AI solution that will use IoT devices to gather data from conference attendees, and then later analyze the data. The IoT devices will connect to an Azure IoT hub.
You need to design a solution to anonymize the data before the data is sent to the IoT hub.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

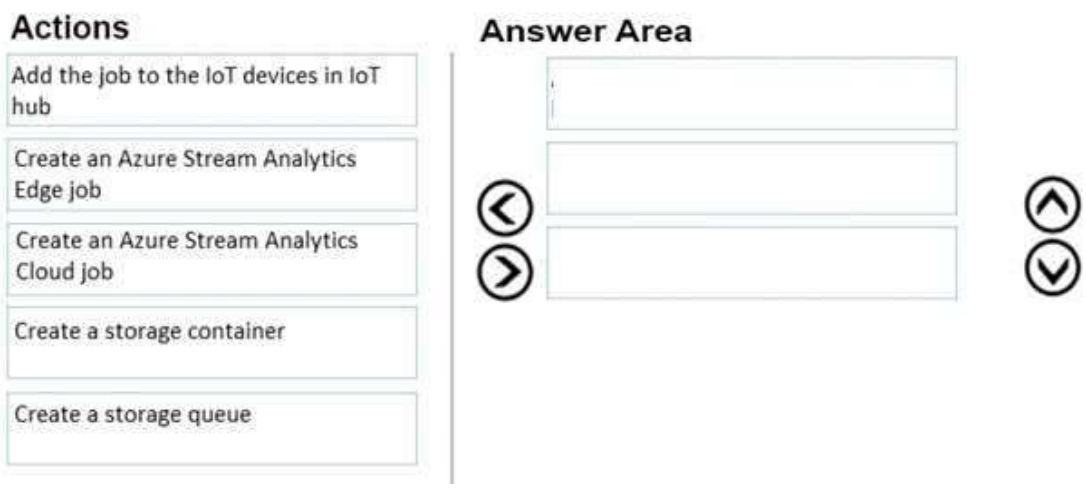

Answer:

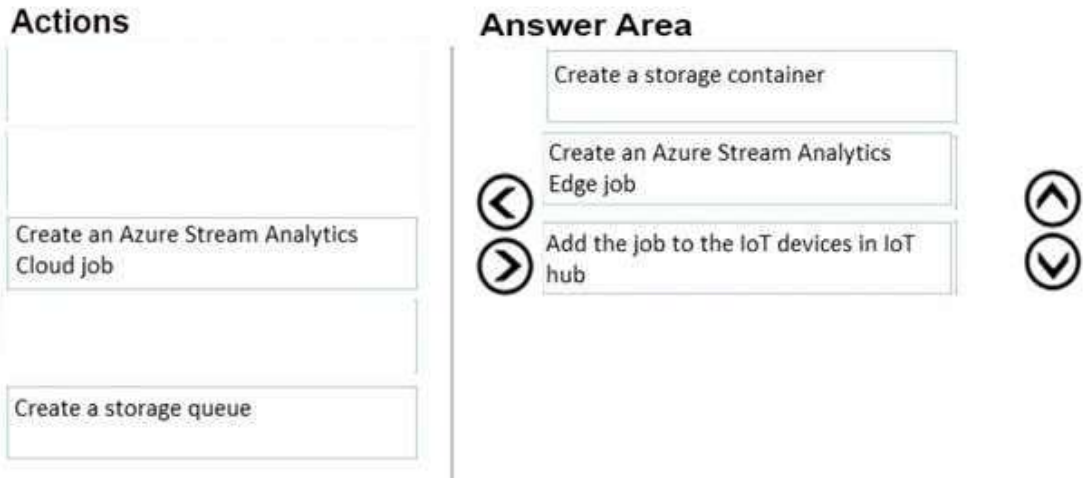

Explanation:
Step 1: Create a storage container
ASA Edge jobs run in containers deployed to Azure IoT Edge devices.
Step 2: Create an Azure Stream Analytics Edge Job
Azure Stream Analytics (ASA) on IoT Edge empowers developers to deploy near-real-time analytical intelligence closer to IoT devices so that they can unlock the full value of device-generated data.
Scenario overview:

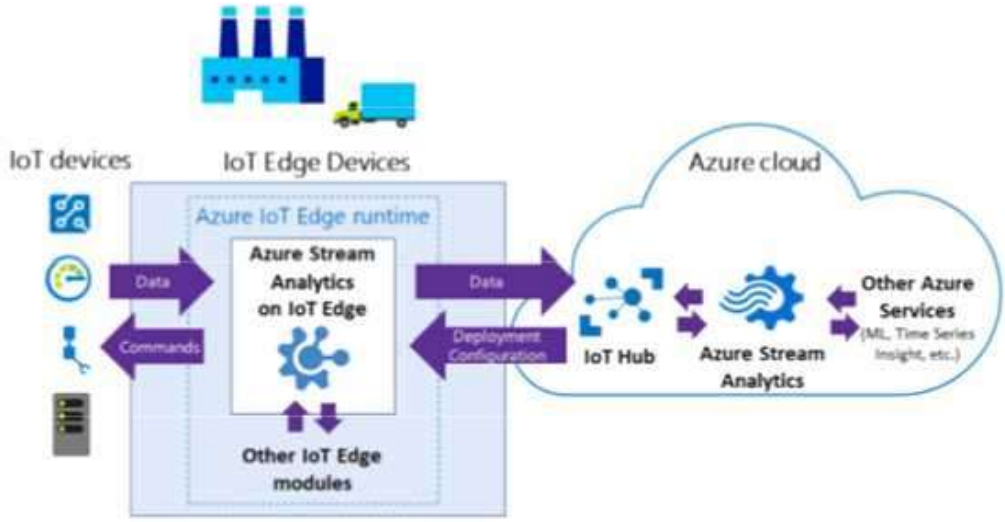

Step 3: Add the job to the IoT devices in IoT
References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-edge

QUESTION 103
Hotspot Question
You are developing an application that will perform clickstream analysis.
The application will ingest and analyze millions of messages in the real time.
You need to ensure that communication between the application and devices is bidirectional.
What should you use for data ingestion and stream processing? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

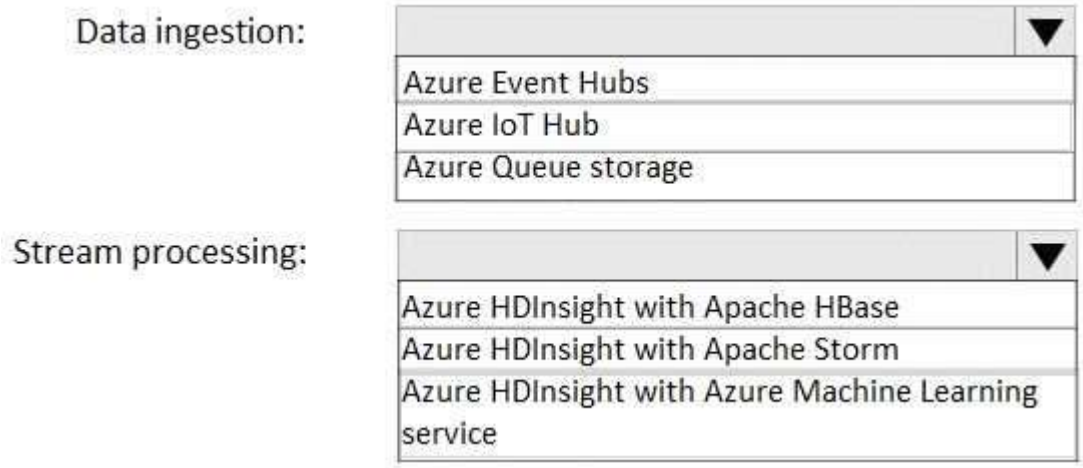

Answer:

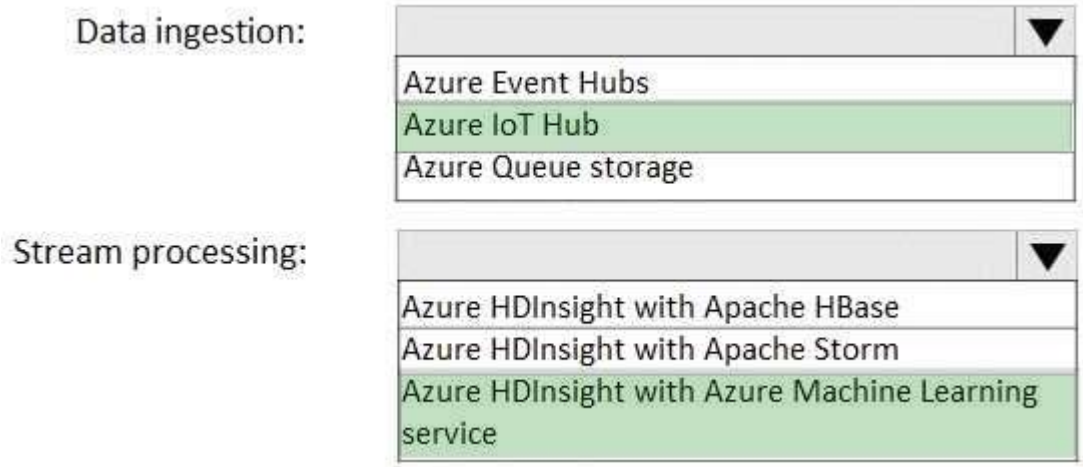

Explanation:
Box 1: Azure IoT Hub
Azure IoT Hub is the cloud gateway that connects IoT devices to gather data and drive business insights and automation. In addition, IoT Hub includes features that enrich the relationship between your devices and your backend systems. Bi-directional communication capabilities mean that while you receive data from devices you can also send commands and policies back to devices. Note on why not Azure Event Hubs: An Azure IoT Hub contains an Event Hub and hence essentially is an Event Hub plus additional features. An important additional feature is that an Event Hub can only receive messages, whereas an IoT Hub additionally can also send messages to individual devices. Further, an Event Hub has access security on hub level, whereas an IoT Hub is aware of the individual devices and can grand and revoke access on device level. Box 2: Azure Hdinsight with Azure Machine Learning service References:
https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-compare-event-hubs https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-machine-learning-overview

QUESTION 104
Drag and Drop Question
You develop a custom application that uses a token to connect to Azure Cognitive Services resources.
A new security policy requires that all access keys are changed every 30 days.
You need to recommend a solution to implement the security policy.
Which three actions should you recommend be performed every 30 days? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

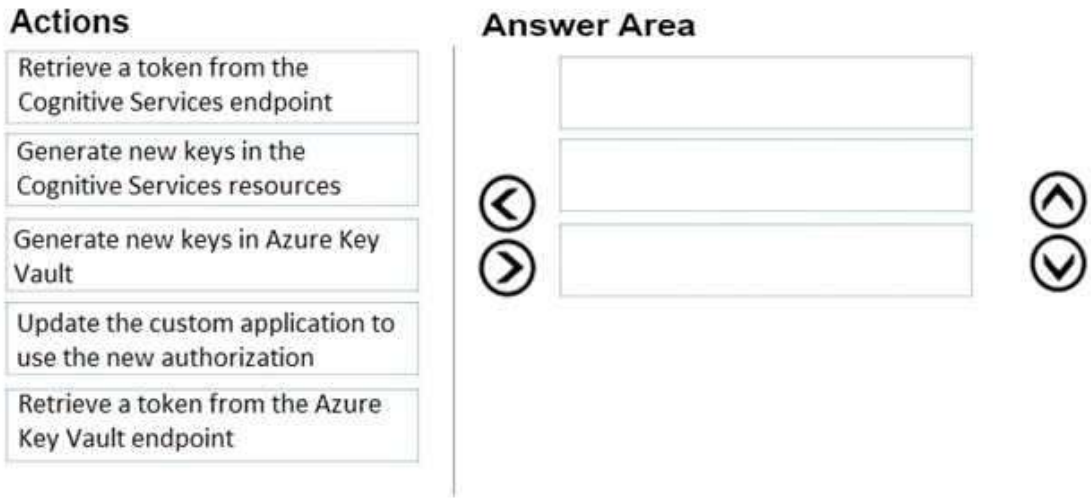

Answer:

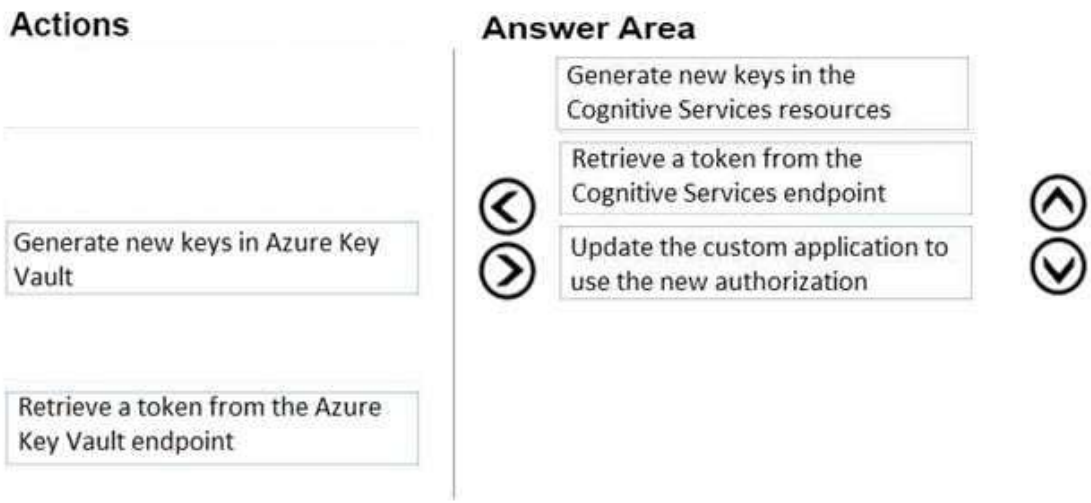

Explanation:
Step 1: Generate new keys in the Cognitive Service resources

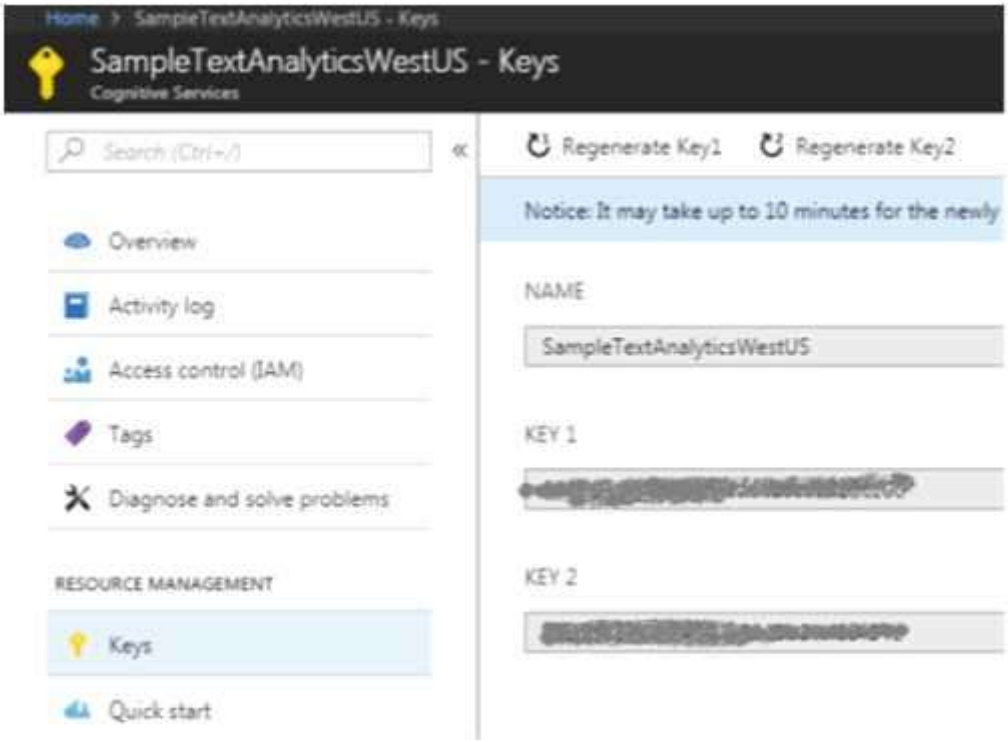

Step 2: Retrieve a token from the Cognitive Services endpoint Step 3: Update the custom application to use the new authorization Each request to an Azure Cognitive Service must include an authentication header. This header passes along a subscription key or access token, which is used to validate your subscription for a service or group of services.
References:
https://docs.microsoft.com/en-us/azure/cognitive-services/authentication

QUESTION 105
Drag and Drop Question
You use an Azure key vault to store credentials for several Azure Machine Learning applications.
You need to configure the key vault to meet the following requirements:
- Ensure that the IT security team can add new passwords and periodically change the passwords.
- Ensure that the applications can securely retrieve the passwords for the applications.
- Use the principle of least privilege.
Which permissions should you grant? To answer, drag the appropriate permissions to the correct targets. Each permission may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

Answer:

Explanation:
Incorrect Answers:
Not Keys as they are used for encryption only.
References:

https://docs.microsoft.com/en-us/azure/key-vault/key-vault-secure-your-key-vault

QUESTION 106
Your company is building custom models that integrate into microservices architecture on Azure Kubernetes Services (AKS).
The model is built by using Python and published to AKS.
You need to update the model and enable Azure Application Insights for the model.
What should you use?

A. the Azure CLI
B. MLNET Model Builder
C. the Azure Machine Learning SDK
D. the Azure portal

Answer: C
Explanation:
You can set up Azure Application Insights for Azure Machine Learning. Application Insights gives you the opportunity to monitor:
Request rates, response times, and failure rates.
Dependency rates, response times, and failure rates.
Exceptions.
Requirements include an Azure Machine Learning workspace, a local directory that contains your scripts, and the Azure Machine Learning SDK for Python installed.
References:
https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/how-to-enable-app-insights

QUESTION 107
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You are developing an application that uses an Azure Kubernetes Service (AKS) cluster.
You are troubleshooting a node issue.
You need to connect to an AKS node by using SSH.
Solution: You run the kubect1 command, and then you create an SSH connection.
Does this meet the goal?

A. Yes
B. No

Answer: B

QUESTION 108
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an app named App1 that uses the Face API.
App1 contains several PersonGroup objects.
You discover that a PersonGroup object for an individual named Ben Smith cannot accept additional entries. The PersonGroup object for Ben Smith contains 10,000 entries.
You need to ensure that additional entries can be added to the PersonGroup object for Ben Smith. The solution must ensure that Ben Smith can be identified by all the entries.
Solution: You modify the custom time interval for the training phase of App1.

Does this meet the goal?

A. Yes
B. No

Answer: B
Explanation:
Instead, use a LargePersonGroup. LargePersonGroup and LargeFaceList are collectively referred to as large-scale operations. LargePersonGroup can contain up to 1 million persons, each with a maximum of 248 faces. LargeFaceList can contain up to 1 million faces. The large-scale operations are similar to the conventional PersonGroup and FaceList but have some differences because of the new architecture.
References:
https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale

QUESTION 109
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You have an app named App1 that uses the Face API.
App1 contains several PersonGroup objects.
You discover that a PersonGroup object for an individual named Ben Smith cannot accept additional entries. The PersonGroup object for Ben Smith contains 10,000 entries.
You need to ensure that additional entries can be added to the PersonGroup object for Ben Smith. The solution must ensure that Ben Smith can be identified by all the entries.
Solution: You create a second PersonGroup object for Ben Smith.
Does this meet the goal?

A. Yes
B. No

Answer: B
Explanation:
Instead, use a LargePersonGroup. LargePersonGroup and LargeFaceList are collectively referred to as large-scale operations. LargePersonGroup can contain up to 1 million persons, each with a maximum of 248 faces. LargeFaceList can contain up to 1 million faces. The large-scale operations are similar to the conventional PersonGroup and FaceList but have some differences because of the new architecture.
References:
https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale