

Microsoft.AI-100.v2026-07-27.q119

Exam Code:	AI-100
Exam Name:	Designing and Implementing an Azure AI Solution
Certification Provider:	Microsoft
Free Question Number:	119
Version:	v2026-07-27
# of views:	108
# of Questions views:	1190
https://www.freeqas.com/qa/Microsoft/AI-100/Microsoft.AI-100.v2026-07-27.q119.html	

NEW QUESTION: 1

You are designing a Computer Vision AI application.
You need to recommend a deployment solution for the application. The solution must ensure that costs scale linearly without any upfront costs.
What should you recommend?

- A. a containerized Computer Vision API on Azure Container Instances
- B. the Computer Vision API as a single resource
- C. an Azure Container Service
- D. a containerized Computer Vision API on Azure Kubernetes Service (AKS) that has virtual nodes configured

Answer: (SHOW ANSWER)

Explanation
Containers enable you to run the Computer Vision APIs in your own environment.
Note: The host is a x64-based computer that runs the Docker container. It can be a computer on your premises or a Docker hosting service in Azure, such as:
* Azure Container Instances.
* Azure Kubernetes Service.
* A Kubernetes cluster deployed to Azure Stack.

References:
<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/computer-vision-how-to-install-conta>

NEW QUESTION: 2

Your company plans to monitor twitter hashtags, and then to build a graph of connected people and places that contains the associated sentiment.
The monitored hashtags use several languages, but the graph will be displayed in English.
You need to recommend the required Azure Cognitive Services endpoints for the planned graph.
Which Cognitive Services endpoints should you recommend?

- A. Language Detection, Content Moderator, and Key Phrase Extraction
- B. Translator Text, Content Moderator, and Key Phrase Extraction
- C. Language Detection, Sentiment Analysis, and Key Phase Extraction
- D. Translator Text, Sentiment Analysis, and Named Entity Recognition

Answer: C ([LEAVE A REPLY](#))

Sentiment analysis, which is also called opinion mining, uses social media analytics tools to determine attitudes toward a product or idea.

Translator Text: Translate text in real time across more than 60 languages, powered by the latest innovations in machine translation.

The Key Phrase Extraction skill evaluates unstructured text, and for each record, returns a list of key phrases.

This skill uses the machine learning models provided by Text Analytics in Cognitive Services.

This capability is useful if you need to quickly identify the main talking points in the record. For example, given input text "The food was delicious and there were wonderful staff", the service returns "food" and "wonderful staff".

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-entity-linking>

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-keyphrases> Analyze solution requirements Testlet 2 Overview Contoso, Ltd. has an office in New York to serve its North American customers and an office in Paris to serve its European customers.

Existing Environment

Infrastructure

Each office has a small data center that hosts Active Directory services and a few off-the-shelf software solutions used by internal users.

The network contains a single Active Directory forest that contains a single domain named contoso.com. Azure Active Directory (Azure AD) Connect is used to extend identity management to Azure.

The company has an Azure subscription. Each office has an Azure ExpressRoute connection to the subscription. The New York office connects to a virtual network hosted in the US East 2 Azure region.

The Paris office connects to a virtual network hosted in the West Europe Azure region.

The New York office has an Azure Stack Development Kit (ASDK) deployment that is used for development and testing.

Current Business Model

Contoso has a web app named Bookings hosted in an App Service Environment (ASE). The ASE is in the virtual network in the East US 2 region. Contoso employees and customers use Bookings to reserve hotel rooms.

Data Environment

Bookings connects to a Microsoft SQL Server database named hotelDB in the New York office.

The database has a view named vwAvailability that consolidates columns from the tables named Hotels, Rooms, and RoomAvailability. The database contains data that was collected during the last 20 years.

Problem Statements

Contoso identifies the following issues with its current business model:

- * European users report that access to Booking is slow, and they lose customers who must wait on the phone while they search for available rooms.
- * Users report that Bookings was unavailable during an outage in the New York data center for more than 24 hours.

Requirements

Contoso identifies the following issues with its current business model:

- * European users report that access to Bookings is slow, and they lose customers who must wait on the phone while they search for available rooms.
- * Users report that Bookings was unavailable during on outage in the New York data center for more than 24 hours.

Business Goals

Contoso wants to provide a new version of the Bookings app that will provide a highly available, reliable service for booking travel packages by interacting with a chatbot named Butler.

Contoso plans to move all production workloads to the cloud.

Technical requirements

Contoso identifies the following technical requirements:

- * Data scientists must test Butler by using ASDK.
- * Whenever possible, solutions must minimize costs.
- * Butler must greet users by name when they first connect.

- * Butler must be able to handle up to 10,000 messages a day.
- * Butler must recognize the users' intent based on basic utterances.
- * All configurations to the Azure Bot Service must be logged centrally.
- * Whenever possible, solutions must use the principle of least privilege.
- * Internal users must be able to access Butler by using Microsoft Skype for Business.
- * The new Bookings app must provide a user interface where users can interact with Butler.
- * Users in an Azure AD group named KeyManagers must be able to manage keys for all Azure Cognitive Services.
- * Butler must provide users with the ability to reserve a room, cancel a reservation, and view existing reservations.
- * The new Bookings app must be available to users in North America and Europe if a single data center or Azure region fails.
- * For continuous improvement, you must be able to test Butler by sending sample utterances and comparing the chatbot's responses to the actual intent.
- * You must maintain relationships between data after migration.

NEW QUESTION: 3

You plan to deploy the Text Analytics and Computer Vision services. The Azure Cognitive Services will be deployed to the West US and East Europe Azure regions.

You need to identify the minimum number of service endpoints and API keys required for the planned deployment.

What should you identify? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Service endpoints: ▼

1
2
4

API keys: ▼

1
2
4

Answer:

Explanation

Service endpoints: ▼

1
2
4

API keys: ▼

1
2
4

Box 1: 2

After creating a Cognitive Service resource in the Azure portal, you'll get an endpoint and a key for authenticating your applications. You can access Azure Cognitive Services through two different resources: A multi-service resource, or a single-service one.

Multi-service resource: Access multiple Azure Cognitive Services with a single key and endpoint.

Note: You need a key and endpoint for a Text Analytics resource. Azure Cognitive Services are represented by Azure resources that you subscribe to.

Each request must include your access key and an HTTP endpoint. The endpoint specifies the region you chose during sign up, the service URL, and a resource used on the request

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/cognitive-services-apis-create-account>

NEW QUESTION: 4

You are designing an AI solution that will use IoT devices to gather data from conference attendees and then analyze the data. The IoT device will connect to an Azure IoT hub.

You need to ensure that data contains no personally identifiable information before it 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.

ACTIONS

Create an Azure Stream Analytics Edge job to process data.

Create a storage container on the device.

Create an Azure Stream Analytics Cloud job.

Add the job to the IoT devices in IoT hub.

Create a storage queue on the device.

Answer Area



Answer:

Actions

Create an Azure Stream Analytics Edge job to process data.

Create a storage container on the device.

Create an Azure Stream Analytics Cloud job.

Add the job to the IoT devices in IoT hub.

Create a storage queue on the device.

Answer Area

Create a storage container on the device.

Create an Azure Stream Analytics Edge job to process data.

Add the job to the IoT devices in IoT hub.

Explanation

Create a storage container on the device.

Create an Azure Stream Analytics Edge job to process data.

Add the job to the IoT devices in IoT hub.

Note:

ASA Edge jobs run in containers deployed to Azure IoT Edge devices. They are composed of two parts:

- * A cloud part that is responsible for job definition: users define inputs, output, query, and other settings (out of order events, etc.) in the cloud.
- * A module running on your IoT devices. It contains the ASA engine and receives the job definition from the cloud.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-edge>

NEW QUESTION: 5

You have a Face API solution that updates in real time. A pilot of the solution runs successfully on a small dataset.

When you attempt to use the solution on a larger dataset that continually changes, the performance degrades, slowing how long it takes to recognize existing faces.

You need to recommend changes to reduce the time it takes to recognize existing faces without increasing costs.

What should you recommend?

- A. Change the solution to use the Computer Vision API instead of the Face API.
- B. Separate training into an independent pipeline and schedule the pipeline to run daily.
- C. Change the solution to use the Bing Image Search API instead of the Face API.

D. Distribute the face recognition inference process across many Azure Cognitive Services instances.

Answer: B ([LEAVE A REPLY](#))

Explanation/Reference:

Incorrect Answers:

A: The purpose of Computer Vision is to inspect each image associated with an incoming article to (1) scrape out written words from the image and (2) determine what types of objects are present in the image.

C: The Bing API provides an experience similar to Bing.com/search by returning search results that Bing determines are relevant to a user's query. The results include Web pages and may also include images, videos, and more.

D: That would increase cost.

References:

<https://github.com/Azure/cognitive-services>

NEW QUESTION: 6

You have several AI applications that use an Azure Kubernetes Service (AKS) cluster. The cluster supports a maximum of 32 nodes.

You discover that occasionally and unpredictably, the application requires more than 32 nodes.

You need to recommend a solution to handle the unpredictable application load.

Which scaling method should you recommend?

A. horizontal pod autoscaler

B. cluster autoscaler

C. manual scaling

D. Azure Container Instances

Answer: B ([LEAVE A REPLY](#))

To keep up with application demands in Azure Kubernetes Service (AKS), you may need to adjust the number of nodes that run your workloads. The cluster autoscaler component can watch for pods in your cluster that can't be scheduled because of resource constraints. When issues are detected, the number of nodes is increased to meet the application demand. Nodes are also regularly checked for a lack of running pods, with the number of nodes then decreased as needed. This ability to automatically scale up or down the number of nodes in your AKS cluster lets you run an efficient, cost-effective cluster.

References:

<https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler>

NEW QUESTION: 7

Which two services should be implemented so that Butler can find available rooms on the technical requirements? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. QnA Maker

B. Bing Entity Search

C. Language Understanding (LUIS)

D. Azure Search

E. Content Moderator

Answer: (SHOW ANSWER)

Explanation/Reference:

References:

<https://azure.microsoft.com/en-in/services/cognitive-services/language-understanding-intelligent-service/>

NEW QUESTION: 8

You plan to use the Microsoft Bot Framework to develop bots that will be deployed by using the Azure Bot Service. You need to configure the Azure Bot Service to support the following types of bots:

- *Bots that use Azure Functions
- *Bots that set a timer

Which template should you use for each bot type? To answer, drag the appropriate templates to the correct bot types. Each template 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.

Templates

Basic

Form

Language Understanding

Proactive

Question and Answer

Answer Area

Bots that use Azure Functions: Template

Bots that set a timer: Template



Answer:

Templates

Basic

Form

Language Understanding

Proactive

Question and Answer


Answer Area

Bots that use Azure Functions: Proactive

Bots that set a timer: Proactive



Explanation



References:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-service-concept-templates?view=azure-bot-service-3.0>

NEW QUESTION: 9

You are designing an AI solution that will provide feedback to teachers who train students over the Internet.

The students will be in classrooms located in remote areas. The solution will capture video and audio data of the students in the classrooms.

You need to recommend Azure Cognitive Services for the AI solution to meet the following requirements:

Alert teachers if a student seems angry or distracted.

Identify each student in the classrooms for attendance purposes.

Allow the teachers to log the text of conversations between themselves and the students.

Which Cognitive Services should you recommend?

- A. Computer Vision, Text Analytics, and Face API
- B. Video Indexer, Face API, and Text Analytics
- C. Computer Vision, Speech to Text, and Text Analytics
- D. Text Analytics, QnA Maker, and Computer Vision
- E. Video Indexer, Speech to Text, and Face API

Answer: E ([LEAVE A REPLY](#))

Explanation

Azure Video Indexer is a cloud application built on Azure Media Analytics, Azure Search, Cognitive Services (such as the Face API, Microsoft Translator, the Computer Vision API, and Custom Speech Service). It enables you to extract the insights from your videos using Video Indexer video and audio models.

Face API enables you to search, identify, and match faces in your private repository of up to 1 million people.

The Face API now integrates emotion recognition, returning the confidence across a set of emotions for each face in the image such as anger, contempt, disgust, fear, happiness, neutral, sadness, and surprise. These emotions are understood to be cross-culturally and universally communicated with particular facial expressions.

Speech-to-text from Azure Speech Services, also known as speech-to-text, enables real-time transcription of audio streams into text that your applications, tools, or devices can consume, display, and take action on as command input. This service is powered by the same recognition technology that Microsoft uses for Cortana and Office products, and works seamlessly with the translation and text-to-speech.

NEW QUESTION: 10

You need to build a reputation monitoring solution that reviews Twitter activity about your company. The solution must identify negative tweets and tweets that contain inappropriate images.

You plan to use Azure Logic Apps to build the solution.

Which two additional Azure services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Computer Vision

- B. Azure Blueprint
- C. Content Moderator
- D. Text Analytics
- E. Azure Machine Learning Service
- F. Form Recognizer

Answer: C,D (LEAVE A REPLY)

C: You can filter your tweets using Azure Logic Apps & Content Moderation. Azure Content Moderator is a cognitive service that checks text, image, and video content for material that is potentially offensive, risky, or otherwise undesirable. When this material is found, the service applies appropriate labels (flags) to the content.

Your app can then handle flagged content in order to comply with regulations or maintain the intended environment for users.

D: You can write an application so that when a user tweets with configured Twitter Hashtag, Logic App gets triggered and passed to Cognitive Text Analytics Connector for detecting the sentiments of the tweet (text). If the tweeted text is found to be harsh or with bad or abusive language, the tweet can be handled appropriately.

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/overview>

<https://www.c-sharpcorner.com/article/role-of-text-analytics-service-as-a-connector-in-azure-logic-apps/>

NEW QUESTION: 11

You are designing a solution that will ingest temperature data from IoT devices, calculate the average temperature, and then take action based on the aggregated data. The solution must meet the following requirements:

*Minimize the amount of uploaded data.


* Take action based on the aggregated data as quickly as possible.

What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



Answer:



Explanation

Service to use:

- Apache Hive
- Azure Data Factory
- Azure Functions
- Azure Stream Analytics

Location to deploy the job:

- A Web Job in Azure
- An Azure IoT Edge device
- Azure Event Hubs
- Azure Notification Hubs

Box 1: Azure Functions

Azure Function is a (serverless) service to host functions (little piece of code) that can be used for e. g. event driven applications.

General rule is always difficult since everything depends on your requirement but if you have to analyze a data stream, you should take a look at Azure Stream Analytics and if you want to implement something like a serverless event driven or timer-based application, you should check Azure Function or Logic Apps.

Note: Azure IoT Edge allows you to deploy complex event processing, machine learning, image recognition, and other high value AI without writing it in-house. Azure services like Azure Functions, Azure Stream Analytics, and Azure Machine Learning can all be run on-premises via Azure IoT Edge.

Box 2: An Azure IoT Edge device

Azure IoT Edge moves cloud analytics and custom business logic to devices so that your organization can focus on business insights instead of data management.

References:

<https://docs.microsoft.com/en-us/azure/iot-edge/about-iot-edge>

NEW QUESTION: 12

You are designing a solution that will ingest temperature data from IoT devices, calculate the average temperature, and then take action based on the aggregated data. The solution must meet the following requirements:

- * Minimize the amount of uploaded data.
- * Take action based on the aggregated data as quickly as possible.

What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Service to use:

- Apache Hive
- Azure Data Factory
- Azure Functions
- Azure Stream Analytics

Location to deploy the job:

- A Web job in Azure
- An Azure IoT Edge device
- Azure Event Hubs
- Azure Notification Hubs

Answer:

Service to use:

- Apache Hive
- Azure Data Factory
- Azure Functions
- Azure Stream Analytics

Location to deploy the job:

- Azure IoT Edge
- Azure IoT Edge device
- Azure Event Hubs
- Azure Notification Hubs

Explanation:

Box 1: Azure Functions

Azure Function is a (serverless) service to host functions (little piece of code) that can be used for e. g. event driven applications.

General rule is always difficult since everything depends on your requirement but if you have to analyze a data stream, you should take a look at Azure Stream Analytics and if you want to implement something like a serverless event driven or timer-based application, you should check Azure Function or Logic Apps.

Note: Azure IoT Edge allows you to deploy complex event processing, machine learning, image recognition, and other high value AI without writing it in-house. Azure services like Azure Functions, Azure Stream Analytics, and Azure Machine Learning can all be run on-premises via Azure IoT Edge.

Box 2: An Azure IoT Edge device

Azure IoT Edge moves cloud analytics and custom business logic to devices so that your organization can focus on business insights instead of data management.

References:

<https://docs.microsoft.com/en-us/azure/iot-edge/about-iot-edge>

NEW QUESTION: 13

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.

Data ingestion:

- Azure Event Hubs
- Azure IoT Hub
- Azure Queue storage

Stream processing:

- Azure HDInsight with Apache HBase
- Azure HDInsight with Apache Storm
- Azure HDInsight with Azure Machine Learning service

Answer:

Data ingestion:

- Azure Event Hubs
- Azure IoT Hub
- Azure Queue storage

Stream processing:

- Azure HDInsight with Apache HBase
- Azure HDInsight with Apache Storm
- Azure HDInsight with Azure Machine Learning service

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 grant and revoke access on device level.

Box 2: Azure HDInsight with Apache Storm

References:

<https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-compare-event-hubs>

<https://www.edureka.co/blog/introduction-to-real-time-analytics-with-apache-storm/>

NEW QUESTION: 14

You are designing a solution that uses drones to monitor remote locations for anomalies. The drones have Azure IoT Edge devices. The solution must meet the following requirements:

- *Email a user the picture and location of an anomaly when an anomaly is detected.
- *Use a video stream to detect anomalies at the location.
- *Send the pictures and location information to Azure.
- *Use the least amount of code possible.

You develop a custom vision Azure Machine Learning module to detect the anomalies.

Which service should you use for each requirement? To answer, drag the appropriate services to the correct requirements. Each service 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.

Services

Azure Functions

Azure IoT Edge

Azure IoT Hub

Azure Logic Apps

Answer Area

Use a video stream to detect anomalies at the location:

Send the pictures and location information to Azure:

Email a user the picture and location of an anomaly when an anomaly is detected:

Service

Service

Service

Answer:

Services

Azure Functions

Azure IoT Edge

Azure IoT Hub

Azure Logic Apps

Answer Area



Use a video stream to detect anomalies at the location:

Send the pictures and location information to Azure:

Email a user the picture and location of an anomaly when an anomaly is detected:

Azure IoT Edge

Azure Functions

Azure Logic Apps

Explanation



Use a video stream to detect anomalies at the location:

Send the pictures and location information to Azure:

Email a user the picture and location of an anomaly when an anomaly is detected:

Azure IoT Edge

Azure Functions

Azure Logic Apps

Box 1: Azure IOT Edge

Example:

You configure the Remote Monitoring solution to respond to anomalies detected by an IoT Edge device. IoT Edge devices let you process telemetry at the edge to reduce the volume of telemetry sent to the solution and to enable faster responses to events on devices.

Box 2: Azure Functions

Box 3: Azure Logic Apps

References:

<https://docs.microsoft.com/en-us/azure/iot-accelerators/iot-accelerators-remote-monitoring-edge>

NEW QUESTION: 15

Your company has a data team of Scala and R experts.

You plan to ingest data from multiple Apache Kafka streams.

You need to recommend a processing technology to broker messages at scale from the Kafka streams to Azure Storage.

What should you recommend?

- A. Azure Databricks
- B. Azure Functions
- C. Azure HDInsight with Apache Storm
- D. Azure HDInsight with Microsoft Machine Learning Server

Answer: C ([LEAVE A REPLY](#))

Explanation

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-streaming-at-scale-overview?toc=https%3A%2F%2F>

NEW QUESTION: 16

You plan to deploy Azure IoT Edge devices. Each device will store more than 10,000 images locally. Each image is approximately 5 MB.

You need to ensure that the images persist on the devices for 14 days.

What should you use?

- A. Azure Stream Analytics on the IoT Edge devices
- B. Azure Database for Postgres SQL
- C. Azure Blob storage on the IoT Edge devices
- D. Microsoft SQL Server on the IoT Edge devices

Answer: C ([LEAVE A REPLY](#))

Explanation

Azure Blob Storage on IoT Edge provides a block blob and append blob storage solution at the edge. A blob storage module on your IoT Edge device behaves like an Azure blob service, except the blobs are stored locally on your IoT Edge device.

This is useful where data needs to be stored locally until it can be processed or transferred to the cloud. This data can be videos, images, finance data, hospital data, or any other unstructured data.

References:

<https://docs.microsoft.com/en-us/azure/iot-edge/how-to-store-data-blob>

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (163 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 17

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 in this section, 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 kubectl command, and then you create an SSH connection.

Does this meet the goal?

- A. No

B. Yes

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 18

You have Azure IoT Edge devices that collect measurements every 30 seconds.

You plan to send the measurements to an Azure IoT hub.

You need to process events in the cloud.

What should you use?

A. Apache Kafka

B. Azure Stream Analytics record functions

C. Azure Stream Analytics windowing functions

D. Azure Machine Learning on the IoT Edge devices

Answer: D ([LEAVE A REPLY](#))

Use Azure Notebooks to develop a machine learning module and deploy it to a Linux device running Azure IoT Edge.

You can use IoT Edge modules to deploy code that implements your business logic directly to your IoT Edge devices.

Reference:

<https://docs.microsoft.com/en-us/azure/iot-edge/tutorial-deploy-machine-learning>

NEW QUESTION: 19

You need to design an application that will analyze real-time data from financial feeds. The data will be ingested into Azure IoT Hub. The data must be processed as quickly as possible in the order in which it is ingested.

Which service should you include in the design?

A. Azure Event Hubs

B. Azure Data Factory

C. Azure Stream Analytics

D. Apache Kafka

Answer: ([SHOW ANSWER](#))

Topic 1, Contoso Case Study: 1

Overview

Contoso. Ltd. has an office in New York to serve its North American customers and an office in Paris to serve its European customers.

Existing Environment

Contoso. Ltd. has an office in New York to serve its North American customers and an office in Paris to serve its European customers.

Infrastructure

Each office has a small data center that hosts Active Directory services and a few off-the-shelf software solutions used by internal users.

The network contains a single Active Directory forest that contains a single domain named contoso.com.

Azure Active Directory (Azure AD) Connect is used to extend identity management to Azure.

The company has an Azure subscription. Each office has an Azure ExpressRoute connection to the subscription. The New York office connects to a virtual network hosted in the US East 2 Azure region.

The Paris office connects to a virtual network hosted in the West Europe Azure region.

The New York office has an Azure Stack Development Kit (ASDK) deployment that is used for development and testing.

Current Business Model

Contoso has a web app named Bookings hosted in an App Service Environment (ASE). The ASE is in the virtual network in the East US 2 region. Contoso employees and customers use Bookings to reserve hotel rooms.

Data Environment

Bookings connects to a Microsoft SQL Server database named hotelDB in the New York office.

The database has a view named vwAvailability that consolidates columns from three tables named Hotels, Rooms, and RoomAvailability. The database contains data that was collected during the last 20 years.

Problem Statements

Contoso identifies the following issues with its current business model:

- * European users report that access to Bookings is slow, and they lose customers who must wait on the phone while they search for available rooms.
- * Users report that Bookings was unavailable during an outage in the New York data center for more than 24 hours.

Requirements

Contoso identifies the following issues with its current business model:

- *European users report that access to Bookings is slow, and they lose customers who must wait on the phone while they search for available rooms.

Business Goals

- *Users report that Bookings was unavailable during an outage in the New York data center for more than 24 hours.

Contoso wants to provide a new version of the Bookings app that will provide a highly available, reliable service for booking travel packages by interacting with a chatbot named Butler.

Contoso plans to move all production workloads to the cloud.

Technical Requirements

Contoso identifies the following technical requirements:

- *Data scientists must test Butler by using ASDK.
- *Whenever possible, solutions must minimize costs.
- *Butler must greet users by name when they first connect.
- *Butler must be able to handle up to 10,000 messages a day.
- *Butler must recognize the users' intent based on basic utterances.
- *All configurations to the Azure Bot Service must be logged centrally.
- *Whenever possible, solutions must use the principle of least privilege.
- *Internal users must be able to access Butler by using Microsoft Skype for Business.
- *The new Bookings app must provide a user interface where users can interact with Butler.
- *Users in an Azure AD group named KeyManagers must be able to manage keys for all Azure Cognitive Services.
- *Butler must provide users with the ability to reserve a room, cancel a reservation, and view existing reservations.
- *The new Bookings app must be available to users in North America and Europe if a single data center or Azure region fails.
- *For continuous improvement, you must be able to test Butler by sending sample utterances and comparing the chatbot's responses to the actual intent.

NEW QUESTION: 20

You are designing an Azure infrastructure to support an Azure Machine Learning solution that will have multiple phases. The solution must meet the following requirements:

- * Securely query an on-premises database once a week to update product lists.
- * Access the data without using a gateway.
- * Orchestrate the separate phases.

What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



Answer:



Explanation

To connect to the on-premises data:

- A point-to-site VPN connection
- A site-to-site VPN connection
- Azure App Service Hybrid Connections

To orchestrate the phases:



- A Machine Learning experiment
- Azure Machine Learning Studio
- Machine Learning pipelines

To control the orchestration:

- Azure Automation
- Azure Databricks
- Azure Notebooks

Box 1: Azure App Service Hybrid Connections

With Hybrid Connections, Azure websites and mobile services can access on-premises resources as if they were located on the same private network. Application admins thus have the flexibility to simply lift-and-shift specific most front-end tiers to Azure with minimal configuration changes, extending their enterprise apps for hybrid scenarios.

Incorrect answer: The VPN connection solution both use gateways.

Box 2: Machine Learning pipelines

Typically when running machine learning algorithms, it involves a sequence of tasks including pre-processing, feature extraction, model fitting, and validation stages. For example, when classifying text documents might involve text segmentation and cleaning, extracting features, and training a classification model with cross-validation. Though there are many libraries we can use for each stage,

connecting the dots is not as easy as it may look, especially with large-scale datasets. Most ML libraries are not designed for distributed computation or they do not provide native support for pipeline creation and tuning.

Box 3: Azure Databricks

References:

<https://azure.microsoft.com/is-is/blog/hybrid-connections-preview/>

<https://databricks.com/glossary/what-are-ml-pipelines>

NEW QUESTION: 21

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 Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models.

Once an

anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You expose a Machine Learning model as an Azure web service.

Does this meet the goal?

A. Yes

B. No

Answer: B ([LEAVE A REPLY](#))

Instead use Azure Stream Analytics and REST API.

Note. Available in both the cloud and Azure IoT Edge, Azure Stream Analytics offers built-in machine learning based anomaly detection capabilities that can be used to monitor the two most commonly occurring anomalies:

temporary and persistent.

Stream Analytics supports user-defined functions, via REST API, that call out to Azure Machine Learning endpoints.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-machine-learning-anomaly-detection>

NEW QUESTION: 22

Your company recently purchased several hundred hardware devices that contains sensors.

You need to recommend a solution to process the sensor data. The solution must provide the ability to write back configuration changes to the devices.

What should you include in the recommendation?

A. Microsoft Azure IoT Hub

B. API apps in Microsoft Azure App Service

C. Microsoft Azure Event Hubs

D. Microsoft Azure Notification Hubs

Answer: A ([LEAVE A REPLY](#))

Explanation

References:

<https://azure.microsoft.com/en-us/resources/samples/functions-js-iot-hub-processing/>

NEW QUESTION: 23


You are building an Azure Analysis Services cube for your AI deployment. The source data for the cube is located in an on premises network in a Microsoft SQL Server database. You need to ensure that the Azure Analysis Services service can access the source data. What should you deploy to your Azure subscription?

- A. a site-to-site VPN
- B. a network gateway
- C. a data gateway
- D. Azure Data Factory

Answer: C ([LEAVE A REPLY](#))

Explanation

From April 2017 onward we can use On-premises Data Gateway for Azure Analysis Services. This means you can connect your Tabular Models hosted in Azure Analysis Services to your on-premises data sources through On-premises Data Gateway.



References:

<https://biinsight.com/on-premises-data-gateway-for-azure-analysis-services/>

NEW QUESTION: 24

You are designing a solution that will ingest data from an Azure IoT Edge device, preprocess the data in Azure Machine Learning, and then move the data to Azure HDInsight for further processing. What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Machine Learning module to use to move the data into HDInsight:

Export Data
Load Trained Model
Partition and Sample
Unpack Zipped Datasets

Query type to use:

Apache Hive
Apache Spark
C#
Transact-SQL


Output the data to:

Azure Cosmos DB
Azure Data Lake
Azure Table storage
HDFS

Answer:



Explanation



Box 1: Export Data

The Export data to Hive option in the Export Data module in Azure Machine Learning Studio. This option is useful when you are working with very large datasets, and want to save your machine learning experiment data to a Hadoop cluster or HDInsight distributed storage.

Box 2: Apache Hive

Apache Hive is a data warehouse system for Apache Hadoop. Hive enables data summarization, querying, and analysis of data. Hive queries are written in HiveQL, which is a query language similar to SQL.

Box 3: Azure Data Lake

Default storage for the HDFS file system of HDInsight clusters can be associated with either an Azure Storage account or an Azure Data Lake Storage.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/export-to-hive-query>

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/hdinsight-use-hive>

NEW QUESTION: 25

Your company develops an API application that is orchestrated by using Kubernetes.

You need to deploy the application.

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

NOTE: Each correct selection is worth one point.

A. Create a Kubernetes cluster.

- B. Create an Azure Container Registry instance.
- C. Create a container image file.
- D. Create a Web App for Containers.
- E. Create an Azure container instance.

Answer: A,B,C (LEAVE A REPLY)

Explanation/Reference:

References:

<https://docs.microsoft.com/en-us/azure/aks/tutorial-kubernetes-prepare-app>

NEW QUESTION: 26

You are designing a real-time speech-to-text AI feature for an Android mobile app. The feature will stream data to the Speech service.

You need to recommend which audio format to use to serialize the audio. The solution must minimize the amount of data transferred to the cloud.

What should you recommend?

- A. MP3
- B. WAV/PCM
- C. MP4a


Answer: A (LEAVE A REPLY)

NEW QUESTION: 27

You need to build a sentiment analysis solution that will use input data from JSON documents and PDF documents. The JSON documents must be processed in batches and aggregated.

Which storage type should you use for each file type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



Answer:



Explanation

JSON documents:

	▼
Azure Blob storage	
Azure Cosmos DB	
Azure Data Lake	
Azure Table storage	

Microsoft
PDF documents:

	▼
Azure Blob storage	
Azure Cosmos DB	
Azure Data Lake	
Azure Table storage	

References:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/batch-processing>

NEW QUESTION: 28

You need to integrate the new Bookings app and the Butler chatbot.

Which four 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.

Actions

On the page where you want Butler to be used, paste the embed code of the new Bookings app.

From the Channels settings of Butler, copy the secret key and embed code for the Skype for Business channel.

Exchange the secret for a token by connecting to <https://directline.botframework.com/api/tokens>.

From the Channels settings of Butler, retrieve the secret key and embed code for the Web Chat channel.

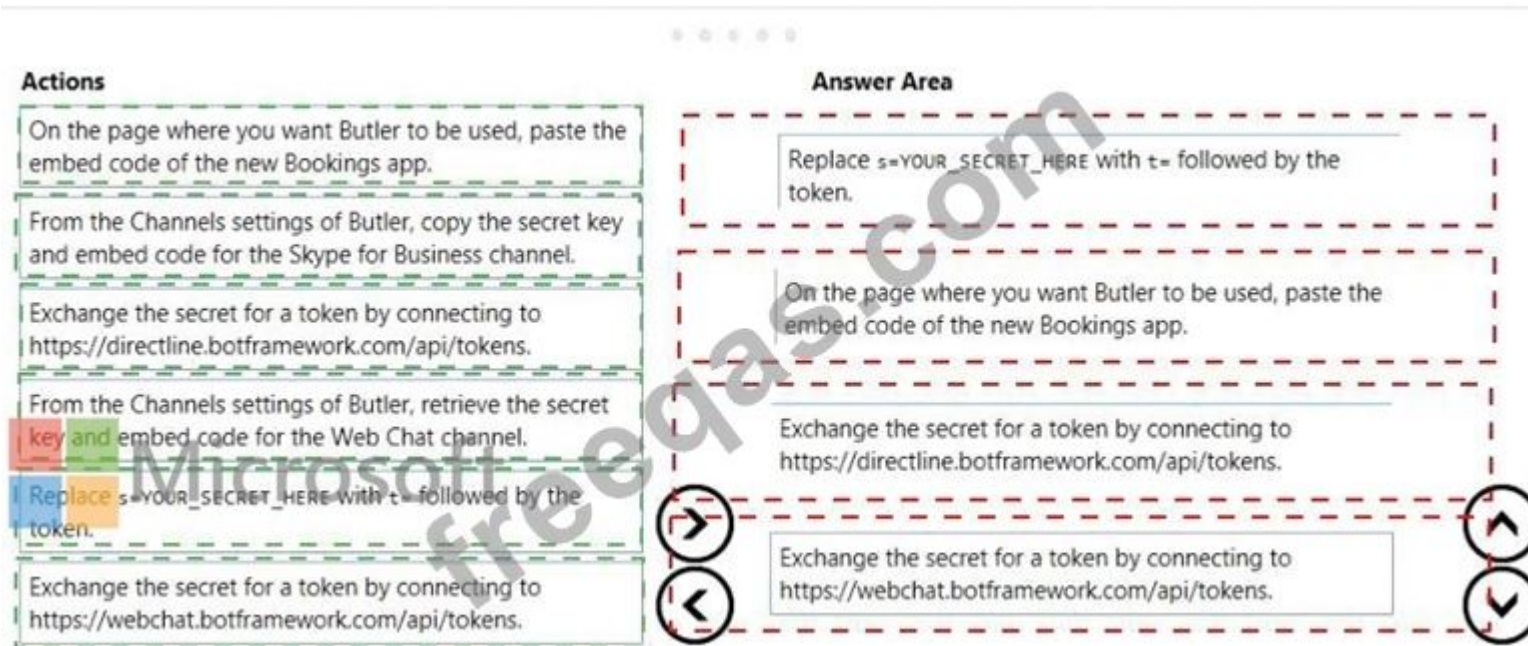
Replace `s=YOUR_SECRET_HERE` with `t=` followed by the token.

Exchange the secret for a token by connecting to <https://webchat.botframework.com/api/tokens>.

Answer Area



Answer:



NEW QUESTION: 29

You are configuring data persistence for a Microsoft Bot Framework application. The application requires a structured NoSQL cloud data store.

You need to identify a storage solution for the application. The solution must minimize costs.

What should you identify?

- A. Azure Blob storage
- B. Azure Cosmos DB
- C. Azure HDInsight
- D. Azure Table storage

Answer: D (LEAVE A REPLY)

Table Storage is a NoSQL key-value store for rapid development using massive semi-structured datasets You can develop applications on Cosmos DB using popular NoSQL APIs.

Both services have a different scenario and pricing model.

While Azure Storage Tables is aimed at high capacity on a single region (optional secondary read only region but no failover), indexing by PK/RK and storage-optimized pricing; Azure Cosmos DB

Tables aims for high throughput (single-digit millisecond latency), global distribution (multiple failover), SLA-backed predictive performance with automatic indexing of each attribute/property and a pricing model focused on throughput.

References:

<https://db-engines.com/en/system/Microsoft+Azure+Cosmos+DB%3BMicrosoft+Azure+Table+Storage>

NEW QUESTION: 30

You need to recommend a data storage solution that meets the technical requirements.

What is the best data storage solution to recommend? More than one answer choice may achieve the goal.

Select the BEST answer.

- A. Azure Databricks
- B. Azure Cosmos DB
- C. Azure Table storage
- D. Azure SQL Database

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 31

You deploy an application that performs sentiment analysis on the data stored in Azure Cosmos DB. Recently, you loaded a large amount of data to the database. The data was for a customer named Contoso, Ltd. You discover that queries for the Contoso data are slow to complete, and the queries slow the entire application. You need to reduce the amount of time it takes for the queries to complete. The solution must minimize costs. What is the best way to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Change the request units.
- B. Change the partitioning strategy.
- C. Change the transaction isolation level.
- D. Migrate the data to the Cosmos DB database.

Answer: B ([LEAVE A REPLY](#))

Throughput provisioned for a container is divided evenly among physical partitions.

Incorrect:

Not A: Increasing request units would also improve throughput, but at a cost.

Reference:

<https://docs.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning>

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (**163** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 32

You are developing a mobile application that will perform optical character recognition (OCR) from photos. The application will annotate the photos by using metadata, store the photos in Azure Blob storage, and then score the photos by using an Azure Machine Learning model. What should you use to process the data?

- A. Azure Event Hubs
- B. Azure Functions
- C. Azure Stream Analytics
- D. Azure Logic Apps
- E. Azure Batch AI

Answer: (SHOW ANSWER)

By using Azure services such as the Computer Vision API and Azure Functions, companies can eliminate the need to manage individual servers, while reducing costs and leveraging the expertise that Microsoft has already developed around processing images with Cognitive Services.

Incorrect:

Not E: The Azure Batch AI service was retired in 2019 and was replaced with Azure Machine Learning Compute.

References: <https://docs.microsoft.com/en-us/azure/architecture/example-scenario/ai/intelligent-apps-image-processing>

NEW QUESTION: 33

You create an image classification model in Azure Machine Learning Studio.

You need to deploy the model as a containerized web service.

Which four 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.

The screenshot shows the Azure Machine Learning Studio interface. On the left, under the heading "Actions", there is a list of six actions in a scrollable container: "Start the container.", "Register the container image.", "Get the http endpoint of the web service.", "Train the model.", "Create a container image.", and "Create an Azure Batch AI account.". To the right of this list is an "Answer Area" which is currently empty. A Microsoft logo is visible in the top right corner of the interface. Navigation arrows (left and right) are located at the bottom of the actions list, and up and down arrows are at the bottom of the answer area.

Answer:

The screenshot shows the same Azure Machine Learning Studio interface, but now the "Answer Area" contains four actions in a specific sequence, indicated by a red dashed border. The actions are: "Train the model.", "Create a container image.", "Register the container image.", and "Get the http endpoint of the web service.". The actions in the "Actions" list on the left are still the same. The Microsoft logo is visible in the top right corner. Navigation arrows are present at the bottom of both the actions list and the answer area.

Explanation



NEW QUESTION: 34

You deploy an application that performs sentiment analysis on the data stored in Azure Cosmos DB. Recently, you loaded a large amount of data to the database. The data was for a customer named Contoso. Ltd. You discover that queries for the Contoso data are slow to complete, and the queries slow the entire application. You need to reduce the amount of time it takes for the queries to complete. The solution must minimize costs. What is the best way to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

- A. Change the partitioning strategy.
- B. Migrate the data to the Cosmos DB database.
- C. Change the transaction isolation level.
- D. Change the requests units.

Answer: C,D ([LEAVE A REPLY](#))

NEW QUESTION: 35

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 ([LEAVE A REPLY](#))

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>

NEW QUESTION: 36

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 migrate all the entries to the LargePersonGroup object for Ben Smith.

Does this meet the goal?

A. Yes

B. No

Answer: A ([LEAVE A REPLY](#))

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>

NEW QUESTION: 37

You need to build a solution to monitor Twitter. The solution must meet the following requirements:

* Send an email message to the marketing department when negative Twitter messages are detected.

* Run sentiment analysis on Twitter messages that mention specific tags.

*Use the least amount of custom code possible.

Which two services should you include in the solution? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Cognitive Services

B. Azure Logic Apps

C. Azure Stream Analytics

D. Azure Databricks

E. Azure Functions

Answer: A,D ([LEAVE A REPLY](#))

NEW QUESTION: 38

You are designing an AI solution that will analyze millions of pictures.

You need to recommend a solution for storing the pictures. The solution must minimize costs.

Which storage solution should you recommend?

A. an Azure Data Lake store

- B. Azure File Storage
- C. Azure Blob storage
- D. Azure Table storage

Answer: C ([LEAVE A REPLY](#))

Explanation/Reference:

Explanation:

Data Lake will be a bit more expensive although they are in close range of each other. Blob storage has more options for pricing depending upon things like how frequently you need to access your data (cold vs hot storage).

References:

<http://blog.pragmaticworks.com/azure-data-lake-vs-azure-blob-storage-in-data-warehousing>

NEW QUESTION: 39

Your company plans to deploy several apps that will use Azure Cognitive Services APIs.

You need to recommend which Cognitive Services APIs must be used to meet the following requirements:

- * Must be able to identify inappropriate text and profanities in multiple languages.
- * Must be able to interpret user requests sent by using text input.
- * Must be able to identify named entities in text.

Which API should you recommend for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Must be able to identify inappropriate text and profanities in multiple languages:

	▼
Text Analytics	
Content Moderator	
Bing Visual Search	
Language Understanding (LUIS)	

Must be able to interpret user requests sent by using text input:

	▼
Text Analytics	
Content Moderator	
Bing Visual Search	
Language Understanding (LUIS)	



Must be able to identify named entities in text:

	▼
Text Analytics	
Content Moderator	
Bing Visual Search	
Language Understanding (LUIS)	

Answer:

Microsoft

Must be able to identify inappropriate text and profanities in multiple languages:

- Text Analytics
- Content Moderator
- Bing Visual Search
- Language Understanding (LUIS)

Must be able to interpret user requests sent by using text input:

- Text Analytics
- Content Moderator
- Bing Visual Search
- Language Understanding (LUIS)

Must be able to identify named entities in text:

- Text Analytics
- Content Moderator
- Bing Visual Search
- Language Understanding (LUIS)

Explanation:

Box 1: Content Moderator

The Azure Content Moderator API is a cognitive service that checks text, image, and video content for material that is potentially offensive, risky, or otherwise undesirable. When such material is found, the service applies appropriate labels (flags) to the content. Your app can then handle flagged content in order to comply with regulations or maintain the intended environment for users.

Box 2: Language Understanding (LUIS)

Designed to identify valuable information in conversations, LUIS interprets user goals (intents) and distills valuable information from sentences (entities), for a high quality, nuanced language model. LUIS integrates seamlessly with the Azure Bot Service, making it easy to create a sophisticated bot.

Box 3: Text Analytics

The Text Analytics API is a cloud-based service that provides advanced natural language processing over raw text, and includes four main functions: sentiment analysis, key phrase extraction, named entity recognition, and language detection.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/cognitive-services/content-moderator/overview>

<https://www.luis.ai/home>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/>

NEW QUESTION: 40

You plan to use the Microsoft Bot Framework to develop bots that will be deployed by using the Azure Bot Service.

You need to configure the Azure Bot Service to support the following types of bots:

* Bots that use Azure Functions

* Bots that set a timer

Which template should you use for each bot type? To answer, drag the appropriate templates to the correct bot types. Each template 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.

Templates

- Basic
- Form
- Language Understanding
- Proactive
- Question and Answer

Answer Area

Bots that use Azure Functions: Template

Bots that set a timer: Template

Answer:

Templates

- Basic
- Form
- Language Understanding
- Proactive
- Question and Answer

Answer Area

Bots that use Azure Functions: Proactive

Bots that set a timer: Proactive

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-service-concept-templates?view=azure-bot-service-3.0>

NEW QUESTION: 41

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.

Dynamically ask questions based on an uploaded image:

	▼
Azure Analysis Services	
Azure Bot Service	
Azure Data Factory	
Azure Linguistic Analysis API	

Analyze and classify an image:

	▼
Bing Image Search	
Bing Visual Search	
Computer Vision	
Video Indexer	

Answer:

Dynamically ask questions based on an uploaded image:

	▼
Azure Analysis Services	
Azure Bot Service	
Azure Data Factory	
Azure Linguistic Analysis API	

Analyze and classify an image:

	▼
Bing Image Search	
Bing Visual Search	
Computer Vision	
Video Indexer	

Explanation

Dynamically ask questions based on an uploaded image:

Analyze and classify an image:

Microsoft

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/>

NEW QUESTION: 42

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.

Location to store the metadata:

Azure Blob storage
Azure Cosmos DB
Azure File Storage

Virtual machine series to run the vision module:

A
F
NV

Virtual machine series to run the mathematical module:

A
F
NV

Answer:

Location to store the metadata:

Azure Blob storage
Azure Cosmos DB
Azure File Storage

Virtual machine series to run the vision module:

A
F
NV

Virtual machine series to run the mathematical module:

A
F
NV

Explanation

Location to store the metadata: ▼
 Azure Blob storage
 Azure Cosmos DB
 Azure File Storage

Virtual machine series to run the vision module: ▼
 A
 F
 NV

Virtual machine series to run the mathematical module: ▼
 A
 F
 NV

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/>

NEW QUESTION: 43

You plan to build an app that will provide users with the ability to dictate messages and convert the messages into text. You need to recommend a solution to meet the following requirements for the app:

- * Must be able to transcribe streaming dictated messages that are longer than 15 seconds.
- * Must be able to upload existing recordings to Azure Blob storage to be transcribed later.

Which solution should you recommend for each requirement? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Must be able to transcribe streaming dictated messages that are longer than 15 seconds:

- The Speech SDK
- The Speech to Text API
- The Batch Transcription API

Must be able to upload existing recordings to Blob storage to be transcribed later:

- The Speech SDK
- The Speech to Text API
- The Batch Transcription API

Answer:

Answer Area

Must be able to transcribe streaming dictated messages that are longer than 15 seconds:

The Speech SDK
The Speech to Text API
The Batch Transcription API

Must be able to upload existing recordings to Blob storage to be transcribed later:

The Speech SDK
The Speech to Text API
The Batch Transcription API

Explanation

Must be able to transcribe streaming dictated messages that are longer than 15 seconds:

The Speech SDK
The Speech to Text API
The Batch Transcription API

Must be able to upload existing recordings to Blob storage to be transcribed later:

The Speech SDK
The Speech to Text API
The Batch Transcription API

Box 1: The Speech SDK

The Speech SDK is not limited to 15 seconds.

Box 2: Batch Transcription API

Batch transcription is a set of REST API operations that enables you to transcribe a large amount of audio in storage. You can point to audio files with a shared access signature (SAS) URI and asynchronously receive transcription results. With the new v3.0 API, you have the choice of transcribing one or more audio files, or process a whole storage container.

Asynchronous speech-to-text transcription is just one of the features.

Reference:

<https://github.com/Azure-Samples/cognitive-services-speech-sdk/issues/13>

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/batch-transcription>

NEW QUESTION: 44

You plan to create a bot that will support five languages. The bot will be used by users located in three different countries. The bot will answer common customer questions. The bot will use Language Understanding (LUIS) to identify which skill to use and to detect the language of the customer.

You need to identify the minimum number of Azure resources that must be created for the planned bot.

How many QnA Maker, LUIS and Language Detection instances should you create? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

QnA Maker:  Microsoft ▼

1
3
5

LUIS: ▼

1
3
5

Language Detection: ▼

1
3
5

Answer:

QnA Maker:  Microsoft ▼

1
3
5

LUIS: ▼

1
3
5

Language Detection: ▼

1
3
5

 Microsoft

Explanation

The screenshot shows three configuration sections in the Azure portal:

- QnA Maker:** A dropdown menu with options 1, 3, and 5. The option '5' is selected and highlighted in dark grey.
- LUIS:** A dropdown menu with options 1, 3, and 5. The option '5' is selected and highlighted in dark grey.
- Language Detection:** A dropdown menu with options 1, 3, and 5. The option '1' is selected and highlighted in dark grey.

QnA Maker: 5

If the user plans to support multiple languages, they need to have a new QnA Maker resource for each language.

LUIS: 5

If you need a multi-language LUIS client application such as a chatbot, you have a few options. If LUIS supports all the languages, you develop a LUIS app for each language. Each LUIS app has a unique app ID, and endpoint log. If you need to provide language understanding for a language LUIS does not support, you can use Microsoft Translator API to translate the utterance into a supported language, submit the utterance to the LUIS endpoint, and receive the resulting scores.

Language detection: 1

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis. This capability is useful for content stores that collect arbitrary text, where language is unknown. You can parse the results of this analysis to determine which language is used in the input document. The response also returns a score that reflects the confidence of the model. The score value is between 0 and 1.

The Language Detection feature can detect a wide range of languages, variants, dialects, and some regional or cultural languages. The exact list of languages for this feature isn't published.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/language-support>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-language-support>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-to/text-analytics-how-to-language>

NEW QUESTION: 45

You have an Azure Machine Learning model that is deployed to a web service.

You plan to publish the web service by using the name ml.contoso.com.

You need to recommend a solution to ensure that access to the web service is encrypted.

Which three actions should you recommend? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Generate a shared access signature (SAS)
- B. Obtain an SSL certificate
- C. Add a deployment slot
- D. Update the web service
- E. Update DNS
- F. Create an Azure Key Vault

Answer: B,D,E (LEAVE A REPLY)

The process of securing a new web service or an existing one is as follows:

1. Get a domain name.
2. Get a digital certificate.
3. Deploy or update the web service with the SSL setting enabled.
4. Update your DNS to point to the web service.

Note: To deploy (or re-deploy) the service with SSL enabled, set the ssl_enabled parameter to True, wherever applicable. Set the ssl_certificate parameter to the value of the certificate file and the ssl_key to the value of the key file.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-secure-web-service>

NEW QUESTION: 46

You create an image classification model in Azure Machine Learning Studio.

You need to deploy the model as a containerized web service.

Which four 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.

Actions

Start the container.

Register the container image.

Get the http endpoint of the web service.


Train the model.

Create a container image.

Create an Azure Batch AI account.

Answer Area

Answer:



Explanation



The following diagram illustrates the complete deployment workflow:

The deployment workflow includes the following steps:

- * Register the model in a registry hosted in your Azure Machine Learning Service workspace
- * Register an image that pairs a model with a scoring script and dependencies in a portable container
- * Deploy the image as a web service in the cloud or to edge devices
- * Monitor and collect data
- * Update a deployment to use a new image.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/concept-model-management-and-deploym>

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (163 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 47

You are designing an AI solution that will analyze media data. The data will be stored in Azure Blob storage.
You need to ensure that the storage account is encrypted by using a key generated by the hardware security module (HSM) of your company.
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.

.....

Actions

Enable encryption that uses customer-managed keys.

Upload a key to an Azure key vault.

Generate an encryption key.

Generate an access key.

Configure a service endpoint for the storage account.

Generate a shared access signature (SAS).

Answer Area

➤

➡

⬅

⬅

⬆

⬇

Answer:

.....

Actions

Enable encryption that uses customer-managed keys.

Upload a key to an Azure key vault.

Generate an encryption key.

Generate an access key.

Configure a service endpoint for the storage account.

Generate a shared access signature (SAS).

Answer Area

Configure a service endpoint for the storage account.

Generate an encryption key.

Enable encryption that uses customer-managed keys.

⬅

⬅

⬆

⬇

Explanation



References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-encryption-keys-portal>

<https://docs.microsoft.com/en-us/azure/key-vault/key-vault-hsm-protected-keys>

NEW QUESTION: 48

Your company plans to monitor Twitter hashtags, and then to build a graph of connected people and places that contains the associated sentiment.

The monitored hashtags use several languages, but the graph will be displayed in English.

You need to recommend the required Azure Cognitive Services endpoints (or the planned graph). Which Cognitive Services endpoints should you recommend?

- A. Translator Text, Sentiment Analysis, and Named Entity Recognition
- B. Language Detection, Sentiment Analysis, and Key Phrase Extraction
- C. Language Detection, Content Moderator, and Key Phrase Extraction
- D. Translator Text, Content Moderator, and Key Phrase Extraction

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 49

Your company has a data team of Transact-SQL experts.

You plan to ingest data from multiple sources into Azure Event Hubs.

You need to recommend which technology the data team should use to move and query data from Event Hubs to Azure Storage. The solution must leverage the data team's existing skills.

What is the best recommendation to achieve the goal? More than one answer choice may achieve the goal.

- A. Azure Notification Hubs
- B. Azure Event Grid
- C. Apache Kafka streams
- D. Azure Stream Analytics


Answer: B ([LEAVE A REPLY](#))

Explanation

Event Hubs Capture is the easiest way to automatically deliver streamed data in Event Hubs to an Azure Blob storage or Azure Data Lake store. You can subsequently process and deliver the data to any other storage destinations of your choice, such as SQL Data Warehouse or Cosmos DB.

You to capture data from your event hub into a SQL data warehouse by using an Azure function triggered by an event grid.

Example:



First, you create an event hub with the Capture feature enabled and set an Azure blob storage as the destination. Data generated by WindTurbineGenerator is streamed into the event hub and is automatically captured into Azure Storage as Avro files.

Next, you create an Azure Event Grid subscription with the Event Hubs namespace as its source and the Azure Function endpoint as its destination.

Whenever a new Avro file is delivered to the Azure Storage blob by the Event Hubs Capture feature, Event Grid notifies the Azure Function with the blob URI. The Function then migrates data from the blob to a SQL data warehouse.

References:

<https://docs.microsoft.com/en-us/azure/event-hubs/store-captured-data-data-warehouse>

NEW QUESTION: 50

DRAG DROP

You have an intelligent edge solution that processes data and outputs the data to an Azure Cosmos DB account that uses the SQL API.

You need to ensure that you can perform full text searches of the data.

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 Area

Actions

Create a cognitive search pipeline.

Create an Azure Search index.

Create an Azure Service Bus.

Create a data source.

Create an Azure Search indexer.

Microsoft

FreeDas.com

Answer:

Actions

- Create a cognitive search pipeline.
- Create an Azure Search index.
- Create an Azure Service Bus.
- Create a data source.
- Create an Azure Search indexer.

Answer Area

- Create an Azure Service Bus.
- Create a cognitive search pipeline.
- Create a data source.

NEW QUESTION: 51

You need to build an A) solution that will be shared between several developers and customers.

You plan to write code, host code, and document the runtime all within a single user experience.

You build the environment to host the solution.

Which three actions should you perform in sequence next? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.


Actions

- Build an experiment.
- Create an Azure Machine Learning Studio workspace.
- Create stream inputs and outputs.
- Implement Transact-SQL for the stream query.
- Create an Azure Stream Analytics service.
- Create a new experiment.
- Create a new notebook.
- Build a notebook.

Answer Area

-
-
-

Answer:



Explanation:

Step 1: Create an Azure Machine Learning Studio workspace

Step 2: Create a notebook

You can manage notebooks using the UI, the CLI, and by invoking the Workspace API.

To create a notebook

Click the Workspace button Workspace Icon or the Home button Home Icon in the sidebar. Do one of the following:

Next to any folder, click the Menu Dropdown on the right side of the text and select Create > Notebook. Create Notebook In the Workspace or a user folder, click Down Caret and select Create > Notebook.

2. In the Create Notebook dialog, enter a name and select the notebook's primary language.

3. If there are running clusters, the Cluster drop-down displays. Select the cluster to attach the notebook to.

4. Click Create.

Step 3: Create a new experiment

Create a new experiment by clicking +NEW at the bottom of the Machine Learning Studio window. Select EXPERIMENT > Blank Experiment.

References:

<https://docs.azuredatabricks.net/user-guide/notebooks/notebook-manage.html>

<https://docs.microsoft.com/en-us/azure/machine-learning/service/quickstart-run-cloud-notebook>

NEW QUESTION: 52

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 in this section, 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 create a managed identity for AKS, and then you create an SSH connection.

Does this meet the goal?

A. Yes

B. No

Answer: B ([LEAVE A REPLY](#))

Instead add an SSH key to the node, and then you create an SSH connection.

References:

<https://docs.microsoft.com/en-us/azure/aks/ssh>

NEW QUESTION: 53

Your company has factories in 10 countries. Each factory contains several thousand IoT devices.

The devices present status and trending data on a dashboard.

You need to ingest the data from the IoT devices into a data warehouse.

Which two Microsoft Azure technologies should you use? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A.** Azure Stream Analytics
- B.** Azure Data Factory
- C.** an Azure HDInsight cluster
- D.** Azure Batch
- E.** Azure Data Lake

Answer: A,E ([LEAVE A REPLY](#))

Explanation

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-define-outputs>

NEW QUESTION: 54

You are designing an AI solution that will use IoT devices to gather data from conference attendees and then analyze the data. The IoT device will connect to an Azure IoT hub.

You need to ensure that data contains no personally identifiable information before it 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.

Actions

Create an Azure Stream Analytics Edge job to process data.

Create a storage container on the device.

Create an Azure Stream Analytics Cloud job.

Add the job to the IoT devices in IoT hub.

Create a storage queue on the device.

Answer Area

←

→

↑

↓

 Microsoft

Answer:

Action

Create an Azure Stream Analytics Edge job to process data.

Create a storage container on the device.

Create an Azure Stream Analytics Cloud job.

Add the job to the IoT devices in IoT hub.

Create a storage queue on the device.


Answer Area

←

→

↑

↓

 Microsoft

Create a storage container on the device.

Create an Azure Stream Analytics Edge job to process data.

Add the job to the IoT devices in IoT hub.

Explanation



Note:

ASA Edge jobs run in containers deployed to Azure IoT Edge devices. They are composed of two parts:

- * A cloud part that is responsible for job definition: users define inputs, output, query, and other settings (out of order events, etc.) in the cloud.
- * A module running on your IoT devices. It contains the ASA engine and receives the job definition from the cloud.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-edge>

NEW QUESTION: 55

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 Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy Azure Functions as an IoT Edge module.

Does this meet the goal?

A. Yes

B. No

Answer: B ([LEAVE A REPLY](#))

Instead use Azure Stream Analytics and REST API.

Note. Available in both the cloud and Azure IoT Edge, Azure Stream Analytics offers built-in machine learning based anomaly detection capabilities that can be used to monitor the two most commonly occurring anomalies:

temporary and persistent.

Stream Analytics supports user-defined functions, via REST API, that call out to Azure Machine Learning endpoints.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-machine-learning-anomaly-detection>

NEW QUESTION: 56

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 Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy Azure Stream Analytics as an IoT Edge module.

Does this meet the goal?

A. Yes

B. No

Answer: (SHOW ANSWER)

Explanation/Reference:

Explanation:

Available in both the cloud and Azure IoT Edge, Azure Stream Analytics offers built-in machine learning based anomaly detection capabilities that can be used to monitor the two most commonly occurring anomalies: temporary and persistent.

Stream Analytics supports user-defined functions, via REST API, that call out to Azure Machine Learning endpoints.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-machine-learning-anomaly-detection>

NEW QUESTION: 57

You are designing a solution that will ingest data from an Azure IoT Edge device, preprocess the data in Azure Machine Learning, and then move the data to Azure HDInsight for further processing.

What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Machine Learning module to use to move the data into HDInsight:

Query type to use:

Output the data to:

Microsoft

Answer:

Machine Learning module to use to move the data into HDInsight:

Query type to use:

Output the data to:

Microsoft

Explanation:

Box 1: Export Data

The Export data to Hive option in the Export Data module in Azure Machine Learning Studio. This option is useful when you are working with very large datasets, and want to save your machine learning experiment data to a Hadoop cluster or HDInsight distributed storage.

Box 2: Apache Hive

Apache Hive is a data warehouse system for Apache Hadoop. Hive enables data summarization, querying, and analysis of data. Hive queries are written in HiveQL, which is a query language similar to SQL.

Box 3: Azure Data Lake

Default storage for the HDFS file system of HDInsight clusters can be associated with either an Azure Storage account or an Azure Data Lake Storage.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/export-to-hive-query>

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/hdinsight-use-hive>

NEW QUESTION: 58

You plan to deploy an application that will perform image recognition. The application will store image data in two Azure Blob storage stores named Blob1 and Blob2. You need to recommend a security solution that meets the following requirements:

- *Access to Blob1 must be controlled by a using a role.
- *Access to Blob2 must be time-limited and constrained to specific operations.


What should you recommend using to control access to each blob store? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.



Answer:



Explanation



References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-auth>

NEW QUESTION: 59

You have a Bing Search service that is used to query a product catalog.

You need to identify the following information:

- * The locale of the query
- * The top 50 query strings
- * The number of calls to the service
- * The top geographical regions of the service

What should you implement?

- A. Bing Statistics
- B. Azure API Management (APIM)
- C. Azure Monitor
- D. Azure Application Insights

Answer: A (LEAVE A REPLY)

Explanation

The Bing Statistics add-in provides metrics such as call volume, top queries, API response, code distribution, and market distribution. The rich slicing-and-dicing capability lets you gather deeper understanding of your users and their usage to inform your business strategy.

References:

<https://www.bingapistatistics.com/>

NEW QUESTION: 60

You plan to use Azure Cognitive Services to provide the development team at your company with the ability to create intelligent apps without having direct AI or data science skills.

The company identifies the following requirements for the planned Cognitive Services deployment:

- * Provide support for the following languages: English, Portuguese, and German.

- * Perform text analytics to derive a sentiment score.

Which Cognitive Service service should you deploy for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Provide support for the following languages:
English, Portuguese, and German:

Perform text analytics to derive a sentiment score:

Microsoft

Answer:

Provide support for the following languages:
English, Portuguese, and German:

Perform text analytics to derive a sentiment score:

Microsoft

Explanation

Provide support for the following languages:
English, Portuguese, and German:

Perform text analytics to derive a sentiment score:

Microsoft

Box 1: Text Analytics

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.

Box 2: Language API

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language>

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-sentiment-analysis-cognitive-services>

NEW QUESTION: 61

You have an app that records meetings by using speech-to-text capabilities from the Speech Services API.

You discover that when action items are listed at the end of each meeting, the app transcribes the text inaccurately.

You need to improve the accuracy of the meeting records.

What should you do?

A. Add a phrase list

B. Create a custom wake word

C. Parse the text by using the Language Understanding (LUIS) API

D. Train a custom model by using Custom Translator

Answer: D ([LEAVE A REPLY](#))

Speech Services API with subscription to the Microsoft Text Translation API enables you to use Custom Translator to use your own data for more accurate translations.

References:

<https://www.microsoft.com/en-us/translator/business/customization/>

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (**163** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 62

You are designing an AI application that will perform real-time processing by using Microsoft Azure Stream Analytics.

You need to identify the valid outputs of a Stream Analytics job.

What are three possible outputs? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. A Hive table in Azure HDInsight

B. Azure SQL Database

C. Azure Cosmos DB

D. Azure Blob storage

E. Azure Redis Cache

Answer: B,C,D ([LEAVE A REPLY](#))

Explanation/Reference:

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-define-outputs>

NEW QUESTION: 63

Your company plans to implement an AI solution that will analyse data from IoT devices.
Data from the devices will be analysed in real time. The results of the analysis will be stored in a SQL database.
You need to recommend a data processing solution that uses the Transact-SQL language.
Which data processing solution should you recommend?

- A. Azure Stream Analytics
- B. SQL Server Integration Services (SSIS)
- C. Azure Event Hubs
- D. Azure Machine Learning

Answer: A (LEAVE A REPLY)

Explanation

References:

<https://www.linkedin.com/pulse/getting-started-azure-iot-services-stream-analytics-rob-tiffany>

NEW QUESTION: 64

You need to build and deploy a real-time media bot for Microsoft Skype on an Azure virtual machine. The bot will use the Azure Bot Service. The solution must minimize custom code.
Which environment and language should you use to develop the bot? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Virtual machine environment:

Linux Ubuntu
Red Hat Enterprise Linux (RHEL)
Windows 10
Windows Server

Language:

C#
C++
Node.js
R

Answer:

Virtual machine environment:

Linux Ubuntu
Red Hat Enterprise Linux (RHEL)
Windows 10
Windows Server

Language:

C#
C++
Node.js
R

Explanation

Virtual machine environment: 

Linux Ubuntu
Red Hat Enterprise Linux (RHEL)
Windows 10
Windows Server

Language: 

C#
C++
Node.js
R

References:

<https://docs.microsoft.com/en-us/microsoftteams/platform/concepts/calls-and-meetings/requirements-considerati>

NEW QUESTION: 65

You have a Face API solution that updates in real time. A pilot of the solution runs successfully on a small dataset.

When you attempt to use the solution on a larger dataset that continually changes, the performance degrades, slowing how long it takes to recognize existing faces.

You need to recommend changes to reduce the time it takes to recognize existing faces without increasing costs.

What should you recommend?

- A. Change the solution to use the Computer Vision API instead of the Face API.
- B. Separate training into an independent pipeline and schedule the pipeline to run daily.
- C. Change the solution to use the Bing Image Search API instead of the Face API.
- D. Distribute the face recognition inference process across many Azure Cognitive Services instances.

Answer: B ([LEAVE A REPLY](#))

Incorrect Answers:

A: The purpose of Computer Vision is to inspect each image associated with an incoming article to (1) scrape out written words from the image and (2) determine what types of objects are present in the image.

C: The Bing API provides an experience similar to Bing.com/search by returning search results that Bing determines are relevant to a user's query. The results include Web pages and may also include images, videos, and more.

D: That would increase cost.

References:

<https://github.com/Azure/cognitive-services>

NEW QUESTION: 66

Your company plan to monitor twitter hashtags. and then to build a graph of connected people and places that contains the associated sentiment.

The monitored hashtags use several languages, but the graph will be displayed in English.

You need to recommend the required Azure Cognitive Services endpoints (or the planned graph Which Cognitive Services endpoints should you recommend?

- A. Translator Text, Content Moderator, and Key Phrase Extraction

- B. Translator Text. Sentiment Analysis, and Named Entity Recognition
- C. Language Detection. Content Moderator, and Key Phrase Extraction
- D. Language Detection. Sentiment Analysis, and Key Phrase Extraction

Answer: D ([LEAVE A REPLY](#))

Explanation

Sentiment analysis, which is also called opinion mining, uses social media analytics tools to determine attitudes toward a product or idea.

Translator Text: Translate text in real time across more than 60 languages, powered by the latest innovations in machine translation.

The Key Phrase Extraction skill evaluates unstructured text, and for each record, returns a list of key phrases.

This skill uses the machine learning models provided by Text Analytics in Cognitive Services.

This capability is useful if you need to quickly identify the main talking points in the record. For example, given input text "The food was delicious and there were wonderful staff", the service returns "food" and

"wonderful staff".

Reference:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-entity-lin>

<https://docs.microsoft.com/en-us/azure/search/cognitive-search-skill-keyphrases>

NEW QUESTION: 67

You have an intelligent edge solution that processes data and outputs the data to an Azure Cosmos DB account that uses the SQL API.

You need to ensure that you can perform full text searches of the data.

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 Area

⇅

Actions

Create a cognitive search pipeline.

Create an Azure Search index.

Create an Azure Service Bus.

Create a data source.

Create an Azure Search indexer.

➤

➡

Microsoft

⬆

⬇

Answer:

Actions

- Create a cognitive search pipeline.
- Create an Azure Search index.
- Create an Azure Service Bus.
- Create a data source.
- Create an Azure Search indexer.

Answer Area

- Create a data source.
- Create an Azure Search indexer.
- Create an Azure Search index.

Explanation

Create a data source.

Create an Azure Search index.

Create an Azure Search indexer.

Reference:

<https://docs.microsoft.com/en-us/azure/search/search-howto-index-cosmosdb>

NEW QUESTION: 68

You are designing an AI solution that must meet the following processing requirements:

- * Use a parallel processing framework that supports the in-memory processing of high volumes of data.
- * Use in-memory caching and a columnar storage engine for Apache Hive queries.

What should you use to meet each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.


Use a parallel processing framework that supports the in-memory processing of high volumes of data.

Use in-memory caching and a columnar storage engine for Apache Hive queries.


Options:

- Apache Kafka
- Apache Spark
- Hive
- Microsoft Machine Learning Server
- Interactive Query

Answer:



Explanation



Box 1: Apache Spark

Apache Spark is a parallel processing framework that supports in-memory processing to boost the performance of big-data analytic applications. Apache Spark in Azure HDInsight is the Microsoft implementation of Apache Spark in the cloud.

Box 2: Interactive Query

Interactive Query provides In-memory caching and improved columnar storage engine for Hive queries.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-overview>

<https://docs.microsoft.com/bs-latn-ba/azure/hdinsight/interactive-query/apache-interactive-query-get-started>

NEW QUESTION: 69

You design an AI workflow that combines data from multiple data sources for analysis. The data sources are composed of:

- * JSON files uploaded to an Azure Storage account
- * On-premises Oracle databases
- * Azure SQL databases

Which service should you use to ingest the data?

- A. Azure Data Factory
- B. Azure SQL Data Warehouse
- C. Azure Data Lake Storage

D. Azure Databricks

Answer: ([SHOW ANSWER](#))

Explanation/Reference:

References:

<https://docs.microsoft.com/en-us/azure/data-factory/introduction>

Design solutions

Testlet 2

Overview

Contoso, Ltd. has an office in New York to serve its North American customers and an office in Paris to serve its European customers.

Existing Environment

Infrastructure

Each office has a small data center that hosts Active Directory services and a few off-the-shelf software solutions used by internal users.

The network contains a single Active Directory forest that contains a single domain named contoso.com. Azure Active Directory (Azure AD) Connect is used to extend identity management to Azure.

The company has an Azure subscription. Each office has an Azure ExpressRoute connection to the subscription. The New York office connects to a virtual network hosted in the US East 2 Azure region.

The Paris office connects to a virtual network hosted in the West Europe Azure region.

The New York office has an Azure Stack Development Kit (ASDK) deployment that is used for development and testing.

Current Business Model

Contoso has a web app named Bookings hosted in an App Service Environment (ASE). The ASE is in the virtual network in the East US 2 region. Contoso employees and customers use Bookings to reserve hotel rooms.

Data Environment

Bookings connects to a Microsoft SQL Server database named hotelDB in the New York office.

The database has a view named vwAvailability that consolidates columns from the tables named Hotels, Rooms, and RoomAvailability. The database contains data that was collected during the last 20 years.

Problem Statements

Contoso identifies the following issues with its current business model:

- * European users report that access to Booking is slow, and they lose customers who must wait on the phone while they search for available rooms.
- * Users report that Bookings was unavailable during an outage in the New York data center for more than 24 hours.

Requirements

Contoso identifies the following issues with its current business model:

- * European users report that access to Bookings is slow, and they lose customers who must wait on the phone while they search for available rooms.
- * Users report that Bookings was unavailable during on outage in the New York data center for more than 24 hours.

Business Goals

Contoso wants to provide a new version of the Bookings app that will provide a highly available, reliable service for booking travel packages by interacting with a chatbot named Butler.

Contoso plans to move all production workloads to the cloud.

Technical requirements

Contoso identifies the following technical requirements:

- * Data scientists must test Butler by using ASDK.
- * Whenever possible, solutions must minimize costs.
- * Butler must greet users by name when they first connect.
- * Butler must be able to handle up to 10,000 messages a day.
- * Butler must recognize the users' intent based on basic utterances.

- * All configurations to the Azure Bot Service must be logged centrally.
- * Whenever possible, solutions must use the principle of least privilege.
- * Internal users must be able to access Butler by using Microsoft Skype for Business.
- * The new Bookings app must provide a user interface where users can interact with Butler.
- * Users in an Azure AD group named KeyManagers must be able to manage keys for all Azure Cognitive Services.
- * Butler must provide users with the ability to reserve a room, cancel a reservation, and view existing reservations.
- * The new Bookings app must be available to users in North America and Europe if a single data center or Azure region fails.
- * For continuous improvement, you must be able to test Butler by sending sample utterances and comparing the chatbot's responses to the actual intent.

NEW QUESTION: 70

DRAG DROP

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.

Actions

- Keys: create
- Keys: get
- Keys: list
- Secrets: all
- Secrets: get
- Secrets: list

Answer Area

IT security team:  Permission

Applications:  Permission

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>

NEW QUESTION: 71

You are designing an AI workflow that will aggregate data stored in Azure as JSON documents.

You expect to store more than 2 TB of new data daily.

You need to choose the data storage service for the data. The solution must minimize costs.

Which data storage service should you choose?

- A. Azure Manage Disks
- B. Azure Blob Storage
- C. Azure File Storage
- D. Azure Data Lake Storage

Answer: B ([LEAVE A REPLY](#))

Explanation/Reference:

Explanation:

Generally, Data Lake will be a bit more expensive although they are in close range of each other. Blob storage has more options for pricing depending upon things like how frequently you need to access your data (cold vs hot storage). Data Lake is priced on volume, so it will go up as you reach certain tiers of volume.

References:

<http://blog.pragmaticworks.com/azure-data-lake-vs-azure-blob-storage-in-data-warehousing>

NEW QUESTION: 72

You have Azure IoT Edge devices that collect measurements every 30 seconds. You plan to send the measurements to an Azure IoT hub. You need to ensure that every event is processed as quickly as possible. What should you use?

- A. Apache Kafka
- B. Azure Stream Analytics record functions
- C. Azure Stream Analytics windowing functions
- D. Azure Machine Learning on the IoT Edge devices

Answer: A ([LEAVE A REPLY](#))

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/kafka/apache-kafka-connector-iot-hub>

NEW QUESTION: 73

You need to create a bot to meet the following requirements:

- * The bot must support multiple bot channels including Direct Line.
- * Users must be able to sign in to the bot by using a Gmail user account and save activities and preferences.

Which four 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions	Answer Area
From the Azure portal, create a client application.	
From the Azure portal, set Guest users permissions are limited to No .	
From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.	
From the Azure portal, enable Azure Multi-Factor Authentication (MFA).	
From the Azure portal, configure an identity provider.	
From the bot code, add the connection settings and OAuthPrompt.	

Answer:

Actions	Answer Area
From the Azure portal, create a client application.	From the Azure portal, configure an identity provider.
From the Azure portal, set Guest users permissions are limited to No .	From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.
From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.	From the Azure portal, create a client application.
From the Azure portal, enable Azure Multi-Factor Authentication (MFA).	From the bot code, add the connection settings and OAuthPrompt.
From the Azure portal, configure an identity provider.	
From the bot code, add the connection settings and OAuthPrompt.	

Explanation:

Step 1: From the Azure portal, configure an identity provider.

The Azure Bot Service and the v4 SDK include new bot authentication capabilities, providing features to make it easier to develop a bot that authenticates users to various identity providers, such as Azure AD (Azure Active Directory), GitHub, Uber, and so on.

Step 2: From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.

Azure Active Directory B2C provides business-to-customer identity as a service. Your customers use their preferred social, enterprise, or local account identities to get single sign-on access to your applications and APIs.

Step 3: From the Azure portal, create a client application

You can enable communication between your bot and your own client application by using the Direct Line API.

Step 4: From the bot code, add the connection settings and OAuthPrompt

Use an OAuth prompt to sign the user in and get a token.

Azure AD B2C uses standards-based authentication protocols including OpenID Connect, OAuth 2.0, and SAML.

References:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-authentication?view=azure-bot-service-4.0>

NEW QUESTION: 74

You are building an Azure Analysis Services cube for your AI deployment.

The source data for the cube is located in an on premises network in a Microsoft SQL Server database.

You need to ensure that the Azure Analysis Services service can access the source data.

What should you deploy to your Azure subscription?

A. a site-to-site VPN


B. a data gateway

C. Azure Data Factory

D. a network gateway

Answer: (SHOW ANSWER)

From April 2017 onward we can use On-premises Data Gateway for Azure Analysis Services. This means you can connect your Tabular Models hosted in Azure Analysis Services to your on-premises data sources through On-premises Data Gateway.



References:

<https://biinsight.com/on-premises-data-gateway-for-azure-analysis-services/>

NEW QUESTION: 75

You plan to use Azure Cognitive Services to provide the development team at your company with the ability to create intelligent apps without having direct AI or data science skills.

The company identifies the following requirements for the planned Cognitive Services deployment:

* Provide support for the following languages: English, Portuguese, and German.

* Perform text analytics to derive a sentiment score.

Which Cognitive Service service should you deploy for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Provide support for the following languages:
English, Portuguese, and German:

Perform text analytics to derive a sentiment score:

Text Analytics
Search services
Speech services

Speech APIs
Decision APIs
Language APIs

Answer:

Provide support for the following languages:
English, Portuguese, and German:

Perform text analytics to derive a sentiment score:

Text Analytics
Search services
Speech services

Speech APIs
Decision APIs
Language APIs

Explanation

Provide support for the following languages:
English, Portuguese, and German:

Perform text analytics to derive a sentiment score:

Text Analytics
Search services
Speech services

Speech APIs
Decision APIs
Language APIs

Box 1: Text Analytics

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.

Box 2: Language API

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language>

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-sentiment-analysis-cognitive-services>

NEW QUESTION: 76

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 Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy Azure Stream Analytics as an IoT Edge module.

Does this meet the goal?

A. Yes

B. No

Answer: A ([LEAVE A REPLY](#))

Available in both the cloud and Azure IoT Edge, Azure Stream Analytics offers built-in machine learning based anomaly detection capabilities that can be used to monitor the two most commonly occurring anomalies:

temporary and persistent.

Stream Analytics supports user-defined functions, via REST API, that call out to Azure Machine Learning endpoints.

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-machine-learning-anomaly-detection>

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (163 Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 77

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.

Location to store the metadata:

Azure Blob storage
Azure Cosmos DB
Azure File Storage

Virtual machine series to run the vision module:

A
F
NV

Virtual machine series to run the mathematical module:

A
F
NV

Answer:

Location to store the metadata:

Azure Blob storage
Azure Cosmos DB
Azure File Storage


Virtual machine series to run the vision module:

A
F
NV

Virtual machine series to run the mathematical module:

A
F
NV

Explanation

Location to store the metadata: 

Virtual machine series to run the vision module: 

Virtual machine series to run the mathematical module: 

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/>

NEW QUESTION: 78

You are designing an AI solution that will analyze millions of pictures by using Azure HDInsight Hadoop cluster.

You need to recommend a solution for storing the pictures. The solution must minimize costs.

Which storage solution should you recommend?

- A. an Azure Data Lake Storage Gen1
- B. Azure File Storage
- C. Azure Blob storage
- D. Azure Table storage

Answer: (SHOW ANSWER)

Data Lake will be a bit more expensive although they are in close range of each other. Blob storage has more options for pricing depending upon things like how frequently you need to access your data (cold vs hot storage).

Reference:

<http://blog.pragmaticworks.com/azure-data-lake-vs-azure-blob-storage-in-data-warehousing>

NEW QUESTION: 79

Your company has recently purchased and deployed 25,000 IoT devices.

You need to recommend a data analysis solution for the devices that meets the following requirements:

- * Each device must use its own credentials for identity.
- * Each device must be able to route data to multiple endpoints.
- * The solution must require the minimum amount of customized code.

What should you include in the recommendation?

- A.** Microsoft Azure Notification Hubs
- B.** Microsoft Azure Event Hubs
- C.** Microsoft Azure IoT Hub
- D.** Microsoft Azure Service Bus

Answer: C ([LEAVE A REPLY](#))

An IoT hub has a default built-in endpoint. You can create custom endpoints to route messages to by linking other services in your subscription to the hub.

Individual devices connect using credentials stored in the IoT hub's identity registry.

References:

<https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-devguide-security>

NEW QUESTION: 80

You need to design an application that will analyze real-time data from financial feeds.

The data will be ingested into Azure IoT Hub. The data must be processed as quickly as possible in the order in which it is ingested.

Which service should you include in the design?

- A.** Azure Data Factory
- B.** Azure Queue storage
- C.** Azure Stream Analytics
- D.** Azure Notification Hubs

Answer: C ([LEAVE A REPLY](#))

References:

<https://docs.microsoft.com/en-us/azure/architecture/data-guide/big-data/real-time-processing>

NEW QUESTION: 81

You are designing an Azure Batch AI solution that will be used to train many different Azure Machine Learning models. The solution will perform the following:

- * Image recognition
- * Deep learning that uses convolutional neural networks

You need to select a compute infrastructure for each model. The solution must minimize the processing time.

What should you use for each model? To answer, drag the appropriate compute infrastructures to the correct models. Each compute infrastructure 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.

Compute Infrastructures

Compute optimized virtual machines

GPU optimized virtual machines

Memory optimized virtual machines

Answer Area

Image recognition:

Compute Infrastructure

Deep learning that uses convolutional neural networks:

Compute Infrastructure



Microsoft

Answer:

Compute Infrastructures

Compute optimized virtual machines

GPU optimized virtual machines

Memory optimized virtual machines

Answer Area

Image recognition: GPU optimized virtual machines

Deep learning that uses convolutional neural networks: GPU optimized virtual machines

The Microsoft logo, consisting of four colored squares (red, green, blue, and yellow) arranged in a 2x2 grid.

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/sizes-gpu>

NEW QUESTION: 82

You are designing an AI solution that will provide feedback to teachers who train students over the Internet. The students will be in classrooms located in remote areas. The solution will capture video and audio data of the students in the classrooms.

You need to recommend Azure Cognitive Services for the AI solution to meet the following requirements:

Alert teachers if a student seems angry or distracted.

Identify each student in the classrooms for attendance purposes.

Allow the teachers to log the text of conversations between themselves and the students.

Which Cognitive Services should you recommend?

A. Computer Vision, Text Analytics, and Face API

B. Video Indexer, Face API, and Text Analytics

C. Computer Vision, Speech to Text, and Text Analytics

D. Text Analytics, QnA Maker, and Computer Vision

E. Video Indexer, Speech to Text, and Face API

Answer: E ([LEAVE A REPLY](#))

Azure Video Indexer is a cloud application built on Azure Media Analytics, Azure Search, Cognitive Services (such as the Face API, Microsoft Translator, the Computer Vision API, and Custom Speech Service). It enables you to extract the insights from your videos using Video Indexer video and audio models.

Face API enables you to search, identify, and match faces in your private repository of up to 1 million people.

The Face API now integrates emotion recognition, returning the confidence across a set of emotions for each face in the image such as anger, contempt, disgust, fear, happiness, neutral, sadness, and surprise.

These emotions are understood to be cross-culturally and universally communicated with particular facial expressions.

Speech-to-text from Azure Speech Services, also known as speech-to-text, enables real-time transcription of audio streams into text that your applications, tools, or devices can consume, display, and take action on as command input. This service is powered by the same recognition technology that Microsoft uses for Cortana and Office products, and works seamlessly with the translation and text-to-speech.

Incorrect Answers:

Computer Vision or the QnA is not required.

References:

<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/video-indexer-overview>

<https://azure.microsoft.com/en-us/services/cognitive-services/face/>

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/speech-to-text>

NEW QUESTION: 83

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 ([LEAVE A REPLY](#))

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>

NEW QUESTION: 84

DRAG DROP

You need to build a pipeline for an Azure Machine Learning experiment.

In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

.....

Actions

Score the model.

Import data into the Machine Learning experiment.

Use the Tune Model Hyperparameters module to find the best combination.

Evaluate the model.

Split data into training and testing sets.

Answer Area

Microsoft

Answer:

.....

Actions

Score the model.

Import data into the Machine Learning experiment.

Use the Tune Model Hyperparameters module to find the best combination.

Evaluate the model.

Split data into training and testing sets.

Answer Area

Score the model.

Use the Tune Model Hyperparameters module to find the best combination.

Import data into the Machine Learning experiment.

Evaluate the model.

Split data into training and testing sets.

Microsoft

NEW QUESTION: 85

You deploy an Azure bot.

You need to collect Key Performance Indicator (KPI) data from the bot. The type of data includes:

- * The number of users interacting with the bot
- * The number of messages interacting with the bot
- * The number of messages on different channels received by the bot

* The number of users and messages continuously interacting with the bot What should you configure?

- A. Bot analytics
- B. Azure Monitor
- C. Azure Analysis Services
- D. Azure Application Insights

Answer: A ([LEAVE A REPLY](#))

Explanation

References:

<https://docs.microsoft.com/en-us/azure/sql-database/saas-multitenantdb-adhoc-reporting>

NEW QUESTION: 86

You need to build and deploy a real-time media bot for Microsoft Skype on an Azure virtual machine. The bot will use the Azure Bot Service. The solution must minimize custom code.

Which environment and language should you use to develop the bot? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Virtual machine environment:

- Linux Ubuntu
- Red Hat Enterprise Linux (RHEL)
- Windows 10
- Windows Server

Language:

- C#
- C++
- Node.js
- R

Answer:

Virtual machine environment:

- Linux Ubuntu
- Red Hat Enterprise Linux (RHEL)
- Windows 10
- Windows Server

Language:

- C#
- C++
- Node.js
- R

Explanation

Virtual machine environment:

Microsoft	▼
Linux Ubuntu	
Red Hat Enterprise Linux (RHEL)	
Windows 10	
Windows Server	

Language:

	▼
C#	
C++	
Node.js	
R	

References:

<https://docs.microsoft.com/en-us/microsoftteams/platform/concepts/calls-and-meetings/requirements-considerati>

NEW QUESTION: 87

You need to recommend a data storage solution that meets the technical requirements.

What is the best data storage solution to recommend? More than one answer choice may achieve the goal.

Select the BEST answer.

- A. Azure Databricks
- B. Azure SQL Database
- C. Azure Table storage
- D. Azure Cosmos DB

Answer: D ([LEAVE A REPLY](#))

Explanation/Reference:

<https://docs.microsoft.com/en-us/azure/architecture/example-scenario/ai/commerce-chatbot>

<https://docs.microsoft.com/en-us/azure/cosmos-db/introduction>

Design AI solutions

Question Set 1

NEW QUESTION: 88

You need to create a bot to meet the following requirements:

- * The bot must support multiple bot channels including Direct Line.
- * Users must be able to sign in to the bot by using a Gmail user account and save activities and preferences.

Which four 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Actions	Answer Area
From the Azure portal, create a client application.	
From the Azure portal, set Guest users permissions are limited to No .	
From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.	
From the Azure portal, enable Azure Multi-Factor Authentication (MFA).	
From the Azure portal, configure an identity provider.	
From the bot code, add the connection settings and OAuthPrompt.	

Answer:

Actions	Answer Area
From the Azure portal, create a client application.	From the Azure portal, configure an identity provider.
From the Azure portal, set Guest users permissions are limited to No .	From the Azure portal, set Guest users permissions are limited to No .
From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.	From the Azure portal, create a client application.
From the Azure portal, enable Azure Multi-Factor Authentication (MFA).	From the bot code, add the connection settings and OAuthPrompt.
From the Azure portal, configure an identity provider.	
From the bot code, add the connection settings and OAuthPrompt.	

Explanation:

Step 1: From the Azure portal, configure an identity provider.

The Azure Bot Service and the v4 SDK include new bot authentication capabilities, providing features to make it easier to develop a bot that authenticates users to various identity providers, such as Azure AD (Azure Active Directory), GitHub, Uber, and so on.

Step 2: From the Azure portal, create an Azure Active Directory (Azure AD) B2C service.

Azure Active Directory B2C provides business-to-customer identity as a service. Your customers use their preferred social, enterprise, or local account identities to get single sign-on access to your applications and APIs.

Step 3: From the Azure portal, create a client application

You can enable communication between your bot and your own client application by using the Direct Line API.

Step 4: From the bot code, add the connection settings and OAuthPrompt

Use an OAuth prompt to sign the user in and get a token.

Azure AD B2C uses standards-based authentication protocols including OpenID Connect, OAuth 2.0, and SAML.

References:


<https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-authentication?view=azure-bot-service-4.0>

NEW QUESTION: 89

You need to build a sentiment analysis solution that will use input data from JSON documents and PDF documents. The JSON documents must be processed in batches and aggregated.

Which storage type should you use for each file type? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



The screenshot shows a Microsoft Azure portal interface for selecting storage types. It has two sections: 'JSON documents:' and 'PDF documents:'. Each section has a dropdown menu with four options: 'Azure Blob storage', 'Azure Cosmos DB', 'Azure Data Lake', and 'Azure Table storage'. The 'JSON documents:' dropdown is currently open, showing the options. The 'PDF documents:' dropdown is also open, showing the same options. The Microsoft logo is visible in the top right corner.

Answer:



The screenshot shows the same Microsoft Azure portal interface as the previous one, but with the correct storage types selected. For 'JSON documents:', 'Azure Table storage' is selected. For 'PDF documents:', 'Azure Blob storage' is selected. The Microsoft logo is visible in the top right corner.

NEW QUESTION: 90

You deploy an application that performs sentiment analysis on the data stored in Azure Cosmos DB.

Recently, you loaded a large amount of data to the database. The data was for a customer named Contoso, Ltd.

You discover that queries for the Contoso data are slow to complete, and the queries slow the entire application.

You need to reduce the amount of time it takes for the queries to complete. The solution must minimize costs.

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

- A. Change the request units.
- B. Change the partitioning strategy.
- C. Change the transaction isolation level.
- D. Migrate the data to the Cosmos DB database.

Answer: B ([LEAVE A REPLY](#))

Explanation/Reference:

References:

<https://docs.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning>

NEW QUESTION: 91

You are designing a solution that uses drones to monitor remote locations for anomalies. The drones have Azure IoT Edge devices. The solution must meet the following requirements:

- *Email a user the picture and location of an anomaly when an anomaly is detected.
- *Use a video stream to detect anomalies at the location.
- *Send the pictures and location information to Azure.
- *Use the least amount of code possible.

You develop a custom vision Azure Machine Learning module to detect the anomalies.

Which service should you use for each requirement? To answer, drag the appropriate services to the correct requirements. Each service 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.

The screenshot shows the Microsoft Learn assessment interface. On the left, under the heading "Services", there are four boxes: "Azure Functions", "Azure IoT Edge" (with a red and green icon), "Azure IoT Hub" (with a blue and orange icon), and "Azure Logic Apps". On the right, under the heading "Answer Area", there are three rows of requirements, each followed by a box labeled "Service":

- Requirement: "Use a video stream to detect anomalies at the location:"
- Requirement: "Send the pictures and location information to Azure:"
- Requirement: "Email a user the picture and location of an anomaly when an anomaly is detected:"


At the top of the interface, there are five blue dots, with the fourth dot from the left being filled, indicating the current question.

Answer:

The screenshot shows the same Microsoft Learn assessment interface, but with the answer applied. The "Services" pane on the left now has dashed green boxes around "Azure IoT Edge", "Azure IoT Hub", and "Azure Logic Apps". The "Answer Area" on the right now has the following services assigned to the requirements:

- Requirement: "Use a video stream to detect anomalies at the location:" is assigned "Azure IoT Edge".
- Requirement: "Send the pictures and location information to Azure:" is assigned "Azure Functions".
- Requirement: "Email a user the picture and location of an anomaly when an anomaly is detected:" is assigned "Azure Logic Apps".

Explanation



Box 1: Azure IOT Edge

Example:

You configure the Remote Monitoring solution to respond to anomalies detected by an IoT Edge device. IoT Edge devices let you process telemetry at the edge to reduce the volume of telemetry sent to the solution and to enable faster responses to events on devices.

Box 2: Azure Functions

Box 3: Azure Logic Apps

References:

<https://docs.microsoft.com/en-us/azure/iot-accelerators/iot-accelerators-remote-monitoring-edge>

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (**163 Q&As** Dumps, **40%OFF** Special Discount: **Exam-Tests**)

NEW QUESTION: 92

Your company has recently purchased and deployed 25,000 IoT devices.

You need to recommend a data analysis solution for the devices that meets the following requirements:

- * Each device must use its own credentials for identity.
- * Each device must be able to route data to multiple endpoints.
- * The solution must require the minimum amount of customized code.

What should you include in the recommendation?

- A. Microsoft Azure Notification Hubs
- B. Microsoft Azure Event Hubs
- C. Microsoft Azure IoT Hub
- D. Microsoft Azure Service Bus

Answer: C (LEAVE A REPLY)

Explanation

An IoT hub has a default built-in endpoint. You can create custom endpoints to route messages to by linking other services in your subscription to the hub.

Individual devices connect using credentials stored in the IoT hub's identity registry.

References:

<https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-devguide-security>

NEW QUESTION: 93

You plan to create an intelligent bot to handle internal user chats to the help desk of your company. The bot has the following requirements:

- * Must be able to interpret what a user means.
- * Must be able to perform multiple tasks for a user.
- * Must be able to answer questions from an existing knowledge base

You need to recommend which solutions meet the requirements.

Which solution should you recommend for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

Must be able to interpret what a user means:

- ☐ The Text Analytics API
- ☐ The QnA Maker service
- ☐ The Speech to Text API
- ☐ The Dispatch tool library
- ☐ The Language Understanding (LUI) service

Must be able to perform multiple tasks for a user:

- ☒ The Text Analytics API
- ☐ The QnA Maker service
- ☐ The Speech to Text API
- ☐ The Dispatch tool library
- ☐ The Language Understanding (LUI) service

Must be able to answer questions from an existing knowledge base:

- ☐ The Text Analytics API
- ☐ The QnA Maker service
- ☐ The Speech to Text API
- ☐ The Dispatch tool library
- ☐ The Language Understanding (LUI) service

Answer:

Answer Area

Must be able to interpret what a user means:

- ☐ The Text Analytics API
- ☐ The QnA Maker service
- ☐ The Speech to Text API
- ☐ The Dispatch tool library
- ☐ The Language Understanding (LUI) service

Must be able to perform multiple tasks for a user:

- ☒ The Text Analytics API
- ☐ The QnA Maker service
- ☐ The Speech to Text API
- ☐ The Dispatch tool library
- ☐ The Language Understanding (LUI) service

Must be able to answer questions from an existing knowledge base:

- ☐ The Text Analytics API
- ☐ The QnA Maker service
- ☐ The Speech to Text API
- ☐ The Dispatch tool library
- ☐ The Language Understanding (LUI) service

Explanation

Must be able to interpret what a user means:

	▼
The Text Analytics API	
The QnA Maker service	
The Speech to Text API	
The Dispatch tool library	
The Language Understanding (LUIS) service	

Must be able to perform multiple tasks for a user:

	▼
The Text Analytics API	
The QnA Maker service	
The Speech to Text API	
The Dispatch tool library	
The Language Understanding (LUIS) service	

Must be able to answer questions from an existing knowledge base:



	▼
The Text Analytics API	
The QnA Maker service	
The Speech to Text API	
The Dispatch tool library	
The Language Understanding (LUIS) service	

Box 1: The Language Understanding (LUIS) service

Language Understanding (LUIS) is a cloud-based API service that applies custom machine-learning intelligence to a user's conversational, natural language text to predict overall meaning, and pull out relevant, detailed information.

Box 2: Text Analytics API

The Text Analytics API is a cloud-based service that provides advanced natural language processing over raw text, and includes four main functions: sentiment analysis, key phrase extraction, named entity recognition, and language detection.

Box 3: The QnA Maker service

QnA Maker is a cloud-based Natural Language Processing (NLP) service that easily creates a natural conversational layer over your data. It can be used to find the most appropriate answer for any given natural language input, from your custom knowledge base (KB) of information.

Reference:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-builder-tutorial-dispatch>

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/overview>

NEW QUESTION: 94

You plan to design a solution for an AI implementation that uses data from IoT devices.

You need to recommend a data storage solution for the IoT devices that meets the following requirements:

*Allow data to be queried in real-time as it streams into the solution.

*Provide the lowest amount of latency for loading data into the solution.

What should you include in the recommendation?

- A. a Microsoft Azure SQL database that has In-Memory OLTP enabled
- B. a Microsoft Azure HDInsight R Server cluster
- C. a Microsoft Azure Table Storage solution
- D. a Microsoft Azure HDInsight Hadoop cluster

Answer: D ([LEAVE A REPLY](#))

You can use HDInsight to process streaming data that's received in real time from a variety of devices.

Internet of Things (IoT)

You can use HDInsight to build applications that extract critical insights from data. You can also use Azure Machine Learning on top of that to predict future trends for your business.

By combining enterprise-scale R analytics software with the power of Apache Hadoop and Apache Spark, Microsoft R Server for HDInsight gives you the scale and performance you need. Multi-threaded math libraries and transparent parallelization in R Server handle up to 1000x more data and up to 50x faster speeds than open-source R, which helps you to train more accurate models for better predictions.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/apache-hadoop-introduction>

NEW QUESTION: 95

Your company has an on-premises datacenter.

You plan to publish an app that will recognize a set of individuals by using the Face API. The model is trained.

You need to ensure that all images are processed in the on-premises datacenter.

What should you deploy to host the Face API?

- A.** a Docker container
- B.** Azure File Sync
- C.** Azure Application Gateway
- D.** Azure Data Box Edge

Answer: A ([LEAVE A REPLY](#))

A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. A Docker container image is a lightweight, standalone, executable package of software that includes everything needed to run an application:

code, runtime, system tools, system libraries and settings.

Incorrect Answers:

D: Azure Data Box Edge is an AI-enabled edge computing device with network data transfer capabilities. This article provides you an overview of the Data Box Edge solution, benefits, key capabilities, and the scenarios where you can deploy this device.

Data Box Edge is a Hardware-as-a-service solution. Microsoft ships you a cloud-managed device with a built-in Field Programmable Gate Array (FPGA) that enables accelerated AI-inferencing and has all the capabilities of a storage gateway.

References:

<https://www.docker.com/resources/what-container>

NEW QUESTION: 96

Your company has a data team of Scala and R experts.

You plan to ingest data from multiple Apache Kafka streams.

You need to recommend a processing technology to broker messages at scale from Kafka streams to Azure Storage.

What should you recommend?

- A.** Azure Databricks
- B.** Azure Functions
- C.** Azure HDInsight with Apache Storm
- D.** Azure HDInsight with Microsoft Machine Learning Server

Answer: C ([LEAVE A REPLY](#))

Explanation/Reference:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-streaming-at-scale-overview?toc=https%3A%2F%2Fdocs.microsoft.com%2Fen-us%2Fazure%2Fhdinsight%2Fhadoop%2FTOC.json&bc=https%3A%2F%2Fdocs.microsoft.com%2Fen-us%2Fazure%2Fbread%2Ftoc.json>

NEW QUESTION: 97

You need to recommend a data storage solution that meets the technical requirements.

What is the best data storage solution to recommend? More than one answer choice may achieve the goal.

Select the BEST answer.

- A.** Azure Databricks
- B.** Azure SQL Database
- C.** Azure Table storage
- D.** Azure Cosmos DB

Answer: ([SHOW ANSWER](#))

Explanation

References:

<https://docs.microsoft.com/en-us/azure/architecture/example-scenario/ai/commerce-chatbot>

NEW QUESTION: 98

You have an app that uses the Language Understanding (LUIS) API as shown in the following exhibit.

Cognitive Services | Language UnderstandingMy apps

MyLUIS(API) ▾DASHBOARDBUILDMANAGETrainTestPublish

App AssetsIntentsEntitiesImprove app performanceReview endpoint utterancesPhrase listsPatterns

Intents ?

+ Create new intent+ Add prebuilt domain intent

Search intents

☐ Name ^	Labeled Utterances
None	0

Microsoft

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
NOTE: Each correct selection is worth one point.

To publish the app, you must first [answer choice] the app.

	▼
test	
train	
create utterances in	

You can add utterances to the model by [answer choice].

	▼
creating intents	
testing the model	
training the model	

The app was [answer choice].

	▼
never published	
already published	
published but never accessed by users	

Answer:

To publish the app, you must first [answer choice] the app.

	▼
test	
train	
create utterances in	

You can add utterances to the model by [answer choice].

	▼
creating intents	
testing the model	
training the model	

The app was [answer choice].

	▼
never published	
already published	
published but never accessed by users	

Explanation

To publish the app, you must first [answer choice] the app.

	▼
test	
train	
create utterances in	

You can add utterances to the model by [answer choice].

	▼
creating intents	
testing the model	
training the model	

The app was [answer choice].

	▼
never published	
already published	
published but never accessed by users	

Box 1: train

Utterances are input from the user that your app needs to interpret. To train LUIS to extract intents and entities from them, it's important to capture a variety of different example utterances for each intent. Active learning, or the process of continuing to train on new utterances, is essential to machine-learned intelligence that LUIS provides.

Box 2: creating intents

Each intent needs to have example utterances, at least 15. If you have an intent that does not have any example utterances, you will not be able to train LUIS. If you have an intent with one or very few example utterances, LUIS will not accurately predict the intent.

Box 3: never published

In each iteration of the model, do not add a large quantity of utterances. Add utterances in quantities of 15.

Train, publish, and test again.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-concept-utteran3ce>

NEW QUESTION: 99

You have thousands of images that contain text.

You need to process the text from the images into a machine-readable character stream.

Which Azure Cognitive Services service should you use?

- A. Translator Text
- B. Text Analytics
- C. Computer Vision
- D. the Image Moderation API

Answer: C ([LEAVE A REPLY](#))

Explanation

With Computer Vision you can detect text in an image using optical character recognition (OCR) and extract the recognized words into a machine-readable character stream.

References:
<https://azure.microsoft.com/en-us/services/cognitive-services/computer-vision/>
<https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/image-moderation-api>

NEW QUESTION: 100

You are designing a solution that will ingest data from an Azure IoT Edge device, preprocess the data in Azure Machine Learning, and then move the data to Azure HDInsight for further processing. What should you include in the solution? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Machine Learning module to use to move the data into HDInsight:

Export Data

Load Trained Model

Partition and Sample

Unpack Zipped Datasets

Query type to use:

Apache Hive

Apache Spark

Q#

Transact-SQL

Output the data to:

Azure Cosmos DB

Azure Data Lake

Azure Table storage

HDFS

Answer:

Machine Learning module to use to move the data into HDInsight:

Export Data

Load Trained Model

Partition and Sample

Unpack Zipped Datasets

Query type to use:

Apache Hive

Apache Spark

Q#

Transact-SQL

Output the data to:


Azure Cosmos DB

Azure Data Lake

Azure Table storage

HDFS

Explanation



Box 1: Export Data

The Export data to Hive option in the Export Data module in Azure Machine Learning Studio. This option is useful when you are working with very large datasets, and want to save your machine learning experiment data to a Hadoop cluster or HDInsight distributed storage.

Box 2: Apache Hive

Apache Hive is a data warehouse system for Apache Hadoop. Hive enables data summarization, querying, and analysis of data. Hive queries are written in HiveQL, which is a query language similar to SQL.

Box 3: Azure Data Lake

Default storage for the HDFS file system of HDInsight clusters can be associated with either an Azure Storage account or an Azure Data Lake Storage.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/export-to-hive-query>

<https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/hdinsight-use-hive>

NEW QUESTION: 101

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.

Actions

Add the job to the IoT devices in IoT hub

Create an Azure Stream Analytics Edge job

Create an Azure Stream Analytics Cloud job

Create a storage container

Create a storage queue

Answer Area



Answer:

Actions

- Add the job to the IoT devices in IoT hub
- Create an Azure Stream Analytics Edge job
- Create an Azure Stream Analytics Cloud job
- Create a storage container
- Create a storage queue

Answer Area

- Create a storage container
- Create an Azure Stream Analytics Edge job
- Add the job to the IoT devices in IoT hub

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:

NEW QUESTION: 102

Your company plans to build an app that will perform the following tasks:

- * Match a user's picture to a picture of a celebrity.
- * Tag a scene from a movie, and then search for movie scenes by using the tags.

You need to recommend which Azure Cognitive Services APIs must be used to perform the tasks.

Which Cognitive Services API should you recommend for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Match a user's picture to a picture of a celebrity.

Tag a scene from a movie, and then search for movie scenes by using the tags:

Both dropdown menus show the following options: Bing Video Search, Video Indexer, Custom Vision, and Computer Vision.

Answer:

Match a user's picture to a picture of a celebrity.

Tag a scene from a movie, and then search for movie scenes by using the tags:

For the first task, 'Computer Vision' is selected. For the second task, 'Bing Video Search' is selected.

Explanation:

Box 1: Computer Vision

Azure's Computer Vision service provides developers with access to advanced algorithms that process images and return information.

Computer Vision Detect Faces: Detect faces in an image and provide information about each detected face. Computer Vision returns the coordinates, rectangle, gender, and age for each detected face.

Computer Vision provides a subset of the Face service functionality. You can use the Face service for more detailed analysis, such as facial identification and pose detection.

Box 2: Bing Video Search

Search for videos and get comprehensive results

With Bing Video Search API v7, find videos across the web. Results provide useful metadata including creator, encoding format, video length, view count, improved & simplified paging, and more.

Incorrect Answers:

Video Indexer:

Automatically extract metadata-such as spoken words, written text, faces, speakers, celebrities, emotions, topics, brands, and scenes-from video and audio files.

Custom Vision:

Easily customize your own state-of-the-art computer vision models for your unique use case. Just upload a few labeled images and let Custom Vision Service do the hard work. With just one click, you can export trained models to be run on device or as Docker containers.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/computer-vision/home>

<https://azure.microsoft.com/en-us/services/cognitive-services/bing-video-search-api/>

NEW QUESTION: 103

Your company plans to create a mobile app that will be used by employees to query the employee handbook.

You need to ensure that the employees can query the handbook by typing or by using speech.

Which core component should you use for the app?

A. Language Understanding (LUIS)

B. QnA Maker

C. Text Analytics

D. Azure Search

Answer: ([SHOW ANSWER](#))

Explanation

Azure Cognitive Search (formerly known as "Azure Search") is a search-as-a-service cloud solution that gives developers APIs and tools for adding a rich search experience over private, heterogeneous content in web, mobile, and enterprise applications. Your code or a tool invokes data ingestion (indexing) to create and load an index. Optionally, you can add cognitive skills to apply AI processes during indexing. Doing so can add new information and structures useful for search and other scenarios.

Incorrect Answwres:

B: QnA Maker is a cloud-based API service that lets you create a conversational question-and-answer layer over your existing data. Use it to build a knowledge base by extracting questions and answers from your semi-structured content, including FAQs, manuals, and documents. Answer users' questions with the best answers from the QnAs in your knowledge base-automatically.

References:

<https://docs.microsoft.com/en-us/azure/search/search-what-is-azure-search>

NEW QUESTION: 104

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.

Location to store the metadata:  ▼

Azure Blob storage
Azure Cosmos DB
Azure File Storage

Virtual machine series to run the vision module: ▼

A
F
NV

Virtual machine series to run the mathematical module: ▼

A
F
NV

Answer:

Location to store the metadata:  ▼

Azure Blob storage
Azure Cosmos DB
Azure File Storage

Virtual machine series to run the vision module: ▼

A
F
NV

Virtual machine series to run the mathematical module: ▼

A
F
NV

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/>

NEW QUESTION: 105

You plan to deploy the Text Analytics and Computer Vision services. The Azure Cognitive Services will be deployed to the West US and East Europe Azure regions. You need to identify the minimum number of service endpoints and API keys required for the planned deployment. What should you identify? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Service endpoints:

▼

1

2

4

API keys:

▼

1

2

4

 Microsoft

Answer:

Service endpoints:

▼

1

2

4

API keys:


▼

1

2

4

Explanation



Box 1: 2

After creating a Cognitive Service resource in the Azure portal, you'll get an endpoint and a key for authenticating your applications. You can access Azure Cognitive Services through two different resources: A multi-service resource, or a single-service one.

Multi-service resource: Access multiple Azure Cognitive Services with a single key and endpoint.

Note: You need a key and endpoint for a Text Analytics resource. Azure Cognitive Services are represented by Azure resources that you subscribe to.

Each request must include your access key and an HTTP endpoint. The endpoint specifies the region you chose during sign up, the service URL, and a resource used on the request

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/cognitive-services-apis-create-account>

NEW QUESTION: 106

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 in this section, 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 kubectl command, and then you create an SSH connection.

Does this meet the goal?

A. Yes

B. No

Answer: ([SHOW ANSWER](#))

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (163 Q&As Dumps, **40%OFF Special Discount: Exam-Tests**)

NEW QUESTION: 107

You are designing an AI solution that will analyze millions of pictures.

You need to recommend a solution for storing the pictures. The solution must minimize costs.

Which storage solution should you recommend?

- A. an Azure Data Lake store
- B. Azure File Storage
- C. Azure Blob storage
- D. Azure Table storage

Answer: C (LEAVE A REPLY)

Generally, Data Lake will be a bit more expensive although they are in close range of each other. Blob storage has more options for pricing depending upon things like how frequently you need to access your data (cold vs hot storage). Data Lake is priced on volume, so it will go up as you reach certain tiers of volume.

References:


<http://blog.practiceworks.com/azure-data-lake-vs-azure-blob-storage-in-data-warehousing>

NEW QUESTION: 108


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/>

NEW QUESTION: 109

You need to configure versioning and logging for Azure Machine Learning models. Which Machine Learning service application should you use?

- A. models
- B. activities
- C. experiments
- D. pipelines
- E. deployments

Answer: (SHOW ANSWER)

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-enable-logging#logging-for-deployed-models>

NEW QUESTION: 110

Your company manages a sports team.

The company sets up a video booth to record messages for the team.

Before replaying the messages on a video screen, you need to generate captions for the messages and check the sentiment of the video to ensure that only positive messages are played.

Which Azure Cognitive Services service should you use?

- A. Language Understanding (LUIS)
- B. Speaker Recognition
- C. Custom Vision
- D. Video Indexer

Answer: D (LEAVE A REPLY)

Video Indexer includes Audio transcription: Converts speech to text in 12 languages and allows extensions.

Supported languages include English, Spanish, French, German, Italian, Mandarin Chinese, Japanese, Arabic, Russian, Portuguese, Hindi, and Korean.

When indexing by one channel, partial result for those models will be available, such as sentiment analysis:

Identifies positive, negative, and neutral sentiments from speech and visual text.

Reference:

<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/video-indexer-overview>

NEW QUESTION: 111

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 deploying an Azure Machine Learning model to an Azure Kubernetes Service (AKS) container.

You need to monitor the accuracy of each run of the model.

Solution: You configure Azure Application Insights.

Does this meet the goal?

A. Yes

B. No

Answer: A ([LEAVE A REPLY](#))

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-enable-data-collection>

<https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-enable-app-insights>

NEW QUESTION: 112

Your company plans to build an app that will perform the following tasks:

* Match a user's picture to a picture of a celebrity.

* Tag a scene from a movie, and then search for movie scenes by using the tags.

You need to recommend which Azure Cognitive Services APIs must be used to perform the tasks.

Which Cognitive Services API should you recommend for each task? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Match a user's picture to a picture of a celebrity:



Tag a scene from a movie, and then search for movie scenes by using the tags:

Bing Video Search
Video Indexer
Custom Vision
Computer Vision

Bing Video Search
Video Indexer
Custom Vision
Computer Vision

Answer:

Explanation

Match a user's picture to a picture of a celebrity:

Bing Video Search
Video Indexer
Custom Vision
Computer Vision

Tag a scene from a movie, and then search for movie scenes by using the tags:

Bing Video Search
Video Indexer
Custom Vision
Computer Vision

Box 1: Computer Vision

Azure's Computer Vision service provides developers with access to advanced algorithms that process images and return information.

Computer Vision Detect Faces: Detect faces in an image and provide information about each detected face.

Computer Vision returns the coordinates, rectangle, gender, and age for each detected face.

Computer Vision provides a subset of the Face service functionality. You can use the Face service for more detailed analysis, such as facial identification and pose detection.

Box 2: Bing Video Search

Search for videos and get comprehensive results

With Bing Video Search API v7, find videos across the web. Results provide useful metadata including creator, encoding format, video length, view count, improved & simplified paging, and more.

NEW QUESTION: 113

Your company plans to deploy several apps that will use Azure Cognitive Services APIs.

You need to recommend which Cognitive Services APIs must be used to meet the following requirements:

- * Must be able to identify inappropriate text and profanities in multiple languages.
- * Must be able to interpret user requests sent by using text input.
- * Must be able to identify named entities in text.

Which API should you recommend for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Must be able to identify inappropriate text and profanities in multiple languages:

	▼
Text Analytics	
Content Moderator	
Bing Visual Search	
Language Understanding (LUIS)	

Must be able to interpret user requests sent by using text input:

	▼
Text Analytics	
Content Moderator	
Bing Visual Search	
Language Understanding (LUIS)	

Must be able to identify named entities in text:

	▼
Text Analytics	
Content Moderator	
Bing Visual Search	
Language Understanding (LUIS)	



Answer:

Must be able to identify inappropriate text and profanities in multiple languages:

 Microsoft ▼

Text Analytics

Content Moderator

Bing Visual Search

Language Understanding (LUIS)

Must be able to interpret user requests sent by using text input:

▼

Text Analytics

Content Moderator

Bing Visual Search

Language Understanding (LUIS)

Must be able to identify named entities in text:

▼

Text Analytics

Content Moderator

Bing Visual Search

Language Understanding (LUIS)

Explanation

Must be able to identify inappropriate text and profanities in multiple languages:

▼

Text Analytics

Content Moderator

Bing Visual Search

Language Understanding (LUIS)

Must be able to interpret user requests sent by using text input:

▼

Text Analytics

Content Moderator

Bing Visual Search

Language Understanding (LUIS)

Must be able to identify named entities in text:

▼

Text Analytics

Content Moderator

Bing Visual Search

Language Understanding (LUIS)

Box 1: Content Moderator

The Azure Content Moderator API is a cognitive service that checks text, image, and video content for material that is potentially offensive, risky, or otherwise undesirable. When such material is found, the service applies appropriate labels (flags) to the content. Your app can then handle flagged content in order to comply with regulations or maintain the intended environment for users.

Box 2: Language Understanding (LUIS)

Designed to identify valuable information in conversations, LUIS interprets user goals (intents) and distills valuable information from sentences (entities), for a high quality, nuanced language model. LUIS integrates seamlessly with the Azure Bot Service, making it easy to create a sophisticated bot.

Box 3: Text Analytics

The Text Analytics API is a cloud-based service that provides advanced natural language processing over raw text, and includes four main functions: sentiment analysis, key phrase extraction, named entity recognition, and language detection.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/cognitive-services/content-moderator/overview>

<https://www.luis.ai/home>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/>

NEW QUESTION: 114

You plan to create a bot that will support five languages. The bot will be used by users located in three different countries. The bot will answer common customer questions. The bot will use Language Understanding (LUIS) to identify which skill to use and to detect the language of the customer.

You need to identify the minimum number of Azure resources that must be created for the planned bot.

How many QnA Maker, LUIS and Language Detection instances should you create? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

QnA Maker  Microsoft ▼

1
3
5

LUIS: ▼

1
3
5

Language Detection: ▼

1
3
5

Answer:

QnA Maker: ▼

1
3
5

LUIS: ▼

1
3
5

Language Detection: ▼

1
3
5

 Microsoft

Explanation

QnA Maker: 1 3 5

LUIS: 1 3 5

Language Detection: 1 3 5

QnA Maker: 5

If the user plans to support multiple languages, they need to have a new QnA Maker resource for each language.

LUIS: 5

If you need a multi-language LUIS client application such as a chatbot, you have a few options. If LUIS supports all the languages, you develop a LUIS app for each language. Each LUIS app has a unique app ID, and endpoint log. If you need to provide language understanding for a language LUIS does not support, you can use Microsoft Translator API to translate the utterance into a supported language, submit the utterance to the LUIS endpoint, and receive the resulting scores.

Language detection: 1

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis. This capability is useful for content stores that collect arbitrary text, where language is unknown. You can parse the results of this analysis to determine which language is used in the input document. The response also returns a score that reflects the confidence of the model. The score value is between 0 and 1.

The Language Detection feature can detect a wide range of languages, variants, dialects, and some regional or cultural languages. The exact list of languages for this feature isn't published.

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/language-support>

<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-language-support>

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language>

NEW QUESTION: 115

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.

Actions

Retrieve a token from the Cognitive Services endpoint

Generate new keys in the Cognitive Services resources

Generate new keys in Azure Key Vault

Update the custom application to use the new authorization

Retrieve a token from the Azure Key Vault endpoint

Answer Area

⬅

➡



⬆

⬇

Answer:

Actions

Retrieve a token from the Cognitive Services endpoint

Generate new keys in the Cognitive Services resources

Generate new keys in Azure Key Vault

Update the custom application to use the new authorization

Retrieve a token from the Azure Key Vault endpoint

Answer Area

Generate new keys in the Cognitive Services resources

Retrieve a token from the Cognitive Services endpoint

Update the custom application to use the new authorization


Explanation

Generate new keys in the Cognitive Services resources

Retrieve a token from the Cognitive Services endpoint

Update the custom application to use the new authorization

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>

NEW QUESTION: 116

Your company recently purchased several hundred hardware devices that contains sensors.

You need to recommend a solution to process the sensor data. The solution must provide the ability to write back configuration changes to the devices.

What should you include in the recommendation?

- A. Microsoft Azure Notification Hubs
- B. API apps in Microsoft Azure App Service
- C. Microsoft Azure Event Hubs
- D. Microsoft Azure IoT Hub

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 117

You are designing an AI solution that will analyze media data. The data will be stored in Azure Blob storage.

You need to ensure that the storage account is encrypted by using a key generated by the hardware security module (HSM) of your company.

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.

Actions

Enable encryption that uses customer-managed keys.

Upload a key to an Azure key vault.

Generate an encryption key.

Generate an access key.

Configure a service endpoint for the storage account.

Generate a shared access signature (SAS).

Answer Area



Microsoft



Answer:

Actions

Enable encryption that uses customer-managed keys.

Upload a key to an Azure key vault.

Generate an encryption key.

Generate an access key.

Configure a service endpoint for the storage account.

Generate a shared access signature (SAS).

Answer Area

Configure a service endpoint for the storage account.

Generate an encryption key.

Enable encryption that uses customer-managed keys.



Microsoft



Explanation

Answer Area

Configure a service endpoint for the storage account.

Generate an encryption key.

Enable encryption that uses customer-managed keys.

Microsoft

References:

<https://docs.microsoft.com/en-us/azure/storage/common/storage-encryption-keys-portal>

<https://docs.microsoft.com/en-us/azure/key-vault/key-vault-hsm-protected-keys>

NEW QUESTION: 118

You plan to use Azure Cognitive Services to provide the development team at your company with the ability to create intelligent apps without having direct AI or data science skills.

The company identifies the following requirements for the planned Cognitive Services deployment:

* Provide support for the following languages: English, Portuguese, and German.

* Perform text analytics to derive a sentiment score.

Which Cognitive Service service should you deploy for each requirement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Provide support for the following languages:
English, Portuguese, and German:


Text Analytics
Search services
Speech services

Perform text analytics to derive a sentiment score:

Speech APIs
Decision APIs
Language APIs

Microsoft

Answer:



Explanation



Box 1: Text Analytics

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.

Box 2: Language API

References:

<https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language>

<https://docs.microsoft.com/en-us/azure/azure-databricks/databricks-sentiment-analysis-cognitive-services>

NEW QUESTION: 119

Your company manages a sports team.

The company sets up a video booth to record messages for the team.

Before replaying the messages on a video screen, you need to generate captions for the messages and check the emotions in the video to ensure that only positive messages are played.

Which Azure Cognitive Services service should you use?

- A. language Understanding (LUIS)
- B. Custom Vision
- C. Speaker Recognition
- D. Video Indexer

Answer: A ([LEAVE A REPLY](#))

Valid AI-100 Dumps shared by PrepPdf.com for Helping Passing AI-100 Exam! PrepPdf.com now offer the **newest AI-100 exam dumps**, the PrepPdf.com AI-100 exam **questions have been updated** and **answers have been corrected** get the **newest** PrepPdf.com AI-100 dumps with Test Engine here: <https://www.preppdf.com/Microsoft/AI-100-prepaway-exam-dumps.html> (**163** Q&As Dumps, **40%OFF** Special Discount: **Exam-Tests**)