

AI-100.44q

Number: AI-100
Passing Score: 800
Time Limit: 120 min

AI-100



Website: <https://vceplus.com>
VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>
Facebook: <https://www.facebook.com/VCE.For.All.VN/>
Twitter : https://twitter.com/VCE_Plus

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

Designing and Implementing an Azure AI Solution

Question Set 1

QUESTION 1

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

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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>

QUESTION 2

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.



<https://www.vceplus.com/> **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

Correct Answer: BDE

Section: [none]

Explanation

Explanation/Reference:

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>

QUESTION 3

Your company recently deployed several hardware devices that contain sensors.

The sensors generate new data on an hourly basis. The data generated is stored on-premises and retained for several years.

During the past two months, the sensors generated 300 GB of data.

You plan to move the data to Azure and then perform advanced analytics on the data.

You need to recommend an Azure storage solution for the data.

Which storage solution should you recommend?

- A. Azure Queue storage
- B. Azure Cosmos DB
- C. Azure Blob storage
- D. Azure SQL Database

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

References: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/data-storage>



QUESTION 4

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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>

QUESTION 5

Your company has 1,000 AI developers who are responsible for provisioning environments in Azure.

You need to control the type, size, and location of the resources that the developers can provision.

What should you use?

- A. Azure Key Vault
- B. Azure service principals
- C. Azure managed identities
- D. Azure Security Center
- E. Azure Policy



Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

When an application needs access to deploy or configure resources through Azure Resource Manager in Azure Stack, you create a service principal, which is a credential for your application. You can then delegate only the necessary permissions to that service principal.

References: <https://docs.microsoft.com/en-us/azure/azure-stack/azure-stack-create-service-principals>

QUESTION 6

You are designing an AI solution in Azure that will perform image classification.

You need to identify which processing platform will provide you with the ability to update the logic over time. The solution must have the lowest latency for inferencing without having to batch.

Which compute target should you identify?

- A. graphics processing units (GPUs)
- B. field-programmable gate arrays (FPGAs)
- C. central processing units (CPUs)
- D. application-specific integrated circuits (ASICs)

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

FPGAs, such as those available on Azure, provide performance close to ASICs. They are also flexible and reconfigurable over time, to implement new logic.

Incorrect Answers:

D: ASICs are custom circuits, such as Google's TensorFlow Processor Units (TPU), provide the highest efficiency. They can't be reconfigured as your needs change.

References: <https://docs.microsoft.com/en-us/azure/machine-learning/service/concept-accelerate-with-fpgas>

QUESTION 7

You have a solution that runs on a five-node Azure Kubernetes Service (AKS) cluster. The cluster uses an N-series virtual machine.

An Azure Batch AI process runs once a day and rarely on demand.

You need to recommend a solution to maintain the cluster configuration when the cluster is not in use. The solution must not incur any compute costs.

What should you include in the recommendation?

- A. Downscale the cluster to one node
- B. Downscale the cluster to zero nodes

C. Delete the cluster

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

An AKS cluster has one or more nodes.

References: <https://docs.microsoft.com/en-us/azure/aks/concepts-clusters-workloads>

QUESTION 8

HOTSPOT

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.

Hot Area:

Answer Area

Location to store the metadata:	<input type="text"/>
	Azure Blob storage
	Azure Cosmos DB
	Azure File Storage
Virtual machine series to run the vision module:	<input type="text"/>
	A
	F
	NV
Virtual machine series to run the mathematical module:	<input type="text"/>
	A
	F
	NV

Correct Answer:

Answer Area

Location to store the metadata:	<div style="border: 1px solid #ccc; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between;">▼</div><div style="padding: 2px;"><p>Azure Blob storage</p><p>Azure Cosmos DB</p><p>Azure File Storage</p></div></div>
Virtual machine series to run the vision module:	<div style="border: 1px solid #ccc; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between;">▼</div><div style="padding: 2px;"><p>A</p><p>F</p><p>NV</p></div></div>
Virtual machine series to run the mathematical module:	<div style="border: 1px solid #ccc; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between;">▼</div><div style="padding: 2px;"><p>A</p><p>F</p><p>NV</p></div></div>

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Blob Storage

Containers and blobs support custom metadata, represented as HTTP headers.

Box 2: NV

The NV-series enables powerful remote visualisation workloads and other graphics-intensive applications backed by the NVIDIA Tesla M60 GPU.

Note: The N-series is a family of Azure Virtual Machines with GPU capabilities. GPUs are ideal for compute and graphics-intensive workloads, helping customers to fuel innovation through scenarios like high-end remote visualisation, deep learning and predictive analytics.

Box 3: F

F-series VMs feature a higher CPU-to-memory ratio. Example use cases include batch processing, web servers, analytics and gaming.

Incorrect:

A-series VMs have CPU performance and memory configurations best suited for entry level workloads like development and test.

References:

<https://azure.microsoft.com/en-in/pricing/details/virtual-machines/series/>

QUESTION 9

Your company has recently deployed 5,000 Internet-connected sensors for a planned AI solution.

You need to recommend a computing solution to perform a real-time analysis of the data generated by the sensors.

Which computing solution should you recommend?

- A. an Azure HDInsight Storm cluster
- B. Azure Notification Hubs
- C. an Azure HDInsight Hadoop cluster
- D. an Azure HDInsight R cluster



Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Azure HDInsight makes it easy, fast, and cost-effective to process massive amounts of data.

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

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

QUESTION 10

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.

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

References:

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

QUESTION 11

You have an AI application that uses keys in Azure Key Vault.

Recently, a key used by the application was deleted accidentally and was unrecoverable.

You need to ensure that if a key is deleted, it is retained in the key vault for 90 days.

Which two features should you configure? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. The expiration date on the keys
- B. Soft delete
- C. Purge protection
- D. Auditors
- E. The activation date on the keys

Correct Answer: BC

Section: [none]

Explanation

Explanation/Reference:

References:

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

QUESTION 12

You plan to build an application that will perform predictive analytics. Users will be able to consume the application data by using Microsoft Power BI or a custom website.

You need to ensure that you can audit application usage.

Which auditing solution should you use?

- A. Azure Storage Analytics
- B. Azure Application Insights
- C. Azure diagnostics logs
- D. Azure Active Directory (Azure AD) reporting

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

References: <https://docs.microsoft.com/en-us/azure/active-directory/reports-monitoring/concept-audit-logs>



QUESTION 13

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:



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 an 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.

QUESTION 1

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

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

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



Question Set 1

QUESTION 1

You have a database that contains sales data.

You plan to process the sales data by using two data streams named Stream1 and Stream2. Stream1 will be used for purchase order data. Stream2 will be used for reference data.

The reference data is stored in CSV files.

You need to recommend an ingestion solution for each data stream.

What two solutions should you recommend? Each correct answer is a complete solution.

NOTE: Each correct selection is worth one point.

- A. an Azure event hub for Stream1 and Azure Blob storage for Stream2
- B. Azure Blob storage for Stream1 and Stream2
- C. an Azure event hub for Stream1 and Stream2
- D. Azure Blob storage for Stream1 and Azure Cosmos DB for Stream2
- E. Azure Cosmos DB for Stream1 and an Azure event hub for Stream2

Correct Answer: AB

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Stream1 - Azure Event

Stream2 - Blob Storage

Azure Event Hubs is a highly scalable data streaming platform and event ingestion service, capable of receiving and processing millions of events per second. Event Hubs can process and store events, data, or telemetry produced by distributed software and devices. Data sent to an event hub can be transformed and stored using any real-time analytics provider or batching/storage adapters. Event Hubs provides publish-subscribe capabilities with low latency at massive scale, which makes it appropriate for big data scenarios.

Stream1, Stream2 - Blob Storage

Stream Analytics has first-class integration with Azure data streams as inputs from three kinds of resources:

Azure Event Hubs

Azure IoT Hub

Azure Blob storage

These input resources can live in the same Azure subscription as your Stream Analytics job or a different subscription.

References: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/real-time-ingestion>

QUESTION 2

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing an application that uses an Azure Kubernetes Service (AKS) cluster.

You are troubleshooting a node issue.

You need to connect to an AKS node by using SSH.

Solution: You create a managed identity for AKS, and then you create an SSH connection.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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>

QUESTION 3

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing an application that uses an Azure Kubernetes Service (AKS) cluster.

You are troubleshooting a node issue.

You need to connect to an AKS node by using SSH.

Solution: You add an SSH key to the node, and then you create an SSH connection.



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



Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

By default, SSH keys are generated when you create an AKS cluster. If you did not specify your own SSH keys when you created your AKS cluster, add your public SSH keys to the AKS nodes.

You also need to create an SSH connection to the AKS node.

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

QUESTION 4

You are developing a Computer Vision application.

You plan to use a workflow that will load data from an on-premises database to Azure Blob storage, and then connect to an Azure Machine Learning service.

What should you use to orchestrate the workflow?

- A. Azure Kubernetes Service (AKS)
- B. Azure Pipelines
- C. Azure Data Factory
- D. Azure Container Instances

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

With Azure Data Factory you can use workflows to orchestrate data integration and data transformation processes at scale. Build data integration, and easily transform and integrate big data processing and machine learning with the visual interface.

References: <https://azure.microsoft.com/en-us/services/data-factory/>



QUESTION 5

DRAG DROP

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.

Select and Place:

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



Correct Answer:

Actions

Create an Azure Stream Analytics Cloud job

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

Section: [none]

Explanation

Explanation/Reference:

Explanation:


Step 1: Create a storage container

ASA Edge jobs run in containers deployed to Azure IoT Edge devices.

Step 2: Create an Azure Stream Analytics Edge Job

Azure Stream Analytics (ASA) on IoT Edge empowers developers to deploy near-real-time analytical intelligence closer to IoT devices so that they can unlock the full value of device-generated data.

Scenario overview:



Step 3: Add the job to the IoT devices in IoT

References:

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

QUESTION 6

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

Correct Answer: B

Section: [none]

Explanation




Explanation/Reference:

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>

QUESTION 7

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

Explanation:

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>

QUESTION 8

You plan to deploy two AI applications named AI1 and AI2. The data for the applications will be stored in a relational database.

You need to ensure that the users of AI1 and AI2 can see only data in each user's respective geographic region. The solution must be enforced at the database level by using row-level security.

Which database solution should you use to store the application data?

- A. Microsoft SQL Server on a Microsoft Azure virtual machine
- B. Microsoft Azure Database for MySQL
- C. Microsoft Azure Data Lake Store
- D. Microsoft Azure Cosmos DB

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Row-level security is supported by SQL Server, Azure SQL Database, and Azure SQL Data Warehouse.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/security/row-level-security?view=sql-server-2017>

QUESTION 9

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing an application that uses an Azure Kubernetes Service (AKS) cluster.

You are troubleshooting a node issue.

You need to connect to an AKS node by using SSH.

Solution: You run the kubectl command, and then you create an SSH connection.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:



QUESTION 10

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

References: <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>

Question Set 1

QUESTION 1

You need to build an API pipeline that analyzes streaming data. The pipeline will perform the following:

- Visual text recognition
- Audio transcription
- Sentiment analysis
- Face detection

Which Azure Cognitive Services should you use in the pipeline?

- A. Custom Speech Service
- B. Face API
- C. Text Analytics
- D. Video Indexer



Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

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 described below:

- Visual text recognition (OCR): Extracts text that is visually displayed in the video.
- Audio transcription: Converts speech to text in 12 languages and allows extensions.
- Sentiment analysis: Identifies positive, negative, and neutral sentiments from speech and visual text.
- Face detection: Detects and groups faces appearing in the video.

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

QUESTION 2

You design an AI solution that uses an Azure Stream Analytics job to process data from an Azure IoT hub. The IoT hub receives time series data from thousands of IoT devices at a factory.

The job outputs millions of messages per second. Different applications consume the messages as they are available. The messages must be purged.

You need to choose an output type for the job.

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

- A. Azure Event Hubs
- B. Azure SQL Database
- C. Azure Blob storage
- D. Azure Cosmos DB

Correct Answer: D

Section: [none]

Explanation



Explanation/Reference:

Explanation:

Stream Analytics can target Azure Cosmos DB for JSON output, enabling data archiving and low-latency queries on unstructured JSON data.

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-documentdb-output>

QUESTION 3

HOTSPOT

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.

Hot Area:

Answer Area

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

	▼
Apache Kafka	
Apache Spark	
Hive	
Microsoft Machine	
Learning Server	

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

	▼
Apache Kafka	
Apache Spark	
Interactive Query	
Microsoft Machine	
Learning Server	

Correct Answer:

Answer Area

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

	▼
Apache Kafka	
Apache Spark	
Hive	
Microsoft Machine	
Learning Server	

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

	▼
Apache Kafka	
Apache Spark	
Interactive Query	
Microsoft Machine	
Learning Server	

Section: [none]

Explanation

Explanation/Reference:

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>

QUESTION 4

You need to deploy cognitive search.

You provision an Azure Search service.

What should you do next?

- A. Search by using the .NET SDK.
- B. Load data.
- C. Search by using the REST API.
- D. Create an index.

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

You create a data source, a skillset, and an index. These three components become part of an indexer that pulls each piece together into a single multi-phased operation.

Note: At the start of the pipeline, you have unstructured text or non-text content (such as image and scanned document JPEG files). Data must exist in an Azure data storage service that can be accessed by an indexer. Indexers can "crack" source documents to extract text from source data.

References: <https://docs.microsoft.com/en-us/azure/search/cognitive-search-tutorial-blob>

QUESTION 5

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 6

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.



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

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

Correct Answer: E

Section: [none]

Explanation

Explanation/Reference:

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.

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>

QUESTION 7

Your company plans to deploy an AI solution that processes IoT data in real-time.

You need to recommend a solution for the planned deployment that meets the following requirements:

- Sustain up to 50 Mbps of events without throttling. ▪
- Retain data for 60 days.

What should you recommend?

A. Apache Kafka

- B. Microsoft Azure IoT Hub
- C. Microsoft Azure Data Factory
- D. Microsoft Azure Machine Learning

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Apache Kafka is an open-source distributed streaming platform that can be used to build real-time streaming data pipelines and applications.

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

QUESTION 8

You are designing a solution that will use the Azure Content Moderator service to moderate user-generated content.

You need to moderate custom predefined content without repeatedly scanning the collected content.

Which API should you use?

- A. Term List API
- B. Text Moderation API
- C. Image Moderation API
- D. Workflow API

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

Explanation:

The default global list of terms in Azure Content Moderator is sufficient for most content moderation needs. However, you might need to screen for terms that are specific to your organization. For example, you might want to tag competitor names for further review.

Use the List Management API to create custom lists of terms to use with the Text Moderation API. The Text - Screen operation scans your text for profanity, and also compares text against custom and shared blacklists.

Incorrect Answers:

B: Use the Text Moderation API in Azure Content Moderator to scan your text content. The operation scans your content for profanity, and compares the content against custom and shared blacklists.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/try-terms-list-api>

QUESTION 9

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

Correct Answer: E

Section: [none]

Explanation



Explanation/Reference:

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

QUESTION 10

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

Correct Answer: A
Section: [none]
Explanation

Explanation/Reference:

References:

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

Question Set 1

QUESTION 1

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 create several AI models in Azure Machine Learning Studio.

You deploy the models to a production environment.

You need to monitor the compute performance of the models.

Solution: You enable AppInsights diagnostics.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B
Section: [none]
Explanation

Explanation/Reference:

Explanation:

You need to enable Model data collection.

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



QUESTION 2

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

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

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>

QUESTION 3

You create an Azure Machine Learning Studio experiment.

You plan to publish the experiment as a Machine Learning Web service.

You need to ensure that you can consume the web service from Microsoft Excel spreadsheets.

What should you use?

- A. a Batch Execution Service (BES) and an API key
- B. a Batch Execution Service (BES) and an Azure managed identity



- C. a Request-Response Service (RRS) and an Azure managed identity
- D. a Request-Response Service (RRS) and an API key

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

Explanation:

Steps to Add a New web service

1. Deploy a web service or use an existing Web service.
2. Click Consume.
3. Look for the Basic consumption info section. Copy and save the Primary Key and the Request-Response URL.
4. In Excel, go to the Web Services section (if you are in the Predict section, click the back arrow to go to the list of web services).
5. Click Add Web Service.
6. Paste the URL into the Excel add-in text box labeled URL.
7. Paste the API/Primary key into the text box labeled API key.
8. Click Add.

References: <https://docs.microsoft.com/en-us/azure/machine-learning/studio/excel-add-in-for-web-services>

QUESTION 4

DRAG DROP

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.

Select and Place:


Actions

- Start the container
- Create a container image
- Create an Azure Batch AI account
- Get the http endpoint of the web service
- Register the container image
- Train the model

Answer Area




Correct Answer:



Section: [none]
Explanation

Explanation/Reference:
Explanation:

The following diagram illustrates the complete deployment workflow:



The deployment workflow includes the following steps:

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

References: <https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/concept-model-management-and-deployment#step-3-deploy-image>

QUESTION 5

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

Correct Answer: B


Section: [none]

Explanation

Explanation/Reference:

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/>

QUESTION 6

You have Azure IoT Edge devices that generate measurement data from temperature sensors. The data changes very slowly.

You need to analyze the data in a temporal two-minute window. If the temperature rises five degrees above a limit, an alert must be raised. The solution must minimize the development of custom code.

What should you use?

- A. A Machine Learning model as a web service

- B. an Azure Machine Learning model as an IoT Edge module
- C. Azure Stream Analytics as an IoT Edge module
- D. Azure Functions as an IoT Edge module

Correct Answer: C

Section: [none]

Explanation

Explanation/Reference:

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

QUESTION 7

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 Monitor for containers.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: [none]

Explanation

Explanation/Reference:

QUESTION 8

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

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>

QUESTION 9

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

Correct Answer: A

Section: [none]

Explanation

Explanation/Reference:

References: <https://docs.microsoft.com/en-us/azure/bot-service/bot-service-manage-analytics?view=azure-bot-service-4.0>

QUESTION 10

You have an Azure Machine Learning experiment that must comply with GDPR regulations. You need to track compliance of the experiment and store documentation about the experiment.

What should you use?

- A. Azure Table storage
- B. Azure Security Center
- C. An Azure Log Analytics workspace
- D. Compliance Manager

Correct Answer: D

Section: [none]

Explanation

Explanation/Reference:

References: <https://azure.microsoft.com/en-us/blog/new-capabilities-to-enable-robust-gdpr-compliance/>



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