

AI-100.VCEplus.premium.exam.56q

Number: AI-100
Passing Score: 800
Time Limit: 120 min
File Version: 1.0



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

AI-100

Designing and Implementing an Azure AI Solution



Sections

1. Analyze solution requirements
2. Design solutions
3. Integrate AI models into solutions
4. Deploy and manage solutions

Exam A

QUESTION 1
HOTSPOT

You are designing an application to parse images of business forms and upload the data to a database. The upload process will occur once a week.

You need to recommend which services to use for the application. The solution must minimize infrastructure costs.

Which services should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Parse the images:

	▼
Azure Bot Service	
Azure Cognitive Services	
Azure Linguistic Analysis API	

Upload the data to the database:

	▼
Azure API Apps	
Azure Batch AI	
Azure Data Factory	
Azure Functions	



Correct Answer:

Answer Area

Parse the images:

	▼
Azure Bot Service	
Azure Cognitive Services	
Azure Linguistic Analysis API	

Upload the data to the database:

	▼
Azure API Apps	
Azure Batch AI	
Azure Data Factory	
Azure Functions	

Section: Analyze solution requirements
Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Cognitive Services

Azure Cognitive Services include image-processing algorithms to smartly identify, caption, index, and moderate your pictures and videos.

Not: Azure Linguistic Analytics API, which provides advanced natural language processing over raw text.

Box 2: Azure Data Factory

The Azure Data Factory (ADF) is a service designed to allow developers to integrate disparate data sources. It is a platform somewhat like SSIS in the cloud to manage the data you have both on-prem and in the cloud.

It provides access to on-premises data in SQL Server and cloud data in Azure Storage (Blob and Tables) and Azure SQL Database.

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

<https://www.jamesserra.com/archive/2014/11/what-is-azure-data-factory/>

QUESTION 2

HOTSPOT

You plan to deploy an Azure Data Factory pipeline that will perform the following:

- Move data from on-premises to the cloud. ▪
- Consume Azure Cognitive Services APIs.

You need to recommend which technologies the pipeline should use. The solution must minimize custom code.

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

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Move data from on-premises to the cloud:	<input type="text"/>	▼
	Azure-SSIS Integration Runtime	
	Common language runtime (CLR)	
	Integration Runtime (IR)	
	Self-hosted integration runtime	
Consume Cognitive Services APIs:	<input type="text"/>	▼
	Azure API Management	
	Azure Logic Apps	
	WebJobs in Azure	

Correct Answer:

Answer Area

Move data from on-premises to the cloud:

	▼
Azure-SSIS Integration Runtime	
Common language runtime (CLR)	
Integration Runtime (IR)	
Self-hosted integration runtime	

Consume Cognitive Services APIs:

	▼
Azure API Management	
Azure Logic Apps	
WebJobs in Azure	

Section: Analyze solution requirements

Explanation

Explanation/Reference:

Explanation:

Box 1: Self-hosted Integration Runtime

A self-hosted IR is capable of running copy activity between a cloud data stores and a data store in private network.

Not Azure-SSIS Integration Runtime, as you would need to write custom code.



Box 2: Azure Logic Apps

Azure Logic Apps helps you orchestrate and integrate different services by providing 100+ ready-to-use connectors, ranging from on-premises SQL Server or SAP to Microsoft Cognitive Services.

Incorrect:

Not Azure API Management: Use Azure API Management as a turnkey solution for publishing APIs to external and internal customers.

References: <https://docs.microsoft.com/en-us/azure/data-factory/concepts-integration-runtime>

<https://docs.microsoft.com/en-us/azure/logic-apps/logic-apps-examples-and-scenarios>

QUESTION 3

HOTSPOT

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.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Analyze solution requirements

Explanation

Explanation/Reference:

Explanation:

Box 1: Azure Bot Service

Box 2: Computer Vision

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

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

QUESTION 4

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

Correct Answer: C

Section: Analyze solution requirements

Explanation

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>

QUESTION 5

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: Analyze solution requirements

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 6 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

Correct Answer: BDE

Section: Analyze solution requirements

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 7 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: Analyze solution requirements

Explanation

Explanation/Reference:

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

QUESTION 8

You plan to design an application that will use data from Azure Data Lake and perform sentiment analysis by using Azure Machine Learning algorithms.

The developers of the application use a mix of Windows- and Linux-based environments. The developers contribute to shared GitHub repositories.

You need all the developers to use the same tool to develop the application.

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

- A. Microsoft Visual Studio Code

- B. Azure Notebooks
- C. Azure Machine Learning Studio
- D. Microsoft Visual Studio

Correct Answer: C

Section: Analyze solution requirements

Explanation

Explanation/Reference:

References: <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/machine-learning/studio/algorithm-choice.md>

QUESTION 9

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: Analyze solution requirements

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 10 You deploy an infrastructure for a big data workload.

You need to run Azure HDInsight and Microsoft Machine Learning Server. You plan to set the RevoScaleR compute contexts to run x function calls in parallel.

What are three compute contexts that you can use for Machine Learning Server? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. SQL
- B. Spark
- C. local parallel
- D. HBase
- E. local sequential

Correct Answer: ABC

Section: Analyze solution requirements

Explanation

Explanation/Reference:

Explanation:

Remote computing is available for specific data sources on selected platforms. The following tables document the supported combinations.

- RxInSqlServer, sqlserver: Remote compute context. Target server is a single database node (SQL Server 2016 R Services or SQL Server 2017 Machine Learning Services). Computation is parallel, but not distributed.
- RxSpark, spark: Remote compute context. Target is a Spark cluster on Hadoop.
- RxLocalParallel, localpar: Compute context is often used to enable controlled, distributed computations relying on instructions you provide rather than a built-in scheduler on Hadoop. You can use compute context for manual distributed computing.

References: <https://docs.microsoft.com/en-us/machine-learning-server/r/concept-what-is-compute-context>

QUESTION 11 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: Analyze solution requirements

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 12 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: Analyze solution requirements

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 13 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: Analyze solution requirements

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 14

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:

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

Correct Answer:

Answer Area

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

Section: Analyze solution requirements
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 15 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: Analyze solution requirements

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 16

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: Design solutions
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 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, 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 change the permissions of the AKS resource group, and then you create an SSH connection.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B
Section: Design solutions **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 18

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.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A
Section: Design solutions

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 19 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: Design solutions 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 20**

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: Design solutions Explanation

Explanation/Reference:

Explanation:

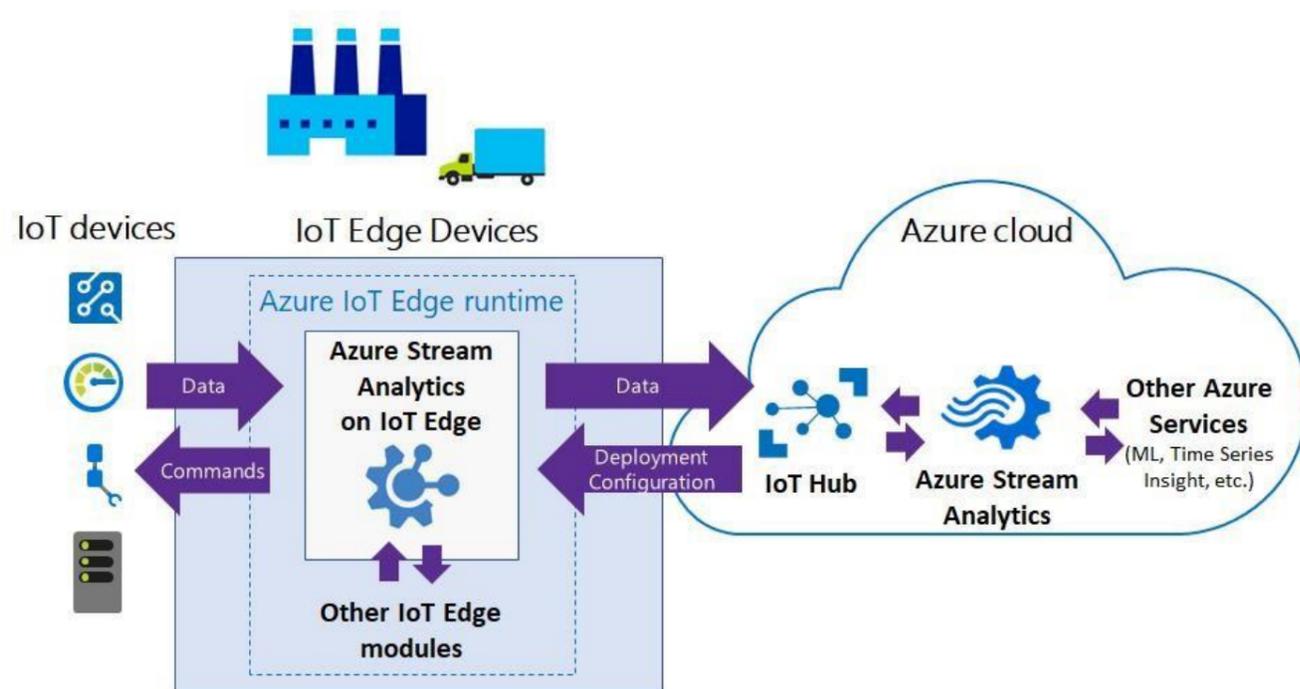
Step 1: Create a storage container

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 21
HOTSPOT

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.

Hot Area:

Answer Area

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



Correct Answer:

Answer Area

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

Section: Design solutions Explanation

Explanation/Reference:

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>

QUESTION 22

DRAG DROP

You need to build an AI 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.

Select and Place:



Actions

Create stream inputs and outputs

Create a new experiment

Build an experiment

Build a notebook

Create an Azure Machine Learning Studio workspace

Create a new notebook

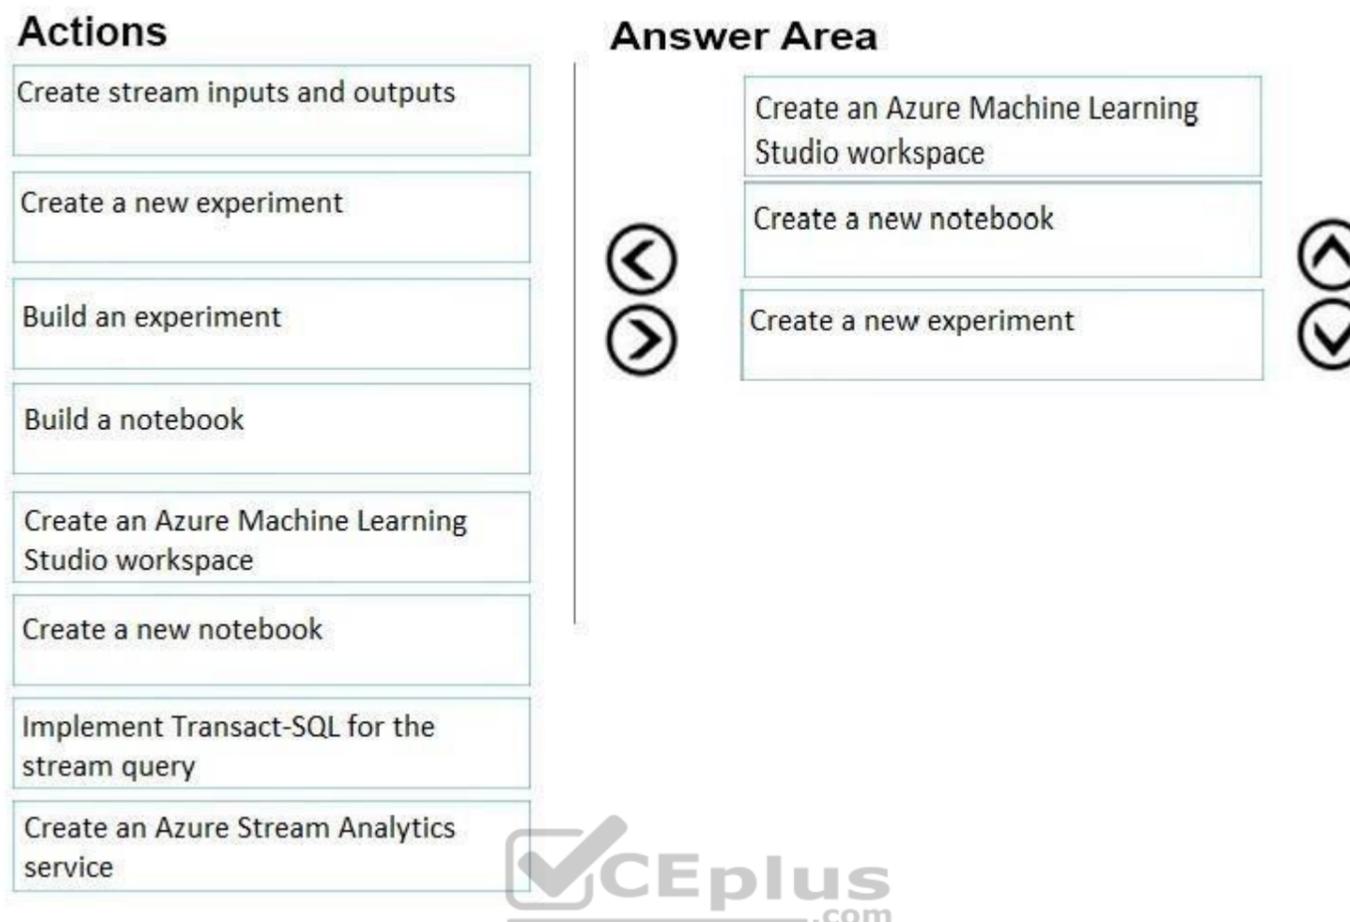
Implement Transact-SQL for the stream query

Create an Azure Stream Analytics service

Answer Area



Correct Answer:



Section: Design solutions Explanation

Explanation/Reference:

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

1. 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>

QUESTION 23

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

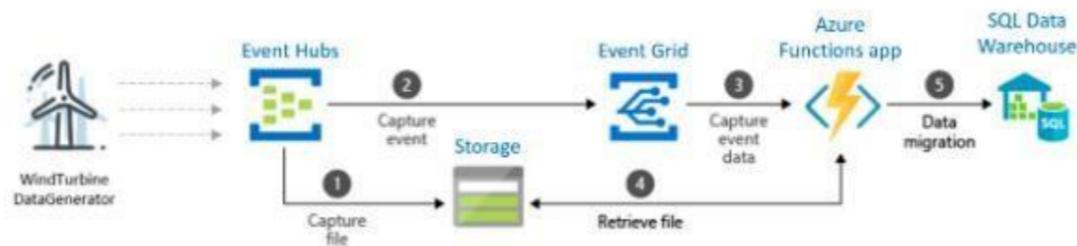
Section: Design solutions Explanation

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 24

HOTSPOT

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.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Design solutions Explanation

Explanation/Reference:

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. Bidirectional 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 Azure Machine Learning service

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

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-machine-learning-overview>

QUESTION 25 HOTSPOT

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.

Hot Area:

Answer Area

To connect to the on-premises data:	<div style="border: 1px solid gray; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between;">▼</div><div style="padding: 2px;"><p>A point-to-site VPN connection</p><p>A site-to-site VPN connection</p><p>Azure App Service Hybrid Connections</p></div></div>
To orchestrate the phases:	<div style="border: 1px solid gray; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between;">▼</div><div style="padding: 2px;"><p>A Machine Learning experiment</p><p>Azure Machine Learning Studio</p><p>Machine Learning pipelines</p></div></div>
To control the orchestration:	<div style="border: 1px solid gray; padding: 2px;"><div style="background-color: #f0f0f0; padding: 2px; display: flex; justify-content: space-between;">▼</div><div style="padding: 2px;"><p>Azure Automation</p><p>Azure Databricks</p><p>Azure Notebooks</p></div></div>

Correct Answer:

Answer Area

To connect to the on-premises data:	<div style="border: 1px solid gray; padding: 2px;">▼</div> <div style="border: 1px solid gray; padding: 2px;">A point-to-site VPN connection</div> <div style="border: 1px solid gray; padding: 2px;">A site-to-site VPN connection</div> <div style="border: 1px solid gray; padding: 2px; background-color: #e0f2f1;">Azure App Service Hybrid Connections</div>
To orchestrate the phases:	<div style="border: 1px solid gray; padding: 2px;">▼</div> <div style="border: 1px solid gray; padding: 2px;">A Machine Learning experiment</div> <div style="border: 1px solid gray; padding: 2px;">Azure Machine Learning Studio</div> <div style="border: 1px solid gray; padding: 2px; background-color: #e0f2f1;">Machine Learning pipelines</div>
To control the orchestration:	<div style="border: 1px solid gray; padding: 2px;">▼</div> <div style="border: 1px solid gray; padding: 2px;">Azure Automation</div> <div style="border: 1px solid gray; padding: 2px; background-color: #e0f2f1;">Azure Databricks</div> <div style="border: 1px solid gray; padding: 2px;">Azure Notebooks</div>

Section: Design solutions Explanation

Explanation/Reference:

Explanation:

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 largescale 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>

QUESTION 26 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: Design solutions 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 27 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



Correct Answer: CE

Section: Design solutions Explanation

Explanation/Reference:

Explanation:

With Azure Data Lake Store (ADLS) serving as the hyper-scale storage layer and HDInsight serving as the Hadoop-based compute engine services. It can be used for prepping large amounts of data for insertion into a Data Warehouse

References: <https://www.blue-granite.com/blog/azure-data-lake-analytics-holds-a-unique-spot-in-the-modern-data-architecture>

QUESTION 28

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: Design solutions 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 29 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

Correct Answer: B

Section: Design solutions Explanation

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>



QUESTION 30

HOTSPOT

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.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

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

Section: Design solutions Explanation

Explanation/Reference:

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>

QUESTION 31 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: Design solutions 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 32 You have thousands of images that contain text.

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

Which Azure Cognitive Services service should you use?

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

Correct Answer: D



Section: Integrate AI models into solutions**Explanation****Explanation/Reference:**

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.

Incorrect Answers:

A: Use Content Moderator's machine-assisted image moderation and human-in-the-loop Review tool to moderate images for adult and racy content. Scan images for text content and extract that text, and detect faces. You can match images against custom lists, and take further action.

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>

QUESTION 33

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: Integrate AI models into solutions**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 34

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: Integrate AI models into solutions

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 35

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: Integrate AI models into solutions

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 36 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: Integrate AI models into solutions

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 37 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: B

Section: Integrate AI models into solutions

Explanation

Explanation/Reference:

QUESTION 38

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

Correct Answer: E

Section: Integrate AI models into solutions

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 39 You create an Azure Cognitive Services resource.

You develop needs to be able to retrieve the keys used by the resource. The solution must use the principle of least privilege.

What is the best role to assign to the developer? More than one answer choice may achieve the goal.

- A. Security Manager
- B. Security Reader
- C. Cognitive Services Contributor
- D. Cognitive Services User

Correct Answer: D

Section: Integrate AI models into solutions

Explanation

Explanation/Reference:

Explanation:

The Cognitive Services User lets you read and list keys of Cognitive Services.

References: <https://docs.microsoft.com/en-us/azure/role-based-access-control/built-in-roles>



QUESTION 40 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: Integrate AI models into solutions

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 41

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: Integrate AI models into solutions

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 42

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

Correct Answer: B

Section: Deploy and manage solutions **Explanation**

Explanation/Reference:

Explanation:

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>

QUESTION 43

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

Correct Answer: A

Section: Deploy and manage solutions **Explanation**

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>

QUESTION 44

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

Correct Answer: B

Section: Deploy and manage solutions **Explanation**

Explanation/Reference:

Explanation:
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>

QUESTION 45

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 Model data collection.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: Deploy and manage solutions **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 46

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: Deploy and manage solutions **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 47

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 write a custom scoring script.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: Deploy and manage solutions 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 48 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: Deploy and manage solutions 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 49 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: Deploy and manage solutions **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 50

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	Answer Area
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	

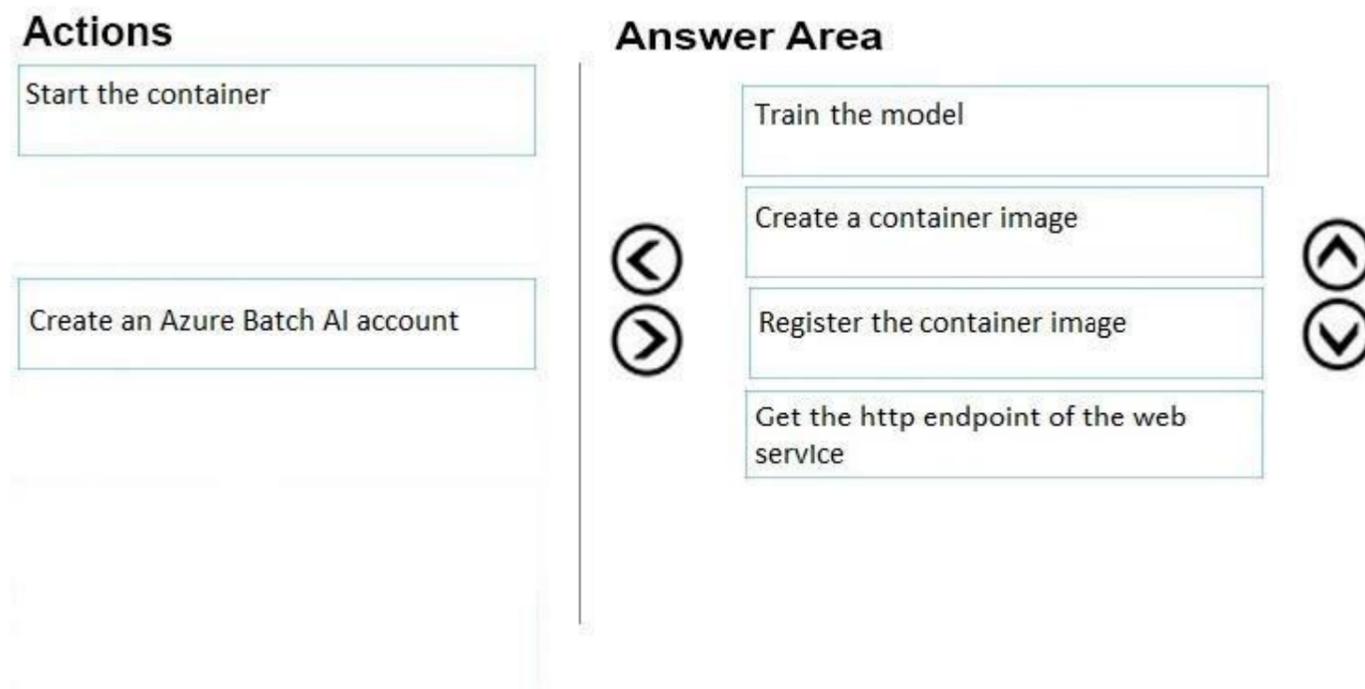
←

→

↑

↓

Correct Answer:

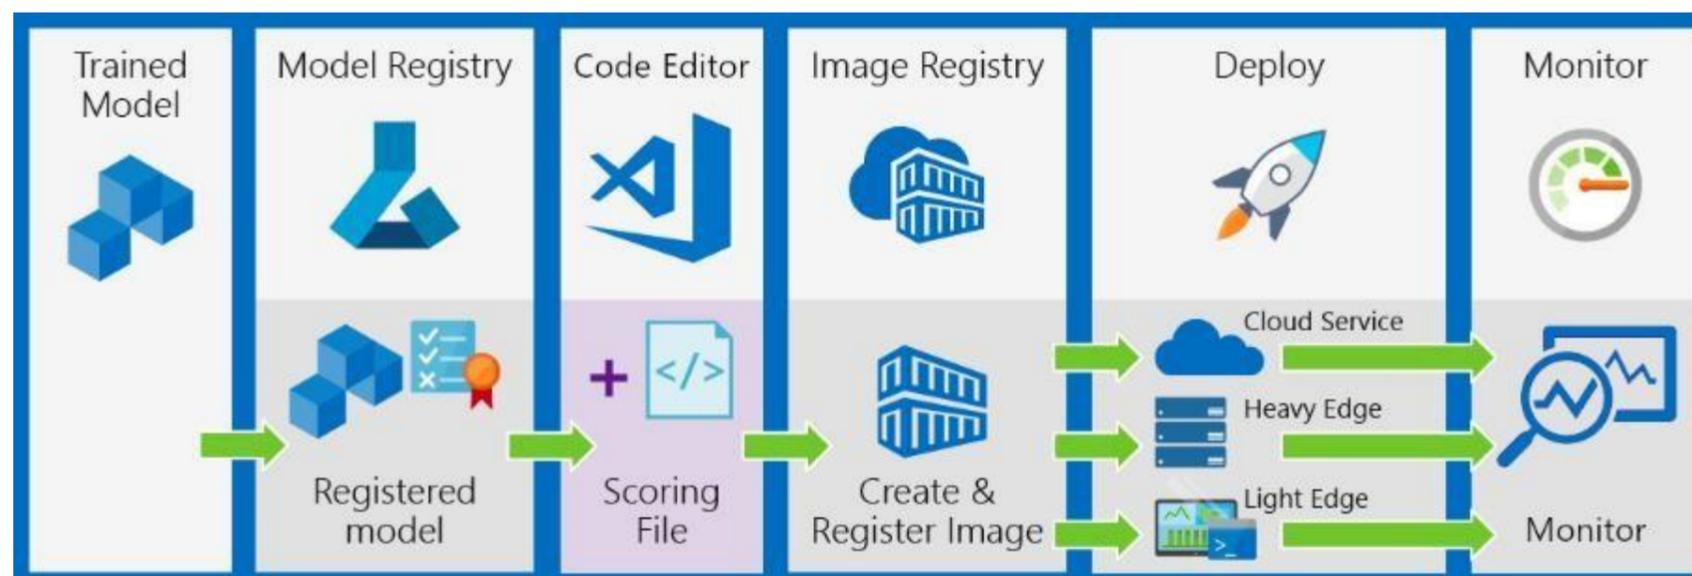


Section: Deploy and manage solutions Explanation

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 51 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?

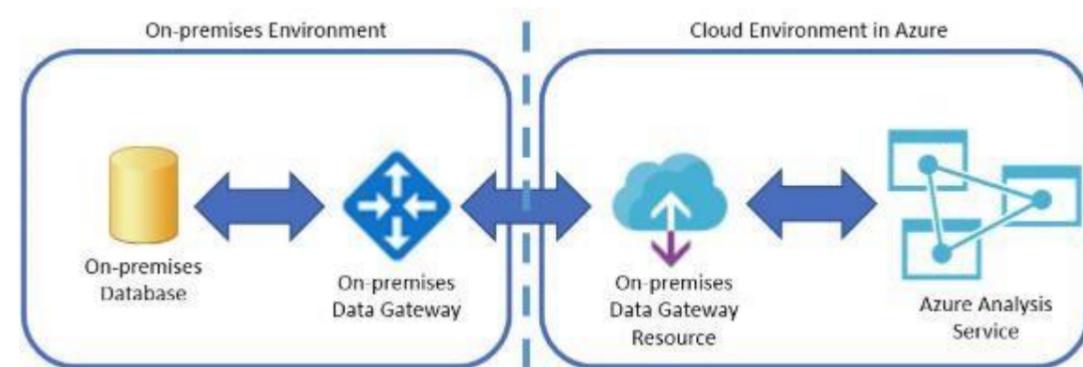
- A. a site-to-site VPN
- B. a data gateway
- C. Azure Data Factory
- D. a network gateway

Correct Answer: B

Section: Deploy and manage solutions 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 52

DRAG DROP

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.

Select and Place:

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



Correct Answer:

Actions

- Generate new keys in the Cognitive Services resources
- Retrieve a token from the Cognitive Services endpoint
- 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

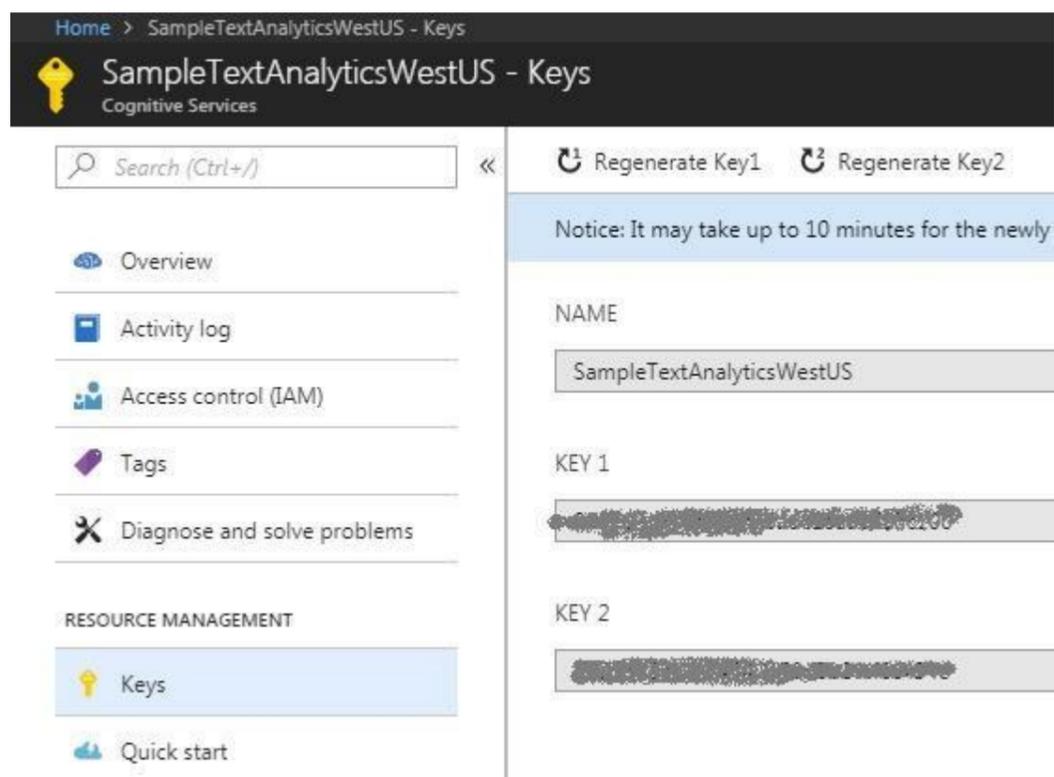


Section: Deploy and manage solutions Explanation

Explanation/Reference:

Explanation:

Step 1: Generate new keys in the Cognitive Service resources



Step 2: Retrieve a token from the Cognitive Services endpoint Step 3: Update the custom application to use the new authorization

Each request to an Azure Cognitive Service must include an authentication header. This header passes along a subscription key or access token, which is used to validate your subscription for a service or group of services.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/authentication>



QUESTION 53

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.

Select and Place:

Actions

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

Answer Area

- IT security team: Permission
- Applications: Permission

Correct Answer:

Actions

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

Answer Area

- IT security team: Secrets: all
- Applications: Secrets: get

Section: Deploy and manage solutions Explanation

Explanation/Reference:

Explanation:

Incorrect Answers:

Not Keys as they are used for encryption only.

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

QUESTION 54 A data scientist deploys a deep learning model on an Fsv2 virtual machine.

Data analysis is slow.

You need to recommend which virtual machine series the data scientist must use to ensure that data analysis occurs as quickly as possible.

Which series should you recommend?

- A. ND
- B. B
- C. DC
- D. Ev3

Correct Answer: A

Section: Deploy and manage solutions **Explanation**

Explanation/Reference:

Explanation:

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.

The ND-series is focused on training and inference scenarios for deep learning. It uses the NVIDIA Tesla P40 GPUs. The latest version - NDv2 - features the NVIDIA Tesla V100 GPUs.

References:

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



QUESTION 55

DRAG DROP

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 latest 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.

Select and Place:

Services

Azure Functions

Azure IoT Hub

Azure IoT Edge

Azure Logic Apps

Answer Area

Use a video stream to detect anomalies at the location:

Service

Send the pictures and location information to Azure:

Service

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

Service

Correct Answer:

Services

Azure IoT Hub

Answer Area

Use a video stream to detect anomalies at the location:

Azure IoT Edge

Send the pictures and location information to Azure:

Azure Functions

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

Azure Logic Apps

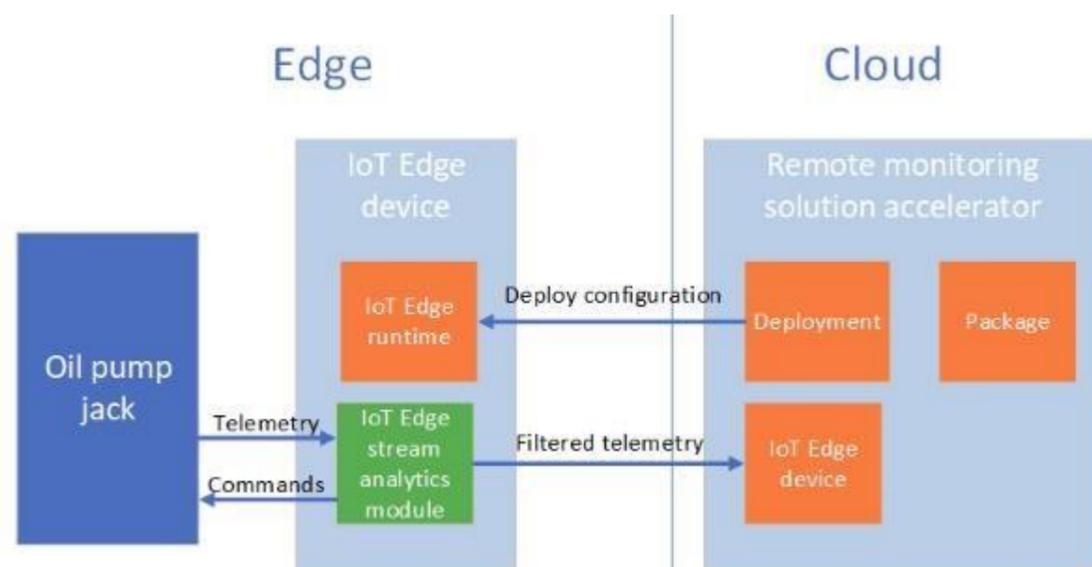
Section: Deploy and manage solutions Explanation

Explanation/Reference:

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>

QUESTION 56

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: Deploy and manage solutions Explanation

Explanation/Reference:

References:

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