

AI-100 Dumps

Designing and Implementing an Azure AI Solution

<https://www.certleader.com/AI-100-dumps.html>



NEW QUESTION 1

- (Exam Topic 1)

Which RBAC role should you assign to the KeyManagers group?

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

Answer: A

Explanation:

References:

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

NEW QUESTION 2

- (Exam Topic 1)

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

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

Actions

- On the page where you want Butler to be used, paste the embed code of the new Bookings app.
- From the Channels settings of Butler, copy the secret key and embed code for the Skype for Business channel.
- Exchange the secret for a token by connecting to <https://directline.botframework.com/api/tokens>.
- From the Channels settings of Butler, retrieve the secret key and embed code for the Web Chat channel.
- Replace s=YOUR_SECRET_HERE with t= followed by the token.
- Exchange the secret for a token by connecting to <https://webchat.botframework.com/api/tokens>.

Answer Area

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-service-channel-connect-webchat?view=azure-bot-servic>

NEW QUESTION 3

- (Exam Topic 2)

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

Answer: D

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>

NEW QUESTION 4

- (Exam Topic 2)

Your company plans to implement an AI solution that will analyse data from IoT devices.

Data from the devices will be analysed in real time. The results of the analysis will be stored in a SQL database.

You need to recommend a data processing solution that uses the Transact-SQL language. Which data processing solution should you recommend?

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

Answer: A

Explanation:

References:

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

NEW QUESTION 5

- (Exam Topic 2)

You create an Azure Cognitive Services resource.

A data scientist needs to call the resource from Azure Logic Apps.

Which two values should you provide to the data scientist? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. endpoint URL
- B. resource name
- C. access key
- D. resource group name
- E. subscription ID

Answer: DE

Explanation:

References:

<https://social.technet.microsoft.com/wiki/contents/articles/36074.logic-apps-with-azure-cognitive-service.aspx>

NEW QUESTION 6

- (Exam Topic 2)

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

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

NOTE: Each correct selection is worth one point.

The screenshot shows a question interface with two dropdown menus. The first dropdown is labeled 'JSON documents:' and the second is labeled 'PDF documents:'. Both dropdowns have a list of options: Azure Blob storage, Azure Cosmos DB, Azure Data Lake, and Azure Table storage. The 'JSON documents:' dropdown is currently open, showing the list of options.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

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

NEW QUESTION 7

- (Exam Topic 2)

You have thousands of images that contain text.

You need to process the text from the images into a machine-readable character stream. Which Azure Cognitive Services service should you use?

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 8

- (Exam Topic 2)

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

You need to recommend a solution for storing the pictures. The solution must minimize costs. Which storage solution should you recommend?

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 9

- (Exam Topic 2)

You plan to deploy Azure IoT Edge devices that will each store more than 10,000 images locally and classify the images by using a Custom Vision Service classifier. Each image is approximately 5 MB.

You need to ensure that the images persist on the devices for 14 days. What should you use?

- A. the device cache
- B. Azure Blob storage on the IoT Edge devices
- C. Azure Stream Analytics on the IoT Edge devices
- D. Microsoft SQL Server on the IoT Edge devices

Answer: B

Explanation:

References:

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

NEW QUESTION 10

- (Exam Topic 2)

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 diagnostic logs
- D. Azure Active Directory (Azure AD) reporting

Answer: D

Explanation:

References:

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

NEW QUESTION 10

- (Exam Topic 2)

You need to design the workflow for an Azure Machine Learning solution. The solution must meet the following requirements:

Retrieve data from file shares, Microsoft SQL Server databases, and Oracle databases that in an on-premises network.

Use an Apache Spark job to process data stored in an Azure SQL Data Warehouse database.

Which service should you use to meet each requirement? To answer, drag the appropriate services to the correct requirements. Each service may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

The screenshot shows a drag-and-drop interface for designing an Azure Machine Learning workflow. On the left, under the heading 'Services', there are four service tiles: 'Azure Data Factory', 'Azure Data Bricks', 'Azure Logic Apps', and 'Azure Stream Analytics'. On the right, under the heading 'Answer Area', there are two requirements listed: 'Retrieve the data:' and 'Process the data:'. Next to each requirement is a box labeled 'Service' where a service can be selected.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio/use-data-from-an-on-premises-sql-server> <https://docs.microsoft.com/en-in/azure/azure-databricks/what-is-azure-databricks>

NEW QUESTION 15

- (Exam Topic 2)

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

Answer: C**Explanation:**

References:

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

NEW QUESTION 19

- (Exam Topic 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 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

Answer: B**NEW QUESTION 20**

- (Exam Topic 2)

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

Which service should you include in the design?

- A. Azure Event Hubs
- B. Azure Data Factory
- C. Azure Stream Analytics
- D. Apache Kafka

Answer: D**NEW QUESTION 21**

- (Exam Topic 2)

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

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

What are three possible outputs? Each correct answer presents a complete solution. NOTE: Each correct selection is worth one point.

- A. a Hive table in Azure HDInsight
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. Azure Blob storage
- E. Azure Redis Cache

Answer: BCD**Explanation:**

References:

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

NEW QUESTION 23

- (Exam Topic 2)

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

Answer: C

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>

NEW QUESTION 28

- (Exam Topic 2)

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

You discover that occasionally and unpredictably, the application requires more than 32 nodes. You need to recommend a solution to handle the unpredictable application load.

Which scaling method should you recommend?

- A. horizontal pod autoscaler
- B. cluster autoscaler
- C. manual scaling
- D. Azure Container Instances

Answer: B

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>

NEW QUESTION 33

- (Exam Topic 2)

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

You need to design a solution to anonymize the data before the data is sent to the IoT hub.

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

Actions	Answer Area
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	

- A. Mastered
- B. Not Mastered

Answer: A


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>

NEW QUESTION 35

- (Exam Topic 2)

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

Answer: D

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>

NEW QUESTION 40

- (Exam Topic 2)

You are developing an application that will perform clickstream analysis. The application will ingest and analyze millions of messages in the real time.

You need to ensure that communication between the application and devices is bidirectional.

What should you use for data ingestion and stream processing? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Data ingestion:

<div>▼</div> <div>Azure Event Hubs</div> <div>Azure IoT Hub</div> <div>Azure Queue storage</div>
--

Stream processing:

<div>▼</div> <div>Azure HDInsight with Apache HBase</div> <div>Azure HDInsight with Apache Storm</div> <div>Azure HDInsight with Azure Machine Learning service</div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure IoT Hub

Azure IoT Hub is the cloud gateway that connects IoT devices to gather data and drive business insights and automation. In addition, IoT Hub includes features that enrich the relationship between your devices and your backend systems. Bi-directional communication capabilities mean that while you receive data from devices you can also send commands and policies back to devices.

Note on why not Azure Event Hubs: An Azure IoT Hub contains an Event Hub and hence essentially is an Event Hub plus additional features. An important additional feature is that an Event Hub can only receive messages, whereas an IoT Hub additionally can also send messages to individual devices. Further, an Event Hub has access security on hub level, whereas an IoT Hub is aware of the individual devices and can grant and revoke access on device level.

Box 2: Azure Hdinsight with 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>

NEW QUESTION 45

- (Exam Topic 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 deploying an Azure Machine Learning model to an Azure Kubernetes Service (AKS) container. You need to monitor the accuracy of each run of the model.
Solution: You configure Azure Application Insights.
Does this meet the goal?

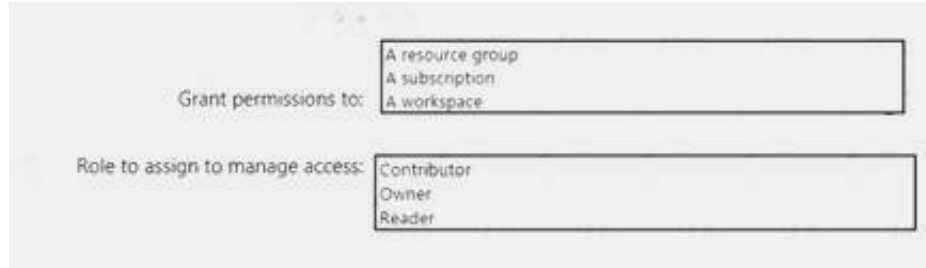
- A. Yes
- B. No

Answer: A

NEW QUESTION 47

- (Exam Topic 2)

You need to configure security for an Azure Machine Learning service used by groups of data scientists. The groups must have access to only their own experiments and must be able to grant permissions to the members of their team.
What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-us/machine-learning-server/operationalize/configure-roles#how-are-roles-assigne> <https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-assign-roles>

NEW QUESTION 50

- (Exam Topic 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 in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You are developing an application that uses an Azure Kubernetes Service (AKS) cluster. You are troubleshooting a node issue. You need to connect to an AKS node by using SSH.
Solution: You run the kubect1 command, and then you create an SSH connection.
Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 55

- (Exam Topic 2)

Your company recently purchased several hundred hardware devices that contains sensors. You need to recommend a solution to process the sensor data. The solution must provide the ability to write back configuration changes to the devices. What should you include in the recommendation?

- A. Microsoft Azure IoT Hub
- B. API apps in Microsoft Azure App Service
- C. Microsoft Azure Event Hubs
- D. Microsoft Azure Notification Hubs

Answer: A

Explanation:

References:

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

NEW QUESTION 60

- (Exam Topic 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 have Azure IoT Edge devices that generate streaming data. On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.
Solution: You deploy Azure Stream Analytics as an IoT Edge module. Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 64

- (Exam Topic 2)

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

Answer: A

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>

NEW QUESTION 68

- (Exam Topic 2)

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

Answer: C

Explanation:

References:

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

NEW QUESTION 70

- (Exam Topic 2)

You are designing an AI solution that will be used to find buildings in aerial pictures.

Users will upload the pictures to an Azure Storage account. A separate JSON document will contain for the pictures.

The solution must meet the following requirements:

- ▶ Store metadata for the pictures in a data store.
- ▶ Run a custom vision Azure Machine Learning module to identify the buildings in a picture and the position of the buildings' edges.
- ▶ Run a custom mathematical module to calculate the dimensions of the buildings in a picture based on the metadata and data from the vision module.

You need to identify which Azure infrastructure services are used for each component of the AI workflow. The solution must execute as quickly as possible.

What should you identify? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Blob Storage

Containers and blobs support custom metadata, represented as HTTP headers. Box 2: NV

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

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

Box 3: F

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

Incorrect:

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

References:

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

NEW QUESTION 75

- (Exam Topic 2)

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)

Answer: B

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.

NEW QUESTION 79

- (Exam Topic 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 have Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream. Solution: You deploy Azure Functions as an IoT Edge module.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

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>

NEW QUESTION 80

- (Exam Topic 2)

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

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

The screenshot shows an exam interface with two main sections: 'Actions' and 'Answer Area'. The 'Actions' section contains a list of six actions: 'Enable encryption that uses customer-managed keys.', 'Upload a key to an Azure key vault.', 'Generate an encryption key.', 'Generate an access key.', 'Configure a service endpoint for the storage account.', and 'Generate a shared access signature (SAS)'. The 'Answer Area' is currently empty. Between the two sections are two sets of circular arrows: a left arrow and a right arrow on the left, and an up arrow and a down arrow on the right, indicating that the actions can be moved and ordered.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

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

NEW QUESTION 81

- (Exam Topic 2)

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

Answer: C

Explanation:

References:

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

NEW QUESTION 84

- (Exam Topic 2)

Your company has factories in 10 countries. Each factory contains several thousand IoT devices. The devices present status and trending data on a dashboard. You need to ingest the data from the IoT devices into a data warehouse. Which two Microsoft Azure technologies should you use? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

Answer: CE

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-dataarchitectur>

NEW QUESTION 86

- (Exam Topic 2)

You have a container image that contains an AI solution. The solution will be used on demand and will only be needed a few hours each month. You plan to use Azure Functions to deploy the environment on-demand. You need to recommend the deployment process. The solution must minimize costs. Which four actions should you recommend Azure Functions perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

The screenshot shows an exam interface with two main sections: 'Actions' and 'Answer Area'. The 'Actions' section contains a list of six actions, each in a box with a right-pointing arrow button next to it. The 'Answer Area' is empty. The actions are:

- Shut down the virtual machine.
- Create an Azure Kubernetes Service (AKS) cluster.
- Pull the container image from the registry.
- Run the AI solution.
- Create an Azure container instance.
- Delete the Azure container instance.

At the bottom of the 'Answer Area', there are two circular buttons with arrows: one pointing left and one pointing right, indicating a sequence of actions.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Actions

Shut down the virtual machine.

Create an Azure Kubernetes Service (AKS) cluster.

Pull the container image from the registry.

Run the AI solution.

Create an Azure container instance.

Delete the Azure container instance.

Answer Area

Create an Azure container instance.

Create an Azure Kubernetes Service (AKS) cluster.

Pull the container image from the registry.

Run the AI solution.

➔

⬅

⬆

⬆

NEW QUESTION 91

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AI-100 Exam with Our Prep Materials Via below:

<https://www.certleader.com/AI-100-dumps.html>