Microsoft

Exam Questions AI-100

Designing and Implementing an Azure AI Solution

NEW QUESTION 1
- (Exam Topic 1)
You need to design the Butler chatbot solution to meet the technical requirements.
What is the best channel and pricing tier to use? More than one answer choice may achieve the goal Select the BEST answer.

A. standard channels that use the S1 pricing tier
B. standard channels that use the Free pricing tier
C. premium channels that use the Free pricing tier
D. premium channels that use the S1 pricing tier

Answer: D

Explanation:
References:
https://azure.microsoft.com/en-in/pricing/details/bot-service/

NEW QUESTION 2
- (Exam Topic 1)
You need to recommend a data storage solution that meets the technical requirements.
What is the best data storage solution to recommend? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: B

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

NEW QUESTION 3
- (Exam Topic 2)
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

Answer: C

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/

NEW QUESTION 4
- (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 5
- (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 create a managed identity for AKS, and then you create an SSH connection. Does this meet the goal?

A. Yes
B. No

Answer: B

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

NEW QUESTION 6
- (Exam Topic 2)
You need to build an A) solution that will be shared between several developers and customers. You plan to write code, host code, and document the runtime all within a single user experience. You build the environment to host the solution.
Which three actions should you perform in sequence next? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

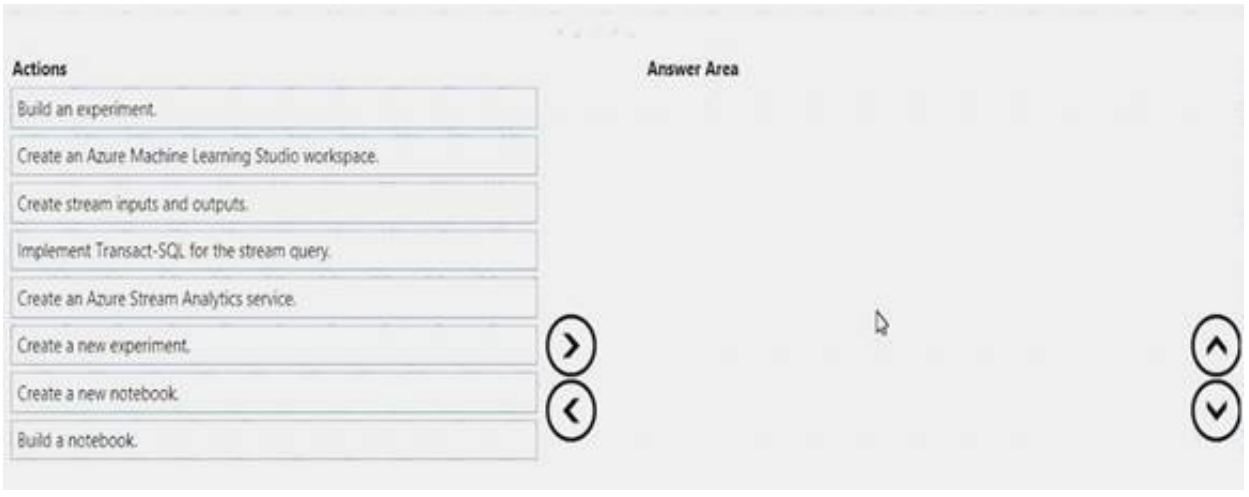

A. Mastered
B. Not Mastered

Answer: A

Explanation:
Step 1: Create an Azure Machine Learning Studio workspace
Step 2: Create a notebook
You can manage notebooks using the UI, the CLI, and by invoking the Workspace API. To create a notebook
Click the Workspace button Workspace Icon or the Home button Home Icon in the sidebar. Do one of the following:
Next to any folder, click the Menu Dropdown on the right side of the text and select Create > Notebook. Create Notebook
In the Workspace or a user folder, click Down Caret and select Create > Notebook.
2. In the Create Notebook dialog, enter a name and select the notebook's primary language.
3. If there are running clusters, the Cluster drop-down displays. Select the cluster to attach the notebook to.
4. Click Create.
Step 3: Create a new experiment
Create a new experiment by clicking +NEW at the bottom of the Machine Learning Studio window. Select EXPERIMENT > Blank Experiment.
References:
https://docs.azuredatabricks.net/user-guide/notebooks/notebook-manage.html https://docs.microsoft.com/en-us/azure/machine-learning/service/quickstart-run-cloud-notebook

NEW QUESTION 7
- (Exam Topic 2)
You design an AI workflow that combines data from multiple data sources for analysis. The data sources are composed of:
• JSON files uploaded to an Azure Storage account
• On-premises Oracle databases
• Azure SQL databases
Which service should you use to ingest the data?

A. Azure Data Factory
B. Azure SQL Data Warehouse
C. Azure Data Lake Storage
D. Azure Databricks

Answer: A

Explanation:
References:
https://docs.microsoft.com/en-us/azure/data-factory/introduction

NEW QUESTION 8
- (Exam Topic 2)
You deploy an application that performs sentiment analysis on the data stored in Azure Cosmos DB.
Recently, you loaded a large amount of data to the database. The data was for a customer named Contoso. Ltd. You discover that queries for the Contoso data are slow to complete, and the queries slow the entire
application.
You need to reduce the amount of time it takes for the queries to complete. The solution must minimize costs. What is the best way to achieve the goal? More than one answer choice may achieve the goal. Select the BEST answer.

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

Answer: B

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

NEW QUESTION 9
- (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 modify the scoring file. Does this meet the goal?

A. Yes
B. No

Answer: B

NEW QUESTION 10
- (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 10
- (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 13
- (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 15
- (Exam Topic 2)
You are designing an Azure Batch AI solution that will be used to train many different Azure Machine Learning models. The solution will perform the following:
• Image recognition
• Deep learning that uses convolutional neural networks
You need to select a compute infrastructure for each model. The solution must minimize the processing time. What should you use for each model? To answer, drag the appropriate compute infrastructures to the correct models. Each compute infrastructure may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

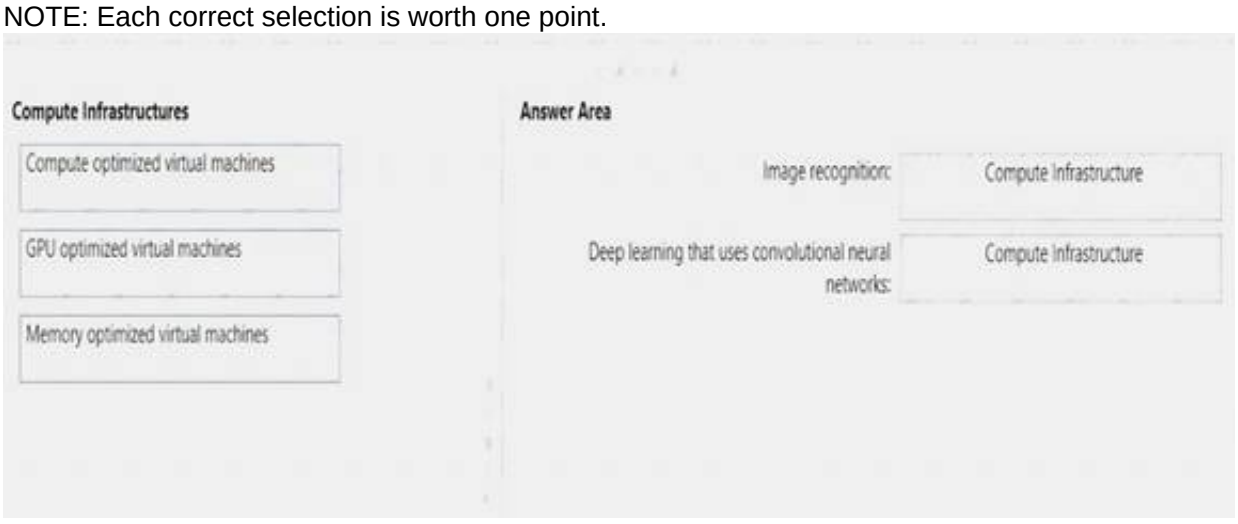

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 18
- (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.

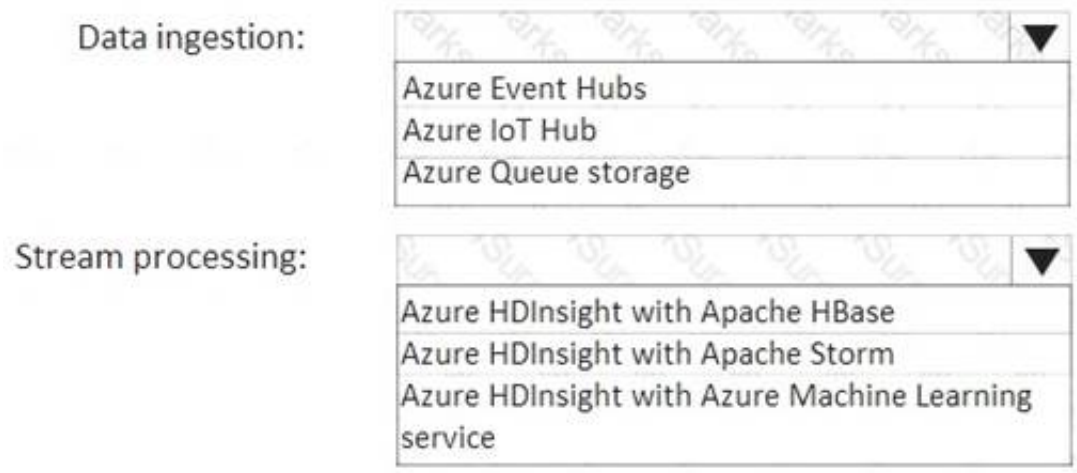

A. Mastered
B. Not Mastered

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 grand and revoke access on device level.

Box 2: Azure Hdinsight with Azure Machine Learning service References:
https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-compare-event-hubs https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-machine-learning-overview

NEW QUESTION 20
- (Exam Topic 2)
You need to build a pipeline for an Azure Machine Learning experiment.
In which order should you perform the actions? To answer, move all actions from the list of actions to the answer area and arrange them in the correct order.

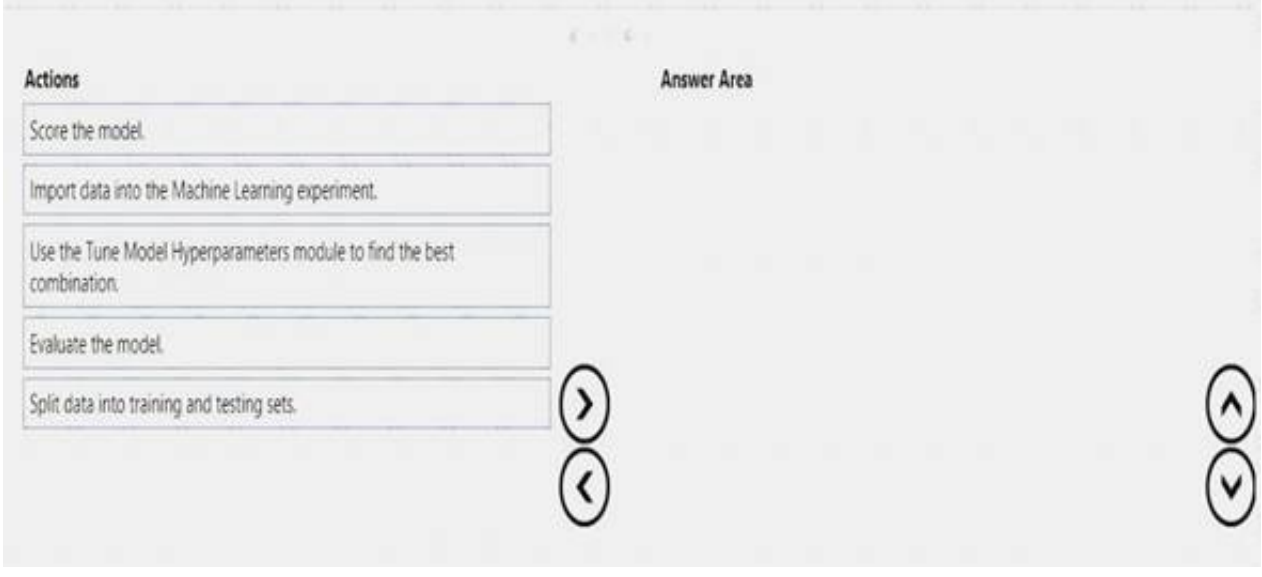

A. Mastered
B. Not Mastered

Answer: A

Explanation:
References:
https://azure.microsoft.com/en-in/blog/experimentation-using-azure-machine-learning/ https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/machine-learning-modules

NEW QUESTION 24
- (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 28
- (Exam Topic 2)
You plan to use the Microsoft 8ot Framework to develop bots that will be deployed by using the Azure Bot Service.
You need to configure the Azure Bot Service to support the following types of bots:
•Bots that use Azure Functions
•Bots that set a timer
Which template should you use for each bot type? To answer, drag the appropriate templates to the correct bot types. Each template may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
NOTE: Each correct selection is worth one point.

A. Mastered
B. Not Mastered

Answer: A

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

NEW QUESTION 29
- (Exam Topic 2)
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.

Answer: D

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

NEW QUESTION 34
- (Exam Topic 2)
You plan to design a solution for an Al implementation that uses data from loT devices.
You need to recommend a data storage solution for the loT devices that meets the following requirements:
•Allow data to be queried in real-time as it streams into the solution.
•Provide the lowest amount of latency for loading data into the solution. What should you include in the recommendation?

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

Answer: D

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

NEW QUESTION 39
- (Exam Topic 2)
You have an Al application that uses keys in Azure Key Vault.
Recently, a key used by the application was deleted accidentally and was unrecoverable. You need to ensure that if a key is deleted, it is retained in the key vault for 90 days. Which two features should you configure? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point

A. the expiration date on the keys
B. soft delete
C. purge protection

D. auditors
E. the activation date on the keys

Answer: BC

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

NEW QUESTION 41
- (Exam Topic 2)
You are configuring data persistence for a Microsoft Bot Framework application. The application requires a structured NoSQL cloud data store.
You need to identify a storage solution for the application. The solution must minimize costs. What should you identify?

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

Answer: D

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

NEW QUESTION 44
- (Exam Topic 2)
You have Azure loT Edge devices that collect measurements every 30 seconds. You plan to send the measurements to an Azure loT hub. You need to ensure that every event is processed as quickly as possible. What should you use?

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

Answer: A

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

NEW QUESTION 49
- (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 53
- (Exam Topic 2)
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

Answer: B

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

NEW QUESTION 57
- (Exam Topic 2)
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

Answer: AB

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 publishsubscribe 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

NEW QUESTION 60
- (Exam Topic 2)
You deploy an Azure bot.
You need to collect Key Performance Indicator (KPI) data from the bot. The type of data includes:
• The number of users interacting with the bot
• The number of messages interacting with the bot
• The number of messages on different channels received by the bot
• The number of users and messages continuously interacting with the bot What should you configure?

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

Answer: A

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

NEW QUESTION 62
- (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 67

- (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.

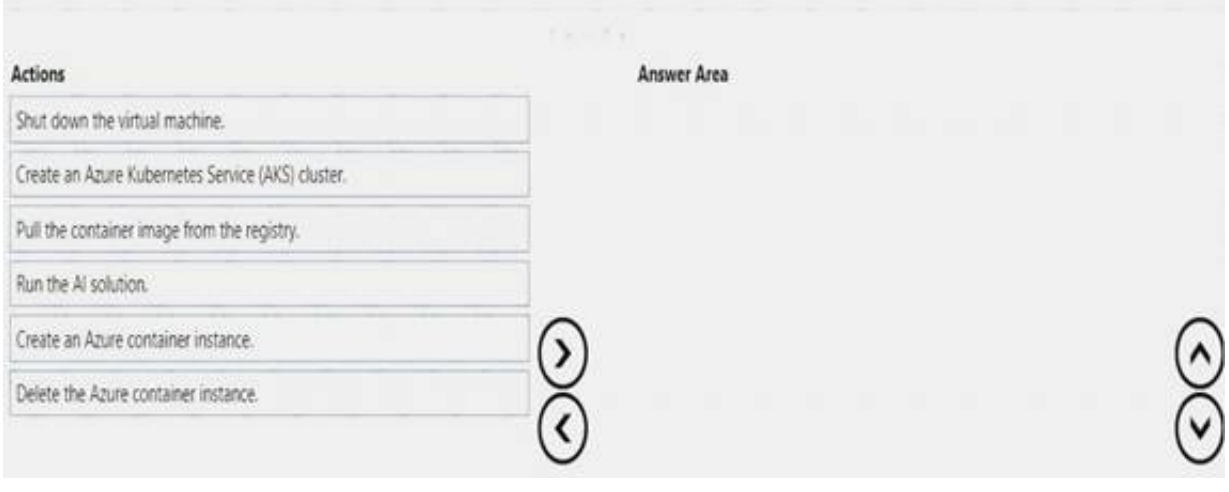

A. Mastered
B. Not Mastered

Answer: A

Explanation:

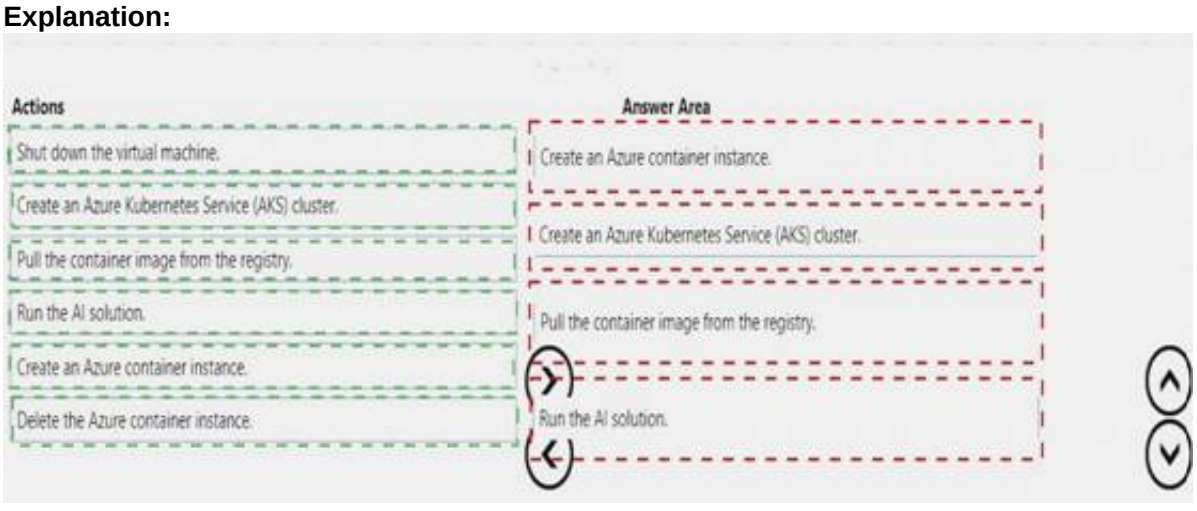

NEW QUESTION 68
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

AI-100 Practice Exam Features:

* AI-100 Questions and Answers Updated Frequently

* AI-100 Practice Questions Verified by Expert Senior Certified Staff

* AI-100 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AI-100 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year