

AI-100 Dumps

Designing and Implementing an Azure AI Solution

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



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 2)

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

Answer: A

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.

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)

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 5

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

Answer: C

Explanation:

References:

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

NEW QUESTION 6

- (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 7

- (Exam Topic 2)

You have an Azure Machine Learning model that is deployed to a web service. You plan to publish the web service by using the name ml.contoso.com.

You need to recommend a solution to ensure that access to the web service is encrypted. Which three actions should you recommend? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

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

Answer: BDE

Explanation:

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

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

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

the key file. References:

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

NEW QUESTION 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)

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 10

- (Exam Topic 2)

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

The source data for the cube is located in an on premises network in a Microsoft SQL Server database. You need to ensure that the Azure Analysis Services service can access the source data.

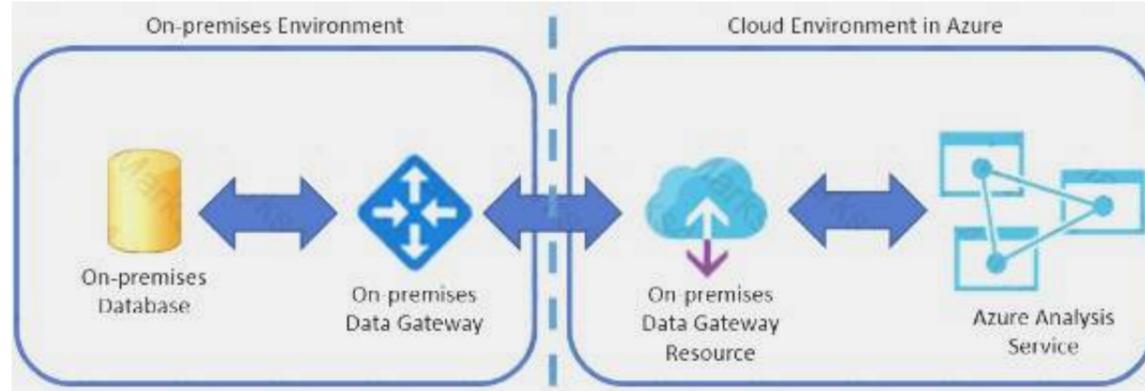
What should you deploy to your Azure subscription?

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

Answer: C

Explanation:

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



References:

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

NEW QUESTION 10

- (Exam Topic 2)

You plan to implement a new data warehouse for a planned AI solution. You have the following information regarding the data warehouse:

- The data files will be available in one week.
- Most queries that will be executed against the data warehouse will be ad-hoc queries.
- The schemas of data files that will be loaded to the data warehouse will change often.
- One month after the planned implementation, the data warehouse will contain 15 TB of data. You need to recommend a database solution to support the planned implementation.

What two solutions should you include in the recommendation? Each correct answer is a complete solution. NOTE: Each correct selection is worth one point.

- A. Apache Hadoop
- B. Apache Spark
- C. a Microsoft Azure SQL database
- D. an Azure virtual machine that runs Microsoft SQL Server

Answer: AB

NEW QUESTION 12

- (Exam Topic 2)

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

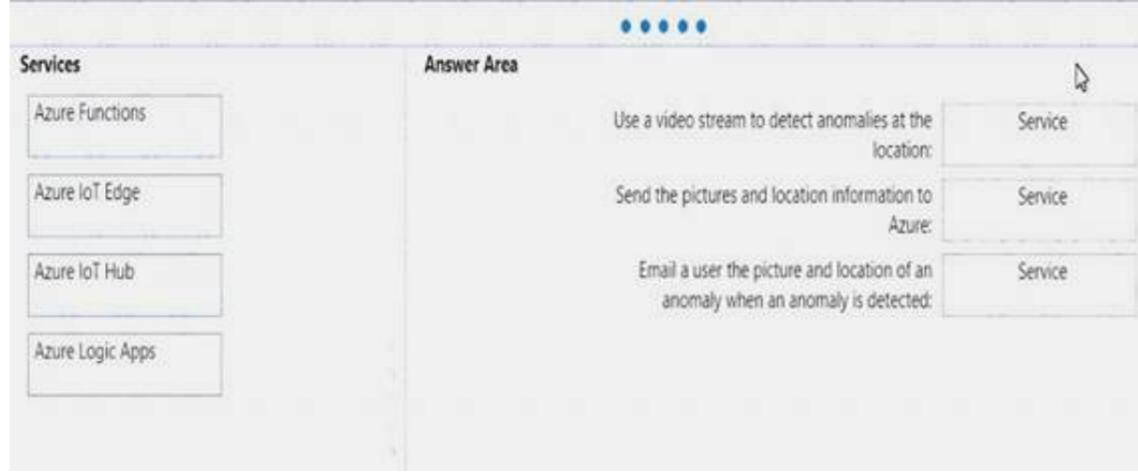
- Email a user the picture and location of an anomaly when an anomaly is detected.
- Use a video stream to detect anomalies at the location.
- Send the pictures and location information to Azure.

•Use the least amount of code possible.

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

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

NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

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>

NEW QUESTION 13

- (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.



- A. Mastered
- B. Not Mastered

Answer: A

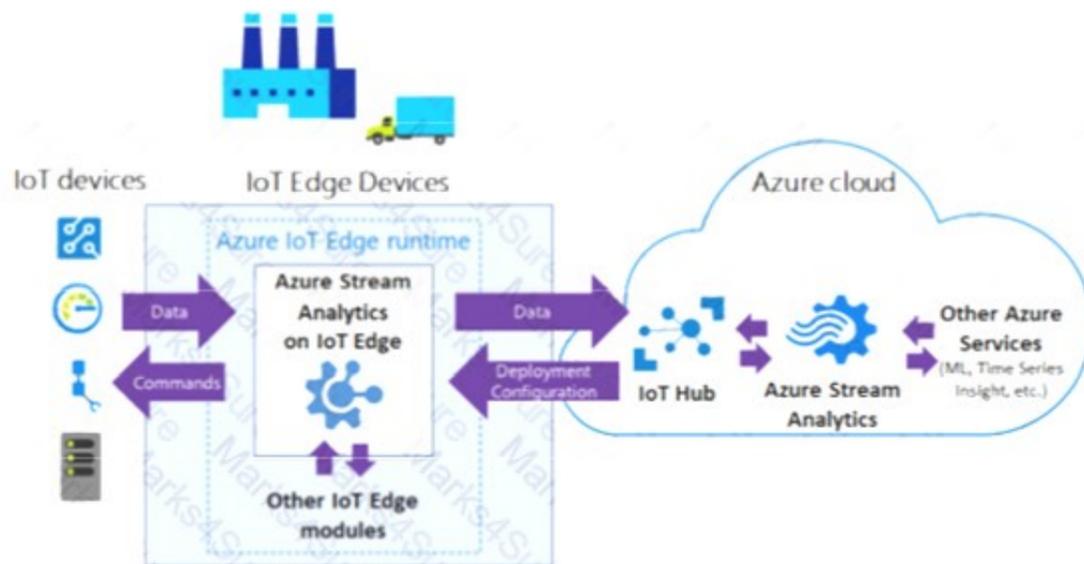
Explanation:

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 17

- (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 18

- (Exam Topic 2)

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.

To connect to the on-premises data:	<input type="checkbox"/> A point-to-site VPN connection <input type="checkbox"/> A site-to-site VPN connection <input type="checkbox"/> Azure App Service Hybrid Connections
To orchestrate the phases:	<input type="checkbox"/> A Machine Learning experiment <input type="checkbox"/> Azure Machine Learning Studio <input type="checkbox"/> Machine Learning pipelines
To control the orchestrations:	<input type="checkbox"/> Azure Automation <input type="checkbox"/> Azure Databricks <input type="checkbox"/> Azure Notebooks

- A. Mastered
- B. Not Mastered

Answer: A

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 Option: The VPN connection solution both use gateways. Box 2: Machine Learning pipelines

Typically when running machine learning algorithms, it involves a sequence of tasks including pre-processing, feature extraction, model fitting, and validation stages. For example, when classifying text documents might involve text segmentation and cleaning, extracting features, and training a classification model with cross-validation. Though there are many libraries we can use for each stage, connecting the dots is not as easy as it may look, especially with large-scale datasets. Most ML libraries are not designed for distributed computation or they do not provide native support for pipeline creation and tuning.

Box 3: Azure Databricks References:

<https://azure.microsoft.com/is-is/blog/hybrid-connections-preview/> <https://databricks.com/glossary/what-are-ml-pipelines>

NEW QUESTION 21

- (Exam Topic 2)

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

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

You need to choose an output type for the job.

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

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

Answer: D

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>

NEW QUESTION 25

- (Exam Topic 2)

You are developing an application that will perform optical character recognition of photos of medical logbooks. You need to recommend a solution to validate the data against a validated set of records. Which service should you include in the recommendation?

- A. Azure Data Catalog
- B. Text Analytics
- C. Bing Autosuggest
- D. Master Data Services (MDS) in Microsoft SQL Server

Answer: D

Explanation:

References:

<https://docs.microsoft.com/en-us/sql/master-data-services/validation-master-data-services?view=sql-server-2017>

NEW QUESTION 27

- (Exam Topic 2)

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

Answer: C

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>

NEW QUESTION 32

- (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 change the permissions of the AKS resource group, 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 37

- (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 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 create environment files.
Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

NEW QUESTION 40

- (Exam Topic 2)

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 rx 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

Answer: ABC

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>

NEW QUESTION 43

- (Exam Topic 2)

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

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

Answer: E

Explanation:

References:

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

NEW QUESTION 47

- (Exam Topic 2)

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

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

Answer: A

Explanation:

References:

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

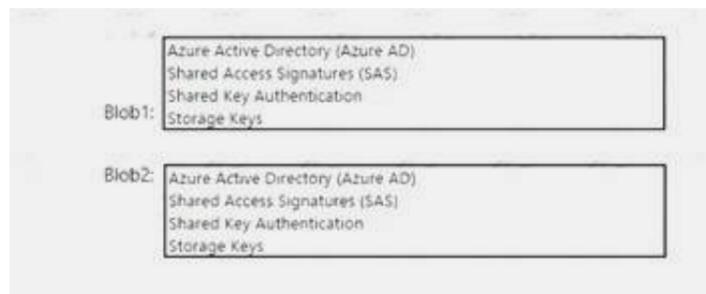
NEW QUESTION 51

- (Exam Topic 2)

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

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

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

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

NEW QUESTION 52

- (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 55

- (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 57

.....

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>