



Microsoft

Exam Questions AI-100

Designing and Implementing an Azure AI Solution

About Exambible

[Your Partner of IT Exam](#)

Found in 1998

Exambible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, Exambible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

- (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 2

- (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 3

- (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 4

- (Exam Topic 2)

You need to build a solution to monitor Twitter. The solution must meet the following requirements:

- Send an email message to the marketing department when negative Twitter messages are detected.
- Run sentiment analysis on Twitter messages that mention specific tags.
- Use the least amount of custom code possible.

Which two services should you include in the solution? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Azure Databricks
- B. Azure Stream Analytics
- C. Azure Functions
- D. Azure Cognitive Services
- E. Azure Logic Apps

Answer: BE

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/stream-analytics/streaming-technologies> <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-twitter-sentiment-analysis-trends>

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:

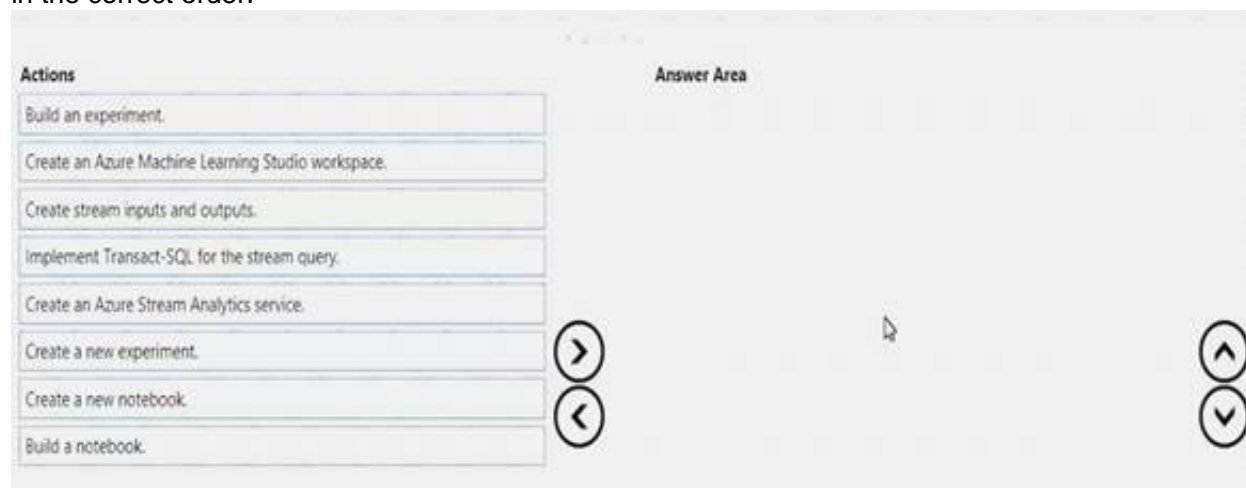
<https://docs.microsoft.com/en-us/azure/aks/ssh>

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

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>

NEW QUESTION 7

- (Exam Topic 2)

Your company has a data team of Scala and R experts.

You plan to ingest data from multiple Apache Kafka streams.

You need to recommend a processing technology to broker messages at scale from the Kafka streams to Azure Storage.

What should you recommend?

A. Azure Databricks

B. Azure Functions

C. Azure HDInsight with Apache Storm

D. Azure HDInsight with Microsoft Machine Learning Server

Answer: C

Explanation:

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-streaming-at-scale-overview?toc=https%3A%2F%2F>

NEW QUESTION 8

- (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 9

- (Exam Topic 2)

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

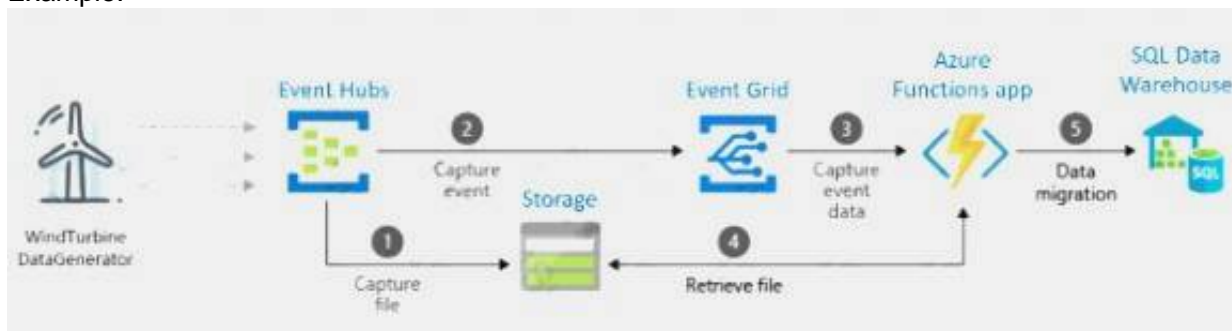
Answer: B

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>

NEW QUESTION 10

- (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 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

Answer: A

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>

NEW QUESTION 10

- (Exam Topic 2)

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.

Actions

Keys: create

Keys: get

Keys: list

Secrets: all

Secrets: get

Secrets: list

Answer Area

IT security team:

Permission

Applications:

Permission

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Actions

Keys: create

Keys: get

Keys: list

Secrets: all

Secrets: get

Secrets: list

Answer Area

IT security team:

Secrets: all

Applications:

Secrets: get

NEW QUESTION 13

- (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.

Compute Infrastructures

Compute optimized virtual machines

GPU optimized virtual machines

Memory optimized virtual machines

Answer Area

Image recognition:

Compute Infrastructure

Deep learning that uses convolutional neural networks:

Compute Infrastructure

- 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 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 20

- (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:

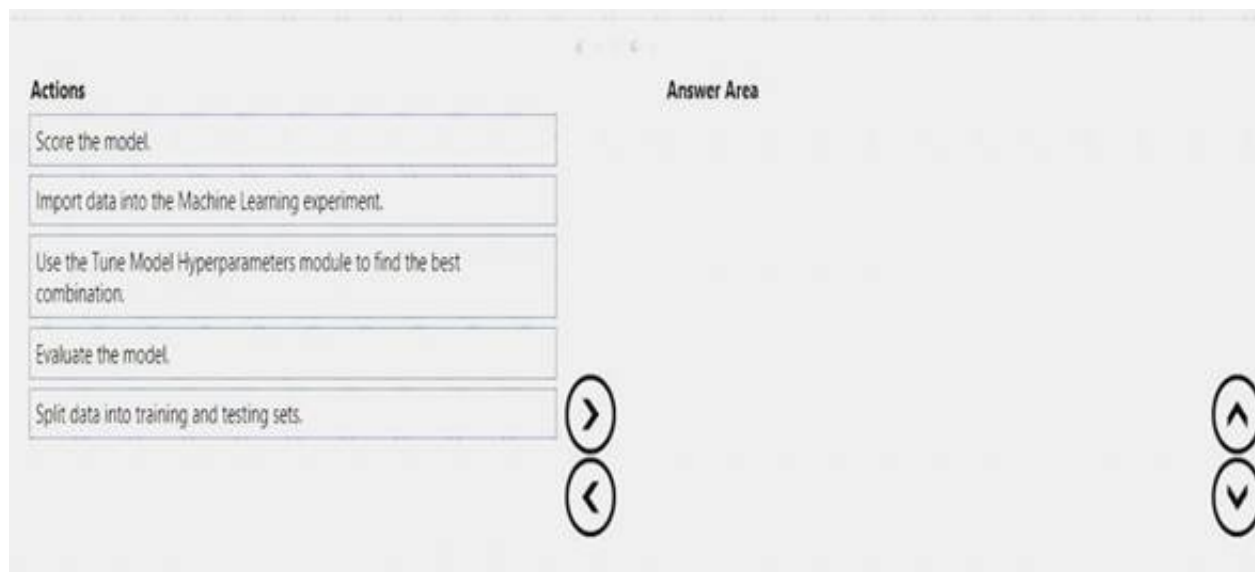
<https://docs.microsoft.com/en-us/azure/aks/cluster-autoscaler>

NEW QUESTION 24

- (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:

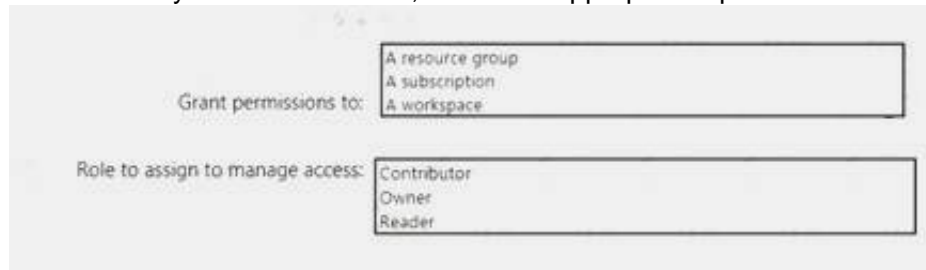
<https://azure.microsoft.com/en-in/blog/experimentation-using-azure-machine-learning/> <https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/machine-learning-modules>

NEW QUESTION 26

- (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 28

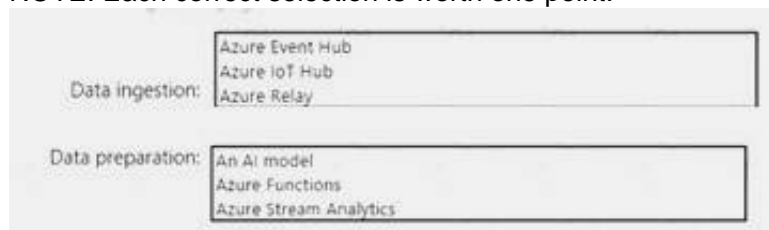
- (Exam Topic 2)

You are designing a solution that will analyze bank transactions in real time. The transactions will be evaluated by using an algorithm and classified into one of five groups. The transaction data will be enriched with information taken from Azure SQL Database before the transactions are sent to the classification process. The enrichment process will require custom code. Data from different banks will require different stored procedures.

You need to develop a pipeline for the solution.

Which components should you use for data ingestion and data preparation? 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/bs-latn-ba/azure/architecture/example-scenario/data/fraud-detection>

NEW QUESTION 31

- (Exam Topic 2)

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 File Storage
- B. Azure Manage Disks
- C. Azure Data Lake Storage
- D. Azure Blob storage

Answer: D

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>

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

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 33

- (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 36

- (Exam Topic 2)

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

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

Answer: D

Explanation:

References:

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

NEW QUESTION 37

- (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 41

- (Exam Topic 2)

You plan to deploy an AI solution that tracks the behavior of 10 custom mobile apps. Each mobile app has several thousand users. You need to recommend a solution for real-time data ingestion for the data originating from the mobile app users. Which Microsoft Azure service should you include in the recommendation?

- A. Azure Event Hubs
- B. Azure Service Bus queues
- C. Azure Service Bus topics and subscriptions
- D. Apache Storm on Azure HDInsight

Answer: A

Explanation:

References:

<https://docs.microsoft.com/en-in/azure/event-hubs/event-hubs-about>

NEW QUESTION 46

- (Exam Topic 2)

You are designing an AI application that will use an Azure Machine Learning Studio experiment. The source data contains more than 200 TB of relational tables. The experiment will run once a month. You need to identify a data storage solution for the application. The solution must minimize compute costs. Which data storage solution should you identify?

- A. Azure Database for MySQL
- B. Azure SQL Database
- C. Azure SQL Data Warehouse

Answer: B

Explanation:

References:

<https://azure.microsoft.com/en-us/pricing/details/sql-database/single/>

NEW QUESTION 47

- (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 52

- (Exam Topic 2)

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.

Service to use:	<div>Apache Hive</div> <div>Azure Data Factory</div> <div>Azure Functions</div> <div>Azure Stream Analytics</div>
Location to deploy the job:	<div>A Web job in Azure</div> <div>An Azure IoT Edge device</div> <div>Azure Event Hubs</div> <div>Azure Notification Hubs</div>

- A. Mastered
- B. Not Mastered

Answer: A

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>

NEW QUESTION 57

- (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 60

- (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:

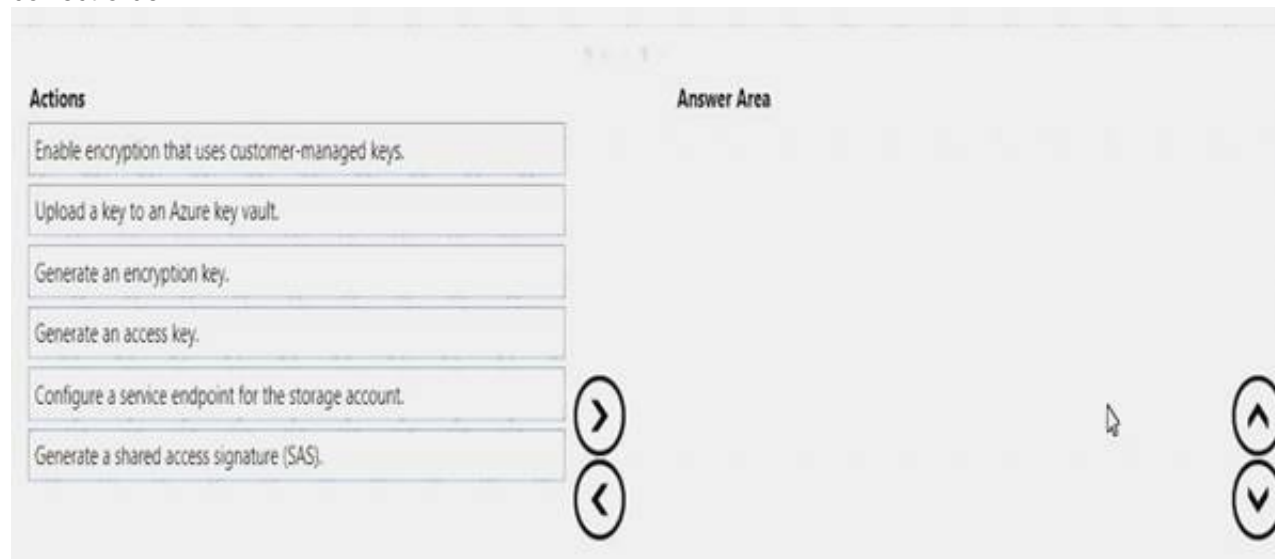
FPGAs, such as those available on Azure, provide performance close to ASICs. They are also flexible and reconfigurable over time, to implement new logic.

NEW QUESTION 62

- (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, each in a box with a right-pointing arrow button to its right. The 'Answer Area' is empty and has two circular buttons with up and down arrows on its right side. The actions are:

- 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.
- Generate a shared access signature (SAS).

- 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>

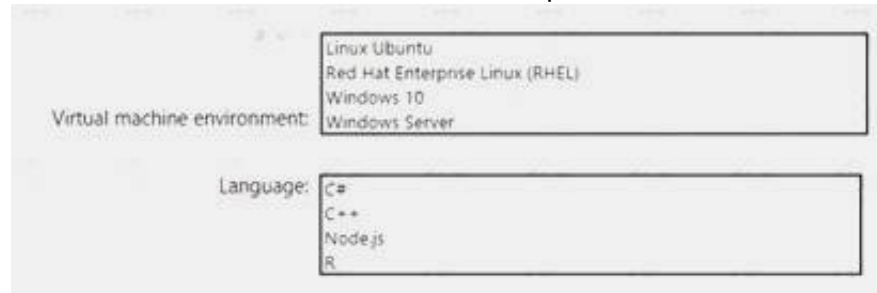
NEW QUESTION 66

- (Exam Topic 2)

You need to build and deploy a real-time media bot for Microsoft Skype on an Azure virtual machine. The bot will use the Azure Bot Service. The solution must minimize custom code.

Which environment and language should you use to develop the bot? 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/microsoftteams/platform/concepts/calls-and-meetings/requirements-considerati>

NEW QUESTION 67

- (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 68

- (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 71

- (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.

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

>

<

- 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.

Create an Azure container instance.

Run the AI solution.

NEW QUESTION 74

.....

Relate Links

100% Pass Your AI-100 Exam with Examible Prep Materials

<https://www.examible.com/AI-100-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.examible.com/>