

## 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 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 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)

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream. Solution: You expose a Machine Learning model as an Azure web service.

Does this meet the goal?

- A. Yes
- B. No

**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 5**

- (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 6**

- (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 7**

- (Exam Topic 2)

Your company plans to deploy an AI solution that processes IoT data in real-time.

You need to recommend a solution for the planned deployment that meets the following requirements: Sustain up to 50 Mbps of events without throttling.

Retain data for 60 days.

What should you recommend?

- A. Apache Kafka
- B. Microsoft Azure IoT Hub
- C. Microsoft Azure Data Factory
- D. Microsoft Azure Machine Learning

**Answer:** A

**Explanation:**

Apache Kafka is an open-source distributed streaming platform that can be used to build real-time streaming data pipelines and applications.

References:

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

**NEW QUESTION 8**

- (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:

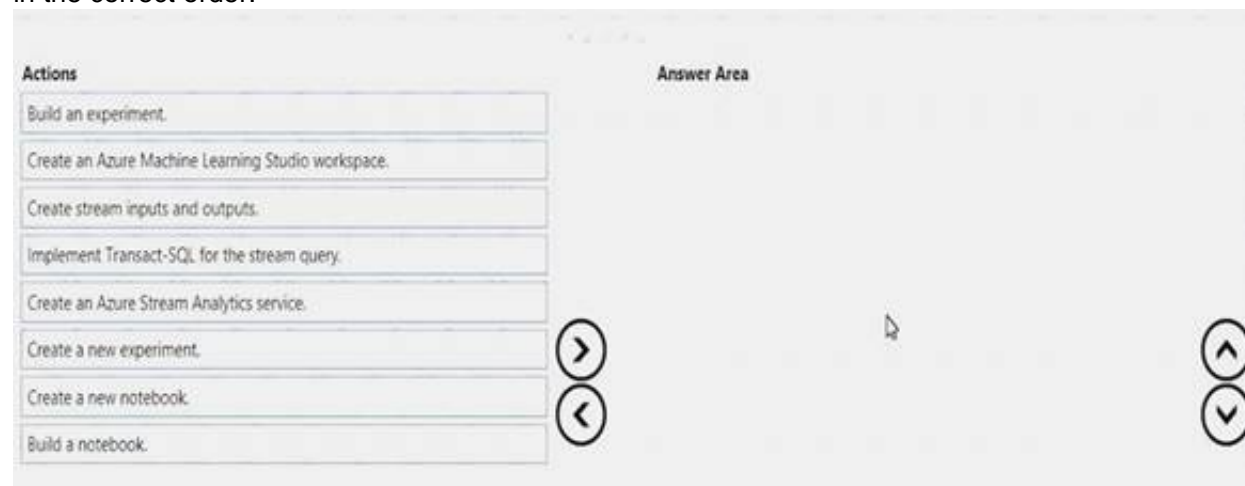
<https://docs.microsoft.com/en-us/azure/aks/ssh>

**NEW QUESTION 9**

- (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 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 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 10**

- (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 14**

- (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 19**

- (Exam Topic 2)

You are developing a mobile application that will perform optical character recognition (OCR) from photos. The application will annotate the photos by using metadata, store the photos in Azure Blob storage, and then score the photos by using an Azure Machine Learning model.

What should you use to process the data?

- A. Azure Event Hubs
- B. Azure Functions
- C. Azure Stream Analytics
- D. Azure Logic Apps

**Answer:** A

**NEW QUESTION 23**

- (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 27**

- (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 create several AI models in Azure Machine Learning Studio. You deploy the models to a production environment.

You need to monitor the compute performance of the models. Solution: You write a custom scoring script.

Does this meet the goal?

- A. Yes
- B. No

**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 32**

- (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 34**

- (Exam Topic 2)

You have a solution that runs on a five-node Azure Kubernetes Service (AKS) cluster. The cluster uses an Nseries virtual machine.

An Azure Batch AI process runs once a day and rarely on demand.

You need to recommend a solution to maintain the cluster configuration when the cluster is not in use. The solution must not incur any compute costs.

What should you include in the recommendation?

- A. Downscale the cluster to one node
- B. Downscale the cluster to zero nodes
- C. Delete the cluster

**Answer:** A

**Explanation:**

An AKS cluster has one or more nodes. References:

<https://docs.microsoft.com/en-us/azure/aks/concepts-clusters-workloads>

**NEW QUESTION 35**

- (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.



- 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 37**

- (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.

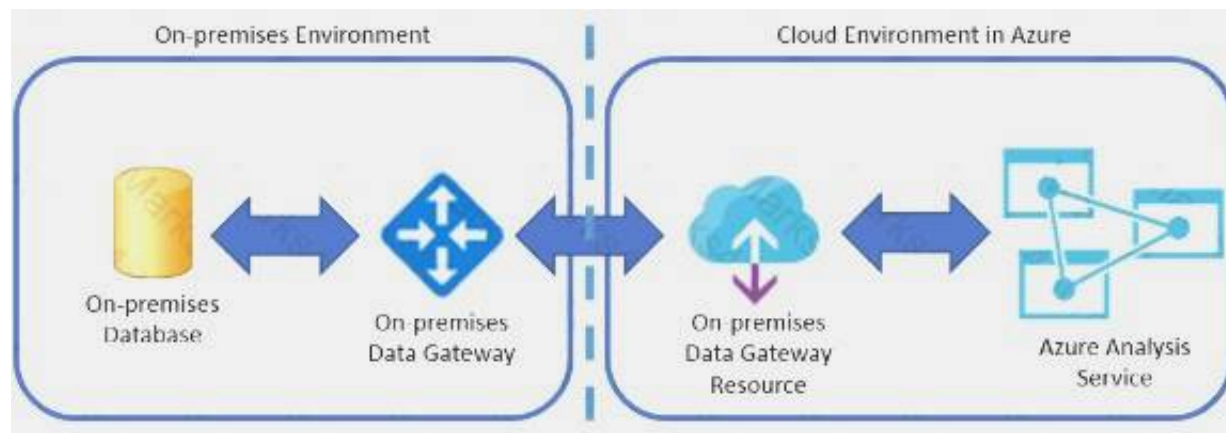
What should you deploy to your Azure subscription?

- 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 39

- (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 42

- (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 47

- (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.

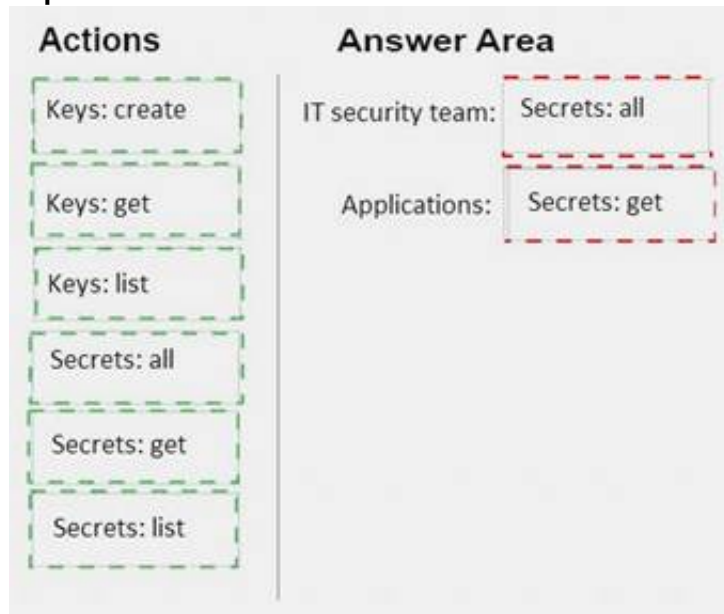
NOTE: Each correct selection is worth one point.

Actions	Answer Area
Keys: create	IT security team: <input type="text" value="Permission"/>
Keys: get	Applications: <input type="text" value="Permission"/>
Keys: list	
Secrets: all	
Secrets: get	
Secrets: list	

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**



**NEW QUESTION 48**

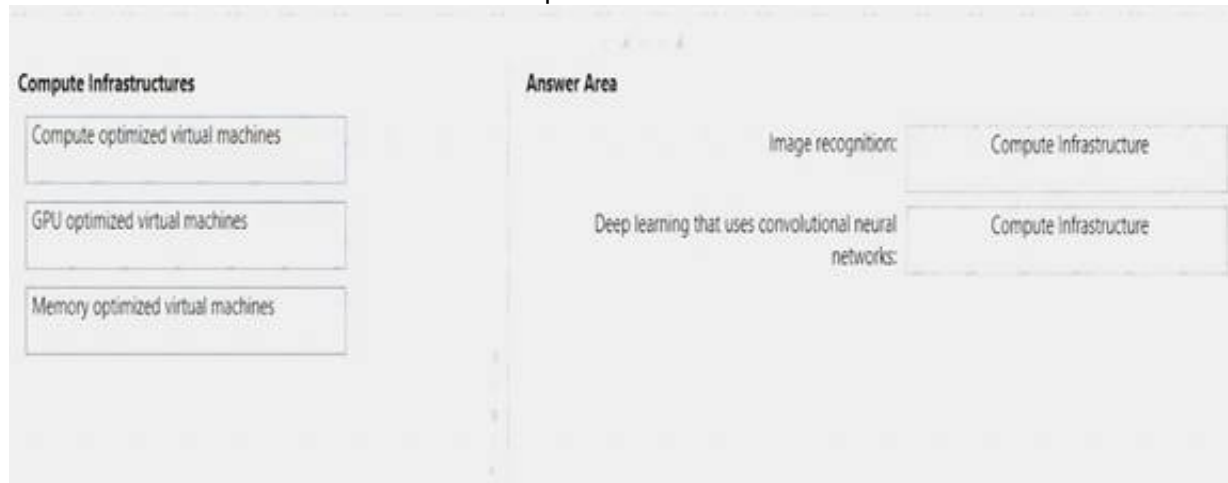
- (Exam Topic 2)

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>

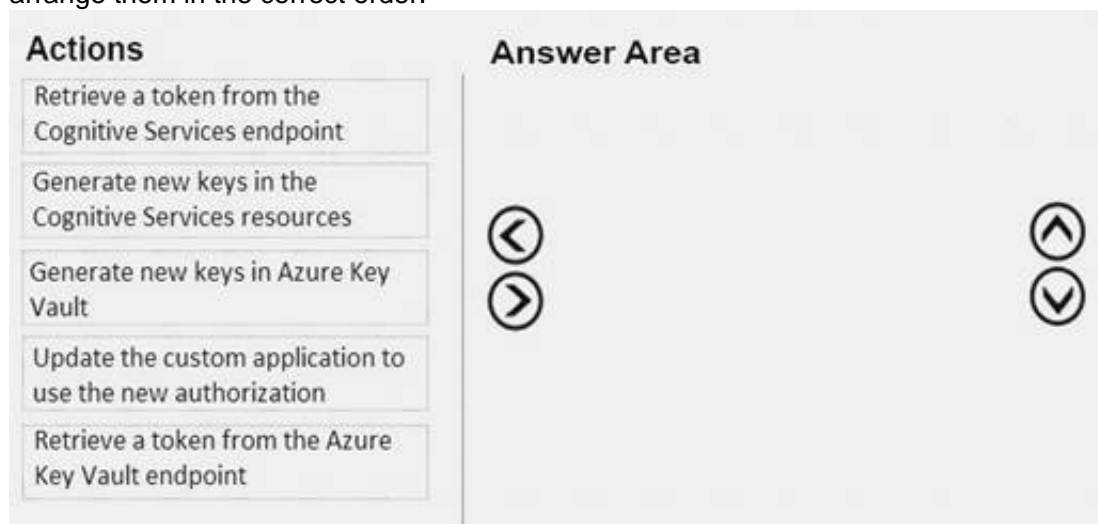
**NEW QUESTION 53**

- (Exam Topic 2)

You develop a custom application that uses a token to connect to Azure Cognitive Services resources. A new security policy requires that all access keys are changed every 30 days.

You need to recommend a solution to implement the security policy.

Which three actions should you recommend be performed every 30 days? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.



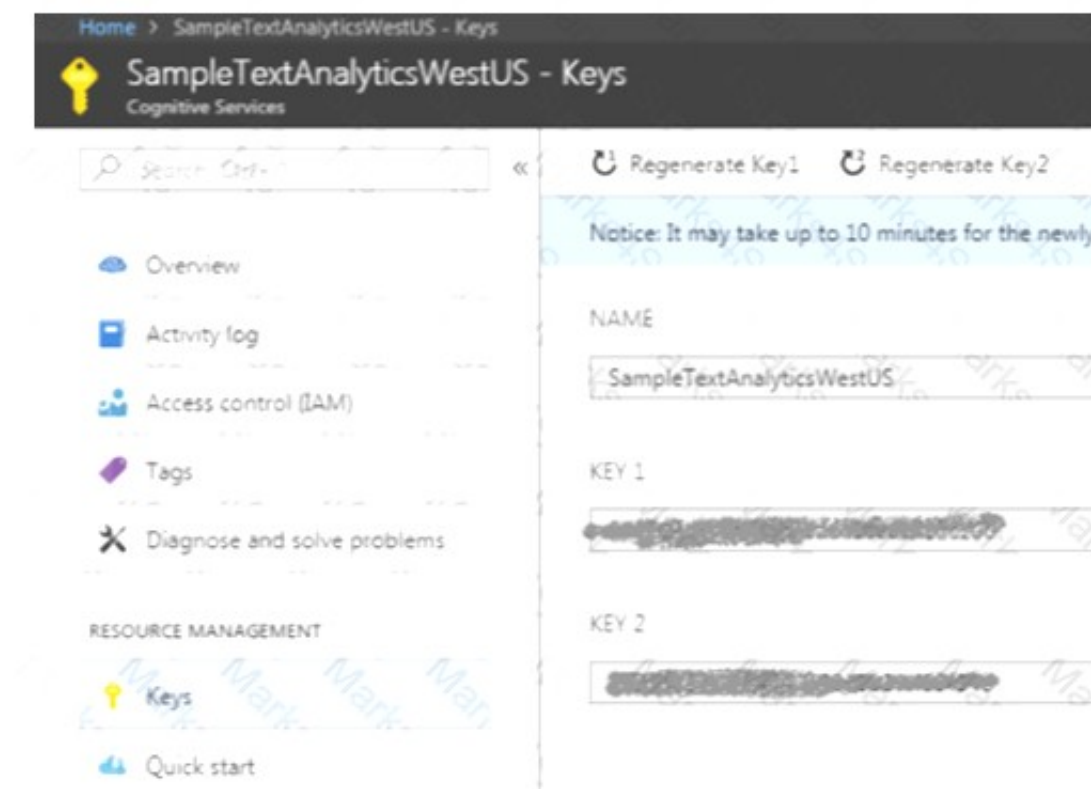
- A. Mastered
- B. Not Mastered



**Answer:** A

**Explanation:**

Step 1: Generate new keys in the Cognitive Service resources



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

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

References:

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

**NEW QUESTION 57**

- (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 60**

- (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 65**

- (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

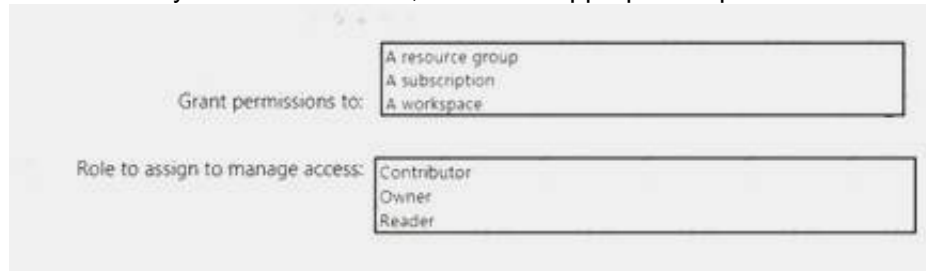
**Answer:** A

**NEW QUESTION 68**

- (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 69**

- (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 71**

- (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 76**

- (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 79**

- (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 84**

- (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 89**

- (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 93**

- (Exam Topic 2)

A data scientist deploys a deep learning model on an Fsv2 virtual machine. Data analysis is slow.

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

Which series should you recommend?

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

**Answer:** A

**Explanation:**

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

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

References:

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

**NEW QUESTION 96**

- (Exam Topic 2)

You have an AI 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 99**

- (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 102**

- (Exam Topic 2)

Your company develops an AI application that is orchestrated by using Kubernetes. You need to deploy the application.

Which three actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Create a Kubernetes cluster.
- B. Create an Azure Container Registry instance.
- C. Create a container image file.
- D. Create a Web App for Containers.
- E. Create an Azure container instance.

**Answer:** ABC

**Explanation:**

References:

<https://docs.microsoft.com/en-us/azure/aks/tutorial-kubernetes-prepare-app>

**NEW QUESTION 105**

- (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 107**

- (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 111**

- (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 114**



- (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 117**

- (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.

The screenshot shows two storage containers, Blob1 and Blob2, in the Azure portal. For each container, there are four checkboxes for authentication methods: 'Azure Active Directory (Azure AD)', 'Shared Access Signatures (SAS)', 'Shared Key Authentication', and 'Storage Keys'. The 'Azure Active Directory (Azure AD)' checkbox is checked for both Blob1 and Blob2.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

References:

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

**NEW QUESTION 120**

- (Exam Topic 2)

Your company has an Azure subscription that contains an Azure Active Directory (Azure AD) tenant. Azure AD contains 500 user accounts for your company's employees. Some temporary employees do NOT have user accounts in Azure AD

You are designing a storage solution for video files and metadata files. You plan to deploy an application to perform analysis of the metadata files.

You need to recommend an authentication solution to provide links to the video files. The solution must provide access to each file for only five minutes.

What should you include in the in the recommendation?

- A. Secondary Storage Key
- B. Primary Storage Key
- C. Shared Access Signature
- D. Azure Active Directory

**Answer:** C

**Explanation:**

References:

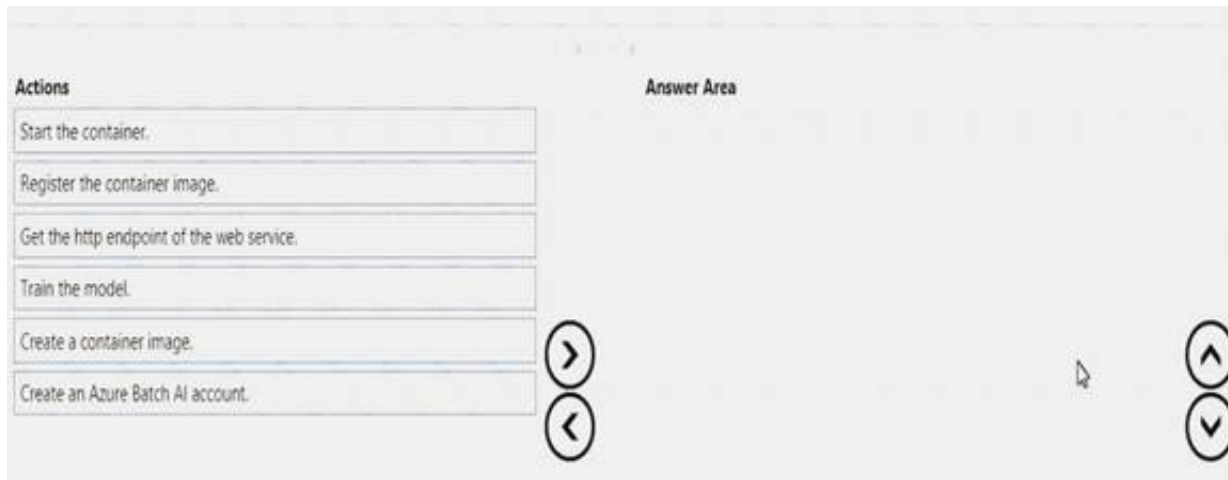
<https://docs.microsoft.com/en-us/azure/storage/common/storage-dotnet-shared-access-signature-part-1>

**NEW QUESTION 121**

- (Exam Topic 2)

You create an image classification model in Azure Machine Learning Studio. You need to deploy the model as a containerized web service.

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



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

The following diagram illustrates the complete deployment workflow: The deployment workflow includes the following steps:

- ▶ Register the model in a registry hosted in your Azure Machine Learning Service workspace
- ▶ Deploy the image as a web service in the cloud or to edge devices
- ▶ Monitor and collect data
- ▶ Update a deployment to use a new image.

References:

<https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/concept-model-management-and-deploym>

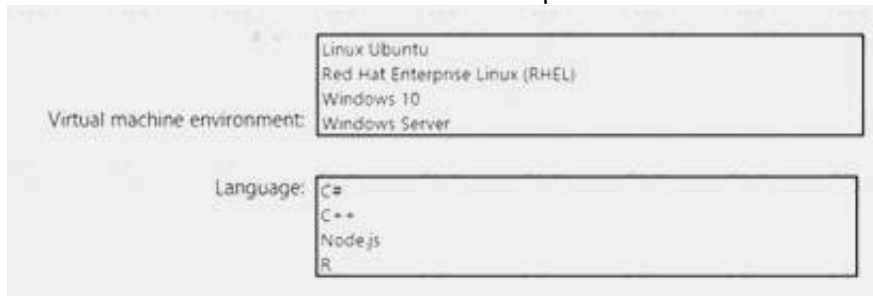
**NEW QUESTION 122**

- (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 126**

- (Exam Topic 2)

You need to build an interactive website that will accept uploaded images, and then ask a series of predefined questions based on each image.

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



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Azure Bot Service

Box 2: Computer Vision

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

References:

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

**NEW QUESTION 130**

.....

## 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>