

## AI-100 Dumps

### Designing and Implementing an Azure AI Solution

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



### NEW QUESTION 1

- (Exam Topic 1)

Which RBAC role should you assign to the KeyManagers group?

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

**Answer:** A

#### Explanation:

References:

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

### NEW QUESTION 2

- (Exam Topic 1)

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

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

**Actions**

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

**Answer Area**

- A. Mastered
- B. Not Mastered

**Answer:** A

#### Explanation:

References:

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

### NEW QUESTION 3

- (Exam Topic 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 4

- (Exam Topic 2)

You are designing a solution that will ingest data from an Azure IoT Edge device, preprocess the data in Azure Machine Learning, and then move the data to Azure HDInsight for further processing.

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.

Machine Learning module to use to move the data into HDInsight:

Export Data  
Load Trained Model  
Partition and Sample  
Unpack Zipped Datasets

Query type to use:

Apache Hive  
Apache Spark  
C#  
Transact-SQL

Output the data to:

Azure Cosmos DB  
Azure Data Lake  
Azure Table storage  
HDFS

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

Box 1: Export Data

The Export data to Hive option in the Export Data module in Azure Machine Learning Studio. This option is useful when you are working with very large datasets, and want to save your machine learning experiment data to a Hadoop cluster or HDInsight distributed storage.

Box 2: Apache Hive

Apache Hive is a data warehouse system for Apache Hadoop. Hive enables data summarization, querying, and analysis of data. Hive queries are written in HiveQL, which is a query language similar to SQL.

Box 3: Azure Data Lake

Default storage for the HDFS file system of HDInsight clusters can be associated with either an Azure Storage account or an Azure Data Lake Storage.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio-module-reference/export-to-hive-query> <https://docs.microsoft.com/en-us/azure/hdinsight/hadoop/hdinsight-use-hive>

**NEW QUESTION 5**

- (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 6**

- (Exam Topic 2)

You have an intelligent edge solution that processes data and outputs the data to an Azure Cosmos DB account that uses the SQL API.

You need to ensure that you can perform full text searches of the data.

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

Actions

Create a cognitive search pipeline.

Create an Azure Search index.

Create an Azure Service Bus.

Create a data source.

Create an Azure Search indexer.


Answer Area

> < ^ v

- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**



### NEW QUESTION 7

- (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:

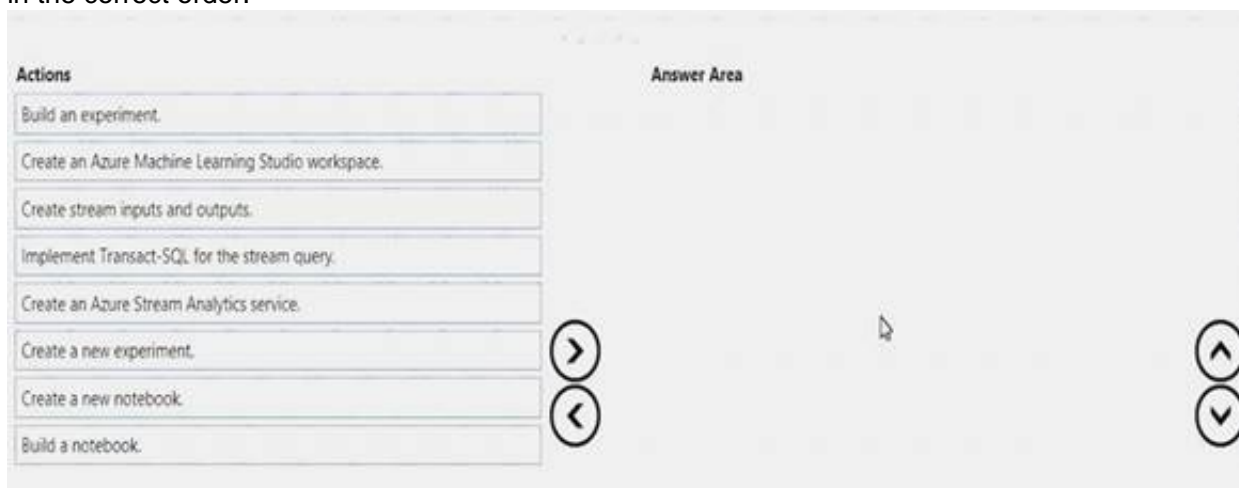
<https://docs.microsoft.com/en-us/azure/media-services/video-indexer/video-indexer-overview>

### NEW QUESTION 8

- (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 9

- (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 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 10

- (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 12

- (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 17**

- (Exam Topic 2)

You plan to build an application that will perform predictive analytics. Users will be able to consume the application data by using Microsoft Power BI or a custom website.

You need to ensure that you can audit application usage. Which auditing solution should you use?

- A. Azure Storage Analytics
- B. Azure Application Insights
- C. Azure diagnostic logs
- D. Azure Active Directory (Azure AD) reporting

**Answer:** D

**Explanation:**

References:

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

**NEW QUESTION 20**

- (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 24**

- (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 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 are deploying an Azure Machine Learning model to an Azure Kubernetes Service (AKS) container. You need to monitor the accuracy of each run of the model.

Solution: You modify the scoring file. Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**NEW QUESTION 31**

- (Exam Topic 2)

You create an Azure Machine Learning Studio experiment.

You plan to publish the experiment as a Machine Learning Web service.

You need to ensure that you can consume the web service from Microsoft Excel spreadsheets. What should you use?

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

**Answer:** C

**Explanation:**

Steps to Add a New web service

1. Deploy a web service or use an existing Web service.
2. Click Consume.
3. Look for the Basic consumption info section. Copy and save the Primary Key and the Request-Response URL.
4. In Excel, go to the Web Services section (if you are in the Predict section, click the back arrow to go to the list of web services).
5. Click Add Web Service.
6. Paste the URL into the Excel add-in text box labeled URL.
7. Paste the API/Primary key into the text box labeled API key.
8. Click Add.

References:

<https://docs.microsoft.com/en-us/azure/machine-learning/studio/excel-add-in-for-web-services>

**NEW QUESTION 34**

- (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.

The screenshot shows the Azure Machine Learning Studio interface with three selection boxes:

- To connect to the on-premises data:**
  - A point-to-site VPN connection
  - A site-to-site VPN connection
  - Azure App Service Hybrid Connections
- To orchestrate the phases:**
  - A Machine Learning experiment
  - Azure Machine Learning Studio
  - Machine Learning pipelines
- To control the orchestration:**
  - Azure Automation
  - Azure Databricks
  - 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 38**

- (Exam Topic 2)

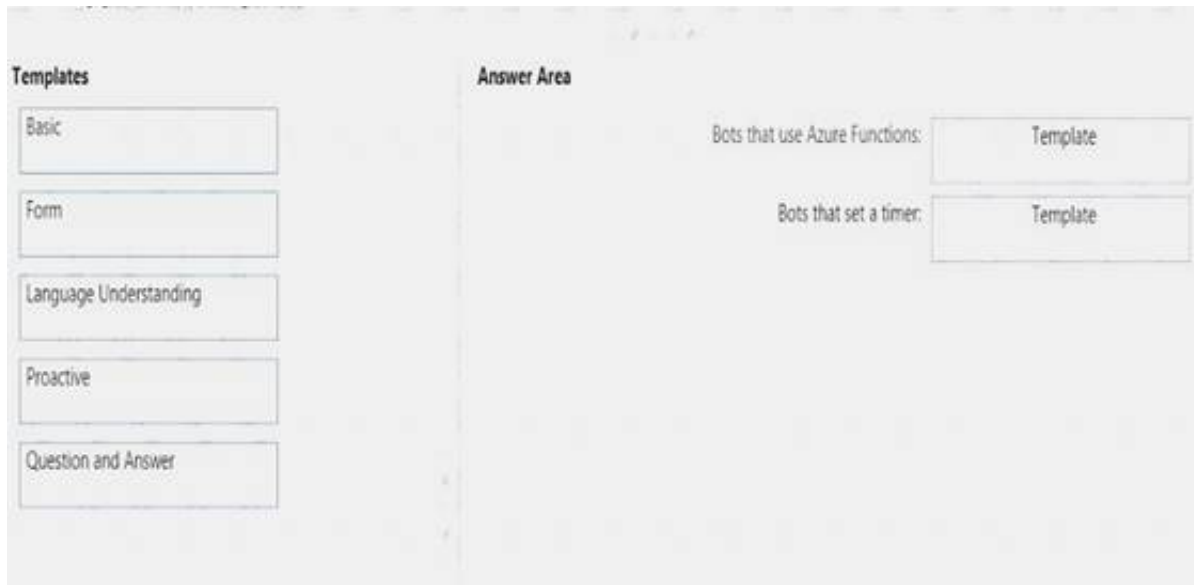
You plan to use the Microsoft Bot Framework to develop bots that will be deployed by using the Azure Bot Service.

You need to configure the Azure Bot Service to support the following types of bots:

- Bots that use Azure Functions
- Bots that set a timer

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

NOTE: Each correct selection is worth one point.



- A. Mastered  
B. Not Mastered

**Answer:** A

**Explanation:**

References:

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

**NEW QUESTION 43**

- (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 45**

- (Exam Topic 2)

You plan to design a solution for an AI implementation that uses data from IoT devices.

You need to recommend a data storage solution for the IoT devices that meets the following requirements:

- Allow data to be queried in real-time as it streams into the solution.
- Provide the lowest amount of latency for loading data into the solution. What should you include in the recommendation?

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

**Answer:** D

**Explanation:**

You can use HDInsight to process streaming data that's received in real time from a variety of devices. Internet of Things (IoT)

You can use HDInsight to build applications that extract critical insights from data. You can also use Azure Machine Learning on top of that to predict future trends for your business.

By combining enterprise-scale R analytics software with the power of Apache Hadoop and Apache Spark, Microsoft R Server for HDInsight gives you the scale and performance you need. Multi-threaded math libraries and transparent parallelization in R Server handle up to 1000x more data and up to 50x faster speeds than open-source R, which helps you to train more accurate models for better predictions.

References:

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

**NEW QUESTION 47**

- (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 48**

- (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 53**

- (Exam Topic 2)

You are configuring data persistence for a Microsoft Bot Framework application. The application requires a structured NoSQL cloud data store.

You need to identify a storage solution for the application. The solution must minimize costs. What should you identify?

A. Azure Blob storage

B. Azure Cosmos DB

C. Azure HDInsight

D. Azure Table storage

**Answer:** D

**Explanation:**

Table Storage is a NoSQL key-value store for rapid development using massive semi-structured datasets You can develop applications on Cosmos DB using popular NoSQL APIs.

Both services have a different scenario and pricing model.

While Azure Storage Tables is aimed at high capacity on a single region (optional secondary read only region but no failover), indexing by PK/RK and storage-optimized pricing; Azure Cosmos DB Tables aims for high throughput (single-digit millisecond latency), global distribution (multiple failover), SLA-backed predictive performance with automatic indexing of each attribute/property and a pricing model focused on throughput.

References:

<https://db-engines.com/en/system/Microsoft+Azure+Cosmos+DB%3BMicrosoft+Azure+Table+Storage>

**NEW QUESTION 55**

- (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 56**

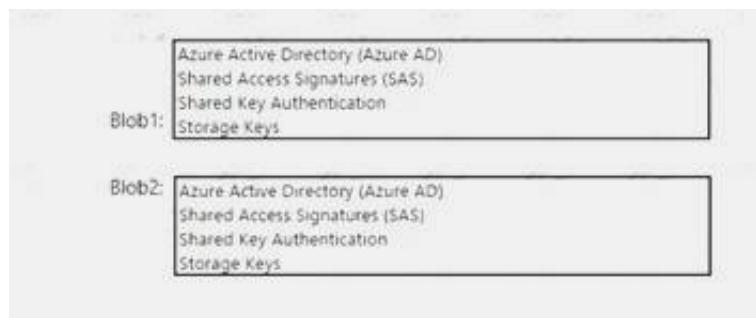
- (Exam Topic 2)

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 59**

- (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 60**

- (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 65**

.....

## 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>