

AI-100.78q

Number: AI-100  
Passing Score: 800  
Time Limit: 120 min

AI-100



**Website:** <https://vceplus.com>  
**VCE to PDF Converter:** <https://vceplus.com/vce-to-pdf/>  
**Facebook:** <https://www.facebook.com/VCE.For.All.VN/>  
**Twitter :** [https://twitter.com/VCE\\_Plus](https://twitter.com/VCE_Plus)

<https://vceplus.com/>

### Designing and Implementing an Azure AI Solution

#### Question Set 1

#### QUESTION 1

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?



<https://vceplus.com/>

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



**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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 storageoptimized 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>

## QUESTION 2

You have an Azure Machine Learning model that is deployed to a web service.

You plan to publish the web service by using the name ml.contoso.com.

You need to recommend a solution to ensure that access to the web service is encrypted.

Which three actions should you recommend? Each correct answer presents part of the solution.

**NOTE:** Each correct selection is worth one point.

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

**Correct Answer:** BDE

**Section:** [none]

**Explanation**



**Explanation/Reference:**

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

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

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

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

## QUESTION 3

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

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://github.com/MicrosoftDocs/azure-docs/blob/master/articles/machine-learning/studio/algorithm-choice.md>

#### QUESTION 4

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 methods should you recommend? (Choose two.)

- A. horizontal pod autoscaler
- B. cluster autoscaler
- C. AKS cluster virtual 32 node autoscaling
- D. Azure Container Instances

**Correct Answer:** AB

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

B: 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.

A: You can also use the horizontal pod autoscaler to automatically adjust the number of pods that run your application.

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

### QUESTION 5

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



**Correct Answer:** ABC

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 6

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

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 7

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

**Correct Answer:** A

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

An AKS cluster has one or more nodes.

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

### **QUESTION 8**

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



**Correct Answer: C**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

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>

### **QUESTION 9**

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

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Throughput provisioned for a container is divided evenly among physical partitions.

Incorrect:

Not A: Increasing request units would also improve throughput, but at a cost.

Reference: <https://docs.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning>

#### **QUESTION 10**

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

**Correct Answer:** BC

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References:

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

**QUESTION 11**

DRAG DROP

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.

**Select and Place:**



Actions	Answer Area
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).	

Correct Answer:

Actions	Answer Area
Enable encryption that uses customer-managed keys.	Configure a service endpoint for the storage account.
Upload a key to an Azure key vault.	Generate an encryption key.
Generate an encryption key.	Enable encryption that uses customer-managed keys.
Generate an access key.	
Configure a service endpoint for the storage account.	
Generate a shared access signature (SAS).	

⤴  
⤵

⤴  
⤵



**Section:** [none]  
**Explanation**

**Explanation/Reference:**  
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> **QUESTION 12**

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



**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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

**QUESTION 13**

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

**Correct Answer:** BE

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

#### **QUESTION 14**

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 diagnostics logs
- D. Azure Active Directory (Azure AD) reporting

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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

**QUESTION 15**

You create an Azure Cognitive Services resource.

A data scientist needs to call the resource from Azure Logic Apps by using the generic HTTP connector.

Which two values should you provide to the data scientist? Each correct answer presents part of the solution.

**NOTE:** Each correct selection is worth one point.

- A. Endpoint URL
- B. Resource name
- C. Access key
- D. Resource group name
- E. Subscription ID

**Correct Answer:** DE

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://social.technet.microsoft.com/wiki/contents/articles/36074.logic-apps-with-azure-cognitive-service.aspx>

**QUESTION 16**

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 queries
- C. Azure Service Bus topics and subscriptions
- D. Apache Storm on Azure HDInsight

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References:

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

**QUESTION 17**

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

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://docs.microsoft.com/en-us/azure/iot-edge/how-to-store-data-blob>

**QUESTION 18**

Your company is building custom models that integrate into microservices architecture on Azure Kubernetes Services (AKS).

The model is built by using Python and published to AKS.

You need to update the model and enable Azure Application Insights for the model.

What should you use?

- A. the Azure CLI
- B. MLNET Model Builder



- C. the Azure Machine Learning SDK
- D. the Azure portal

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You can set up Azure Application Insights for Azure Machine Learning. Application Insights gives you the opportunity to monitor:

- Request rates, response times, and failure rates.
  - Dependency rates, response times, and failure rates. ▪
- Exceptions.

Requirements include an Azure Machine Learning workspace, a local directory that contains your scripts, and the Azure Machine Learning SDK for Python installed.

References: <https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/how-to-enable-app-insights>

#### **QUESTION 19**

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 should you do? More than one answer choice may achieve the goal. (Choose two.)

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

**Correct Answer:** AB

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Increasing request units would improve throughput, but at a cost.

Throughput provisioned for a container is divided evenly among physical partitions.

References:

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

### **QUESTION 20**

Your company has several AI solutions and bots.

You need to implement a solution to monitor the utilization of the bots. The solution must ensure that analysts at the company can generate dashboards to review the utilization.

What should you include in the solution?

- A. Azure Application Insights
- B. Azure Data Explorer
- C. Azure Logic Apps
- D. Azure Monitor



**Correct Answer: A**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Bot Analytics.

Analytics is an extension of Application Insights. Application Insights provides service-level and instrumentation data like traffic, latency, and integrations. Analytics provides conversation-level reporting on user, message, and channel data.

References:

<https://docs.microsoft.com/en-us/azure/bot-service/bot-service-manage-analytics>

**Question Set 1**

### QUESTION 1

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

**Correct Answer: C**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 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

**Correct Answer:** CE

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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-data-architecture>

### QUESTION 3

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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

**QUESTION 4**

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

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

**QUESTION 5**

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 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

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://docs.microsoft.com/en-us/azure/hdinsight/hdinsight-streaming-at-scale-overview?toc=https%3A%2F%2Fdocs.microsoft.com%2Fen-us%2Fazure%2Fhdinsight%2Fhadoop%2FTOC.json&bc=https%3A%2F%2Fdocs.microsoft.com%2Fen-us%2Fazure%2Fbread%2Ftoc.json>

### QUESTION 6

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

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References:

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

### QUESTION 7

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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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

#### **QUESTION 8**

HOTSPOT



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.

**Hot Area:**

## Answer Area

Service to use:

	▼
Apache Hive	
Azure Data Factory	
Azure Functions	
Azure Stream Analytics	

Location to deploy the job:

	▼
A Web Job in Azure	
An Azure IoT Edge device	
Azure Event Hubs	
Azure Notification Hubs	

Correct Answer:

## Answer Area

Service to use:

	▼
Apache Hive	
Azure Data Factory	
Azure Functions	
Azure Stream Analytics	

Location to deploy the job:

	▼
A Web Job in Azure	
An Azure IoT Edge device	
Azure Event Hubs	
Azure Notification Hubs	

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 9

**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 an app named App1 that uses the Face API.

App1 contains several PersonGroup objects.

You discover that a PersonGroup object for an individual named Ben Smith cannot accept additional entries. The PersonGroup object for Ben Smith contains 10,000 entries.

You need to ensure that additional entries can be added to the PersonGroup object for Ben Smith. The solution must ensure that Ben Smith can be identified by all the entries.

Solution: You modify the custom time interval for the training phase of App1.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Instead, use a LargePersonGroup. LargePersonGroup and LargeFaceList are collectively referred to as large-scale operations. LargePersonGroup can contain up to 1 million persons, each with a maximum of 248 faces. LargeFaceList can contain up to 1 million faces. The large-scale operations are similar to the conventional PersonGroup and FaceList but have some differences because of the new architecture.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale>

### QUESTION 10

**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 an app named App1 that uses the Face API.

App1 contains several PersonGroup objects.

You discover that a PersonGroup object for an individual named Ben Smith cannot accept additional entries. The PersonGroup object for Ben Smith contains 10,000 entries.

You need to ensure that additional entries can be added to the PersonGroup object for Ben Smith. The solution must ensure that Ben Smith can be identified by all the entries.

Solution: You create a second PersonGroup object for Ben Smith.

Does this meet the goal?

- A. Yes
- B. No



**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Instead, use a LargePersonGroup. LargePersonGroup and LargeFaceList are collectively referred to as large-scale operations. LargePersonGroup can contain up to 1 million persons, each with a maximum of 248 faces. LargeFaceList can contain up to 1 million faces. The large-scale operations are similar to the conventional PersonGroup and FaceList but have some differences because of the new architecture.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale>

#### **QUESTION 11**

**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 an app named App1 that uses the Face API.

App1 contains several PersonGroup objects.

You discover that a PersonGroup object for an individual named Ben Smith cannot accept additional entries. The PersonGroup object for Ben Smith contains 10,000 entries.

You need to ensure that additional entries can be added to the PersonGroup object for Ben Smith. The solution must ensure that Ben Smith can be identified by all the entries.

Solution: You migrate all the entries to the LargePersonGroup object for Ben Smith.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

LargePersonGroup and LargeFaceList are collectively referred to as large-scale operations. LargePersonGroup can contain up to 1 million persons, each with a maximum of 248 faces. LargeFaceList can contain up to 1 million faces. The large-scale operations are similar to the conventional PersonGroup and FaceList but have some differences because of the new architecture.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/face/face-api-how-to-topics/how-to-use-large-scale>

## QUESTION 12

Your company plans to develop a mobile app to provide meeting transcripts by using speech-to-text. Audio from the meetings will be streamed to provide real-time transcription.

You need to recommend which task each meeting participant must perform to ensure that the transcripts of the meetings can identify all participants.

Which task should you recommend?

- A. Record the meeting as an MP4.
- B. Create a voice signature.

- C. Sign up for Azure Speech Services.
- D. Sign up as a guest in Azure Active Directory (Azure AD)

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

The first step is to create voice signatures for the conversation participants. Creating voice signatures is required for efficient speaker identification.

Note: In addition to the standard baseline model used by the Speech Services, you can customize models to your needs with available data, to overcome speech recognition barriers such as speaking style, vocabulary and background noise.

References: <https://docs.microsoft.com/bs-latn-ba/azure/cognitive-services/speech-service/how-to-use-conversation-transcription-service>

### QUESTION 13

You are designing an AI solution that will provide feedback to teachers who train students over the Internet. The students will be in classrooms located in remote areas. The solution will capture video and audio data of the students in the classrooms.

You need to recommend Azure Cognitive Services for the AI solution to meet the following requirements:

- Alert teachers if a student facial expression indicates the student is angry or scared.
- Identify each student in the classrooms for attendance purposes. ▪

Allow the teachers to log voice conversations as text.

Which Cognitive Services should you recommend?

- A. Face API and Text Analytics
- B. Computer Vision and Text Analytics
- C. QnA Maker and Computer Vision
- D. Speech to Text and Face API

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Speech-to-text from Azure Speech Services, also known as speech-to-text, enables real-time transcription of audio streams into text that your applications, tools, or devices can consume, display, and take action on as command input.

Face detection: Detect one or more human faces in an image and get back face rectangles for where in the image the faces are, along with face attributes which contain machine learning-based predictions of facial features. The face attribute features available are: Age, Emotion, Gender, Pose, Smile, and Facial Hair along with 27 landmarks for each face in the image.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/speech-to-text> <https://azure.microsoft.com/en-us/services/cognitive-services/face/>

#### QUESTION 14

You need to evaluate trends in fuel prices during a period of 10 years. The solution must identify unusual fluctuations in prices and produce visual representations.

Which Azure Cognitive Services API should you use?

- A. Anomaly Detector
- B. Computer Vision
- C. Text Analytics
- D. Bing Autosuggest



**Correct Answer:** A

**Section:** [none]

**Explanation**

#### Explanation/Reference:

Explanation:

The Anomaly Detector API enables you to monitor and detect abnormalities in your time series data with machine learning. The Anomaly Detector API adapts by automatically identifying and applying the best-fitting models to your data, regardless of industry, scenario, or data volume. Using your time series data, the API determines boundaries for anomaly detection, expected values, and which data points are anomalies.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/anomaly-detector/overview>

#### QUESTION 15

You plan to perform analytics of the medical records of patients located around the world.

You need to recommend a solution that avoids storing and processing data in the cloud.

What should you include in the recommendation?

- A. Azure Machine Learning Studio
- B. the Text Analytics API that has container support
- C. Azure Machine Learning services
- D. an Apache Spark cluster that uses MMLSpark

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

The Microsoft Machine Learning Library for Apache Spark (MMLSpark) assists in provisioning scalable machine learning models for large datasets, especially for building deep learning problems. MMLSpark works with SparkML pipelines, including Microsoft CNTK and the OpenCV library, which provide end-to-end support for the ingress and processing of image input data, categorization of images, and text analytics using pre-trained deep learning algorithms.

References:

[https://subscription.packtpub.com/book/big\\_data\\_and\\_business\\_intelligence/9781789131956/10/ch10lv1sec61/an-overview-of-the-microsoft-machine-learning-library-for-apache-spark-mmlspark](https://subscription.packtpub.com/book/big_data_and_business_intelligence/9781789131956/10/ch10lv1sec61/an-overview-of-the-microsoft-machine-learning-library-for-apache-spark-mmlspark)

**QUESTION 16**

Your company has an on-premises datacenter.

You plan to publish an app that will recognize a set of individuals by using the Face API. The model is trained.

You need to ensure that all images are processed in the on-premises datacenter.

What should you deploy to host the Face API?

- A. a Docker container
- B. Azure File Sync
- C. Azure Application Gateway
- D. Azure Data Box Edge

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. A Docker container image is a lightweight, standalone, executable package of software that includes everything needed to run an application: code, runtime, system tools, system libraries and settings.

Incorrect Answers:

D: Azure Data Box Edge is an AI-enabled edge computing device with network data transfer capabilities. This article provides you an overview of the Data Box Edge solution, benefits, key capabilities, and the scenarios where you can deploy this device.

Data Box Edge is a Hardware-as-a-service solution. Microsoft ships you a cloud-managed device with a built-in Field Programmable Gate Array (FPGA) that enables accelerated AI-inferencing and has all the capabilities of a storage gateway.

References: <https://www.docker.com/resources/what-container>

**QUESTION 17**

You have a Bing Search service that is used to query a product catalog.

You need to identify the following information:

- The locale of the query
- The top 50 query strings
- The number of calls to the service
- The top geographical regions of the service What

should you implement?

- A. Bing Statistics
- B. Azure API Management (APIM)
- C. Azure Monitor
- D. Azure Application Insights

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:



The Bing Statistics add-in provides metrics such as call volume, top queries, API response, code distribution, and market distribution. The rich slicing-and-dicing capability lets you gather deeper understanding of your users and their usage to inform your business strategy.

References:

<https://www.bingapistatistics.com/>

### QUESTION 18

You have a Face API solution that updates in real time. A pilot of the solution runs successfully on a small dataset.

When you attempt to use the solution on a larger dataset that continually changes, the performance degrades, slowing how long it takes to recognize existing faces.

You need to recommend changes to reduce the time it takes to recognize existing faces without increasing costs.

What should you recommend?

- A. Change the solution to use the Computer Vision API instead of the Face API.
- B. Separate training into an independent pipeline and schedule the pipeline to run daily.
- C. Change the solution to use the Bing Image Search API instead of the Face API.
- D. Distribute the face recognition inference process across many Azure Cognitive Services instances.

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Incorrect Answers:

A: The purpose of Computer Vision is to inspect each image associated with an incoming article to (1) scrape out written words from the image and (2) determine what types of objects are present in the image.

C: The Bing API provides an experience similar to Bing.com/search by returning search results that Bing determines are relevant to a user's query. The results include Web pages and may also include images, videos, and more.

D: That would increase cost.

References: <https://github.com/Azure/cognitive-services>

### QUESTION 19

HOTSPOT

You plan to create a bot that will support five languages. The bot will be used by users located in three different countries. The bot will answer common customer questions. The bot will use Language Understanding (LUIS) to identify which skill to use and to detect the language of the customer.

You need to identify the minimum number of Azure resources that must be created for the planned bot.

How many QnA Maker, LUIS and Language Detection instances should you create? To answer, select the appropriate options in the answer area.

**NOTE:** Each correct selection is worth one point.

**Hot Area:**

## Answer Area

QnA Maker:	<input type="text" value="1"/>	▼
	<input type="text" value="3"/>	
	<input type="text" value="5"/>	
LUIS:	<input type="text" value="1"/>	▼
	<input type="text" value="3"/>	
	<input type="text" value="5"/>	
Language Detection:	<input type="text" value="1"/>	▼
	<input type="text" value="3"/>	
	<input type="text" value="5"/>	

**Correct Answer:**

## Answer Area

QnA Maker: 

1	▼
3	
5	

LUIS: 

1	▼
3	
5	

Language Detection: 

1.com	▼
3	
5	

**Section:** [none]

**Explanation**

**Explanation/Reference:**

QnA Maker: 5

If the user plans to support multiple languages, they need to have a new QnA Maker resource for each language.

LUIS: 5

If you need a multi-language LUIS client application such as a chatbot, you have a few options. If LUIS supports all the languages, you develop a LUIS app for each language. Each LUIS app has a unique app ID, and endpoint log. If you need to provide language understanding for a language LUIS does not support, you can use Microsoft Translator API to translate the utterance into a supported language, submit the utterance to the LUIS endpoint, and receive the resulting scores.

Language detection: 1

The Language Detection feature of the Azure Text Analytics REST API evaluates text input for each document and returns language identifiers with a score that indicates the strength of the analysis.

This capability is useful for content stores that collect arbitrary text, where language is unknown. You can parse the results of this analysis to determine which language is used in the input document. The response also returns a score that reflects the confidence of the model. The score value is between 0 and 1.

The Language Detection feature can detect a wide range of languages, variants, dialects, and some regional or cultural languages. The exact list of languages for this feature isn't published.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/qnamaker/overview/language-support>  
<https://docs.microsoft.com/en-us/azure/cognitive-services/luis/luis-language-support> <https://docs.microsoft.com/en-us/azure/cognitive-services/text-analytics/how-tos/text-analytics-how-to-language-detection>

#### QUESTION 20

You have a database that contains sales data.

You plan to process the sales data by using two data streams named Stream1 and Stream2. Stream1 will be used for purchase order data. Stream2 will be used for reference data.

The reference data is stored in CSV files.

You need to recommend an ingestion solution for each data stream.

What two solutions should you recommend? Each correct answer is a complete solution.

**NOTE:** Each correct selection is worth one point.

- A. an Azure event hub for Stream1 and Azure Blob storage for Stream2
- B. Azure Blob storage for Stream1 and Stream2
- C. an Azure event hub for Stream1 and Stream2
- D. Azure Blob storage for Stream1 and Azure Cosmos DB for Stream2
- E. Azure Cosmos DB for Stream1 and an Azure event hub for Stream2

**Correct Answer:** AB

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Stream1 - Azure Event

Stream2 - Blob Storage

Azure Event Hubs is a highly scalable data streaming platform and event ingestion service, capable of receiving and processing millions of events per second. Event Hubs can process and store events, data, or telemetry produced by distributed software and devices. Data sent to an event hub can be transformed and stored using any real-time analytics provider or batching/storage adapters. Event Hubs provides publish-subscribe capabilities with low latency at massive scale, which makes it appropriate for big data scenarios.

Stream1, Stream2 - Blob Storage

Stream Analytics has first-class integration with Azure data streams as inputs from three kinds of resources:

Azure Event Hubs

Azure IoT Hub

Azure Blob storage

These input resources can live in the same Azure subscription as your Stream Analytics job or a different subscription.

References: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/real-time-ingestion>

## QUESTION 21

**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 create a managed identity for AKS, and then you create an SSH connection.



<https://vceplus.com/>

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

#### QUESTION 22

**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 change the permissions of the AKS resource group, and then you create an SSH connection.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 23

**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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

**QUESTION 24**

**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 need to create an IoT solution that performs the following tasks:

- Identifies hazards
- Provides a real-time online dashboard
- Takes images of an area every minute
- Counts the number of people in an area every minute

Solution: You configure the IoT devices to send the images to an Azure IoT hub, and then you configure an Azure Functions call to Azure Cognitive Services that sends the results to an Azure event hub. You configure Microsoft Power BI to connect to the event hub by using Azure Stream Analytics.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Instead use Cognitive Services containers on the IoT devices.

References: <https://azure.microsoft.com/es-es/blog/running-cognitive-services-on-iot-edge/> <https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-live-data-visualization-in-power-bi>

### QUESTION 25

**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 need to create an IoT solution that performs the following tasks:

- Identifies hazards
- Provides a real-time online dashboard
- Takes images of an area every minute
- Counts the number of people in an area every minute

Solution: You configure the IoT devices to send the images to an Azure IoT hub, and then you configure an Azure Automation call to Azure Cognitive Services that sends the results to an Azure event hub. You configure Microsoft Power BI to connect to the event hub by using Azure Stream Analytics.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Instead use Cognitive Services containers on the IoT devices.

References: <https://azure.microsoft.com/es-es/blog/running-cognitive-services-on-iot-edge/> <https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-live-data-visualization-in-power-bi> **Testlet 2**

**Overview**

Contoso, Ltd. has an office in New York to serve its North American customers and an office in Paris to serve its European customers.

## Existing Environment

### Infrastructure

Each office has a small data center that hosts Active Directory services and a few off-the-shelf software solutions used by internal users.

The network contains a single Active Directory forest that contains a single domain named contoso.com. Azure Active Directory (Azure AD) Connect is used to extend identity management to Azure.

The company has an Azure subscription. Each office has an Azure ExpressRoute connection to the subscription. The New York office connects to a virtual network hosted in the US East 2 Azure region. The Paris office connects to a virtual network hosted in the West Europe Azure region.

The New York office has an Azure Stack Development Kit (ASDK) deployment that is used for development and testing.

### Current Business Model

Contoso has a web app named Bookings hosted in an App Service Environment (ASE). The ASE is in the virtual network in the East US 2 region. Contoso employees and customers use Bookings to reserve hotel rooms.

### Data Environment

Bookings connects to a Microsoft SQL Server database named hotelDB in the New York office.

The database has a view named vwAvailability that consolidates columns from the tables named Hotels, Rooms, and RoomAvailability. The database contains data that was collected during the last 20 years.

### Problem Statements

Contoso identifies the following issues with its current business model:

- European users report that access to Booking is slow, and they lose customers who must wait on the phone while they search for available rooms.
- Users report that Bookings was unavailable during an outage in the New York data center for more than 24 hours.

### Requirements

Contoso identifies the following issues with its current business model:

- European users report that access to Bookings is slow, and they lose customers who must wait on the phone while they search for available rooms.
- Users report that Bookings was unavailable during on outage in the New York data center for more than 24 hours.

## Business Goals

Contoso wants to provide a new version of the Bookings app that will provide a highly available, reliable service for booking travel packages by interacting with a chatbot named Butler.

Contoso plans to move all production workloads to the cloud.

## Technical requirements

Contoso identifies the following technical requirements:

- Data scientists must test Butler by using ASDK.
  - Whenever possible, solutions must minimize costs.
  - Butler must greet users by name when they first connect.
  - Butler must be able to handle up to 10,000 messages a day.
  - Butler must recognize the users' intent based on basic utterances.
  - All configurations to the Azure Bot Service must be logged centrally.
  - Whenever possible, solutions must use the principle of least privilege.
  - Internal users must be able to access Butler by using Microsoft Skype for Business.
  - The new Bookings app must provide a user interface where users can interact with Butler.
  - Users in an Azure AD group named KeyManagers must be able to manage keys for all Azure Cognitive Services.
  - Butler must provide users with the ability to reserve a room, cancel a reservation, and view existing reservations.
  - The new Bookings app must be available to users in North America and Europe if a single data center or Azure region fails.
  - For continuous improvement, you must be able to test Butler by sending sample utterances and comparing the chatbot's responses to the actual intent. ▪
- You must maintain relationships between data after migration.

## QUESTION 1

You need to design the Butler chatbot solution to meet the technical requirements.

What is the best channel and pricing tier to use? More than one answer choice may achieve the goal. Select the **BEST** answer.

- A. Standard channels that use the S1 pricing tier
- B. Standard channels that use the Free pricing tier
- C. Premium channels that use the Free pricing tier
- D. Premium channels that use the S1 pricing tier

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References:

<https://azure.microsoft.com/en-in/pricing/details/bot-service/>



## Question Set 1

### QUESTION 1

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



<https://vceplus.com/>

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Use Azure Notebooks to develop a machine learning module and deploy it to a Linux device running Azure IoT Edge. You can use IoT Edge modules to deploy code that implements your business logic directly to your IoT Edge devices.

References: <https://docs.microsoft.com/en-us/azure/iot-edge/tutorial-deploy-machine-learning>

## QUESTION 2

Your company recently purchased several hundred hardware devices that contain 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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://azure.microsoft.com/en-us/resources/samples/functions-js-iot-hub-processing/>

## QUESTION 3

You have thousands of images that contain text.

You need to process the text from the images to a machine-readable character stream.

Which Azure Cognitive Services service should you use?

- A. the Image Moderation API
- B. Text Analytics
- C. Translator Text
- D. Computer Vision

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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.

Incorrect Answers:

A: Use Content Moderator's machine-assisted image moderation and human-in-the-loop Review tool to moderate images for adult and racy content. Scan images for text content and extract that text, and detect faces. You can match images against custom lists, and take further action.

Reference: <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>

#### QUESTION 4

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 process events in the cloud.

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



**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Use Azure Notebooks to develop a machine learning module and deploy it to a Linux device running Azure IoT Edge. You can use IoT Edge modules to deploy code that implements your business logic directly to your IoT Edge devices.

Reference: <https://docs.microsoft.com/en-us/azure/iot-edge/tutorial-deploy-machine-learning>

#### QUESTION 5

**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 scoring accuracy of each run of the model.

Solution: You modify the Config.json file.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Instead update the manifest file.

Reference: [https://azure.github.io/learnAnalytics-UsingAzureMachineLearningforAIWorkloads/lab07-deploying\\_a\\_scoring\\_service\\_to\\_aks/0\\_README.html](https://azure.github.io/learnAnalytics-UsingAzureMachineLearningforAIWorkloads/lab07-deploying_a_scoring_service_to_aks/0_README.html)

### **QUESTION 6**

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

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 7

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

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

**QUESTION 8**

HOTSPOT

You are designing an AI solution that must meet the following processing requirements:

- Use a parallel processing framework that supports the in-memory processing of high volumes of data. ▪
- Use in-memory caching and a columnar storage engine for Apache Hive queries.

What should you use to meet each requirement? To answer, select the appropriate options in the answer area.

**NOTE:** Each correct selection is worth one point.

**Hot Area:**



## Answer Area

Use a parallel processing framework that supports the in-memory processing of high volumes of data:

	▼
Apache Kafka	
Apache Spark	
Hive	
Microsoft Machine	
Learning Server	

Use in-memory caching and a columnar storage engine for Hive queries:

	▼
Apache Kafka	
Apache Spark	
Interactive Query	
Microsoft Machine	
Learning Server	

Correct Answer:

## Answer Area

Use a parallel processing framework that supports the in-memory processing of high volumes of data:

	▼
Apache Kafka	
Apache Spark	
Hive	
Microsoft Machine	
Learning Server	

Use in-memory caching and a columnar storage engine for Hive queries:

	▼
Apache Kafka	
Apache Spark	
Interactive Query	
Microsoft Machine	
Learning Server	

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Box 1: Apache Spark

Apache Spark is a parallel processing framework that supports in-memory processing to boost the performance of big-data analytic applications. Apache Spark in Azure HDInsight is the Microsoft implementation of Apache Spark in the cloud.

Box 2: Interactive Query

Interactive Query provides In-memory caching and improved columnar storage engine for Hive queries.

References: <https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-overview> <https://docs.microsoft.com/bs-latn-ba/azure/hdinsight/interactive-query/apache-interactive-query-get-started>

### QUESTION 9

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.

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 10

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
- E. Apache Kafka
- F. Azure Event Hubs

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

Stream processing can be handled by Azure Stream Analytics. Azure Stream Analytics can run perpetual queries against an unbounded stream of data. These queries consume streams of data from storage or message brokers, filter and aggregate the data based on temporal windows, and write the results to sinks such as storage, databases, or directly to reports in Power BI. Stream Analytics uses a SQL-based query language that supports temporal and geospatial constructs, and can be extended using JavaScript.

Incorrect Answers:

E: Apache Kafka is used for ingestion, not for stream processing.

F: Azure Event Hubs is used for ingestion, not for stream processing.

Reference:

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

### QUESTION 11

You are designing an AI solution that will provide feedback to teachers who train students over the Internet. The students will be in classrooms located in remote areas. The solution will capture video and audio data of the students in the classrooms.

You need to recommend Azure Cognitive Services for the AI solution to meet the following requirements:

- Alert teachers if a student seems angry or distracted.
- Identify each student in the classrooms for attendance purposes.
- Allow the teachers to log the text of conversations between themselves and the students.

Which Cognitive Services should you recommend?

- A. Computer Vision, Text Analytics, and Face API
- B. Video Indexer, Face API, and Text Analytics
- C. Computer Vision, Speech to Text, and Text Analytics

- D. Text Analytics, QnA Maker, and Computer Vision
- E. Video Indexer, Speech to Text, and Face API

**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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.

Face API enables you to search, identify, and match faces in your private repository of up to 1 million people.

The Face API now integrates emotion recognition, returning the confidence across a set of emotions for each face in the image such as anger, contempt, disgust, fear, happiness, neutral, sadness, and surprise. These emotions are understood to be cross-culturally and universally communicated with particular facial expressions.

Speech-to-text from Azure Speech Services, also known as speech-to-text, enables real-time transcription of audio streams into text that your applications, tools, or devices can consume, display, and take action on as command input. This service is powered by the same recognition technology that Microsoft uses for Cortana and Office products, and works seamlessly with the translation and text-to-speech.

Incorrect Answers:

Computer Vision or the QnA is not required.

References:

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

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

<https://docs.microsoft.com/en-us/azure/cognitive-services/speech-service/speech-to-text>

## QUESTION 12

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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 13

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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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.

Incorrect Answers:

B: Use the Text Moderation API in Azure Content Moderator to scan your text content. The operation scans your content for profanity, and compares the content against custom and shared blacklists.

References: <https://docs.microsoft.com/en-us/azure/cognitive-services/content-moderator/try-terms-list-api>

#### QUESTION 14

You need to configure versioning and logging for Azure Machine Learning models.

Which Machine Learning service application should you use?

- A. Models
- B. Activities
- C. Experiments
- D. Pipelines
- E. Deployments



**Correct Answer:** E

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://docs.microsoft.com/en-us/azure/machine-learning/service/how-to-enable-logging#logging-for-deployed-models> **Testlet 2**

#### Overview

Contoso, Ltd. has an office in New York to serve its North American customers and an office in Paris to serve its European customers.

#### Existing Environment

##### Infrastructure

Each office has a small data center that hosts Active Directory services and a few off-the-shelf software solutions used by internal users.

The network contains a single Active Directory forest that contains a single domain named contoso.com. Azure Active Directory (Azure AD) Connect is used to extend identity management to Azure.

The company has an Azure subscription. Each office has an Azure ExpressRoute connection to the subscription. The New York office connects to a virtual network hosted in the US East 2 Azure region. The Paris office connects to a virtual network hosted in the West Europe Azure region.

The New York office has an Azure Stack Development Kit (ASDK) deployment that is used for development and testing.

### **Current Business Model**

Contoso has a web app named Bookings hosted in an App Service Environment (ASE). The ASE is in the virtual network in the East US 2 region. Contoso employees and customers use Bookings to reserve hotel rooms.

### **Data Environment**

Bookings connects to a Microsoft SQL Server database named hotelDB in the New York office.

The database has a view named vwAvailability that consolidates columns from the tables named Hotels, Rooms, and RoomAvailability. The database contains data that was collected during the last 20 years.

### **Problem Statements**

Contoso identifies the following issues with its current business model:

- European users report that access to Booking is slow, and they lose customers who must wait on the phone while they search for available rooms.
- Users report that Bookings was unavailable during an outage in the New York data center for more than 24 hours.

### **Requirements**

Contoso identifies the following issues with its current business model:

- European users report that access to Bookings is slow, and they lose customers who must wait on the phone while they search for available rooms.
- Users report that Bookings was unavailable during on outage in the New York data center for more than 24 hours.

### **Business Goals**

Contoso wants to provide a new version of the Bookings app that will provide a highly available, reliable service for booking travel packages by interacting with a chatbot named Butler.

Contoso plans to move all production workloads to the cloud.

### **Technical requirements**

Contoso identifies the following technical requirements:

- Data scientists must test Butler by using ASDK.
  - Whenever possible, solutions must minimize costs.
  - Butler must greet users by name when they first connect.
  - Butler must be able to handle up to 10,000 messages a day.
  - Butler must recognize the users' intent based on basic utterances.
  - All configurations to the Azure Bot Service must be logged centrally.
  - Whenever possible, solutions must use the principle of least privilege.
  - Internal users must be able to access Butler by using Microsoft Skype for Business.
  - The new Bookings app must provide a user interface where users can interact with Butler.
  - Users in an Azure AD group named KeyManagers must be able to manage keys for all Azure Cognitive Services.
  - Butler must provide users with the ability to reserve a room, cancel a reservation, and view existing reservations.
  - The new Bookings app must be available to users in North America and Europe if a single data center or Azure region fails.
  - For continuous improvement, you must be able to test Butler by sending sample utterances and comparing the chatbot's responses to the actual intent. ▪
- You must maintain relationships between data after migration.

#### QUESTION 1

Which two services should be implemented so that Butler can find available rooms on the technical requirements? Each correct answer presents part of the solution.

**NOTE:** Each correct selection is worth one point.

- A. QnA Maker
- B. Bing Entity Search
- C. Language Understanding (LUIS)
- D. Azure Search
- E. Content Moderator

**Correct Answer:** AC

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References:

<https://azure.microsoft.com/en-in/services/cognitive-services/language-understanding-intelligent-service/>

## Question Set 1

### QUESTION 1

Your company develops an API 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.

**Correct Answer:** ABC

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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



### QUESTION 2

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

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

You have Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy an Azure Machine Learning model as an IoT Edge module.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** A  
**Section:** [none]  
**Explanation**

**Explanation/Reference:**

Explanation:

You can use IoT Edge modules to deploy code that implements your business logic directly to your IoT Edge devices. For example, you can deploy an Azure Machine Learning module that predicts when a device fails based on simulated machine temperature data.

References: <https://docs.microsoft.com/bs-latn-ba/azure/iot-edge/tutorial-deploy-machine-learning>

**QUESTION 3**

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

**Correct Answer:** D  
**Section:** [none]  
**Explanation**

**Explanation/Reference:**

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

**QUESTION 4**

You are developing an application that will perform optical character recognition of photos of medical logbooks.

You need to recommend a solution to validate the data against a validated set of records.

Which service should you include in the recommendation?

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

**Correct Answer:** D

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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

#### QUESTION 5

You are designing an AI application that will perform real-time processing by using Microsoft Azure Stream Analytics.

You need to identify the valid outputs of a Stream Analytics job.

What are three possible outputs? Each correct answer presents a complete solution.

**NOTE:** Each correct selection is worth one point.

- A. A Hive table in Azure HDInsight
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. Azure Blob storage
- E. Azure Redis Cache

**Correct Answer:** BCD

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-define-outputs>

#### QUESTION 6

Your company plans to implement an AI solution that will analyze data from IoT devices. Data from the devices will be analyzed 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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

References:

<https://www.linkedin.com/pulse/getting-started-azure-iot-services-stream-analytics-rob-tiffany>

#### QUESTION 7

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

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

You have Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy Azure Functions as an IoT Edge module.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

**Section: [none]**

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 8

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

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

You have Azure IoT Edge devices that generate streaming data.

On the devices, you need to detect anomalies in the data by using Azure Machine Learning models. Once an anomaly is detected, the devices must add information about the anomaly to the Azure IoT Hub stream.

Solution: You deploy Azure Stream Analytics as an IoT Edge module.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

Available in both the cloud and Azure IoT Edge, Azure Stream Analytics offers built-in machine learning based anomaly detection capabilities that can be used to monitor the two most commonly occurring anomalies: temporary and persistent.

Stream Analytics supports user-defined functions, via REST API, that call out to Azure Machine Learning endpoints.

References: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-machine-learning-anomaly-detection>

### QUESTION 9

**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



**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

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>

### QUESTION 10

**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 enable AppInsights diagnostics.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: B**

**Section: [none]**

**Explanation**

**Explanation/Reference:**

Explanation:

You need to enable Model data collection.

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

#### **QUESTION 11**

**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

**Correct Answer:** B

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

You need to enable Model data collection.

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

#### **QUESTION 12**

Your company has recently purchased and deployed 25,000 IoT devices.

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

- Each device must use its own credentials for identity.
- Each device must be able to route data to multiple endpoints.
- The solution must require the minimum amount of customized code.

What should you include in the recommendation?

- A. Microsoft Azure Notification Hubs
- B. Microsoft Azure Event Hubs
- C. Microsoft Azure IoT Hub
- D. Microsoft Azure Service Bus

**Correct Answer:** C

**Section:** [none]

**Explanation**

**Explanation/Reference:**

Explanation:

An IoT hub has a default built-in endpoint. You can create custom endpoints to route messages to by linking other services in your subscription to the hub. Individual devices connect using credentials stored in the IoT hub's identity registry.

References: <https://docs.microsoft.com/en-us/azure/iot-hub/iot-hub-devguide-security>

### QUESTION 13

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

**Correct Answer:** D

**Section:** [none]

**Explanation**

#### **Explanation/Reference:**

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>

**QUESTION 14**  
DRAG DROP

You create an image classification model in Azure Machine Learning Studio.

You need to deploy the model as a containerized web service.

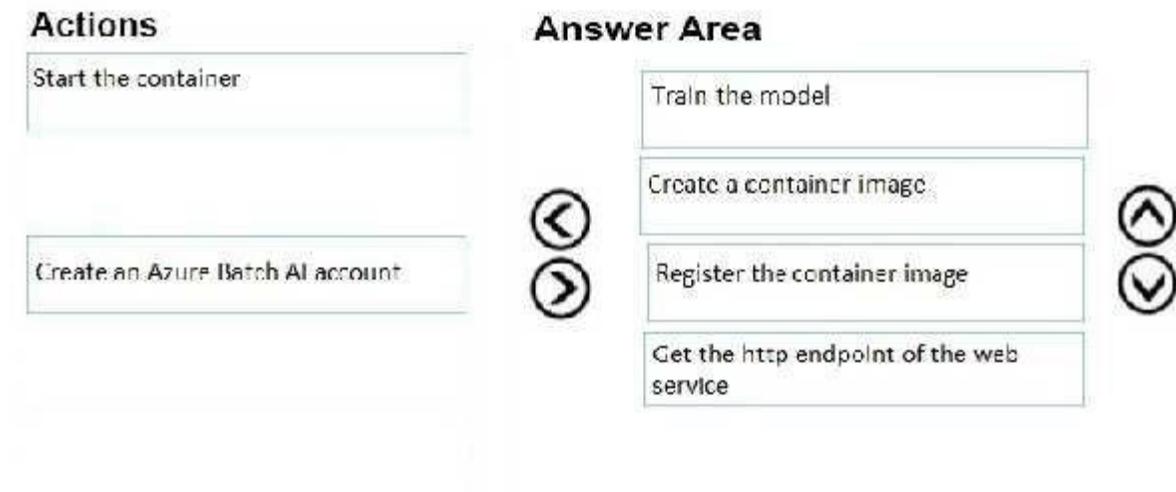
Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

**Select and Place:**

Actions	Answer Area
Start the container	
Create a container image	
Create an Azure Batch AI account	
Get the http endpoint of the web service	
Register the container image	
Train the model	



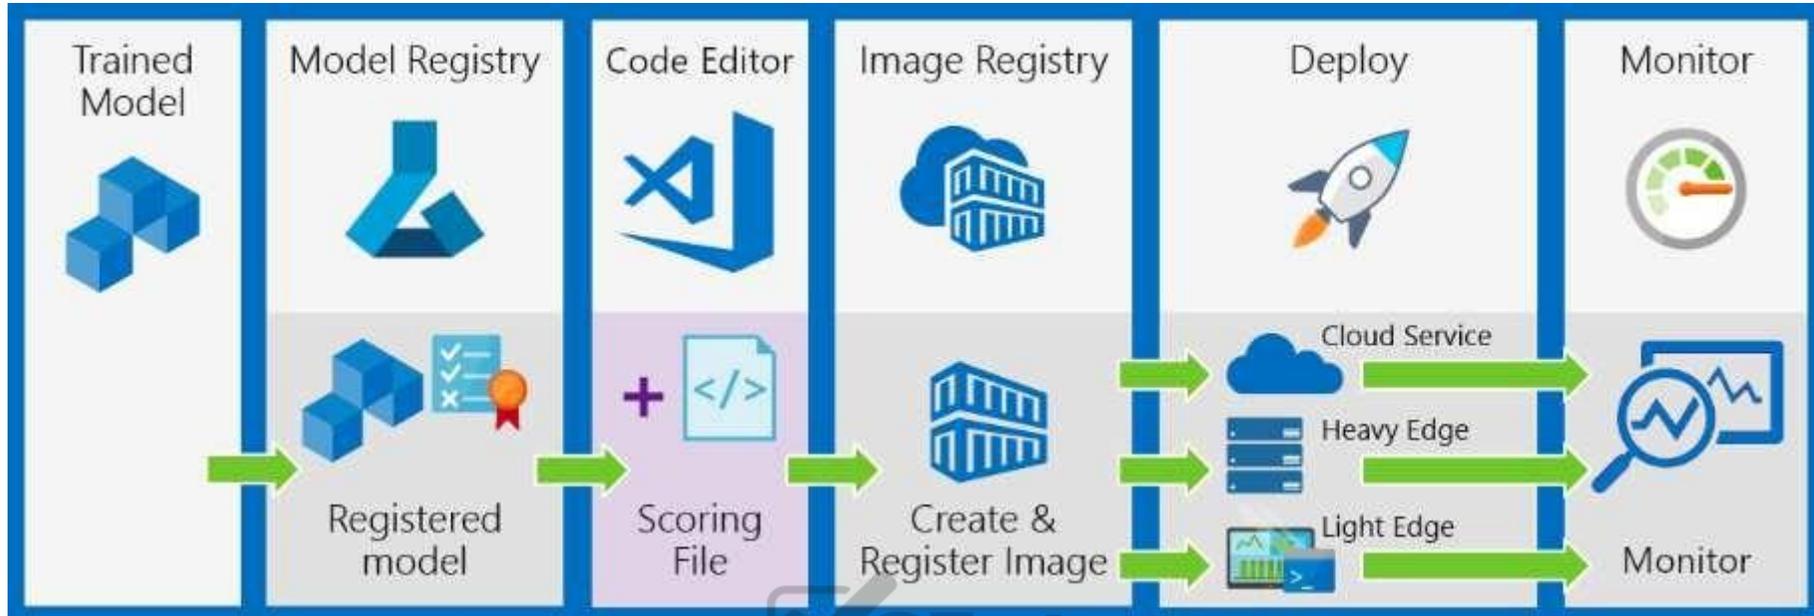
**Correct Answer:**



**Section:** [none]  
**Explanation**

**Explanation/Reference:**  
Explanation:

The following diagram illustrates the complete deployment workflow:



The deployment workflow includes the following steps:

1. Register the model in a registry hosted in your Azure Machine Learning Service workspace
2. Register an image that pairs a model with a scoring script and dependencies in a portable container
3. Deploy the image as a web service in the cloud or to edge devices
4. Monitor and collect data
5. Update a deployment to use a new image.

References: <https://docs.microsoft.com/bs-latn-ba/azure/machine-learning/service/concept-model-management-and-deployment#step-3-deploy-image>

### QUESTION 15

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 data gateway
- C. Azure Data Factory
- D. a network gateway

**Correct Answer:** B

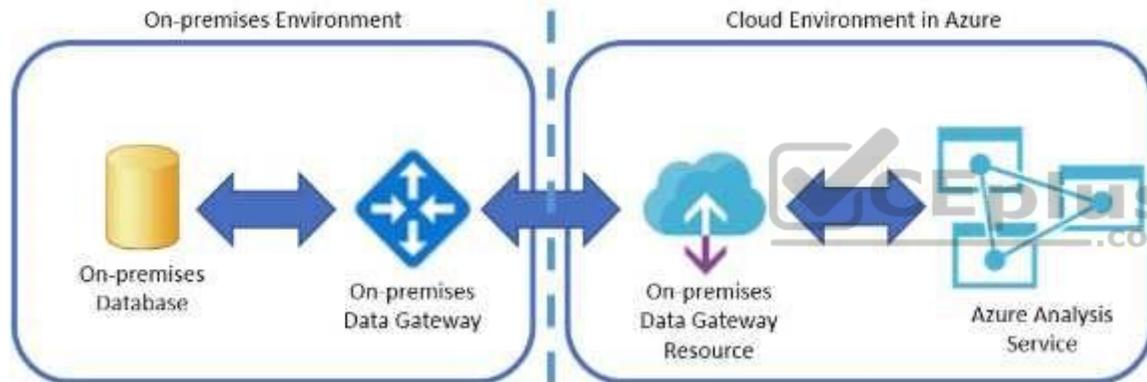
**Section:** [none]

**Explanation**

**Explanation/Reference:**

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/>

#### QUESTION 16

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

**Correct Answer:** A

**Section:** [none]

**Explanation**

**Explanation/Reference:**

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/>

#### QUESTION 17

**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 scoring accuracy of each run of the model.

Solution: You modify the scoring file.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer:** B

Section: [none]

Explanation

Explanation/Reference:



<https://vceplus.com/>

