AI-100.67q

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

AI-100

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?

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
Your company recently deployed several hardware devices that contain sensors.

The sensors generate new data on an hourly basis. The data generated is stored on-premises and retained for several years.

During the past two months, the sensors generated 300 GB of data.

You plan to move the data to Azure and then perform advanced analytics on the data.

You need to recommend an Azure storage solution for the data.

Which storage solution should you recommend?

A. Azure Queue storage
B. Azure Cosmos DB
C. Azure Blob storage
D. Azure SQL Database

Correct Answer: C
Section: [none]
Explanation

Explanation/Reference: References: https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/data-storage

QUESTION 4
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 5
You have several AI applications that use an Azure Kubernetes Service (AKS) cluster. The cluster supports a maximum of 32 nodes.

You discover that occasionally and unpredictably, the application requires more than 32 nodes.

You need to recommend a solution to handle the unpredictable application load.

Which scaling method should you recommend?

A. horizontal pod autoscaler
B. cluster autoscaler
C. manual scaling
D. Azure Container Instances

Correct Answer: B
Section: [none]
Explanation

Explanation/Reference:
Explanation:
To keep up with application demands in Azure Kubernetes Service (AKS), you may need to adjust the number of nodes that run your workloads. The cluster autoscaler component can watch for pods in your cluster that can't be scheduled because of resource constraints. When issues are detected, the number of nodes is increased to meet the application demand. Nodes are also regularly checked for a lack of running pods, with the number of nodes then decreased as needed. This ability to automatically scale up or down the number of nodes in your AKS cluster lets you run an efficient, cost-effective cluster.

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

QUESTION 6
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 7
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 8
You are designing an AI solution in Azure that will perform image classification.

You need to identify which processing platform will provide you with the ability to update the logic over time. The solution must have the lowest latency for inferencing without having to batch.

Which compute target should you identify?

A. graphics processing units (GPUs)
B. field-programmable gate arrays (FPGAs)
C. central processing units (CPUs)
D. application-specific integrated circuits (ASICs)

Correct Answer: B
Section: [none]
Explanation

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

Incorrect Answers:
D: ASICs are custom circuits, such as Google's TensorFlow Processor Units (TPU), provide the highest efficiency. They can't be reconfigured as your needs change.

References: https://docs.microsoft.com/en-us/azure/machine-learning/service/concept-accelerate-with-fpgas

QUESTION 9
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 10
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 11
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:
References:
https://docs.microsoft.com/en-us/azure/architecture/best-practices/data-partitioning

QUESTION 12
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 13
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:

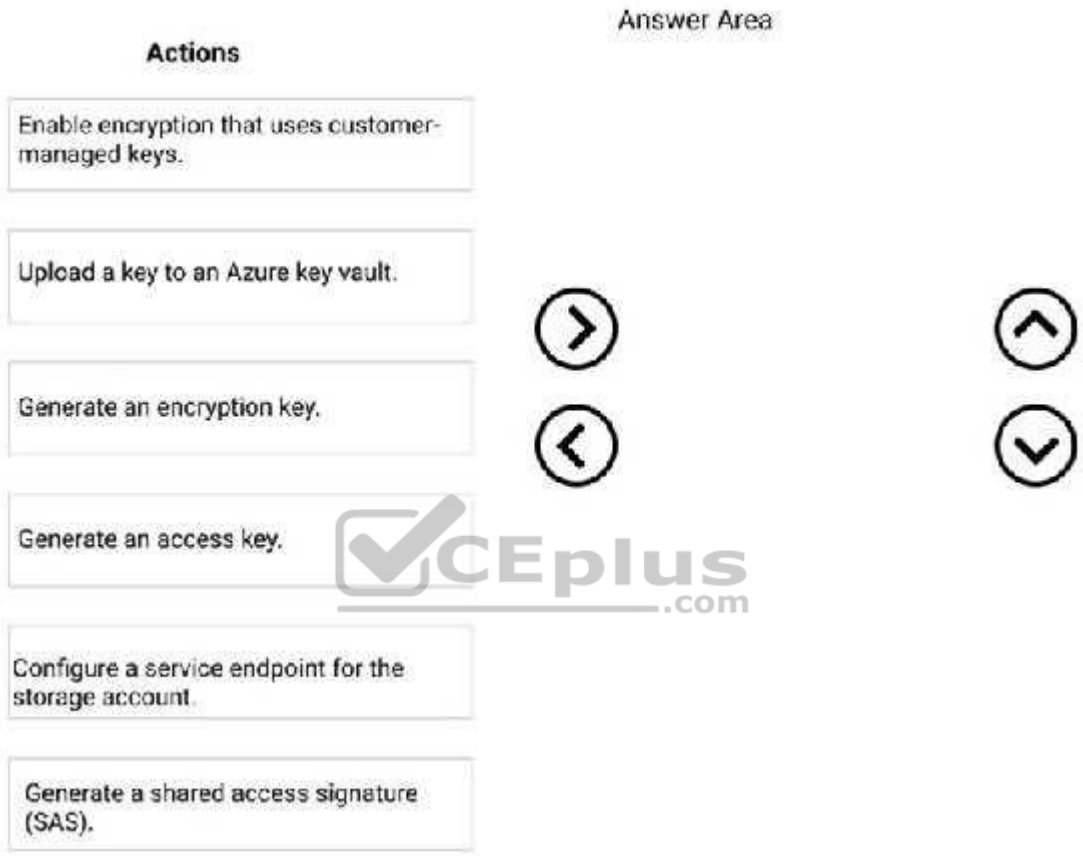

Correct Answer:

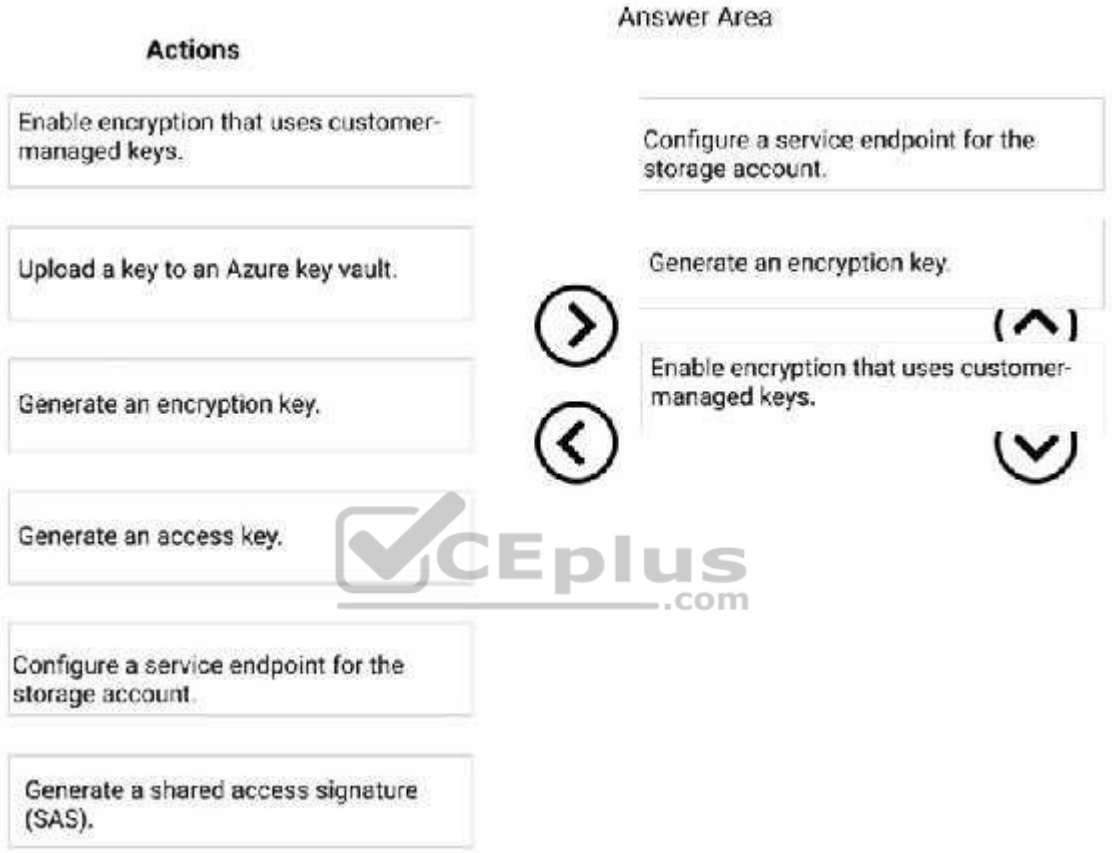

Section: [none]
Explanation

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 14**

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 15
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 16
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 17
You are developing a mobile application that will perform optical character recognition (OCR) from photos.

The application will annotate the photos by using metadata, store the photos in Azure Blob storage, and then score the photos by using an Azure Machine Learning model.

What should you use to process the data?

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

Correct Answer: A
Section: [none]
Explanation

Explanation/Reference:

QUESTION 18
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 19
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 Set 1

QUESTION 1
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 2
You are developing a Computer Vision application.

You plan to use a workflow that will load data from an on-premises database to Azure Blob storage, and then connect to an Azure Machine Learning service.

What should you use to orchestrate the workflow?

A. Azure Kubernetes Service (AKS)
B. Azure Pipelines
C. Azure Data Factory

D. Azure Container Instances

Correct Answer: C
Section: [none]
Explanation

Explanation/Reference:
Explanation:
With Azure Data Factory you can use workflows to orchestrate data integration and data transformation processes at scale.
Build data integration, and easily transform and integrate big data processing and machine learning with the visual interface.

References: https://azure.microsoft.com/en-us/services/data-factory/

QUESTION 3
Your company has a data team of Transact-SQL experts.

You plan to ingest data from multiple sources into Azure Event Hubs.

You need to recommend which technology the data team should use to move and query data from Event Hubs to Azure Storage. The solution must leverage the data team's existing skills.

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

A. Azure Notification Hubs
B. Azure Event Grid
C. Apache Kafka streams
D. Azure Stream Analytics

Correct Answer: B
Section: [none]
Explanation

Explanation/Reference:
Explanation:
Event Hubs Capture is the easiest way to automatically deliver streamed data in Event Hubs to an Azure Blob storage or Azure Data Lake store. You can subsequently process and deliver the data to any other storage destinations of your choice, such as SQL Data Warehouse or Cosmos DB.
You to capture data from your event hub into a SQL data warehouse by using an Azure function triggered by an event grid.

Example:

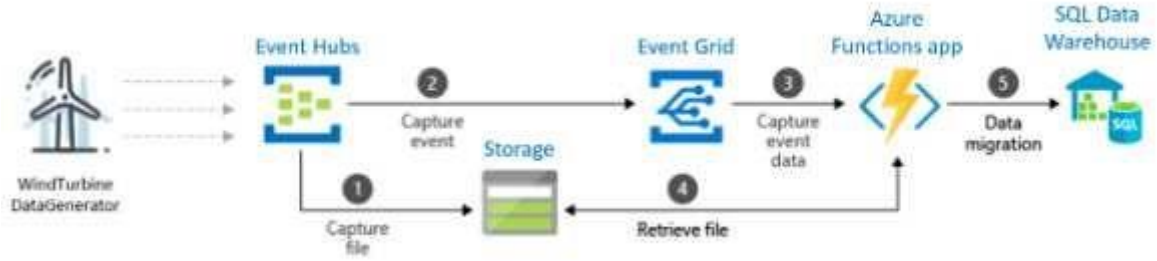

First, you create an event hub with the Capture feature enabled and set an Azure blob storage as the destination. Data generated by WindTurbineGenerator is streamed into the event hub and is automatically captured into Azure Storage as Avro files.
Next, you create an Azure Event Grid subscription with the Event Hubs namespace as its source and the Azure Function endpoint as its destination.
Whenever a new Avro file is delivered to the Azure Storage blob by the Event Hubs Capture feature, Event Grid notifies the Azure Function with the blob URI. The Function then migrates data from the blob to a SQL data warehouse.

References: https://docs.microsoft.com/en-us/azure/event-hubs/store-captured-data-data-warehouse

QUESTION 4
HOTSPOT

You are designing an Azure infrastructure to support an Azure Machine Learning solution that will have multiple phases. The solution must meet the following requirements:

- Securely query an on-premises database once a week to update product lists.
- Access the data without using a gateway. ▪
Orchestrate the separate phases.

What should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

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

Correct Answer:

Answer Area

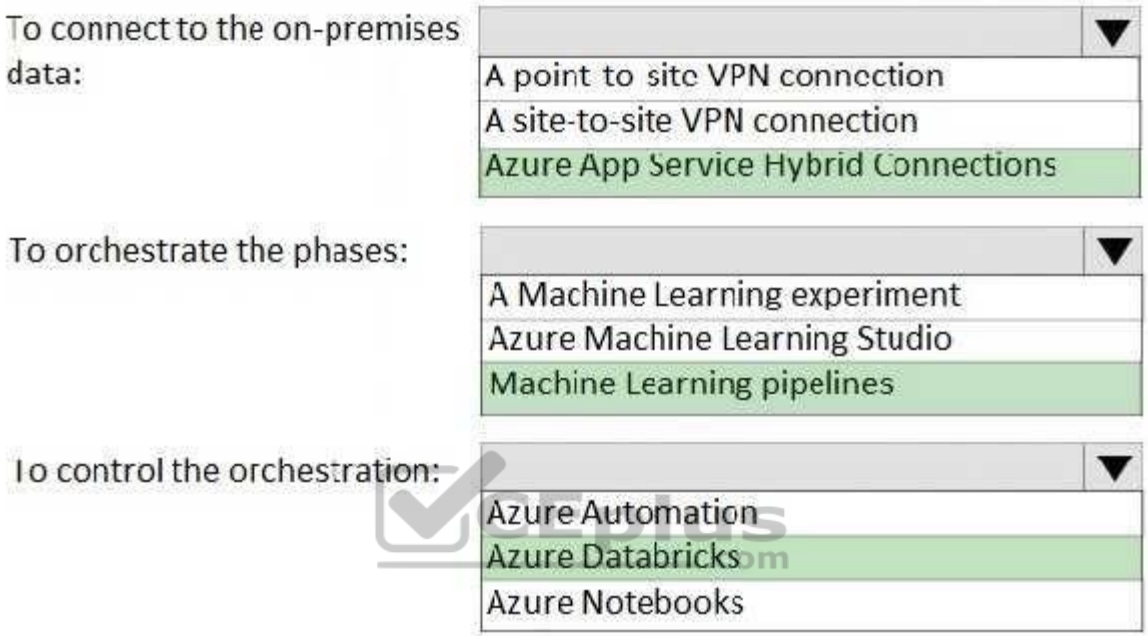

Section: [none]
Explanation

Explanation/Reference:
Explanation:

Box 1: Azure App Service Hybrid Connections
With Hybrid Connections, Azure websites and mobile services can access on-premises resources as if they were located on the same private network. Application admins thus have the flexibility to simply lift-and-shift specific most front-end tiers to Azure with minimal configuration changes, extending their enterprise apps for hybrid scenarios.

Incorrect Answer: The VPN connection solution both use gateways.

Box 2: Machine Learning pipelines

Typically when running machine learning algorithms, it involves a sequence of tasks including pre-processing, feature extraction, model fitting, and validation stages. For example, when classifying text documents might involve text segmentation and cleaning, extracting features, and training a classification model with cross-validation. Though there are many libraries we can use for each stage, connecting the dots is not as easy as it may look, especially with large-scale datasets. Most ML libraries are not designed for distributed computation or they do not provide native support for pipeline creation and tuning. Box 3: Azure Databricks

References:
https://azure.microsoft.com/is-is/blog/hybrid-connections-preview/

https://databricks.com/glossary/what-are-ml-pipelines

QUESTION 5
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 6
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 7
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 8
You are designing an AI workflow that will aggregate data stored in Azure as JSON documents.

You expect to store more than 2 TB of new data daily.

You need to choose the data storage service for the data. The solution must minimize costs.

Which data storage service should you choose?

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

Correct Answer: B
Section: [none]
Explanation

Explanation/Reference:
Explanation:
Generally, Data Lake will be a bit more expensive although they are in close range of each other. Blob storage has more options for pricing depending upon things like how frequently you need to access your data (cold vs hot storage). Data Lake is priced on volume, so it will go up as you reach certain tiers of volume.

References:

http://blog.pragmaticworks.com/azure-data-lake-vs-azure-blob-storage-in-data-warehousing

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 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 kubect1 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 10
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 11
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 12
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 13
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 14
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 15
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 16
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 17
You need to create a prototype of a bot to demonstrate a user performing a task. The demonstration will use the Bot Framework Emulator.

Which botbuilder CLI tool should you use to create the prototype?

A. Chatdown
B. QnAMaker
C. Dispatch
D. LuDown

Correct Answer: A
Section: [none]
Explanation
Explanation/Reference:
Explanation:
Use Chatdown to produce prototype mock conversations in markdown and convert the markdown to transcripts you can load and view in the new V4 Bot Framework Emulator.

Incorrect Answers:
B: QnA Maker is a cloud-based API service that lets you create a conversational question-and-answer layer over your existing data. Use it to build a knowledge base by extracting questions and answers from your semi-structured content, including FAQs, manuals, and documents. Answer users' questions with the best answers from the QnAs in your knowledge base—automatically. Your knowledge base gets smarter, too, as it continually learns from user behavior.

C: Dispatch lets you build language models that allow you to dispatch between disparate components (such as QnA, LUIS and custom code).

D: LuDown build LUIS language understanding models using markdown files

References:
https://github.com/microsoft/botframework/blob/master/README.md

QUESTION 18
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 19**
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 processsed 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 20
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/

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.

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 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 2
You design an AI solution that uses an Azure Stream Analytics job to process data from an Azure IoT hub. The IoT hub receives time series data from thousands of IoT devices at a factory.

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

You need to choose an output type for the job.

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

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

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 3
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 4
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 5
You create an Azure Cognitive Services resource.

You develop needs to be able to retrieve the keys used by the resource. The solution must use the principle of least privilege.

What is the best role to assign to the developer? More than one answer choice may achieve the goal.

A. Security Manager
B. Security Reader
C. Cognitive Services Contributor
D. Cognitive Services User

Correct Answer: D
Section: [none]
Explanation

Explanation/Reference:
Explanation:
The Cognitive Services User lets you read and list keys of Cognitive Services.

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

QUESTION 6
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 7
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 8
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

QUESTION 9
DRAG DROP

You need to design the workflow for an Azure Machine Learning solution. The solution must meet the following requirements:

▪ Retrieve data from file shares, Microsoft SQL Server databases, and Oracle databases that in an on-premises network.▪
Use an Apache Spark job to process data stored in an Azure SQL Data Warehouse database.

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

NOTE: Each correct selection is worth one point.

Select and Place:

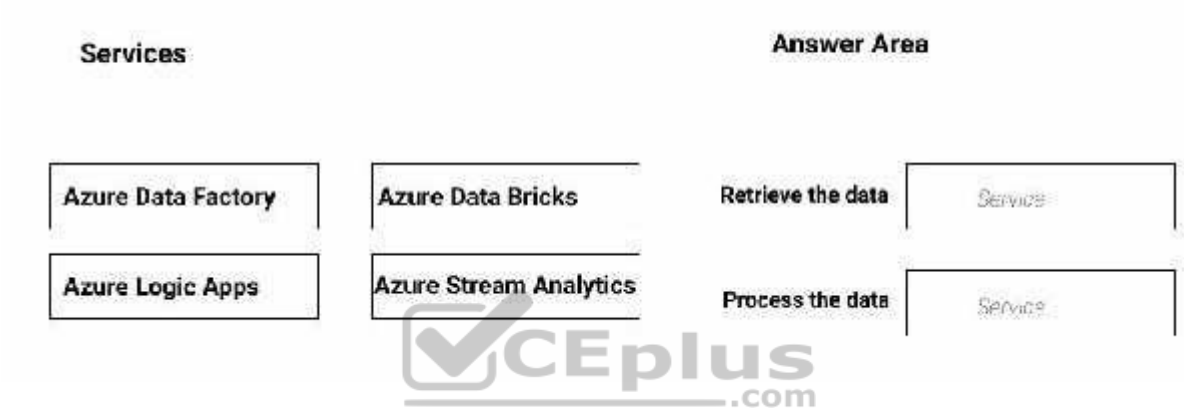

Correct Answer:

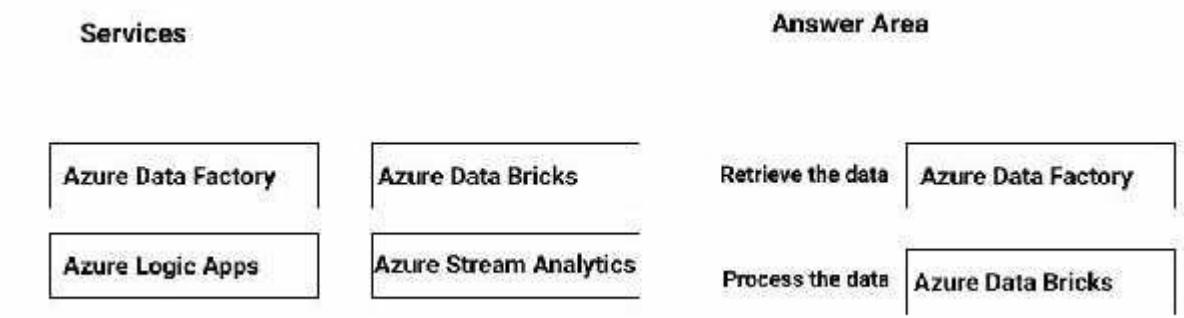

Section: [none]
Explanation

Explanation/Reference:
References:
https://docs.microsoft.com/en-us/azure/machine-learning/studio/use-data-from-an-on-premises-sql-server https://docs.microsoft.com/en-in/azure/azure-databricks/what-is-azure-databricks

QUESTION 10
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

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 11
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 12
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 **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.

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 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 2
You create an Azure Machine Learning Studio experiment.

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

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

What should you use?

A. a Batch Execution Service (BES) and an 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 3
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.

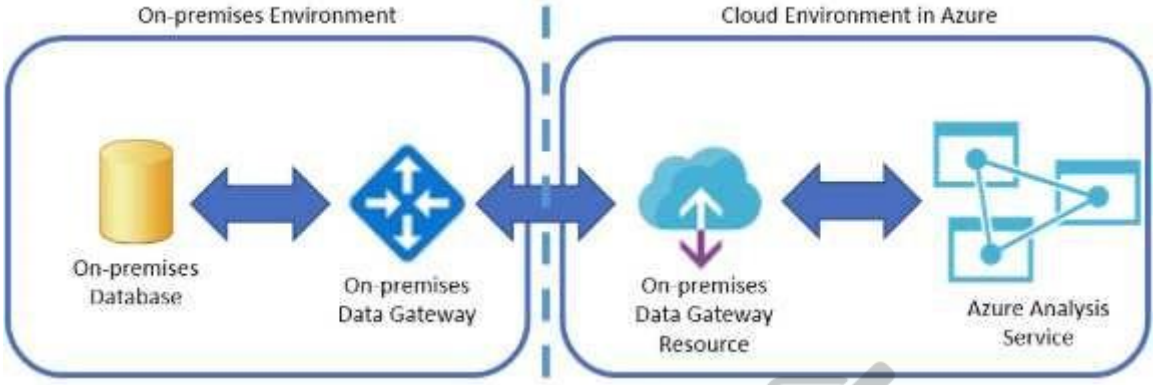

References: https://biinsight.com/on-premises-data-gateway-for-azure-analysis-services/

QUESTION 4
DRAG DROP

You use an Azure key vault to store credentials for several Azure Machine Learning applications.

You need to configure the key vault to meet the following requirements:

- Ensure that the IT security team can add new passwords and periodically change the passwords.
- Ensure that the applications can securely retrieve the passwords for the applications.
- Use the principle of least privilege.

Which permissions should you grant? To answer, drag the appropriate permissions to the correct targets. Each permission may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

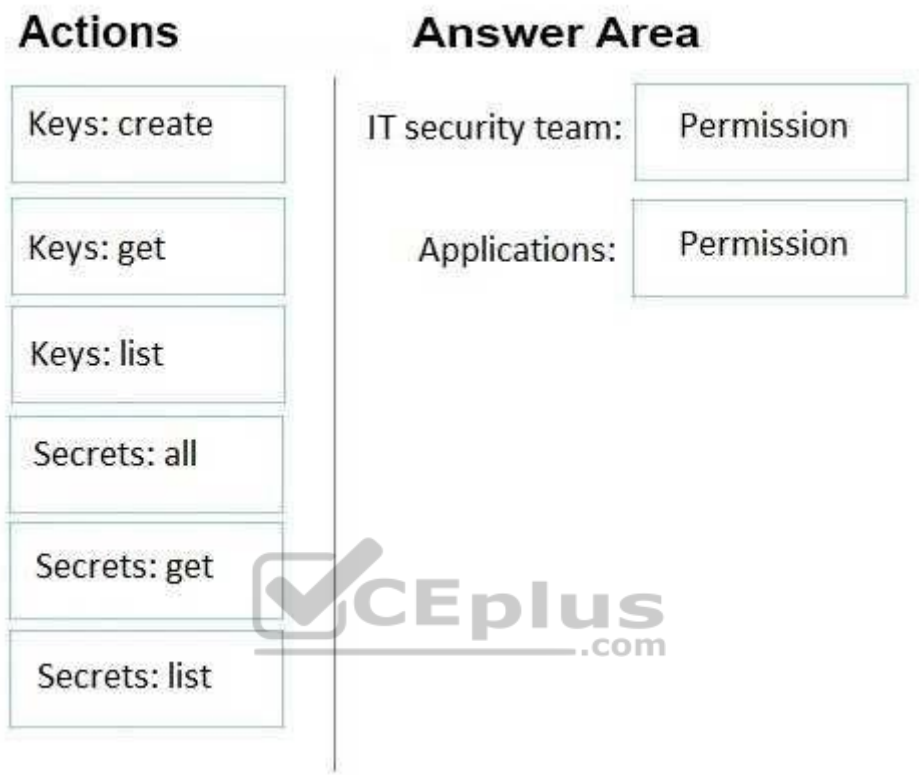

Correct Answer:

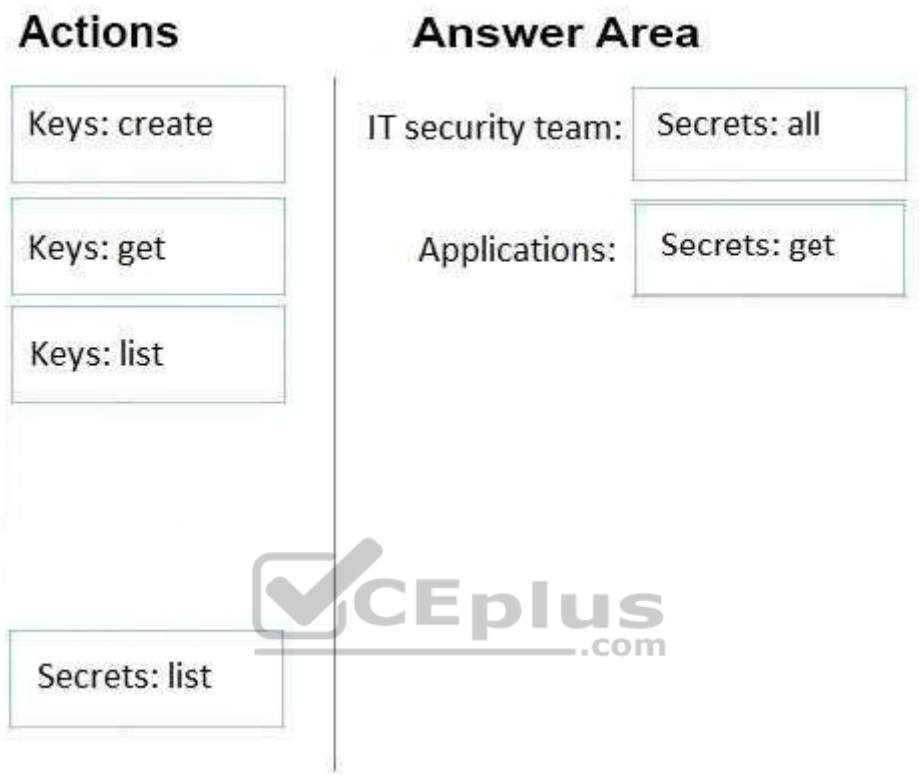

Section: [none]
Explanation

Explanation/Reference:
Explanation:

Incorrect Answers:
Not Keys as they are used for encryption only.

References: https://docs.microsoft.com/en-us/azure/key-vault/key-vault-secure-your-key-vault

QUESTION 5

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 6
You have Azure IoT Edge devices that generate measurement data from temperature sensors. The data changes very slowly.

You need to analyze the data in a temporal two-minute window. If the temperature rises five degrees above a limit, an alert must be raised. The solution must minimize the development of custom code.

What should you use?

A. A Machine Learning model as a web service
B. an Azure Machine Learning model as an IoT Edge module
C. Azure Stream Analytics as an IoT Edge module

D. Azure Functions as an IoT Edge module

Correct Answer: C
Section: [none]
Explanation

Explanation/Reference:
References: https://docs.microsoft.com/en-us/azure/iot-edge/tutorial-deploy-stream-analytics

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 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:

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 create several API models in Azure Machine Learning Studio.

You deploy the models to a production environment.

You need to monitor the compute performance of the models.

Solution: You create environment files.

Does this meet the goal?

A. Yes
B. No

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 9
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 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 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 11
You plan to deploy a global healthcare app named App1 to Azure.

App1 will use Azure Cognitive Services APIs. Users in Germany, Canada, and the United States will connect to App1.

You need to recommend an app deployment solution to ensure that all the personal data of the users remain in their country or origin only.

Which three Azure services should you recommend deploying to each Azure region? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

A. Azure Key Vault
B. Azure Traffic Manager
C. Azure Kubernetes Service (AKS)
D. App1
E. the Cognitive Services resources
F. an Azure Storage resource

Correct Answer: ADF
Section: [none]
Explanation

Explanation/Reference:
References:
https://github.com/microsoft/computerscience/blob/master/Labs/Azure%20Services/Azure%20Storage/Azure%20Storage%20and%20Cognitive%20Services%20(MVC).md

QUESTION 12
DRAG DROP

You are designing an AI solution that will use IoT devices to gather data from conference attendees and then analyze the data. The IoT device will connect to an Azure IoT hub.

You need to ensure that data contains no personally identifiable information before it is sent to the IoT hub.

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

Select and Place:

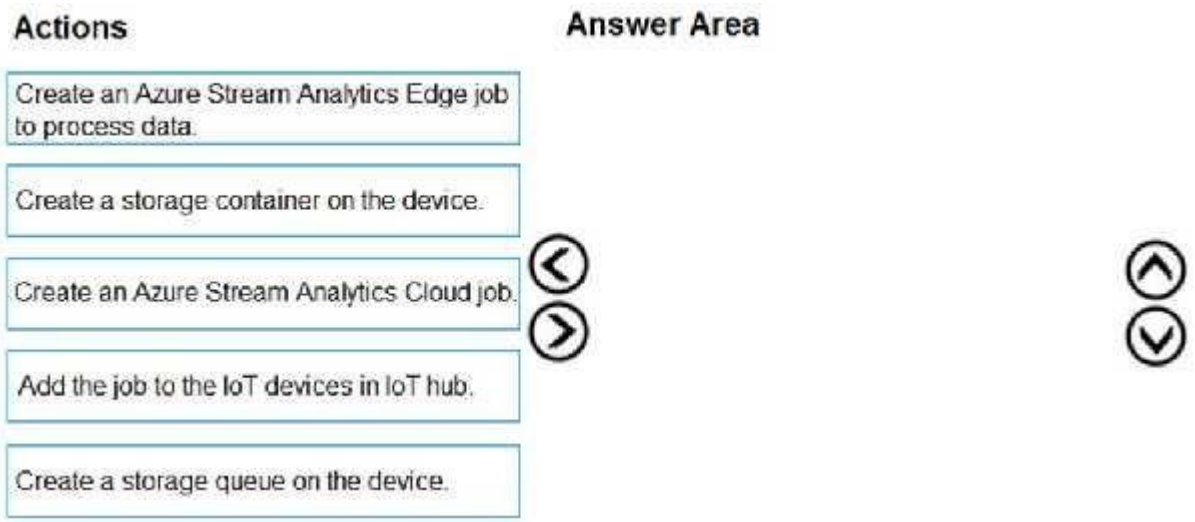

Correct Answer:

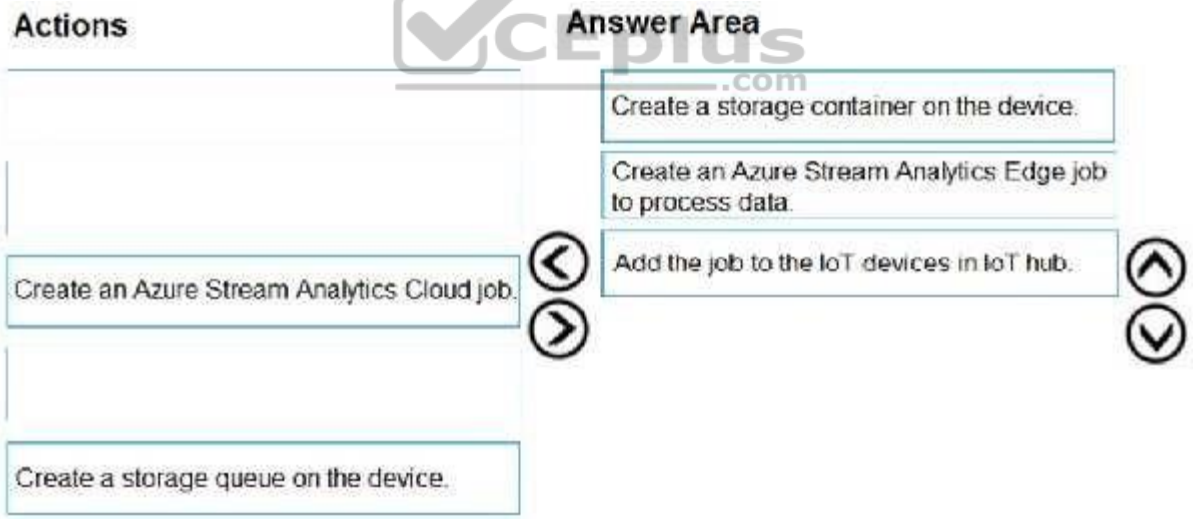

Section: [none]
Explanation

Explanation/Reference:
Note:
ASA Edge jobs run in containers deployed to Azure IoT Edge devices. They are composed of two parts:

1. A cloud part that is responsible for job definition: users define inputs, output, query, and other settings (out of order events, etc.) in the cloud.
2. A module running on your IoT devices. It contains the ASA engine and receives the job definition from the cloud.

References: https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-edge

QUESTION 13
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 14
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