QUESTION 137
Case Study 6 - Litware, Inc

Overview
Litware, Inc. owns and operates 300 convenience stores across the US. The company sells a variety of packaged foods and drinks, as well as a variety of prepared foods, such as sandwiches and pizzas.

Litware has a loyalty club whereby members can get daily discounts on specific items by providing their membership number at checkout.

Litware employs business analysts who prefer to analyze data by using Microsoft Power BI, and data scientists who prefer analyzing data in Azure Databricks notebooks.

Requirements. Business Goals
Litware wants to create a new analytics environment in Azure to meet the following requirements:

- See inventory levels across the stores. Data must be updated as close to real time as possible.
- Execute ad hoc analytical queries on historical data to identify whether the loyalty club discounts increase sales of the discounted products.
- Every four hours, notify store employees about how many prepared food items to produce based on historical demand from the sales data.

Requirements. Technical Requirements
Litware identifies the following technical requirements:

- Minimize the number of different Azure services needed to achieve the business goals
- Use platform as a service (PaaS) offerings whenever possible and avoid having to provision virtual machines that must be managed by Litware.
- Ensure that the analytical data store is accessible only to the company's on-premises network and Azure services.
- Use Azure Active Directory (Azure AD) authentication whenever possible.
- Use the principle of least privilege when designing security.
- Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
- Limit the business analysts' access to customer contact information, such as phone numbers, because this type of data is not analytically relevant.
- Ensure that you can quickly restore a copy of the analytical data store within one hour in the event of corruption or accidental deletion.

Requirements. Planned Environment
Litware plans to implement the following environment:

- The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
- Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Daily inventory data comes from a Microsoft SQL server located on a private network.
- Litware currently has 5 TB of historical sales data and 100 GB of customer data. The company expects approximately 100 GB of new data per month for the next year.
- Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
- Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

Inventory levels must be calculated by subtracting the current day's sales from the previous day's final inventory.

Which two options provide Litware with the ability to quickly calculate the current inventory levels by store and product? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Consume the output of the event hub by using Azure Stream Analytics and aggregate the data by store and product. Output the resulting data directly to Azure Synapse Analytics. Use Transact-SQL to calculate the inventory levels.
B. Output Event Hubs Avro files to Azure Blob storage. Use Transact-SQL to calculate the inventory levels by using PolyBase in Azure Synapse Analytics.
C. Consume the output of the event hub by using Databricks. Use Databricks to calculate the inventory levels and output the data to Azure Synapse Analytics.
D. Consume the output of the event hub by using Azure Stream Analytics and aggregate the data by store and product. Output the resulting data into Databricks. Calculate the inventory levels in Databricks and output the data to Azure Blob storage.
E. Output Event Hubs Avro files to Azure Blob storage. Trigger an Azure Data Factory copy activity to run every 10 minutes to load the data into Azure Synapse Analytics. Use Transact-SQL to aggregate the data by store and product.

Correct Answer: AE
Section: (none)
Explanation

Explanation/Reference:
Explanation:
A: Azure Stream Analytics is a fully managed service providing low-latency, highly available, scalable complex event processing over streaming data in the cloud. You can use your Azure SQL Data Warehouse database as an output sink for your Stream Analytics jobs.
E: Event Hubs Capture is the easiest way to get data into Azure. Using Azure Data Lake, Azure Data Factory, and Azure HDInsight, you can perform batch processing and other analytics using familiar tools and platforms of your choosing, at any scale you need.
Note: Event Hubs Capture creates files in Avro format.
Captured data is written in Apache Avro format: a compact, fast, binary format that provides rich data structures with inline schema. This format is widely used in the Hadoop ecosystem, Stream Analytics, and Azure Data Factory.
Scenario: The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
Reference:
https://docs.microsoft.com/bs-latn-ba/azure/sql-data-warehouse/sql-data-warehouse-integrate-azure-stream-analytics
https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-capture-overview

QUESTION 138
Case Study 6 - Litware, Inc

Overview
Litware, Inc. owns and operates 300 convenience stores across the US. The company sells a variety of packaged foods and drinks, as well as a variety of prepared foods, such as sandwiches and pizzas.

Litware has a loyalty club whereby members can get daily discounts on specific items by providing their membership number at checkout.

Litware employs business analysts who prefer to analyze data by using Microsoft Power BI, and data scientists who prefer analyzing data in Azure Databricks notebooks.

Requirements. Business Goals
Litware wants to create a new analytics environment in Azure to meet the following requirements:

- See inventory levels across the stores. Data must be updated as close to real time as possible.
- Execute ad hoc analytical queries on historical data to identify whether the loyalty club discounts increase sales of the discounted products. Every
- four hours, notify store employees about how many prepared food items to produce based on historical demand from the sales data.

Requirements. Technical Requirements
Litware identifies the following technical requirements:

- Minimize the number of different Azure services needed to achieve the business goals
- Use platform as a service (PaaS) offerings whenever possible and avoid having to provision virtual machines that must be managed by Litware. Ensure
- that the analytical data store is accessible only to the company's on-premises network and Azure services.
- Use Azure Active Directory (Azure AD) authentication whenever possible.
- Use the principle of least privilege when designing security.
- Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
- Limit the business analysts' access to customer contact information, such as phone numbers, because this type of data is not analytically relevant.
- Ensure that you can quickly restore a copy of the analytical data store within one hour in the event of corruption or accidental deletion.

Requirements. Planned Environment
Litware plans to implement the following environment:

- The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
- Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Daily inventory data comes from a Microsoft SQL server located on a private network.
- Litware currently has 5 TB of historical sales data and 100 GB of customer data. The company expects approximately 100 GB of new data per month for the next year.
- Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
- Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

What should you do to improve high availability of the real-time data processing solution?

A. Deploy identical Azure Stream Analytics jobs to paired regions in Azure.
B. Deploy a High Concurrency Databricks cluster.
C. Deploy an Azure Stream Analytics job and use an Azure Automation runbook to check the status of the job and to start the job if it stops.
D. Set Data Lake Storage to use geo-redundant storage (GRS).

Correct Answer: A
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Guarantee Stream Analytics job reliability during service updates Part of being a fully managed service is the capability to introduce new service functionality and improvements at a rapid pace. As a result, Stream Analytics can have a service update deploy on a weekly (or more frequent) basis. No matter how much testing is done there is still a risk that an existing, running job may break due to the introduction of a bug. If you are running mission critical jobs, these risks need to be avoided. You can reduce this risk by following Azure's paired region model.
Scenario: The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure
Reference:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-job-reliability

QUESTION 139
Case Study 6 - Litware, Inc

Overview
Litware, Inc. owns and operates 300 convenience stores across the US. The company sells a variety of packaged foods and drinks, as well as a variety of prepared foods, such as sandwiches and pizzas.

Litware has a loyalty club whereby members can get daily discounts on specific items by providing their membership number at checkout.

Litware employs business analysts who prefer to analyze data by using Microsoft Power BI, and data scientists who prefer analyzing data in Azure Databricks notebooks.

Requirements. Business Goals
Litware wants to create a new analytics environment in Azure to meet the following requirements:

- See inventory levels across the stores. Data must be updated as close to real time as possible.
- Execute ad hoc analytical queries on historical data to identify whether the loyalty club discounts increase sales of the discounted products. Every
- four hours, notify store employees about how many prepared food items to produce based on historical demand from the sales data.

Requirements. Technical Requirements
Litware identifies the following technical requirements:

- Minimize the number of different Azure services needed to achieve the business goals
- Use platform as a service (PaaS) offerings whenever possible and avoid having to provision virtual machines that must be managed by Litware. Ensure
- that the analytical data store is accessible only to the company's on-premises network and Azure services.
- Use Azure Active Directory (Azure AD) authentication whenever possible.
- Use the principle of least privilege when designing security.
- Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
- Limit the business analysts' access to customer contact information, such as phone numbers, because this type of data is not analytically relevant.
- Ensure that you can quickly restore a copy of the analytical data store within one hour in the event of corruption or accidental deletion.

Requirements. Planned Environment
Litware plans to implement the following environment:

- The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
- Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Daily inventory data comes from a Microsoft SQL server located on a private network.
- Litware currently has 5 TB of historical sales data and 100 GB of customer data. The company expects approximately 100 GB of new data per month for the next year.
- Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
- Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

Hotspot Question

Which Azure Data Factory components should you recommend using together to import the customer data from Salesforce to Data Lake Storage? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

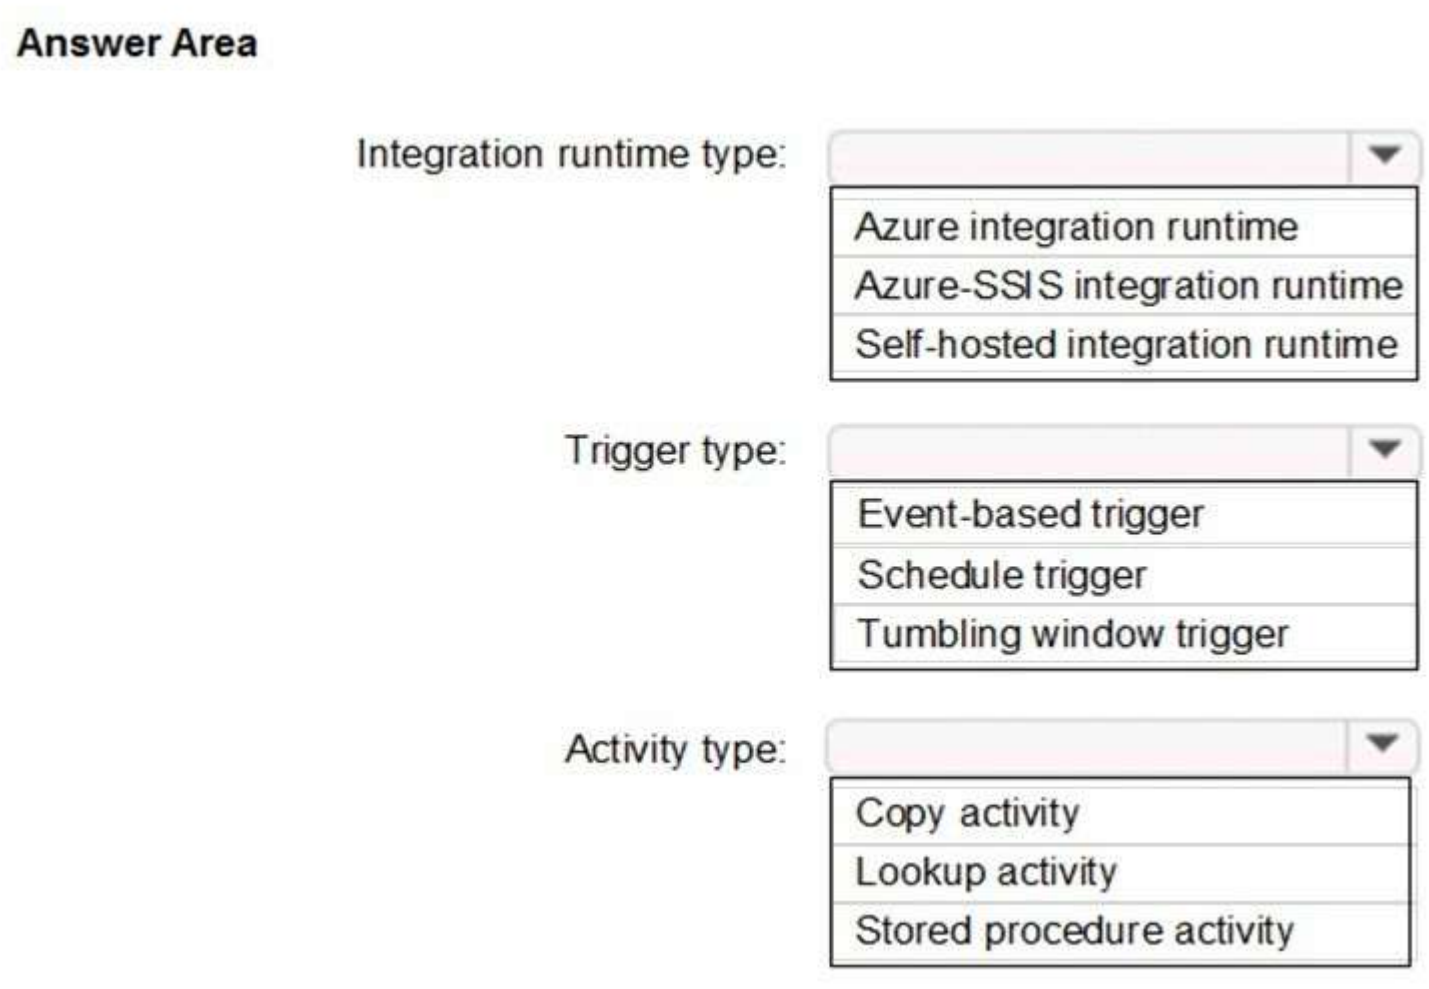

Correct Answer:

Answer Area

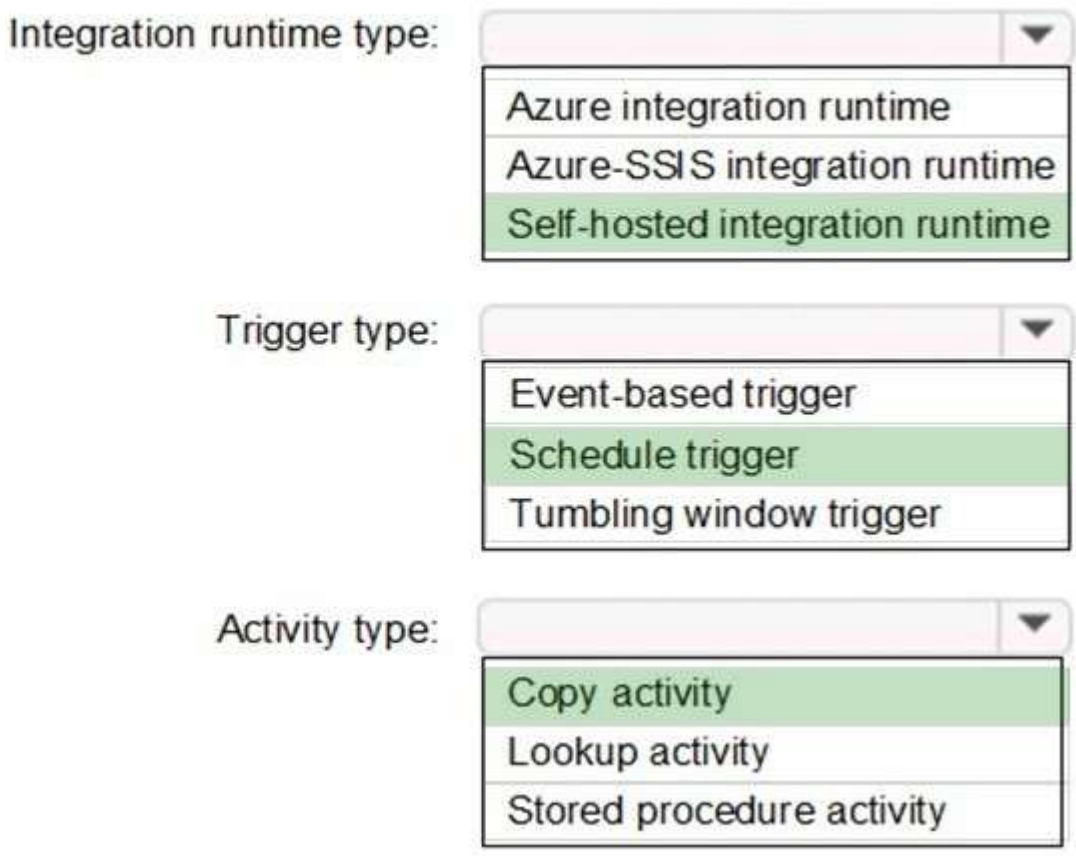

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Box 1: Self-hosted integration runtime
A self-hosted IR is capable of nunning copy activity between a cloud data stores and a data store in private network.
Box 2: Schedule trigger
Schedule every 8 hours
Box 3: Copy activity
Scenario:
Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.

QUESTION 140
Case Study 6 - Litware, Inc

Overview
Litware, Inc. owns and operates 300 convenience stores across the US. The company sells a variety of packaged foods and drinks, as well as a variety of prepared foods, such as sandwiches and pizzas.

Litware has a loyalty club whereby members can get daily discounts on specific items by providing their membership number at checkout.

Litware employs business analysts who prefer to analyze data by using Microsoft Power BI, and data scientists who prefer analyzing data in Azure Databricks notebooks.

Requirements. Business Goals
Litware wants to create a new analytics environment in Azure to meet the following requirements:

- See inventory levels across the stores. Data must be updated as close to real time as possible.
- Execute ad hoc analytical queries on historical data to identify whether the loyalty club discounts increase sales of the discounted products. Every
- four hours, notify store employees about how many prepared food items to produce based on historical demand from the sales data.

Requirements. Technical Requirements
Litware identifies the following technical requirements:

- Minimize the number of different Azure services needed to achieve the business goals
- Use platform as a service (PaaS) offerings whenever possible and avoid having to provision virtual machines that must be managed by Litware. Ensure
- that the analytical data store is accessible only to the company's on-premises network and Azure services.
- Use Azure Active Directory (Azure AD) authentication whenever possible.
- Use the principle of least privilege when designing security.
- Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
- Limit the business analysts' access to customer contact information, such as phone numbers, because this type of data is not analytically relevant.
- Ensure that you can quickly restore a copy of the analytical data store within one hour in the event of corruption or accidental deletion.

Requirements. Planned Environment
Litware plans to implement the following environment:

- The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty

number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
- Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Daily inventory data comes from a Microsoft SQL server located on a private network.
- Litware currently has 5 TB of historical sales data and 100 GB of customer data. The company expects approximately 100 GB of new data per month for the next year.
- Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
- Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

Hotspot Question

Which Azure Data Factory components should you recommend using together to import the daily inventory data from SQL to Data Lake Storage? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

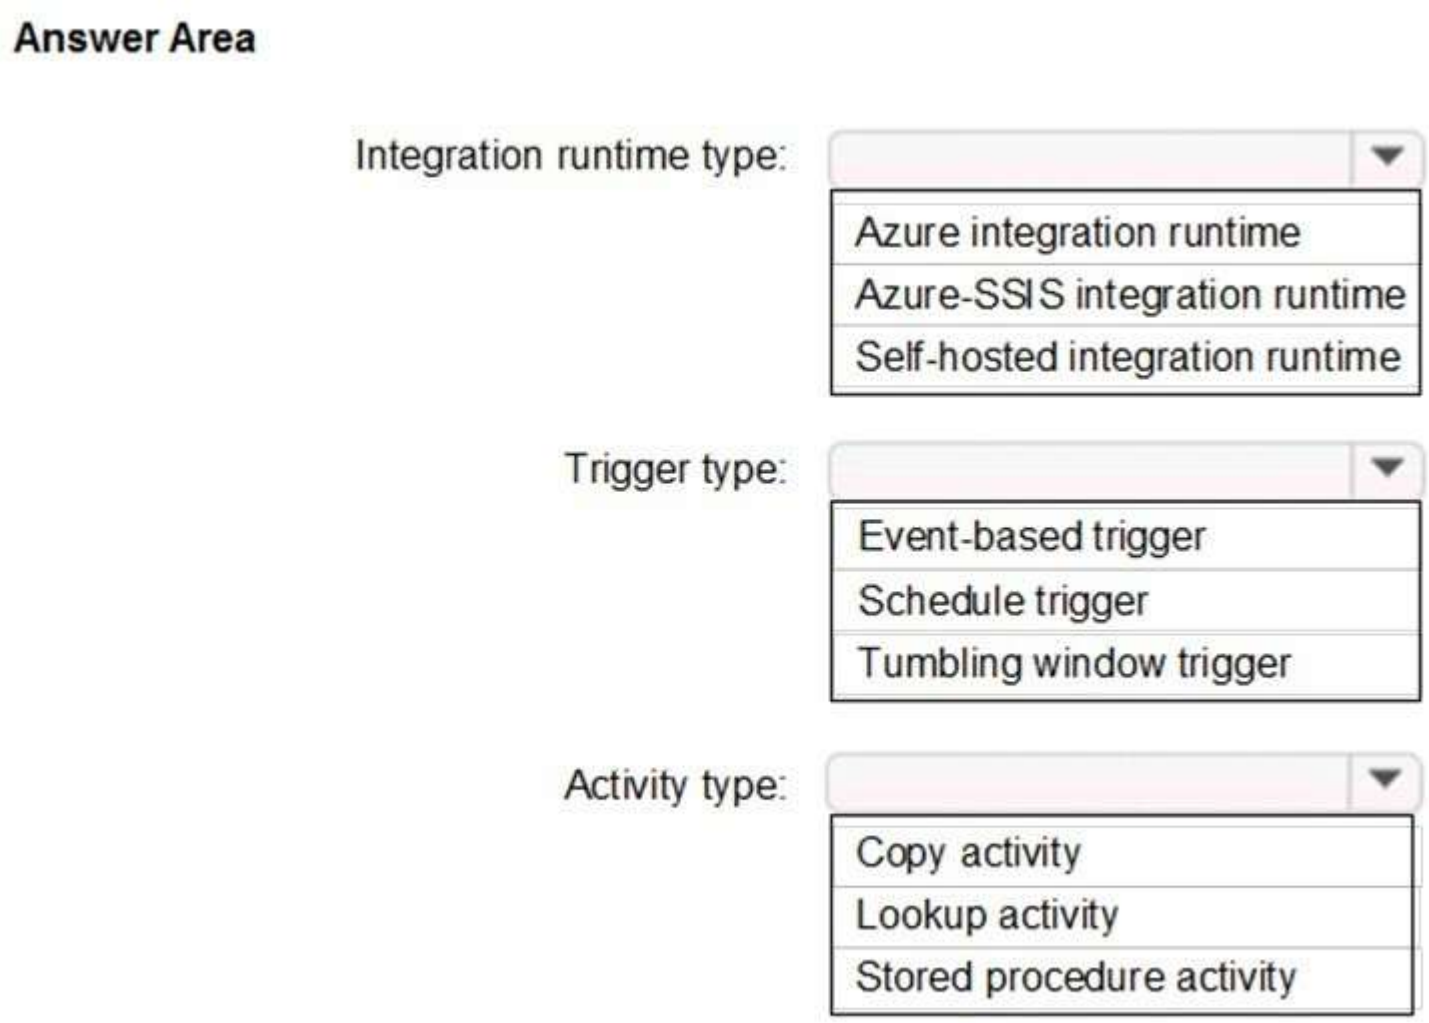

Correct Answer:

Answer Area

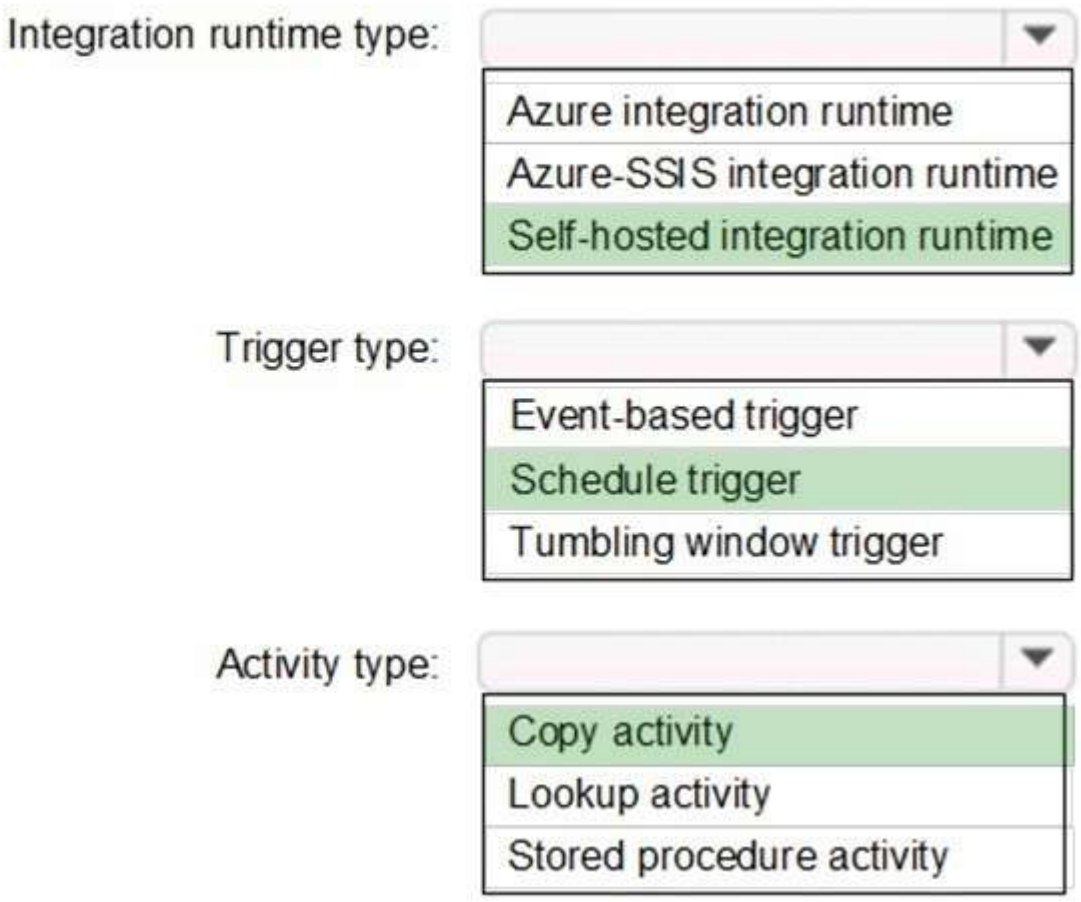

Section: (none)

Explanation/Reference:

Explanation:

Box 1: Self-hosted integration runtime
A self-hosted IR is capable of nunning copy activity between a cloud data stores and a data store in private network.
Scenario: Daily inventory data comes from a Microsoft SQL server located on a private network.
Box 2: Schedule trigger
Daily schedule
Box 3: Copy activity
Scenario:
Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.

QUESTION 141
Case Study 6 - Litware, Inc

Overview
Litware, Inc. owns and operates 300 convenience stores across the US. The company sells a variety of packaged foods and drinks, as well as a variety of prepared foods, such as sandwiches and pizzas.

Litware has a loyalty club whereby members can get daily discounts on specific items by providing their membership number at checkout.

Litware employs business analysts who prefer to analyze data by using Microsoft Power BI, and data scientists who prefer analyzing data in Azure Databricks notebooks.

Requirements. Business Goals
Litware wants to create a new analytics environment in Azure to meet the following requirements:

- See inventory levels across the stores. Data must be updated as close to real time as possible.
- Execute ad hoc analytical queries on historical data to identify whether the loyalty club discounts increase sales of the discounted products. Every
- four hours, notify store employees about how many prepared food items to produce based on historical demand from the sales data.

Requirements. Technical Requirements
Litware identifies the following technical requirements:

- Minimize the number of different Azure services needed to achieve the business goals
- Use platform as a service (PaaS) offerings whenever possible and avoid having to provision virtual machines that must be managed by Litware. Ensure
- that the analytical data store is accessible only to the company's on-premises network and Azure services.
- Use Azure Active Directory (Azure AD) authentication whenever possible.
- Use the principle of least privilege when designing security.
- Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
- Limit the business analysts' access to customer contact information, such as phone numbers, because this type of data is not analytically relevant.
- Ensure that you can quickly restore a copy of the analytical data store within one hour in the event of corruption or accidental deletion.

Requirements. Planned Environment
Litware plans to implement the following environment:

- The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
- Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.

- Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Daily inventory data comes from a Microsoft SQL server located on a private network.
- Litware currently has 5 TB of historical sales data and 100 GB of customer data. The company expects approximately 100 GB of new data per month for the next year.
- Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
- Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

Hotspot Question

Which Azure service and feature should you recommend using to manage the transient data for Data Lake Storage? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

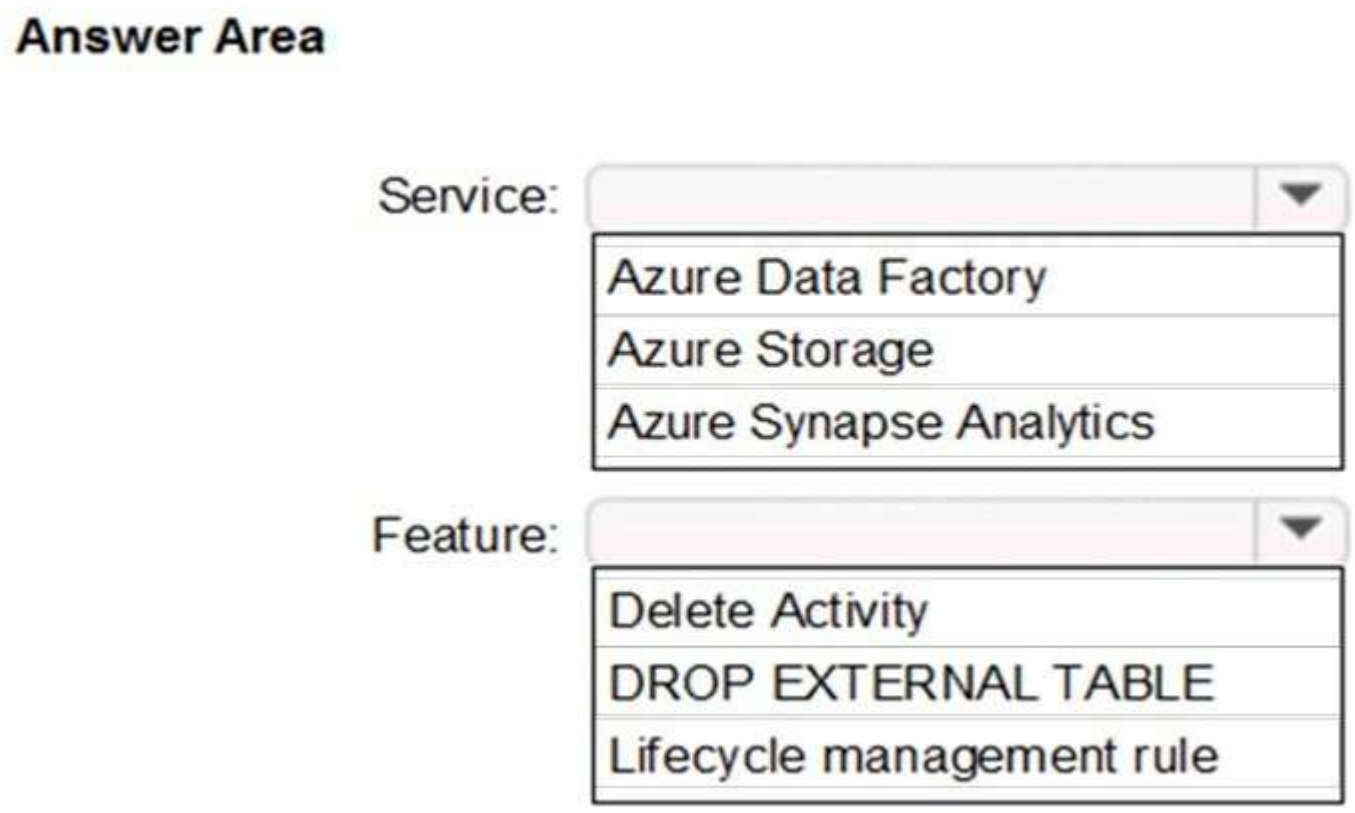

Correct Answer:

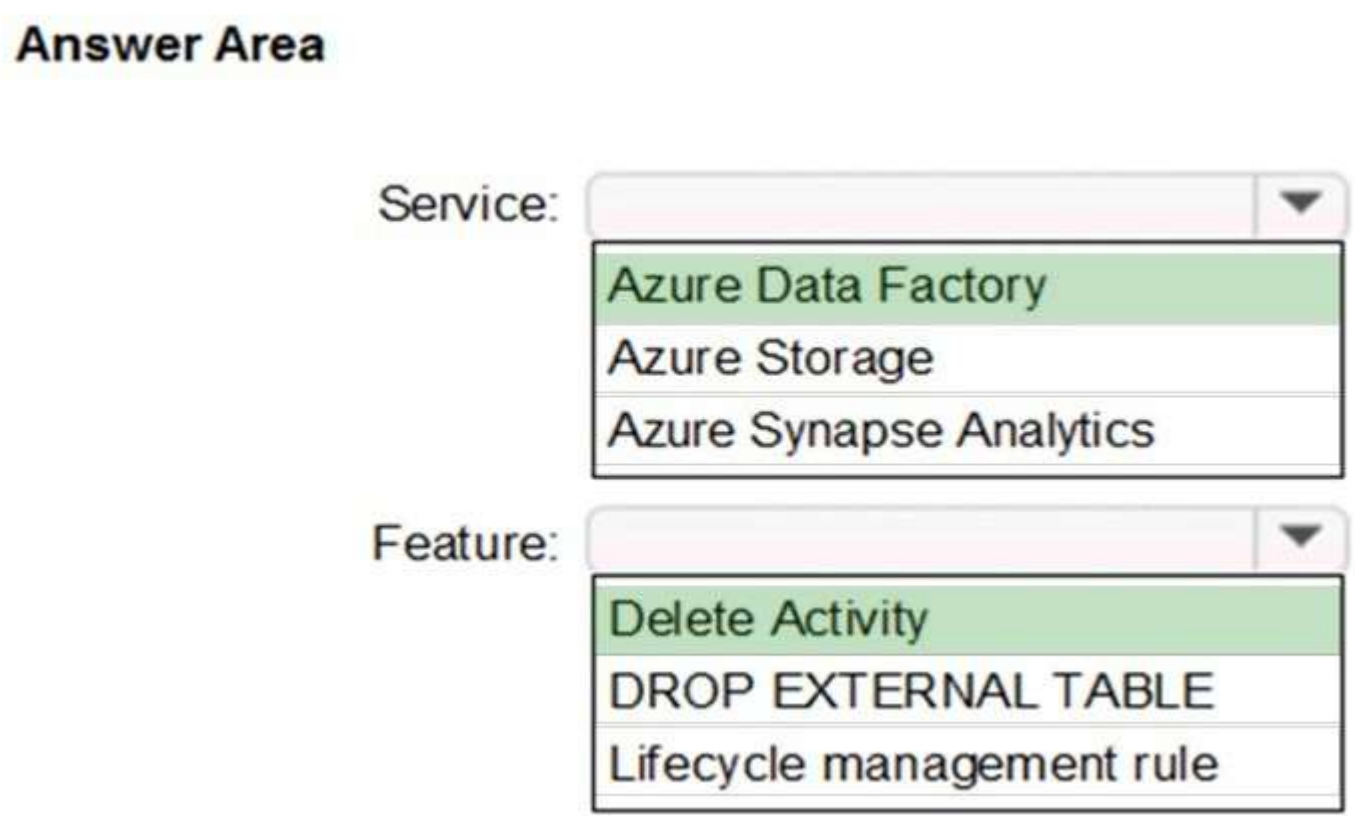

Section: (none)
Explanation

Explanation/Reference:
Explanation:
Scenario: Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
Service: Azure Data Factory
Clean up files by built-in delete activity in Azure Data Factory (ADF). ADF built-in delete activity, which can be part of your ETL workflow to deletes undesired files without writing code. You can use ADF to delete folder or files from Azure Blob Storage, Azure Data Lake Storage Gen1, Azure Data Lake Storage Gen2, File System, FTP Server, sFTP Server, and Amazon S3.
You can delete expired files only rather than deleting all the files in one folder. For example, you may want to only delete the files which were last modified more

than 13 days ago.
Feature: Delete Activity
Reference:
https://azure.microsoft.com/sv-se/blog/clean-up-files-by-built-in-delete-activity-in-azure-data-factory/

QUESTION 142
Case Study 6 - Litware, Inc

Overview
Litware, Inc. owns and operates 300 convenience stores across the US. The company sells a variety of packaged foods and drinks, as well as a variety of prepared foods, such as sandwiches and pizzas.

Litware has a loyalty club whereby members can get daily discounts on specific items by providing their membership number at checkout.

Litware employs business analysts who prefer to analyze data by using Microsoft Power BI, and data scientists who prefer analyzing data in Azure Databricks notebooks.

Requirements. Business Goals
Litware wants to create a new analytics environment in Azure to meet the following requirements:

- See inventory levels across the stores. Data must be updated as close to real time as possible.
- Execute ad hoc analytical queries on historical data to identify whether the loyalty club discounts increase sales of the discounted products. Every
- four hours, notify store employees about how many prepared food items to produce based on historical demand from the sales data.

Requirements. Technical Requirements
Litware identifies the following technical requirements:

- Minimize the number of different Azure services needed to achieve the business goals
- Use platform as a service (PaaS) offerings whenever possible and avoid having to provision virtual machines that must be managed by Litware. Ensure
- that the analytical data store is accessible only to the company's on-premises network and Azure services.
- Use Azure Active Directory (Azure AD) authentication whenever possible.
- Use the principle of least privilege when designing security.
- Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.
- Limit the business analysts' access to customer contact information, such as phone numbers, because this type of data is not analytically relevant.
- Ensure that you can quickly restore a copy of the analytical data store within one hour in the event of corruption or accidental deletion.

Requirements. Planned Environment
Litware plans to implement the following environment:

- The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure.
- Customer data, including name, contact information, and loyalty number, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Product data, including product ID, name, and category, comes from Salesforce and can be imported into Azure once every eight hours. Row modified dates are not trusted in the source table.
- Daily inventory data comes from a Microsoft SQL server located on a private network.
- Litware currently has 5 TB of historical sales data and 100 GB of customer data. The company expects approximately 100 GB of new data per month for the next year.
- Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
- Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

Drag and Drop Question

Which three actions should you perform in sequence to allow FoodPrep access to the analytical data store? 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

- Create a user-defined database role that grants access to execute the appropriate stored procedures.
- Register the FoodPrep application in Azure AD.
- Grant SELECT permissions on all the user tables to the users of FoodPrep.
- Create a contained user in the database.
- Map the login to the user in the database.
- Create a login for the service principal on the Azure SQL server.

Answer Area

Correct Answer:

Actions

- Map the login to the user in the database.
- Create a contained user in the database.
- Grant SELECT permissions on all the user tables to the users of FoodPrep.

Answer Area

- Register the FoodPrep application in Azure AD.
- Create a login for the service principal on the Azure SQL server.
- Create a user-defined database role that grants access to execute the appropriate stored procedures.

Section: (none) Explanation

Explanation/Reference:
Explanation:
Scenario: Litware will build a custom application named FoodPrep to provide store employees with the calculation results of how many prepared food items to produce every four hours.
Step 1: Register the FoodPrep application in Azure AD
You create your Azure AD application and service principal.
Step 2: Create a login for the service principal on the Azure SQL Server
Step 3: Create a user-defined database role that grant access.. To access resources in your subscription, you must assign the application to a role. You can then assign the required permissions to the service principal.
Reference:
https://docs.microsoft.com/en-us/azure/active-directory/develop/howto-create-service-principal-portal

QUESTION 143

A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices. The company must be able to monitor the devices in real-time.

You need to design the solution.

What should you recommend?

A. Azure Data Factory instance using Azure PowerShell
B. Azure Analysis Services using Microsoft Visual Studio
C. Azure Stream Analytics cloud job using Azure PowerShell
D. Azure Data Factory instance using Microsoft Visual Studio

Correct Answer: C **Section:**
(none) Explanation

Explanation/Reference:
Explanation:
Stream Analytics is a cost-effective event processing engine that helps uncover real-time insights from devices, sensors, infrastructure, applications and data quickly and easily.
Monitor and manage Stream Analytics resources with Azure PowerShell cmdlets and powershell scripting that execute basic Stream Analytics tasks.
Reference:
https://cloudblogs.microsoft.com/sqlserver/2014/10/29/microsoft-adds-iot-streaming-analytics-data-production-and-workflow-services-to-azure/

QUESTION 144
A company purchases IoT devices to monitor manufacturing machinery. The company uses an IoT appliance to communicate with the IoT devices. The company must be able to monitor the devices in real-time.

You need to design the solution.
What should you recommend?

A. Azure Data Factory instance using Azure PowerShell
B. Azure Analysis Services using Microsoft Visual Studio
C. Azure Stream Analytics Edge application using Microsoft Visual Studio
D. Azure Analysis Services using Azure PowerShell

Correct Answer: C
Section: (none)
Explanation

Explanation/Reference:
Explanation:
Azure Stream Analytics (ASA) on IoT Edge empowers developers to deploy near-real-time analytical intelligence closer to IoT devices so that they can unlock the full value of device-generated data.
References:
https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-edge

QUESTION 145
You plan to ingest streaming social media data by using Azure Stream Analytics. The data will be stored in files in Azure Data Lake Storage, and then consumed by using Azure Databricks and PolyBase in Azure SQL Data Warehouse.

You need to recommend a Stream Analytics data output format to ensure that the queries from Databricks and PolyBase against the files encounter the fewest possible errors. The solution must ensure that the files can be queried quickly and that the data type information is retained.

What should you recommend?

A. Avro
B. CSV
C. Parquet
D. JSON

Correct Answer: A **Section:**
(none) Explanation

Explanation/Reference:
Explanation:
The Avro format is great for data and message preservation.
Avro schema with its support for evolution is essential for making the data robust for streaming architectures like Kafka, and with the metadata that schema provides, you can reason on the data. Having a schema provides robustness in providing meta-data about the data stored in Avro records which are self-documenting the data.
References:
http://cloudurable.com/blog/avro/index.html

QUESTION 146
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

You have an Azure SQL database that has columns. The columns contain sensitive Personally Identifiable Information (PII) data.

You need to design a solution that tracks and stores all the queries executed against the PII data. You must be able to review the data in Azure Monitor, and the data must be available for at least 45 days.

Solution: You add classifications to the columns that contain sensitive data. You turn on Auditing and set the audit log destination to use Azure Log Analytics.

Does this meet the goal?

A. Yes
B. No

Correct Answer: B **Section:**
(none) Explanation

Explanation/Reference:
Explanation:
The default retention for Log Analytics is 31 days only. The Log Analytics retention settings allow you to configure a minimum of 31 days (if not using a free tier) up to 730 days.
You would need to reconfigure to at least 45 days, or, for example, use Azure Blob Storage as destination. Reference:
https://docs.microsoft.com/en-us/azure/sql-database/sql-database-auditing https://blogs.msdn.microsoft.com/canberrapfe/2017/01/25/change-oms-log-analytics-retention-period-in-the-azure-portal/

QUESTION 147
You have an Azure Storage account.

You plan to copy one million image files to the storage account.

You plan to share the files with an external partner organization. The partner organization will analyze the files during the next year.

You need to recommend an external access solution for the storage account. The solution must meet the following requirements:

- Ensure that only the partner organization can access the storage account.
- Ensure that access of the partner organization is removed automatically after 365 days.

What should you include in the recommendation?

A. shared keys
B. Azure Blob storage lifecycle management policies
C. Azure policies
D. shared access signature (SAS)

Correct Answer: D **Section:**
(none) Explanation

Explanation:
A shared access signature (SAS) is a URI that grants restricted access rights to Azure Storage resources.
You can provide a shared access signature to clients who should not be trusted with your storage account key but to whom you wish to delegate access to certain storage account resources. By distributing a shared access signature URI to these clients, you can grant them access to a resource for a specified period of time, with a specified set of permissions.
Reference:
https://docs.microsoft.com/en-us/rest/api/storageservices/delegate-access-with-shared-access-signature