Latest Version: 11.0

Question: 1

DRAG DROP
You need to ensure that the Twitter feed data can be analyzed in the dedicated SQL pool. The solution must meet the customer sentiment analytics requirements.
Which three Transaction-SQL DDL commands should you run in sequence? To answer, move the appropriate commands from the list of commands to the answer area and arrange them in the correct order.
NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Commands

CREATE EXTERNAL DATA SOURCE
CREATE EXTERNAL FILE FORMAT
CREATE EXTERNAL TABLE
CREATE EXTERNAL TABLE AS SELECT
CREATE DATABASE SCOPED CREDENTIAL

Answer Area

Answer:

CREATE EXTERNAL DATA SOURCE
CREATE EXTERNAL FILE FORMAT
CREATE EXTERNAL TABLE AS SELECT

Explanation:
Scenario: Allow Contoso users to use PolyBase in an Azure Synapse Analytics dedicated SQL pool to query the content of the data records that host the Twitter feeds. Data must be protected by using rowlevel
security (RLS). The users must be authenticated by using their own Azure AD credentials.
Box 1: CREATE EXTERNAL DATA SOURCE
External data sources are used to connect to storage accounts.
Box 2: CREATE EXTERNAL FILE FORMAT
CREATE EXTERNAL FILE FORMAT creates an external file format object that defines external data stored in Azure Blob Storage or Azure Data Lake Storage. Creating an external file format is a prerequisite for creating an external table.
Box 3: CREATE EXTERNAL TABLE AS SELECT

When used in conjunction with the CREATE TABLE AS SELECT statement, selecting from an external table imports data into a table within the SQL pool. In addition to the COPY statement, external tables are useful for loading data.
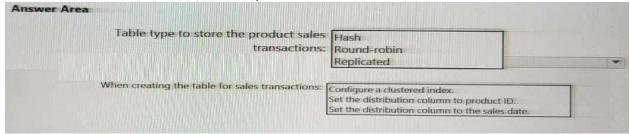Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-tables-external-tables
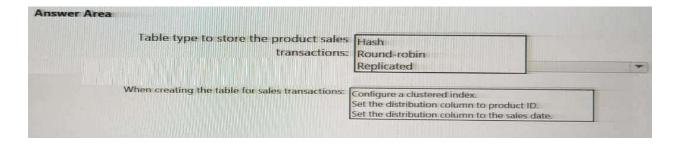
Question: 2

HOTSPOT
You need to design a data storage structure for the product sales transactions. The solution must meet the sales transaction dataset requirements.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Answer Area

Table type to store the product sales transactions:
- Hash
- Round-robin
- Replicated

When creating the table for sales transactions:
- Configure a clustered index.
- Set the distribution column to product ID.
- Set the distribution column to the sales date.

Answer:

Answer Area

Table type to store the product sales transactions:
- Hash
- Round-robin
- Replicated

When creating the table for sales transactions:
- Configure a clustered index.
- Set the distribution column to product ID.
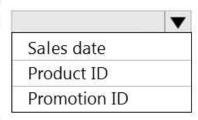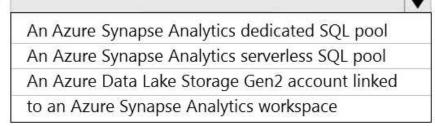- Set the distribution column to the sales date.

Question: 3

HOTSPOT
You need to design the partitions for the product sales transactions. The solution must meet the sales transaction dataset requirements.
What should you include in the solution? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
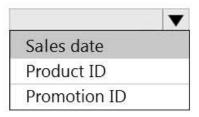
Partition product sales
transactions data by:

▼
Sales date
Product ID
Promotion ID

Store product sales
transactions data in:

▼
An Azure Synapse Analytics dedicated SQL pool
An Azure Synapse Analytics serverless SQL pool
An Azure Data Lake Storage Gen2 account linked to an Azure Synapse Analytics workspace

Answer:

Partition product sales
transactions data by:

▼
Sales date
Product ID
Promotion ID

Store product sales
transactions data in:

▼
An Azure Synapse Analytics dedicated SQL pool
An Azure Synapse Analytics serverless SQL pool
An Azure Data Lake Storage Gen2 account linked to an Azure Synapse Analytics workspace

Explanation:
Box 1: Sales date
Scenario: Contoso requirements for data integration include:
Partition data that contains sales transaction records. Partitions must be designed to provide efficient
loads by month. Boundary values must belong to the partition on the right.
Box 2: An Azure Synapse Analytics Dedicated SQL pool
Scenario: Contoso requirements for data integration include:
Ensure that data storage costs and performance are predictable.
The size of a dedicated SQL pool (formerly SQL DW) is determined by Data Warehousing Units (DWU).
Dedicated SQL pool (formerly SQL DW) stores data in relational tables with columnar storage. This

format significantly reduces the data storage costs, and improves query performance.
Synapse analytics dedicated sql pool
Reference:
https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouseoverview-what-is
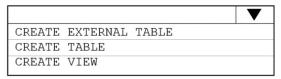
Question: 4

HOTSPOT
You need to implement an Azure Synapse Analytics database object for storing the sales transactions data. The solution must meet the sales transaction dataset requirements.
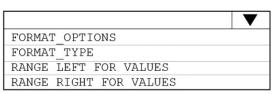What solution must meet the sales transaction dataset requirements.
What should you do? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.
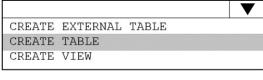
Transact-SQL DDL command to use: ▼

| CREATE EXTERNAL TABLE |
| CREATE TABLE |
| CREATE VIEW |

Partitioning option to use in the WITH clause of the DDL statement: ▼

| FORMAT_OPTIONS |
| FORMAT_TYPE |
| RANGE LEFT FOR VALUES |
| RANGE RIGHT FOR VALUES |

Answer:

Transact-SQL DDL command to use: ▼

| CREATE EXTERNAL TABLE |
| CREATE TABLE |
| CREATE VIEW |

Partitioning option to use in the WITH clause of the DDL statement: ▼

| FORMAT_OPTIONS |
| FORMAT_TYPE |
| RANGE LEFT FOR VALUES |
| RANGE RIGHT FOR VALUES |

Explanation:
Box 1: Create table
Scenario: Load the sales transaction dataset to Azure Synapse Analytics
Box 2: RANGE RIGHT FOR VALUES
Scenario: Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.
RANGE RIGHT: Specifies the boundary value belongs to the partition on the right (higher values).

FOR VALUES (boundary_value [,...n]): Specifies the boundary values for the partition.
Scenario: Load the sales transaction dataset to Azure Synapse Analytics.
Contoso identifies the following requirements for the sales transaction dataset:
Partition data that contains sales transaction records. Partitions must be designed to provide efficient loads by month. Boundary values must belong to the partition on the right.
Ensure that queries joining and filtering sales transaction records based on product ID complete as quickly as possible.
Implement a surrogate key to account for changes to the retail store addresses.
Ensure that data storage costs and performance are predictable.
Minimize how long it takes to remove old records.
Reference:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-table-azure-sql-data-warehouse

Question: 5

You need to integrate the on-premises data sources and Azure Synapse Analytics. The solution must meet the data integration requirements.
Which type of integration runtime should you use?

A. Azure-SSIS integration runtime
B. self-hosted integration runtime
C. Azure integration runtime

Answer: C