

- **Vendor: Microsoft**
- **Exam Code: DP-203**
- **Exam Name: Data Engineering on Microsoft Azure**
- **New Updated Questions from [Braindump2go](#)**
- **(Updated in [March/2021](#))**

[Visit Braindump2go and Download Full Version DP-203 Exam Dumps](#)

Topic 1, Contoso Case Study

Transactional Data

Contoso has three years of customer, transactional, operation, sourcing, and supplier data comprised of 10 billion records stored across multiple on-premises Microsoft SQL Server servers. The SQL server instances contain data from various operational systems. The data is loaded into the instances by using SQL server integration Services (SSIS) packages.

You estimate that combining all product sales transactions into a company-wide sales transactions dataset will result in a single table that contains 5 billion rows, with one row per transaction.

Most queries targeting the sales transactions data will be used to identify which products were sold in retail stores and which products were sold online during different time period. Sales transaction data that is older than three years will be removed monthly.

You plan to create a retail store table that will contain the address of each retail store. The table will be approximately 2 MB. Queries for retail store sales will include the retail store addresses.

You plan to create a promotional table that will contain a promotion ID. The promotion ID will be associated to a specific product. The product will be identified by a product ID. The table will be approximately 5 GB.

Streaming Twitter Data

The ecommerce department at Contoso develops and Azure logic app that captures trending Twitter feeds referencing the company's products and pushes the products to Azure Event Hubs.

Planned Changes

Contoso plans to implement the following changes:

- * Load the sales transaction dataset to Azure Synapse Analytics.
- * Integrate on-premises data stores with Azure Synapse Analytics by using SSIS packages.
- * Use Azure Synapse Analytics to analyze Twitter feeds to assess customer sentiments about products.

Sales Transaction Dataset Requirements

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.

[DP-203 Exam Dumps](#) [DP-203 Exam Questions](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#)

<https://www.braindump2go.com/dp-203.html>

- Ensure that data storage costs and performance are predictable.
- Minimize how long it takes to remove old records.

Customer Sentiment Analytics Requirement

Contoso identifies the following requirements for customer sentiment analytics:

- 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 row-level security (RLS). The users must be authenticated by using their own Azure AD credentials.
 - Maximize the throughput of ingesting Twitter feeds from Event Hubs to Azure Storage without purchasing additional throughput or capacity units.
 - Store Twitter feeds in Azure Storage by using Event Hubs Capture. The feeds will be converted into Parquet files.
 - Ensure that the data store supports Azure AD-based access control down to the object level.
 - Minimize administrative effort to maintain the Twitter feed data records.
 - Purge Twitter feed data records that are older than two years.

Data Integration Requirements

Contoso identifies the following requirements for data integration:

Use an Azure service that leverages the existing SSIS packages to ingest on-premises data into datasets stored in a dedicated SQL pool of Azure Synapse Analytics and transform the data.

Identify a process to ensure that changes to the ingestion and transformation activities can be

version controlled and developed independently by multiple data engineers.

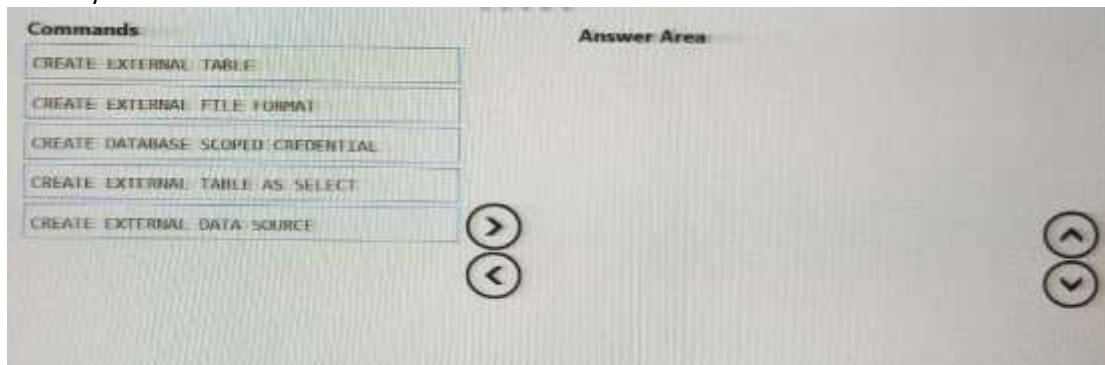
Question: 1

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.



Answer:

1. Create External Table

- 2. Create External Data Source
- 3. Create External Data Source

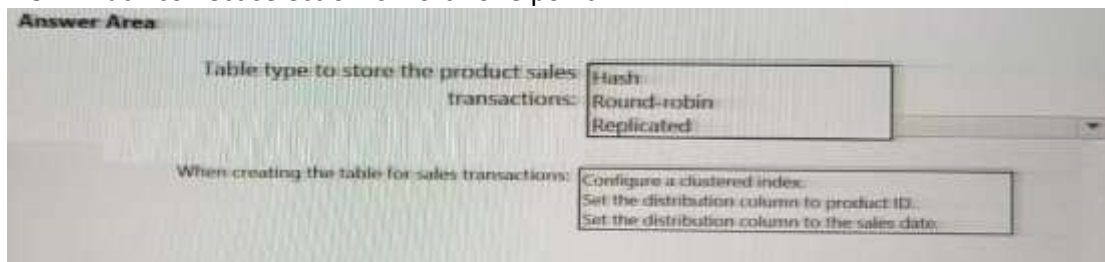
Question:2

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:

- 1. Replicated
- 2. Set the distribution column to product ID.

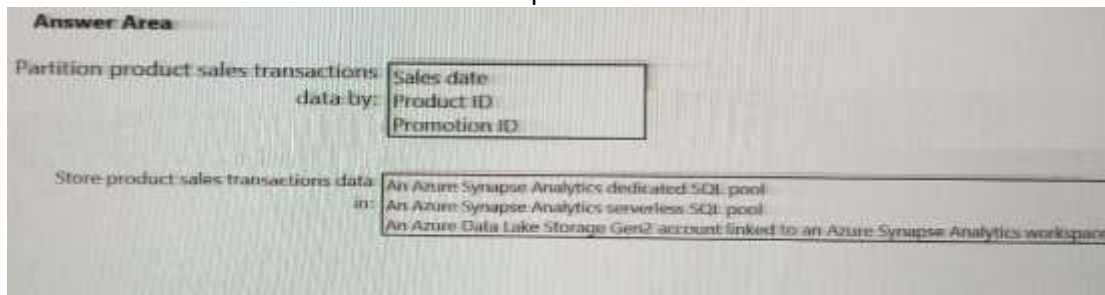
Question:3

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.



Answer:

- 1. Sales date
- 2. An Azure Synapse Analytics dedicated SQL pool.

Question:4

HOTSPOT

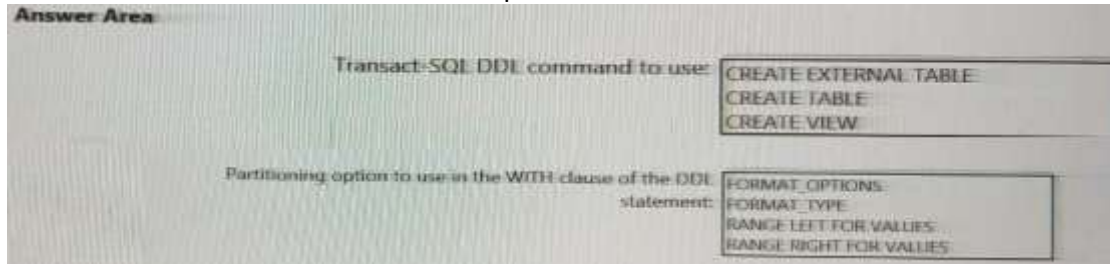
You need to implement an Azure Synapse Analytics database object for storing the sales transactions data

a. 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.



The screenshot shows an 'Answer Area' with two dropdown menus. The first dropdown is labeled 'Transact-SQL DDL command to use:' and contains three options: 'CREATE EXTERNAL TABLE', 'CREATE TABLE', and 'CREATE VIEW'. The second dropdown is labeled 'Partitioning option to use in the WITH clause of the DDL statement:' and contains three options: 'FORMAT OPTIONS', 'FORMAT TYPE', and 'RANGE RIGHT FOR VALUES'.

Answer:

1. Create Table
2. Range Right for values

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: A

Topic 2, Mix Questions

Question:6

You have an Azure Data Lake Storage account that has a virtual network service endpoint configured.

You plan to use Azure Data Factory to extract data from the Data Lake Storage account. The data will then be loaded to a data warehouse in Azure Synapse Analytics by using PolyBase.

Which authentication method should you use to access Data Lake Storage?

- A. shared access key authentication

[DP-203 Exam Dumps](#) [DP-203 Exam Questions](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#)

<https://www.braindump2go.com/dp-203.html>

- B. managed identity authentication
- C. account key authentication
- D. service principal authentication

Answer: B

Reference:

<https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse#use-polybase-to-load-data-into-azure-sql-data-warehouse>

Question: 7

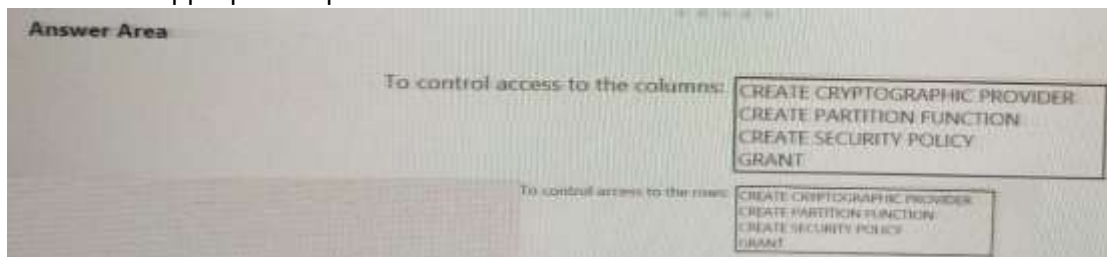
HOTSPOT

You have an Azure subscription that contains the following resources:

- * An Azure Active Directory (Azure AD) tenant that contains a security group named Group1.
- * An Azure Synapse Analytics SQL pool named Pool1.

You need to control the access of Group1 to specific columns and rows in a table in Pool1. Which Transact-SQL commands should you use? To answer, select the appropriate options in the answer area.

NOTE: Each appropriate options in the answer area.



Answer:

1. Create Security Policy
2. Create Cryptographic Provider

Question: 8

HOTSPOT

You need to implement an Azure Databricks cluster that automatically connects to Azure Data Lake Storage Gen2 by using Azure Active Directory (Azure AD) integration.

How should you configure the new cluster? To answer, select the appropriate options in the answer area.

NOTE: Each correct election is worth one point.

Cluster Mode:

High Concurrency
Premium
Standard

Advanced option to enable:

Azure Data Lake Storage Gen1 Credential Passthrough
Table Access Control

Answer:

Cluster Mode:

	▼
High Concurrency	
Premium	
Standard	

Advanced option to enable:

	▼
Azure Data Lake Storage Gen1 Credential Passthrough	
Table Access Control	

Box 1: High Concurrency

Enable Azure Data Lake Storage credential passthrough for a high-concurrency cluster.

Incorrect:

Support for Azure Data Lake Storage credential passthrough on standard clusters is in Public Preview.

Standard clusters with credential passthrough are supported on Databricks Runtime 5.5 and above and are limited to a single user.

Box 2: Azure Data Lake Storage Gen1 Credential Passthrough

You can authenticate automatically to Azure Data Lake Storage Gen1 and Azure Data Lake Storage Gen2 from Azure Databricks clusters using the same Azure Active Directory (Azure AD) identity that you use to log into Azure Databricks. When you enable your cluster for Azure Data Lake Storage credential passthrough, commands that you run on that cluster can read and write data in Azure Data Lake Storage without requiring you to configure service principal credentialsforaccesstostorage.

References:

<https://docs.azuredatabricks.net/spark/latest/data-sources/azure/adls-passthrough.html>**Question: 9**

You have an Azure Synapse Analytics dedicated SQL pool that contains a table named Contacts. Contacts contains a column named Phone.

You need to ensure that users in a specific role only see the last four digits of a phone number when querying the Phone column.

What should you include in the solution?

- A. a default value
- B. dynamic data masking
- C. row-level security (RLS)
- D. column encryption
- E. table partitions

Answer: C

Question: 10

HOTSPOT

You have an Azure Synapse Analytics dedicated SQL pool that contains the users shown in the following table.

Name	Role
User1	Server admin
User2	db_datareader

User1 executes a query on the database, and the query returns the results shown in the following exhibit.

```
1: SELECT c.name,
2:     tbl.name as table_name,
3:     typ.name as datatype,
4:     c.is_masked,
5:     c.masking_function
6: FROM sys.masked_columns AS c
7: INNER JOIN sys.tables AS tbl ON c.[object_id] = tbl.[object_id]
8: INNER JOIN sys.types typ ON c.user_type_id = typ.user_type_id
9: WHERE is_masked = 1;
10
```

Results		Messages			
name	table_name	datatype	is_masked	masking_function	
1	BirthDate	DimCustomer	date	1	default()
2	Gender	DimCustomer	nvarchar	1	default()
3	EmailAddress	DimCustomer	nvarchar	1	email()
4	YearlyIncome	DimCustomer	money	1	default()

[DP-203 Exam Dumps](#) [DP-203 Exam Questions](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#)

<https://www.braindump2go.com/dp-203.html>

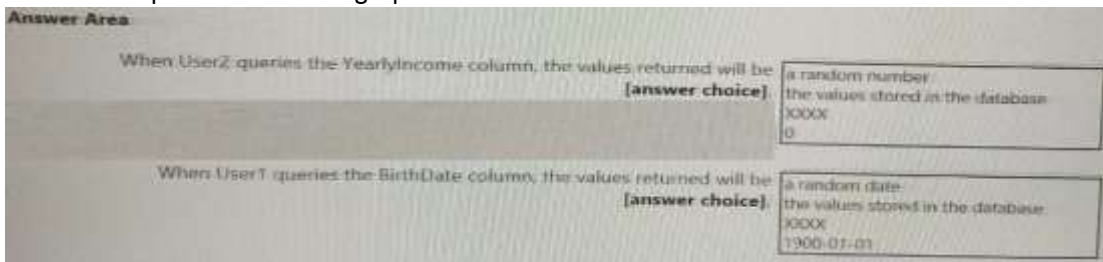
User1 is the only user who has access to the unmasked data.

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

Answer Area

When User2 queries the YearlyIncome column, the values returned will be [answer choice]

When User1 queries the BirthDate column, the values returned will be [answer choice]



Answer:

1. A Random number
2. 1900-01-01

Question: 11

You develop data engineering solutions for a company.

A project requires the deployment of data to Azure Data Lake Storage.

You need to implement role-based access control (RBAC) so that project member can manage the Azure Data Lake Storage resources.

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

NOTE: Each correct selection is worth one point.

- A. Assign Azure AD security groups to Azure Data Lake Storage.
- B. Configure end-user authentication for the Azure Data Lake Storage account.
- C. Configure service-to-service authentication for the Azure Data Lake Storage account.
- D. Create security groups in Azure Active Directory (Azure AD) and add project members.
- E. Configure access control lists (ACL) for the Azure Data Lake Storage account.

Answer: ADE

References:

<https://docs.microsoft.com/en-us/azure/data-lake-store/data-lake-store-secure-data>

Question: 12

You are designing an Azure Synapse Analytics dedicated SQL pool.

You need to ensure that you can audit access to Personally Identifiable information (PII).

What should you include in the solution?

- A. dynamic data masking
- B. row-level security (RLS)
- C. sensitivity classifications
- D. column-level security

Answer: D

Question: 13

You are designing a sales transactions table in an Azure Synapse Analytics dedicated SQL pool. The table will contain approximately 60 million rows per month and will be partitioned by month. The table will use a clustered columnstore index and round-robin distribution.

Approximately how many rows will there be for each combination of distribution and partition?

- A. 1 million
- B. 5 million
- C. 20 million
- D. 60 million

Answer: D

<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-tables-partition>

Question: 14

You are designing a dimension table for a data warehouse. The table will track the value of the dimension attributes over time and preserve the history of the data by adding new rows as the data changes.

Which type of slowly changing dimension (SCD) should use?

- A. Type 0
- B. Type 1
- C. Type 2
- D. Type 3

Answer: C

Type 2 - Creating a new additional record. In this methodology all history of dimension changes is kept in the database. You capture attribute change by adding a new row with a new surrogate key to the dimension table. Both the prior and new rows contain as attributes the natural key (or other durable identifier). Also 'effective date' and 'current indicator' columns are used in this method. There could be only one record with current indicator set to 'Y'. For 'effective date' columns, i.e. start_date and end_date, the end_date for current record usually is set to value 9999-12-31. Introducing changes to the dimensional model in type 2 could be very expensive database operation so it is not recommended to use it in dimensions where a new attribute could be added in the future.

<https://www.datawarehouse4u.info/SCD-Slowly-Changing-Dimensions.html>

Question: 15

[DP-203 Exam Dumps](#) [DP-203 Exam Questions](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#)

<https://www.braindump2go.com/dp-203.html>

You are designing an inventory updates table in an Azure Synapse Analytics dedicated SQL pool. The table will have a clustered columnstore index and will include the following columns:

- EventDate: 1 million per day
- EventTypeID: 10 million per event type
- WarehouseID: 100 million per warehouse
- ProductCategoryTypeID: 25 million per product category type

You identify the following usage patterns:

Analyst will most commonly analyze transactions for a warehouse.

Queries will summarize by product category type, date, and/or inventory event type.

You need to recommend a partition strategy for the table to minimize query times.

On which column should you recommend partitioning the table?

- A. ProductCategoryTypeID
- B. EventDate
- C. WarehouseID
- D. EventTypeID

Answer: D

Question: 16

HOTSPOT

You are designing an application that will store petabytes of medical imaging data

When the data is first created, the data will be accessed frequently during the first week. After one month, the data must be accessible within 30 seconds, but files will be accessed infrequently. After one year, the data will be accessed infrequently but must be accessible within five minutes.

You need to select a storage strategy for the data

- a. The solution must minimize costs

Which storage tier should you use for each time frame? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

First week:

▼
Archive
Cool
Hot

After one month:

▼
Archive
Cool
Hot

After one year:

▼
Archive
Cool
Hot

Answer:

First week:

▼
Archive
Cool
Hot

After one month:

▼
Archive
Cool
Hot

After one year:

▼
Archive
Cool
Hot

Explanation:

First week: Hot

[DP-203 Exam Dumps](#) [DP-203 Exam Questions](#) [DP-203 PDF Dumps](#) [DP-203 VCE Dumps](#)

<https://www.braindump2go.com/dp-203.html>

Hot - Optimized for storing data that is accessed frequently.

After one month: Cool

Cool - Optimized for storing data that is infrequently accessed and stored for at least 30 days.

After one year: Cool

Incorrect Answers:

Archive: Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements (on the order of hours).

References:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>