



DP-203^{Q&As}

Data Engineering on Microsoft Azure

Pass Microsoft DP-203 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

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

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





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 in this scenario, 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 Storage account that contains 100 GB of files. The files contain text and numerical values. 75% of the rows contain description data that has an average length of 1.1 MB.

You plan to copy the data from the storage account to an Azure SQL data warehouse.

You need to prepare the files to ensure that the data copies quickly.

Solution: You modify the files to ensure that each row is less than 1 MB.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Polybase loads rows that are smaller than 1 MB.

Note on Polybase Load: PolyBase is a technology that accesses external data stored in Azure Blob storage or Azure Data Lake Store via the T-SQL language.

Extract, Load, and Transform (ELT)

Extract, Load, and Transform (ELT) is a process by which data is extracted from a source system, loaded into a data warehouse, and then transformed.

The basic steps for implementing a PolyBase ELT for dedicated SQL pool are:

Extract the source data into text files.

Land the data into Azure Blob storage or Azure Data Lake Store.

Prepare the data for loading.

Load the data into dedicated SQL pool staging tables using PolyBase.

Transform the data.

Insert the data into production tables.

Reference: <https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data-warehouse/sql-data-warehouse-service-capacity-limits>



<https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/load-data-overview>

QUESTION 2

HOTSPOT

You are designing an application that will use an Azure Data Lake Storage Gen 2 account to store petabytes of license plate photos from toll booths. The account will use zone- redundant storage (ZRS).

You identify the following usage patterns:

1.

The data will be accessed several times a day during the first 30 days after the data is created.

2.

The data must meet an availability SU of 99.9%.

3.

After 90 days, the data will be accessed infrequently but must be available within 30 seconds.

4.

After 365 days, the data will be accessed infrequently but must be available within five minutes.

Hot Area:



Frist 30 days

	▼
Archive	
Cool	
Hot	

Frist 90 days

	▼
Archive	
Cool	
Hot	

After 365 days

	▼
Archive	
Cool	
Hot	

Correct Answer:



Frist 30 days

	▼
Archive	
Cool	
Hot	

Frist 90 days

	▼
Archive	
Cool	
Hot	

After 365 days

	▼
Archive	
Cool	
Hot	

QUESTION 3

You are designing a streaming data solution that will ingest variable volumes of data. You need to ensure that you can change the partition count after creation.

Which service should you use to ingest the data?

- A. Azure Event Hubs Dedicated
- B. Azure Stream Analytics
- C. Azure Data Factory
- D. Azure Synapse Analytics



Correct Answer: A

You can\\'t change the partition count for an event hub after its creation except for the event hub in a dedicated cluster.

Reference: <https://docs.microsoft.com/en-us/azure/event-hubs/event-hubs-features>

QUESTION 4

HOTSPOT

You use Azure Data Factory to prepare data to be queried by Azure Synapse Analytics serverless SQL pools.

Files are initially ingested into an Azure Data Lake Storage Gen2 account as 10 small JSON files. Each file contains the same data attributes and data from a subsidiary of your company.

You need to move the files to a different folder and transform the data to meet the following requirements:

1.

Provide the fastest possible query times.

2.

Automatically infer the schema from the underlying files.

How should you configure the Data Factory copy activity? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

Copy behavior:

	▼
Flatten hierarchy	
Merge files	
Preserve hierarchy	

Sink file type:

	▼
CSV	
JSON	
Parquet	
TXT	

Correct Answer:



Answer Area

Copy behavior:

	▼
Flatten hierarchy	
Merge files	
Preserve hierarchy	

Sink file type:

	▼
CSV	
JSON	
Parquet	
TXT	

Box 1: Preserver herarchy

Compared to the flat namespace on Blob storage, the hierarchical namespace greatly improves the performance of directory management operations, which improves overall job performance.

Box 2: Parquet

Azure Data Factory parquet format is supported for Azure Data Lake Storage Gen2.

Parquet supports the schema property.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-introduction>

<https://docs.microsoft.com/en-us/azure/data-factory/format-parquet>

QUESTION 5

What should you recommend to prevent users outside the Litware on-premises network from accessing the analytical data store?

A. a server-level virtual network rule



B. a database-level virtual network rule

C. a server-level firewall IP rule

D. a database-level firewall IP rule

Correct Answer: C

Scenario:

1.

Ensure that the analytical data store is accessible only to the company's on-premises network and Azure services.

2.

Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure.

Since Litware does not plan to implement Azure ExpressRoute or a VPN between the on-premises network and Azure, they will have to create firewall IP rules to allow connection from the IP ranges of the on-premise network. They can also use the firewall rule 0.0.0.0 to allow access from Azure services.

Reference: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-vnet-service-endpoint-rule-overview>

[DP-203 PDF Dumps](#)

[DP-203 VCE Dumps](#)

[DP-203 Practice Test](#)