

Pass Guaranteed Quiz Perfect DP-203 - Data Engineering on Microsoft Azure Reliable Exam Sample



DOWNLOAD the newest ITPassLeader DP-203 PDF dumps from Cloud Storage for free: https://drive.google.com/open?id=1foq2PP3327UMSSGEZJEVpZZqaobm_b1_

Up to now, there are three versions of DP-203 exam materials for your choice. So high-quality contents and flexible choices of DP-203 learning mode will bring about the excellent learning experience for you. Though the content of these three versions of our DP-203 study questions is the same, their displays are totally different. And you can be surprised to find that our DP-203 learning quiz is developed with the latest technologies as well.

Microsoft DP-203: Data Engineering on Microsoft Azure exam is an important certification for professionals working in the field of data engineering. It tests the candidate's skills and knowledge in designing and implementing data solutions using Azure services. Candidates can prepare for the exam by taking the official Microsoft training course or by using study materials available online. Upon passing the exam, candidates will earn the Microsoft Certified: Azure Data Engineer Associate certification, which can lead to better job opportunities and higher salaries.

[**>> DP-203 Reliable Exam Sample <<**](#)

Exam DP-203 Course - DP-203 Reliable Test Notes

It never needs an internet connection. ITPassLeader's Data Engineering on Microsoft Azure practice exam software has several mock exams, designed just like the real exam. Microsoft DP-203 practice exam software contains all the important questions which have a greater chance of appearing in the final exam. ITPassLeader always tries to ensure that you are provided with the most updated Data Engineering on Microsoft Azure (DP-203) Exam Questions to pass the exam on the first attempt.

By passing the Microsoft DP-203 certification exam, you can demonstrate your proficiency in data engineering on the Azure platform and increase your chances of getting hired or promoted. Data Engineering on Microsoft Azure certification also provides you with access to Microsoft's vast community of data professionals, which can help you stay current with the latest trends and technologies in the field. Overall, the Microsoft DP-203 Certification is an excellent investment for anyone looking to build a career in data engineering on the Azure platform.

Microsoft Data Engineering on Microsoft Azure Sample Questions (Q283-Q288):

NEW QUESTION # 283

You are designing a fact table named FactPurchase in an Azure Synapse Analytics dedicated SQL pool. The table contains purchases from suppliers for a retail store. FactPurchase will contain the following columns.

Name	Data type	Nullable
PurchaseKey	Bigint	No
DateKey	Int	No
SupplierKey	Int	No
StockItemKey	Int	No
PurchaseOrderID	Int	Yes
OrderedQuantity	Int	No
OrderedOuters	Int	No
ReceivedOuters	Int	No
Package	Nvarchar(50)	No
IsOrderFinalized	Bit	No
LineageKey	Int	No

FactPurchase will have 1 million rows of data added daily and will contain three years of data.

Transact-SQL queries similar to the following query will be executed daily.

```
SELECT
    SupplierKey, StockItemKey, COUNT(*)
FROM FactPurchase
WHERE DateKey >= 20210101
AND DateKey <= 20210131
GROUP By SupplierKey, StockItemKey
```

Which table distribution will minimize query times?

- A. hash-distributed on DateKey
- B. round-robin
- C. replicated
- D. hash-distributed on PurchaseKey

Answer: D

Explanation:

Explanation

Hash-distributed tables improve query performance on large fact tables, and are the focus of this article.

Round-robin tables are useful for improving loading speed.

Reference:

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

NEW QUESTION # 284

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.

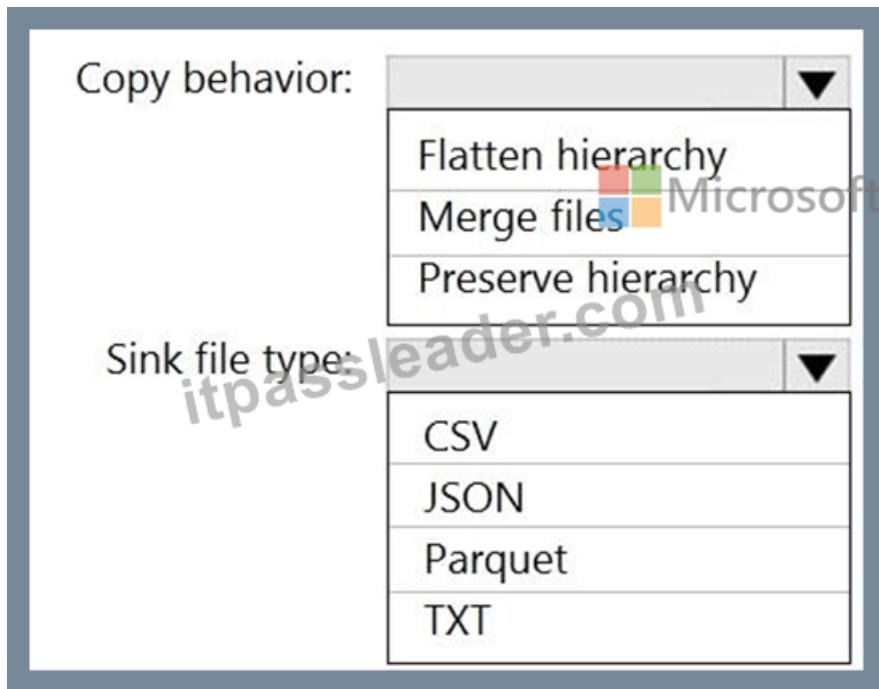
You need to move the files to a different folder and transform the data to meet the following requirements:

Provide the fastest possible query times.

Automatically infer the schema from the underlying files.

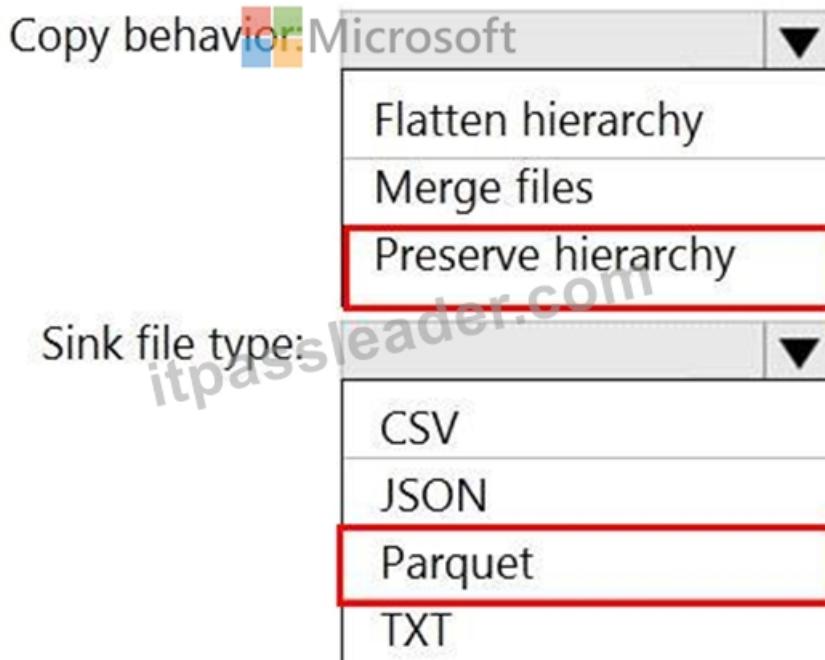
How should you configure the Data Factory copy activity? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.



Answer:

Explanation:



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>

NEW QUESTION # 285

You are designing an Azure Synapse Analytics dedicated SQL pool.

Groups will have access to sensitive data in the pool as shown in the following table.

Name	Enhanced access
Executives	No access to sensitive data
Analysts	Access to in-region sensitive data
Engineers	Access to all numeric sensitive data

You have policies for the sensitive data. The policies vary by region as shown in the following table.

Region	Data considered sensitive
RegionA	Financial, Personally Identifiable Information (PII)
RegionB	Financial, Personally Identifiable Information (PII), medical
RegionC	Financial, medical

You have a table of patients for each region. The tables contain the following potentially sensitive columns.

Name	Sensitive data	Description
CardOnFile	Financial	Debit/credit card number for charges
Height	Medical	Patient's height in cm
ContactEmail	PII	Email address for secure communications

You are designing dynamic data masking to maintain compliance.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Statements	Yes	No
Analysts in RegionA require dynamic data masking rules for [Patients_RegionA].	<input type="radio"/>	<input type="radio"/>
Engineers in RegionC require a dynamic data masking rule for [Patients_RegionA], [Height]	<input type="radio"/>	<input type="radio"/>
Engineers in RegionB require a dynamic data masking rule for [Patients_RegionB], [Height]	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation:

Statements	Yes	No
Analysts in RegionA require dynamic data masking rules for [Patients_RegionA].	<input checked="" type="checkbox"/>	<input type="radio"/>
Engineers in RegionC require a dynamic data masking rule for [Patients_RegionA], [Height]	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Engineers in RegionB require a dynamic data masking rule for [Patients_RegionB], [Height]	<input checked="" type="checkbox"/>	<input type="radio"/>

Explanation

Text Description automatically generated



Yes No

Analysts in RegionA require dynamic data masking rules for [Patients_RegionA].

Engineers in RegionC require a dynamic data masking rule for [Patients_RegionA], [Height]

Engineers in RegionB require a dynamic data masking rule for [Patients_RegionB], [Height]

Reference:

<https://docs.microsoft.com/en-us/azure/azure-sql/database/dynamic-data-masking-overview>

NEW QUESTION # 286

You have a table named SalesFact in an enterprise data warehouse in Azure Synapse Analytics. SalesFact contains sales data from the past 36 months and has the following characteristics:

Is partitioned by month

Contains one billion rows

Has clustered columnstore indexes

At the beginning of each month, you need to remove data from SalesFact that is older than 36 months as quickly as possible.

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

ACTIONS

ANSWER AREA

Switch the partition containing the stale data from SalesFact to SalesFact_Work.

Truncate the partition containing the stale data.

Drop the SalesFact_Work table.

Create an empty table named SalesFact_Work that has the same schema as SalesFact.

Execute a DELETE statement where the value in the Date column is more than 36 months ago.

Copy the data to a new table by using CREATE TABLE AS SELECT (CTAS).

Answer:

Explanation:

Actions	Answer Area
Switch the partition containing the stale data from SalesFact to SalesFact_Work.	Create an empty table named SalesFact_Work that has the same schema as SalesFact.
Truncate the partition containing the stale data.	Switch the partition containing the stale data from SalesFact to SalesFact_Work.
Drop the SalesFact_Work table.	Drop the SalesFact_Work table.
Create an empty table named SalesFact_Work that has the same schema as SalesFact.	
Execute a DELETE statement where the value in the Date column is more than 36 months ago.	
Copy the data to a new table by using CREATE TABLE AS SELECT (CTAS).	

Explanation

Create an empty table named SalesFact_Work that has the same schema as SalesFact.

Switch the partition containing the stale data from SalesFact to SalesFact_Work.

Drop the SalesFact_Work table.

Step 1: Create an empty table named SalesFact_Work that has the same schema as SalesFact.

Step 2: Switch the partition containing the stale data from SalesFact to SalesFact_Work.

SQL Data Warehouse supports partition splitting, merging, and switching. To switch partitions between two tables, you must ensure that the partitions align on their respective boundaries and that the table definitions match.

Loading data into partitions with partition switching is a convenient way stage new data in a table that is not visible to users the switch in the new data.

Step 3: Drop the SalesFact_Work table.

Reference:

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

NEW QUESTION # 287

You have an Azure Synapse Analytics dedicated SQL pool.

You need to create a copy of the data warehouse and make the copy available for 28 days. The solution must minimize costs.

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.

Actions	Answer Area
Restore the latest automatic restore point to a new data warehouse.	
Restore the copy from the latest automatic restore point to the current data warehouse.	
Create a new user-defined restore point.	
Restore the copy from the new user-defined restore point to a new data warehouse.	
Pause the restored data warehouse.	

Answer:

Explanation:

Actions	Answer Area
Restore the latest automatic restore point to a new data warehouse.	
Restore the copy from the latest automatic restore point to the current data warehouse.	Create a new user-defined restore point.
Create a new user-defined restore point.	
Restore the copy from the new user-defined restore point to a new data warehouse.	Restore the copy from the new user-defined restore point to a new data warehouse.
Pause the restored data warehouse.	Pause the restored data warehouse.

Explanation:

Actions	Answer Area
:: Restore the latest automatic restore point to a new data warehouse.	1 :: Create a new user-defined restore point.
:: Restore the copy from the latest automatic restore point to the current data warehouse.	2 :: Restore the copy from the new user-defined restore point to a new data warehouse.
	3 :: Pause the restored data warehouse.

NEW QUESTION # 288

Exam DP-203 Course: <https://www.itpassleader.com/Microsoft/DP-203-dumps-pass-exam.html>

P.S. Free 2026 Microsoft DP-203 dumps are available on Google Drive shared by ITPassLeader: https://drive.google.com/open?id=1fqo2PP3327UMSSGEZJEVpZZqaobm_b1