



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.pass4itsure.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

You plan to create a dimension table in Azure Synapse Analytics that will be less than 1 GB.

You need to create the table to meet the following requirements:

1.

Provide the fastest Query time.

2.

Minimize data movement during queries.

Which type of table should you use?

A. hash distributed

B. heap

C. replicated

D. round-robin

Correct Answer: C

A replicated table has a full copy of the table accessible on each Compute node. Replicating a table removes the need to transfer data among Compute nodes before a join or aggregation. Since the table has multiple copies, replicated tables work best when the table size is less than 2 GB compressed. 2 GB is not a hard limit.

Reference: https://docs.microsoft.com/en-us/azure/synapse-analytics/sql-data- warehouse/design- guidance-for-replicated-tables

QUESTION 2

HOTSPOT

You have an Azure Synapse Analytics dedicated SQL pool named Pool1 that contains an external table named Sales. Sales contains sales data. Each row in Sales contains data on a single sale, including the name of the salesperson.

You need to implement row-level security (RLS). The solution must ensure that the salespeople can access only their respective sales.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Create:		-
	A materialized view in Pool1	
	A security policy for Sales	
	Database scoped credentials in Pool1	
Restrict row access by using:	▼	
Restrict row access by using:	A masking rule	
Restrict row access by using:	A masking rule A table-valued function	

Correct Answer:

Re

Create:		-
	A materialized view in Pool1	
	A security policy for Sales	
	Database scoped credentials in Pool1	
trict row access by using:		
	A masking rule	
	A masking rule	
	A table-valued function	

Box 1: A security policy for sale

Here are the steps to create a security policy for Sales:

Create a user-defined function that returns the name of the current user:

CREATE FUNCTION dbo.GetCurrentUser()

RETURNS NVARCHAR(128)

AS

BEGIN

RETURN SUSER_SNAME();

END;

Create a security predicate function that filters the Sales table based on the

current user:



CREATE FUNCTION dbo.SalesPredicate(@salesperson NVARCHAR(128)) RETURNS TABLE

WITH SCHEMABINDING

AS

RETURN SELECT 1 AS access_result

WHERE @salesperson = SalespersonName;

Create a security policy on the Sales table that uses the SalesPredicate function to filter the data:

CREATE SECURITY POLICY SalesFilter

ADD FILTER PREDICATE dbo.SalesPredicate(dbo.GetCurrentUser()) ON dbo.Sales

WITH (STATE = ON);

By creating a security policy for the Sales table, you ensure that each salesperson can only access their own sales data. The security policy uses a user-defined function to get the name of the current user and a security predicate function to

filter the Sales table based on the current user.

Box 2: table-value function

to restrict row access by using row-level security, you need to create a table-valued function that returns a table of values that represent the rows that a user can access. You then use this function in a security policy that applies a predicate on

the table.

QUESTION 3

You have an Azure Data Lake Storage Gen2 account named account1 that contains a container named container1.

You plan to create lifecycle management policy rules for container1.

You need to ensure that you can create rules that will move blobs between access tiers based on when each blob was accessed last.

What should you do first?

- A. Configure object replication
- B. Create an Azure application
- C. Enable access time tracking
- D. Enable the hierarchical namespace

Correct Answer: C

Generally available: Access time-based lifecycle management rules for Data Lake Storage Gen2

Some data in Azure Storage is written once and read many times. To effectively manage the lifecycle of such data and



optimize your storage costs, it is important to know the last time of access for the data. When access time tracking is

enabled for a storage account, the last access time property on the file is updated when it is read. You can then define lifecycle management policies based on last access time:

Transition objects from hotter to cooler access tiers if the file has not been accessed for a specified duration.

Automatically transition objects from cooler to hotter access tiers when a file is accessed again.

Delete objects if they have not been accessed for an extended duration.

Access time tracking is only available for files in Data Lake Storage Gen2.

Reference:

https://azure.microsoft.com/en-us/updates/access-time-based-lifecycle-management-adls-gen2/

QUESTION 4

HOTSPOT

You have an Azure Data Factory instance named ADF1 and two Azure Synapse Analytics workspaces named WS1 and WS2.

ADF1 contains the following pipelines:

1.

P1: Uses a copy activity to copy data from a nonpartitioned table in a dedicated SQL pool of WS1 to an Azure Data Lake Storage Gen2 account

2.

P2: Uses a copy activity to copy data from text-delimited files in an Azure Data Lake Storage Gen2 account to a nonpartitioned table in a dedicated SQL pool of WS2

You need to configure P1 and P2 to maximize parallelism and performance.

Which dataset settings should you configure for the copy activity if each pipeline? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Answer Area

P1:

Set the Copy method to Bulk insert Set the Copy method to PolyBase Set the Isolation level to Repeatable read

Set the Partition option to Dynamic range

P2:

Set the Copy method to Bulk insert	
Set the Copy method to PolyBase	
Set the Isolation level to Repeatable read	
Set the Partition option to Dynamic range	

Correct Answer:



Answer Area

P1:

Set the Copy method to Bulk insert Set the Copy method to PolyBase Set the Isolation level to Repeatable read Set the Partition option to Dynamic range

P2:

Set the Copy method to Bulk insert Set the Copy method to PolyBase Set the Isolation level to Repeatable read Set the Partition option to Dynamic range

Box 1: Set the Copy method to PolyBase

While SQL pool supports many loading methods including non-Polybase options such as BCP and SQL BulkCopy API, the fastest and most scalable way to load data is through PolyBase. PolyBase is a technology that accesses external data

stored in Azure Blob storage or Azure Data Lake Store via the T-SQL language.

Box 2: Set the Copy method to Bulk insert

Polybase not possible for text files. Have to use Bulk insert.

Reference:

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

QUESTION 5

You have an Azure data factory that connects to a Microsoft Purview account. The data factory is registered in Microsoft Purview.

You update a Data Factory pipeline.



You need to ensure that the updated lineage is available in Microsoft Purview.

What should you do first?

- A. Locate the related asset in the Microsoft Purview portal.
- B. Execute the pipeline.
- C. Disconnect the Microsoft Purview account from the data factory.
- D. Execute an Azure DevOps build pipeline.

Correct Answer: B

Run pipeline and push lineage data to Microsoft Purview

Step 1: Connect Data Factory to your Microsoft Purview account

Step 2: Run pipeline in Data Factory

You can create pipelines, Copy activities and Dataflow activities in Data Factory. You don///t need any additional configuration for lineage data capture. The lineage data will automatically be captured during the activities execution.

Step 3: Monitor lineage reporting status

After you run the pipeline, in the pipeline monitoring view, you can check the lineage reporting status by clicking the following Lineage status button.

Step 4: View lineage information in your Microsoft Purview account

On Microsoft Purview UI, you can browse assets and choose type "Azure Data Factory". You can also search the Data Catalog using keywords.

Reference:

https://learn.microsoft.com/en-us/azure/data-factory/tutorial-push-lineage-to-purview

Latest DP-203 Dumps

DP-203 Practice Test

DP-203 Study Guide