DP-300 Q&As

Administering Relational Databases on Microsoft Azure

Pass Microsoft DP-300 Exam with 100% Guarantee

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

https://www.geekcert.com/dp-300.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

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 of each pipeline? To answer, select the appropriate options in the answer area.

Hot Area:

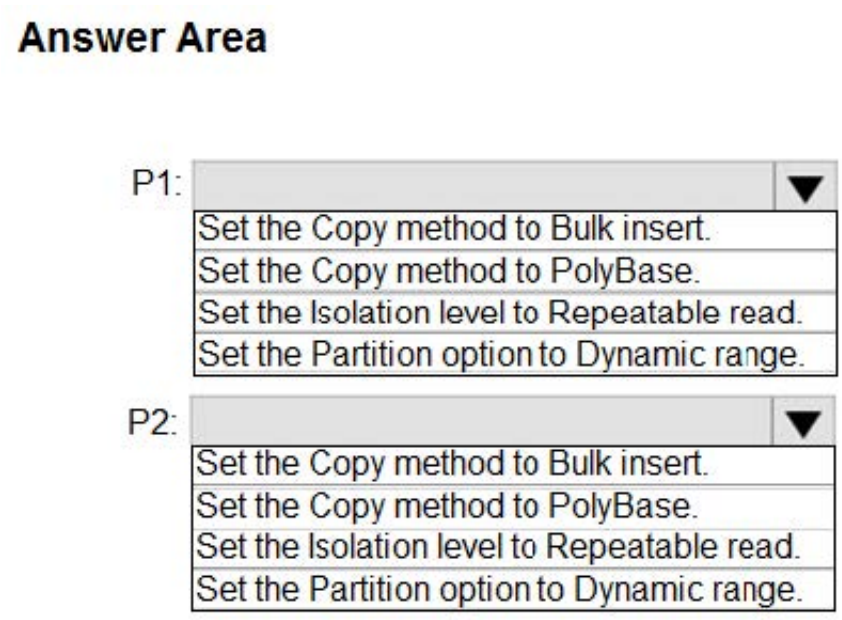

Correct Answer:

Answer Area

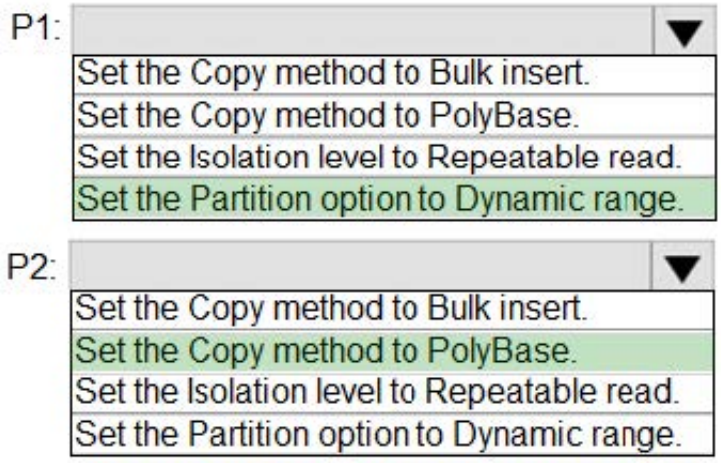

P1: Set the Partition option to Dynamic Range.

The SQL Server connector in copy activity provides built-in data partitioning to copy data in parallel.

P2: Set the Copy method to PolyBase

Polybase is the most efficient way to move data into Azure Synapse Analytics. Use the staging blob feature to achieve high load speeds from all types of data stores, including Azure Blob storage and Data Lake Store. (Polybase supports

Azure Blob storage and Azure Data Lake Store by default.)

Reference:

https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-sql-data-warehouse

https://docs.microsoft.com/en-us/azure/data-factory/load-azure-sql-data-warehouse

QUESTION 2

HOTSPOT

You are performing exploratory analysis of bus fare data in an Azure Data Lake Storage Gen2 account by using an Azure Synapse Analytics serverless SQL pool.

You execute the Transact-SQL query shown in the following exhibit.

```
SELECT
   payment_type,
   SUM(fare_amount) AS fare_total
FROM OPENROWSET(
          BULK 'csv/busfare/tripdata_2020*.csv',
          DATA_SOURCE = 'BusData',
          FORMAT = 'CSV', PARSER_VERSION = '2.0',
          FIRSTROW = 2
   )
   WITH (
          payment_type INT 10,
          fare_amount FLOAT 11
   ) AS nyc
GROUP BY payment_type
ORDER BY payment_type;
```

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

Hot Area:

Answer Area

The query results include only **[answer choice]** in the csv/busfare folder.

- CSV files in the tripdata_2020 subfolder
- files that have files names beginning with "tripdata_2020"
- CSV files that have file names containing "tripdata_202'
- CSV files that have file named beginning with "tripdata_2020"

The query assumes that the first row in a CSV file is **[answer choice]** row.

- a header
- a data
- an empty

Correct Answer:

Answer Area

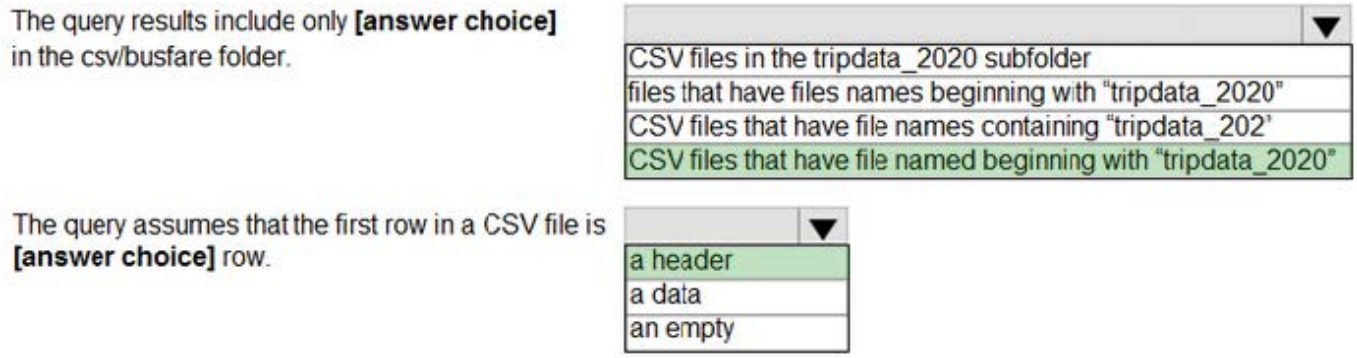

Box 1: CSV files that have file named beginning with "tripdata_2020"

Box 2: a header

FIRSTROW = \\'first_row\\'

Specifies the number of the first row to load. The default is 1 and indicates the first row in the specified data file. The row numbers are determined by counting the row terminators. FIRSTROW is 1-based.

Reference:

https://docs.microsoft.com/en-us/azure/synapse-analytics/sql/develop-openrowset

---

QUESTION 3

You have a SQL pool in Azure Synapse that contains a table named dbo.Customers. The table contains a column name Email.

You need to prevent nonadministrative users from seeing the full email addresses in the Email column. The users must see values in a format of aXXX@XXXX.com instead.

What should you do?

A. From the Azure portal, set a mask on the Email column.

B. From the Azure portal, set a sensitivity classification of Confidential for the Email column.

C. From Microsoft SQL Server Management Studio, set an email mask on the Email column.

D. From Microsoft SQL Server Management Studio, grant the SELECT permission to the users for all the columns in the dbo.Customers table except Email.

Correct Answer: A

QUESTION 4

You have an Azure SQL managed instance.

You need to enable SQL Agent Job email notifications.

What should you do?

A. Use the Agent XPs option.

B. Enable the SQL Server Agent.

C. Run the sp_configurecommand.

D. Run the sp_set_agent_propertiescommand.

Correct Answer: C

Reference: https://techcommunity.microsoft.com/t5/azure-sql-blog/sending-emails-in-azure-sql-managed-instance/ba-p/386235

QUESTION 5

DRAG DROP

You have an Azure SQL database named DB1.

You need to create a partitioned table in DB1.

Which three objects should you create in sequence? To answer, move the appropriate objects from the list of objects to the answer area and arrange them in the correct order.

Select and Place:

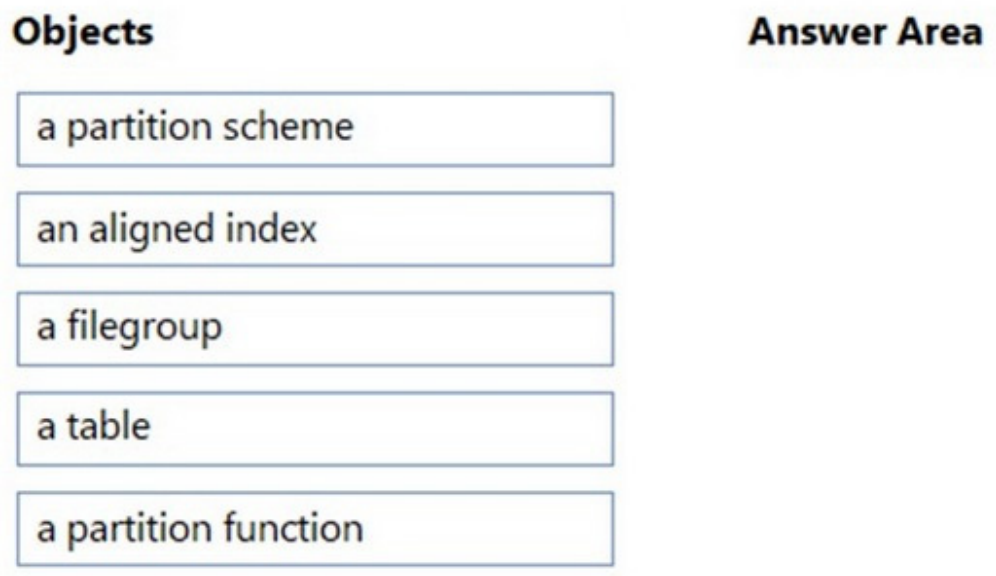

Correct Answer:

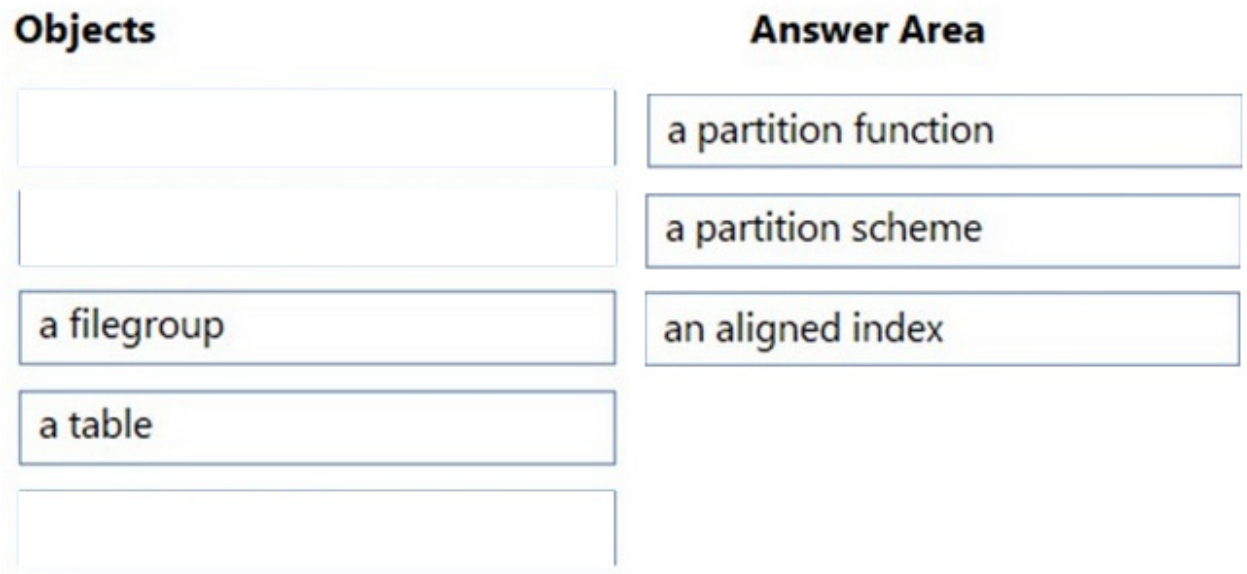

Create partitioned tables and indexes

Creating a partitioned table or index typically happens in three or four parts:

1.

Optionally create a filegroup or filegroups and corresponding data files that will hold the partitions specified by the partition scheme. The main reason to place partitions on multiple filegroups is to ensure you can independently perform backup and restore operations on filegroups. If this is not required, you may choose to assign all partitions to a single filegroup, using either an existing filegroup, such as PRIMARY, or a new filegroup with related data file(s). In nearly all scenarios, you will achieve all benefits of partitioning whether or not you use multiple filegroups.

2.

Create a partition function that maps the rows of a table or index into partitions based on the values of a specified column. You can use a single partition function to partition multiple objects. (Step 1)

3.

Create a partition scheme that maps the partitions of a partitioned table or index to one filegroup or to multiple filegroups. You can use a single partition scheme to partition multiple objects. (Step 2)

4.

Create or alter a table or index and specify the partition scheme as the storage location, along with the column that will serve as the partitioning column. (Step 3)

Aligns all indexes of the table that are partitioned with the same partition scheme. When a table and its indexes are aligned, you can move partitions in and out of partitioned tables more effectively, because your data is partitioned in the same way. Step 1: a partition function Step 2: a partition scheme Step 3: an aligned index Reference:

https://learn.microsoft.com/en-us/sql/relational-databases/partitions/create-partitioned-tables-and-indexes