



DP-600^{Q&As}

Implementing Analytics Solutions Using Microsoft Fabric

Pass Microsoft DP-600 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/dp-600.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 have a Fabric tenant that contains a warehouse. The warehouse uses row-level security (RLS). You create a Direct Lake semantic model that uses the Delta tables and RLS of the warehouse. When users interact with a report built from the model, which mode will be used by the DAX queries?

- A. DirectQuery
- B. Dual
- C. Direct Lake
- D. Import

Correct Answer: A

Explanation: When users interact with a report built from a Direct Lake semantic model that uses row-level security (RLS), the DAX queries will operate in DirectQuery mode (A). This is because the model directly queries the underlying data source without importing data into Power BI. References = The Power BI documentation on DirectQuery provides detailed explanations of how RLS and DAX queries function in this mode.

QUESTION 2

You have a Fabric tenant that contains two lakehouses.

You are building a dataflow that will combine data from the lakehouses. The applied steps from one of the queries in the dataflow is shown in the following exhibit.



Query settings >

▼ Properties

Name

Customers1

Entity type ⓘ

Custom

▼ Applied steps

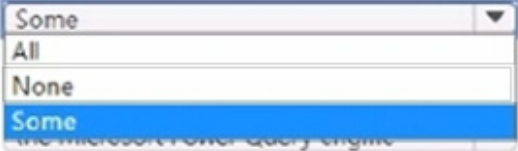
- Source
- Navigation 1
- Capitalized each word
- Appended query
- Changed column type
- Filtered rows**

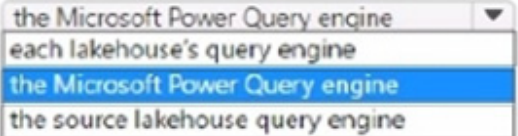
- Edit settings
- Rename
- Delete
- Delete until end
- Insert step after
- Move before
- Move after
- Extract previous...
- View data source query
- View query plan
- Properties...

Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

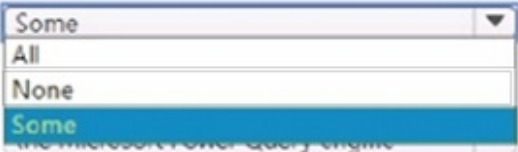


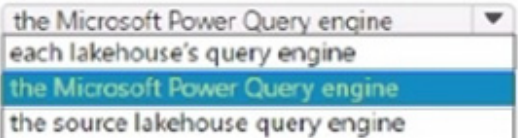
Hot Area:

[Answer choice] of the transformation steps in the query will fold. 

The Added custom step will be performed in [answer choice]. 

Correct Answer:

[Answer choice] of the transformation steps in the query will fold. 

The Added custom step will be performed in [answer choice]. 

Folding in Power Query refers to operations that can be translated into source queries. In this case, "some" of the steps can be folded, which means that some transformations will be executed at the data source level. The steps that cannot be folded will be executed within the Power Query engine. Custom steps, especially those that are not standard query operations, are usually executed within Power Query engine rather than being pushed down to the source system. References = Query folding in Power Query Power Query M formula language

QUESTION 3


You have a Fabric tenant that contains a warehouse named Warehouse1. Warehouse1 contains three schemas named schemaA, schemaB, and schemaC.

You need to ensure that a user named User1 can truncate tables in schemaA only.


How should you complete the T-SQL statement? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Correct Answer:



GRANT ALTER ON SCHEMA::schemaA TO User1;

The ALTER permission allows a user to modify the schema of an object, and granting ALTER on a schema will allow the user to perform operations like TRUNCATE TABLE on any object within that schema. It is the correct permission to

grant to User1 for truncating tables in schemaA.

References =

GRANT Schema Permissions

Permissions That Can Be Granted on a Schema

QUESTION 4

You have a Fabric tenant that contains a lakehouse named Lakehouse1

Readings from 100 IoT devices are appended to a Delta table in Lakehouse1. Each set of readings is approximately 25 KB. Approximately 10 GB of data is received daily.

All the table and SparkSession settings are set to the default.

You discover that queries are slow to execute. In addition, the lakehouse storage contains data and log files that are no longer used.

You need to remove the files that are no longer used and combine small files into larger files with a target size of 1 GB per file.

What should you do? To answer, drag the appropriate actions to the correct requirements.



Each action may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Select and Place:

Actions	Answer Area
<input type="text" value="Set the autoCompact table setting."/>	Remove the files: <input type="text"/>
<input type="text" value="Set the optimizeWrite table setting."/>	Combine the files: <input type="text"/>
<input type="text" value="Run the VACUUM command on a schedule."/>	
<input type="text" value="Set the autoCompact SparkSession setting."/>	
<input type="text" value="Run the OPTIMIZE command on a schedule."/>	
<input type="text" value="Set the parallelDelete SparkSession setting."/>	

Correct Answer:

Actions	Answer Area
<input type="text" value="Set the autoCompact table setting."/>	Remove the files: <input type="text" value="Run the VACUUM command on a schedule."/>
<input type="text" value="Set the autoCompact SparkSession setting."/>	Combine the files: <input type="text" value="Set the optimizeWrite table setting."/>
<input type="text" value="Run the OPTIMIZE command on a schedule."/>	
<input type="text" value="Set the parallelDelete SparkSession setting."/>	

Remove the files: Run the VACUUM command on a schedule.

Combine the files: Set the optimizeWrite table setting. or Run the OPTIMIZE command on a schedule.

To remove files that are no longer used, the VACUUM command is used in Delta Lake to clean up invalid files from a table. To combine smaller files into larger ones, you can either set the optimizeWrite setting to combine files during write operations or use the OPTIMIZE command, which is a Delta Lake operation used to compact small files into larger ones.

QUESTION 5

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Hot Area:



Statements	Yes	No
The Spark engine will read only the 'SalesOrderNumber', 'OrderDate','CustomerName', 'UnitPrice' columns from Sales_raw.csv.	<input type="radio"/>	<input type="radio"/>
Removing the partition will reduce the execution time of the query.	<input type="radio"/>	<input type="radio"/>
Adding <code>inferSchema='true'</code> to the options will increase the execution time of the query.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Statements	Yes	No
The Spark engine will read only the 'SalesOrderNumber', 'OrderDate','CustomerName', 'UnitPrice' columns from Sales_raw.csv.	<input checked="" type="radio"/>	<input type="radio"/>
Removing the partition will reduce the execution time of the query.	<input type="radio"/>	<input checked="" type="radio"/>
Adding <code>inferSchema='true'</code> to the options will increase the execution time of the query.	<input checked="" type="radio"/>	<input type="radio"/>

The Spark engine will read only the `\\SalesOrderNumber\\`, `\\OrderDate\\`, `\\CustomerName\\`, `\\UnitPrice\\` columns from Sales_raw.csv. - Yes
Removing the partition will reduce the execution time of the query. - No
Adding `inferSchema=\\true\\` to the options will increase the execution time of the query. - Yes

The code specifies the selection of certain columns, which means only those columns will be read into the DataFrame. Partitions in Spark are a way to optimize the execution of queries by organizing the data into parts that can be processed in parallel. Removing the partition could potentially increase the execution time because Spark would no longer be able to process the data in parallel efficiently. The `inferSchema` option allows Spark to automatically detect the column data types, which can increase the execution time of the initial read operation because it requires Spark to read through the data to infer the schema.

QUESTION 6

You have a Fabric tenant that contains a lakehouse.

You plan to query sales data files by using the SQL endpoint. The files will be in an Amazon Simple Storage Service (Amazon S3) storage bucket.

You need to recommend which file format to use and where to create a shortcut.

Which two actions should you include in the recommendation? Each correct answer presents part of the solution.

NOTE: Each correct answer is worth one point.

- A. Create a shortcut in the Files section.
- B. Use the Parquet format



- C. Use the CSV format.
- D. Create a shortcut in the Tables section.
- E. Use the delta format.

Correct Answer: BD

Explanation: You should use the Parquet format (B) for the sales data files because it is optimized for performance with large datasets in analytical processing and create a shortcut in the Tables section (D) to facilitate SQL queries through the lakehouse's SQL endpoint. References = The best practices for working with file formats and shortcuts in a lakehouse environment are covered in the lakehouse and SQL endpoint documentation provided by the cloud data platform services.

QUESTION 7

You have a Fabric tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a subfolder named Subfolder1 that contains CSV files. You need to convert the CSV files into the delta format that has V-Order optimization enabled. What should you do from Lakehouse explorer?

- A. Use the Load to Tables feature.
- B. Create a new shortcut in the Files section.
- C. Create a new shortcut in the Tables section.
- D. Use the Optimize feature.

Correct Answer: D

Explanation: To convert CSV files into the delta format with Z-Order optimization enabled, you should use the Optimize feature (D) from Lakehouse Explorer. This will allow you to optimize the file organization for the most efficient querying. References = The process for converting and optimizing file formats within a lakehouse is discussed in the lakehouse management documentation.

QUESTION 8

You have source data in a folder on a local computer.

You need to create a solution that will use Fabric to populate a data store. The solution must meet the following requirements:

Support the use of dataflows to load and append data to the data store.

Ensure that Delta tables are V-Order optimized and compacted automatically.

Which type of data store should you use?

- A. a lakehouse
- B. an Azure SQL database
- C. a warehouse



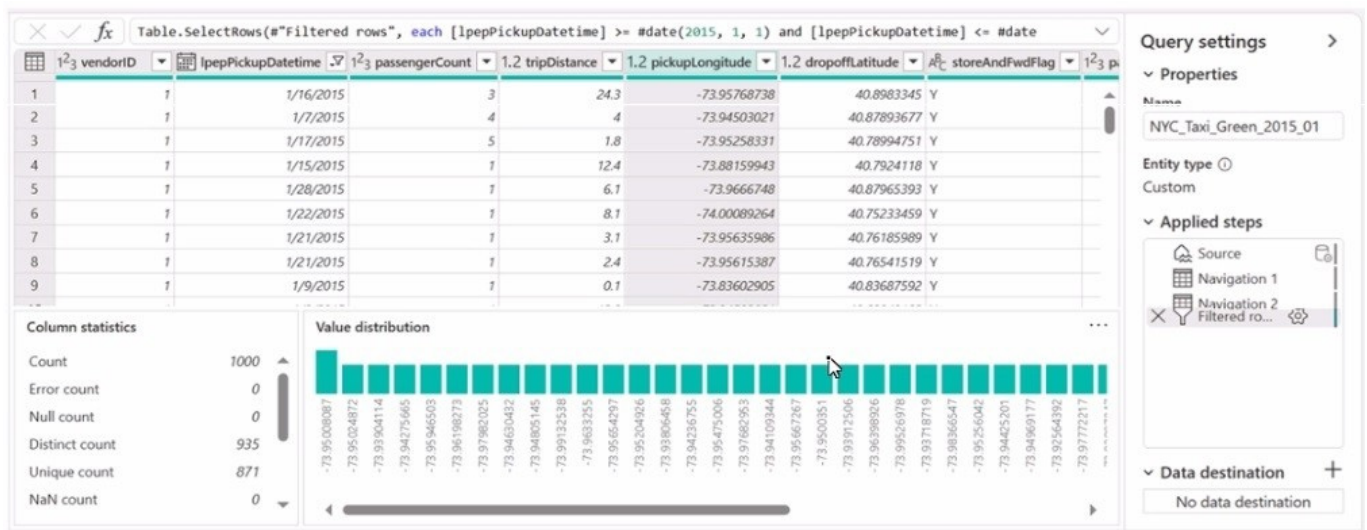
D. a KQL database

Correct Answer: A

Explanation: A lakehouse (A) is the type of data store you should use. It supports dataflows to load and append data and ensures that Delta tables are Z-Order optimized and compacted automatically. References = The capabilities of a lakehouse and its support for Delta tables are described in the lakehouse and Delta table documentation.

QUESTION 9

You have a Fabric workspace named Workspace 1 that contains a dataflow named Dataflow1. Dataflow1 has a query that returns 2,000 rows. You view the query in Power Query as shown in the following exhibit.



What can you identify about the pickupLongitude column?

- A. The column has duplicate values.
- B. All the table rows are profiled.
- C. The column has missing values.
- D. There are 935 values that occur only once.

Correct Answer: A

Explanation: The pickupLongitude column has duplicate values. This can be inferred because the 'Distinct count' is 935 while the 'Count' is 1000, indicating that there are repeated values within the column. References = Microsoft Power BI documentation on data profiling could provide further insights into understanding and interpreting column statistics like these.

QUESTION 10

You have a Fabric tenant that contains 30 CSV files in OneLake. The files are updated daily.

You create a Microsoft Power BI semantic model named Modell that uses the CSV files as a data source. You configure



incremental refresh for Model 1 and publish the model to a Premium capacity in the Fabric tenant.

When you initiate a refresh of Model1, the refresh fails after running out of resources.

What is a possible cause of the failure?

- A. Query folding is occurring.
- B. Only refresh complete days is selected.
- C. XMLA Endpoint is set to Read Only.
- D. Query folding is NOT occurring.
- E. The data type of the column used to partition the data has changed.

Correct Answer: D

Explanation: A possible cause for the failure is that query folding is NOT occurring (D). Query folding helps optimize refresh by pushing down the query logic to the source system, reducing the amount of data processed and transferred, hence conserving resources. References = The Power BI documentation on incremental refresh and query folding provides detailed information on this topic.

QUESTION 11

You have a Fabric tenant that contains a semantic model. The model uses Direct Lake mode.

You suspect that some DAX queries load unnecessary columns into memory.

You need to identify the frequently used columns that are loaded into memory.

What are two ways to achieve the goal? Each correct answer presents a complete solution.

NOTE: Each correct answer is worth one point.

- A. Use the Analyze in Excel feature.
- B. Use the Vertipaq Analyzer tool.
- C. Query the \$system.discovered_STORAGE_TABLE_COLUMNS_IN_SEGMENTS dynamic management view (DMV).
- D. Query the discover_hehory6Rant dynamic management view (DMV).

Correct Answer: BC

Explanation: The Vertipaq Analyzer tool (B) and querying the \$system.discovered_STORAGE_TABLE_COLUMNS_IN_SEGMENTS dynamic management view (DMV) (C) can help identify which columns are frequently loaded into memory. Both methods provide insights into the storage and retrieval aspects of the semantic model. References = The Power BI documentation on Vertipaq Analyzer and DMV queries offers detailed guidance on how to use these tools for performance analysis.

QUESTION 12



You have a Fabric tenant that contains a warehouse named Warehouse1. Warehouse1 contains a fact table named FactSales that has one billion rows. You run the following TSQL statement.

```
CREATE TABLE test.FactSales AS CLONE OF Dbo.FactSales;
```

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Hot Area:

Statements	Yes	No
A replica of <code>dbo.Sales</code> is created in the test schema by copying the metadata only.	<input type="radio"/>	<input type="radio"/>
Additional schema changes to <code>dbo.FactSales</code> will also apply to <code>test.FactSales</code> .	<input type="radio"/>	<input type="radio"/>
Additional data changes to <code>dbo.FactSales</code> will also apply to <code>test.FactSales</code> .	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Statements	Yes	No
A replica of <code>dbo.Sales</code> is created in the test schema by copying the metadata only.	<input type="radio"/>	<input checked="" type="radio"/>
Additional schema changes to <code>dbo.FactSales</code> will also apply to <code>test.FactSales</code> .	<input type="radio"/>	<input checked="" type="radio"/>
Additional data changes to <code>dbo.FactSales</code> will also apply to <code>test.FactSales</code> .	<input checked="" type="radio"/>	<input type="radio"/>

A replica of `dbo.Sales` is created in the test schema by copying the metadata only. - No Additional schema changes to `dbo.FactSales` will also apply to `test.FactSales`. - No Additional data changes to `dbo.FactSales` will also apply to `test.FactSales`. - Yes The CREATE TABLE AS CLONE statement creates a copy of an existing table, including its data and any associated data structures, like indexes. Therefore, the statement does not merely copy metadata; it also copies the data. However, subsequent schema changes to the original table do not automatically propagate to the cloned table. Any data changes in the original table after the clone operation will not be reflected in the clone unless explicitly updated. References = CREATE TABLE AS SELECT (CTAS) in SQL Data Warehouse

QUESTION 13

You have a Fabric tenant that contains a new semantic model in OneLake.

You use a Fabric notebook to read the data into a Spark DataFrame.

You need to evaluate the data to calculate the min, max, mean, and standard deviation values for all the string and numeric columns.

Solution: You use the following PySpark expression:

```
df.explain()
```



Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Explanation: The `df.explain()` method does not meet the goal of evaluating data to calculate statistical functions. It is used to display the physical plan that Spark will execute. References = The correct usage of the `explain()` function can be found in the PySpark documentation.

[Latest DP-600 Dumps](#)

[DP-600 PDF Dumps](#)

[DP-600 VCE Dumps](#)