



Microsoft DP-700 Mock Exam

Shared by Hurst on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page



Question 1

Question Type: MultipleChoice

You have a Fabric workspace named Workspace1. Your company acquires GitHub licenses.

You need to configure source control for Workspace1 to use GitHub. The solution must follow the principle of least privilege. Which permissions do you require to ensure that you can commit code to GitHub?

Options:

- A- Actions (Read and write) and Contents (Read and write)
- B- Actions (Read and write) only
- C- Contents (Read and write) only
- D- Contents (Read) and Commit statuses (Read and write)

Answer:

C

Question 2

Question Type: MultipleChoice

You have a Fabric workspace that contains a data pipeline named Pipeline1 as shown in the exhibit.

The screenshot displays the Microsoft Fabric Data Pipeline console. At the top, there is a navigation bar with tabs for Home, Activities, Run, and View. Below this is a toolbar with icons for Validate, Run, Schedule, Trigger (preview), View run history, Copy data, Dataflow, Notebook, and Lookup. The main area shows a data pipeline named 'Pipeline1' with two activities: 'Execute procedure1' and 'Copy_kdi'. Both activities have a green checkmark indicating they are successful. Below the pipeline, there is a section for 'Pipeline run ID: 77c397af-ba17-48c2-9242-4b259aecdb3d' and 'Pipeline status: Succeeded'. A table shows the run details for the two activities.

Activity name	Activity status	Run start	Duration	Input
Copy_kdi	Succeeded	8/8/2024, 2:36:27 PM	33s	
Execute procedure1	Inactive	8/8/2024, 2:36:27 PM	Less than 1s	

Options:

- A- Both activities will run simultaneously.
- B- Both activities will be skipped.
- C- Execute procedure1 will run and Copy_kdi will be skipped.
- D- Copy.kdi will run and Execute procedure1 will be skipped.
- E- Execute procedure1 will run first, and then Copy_kdi will run.
- F- Copy.kdi will run first, and then Execute procedure1 will run.

Answer:

B

Question 3

Question Type: MultipleChoice

You have a Fabric workspace that contains an eventstream named EventStream1. EventStream1 outputs events to a table in a lakehouse.

You need to remove files that are older than seven days and are no longer in use.

Which command should you run?

Options:

- A- VACUUM
- B- COMPUTE
- C- OPTIMIZE
- D- CLONE

Answer:

A

Explanation:

VACUUM is used to clean up storage by removing files no longer in use by a Delta table. It removes old and unreferenced files from Delta tables. For example, to remove files older than 7 days:

```
VACUUM delta.`/path_to_table` RETAIN 7 HOURS;
```

Question 4

Question Type: MultipleChoice

You need to ensure that processes for the bronze and silver layers run in isolation. How should you configure the Apache Spark settings?

Options:

- A- Modify the number of executors.
- B- Disable high concurrency.
- C- Create a custom pool.
- D- Set the default environment.

Answer:

C

Question 5

Question Type: MultipleChoice

You have an Azure SQL database named DB1.

In a Fabric workspace, you deploy an eventstream named EventStreamDBI to stream record changes from DB1 into a lakehouse.

You discover that events are NOT being propagated to EventStreamDBI.

You need to ensure that the events are propagated to EventStreamDBI.

What should you do?

Options:

- A- Create a read-only replica of DB1.
- B- Create an Azure Stream Analytics job.
- C- Enable Extended Events for DB1.
- D- Enable change data capture (CDC) for DB1.

Answer:

D

Question 6

Question Type: MultipleChoice

What should you do to optimize the query experience for the business users?

Options:

- A- Enable V-Order.
- B- Create and update statistics.
- C- Run the VACUUM command.
- D- Introduce primary keys.

Answer:

B

Question 7

Question Type: MultipleChoice

You have a Fabric workspace named Workspace1 that contains a warehouse named Warehouse1.

You plan to deploy Warehouse1 to a new workspace named Workspace2.

As part of the deployment process, you need to verify whether Warehouse1 contains invalid references. The solution must minimize development effort.

What should you use?

Options:

- A- a database project
- B- a deployment pipeline
- C- a Python script
- D- a T-SQL script

Answer:

C

Explanation:

A deployment pipeline in Fabric allows you to deploy assets like warehouses, datasets, and reports between different workspaces (such as from Workspace1 to Workspace2). One of the key features of a deployment pipeline is the ability to check for invalid references before deployment. This can help identify issues with assets, such as broken links or dependencies, ensuring the deployment is successful without introducing errors. This is the most efficient way to verify references and manage the deployment with minimal development effort.

Question 8

Question Type: MultipleChoice

You have a Fabric workspace that contains a warehouse named Warehouse1. Data is loaded daily into Warehouse1 by using data pipelines and stored procedures.

You discover that the daily data load takes longer than expected.

You need to monitor Warehouse1 to identify the names of users that are actively running queries.

Which view should you use?

Options:

- A- sys.dm_exec_connections
- B- sys.dm_exec_requests
- C- queryinsights.long_running_queries
- D- queryinsights.frequently_run_queries
- E- sys.dm_exec_sessions

Answer:

E

Explanation:

sys.dm_exec_sessions provides real-time information about all active sessions, including the

user, session ID, and status of the session. You can filter on session status to see users actively running queries.

Question 9

Question Type: MultipleChoice

You have a Fabric workspace named Workspace1 that contains a data pipeline named Pipeline1 and a lakehouse named Lakehouse1.

You have a deployment pipeline named deployPipeline1 that deploys Workspace1 to Workspace2.

You restructure Workspace1 by adding a folder named Folder1 and moving Pipeline1 to Folder1.

You use deployPipeline1 to deploy Workspace1 to Workspace2.

What occurs to Workspace2?

Options:

- A- Folder1 is created, Pipeline1 moves to Folder1, and Lakehouse1 is deployed.
- B- Only Pipeline1 and Lakehouse1 are deployed.
- C- Folder1 is created, and Pipeline1 and Lakehouse1 move to Folder1.
- D- Only Folder1 is created and Pipeline1 moves to Folder1.

Answer:

A

Explanation:

When you restructure Workspace1 by adding a new folder (Folder1) and moving Pipeline1 into it, deployPipeline1 will deploy the entire structure of Workspace1 to Workspace2, preserving the changes made in Workspace1. This includes:

Folder1 will be created in Workspace2, mirroring the structure in Workspace1.

Pipeline1 will be moved into Folder1 in Workspace2, maintaining the same folder structure.

Lakehouse1 will be deployed to Workspace2 as it exists in Workspace1.

Question 10

Question Type: MultipleChoice

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 section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a Fabric eventstream that loads data into a table named Bike_Location in a KQL database. The table contains the following columns:

BikepointID

Street

Neighbourhood

No_Bikes

No_Empty_Docks

Timestamp

You need to apply transformation and filter logic to prepare the data for consumption. The solution must return data for a neighbourhood named Sands End when No_Bikes is at least 15. The results must be ordered by No_Bikes in ascending order.

Solution: You use the following code segment:

```
bike_location
| filter Neighbourhood == "Sands End" and No_Bikes >= 15
| sort by No_Bikes asc
| project BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
```

Does this meet the goal?

Options:

A- Yes

B- no

Answer:

A

Explanation:

Filter Condition: It correctly filters rows where Neighbourhood is 'Sands End' and No_Bikes is greater than or equal to 15.

Sorting: The sorting is explicitly done by No_Bikes in ascending order using sort by No_Bikes asc.

Projection: It projects the required columns (BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp), which minimizes the data returned for consumption.



To Get Premium Files for DP-700 Visit

<https://www.p2pexams.com/products/dp-700>

For More Free Questions Visit

<https://www.p2pexams.com/microsoft/pdf/dp-700>

20%
DISCOUNT

P2P
exams